System Administration

Configure, Deploy, Maintain and Audit

PDF generated using the open source mwlib toolkit. See http://code.pediapress.com/ for more information.
PDF generated at: Mon, 30 Jul 2012 17:16:37 UTC

Contents

Articles

Introduction and Basic Host Configuration

System administrator

A **system administrator**, **IT systems administrator**, **systems administrator**, or **sysadmin** is a person employed to maintain and operate a computer system and/or network. System administrators may be members of an information technology (IT) or Electronics and Communication Engineering department.

The duties of a system administrator are wide-ranging, and vary widely from one organization to another. Sysadmins are usually charged with installing, supporting and maintaining servers or other computer systems, and planning for and responding to service outages and other problems. Other duties may include scripting or light programming, project management for systems-related projects, supervising or training computer operators, and being the consultant for computer problems beyond the knowledge of technical support staff. To perform his or her job well, a system administrator must demonstrate a blend of technical skills and responsibility.

Related fields

Many organizations staff other jobs related to system administration. In a larger company, these may all be separate positions within a computer support or Information Services (IS) department. In a smaller group they may be shared by a few sysadmins, or even a single person.

- A database administrator *(DBA) maintains* a database system, and is responsible for the integrity of the data and the efficiency and performance of the system.
- A network administrator maintains network infrastructure such as switches and routers, and diagnoses problems with these or with the behavior of network-attached computers.
- A security administrator is a specialist in computer and network security, including the administration of security devices such as firewalls, as well as consulting on general security measures.
- A web administrator maintains web server services (such as Apache or IIS) that allow for internal or external access to web sites. Tasks include managing multiple sites, administering security, and configuring necessary components and software. Responsibilities may also include software change management.
- Technical support staff respond to individual users' difficulties with computer systems, provide instructions and sometimes training, and diagnose and solve common problems.
- A computer operator performs routine maintenance and upkeep, such as changing backup tapes or replacing failed drives in a RAID. Such tasks usually require physical presence in the room with the computer; and while less skilled than sysadmin tasks require a similar level of trust, since the operator has access to possibly sensitive data.
- A postmaster is the administrator of a mail server.

In some organizations, a person may begin as a member of technical support staff or a computer operator, then gain experience on the job to be promoted to a sysadmin position.

Training

Unlike many other professions, there is no single path to becoming a system administrator. Many system administrators have a degree in a related field: computer science, information technology, computer engineering, information system management, or even a trade school program. Other schools have offshoots of their Computer Science program specifically for system administration.

Some schools have started offering undergraduate degrees in System Administration. The first, Rochester Institute of Technology[1] started in 1992. Others such as Rensselaer Polytechnic Institute, University of New Hampshire[2], Marist College, and Drexel University have more recently offered degrees in Information Technology. Symbiosis Institute of Computer Studies and Research (SICSR) in Pune, India offers Masters degree in Computers Applications with a specialization in System Administration. The University of South Carolina[3] offers an Integrated Information Technology B.S. degree specializing in Microsoft product support.

As of 2011, only five U.S. universities, Rochester Institute of Technology [4], New York City College of Technology, Tufts, Michigan Tech, and Florida State University [5] have graduate programs in system administration. In Norway, there is a special English-taught MSc program organized by Oslo University College [6] in cooperation with Oslo University, named "Masters programme in Network and System Administration." University of Amsterdam (UvA) offers a similar program in cooperation with Hogeschool van Amsterdam (HvA) named "Master System and Network Engineering". In Israel, the IDF's ntmm course in considered a prominent way to train System administrators. [7] However, many other schools offer related graduate degrees in fields such as network systems and computer security.

One of the primary difficulties with teaching system administration as a formal university discipline, is that the industry and technology changes much faster than the typical textbook and coursework certification process. By the time a new textbook has spent years working through approvals and committees, the specific technology for which it is written may have changed significantly or become obsolete.

In addition, because of the practical nature of system administration and the easy availability of open-source server software, many system administrators enter the field self-taught. Some learning institutions are reluctant to, what is in effect, teach hacking to undergraduate level students.

Generally, a prospective will be required to have some experience with the computer system he or she is expected to manage. In some cases, candidates are expected to possess industry certifications such as the Microsoft MCSA, MCSE, MCITP, Red Hat RHCE, Novell CNA, CNE, Cisco CCNA or CompTIA's A+ or Network+, Sun Certified SCNA, Linux Professional Institute among others.

Sometimes, almost exclusively in smaller sites, the role of system administrator may be given to a skilled user in addition to or in replacement of his or her duties. For instance, it is not unusual for a mathematics or computing teacher to serve as the system administrator of a secondary school.

Skills

Some of this section is from the Occupational Outlook Handbook [8]*, 2010-11 Edition, which is in the public domain as a work of the United States Government.* The *subject matter* of system administration includes computer systems and the ways people use them in an organization. This entails a knowledge of operating systems and applications, as well as hardware and software troubleshooting, but also knowledge of the purposes for which people in the organization use the computers.

Perhaps the most important skill for a system administrator is problem solving -- frequently under various sorts of constraints and stress. The sysadmin is on call when a computer system goes down or malfunctions, and must be able to quickly and correctly diagnose what is wrong and how best to fix it.

System administrators are not software engineers or developers. It is not usually within their duties to design or write new application software. However, sysadmins must understand the behavior of software in order to deploy it and to troubleshoot problems, and generally know several programming languages used for scripting or automation of routine tasks.

Particularly when dealing with Internet-facing or business-critical systems, a sysadmin must have a strong grasp of computer security. This includes not merely deploying software patches, but also preventing break-ins and other security problems with preventive measures. In some organizations, computer security administration is a separate role responsible for overall security and the upkeep of firewalls and intrusion detection systems, but all sysadmins are generally responsible for the security of computer systems.

Duties of a system administrator

A system administrator's responsibilities might include:

- Analyzing system logs and identifying potential issues with computer systems.
- Introducing and integrating new technologies into existing data center environments.
- Performing routine audits of systems and software.
- Performing backups.
- Applying operating system updates, patches, and configuration changes.
- Installing and configuring new hardware and software.
- Adding, removing, or updating user account information, resetting passwords, etc.
- Answering technical queries and dealing with often frustrated users.
- Responsibility for security.
- Responsibility for documenting the configuration of the system.
- Troubleshooting any reported problems.
- System performance tuning.
- Ensuring that the network infrastructure is up and running.

In larger organizations, some tasks listed above may be divided among different system administrators or members of different organizational groups. For example, a dedicated individual(s) may apply all system upgrades, a Quality Assurance (QA) team may perform testing and validation, and one or more technical writers may be responsible for all technical documentation written for a company.

In smaller organizations, the system administrator can also perform any number of duties elsewhere associated with other fields:

- Technical support
- Database administrator (DBA)
- Network administrator/analyst/specialist
- Application analyst
- Security administrator
- Programmer

System administrators, in larger organizations, tend not to be system architects, system engineers, or system designers. However, like many roles in this field, demarcations between system administration and other technical roles often are not well defined in smaller organizations. Even in larger organizations, senior system administrators often have skills in these other areas as a result of their working experience.

In smaller organizations, IT/computing specialties are less often discerned in detail, and the term *system administrator* is used in a rather generic way — they are the people who know how the computer systems work and can respond when something fails.

System Administrator privileges

The term "system administrator" may also be used to describe a security privilege which is assigned to a user or users of a specific computer, server, network or other IT System.

The Administrator level of system access permits that user to gain access to, and perform high level configuration features of the system.

This user privilege level is more commonly referred to within a computer or IT system as "administrator" (without the epithet "system"). It may also be called superuser or root.

For example a computer may have a user named "Administrator" or "Root" which has a security level sufficient to install software, or give other users access to the system. Alternatively a user of a system may be assigned to an "Administrators" group, membership of which grants them the same privilege as the Administrator user. These users may be referred to as System Administrators, referring only to the system privilege level, rather than the job function.

For security reasons, the name of an Administrator user or Administrators security group is often changed locally so that it is less easy to guess, in order to reduce system vulnerability to access by hackers.

References

[1] http://nssa.rit.edu/~nssa/?q=node/8
[2] http://www.cs.unh.edu/bsit.htm
[3] http://www.hrsm.sc.edu/iit/
[4] http://nssa.rit.edu/~nssa/nssa/grad/index.maml
[5] http://www.cs.fsu.edu/current/grad/cnsa_ms.php
[6] http://www.hio.no/Studietilbud/Masterstudier/Master-Programme-in-Network-and-System-Administration
[7] UvA Master SNE homepage (https://www.os3.nl)
[8] http://www.bls.gov/oco/ocos305.htm

Further reading

- Essential System Administration (O'Reilly), 3rd Edition, 2001, by Æleen Frisch
- Essential Linux Administration (Cengage Press): A Comprehensive Guide for Beginners, 2011 by Chuck Easttom
- Principles of Network and System Administration (J. Wiley & Sons), 2000,2003(2ed), by Mark Burgess
- The Practice of System and Network Administration (Addison-Wesley), 2nd Edition (July 5, 2007), by Thomas A. Limoncelli, Christine Hogan and Strata R. Chalup
- Time Management for System Administrators (O'Reilly), 2005, by Thomas A. Limoncelli
- UNIX and Linux System Administration Handbook (Prentice Hall), 4th Edition, 2010, by Evi Nemeth, Garth Snyder, Trent R. Hein, Ben Whaley

External links

- "The Future of IT Jobs in America" article (http://www.ideosphere.com/fx-bin/Claim?claim=ITJOBS)
- The US Department of Labor's description of "Computer Network" (http://www.bls.gov/oco/ocos305.htm) and "Systems, and Database Administrators" (http://www.bls.gov/oco/ocos305.htm) and statistics for employed "System administrator" (http://www.bls.gov/oco/oco2001.htm#emply)
- BSD Certification (http://www.bsdcertification.org/)
- ITIL (http://www.itil.co.uk/) for ITIL certification (part of Office of Government Commerce)
- BigAdmin Newsletter (http://www.sun.com/bigadmin/newsletter/)
- Art du web.com : the website of system administrator (http://www.artduweb.com/)
- Administration système : Vie et mort des sysadm ;) (http://www.administration-systeme.com/)
- Linux System Administration Information (http://linux-administration-pro.com)

- A System Admin site to share Tools, Script, Plugins and Books (http://www.isystemadmin.com)
- Guide to the System Administration Body of Knowledge (https://lopsa.org/content/sabok)
- Info Blog For System Administrator (http://www.xiitec.com/blog/)

Usage share of operating systems

The **usage share of operating systems** is the percentage market share of the operating systems used in computers. Different categories of computers use a wide variety of operating systems, so the total usage share varies enormously from one category to another.

In some categories, one family of operating systems dominates. For example, most desktop and laptop computers use Microsoft Windows and most supercomputers use Linux. In other categories, such as smartphones and servers, there is more diversity and competition.

Information about operating system share is difficult to obtain, since in most of the categories below there are no reliable primary sources or methodology for its collection.

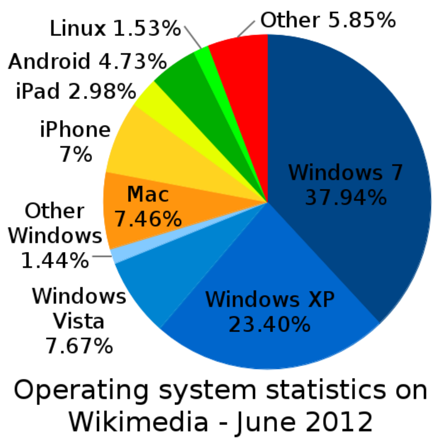

Piechart of OS share[1]

Desktop and laptop computers

There is little openly published information on the usage share of desktop and laptop computers. Gartner publishes estimates, but the way the estimates are calculated is not openly published. Also, sales may overstate usage. Most computers are sold with a pre-installed OS; some users replace that OS with a different one due to personal preference. Conversely, sales underestimate usage, by not counting pirated copies. For example, in 2009, "U.S. research firm IDC estimated that 80% of software sold in China last year was pirated." (Windows was mentioned, but no specific estimate for Windows was given.)[2] As another example, in 2007, the automated push of IE7 update onto legal copies of Windows, contrasted with web browser share statistics, led one author to "estimate that 25%–35% of all Windows XP machines are illegal".[3]

Estimates for 2011

These are current-year sales estimates, not accumulative usage share for all PCs that are in use.

In August 2011, Gartner estimated Apple's PC market share in US as 10.7% for Q2 2011. Apple's worldwide market share is not listed, because it is not in the list of top 5 computer manufacturers, and is inferred to be 5% or lower. Gartner's numbers include netbooks, but not media tablets such as the iPad. Total units in Q2 2011 from all vendors, ~85 million.[4]

"A Gartner forecast calls for Mac OS to ship on 4.5 percent of new PCs worldwide in 2011 and 5.2 percent in 2012. Gartner does not expect Google Chrome OS, Google Android or HP's webOS to get 'any significant market share' on PCs in the next few years, and expects Linux operating systems to remain at less than 2 percent share over the next several years."[5]

Analyst Charlie Wolf of Needham & Co. has stated that in 2011 Apple broke 5% of global desktop market share for the first time.[6]

Net Applications has estimated global desktop market share of Windows 92.2%, Apple 6.36%, Linux 1.41%.[7][8]

Older information

In a speech to investors in February 2009, Steve Ballmer of Microsoft presented a slide based on Microsoft's research; while it showed no figures, the pie chart depicted Linux and Apple as each having roughly 5–6% of home and business PCs.[9]

Web client usage share (see below) is often used as a proxy for desktop share, but many desktops are not used for web access so do not get counted in these figures. The correlation between desktop share and web client share is also being challenged by mobile web access, which rose through 1% in 2009 and 4% in 2010.[10]

Microsoft's CFO Peter Klein stated in July 2010 that Windows 7 now runs on more than 15% of all PCs worldwide.[11] A Forrester Research study of desktop operating systems used in North American and European companies in 2010 found Windows 7 on 10% of all commercial desktops, Windows XP on 75% and Vista on 7%.[12]

In October 2010 Tim Cook of Apple claimed that 1 in 5 of desktop/laptop computers sold in the United States is a Mac.[13]

Web clients

The following information on web clients is obtained from the user agent information supplied to web servers by web browsers. This is an inexact science for a variety of reasons. For a discussion on the shortcomings see Usage share of web browsers.

The most recent data from various sources published during the last six months is summarized in the table below. All of these sources monitor a substantial number of web sites; statistics related to one web site only are excluded.

Source	Date	Microsoft Windows				Apple		Linux kernel based		Symbian	Black-Berry OS	Other
		7	Vista	XP	All versions	OS X	iOS	GNU Linux	Android			
AT Internet[14]	Aug. 2011	33.0%	14.7%	38.4%	87.1%	6.7%	3.7%	1.0%	0.9%	—	—	0.6%
Clicky Web Analytics[15]	Oct. 2011	33.67%	10.17%	26.11%	79.53%	13.25%	3.81%	1.16%	1.77%	0.12%	0.32%	0.04%
Chitika[16]	Oct. 2011	—	—	—	77.2%	10.2%	5.4%	2.9%	3.3%	—	—	1.0%
Net Market Share[17][18]	May. 2012	37.35%	6.28%	40.93%	84.83%	6.09%	4.87%	1.03%	1.53%	0.14%	0.16%	1.35%
StatCounter Global Stats[19][20][21]	Apr. 2012	43.02%	8.07%	28.74%	82.21%	6.92%	2.97%	0.77%	2.28%	2.73%	0.59%	1.53%
StatOwl[22]	Oct. 2011	33.24%	12.91%	33.20%	78.94%	14.04%	—	0.70%	—	—	—	6.32%
W3Counter[23]	Jun. 2012	43.84%	7.36%	27.59%	78.89%	8.78%	6.69%	1.80%	2.15%	0.20%	0.66%	3.49%
Webmasterpro[24]	Nov. 2011	38.3%	14.1%	32.1%	85.8%	6.3%	4.4%	1.4%	1.6%	0.2%	0.02%	0.28%
Wikimedia[25]	May. 2012	37.59%	7.64%	24.60%	71.30%	8.20%	9.13%	1.61%	4.22%	0.16%	0.83%	4.55%
Notes:												

- The 'Other' column is obtained by summing Windows 'all versions' through BlackBerry OS and subtracting from 100%. ('Other' may be quite different between different sources. If a source does not show a value for a given OS, that OS is included in the source's 'Other' column.)
- AT Internet measures 23 European countries.
- Clicky Web Analytics does not publish desktop/mobile split so mean of Net Market Share and StatCounter figures (6.035% mobile) used in lieu. Figures are averages over last 7 days of month.
- StatOwl measures desktop share and predominantly US web sites with "broad appeal".[26] Figure for XP includes Server 2003. Figures reduces by 6.035% (see above) to allow for missing mobile usage.
- W3Counter shows only the top ten operating systems and is based on the last 15,000 page views to each of over 50,000 web sites tracked.
- Webmasterpro samples over 100,000 predominantly German-language sites. Figures are seven-day averages over period Nov 9-15.
- Wikimedia uses 1:1000 sampling of its logs when deriving the usage numbers. Figure for Vista includes Server 2008; XP includes Server 2003.
- iOS figures include iPhone, iPod and iPad.
- OS X is broken down by four of the sources listed above and all of them show that version 10.6 (Snow Leopard) is the most widely used.

Netbooks

The netbook market has been dominated by Microsoft Windows, with Linux in second place.

Initially, Linux dominated the netbook market when Asus started it with the Eee PC in October 2007, but this lead did not last long. Asus and Acer, which accounted for 90% of the early netbook market, installed Linux on 30% of their machines.[27]

Microsoft responded by extending the sales of Windows XP Home Edition. By February 2009, Microsoft cited data from NPD Retail Tracking Service which showed that US market share of Windows on netbooks went from under 10% to 96%.[28]

In November 2009, an analyst at ABI said that of the 35 million netbooks to ship globally in 2009, 68% would have Windows and 32% Linux.[29]

According to DisplaySearch, netbooks and tablets rose from just under a 14% share of the overall portable computer market in third quarter of 2008 to around 20% in the second quarter of 2009, and remained at around 20% until the middle of 2010. During 2010, Apple's iPad tablet computer gained a 6.5% share of this market sector in the first quarter and DisplaySearch forecast this will rise to 30% in the second.[30]

Tablet computers

Tablet computers, or simply *tablets*, became a significant OS market share category starting with Apple's iOS-based iPad in 2010. As of June 2011, ~29 million iPads sold.

2011 Sales and estimates (millions of units):

iOS: Q1: 7.3, Q2: 4.7 (limited by supply shortages)[31]

Android: "Android media tablets have collectively taken 20% market share away from the iPad in the last 12 months."[32]

"Apple is set to increase its iPad shipments at a faster rate than previously expected in 2011 and beyond, causing the global tablet market to exceed growth expectations during the next few years, according to data from information and analysis provider IHS. Apple will ship 44 million iPads in 2011; shipments expected to reach 120 million units in 2015. Apple is expected to account for 74 percent of tablet shipments in 2011 and 43-44 percent in 2015."[33]

Top vendors: In Q1 2011, Apple's iOS sold 7.3 million tablets. (Q2 number not used, because was limited by supply shortages.) Google's Android shipped on 1–2 million tablets per quarter (20% estimate by ABI).

Source	Date	iOS	Android	Microsoft Windows	Other
Strategy Analysis[34]	Q4 2011	57.6%	39.1%	1.5%	1.9%

Mobile devices

Mobile operating systems that can be found on smartphones include Nokia's Symbian, Apple's iOS, Research in Motion's BlackBerry OS, Microsoft's Windows Mobile and its successor Windows Phone 7, Google's Android, Samsung's Bada, and HP's webOS. Android and webOS are in turn built on top of Linux, and iOS is derived from OS X which in turn is built upon the BSD and NeXTSTEP operating systems. Linux, BSD, and NeXTSTEP are all related to Unix.

Gartner's Q3 2011 unit numbers total 115 million with Google's Android shipping on 60 million smartphones, Nokia's Symbian on 19 million and Apple's iOS on 17 million.[35]

Canalys' Q2 2011 unit numbers total 108 million with Google's Android shipping on 52 million smartphones, Apple's iOS on 20 million, and Nokia's Symbian on < 20 million.[36] Nokia confirms its Q2 smartphone sales were 16.7 million units.[37][38]

Total unit forecasts for 2011: (IDC) 472 million;[39] (Gartner) 468 million.[40]

Forecasts for 2012: (Gartner): 630 million units; Android 49%, iOS 19%, BlackBerry 13%, Windows 11%, Symbian 5%, Other 3%.[40] (Taiwan/Market Intelligence Center): Android 40%, iOS 19%, Windows 17%, Other 24%.[41] (IDC) 582 million units total.[42]

Forecasts for 2015: (Gartner): 1105 million units; Android 49%, Windows 20%, iOS 17%, BlackBerry 11%, Other 3%.[40] (IDC): 982 million units; Android 44%, Windows 20%, iOS 17%, BlackBerry 13%, Other 6%.[39]

Source	Date	Symbian	BlackBerry	iOS	Android	Bada	Windows	Other
Gartner[35]	Q3 2011	16.9%	11.0%	15.0%	52.5%	2.2%	1.5%	0.9%
Canalys[36]	Q2 2011	**	**	19%	48%	**	1%	**
Canalys[43]	Q4 2010	31.0%	14.6%	16.2%	33.3%	*	3.1%	3.0%
Kantar Worldpanel[44]	Jan 2012	13.1%	18.1%	28.5%	36.9%	0.7%	2.3%	0.4%
Stat Counter[45]	April 2012	28.45%	6.1%	23.85%	23.79%	0.57%		1.90%
Wikimedia[25]	August 2011	2.16%	6.91%	67.09%	22.94%		0.90%	
NPD Group (*US only*)[46]	Q2 2011		11%	29%	52%			
Millennial Media (*US only*)[47][48]	August 2011	1%	13%	28%	54%	*	1%	3%
Nielsen Company (*US only*)[49]	Q4 2011	1.4%	14.9%	30%	46.3%		5.9%	
Comscore (*US only*)[50]	December 2011	1.4%	16.0%	29.6%	47.3%		4.7%	

Notes:

- The above table is share of smartphone OSs - not overall marketshare.
- StatCounter measures share of mobile web browsing
- Millennial Media measures share of ad impressions

(*) Included in "Other". (**) Source numbers are per device vendor; breakdown by OS is incomplete.

Servers

Server market share can be measured with statistical surveys of publicly accessible servers, such as web servers, mail servers[51] or DNS servers on the Internet: the operating system powering such servers is found by inspecting raw response messages. This method gives insight only into market share of operating systems that are publicly accessible on the Internet.

A typical server "rack"

There will be differences in the result depending on how the sample is done and observations weighted. Usually the surveys are not based on a random sample of all ip numbers, domain names, hosts or organisations, but on servers found by some other method. Additionally many domains and ip numbers may be served by one host and some domains may be served by several hosts or by one host with several ip numbers.

Source	Date	Method	Unix, Unix like				Microsoft Windows	References
			All	Linux	BSD	Solaris		
W3Techs	August 2011	Units (Web)	63.90%				36.10%	[52][53]
Security Space	August 2009	Units (Web)		>60.00%				[54][55]

Notes:

- W3Techs survey in August 2011 checked the top 1 million Web servers (according to Alexa).
- Security Space survey in August 2009 checked 38,549,333 publicly accessible Web servers.
- Netcraft SSL survey[56] in January 2009 also checked 1,014,301 publicly accessible Web servers, but the survey is only valid for SSL Web servers and it is not a good measure for our purpose.

A method to measure the overall server market, rather than subsets like publicly accessible web servers, is to count server hardware sales, using data from server manufacturers. Using this method, market share can be measured either in units or in revenue. In either case, the measure refers to server hardware, not to software. Units refers to the number of physical servers running a given OS, and revenue refers to hardware revenue for physical servers running that OS. It does not refer to software licensing or support revenue, which often varies considerably from one OS to another.[57]

Source	Date	Method	Microsoft Windows	Unix, Unix like		References
				All	Linux	
IDC	Q1 2012	Revenue	50.2%	38.9%	20.7%	[57]

Mainframes

Nearly 95% of Fortune 1000 companies use IBM's Information Management System.[58]

Operating systems for IBM System z generation hardware include IBM's bundled proprietary z/OS,[58] Linux on System z and as at October 7, 2008 the prototype OpenSolaris for System z.

Gartner reported on December 23, 2008 that Linux on System z was used on approximately 28% of the "customer z base" and that they expected this to increase to over 50% in the following five years.[59]

Of Linux on System z, Red Hat and Novell compete to sell RHEL and SLES respectively.

- Prior to 2006, Novell claimed a market share of 85% or more.
- Red Hat has since claimed 18.4% in 2007 and 37% in 2008.[60]
- Gartner reported at the end of 2008 that Novell had an 80% share of mainframe Linux.[59]

Supercomputers

The TOP500 project lists and ranks the 500 fastest supercomputers that benchmark results are submitted for. It then publishes the collected data twice a year. The November 2011 figures are below.

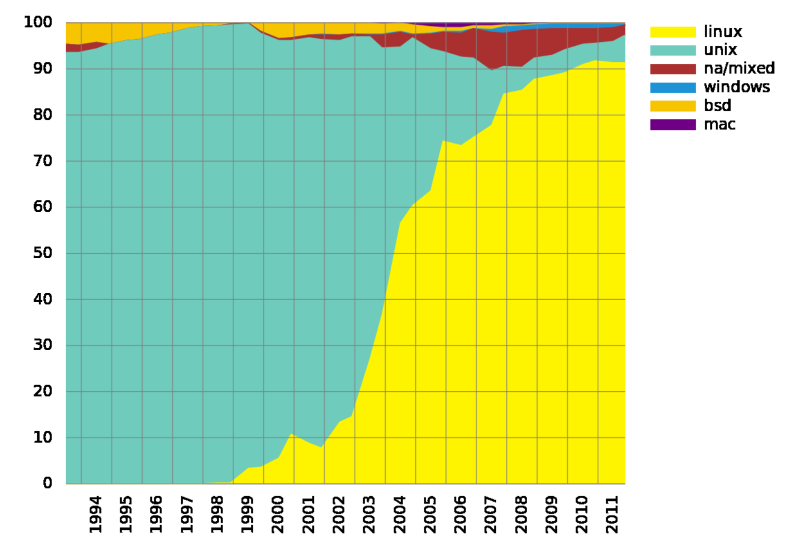

Graph of supercomputer OS market share according to TOP500.[61]

Source	Date	Linux	IBM AIX	Other Unix	Microsoft HPCS 2008	Other	References
TOP500	November 2011	91.4%	5.6%	0.4%	0.2%	2.4%.	[62][63]

References

[1] *Wikimedia Traffic Analysis Report - Operating Systems, 2012-04* (http://stats.wikimedia.org/archive/squid_reports/2012-04/SquidReportOperatingSystems.htm), Wikimedia,

[2] "Software Pirates in China Beat Microsoft to the Punch" (http://www.nytimes.com/2009/10/19/business/global/19iht-windows.html). *The New York Times*. 19 October 2009. .

[3] "Around 25-35% of Windows XP systems are pirated (calculations included)" (http://www.seopher.com/articles/around_25_35_of_windows_xp_systems_are_pirated_calculations_included_). .

[4] "Gartner Says Worldwide PC Shipments Increased 2.3 Percent in Second Quarter of 2011" (http://www.gartner.com/it/page.jsp?id=1744216). 13 July 2011. .

[5] "Which operating system will be 2011's bestseller?" (http://gcn.com/Articles/2011/08/10/ECG-Windows-7-Top-Selling-OS-by-End-of-2011.aspx?Page=2). 11 August 2011. .

[6] "Mac platform reaches 15-year high with 5% worldwide market share" (http://www.appleinsider.com/articles/11/11/16/macs_worldwide_market_share_reaches_15_year_high_at_5/). 2011-11-16. .

[7] "Linux gains share on enterprise servers -- and desktops, too" (http://www.linuxfordevices.com/c/a/News/Linux-Foundation-enterprise-Linux-survey-plus-Net-Applications-desktop-stats/). 2012-01-19. .

[8] "Is the Linux Desktop actually growing?" (http://www.zdnet.com/blog/open-source/is-the-linux-desktop-actually-growing/10160). 2012-01-18. .

[9] Holwerda, Thom (25 February 2009). "Ballmer: Linux Bigger Competitor than Apple" (http://www.osnews.com/story/21035/Ballmer_Linux_Bigger_Competitor_than_Apple). osnews.com. .

[10] http://gs.statcounter.com/#mobile_vs_desktop-ww-monthly-200812-201011

[11] "Windows 7, Office Drive Record Microsoft Revenue" (http://www.toptechnews.com/story.xhtml?story_id=11300CM9DYVG). Top Tech News. 23 July 2010. .

[12] "Updated 2010: Windows 7 Commercial Adoption Outlook" (http://www.forrester.com/rb/Research/updated_2010_windows_7_commercial_adoption_outlook/q/id/58012/t/2). Forrester Research. 2 November 2010.. .

[13] "App Store comes to Mac in 90 days, new iLife Suite and trimmer MacBook Air available now" (http://www.betanews.com/joewilcox/article/App-Store-comes-to-Mac-in-90-days-new-iLife-Suite-and-trimmer-MacBook-Air-available-now/1287600183). Betanews Inc.. 20 October 2010. .

[14] *Operating systems, August 2011* (http://en.atinternet.com/Resources/Surveys/Internet-user-equipment/Operating-systems-August-2011/index-1-2-7-246.aspx),

[15] *Operating Systems* (http://getclicky.com/marketshare/global/operating-systems/), Get Clicky,

[16] *Operating System Market Report: October, 2011 Update* (http://insights.chitika.com/2011/operating-system-market-report-october-2011-update/), Chitika,

[17] *Operating System Market Share* (http://marketshare.hitslink.com/operating-system-market-share.aspx?qprid=8), Hits link,

[18] *Operating System Market Share* (http://marketshare.hitslink.com/operating-system-market-share.aspx?qprid=10), Hits link,

[19] *OS monthly 2012-04* (http://gs.statcounter.com/#os-ww-monthly-201204-201204-bar), StatCounter,

[20] *Mobile OS monthly 2012-04* (http://gs.statcounter.com/#mobile_os-ww-monthly-201204-201204-bar), StatCounter,

[21] *Mobile vs. desktop monthly 2012-04* (http://gs.statcounter.com/#mobile_vs_desktop-ww-monthly-201204-201204-bar), StatCounter,

[22] *Operating system market share* (http://statowl.com/operating_system_market_share.php), StatOwl,

[23] *Global stats* (http://www.w3counter.com/globalstats.php?year=2012&month=6), W3 Counter,

[24] (in German) *Webanalyse - Betriebssysteme und Geräte* (http://www.webmasterpro.de/portal/webanalyse-systeme.html), DE: Webmaster Pro,

[25] *Squid report: Operating systems 2011-10* (http://stats.wikimedia.org/archive/squid_reports/2011-12/SquidReportOperatingSystems.htm), Wikimedia,

[26] *About our data* (http://statowl.com/about_our_data.php), StatOwl,

[27] Culpan, Tim; Bass, Dina (6 November 2008). "Microsoft Missing Netbook Growth as Linux Wins Sales (Update2)" (http://www.bloomberg.com/apps/news?pid=20601103&sid=a3VyE_ofSwwE&refer=news). Bloomberg.com. .

[28] "Windows on Netbook PCs: A Year in Review" (http://windowsteamblog.com/windows/b/windowsexperience/archive/2009/04/03/windows-on-netbook-pcs-a-year-in-review.aspx). Windows Experience Blog. .

[29] Lai, Eric (4 November 2009). "Linux's share of netbooks surging, not sagging, says analyst" (http://www.computerworld.com/s/article/9140343/Linux_s_share_of_netbooks_surging_not_sagging_says_analyst). computerworld.com. .

[30] Smith, Tony (17 June 2010). "iPad gouges netbook sales" (http://www.reghardware.com/2010/06/17/netbooks_vs_tablets/). The Register. . Retrieved 18 January 2011.

[31] "Apple unveils fiscal 2011 Q2 results" (http://www.unwiredview.com/2011/04/20/apple-unveils-fiscal-2011-q2-results-18-65-million-iphones-4-69-million-ipads-sold/). .

[32] "Android Takes 20% Media Tablet Market Share from iPad in Last 12 Months" (http://www.abiresearch.com/press/3753-Android+Takes+20%+Media+Tablet+Market+Share+from+iPad+in+Last+12+Months). 11 August 2011. .

[33] "Media Tablet Forecast Increased as Apple’s Dominance Grows" (http://www.isuppli.com/Display-Materials-and-Systems/News/Pages/Media-Tablet-Forecast-Increased-as-Appleâs-Dominance-Grows.aspx). 24 August 2011. .

[34] "Android Captures Record 39 Percent Share of Global Tablet Shipments in Q4 2011" (http://www.strategyanalytics.com/default.aspx?mod=pressreleaseviewer&a0=5167). .
[35] "Gartner Says Sales of Mobile Devices Grew 5.6 Percent in Third Quarter of 2011; Smartphone Sales Increased 42 Percent" (http://www.gartner.com/it/page.jsp?id=1848514). Gartner. 15 November 2011. . Retrieved 21 December 2011.
[36] "Android takes almost 50% share of worldwide smart phone market" (http://www.canalys.com/newsroom/android-takes-almost-50-share-worldwide-smart-phone-market). Canalys. 1 August 2011. .
[37] "Nokia Q2 2011 results - smartphone sales fall to 16.7m" (http://www.allaboutsymbian.com/news/item/13121_Nokia_Q2_2011_results-smartpho.php). All About Symbian. 21 July 2011. .
[38] "Nokia Q2 2011 net sales EUR 9.3 billion, non-IFRS EPS EUR 0.06 (reported EPS EUR -0.10)" (http://press.nokia.com/2011/07/21/nokia-q2-2011-net-sales-eur-9-3-billion-non-ifrs-eps-eur-0-06-reported-eps-eur-0-10/). Nokia. 21 July 2011. .
[39] "Worldwide Smartphone Market Expected to Grow 55% in 2011 and Approach Shipments of One Billion in 2015, According to IDC" (http://www.idc.com/getdoc.jsp?containerId=prUS22871611). IDC. 9-Jun-2011. .
[40] "Gartner: Microsoft Windows Phone market share to surpass Apple’s iOS in 2015" (http://macdailynews.com/2011/04/07/gartner-microsoft-windows-phone-market-share-to-surpass-apples-ios-in-2015/). Gartner. 9-Sep-2011. .
[41] "HTC expected to overtake RIM in '12: MIC" (http://www.chinapost.com.tw/business/company-focus/2011/06/30/308129/HTC-expected.htm). The China Post. 30-Jun-2011. .
[42] "The Next Smart Revolution - Smart Enterprise" (http://www.accenture.com/microsites/gcf-2011/Documents/Jongshin_Kim_Samsung_GCF2011.pdf). Samsung. . Retrieved 11 September 2011.
[43] "Google’s Android becomes the world’s leading smart phone platform" (http://www.canalys.com/newsroom/google’s-android-becomes-world’s-leading-smart-phone-platform). 2011-06-30. .
[44] "Android is most-used smartphone OS in UK, overtaking Apple, data shows" (http://www.guardian.co.uk/technology/2012/feb/21/android-smartphones-os-uk-apple). .
[45] "StatCounter Global Stats: Top 8 Mobile OSs on 12 April" (http://gs.statcounter.com/#mobile_os-ww-monthly-201204-201204-bar). .
[46] "The NPD Group: As Android Solidifies Lead, Google Acquisition Has Potential to Revitalize Flagging Motorola" (http://www.npd.com/press/releases/press_110822a.html). 22 August 2011. . Retrieved 24 August 2011.
[47] http://www.millennialmedia.com/wp-content/images/mobilemix/MM-MobileMix-August2011.pdf
[48] "Millennial: Android top OS in August" (http://www.teamphones.com/apple-ios/millennial-android-top-os-in-august-with-nearly-double-of-appleâ€™s-ios-share-314.html). Team Phones. 23 September 2011. .
[49] "More US Consumers Choosing Smartphones as Apple Closes the Gap on Android" (http://blog.nielsen.com/nielsenwire/consumer/more-us-consumers-choosing-smartphones-as-apple-closes-the-gap-on-android/). 18 January 2012. . Retrieved 16 February 2012.
[50] "comScore Reports December 2011 U.S. Mobile Subscriber Market Share" (http://www.comscore.com/Press_Events/Press_Releases/2012/2/comScore_Reports_December_2011_U.S._Mobile_Subscriber_Market_Share). comScore. 2 February 2012. . Retrieved 16 February 2012.
[51] "Mail Server Survey" (http://www.securityspace.com/s_survey/data/man.201007/mxsurvey.html). Security Space. August 2011. .
[52] "Usage of operating systems for websites" (http://w3techs.com/technologies/overview/operating_system/all). W3Techs. August 2011. .
[53] "Usage of Unix for websites" (http://w3techs.com/technologies/details/os-unix/all/all). W3Techs. August 2011. .
[54] "Web Server Survey" (https://secure1.securityspace.com/s_survey/data/200907/index.html). Security Space. August 2009. .
[55] "OS/Linux Distributions using Apache" (https://secure1.securityspace.com/s_survey/data/man.200907/apacheos.html). Security Space. August 2009. .
[56] "Operating System Share by Groups for Sites in All Locations January 2009" (http://news.netcraft.com/SSL-Survey/CMatch/osdv_all). .
[57] "Worldwide Server Market Revenues Decline 2.4% in First Quarter as Market Growth Slows in Face of Market Transitions, According to IDC" (http://www.idc.com/getdoc.jsp?containerId=prUS23513412). IDC. 30 May 2012. .
[58] "IBM Tightens Stranglehold Over Mainframe Market; Gets Hit with Antitrust Complaint in Europe" (http://www.ccianet.org/index.asp?sid=5&artid=62&evtflg=False). Computer & Communications Industry Association. 2 July 2008. .
[59] "Vendor Rating: Novell, 2008" (http://mediaproducts.gartner.com/reprints/novell/vol3/article1/article1.html). *Gartner RAS Core Research Note G00162399*. 23 December 2008. . Retrieved 23 November 2009.
[60] Claybrook, Bill (1 September 2009). "Red Hat bolsters Linux for mainframes, tries to catch Novell" (http://searchdatacenter.techtarget.com/tip/0,289483,sid80_gci1366811,00.html). SearchDataCenter.com. . Retrieved 23 November 2009.
[61] "Top500 OS chart" (http://www.top500.org/overtime/list/32/os). . Retrieved 18 January 2011.
[62] "Operating system Family share for 11/2011" (http://top500.org/stats/list/38/osfam). Top 500 project. .
[63] "Operating System share for 11/2011" (http://top500.org/stats/list/38/os). Top 500 project. .

Linux distribution

A **Linux distribution** is a member of the family of Unix-like operating systems built on top of the Linux kernel. Such distributions (often called *distros* for short) are operating systems including a large collection of software applications such as word processors, spreadsheets, media players, and database applications. These operating systems consist of the Linux kernel and, usually, a set of libraries and utilities from the GNU Project, with graphics support from the X Window System. Distributions optimized for size may not contain X and tend to use more compact alternatives to the GNU utilities, such as BusyBox, uClibc, or dietlibc. There are currently over six hundred Linux distributions. Over three hundred of those are in active development, constantly being revised and improved.

Because most of the kernel and supporting packages are free and open source software, Linux distributions have taken a wide variety of forms—from fully featured desktop, server, laptop, netbook, mobile phone, and tablet operating systems as well as minimal environments (typically for use in embedded systems or for booting from a floppy disk). Aside from certain custom software (such as installers and configuration tools), a distribution is most simply described as a particular assortment of applications installed on top of a set of libraries married with a version of the kernel, such that its "out-of-the-box" capabilities meet most of the needs of its particular end-user base.

One can distinguish between commercially-backed distributions, such as Fedora (Red Hat), openSUSE (Novell), Ubuntu (Canonical Ltd.), and Mandriva Linux (Mandriva), and entirely community-driven distributions, such as Debian and Gentoo.

History

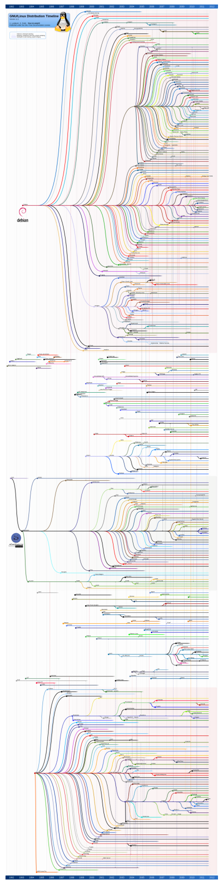

GNU/Linux Distro Timeline [1], timeline representing the development of various Linux distributions.

Before the first Linux distributions, a would-be Linux user was required to be something of a Unix expert, needing to know not only what libraries and executables were required to successfully get the system to boot and run, but also important details concerning configuration and placement of files in the system.

Linux distributions began to appear very soon after the Linux kernel was first used by individuals other than the original Linux programmers who were more interested in developing the operating system than developing application programs, the user interface, or convenient packaging.

Early distributions included:

- H J Lu's "Boot-root" a two disk pair with the kernel and the absolute minimal tools to get started
- MCC Interim Linux, which was made available to the public for download on the FTP server of University of Manchester in February 1992
- TAMU, created by individuals at Texas A&M University about the same time
- SLS (Softlanding Linux System)
- Yggdrasil Linux/GNU/X, the first CD-ROM based Linux distribution

SLS was not well maintained, so Patrick Volkerding released a distribution based on SLS, which he called Slackware, released in 1993.[2] This is the oldest distribution still in active development.

Users were attracted to Linux distributions as alternatives to the DOS and Microsoft Windows operating systems on the PC, Mac OS on the Apple Macintosh, and proprietary versions of Unix. Most early adopters were familiar with Unix from work or school. They embraced Linux for its stability, low (if any) cost, and availability of the source code for most or all of the software included.

Originally, the distributions were simply a convenience, but today, they have become the usual choice even for Unix or Linux experts. To date, Linux has proven more popular in the server market, primarily for Web and database servers (*see also* LAMP), and in embedded devices market than in the desktop market.

Components

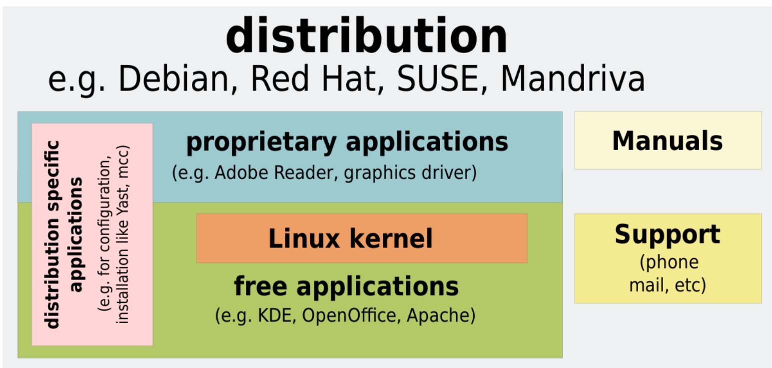

components of a Linux distribution

A typical desktop Linux distribution comprises a Linux kernel, GNU tools and libraries, additional software, documentation, a window system, window manager, and a desktop environment. Most of the included software is free software/open-source software which is distributed by its maintainers both as compiled binaries and in source code form, allowing users to modify and compile the original source code if they wish. Other software included with some distributions may be proprietary and may not be available in source code form.

Many distributions provide an installation system akin to that provided with other modern operating systems. Some distributions like Gentoo Linux, T2, and Linux From Scratch include only binaries of a basic kernel, compilation tools, and an installer; the installer compiles all the requested software for the specific microarchitecture of the user's machine, using these tools and the provided source code.

Package management

Distributions are normally segmented into **packages**. Each package contains a specific application or service. Examples of packages are a library for handling the PNG image format, a collection of fonts or a web browser.

The package is typically provided as compiled code, with installation and removal of packages handled by a package management system (PMS) rather than a simple file archiver. Each package intended for such a PMS contains meta-information such as a package description, version, and "dependencies". The package management system can evaluate this meta-information to allow package searches, to perform an automatic upgrade to a newer version, to check that all dependencies of a package are fulfilled, and/or to fulfill them automatically.

Although Linux distributions typically contain much more software than proprietary operating systems, it is normal for local administrators to also install software not included in the distribution. An example would be a newer version of a software application than that supplied with a distribution, or an alternative to that chosen by the distribution (*e.g.*, KDE rather than GNOME or vice versa for the user interface layer). If the additional software is distributed in source-only form, this approach requires local compilation. However, if additional software is locally added, the 'state' of the local system may fall out of synchronization with the state of the package manager's database. If so, the local administrator will be required to take additional measures to ensure the entire system is kept up to date. The package manager may no longer be able to do so automatically.

Most distributions install packages, including the kernel and other core operating system components, in a predetermined configuration. Few now require or even permit configuration adjustments at first install time. This makes installation less daunting, particularly for new users, but is not always acceptable. For specific requirements, much software must be carefully configured to be useful, to work correctly with other software, or to be secure, and local administrators are often obliged to spend time reviewing and reconfiguring assorted software.

Some distributions go to considerable lengths to specifically adjust and customize most or all of the software included in the distribution. Not all do so. Some distributions provide configuration tools to assist in this process.

By replacing *everything* provided in a distribution, an administrator may reach a "distribution-less" state: everything was retrieved, compiled, configured, and installed locally. It is possible to build such a system from scratch, avoiding a distribution altogether. One needs a way to generate the first binaries until the system is *self-hosting*. This can be done via compilation on another system capable of building binaries for the intended target (possibly by cross-compilation). For example, see Linux From Scratch.

Types and trends

Broadly, Linux distributions may be:

- Commercial or non-commercial;
- Designed for enterprise users, power users, or for home users;
- Supported on multiple types of hardware, or platform-specific, even to the extent of certification by the platform vendor;
- Designed for servers, desktops, or embedded devices;
- General purpose or highly specialized toward specific machine functionalities (e.g. firewalls, network routers, and computer clusters);
- Targeted at specific user groups, for example through language internationalization and localization, or through inclusion of many music production or scientific computing packages;

- Built primarily for security, usability, portability, or comprehensiveness.

The diversity of Linux distributions is due to technical, organizational, and philosophical variation among vendors and users. The permissive licensing of free software means that any user with sufficient knowledge and interest can customize an existing distribution or design one to suit his or her own needs.

Installation-free distributions (Live CDs)

A Live Distro or Live CD is a Linux distribution that can be booted from a compact disc or other removable medium (such as a DVD or USB flash drive) instead of the conventional hard drive. Some minimal distributions such as tomsrtbt can be run directly from as little as one floppy disk without needing to change the system's hard drive contents.

When the operating system is booted from a read-only device such as a CD or DVD, if user data needs to be retained between sessions, it cannot be stored on the boot device but must be written to some other media such as a USB flash drive or an installed hard drive. Temporary operating system data is usually kept solely in RAM.

The portability of installation-free distributions makes them advantageous for applications such as demonstrations, borrowing someone else's computer, rescue operations, or as installation media for a standard distribution. Many popular distributions come in both "Live" and conventional forms (the conventional form being a network or removable media image which is intended to be used for installation only). This includes SUSE, Ubuntu, Linux Mint, MEPIS, Sidux, and Fedora. Some distributions, such as Knoppix, Devil-Linux, SuperGamer, and dyne:bolic are designed primarily for Live CD, Live DVD, or USB flash drive use.

Examples

Popular distributions

Well-known Linux distributions include:

- Arch Linux, a minimalist rolling release distribution targeted at experienced Linux users, maintained by a volunteer community and primarily based on binary packages in the tar.gz and tar.xz format.
- Debian, a non-commercial distribution maintained by a volunteer developer community with a strong commitment to free software principles
  - Knoppix, the first Live CD distribution to run completely from removable media without installation to a hard disk, derived from Debian
  - Linux Mint Debian Edition (LMDE) is based directly on Debian's *testing* distribution.
  - Ubuntu, a popular desktop and server distribution derived from Debian, maintained by British company Canonical Ltd.
    - BackTrack, based on the Ubuntu operating system. Used for digital forensics and penetration testing.
    - Kubuntu, the KDE version of Ubuntu.
    - Linux Mint, a distribution based on and compatible with Ubuntu. Currently using Gnome 3 shell, "Cinnamon", or optionally Gnome 2 fork, MATE.
    - Xubuntu is the Xfce version of Ubuntu.
    - Lubuntu, the LXDE version of Ubuntu, advertised as a lightweight distribution.
- Fedora, a community distribution sponsored by American company Red Hat
  - Red Hat Enterprise Linux, which is a derivative of Fedora, maintained and commercially supported by Red Hat.
    - CentOS, a distribution derived from the same sources used by Red Hat, maintained by a dedicated volunteer community of developers with both 100% Red Hat-compatible versions and an upgraded version that is not always 100% upstream compatible

    - Oracle Enterprise Linux, which is a derivative of Red Hat Enterprise Linux, maintained and commercially supported by Oracle.
  - Mandriva, a Red Hat derivative popular in several European countries and Brazil, today maintained by the French company of the same name.
    - PCLinuxOS, a derivative of Mandriva, grew from a group of packages into a community-spawned desktop distribution.
- Gentoo, a distribution targeted at power users, known for its FreeBSD Ports-like automated system for compiling applications from source code
- openSUSE a community distribution mainly sponsored by American company Novell.
  - SUSE Linux Enterprise, derived from openSUSE, maintained and commercially supported by Novell.
- Slackware, one of the first Linux distributions, founded in 1993, and since then actively maintained by Patrick J. Volkerding.
- Damn Small Linux, "DSL" is a Biz-card Desktop OS

DistroWatch attempts to include every known distribution of Linux, whether currently active or not; it also maintains a ranking of distributions based on page views, as a measure of relative popularity.

Niche distributions

Other distributions are targeted at other specific niches, such as the tiny embedded router distribution OpenWrt, distributions for bioinformatics, the Ubuntu project to create Edubuntu for educational users, and KnoppMyth, which wraps MythTV around Knoppix to ease building Linux-powered DVRs. Similarly, there is the XBMC Live distro which wraps Ubuntu around XBMC Media Center ease building Linux-powered HTPC (Home Theater PC). Others target the Apple Macintosh platform, including mkLinux, Yellow Dog Linux, and Black Lab Linux. Karoshi is a server system based on PCLinuxOS and aimed at educational users. SuperGamer is one of the few distributions focused solely on gaming. Trisquel is a Debian-based distribution that is composed entirely of Free Software and is endorsed by the Free Software Foundation. Scientific Linux is commonly used for scientific computer servers and workstations.

Interdistribution issues

The Free Standards Group is an organization formed by major software and hardware vendors that aims to improve interoperability between different distributions. Among their proposed standards are the Linux Standard Base, which defines a common ABI and packaging system for Linux, and the Filesystem Hierarchy Standard which recommends a standard filenaming chart, notably the basic directory names found on the root of the tree of any Linux filesystem. Those standards, however, see limited use, even among the distributions developed by members of the organization.

The diversity of Linux distributions means that not all software runs on all distributions, depending on what libraries and other system attributes are required. Packaged software and software repositories are usually specific to a particular distribution, though cross-installation is sometimes possible on closely related distributions.

Tools for choosing a distribution

There are tools available to help people select an appropriate distribution, such as several different versions of the Linux Distribution Chooser,[3][4] and the universal package search tool *whohas*.[5] There are easy ways to try out several Linux distributions before deciding on one: Multi Distro is a Live CD that contains nine space-saving distributions.[6] Tools are available to make such CDs and DVDs, among them Nautopia.[7]

Virtual machines such as VirtualBox and VMware Workstation permit booting of Live CD image files without actually burning a CD.

Details and interest rankings of Linux distributions are available on DistroWatch and a fairly comprehensive list of live CDs is available at livecdlist.com. Some websites such as OSDir.com and www.osvids.com offer screenshots and videos as a way to get a first impression of various distributions.

Workspot provides online Linux desktop demos using Virtual Network Computing (VNC).

Installation

There are many ways to install a Linux distribution. The most common method of installing Linux is by booting from a CD-ROM or DVD that contains the installation program and installable software. Such a CD can be burned from a downloaded ISO image, purchased alone for a low price, provided as a cover disk with a magazine, shipped for free by request, or obtained as part of a box set that may also include manuals and additional commercial software. New users tend to begin by partitioning a hard drive in order to keep their previously-installed operating system. The Linux distribution can then be installed on its own separate partition without affecting previously saved data.

Early Linux distributions were installed using sets of floppies but this has been abandoned by all major distributions. Nowadays most distributions offer CD and DVD sets with the vital packages on the first disc and less important packages on later ones. They usually also allow installation over a network after booting from either a set of floppies or a CD with only a small amount of data on it.[8]

Still another mode of installation is to install on a powerful computer to use as a servers and to use less powerful machines (perhaps without hard drives, with less memory and slower CPUs) as thin clients over the network. Clients can boot over the network from the server and display results and pass information to the server where all the applications run. The clients can be ordinary PCs with the addition of a network bootloader on a drive or network interface controller; hard disk space and processor power can be offloaded onto the client machine if desired. The cost savings achieved by using thin clients can be invested in greater computing power or storage on the server.

In a Live CD setup, the computer boots the entire operating system from CD without first installing it on the computer's hard disk. Some distributions have a Live CD *installer*, where the computer boots the operating system from the disk, and then proceeds to install it onto the computer's hard disk, providing a seamless transition from the OS running from the CD to the OS running from the hard disk.

Both servers and personal computers that come with Linux already installed are available from vendors including Hewlett-Packard, Dell and System76.

On embedded devices, Linux is typically held in the device's firmware and may or may not be consumer-accessible.

Anaconda, one of the more popular installers, is used by Red Hat Enterprise Linux, Fedora and other distributions to simplify the installation process.

Installation via an existing operating system

Some distributions let the user install Linux on top of their current system, such as WinLinux or coLinux. Linux is installed to the Windows hard disk partition, and can be started from inside Windows itself.

Virtual machines (such as VirtualBox or VMware) also make it possible for Linux to be run inside another OS. The VM software simulates a separate computer onto which the Linux system is installed. After installation, the virtual machine can be booted as if it were an independent computer.

Various tools are also available to perform full dual-boot installations from existing platforms without a CD, most notably:

- The Wubi installer, which allows Windows users to download and install Ubuntu or its derivatives into a FAT32 or a NTFS partition without an installation CD, allowing users to easily dual boot between either operating system on the same hard drive without losing data
- Win32-loader, which is in the process of being integrated in official Debian CDs/DVDs, and allows Windows users to install Debian without a CD, though it performs a network installation and thereby requires repartitioning[9]
- UNetbootin, which allows Windows and Linux users to perform similar no-CD network installations for a wide variety of Linux distributions and additionally provides live USB creation support

Proprietary software

Some specific proprietary software products are not available in any form for Linux. This includes many popular computer games, although in recent years some game manufacturers have begun making their software available for Linux. Emulation and API-translation projects like Wine and CrossOver make it possible to run non-Linux-based software on Linux systems, either by emulating a proprietary operating system or by translating proprietary API calls (e.g., calls to Microsoft's Win32 or DirectX APIs) into native Linux API calls. A virtual machine can also be used to run a proprietary OS (like Microsoft Windows) on top of Linux.

OEM contracts

Computer hardware is usually sold with an operating system other than Linux already installed by the original equipment manufacturer (OEM). In the case of IBM PC compatibles the OS is usually Microsoft Windows; in the case of Apple Macintosh computers it has always been a version of Apple's OS, currently Mac OS X; Sun Microsystems sells SPARC hardware with Solaris installed; video game consoles such as the Xbox, PlayStation, and Wii each have their own proprietary OS. This limits Linux's market share: consumers are unaware that an alternative exists, they must make a conscious effort to use a different operating system, and they must either perform the actual installation themselves, or depend on support from a friend, relative, or computer professional.

However, it is possible to buy hardware with Linux already installed. Lenovo, Hewlett-Packard, Dell, Affordy,[10] and System76 all sell general-purpose Linux laptops,[11] and custom-order PC manufacturers will also build Linux systems (but possibly with the Windows key on the keyboard). Fixstars Solutions (formerly Terra Soft) sells Macintosh computers and PlayStation 3 consoles with Yellow Dog Linux installed.

It is more common to find embedded devices sold with Linux as the default manufacturer-supported OS, including the Linksys NSLU2 NAS device, TiVo's line of personal video recorders, and Linux-based cellphones (including Android smartphones), PDAs, and portable music players.

The end user license agreement (EULA) for Apple gives the consumer the opportunity to reject the license and obtain a refund. The current Microsoft Windows license lets the manufacturer determine the refund policy.[12] With previous versions of Windows, it was possible to obtain a refund if the manufacturer failed to provide the refund by litigation in the small claims courts.[13] On 15 February 1999, a group of Linux users in Orange County, California held a "Windows Refund Day" protest in an attempt to pressure Microsoft into issuing them refunds.[14] In France, the Linuxfrench and AFUL organizations along with free software activist Roberto Di Cosmo started a "Windows Detax" movement,[15] which led to a 2006 petition against "racketiciels"[16] (translation: Racketware) and the DGCCRF branch of the French government filing several complaints against bundled software.

Screenshots of common distributions

Note: These images do not display a preview of what the distribution looks like, but what the desktop environment or window manager running on the distribution is.

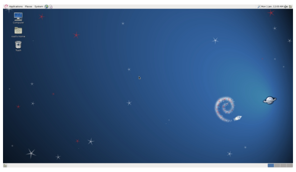
Debian 6.0 "Squeeze"


Fedora 16 "Verne"


Gentoo Linux 10.1

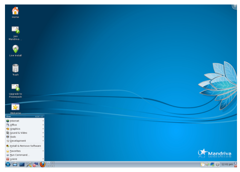
Mandriva Linux 2010.0


openSUSE 12.1


Slackware 13.1

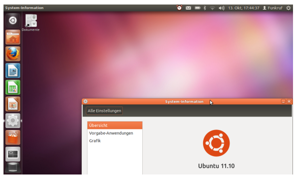
Ubuntu 11.10


Puppy Linux 5.2.5

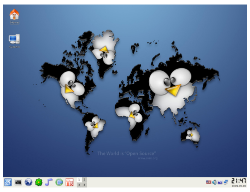
Slax 6.1.2

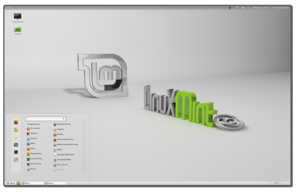
Linux Mint 12 "Lisa"

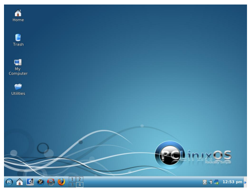
PCLinuxOS 2009.2

References

[1] http://futurist.se/gldt/
[2] "The Slackware Linux Project: Slackware Release Announcement" (http://www.slackware.com/announce/1.0.php). Slackware.com. 16 July 1993. . Retrieved 29 July 2011.
[3] "Zegenie Studios Linux Distribution Chooser" (http://www.zegeniestudios.net/ldc/). Zegeniestudios.net. 14 May 2011. . Retrieved 29 July 2011.
[4] "Distro Selector" (http://desktoplinuxathome.com/distro.html). Desktop Linux At Home. . Retrieved 29 July 2011.
[5] "Philipp's Homepage: whohas" (http://www.philippwesche.org/200811/whohas/intro.html). Philippwesche.org. 11 February 2010. . Retrieved 29 July 2011.
[6] "Multi Distro is Linux times 9 on a single CD-R" (http://www.linux.com/article.pl?sid=06/05/22/1353239). Linux.com. . Retrieved 29 July 2011.
[7] Linux.com :: Multiple live CDs in one DVD (http://applications.linux.com/article.pl?sid=06/03/16/1724219&tid=47)
[8] "Network install from a minimal CD" (http://www.debian.org/CD/netinst/). Debian. . Retrieved 29 July 2011.
[9] Debian Webmaster, webmaster@debian.org. "Debian - Details of package win32-loader in Lenny" (http://packages.debian.org/testing/win32-loader). Packages.debian.org. . Retrieved 29 July 2011.
[10] "Affordy - TITAN Computers" (http://shop.affordy.com/?lang=en). Shop.affordy.com. . Retrieved 29 July 2011.
[11] "Laptops/Notebooks with Linux Preinstalled" (http://mcelrath.org/laptops.html). Mcelrath.org. 10 June 2010. . Retrieved 29 July 2011.
[12] "Microsoft Software License Terms: Windows 7 Professional" (http://download.microsoft.com/Documents/UseTerms/Windows 7_Professional_English_b7a7153f-1a6c-498c-9350-c86926bb1aa9.pdf). .
[13] "Getting a Windows Refund in California Small Claims Court" (http://www.linuxjournal.com/article/7040). Linuxjournal.com. . Retrieved 29 July 2011.

[14] "Windows Refund Day" (http://marc.merlins.org/linux/refundday/). Marc.merlins.org. 15 February 1999. . Retrieved 29 July 2011.
[15] Detaxe.org (http://www.detaxe.org/) (French) Say no to bundled software - Say yes to informed consumers
[16] "(fr) Petition against software racketeers" (http://www.racketiciel.info/). Racketiciel.info. . Retrieved 29 July 2011.

Debian

Debian

debian

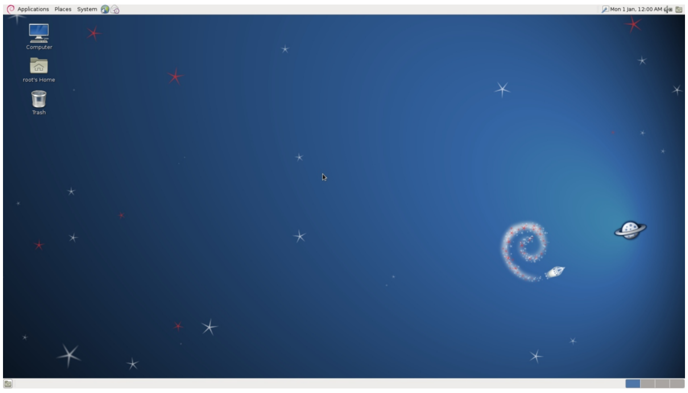

Debian GNU/Linux 6.0 ("Squeeze") with GNOME

Company / developer	Debian Project
OS family	Unix-like
Working state	Current
Source model	Free and open source software
Initial release	16 August 1993
Latest stable release	(12 May 2012)[1] [2] [±]
Latest unstable release	(13 May 2012)[3] [4] [±]
Available language(s)	over 65[5]
Update method	APT (several front-ends available)
Package manager	dpkg
Supported platforms	i386, AMD64, PowerPC, SPARC, ARM, MIPS, S390, IA-64
Kernel type	Monolithic (Linux, FreeBSD), Micro (Hurd)
Userland	GNU
Default user interface	GNOME
License	Free software, mainly the GNU GPL, and other licenses[6]
Official website	[www.debian.org www.debian.org]

Debian (/ˈdɛbiən/) is a computer operating system composed of software packages released as free and open source software primarily under the GNU General Public License along with other free software licenses.[7] **Debian GNU/Linux**, which includes the GNU OS tools and Linux kernel,[8] is a popular and influential Linux distribution.[9][10] It is distributed with access to repositories containing thousands of software packages ready for installation and use. Debian is known for relatively strict adherence to the philosophies of Unix and free software[11] as well as using collaborative software development and testing processes.[12] Debian can be used on a variety of hardware, from laptops and desktops to NAS devices, phones, and servers. It focuses on stability and security and is used as a base for many other distributions.

Organization

The Debian Project is governed by the Debian Constitution and the Social Contract which set out the governance structure of the project and explicitly states that the goal of the project is the development of a free operating system.[13][14] Debian is developed by over three thousand volunteers[15] from around the world and supported by donations through several nonprofit organizations around the world. Most important of these is Software in the Public Interest,[16] the owner of the Debian trademark and umbrella organization for various other community free software projects.[17]

Thus, the Debian Project is an independent decentralized organization; it is not backed by a company like Linux distributions such as Ubuntu, openSUSE, Fedora, and Mandriva. The cost of developing all the packages included in Debian 5.0 *lenny* (323 million lines of code), using the COCOMO model, has been estimated to be about US$ 8 billion.[18] Ohloh estimates that the codebase (54 million lines of code), using the COCOMO model, would cost about US$ 1 billion to develop.[19]

Features

Debian is known for an abundance of options. The current stable release includes over 29,000 software packages for 11 different computer architectures[20] using the Linux kernel. There are also packages for architectures using the FreeBSD kernel (kfreebsd-i386 and kfreebsd-amd64). These architectures range from the Intel/AMD 32-bit/64-bit architectures commonly found in personal computers to the ARM architecture commonly found in embedded systems and the IBM eServer zSeries mainframes.[21] The Debian *standard install* makes use of the GNOME desktop environment. It includes popular programs such as LibreOffice,[22] Iceweasel (a rebranding of Firefox), Evolution mail, CD/DVD writing programs, music and video players, image viewers and editors, and PDF viewers. There are pre-built CD images for KDE Software Compilation, Xfce and LXDE also.[23] The remaining discs, which span five DVDs or over thirty CDs, contain all packages currently available and are not necessary for a *standard install*. Another install method is via a net install CD, which is much smaller than a normal install CD/DVD. It contains only the bare essentials needed to start the installer and downloads the packages selected during installation via APT.[24] These CD/DVD images can be freely obtained by web download, BitTorrent, jigdo or from online retailers.[25]

Package management

Debian was one of the earlier Linux distributions to compose itself from packages,[26] and robust package management is perhaps Debian's most prominent feature. The APT package management system, repositories with large numbers of packages, and strict policies regarding packages, promote high quality releases,[27] easy upgrades between releases, and automated installation and removal of packages.

dpkg, installing local .deb packages

dpkg is the base of the Debian package management system. dpkg is a command-line utility to install, remove, and provide information about local .deb packages.[28] GDebi is a tool that expands the functionality of dpkg by not only installing local .deb packages but also fetching and installing dependencies from online repositories. GDebi can be used both in a graphical interface and by command-line.

APT, installing packages from online repositories

APT expands the functionality of dpkg by searching, fetching and installing packages from online repositories along with their dependencies, either from binary files or by compiling source code. It can also upgrade packages and upgrade the whole distribution to a new release.

Graphical front-ends

- Software Center, originally developed for Ubuntu, is a GNOME program to discover and install applications.
- Synaptic is a GTK+ front-end for APT.[29]
- Apper is KDE's front-end for Packagekit.
- Adept is a discontinued KDE front-end for APT.
- *Gnome Application Installer* was a GNOME program to discover and install applications, now removed in favor of Software Center

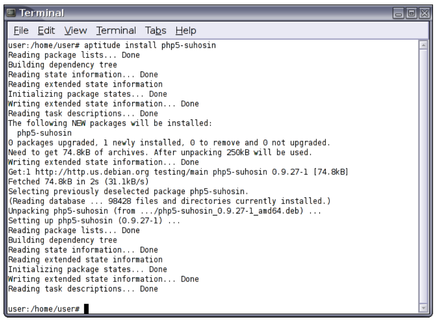

Package installed with aptitude.

Command-line front-ends

- apt-get is the most basic package management tool and the preferred APT front-end for non-interactive package management. It offers a command line user interface and it is most suitable for upgrades between major releases.[30][31]
- aptitude is the most versatile package management tool and the preferred APT front-end for interactive package management. It offers a full screen interactive text user interface, an enhanced package resolver and an enhanced search function. It is most suitable for daily package management.[30][31]
- dselect is an old front-end for APT, largely superseded by other front-ends.[32]
- wajig is a simplified front end, providing the functionality of apt-get, dpkg, dpkg-deb, apt-cache, and other tools.

Debian Live

A Debian Live system is a version of Debian that can be booted directly from removable media (CDs, DVDs, USB keys) or via netboot without having to install it on the hard drive.[33] This allows the user to try out Debian before installing it or use it as a boot-disk. There are prebuilt Debian Live images for *rescue*, *standard*, GNOME, KDE Plasma Workspaces, Xfce and LXDE for several architectures. A hard disk installation can be achieved using the Debian Installer included in the live image. Most of the live ISO images for the current Squeeze release no longer fit on a 700MB CD. Customized CD images can be built using live-build.[34] Live-build can not only generate CD Images, but also bootable DVDs, images for USB thumb drives, or netboot images. Live-magic is a GUI for live-build. Ging is a Debian GNU/kFreeBSD Live CD.[35]

Desktop environments

Debian offers *stable* and *testing* CD images specifically built for GNOME (the default), KDE Plasma Workspaces, Xfce and LXDE.[36] Less common window managers such as Enlightenment, Openbox, Fluxbox, GNUstep, IceWM, Window Maker and others can also be installed.

Embedded systems

Recent releases of Debian support an increasing number of ARM-based NAS devices. The cheap NSLU2 was supported by Debian 4.0 and 5.0[37] and can be upgraded to Debian 6.0 although there are problems with a 6.0 clean install.[38] Debian 5.0 added support for the Buffalo Kurobox Pro,[39] and Debian 6.0 for the SheevaPlug.[40]

Other NAS devices supported by Debian, but perhaps not so widely used by home users, include GLAN Tank[41] and Thecus N2100 as of Debian 4.0,[42] QNAP Turbo Station (TS-109, TS-209, TS-409) and HP mv2120 as of Debian 5.0,[39] and QNAP Turbo NAS TS-11x, TS-21x and TS-41x, OpenRD, Lanner EM7210 and Intel SS4000-e as of Debian 6.0.[40]

History

1993–1998

Debian was first announced on 16 August 1993 by Ian Murdock,[43] who initially called the system "the Debian Linux Release".[44] The word "Debian" was formed as a combination of the first name of his then-girlfriend Debra Lynn and his own first name.[45] Prior to Debian's release, the Softlanding Linux System (SLS) had been the first Linux distribution compiled from various software packages, and was a popular basis for other distributions in 1993-1994.[26] The perceived poor maintenance and prevalence of bugs in SLS[46] motivated Murdock to launch a new distribution.

In 1993 Murdock also released the *Debian Manifesto*,[47] outlining his view for the new operating system. In it he called for the creation of a distribution to be maintained in an open manner, in the spirit of Linux and GNU.

The Debian Project grew slowly at first and released the first 0.9x versions in 1994 and 1995. During this time it was sponsored by the Free Software Foundation's GNU Project.[48] The first ports to other, non-i386 architectures began in 1995, and the first 1.x version of Debian was released in 1996.

In 1996, Bruce Perens replaced Ian Murdock as the project leader. In the same year, Perens read a discussion between fellow developer Ean Schuessler and Donnie Barnes of Red Hat, suggesting that Red Hat had no stated social contract with its users which guaranteed the future freedom of the system's contents. Perens decided to create a social contract for Debian, created a first draft, and edited suggestions from a month-long discussion on the Debian mailing lists into the Debian Social Contract and the Debian Free Software Guidelines, defining fundamental commitments for the development of the distribution. He also initiated the creation of the legal umbrella organization, Software in the Public Interest.[16] Perens developed the project from 40 to 200 developers. He broke apart the "base system", the core packages of Debian, which had been maintained by Murdock alone, and distributed them to many maintainers. He led the conversion of the project from a.out to ELF. He created the BusyBox program to make it possible to run a Debian installer on a single floppy, and wrote a new installer. Perens was also responsible for many policy and design elements of Debian that persist to this day. Perens left the project in 1998.

1999–2004

The Project elected new leaders and made two more 2.x releases, each including more ports and packages. The Advanced Packaging Tool was deployed during this time and the first port to a non-Linux kernel, Debian GNU/Hurd, was started. The first Linux distributions based on Debian, namely Libranet, Corel Linux and Stormix's Storm Linux, were started in 1999.[16] The 2.2 release in 2000 was dedicated to Joel Klecker, a developer who died of Duchenne muscular dystrophy.[49]

In late 2000, the project made major changes to archive and release management, reorganizing software archive processes with new "package pools" and creating a testing distribution as an ongoing, relatively stable staging area for the next release. In the same year, developers began holding an annual conference called *DebConf* with talks and workshops for developers and technical users.[16]

In July 2002, the Project released version 3.0, codenamed *woody*, (each Debian release is named after a character in the movie Toy Story), a stable release which would see relatively few updates until the following release.[16]

2005–present

The 3.1 *sarge* release was made in June 2005. There were many major changes in this release, mostly due to the long time it took to freeze and release the distribution. Not only did this release update over 73% of the software shipped in the prior version, but it also included much more software than prior releases, almost doubling in size with more than 9,000 new packages.[50] A new installer replaced the aging boot-floppies installer with a modular design. This allowed advanced installations (with RAID, XFS and LVM support) including hardware detection, making installations easier for novice users. The installation system also boasted full internationalization support as the software was translated into almost forty languages. An installation manual and comprehensive release notes were released in ten and fifteen different languages respectively. This release included the efforts of the Debian-Edu/Skolelinux, Debian-Med and Debian-Accessibility sub-projects which raised the number of packages that were educational, had a medical affiliation, and ones made for people with disabilities.[16]

In 2006, as a result of a much-publicized dispute, Mozilla software was rebranded in Debian, with Firefox becoming Iceweasel, Thunderbird becoming Icedove, along with other Mozilla programs. The Mozilla Corporation stated that Debian may not use the Firefox trademark if it distributes Firefox with modifications which have not been approved by the Mozilla Corporation. Two prominent reasons that Debian modifies the Firefox software are to change the artwork, and to provide security patches. Debian Free Software Guidelines consider Mozilla's artwork *non-free*. Debian provides long term support for older versions of Firefox in the *stable* release, where Mozilla prefers that old versions are not supported. These software programs developed largely by the Mozilla Corporation were rebranded despite having only minor differences in the source code.[51]

Debian 4.0 (*etch*) was released April 8, 2007 for the same number of architectures as in *sarge*. It included the AMD64 port but dropped support for m68k. The m68k port was, however, still available in the unstable distribution. There were approximately 18,200 binary packages maintained by more than 1,030 Debian developers.[16]

Debian 5.0 (*lenny*) was released February 14, 2009 after 22 months of development. It includes more than 25,000 software packages. Support was added for Marvell's Orion platform and for netbooks such as the Asus Eee PC, but support was dropped for 32-bit SPARC machines.[52] The release was dedicated to Thiemo Seufer, an active developer and member of the community who died in a car accident on December 26, 2008.[53]

On September 5, 2010, Debian officially acquired the backports service, which provides more recent versions of some software for the stable release of Debian.[54]

Debian 6.0 (*squeeze*) was released February 6, 2011 after 24 months of development. For the first time, Debian GNU/kFreeBSD was introduced with this version as a technology preview.[55]

Project organization

The Debian Project is a volunteer organization with three foundational documents:

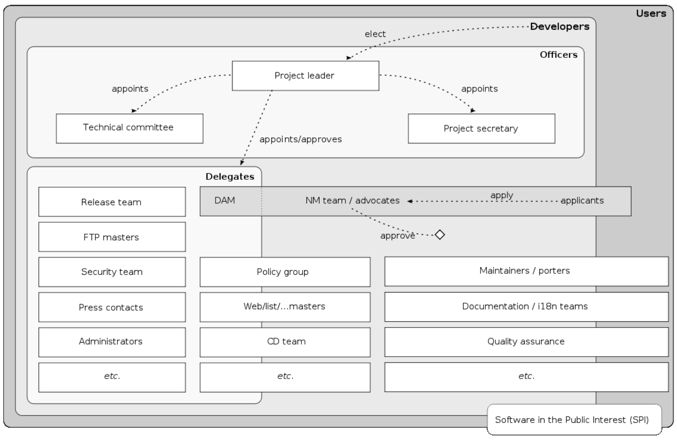

Diagram of the organizational structure of the project

- The Debian Social Contract defines a set of basic principles by which the project and its developers conduct affairs.[14]
- The Debian Free Software Guidelines define the criteria for "free software" and thus what software is permissible in the distribution, as referenced in the Social Contract. These guidelines have also been adopted as the basis of the Open Source Definition. Although it can be considered a separate document for all practical purposes, it formally is part of the Social Contract.[14]
- **The Debian Constitution** describes the organizational structure for formal decision-making within the Project, and enumerates the powers and responsibilities of the Debian Project Leader, the Debian Project Secretary, and the Debian Developers generally.[13]

Currently, the project includes more than a thousand developers. Each of them sustains some niche in the project, be it package maintenance, software documentation, maintaining the project infrastructure, quality assurance, or release coordination. Package maintainers have jurisdiction over their own packages, although packages are increasingly co-maintained. Other tasks are usually handled by the domain of smaller, more collaborative groups of developers.

The project maintains official mailing lists and conferences for communication and coordination between developers.[56] For issues with single packages or domains, a public bug tracking system is used by developers and end-users. Informally, Internet Relay Chat channels (primarily on the OFTC and freenode networks) are used for communication among developers and users also.

Together, the Developers may make binding general decisions by way of a General Resolution or election. All voting is conducted by Cloneproof Schwartz Sequential Dropping, a Condorcet method of voting. A Project Leader is elected once per year by a vote of the Developers; in April 2010, Stefano Zacchiroli was voted into this position, succeeding Steve McIntyre. The Debian Project Leader has several special powers, but this power is far from absolute and is rarely used. Under a General Resolution, the Developers may, among other things, recall the leader, reverse a decision by him or his delegates, and amend the constitution and other foundational documents.

The Leader sometimes delegates authority to other developers in order for them to perform specialized tasks. Generally this means that a leader delegates someone to start a new group for a new task, and gradually a team gets formed that carries on doing the work and regularly expands or reduces their ranks as they think is best and as the circumstances allow.

A role in Debian with a similar importance to the Project Leader's is that of a Release Manager. Release Managers set goals for the next release, supervise the processes, and make the final decision as to when to release.[57][58]

Project leaders

The Debian Project Leader (DPL) is the public face of the project and defines the current direction of the project.[59] The project has had the following leaders:[60]

- Ian Murdock (August 1993 – March 1996), founder of the Debian Project
- Bruce Perens (April 1996 – December 1997)
- Ian Jackson (January 1998 – December 1998)
- Wichert Akkerman (January 1999 – March 2001)
- Ben Collins (April 2001 – April 2002)
- Bdale Garbee (April 2002 – April 2003)
- Martin Michlmayr (March 2003 – March 2005)
- Branden Robinson (April 2005 – April 2006)
- Anthony Towns (April 2006 – April 2007)
- Sam Hocevar (April 2007 – April 2008)
- Steve McIntyre (April 2008 – April 2010)
- Stefano Zacchiroli (April 2010 – present)

A supplemental position, *Debian Second in Charge* (2IC), was created by Anthony Towns. Steve McIntyre held the position between April 2006 and April 2007. From April 2009 to April 2010 this position was held by Luk Claes. Stefano Zacchiroli — the current project leader — abandoned this unofficial position when elected in April 2010.[61]

Release managers

- Brian C. White (1997–1999)
- Richard Braakman (1999–2000)
- Anthony Towns (2000–2004)
- Steve Langasek, Andreas Barth and Colin Watson (2004–2007)
- Andreas Barth and Luk Claes (2007–2008)
- Luk Claes and Marc Brockschmidt (2008–2009)
- Luk Claes and Adeodato Simó (2009–2010)
- Adam D. Barratt and Neil McGovern (2010–present)[62]

Note that this list includes the active release managers; it does not include the release assistants (first introduced in 2003) and the retiring managers ("release wizards").[57]

Developer recruitment, motivation, and resignation

The Debian project has a steady influx of applicants wishing to become developers. These applicants must undergo an elaborate vetting process which establishes their identity, motivation, understanding of the project's goals (embodied in the Social Contract), and technical competence.[63]

Debian Developers join the Project for a number of reasons; some that have been cited in the past include:[64]

- A desire to contribute back to the free-software community (practically all applicants are users of free software)
- A desire to see some specific software task accomplished (some view the Debian user community as a valuable testing or proving ground for new software)
- A desire to make, or keep, free software competitive with proprietary alternatives
- A desire to work closely with people who share some of their aptitudes, interests, and goals (there is a very strong sense of community within the Debian project which some applicants do not experience in their paid jobs)
- A simple enjoyment of the iterative process of software development and maintenance

Debian Developers may resign their positions at any time by orphaning the packages they were responsible for and sending a notice to the developers and the keyring maintainer (so that their upload authorization can be revoked).

Development procedures

Software packages in development are either uploaded to the project distribution named *unstable* (also known as *sid*), or to the *experimental* repository. Software packages uploaded to *unstable* are normally versions stable enough to be released by the original upstream developer, but with the added Debian-specific packaging and other modifications introduced by Debian developers. These additions may be new and untested. Software not ready yet for the *unstable* distribution is typically placed in the *experimental* repository.[65]

After a version of a software package has remained in *unstable* for a certain length of time (depending on the urgency of the software's changes), that package is automatically migrated to the *testing* distribution. The package's migration to testing occurs only if no serious (*release-critical*) bugs in the package are reported and if other software needed for package functionality qualifies for inclusion in *testing*.[65]

Since updates to Debian software packages between official releases do not contain new features, some choose to use the *testing* and *unstable* distributions for their newer packages. However, these distributions are less tested than *stable*, and *unstable* does not receive timely security updates. In particular, incautious upgrades to working *unstable* packages can sometimes seriously break software functionality.[66] Since September 9, 2005[67] the *testing* distribution's security updates have been provided by the *testing* security team.[68]

After the packages in *testing* have matured and the goals for the next release are met, the *testing* distribution becomes the next stable release. The timing of the release is decided by the Release Managers, and in the past the exact date was rarely announced earlier than a couple of weeks beforehand.[69]

Package maintenance

Each Debian software package has a maintainer who keeps track of releases by the "upstream" authors of the software and ensures that the package is compliant with Debian Policy, coheres with the rest of the distribution, and meets the standards of quality of Debian. In relations with users and other developers, the maintainer uses the bug tracking system to follow up on bug reports and fix bugs. Typically, there is only one maintainer for a single package, but, increasingly, small teams of developers "co-maintain" larger and more complex packages and groups of packages.[70]

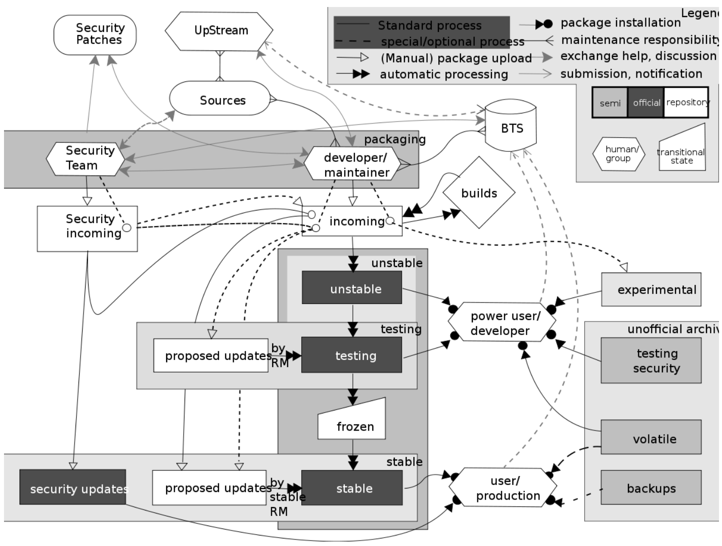

Flowchart of the life cycle of a Debian package

Periodically, a package maintainer makes a release of a package by uploading it to the "incoming" directory of the Debian package archive (or an "upload queue" which periodically batch-transmits packages to the incoming directory). Package uploads are automatically processed to ensure that they are well-formed (all the requisite files are in place) and that the package is digitally signed by a Debian developer using OpenPGP-compatible software. All Debian developers have individual cryptographic key pairs.[71] Packages are signed to be able to reject uploads from hostile outsiders to the project, and to permit accountability in the event that a package contains a serious bug, a violation of policy, or malicious code.

If the package in incoming is found to be validly signed and well-formed, it is installed into the archive into an area called the "pool" and distributed every day to hundreds of mirrors worldwide. Initially, all package uploads accepted into the archive are only available in the "unstable" suite of packages, which contains the most up-to-date version of each package.

However, new code is also untried code, and those packages are only distributed with clear disclaimers. For packages to become candidates for the next "stable" release of the Debian distribution, they first need to be included in the "testing" suite. For a package to be included in *testing*:[72][73]

- It must have been in *unstable* for the appropriate length of time (the exact duration depends on the "urgency" of the upload)
- It must not have a greater number of "release-critical" bugs filed against it than the current version in testing. Release-critical bugs are those bugs which are considered serious enough that they make the package unsuitable for release.
- It must be compiled for all release architectures the package claims to support (e.g.: the i386-specific package gmod can be included in "testing")
- All of its dependencies must either be satisfiable by packages already in *testing*, or be satisfiable by the group of packages which are going to be installed at the same time.
- The operation of installing the package into *testing* must not break any packages currently in *testing*.

Thus, a release-critical bug in a package on which many packages depend, such as a shared library, may prevent many packages from entering the *testing* area, because that library is considered deficient.

Periodically, the Release Manager publishes guidelines to the developers in order to ready the release, and in accordance with them eventually decides to make a release. This occurs when all important software is reasonably up-to-date in the release-candidate suite for all architectures for which a release is planned, and when any other goals set by the Release Manager have been met. At that time, all packages in the release-candidate suite ("testing") become part of the released suite ("stable").

It is possible for a package, particularly an old, stable, and seldom-updated one, to belong to more than one suite at the same time. The suites are simply collections of pointers into the package "pool" mentioned above.

Security information and policy

The Debian Project, being free software, handles security policy through public disclosure rather than through security through obscurity. Many advisories are coordinated with other free software vendors (Debian is a member of vendor-sec) and are published the same day a vulnerability is made public. Debian has a security audit team that reviews the archive looking for new or unfixed security bugs. Debian also participates in security standardization efforts: the Debian security advisories are compatible with the Common Vulnerabilities and Exposures (CVE) dictionary, and Debian is represented in the Board of the Open Vulnerability and Assessment Language (OVAL) project.[74]

The Debian Project offers extensive documentation and tools to harden a Debian installation both manually and automatically.[75] SELinux (Security-Enhanced Linux) packages are installed by default though not enabled.[32] Debian provides an optional hardening wrapper but does not compile their packages by default using gcc features such as PIE and buffer overflow protection to harden their software, unlike Ubuntu, Fedora and Hardened Gentoo among others.[76] These extra features greatly increase security at a performance cost of 1% in 32-bit and 0.01% in 64-bit.[77]

It is a release goal for Debian 7.0 (*wheezy*) "to update as many packages as possible to use security hardening build flags via dpkg-buildflags. These flags enable various protections against security issues such as stack smashing, predictable locations of values in memory, etc."[78]

Releases

Debian Installer

As of February 2011, the latest stable release is version 6.0, code name *squeeze*.[79] When a new version is released, the prior stable version becomes *oldstable*. As of 2011, this is version 5.0, code name *lenny*.

In addition, a stable release gets minor updates (called *point releases*). The numbering scheme for the point releases up to Debian 4.0 was to include the letter *r* (for *release*) after the main version number (e.g. 4.0) and then the number of the point release; for example, the latest point release of version 4.0 (*etch*) as of 8 December 2010 is 4.0r9.[80] From Debian 5.0 (*lenny*), the numbering scheme of point releases has been changed and conforms to the GNU version numbering standard; so, for example, the first point release of Debian 5.0 was 5.0.1 (instead of 5.0r1).[81]

The Debian security team releases security updates for the latest stable major release, and for the prior stable release for one year.[66] Version 4.0 *etch* was released on 8 April 2007, and the security team supported version 3.1 *Sarge* until 21 March 2008. For most uses it is strongly recommended to run a system which receives security updates. The *testing* distribution also receives security updates, but not in as timely a manner as stable.[82]

For Debian 6.0 (*squeeze*) a new policy of time-based development freezes on a two-year cycle was announced. Time-based freezes are intended to allow the Debian Project to blend the predictability of time based releases with its policy of feature based releases. The new freeze policy aims to provide better predictability of releases for users of the Debian distribution, and to allow Debian developers to do better long-term planning. Debian developers expect that a two-year release cycle will give more time for disruptive changes, reducing inconveniences caused for users. Having predictable freezes was expected to reduce overall freeze time. The squeeze cycle was intended to be especially short to "get into the new cycle".[83] However this short freeze cycle for squeeze was abandoned.[84]

The code names of Debian releases are names of characters from the film *Toy Story*. The unstable, development distribution is permanently nicknamed *sid*, after the emotionally unstable next-door neighbor boy who regularly destroyed toys.[65] The release after *squeeze* will be named *wheezy*, after the rubber toy penguin in *Toy Story 2*.[85]The release after *wheezy* will be named *jessie*, after the cowgirl in *Toy Story 2* and *Toy Story 3*.[86]

Release history

Debian has made eleven major stable releases:[16]

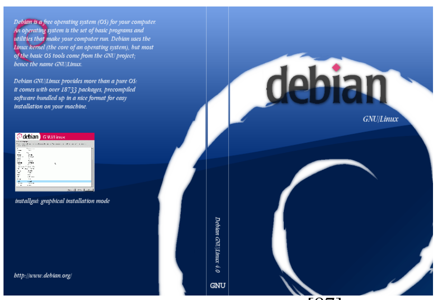

A Debian 4.0 Box Cover[87]

Legend
Release no longer supported
Release still supported
Future release

TBA stands for *to be announced*.

Version	Code name	Release date	Ports	Packages	Supported until	Notes
1.1	*buzz*	1996-06-17	1	474	1996-09 [88]	`dpkg`, ELF transition, Linux 2.0[16]
1.2	*rex*	1996-12-12	1	848	1996	-
1.3	*bo*	1997-06-05	1	974	1997	-
2.0	*hamm*	1998-07-24	2	≈ 1,500	1998	glibc transition, new architecture: `m68k`[16]
2.1	*slink*	1999-03-09	4	≈ 2,250	2000-12	`APT`, new architectures: `alpha`, `sparc`[16]
2.2	*potato*	2000-08-15	6	≈ 3,900	2003-04	New architectures: `arm`, `powerpc`[89]
3.0	*woody*	2002-07-19	11	≈ 8,500	2006-08	New architectures: `hppa`, `ia64`, `mips`, `mipsel`, `s390`[16]
3.1	*sarge*	2005-06-06	11	≈ 15,400	2008-04[66]	Modular installer, semi-official `amd64` support.
4.0	*etch*	2007-04-08	11	≈ 18,000	2010-02-15[90]	New architecture: `amd64`, dropped architecture: `m68k`.[91] Graphical installer, udev transition, modular X.Org transition. Final update 4.0r9 was released 2010-05-22.[92]
5.0[93]	*lenny*[94]	2009-02-14[69]	12	≈ 23,000[52]	2012-02-06[95]	New architecture/binary ABI: `armel`.[96] SPARC 32-bit hardware support dropped.[97] Full Eee PC support.[98] Final update 5.0.10 was released 2012-03-10.[99]
6.0[100]	*squeeze*[101]	2011-02-06[55]	9+2[A]	≈ 29,000[55]	TBA	New architectures/kernels: `kfreebsd-i386`, `kfreebsd-amd64`, dropped architectures: `alpha`, `hppa`, OABI[102] `arm`.[55] eglibc instead of glibc.[103] Dependency-based boot sequence, which allows for parallel init script processing.[104] Remove old libraries such as GTK 1.[105] Default Linux kernel purged of non-free firmware.[15]
7.0[85]	*wheezy*[85]	TBA (Freeze: June 30th, 2012)[106]	TBA	TBA	TBA	New architectures: `armhf`, `s390x`.[107] Remove old libraries such as Qt3.[108] Introduce multiarch support.[109]
8.0[86]	*jessie*[86]	TBA	TBA	TBA	TBA	TBA

A 9 architectures with Linux kernel + 2 architectures with kernel of FreeBSD[55]

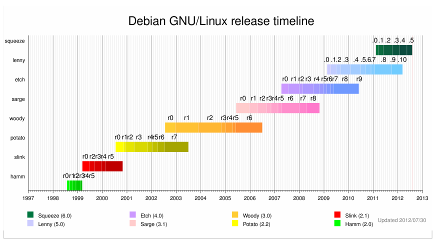

Due to an incident involving a CD vendor who made an unofficial and broken release labeled *1.0*, an official 1.0 release was never made.[16]

Kernels

Debian release	Date	Debian kernel	Latest kernel version just before the Debian release
1.1 *buzz*	1996-06-17	2.0 on 9 June 1996[110]	
1.2 *rex*	1996-12-12	2.0.27 on 1 December 1996[111]	
1.3 *bo*	1997-06-05	2.0.29 2.0.30 [112] for 1.3.1 also 2.0.33[112]	2.0.30 on 8 April 1997 [111] 2.1.42 on 29 May 1997[113]
2.0 *hamm*	1998-07-24	2.0.33 2.0.34[114]	2.0.35 on 13 July 1998[111] 2.1.110 on 21 July 1998[113]
2.1 *slink*	1999-03-09	2.0.35-3 2.0.36-3 2.1.125-1 2.2.1-1[115]	2.2.3 on 9 March 1999[116]
2.2 *potato*	2000-08-15	2.2.16[89]	2.2.16 on 7 June 2000[116] 2.3.99-pre9 on 23 May 2000[117]
3.0 *woody*	2002-07-19	2.2.20 2.4.18[118]	2.2.21 on 20 May 2002[116] 2.4.18 on 25 February 2002[119] 2.5.26 on 16 July 2002[120]
3.1 *sarge*	2005-06-06	2.4.30 on 4 April 2005 [119] 2.6.11.11 on 27 May 2005[121]	
4.0 *etch*	2007-04-08	2.6.18[91]	2.6.20.6 on 6 April 2007[121]
5.0 *lenny*	2009-02-14	2.6.26[52]	2.6.28.5 on 12 February 2009[121]
6.0 *squeeze*	2011-02-06	Linux 2.6.32[122] kFreeBSD 8.1	2.6.37 on 5 January 2011[121] 8.1 on 19 July 2010

Distributions

The Debian Project offers three distributions, each with different characteristics. The distributions include packages which comply with the Debian Free Software Guidelines (DFSG), which are included inside the *main* repositories.[123]

- **stable**, *currently* aliased ***squeeze***, is the current release that has stable and well-tested software. *Stable* is made by freezing *testing* for a few months where bugs are fixed to make the distribution as stable as possible; then the resulting system is released as *stable*. It is updated only if major security or usability fixes are incorporated. After Debian 6.0, new releases will be made every two years.[83] *Stable's* CDs and DVDs can be found in the Debian web site.[123]
- **testing**, *currently* aliased ***wheezy***, is what the next major release will be and is currently being tested. The packages included in this distribution have had some testing in *unstable* but they may not be completely fit for release yet. It contains more modern packages than *stable* but older than *unstable*. This distribution is updated continually until it enters the "frozen" state. Security updates for *testing* distribution are provided by Debian testing security team. *Testing's* CDs and DVDs can be found on the Debian web site.[123]
- **unstable**, *permanently* aliased ***sid***, repository contains packages currently under development; it is updated continually. This repository is designed for Debian developers who participate in a project and need the latest libraries available, or for those who like to "live on the edge", so it will not be as stable as the other distributions.

There are no official CDs/DVDs because it is rapidly changing and the project does not support it, although CD and DVD images of sid are built quarterly by aptosid. Additionally, the other two distributions can be upgraded to *unstable*.[123]

Additional repositories

The Debian Free Software Guidelines (DFSG) adhere to a relatively strict interpretation of free and open source software (FOSS), although it is still not endorsed by the Free Software Foundation as it includes and supports a proprietary repository and documentation that recommends non-free software.[124][125] In accordance with its guidelines, a relatively small number of packages are excluded from the distributions' *main* repositories and included inside the *non-free* and *contrib* repositories. These two repositories are not officially part of Debian GNU/Linux.

- **non-free**: repositories include packages which do not comply with the DFSG (this does usually not include legally questionable packages, like libdvdcss).[123]
- **contrib**: repositories include packages which do comply with the DFSG, but may fail other requirements. For instance, they may depend on packages which are in *non-free* or requires such for building them.[123]

These are other repositories available in Debian:

- **experimental**: is not actually a full (self-contained) development distribution, it is meant to be a temporary staging area for highly experimental software. Dependencies missing are most likely found in *unstable*. Debian warns that these packages are likely unstable or buggy and are to be used at the user's own risk.[123]
- **backports**: This repository provides more recent versions than *stable* for some software. It is mainly intended for users of *stable* who need a newer version of a particular package.
- **oldstable**, presently aliased ***lenny***, is the prior *stable* release. It is supported until 1 year after a new *stable* is released. Debian recommends to update to the new *stable* once it has been released.
- **snapshot**: The snapshot repositories provide older versions of other repositories. They may be used to install a specific older version of some software.

Third-party repositories

These repositories are not part of the Debian Project, they are maintained by third party organizations. They contain packages that are either more modern than the ones found in *stable* or include packages that are not included in the Debian Project for a variety of reasons such as: e.g. alleged possible patent infringement, binary-only/no sources, or special too restrictive licenses. Their use requires precise configuration of the priority of the repositories to be merged; otherwise these packages may not integrate correctly into the system, and may cause problems upgrading or conflicts between packages from different sources. The Debian Project discourages the use of these repositories as they are not part of the project.

Repository access

Debian provides repositories of more than 29,000 installable packages. Any of the repositories can be added or modified by directly editing the files in /etc/apt/sources.list or modifying the settings in APT front-ends.[126] This is an example of the contents of this file:

```
deb http://ftp.us.debian.org/debian squeeze main contrib non-free
deb-src http://ftp.us.debian.org/debian squeeze main contrib non-free

deb http://security.debian.org/ squeeze/updates main
deb-src http://security.debian.org/ squeeze/updates main

deb http://ftp.us.debian.org/debian squeeze-updates main contrib non-free
```

```
deb-src http://ftp.us.debian.org/debian squeeze-updates main contrib non-free

# backports
deb http://backports.debian.org/debian-backports squeeze-backports main
deb-src http://backports.debian.org/debian-backports squeeze-backports main
```

Hardware support

Hardware requirements

Debian has no hardware requirements beyond those of the Linux kernel and the GNU tool-sets (gcc, coreutils, bash, etc.). Therefore, any architecture or platform to which these packages have been ported, and for which a Debian port exists, can run Debian.[127]

Linux, and therefore Debian, supports the use of multiple processors in a system (symmetric multiprocessing). This does not inhibit support for single-processor systems.[127]

Debian's recommended system requirements differ depending on the level of installation, which corresponds to increased numbers of installed components:[128]

Install desktop	RAM minimum[128]	RAM recommended[128]	Hard drive space used[128]
No	64 MB	256 MB	1 GB
Yes	128 MB	512 MB	5 GB

A 1 GHz processor is the minimum recommended for desktop systems.[128]

The real minimum memory requirements are much less than the numbers listed in this table. Depending on the architecture, it is possible to install Debian with as little as 20 MB RAM for s390 or 48 MB RAM for i386 and AMD64. Similarly, disk space requirements, which depend on the packages to be installed, can also be reduced.[128] Emdebian [129] (embedded Debian) improves installation to devices with minimal disk space, partially by removing documentation and installing only needed translations. In its Grip and Baked form it is binary compatible.

It is possible to run graphical user interfaces on older or low-end systems, but the installation of window managers instead of desktop environments is recommended, as desktop environments are more resource-intensive.[128] For example, the LXDE desktop environment was released with *lenny* and has much lower processor and memory usage compared with GNOME or KDE Plasma Desktop.[32]

Depending on the nature of the server, RAM and disk space requirements can vary widely.[128]

Ports

Stable ports

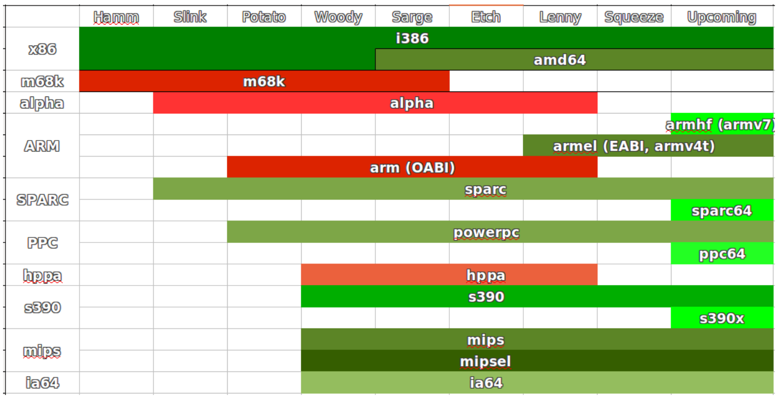

Debian GNU/Linux ports timeline.

As of the current stable release, the official ports are:[130]

- i386: x86 architecture designed for Intel/AMD 32-bit PCs. Also compatible but not recommended on Intel/AMD 64-bit single/multi core PCs[131]
- amd64: x86-64 architecture designed for Intel/AMD 64-bit single/multi core PCs
- armel: little-endian ARM architecture on RiscPC and various embedded systems (EABI)
- sparc: Sun SPARC architecture on sun4u/v systems
- powerpc: PowerPC architecture
- ia64: Intel Itanium (IA-64) architecture
- mips, mipsel: MIPS architecture (big-endian and little-endian)
- s390: IBM ESA/390 architecture and z/Architecture

and as a "technology preview":

- kfreebsd-i386: Kernel of FreeBSD on x86 architecture
- kfreebsd-amd64: Kernel of FreeBSD on x86-64 architecture

Unstable ports

In addition to the stable ports, in the official *unstable* distribution are:

- hurd-i386: GNU Hurd kernel on x86 architecture
- armhf[132]: ARM hard-float architecture requiring hardware with a floating-point unit (FPU)
- s390x: IBM ESA/390 architecture and z/Architecture with 64-bit userland

Unofficial ports are also available as part of the *unstable* distribution at http://www.debian-ports.org [133]:

- alpha: DEC Alpha architecture
- hppa: HP PA-RISC architecture
- m68k: Motorola 68k architecture on Amiga, Atari, Macintosh and various embedded VME systems
- ppc64: PowerPC64 architecture supporting 64-bit PowerPC CPUs with VMX
- powerpcspe: PowerPCSPE architecture (binary-incompatible variant of the PowerPC)
- sh4: Hitachi SuperH architecture
- sparc64: Sun SPARC architecture with 64-bit userland

The m68k port was the second official one in Debian, and has been part of five stable Debian releases. Due to its failure to meet the release criteria, it was dropped before the release of *etch*. The arm (OABI[102], <armv4t), alpha and hppa ports were dropped before the release of *squeeze*.

Derivatives

Many distributions use one of Debian's branches as a base.

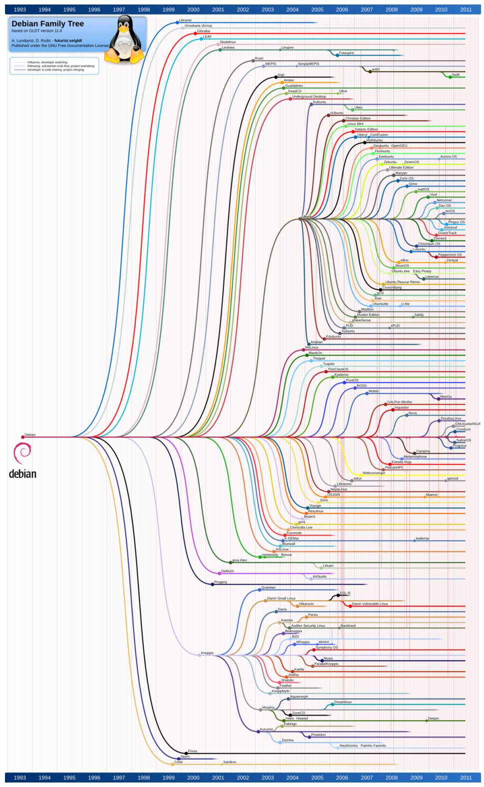

Debian family tree

Debian distribution branch	Distribution
Stable	MEPIS (SimplyMEPIS), Kanotix, CrunchBang Linux, Floppix
Testing	Ubuntu (long term release),[134] MEPIS (antiX), Parsix, Linux Mint (Debian edition)
Unstable	Ubuntu, aptosid (former sidux), Semplice Linux
Unknown	Damn Small Linux, Xandros, Knoppix, BackTrack, Linspire, LinEx, Vyatta, others [27]

Reception

Debian was ranked second only to Ubuntu (which is derived from Debian) for *Most Used Linux Distribution* for both personal and organizational use in a 2007 survey by *SurveyMonkey.com*.[135] Debian won the 2007 poll on *Server Distribution of the Year* by *LinuxQuestions.org*.[136]

Both the Debian distribution and their website have won various awards from different organizations. Debian was awarded the 2004 Readers' Choice Award for *Favorite Linux Distribution* by the *Linux Journal*.[137] A total of fifteen other awards have been awarded throughout Debian's lifetime including *Best Linux Distribution*.[138]

Debian has also received negative assessments. In May 2008, a Debian Developer revealed his discovery that changes made in 2006 to the random number generator in the version of the openSSL package distributed with Debian and other Debian-based distributions such as Ubuntu or Knoppix, made a variety of security keys vulnerable to a random number generator attack.[139][140] The security weakness was caused by changes made to the openssl code by another Debian Developer in response to memory debugger warnings.[141] The security hole was soon patched by Debian and others, but the complete resolution procedure was cumbersome for users because it involved regenerating all affected keys, and it drew criticism to Debian's practice of making Debian-specific changes to software.

Richard Stallman and the Free Software Foundation (FSF) have criticized the Debian Project for providing the *non-free* repository, rather than excluding this type of software entirely,[142] an opinion also echoed by some in Debian including the then-Project Leader Wichert Akkerman.[143] The internal dissent in the Debian Project regarding the *non-free* section has persisted, but the last time it came to a vote in 2006, a large majority decided to keep it.[144]

During the release cycles of *Woody* and *Sarge*, the Debian Project drew considerable criticism from the free software community because of the long time between *stable* releases. This triggered the creation of Ubuntu in 2004. Ubuntu has releases every 6 months which are forks of Debian's *unstable* distribution with bug fixes and other modifications. However, it may be more appropriate to compare Debian releases (which continue to be supported after the release of subsequent versions) to Ubuntu's Long Term Support releases (which are supported for five years for servers and also for five years for desktops starting with Ubuntu 12.04 LTS);[145] Ubuntu produces a new LTS release every two years, which is therefore similar to Debian's new two-year release cycle for post-Debian 6.0 releases.[83]

When in need of updated versions of software, it is possible to use Debian *testing* instead of *stable* as it usually contains more modern, though slightly less stable packages. Another alternative is to use Debian backports, which are "recompiled packages from testing (mostly) and unstable (in a few cases only, e.g. security updates), so they will run without new libraries (wherever it is possible) on a stable Debian distribution".[146]

References

[1] *Updated Debian 6.0: 6.0.5 released* (http://www.debian.org/News/2012/20120512), debian.org, May 12, 2012,

[2] http://en.wikipedia.org/wiki/Template%3Alatest_stable_software_release%2Fdebian?action=edit&preload=Template:LSR/syntax

[3] *Debian Installer 7.0 Alpha1 release* (http://www.debian.org/devel/debian-installer/News/2012/20120513), debian.org, May 13, 2012,

[4] http://en.wikipedia.org/wiki/Template%3Alatest_preview_software_release%2Fdebian?action=edit&preload=Template:LSR/syntax

[5] "Debian Installer 6.0 Release Candidate 1 published" (http://debian.org/News/2011/20110113). Debian. 2011-01-13. . Retrieved 2011-01-13.

[6] "License information" (http://www.debian.org/legal/licenses/). Debian. . Retrieved 2009-02-28.

[7] "What Does Free Mean? or What do you mean by Free Software?" (http://debian.org/intro/free). *Debian official website*. The Debian Project. . Retrieved 2009-12-01.

[8] "What is Debian?" (http://debian.org/index.en.html). *Debian official website*. The Debian Project. . Retrieved 2009-12-01.

[9] "Linux Distributions: Facts and Figures" (http://distrowatch.com/stats.php?section=popularity). distrowatch.com. . Retrieved 2008-10-31.

[10] "Usage statistics and market share of Linux for websites" (http://w3techs.com/technologies/details/os-linux/all/all). W3Techs.com. . Retrieved 2012-05-01.

[11] Debian Social Contract (http://www.debian.org/social_contract) on Debian homepage

[12] "The Debian GNU/Linux FAQ: Definitions and overview" (http://www.debian.org/doc/manuals/debian-faq/ch-basic_defs.en.html#s-whatisdebian). Debian. . Retrieved 2008-05-12.

[13] "The Debian Constitution" (http://www.debian.org/devel/constitution). Debian. . Retrieved 2008-05-18.

[14] "Debian Social Contract" (http://www.debian.org/social_contract). Debian. . Retrieved 2008-11-21.

[15] "Debian 6.0 "Squeeze" to be released with completely free Linux Kernel" (http://www.debian.org/News/2010/20101215). Debian. . Retrieved 2010-12-15.

[16] "A Brief History of Debian: Debian Releases" (http://www.debian.org/doc/manuals/project-history/ch-releases.en.html). Debian. . Retrieved 2008-10-31.

[17] "SPI Projects: Welcome to SPI" (http://www.spi-inc.org/projects/spi-projects). www.spi-inc.org. . Retrieved 2008-05-12.

[18] "Measuring Lenny: the size of Debian 5.0" (http://gsyc.es/~frivas/paper.pdf). libresoft. . Retrieved 2011-02-28.

[19] "Debian GNU/Linux" (http://www.ohloh.net/p/debian). Ohloh.net. . Retrieved 2011-02-28.

[20] "Debian 6.0 "Squeeze" released" (http://www.debian.org/News/2011/20110205a). Debian. . Retrieved 2011-02-06.

[21] "The Debian GNU/Linux FAQ: Compatibility issues" (http://www.debian.org/doc/manuals/debian-faq/ch-compat.en.html#s-arches). Debian. . Retrieved 2008-05-12.

[22] "Debian Moves to LibreOffice" (http://www.debian.org/News/2011/20110623). Debian. . Retrieved 2012-03-05.

[23] "Downloading Debian CD/DVD images via HTtp://FTP" (http://www.debian.org/CD/http-ftp/). Debian. . Retrieved 2008-12-11.

[24] "Installing Debian GNU/Linux via the Internet" (http://www.debian.org/distrib/netinst). Debian. . Retrieved 2008-12-11.

[25] "Debian GNU/Linux on CDs" (http://www.debian.org/CD/). Debian. . Retrieved 2009-01-06.

[26] Hillesley, Richard (2007-11-05). "Debian and the grass roots of Linux" (http://www.itpro.co.uk/applications/features/135084/debian-and-the-grass-roots-of-linux.html). . Retrieved 2008-10-31.

[27] "The Debian GNU/Linux FAQ: Choosing a Debian distribution" (http://www.debian.org/doc/manuals/debian-faq/ch-choosing.en.html#s3.2). Debian. . Retrieved 2008-05-12.
[28] "Debian Wiki dpkg" (http://wiki.debian.org/dpkg). Debian. . Retrieved 2008-12-10.
[29] "Debian Wiki Synaptic" (http://wiki.debian.org/Synaptic). Debian. . Retrieved 2008-12-10.
[30] "Chapter 2. What's new in Debian GNU/Linux 6.0" (http://www.debian.org/releases/squeeze/i386/release-notes/ch-whats-new.en.html). Debian. . Retrieved 2011-01-08.
[31] "Chapter 2. Debian package management. Section 2.2.1 apt-get / apt-cache vs. aptitude" (http://www.debian.org/doc/manuals/debian-reference/ch02.en.html#_literal_apt_get_literal_literal_apt_cache_literal_vs_literal_aptitude_literal). Debian. . Retrieved 2011-02-24.
[32] "Chapter 2. What's new in Debian GNU/Linux 5.0" (http://www.debian.org/releases/stable/amd64/release-notes/ch-whats-new.en.html). Debian. . Retrieved 2009-02-18.
[33] "DebianLive: Debian Wiki" (http://wiki.debian.org/DebianLive/). Debian. . Retrieved 2008-10-31.
[34] "DebianLive/ live-build: Debian Wiki" (http://live.debian.net/devel/live-build/). Debian. . Retrieved 2011-08-15.
[35] "The Ging FAQ" (http://glibc-bsd.alioth.debian.org/ging/). Debian. . Retrieved 2008-10-31.
[36] "Debian GNU/Linux 4.0 released" (http://www.debian.org/News/2007/20070408). Debian. 2008-04-08. . Retrieved 2008-10-31.
[37] "Upgrading your Slug LG #161" (http://linuxgazette.net/161/brownss.html). Linuxgazette.net. . Retrieved 2011-07-27.
[38] "Installing Debian on NSLU2" (http://www.cyrius.com/debian/nslu2/install.html). Cyrius.com. 2011-02-24. . Retrieved 2011-07-27.
[39] "Chapter 2. What's new in Debian GNU/Linux 5.0" (http://www.debian.org/releases/lenny/arm/release-notes/ch-whats-new.en.html#newforarm). Debian.org. . Retrieved 2011-07-27.
[40] "Chapter 2. What's new in Debian GNU/Linux 6.0" (http://www.debian.org/releases/squeeze/armel/release-notes/ch-whats-new.en.html#newforarm). Debian.org. . Retrieved 2011-07-27.
[41] The GLAN Tank is a Japanese NAS device, described on Japanese Wikipedia GLAN Tank **(Japanese)**
[42] "What's new in Debian GNU/Linux 4.0" (http://www.debian.org/releases/etch/arm/release-notes/ch-whats-new.en.html). *Release Notes for Debian GNU/Linux 4.0 ("etch"), ARM*. Debian.org. 2007-08-16. . Retrieved 2011-07-27.
[43] "A Brief History of Debian: Introduction: What is the Debian Project?" (http://www.debian.org/doc/manuals/project-history/ch-intro.en.html). Debian. . Retrieved 2008-05-12.
[44] Ian A Murdock (1993-08-16). "usDD.MIK@unix.portal.com New release under development; suggestions requested (news:CB)". [news:comp.os.linux.development comp.os.linux.development]. (Web link) (https://groups.google.com/group/comp.os.linux.development/msg/a32d4e2ef3bcdcc6?dmode=source&output=gplain&noredirect). Retrieved 2012-06-13.
[45] Nixon, Robin (2010). *Ubuntu: Up and Running* (http://books.google.com/books?id=badSTnfeOoAC). O'Reilly Media. p. 3. ISBN 978-0-596-80484-8. .
[46] Murdock, Ian A (1993-08-16). "usDD.MIK@unix.portal.com NNTP Subject: New release under development; suggestions requested (news:CB)". [news:comp.oslinux.development comp.oslinux.development]. (Web link) (http://groups.google.com/group/comp.os.linux.development/msg/a32d4e2ef3bcdcc6). Retrieved 2007-08-17.
[47] "Appendix A: The Debian Manifesto" (http://www.debian.org/doc/manuals/project-history/ap-manifesto.en.html). Debian. . Retrieved 2008-08-13.
[48] "What is the Debian Project?" (http://www.debian.org/doc/manuals/project-history/ch-intro.en.html#s1.1). *A Brief History of Debian*. Debian.org. . Retrieved 2011-07-27.
[49] "Debian GNU/Linux 2.2, the "Joel 'Espy' Klecker" release, is officially released" (http://www.debian.org/News/2000/20000815). Debian.org. 2000-08-15. . Retrieved 2011-07-27.
[50] "Chapter 2 - What's new in Debian GNU/Linux 3.1" (http://www.debian.org/releases/sarge/i386/release-notes/ch-whats-new.en.html). Debian. . Retrieved 2010-08-05.
[51] Hoover, Lisa (2006-10-10). "Behind the Debian and Mozilla dispute over use of Firefox" (http://www.linux.com/feature/57675). linux.com. . Retrieved 2009-02-09.
[52] "Debian GNU/Linux 5.0 released" (http://www.debian.org/News/2009/20090214.en.html). Debian. 2009-02-14. . Retrieved 2009-02-15.
[53] "Appendix C. Lenny dedicated to Thiemo Seufer" (http://www.debian.org/releases/stable/i386/release-notes/apc.en.html). Debian. 2009-02-14. . Retrieved 2009-02-22.
[54] "Debian -- News -- Backports service becoming official" (http://www.debian.org/News/2010/20100905). . Retrieved 13 September 2010.
[55] "Debian 6.0 "Squeeze" released" (http://www.debian.org/News/2011/20110205a). Debian -- News. 2011-02-06. . Retrieved 2011-02-06.
[56] "The Debian GNU/Linux FAQ: Getting support for Debian GNU/Linux" (http://www.debian.org/doc/manuals/debian-faq/ch-support.en.html#s12.2.1). Debian. . Retrieved 2008-05-12.
[57] "The Debian organization web page" (http://www.debian.org/intro/organization). Debian. . Retrieved 2008-11-01.
[58] O'Mahony, Siobhan; Fabrizio Ferraro (2008-04). "The Emergence of Governance in an Open Source Community" (http://web.archive.org/web/20080529081303/http://www.business.ualberta.ca/tcc/documents/TII_3_OMahoney_Ferraro_final.pdf). ualberta.ca. Archived from the original (http://www.business.ualberta.ca/tcc/documents/TII_3_OMahoney_Ferraro_final.pdf) on 2008-05-29. . Retrieved 2008-11-01.
[59] What does a Debian Project Leader do (http://www.debian.org/devel/leader) www.debian.org

[60] "A Brief History of Debian Chapter 2: Leadership" (http://www.debian.org/doc/manuals/project-history/ch-leaders). Debian. . Retrieved 2008-11-01.
[61] "DPL Platform — Stefano Zacchiroli" (http://www.debian.org/vote/2010/platforms/zack). Debian. 2010-03-12. section 2.2.3 No DPL board. . "I do not believe in the utility of a DPL board, so I will not have one; nor I will have a second in charge (2IC)."
[62] "Bits from the (chilly) release team" (http://lists.debian.org/debian-devel-announce/2010/08/msg00000.html). .
[63] "Debian New Maintainers" (http://www.debian.org/devel/join/newmaint). Debian. . Retrieved 2008-10-31.
[64] "How You Can Join" (http://www.debian.org/devel/join/). Debian. . Retrieved 2008-10-31.
[65] "The Debian GNU/Linux FAQ Chapter 6: The Debian FTP archives" (http://www.us.debian.org/doc/FAQ/ch-ftparchives). Debian. . Retrieved 2007-05-24.
[66] "Debian security FAQ" (http://www.debian.org/security/faq). Debian. 2007-02-28. . Retrieved 2008-10-21.
[67] Hess, Joey (2005-09-05). "announcing the beginning of security support for testing" (http://lists.debian.org/debian-devel-announce/2005/09/msg00006.html). *debian-devel-announce mailing list*. . Retrieved 2007-04-20.
[68] "Debian testing security team" (http://testing-security.debian.net). Debian. . Retrieved 2008-10-31.
[69] Simó, Adeodato (2009-02-01). "Release update: deep freeze, planned dates, and remaining bugs" (http://lists.debian.org/debian-devel-announce/2009/02/msg00000.html). *debian-devel-announce@lists.debian.org mailing list*. . Retrieved 2009-02-07.
[70] "General Resolution: Endorse the concept of Debian Maintainers" (http://www.debian.org/vote/2007/vote_003). Debian. . Retrieved 2008-12-13.
[71] "Debian Developer's Reference" (http://www.debian.org/doc/developers-reference/new-maintainer.html#registering). Debian. . Retrieved 2010-10-09.
[72] "Debian Developer Reference" (http://www.debian.org/doc/developers-reference/pkgs.html#testing). Debian. . Retrieved 2008-10-31.
[73] "Debian "testing" distribution" (http://www.debian.org/devel/testing). Debian. . Retrieved 2008-11-24.
[74] "Security Information" (http://www.debian.org/security/). Debian. . Retrieved 2008-12-13.
[75] "Securing Debian Manual" (http://www.debian.org/doc/user-manuals#securing). Debian. . Retrieved 2008-12-13.
[76] "Ubuntu Wiki CompilerFlags" (https://wiki.ubuntu.com/CompilerFlags). Ubuntu. . Retrieved 2011-01-31.
[77] "Debian Secure by Default" (http://d-sbd.alioth.debian.org/www/). Debian: Secure by Default Project. . Retrieved 2011-01-31.
[78] Release Goal Security Hardening Build Flags (http://wiki.debian.org/ReleaseGoals/SecurityHardeningBuildFlags)
[79] "Debian Releases" (http://www.us.debian.org/releases/). Debian. . Retrieved 2011-02-05.
[80] "Debian GNU/Linux 4.0 updated" (http://www.debian.org/News/2009/20090408). Debian. 2009-04-08. . Retrieved 2009-07-21.
[81] Brockschmidt, Marc (2009-02-15). "Debian squeeze waiting for development" (http://lists.debian.org/debian-devel-announce/2009/02/msg00003.html). Debian. . Retrieved 2009-02-15.
[82] "Debian testing security team" (http://secure-testing-master.debian.net/). Debian. . Retrieved 2008-10-31.
[83] "Debian decides to adopt time-based release freezes" (http://www.debian.org/News/2009/20090729). Debian. 2009-07-29. . Retrieved 2009-09-12.
[84] "Debian GNU/Linux 6.0 "Squeeze" release goals" (http://www.debian.org/News/2009/20090730). Debian. 2009-07-30. . Retrieved 2009-12-01.
[85] "Release Update: freeze guidelines, transitions, BSP, rc bug fixes" (http://lists.debian.org/debian-devel-announce/2010/09/msg00000.html). *debian-devel-announce mailing list*. 2010-09-03. . Retrieved 2010-09-03.
[86] "Bits from the nippy Release Team" (https://lists.debian.org/debian-devel-announce/2012/07/msg00004.html). *debian-devel-announce mailing list*. 2012-07-27. . Retrieved 2010-07-27.
[87] "Artwork for Debian CDs" (http://www.debian.org/CD/artwork/). Debian. . Retrieved 2009-01-11.
[88] *date of file packages* (http://archive.debian.org/debian/dists/Debian-1.1/main/binary-i386/Packages.gz),
[89] Schulze, Martin (2000-08-15). "Debian GNU/Linux 2.2, the "Joel 'Espy' Klecker" release" (http://lists.debian.org/debian-announce/debian-announce-2000/msg00009.html). *debian-announce mailing list*. .
[90] "Debian Wiki: Debian Releases > Debian Etch" (http://wiki.debian.org/DebianEtch). Debian. . Retrieved 2010-02-16.
[91] Schmehl, Alexander (2007-04-08). "Debian GNU/Linux 4.0 released" (http://lists.debian.org/debian-announce/debian-announce-2007/msg00002.html). *debian-announce mailing list*. . Retrieved 2008-11-01.
[92] "Debian -- News -- Debian GNU/Linux 4.0 updated" (http://www.debian.org/News/2010/20100522). Debian. . Retrieved 2012-03-12.
[93] Brockschmidt, Marc (2008-03-02). "Release Update: Release numbering, goals, armel architecture, BSPs" (http://lists.debian.org/debian-devel-announce/2008/03/msg00001.html). *debian-announce mailing list*. . Retrieved 2008-11-01.
[94] Langasek, Steve (2006-11-16). "testing d-i Release Candidate 1 and more release adjustments" (http://lists.debian.org/debian-devel-announce/2006/11/msg00004.html). *debian-devel-announce mailing list*. . Retrieved 2008-11-01.
[95] Muehlenhoff, Moritz (2011-12-06). "[SECURITY (http://lists.debian.org/debian-security-announce/2011/msg00238.html) [DSA 2360-1] Two month advance notification for upcoming end-of-life for Debian oldstable"]. . Retrieved 2011-12-07.
[96] Brockschmidt, Marc (2008-06-02). "Release Update: arch status, major transitions finished, freeze coming up" (http://lists.debian.org/debian-devel-announce/2008/06/msg00000.html). *debian-devel-announce mailing list*. . Retrieved 2008-11-01.
[97] Smakov, Jurij (2007-07-18). "Retiring the sparc32 port" (http://lists.debian.org/debian-devel-announce/2007/07/msg00006.html). *debian-devel-announce mailing list*. . Retrieved 2008-10-31.
[98] Armstrong, Ben (2008-08-03). "Bits from the Debian Eee PC team, summer 2008" (http://lists.debian.org/debian-devel-announce/2008/08/msg00002.html). *debian-devel-announce mailing list*. . Retrieved 2008-10-31.

[99] "Debian -- News -- Updated Debian 5.0: 5.0.10 released" (http://www.debian.org/News/2012/20120310). Debian. 2012-03-10. . Retrieved 2012-03-12.

[100] "Debian GNU/Linux 6.0: Release Notes" (http://www.debian.org/releases/squeeze/releasenotes). Debian. . Retrieved 2009-02-15.

[101] Claes, Luk (2008-09-01). "Release Update: freeze guidelines, testing, BSP, rc bug fixes" (http://lists.debian.org/debian-devel-announce/2008/09/msg00000.html). *debian-devel-announce mailing list*. . Retrieved 2008-10-31.

[102] Wiki.debian.org (http://wiki.debian.org/ArmEabiPort#Terminology)

[103] "Aurelien's weblog: Debian is switching to EGLIBC" (http://blog.aurel32.net/?p=47). Aurélien Jarno. . Retrieved 2009-05-21.

[104] "LSBInitScripts/DependencyBasedBoot" (http://wiki.debian.org/LSBInitScripts/DependencyBasedBoot). Debian Wiki. . Retrieved 2011-07-27.

[105] "ReleaseGoals/RemoveOldLibs" (http://wiki.debian.org/ReleaseGoals/RemoveOldLibs). Debian Wiki. . Retrieved 2011-07-27.

[106] "Bits from the Release Team: Final countdown!" (http://lists.debian.org/debian-devel-announce/2012/06/msg00006.html). Debian.org. 2012-06-21. . Retrieved 2012-06-23.

[107] "Bit from the Release Team: armhf and s390x" (https://lists.debian.org/debian-devel-announce/2012/06/msg00003.html). debian.org. 2012-06-07. . Retrieved 2011-06-14.

[108] "Ana’s blog » Blog Archive » post and pre-release fun" (http://ekaia.org/blog/2011/02/07/post-and-pre-release-fun/). Ekaia.org. 2011-02-07. . Retrieved 2011-07-27.

[109] "Debian 7 Wheezy to introduce multiarch support" (http://www.debian.org/News/2011/20110726b). *Debian*. . Retrieved 27 July 2011.

[110] "Index of /pub/linux/kernel/v1.3" (http://www.kernel.org/pub/linux/kernel/v1.3/). Kernel.org. . Retrieved 2011-07-27.

[111] "Index of /pub/linux/kernel/v2.0" (http://www.kernel.org/pub/linux/kernel/v2.0/). Kernel.org. . Retrieved 2011-07-27.

[112] kernel-image-* packages from Debian 1.3 i386 Packages file (http://archive.debian.org/debian/dists/Debian-1.3.1/main/binary-i386/Packages)

[113] "Index of /pub/linux/kernel/v2.1" (http://www.kernel.org/pub/linux/kernel/v2.1/). Kernel.org. . Retrieved 2011-07-27.

[114] kernel-image-* packages from Debian 2.0 i386 Packages file (http://archive.debian.org/debian/dists/Debian-2.0/main/binary-i386/Packages)

[115] "Packages in slink(2.1) only" (http://www.debian.org/releases/slink/i386/reps_hamm_slink/rep.0.html). Debian.org. 1998-02-11. . Retrieved 2011-07-27.

[116] "Index of /pub/linux/kernel/v2.2" (http://www.kernel.org/pub/linux/kernel/v2.2/). Kernel.org. . Retrieved 2011-07-27.

[117] "Index of /pub/linux/kernel/v2.3" (http://www.kernel.org/pub/linux/kernel/v2.3/). Kernel.org. . Retrieved 2011-07-27.

[118] "Debian GNU/Linux 3.0 released" (http://lists.debian.org/debian-announce/2002/msg00004.html). Lists.debian.org. 2002-07-19. . Retrieved 2011-07-27.

[119] "Index of /pub/linux/kernel/v2.4" (http://www.kernel.org/pub/linux/kernel/v2.4/). Kernel.org. . Retrieved 2011-07-27.

[120] "Index of /pub/linux/kernel/v2.5" (http://www.kernel.org/pub/linux/kernel/v2.5/). Kernel.org. . Retrieved 2011-07-27.

[121] "Index of /pub/linux/kernel/v2.6" (http://www.kernel.org/pub/linux/kernel/v2.6/). Kernel.org. . Retrieved 2011-07-27.

[122] "Debian 6.0 "Squeeze" Frozen" (http://www.debian.org/News/2010/20100806). Debian. . Retrieved 2010-08-16.

[123] "Chapter 4. Resources for Debian Developers" (http://www.debian.org/doc/developers-reference/resources.html). Debian. . Retrieved 2008-10-31.

[124] "Explaining Why We Don't Endorse Other Systems" (http://www.gnu.org/philosophy/common-distros.html#Debian). FSF. . Retrieved 2010-02-21.

[125] "Debian permits binary blobs" (http://www.desktoplinux.com/news/NS8731622728.html). DesktopLinux.com. . Retrieved 2009-12-29.

[126] "Chapter 2: Basic Configuration" (http://www.debian.org/doc/manuals/apt-howto/ch-basico.en.html). Debian. . Retrieved 2008-10-31.

[127] "2.1. Supported Hardware Chapter 2. System Requirements" (http://www.debian.org/releases/stable/i386/ch02s01.html.en). Debian. . Retrieved 2008-11-02.

[128] "Meeting Minimum Hardware Requirements" (http://www.debian.org/releases/stable/i386/ch03s04.html.en). Debian. . Retrieved 2008-10-31.

[129] http://www.emdebian.org/

[130] "Debian Ports" (http://www.debian.org/ports/). Debian. . Retrieved 2009-02-15.

[131] "Supported Hardware" (http://www.debian.org/releases/stable/i386/ch02s01.html.en). Debian. . Retrieved 2008-10-11.

[132] "Debian Wiki: armhf - Arm Hard Float Port" (http://wiki.debian.org/ArmHardFloatPort). .

[133] http://www.debian-ports.org

[134] "LTS - Ubuntu Wiki" (https://wiki.ubuntu.com/LTS). *Wiki.ubuntu.com*. 2012-03-02. . Retrieved 2012-04-26.

[135] "2007 Linux Desktop/Client Survey Results" (http://www.surveymonkey.com/sr.aspx?sm=WE9PoU_2bO_2b9Incc1aPjQTMvqIHQ7sZ6Y5sA1yP920aDE_3d). surveymonkey.com. . Retrieved 2008-11-02.

[136] "2007 LinuxQuestions.org Members Choice Awards" (http://www.linuxquestions.org/questions/2007-linuxquestions.org-members-choice-awards-79/server-distribution-of-the-year-610200/). LinuxQuestions.org. . Retrieved 2008-11-02.

[137] "2004 Readers' Choice Awards" (http://www.linuxjournal.com/article/7724). linuxjournal.com. 2004-11-01. . Retrieved 2008-11-02.

[138] "Awards" (http://www.debian.org/misc/awards). Debian. . Retrieved 2008-11-02.

[139] "DSA-1571-1 openssl: predictable random number generator" (http://www.debian.org/security/2008/dsa-1571). Debian. . Retrieved 2008-10-31.
[140] "CVE-2008-0166 (under review)" (http://cve.mitre.org/cgi-bin/cvename.cgi?name=CVE-2008-0166). mitre.org. . Retrieved 2008-10-31.
[141] "Debian OpenSSL Security Flaw" (http://cryptogon.com/?p=2635). cryptogon.com. 2008-05-26. . Retrieved 2008-10-31.
[142] "Explaining Why We Don't Endorse Other Systems" (http://www.gnu.org/philosophy/common-distros.html). Free Software Foundation. . Retrieved 2009-09-28.
[143] "Moving contrib and non-free [out] of master.debian.org" (http://lists.debian.org/debian-vote/1999/06/msg00043.html). *debian-devel mailing list*. 1999-06-21. .
[144] "General Resolution: Status of the non-free section" (http://www.debian.org/vote/2004/vote_002). Debian. . Retrieved 2009-09-28.
[145] Ubuntu Wiki: LTS (https://wiki.ubuntu.com/LTS)
[146] "Debian Backports" (http://www.backports.org/dokuwiki). Debian Project. . Retrieved 2010-04-10.

Further reading

- Krafft, Martin F. (2005). *The Debian System* (http://www.debiansystem.info). U.S.A.: No Starch Press. p. 608. ISBN 1-59327-069-0.
- Mako Hill, Benjamin (2005). *Debian GNU/Linux 3.1 Bible*. John Wiley & Sons. p. 672. ISBN 0-7645-7644-5.
- Jurzik, Heike (2005). *Debian GNU / Linux*. Galileo Press GmbH. p. 630. ISBN 3-89842-627-0.
- Schoblick, Robert (2005). *Debian / GNU Linux*. BHV Verlag. p. 767. ISBN 3-8266-8151-7.
- Bellomo, Michael (2000). *Debian GNU/Linux for Dummies*. Hungry Minds. p. 324. ISBN 0-7645-0713-3.
- McCarty, Bill (1999). *Learning Debian GNU/Linux*. O'Reilly. p. 360. ISBN 1-56592-705-2.
- Goerzen, John (1999). *Debian Gnu/Linux: Guide to Installation and Usage*. New Riders Pub. p. 158. ISBN 0-7357-0914-9.
- Scheetz, Dale (1998). *The Debian Linux User's Guide*. Linux Press. p. 268. ISBN 0-9659575-1-9.

External links

- Official website for Debian (http://www.debian.org/)
- Official website for Debian Live (http://live.debian.net/)
- Official website for the Debian Wiki (http://wiki.debian.org/)
- *Debian GNU/Linux* (http://distrowatch.com/table.php?distribution=debian) at DistroWatch

Ubuntu (operating system)

Ubuntu

Ubuntu 12.04 LTS desktop (Precise Pangolin)

Company / developer	Canonical Ltd. / Ubuntu Foundation
OS family	Unix-like
Working state	Current
Source model	Open source
Initial release	20 October 2004
Latest stable release	12.04 (Precise Pangolin) / 26 April 2012[1]
Available language(s)	Multilingual (more than 55)
Update method	APT (front-ends available)
Package manager	dpkg (front-ends like Synaptic available)
Supported platforms	i386, AMD64[2], ARM[2][3][4][5][6]
Kernel type	Monolithic (Linux kernel)
Userland	GNU
Default user interface	Unity shell on top of GNOME 3.x
License	Mainly the GNU GPL and various other free software licenses
Official website	[www.ubuntu.com www.ubuntu.com]

Ubuntu (/ʊˈbʊntuː/ *oo-**BOON**-too*)[7] [8] is a computer operating system based on the Debian Linux distribution and distributed as free and open source software, using its own desktop environment. It is named after the Southern African philosophy of *ubuntu* ("humanity towards others").[9]

As of 2012, according to online surveys, Ubuntu is the most popular Linux distribution on desktop/laptop personal computers[10][11][12][13], and most Ubuntu coverage focuses on its use in that market. However, it is also popular on servers and for cloud computing.

Ubuntu is sponsored by the UK-based company Canonical Ltd., owned by South African entrepreneur Mark Shuttleworth. Canonical generates revenue by selling technical support and services related to Ubuntu, while the operating system itself is entirely free of charge. The Ubuntu project is committed to the principles of free software development; people are encouraged to use free software, improve it, and distribute it.[14]

History and development process

Ubuntu is a fork of the Debian project's codebase. The original aim of the Ubuntu developers was to create an easy-to-use Linux desktop with new releases scheduled on a predictable six-month basis, resulting in a frequently updated system.[8][15]

Ubuntu's first release was on 20 October 2004. Since then, Canonical has released new versions of Ubuntu every six months[16] with commitment to support each release for eighteen months by providing security fixes, patches to critical bugs and minor updates to programs. It was decided that every fourth release, issued on a two-year basis, would receive long-term support (LTS).[8] LTS releases were traditionally supported for three years on the desktop and five years on the server.[16] However with the release of Ubuntu 12.04 LTS, desktop support for LTS releases was extended to five years (for example, Ubuntu 12.04 LTS is scheduled to be supported until April 2017). Support was extended to better accommodate business and corporate IT users of Ubuntu who operate on longer release cycles and are more conscious of the costs associated with frequent software upgrades.[17]

Ubuntu packages are based on packages from Debian's unstable branch: both distributions use Debian's deb package format and package management tools (APT and Ubuntu Software Center). Debian and Ubuntu packages are not necessarily binary compatible with each other, however, and sometimes .deb packages may need to be rebuilt from source to be used in Ubuntu.[18] Many Ubuntu developers are also maintainers of key packages within Debian. Ubuntu cooperates with Debian by pushing changes back to Debian,[19] although there has been criticism that this does not happen often enough. In the past, Ian Murdock, the founder of Debian, has expressed concern about Ubuntu packages potentially diverging too far from Debian to remain compatible.[20] Before release, packages are imported from Debian Unstable continuously and merged with Ubuntu-specific modifications. A month before release, imports are frozen, and packagers then work to ensure that the frozen features interoperate well together.

Ubuntu is currently funded by Canonical Ltd. On 8 July 2005, Mark Shuttleworth and Canonical Ltd. announced the creation of the Ubuntu Foundation and provided an initial funding of US$10 million. The purpose of the foundation is to ensure the support and development for all future versions of Ubuntu. Mark Shuttleworth describes the foundation as an "emergency fund" (in case Canonical's involvement ends).[21]

On 12 March 2009, Ubuntu announced developer support for 3rd party cloud management platforms, such as for those used at Amazon EC2.[22]

The latest release is Ubuntu 12.04 LTS (Precise Pangolin), released on 26 April 2012.

Mark Shuttleworth announced on 31 October 2011 that Ubuntu's support for smartphones, tablets, TVs and smart screens is scheduled to be added by Ubuntu 14.04.[23] On 9 January 2012, Canonical announced Ubuntu TV at the Consumer Electronics Show.[24]

Features

Ubuntu is composed of many software packages, the vast majority of which are distributed under a free software license. The only exceptions are some proprietary hardware drivers.[25] The main license used is the GNU General Public License (GNU GPL) which, along with the GNU Lesser General Public License (GNU LGPL), explicitly declares that users are free to run, copy, distribute, study, change, develop and improve the software. On the other hand, there is also proprietary software available that can run on Ubuntu. The Ubiquity installer allows Ubuntu to be installed to the hard disk from within the Live CD environment, without the need for restarting the computer prior to installation. Ubuntu also emphasizes accessibility and internationalization to reach as many people as possible. Beginning with 5.04, UTF-8 became the default character encoding,[26] which allows for support of a variety of non-Roman scripts. As a security feature, the sudo tool is used to assign temporary privileges for performing administrative tasks, allowing the root account to remain locked, and preventing inexperienced users from inadvertently making catastrophic system changes or opening security holes.[27] PolicyKit is also being widely implemented into the desktop to further harden the system through the principle of least privilege.

Ubuntu Desktop includes a graphical desktop environment. In versions prior to 11.04 the default GUI was GNOME Panel but it was dropped in favor of Unity, a graphical interface Canonical first developed for the Ubuntu Netbook Edition.[28]

Ubuntu comes installed with a wide range of software that includes LibreOffice, Firefox, Thunderbird, Empathy, Transmission, and several lightweight games (such as Sudoku and chess). Additional software that is not installed by default (including software that used to be in the default installation such as Evolution, GIMP, Pidgin, and Synaptic) can be downloaded and installed using the Ubuntu Software Center[29] or other apt based package management tools. Programs in the Software Center are mostly free, but there are also paid for products, including applications and magazines.

Ubuntu allows networking ports to be closed using its firewall, with customized port selection available. End-users can install Gufw (GUI for Uncomplicated Firewall) and keep it enabled.[30] GNOME (the former default desktop) offers support for more than 46 languages.[31] Ubuntu can also run many programs designed for Microsoft Windows (such as Microsoft Office), through Wine or using a Virtual Machine (such as VMware Workstation or VirtualBox).

Ubuntu compiles their packages using gcc features such as PIE and Buffer overflow protection to harden their software.[32] These extra features greatly increase security at the performance expense of 1% in 32 bit and 0.01% in 64 bit.[33]

System requirements

As of version 12.04, Ubuntu supports the ARM[2][3][4][5] and x86 (32 bit and 64 bit) architectures. There is unofficial support for PowerPC.[34][35][2]

The system requirements vary among Ubuntu products. For the main Ubuntu desktop product, the official Ubuntu Documentation recommends a 1GHz Pentium 4 with 512 megabytes of RAM and 5 gigabytes of Hard Drive space, or better.[36] For slower computers, there are other Ubuntu products such as Lubuntu and Xubuntu.

Installation

Installation of Ubuntu is generally performed with the Live CD or can be installed via a Live USB drive. The Ubuntu OS can run directly from the CD (although this is usually slower than running Ubuntu from an HDD), allowing a user to "test-drive" the OS for hardware compatibility and driver support. The CD also contains the Ubiquity installer,[37] which can then guide the user through the permanent installation process. CD images of all current and past versions are available for download at the Ubuntu web site.[38] Installing from the CD requires a minimum of 256 MiB of RAM.

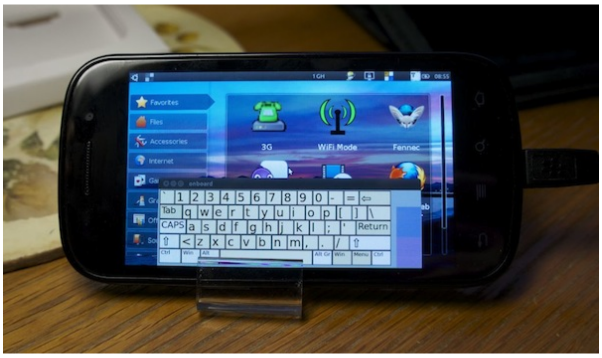

Ubuntu running on the Nexus S, an Android smartphone

Users can download a disk image (.iso) of the CD, which can then either be written to a physical medium (CD or DVD), or optionally run directly from a hard drive (via UNetbootin or GRUB). Ubuntu is also available on ARM, PowerPC, SPARC, and IA-64 platforms, although none are officially supported.[39]

Canonical offered Ubuntu[40] and Kubuntu[41] Live installation CDs of the latest distribution of the operating system at no cost (though they now charge £5.00 for 5), including paid postage for destinations in most countries around the world (via a service called ShipIt). Various third-party programs such as remastersys and Reconstructor are available to create customised copies of the Ubuntu Live CDs.

A Microsoft Windows migration tool, called Migration Assistant (introduced in April 2007),[42] can be used to import bookmarks, desktop background (wallpaper), and various settings from an existing Windows installation into a new Ubuntu installation.[43]

Ubuntu and Kubuntu can be booted and run from a USB Flash drive (as long as the BIOS supports booting from USB), with the option of saving settings to the flashdrive. This allows a portable installation that can be run on any PC which is capable of booting from a USB drive.[44] In newer versions of Ubuntu, the USB creator program is available to install Ubuntu on a USB drive (with or without a LiveCD disc).

Wubi, which is included as an option on the Live CD,[45] allows Ubuntu to be installed and run from within a virtual Windows loop device (as a large image file that is managed like any other Windows program via the Windows Control Panel). This method requires no partitioning of a Windows user's hard drive. Wubi also makes use of the Migration Assistant to import users' settings. It also incurs a slight performance loss. Hibernation is not supported and the filesystem is more vulnerable to hard reboots.

Package classification and support

Ubuntu divides all software into four domains to reflect differences in licensing and the degree of support available.[25] Some unsupported applications receive updates from community members, but not from Canonical Ltd.

	Free software	Non-free software
Supported	Main	Restricted
Unsupported	Universe	Multiverse

Free software includes only software that has met the Ubuntu licensing requirements,[25] which roughly correspond to the Debian Free Software Guidelines. Exceptions, however, include firmware and fonts, in the Main category, because although they are not allowed to be modified, their distribution is otherwise unencumbered.

Non-free software is usually unsupported (Multiverse), but some exceptions (Restricted) are made for important non-free software. Supported non-free software includes device drivers that can be used to run Ubuntu on some current hardware, such as binary-only graphics card drivers. The level of support in the Restricted category is more limited than that of Main, because the developers may not have access to the source code. It is intended that Main and Restricted should contain all software needed for a complete desktop environment.[25] Alternative programs for the same tasks and programs for specialized applications are placed in the Universe and Multiverse categories.

In addition to the above, in which the software does not receive new features after an initial release, *Ubuntu Backports* is an officially recognized repository for backporting newer software from later versions of Ubuntu.[46] The repository is not comprehensive; it consists primarily of user-requested packages, which are approved if they meet quality guidelines. Backports receives no support at all from Canonical, and is entirely community-maintained.

The *-updates* repository provides stable release updates (SRU) of Ubuntu and are generally installed through update-manager. Each release is given its own -updates repository (e.g. intrepid-updates). The repository is supported by Canonical Ltd. for packages in main and restricted, and by the community for packages in universe and multiverse. All updates to the repository must meet certain requirements and go through the *-proposed* repository before being made available to the public.[47] Updates are scheduled to be available until the end of life for the release.

In addition to the -updates repository, the unstable *-proposed* repository contains uploads which must be confirmed before being copied into -updates. All updates must go through this process to ensure that the patch does truly fix the bug and there is no risk of regression.[48] Updates in -proposed are confirmed by either Canonical or members of the community.

Canonical's *partner* repository lets vendors of proprietary software deliver their products to Ubuntu users at no cost through the same familiar tools for installing and upgrading software.[49] The software in the partner repository is officially supported with security and other important updates by its respective vendors. Canonical supports the packaging of the software for Ubuntu[][50] and provides guidance to vendors.[49] The partner repository is disabled by default and can be enabled by the user.[51] Some popular products distributed via the partner repository as of November 2011 are Adobe Flash Player, Adobe Reader, Braid and Oil Rush.[52]

Availability of third-party software

Ubuntu has a certification system for third party software.[53] Some third-party software that does not limit distribution is included in Ubuntu's multiverse component. The package *ubuntu-restricted-extras* additionally contains software that may be legally restricted, including support for MP3 and DVD playback, Microsoft TrueType core fonts, Sun's Java runtime environment, Adobe's Flash Player plugin, many common audio/video codecs, and unrar, an unarchiver for files compressed in the RAR file format.

Additionally, third party application suites are available for purchase through the Canonical web-based store,[54] including software for DVD playback and media codecs.

Releases

<table>
<tr><th rowspan="2">Version</th><th rowspan="2">Code name</th><th rowspan="2">Release date</th><th colspan="2">Supported until</th></tr>
<tr><th>Desktop</th><th>Server</th></tr>
<tr><td>4.10</td><td>Warty Warthog</td><td>2004-10-20</td><td colspan="2">2006-04-30</td></tr>
<tr><td>5.04</td><td>Hoary Hedgehog</td><td>2005-04-08</td><td colspan="2">2006-10-31</td></tr>
<tr><td>5.10</td><td>Breezy Badger</td><td>2005-10-13</td><td colspan="2">2007-04-13</td></tr>
<tr><td>6.06 LTS</td><td>Dapper Drake</td><td>2006-06-01</td><td>2009-07-14</td><td>2011-06-01</td></tr>
<tr><td>6.10</td><td>Edgy Eft</td><td>2006-10-26</td><td colspan="2">2008-04-25</td></tr>
<tr><td>7.04</td><td>Feisty Fawn</td><td>2007-04-19</td><td colspan="2">2008-10-19</td></tr>
<tr><td>7.10</td><td>Gutsy Gibbon</td><td>2007-10-18</td><td colspan="2">2009-04-18</td></tr>
<tr><td>8.04 LTS</td><td>Hardy Heron</td><td>2008-04-24</td><td>2011-05-12</td><td>2013-04</td></tr>
<tr><td>8.10</td><td>Intrepid Ibex</td><td>2008-10-30</td><td colspan="2">2010-04-30</td></tr>
<tr><td>9.04</td><td>Jaunty Jackalope</td><td>2009-04-23</td><td colspan="2">2010-10-23</td></tr>
<tr><td>9.10</td><td>Karmic Koala</td><td>2009-10-29</td><td colspan="2">2011-04-30</td></tr>
<tr><td>10.04 LTS</td><td>Lucid Lynx</td><td>2010-04-29</td><td>2013-04</td><td>2015-04</td></tr>
<tr><td>10.10</td><td>Maverick Meerkat</td><td>2010-10-10</td><td colspan="2">2012-04-10</td></tr>
<tr><td>11.04</td><td>Natty Narwhal</td><td>2011-04-28</td><td colspan="2">2012-10</td></tr>
<tr><td>11.10</td><td>Oneiric Ocelot</td><td>2011-10-13</td><td colspan="2">2013-04</td></tr>
<tr><td>12.04 LTS</td><td>Precise Pangolin</td><td>2012-04-26</td><td colspan="2">2017-04</td></tr>
<tr><td>12.10</td><td>Quantal Quetzal</td><td>2012-10-18</td><td colspan="2">2014-04</td></tr>
<tr><th colspan="2">Colour</th><th colspan="3">Meaning</th></tr>
<tr><td colspan="2">Red</td><td colspan="3">Release no longer supported</td></tr>
<tr><td colspan="2">Green</td><td colspan="3">Release still supported</td></tr>
<tr><td colspan="2">Blue</td><td colspan="3">Future release</td></tr>
</table>

Each Ubuntu release has a version number that consists of the year and month number of the release.[55] For example, the first release was Ubuntu 4.10 as it was released on 20 October 2004. Version numbers for future versions are provisional; if the release is delayed the version number changes accordingly.

Ubuntu releases are also given alliterative code names, using an adjective and an animal (e.g., "Dapper Drake" and "Intrepid Ibex"). With the exception of the first three releases, code names are in consecutive alphabetical order, allowing a quick determination of which release is newer. "We might skip a few letters, and we'll have to wrap eventually." says Mark Shuttleworth while describing the naming scheme.[56] Commonly, Ubuntu releases are referred to using only the adjective portion of the code name.

Releases are timed to be approximately one month after GNOME releases (which in turn are about one month after releases of X.org). Consequently, every Ubuntu release comes with an updated version of both GNOME and X.

Upgrades between releases have to be done from one release to the next release (e.g. Ubuntu 10.04 to Ubuntu 10.10) or from one LTS release to the next LTS release (e.g. Ubuntu 8.04 LTS to Ubuntu 10.04 LTS).[57]

Ubuntu 10.10 (Maverick Meerkat), was released on 10 October 2010 (10/10/10). This is a departure from the traditional schedule of releasing at the end of October to get "the perfect 10",[58] and a playful reference to *The Hitchhiker's Guide to the Galaxy*, since, in binary, 101010 is equal to the number 42, the "Answer to the Ultimate Question of Life, the Universe and Everything" within the series.[59]

Ubuntu 11.04 was released on 28 April 2011, and is code named "Natty Narwhal".[60] The desktop interface of this release significantly differs from the previous releases because Unity was introduced as the default GUI. It is easy to switch into "classic" GUI (GNOME Panel).[61] The new GUI has received strong criticism from some users as too different from and less capable than the previous Gnome Panel,[62][63] while other users have found they prefer the new approach and the minimalism compared to the older desktop paradigm.[64] However, those positive about Unity also believed there was much room for improvement.[65]

Variants

Official Ubuntu editions, which are created and maintained by Canonical and the Ubuntu community and receive full support from Canonical, its partners and the Community, are the following:[66][67]

- **Ubuntu Desktop**, designed for desktop and laptop PCs. (Formerly there was also Ubuntu Netbook Edition, designed for netbooks and other ultra-portables with screens up to 10", but it was discontinued as its user interface and functionality was integrated into the desktop edition.) The desktop edition can be also installed using the **alternative install CD** which uses the debian-installer and allows performing certain specialist installations of Ubuntu: setting up automated deployments, upgrading from older installations without network access, LVM and/or RAID partitioning, installs on systems with less than about 256 MiB of RAM (although low-memory systems may not be able to run a full desktop environment reasonably).[68]
- **Ubuntu Server**, made for use in servers.[69] The server install CD allows the user to install Ubuntu permanently on a computer for use as a server. It does not install a graphical user interface.

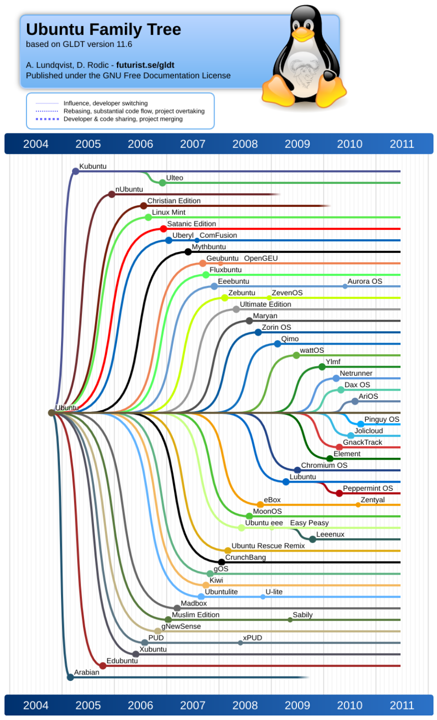

Ubuntu family tree

- **Ubuntu Business Desktop Remix**, a release meant for business users that comes with special enterprise software including Adobe Flash, Canonical Landscape, OpenJDK 6 and VMware View, while removing social networking and file sharing applications, games and development/sysadmin tools.[70] The goal of the Business Desktop Remix is not to copy other enterprise-oriented distributions, such as Red Hat Enterprise Linux, but to make it, according to Mark Shuttleworth's blog, "easier for institutional users to evaluate Ubuntu Desktop for their specific needs."[71]
- **Ubuntu TV**, labeled "TV for human beings" by Canonical, was introduced at the 2012 Consumer Electronics Show.[72] Created for SmartTVs, Ubuntu TV provides access to popular Internet services and stream content to mobile devices running Android, iOS and Ubuntu.[73]
- **Ubuntu for Android**, variant of Ubuntu designed to run on Android phones. It is expected to come pre-loaded on several phones.[74] *Ubuntu for Android* is expected to be shown at Mobile World Congress 2012.[75]

Ubuntu TV

By *Precise Pangolin* (12.04), Kubuntu is a community-supported variant of the Ubuntu distribution which uses the KDE Plasma Workspaces.

There are many Ubuntu variants (or derivatives) based on the official Ubuntu editions. These Ubuntu variants install a set of packages that differ from the official Ubuntu distributions.

The variants recognized by Canonical as contributing significantly towards the Ubuntu project are the following:[66]

- Edubuntu, a GNOME-based subproject and add-on for Ubuntu, designed for school environments and home users.[76]
- Kubuntu, a desktop distribution using the KDE Plasma Workspaces desktop environment.
- Lubuntu, a lightweight distribution using the LXDE desktop environment.
- Mythbuntu, designed for creating a home theater PC with MythTV and uses the Xfce desktop environment.
- Ubuntu Studio, a distribution made for professional video and audio editing, comes with higher-end free editing software and is a DVD .iso image unlike the Live CD the other Ubuntu distributions use.
- Xubuntu, a distribution based on the Xfce desktop environment, designed to run more efficiently on low-specification computers.

Kubuntu, Mythbuntu, Ubuntu Studio, Xubuntu and Gobuntu are not commercially supported by Canonical.[77]

Other variants are created and maintained by individuals and organizations outside of Canonical, and they are self-governed projects that work more or less closely with the Ubuntu community.[67]

Ubuntu Server Edition

Ubuntu also offers its operating system in a server edition. The Ubuntu 10.04 Long Term Support (LTS) release is scheduled to receive continued updates until April 2015. Starting with 12.04 the support of the LTS deskop variants has been extended to match the 5 years of the server variant. Long term support include updates to support new features of the latest computing hardware, security patches and updates to the 'Ubuntu stack' (cloud computing infrastructure).[78]

Ubuntu 10.04 Server Edition can also run on VMware ESX Server, Oracle's VirtualBox and VM, Citrix Systems XenServer hypervisors, as well as Kernel-based Virtual Machine. Ubuntu uses AppArmor security module for the Linux kernel which is turned on by default on key software packages, and the firewall is extended to common services used by the operating system. The home and Private directories are also encrypted. The 10.04 server version includes MySQL 5.1, Tomcat 6, OpenJDK 6, Samba 3.4, Nagios 3, PHP 5.3, Python 2.6. Many of its services only take 30 minutes to configure.[78]

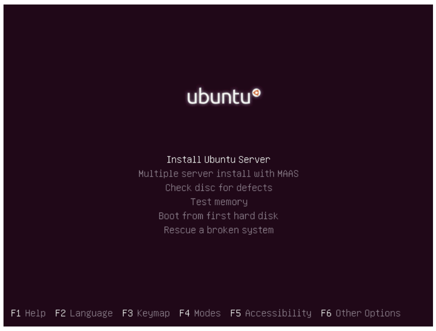

A screenshot of the Ubuntu 12.04 Server installation boot menu

Ubuntu 10.04 LTS Server Edition supports two major architectures: Intel x86 and AMD64. The server edition provides features such as file/print services, web hosting, email hosting, etc. There are a few differences between the Ubuntu Server Edition and the Ubuntu Desktop Edition although both use the same apt repositories. The main difference between the two editions is the lack of a default installation of a X window environment in the server edition, although GUIs can be installed like GNOME/Unity (Ubuntu 11.04), KDE (Kubuntu 11.04), XFCE, (Xubuntu 11.04), as well as more resource-economical GUIs such as Fluxbox, Openbox and Blackbox. Kernel versions also differ. The server edition uses a screen mode character-based interface for the installation, instead of a graphical installation process. The server CD also has the option of installing Ubuntu enterprise cloud.[79]

Ubuntu Server is also distributed free of charge. Users can choose to pay for consulting and technical support. Annual support contract with 9x5 business hour support is about $750 per server, and a contract covering 24x7 over a year costs $1,200.[78]

Cloud computing

Ubuntu Server Edition offers technology and resources to make a private or public cloud. Ubuntu Enterprise Cloud (UEC) provides virtualization capability, applications and flexibility to help deploy a cloud within an organization. It consists of the open core Eucalyptus, libvirt, KVM or Xen virtualization technology.[80]

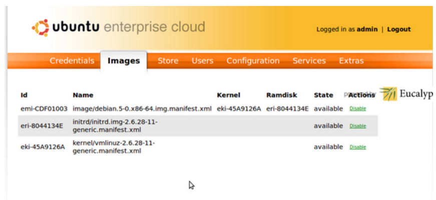

Eucalyptus full

Ubuntu 11.04 added support for OpenStack, with Eucalyptus to OpenStack migration tools to be released by Canonical in Ubuntu Server 11.10.[81] Ubuntu 11.10 is expected to focus on OpenStack as the Ubuntu's preferred IaaS offering though Eucalyptus is also expected to be supported. Another major focus is Ubuntu Orchestra for provisioning, deploying, hosting, managing, and orchestrating enterprise data center infrastructure services, by, with, and for the Ubuntu Server.[82]

Development

The Ubuntu Developer Summit (UDS) is a gathering of software developers which occurs prior to the release of a new public version of Ubuntu.[83]

At the beginning of a new development cycle, Ubuntu developers from around the world gather to help shape and scope the next release of Ubuntu. The summit is open to the public, but it is not a conference, exhibition or other audience-oriented event. Rather, it is an opportunity for Ubuntu developers, who usually collaborate online, to work together in person on specific tasks.

UDS group photo for the Ubuntu 12.04 LTS release

Adoption and reception

Installed base

Chris Kenyon, vice president for OEM at Canonical Ltd., said that because of a lack of registration, any number provided for Ubuntu usage is a "guesstimate".[84] In June 2009 ZDNet reported, "Worldwide, there are 13 million active Ubuntu users with use growing faster than any other distribution".[85] In fall 2011 Canonical estimated that Ubuntu had more than 20 million users worldwide.[86]

In a 2012 online poll for Linux distributions, Ubuntu and its variants received more than 51% of the overall vote, followed by Linux Mint with less than 16%.[10]

User agent counting suggests Ubuntu is the most popular Linux distribution for web clients, generating between 0.5%[13] and 0.72%[11][12] of Internet traffic. However, these measurements are subject to sampling error and bias.

Ubuntu's popularity in web servers is rapidly increasing.[87][88][89][90] As of July 2012, Ubuntu is the third most popular Linux distribution in that market, behind CentOS and Debian (on which Ubuntu is based)[91],

As of 2012, Ubuntu's page on DistroWatch is the second most accessed among Linux distribution pages there, behind the page of Linux Mint.[92][93][94] However, DistroWatch is not an indication of market share or quality.

Publicized large-scale deployments

The public sector has also adopted Ubuntu. As of January 2009, the Ministry of Education and Science of Republic of Macedonia deployed more than 180,000[95] Ubuntu based classroom desktops, and has encouraged every student in the country to use Ubuntu-powered computer workstations[96]; the Spanish school system has 195,000 Ubuntu desktops.[95] The French police, having already started using open source software in 2005 by replacing Microsoft Office with OpenOffice.org, decided to transition to Ubuntu from Windows XP after the release of Windows Vista in 2006.[97] By March 2009, the Gendarmerie Nationale had already switched 5000 workstations to Ubuntu.[97] Based on the success of that transition, it planned to switch 15,000 more over by the end of 2009 and to have switched all 90,000 workstations over by 2015.[97] Lt. Colonel Guimard announced that the move was very easy and allowed for a 70% saving on the IT budget without having to reduce its capabilities.[97]

In 2011, Ubuntu 10.04 was adopted by the Indian Justice system.[98]

The city of Munich, Germany has forked Ubuntu 10.04 LTS and created LiMux for use on the city's computers.[99] Munich expects to have all city computers using LiMux by 2013.

In March 2012, the government of Iceland launched a project to get all public institutions using free and open-source software. Already several government agencies and schools have adopted Ubuntu. The government cited cost savings as a big factor for the decision, and also stated that open source software avoids vendor lock-in. A 12-month project has launched to migrate the biggest public institutions in Iceland to open-source, and help ease the migration for others.[100]

Critical reception

Ubuntu was awarded the Reader Award for best Linux distribution at the 2005 LinuxWorld Conference and Expo in London,[101] received favorable reviews in online and print publications,[102] and has won InfoWorld's 2007 Bossie Award for *Best Open Source Client OS*.[103] In early 2008 *PC World* named Ubuntu the "best all-around Linux distribution available today", though it criticized the lack of an integrated desktop effects manager.[104] Chris DiBona, the program manager for open-source software at Google, said “I think Ubuntu has captured people’s imaginations around the Linux desktop,” and “If there is a hope for the Linux desktop, it would be them”. As of January 2009, almost half of Google’s 20,000 employees used a slightly modified version of Ubuntu.[95]

Ubuntu 10.04LTS has also been criticized for its poor battery life on Laptops and Netbooks, even as OEM on devices such as Asus's eeePC, when compared to Microsoft Windows 7, with Ubuntu having been shown to use between 14-56% more power.[105] Ubuntu's developers have acknowledged and sought to solve the issues of power consumption in the 12.04LTS release.[106]

In 2008, Jamie Hyneman, co-host of the American television series *Mythbusters*, advocated Linux (giving the example of Ubuntu) as a solution to software bloat.[107] Other celebrity users of Ubuntu include science fiction writer and open content proponent Cory Doctorow.[108]

Local Communities (LoCos)

In an effort to reach out to users who are less technical, and to foster a sense of community around the distribution, Local Communities,[109] better known as "LoCos", have been established throughout the world. Originally, each country had one LoCo Team. However, in some areas, most notably the United States, each state or province may establish a team. A LoCo Council approves teams based upon their efforts to aid in either the development or the promotion of Ubuntu.

Vendor support

A number of vendors offer computers with Ubuntu pre-installed, including Dell,[110] Gliese IT, Hasee, Lotus Computers,[111] Ohava Computers,[112] Sharp Corporation[113], System76,[114] and Tesco. System76 PCs come with Ubuntu exclusively. Dell and System76 customers are able to choose between 30-day, three-month, and yearly Ubuntu support plans through Canonical.[115] Dell computers (running Ubuntu 10.04) include extra support for ATI Video Graphics, Dell Wireless, Fingerprint Readers, HDMI, Bluetooth, DVD playback (using LinDVD), and MP3/WMA/WMV.[116] Asus is also selling some Asus Eee PCs with Ubuntu pre-installed and announced that "many more" Eee PC models running Ubuntu for 2011.[117] Vodafone has made available a notebook for the South-African market called "Webbook".[118]

Dell sell computers (initially Inspiron 14R and 15R laptops) pre-loaded with Ubuntu in India and China, with 850 and 350 retail outlets respectively.[119][120]

References

[1] "Ubuntu 12.04 is released" (https://wiki.ubuntu.com/PrecisePangolin). *Ubuntu-Wiki*. 26 April 2012. . Retrieved 26 April 2012.

[2] "Supported Hardware" (https://help.ubuntu.com/12.04/installation-guide/i386/hardware-supported.html#idp3246048). *Official Ubuntu Documentation*. . Retrieved 7 July 2012.

[3] "Ubuntu 11.10 will support ARM processors to take on Red Hat" (http://www.theinquirer.net/inquirer/news/2115929/ubuntu-1110-support-arm-chips-fight-red-hat). *The Inquirer*. 10 October 2011. . Retrieved 20 October 2011.

[4] Paul, Ryan (26 April 2012). "Precise Pangolin rolls out: Ubuntu 12.04 released, introduces Unity HUD" (http://arstechnica.com/business/2012/04/precise-pangolin-rolls-out-ubuntu-1204-released-introduces-unity-hud/). *Ars Technica*. Condé Nast. . Retrieved 7 July 2012.

[5] Larabel, Michael (23 January 2012). "Ubuntu's Already Making Plans For ARM In 2014, 2015" (http://www.phoronix.com/scan.php?page=news_item&px=MTA0NzA). *Phoronix*. . Retrieved 7 July 2012.

[6] Steven J. Vaughan-Nichols (22 August 2011). "Ubuntu Linux bets on the ARM server" (http://www.zdnet.com/blog/open-source/ubuntu-linux-bets-on-the-arm-server/9445). ZDNet. . Retrieved 20 October 2011.

[7] Nelson Mandela (11 January 2006). *The Ubuntu Experience (Nelson Mandela Interview)* (http://www.youtube.com/watch?v=ODQ4WiDsEBQ) (Motion picture). .

[8] ""About Ubuntu" (http://www.ubuntu.com/project/about-ubuntu). Canonical. . Retrieved 27 May 2011.

[9] "About the Name" (https://help.ubuntu.com/10.10/about-ubuntu/C/about-ubuntu-name.html). *Official Ubuntu Documentation*. Canonical. . Retrieved 8 November 2010.

[10] Noyes. "Which Linux Distro Is Fairest of Them All? Ubuntu, Survey Says" (http://www.pcworld.com/businesscenter/article/254516/which_linux_distro_is_fairest_of_them_all_ubuntu_survey_says.html). *PCWorld*. . Retrieved 8 July 2012.

[11] Zachte, Eric. "Wikimedia Traffic Analysis Report - Operating Systems" (http://stats.wikimedia.org/archive/squid_reports/2012-04/SquidReportOperatingSystems.htm). *Wikimedia Traffic Analysis Report*. Wikimedia Foundation. . Retrieved 8 July 2012.

[12] Relph-Knight, Terry (10 February 2012). "A tale of two distros: Ubuntu and Linux Mint" (http://www.zdnet.com/a-tale-of-two-distros-ubuntu-and-linux-mint-3040095012/). *ZDNet*. . Retrieved 8 July 2012.

[13] "Linux (Global marketshare)" (http://getclicky.com/marketshare/global/operating-systems/linux). *Clicky Web Analytics*. Roxr Software Limited. . Retrieved 8 July 2012.

[14] "The project" (http://www.ubuntu.com/project). ubuntu.com. . Retrieved 13 October 2011.

[15] "Time Based Releases" (https://wiki.ubuntu.com/TimeBasedReleases). *Ubuntu Wiki*. . Retrieved 2 June 2008.

[16] "Releases" (https://wiki.ubuntu.com/Releases). *Ubuntu Wiki*. . Retrieved 24 November 2010.

[17] "5yearDesktop" (http://www.canonical.com/content/ubuntu-1204-feature-extended-support-period-desktop-users). Canonical. . Retrieved 21 October 2011.

[18] "Mark Shuttleworth: What about binary compatibility between distributions?" (https://wiki.ubuntu.com/MarkShuttleworth#What about binary compatibility between distributions?). *Ubuntu Team Wiki*. . Retrieved 4 February 2011.

[19] "Website does not reference Debian visibly" (https://bugs.launchpad.net/ubuntu-website/+bug/154274). *Ubuntu in Launchpad*. . Retrieved 31 August 2010.

[20] "Ubuntu vs. Debian, reprise" (http://ianmurdock.com/debian/ubuntu-vs-debian-reprise/). 20 April 2005. . Retrieved 21 October 2007.

[21] Benjamin Mako Hill (8 July 2005). "Announcing Launch of ($10 m) Ubuntu Foundation" (https://lists.ubuntu.com/archives/ubuntu-announce/2005-July/000025.html). . Retrieved 19 August 2008.

[22] "RightScale Adds Full Support for Ubuntu Server to its Cloud Management Platform" (http://www.ubuntu.com/news/rightscale-cloud-management). Canonical. 12 March 2009. . Retrieved 4 February 2011.

[23] Shuttleworth, Mark (31 October 2011). "Ubuntu on phones, tablets, TV's and smart screens everywhere" (http://www.markshuttleworth.com/archives/820). .

- "Canonical outs Ubuntu TV: Brave or stupid?" (http://www.extremetech.com/computing/112601-canonical-outs-ubuntu-tv-brave-or-stupid). Extremetech.com. 9 January 2012. . Retrieved 13 January 2012.
- Schofield, Jack (09 January 2012). "CES 2012: free Ubuntu TV has service and revenue fees" (http://www.zdnet.co.uk/blogs/jacks-blog-10017212/ces-2012-free-ubuntu-tv-has-service-and-revenue-fees-10025161/?s_cid=938). Zdnet.co.uk. . Retrieved 13 January 2012.
- "Canonical Demonstrates Ubuntu TV" (http://www.pcworld.com/businesscenter/article/247571/canonical_demonstrates_ubuntu_tv.html). Pcworld.com. 2012-01-09. . Retrieved 13 January 2012.
- "CES: Canonical shows off Ubuntu TV" (http://www.pcadvisor.co.uk/new-product/ces/3328531/ces-canonical-shows-off-ubuntu-tv/). Pcadvisor.co.uk. 19 December 2011. . Retrieved 13 January 2012.
- "Canonical's Ubuntu TV Surfaces at CES 2012" (http://www.itproportal.com/2012/01/10/canonical-ubuntu-tv-surfaces-ces-2012/). Itproportal.com. . Retrieved 13 January 2012.

[25] "About Ubuntu. Licensing" (http://www.ubuntu.com/project/about-ubuntu/licensing). Canonical. . Retrieved 26 July 2011.

[26] "HoaryGoals" (https://wiki.ubuntu.com/HoaryGoals). *Ubuntu Wiki*. Canonical. . Retrieved 3 September 2010.

[27] "RootSudo" (https://wiki.ubuntu.com/RootSudo?highlight=(sudo)). . Retrieved 19 August 2008.

[28] Katherine Noyes (26 October 2010). "Is Unity the Right Interface for Desktop Ubuntu?" (http://www.networkworld.com/news/2010/102610-is-unity-the-right-interface.html). *PC World*. . Retrieved 28 October 2010.

[29] "Ubuntu Desktop Features" (http://www.ubuntu.com/desktop/features). . Retrieved 3 September 2010.

[30] "Gufw" (https://help.ubuntu.com/community/Gufw). *Ubuntu Documentation*. . Retrieved 4 February 2011.
[31] "GNOME 2.22 Release Notes – Internationalization" (http://library.gnome.org/misc/release-notes/2.22/#rni18). GNOME Foundation. . Retrieved 2 June 2008.
[32] "Ubuntu Wiki CompilerFlags" (https://wiki.ubuntu.com/CompilerFlags). *Ubuntu Wiki*. . Retrieved 31 January 2011.
[33] "Debian Secure by Default" (http://d-sbd.alioth.debian.org/www/). Debian: Secure by Default Project. . Retrieved 31 January 2011.
[34] Larabel, Michael (14 March 2012). "Ubuntu Plans To Drop Non-SMP PowerPC Support" (http://www.phoronix.com/scan.php?page=news_item&px=MTA3MDc). *Phoronix*. . Retrieved 7 July 2012.
[35] "Technical Board Decision" (https://lists.ubuntu.com/archives/ubuntu-announce/2007-February/000098.html). February 2007. . Retrieved 13 June 2008.
[36] "Meeting Minimum Hardware Requirements" (https://help.ubuntu.com/12.04/installation-guide/i386/minimum-hardware-reqts.html). *Official Ubuntu Documentation*. . Retrieved 08 July 2012.
[37] "Installing Ubuntu from the Live CD" (http://www.easy-ubuntu-linux.com/ubuntu-installation-606-2.html). Integrity Enterprises. . Retrieved 19 August 2008.
[38] "Get Ubuntu Desktop Edition Download" (http://www.ubuntu.com/desktop/get-ubuntu/download). . Retrieved 4 February 2011.
[39] "Ubuntu 10.04 LTS (Lucid Lynx)" (http://cdimage.ubuntu.com/ports/releases/lucid/release/). cdimage.ubuntu.com. . Retrieved 24 July 2010.
[40] "Requesting an Ubuntu CD" (https://shipit.ubuntu.com/). . Retrieved 17 April 2009. "A CD of the latest version (9.10 (Karmic Koala)) with no extra cost, but delivery may take up to ten weeks."
[41] "Requesting CDs from ShipIt" (https://shipit.kubuntu.org/). . Retrieved 17 April 2009. "Pre-order CDs of Kubuntu 9.10 (Jaunty Jackalope)."
[42] "Ubuntu 7.04 Adds a Migration Tool" (http://www.easy-ubuntu-linux.com/ubuntu-feisty-installation.html). Integrity Enterprises. . Retrieved 19 August 2008.
[43] "Migration Assistant in Launchpad" (https://launchpad.net/migration-assistant). Launchpad. . Retrieved 2 June 2008.
[44] "Ubuntu 8.10 Persistent Flash Drive Installation" (http://www.pendrivelinux.com/ubuntu-810-persistent-flash-drive-install-from-live-cd/). Pendrivelinux. . Retrieved 5 September 2009.
[45] "Wubi – Ubuntu Installer for Windows" (http://wubi.sourceforge.net/faq.php). wubi.sourceforge.net. . Retrieved 14 May 2008.
[46] "UbuntuBackports" (https://help.ubuntu.com/community/UbuntuBackports). *Ubuntu Documentation*. . Retrieved 24 September 2010.
[47] "StableReleaseUpdates" (https://wiki.ubuntu.com/StableReleaseUpdates). Canonical. . Retrieved 2 April 2009.
[48] "SRU Verification" (https://wiki.ubuntu.com/StableReleaseUpdates#Verification). Canonical. . Retrieved 2 April 2009.
[49] "Application packaging" (http://www.canonical.com/engineering-services/certification/application-packaging). Canonical. . Retrieved 15 August 2010.

- Brian Thomason. "Partner Repository Forum FAQ" (http://ubuntuforums.org/showthread.php?t=723804). . Retrieved 15 August 2010.
- "Desktop support features" (http://www.canonical.com/support/services/support-features). Canonical. . Retrieved 15 August 2010.

[51] "RepositoriesUbuntu" (https://help.ubuntu.com/community/Repositories/Ubuntu#Adding Canonical Partner Repositories). . Retrieved 4 February 2011.
[52] David Planella. "Top 10 Ubuntu Software Centre app downloads for November | Ubuntu App Developer" (http://developer.ubuntu.com/2011/12/top-10-ubuntu-software-centre-app-downloads-for-november/). Developer.ubuntu.com. . Retrieved 29 December 2011.
[53] "Partners/certification/software" (http://www.ubuntu.com/partners/softwareprogramme). Canonical. . Retrieved 16 March 2006.
[54] "Canonical Store" (http://shop.canonical.com/index.php?cPath=19). Shop.canonical.com. . Retrieved 27 May 2011.
[55] "CommonQuestions. Ubuntu Releases and Version Numbers." (https://help.ubuntu.com/community/CommonQuestions#Ubuntu Releases and Version Numbers). *Ubuntu Community Documentation*. Canonical. . Retrieved 24 November 2010.
[56] "DevelopmentCodeNames" (https://wiki.ubuntu.com/DevelopmentCodeNames). ubuntu.com. . Retrieved 8 April 2011.
[57] "UpgradeNotes. General Upgrade Information" (https://help.ubuntu.com/community/UpgradeNotes). *Ubuntu Community Documentation*. Canonical. . Retrieved 26 October 2010.
[58] "Shooting for the Perfect 10.10 with Maverick Meerkat" (http://www.markshuttleworth.com/archives/336). Mark Shuttleworth. 2 April 2010. . Retrieved 8 June 2010.
[59] Mark Shuttleworth (11 May 2010). "ubuntu-marketing: 10.10.10" (https://lists.ubuntu.com/archives/ubuntu-marketing/2010-May/003987.html). *lists.ubuntu.com*. . Retrieved 4 February 2011.

- Mark Shuttleworth (17 August 2010). "N-imal?" (http://www.markshuttleworth.com/archives/478). . Retrieved 17 August 2010.
- "Ubuntu 11.04 Natty Narwhal Alpha released" (http://digitalanalog.in/2010/12/03/ubuntu-11-04-natty-narwhal-alpha-released/). *DigitalAnalog*. 3 December 2010. . Retrieved 4 February 2011.

[61] http://scottlinux.com/ (http://scottlinux.com/2011/03/05/ubuntu-11-04-change-from-unity-to-classic-gnome/?replytocom=76) Ubuntu 11.04 Change From Unity To Classic Gnome
[62] Scott Gilbertson (2011) (http://www.theregister.co.uk/2011/04/01/ubuntu1004_beta_review/) *Natty Narwhal with Unity: Worst Ubuntu beta ever. Nightmare KDE 4 scenario replayed*. The Register.
[63] Lynch, Jim (May 2011). "Ubuntu 11.04" (http://desktoplinuxreviews.com/2011/05/01/ubuntu-11-04/6/). *Desktop Linux Reviews*. . Retrieved 02 May 2011.
[64] (http://ubuntuforums.org/showthread.php?t=1735385) *Unity – I love it, so far.*. Ubuntu Forums.

[65] Humphrey, Benjamin (March 2011). "What's wrong with Unity & how we can fix it" (http://www.omgubuntu.co.uk/2011/03/unity/). *OMG Ubuntu*. . Retrieved 14 March 2011.
[66] "Ubuntu derivatives" (http://www.ubuntu.com/project/derivatives). Canonical. . Retrieved 10 September 2010.
[67] "DerivativeTeam/Derivatives" (https://wiki.ubuntu.com/DerivativeTeam/Derivatives). *Ubuntu Wiki*. . Retrieved 10 September 2010.
[68] "Ubuntu 11.04 (Natty Narwhal). Alternate install CD" (http://releases.ubuntu.com/natty/). *Ubuntu Releases*. . Retrieved 02 July 2011.
[69] "Ubuntu Server Edition" (http://www.ubuntu.com/server). Canonical. . Retrieved 10 September 2010.
[70] "Introducing Ubuntu Business Desktop Remix" (http://www.ubuntu.com/business/desktop/remix). Ubuntu. . Retrieved 11 February 2012.
[71] "Remixing Ubuntu for the Enterprise Desktop" (http://www.markshuttleworth.com/archives/1002). . Retrieved 11 February 2012.
[72] "Ubuntu TV readies for battle with Google and Apple" (http://news.techworld.com/operating-systems/3329027/ubuntu-tv-readies-for-battle-with-google-apple/). . Retrieved 11 February 2012.
[73] "Features and Specs" (http://www.ubuntu.com/tv/features-and-specs). . Retrieved 11 February 2012.
[74] Noyes, Katherine (Feb 21 2012). "Ubuntu for Android Will Bring the Desktop to Your Phone | PCWorld Business Center" (http://www.pcworld.com/businesscenter/article/250356/ubuntu_for_android_will_bring_the_desktop_to_your_phone.html). Pcworld.com. . Retrieved 23 February 2012.
[75] «. "Blog Archive » Ubuntu in your pocket" (http://www.markshuttleworth.com/archives/1011). Mark Shuttleworth. . Retrieved 23 February 2012.
[76] "About Edubuntu" (http://www.edubuntu.org/about). . Retrieved 10 September 2010.

- "Ubuntu Releases" (http://releases.ubuntu.com/). Canonical. . Retrieved 13 September 2010.
- "Ubuntu Mailing List" (https://lists.ubuntu.com/archives/kubuntu-devel/2012-February/005782.html). Canonical. . Retrieved 6 February 2012.

[78] Morgan, Timothy. "Ubuntu Server Primed For The Bigtime" (http://www.theregister.co.uk/2010/04/20/ubuntu_server_10_04/). The Register. . Retrieved 10 April 2011.
[79] "Why use Ubuntu for servers?" (http://www.ubuntu.com/business/server/overview). Ubuntu. . Retrieved 5 April 2011.
[80] "Cloud" (http://www.ubuntu.com/business/server/cloud). Ubuntu. . Retrieved 15 April 2011.

- "Canonical switches to OpenStack for Ubuntu Linux cloud" (http://www.zdnet.com/blog/open-source/canonical-switches-to-openstack-for-ubuntu-linux-cloud/8875). Zdnet.com. 10 May 2011. . Retrieved 10 October 2011.
- Prickett, Timothy (10 May 2011). "Ubuntu eats OpenStack for clouds" (http://www.theregister.co.uk/2011/05/10/canonical_ubuntu_openstack/). Theregister.co.uk. . Retrieved 10 October 2011.
- "Dustin Kirkland of Canonical" (http://www.youtube.com/watch?v=uUv5vwHMPk0). YouTube. 07 June 2011. . Retrieved 13 January 2012.
- "ServerTeam/Orchestra – Ubuntu Wiki" (https://wiki.ubuntu.com/ServerTeam/Orchestra). Wiki.ubuntu.com. 4 January 2012. . Retrieved 13 January 2012.

[83] "Ubuntu Developer Summit" (http://summit.ubuntu.com/). Summit.ubuntu.com. . Retrieved 27 May 2011.
[84] Sean Michael Kerner (7 April 2010). "Ubuntu Claims 12 Million Users as Lucid Linux Desktop Nears" (http://www.linuxplanet.com/linuxplanet/reports/7032/1/). Linux Planet. . Retrieved 7 April 2010.
[85] Christopher Dawson (17 June 2009). "Ubuntu a minor player? Not outside the States." (http://education.zdnet.com/?p=2709). . Retrieved 18 June 2009.
[86] "Canonical launches Ubuntu 11.10 update" (http://www.computerworlduk.com/news/operating-systems/3310952/canonical-launches-ubuntu-1110-update/?page=0,0). Computerworld UK. 14 October 2011. . Retrieved 18 October 2011.
[87] Gelbmann, Matthias (1 February 2011). "Highlights of web technology surveys, February 2011: Ubuntu increases its market share on web servers" (http://w3techs.com/blog/entry/highlights_of_web_technology_surveys_february_2011). w3Techs. . Retrieved 1 July 2011.
[88] "Usage Statistics and Market Share of Ubuntu for Websites, July 2012" (http://w3techs.com/technologies/details/os-ubuntu/all/all). *W3Techs*. . Retrieved 13 July 2012.
[89] Gelbmann (3 January 2012). "Top 10 rising web technologies in 2011" (http://w3techs.com/blog/entry/top_10_rising_web_technologies_in_2011). *W3Techs*. . Retrieved 13 July 2012.
[90] "Ubuntu Usage Statistics" (http://trends.builtwith.com/Server/Ubuntu). builtwith.com. . Retrieved 19 April 2010.
[91] "Usage Statistics and Market Share of Linux for Websites, July 2012" (http://w3techs.com/technologies/details/os-linux/all/all). *W3Techs*. . Retrieved 13 July 2012.
[92] http://www.pcworld.com/article/244950/ubuntu_shows_distrowatch_decline_as_mint_soars.html
[93] http://www.muktware.com/3193/linux-mint-touches-all-time-high-distrowatch-will-ubuntu-recover
[94] http://www.theinquirer.net/inquirer/news/2127579/ubuntu-popularity-falls-linux-mint-flourishes
[95] Ashlee Vance (10 January 2009). "A Software Populist Who Doesn't Do Windows" (http://www.nytimes.com/2009/01/11/business/11ubuntu.html?_r=2&th&emc=th). *The New York Times*. . Retrieved 22 February 2009.
[96] "Every Student in the Republic of Macedonia to Use Ubuntu-Powered Computer Workstations" (http://www.ubuntu.com/news/macedonia-school-computers). Ubuntu. 20 November 2007. . Retrieved 2 December 2010.
[97] Paul, Ryan (11 March 2009). "French police: we saved millions of euros by adopting Ubuntu" (http://arstechnica.com/open-source/news/2009/03/french-police-saves-millions-of-euros-by-adopting-ubuntu.ars). *Ars Technica*. Condé Nast. . Retrieved 2 December 2010.

[98] "India's Justice Sytem(sic) Switches to Ubuntu 10.04 – Softpedia" (http://news.softpedia.com/news/India-s-Justice-Sytem-Switches-to-Ubuntu-10-04-228451.shtml). News.softpedia.com. 18 October 2011. . Retrieved 21 October 2011.

[99] "Landeshauptstadt München - Das Projekt LiMux [City of Munich - The project LiMux (Google translation)]" (http://www.muenchen.de/rathaus/Stadtverwaltung/Direktorium/LiMux.html) (in German). . Retrieved 9 July 2012.

[100] Brown, Mark (23 March 2012). "Icelandic government makes a push for open-source software" (http://www.wired.co.uk/news/archive/2012-03/23/iceland-open-source-software). *Wired UK*. . Retrieved 9 July 2012.

[101] John Masters (June 2005). "LinuxWorld Expo UK 2005" (http://w3.linux-magazine.com/issue/61/Linuxworld_London_UK_2005.pdf) (PDF). *Linux Magazine*. Linux New Media. . Retrieved 19 June 2008.

- Tom Adelstein (19 April 2005). "Linux in Government: Linux Desktop Reviews, Part 6 – Ubuntu" (http://www.linuxjournal.com/article/8253). Linux Journal. . Retrieved 2 December 2010.
- Neil McAllister (January 2008), *Gutsy Gibbon: Desktop Linux OS Made Easy*, **26**, PC World, p. 84

[103] Paul Venenzia (10 September 2007). "Best of open source in platforms and middleware" (http://www.infoworld.com/t/platforms/best-open-source-in-platforms-and-middleware-129). *InfoWorld*. IDG. . Retrieved 2 December 2010.

[104] Robert Strohmeyer (2 June 2008). "Desktop Linux Face-Off: Ubuntu 8.04 vs. Fedora 9" (http://www.pcworld.com/article/146400/desktop_linux_faceoff_ubuntu_804_vs_fedora_9.html). *PC World*. IDG. . Retrieved 19 August 2008.

[105] Farrel, Nick. "New Ubuntu eats more power than Windows 7" (http://www.theinquirer.net/inquirer/news/1610079/new-ubuntu-eats-power-windows). The Inquirer. . Retrieved 5 May 2010.

[106] King, Colin. "Improving Battery Life in Ubuntu Precise 12.04 LTS" (http://smackerelofopinion.blogspot.com/2011/12/improving-battery-life-in-ubuntu.html). . Retrieved 13 December 2011.

[107] Hyneman, Jamie (18 February 2008). "MythBusters: 7 Tech Headaches—and How to Fix Them" (http://www.popularmechanics.com/technology/upgrade/4243994.html?page=3). *Popular Mechanics*. Hearst. . Retrieved 2 December 2010.

[108] Thomas, K.; A. Channelle, J. Sicam (2009). *Beginning Ubuntu Linux*. Apress. p. xxxii. ISBN 978-1-4302-1999-6.

[109] "Ubuntu Local Community Teams" (https://wiki.ubuntu.com/LoCoTeams). .

[110] "Ubuntu on Dell" (http://www.ubuntu.com/dell). Canonical. . Retrieved 24 February 2008.

[111] "Lotus Computers" (http://www.lotuspc.net/us/index.php). Lotus Computers. . Retrieved 16 November 2011.

[112] "Amazing Linux" (http://www.ohava.com/linux). Ohava Computers. . Retrieved 16 November 2011.

[113] "Sharp NetWalker PC-Z1: What you get when you shrink a netbook" (http://www.liliputing.com/2009/08/sharp-netwalker-pc-z1-what-you-get-when-you-shrink-a-netbook.html). Liliputing. . Retrieved 26 March 2010.

[114] "System76: About Ubuntu" (https://www.system76.com/laptops/ubuntu/). System76. . Retrieved O2 May 2012.

[115] "System76 announces servers with Ubuntu 7.10 and Canonical support services" (http://www.ubuntu.com/news/system76-certified-servers). Canonical. . Retrieved 5 March 2008.

[116] "Your Blog " Blog Archive " Dell Upgrades Consumer Linux PCs to Ubuntu 8.04" (http://yourblog.dell.com/2008/07/18/dell-upgrades-consumer-linux-pcs-to-ubuntu-804/). Dell. . Retrieved 13 September 2008.

- "Asus will preload ubuntu linux on three eee pcs" (http://www.theinquirer.net/inquirer/news/2075819/asus-preload-ubuntu-linux-eee-pcs). The Inquirer. . Retrieved 15 June 2011.
- Woods, Ben (3 June 2011). "Asus preloads Eee PC models with Ubuntu | ZDNet UK" (http://www.zdnet.co.uk/blogs/tech-tech-boom-10017860/asus-preloads-eee-pc-models-with-ubuntu-10022619/). Zdnet.co.uk. . Retrieved 15 June 2011.
- "Asus Launching Eee PC Netbooks with Ubuntu" (http://www.tomshardware.com/news/Ubuntu-Eee-PC-Seashell-Netbooks-Canonical,12878.html). Tomshardware.com. . Retrieved 15 June 2011.

- "Vodafone brings ARM and Ubuntu together for South African Webbook" (http://www.engadget.com/2011/10/19/vodafone-brings-arm-and-ubuntu-together-for-south-african-webboo/). Engadget. . Retrieved 20 October 2011.
- "The Ubuntu Powered 'Vodafone Webbook' Launched" (http://www.omgubuntu.co.uk/2011/10/the-ubuntu-powered-vodafone-webbook-launches-tomorrow/). Omgubuntu.co.uk. . Retrieved 20 October 2011.
- "Ubuntu 11.10 Powered Webbook Sells at $190 – Softpedia" (http://news.softpedia.com/news/Ubuntu-11-10-Powered-Webbook-Sells-at-190-229168.shtml). News.softpedia.com. . Retrieved 21 October 2011.

[119] "Dell launch with Ubuntu at retail in India" (http://www.canonical.com/content/canonical-and-dell-announce-today-dell-laptops-pre-loaded-ubuntu-will-be-sold-850-retail-out) (Press release). Canonical. 18 June 2012. . Retrieved 18 June 2012.

[120] Mark Murphy (18 June 2012). "Dell Extends Ubuntu Retail into India" (http://blog.canonical.com/2012/06/18/dell-extends-ubuntu-retail-into-india/). Canonical. . Retrieved 18 June 2012.

Bibliography

- Gagne, Marcel (27 August 2006). *Moving to Ubuntu Linux* (http://www.informit.com/store/product.aspx?isbn=032142722X) (1st ed.). Addison-Wesley Professional. p. 496. ISBN 0-321-42722-X.
- Grant, Rickford; Bull, Phil (7 July 2010). *Ubuntu for Non-Geeks: A Pain-Free, Get-Things-Done Guide* (http://www.nostarch.com/ubuntu4.htm) (4th ed.). No Starch Press. p. 496. ISBN 978-1-59327-257-9.
- Benjamin Mako Hill; Jono Bacon, Corey Burger, Jonathan Jesse, Ivan Krstic (21 August 2006). *The Official Ubuntu Book* (http://www.informit.com/store/product.aspx?isbn=0132435942) (1st ed.). Prentice Hall. p. 448. ISBN 0-13-243594-2.
- Hudson, Andrew; Hudson, Paul; Helmke, Matthew; Troy, Ryan (25 December 2009). *Ubuntu Unleashed 2010 Edition: Covering 9.10 and 10.4* (http://www.informit.com/store/product.aspx?isbn=0672331098) (5th ed.). Sams. p. 864. ISBN 0-672-33109-8.
- Keir, Thomas (15 March 2006). *Beginning Ubuntu Linux: From Novice to Professional* (http://apress.com/book/view/1590596277). Apress. p. 608. ISBN 1-59059-627-7.
- Jonathan Oxer; Kyle Rankin, Bill Childers (14 June 2006). *Ubuntu Hacks: Tips & Tools for Exploring, Using, and Tuning Linux* (http://oreilly.com/catalog/9780596527204) (1st ed.). O'Reilly Media. p. 448. ISBN 0-596-52720-9.
- von Hagen, William (3 January 2007). *Ubuntu Linux Bible* (http://eu.wiley.com/WileyCDA/WileyTitle/productCd-0470038993.html) (1st ed.). Wiley. p. 936. ISBN 0-470-03899-3.

External links

- Official website (http://www.ubuntu.com/)
- *Ubuntu* (http://distrowatch.com/table.php?distribution=Ubuntu) at DistroWatch
- Ubuntu (http://www.dmoz.org/Computers/Software/Operating_Systems/Linux/Distributions/Ubuntu/) at the Open Directory Project

Microsoft Windows

Microsoft Windows

The latest Windows release, Windows 7, showing desktop and start menu	
Company / developer	Microsoft
Programmed in	C, C++ and Assembly language[1]
OS family	Windows 9x, Windows CE and Windows NT
Working state	Publicly released
Source model	Closed source / Shared source
Initial release	November 20, 1985 (as Windows 1.0)
Latest stable release	], Windows Server 2008 R2 NT 6.1 (Build 7601: Service Pack 1)] (February 22, 2011) [2] [±]
Latest unstable release	] NT 6.2 (Build 8400)] (May 31, 2012) [3] [±]
Marketing target	Personal computing
Available language(s)	Multilingual (listing of available Windows 7 language packs [4])
Update method	Windows Update, Windows Anytime Upgrade
Supported platforms	ARM, IA-32, x86-64 and Itanium
Kernel type	Hybrid (Windows NT family), DOS (16-bit Windows and Windows 9x/ME series)
Default user interface	Graphical (Windows Shell)
License	Proprietary commercial software
Official website	[windows.microsoft.com windows.microsoft.com]

Microsoft Windows is a series of graphical interface operating systems developed, marketed, and sold by Microsoft.

Microsoft introduced an operating environment named *Windows* on November 20, 1985 as an add-on to MS-DOS in response to the growing interest in graphical user interfaces (GUIs).[5] Microsoft Windows came to dominate the world's personal computer market with over 90% market share, overtaking Mac OS, which had been introduced in 1984.

The most recent client version of Windows is Windows 7; the most recent server version is Windows Server 2008 R2; the most recent mobile version is Windows Phone 7.5.

Versions

The term *Windows* collectively describes any or all of several generations of Microsoft operating system products. These products are generally categorized as follows:

Early versions

The history of Windows dates back to September 1981, when Chase Bishop, a computer scientist, designed the first model of an electronic device and project "Interface Manager" was started. It was announced in November 1983 (after the Apple Lisa, but before the Macintosh) under the name "Windows", but Windows 1.0 was not released until November 1985.[6] The shell of Windows 1.0 was a program known as the MS-DOS Executive. Other supplied programs were Calculator, Calendar, Cardfile, Clipboard viewer, Clock, Control Panel, Notepad, Paint, Reversi, Terminal, and Write. Windows 1.0 did not allow overlapping windows. Instead all windows were tiled. Only dialog boxes could appear over other windows.

Windows 2.0 was released in October 1987 and featured several improvements to the user interface and memory management.[6] Windows 2.0 allowed application windows to overlap each other and also introduced more sophisticated keyboard shortcuts. It could also make use of expanded memory.

Windows 2.1 was released in two different versions: Windows/386 employed the 386 virtual 8086 mode to multitask several DOS programs, and the paged memory model to emulate expanded memory using available extended memory. Windows/286 (which, despite its name, would run on the 8086) still ran in real mode, but could make use of the high memory area.

The early versions of Windows were often thought of as simply graphical user interfaces, mostly because they ran on top of MS-DOS and used it for file system services.[7] However, even the earliest 16-bit Windows versions already assumed many typical operating system functions; notably, having their own executable file format and providing their own device drivers (timer, graphics, printer, mouse, keyboard and sound) for applications. Unlike MS-DOS, Windows allowed users to execute multiple graphical applications at the same time, through cooperative multitasking. Windows implemented an elaborate, segment-based, software virtual memory scheme, which allowed it to run applications larger than available memory: code segments and resources were swapped in and thrown away when memory became scarce, and data segments moved in memory when a given application had relinquished processor control.

Windows 3.0 and 3.1

Windows 3.0 (1990) and Windows 3.1 (1992) improved the design, mostly because of virtual memory and loadable virtual device drivers (VxDs) that allowed them to share arbitrary devices between multitasked DOS windows. Also, Windows applications could now run in protected mode (when Windows was running in Standard or 386 Enhanced Mode), which gave them access to several megabytes of memory and removed the obligation to participate in the software virtual memory scheme. They still ran inside the same address space, where the segmented memory provided a degree of protection, and multi-tasked cooperatively. For Windows 3.0, Microsoft also rewrote critical operations from C into assembly.

Windows 95, 98, and Me

Windows 95 was released in August 1995, featuring a new object oriented user interface, support for long file names of up to 255 characters, and the ability to automatically detect and configure installed hardware (plug and play). It could natively run 32-bit applications, and featured several technological improvements that increased its stability over Windows 3.1. There were several OEM Service Releases (OSR) of Windows 95, each of which was roughly equivalent to a service pack.

Microsoft's next release was Windows 98 in June 1998. Microsoft released a second version of Windows 98 in May 1999, named Windows 98 Second Edition (often shortened to Windows 98 SE).

In February 2000, Windows 2000 (in the NT family) was released, followed by Windows Me in September 2000 (*Me* standing for *Millennium Edition*). Windows Me updated the core from Windows 98, but adopted some aspects of Windows 2000 and removed the "boot in DOS mode" option. It also added a new feature called System Restore, allowing the user to set the computer's settings back to an earlier date.

Windows Millennium Edition is often confused with Windows 2000 (because of its name), and has been said to be one of the worst operating systems Microsoft ever released.[8]

Windows NT family

The NT family of Windows systems was fashioned and marketed for higher reliability business use. The first release was NT 3.1 (1993), numbered "3.1" to match the consumer Windows version, which was followed by NT 3.5 (1994), NT 3.51 (1995), NT 4.0 (1996), and Windows 2000, which is the last NT-based Windows release that does not include Microsoft Product Activation. Windows NT 4.0 was the first in this line to implement the "Windows 95" user interface (and the first to include Windows 95's built-in 32-bit runtimes).

Microsoft then moved to combine their consumer and business operating systems with Windows XP that was released on October 25, 2001. It came both in home and professional versions (and later niche market versions for tablet PCs and media centers); they also diverged release schedules for server operating systems. Windows Server 2003, released a year and a half after Windows XP, brought Windows Server up to date with Windows XP. After a lengthy development process, Windows Vista was released on November 30, 2006 for volume licensing and January 30, 2007 for consumers. And its server counterpart, Windows Server 2008 was released in early 2008. On July 22, 2009, Windows 7 and Windows Server 2008 R2 were released as RTM (release to manufacturing) while the former was released to the public 3 months later on October 22, 2009.

64-bit operating systems

Windows NT included support for several different platforms before the x86-based personal computer became dominant in the professional world. Versions of NT from 3.1 to 4.0 variously supported PowerPC, DEC Alpha and MIPS R4000, some of which were 64-bit processors, although the operating system treated them as 32-bit processors.

With the introduction of the Intel Itanium architecture (also known as IA-64), Microsoft released new versions of Windows to support it. Itanium versions of Windows XP and Windows Server 2003 were released at the same time as their mainstream x86 (32-bit) counterparts. On April 25, 2005, Microsoft released Windows XP Professional x64 Edition and Windows Server 2003 x64 Editions to support the x86-64 (or *x64* in Microsoft terminology) architecture. Microsoft dropped support for the Itanium version of Windows XP in 2005. Windows Vista was the first end-user version of Windows that Microsoft released simultaneously in x86 and x64 editions. Windows Vista does not support the Itanium architecture. The modern 64-bit Windows family comprises AMD64/Intel64 versions of Windows 7 and Windows Server 2008, in both Itanium and x64 editions. Windows Server 2008 R2 drops the 32-bit version, although Windows 7 does not.

Windows CE

Windows CE (officially known as *Windows Embedded Compact*), is an edition of Windows that runs on minimalistic computers, like satellite navigation systems and some mobile phones. Windows Embedded Compact is based on its own dedicated kernel, dubbed Windows CE kernel. Microsoft licenses Windows CE to OEMs and device makers. The OEMs and device makers can modify and create their own user interfaces and experiences, while Windows CE provides the technical foundation to do so.

Windows CE was used in the Dreamcast along with Sega's own proprietary OS for the console. Windows CE is the core from which Windows Mobile is derived. Microsoft's latest mobile OS, Windows Phone, is based on components from both Windows CE 6.0 R3 and the current Windows CE 7.0.

Windows Embedded Compact is not to be confused with Windows XP Embedded or Windows NT 4.0 Embedded, modular editions of Windows based on Windows NT kernel.

Future of Windows

Windows 8, the successor to Windows 7, is currently in development. Microsoft posted a blog entry in Dutch on October 22, 2010 hinting that Windows 8 would be released in roughly two years.[9] Also, during the pre-Consumer Electronics Show keynote, Microsoft's CEO announced that Windows 8 will also run on ARM CPUs. This Windows version will also be more suitable for tablets and netbooks, featuring a more touch-friendly interface and a Microsoft 'App' market where the ARM versions of the software will download apps. Several new features will also be introduced, such as support for USB 3.0 and the ability to run Windows from USB devices (like USB Hard Disks or USB Flash drives) with Windows To Go.

Microsoft has also announced a version of Windows 8 for ARM devices, named Windows RT.[10]

History

The first version of Microsoft Windows, version 1.0, released in November 1985, lacked a degree of functionality, achieved little popularity and was to compete with Apple's own operating system. Windows 1.0 is not a complete operating system; rather, it extends MS-DOS. Microsoft Windows version 2.0 was released in November 1987 and was slightly more popular than its predecessor. Windows 2.03 (release date January 1988) had changed the OS from tiled windows to overlapping windows. The result of this change led to Apple Computer filing a suit against Microsoft alleging infringement on Apple's copyrights.[11][12]

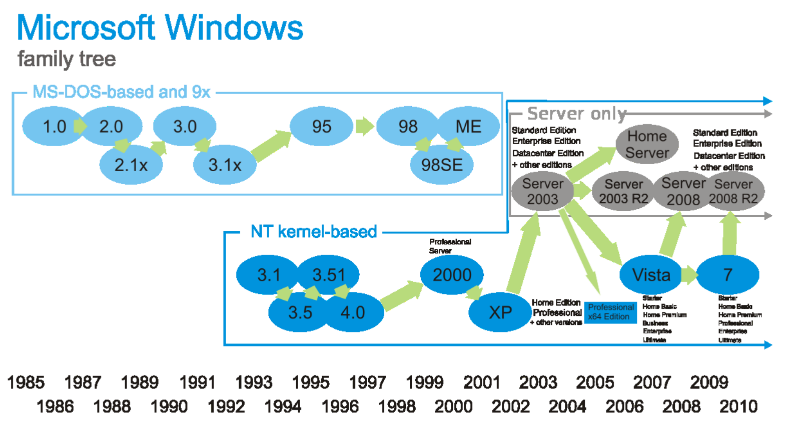

The Windows family tree.

Microsoft Windows version 3.0, released in 1990, was the first Microsoft Windows version to achieve broad commercial success, selling 2 million copies in the first six months.[13][14] It featured improvements to the user interface and to multitasking capabilities. It received a facelift in Windows 3.1, made generally available on March 1, 1992. Windows 3.1 support ended on December 31, 2001.[15]

In July 1993, Microsoft released Windows NT based on a new kernel. Windows NT 3.1 was the first release of Windows NT. NT was considered to be the professional OS and was the first Windows version to utilize preemptive multitasking. Windows NT would later be retooled to also function as a home operating system, with Windows XP.

On August 24, 1995, Microsoft released Windows 95, a new, and major, consumer version that made further changes to the user interface, and also used preemptive multitasking. Windows 95 was designed to replace not only Windows 3.1, but also Windows for Workgroups, and MS-DOS. It was also the first Windows operating system to include object oriented document management and use Plug and Play capabilities. The changes Windows 95 brought to the desktop were revolutionary, as opposed to evolutionary, such as those in Windows 98 and Windows Me. Mainstream support for Windows 95 ended on December 31, 2000 and extended support for Windows 95 ended on

December 31, 2001.[16]

Next in the consumer line was Microsoft Windows 98 released on June 25, 1998. It was followed with the release of Windows 98 Second Edition (Windows 98 SE) in 1999. Mainstream support for Windows 98 ended on June 30, 2002 and extended support for Windows 98 ended on July 11, 2006.[17]

As part of its "professional" line, Microsoft released Windows 2000 in February 2000. During 2004 part of the Source Code for Windows 2000 was leaked onto the Internet. This was bad for Microsoft as the same kernel used in Windows 2000 was used in Windows XP. The consumer version following Windows 98 was Windows Me (Windows Millennium Edition). Released in September 2000, Windows Me implemented a number of new technologies for Microsoft: most notably publicized was "Universal Plug and Play". Windows Me was heavily criticized due to slowness, freezes and hardware problems.

In October 2001, Microsoft released Windows XP, a version built on the Windows NT kernel that also retained the consumer-oriented usability of Windows 95 and its successors. This new version was widely praised in computer magazines.[18] It shipped in two distinct editions, "Home" and "Professional", the former lacking many of the superior security and networking features of the Professional edition. Additionally, the first "Media Center" edition was released in 2002,[19] with an emphasis on support for DVD and TV functionality including program recording and a remote control. Mainstream support for Windows XP ended on April 14, 2009. Extended support will continue until April 8, 2014.[20]

In April 2003, Windows Server 2003 was introduced, replacing the Windows 2000 line of server products with a number of new features and a strong focus on security; this was followed in December 2005 by Windows Server 2003 R2.

On January 30, 2007, Microsoft released Windows Vista. It contains a number of new features, from a redesigned shell and user interface to significant technical changes, with a particular focus on security features. It is available in a number of different editions, and has been subject to some criticism.

On October 22, 2009, Microsoft released Windows 7. Unlike its predecessor, Windows Vista, which introduced a large number of new features, Windows 7 was intended to be a more focused, incremental upgrade to the Windows line, with the goal of being compatible with applications and hardware which Windows Vista was not at the time.[21] Windows 7 has multi-touch support, a redesigned Windows shell with a new taskbar, referred to as the Superbar, a home networking system called HomeGroup,[22] and performance improvements.

On February 29, 2012, Microsoft released Windows 8 Consumer Preview, the beta version of Windows 8, build 8250. For the first time since Windows 95, the Start button is no longer available on the taskbar, though the Start screen is still triggered by clicking the bottom-left corner of the screen and by clicking Start in the Charm. Windows president Steven Sinofsky said more than 100,000 changes had been made since the developer version went public. In the first day of its release, Windows 8 Consumer Preview was downloaded over one million times. Microsoft released the Windows 8 Release Preview, Build 8400 on June 1, 2012. Like the Developer Preview, the Consumer Preview and the Release Preview are both set to expire on January 15, 2013.

Timeline of releases

Timeline of releases				
Release date	**Product name**	**Current Version / Build**	**Notes**	**Last IE**
November 1985	Windows 1.01	1.01	Unsupported	—
November 1987	Windows 2.03	2.03	Unsupported	—
May 1988	Windows 2.10	2.10	Unsupported	—
March 1989	Windows 2.11	2.11	Unsupported	—
May 1990	Windows 3.0	3.0	Unsupported	—
March 1992	Windows 3.1x	3.1	Unsupported	5
October 1992	Windows For Workgroups 3.1	3.1	Unsupported	5
July 1993	Windows NT 3.1	NT 3.1	Unsupported	5
December 1993	Windows For Workgroups 3.11	3.11	Unsupported	5
January 1994	Windows 3.2 (released in Simplified Chinese only)	3.2	Unsupported	5
September 1994	Windows NT 3.5	NT 3.5	Unsupported	5
May 1995	Windows NT 3.51	NT 3.51	Unsupported	5
August 1995	Windows 95	4.0.950	Unsupported[16]	5.5
July 1996	Windows NT 4.0	NT 4.0.1381	Unsupported[23]	6
June 1998	Windows 98	4.10.1998	Unsupported[17]	6
May 1999	Windows 98 SE	4.10.2222	Unsupported[24]	6
February 2000	Windows 2000	NT 5.0.2195	Unsupported[25]	6
September 2000	Windows Me	4.90.3000	Unsupported[26]	6
October 2001	Windows XP	NT 5.1.2600	Extended Support for SP3 until April 8, 2014. (RTM, SP1 and SP2 unsupported).	8
March 2003	Windows XP 64-bit Edition (IA-64)	NT 5.2.3790	Unsupported	6
April 2003	Windows Server 2003	NT 5.2.3790	Extended Support for SP2 until July 14, 2015. (RTM and SP1 unsupported).	8
April 2005	Windows XP Professional x64 Edition (x86-64)	NT 5.2.3790	Extended Support for SP2 until April 8, 2014. (RTM and SP1 unsupported).	8
July 2006	Windows Fundamentals for Legacy PCs	NT 5.1.2600	Current	8
November 2006 (volume licensing) January 2007 (retail)	Windows Vista	NT 6.0.6002	Current (RTM unsupported). Version changed to NT 6.0.6001 with SP1 (February 4, 2008) and to NT 6.0.6002 with SP2 (April 28, 2009).	9
July 2007	Windows Home Server	NT 5.2.4500	Current	8
February 2008	Windows Server 2008	NT 6.0.6002	Current Version changed to NT 6.0.6002 with SP2 (April 28, 2009).	9
October 2009[27]	Windows 7 and Windows Server 2008 R2	NT 6.1.7601	Current Version changed to NT 6.1.7601 with SP1 (February 22, 2011).	9

April 2011	Windows Home Server 2011	NT 6.1.8400	Current	9
31st May 2012	Windows 8	NT 6.2.8400	Pre-Release/Limited Support	10

Usage share

Source	Net Market Share[28]	W3Counter[29]	Global Stats[30]
Date	June 2012	June 2012	June 2012
All versions	92.23%	78.89%	88.43%
Windows 7	41.59%	43.84%	50.2%
Windows XP	43.61%	27.59%	29.91%
Windows Vista	6.72%	7.36%	8.32%
Windows 8	0.18%	0.1%	—
Windows 2000	0.08%	—	—
Windows NT 4.0	0.04%	—	—
Windows 98	0.01%	—	—

Security

Consumer versions of Windows were originally designed for ease-of-use on a single-user PC without a network connection, and did not have security features built in from the outset.[31] However, Windows NT and its successors are designed for security (including on a network) and multi-user PCs, but were not initially designed with Internet security in mind as much, since, when it was first developed in the early 1990s, Internet use was less prevalent.[32]

These design issues combined with programming errors (e.g. buffer overflows) and the popularity of Windows means that it is a frequent target of computer worm and virus writers. In June 2005, Bruce Schneier's *Counterpane Internet Security* reported that it had seen over 1,000 new viruses and worms in the previous six months.[33] In 2005, Kaspersky Lab found around 11,000 malicious programs—viruses, Trojans, back-doors, and exploits written for Windows.[34]

Microsoft releases security patches through its Windows Update service approximately once a month (usually the second Tuesday of the month), although critical updates are made available at shorter intervals when necessary.[35] In versions of Windows after and including Windows 2000 SP3 and Windows XP, updates can be automatically downloaded and installed if the user selects to do so. As a result, Service Pack 2 for Windows XP, as well as Service Pack 1 for Windows Server 2003, were installed by users more quickly than it otherwise might have been.[36]

While the Windows 9x series offered the option of having profiles for multiple users, they had no concept of access privileges, and did not allow concurrent access; and so were not true multi-user operating systems. In addition, they implemented only partial memory protection. They were accordingly widely criticised for lack of security.

The Windows NT series of operating systems, by contrast, are true multi-user, and implement absolute memory protection. However, a lot of the advantages of being a true multi-user operating system were nullified by the fact that, prior to Windows Vista, the first user account created during the setup process was an administrator account, which was also the default for new accounts. Though Windows XP did have limited accounts, the majority of home users did not change to an account type with fewer rights – partially due to the number of programs which unnecessarily required administrator rights – and so most home users ran as administrator all the time.

Windows Vista changes this[37] by introducing a privilege elevation system called User Account Control. When logging in as a standard user, a logon session is created and a token containing only the most basic privileges is assigned. In this way, the new logon session is incapable of making changes that would affect the entire system. When logging in as a user in the Administrators group, two separate tokens are assigned. The first token contains all privileges typically awarded to an administrator, and the second is a restricted token similar to what a standard user would receive. User applications, including the Windows Shell, are then started with the restricted token, resulting in a reduced privilege environment even under an Administrator account. When an application requests higher privileges or "Run as administrator" is clicked, UAC will prompt for confirmation and, if consent is given (including administrator credentials if the account requesting the elevation is not a member of the administrators group), start the process using the unrestricted token.[38]

File permissions

All Windows versions from Windows NT 3 have been based on a file system permission system referred to as AGLP (Accounts, Global, Local, Permissions) AGDLP which in essence where file permissions are applied to the file/folder in the form of a 'local group' which then has other 'global groups' as members. These global groups then hold other groups or users depending on different Windows versions used. This system varies from other vendor products such as Linux and NetWare due to the 'static' allocation of permission being applied directory to the file or folder. However using this process of AGLP/AGDLP/AGUDLP allows a small number of static permissions to be applied and allows for easy changes to the account groups without reapplying the file permissions on the files and folders.

Windows Defender

On January 6, 2005, Microsoft released a Beta version of Microsoft AntiSpyware, based upon the previously released Giant AntiSpyware. On February 14, 2006, Microsoft AntiSpyware became Windows Defender with the release of Beta 2. Windows Defender is a freeware program designed to protect against spyware and other unwanted software. Windows XP and Windows Server 2003 users who have genuine copies of Microsoft Windows can freely download the program from Microsoft's web site, and Windows Defender ships as part of Windows Vista and 7.[39] In Windows 8, Windows Defender and Microsoft Security Essentials have been combined into a single program, named Windows Defender. It is based on Microsoft Security Essentials borrowing its features and user interface. Although it is enabled by default, it can be turned off to use another anti-virus solution.[40]

Third-party analysis

In an article based on a report by Symantec,[41] internetnews.com has described Microsoft Windows as having the "fewest number of patches and the shortest average patch development time of the five operating systems it monitored in the last six months of 2006."[42]

A study conducted by Kevin Mitnick and marketing communications firm Avantgarde in 2004 found that an unprotected and unpatched Windows XP system with Service Pack 1 lasted only 4 minutes on the Internet before it was compromised, and an unprotected and also unpatched Windows Server 2003 system was compromised after being connected to the internet for 8 hours.[43] This study does not apply to Windows XP systems running the Service Pack 2 update (released in late 2004), which vastly improved the security of Windows XP. The computer that was running Windows XP Service Pack 2 was not compromised. The AOL National Cyber Security Alliance Online Safety Study of October 2004 determined that 80% of Windows users were infected by at least one spyware/adware product. Much documentation is available describing how to increase the security of Microsoft Windows products. Typical suggestions include deploying Microsoft Windows behind a hardware or software firewall, running anti-virus and anti-spyware software, and installing patches as they become available through Windows Update.[44]

Emulation software

Emulation allows the use of some Windows applications without using Microsoft Windows. These include:

- Wine – a free and open source software implementation of the Windows API, allowing one to run many Windows applications on x86-based platforms, including Linux and Mac OS X. Wine developers refer to it as a "compatibility layer";[45] and make use of Windows-style APIs to emulate the Windows environment.
  - CrossOver – A Wine package with licensed fonts. Its developers are regular contributors to Wine, and focus on Wine running officially supported applications.
  - Cedega – TransGaming Technologies' proprietary fork of Wine, designed specifically for running games written for Microsoft Windows under Linux. A version of Cedega known as Cider is used by some video game publishers to allow Windows games to run on Mac OS X. Since Wine was licensed under the LGPL, Cedega has been unable to port the improvements made to Wine to their proprietary codebase. Cedega ceased its service in February 2011.
  - Darwine – A bundling of Wine to the PowerPC Macs running OS X by running Wine on top of QEMU. Intel Macs use the same Wine as other *NIX x86 systems.
- ReactOS – An open-source OS that is intended to run the same software as Windows, originally designed to simulate Windows NT 4.0, now aiming at Windows XP and Vista/7 compatibility. It has been in the development stage since 1996.

References

[1] "NT Server Training: Architectural Overview. Lesson 2 – Windows NT System Overview." (http://www.microsoft.com/technet/archive/winntas/training/ntarchitectoview/ntarc_2.mspx). *Microsoft TechNet*. Microsoft. . Retrieved December 9, 2010.

[2] http://en.wikipedia.org/wiki/Template%3Alatest_stable_software_release%2Fmicrosoft_windows?action=edit&preload=Template:LSR/syntax

[3] http://en.wikipedia.org/wiki/Template%3Alatest_preview_software_release%2Fmicrosoft_windows?action=edit&preload=Template:LSR/syntax

[4] http://msdn.microsoft.com/goglobal/ee461121#AvailableLanguagePacks

[5] "The Unusual History of Microsoft Windows" (http://inventors.about.com/od/mstartinventions/a/Windows.htm?rd=1). . Retrieved April 22, 2007.

[6] Petzold

[7] "Windows Evolution" (http://news.soft32.com/windows-evolution_1629.html). Soft32.com News. .

[8] "The 25 Worst Tech Products of All Time" (http://www.pcworld.com/article/125772-2/the_25_worst_tech_products_of_all_time.html). PC World. . Retrieved February 10, 2012.

[9] "Microsoft says Windows 8 roughly two years away" (http://news.cnet.com/8301-13860_3-20020544-56.html?tag=mncol;title). *CNET News* (CNET). October 24, 2010. . Retrieved December 9, 2010.

[10] "Windows Announcing the Windows 8 Editions" (http://windowsteamblog.com/windows/b/bloggingwindows/archive/2012/04/16/announcing-the-windows-8-editions.aspx). The Windows Blog. .

[11] "The Apple vs. Microsoft GUI Lawsuit" (http://lowendmac.com/orchard/06/apple-vs-microsoft.html). 2006. . Retrieved March 12, 2008.

[12] "Apple Computer, Inc. v. MicroSoft Corp., 35 F.3d 1435 (9th Cir. 1994)" (http://home.earthlink.net/~mjohnsen/Technology/Lawsuits/appvsms.html). . Retrieved March 12, 2008.

[13] "Chronology of Personal Computer Software" (http://www.islandnet.com/~kpolsson/compsoft/soft1991.htm). .

[14] "Microsoft Company" (http://www.thocp.net/companies/microsoft/microsoft_company.htm). .

[15] "Windows 3.1 Standard Edition Support Lifecycle" (http://support.microsoft.com/lifecycle/?LN=en-us&p1=3078&x=10&y=11). . Retrieved January 3, 2011.

[16] "Windows 95 Support Lifecycle" (http://support.microsoft.com/lifecycle/?p1=7864). Microsoft. . Retrieved January 3, 2011.

[17] "Windows 98 Standard Edition Support Lifecycle" (http://support.microsoft.com/lifecycle/?p1=6513). Microsoft. . Retrieved January 3, 2011.

[18] David Coursey (August 31, 2001). "Your top Windows XP questions answered! (Part One)" (http://web.archive.org/web/20071219121319/http://review.zdnet.com/4520-6033_16-4206367.html). *ZDNet*. CNET. Archived from the original (http://review.zdnet.com/4520-6033_16-4206367.html) on December 19, 2007. . Retrieved January 3, 2011.

[19] "A Look at Freestyle and Mira" (http://www.winsupersite.com/article/showcase/a-look-at-freestyle-and-mira.aspx). *Paul Thurrott's SuperSite for Windows*. Penton. September 3, 2002. . Retrieved January 3, 2011.

[20] "Windows XP Professional Lifecycle Support" (http://support.microsoft.com/lifecycle/?p1=3223). . Retrieved January 3, 2011.

[21] Mike Nash (October 28, 2008). "Windows 7 Unveiled Today at PDC 2008" (http://windowsteamblog.com/blogs/windows7/archive/2008/10/28/windows-7-unveiled-today-at-pdc-2008.aspx). *Windows Team Blog*. Microsoft. . Retrieved November 11, 2008.

[22] Brandon LeBlanc (October 28, 2008). "How Libraries & HomeGroup Work Together in Windows 7" (http://windowsteamblog.com/blogs/windowsexperience/archive/2008/10/28/how-libraries-amp-homegroup-work-together-in-windows-7.aspx). *Windows Team Blog*. Microsoft. . Retrieved November 11, 2008.

[23] "Windows NT Embedded 4.0 Support Lifecycle" (http://support.microsoft.com/lifecycle/?p1=3185). Microsoft. . Retrieved January 3, 2011.

[24] "Windows 98 Second Edition Support Lifecycle" (http://support.microsoft.com/lifecycle/?p1=6898). Microsoft. . Retrieved January 3, 2011.

[25] "Windows 2000 Professional Edition Support Lifecycle" (http://support.microsoft.com/lifecycle/?p1=3071). Microsoft. . Retrieved January 3, 2011.

[26] "Windows Millennium Edition Support Lifecycle" (http://support.microsoft.com/lifecycle/?p1=6519). Microsoft. . Retrieved January 3, 2011.

[27] "Microsoft Delivers New Wave of Technologies to Help Businesses Thrive in Today's Economy" (http://www.microsoft.com/presspass/press/2009/May09/05-11TechEd09PR.mspx?rss_fdn=Press Releases). Microsoft. May 11, 2009. . Retrieved May 22, 2009.

[28] "Operating System Market Share" (http://marketshare.hitslink.com/operating-system-market-share.aspx?qprid=10&qpcustomd=0&qptimeframe=M&qpsp=161&qpnp=1). *Net Market Share*. Net Applications. June 2012. . Retrieved July 8, 2012.

[29] "Global Web Stats" (http://www.w3counter.com/globalstats.php?year=2012&month=6). *W3Counter*. Awio Web Services. June 2012. . Retrieved July 8, 2012.

[30] "StatCounter Global Stats" (http://gs.statcounter.com/#os-ww-monthly-201206-201206-bar). *Global Stats*. StatCounter. June 2012. . Retrieved July 8, 2012.

[31] Multi-user memory protection was not introduced until Windows NT and XP, and a computer's default user was an administrator until Windows Vista. Source: UACBlog (http://blogs.msdn.com/uac/).

[32] "Telephones and Internet Users by Country, 1990 and 2005" (http://www.infoplease.com/ipa/A0883396.html). Information Please Database. . Retrieved June 9, 2009.

[33] Bruce Schneier (June 15, 2005). "Crypto-Gram Newsletter" (http://www.schneier.com/crypto-gram-0506.html). Counterpane Internet Security, Inc.. . Retrieved April 22, 2007.

[34] Andy Patrizio (April 27, 2006). "Linux Malware On The Rise" (http://www.internetnews.com/dev-news/article.php/3601946). *InternetNews*. QuinStreet. . Retrieved January 3, 2011.

[35] Ryan Naraine (June 8, 2005). "Microsoft's Security Response Center: How Little Patches Are Made" (http://www.eweek.com/c/a/Windows/Microsofts-Security-Response-Center-How-Little-Patches-Are-Made/). *eWeek*. Ziff Davis Enterprise. . Retrieved January 3, 2011.

[36] John Foley (October 20, 2004). "Windows XP SP2 Distribution Surpasses 100 Million" (http://www.informationweek.com/news/security/vulnerabilities/showArticle.jhtml?articleID=50900297). *InformationWeek*. UBM TechWeb. . Retrieved January 3, 2011.

[37] Microsoft describes in detail the steps taken to combat this in a TechNet bulletin. Source: Windows Vista Security and Data Protection Improvements (http://technet.microsoft.com/en-us/windowsvista/aa905073.aspx).

[38] Kenny Kerr (September 29, 2006). "Windows Vista for Developers – Part 4 – User Account Control" (http://weblogs.asp.net/kennykerr/archive/2006/09/29/Windows-Vista-for-Developers-_1320_-Part-4-_1320_-User-Account-Control.aspx). . Retrieved March 15, 2007.

[39] "Windows Vista: Security & Safety" (http://windows.microsoft.com/en-US/windows-vista/products/features/security-safety). Microsoft. . Retrieved April, 16 2012.

[40] "Microsoft Answers: How do I keep Windows 8 Consumer Preview secure from malware?" (http://answers.microsoft.com/en-us/windows/forum/windows_8-security/how-do-i-keep-windows-8-consumer-preview-secure/acd2bfea-ed36-401e-9050-f2fe4212ecf3). Microsoft. . Retrieved April 16, 2012.

[41] "Symantec Internet Security Threat Report Trends for July – December 2006" (http://eval.symantec.com/mktginfo/enterprise/white_papers/ent-whitepaper_internet_security_threat_report_xi_03_2007.en-us.pdf) (PDF). *Internet Security Threat Report Volume XI*. Symantec. March 2007. . Retrieved January 3, 2011.

[42] Andy Patrizio (March 21, 2007). "Report Says Windows Gets The Fastest Repairs" (http://www.internetnews.com/security/article.php/3667201). *InternetNews*. QuinStreet. . Retrieved January 3, 2011.

[43] "Automated "Bots" Overtake PCs Without Firewalls Within 4 Minutes" (http://www.avantgarde.com/ttln113004.html). *Avantgarde*. Avantgarde. November 30, 2004. . Retrieved January 3, 2011.

[44] Richard Rogers (September 21, 2009). "5 Steps To Securing Your Windows XP Home Computer" (http://www.computer-security-news.com/0969/5-steps-to-securing-your-windows-xp-home-computer). *Computer Security News*. Computer Security News. . Retrieved January 3, 2011.

[45] Wine (http://www.winehq.org/)

External links

- Official website (http://windows.microsoft.com/en-US/windows/home)
- Microsoft Developer Network (http://www.msdn.com/)
- Windows Client Developer Resources (http://msdn.microsoft.com/en-au/windows/default.aspx?WT.mc_id=soc-c-au-loc--2010oct)
- Microsoft Windows History Timeline (http://www.microsoft.com/Windows/WinHistoryIntro.mspx)
- Pearson Education, InformIT (http://www.informit.com/articles/article.aspx?p=1358665&rll=1) – History of Microsoft Windows
- Microsoft Windows 7 for Government (http://www.microsoft.com/industry/government/products/windows7/default.aspx)

Windows Server 2008

Windows Server 2008	
Part of the Microsoft Windows family	
Screenshot of Windows Server 2008	
Developer	
Microsoft	
Website	www.microsoft.com/windowsserver2008/en/us/default.aspx [1]
Releases	
Release date	**RTM:** February 4, 2008; **Retail:** February 27, 2008 [info [2]]
Current version	6.0 (Build 6002: Service Pack 2) (July 22, 2009) [info [3]]
Source model	Closed source / Shared source
License	Proprietary commercial software
Kernel type	Hybrid
Update method	Windows Update, Windows Server Update Services, SCCM
Platform support	IA-32, x86-64, Itanium
Preceded by	Windows Server 2003
Succeeded by	Windows Server 2008 R2
Support status	
Mainstream support until 9 July 2013.[4] Extended support until 10 July 2018.	
Further reading	
• Features new to Windows Vista • Features removed from Windows Vista • Development of Windows Vista	

Windows Server 2008 (sometimes abbreviated as "**Win2K8**"[5] or "**W2K8**") is one of Microsoft Windows' server line of operating systems. Released to manufacturing on February 4, 2008, and officially released on February 27, 2008, it is the successor to Windows Server 2003, released nearly five years earlier. A second release, named Windows Server 2008 R2, was released to manufacturing on July 22, 2009.[6] Like Windows Vista and Windows 7, Windows Server 2008 is based on Windows NT 6.x.

History

Originally known as **Windows Server Codename "Longhorn"**, Microsoft chairman Bill Gates announced its official title (Windows Server 2008) during his keynote address at WinHEC 16 May 2007.[7]

Beta 1 was released on 27 July 2005, Beta 2 was announced and released on 23 May 2006 at WinHEC 2006 and Beta 3 was released publicly on 25 April 2007.[8] Release Candidate 0 was released to the general public on 24 September 2007[9] and Release Candidate 1 was released to the general public on 5 December 2007. Windows Server 2008 was released to manufacturing on 4 February 2008 and officially launched on 27 February 2008.[10]

Features

Windows Server 2008 is built from the same code base as Windows Vista; therefore, it shares much of the same architecture and functionality. Since the code base is common, it automatically comes with most of the technical, security, management and administrative features new to Windows Vista such as the rewritten networking stack (native IPv6, native wireless, speed and security improvements); improved image-based installation, deployment and recovery; improved diagnostics, monitoring, event logging and reporting tools; new security features such as BitLocker and ASLR (address space layout randomization); improved Windows Firewall with secure default configuration; .NET Framework 3.0 technologies, specifically Windows Communication Foundation, Microsoft Message Queuing and Windows Workflow Foundation; and the core kernel, memory and file system improvements. Processors and memory devices are modeled as Plug and Play devices, to allow hot-plugging of these devices. This allows the system resources to be partitioned dynamically using *Dynamic Hardware Partitioning*; each partition has its own memory, processor and I/O host bridge devices independent of other partitions.[11]

Server Core

Windows Server 2008 includes a variation of installation called *Server Core*. *Server Core* is a significantly scaled-back installation where no Windows Explorer shell is installed. All configuration and maintenance is done entirely through command line interface windows, or by connecting to the machine remotely using Microsoft Management Console. However, Notepad and some control panel applets, such as Regional Settings, are available.

Server Core does not include the .NET Framework, Internet Explorer, Windows PowerShell or many other features not related to core server features. A *Server Core* machine can be configured for several basic roles: Domain controller/Active Directory Domain Services, ADLDS (ADAM), DNS Server, DHCP Server, file server, print server, Windows Media Server, IIS 7 Web server and Hyper-V virtual server. Server Core can also be used to create a cluster with high availability using failover clustering or network load balancing.

Andrew Mason, a program manager on the Windows Server team, noted that a primary motivation for producing a Server Core variant of Windows Server 2008 was to reduce the attack surface of the operating system, and that about 70% of the security vulnerabilities in Microsoft Windows from the prior five years would not have affected Server Core.[12]

Active Directory roles

Active Directory roles are expanded with identity, certificate, and rights management services. Active Directory, until Windows Server 2003, allowed network administrators to centrally manage connected computers, to set policies for groups of users, and to centrally deploy new applications to multiple computers. This role of Active Directory is being renamed as Active Directory Domain Services (ADDS).[13] A number of other additional services are being introduced, including Active Directory Federation Services (ADFS), Active Directory Lightweight Directory Services (AD LDS), (formerly Active Directory Application Mode, or ADAM), Active Directory Certificate Services (ADCS), and Active Directory Rights Management Services (ADRMS). Identity and certificate services allow administrators to manage user accounts and the digital certificates that allow them to access certain services and systems. Federation management services enable enterprises to share credentials with trusted partners and customers, allowing a consultant to use his company user name and password to log in on a client's network. *Identity Integration Feature Pack* is included as *Active Directory Metadirectory Services*. Each of these services represents a server role.

Failover Clustering

Windows Server 2008 offers high-availability to services and applications through Failover Clustering. Most server features and roles can be kept running with little to no downtime.

In Windows Server 2008 and Windows Server 2008 R2, the way clusters are qualified changed significantly with the introduction of the cluster validation wizard.[14] The cluster validation wizard is a feature that is integrated into failover clustering in Windows Server 2008 and Windows Server 2008 R2. With the cluster validation wizard, an administrator can run a set of focused tests on a collection of servers that are intended to use as nodes in a cluster. This cluster validation process tests the underlying hardware and software directly, and individually, to obtain an accurate assessment of how well failover clustering can be supported on a given configuration.

Note: This feature is only available in Enterprise and Datacenter editions of Windows Server.

Self-healing NTFS

In Windows versions prior to Windows Vista, if the operating system detected corruption in the file system of an NTFS volume, it marked the volume "dirty"; to correct errors on the volume, it had to be taken offline. With self-healing NTFS, an NTFS worker thread is spawned in the background which performs a localized fix-up of damaged data structures, with only the corrupted files/folders remaining unavailable without locking out the entire volume and needing the server to be taken down. The operating system now features S.M.A.R.T. detection techniques to help determine when a hard disk may fail.[15]

Hyper-V

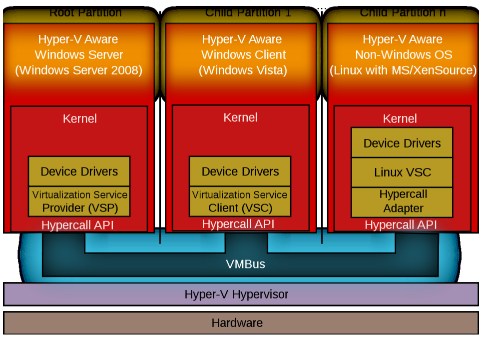

Hyper-V architecture

Hyper-V is hypervisor-based virtualization software, forming a core part of Microsoft's virtualization strategy. It virtualizes servers on an operating system's kernel layer. It can be thought of as partitioning a single physical server into multiple small computational partitions. Hyper-V includes the ability to act as a Xen virtualization hypervisor host allowing Xen-enabled guest operating systems to run virtualized.[16] A beta version of Hyper-V shipped with certain x86-64 editions of Windows Server 2008, prior to Microsoft's release of the final version of Hyper-V on 26 June 2008 as a free download. Also, a standalone version of Hyper-V exists; this version supports only x86-64 architecture.[17] While the IA-32 editions of Windows Server 2008 cannot run or install Hyper-V, they can run the MMC snap-in for managing Hyper-V.

Windows System Resource Manager

Windows System Resource Manager (WSRM) is integrated into Windows Server 2008. It provides resource management and can be used to control the amount of resources a process or a user can use based on business priorities. *Process Matching Criteria*, which is defined by the name, type or owner of the process, enforces restrictions on the resource usage by a process that matches the criteria. CPU time, bandwidth that it can use, number of processors it can be run on, and allocated to a process can be restricted. Restrictions can be set to be imposed only on certain dates as well.

Server Manager

Server Manager is a new roles-based management tool for Windows Server 2008.[18] It is a combination of *Manage Your Server* and *Security Configuration Wizard* from Windows Server 2003. Server Manager is an improvement of the *Configure my server* dialog that launches by default on Windows Server 2003 machines. However, rather than serve only as a starting point to configuring new roles, *Server Manager* gathers together all of the operations users would want to conduct on the server, such as, getting a remote deployment method set up, adding more server roles etc., and provides a consolidated, portal-like view about the status of each role.[19]

Other features

Other new or enhanced features include:

Core OS improvements

- Fully multi-componentized operating system.
- Improved hot patching, a feature that allows non-kernel patches to occur without the need for a reboot.
- Support for being booted from Extensible Firmware Interface (EFI)-compliant firmware on x86-64 systems.
- Dynamic Hardware Partitioning
  - Support for the hot-addition or replacement of processors and memory, on capable hardware.

Active Directory improvements

- Read-only domain controllers (RODCs) in Active Directory, intended for use in branch office or other scenarios where a domain controller may reside in a low physical security environment. The RODC holds a non-writeable copy of Active Directory, and redirects all write attempts to a Full Domain Controller. It replicates all accounts except sensitive ones. In RODC mode, credentials are not cached by default. Moreover, only the replication partner of the RODC needs to run Windows Server 2008. Also, local administrators can log on to the machine to perform maintenance tasks without requiring administrative rights on the domain.
- Restartable Active Directory allows ADDS to be stopped and restarted from the *Management Console* or the command-line without rebooting the domain controller. This reduces downtime for offline operations and reduces overall DC servicing requirements with *Server Core*. ADDS is implemented as a *Domain Controller Service* in Windows Server 2008.

Policy related improvements

- All of the Group Policy improvements from Windows Vista are included. Group Policy Management Console (GPMC) is built-in. The Group Policy objects are indexed for search and can be commented on.[20]
- Policy-based networking with Network Access Protection, improved branch management and enhanced end user collaboration. Policies can be created to ensure greater Quality of Service for certain applications or services that require prioritization of network bandwidth between client and server.
- Granular password settings within a single domain - ability to implement different password policies for administrative accounts on a "group" and "user" basis, instead of a single set of password settings to the whole domain.

Disk management and file storage improvements

- The ability to resize hard disk partitions without stopping the server, even the system partition. This applies only to simple and spanned volumes, not to striped volumes.
- Shadow Copy based block-level backup which supports optical media, network shares and Windows Recovery Environment.
- DFS enhancements - SYSVOL on DFS-R, Read-only Folder Replication Member. There is also support for domain-based DFS namespaces that exceed the previous size recommendation of 5,000 folders with targets in a namespace.[21]
- Several improvements to Failover Clustering (High-availability clusters).[22]
- Internet Storage Naming Server (iSNS) enables central registration, deregistration and queries for iSCSI hard drives.

Protocol and cryptography improvements

- Support for 128- and 256-bit AES encryption for the Kerberos authentication protocol.
- New cryptography (CNG) API which supports elliptic curve cryptography and improved certificate management.
- Secure Socket Tunneling Protocol, a new Microsoft proprietary VPN protocol.
- AuthIP, a Microsoft proprietary extension of the IKE cryptographic protocol used in IPsec VPN networks.
- Server Message Block 2.0 protocol in the new TCP/IP stack provides a number of communication enhancements, including greater performance when connecting to file shares over high-latency links and better security through the use of mutual authentication and message signing.

Improvements due to client-side (Windows Vista) enhancements

- Searching Windows Server 2008 servers from Windows Vista clients delegates the query to the server, which uses the Windows Search technology to search and transfer the results back to the client.
- In a networked environment with a print server running Windows Vista, clients can render print jobs locally before sending them to print servers to reduce the load on the server and increase its availability.
- Event forwarding aggregates and forwards logs of subscribed Windows Vista client computers back to a central console. Event forwarding can be enabled on the client subscribers from the central server directly from the event management console.

Miscellaneous improvements

- Windows Deployment Services replacing Automated Deployment Services and Remote Installation Services. Windows Deployment Services (WDS) support an enhanced multicast feature when deploying operating system images.[23]
- Internet Information Services 7 - Increased security, Robocopy deployment, improved diagnostic tools, delegated administration.
- Windows Internal Database, a variant of SQL Server Express 2005, which serves as a common storage back-end for several other components such as Windows System Resource Manager, Windows SharePoint Services and Windows Server Update Services. It is not intended to be used by third-party applications.
- An optional "Desktop Experience" component provides the same Windows Aero user interface as Windows Vista, both for local users, as well as remote users connecting through Remote Desktop.

Removed features

- The Open Shortest Path First (OSPF) routing protocol component in Routing and Remote Access Service was removed.[24]
- Services for Macintosh, which provided file and print sharing via the now deprecated AppleTalk protocol, has been removed. Services for Macintosh were removed in Windows XP from client operating systems but were available in Windows Server 2003.[24]
- NTBackup is replaced by Windows Server Backup, and no longer supports backing up to tape drives.[25] As a result of NTBackup removal, Exchange Server 2007 does not have volume snapshot backup functionality; however Exchange Server 2007 SP2 adds back an Exchange backup plug-in for Windows Server Backup which restores partial functionality.[26] Windows Small Business Server and Windows Essential Business Server both include this Exchange backup component.[27]
- The POP3 service has been removed from Internet Information Services 7.0.[28] The SMTP (Simple Mail Transfer Protocol) service is not available as a server role in IIS 7.0, it is a server feature managed through IIS 6.0.
- NNTP (Network News Transfer Protocol) is no longer part of Internet Information Services 7.0.

Editions

Most editions of Windows Server 2008 are available in x86-64 and IA-32 versions. Windows Server 2008 for Itanium-based Systems supports IA-64 processors. Microsoft has optimized the IA-64 version for high-workload scenarios like database servers and Line of Business (LOB) applications. As such it is not optimized for use as a file server or media server. Microsoft has announced that Windows Server 2008 is the last 32-bit Windows server operating system.[29] Windows Server 2008 is available in the editions listed below,[30] similar to Windows Server 2003.

- Windows Server 2008 Standard (IA-32 and x86-64)
- Windows Server 2008 Enterprise (IA-32 and x86-64)
- Windows Server 2008 Datacenter (IA-32 and x86-64)
- Windows HPC Server 2008 (Codenamed "Socrates") (replacing Windows Compute Cluster Server 2003)
- Windows Web Server 2008 (IA-32 and x86-64)
- Windows Storage Server 2008 (Codenamed "Magni") (IA-32 and x86-64)
- Windows Small Business Server 2008 (Codenamed "Cougar") (x86-64) for small businesses
- Windows Essential Business Server 2008 (Codenamed "Centro") (x86-64) for medium-sized businesses[31] (Discontinued) [32]
- Windows Server 2008 for Itanium-based Systems
- Windows Server 2008 Foundation (Codenamed "Lima")

Server Core is available in the Web, Standard, Enterprise and Datacenter editions. It is not available in the Itanium edition. *Server Core* is simply an alternate installation option supported by some of the editions, and not a separate edition by itself. Each architecture has a separate installation DVD. The 32-bit version of Windows Server 2008 Standard Edition is available to verified students for free through Microsoft's DreamSpark program.

Service Packs

Microsoft occasionally releases service packs for its Windows operating systems to fix bugs and also add new features.

Service Pack 2

Because Windows Server 2008 is based on the Windows NT 6.0 Service Pack 1 kernel, the RTM release is considered to be Service Pack 1; accordingly, the first service pack is called Service Pack 2. Announced on October 24, 2008,[33] this service pack contains the same changes and improvements as the Windows Vista Service Pack 2, as well as the final release of Hyper-V 1.0, and an approximate 10% reduction in power usage.

The first SP2 beta build was sent out in in October 2008 [34], a public beta arrived in December 2008 [35], and an RC-escrow build was given to testers in January 2009 [36]. Windows Vista and Windows Server 2008 share a single service pack binary, reflecting the fact that their code bases were joined with the release of Server 2008. On May 26, 2009, Service Pack 2 was ready for release. It is now available in Windows Update.

Windows Server 2008 R2

A second release, Windows Server 2008 R2, was released on October 22, 2009.[37] Retail availability began September 14, 2009.[38] Windows Server 2008 R2 reached the RTM milestone on July 22, 2009.[39] Like Windows 7, it is built on Windows NT 6.1. New features include new virtualization features, new Active Directory features, IIS 7.5, and support for 256 logical processors. Support for 32-bit-only processors (IA-32) has been removed. On July 22, 2009, Microsoft officially announced that they had released both Windows Server 2008 R2 and Windows 7 to manufacturing. Windows Server 2008 R2 was generally available for download from MSDN and Technet on August 19 and for retail purchase from October 22, 2009.

System requirements

System requirements for Windows Server 2008 are as follows:

	Minimum for Windows Server 2008[40]	**Recommended for Windows Server 2008[40]**	**Minimum for Windows Server 2008 R2[41]**	**Recommended for Windows Server 2008 R2[41]**
Processor	1 GHz (IA-32) or 1.4 GHz (x86-64) or Intel Itanium 2	2 GHz or faster	1.4 GHz (x86-64 processor) or Intel Itanium 2	2 GHz or faster
Memory	512 MB RAM (may limit performance and some features)	2 GB RAM or higher • Maximum (32-bit systems): 4 GB RAM (Standard) or 64 GB RAM (Enterprise, Datacenter) • Maximum (64-bit systems): 8 GB (Foundation) or 32 GB RAM (Standard) or 2 TB RAM (Enterprise, Datacenter and Itanium-Based Systems)	512 MB RAM	Maximum: 8 GB (Foundation) or 32 GB (Standard) or 2 TB (Enterprise, Datacenter, and Itanium-Based Systems)
Video adapter and monitor	Super VGA (800 × 600)	Super VGA (800 × 600) or higher resolution	Super VGA (800 × 600)	Super VGA (800 × 600) or higher resolution

Hard drive disk free space	• Minimum (Non-Foundation 32-bit systems): 20 GB or greater • Minimum (Non-Foundation 64-bit systems): 32 GB or greater • Foundation: 10 GB or greater.[42] • Computers with more than 16 GB of RAM require more disk space for paging, hibernation, and dump files[41]	40 GB or higher	• 32 GB or greater for editions other than Foundation • Foundation: 10 GB or more • Computers with more than 16 GB of RAM require more disk space for paging, hibernation, and dump files[41]	32 GB or greater for editions other than Foundation • Foundation: 10 GB or more • Computers with more than 16 GB of RAM require more disk space for paging, hibernation, and dump files[41]
Optical drive	DVD-ROM			
Devices	Super VGA (800 × 600) or higher-resolution monitor, keyboard and mouse			

References

[1] http://www.microsoft.com/windowsserver2008/en/us/default.aspx

[2] http://www.microsoft.com/presspass/features/2008/feb08/02-04WS2008.mspx

[3] http://blogs.technet.com/windowsserver/archive/2009/07/22/windows-server-2008-r2-rtm.aspx

[4] Microsoft. "Windows Server 2008 Lifecycle Policy" (http://support.microsoft.com/lifecycle/search/default.aspx?sort=PN&alpha=Windows+Server+2008&Filter=FilterNO). Microsoft. . Retrieved 2009-10-22.

[5] Crowley, Matthew (2010). *Pro Internet Explorer 8 & 9 Development: Developing Powerful Applications for the Next Generation of Ie*. Apress Series. Apress. p. 470. ISBN 978-1-4302-2853-0. "On XP SP3, Vista SP1+, Win2K8, Windows 7, and Win2K8R2, IE opts onto DEP/NX by default"

[6] Windows Server 2008 R2 Reaches the RTM Milestone! (http://blogs.technet.com/windowsserver/archive/2009/07/22/windows-server-2008-r2-rtm.aspx)

[7] Miller, Michael J. (2007-05-15). "Gates at WinHec 2007: Windows Server 2008, Rally, Home Server and More" (http://blog.pcmag.com/blogs/miller/archive/2007/05/15/1883.aspx). *Forward Thinking*. . Retrieved 2007-07-09.

[8] Lowe, David (2007-04-25). "Beta 3 is Go!" (http://blogs.technet.com/windowsserver/archive/2007/04/25/beta-3-is-go.aspx). *Windows Server Division WebLog*. Microsoft. . Retrieved 2007-04-25.

[9] Ralston, Ward (2007-09-24). "Windows Server 2008 Rc0 Released!" (http://blogs.technet.com/windowsserver/archive/2007/09/24/windows-server-2008-rc0-released.aspx). *Windows Server Division WebLog*. Microsoft. . Retrieved 2007-09-24.

[10] Nate Mook. "New Windows Server, Visual Studio, SQL Server to Launch in February" (http://www.betanews.com/article/New_Windows_Server_Visual_Studio_SQL_Server_to_Launch_in_February/1184080223). BetaNews. . Retrieved 2007-07-11. It is also commonly referred to as Vista Server

[11] "Dynamic Hardware Partitioning Architecture" (http://msdn2.microsoft.com/en-us/library/aa938679.aspx). MSDN. . Retrieved 2007-07-23.

[12] "Iain McDonald and Andrew Mason show off the new Windows Server OS" (http://channel9.msdn.com/posts/Duncanma/Iain-McDonald-and-Andrew-Mason-show-off-the-new-Windows-Server-OS/). *Channel 9*. Microsoft. May 24, 2006. . Retrieved 2008-11-01. "18:55"

[13] Hynes, Byron (November 2006). "The Future of Windows: Directory Services in Windows Server 2008" (http://www.microsoft.com/technet/technetmag/issues/2006/11/FutureOfWindows). *TechNet Magazine*. . Retrieved 2007-05-02.

[14] "Failover Cluster Validation Error 80070005 on Windows Server 2008 R2 x64" (http://capitalhead.com/articles/failover-cluster-validation-error-80070005-on-windows-server-2008-r2-x64.aspx). *Failover Cluster Validation Error 80070005 on Windows Server 2008 R2 x64*. 1900-1-0. . Retrieved 2009-11-05.

[15] Loveall, John (2006). "Storage improvements in Windows Vista and Windows Server 2008" (http://download.microsoft.com/download/5/b/9/5b97017b-e28a-4bae-ba48-174cf47d23cd/STO123_WH06.ppt) (PowerPoint). Microsoft Corporation. . Retrieved 2007-07-09.

[16] "Benchmarking Hyper-V on Windows Server 2008 R2 x64" (http://capitalhead.com/articles/benchmarking-hyper-v-on-windows-server-2008-r2-x64.aspx). *Benchmarking Hyper-V on Windows Server 2008 R2 x64*. 2010-01-20. . Retrieved 2010-01-28.

[17] "Microsoft Extends Virtualization Strategy, Outlines Product Road Map" (http://www.microsoft.com/presspass/features/2006/may06/05-22Virtualization.mspx). Microsoft. 2006-05-22. . Retrieved 2007-07-09.

[18] "Server Manager" (http://technet2.microsoft.com/windowsserver2008/en/library/b3274a34-7574-4ea6-aec0-e05ba297481e1033.mspx?mfr=true). *Windows Server 2008 Technical Library*. Microsoft TechNet. 2007-06-25. . Retrieved 2007-05-02.

[19] "Unexpected error refreshing Server Manager-0x800706BE and 1601 on Window Server 2008 R2" (http://capitalhead.com/articles/unexpected-error-refreshing-server-manager-errors-0x800706be-and-1601-on-window-server-2008-r2.aspx). *Unexpected error refreshing Server Manager-0x800706BE and 1601 on Window Server 2008 R2*. . Retrieved 2010-11-05.
[20] Keith Ward. "Top 10 Overlooked Windows Server 2008 Features, Part 2" (http://reddevnews.com/news/article.aspx?editorialsid=9130#1). Redmond Developer News. . Retrieved 2007-10-09.
[21] Breaking the 5K Folder "Barrier" in Domain-Based Namespaces: Filing Cabinet blog (http://blogs.technet.com/filecab/archive/2007/07/24/new-in-windows-server-2008-breaking-the-5k-folder-barrier-in-domain-based-namespaces.aspx)
[22] "Failover Clustering with Windows Server 2008 including http://en.wikipedia.org/wiki/Cluster_shared_volumes" (http://www.microsoft.com/windowsserver2008/failover-clusters.mspx). Microsoft. 2007-01-17. . Retrieved 2007-07-09.
[23] Multicasting OS deployments with Windows Server 2008 (http://blogs.technet.com/kevinsul_blog/archive/2007/08/28/multicast-os-deployments-at-last-and-in-the-box-with-windows-server-2008.aspx)
[24] Removed technologies in Routing and Remote Access in Windows Server 2008 (http://technet.microsoft.com/en-us/library/cc730711(WS.10).aspx#BKMK_6)
[25] Windows Server Backup Step-by-Step Guide for Windows Server 2008 (http://technet.microsoft.com/en-us/library/cc770266(WS.10).aspx)
[26] Microsoft Exchange Team Blog - Exchange Server 2007 Service Pack 2 (http://msexchangeteam.com/archive/2009/05/11/451281.aspx)
[27] To Backup or Not to Backup? Yes! To backup!! (http://msexchangeteam.com/archive/2008/06/18/449031.aspx)
[28] IIS 7.0 Protocols (http://technet.microsoft.com/en-us/library/cc268242.aspx)
[29] Heaton, Alex (2007-05-18). "On 64-bit and Windows Client" (http://windowsvistablog.com/blogs/windowsvista/archive/2007/05/18/on-64-bit-and-windows-client.aspx). *Windows Vista Team Blog*. . Retrieved 2007-07-09.
[30] "Windows Server 2008 Product Editions" (http://www.microsoft.com/windowsserver2008/evaleditions.mspx). Microsoft. 2007-04-25. . Retrieved 2007-07-09.
[31] Announcing Windows Essential Business Server (http://blogs.msdn.com/mssmallbiz/archive/2007/11/07/5955270.aspx)
[32] http://technet.microsoft.com/en-us/ebs/cc164281
[33] Justin Graham (October 24, 2008). "Windows Server 2008 Service Pack 2 beta" (http://blogs.technet.com/windowsserver/archive/2008/10/24/windows-server-2008-service-pack-2-beta.aspx). Microsoft. . Retrieved 2008-10-29.
[34] http://arstechnica.com/microsoft/news/2008/10/first-vista-sp2-and-server-2008-sp2-beta-build-sent-out.ars
[35] http://arstechnica.com/microsoft/news/2008/12/vista-sp2-and-server-2008-sp2-public-betas-arrive.ars
[36] http://arstechnica.com/microsoft/news/2009/01/testers-get-rc-escrow-build-of-sp2-for-vista-and-server-2008.ars
[37] "Windows Server 2008 R2: Getting Started" (http://www.microsoft.com/windowsserver2008/en/us/R2.aspx). Microsoft. . Retrieved 2009-07-13.
[38] http://blogs.technet.com/windowsserver/archive/2009/07/22/when-to-expect-windows-server-2008-r2-rtm.aspx
[39] http://blogs.technet.com/windowsserver/archive/2009/07/22/windows-server-2008-r2-rtm.aspx
[40] "Windows Server 2008 System Requirements" (http://msdn2.microsoft.com/en-us/windowsserver/cc196364.aspx). 31 March 2008. . Retrieved 2008-03-31.
[41] http://www.microsoft.com/windowsserver2008/en/us/system-requirements.aspx
[42] http://www.microsoft.com/windowsserver2008/en/us/WS08-system-requirements.aspx

External links

Microsoft

- Product Evaluation site for Windows Server 2008 R2 (http://www.microsoft.com/windowsserver2008/en/us/default.aspx)
- Microsoft TechCenter for Windows Server 2008 R2 and Windows Server 2008 (http://technet.microsoft.com/en-us/windowsserver/bb310558.aspx)
- MSDN Developer Center for Windows Server (http://msdn.microsoft.com/en-us/windowsserver/default.aspx)
- What's New in Networking for Windows Server 2008 R2 and Windows 7 (http://technet.microsoft.com/en-us/library/dd391869(WS.10).aspx) — lengthy article from Microsoft Technet covering new networking features in detail
- Changes in Functionality from Windows Server 2003 with SP1 to Windows Server 2008 (http://go.microsoft.com/fwlink/?LinkId=87080)

- Windows Server 2008: Compare Server Core Installation Options (http://www.microsoft.com/windowsserver2008/en/us/compare-core-installation.aspx)
- Windows Server Performance Team Blog (http://blogs.technet.com/winserverperformance/)
- List of server applications that are supported, will be supported, or are incompatible (http://support.microsoft.com/kb/948680)
- Official System Requirements page (http://www.microsoft.com/windowsserver2008/en/us/system-requirements.aspx)
- Windows Server 2008 for government (http://www.microsoft.com/industry/government/products/server/windows2008.aspx)

Miscellaneous

- Early review from Network World (http://www.networkworld.com/reviews/2008/022108-windows-2008-server-test.html)
- Screenshots of Server Components from Latest Longhorn Builds (http://neosmart.net/gallery/v/os/Longhorn+Server/Beta3/)
- Windows Server 2008 Build and Revision Numbers (http://capitalhead.com/articles/how-to-find-build-and-revision-number-of-windows-vista-or-windows-server-2008-installed.aspx)
- Microsoft lays out server road map (http://news.com.com/2100-7343_3-5211424.html)

Windows 7

Windows 7	
Part of the Microsoft Windows family	
Screenshot of Windows 7 Ultimate	
Developer	
Microsoft Corporation	
Website	Official Website [1]
Releases	
Release date	**RTM version:** July 22, 2009 **Retail version:** October 22, 2009 [info [2]]
Current version	6.1[3] (Build 7601: Service Pack 1)[4] (February 22, 2011) [info [2]]
Source model	Closed source / Shared source
License	Proprietary commercial software
Kernel type	Hybrid
Update method	Windows Update
Platform support	IA-32 and x86-64
Preceded by	Windows Vista[5]
Succeeded by	Windows 8 (*to be released on October 26, 2012*)[6]
Support status	
Mainstream support until January 13, 2015.[7] Extended support until January 14, 2020.	
Further reading	
• Development of Windows 7 • Features new to Windows 7 • List of features removed in Windows 7 • Windows 7 editions	

Windows 7 is the current release of Microsoft Windows, a series of operating systems produced by Microsoft for use on personal computers, including home and business desktops, laptops, netbooks, tablet PCs, and media center PCs.[8] Windows 7 was released to manufacturing on July 22, 2009,[9] and reached general retail availability worldwide on October 22, 2009,[10] less than three years after the release of its predecessor, Windows Vista. Windows 7's server counterpart, Windows Server 2008 R2, was released at the same time.

Unlike Windows Vista, which introduced a large number of new features, Windows 7 was intended to be a more focused, incremental upgrade to the Windows line, with the goal of being compatible with applications and hardware with which Windows Vista was already compatible.[11] Presentations given by Microsoft in 2008 focused on multi-touch support, a redesigned Windows shell with a new taskbar, referred to as the Superbar, a home networking system called HomeGroup,[12] and performance improvements. Some standard applications that have been included with prior releases of Microsoft Windows, including Windows Calendar, Windows Mail, Windows Movie Maker, and Windows Photo Gallery, are not included in Windows 7;[13][14] most are instead offered separately at no charge as part of the Windows Live Essentials suite.[15]

Features

New and changed features

Windows 7 includes a number of new features, such as advances in touch and handwriting recognition, support for virtual hard disks, improved performance on multi-core processors,[16][17][18][19] improved boot performance, DirectAccess, and kernel improvements. Windows 7 adds support for systems using multiple heterogeneous graphics cards from different vendors (Heterogeneous Multi-adapter), a new version of Windows Media Center,[20] a Gadget for Windows Media Center, improved media features, the XPS Essentials Pack and Windows PowerShell being included, and a redesigned Calculator with multiline capabilities including *Programmer* and *Statistics* modes along with unit conversion for length, weight, temperature, and several others. Many new items have been added to the Control Panel, including ClearType Text Tuner, Display Color Calibration Wizard, Gadgets, Recovery, Troubleshooting, Workspaces Center, Location and Other Sensors, Credential Manager, Biometric Devices, System Icons, and Display.[21] Windows Security Center has been renamed to Windows Action Center (Windows Health Center and Windows Solution Center in earlier builds), which encompasses both security and maintenance of the computer. ReadyBoost on 32-bit editions now supports up to 256 gigabytes of extra allocation. The default setting for User Account Control in Windows 7 has been criticized for allowing untrusted software to be launched with elevated privileges without a prompt by exploiting a trusted application.[22] Microsoft's Windows kernel engineer Mark Russinovich acknowledged the problem, but noted that malware can also compromise a system when users agree to a prompt.[23] Windows 7 also supports images in RAW image format through the addition of Windows Imaging Component-enabled image decoders, which enables raw image thumbnails, previewing and metadata display in Windows Explorer, plus full-size viewing and slideshows in Windows Photo Viewer and Windows Media Center.[24]

The taskbar has seen the biggest visual changes, where the Quick Launch toolbar has been replaced with the ability to pin applications to the taskbar. Buttons for pinned applications are integrated with the task buttons. These buttons also enable the *Jump Lists* feature to allow easy access to common tasks.[25] The revamped taskbar also allows the reordering of taskbar buttons. To the far right of the system clock is a small rectangular button that serves as the *Show desktop* icon. This button is part of the new feature in Windows 7 called *Aero Peek*. Hovering over this button makes all visible windows transparent for a quick look at the desktop.[26] In touch-enabled displays such as touch screens, tablet PCs, etc., this button is slightly wider to accommodate being pressed with a finger.[27] Clicking this button minimizes all windows, and clicking it a second time restores them. Additionally, there is a feature named *Aero Snap*, that automatically maximizes a window when it is dragged to the top of the screen.[28] Dragging windows to the left/right edges of the screen allows users to snap documents or files on either side of the screen for comparison between windows, such that the windows vertically take up half the screen. When a user moves windows that were maximized using Aero Snap, the system restores their previous state automatically. This functionality is also accomplished with keyboard shortcuts. Unlike in Windows Vista, window borders and the taskbar do not turn opaque when a window is maximized with Windows Aero applied. Instead, they remain translucent.

For developers, Windows 7 includes a new networking API with support for building SOAP-based web services in native code (as opposed to .NET-based WCF web services),[29] new features to shorten application install times, reduced UAC prompts, simplified development of installation packages,[30] and improved globalization support through a new Extended Linguistic Services API.[31] At WinHEC 2008 Microsoft announced that color depths of 30-bit and 48-bit would be supported in Windows 7 along with the wide color gamut scRGB (which for HDMI 1.3 can be converted and output as xvYCC). The video modes supported in Windows 7 are 16-bit sRGB, 24-bit sRGB, 30-bit sRGB, 30-bit with extended color gamut sRGB, and 48-bit scRGB.[32][33] Microsoft has also implemented better support for solid-state drives,[34] including the new TRIM command, and Windows 7 is able to identify a solid-state drive uniquely. Microsoft is planning to support USB 3.0 in a subsequent patch, support not being included in the initial release due to delays in the finalization of the standard.[35]

Internet Spades, Internet Backgammon and Internet Checkers, which were removed from Windows Vista, were restored in Windows 7. Windows 7 includes Internet Explorer 8 and Windows Media Player 12. Users are also able to disable many more Windows components than was possible in Windows Vista. New additions to this list of components include Internet Explorer, Windows Media Player, Windows Media Center, Windows Search, and the Windows Gadget Platform.[36] Windows 7 includes 13 additional sound schemes, titled Afternoon, Calligraphy, Characters, Cityscape, Delta, Festival, Garden, Heritage, Landscape, Quirky, Raga, Savanna, and Sonata.[37] A new version of Microsoft Virtual PC, newly renamed as Windows Virtual PC was made available for Windows 7 Professional, Enterprise, and Ultimate editions.[38] It allows multiple Windows environments, including *Windows XP Mode*, to run on the same machine. Windows XP Mode runs Windows XP in a virtual machine and redirects displayed applications running in Windows XP to the Windows 7 desktop.[39] Furthermore, Windows 7 supports the mounting of a virtual hard disk (VHD) as a normal data storage, and the bootloader delivered with Windows 7 can boot the Windows system from a VHD; however, this ability is only available in the Enterprise and Ultimate editions.[40] The Remote Desktop Protocol (RDP) of Windows 7 is also enhanced to support real-time multimedia application including video playback and 3D games, thus allowing use of DirectX 10 in remote desktop environments.[41] The three application limit, previously present in the Windows Vista Starter Edition, has been removed from Windows 7.[42] Recommendations for Windows 7 is to be on Windows Vista (Longhorn) before upgrading to any version of Windows 7, and to have 16GB on the hard drive.

Removed features

Certain capabilities and programs that were a part of Windows Vista are no longer present or have been changed, resulting in the removal of certain functionalities. These include the classic Start Menu user interface, some taskbar features, Windows Explorer features, Windows Media Player features, Windows Ultimate Extras and InkBall. Four applications bundled with Windows Vista – Windows Photo Gallery, Windows Movie Maker, Windows Calendar[43] and Windows Mail – are not included with Windows 7, but applications with close functionality are instead available for free in a separate package called Windows Live Essentials which can be downloaded on the Microsoft website. Although Windows Ultimate Extras was removed, many of the extras can be installed separately.[44] Most popular extras were Microsoft Texas Hold 'em, Microsoft Tinker, and Windows DreamScene.[45] InkBall may also be installed into Windows 7.

Development

Originally, a version of Windows codenamed *Blackcomb* was planned as the successor to Windows XP (codename Whistler) and Windows Server 2003. Major features were planned for Blackcomb, including an emphasis on searching and querying data and an advanced storage system named WinFS to enable such scenarios. However, an interim, minor release, codenamed "Longhorn," was announced for 2003, delaying the development of Blackcomb.[46] By the middle of 2003, however, Longhorn had acquired some of the features originally intended for Blackcomb. After three major viruses exploited flaws in Windows operating systems within a short time period in 2003, Microsoft changed its development priorities, putting some of Longhorn's major development work on hold while developing new service packs for Windows XP and Windows Server 2003. Development of Longhorn (Windows Vista) was also restarted, and thus delayed, in August 2004. A number of features were cut from Longhorn.[47]

Blackcomb was renamed *Vienna* in early 2006[48] and again renamed *Windows 7* in 2007.[49] In 2008, it was announced that *Windows 7* would also be the official name of the operating system.[50][51] There has been some confusion over naming the product Windows 7,[52] while versioning it as 6.1 to indicate its similar build to Vista and increase compatibility with applications that only check major version numbers, similar to Windows 2000 and Windows XP both having 5.x version numbers.[53]

The first external release to select Microsoft partners came in January 2008 with Milestone 1, build 6519.[54] At PDC 2008, Microsoft demonstrated Windows 7 with its reworked taskbar. Copies of Windows 7 build 6801 were distributed at the end of the conference; however, the demonstrated taskbar was disabled in this build.

On December 27, 2008, the Windows 7 Beta was leaked onto the Internet via BitTorrent.[55] According to a performance test by ZDNet,[56] Windows 7 Beta beat both Windows XP and Vista in several key areas; including boot and shutdown time and working with files, such as loading documents. Other areas did not beat XP; including PC Pro benchmarks for typical office activities and video editing, which remain identical to Vista and slower than XP.[57] On January 7, 2009, the 64-bit version of the Windows 7 Beta (build 7000) was leaked onto the web, with some torrents being infected with a trojan.[58][59] At CES 2009, Microsoft CEO Steve Ballmer announced the Windows 7 Beta, build 7000, had been made available for download to MSDN and TechNet subscribers in the format of an ISO image.[60] The Beta was to be publicly released January 9, 2009, and Microsoft initially planned for the download to be made available to 2.5 million people on this date. However, access to the downloads was delayed because of high traffic.[61] The download limit was also extended, initially until January 24, then again to February 10. People who did not complete downloading the beta had two extra days to complete the download. After February 12, unfinished downloads became unable to complete. Users could still obtain product keys from Microsoft to activate their copies of Windows 7 Beta, which expired on August 1, 2009.

The release candidate, build 7100, became available for MSDN and TechNet subscribers and Connect Program participants on April 30, 2009. On May 5, 2009 it became available to the general public, although it had also been leaked onto the Internet via BitTorrent.[62] The release candidate was available in five languages and expired on June 1, 2010, with shutdowns every two hours starting March 1, 2010.[63] Microsoft stated that Windows 7 would be released to the general public on October 22, 2009. Microsoft released Windows 7 to MSDN and Technet subscribers on August 6, 2009, at 10:00 am PDT.[64] Microsoft announced that Windows 7, along with Windows Server 2008 R2, was released to manufacturing on July 22, 2009. Windows 7 RTM is build 7600.16385.090713-1255, which was compiled on July 13, 2009, and was declared the final RTM build after passing all Microsoft's tests internally.[9]

An estimated 1000 developers worked on Windows 7. These were broadly divided into "core operating system" and "Windows client experience", in turn organized into 25 teams of around 40 developers on average.[65]

Goals

Bill Gates, in an interview with *Newsweek*, suggested that this version of Windows would be more "user-centric".[66] Gates later said that Windows 7 would also focus on performance improvements.[67] Steven Sinofsky later expanded on this point, explaining in the *Engineering Windows 7* blog that the company was using a variety of new tracing tools to measure the performance of many areas of the operating system on an ongoing basis, to help locate inefficient code paths and to help prevent performance regressions.[68]

Senior Vice President Bill Veghte stated that Windows Vista users migrating to Windows 7 would not find the kind of device compatibility issues they encountered migrating from Windows XP.[69] Speaking about Windows 7 on October 16, 2008, Microsoft CEO Steve Ballmer confirmed compatibility between Windows Vista and Windows 7, indicating that Windows 7 would be a refined version of Windows Vista.[70]

Antitrust regulatory attention

As with other Microsoft operating systems, Windows 7 is being studied by United States federal regulators who oversee the company's operations following the 2001 *United States v. Microsoft* settlement. According to status reports filed, the three-member panel began assessing prototypes of the new operating system in February 2008. Michael Gartenberg, an analyst at Jupiter Research said that, "[Microsoft's] challenge for Windows 7 will be how can they continue to add features that consumers will want that also don't run afoul of regulators."[71]

In order to comply with European antitrust regulations, Microsoft has proposed the use of a "ballot" screen, allowing users to download a competing browser, thus removing the need for a version of Windows completely without Internet Explorer, as previously planned.[72] In response to criticism involving Windows 7 E and concerns from manufacturers about possible consumer confusion if a version of Windows 7 with Internet Explorer were shipped later after one without Internet Explorer, Microsoft announced that it would scrap the separate version for Europe and ship the standard upgrade and full packages worldwide.[73]

As with the previous version of Windows, an N version, which does not come with Windows Media Player, has been released in Europe, but only for sale directly from Microsoft sales websites and selected others.[74]

Reception

In July 2009, in only eight hours, pre-orders of Windows 7 at amazon.co.uk surpassed the demand which Windows Vista had had in its first 17 weeks.[75] It became the highest-grossing pre-order in Amazon's history, surpassing sales of the previous record holder, the seventh Harry Potter book.[76] After 36 hours, 64-bit versions of Windows 7 Professional and Ultimate editions sold out in Japan.[77] Two weeks after its release its market share had surpassed that of Snow Leopard, released two months previously as the most recent update to Apple's Mac OS X operating system.[78][79] According to Net Applications, Windows 7 reached a 4% market share in less than three weeks. (In comparison, it took Windows Vista seven months to reach the same mark.[80])

On March 4, 2010, Microsoft announced that it had sold more than 90 million Windows 7 licenses.[81] By April 23, 2010, Windows 7 had sold more than 100 million copies in six months, which made it Microsoft's fastest-selling operating-system.[82][83] As of June 23, 2010, Windows 7 has sold 150 million copies which made it the fastest selling operating system in history with seven copies sold every second.[83][84] Based on worldwide data taken during June 2010 from Windows Update 46% of Windows 7 PCs run the 64-bit edition of Windows 7.[85] According to Stephen Baker of the NPD Group during April 2010 in the United States 77% of PCs sold at retail were pre-installed with the 64-bit edition of Windows 7.[85][86] As of July 22, 2010, Windows 7 had sold 175 million copies.[87] On October 21, 2010, Microsoft announced that more than 240 million copies of Windows 7 had been sold.[88] Three months later, on January 27, 2011, Microsoft announced total sales of 300 million copies of Windows 7.[89] On July 12, 2011, the sales figure was refined to over 400 million end-user licenses and business installations.[90] As of January 19, 2012, over 525 million copies have been sold.[91]

Reviews of Windows 7 have been mostly positive, noting the increased usability and functionality when compared to its predecessor, Windows Vista. CNET gave Windows 7 Home Premium a rating of 4.5 out of 5 stars,[92] stating that it "is more than what Vista should have been, [and] it's where Microsoft needed to go". PC Magazine rated it a 4 out of 5 saying that Windows 7 is a "big improvement" over Windows Vista, with fewer compatibility problems, a retooled taskbar, simpler home networking and faster start-up.[93] Maximum PC gave Windows 7 a rating of 9 out of 10 and called Windows 7 a "massive leap forward" in usability and security, and praised the new Taskbar as "worth the price of admission alone".[94] PC World called Windows 7 a "worthy successor" to Windows XP and said that speed benchmarks showed Windows 7 to be slightly faster than Windows Vista.[95] PC World also named Windows 7 one of the best products of the year.[96] In its review of Windows 7, Engadget said that Microsoft had taken a "strong step forward" with Windows 7 and reported that speed is one of Windows 7's major selling points – particularly for the netbook sets.[97] LAPTOP Magazine gave Windows 7 a rating of 4 out of 5 stars and said that Windows 7 makes computing more intuitive, offered better overall performance including a "modest to dramatic" increase in battery life on laptop computers.[98] Techradar gave it a 5 star rating calling it the best version of Windows yet.[99] The *New York Times*,[100] *USA Today*,[101] *The Wall Street Journal*,[102] and *The Telegraph*[103] also gave Windows 7 favorable reviews.

Some Windows Vista Ultimate users have expressed concerns over Windows 7 pricing and upgrade options.[104][105] Windows Vista Ultimate users wanting to upgrade from Windows Vista to Windows 7 must either pay $219.99[106] to upgrade to Windows 7 Ultimate or perform a clean install, which requires them to reinstall all of their programs.[107]

Editions

Windows 7 is available in six different editions, of which the Home Premium, Professional, and Ultimate editions are available for retail sale to consumers in most countries.[108] The other editions are not available in retail.[108] The Starter edition is only available preinstalled by OEMs on new PCs, the Enterprise edition only by volume licensing, and Home Basic only to certain developing countries' markets. Each edition of Windows 7 includes all of the capabilities and features of the edition below it.[108][109][110][111][112] All editions support the IA-32 processor architecture and all editions except Starter support the x86-64 processor architecture. The installation medium is the same for all the consumer editions of Windows 7 that have the same processor architecture, with the license determining the features that are activated; license upgrades permit the subsequent unlocking of features without re-installation of the operating system.[113] This is the first time Microsoft has distributed 2 DVDs (1 DVD for IA-32 processor architecture, the other DVD for x86-64 processor architecture) for each edition of Windows 7 (Except for Starter and Home Basic; some OEM copies have only DVD for IA-32 architecture; the installation DVD of Windows 7 Home Basic 64-bit edition is not included but can be obtained from Microsoft.). Users who wish to upgrade to an edition of Windows 7 with more features can then use Windows Anytime Upgrade to purchase the upgrade, and unlock the features of those editions.[109][113][114] Some copies of Windows 7 have restrictions, in which it must be distributed, sold, or bought and activated in the geographical region[115] specified in its front cover box.[116]

Microsoft is offering a family pack of Windows 7 Home Premium (in select markets) that allows installation on up to three PCs.[117] The "Family Pack" costs US$149.99 in the United States. On September 18, 2009, Microsoft said they were to offer temporary student discounts for Windows 7. The offer ran in the US and the United Kingdom, with similar schemes available in Canada, Australia, Korea, Mexico, France and India. Students with a valid .edu or .ac.uk email address could apply for either Windows 7 Home Premium or Professional, priced at $30 or £30.[118][119]

Windows 7 is also currently available as an embedded version to developers (previously Windows Embedded 2011).[120]

Marketing

The different editions of Windows 7 have been designed and marketed toward people with different needs.[121] Out of the different editions (Starter, Home Basic, Home Premium, Professional, Enterprise, and Ultimate), the Starter edition has been designed and marketed for lower cost notebooks, Home Basic for emerging markets, Home Premium for normal home users, Professional for businesses, Enterprise for larger businesses and corporations, and Ultimate for enthusiasts.[121]

Hardware requirements

Microsoft has published the minimum specifications for a system to run Windows 7.[122] Requirements for the 32-bit version are similar to that of premium editions of Vista, but are higher for 64-bit versions. Microsoft has released an upgrade advisor that determines if a computer is compatible with Windows 7.

Minimum hardware requirements for Windows 7[122]

Architecture	32-bit	64-bit
Processor	1 GHz IA-32 processor	1 GHz x86-64 processor
Memory (RAM)	1 GB	2 GB
Graphics card	DirectX 9 graphics processor with WDDM driver model 1.0 (Not absolutely necessary; only required for Aero)	
HDD free space	16 GB of free disk space	20 GB of free disk space
Optical drive	DVD-ROM drive[123] (Only to install from DVD-ROM media)	
SATA AHCI support	See below	

Additional requirements to use certain features:[122]

- SATA AHCI support was not added to Windows until XP Service Pack 1. As a result, in most motherboards the BIOS default for SATA support is to (emulate) IDE (ATA) rather than use AHCI (SATA). As explained here, this setting needs to be changed before installation, and any chipset-specific AHCI or RAID drivers need to be loaded (from a USB Flash drive, for example) at installation time.
- Windows XP Mode (Professional, Ultimate and Enterprise): Requires an additional 1 GB of RAM and additional 15 GB of available hard disk space. The requirement for a processor capable of hardware virtualization has been lifted.[124]
- Windows Media Center (included in Home Premium, Professional, Ultimate and Enterprise), requires a TV tuner to receive and record TV.

Physical memory limits

Maximum limits on physical memory (RAM) that Windows 7 can address vary depending on both the Windows version and between 32-bit and 64-bit versions.[125] The following table specifies the maximum physical memory limits supported:

Physical memory limits for Windows 7[125]

Version	Limit in 32-bit Windows	Limit in 64-bit Windows
Windows 7 Ultimate	4 GB	192 GB
Windows 7 Enterprise		
Windows 7 Professional		
Windows 7 Home Premium		16 GB
Windows 7 Home Basic		8 GB
Windows 7 Starter	2 GB	N/A

Processor limits

The maximum total number of logical processors[126] in a PC that Windows 7 supports is: 32[127] for 32-bit, 256[128] for 64-bit.

The maximum number of physical processors in a PC that Windows 7 supports is: 2 for Professional, Enterprise, and Ultimate, and 1 for Starter, Home Basic, and Home Premium.[129]

Service Pack 1

Windows 7 Service Pack 1 (SP1) was announced on March 18, 2010. A beta was released on July 12, 2010.[130][131][132] The final version was released to the public on February 9, 2011.[133] At the time of release, it was not made mandatory. Technet has information on blocking it, or getting it via Windows Update, direct download, or by ordering the Windows 7 SP1 DVD.[134] Microsoft confirmed that the service pack is to be on a much smaller scale than those released for previous versions of Windows, particularly Windows Vista.[135]

Windows 7 Service Pack 1 adds support for Advanced Vector Extensions (AVX), a 256-bit instruction set extension for processors, and improves IKEv2 by adding additional identification fields such as E-mail ID to it. In addition, it adds support for Advanced Format 512e as well as additional Identity Federation Services.[136][137] Windows 7 Service Pack 1 also resolves a bug related to HDMI audio and another related to printing XPS documents.[136]

Some programs have compatibility issues with SP1 and a limited number of programs may experience a loss of functionality.[138]

References

[1] http://www.microsoft.com/windows/windows-7/default.aspx

[2] http://windowsteamblog.com/blogs/windows7/archive/2009/07/22/windows-7-has-been-released-to-manufacturing.aspx

[3] Mike Nash (October 14, 2008). "Why 7?" (http://windowsteamblog.com/windows/archive/b/windowsvista/archive/2008/10/14/why-7.aspx). *The Windows Blog*. Microsoft. . Retrieved February 1, 2011.

[4] "Announcing Availability of Windows 7 and Windows Server 2008 R2 SP1" (http://windowsteamblog.com/windows/b/bloggingwindows/archive/2011/02/09/announcing-availability-of-windows-7-and-windows-server-2008-r2-sp1.aspx). Microsoft. . Retrieved February 9, 2011.

[5] Thadani, Rahul (September 6, 2010). "Windows 7 System Requirements" (http://www.buzzle.com/articles/windows7-system-requirements.html). *Buzzle*. . Retrieved February 25, 2011.

[6] LeBlanc, Brandon. "Windows 8 will be available on…" (http://windowsteamblog.com/windows/b/bloggingwindows/archive/2012/07/18/windows-8-will-be-available-on.aspx). The Windows Blog. . Retrieved 18 July 2012.

[7] Microsoft. "Windows 7 Lifecycle Policy" (http://support.microsoft.com/lifecycle/search/default.aspx?sort=PN&alpha=Windows+7&Filter=FilterNO). Microsoft. . Retrieved October 22, 2009.

[8] Ricciuti, Mike (July 20, 2007). "Next version of Windows: Call it 7" (http://www.news.com/2100-1016_3-6197943.html). CNET News. .

[9] Brandon LeBlanc. "Windows 7 Has Been Released to Manufacturing" (http://windowsteamblog.com/blogs/windows7/archive/2009/07/22/windows-7-has-been-released-to-manufacturing.aspx). .

[10] name="bott20090511"> "Windows 7 and Windows Server 2008 R2 Timelines Shared at Computex" (http://www.microsoft.com/presspass/features/2009/Jun09/06-02SteveGuggenheimer.mspx). Microsoft. June 3, 2009. . Retrieved June 3, 2009.

[11] Nash, Mike (October 28, 2008). "Windows 7 Unveiled Today at PDC 2008" (http://windowsteamblog.com/blogs/windows7/archive/2008/10/28/windows-7-unveiled-today-at-pdc-2008.aspx). *Windows Team Blog*. Microsoft. . Retrieved November 11, 2008.

[12] LeBlanc, Brandon (October 28, 2008). "How Libraries & HomeGroup Work Together in Windows 7" (http://windowsteamblog.com/blogs/windowsexperience/archive/2008/10/28/how-libraries-amp-homegroup-work-together-in-windows-7.aspx). *Windows Team Blog*. Microsoft. . Retrieved November 11, 2008.

[13] "Windows 7 to Skip Photo, Mail, Calendar and Movie Editing tools" (http://www.techpluto.com/softwares-missing-in-windows-7/). .

[14] "E-mail, photos, movie making will not be included in Windows 7" (http://www.networkworld.com/community/node/33084). .

[15] LeBlance, Brandon (October 28, 2008). "The Complete Windows Experience – Windows 7 + Windows Live" (http://windowsteamblog.com/blogs/windowsexperience/archive/2008/10/28/the-complete-windows-experience-windows-7-windows-live.aspx). *Windows Team Blog*. Microsoft. . Retrieved November 11, 2008.

[16] "Windows 7 Takes More Advantage of Multi-Core CPUs – Windows 7" (http://www.windowsvienna.com/windows-7-takes-more-advantage-of-multi-core-cpus-a24.html). Windowsvienna.com. . Retrieved May 25, 2009.

[17] "Windows 7 to get parallel-processing tweaks | All about Microsoft | ZDNet.com" (http://blogs.zdnet.com/microsoft/?p=1612). Blogs.zdnet.com. September 30, 2008. . Retrieved May 25, 2009.

[18] "Windows 7 to get parallel-processing tweaks | PC Tips" (http://www.pctipsbox.com/windows-7-to-get-parallel-processing-tweaks/). Pctipsbox.com. October 5, 2008. . Retrieved May 25, 2009.

[19] "Windows 7 enters parallel universe" (http://vista.blorge.com/2008/09/30/windows-7-enters-parallel-universe/). Vista.Blorge. September 30, 2008. . Retrieved May 25, 2009.

[20] Gruener, Wolfgang (January 16, 2008). "TG Daily — Windows Vista successor scheduled for a H2 2009 release?" (http://www.tgdaily.com/content/view/35641/118/). TG Daily. . Retrieved January 17, 2008.

[21] "ThinkNext.net: Screenshots from a blogger with Windows 7 M1" (http://www.thinknext.net/archives/2150). .

[22] Zack Whittaker (June 12, 2009). "Windows 7 UAC flaw: "Pandora's box of all vulnerabilities"" (http://blogs.zdnet.com/igeneration/?p=1826). . Retrieved June 14, 2009.

[23] Russinovich, Mark. "User Account Control Inside Windows 7 User Account Control" (http://technet.microsoft.com/en-us/magazine/2009.07.uac.aspx). Microsoft Corporation. . Retrieved June 14, 2009.

[24] "How to add Mac-like RAW image support to Windows 7, Vista, XP" (http://www.downloadsquad.com/2009/10/21/raw-image-support-windows/). .

[25] Softpedia (November 2008). "Windows 7 User Interface – The Superbar (Enhanced Taskbar)" (http://news.softpedia.com/news/Windows-7-User-Interface-The-Superbar-Enhanced-Taskbar-97143.shtml). . Retrieved November 12, 2008.

[26] "Windows 7: Some Minor Improvements, No Game Changer" (http://money.cnn.com/news/newsfeeds/siliconalley/big-tech/2008_12_windows_7_some_minor_improvements_no_game_changer_msft.html). .

[27] "Touching Windows 7 (Engineering Windows 7 Blog)" (http://blogs.msdn.com/e7/archive/2009/03/25/touching-windows-7.aspx). .

[28] "Engineering Windows 7 : Designing Aero Snap" (http://blogs.msdn.com/e7/archive/2009/03/17/designing-aero-snap.aspx). Steven Sinofsky/Microsoft. . Retrieved June 8, 2009.

[29] "Windows 7: Web Services in Native Code" (http://channel9.msdn.com/pdc2008/PC01/). PDC 2008. . Retrieved September 26, 2008.

[30] "Windows 7: Deploying Your Application with Windows Installer (MSI) and ClickOnce" (http://channel9.msdn.com/pdc2008/PC42/). PDC 2008. . Retrieved September 26, 2008.

[31] "Windows 7: Writing World-Ready Applications" (http://channel9.msdn.com/pdc2008/PC52/). PDC 2008. . Retrieved September 26, 2008.

[32] "WinHEC 2008 GRA-583: Display Technologies" (http://download.microsoft.com/download/5/E/6/5E66B27B-988B-4F50-AF3A-C2FF1E62180F/GRA-T583_WH08.pptx). Microsoft. November 6, 2008. . Retrieved December 4, 2008.

[33] "Windows 7 High Color Support" (http://news.softpedia.com/news/Windows-7-High-Color-Support-98741.shtml). Softpedia. November 26, 2008. . Retrieved December 5, 2008.

[34] "Support and Q&A for Solid-State Drives" (http://blogs.msdn.com/e7/archive/2009/05/05/support-and-q-a-for-solid-state-drives-and.aspx). *Engineering Windows 7*. Microsoft. May 5, 2009. . Retrieved May 9, 2009.

[35] Crothers, Brooke (November 6, 2008). "Microsoft describes USB 3.0 delays" (http://news.cnet.com/8301-13924_3-10083822-64.html). CNet. . Retrieved November 13, 2008.

[36] "Beta to RC Changes — Turning Windows Features On or Off" (http://blogs.msdn.com/e7/archive/2009/03/06/beta-to-rc-changes-turning-windows-features-on-or-off.aspx). .

[37] Thurrott, Paul (March 8, 2009). "Paul Thurrott's SuperSite for Windows: Windows 7 Build 7048 Notes" (http://www.winsupersite.com/win7/win7_7048_02.asp). Paul Thurrott's SuperSite for Windows. . Retrieved April 24, 2009.

[38] "Windows Virtual PC" (http://www.microsoft.com/windows/virtual-pc/). Microsoft. . Retrieved May 6, 2009.

[39] "Windows XP Mode for Windows 7 brochure" (http://download.microsoft.com/download/7/5/A/75A2C993-BFCC-47D0-8B6C-7C8CE2BA9833/Windows XP Mode for Windows 7_brochure.pdf). Microsoft. . Retrieved May 6, 2009.

[40] "Demonstration: Windows 7 VHD Boot" (http://www.microsoft.com/downloads/info.aspx?na=90&p=&SrcDisplayLang=en&SrcCategoryId=&SrcFamilyId=80ede31d-3509-407b-a896-0beea8705589&u=http://download.microsoft.com/download/7/0/F/70FE9C38-08D1-4FCC-BEF8-42B47DD968FE/Windows7VHDBoot.wmv). Microsoft. . Retrieved April 29, 2009.

[41] "Windows 7 Presentation Virtualization: Graphics Remoting (RDP) Today and Tomorrow" (http://channel9.msdn.com/pdc2008/ES21/). Microsoft. . Retrieved 2008.

[42] "Let's talk about Windows 7 Starter" (http://windowsteamblog.com/blogs/windows7/archive/2009/05/29/let-s-talk-about-windows-7-starter.aspx). Windows 7 Team. . Retrieved May 29, 2009.

[43] Windows Live team (October 22, 2009). "Finding your applications in Windows 7" (http://download.live.com/windows7). Microsoft. . Retrieved October 23, 2009.

[44] Thurrott, Paul (February 10, 2009). "No Ultimate Extras in Windows 7. But wait, it gets worse" (http://www.winsupersite.com/blogs/tabid/3256/entryid/75235/no-ultimate-extras-in-windows-7-but-wait-it-gets-worse.aspx). *Penton Media* (Windows IT Pro). . Retrieved March 24, 2011.

[45] Protalinski, Emil (September 24, 2008). "Three new Ultimate Extras arrive, including Microsoft Tinker" (http://arstechnica.com/microsoft/news/2008/09/three-new-ultimate-extras-arrive-including-microsoft-tinker.ars). *Condé Nast Publications* (Ars Technica). . Retrieved March 24, 2011.

[46] Lettice, John (October 24, 2001). "Gates confirms Windows Longhorn for 2003" (http://www.theregister.co.uk/2001/10/24/gates_confirms_windows_longhorn). The Register. . Retrieved March 5, 2008.

[47] "Microsoft cuts key Longhorn feature" (http://www.seattlepi.com/business/188339_msftcuts28.html). *Todd Bishop*. Hearst Seattle Media, LLC. August 28, 2004. . Retrieved March 25, 2009.

[48] Thurrott, Paul (February 14, 2007). "Windows "7" FAQ" (http://www.winsupersite.com/faq/windows_7.asp). Paul Thurrott's SuperSite for Windows. . Retrieved January 5, 2008.

[49] Foley, Mary J (July 20, 2007). "Windows Seven: Think 2010" (http://blogs.zdnet.com/microsoft/?p=592). ZDNet. . Retrieved September 19, 2007.

[50] Fried, Ina (October 13, 2008). "Microsoft makes Windows 7 name final" (http://news.cnet.com/8301-13860_3-10064971-56.html). CNET. . Retrieved October 13, 2008.

[51] Canadian Broadcasting Corporation (October 2008). "For Microsoft's Windows, 7th time's a charm" (http://www.cbc.ca/technology/story/2008/10/14/tech-windows.html). *CBC News*. . Retrieved October 27, 2008.

[52] Alex Castle (October 15, 2008). "Microsoft Justifies Its Windows 7 Naming Decision" (http://www.maximumpc.com/article/news/microsoft_justifies_its_windows_7_naming_decision). Maximum PC. . Retrieved November 18, 2009.

[53] Andrew. "Why Call it Windows 7?" (http://www.worldstart.com/tips/tips.php/why-call-it-windows-7). worldstart.com. . Retrieved November 20, 2009.

[54] Ian Cunningham (December 3, 2008). "Windows 7 Build Numbers" (http://www.w7forums.com/windows-7-build-numbers-t58.html). .

[55] "OSNews.com" (http://www.osnews.com/story/20703/Windows_7_Beta_1_Leaked). OSNews.com. . Retrieved May 25, 2009.

[56] "Windows 7 beta 1 performance - How does the OS compare to Vista and XP? | Hardware 2.0 | ZDNet.com" (http://blogs.zdnet.com/hardware/?p=3236&page=2). Blogs.zdnet.com. January 1, 2009. . Retrieved May 25, 2009.

[57] Graham-Smith, Darien (January 2009). "Follow-up: Benchmarking Windows 7" (http://www.pcpro.co.uk/blogs/2008/11/12/follow-up-benchmarking-windows-7/). . Retrieved January 29, 2009.

[58] "Leaked Windows 7 RC torrents infected with trojan" (http://www.slashgear.com/leaked-windows-7-rc-torrents-infected-with-trojan-2842048/). . SlashGear

[59] Pennington, Kenneth (January 2009). "Windows 7 64-Bit Beta Hits the Web" (http://w7info.com/articles/2009/01/windows-7-64-bit-beta-hits-the-web/). . Retrieved January 7, 2009.

[60] Johnson, Bobbie (January 8, 2009). "CES: Steve Ballmer unveils Microsoft's Windows 7 | Technology | guardian.co.uk" (http://www.guardian.co.uk/technology/2009/jan/08/steve-ballmer-ces-keynote). *The Guardian* (London). . Retrieved May 25, 2009.

[61] "Microsoft delays first Windows 7 public beta" (http://www.theregister.co.uk/2009/01/10/winows_7_beta_delay/). *Gavin Clarke*. The Register. January 10, 2009. . Retrieved March 25, 2009.

[62] Miller, Paul (April 24, 2009). "Windows 7 RC 7100 making its way to OEMs, a torrent tracker near you" (http://www.engadget.com/2009/04/24/windows-7-rc-7100-making-its-way-to-oems-a-torrent-tracker-near/). Engadget.com. . Retrieved May 25, 2009.

[63] "Windows 7 Release Candidate Customer Preview Program" (http://www.microsoft.com/windows/windows-7/download.aspx). Microsoft. . Retrieved May 5, 2009.

[64] "The Windows Blog" (http://windowsteamblog.com/blogs/windows7/archive/2009/07/21/when-will-you-get-windows-7-rtm.aspx). .

[65] Sinovsky, Steven (August 18, 2008). "The Windows 7 Team" (http://blogs.msdn.com/e7/archive/2008/08/18/windows_5F00_7_5F00_team.aspx). Engineering Windows 7, MSDN. . Retrieved December 9, 2009.

[66] Steven Levy (February 3, 2007). "Bill Gates on Vista and Apple's 'Lying' Ads" (http://www.msnbc.msn.com/id/16934083/site/newsweek/page/4/print/1/displaymode/1098/). .

[67] Bill Gates (May 12, 2007). "Bill Gates: Japan—Windows Digital Lifestyle Consortium" (http://www.microsoft.com/presspass/exec/billg/speeches/2008/05-07japanwdlc.mspx). .

[68] Sinofsky, Steven (December 15, 2008). "Continuing our discussion on performance" (http://blogs.msdn.com/e7/archive/2008/12/15/continuing-our-discussion-on-performance.aspx). *Engineering Windows 7*. Microsoft. . Retrieved December 18, 2008.

[69] Oiaga, Marius (June 24, 2008). "Windows 7 Will Not Inherit the Incompatibility Issues of Vista" (http://news.softpedia.com/news/Windows-7-Will-Not-Inherent-the-Incompatibility-Issues-of-Vista-88625.shtml). .

[70] Dignan, Larry (October 2008). "Ballmer: It's ok to wait until Windows 7; Yahoo still 'makes sense'; Google Apps 'primitive'" (http://blogs.zdnet.com/BTL/?p=10464). . Retrieved October 17, 2008.

[71] Keizer, Gregg F. (March 2008). "Windows 7 eyed by antitrust regulators" (http://web.archive.org/web/20080423070835/http://www.computerworld.com/action/article.do?command=viewArticleBasic&articleId=9068339&taxonomyId=14&intsrc=kc_top). Archived from the original (http://www.computerworld.com/action/article.do?command=viewArticleBasic&articleId=9068339&taxonomyId=14&intsrc=kc_top) on April 23, 2008. . Retrieved March 19, 2008.

[72] "Microsoft proposes "Browser Ballot Screen" to the EU" (http://www.neowin.net/news/main/09/07/24/microsoft-propose-browser-ballot-screen-to-the-eu). Neowin. July 24, 2009. . Retrieved August 1, 2009.

[73] Warren, Tom (August 1, 2009). "Microsoft scraps Windows 7 'E' version for Europe" (http://www.neowin.net/news/main/09/08/01/microsoft-scraps-windows-7-e-version-for-europe). Neowin. . Retrieved August 1, 2009.

[74] "Microsoft online Windows 7 store page" (http://emea.microsoftstore.com/UK/Microsoft/Windows/Windows-7/?WT.mc_id=MSCOMUK_HomePage_Buy_Windowstab). . Retrieved September 9, 2009.

[75] "Windows 7 flies off virtual shelf" (http://news.bbc.co.uk/1/hi/technology/8151342.stm). BBC News. July 15, 2009. . Retrieved July 15, 2009.

[76] Johnson, Bobbie (October 21, 2009). "Windows 7 set to break retail records" (http://www.guardian.co.uk/technology/2009/oct/21/windows-7-launch). *The Guardian* (London). . Retrieved October 21, 2009.

[77] "64bit版Windows 7は人気でやや品薄、週明けには回復？" (http://akiba-pc.watch.impress.co.jp/hotline/20091024/etc_win75.html). October 24, 2009. .

[78] "October 2009 OS stats: Windows 7 passes Snow Leopard, Linux" (http://arstechnica.com/microsoft/news/2009/11/october-2009-os-stats-windows-7-passes-snow-leopard-linux-1.ars?utm_source=microblogging&utm_medium=arstch&utm_term=Main Account&utm_campaign=microblogging). ars technica. November 6, 2009. . Retrieved November 7, 2009.

[79] "Windows 7 surpasses Snow Leopard in under two weeks" (http://www.neowin.net/news/main/09/11/06/windows-7-surpasses-snow-leopard-in-under-two-weeks). Neowin. November 7, 2009. . Retrieved November 7, 2009.

[80] Ina Fried (November 10, 2009). "Windows 7 use continues to climb" (http://news.cnet.com/8301-10805_3-10394517-75.html). CNET. . Retrieved November 13, 2009.

[81] Warman, Matt (March 5, 2010). "Microsoft sells more than 90 million copies of Windows 7" (http://www.telegraph.co.uk/technology/microsoft/7374148/Microsoft-sells-more-than-90-million-copies-of-Windows-7.html). *The Daily Telegraph* (UK: Telegraph Media Group Limited). . Retrieved April 30, 2011.

[82] "Windows 7 now fastest-selling Windows OS" (http://www.zdnet.com/news/windows-7-now-fastest-selling-windows-os/417718). ZDNet. April 27, 2010. . Retrieved June 24, 2010.

[83] "Windows 7: 150 Million Copies Sold" (http://www.windowsitpro.com/article/paul-thurrotts-wininfo/Windows-7-150-Million-Copies-Sold.aspx). Windows IT Pro. June 23, 2010. . Retrieved June 24, 2010.

[84] "150 Million Licenses of Windows 7 Sold, Windows Live Betas Announced" (http://windowsteamblog.com/windows/b/bloggingwindows/archive/2010/06/23/150-million-licenses-of-windows-7-sold-windows-live-betas-announced.aspx). Microsoft. June 23, 2010. . Retrieved June 24, 2010.

[85] "64-Bit Momentum Surges with Windows 7" (http://windowsteamblog.com/windows/b/bloggingwindows/archive/2010/07/08/64-bit-momentum-surges-with-windows-7.aspx). Microsoft. July 8, 2010. . Retrieved August 2, 2010.

[86] "Microsoft: Windows 7 makes 64-bit headway" (http://news.cnet.com/8301-10805_3-20010092-75.html). cnet. July 9, 2010. . Retrieved August 2, 2010.

[87] "Windows 7 Momentum Continues: 175 Million Licenses Sold" (http://windowsteamblog.com/windows/b/bloggingwindows/archive/2010/07/22/windows-7-momentum-continues-175-million-licenses-sold.aspx). Microsoft. July 22, 2010. . Retrieved July 27, 2010.

[88] "Celebrating Windows 7 at 1 Year – More than 240 Million Licenses Sold" (http://windowsteamblog.com/windows/b/bloggingwindows/archive/2010/10/21/celebrating-windows-7-at-1-year-more-than-240-million-licenses-sold.aspx). Microsoft. October 21, 2010. . Retrieved October 22, 2010.

[89] "Windows 7: 300 Million Licenses Sold" (http://windowsteamblog.com/windows/b/bloggingwindows/archive/2011/01/27/windows-7-300-million-licenses-sold.aspx). Microsoft. January 27, 2011. . Retrieved February 20, 2010.

[90] "Microsoft Reports Record Fourth-Quarter and Full-Year Results" (http://www.microsoft.com/presspass/press/2011/jul11/07-21fy11Q4earningsPR.mspx?rss_fdn=Press Releases). Microsoft. July 12, 2011. . Retrieved August 30, 2011.

[91] "Earnings Release FY12 Q2" (http://www.microsoft.com/investor/EarningsAndFinancials/Earnings/PressReleaseAndWebcast/FY12/Q2/default.aspx). Microsoft. 2012-01-19. . Retrieved 2012-01-21.

[92] "Microsoft Windows 7 (Home Premium) Review – CNet" (http://reviews.cnet.com/windows/microsoft-windows-7-home/4505-3672_7-33704139.html). CNet. July 31, 2009. .

[93] Michael Muchmore (October 22, 2009). "Microsoft Windows 7" (http://www.pcmag.com/article2/0,2817,2348899,00.asp). PC Magazine. . Retrieved November 13, 2009.

[94] Will Smith (October 19, 2009). "Windows 7 Review: XP vs Vista vs 7 in 80+ Benchmarks" (http://www.maximumpc.com/article/reviews/windows_7_review). Maximum PC. . Retrieved November 13, 2009.

[95] Harry McCracken (October 19, 2009). "Windows 7 Review" (http://www.pcworld.com/article/172602/windows_7_review.html). PC World. . Retrieved November 13, 2009.

[96] The PC World Editorial Team (October 19, 2009). "The PC World 100: Best Products of 2009" (http://www.pcworld.com/article/174171-2/the_pc_world_100_best_products_of_2009.html). PC World. . Retrieved November 13, 2009.

[97] Paul Miller (August 12, 2009). "Windows 7 review" (http://www.engadget.com/2009/08/12/windows-7-review/). Engadget. . Retrieved November 13, 2009.

[98] Dana Wollman (August 21, 2009). "Windows 7" (http://www.laptopmag.com/review/software/windows-7.aspx). LAPTOP Magazine. . Retrieved November 13, 2009.

[99] Mary Branscombe (August 7, 2009). "Windows 7 review" (http://www.techradar.com/reviews/pc-mac/software/operating-systems/microsoft-windows-7-622923/review). TechRadar. . Retrieved November 13, 2009.

[100] David Pogue (October 21, 2009). "Windows 7 Keeps the Good, Tries to Fix Flaws" (http://www.nytimes.com/2009/10/22/technology/personaltech/22pogue.html). *New York Times*. . Retrieved November 13, 2009.

[101] Edward C. Baig (October 21, 2009). "After Vista, Windows 7 is a giant leap for Microsoft" (http://www.usatoday.com/tech/products/2009-10-16-baig16_CV_N.htm). *USA Today*. . Retrieved November 13, 2009.

[102] Walter S. Mossberg (October 8, 2009). "A Windows to Help You Forget" (http://online.wsj.com/article/SB10001424052748703298004574459293141191728.html). *Wall Street Journal*. . Retrieved November 13, 2009.

[103] Matt Warman (October 20, 2009). "Microsoft Windows 7 review" (http://www.telegraph.co.uk/technology/microsoft/6384509/Microsoft-Windows-7-review.html). *The Daily Telegraph* (London: The Daily Telegraph). . Retrieved November 13, 2009.

[104] "Some Vista users say they're getting the Ultimate shaft" (http://news.cnet.com/8301-13860_3-10277506-56.html). July 2, 2009. .

[105] "Vista Ultimate users fume, rant over Windows 7 deals" (http://www.computerworld.com/s/article/9135121/Vista_Ultimate_users_fume_rant_over_Windows_7_deals). July 2, 2009. .

[106] "Shop: Windows 7" (http://www.microsoft.com/windows/buy/default.aspx). October 22, 2009. .
[107] "Windows 7 Upgrade Considerations" (http://www.microsoft.com/windows/windows-7/get/upgrade-considerations.aspx). October 22, 2009. .
[108] "All Windows 7 Versions—What You Need to Know" (http://www.extremetech.com/article2/0,2845,2340431,00.asp). ExtremeTech. February 5, 2009. . Retrieved February 5, 2009.
[109] Thurrott, Paul (February 3, 2009). "Windows 7 Product Editions" (http://www.winsupersite.com/win7/win7_skus.asp). . Retrieved February 3, 2009.
[110] "Windows 7 will come in many flavors" (http://news.cnet.com/8301-13860_3-10155193-56.html). CNET News. February 3, 2009. . Retrieved February 3, 2009.
[111] "Windows 7 Editions – Features on Parade" (http://news.softpedia.com/news/Windows-7-Editions-Features-on-Parade-103766.shtml). Softpedia. February 5, 2009. . Retrieved February 5, 2009.
[112] "Windows 7: Which Edition is Right For You?" (http://www.pcworld.com/article/158870/microsoft_announces_six_windows_7_editions.html). PCWorld. February 3, 2009. . Retrieved February 5, 2009.
[113] LeBlanc, Brandon (February 9, 2009). "A closer look at the Windows 7 SKUs" (http://windowsteamblog.com/blogs/windows7/archive/2009/02/04/a-closer-look-at-the-windows-7-skus.aspx). *Windows Team Blog*. Microsoft. . Retrieved February 9, 2009.
[114] "All Windows 7 Versions—What You Need to Know – Release Date, Cost, and Upgrades" (http://www.extremetech.com/article2/0,2845,2340432,00.asp). ExtremeTech. February 5, 2009. . Retrieved February 16, 2009.
[115] One of the geographical regions can be either: Southeast Asia; SouthEast Asia excluding Singapore; India; Middle East and Africa; or Latin America and the Caribbean.
[116] "MICROSOFT SOFTWARE LICENSE TERMS: WINDOWS 7." (http://download.microsoft.com/Documents/UseTerms/Windows 7_Ultimate_English_c44ca3df-8338-4a2f-a176-39d2e68986c4.pdf) (PDF). . Retrieved January 7, 2012.
[117] "Microsoft Announces "Family Pack" For Windows 7" (http://windowsteamblog.com/blogs/windows7/archive/2009/07/21/when-will-you-get-windows-7-rtm.aspx). Microsoft. July 21, 2009. . Retrieved July 21, 2009.
[118] Microsoft. "Windows: Student Offer" (http://www.microsoft.com/uk/windows/studentoffer/default.aspx). . Retrieved September 18, 2009.
[119] Warren, Tom. "Microsoft: Students to get Windows 7 for £30/$30" (http://www.neowin.net/news/main/09/09/17/microsoft-uk-students-to-get-windows-7-for-30). . Retrieved September 18, 2009.
[120] "Windows Embedded Standard 7" (http://www.microsoft.com/windowsembedded/en-us/products/westandard/default.mspx). Microsoft. . Retrieved June 19, 2010.
[121] "All Windows 7 Versions—What You Need to Know" (http://www.extremetech.com/article2/0,2845,2340431,00.asp). ExtremeTech. February 5, 2009. . Retrieved August 2, 2010.
[122] "Windows 7 system requirements" (http://windows.microsoft.com/en-us/windows7/products/system-requirements). Microsoft. .
[123] Any optical drive that can read DVD-ROM media.
[124] "Windows Virtual PC – no hardware virtualization update now available for download" (http://blogs.msdn.com/virtual_pc_guy/archive/2010/03/18/windows-virtual-pc-no-hardware-virtualization-update-now-available-for-download.aspx). .
[125] "Memory Limits for Windows Releases" (http://msdn.microsoft.com/en-us/library/aa366778(VS.85).aspx#physical_memory_limits_windows_7). Microsoft. . Retrieved June 25, 2010.
[126] A logical processor is either: 1) One of the two handlers of thread of instructions of one of the number of cores of one of the number of physical processors with support for HyperThreading; or 2) One of the number of cores of one of the number of physical processors without support for HyperThreading.
[127] 32 cores without support for HyperThreading, 16 cores with support for HyperThreading.
[128] 256 cores without support for HyperThreading, 128 cores with support for HyperThreading.
[129] "Processor limits for Windows 7" (http://www.microsoft.com/windows/windows-7/get/system-requirements.aspx). Microsoft.com. . Retrieved November 29, 2010.
[130] "Virtualization Updates at TechEd" (http://windowsteamblog.com/windows/b/business/archive/2010/06/07/virtualization-updates-at-teched.aspx). *Windows Team Blog*. Microsoft. June 7, 2010. .
[131] "Microsoft Announces Windows 7 Service Pack 1" (http://www.computerworld.com/s/article/9172738/Microsoft_announces_Windows_7_SP1). Computer World. .
[132] "Microsoft Releases SP1 Beta for Windows Server 2008 R2" (http://redmondmag.com/articles/2010/07/12/microsoft-releases-sp1-beta-for-windows-server-2008-r2.aspx). Redmondmag.com. .
[133] "Announcing Availability of Windows 7 and Windows Server 2008 R2 SP1" (http://windowsteamblog.com/windows/b/bloggingwindows/archive/2011/02/09/announcing-availability-of-windows-7-and-windows-server-2008-r2-sp1.aspx). windowsteamblog. .
[134] "Service Pack 1 (SP1) for Windows Server 2008 R2 and Windows 7" (http://technet.microsoft.com/en-us/windows/gg635126.aspx). Technet.microsoft.com. . Retrieved June 21, 2011.
[135] "Microsoft: Few issues to address with Windows 7 Service Pack 1" (http://www.techradar.com/news/software/operating-systems/microsoft-few-issues-to-address-with-windows-7-service-pack-1-665048). TechRadar UK. .
[136] "Notable Changes in Windows 7 and Windows Server 2008 R2 Service Pack 1" (http://download.microsoft.com/download/E/B/A/EBA55FE2-373A-4351-9346-6D762B79AA69/Notable Changes in Windows 7 and Windows Server 2008 R2 Service Pack 1.doc). *Microsoft Download Center*. Microsoft Corporation. February 9, 2011. . Retrieved March 10, 2011.

[137] "Information about Service Pack 1 for Windows 7 and for Windows Server 2008 R2 (Revision 3.1)" (http://support.microsoft.com/?kbid=976932). *Microsoft Support Center*. Microsoft Corporation. February 22, 2011. . Retrieved March 10, 2011.
[138] "Some programs have compatibility issues with Service Pack 1 for Windows 7 and for Windows Server 2008 R2" (http://support.microsoft.com/kb/2492938). Microsoft. February 22, 2011. . Retrieved April 28, 2011.

External links

- Official website (http://www.microsoft.com/windows/windows-7/)
- The Windows 7 Blog for Developers (http://windowsteamblog.com/blogs/developers/default.aspx)
- The Windows 7 Team Blog – Windows Team Blog (http://windowsteamblog.com/blogs/windows7/default.aspx)

Command-line interface

A **command-line interface (CLI)** is a means of interaction with a computer program where the user (or client) passes commands in the form of a line of text (a command line).

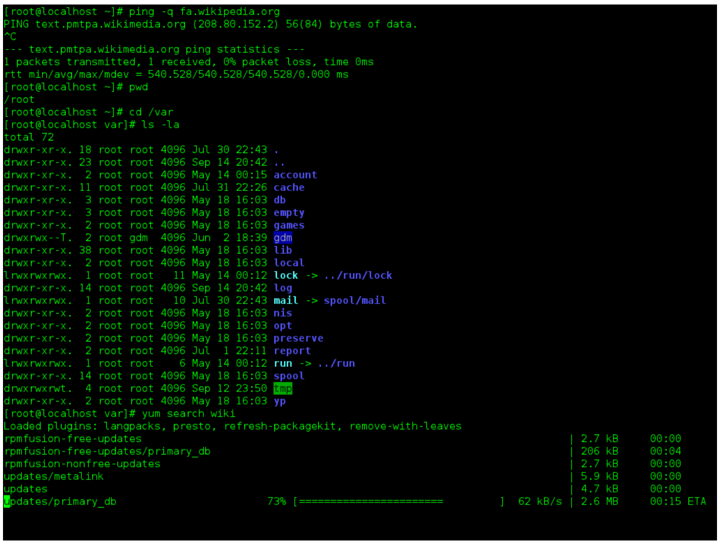

Screenshot of a sample Bash session. GNOME Terminal 3, Fedora 15

The command-line interface evolved from a form of dialog once conducted by humans over teleprinter machines, in which human operators remotely exchanged information, usually one line of text at a time. Early computer systems often used teleprinter machines as the means of interaction with a human operator. The computer became one end of the human-to-human teleprinter model. So instead of a human communicating with another human over a teleprinter, a human communicated with a computer.

In time, the actual mechanical teleprinter was replaced by a glass tty (keyboard and screen, but emulating the teleprinter), and then by a terminal (where the computer software could address all of the screen, rather than only print successive lines).

Due to its text-based nature, a command-line interface is sometimes confused with the text-based user interface, a kind of user interface which uses only text, but not necessarily presented in successive lines. Text may be formatted and appear in fixed locations on a computer terminal display, as opposed to only appearing in successive lines.

The CLI was the primary means of human interaction with most early operating systems,*Talk:Command-line interface#early operating systems (sic)* including MS-DOS, CP/M, Unix, and Apple DOS. The interface is usually implemented with a command line shell, which is a program that accepts commands as text input and converts commands to appropriate operating system functions.

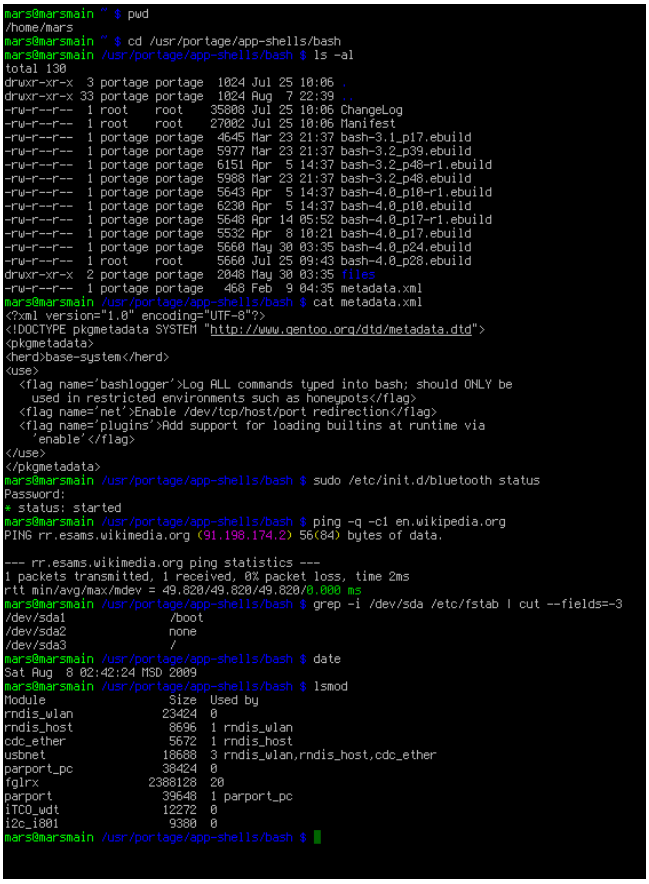

Screenshot of a sample Bash session, taken on an old release of Gentoo Linux.

Alternatives to the command line include, but are not limited to menus and various desktop metaphors centered on the pointer (usually controlled with a mouse).

Today, command-line interfaces to a computer operating system are less widely used by casual computer users, who favor graphical user interfaces.

Command-line interfaces are still often preferred by more advanced computer users, as they often provide a more concise and powerful means to control a program or operating system.

Programs with command-line interfaces are often much easier to automate via scripting.

Operating System Command-Line Interfaces

Operating system command line interfaces are usually distinct programs supplied with the operating system.

A program that implements such a text interface is often called a command-line interpreter, command processor or shell. The term 'shell', often used to describe a command-line interpreter, can be in principle any program that constitutes the user-interface, including fully graphically oriented ones—for example, the default

Windows GUI is created by a shell program named EXPLORER.EXE, as defined in the SHELL=EXPLORER.EXE line in the WIN.INI configuration file.

Examples of command-line interpreters include the various Unix shells (sh, ksh, csh, tcsh, bash, etc.), the historical CP/M CCP, and MS-DOS/IBM-DOS/DR-DOS's COMMAND.COM, as well as the OS/2 and the Windows CMD.EXE programs, the latter groups being based heavily on DEC's RSX and RSTS CLIs. Under most operating systems, it is possible to replace the default shell program by more specialized or powerful alternatives; some widespread examples include 4DOS for DOS, 4OS2 for OS/2, and 4NT or Take Command for Windows.

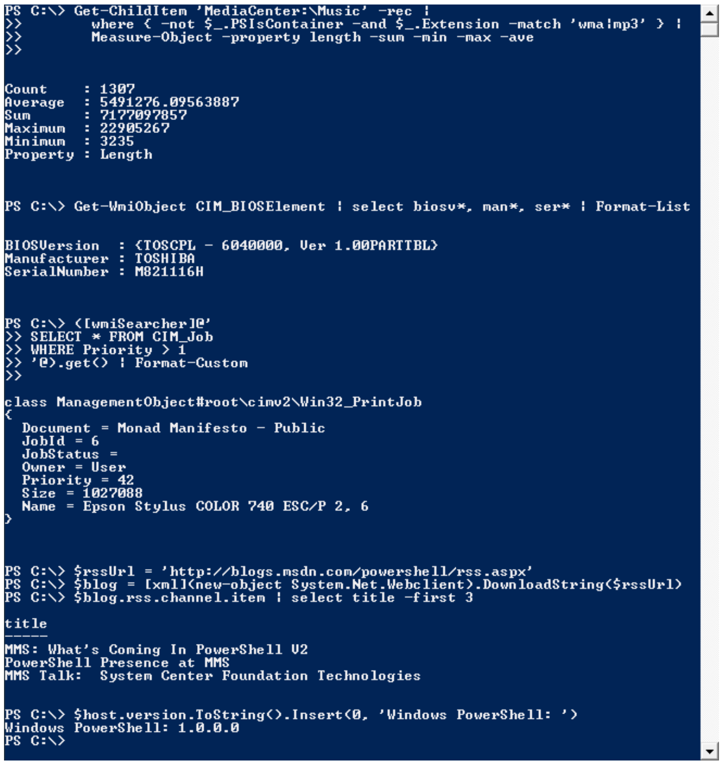

Screenshot of Windows PowerShell 1.0, running under Windows Vista

In November 2006, Microsoft released version 1.0 of Windows PowerShell (formerly codenamed *Monad*), which combined features of traditional Unix shells with their object-oriented .NET Framework. MinGW and Cygwin are open-source packages for Windows that offer a Unix-like CLI. Microsoft provides MKS Inc.'s ksh implementation *MKS Korn shell* for Windows through their Services for UNIX add-on.

The latest versions of the Macintosh operating system are based on a variation of Unix called Darwin. On these computers, users can access a Unix-like command-line interface called Terminal found in the Applications Utilities folder. (This terminal uses bash by default.)

Application Command-Line Interfaces

Application programs (as opposed to operating systems) may also have command line interfaces.

An application program may support none, any, or all of these three major types of command line interface mechanisms:

1. Commands are specified as parameters upon invocation of the program.
   - Most operating systems support a means to pass additional information to a program when it is launched. When a program is launched from an OS command line shell, additional text provided at the shell command line along with the program name is passed to the launched program via command line arguments.
2. Commands are entered via an interactive command line session controlled by the application.
   - After launch, a program may provide an operator with an independent means to enter commands in the form of text.
3. Commands are entered via the operating system STDIN/STDOUT mechanism.

Some applications provide both a CLI and a GUI. In some cases, the GUI is a wrapper around a CLI application; other times, there is a CLI to control a GUI application. The engineering/scientific numerical computation package MATLAB provides no GUI for some calculations, but the CLI can handle any calculation. The three-dimensional-modelling program Rhinoceros 3D provides a CLI as well as a distinct scripting language. In some computing environments, such as the Oberon or Smalltalk user interface, most of the text which appears on the screen may be used for giving commands.

There are command-line interpreters for editing text files like ED and EDLIN, DEBUG, for disk management DISKPART, DFSEE, calculators (PC-DOS ACALC), all of which present a usable command prompt.

Some programs have the command line separate from the action window (where the output would appear). The early Sierra games, like KingsQuest 1-3, use commands from an internal command line to move the character around in the graphic window. Editors like Vi and IBM PC DOS E Editor use commands typed on the command line, switching back and forth to edit mode, a kind of graphical ED or EDLIN, as it were. (ED and EDLIN are pure teletype programs).

The Command-Line Interface continues to co-evolve with GUIs like those provided by Microsoft Windows, Mac OS and the X Window System. Programs that make use of external helper programs, often make use of command lines embedded in the GUI interface or configuration. In some applications, such as MATLAB, AutoCAD or EAGLE, a CLI is integrated with the GUI, with some benefits of both.

Usage

A CLI is used whenever a large vocabulary of commands or queries, coupled with a wide (or arbitrary) range of options, can be entered more rapidly as text than with a pure GUI. This is typically the case with operating system command shells. CLIs are also used by systems with insufficient resources to support a graphical user interface. Some computer language systems (such as Python, Forth, LISP and many dialects of BASIC) provide an interactive command-line mode to allow for experimentation.

CLIs are often used by programmers and system administrators, in engineering and scientific environments, and by technically advanced personal computer users. CLIs are also popular among people with visual disability, since the commands and responses can be displayed using Refreshable Braille displays.

Anatomy of a shell CLI

A CLI can generally be considered as consisting of syntax and semantics. The *syntax* is the grammar that all commands must follow. In the case of operating systems (OS), MS-DOS and Unix each define their own set of rules that all commands must follow. In the case of embedded systems, each vendor, such as Nortel, Juniper Networks or Cisco Systems, defines their own proprietary set of rules that all commands within their CLI conform to. These rules also dictate how a user navigates through the system of commands. The *semantics* define what sort of operations are possible, on what sort of data these operations can be performed, and how the grammar represents these operations and data—the symbolic meaning in the syntax.

Two different CLIs may agree on either syntax or semantics, but it is only when they agree on both that they can be considered sufficiently similar to allow users to use both CLIs without needing to learn anything, as well as to enable re-use of scripts.

A simple CLI will display a prompt, accept a "command line" typed by the user terminated by the Enter key, then execute the specified command and provide textual display of results or error messages. Advanced CLIs will validate, interpret and parameter-expand the command line before executing the specified command, and optionally capture or redirect its output.

Unlike a button or menu item in a GUI, a command line is typically self-documenting, stating exactly what the user wants done. In addition, command lines usually include many defaults that can be changed to customize the results. Useful command lines can be saved by assigning a character string or alias to represent the full command, or several commands can be grouped to perform a more complex sequence – for instance, compile the program, install it, and run it — creating a single entity, called a command procedure or script which itself can be treated as a command.

These advantages mean that a user must figure out a complex command or series of commands only once, because they can be saved, to be used again.

The commands given to a CLI shell are often in one of the following forms:

- *doSomething how toFiles*
- *doSomething how sourceFile destinationFile*
- *doSomething how < inputFile > outputFile*
- *doSomething how | doSomething how | do Something how > outputFile*

where *doSomething* is, in effect, a verb, *how* an adverb (for example, should the command be executed "verbosely" or "quietly") and *toFiles* an object or objects (typically one or more files) on which the command should act. The '>' in the third example is a redirection operator, telling the command-line interpreter to send the output of the command not to the screen but to the file named on the right of the '>'. This will overwrite the file. Using '>>' will redirect the output and append it to the file. Another redirection operator is the vertical bar ('|'), which creates a pipeline where the output of one command becomes the input to the next command.

CLI and resource protection

A command interpreter running in an operating system, is bound by the same restrictions as any other program provides. These come from the operating system, file-system drivers, and network servers.

A command processor might provide further protection in preventing redirection overwriting files, or writing to invalid types. However, this is a feature of an individual program, and is not part of the interface.

One can modify the set of available commands, by modifying which paths appear in the PATH environment. Under UNIX, commands further need to have the 'x' attribute set to mark them as executable files. Since the path is searched for in order, by re-ordering the path, one can run eg \OS2\MDOS\E.EXE instead of \OS2\E.EXE, when the default is the opposite. Renaming of the executables also works: people often rename their favourite editor to EDIT, for example.

The command line does not have a 'heirachy of commands' outside of that provided by PATH. This is an idea associated with menu commands. Even the graphical user interface provides more commands than are present on the menu. The tools menu usually contains a command to edit the menus, button-bars, and hot-keys to alter the heirachy.

Both the command line and the GUI allow one to restrict available commands, such as access to advanced internal commands. The Windows CMD.EXE does this. Often, shareware programs will limit the range of commands, including greying out commands from use, or printing a command 'your administrator has disabled running batch files' from the prompt.

Although the command line interface does not provide for GUI features, one can add this to the console it runs under. BASICA for example, has a list of function keys at the bottom, each of which equates to typing the text at the prompt. There are programs that add buttons to the console in Windows, and alternate consoles, which do exactly the same thing (TCBUTTON in TCMD32, for example).

It might be noted that the interface, either command or graphic, can reduce access to available commands. For example, a number of GUI programs have an 'expert' mode, which makes more commands accessable from the menus. Also, the user can configure the menus to add or remove commands to menus, button bars etc.

Command prompt

A command prompt (or just *prompt*) is a sequence of (one or more) characters used in a command-line interface to indicate readiness to accept commands. Its intent is to literally prompt the user to take action. A prompt usually ends with one of the characters $, %, #, :, > and often includes other information, such as the path of the current working directory.

On many Unix system and derivative systems, it is common for the prompt to end in a $ or % character if the user is a normal user, but in a # character if the user is a superuser ("root" in Unix terminology).

It is common for prompts to be modifiable by the user. Depending on the environment, they may include colors, special characters, and other elements like variables and functions for the current time, user, shell number or working directory, in order, for instance, to make the prompt more informative or visually pleasing, to distinguish sessions on various machines, or to indicate the current level of nesting of commands. On some systems, special tokens in the definition of the prompt can be used to cause external programs to be called by the command-line interpreter while displaying the prompt.

In DOS's COMMAND.COM and in the Windows NT's command-line interpreter cmd.exe the prompt is modifiable by issuing a `prompt` command or by directly changing the value of the corresponding `%PROMPT%` environment variable. The default of most modern systems, the `C:\>` style is obtained, for instance, with "`prompt $P$G`". The default of older DOS systems, `C>` is obtained by just "`prompt`", although on some systems this produces the newer `C:\>` style, unless used on floppy drives A: or B:; on those systems "`prompt $N$G`" can be used to override the automatic default and explicitly switch to the older style.

On many Unix systems, the `$PS1` variable can be used, although other variables also may have an impact on the prompt (depending on what shell is being used). In the bash shell, a prompt of the form

```
[time] user@host: work_dir $
```

could be set by issuing the command

```
export PS1='[\t] \u@\H: \W $'
```

In zsh the `$RPROMPT` variable controls an optional "prompt" on the right hand side of the display. It is not a real prompt in that the location of text entry does not change. It is used to display information on the same line as the prompt, but right justified.

In RISC OS, the command prompt is a '`*`' symbol, and thus (OS)CLI commands are often referred to as "star commands".[1] It is also possible to access the same commands from other command lines (such as the BBC BASIC command line), by preceding the command with a '`*`'.

Arguments

A **command-line argument** or **parameter** is an argument sent to a program being called. In principle a program can take many command-line arguments, the meaning and importance of which depend entirely upon the program.

When a command processor is active a program is typically invoked by typing its name followed by command-line arguments (if any). For example, in Unix and Unix-like environments, an example of a command-line argument is:

```
rm file.s
```

"file.s" is a command-line argument which tells the program rm to remove the file "file.s".

Some programming languages, such as C, C++ and Java, allow a program to interpret the command-line arguments by handling them as string parameters in the main function. Other languages, such as Python, expose these arguments as global variables.

Command-line option

A **command-line option** or simply **option** (also known as a **flag** or **switch**) modifies the operation of a command; the effect is determined by the command's program. Options follow the command name on the command line, separated by spaces. A space before the first option is not always required.

For example, in the OpenVMS operating system, the command *directory* is used to list the files inside a directory. By default—that is, when the user simply types *directory*—it will list only the names of the files. By appending the */owner* option (to form the command *directory/owner*), the user can instruct the *directory* command to also display the ownership of the files.

The format of options varies widely between operating systems. In most cases the syntax is by convention rather than an operating system requirement; the entire command line is simply a string passed to a program, which can process it in any way the programmer wants.

Uniquely, in Multics, unique abbreviations are associated command-line options and subsystem keywords, an idea that appears to derive from the PL/I programming language with its shortened keywords (*e.g.*, STRGE for STRINGRANGE and DCL for DECLARE). For example, in the Multics "forum" subsystem, the **-long_subject** parameter can be abbreviated **-lgsj**. It is also common for Multics commands to sport unique abbreviations, typically corresponding to the initial letters of the words that are strung together with underscores to form the commands' names, such as the use of **did delete_iacl_dir**.

Sometimes different programs use different syntax in the same operating system. For example:

- Options may be indicated by `-`, `/`, or either.
- They may or may not be case-sensitive.
- Sometimes options and their arguments are run together, sometimes separated by whitespace, and sometimes by a character, typically `:` or =. Thus "Prog -fFilename", "Prog -f Filename", "Prog -f:Filename", "Prog -f=Filename".
- Some programs allow single-character options to be combined; others do not. The switch "-fA" may mean the same as "-f -A", or it may be incorrect, or it may even be a valid but different parameter.

In Unix-like systems, the ASCII hyphen-minus is commonly used to specify options. The character is usually followed by one or more letters. An argument that is a single hyphen-minus by itself without any letters usually specifies that a program should handle data coming from the standard input or send data to the standard output. Two hyphen-minus characters (--) are used on some programs to specify "long options" where more descriptive option names are used. This is a common feature of GNU software.

In DOS, OS/2 and Windows the forward slash ('/') is more prevalent, although the hyphen-minus is also sometimes used there. FlexOS, 4680 OS and 4690 OS use '-'. In many versions of DOS (MS-DOS/PC-DOS 2.xx and higher, all versions of DR-DOS since 5.0, as well as in PTS-DOS and FreeDOS) the current switch character to be used is defined by a value returned from a system call (INT 21/AH=37h). The default character returned by this API is '/', but can be changed to a hyphen-minus on all above-mentioned systems (except for MS-DOS/PC-DOS 5.0 and higher). In some of these systems (MS-DOS/PC-DOS 2.xx/3.xx, DR-DOS 7.02 and higher, and FreeDOS), the setting can not only be controlled by programs, but also pre-configured by a SWITCHAR directive in CONFIG.SYS. While applications should retrieve this setting before parsing command line arguments, many existing programs don't adhere to this standard and are hardwired to use '/' only. Under DR-DOS, if the setting has been changed from '/', the first directory separator '\' in the display of the PROMPT parameter $G will change to a forward slash '/' (also a valid directory separator in DOS, FlexOS, 4680 OS, 4690 OS, OS/2 and Windows) as a visual clue to indicate the change. Some command-line interpreters (including newer versions of DR-DOS COMMAND.COM and 4DOS) also provide pseudo-environment variables named %/% or %SWITCHAR% to allow portable batchjobs to be written.

It is impossible to know what arguments a program can recognise, and what syntax to use without consulting program documentation. It is usual for a program to display a brief summary of its parameters when invoked with a command-line which is typically one of: no parameters; `?`; `-?`; `-h`; `/?`; `/h`; `-help`; or `--help`. Entering a program name without parameters in the hope that it will display parameters can be hazardous, as some programs and scripts execute without further ado.

The space character

In many areas of computing, but particularly in the command line, the space character can cause problems as it has two distinct and incompatible functions: as part of a command or parameter, or as a parameter or name separator. Ambiguity can be prevented either by prohibiting embedded spaces in file- and directory names in the first place (for example, by substituting them with underscores '_'), or, if supported by the command-line interpreter and the programs taking these parameters as arguments, by enclosing a name with embedded spaces between quote characters or using a escape character before the space, usually a backslash ('\'). For example

`Long path/Long program name Parameter one Parameter two` ...

is ambiguous (is "program name" part of the program name, or two parameters?); however

`Long_path/Long_program_name Parameter_one Parameter_two` ...,

`LongPath/LongProgramName ParameterOne ParameterTwo` ...,

`"Long path/Long program name" "Parameter one" "Parameter two"` ...

and

`Long\ path/Long\ program\ name Parameter\ one Parameter\ two` ...

are not ambiguous. Unix-based operating systems minimize the use of embedded spaces to minimize the need for quotes. In Microsoft Windows, one often has to use quotes because embedded spaces (such as in directory names) are common.

Command-line interpreter

The terms **command-line interpreter**, **command line shell**, **command language interpreter**, or identical abbreviation **CLI**, are applied to computer programs designed to interpret a sequence of lines of text which may be entered by a user, read from a file or another kind of data stream. The context of interpretation is usually one of a given operating system or programming language.

Command-line interpreters allow users to issue various commands in a very efficient (and often terse) way. This requires the user to know the names of the commands and their parameters, and the syntax of the language that is interpreted.

The unix /# and OS/2 EXTPROC commands facilitate the passing of batch files to external processors. One can use this to write specific command processors for dedicated uses, and process external data files which reside in batch files.

Many graphical interfaces, such as the OS/2 Presentation Manager and early versions of Microsoft Windows use command-lines to call helper programs to open documents and programs. The commands are stored in the graphical shell or in files like the registry or the OS/2 `os2user.ini` file.

Early history

From the 1960s onwards, user interaction with computers was primarily by means of command-line interfaces, initially on machines like the Teletype Model 33 ASR, but then on early CRT-based computer terminals such as the VT52.

All of these devices were purely text based, with no ability to display graphic or pictures.[2] For business application programs, text-based menus were used, but for more general interaction the command line was the interface.

From the early 1970s the Unix operating system on minicomputers pioneered the concept of a powerful command-line environment, which Unix called the "shell", with the ability to "pipe" the output of one command in as input to another, and to save and re-run strings of commands as "shell scripts" which acted like custom commands.

The command-line was also the main interface for the early home computers such as the Commodore PET, Apple II and BBC Micro – almost always in the form of a BASIC interpreter. When more powerful business oriented microcomputers arrived with CP/M and later MSDOS computers such as the IBM PC, the command-line began to borrow some of the syntax and features of the Unix shells such as globbing and piping of output.

The command-line was first seriously challenged by the PARC GUI approach used in the 1983 Apple Lisa and the 1984 Apple Macintosh. The majority of IBM PC users did not replace their command.com shell with a GUI until Windows 95 was released in 1995.

Modern usage as an operating system shell

While most computer users now use a GUI almost exclusively, more advanced users have access to powerful command-line environments:

- IBM OS/2 has the cmd.exe processor. This copies the command.com commands, with extensions to REXX.
- MS Windows users have a CLI environment named Command Prompt, which might use the CScript interface to alternate programs. The new PowerShell program provides a command-line interface, but its applets are not written in shell-script. Implementations of the Unix shell are also available as part of the POSIX sub-system[3], Cygwin, and other software packages.
- Apple Mac OS X[4] and GNU/Linux have the Bash implementation of the Unix shell.
- Embedded Linux (and other embedded Unix-like) devices often use the Ash implementation of the Unix shell, as part of Busybox.
- Android uses a Unix shell derived from Ash[5] with commands from the separate *toolbox*[6].
- Routers with Cisco IOS[7], Junos[8] and many others are commonly configured from the command line.

Scripting

Most command-line interpreters support scripting, to various extents. (They are, after all, interpreters of an interpreted programming language, albeit in many cases the language is unique to the particular command-line interpreter.) They will interpret scripts (variously termed shell scripts or batch files) written in the language that they interpret. Some command-line interpreters also incorporate the interpreter engines of other languages, such as REXX, in addition to their own, allowing the executing of scripts, in those languages, directly within the command-line interpreter itself.

Conversely, scripting programming languages, in particular those with an eval function (such as REXX, Perl, Python, Ruby or Jython), can be used to implement command-line interpreters and filters. For a few operating systems, most notably DOS, such a command interpreter provides a more flexible command line interface than the one supplied. In other cases, such a command interpreter can present a highly customised user interface employing the user interface and input/output facilities of the language.

Other Command-line interfaces

The command line provides an interface between programs as well as the user. In this sense, a command line is an alternative to a dialog. Editors and data-bases present a command line, in which alternate command processors might run. On the other hand, one might have options on the command line which opens a dialog box. The latest version of 'Take Command' has this feature. DBase used a dialog box to construct command lines, which could be further edited before use.

Programs like Basic, Diskpart, Edlin, and QBasic all provide command-line interfaces, some of which use the system shell. Basic is modeled on the default interface for 8-bit Intel computers. Calculators can be run as command-line or dialog interfaces.

There are a number of games from the pre-mouse games, (like Kings Quest 1-3), which relied on the user typing commands at the bottom of the screen. One controls the character by typing commands like 'get ring' or 'look'. The program returns a dialog which describes how the character sees it, or makes the action happen. The game 'h2g2' (hitchhiker's guide to the galaxy, based on Douglas Adam's book of the same name), is a teletype-style command-line game.

The most notable of these interfaces is the STDIO interface, which allows the output of one command to be passed to the input of another. Text files can serve either purpose as well. This provides the interfaces of piping, filters and redirection. Under Unix, devices are files too, so the normal type of file for the shell used for stdin,stdout and stderr is a tty device file.

Another command-line interface allows a shell program to launch helper programs, either to launch documents or start a program. The command is porcessed internally by the shell, and then passed on to another program to launch the document. The graphical interface of Windows and OS/2 rely heavily on command-lines passed through to other programs – console or graphical, which then usually process the command line without presenting a user-console.

Programs like the OS/2 E editor and some other IBM editors, can process command-lines normally meant for the shell, the output being placed directly in the document window.

One should remember that the 'address line' in a web browser is in fact a command line, even to the thing about about:blank interface. One can program extra commands for a web browser, say to open the registry at a given item, by using the reg: keyword.

On the other hand, features like command history and editing are not nesseccarily part of the command-line interface.

Quotes

> Although most users think of the shell as an interactive command interpreter, it is really a programming language in which each statement runs a command. Because it must satisfy both the interactive and programming aspects of command execution, it is a strange language, shaped as much by history as by design.
>
> — Brian Kernighan & Rob Pike [9]

References

[1] RISC OS 3 User Guide, Part 3: Outside the desktop (page 125) Published by Acorn Computers Limited, ISBN 1-85250-124-2 , Edition 2 Part number 0496,075 Issue 1 March 1992

[2] With the exception of ASCII art

[3] "Windows POSIX shell compatibility" (http://technet.microsoft.com/en-us/library/cc754351.aspx). .

[4] via *Finder, Applications, Utilities, Terminal*

[5] "Android sh source" (https://github.com/android/platform_system_core/tree/master/sh). .

[6] "Android toolbox source" (https://github.com/android/platform_system_core/tree/master/toolbox). .

[7] *"The Cisco IOS command-line interface (CLI) is the primary user interface..."* (http://www.cisco.com/en/US/docs/ios/12_0/configfun/configuration/guide/fcui.html#wp7687)

[8] *"...the software that you use whenever you access the router..."* (http://www.juniper.net/techpubs/software/junos/junos55/swconfig55-getting-started/html/cli-overview-getting-started.html)
[9] Brian W. Kernighan and Rob Pike, "The UNIX Programming Environment", *Prentice-Hall* (1984).

External links

- In the Beginning... Was the Command Line (http://books.google.com/books?id=OmnF5MGRNn8C&dq=in+the+beginning+comand+line&pg=PP1&ots=XTN1PFZuJw&sig=S-k5FfNtMXcbbqjqUiqSaMyII8Q&hl=en) — Short book about CLIs by Neal Stephenson.
- The Roots of DOS (http://www.patersontech.com/Dos/Softalk/Softalk.html) David Hunter, *Softalk for the IBM Personal Computer* March 1983. Archived at Patersontech.com since 2000 (http://web.archive.org/web/20001003150623/http://www.patersontech.com/Dos/Softalk/Softalk.html).

Bash (Unix shell)

Bash

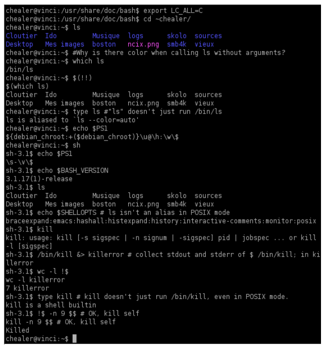

Screenshot of Bash and sh sessions demonstrating some features

Original author(s)	Brian Fox
Initial release	June 7, 1989
Stable release	(February 13, 2011)[1] [2] [±]
Development status	Active
Programming language used	C
Operating system	Cross-platform
Platform	GNU
Translation available	English, multilingual (gettext)
Type	Unix shell
License	GNU General Public License version 3+[3]
Website	www.gnu.org/software/bash/ [4]

Bash is a Unix shell written by Brian Fox for the GNU Project as a free software replacement for the Bourne shell (sh).[5][6] Released in 1989,[7] it has been distributed widely as the shell for the GNU operating system and as the default shell on Linux, Mac OS X and Darwin. It has been ported to Microsoft Windows and distributed with Cygwin and MinGW, to DOS by the DJGPP project, to Novell NetWare and to Android via various terminal emulation applications.

Bash is a command processor, typically run in a text window, allowing the user to type commands which cause actions. Bash can also read commands from a file, called a script. Like all Unix shells, it supports filename wildcarding, piping, here documents, command substitution, variables and control structures for condition-testing and iteration.[8] The keywords, syntax and other basic features of the language were all copied from sh. Other features, e.g., history, were copied from csh and ksh. Bash is a POSIX shell but with a number of extensions.

The name itself is an acronym, a pun and descriptive. As an acronym, it stands for *Bourne-again shell*, referring to its objective as a free replacement for the Bourne shell.[9] As a pun, it expressed that objective in a phrase that sounds similar to *born again*, a term for spiritual rebirth.[10][11] The name is also descriptive of what it did, *bashing together* the features of sh, csh and ksh.[12]

History

Brian Fox began coding Bash on January 10, 1988,[13] after Richard Stallman became dissatisfied with the lack of progress being made by a prior developer.[5] Stallman and the Free Software Foundation (FSF) considered a free shell that could run existing sh scripts so strategic to a completely free system built from BSD and GNU code that this was one of the few projects they funded themselves, with Fox undertaking the work as an employee of FSF.[5][14] Fox released Bash as a beta, version .99, on June 7, 1989[7] and remained the primary maintainer until sometime between mid-1992[15] and mid-1994,[16] when he was laid off from FSF[17] and his responsibility was transitioned to another early contributor, Chet Ramey.[18][19][20]

Features

The Bash command syntax is a superset of the Bourne shell command syntax. The vast majority of Bourne shell scripts can be executed by Bash without modification, with the exception of Bourne shell scripts stumbling into fringe syntax behavior interpreted differently in Bash or attempting to run a system command matching a newer Bash builtin, etc. Bash command syntax includes ideas drawn from the Korn shell (ksh) and the C shell (csh) such as command line editing, command history, the directory stack, the `$RANDOM` and `$PPID` variables, and POSIX command substitution syntax `$(…)`. When used as an interactive command shell and pressing the tab key, Bash automatically uses command line completion to match partly typed program names, filenames and variable names. The Bash command-line completion system is very flexible and customizable, and is often packaged with functions that complete arguments and filenames for specific programs and tasks.

Bash's syntax has many extensions which the Bourne shell lacks. Bash can perform integer calculations without spawning external processes, unlike the Bourne shell. Bash uses the `((…))` command and the `$((…))` variable syntax for this purpose. Bash syntax simplifies I/O redirection in ways that are not possible in the traditional Bourne shell. For example, Bash can redirect standard output (stdout) and standard error (stderr) at the same time using the `&>` operator. This is simpler to type than the Bourne shell equivalent '`command > file 2>&1`'. Bash supports process substitution using the `<(command)` syntax, which substitutes the output of (or input to) a command where a filename is normally used.

When using the 'function' keyword, Bash function declarations are not compatible with Bourne/Korn/POSIX scripts (the Korn shell has the same problem when using 'function'), but Bash accepts the same function declaration syntax as the Bourne and Korn shells, and is POSIX conformant. Due to these and other differences, Bash shell scripts are rarely runnable under the Bourne or Korn shell interpreters unless deliberately written with that compatibility in mind, which is becoming less common as Linux becomes more widespread. But in POSIX mode,[21] Bash conformance with POSIX is nearly perfect.

Bash supports here documents just as the Bourne shell always has. In addition, since version 2.05b Bash can redirect standard input (stdin) from a "here string" using the `<<<` operator.

Bash 3.0 supports in-process regular expression matching using a syntax reminiscent of Perl.[22]

Bash 4.0 introduced support for associative arrays.[21][23] Associative arrays allow a fake support for multi-dimensional (indexed) arrays, in a similar way to AWK:

```
declare -A a          # declare an associative array 'a' faking a
bi-dimensional indexed array
i=1; j=2              # initialize some indices
a[$i,$j]=5            # associate value "5" to key "$i,$j" (i.e. "1,2")
echo ${a[$i,$j]}      # print the stored value at key "$i,$j"
```

Brace expansion

Brace expansion, also called alternation, is a feature copied from the C shell that generates a set of alternative combinations. The generated results need not exist as files. The results of each expanded string are not sorted and left to right order is preserved:

```
echo a{p,c,d,b}e # ape ace ade abe
echo {a,b,c}{d,e,f} # ad ae af bd be bf cd ce cf
```

Brace expansions should not be used in portable shell scripts, because the Bourne shell will not produce the same output.

```
#!/bin/sh

# A traditional shell does not produce the same output
echo a{p,c,d,b}e # a{p,c,d,b}e
```

When brace expansion is combined with wildcards, the braces are expanded first, and then the resulting wildcards are substituted normally. Hence, a listing of JPEG and PNG images in the current directory could be obtained with:

```
ls *.{jpg,jpeg,png}    # expands to *.jpg *.jpeg *.png - after which,
                       # the wildcards are processed
```

Startup scripts

When Bash starts, it executes the commands in a variety of different scripts.

When started as an interactive login shell:

- Bash reads and executes the `/etc/profile` (if it exists).
  - *(Not infrequently, /etc/profile calls /etc/bash.bashrc.)*
- After reading that file, it looks for `~/.bash_profile`, `~/.bash_login`, and `~/.profile` *in that order*, and reads and executes the first one (that exists and is readable).

When a login shell exits:

- Bash reads and executes `~/.bash_logout` (if it exists).

When started as an interactive shell (but not a login shell):

- Bash reads and executes `~/.bashrc` (if it exists). This may be inhibited by using the `--norc` option. The `--rcfile file` option will force Bash to read and execute commands from `file` instead of `~/.bashrc`.

Portability

Invoking Bash with the `--posix` option or stating `set -o posix` in a script causes Bash to conform very closely to the POSIX 1003.2 standard.[24] Bash shell scripts intended for portability should at least take into account the Bourne shell it intends to replace. Bash has certain features that the traditional Bourne shell lacks. Among these are[24]:

- Certain extended invocation options
- Command substitution using $( ) notation (this feature is part of the POSIX 1003.2 standard though)
- Brace expansion
- Certain array operations, and associative arrays
- The double brackets extended test construct
- The double-parentheses arithmetic-evaluation construct
- Certain string manipulation operations

- Process substitution
- A Regular Expression matching operator
- Bash-specific builtins
- Coprocesses

Keyboard shortcuts

The following shortcuts work when using default (Emacs) key bindings. Vi-bindings can be enabled by running `set -o vi`.[25]

Note: For shortcuts involving `Alt`, you may be able to use `Esc` instead.

Note: Sometimes, you must use `Esc` instead of `Alt`, because the `Alt` shortcut conflicts with another shortcut. For example, in Trisquel 5.0 (a distribution of Linux), pressing `Alt+f` will not move the cursor forward one word, but will activate "File" in the menu of the terminal window.

- `Tab` ⇥ : Autocompletes from the cursor position.
- `Ctrl+a` : Moves the cursor to the line start (equivalent to the key `Home`).
- `Ctrl+b` : Moves the cursor back one character (equivalent to the key ←).
- `Ctrl+c` : Sends the signal SIGINT to the current task, which aborts and closes it.
- `Ctrl+d`
  - Sends an EOF marker, which (unless disabled by an option) closes the current shell (equivalent to the command `exit`). (Only if there is no text on the current line)
  - If there is text on the current line, deletes the current character (then equivalent to the key `Delete`).
- `Ctrl+e` : (end) moves the cursor to the line end (equivalent to the key `End`).
- `Ctrl+f` : Moves the cursor forward one character (equivalent to the key →).
- `Ctrl+g` : Abort the research and restore the original line.
- `Ctrl+h` : Deletes the previous character (same as backspace).
- `Ctrl+i` : Equivalent to the tab key.
- `Ctrl+j` : Equivalent to the enter key.
- `Ctrl+k` : Clears the line content after the cursor and copies it into the clipboard.
- `Ctrl+l` : Clears the screen content (equivalent to the command `clear`).
- `Ctrl+n` : (next) recalls the next command (equivalent to the key ↓).
- `Ctrl+o` : Executes the found command from history, and fetch the next line relative to the current line from the history for editing.
- `Ctrl+p` : (previous) recalls the prior command (equivalent to the key ↑).
- `Ctrl+r` : (research) recalls the last command including the specified character(s). A second `Ctrl+r` recalls the next anterior command which corresponds to the research
- `Ctrl+s` : Go back to the next more recent command of the research (beware to not execute it from a terminal because this command also launches its XOFF). If you changed that XOFF setting, use `Ctrl+q` to return.
- `Ctrl+t` : Transpose the previous two characters.
- `Ctrl+u` : Clears the line content before the cursor and copies it into the clipboard.
- `Ctrl+v` : If the next input is also a control sequence, type it literally (e. g. * `Ctrl+v` `Ctrl+h` types "^H", a literal backspace.)
- `Ctrl+w` : Clears the word before the cursor and copies it into the clipboard.
- `Ctrl+x` `Ctrl+e` : Edits the current line in the $EDITOR program, or vi if undefined.
- `Ctrl+x` `Ctrl+r` : Read in the contents of the inputrc file, and incorporate any bindings or variable assignments found there.
- `Ctrl+x` `Ctrl+u` : Incremental undo, separately remembered for each line.
- `Ctrl+x` `Ctrl+v` : Display version information about the current instance of bash.

- `Ctrl+x Ctrl+x` : Alternates the cursor with its old position. (C-x, because x has a crossing shape).
- `Ctrl+y` : (yank) adds the clipboard content from the cursor position.
- `Ctrl+z` : Sends the signal SIGTSTP to the current task, which suspends it. To execute it in background one can enter `bg`. To bring it back from background or suspension `fg ['process name or job id']` (foreground) can be issued.
- `Ctrl+_` : Incremental undo, separately remembered for each line.
- `Alt+b` : (backward) moves the cursor backward one word.
- `Alt+c` : Capitalizes the character under the cursor and moves to the end of the word.
- `Alt+d` : Cuts the word after the cursor.
- `Alt+f` : (forward) moves the cursor forward one word.
- `Alt+l` : Lowers the case of every character from the cursor's position to the end of the current word.
- `Alt+r` : Cancels the changes and puts back the line as it was in the history.
- `Alt+u` : Capitalizes every character from the cursor's position to the end of the current word.
- `Alt+.` : Insert the last argument to the previous command (the last word of the previous history entry).

References

[1] Ramey, Chet (2011-02-16). "Bash-4.2 available for FTP" (http://www.mail-archive.com/info-gnu@gnu.org/msg01139.html). *info-gnu mailing list*. . Retrieved 2011-02-20.

[2] http://en.wikipedia.org/wiki/Template%3Alatest_stable_software_release%2Fbash?action=edit&preload=Template:LSR/syntax

[3] GNU Project. "README file" (http://www.gnu.org/software/bash/). . "Bash is free software, distributed under the terms of the [GNU] General Public License as published by the Free Software Foundation, version 3 of the License (or any later version)."

[4] http://www.gnu.org/software/bash/

[5] Richard Stallman (forwarded with comments by Chet Ramey) (February 10, 1988). "GNU + BSD = ?". [news:comp.unix.questions comp.unix.questions]. (Web link) (http://groups.google.com/group/comp.unix.questions/msg/e44ad620f76be7c9?hl=en). "For a year and a half, the GNU shell was "just about done". The author made repeated promises to deliver what he had done, and never kept them. Finally I could no longer believe he would ever deliver anything. So Foundation staff member Brian Fox is now implementing an imitation of the Bourne shell.". Retrieved Mar 22, 2011.

[6] [|Hamilton, Naomi (http://www.computerworld.com.au/author/97635374/naomi-hamilton/articles)] (May 30, 2008), "The A-Z of Programming Languages: BASH/Bourne-Again Shell" (http://www.computerworld.com.au/article/222764/a-z_programming_languages_bash_bourne-again_shell/?pp=2&fp=16&fpid=1), *Computerworld*: 2, , retrieved Mar 21, 2011, "When Richard Stallman decided to create a full replacement for the then-encumbered Unix systems, he knew that he would eventually have to have replacements for all of the common utilities, especially the standard shell, and those replacements would have to have acceptable licensing."

[7] Brian Fox (forwarded by Leonard H. Tower Jr.) (June 8, 1989). "Bash is in beta release!". [news:gnu.announce gnu.announce]. (Web link) (http://groups.google.com/group/gnu.announce/msg/a509f48ffb298c35?hl=en). Retrieved Oct 28 2010.

[8] Bourne, S.R. (July/August 1978). "The UNIX Shell". *The Bell System Technical Journal* (Short Hills, NJ: American Telephone and Telegraph Company) **57** (6): 1971–1990. ISSN 00005-8580. "The lines between `<<!` and `!` are called a *here* document; they are read by the shell and made available as the standard input."

[9] C Programming (http://www.ddj.com/cpp/184404693) by Al Stevens, Dr. Dobb's Journal, July 1, 2001

[10] Richard Stallman (Nov 12, 2010). "About the GNU Project" (http://www.gnu.org/gnu/thegnuproject.html). Free Software Foundation. . Retrieved Mar 13, 2011. ""Bourne Again Shell" is a play on the name "Bourne Shell", which was the usual shell on Unix."

[11] Gattol, Markus (Mar 13, 2011), *Bourne-again Shell* (http://www.markus-gattol.name/ws/bash.html), , retrieved Mar 13, 2011, "The name is a pun on the name of the Bourne shell (sh), an early and important Unix shell written by Stephen Bourne and distributed with Version 7 Unix circa 1978, and the concept of being "born again"."

[12] Ian Darwin (June 13, 1989). "at&t-free ksh (was: job control is a bug, not a feature)". [news:comp.os.minix comp.os.minix]. (Web link) (http://groups.google.com/group/comp.os.minix/msg/63c036d82ceca4d6?hl=en). "Yup, the gnu project's Born Again Shell ("bash") is an attempt at bashing all the features of sh together with many of those from both csh and ksh.". Retrieved Mar 21, 2011.

[13] Brian Fox (August 29, 1996), *shell.c* (http://ftp.gnu.org/gnu/bash/bash-1.14.7.tar.gz), Free Software Foundation, , "Birthdate: Sunday, January 10th, 1988. Initial author: Brian Fox"

[14] Richard Stallman (October 3, 2010). "About the GNU Project" (http://www.gnu.org/gnu/thegnuproject.html). Free Software Foundation. . Retrieved Mar 21, 2011. "Free Software Foundation employees have written and maintained a number of GNU software packages. Two notable ones are the C library and the shell. ... We funded development of these programs because the GNU Project was not just about tools or a development environment. Our goal was a complete operating system, and these programs were needed for that goal."

[15] len (g...@prep.ai.mit.edu) (April 20, 1993). "[news:gnusenet930421bulletin@prep.ai.mit.edu January 1993 GNU's Bulletin]". [news:gnu.announce gnu.announce]. (Web link) (http://groups.google.com/group/gnu.misc.discuss/msg/4f42c739cd7e8bd8). Retrieved Oct 28 2010.

[16] Ramey, Chet (1994-08-01). "Bash - the GNU shell (Reflections and Lessons Learned)" (http://www.linuxjournal.com/article/2800#N0xa50890.0xb46380). *Linux Journal*. . Retrieved 2008-11-13.

[17] Chet Ramey (October 31, 2010), *Dates in your Computerworld interview* (http://www.scribd.com/doc/40556434/2010-10-31-Chet-Ramey-Early-Bash-Dates), , retrieved Oct 31 2010

[18] Chet Ramey (June 12, 1989). "Bash 0.99 fixes & improvements". [news:gnu.bash.bug gnu.bash.bug]. (Web link) (http://groups.google.com/group/gnu.bash.bug/msg/1fc7b688f5d44438?hl=en). Retrieved Nov 1 2010.

[19] Chet Ramey (July 24, 1989). "Some bash-1.02 fixes". [news:gnu.bash.bug gnu.bash.bug]. (Web link) (http://groups.google.com/group/gnu.bash.bug/msg/072a03645663caea?hl=en). Retrieved Oct 30 2010.

[20] Brian Fox (March 2, 1990). "Availability of bash 1.05". [news:gnu.bash.bug gnu.bash.bug]. (Web link) (http://groups.google.com/group/gnu.bash.bug/msg/e6112ccc8866e2f4?hl=en). Retrieved Oct 30 2010.

[21] "6.11 Bash POSIX Mode" (http://www.gnu.org/software/bash/manual/html_node/Bash-POSIX-Mode.html), *The GNU Bash Reference Manual, for Bash, Version 4.1* (http://www.gnu.org/software/bash/manual/html_node/index.html), December 23, 2009, , retrieved Oct 26 2010

[22] The syntax matches that shown on the `regex(7)` (http://www.tin.org/bin/man.cgi?section=7&topic=regex) man page.

[23] "The shell provides associative array variables, with the appropriate support to create, delete, assign values to, and expand them." http://tiswww.case.edu/php/chet/bash/NEWS

[24] Mendel Cooper. "Portability Issues" (http://tldp.org/LDP/abs/html/portabilityissues.html). *The Linux Documentation Project*. ibiblio.org. .

[25] http://www.hypexr.org/bash_tutorial.php#emacs

External links

- Official website (http://www.gnu.org/software/bash/bash.html)
- Bash Reference Manual (http://www.gnu.org/software/bash/manual/) (HTML (http://www.gnu.org/software/bash/manual/html_node/index.html) PS (http://www.gnu.org/software/bash/manual/bash.ps.gz) PDF (http://www.gnu.org/software/bash/manual/bash.pdf))
- Bash Guide for Beginners (http://tldp.org/LDP/Bash-Beginners-Guide/html/index.html) article at The Linux Documentation Project
- Advanced Bash Scripting Guide (http://tldp.org/LDP/abs/html/index.html) article at The Linux Documentation Project
- Linux Shell Scripting Tutorial (LSST) v2.0 wiki (http://bash.cyberciti.biz/guide/Main_Page)
- "The Comprehensive List of bash Reference Documentation and Examples" (http://www.bashcookbook.com/bashinfo/)
- Useful Bash History Tips and Tricks (http://spsneo.com/blog/2009/09/19/bash-history-tips-and-tricks/)
- 2008 interview with GNU Bash's maintainer, Chet Ramey (http://www.computerworld.com.au/index.php/id;1591223321;fp;16;fpid;1;pf;1)
- Working with BASH environment variables (http://www.geeksww.com/tutorials/operating_systems/linux/tips_and_tricks/working_bash_environment_variables_beginners_linux.php)
- Video Tutorial for Creating a Bash Script (http://www.galatech.co.uk/index.php/tutorials/viewpost/100)
- Bash commands and examples (http://www.shell-fu.org/lister.php?tag=bash)
- Colorized Bash prompt (http://www.markus-gattol.name/ws/bash.html#colorized_shell_prompt) - how to set up a colorized Bash prompt based on the current connection method (SSH, telnet, etc.).
- jBash Project is a Java Parser for the Bourne Again Shell (http://code.google.com/p/jbash/)
- The 'official' channel FAQ for freenode's #bash channel is BashFAQ. (http://mywiki.wooledge.org/EnglishFrontPage)
- BASHDB - Bash with a built-in debugger. (http://bashdb.sourceforge.net/)
- Bash Quick Reference card (http://www.digilife.be/quickreferences/QRC/Bash Quick Reference.pdf)
- List of Bash commands and environment variables (Built-in Commands (http://mitblog.com/posts/bash-built-in-commands/) External Commands (http://mitblog.com/posts/bash-shell-external-commands/) Environment Variables (http://mitblog.com/posts/bash-shell-environment-variables/))

Batch file

Batch file

Filename extension	`.bat .cmd .btm`
Type of format	Scripting
Container for	Shell scripts

In DOS, OS/2, and also Microsoft Windows, **batch file** is the name given to a type of script file, a text file containing a series of commands to be executed by the command interpreter.

The commands may be built into the command processor (`COPY`), supplied with the operating system but not built into it (`XCOPY` invokes the Microsoft DOS program `XCOPY.EXE`), or may be any program (`cp` invokes the program `cp.exe` if present, an .EXE port of the Unix `cp` command, with essentially the same functionality as `XCOPY.EXE`).

Similar to job control language and other systems on mainframe and minicomputer systems, batch files were added to ease the work required for certain regular tasks by allowing the user to set up a script to automate them. When a batch file is run, the shell program (usually COMMAND.COM or cmd.exe) reads the file and executes its commands, normally line-by-line.[1] Unix-like operating systems (such as Linux) have a similar type of file called a shell script.[2]

The filename extension **.bat** was used in DOS, and the Windows 9x family of operating systems. The Microsoft Windows NT-family of operating systems and OS/2 added **.cmd**. Batch files for other environments may have different extensions, e.g. **.btm** in 4DOS and 4NT related shells.

There have been changes to the detailed handling of batch files; some of the detail in this article is applicable to all batch files, while other details apply only to certain versions.

Variants

Brief information on the function and parameters of commands are usually displayed by typing the command at the command prompt followed by "/?" and pressing the Enter key. In some cases "-?", "?", or just the command name without parameters (if parameters are required) will also elicit information. Some commands ported from Unix require "--help".

DOS

In MS-DOS, a batch file can be started from the command line by typing its name followed by any required parameters and pressing the "enter" key. When MS-DOS loads, the file AUTOEXEC.BAT is automatically executed, so any commands that need to be run to set up the MS-DOS environment for use could be placed in this file. Computer users would have the autoexec file set up the system date and time, initialize the MS-DOS environment, load any resident programs or device drivers, or initialize network connections and assignments.

In MS-DOS, the extension ".BAT" identified a file containing commands which could be executed by the command interpreter COMMAND.COM line by line as if it was a list of commands to be entered, with some extra batch-file-specific commands for basic programming functionality, including a `GOTO` command for changing flow of line execution.

Early Windows

Microsoft Windows was introduced in 1985 as a GUI Operating System alternative to text-based operating and was designed to run on MS-DOS. In order to start it, the `WIN` command was used and could be added to the end of the AUTOEXEC.BAT file to allow automatic loading of Windows. In the earlier versions one could run a .bat type file from Windows in the MS-DOS Prompt.

Windows was run from MS-DOS and used COMMAND.COM to run .bat files on the following operating systems:

- Windows 1, 2 and 3.
- Windows 95 and 98.
- Windows ME (access to real mode MS-DOS was restricted).

OS/2

The IBM OS/2 operating system supported DOS-style batch files. It also included a version of REXX, which was a more advanced batch-file scripting language. IBM and Microsoft started developing this system, but during the construction of it broke up after a dispute; as a result of this, IBM referred to their MS-DOS-like console shell without mention of Microsoft, naming it just DOS, although this seemingly made no difference on the way batch files worked from COMMAND.COM.

OS/2's batch file interpreter also supports an EXTPROC command. This passes the batch file to the program named on the EXTPROC file as a data file. The named program can be a script file; this is similar to the #! mechanism

Windows NT

Windows versions other than the NT line of operating systems were run from MS-DOS and used the same command interpreter, COMMAND.COM, to execute batch files. However, the operating systems in the Windows NT series run directly from booting the hard drive; they are true operating systems, not graphical user interfaces for underlying MS-DOS. An enhanced 32-bit command processor, cmd.exe, was introduced; it could execute scripts with either the .CMD or .BAT extension. Cmd.exe added additional commands, and implemented existing ones in a slightly different way, so that the same batch file (with different extension) might work differently with cmd.exe and COMMAND.COM. In most cases operation is identical if the few unsupported commands are not used. Cmd.exe's extensions to COMMAND.COM can be disabled for compatibility.

Microsoft released a version of cmd.exe for Windows 9x and ME called WIN95CMD to allow users of older versions of Windows to use certain cmd.exe-style batch files.

As of Windows 8, cmd.exe is the normal command interpreter for batch files; the older COMMAND.COM can be run from within a cmd.exe window in 32-bit versions of Windows able to run 16-bit programs[3].

Filename extensions

- **.bat**: The first extension used by Microsoft for batch files. This extension runs with MS-DOS and all versions of Windows, under COMMAND.COM or cmd.exe, despite the different ways the two command interpreters execute batch files.
- **.cmd**: The extension used by operating systems in the Windows NT family and sent to cmd.exe for interpretation. It does not work on computers relying on COMMAND.COM so prevents cmd.exe scripts from being executed in the wrong Windows environment. It is also used by IBM's OS/2 for batch files.
- **.btm**: The extension used by 4DOS and 4NT. The scripts that run on 4DOS and 4NT are faster, especially with longer ones, as the script is loaded entirely ready for execution, rather than line-by-line.[4]

Differences between .cmd and .bat execution in the Windows NT family

The only known difference between .cmd and .bat file execution is that in a .cmd file the `ERRORLEVEL` variable changes even on a successful command that is affected by Command Extensions (when Command Extensions are enabled), whereas in .bat files the `ERRORLEVEL` variable changes only upon errors.[5]

Example

This example batch file displays "Hello World!", prompts and waits for the user to press a key, and terminates.

```
@ECHO off
ECHO Hello World!
PAUSE
```

To execute the file it must be saved with the extension .bat (or .cmd for Windows-NT type operating systems) in plain text format, typically created by using a text editor such as Notepad or a word processor in text mode.

Result

When executed (either from Windows Explorer or Command Prompt) this is displayed:

```
Hello World!
Press any key to continue . . .
```

Explanation

The interpreter executes each line in turn, starting with the first. The `@` symbol at the start of the line turns off the prompt from displaying that command. The command `ECHO off` turns off the prompt permanently, or until it is turned on again. Then the next line is executed, the `ECHO Hello World!` command outputs `Hello World!`, as only `off` and `on` have special functions. Then the next line is executed, the `PAUSE` command displays `Press any key to continue . . .` and pauses the script's execution until a key is pressed, when the script terminates as there are no more commands. In Windows, if the script is run within a Command Prompt window, the window remains open at the prompt as in MS-DOS, otherwise the command prompt windows closes on termination (unless the batch file has a command to prevent this).

Advanced Windows batch example - conditional shutdown

```
@echo off
color 0A
title Conditional Shutdown.

:start
echo Welcome, %USERNAME%
echo What would you like to do?
echo.
echo 1. Shutdown in specified time
echo 2. Shutdown at a specified time
echo 3. Shutdown now
echo 4. Restart now
echo 5. Log off now
echo 6. Hibernate now
echo.
```

```
echo 0. Quit
echo.

set /p choice="Enter your choice: "
if "%choice%"=="1" goto shutdown
if "%choice%"=="2" goto shutdown-clock
if "%choice%"=="3" shutdown.exe -s -f
if "%choice%"=="4" shutdown.exe -r -f
if "%choice%"=="5" shutdown.exe -l -f
if "%choice%"=="6" shutdown.exe -h -f
if "%choice%"=="0" exit
echo Invalid choice: %choice%
echo.
pause
cls
goto start

:shutdown
cls
set /p sec="Minutes until shutdown: "
set /a min=60*%sec%
shutdown.exe -s -f -t %min%
echo Shutdown initiated at %time%
echo.
goto cancel

:shutdown-clock
echo.
echo the time format is HH:MM:SS (24 hour time)
echo example: 14:30:00 for 2:30 PM
echo.
set /p tmg=enter the time that you wish the computer to shutdown on:
schtasks.exe /create /sc ONCE /tn shutdown /st %tmg% /tr "shutdown.exe -s -t 00"
echo shutdown initiated at %tmg%
echo.

:cancel
set /p cancel="Type cancel to stop shutdown: "
if not "%cancel%"=="cancel" exit
shutdown.exe -a
cls
schtasks.exe /end /tn shutdown
cls
schtasks.exe /delete /tn shutdown
cls
echo Shutdown is cancelled.
echo.
```

```
pause
exit
```

When doing conditions with IF command, batch commands can use:

```
  EQU : Equal (=)
  NEQ : Not equal (≠)
```

```
  LSS : Less than (<)
  LEQ : Less than or Equal (≤)
```

```
GTR : Greater than (>)
  GEQ : Greater than or Equal (≥)
```

Limitations and exceptions

Null values in variables

Variable expansions are substituted textually into the command, and thus variables which contain nothing simply disappear from the syntax, and variables which contain spaces turn into multiple tokens. This leads to syntax errors or bugs.

For example:

```
IF %foo%==bar ECHO Equal
```

if %foo% is empty, parses as the erroneous construct:

```
IF ==bar ECHO Equal
```

and if %foo% contains "abc def", then the syntax is also wrong:

```
IF abc def==bar ECHO Equal
```

The usual way to prevent this problem is to surround variable expansions in quotes so that an empty variable expands into the valid expression `IF ""=="bar"` instead of the invalid `IF ==bar`. The text that is being compared to the variable must also be enclosed in quotes, because the quotes are not special delimiting syntax; these characters represent themselves.

```
IF "%foo%"=="bar" ECHO Equal
```

The delayed !VARIABLE! expansion available in Windows 2000/XP/Vista/7 may be used to avoid these syntactical errors. In this case, null or multi-word variables will not fail syntactically because the value will be expanded after the IF command is parsed:

```
IF !foo!==bar ECHO Equal
```

Quotation marks and spaces in passed strings

- For some commands, spaces are treated as delimiters in commands, unless those spaces are enclosed by quotation marks. A single quotation mark (") is not included as part of the string. However, an escaped quotation mark (""") can be part of the string.
- For other commands, spaces are not treated as delimiters and do not need quotation marks. If quotes are included they become part of the string.

This can cause conflicts where a string contains quotation marks, and is to be inserted into another line of text that must also be enclosed in quotation marks:

```
C:\> Set foo="this string is enclosed in quotation marks"

C:\> Echo "test 1 %foo%"
"test 1 "this string is enclosed in quotation marks""

C:\> Eventcreate /T Warning /ID 1 /L System /SO "Source" /D "Example: %foo%"
ERROR: Invalid Argument/Option - 'string'.
Type "EVENTCREATE /?" for usage.
```

Under Windows 2000/XP/Vista/7, the solution is to replace all occurrences of one quote characters by three quotes:

```
C:\> Set foo="this string is enclosed in quotes"

C:\> Set foo=%foo:"="""%

C:\> Echo "test 1 %foo%"
"test 1 """this string is enclosed in quotes""""

C:\> Eventcreate /T Warning /ID 1 /L System /SO "Source" /D "Example: %foo%"
SUCCESS: A 'Warning' type event is created in the 'Source' log/source.
```

Escaped characters in strings

Some characters have special meaning to the command line, such as the pipe | character. These cannot be printed as text using the ECHO command unless escaped using the caret ^ symbol:

```
C:\> Echo foo | bar
'bar' is not recognized as an internal or external command,
operable program or batch file.

C:\> Echo foo ^| bar
foo | bar
```

However, escaping does not work as expected when inserting the escaped character into an environment variable, and the variable ends up containing a live pipe command when merely echoed. It is necessary to escape both the caret itself and the escaped character for the character display as text in the variable:

```
C:\> set foo=bar | baz
'baz' is not recognized as an internal or external command,
operable program or batch file.

C:\> set foo=bar ^| baz
```

```
C:\> echo %foo%
'baz' is not recognized as an internal or external command,
operable program or batch file.

C:\> set foo=bar ^^^| baz
C:\> echo %foo%
bar | baz
```

The delayed !VARIABLE! expansion available in Windows 2000/XP/Vista/7 may be used to show special characters stored in environment variables because the variable value will be expanded after the command was parsed:

```
C:\> set foo=bar ^| baz
C:\> echo !foo!
bar | baz
```

Sleep or scripted delay

The PAUSE command halts script activity indefinitely until a key is pressed; small programs and workarounds were written to implement a timed pause[6]. Many workarounds using scripting commands only worked in some environments, depending upon the `CHOICE` function not available in older command interpreters, `PING` only available if TCP/IP was installed, and so on. Simple small programs were readily available; a typical example is the 94-byte WAIT.COM[7] executable; `WAIT 5` would wait for 5 seconds, then return control to the script. Most such programs are 16-bit .COM files incompatible with 64-bit Windows, but are not needed since Windows Vista and later introduced the TIMEOUT command.

Text output with stripped CR/LF

Normally all printed text automatically has the control characters for "carriage return" and "line feed" appended to the end of each line.

```
batchtest.bat:
@echo foo
@echo bar

C:\>batchtest.bat
foo
bar
```

It does not matter if the two echo commands share the same command line; the CR/LF codes are inserted to break the output onto separate lines:

```
C:\> @echo foo&@echo bar
foo
bar
```

A trick discovered with Windows 2000/XP/Vista/7 is to use the special prompt for input to output text without CR/LF trailing the text. In this example, the CR/LF does not follow Line 1, but does follow Line 2 and Line 3:

```
batchtest.bat:
@echo off
set foo=Line 1
```

```
echo y | set /p tmp="%foo%"
echo Line 2
echo Line 3

C:\>batchtest.bat
Line 1Line 2
Line 3

C:\>
```

This can be used to output data to a text file without CR/LF appended to the end:

```
C:\> echo y | set /p tmp="Line 1"> data.txt
C:\> echo y | set /p tmp="Line 2">> data.txt
C:\> echo y | set /p tmp="Line 3">> data.txt
C:\> type data.txt
Line 1Line 2Line 3
```

However, there is no way to inject this stripped CR/LF prompt output directly into an environment variable.

Setting a UNC working directory from a shortcut

It is not possible to have a command prompt that uses a UNC file path as the current working directory, like this:

```
\\server\share\directory\>
```

The command prompt requires the use of drive letters to assign a working directory, which makes running complex batch files stored on a server UNC share more difficult. While a batch file can be run from a UNC file path, the working directory will default to "C:\windows\system32\"

In Windows 2000/XP/Vista/7, a workaround is to use the **PUSHD** and **POPD** command with command extensions. Quoting the help for PUSHD in Windows 7, *If Command Extensions are enabled the PUSHD command accepts network paths in addition to the normal drive letter and path. If a network path is specified, PUSHD will create a temporary drive letter that points to that specified network resource and then change the current drive and directory, using the newly defined drive letter. Temporary drive letters are allocated from Z: on down, using the first unused drive letter found.*

If not enabled by default, command extensions can be temporarily enabled using the "/E:ON" switch for the command interpreter.

So to run a batch file on a UNC share, assign a temporary drive letter to the UNC share, and use the UNC share as the working directory of the batch file, a Windows shortcut can be constructed that looks like this:

```
Target: %COMSPEC% /E:ON /C "PUSHD """\\SERVER\SHARE\DIR1\DIR2\""" & BATCHFILE.BAT & POPD"
```

The working directory attribute of this shortcut is ignored.

The following syntax does correctly expand to the path of the current batch script.

```
 %~dp0
```

Character set

Batch files use a DOS character set, as defined by the computer, e.g. Code page 437. The non-ASCII parts of these are incompatible with the Unicode or Windows character sets otherwise used in Windows so care needs to be taken.[8] Non-English file names work only if entered through a DOS character set compatible editor. File names with characters outside this set won't work in batch files.

To get output in Unicode into file pipes from an internal command such as `dir`, one can use the `cmd /U` command. For example `cmd /U /C dir > files.txt` will create a file containing a directory listing with correct Windows characters, in the UTF-16LE encoding.

Other Windows scripting languages

The cmd.exe command processor that interprets .cmd files is supported in all 32- and 64-bit versions of Windows up to at least the 2011 Windows 8 preview. COMMAND.EXE, which interprets .BAT files, was supported in all 16- and 32-bit versions up to at least Windows 8 preview.[9]

While the more powerful 2006 Windows PowerShell is favored in later versions of Windows that support it, Microsoft was also using .cmd files as far as, at least, Window Server 2008. An example is *servermanagercmd.exe*[10] which incorporates the entire set of Server Manager functions for Windows Server 2008.

There are other, later and more powerful, scripting languages produced by Microsoft for Windows:

- KiXtart (**.kix**) - developed by a Microsoft employee in 1991, specifically to meet the need for commands useful in a network logon script while retaining the simple 'feel' of a .cmd file.
- Windows Script Host (**.vbs** and **.js**) - released in 1998, (consisting of cscript.exe and wscript.exe) runs scripts written in VBScript or JScript. It can run them in windowed mode (with the wscript.exe host) or in console-based mode (with the cscript.exe host). They have been a part of Windows since Windows 98.
- Windows PowerShell (**.ps1**) - released in 2006 by Microsoft and can operate with Windows XP (SP2/SP3) and later versions. PowerShell can operate both interactively (from a command-line interface) and also via saved scripts, and has a strong resemblance to Unix shells.[11]

Cross-platform scripting tools including Perl, Python, Ruby and Rexx are available for Windows.

Script files will run if the filename without extension is entered. There are rules of precedence governing interpretation of, say `DoThis` if several of `DoThis.cmd`, `DoThis.bat`, `DoThis.exe`, etc. exist; by default `DoThis.com` has highest priority. This default order may be modified in newer operating systems by the user-settable PATHEXT environment variable.

References

[1] http://technet.microsoft.com/en-us/library/cc758944(WS.10).aspx
[2] http://www.file-extensions.org/bat-file-extension
[3] To verify that COMMAND.COM remains available (in the \WINDOWS\SYSTEM32 directory), type "COMMAND.COM" in the 32-bit Windows 7 command prompt.
[4] http://www.cryer.co.uk/filetypes/b/btm.htm
[5] http://groups.google.com/group/microsoft.public.win2000.cmdprompt.admin/msg/ad9066638815812c
[6] "How to do a delay" (http://www.ericphelps.com/batch/samples/sleep.txt), ericphelps.com
[7] Utilities for DOS, linking to WAIT.ZIP (archive of WAIT.COM) and other programs (http://www.uwe-sieber.de/util_e.html)
[8] "Keep your eye on the code page" (http://blogs.msdn.com/b/oldnewthing/archive/2005/03/08/389527.aspx), Raymond Chen, Microsoft
[9] Availability of cmd.exe and COMMAND.COM can be confirmed by invoking them in any version of WIndows (COMMAND.COM not in 64-bit versions; probably only available in Win 8/32 if installed with option to support 16-bit programs)
[10] http://www.winsupersite.com/showcase/win2008_cli.asp
[11] http://geekswithblogs.net/sdorman/archive/2006/06/18/82258.aspx

External links

- Microsoft Windows XP Batch file reference (http://www.microsoft.com/resources/documentation/windows/xp/all/proddocs/en-us/batch.mspx?mfr=true)
- How Windows batch files work (http://commandwindows.com/batch.htm)
- Windows batch file command overview (http://www.ss64.com/nt/)
- FreeDOS' FreeCOM : complete feature list (http://www.freedos.org/freecom/)
- MS-DOS+Win../95/98/ME batch programming links (http://www.netikka.net/tsneti/http/tsnetihttpprog.php#batch)
- Windows Command Line Interface script programming links (http://www.netikka.net/tsneti/http/tsnetihttpprog.php#cmdscript)
- scripting related information (also command line) (http://www.robvanderwoude.com/)

Windows PowerShell

Windows PowerShell

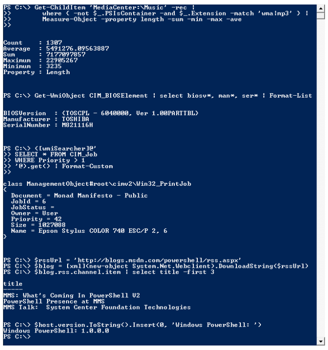

Screenshot of a PowerShell session

Developer(s)	Microsoft Corporation
Initial release	November 14, 2006
Stable release	2.0 / July 22, 2009
Preview release	3.0 / 19 September 2011
Development status	Active
Operating system	Windows XP Windows Server 2003 Windows Vista Windows Server 2008 Windows 7 Windows Server 2008 R2
Platform	.NET Framework
Type	Operating system shell
License	Proprietary software; a component of Microsoft Windows
Website	www.microsoft.com/powershell [1]

PowerShell

Paradigm(s)	Multi-paradigm: Imperative, pipeline, object-oriented, functional and reflective
Appeared in	2006
Designed by	Jeffrey Snover, Bruce Payette, James Truher (et al.)
Developer	Microsoft Corporation
Typing discipline	Strong, safe, implicit and dynamic
Influenced by	Ksh, Perl, C#, CL, DCL, SQL, Tcl and Tk[2]
Platform	.NET Framework
OS	Microsoft Windows

Windows PowerShell is Microsoft's task automation framework, consisting of a command-line shell and associated scripting language built on top of, and integrated with the .NET Framework. PowerShell provides full access to COM and WMI, enabling administrators to perform administrative tasks on both local and remote Windows systems.

In PowerShell, administrative tasks are generally performed by *cmdlets* (pronounced *command-lets*), specialized .NET classes implementing a particular operation. Sets of cmdlets may be combined together in *scripts*, *executables* (which are standalone applications), or by instantiating regular .NET classes (or WMI/COM Objects).[3][4] These work by accessing data in different data stores, like the filesystem or registry, which are made available to the PowerShell runtime via Windows PowerShell *providers*.

Windows PowerShell also provides a hosting mechanism with which the Windows PowerShell runtime can be embedded inside other applications. These applications then leverage Windows PowerShell functionality to implement certain operations, including those exposed via the graphical interface. This capability has been utilized by Microsoft Exchange Server 2007[3][5] to expose its management functionality as PowerShell cmdlets and providers and implement the graphical management tools as PowerShell hosts which invoke the necessary cmdlets. Other Microsoft applications including Microsoft SQL Server 2008[6] also expose their management interface via PowerShell cmdlets. With PowerShell, graphical interface-based management applications on Windows are layered on top of Windows PowerShell. A PowerShell scripting interface for Windows products is mandated by the Common Engineering Criteria.[7]

Windows PowerShell includes its own extensive, console-based help, similar to man pages in Unix shells via the `Get-Help` cmdlet.

Background

Every released version of Microsoft DOS and Microsoft Windows for personal computers has included a command-line interface tool (shell). These are `COMMAND.COM` (in installations relying on MS-DOS, including Windows 9x) and `cmd.exe` (in Windows NT family operating systems). The shell is a command line interpreter that supports a few basic commands. For other purposes, a separate console application must be invoked from the shell. The shell also includes a scripting language (batch files), which can be used to automate various tasks. However, the shell cannot be used to automate all facets of GUI functionality, in part because command-line equivalents of operations exposed via the graphical interface are limited, and the scripting language is elementary and does not allow the creation of complex scripts. In Windows Server 2003, the situation was improved,[8] but scripting support was still considered unsatisfactory.

Microsoft attempted to address some of these shortcomings by introducing the Windows Script Host in 1998 with Windows 98, and its command-line based host: `cscript.exe`. It integrates with the Active Script engine and allows scripts to be written in compatible languages, such as JScript and VBScript, leveraging the APIs exposed by applications via COM. However, it too has its own deficiencies: it is not integrated with the shell, its documentation is not very accessible, and it quickly gained a reputation as a system vulnerability vector after several high-profile computer viruses exploited weaknesses in its security provisions. Different versions of Windows provided various special-purpose command line interpreters (such as netsh and WMIC) with their own command sets. None of them were integrated with the command shell; nor were they interoperable.

By 2002 Microsoft had started to develop a new approach to command line management, including a shell called Monad (aka Microsoft Shell or MSH). The shell and the ideas behind it were published in August 2002 in a white paper entitled Monad Manifesto.[9] Monad was to be a new extensible command shell with a fresh design that would be capable of automating a full range of core administrative tasks. Microsoft first showed off Monad at the Professional Development Conference in Los Angeles in September 2003. A private beta program began a few months later which eventually led to a public beta program.

Microsoft published the first Monad public beta release on June 17, 2005, Beta 2 on September 11, 2005, and Beta 3 on January 10, 2006. Not much later, on April 25, 2006 Microsoft formally announced that Monad had been renamed *Windows PowerShell*, positioning it as a significant part of their management technology offerings.[10] Release Candidate 1 of PowerShell was released at the same time. A significant aspect of both the name change and the RC was that this was now a component of Windows, and not an add-on product.

Release Candidate 2 of PowerShell version 1 was released on September 26, 2006 with final Release to the web (RTW) on November 14, 2006 and announced at TechEd Barcelona. PowerShell for earlier versions of Windows was released on January 30, 2007.[11]

PowerShell v2.0 development began when PowerShell v1.0 shipped. There were three community technology preview (CTP) releases that Microsoft made available to the public. The last CTP release of Windows PowerShell v2.0 was made available in December 2008.

PowerShell v2.0 was completed and released to manufacturing in August 2009, as an integral part of Windows 7 and Windows Server 2008 R2. Versions of PowerShell for Windows XP, Windows Server 2003, Windows Vista and Windows Server 2008 were released in October 2009 and are available for download for both 32-bit and 64-bit platforms.[12]

Overview

There are four kinds of commands Windows PowerShell can execute:

- *cmdlets*, which are .NET programs designed to interact with PowerShell
- PowerShell scripts (files suffixed by `.ps1`)
- PowerShell functions
- standalone executable programs

If a command is a standalone executable program, `PowerShell.exe` launches it in a separate process; if it is a cmdlet, it is executed in the PowerShell process. PowerShell also provides an interactive command line interface, wherein the commands can be entered and their output displayed. The user interface, based on the Win32 console, offers customizable tab completion but lacks syntax highlighting. PowerShell also enables the creation of *aliases* for cmdlets, which are textually translated by PowerShell into invocations of the original commands. PowerShell also supports both named and positional parameters for commands. In executing a cmdlet, the job of binding the argument value to the parameter is done by PowerShell itself, but, for external executables, arguments are passed via the argv (or equivalent) variable array to be parsed by the executable.

Another concept used by PowerShell is that of a *pipeline*. Like Unix pipelines, PowerShell pipelines are used to compose complex commands, allowing the output of one command to be passed as input to another, using the `|` operator. But unlike its Unix counterpart, the PowerShell pipeline is an object pipeline; that is, the data passed between cmdlets are fully typed objects, rather than character streams. When data is piped as objects, the elements they encapsulate retain their structure and types across cmdlets, without the need for any serialization or explicit parsing of the stream, as would be the need if only character streams were shared. An object can also encapsulate certain functions that work on the contained data. These also become available to the recipient command for use.[13][14] For the last cmdlet in a pipeline, PowerShell automatically pipes its output object to the `Out-Default` cmdlet, which transforms the objects into a stream of format objects and then renders those to the screen.[15][16]

Because all PowerShell objects are .NET objects, they share a `.ToString()` method, which retrieves the text representation of the data in an object. Windows PowerShell uses this method to convert an object to text. In addition, it also allows formatting definitions to be specified, so the text representation of objects may be customized by choosing which data elements to display, and how. However, in order to maintain backwards compatibility, if an external executable is used in a pipeline, it receives a text stream representing the object, and does not integrate with the PowerShell type system.

The PowerShell *Extended Type System* (*ETS*) is based on the .NET type system, but with extended semantics (e.g. propertySets and 3rd party extensibility) . For example, it enables the creation of different views of objects by exposing only a subset of the data fields, properties, and methods, as well as specifying custom formatting and sorting behavior. These views are mapped to the original object using XML-based configuration files.[17]

Cmdlets

Cmdlets are specialized commands in the PowerShell environment that implement specific functions. These are the native commands in the PowerShell stack. Cmdlets follow a <*verb*>-<*noun*> naming pattern, such as *Get-ChildItem*, helping to make them self-descriptive.[18] Cmdlets output their results as objects, or collections thereof (including arrays), and can optionally receive input in that form, making them suitable for use as recipients in a pipeline. But, whereas PowerShell allows arrays and other collections of objects to be written to the pipeline, cmdlets always process objects individually. For collections of objects, PowerShell invokes the cmdlet on each object in the collection, in sequence.[18]

Cmdlets are specialized .NET classes, which the PowerShell runtime instantiates and invokes when they are run. Cmdlets derive either from `Cmdlet` or from `PSCmdlet`, the latter being used when the cmdlet needs to interact with the PowerShell runtime.[18] These base classes specify certain methods - `BeginProcessing(), ProcessRecord() and EndProcessing() - which the cmdlet's implementation overrides to provide the functionality. Whenever a cmdlet is run, these methods are invoked by PowerShell in sequence, with ProcessRecord() being called if it receives pipeline input.`[19] `If a collection of objects is piped, the method is invoked for each object in the collection. The class implementing the Cmdlet must have one .NET attribute - CmdletAttribute - which specifies the verb and the noun that make up the name of the cmdlet. Common verbs are provided as an enum.`

If a cmdlet receives either pipeline input or command-line parameter input, there must be a corresponding property in the class, with a mutator implementation. PowerShell invokes the mutator with the parameter value or pipeline input, which is saved by the mutator implementation in class variables. These values are then referred to by the methods which implement the functionality. Properties that map to command-line parameters are marked by `ParameterAttribute`[20] and are set before the call to `BeginProcessing()`. Those which map to pipeline input are also flanked by `ParameterAttribute`, but with the `ValueFromPipeline` attribute parameter set.[21]

The implementation of these cmdlet classes can refer to any .NET API and may be in any .NET language. In addition, PowerShell makes certain APIs available, such as `WriteObject()`, which is used to access PowerShell-specific functionality, such as writing resultant objects to the pipeline. Cmdlets can use .NET data access APIs directly or use the PowerShell infrastructure of PowerShell *Providers*, which make data stores addressable using unique paths. Data stores are exposed using drive letters, and hierarchies within them, addressed as directories. Windows PowerShell ships with providers for the file system, registry, the certificate store, as well as the namespaces for command aliases, variables, and functions.[22] Windows PowerShell also includes various cmdlets for managing various Windows systems, including the file system, or using Windows Management Instrumentation to control Windows components. Other applications can register cmdlets with PowerShell, thus allowing it to manage them, and, if they enclose any datastore (such as databases), they can add specific providers as well.

In PowerShell V2, a more portable version of Cmdlets called Modules have been added. The PowerShell V2 release notes state, "Modules allow script developers and administrators to partition and organize their Windows PowerShell code in self-contained, reusable units. Code from a module executes in its own self-contained context and does not affect the state outside of the module. Modules also enable you to define a restricted runspace environment by using a script."

Pipeline

PowerShell, like Unix/Linux based shells, implements a pipeline. This pipeline enables the output of one cmdlet to be piped as input to another cmdlet. For example, the output of the Get-Process cmdlet can be piped to the Sort-Object cmdlet (e.g. to sort the objects by handle count) and then to the Where-Object to filter any process that has, say, less than 1 MB of paged memory, then finally to the Select-Object cmdlet to select just the first 10 (i.e. the 10 processes based on handle count).

PowerShell differs from Unix/Linux in that structured .NET objects are passed between stages in the pipeline instead of typically unstructured text. Using objects eliminates the need to explicitly parse text output to extract data.[23]

Scripting

Windows PowerShell includes a dynamically typed scripting language which can implement complex operations using cmdlets imperatively. The scripting language supports variables, functions, branching (`if-then-else`), loops (`while`, `do`, `for`, and `foreach`), structured error/exception handling and closures/lambda expressions,[24] as well as integration with .NET. Variables in PowerShell scripts have names that start with `$`; they can be assigned any value, including the output of cmdlets. Strings can be enclosed either in single quotes or in double quotes: when using double quotes, variables will be expanded even if they are inside the quotation marks. According to the variable syntax, if the path to a file is enclosed in braces preceded by a dollar sign (as in `${C:\foo.txt}`), it refers to the contents of the file. If it is used as an L-value, anything assigned to it will be written to the file. When used as an R-value, it will be read from the file. If an object is assigned, it is serialized before storing it.

Object members can be accessed using `.` notation, as in C# syntax. PowerShell provides special variables, such as `$args`, which is an array of all the command line arguments passed to a function from the command line, and `$_`, which refers to the current object in the pipeline.[25] PowerShell also provides arrays and associative arrays. The PowerShell scripting language also evaluates arithmetic expressions entered on the command line immediately, and it parses common abbreviations, such as GB, MB, and KB.

Using the `function` keyword, PowerShell provides for the creation of functions, which can take parameters. A common problem for people new to PowerShell is that function arguments are separated by spaces, not commas:

1. *<function> <param1> <param2>*: Calls the function with two arguments. (These arguments may be bound to parameters declared in the function definition or accessed by position from the $args array.)
2. *<function>(<param1>, <param2>)*: Calls the function with a single argument, a two element array.

PowerShell allows any .NET methods to be called by providing their namespaces enclosed in brackets (`[]`), and then using a pair of colons (`::`) to indicate the static method.[26] For example, `[System.Console]::WriteLine("PowerShell")`. Objects are created using the `New-Object` cmdlet. Calling methods of .NET objects is accomplished by using the regular `.` notation.[26]

For error handling, PowerShell provides a .NET-based exception handling mechanism. In case of errors, objects containing information about the error (`Exception` object) are thrown, which are caught using the `trap` keyword. However, the action-or-error is configurable; in case of an error, PowerShell can be configured to silently resume execution, without trapping the exception.[27]

Scripts written using PowerShell can be made to persist across sessions in a `.ps1` file. Later, either the entire script or individual functions in the script can be used. Scripts and functions are used analogously with cmdlets, in that they can be used as commands in pipelines, and parameters can be bound to them. Pipeline objects can be passed between functions, scripts, and cmdlets seamlessly. However, script execution is disabled by default and must be enabled explicitly.[28] PowerShell scripts can be signed to verify their integrity, and are subject to Code Access Security.

The PowerShell scripting language supports binary prefix notation similar to the scientific notation supported by many programming languages in the C-family.

Hosting

Another use of PowerShell is being embedded in a management application, which then uses the PowerShell runtime to implement the management functionality. For this, PowerShell provides a managed hosting API. Via the APIs, the application can instantiate a *runspace* (one instantiation of the PowerShell runtime), which runs in the application's process and is exposed as a `Runspace` object.[3] The state of the runspace is encased in a `SessionState` object. When the runspace is created, the Windows PowerShell runtime initializes the instantiation, including initializing the providers and enumerating the cmdlets, and updates the `SessionState` object accordingly. The Runspace then must be opened for either synchronous processing or asynchronous processing. After that it can be used to execute commands.

To execute a command, a pipeline (represented by a `Pipeline` object) must be created and associated with the runspace. The pipeline object is then populated with the cmdlets that make up the pipeline. For sequential operations (as in a PowerShell script), a Pipeline object is created for each statement and nested inside another Pipeline object.[3] When a pipeline is created, Windows PowerShell invokes the pipeline processor, which resolves the cmdlets into their respective assemblies (the *command processor*) and adds a reference to them to the pipeline, and associates them with an `InputPipe`, `Outputpipe` and `ErrorOutputPipe` objects, to represent the connection with the pipeline. The types are verified and parameters bound using reflection.[3] Once the pipeline is set up, the host calls the `Invoke()` method to run the commands, or its asynchronous equivalent - `InvokeAsync()`. If the pipeline has the `Write-Host` cmdlet at the end of the pipeline, it writes the result onto the console screen. If not, the results are handed over to the host, which might either apply further processing or display it itself.

The hosting APIs are used by Microsoft Exchange Server 2007 to provide its management GUI. Each operation exposed in the GUI is mapped to a sequence of PowerShell commands (or pipelines). The host creates the pipeline and executes them. In fact, the interactive PowerShell console itself is a PowerShell host, which interprets the scripts entered at command line and creates the necessary `Pipeline` objects and invokes them.

Versions

Initially using the code name "Monad", PowerShell was first shown publicly at the Professional Developers Conference in September 2003. There are currently three versions of PowerShell supported by Microsoft.[29]

Version 1.0

Version 1.0 was released in 2006 for Windows XP SP2/SP3, Windows Server 2003, and Windows Vista. For Windows Server 2008, it is included as an optional feature.

Version 2.0

Version 2.0 is integrated with Windows 7 and Windows Server 2008 R2 and is released for Windows XP with Service Pack 3, Windows Server 2003 with Service Pack 2 and Windows Vista with Service Pack 1.[30]

Microsoft released PowerShell 2.0 with Windows 7 and Windows 2008 R2. Windows PowerShell 2.0 is installed by default on Windows Server 2008 R2 (except on Core installations where it is optional) and Windows 7. [31] For older platforms it is available via the Windows Management Framework.[32] PowerShell V2 includes changes to the scripting language and hosting API, in addition to including more than 240 new cmdlets.[33][34]

A non-exhaustive list of the new features included in PowerShell V2 is: [35] [36] [37]

- **PowerShell Remoting**: Using WS-Management, PowerShell 2.0 allows scripts and cmdlets to be invoked on a remote machine or a large set of remote machines.
- **Background Jobs**: Also called a *PSJob*, it allows a command sequence (script) or pipeline to be invoked asynchronously. Jobs can be run on the local machine or on multiple remote machines. An interactive cmdlet in a PSJob blocks the execution of the job until user input is provided.
- **Transactions**: Enable cmdlet and developers can perform transactional operations. PowerShell 2.0 includes transaction cmdlets for starting, committing, and rolling back a *PSTransaction* as well as features to manage and direct the transaction to the participating cmdlet and provider operations. The PowerShell Registry provider supports transactions.
- **ScriptCmdlets**: These are cmdlets written using the PowerShell scripting language. **NOTE:** The preferred name for script cmdlets is now Advanced Functions.
- **SteppablePipelines**: This allows the user to control when the `BeginProcessing()`, `ProcessRecord()` and `EndProcessing()` functions of a cmdlet are called.
- **Modules**: This allows script developers and administrators to organize and partition PowerShell scripts in self-contained, reusable units. Code from a module executes in its own self-contained context and does not affect the state outside of the module. Modules can define a restricted runspace environment by using a script. They have a persistent state as well as public and private members.
- **Data Language**: A domain-specific subset of the PowerShell scripting language that allows data definitions to be decoupled from the scripts and allows localized string resources to be imported into the script at runtime (*Script Internationalization*).
- **Script Debugging**: It allows breakpoints to be set in a PowerShell script or function. Breakpoints can be set on lines, line & columns, commands and read or write access of variables. It includes a set of cmdlets to control the breakpoints via script.
- **Eventing**: This feature allows listening, forwarding, and acting on management and system events. Eventing allows PowerShell hosts to be notified about state changes to their managed entities. It also enables PowerShell scripts to subscribe to *ObjectEvents*, *PSEvents*, and *WmiEvents* and process them synchronously and asynchronously.
- **Windows PowerShell Integrated Scripting Environment (ISE)**: PowerShell 2.0 includes a GUI-based PowerShell host (formerly known as *Graphical Windows PowerShell*) that provides integrated debugger, syntax highlighting, tab completion and up to 8 PowerShell Unicode-enabled consoles (Runspaces) in a tabbed UI, as

well as the ability to run only the selected parts in a script.

- **Network File Transfer**: Native support for prioritized, throttled, and asynchronous transfer of files between machines using the Background Intelligent Transfer Service (BITS).[38]
- **New Cmdlets**: Including `Out-GridView`, which displays tabular data in the WPF GridView object.
- **New Operators**: `-Split`, `-Join`, and Splatting (`@`) operators.
- **Exception Handling with Try-Catch-Finally**: Unlike other .NET languages, this allows multiple exception types for a single catch block.
- **Nestable Here-Strings**: PowerShell Here-Strings have been improved and can now nest.[39]
- **Block Comments**: PowerShell 2.0 supports block comments using `<#` and `#>` as delimiters.[40]
- **New APIs**: The new APIs range from handing more control over the PowerShell parser and runtime to the host, to creating and managing collection of Runspaces (*RunspacePools*) as well as the ability to create *Restricted Runspaces* which only allow a configured subset of PowerShell to be invoked. The new APIs also support participation in a Windows PowerShell managed transaction.

Version 3.0

Version 3.0 is the next major release of Windows PowerShell, slated to be released with Windows 8 and Windows Server 2012. Microsoft is also planning on releasing down-level versions of V3 for Windows 7 plus the server versions (Server 2008 and Server 2008R2).[41] PowerShell V3 is part of a greater package, Windows Management Framework 3.0, which additionally contains the WINRM service to support remoting.

There have been several Community Technology Preview versions of WMF 3. An early community technology preview 2 (CTP 2) version of Windows Management Framework 3.0 was released by Microsoft in December, 2011.[42] Since then, updated versions have been shipped in the Windows 8/Server 2012 Beta and Release Candidate versions. Down-level versions of those milestone releases are also available.

There are 2 separate types of changes in V3. First there are the changes to PowerShell itself, including a large number of minor improvements (e.g. better Where-Object syntax), as well as major upgrades (Workflows, Cmdlets from WMI Classes, etc.).

Comparison of cmdlets with similar commands

The following table contains a selection of the cmdlets that ship with PowerShell, noting the most similar commands in other well-known command line interpreters.

PowerShell (Cmdlet)	PowerShell (Alias)	CMD.EXE / COMMAND.COM	Unix shell	Description
Get-ChildItem	gci, dir, ls	dir	ls	List all files / directories in the (current) directory
Get-Content	gc, type, cat	type	cat	Get the content of a file
Get-Command	gcm	help	help, which	List available commands
Get-Help	help, man	help	man	Help on commands
Clear-Host	cls, clear	cls	clear	Clear the screen[43]
Copy-Item	cpi, copy, cp	copy	cp	Copy one or several files / a whole directory tree
Move-Item	mi, move, mv	move	mv	Move a file / a directory to a new location
Remove-Item	ri, del, erase, rmdir, rd, rm	del, erase, rmdir, rd	rm, rmdir	Delete a file / a directory

Rename-Item	rni, ren, mv	ren, rename	mv	Rename a file / a directory
Get-Location	gl, pwd	cd	pwd	Display the current directory/present working directory.
Pop-Location	popd	popd	popd	Change the current directory to the directory most recently pushed onto the stack
Push-Location	pushd	pushd	pushd	Push the current directory onto the stack
Set-Location	sl, cd, chdir	cd, chdir	cd	Change the current directory
Tee-Object	tee	n/a	tee	Pipe input to a file or variable, then pass the input along the pipeline
Write-Output	echo, write	echo	echo	Print strings, variables etc. to standard output
Get-Process	gps, ps	tlist,[44] tasklist[45]	ps	List all currently running processes
Stop-Process	spps, kill	kill,[44] taskkill[45]	kill[46]	Stop a running process
Select-String	—	find, findstr	grep	Print lines matching a pattern
Set-Variable	sv, set	set	env, export, set, setenv	Set the value of a variable / create a variable

Notes

[1] http://www.microsoft.com/powershell
[2] Snover, Jeffrey (2008-05-25). "PowerShell and WPF: WTF" (http://blogs.msdn.com/b/powershell/archive/2008/05/25/powershell-and-wpf-wtf.aspx). . Retrieved 2010-10-24.
[3] "How Windows PowerShell works" (http://msdn2.microsoft.com/en-us/library/ms714658.aspx). MSDN. . Retrieved 2007-11-27.
[4] "Extend Windows PowerShell With Custom Commands" (http://msdn.microsoft.com/msdnmag/issues/07/12/PowerShell/default.aspx). MSDN. . Retrieved 2007-11-27.
[5] Scott Lowe (January 5, 2007). "Exchange 2007: Get used to the command line" (http://articles.techrepublic.com.com/5100-10878_11-6147304.html). *TechRepublic*. . Retrieved 2007-11-28.
[6] "SQL Server Support for PowerShell!" (http://blogs.msdn.com/powershell/archive/2007/11/13/sql-server-support-for-powershell.aspx). . Retrieved 2007-11-28.
[7] "Common Engineering Criteria Overview" (http://www.microsoft.com/cec/en/us/cec-overview.aspx#man-windows). . Retrieved 2011-02-12.
[8] Dragan, Richard V. (April 23, 2003). "Windows Server 2003 Delivers Improvements All Around" (http://www.pcmag.com/article2/0,2704,1040410,00.asp). PC Magazine. . Retrieved 2007-11-02. "A standout feature here is that virtually all admin utilities now work from the command line (and most are available through telnet)."
[9] Snover, Jeffrey (August 2, 2002). "Monad Manifesto" (http://blogs.msdn.com/cfs-file.ashx/__key/CommunityServer-Components-PostAttachments/00-01-91-05-67/Monad-Manifesto-_2D00_-Public.doc). *Windows PowerShell Team*. Microsoft. . Retrieved 2002-08-09..}}
[10] Snover, Jeffrey (April 25, 2006). "Windows PowerShell (Monad) Has Arrived" (http://blogs.msdn.com/powershell/archive/2006/04/25/583344.aspx). *Windows PowerShell team blog*. MSDN. . Retrieved 2006-04-26.
[11] Snover, Jeffrey (November 15, 2006). "Windows PowerShell : Windows PowerShell & Windows Vista" (http://blogs.msdn.com/powershell/archive/2006/11/15/windows-powershell-windows-vista.aspx). *Windows PowerShell team blog*. MSDN. . Retrieved 2007-01-26.
[12] "Windows PowerShell v2.0 RTM download for Downlevel OSs" (http://support.microsoft.com/kb/968929). .
[13] Don Jones (2008). "Windows PowerShell - Rethinking the Pipeline" (http://www.microsoft.com/technet/technetmag/issues/2007/07/PowerShell/default.aspx). *Microsoft TechNet*. . Retrieved 2007-11-28.
[14] "Windows PowerShell Object Concepts" (http://msdn2.microsoft.com/en-us/library/aa347685.aspx). . Retrieved 2007-11-28.
[15] "How PowerShell Formatting and Outputting REALLY works" (http://blogs.msdn.com/powershell/archive/2006/04/30/586973.aspx). . Retrieved 2007-11-28.
[16] "More - How does PowerShell formatting really work?" (http://blogs.msdn.com/powershell/archive/2006/06/21/641738.aspx). . Retrieved 2007-11-28.

[17] "Windows PowerShell Extended Type System" (http://msdn2.microsoft.com/en-us/library/ms714419.aspx). . Retrieved 2007-11-28.
[18] "Windows PowerShell Cmdlets" (http://msdn2.microsoft.com/en-us/library/ms714395.aspx). . Retrieved 2007-11-28.
[19] "Creating Your First Cmdlet" (http://msdn2.microsoft.com/en-us/library/ms714622.aspx). . Retrieved 2007-11-28.
[20] "Adding parameters That Process Command Line Input" (http://msdn2.microsoft.com/en-us/library/ms714663.aspx). . Retrieved 2007-11-28.
[21] "Adding parameters That Process Pipeline Input" (http://msdn2.microsoft.com/en-us/library/ms714597.aspx). . Retrieved 2007-11-28.
[22] "Windows PowerShell Providers" (http://technet.microsoft.com/en-us/library/dd347723.aspx). . Retrieved 2010-10-14.
[23] "Windows PowerShell Owner's Manual: Piping and the Pipeline in Windows PowerShell" (http://technet.microsoft.com/en-us/library/ee176927.aspx). Technet.microsoft.com. . Retrieved 2011-09-27.
[24] "Anonymous Functions and Code Blocks in PowerShell" (http://defndo.com/powershell-code-blocks-and-anonymous-functions/). . Retrieved 2012-01-21.
[25] "Introduction to Windows PowerShell's Variables" (http://www.computerperformance.co.uk/powershell/powershell_variables.htm). . Retrieved 2007-11-28.
[26] "Lightweight Testing with Windows PowerShell" (http://msdn.microsoft.com/msdnmag/issues/07/05/TestRun/default.aspx). . Retrieved 2007-11-28.
[27] "Trap [Exception] { “In PowerShell” }" (http://huddledmasses.org/trap-exception-in-powershell). . Retrieved 2007-11-28.
[28] "Running Windows PowerShell Scripts" (http://www.microsoft.com/technet/scriptcenter/topics/winpsh/manual/run.mspx). . Retrieved 2007-11-28.
[29] "Download Windows PowerShell" (http://blogs.msdn.com/b/powershell/archive/2008/12/30/download-windows-powershell.aspx). 2008-12-30. . Retrieved 2010-06-27.
[30] "Windows Management Framework is here!" (http://blogs.msdn.com/powershell/archive/2009/10/27/windows-management-framework-is-here.aspx). 2009-10-27. . Retrieved 2009-10-30.
[31] PowerShell Team (2008-10-28). "PowerShell will be installed by default on Windows Server 08 R2 (WS08R2) and Windows 7 (W7)! - Windows PowerShell Blog - Site Home - MSDN Blogs" (http://blogs.msdn.com/powershell/archive/2008/10/28/powershell-will-be-installed-by-default-on-windows-server-08-r2-ws08r2-and-windows-7-w7.aspx). Blogs.msdn.com. . Retrieved 2011-09-27.
[32] "Microsoft Support Knowledge Base: Windows Management Framework (Windows PowerShell 2.0, WinRM 2.0, and BITS 4.0)" (http://support.microsoft.com/kb/968929). Support.microsoft.com. 2011-09-23. . Retrieved 2011-09-27.
[33] PowerShell Team (2008-10-29). "574 Reasons Why We Are So Proud and Optimistic About W7 and WS08R2 - Windows PowerShell Blog - Site Home - MSDN Blogs" (http://blogs.msdn.com/powershell/archive/2008/10/29/574-reasons-why-we-are-so-proud-and-optimistic-about-w7-and-ws08r2.aspx). Blogs.msdn.com. . Retrieved 2011-09-27.
[34] (http://channel9.msdn.com/pdc2008/ES24/)
[35] "What's New in CTP of PowerShell 2.0" (http://blogs.msdn.com/powershell/archive/2007/11/06/what-s-new-in-ctp-of-powershell-2-0.aspx). . Retrieved 2007-11-28.
[36] "Windows PowerShell V2 Community Technology Preview 2 (CTP2) - releaseNotes" (http://www.microsoft.com/downloads/details.aspx?FamilyId=7C8051C2-9BFC-4C81-859D-0864979FA403&displaylang=en). . Retrieved 2008-05-05.
[37] "Differences between PowerShell 1.0 and PowerShell 2.0" (http://activexperts.com/admin/powershell/ps1vs2/). . Retrieved 2010-06-26.
[38] PowerShell Team (2008-10-14). "GoGrid Snap-in - Managing Cloud Services with PowerShell - Windows PowerShell Blog - Site Home - MSDN Blogs" (http://blogs.msdn.com/powershell/archive/2008/10/14/gogrid-snap-in-managing-cloud-services-with-powershell.aspx). Blogs.msdn.com. . Retrieved 2011-09-27.
[39] PowerShell Team (2008-10-17). "Emit-XML - Windows PowerShell Blog - Site Home - MSDN Blogs" (http://blogs.msdn.com/powershell/archive/2008/10/18/emit-xml.aspx). Blogs.msdn.com. . Retrieved 2011-09-27.
[40] PowerShell Team (2008-06-14). "Block Comments in V2 - Windows PowerShell Blog - Site Home - MSDN Blogs" (http://blogs.msdn.com/powershell/archive/2008/06/14/block-comments-in-v2.aspx). Blogs.msdn.com. . Retrieved 2011-09-27.
[41] "Upgrading to Powershell 3.0 on Windows 2008 R2 SP1" (http://www.happysysadm.com/2012/06/upgrading-to-powershell-30-on-windows.html). .
[42] "Windows Management Framework 3.0 Community Technology Preview (CTP) #2 Available for Download" (http://blogs.msdn.com/b/powershell/archive/2011/12/02/windows-management-framework-3-0-community-technology-preview-ctp-2-available-for-download.aspx). .
[43] Clear-Host is implemented as a predefined PowerShell function.
[44] Available in Windows NT4, Windows 98 Resource Kit, Windows 2000 Support Tools
[45] Available in Windows XP Professional Edition and later
[46] Also used in UNIX to send a process any signal, the "Terminate" signal is merely the default

Examples

Examples are provided first using the long-form canonical syntax and then using more terse UNIX-like and DOS-like aliases that are set up in the default configuration. For a list of all aliases, use the Get-Alias Cmdlet.

- Stop all processes that begin with the letter p:

```
PS> Get-Process p* | Stop-Process
```

```
PS> ps p* | kill
```

- Find the processes that use more than 1000 MB of memory and kill them:

```
PS> Get-Process | Where-Object { $_.WS -gt 1000MB } | Stop-Process
```

```
PS> ps | ? { $_.WS -gt 1000MB } | kill
```

- Calculate the number of bytes in the files in a directory:

```
PS> Get-Childitem | Measure-Object -Property Length -Sum
```

```
PS> ls  | measure length -s
```

```
PS> dir | measure length -s
```

- Determine whether a specific process is no longer running:

```
PS> $processToWatch = Get-Process Notepad
PS> $processToWatch.WaitForExit()
```

```
PS> (ps notepad).WaitForExit()
```

- Change the case of a string from lower to upper:

```
PS> 'hello, world!'.ToUpper()
```

- Insert the string 'ABC' after the first character in the word 'string' to have the result 'sABCtring':

```
PS> 'string'.Insert(1, 'ABC')
```

- Download a specific RSS feed and show the titles of the 8 most recent entries:

```
PS> $rssUrl = 'http://blogs.msdn.com/powershell/rss.aspx'
PS> $blog = [xml](new-object System.Net.WebClient).DownloadString($rssUrl)
PS> $blog.rss.channel.item | select title -first 8
```

```
PS> $x = new-object xml
PS> $x.load('http://blogs.msdn.com/powershell/rss.aspx')
PS> $x.rss.channel.item | select title -f 8
```

- Sets $UserProfile to the value of the UserProfile environment variable

```
PS> $UserProfile = $env:UserProfile
```

- Call a static method of a .Net object

```
PS> [System.Math]::Sqrt(16)
4
```

- Run a command line executable with arguments:

```
PS> [Array]$arguments = '-h', '15', 'www.Wikipedia.com'
PS> tracert $arguments
```

- Get the serial number of a remote computer from WMI:

```
PS> Get-WmiObject -ComputerName MyServer -Class Win32_BIOS | Select-Object SerialNumber
```

```
PS> gwmi -co MyServer Win32_BIOS | select SerialNumber
```

File extensions

- PS1 – Windows PowerShell shell script
- PS1XML – Windows PowerShell format and type definitions
- PSC1 – Windows PowerShell console file
- PSD1 – Windows PowerShell data file (for Version 2)
- PSM1 – Windows PowerShell module file (for Version 2)

Application support

Snap-ins and hosts

Application	Version	Cmdlets	Provider	Management GUI
Exchange Server	2007	402	Yes	Yes
Windows Server	2008	Yes	Yes	No
Microsoft SQL Server	2008	Yes	Yes	No
System Center Operations Manager	2007	74	Yes	No
System Center Virtual Machine Manager	2007	Yes	Yes	Yes
System Center Data Protection Manager	2007	Yes	No	No
Windows Compute Cluster Server	2007	Yes	Yes	No
Microsoft Transporter Suite for Lotus Domino[1]	08.02.0012	47	No	No
Microsoft PowerTools for Open XML[2]	1.0	33	No	No
IBM WebSphere MQ[3]	6.0.2.2	44	No	No
Quest Management Shell for Active Directory[4]	1.1	40	No	No
Special Operations Software Specops Command[5]	1.0	Yes	No	Yes
VMware PowerCLI[6]	5.0.1	293	No	Yes
Internet Information Services[7]	7.0	54	Yes	No
Windows 7 Troubleshooting Center[8]	6.1	Yes	No	Yes
Microsoft Deployment Toolkit	2010	Yes	No	No
NetApp Data ONTAP PowerShell Toolkit[9]	2.0	1059	Yes	Yes
JAMS Scheduler - Job Access & Management System[10]	5.0	52	Yes	Yes
UIAutomation[11]	0.7	325	No	No
Dell Equallogic[12]	3.5	55	No	No

LOGINventory[13]	5.8	Yes	Yes	Yes

Alternative implementation

A project named **Pash** was attempted to create an open source and cross-platform implementation of PowerShell via the Mono framework. Pash was maintained by Igor Moochnick, written in C# and was released under the GNU General Public License. Pash development ceased in 2008 and is currently only partially complete.[14]

References

[1] "Microsoft Transporter Suite for Lotus Domino" (http://www.microsoft.com/downloads/details.aspx?familyid=35fc4205-792b-4306-8e4b-0de9cce72172&displaylang=en). . Retrieved 2008-03-07.
[2] "PowerTools for Open XML" (http://www.codeplex.com/PowerTools). . Retrieved 2008-06-20.
[3] "MO74: WebSphere MQ - Windows PowerShell Library" (http://www-1.ibm.com/support/docview.wss?rs=171&uid=swg24017698). . Retrieved 2007-12-05.
[4] "PowerShell Commands for Active Directory by Quest Software" (http://www.quest.com/powershell/activeroles-server.aspx). . Retrieved 2008-07-02.
[5] "PowerShell Remoting through Group Policy" (http://www.specopssoft.com/powershell/). . Retrieved 2007-12-07.
[6] "VMware PowerCLI" (http://www.vmware.com/sdk/vitk_win/). . Retrieved 2012-01-01.
[7] "Windows PowerShell : IIS7 PowerShell Provider Tech Preview 2" (http://blogs.msdn.com/powershell/archive/2008/07/03/iis7-powershell-provider-tech-preview-2.aspx). . Retrieved 2008-07-03.
[8] "Kudos to the Win7 Diagnostics Team" (http://blogs.msdn.com/powershell/archive/2009/06/14/kudos-to-the-win7-diagnostics-team.aspx). . Retrieved 2009-06-15.
[9] "Kudos to NetApp for Data ONTAP PowerShell ToolKit" (http://blogs.msdn.com/b/powershell/archive/2010/06/16/kudos-to-netapp-for-data-ontap-powershell-toolkit.aspx). . Retrieved 2010-06-15.
[10] "Heterogeneous Job Scheduling With PowerShell" (http://blogs.msdn.com/b/powershell/archive/2007/06/28/heterogeneous-job-scheduling-with-powershell.aspx). . Retrieved 2010-09-15.
[11] "UIAutomation PowerShell Extensions" (http://uiautomation.codeplex.com). . Retrieved 2012-02-16.
[12] "EqualLogic HIT-ME with PowerShell" (http://en.community.dell.com/techcenter/storage/w/wiki/2688.aspx). . Retrieved 2012-03-09.
[13] "LOGINventory" (http://de.wikipedia.org/wiki/LOGINventory). .
[14] "igor.moochnick - Pash" (http://pash.sourceforge.net/). Pash.sourceforge.net. . Retrieved 2011-09-27.

Further reading

- Holmes, Lee. *Windows PowerShell Quick Reference* (http://www.oreilly.com/catalog/windowspowershell/). ISBN 0-596-52813-2.
- Holmes, Lee. *Windows PowerShell Cookbook* (http://www.oreilly.com/catalog/9780596528492/index.html). ISBN 0-596-52849-3. Retrieved 2010-09-19.
- Johnson, Willis. *Scripting SQL Management Objects in Windows PowerShell* (http://www.wrox.com/WileyCDA/WroxTitle/Scripting-SQL-Management-Objects-in-Windows-PowerShell.productCd-0470279354.html). ISBN 978-0-470-27935-9. Retrieved 2010-09-19.
- Don Jones, Jeffery Hicks. *Windows PowerShell 2.0: TFM* (http://www.sapienpress.com/PowerShell3.asp). ISBN 978-0-9821314-2-8.
- Kopczynski, Tyson. *Windows PowerShell Unleashed* (http://www.informit.com/title/0672329530). ISBN 0-672-32953-0.
- Kumaravel, Arul; Jon White, Michael Naixin Li, Scott Happell, Guohui Xie, Krishna C. Vutukuri. *Professional Windows PowerShell Programming: Snapins, Cmdlets, Hosts and Providers* (http://www.wrox.com/WileyCDA/WroxTitle/productCd-0470173939.html). ISBN 0-470-17393-9. Retrieved 2010-09-19.
- Oakley, Andy. *Monad (AKA PowerShell)* (http://www.oreilly.com/catalog/powershell/). ISBN 0-596-10009-4.

- Rottenberg, Hal. *Managing VMware Infrastructure with Windows PowerShell: TFM* (http://sapienpress.com/vmware.asp). ISBN 0-9821314-0-2. Retrieved 2010-09-19.
- Tyson Kopczynski, Pete Handley, Marco Shaw. *Windows PowerShell Unleashed* (http://www.informit.com/title/0672329883). ISBN 978-0-672-32988-3.
- Watt, Andrew. *Professional Windows PowerShell* (http://www.wrox.com/WileyCDA/WroxTitle/productCd-0471946931.html). ISBN 0-471-94693-1. Retrieved 2010-09-19.
- Wilson, Ed. *Microsoft Windows PowerShell Step By Step* (http://www.microsoft.com/MSPress/books/10329.aspx). ISBN 0-7356-2395-3.

External links

- Microsoft's Windows PowerShell webpage (http://technet.microsoft.com/en-us/library/bb978526.aspx)
- Windows PowerShell Team Blog (http://blogs.msdn.com/powershell/)
- Windows PowerShell Community Forum (http://social.technet.microsoft.com/Forums/en-US/winserverpowershell/threads)
- Windows PowerShell Quick Reference Materials (http://social.technet.microsoft.com/wiki/contents/articles/183.windows-powershell-survival-guide-en-us.aspx#Quick_Reference_Materials)

Environment variable

Environment variables are a set of dynamic named values that can affect the way running processes will behave on a computer.

They can be said in some sense to create the operating environment in which a process runs. For example, an environment variable with a standard name can designate the location that a particular computer system uses to store temporary files – this may vary from one computer system to another. A process which invokes the environment variable by (standard) name can be sure that it is storing temporary information in a directory (folder) that exists and is expected to have sufficient space.

Synopsis

In all Unix and Unix-like systems, each process has its own separate set of environment variables. By default, when a process is created, it inherits a duplicate environment of its parent process, except for explicit changes made by the parent when it creates the child. At API level, these changes must be done between running `fork` and `exec`. Alternatively, from command shells such as bash, a user can change environment variables for a particular command invocation by indirectly invoking it via `env` or using the `ENVIRONMENT_VARIABLE=VALUE <command>` notation. All Unix operating system flavors, MS-DOS, and Microsoft Windows have environment variables; however, they do not all use the same variable names. A running program can access the values of environment variables for configuration purposes.

Examples of environment variables include:

- PATH - lists directories the shell searches, for the commands the user may type without having to provide the full path.
- HOME (Unix-like) and USERPROFILE (Microsoft Windows) - indicate where a user's home directory is located in the file system.
- HOME/{.AppName} (Unix-like) and APPDATA\{DeveloperName\AppName} (Microsoft Windows) - for storing application settings. Many open source programs incorrectly use USERPROFILE for application settings in Windows - USERPROFILE should only be used in dialogs that allow user to choose between paths like Documents/Pictures/Downloads/Music, for programmatic purposes APPDATA (roaming), LOCALAPPDATA or

PROGRAMDATA (shared between users) is used.

- TERM (Unix-like) - specifies the type of computer terminal or terminal emulator being used (e.g., **vt100** or **dumb**).
- PS1 (Unix-like) - specifies how the prompt is displayed in the Bourne shell and variants.
- MAIL (Unix-like) - used to indicate where a user's mail is to be found.
- TEMP - location where processes can store temporary files

Shell scripts and batch files use environment variables to communicate data and preferences to child processes. They can also be used to store temporary values for reference later in a shell script. However, in Unix, other variables are usually used for this.

In Unix, an environment variable that is changed in a script or compiled program will only affect that process and possibly child processes. The parent process and any unrelated processes will not be affected. In MS-DOS, changing or removing a variable's value inside a BATCH file will change the variable for the duration of command.com's existence.

In Unix, the environment variables are normally initialized during system startup by the system init scripts, and hence inherited by all other processes in the system. Users can, and often do, augment them in the profile script for the command shell they are using. In Microsoft Windows, each environment variable's default value is stored in the Windows registry or set in the autoexec.bat file.

Getting and setting environment variables

The variables can be used both in scripts and on the command line. They are usually referenced by putting special symbols in front of or around the variable name. For instance, to display the user home directory, in most scripting environments, the user has to type:

```
echo $HOME
```

On DOS, OS/2 or Windows systems, the user has to type this:

```
echo %HOME%
```

In Windows PowerShell, the user has to type this:

```
Write-Output $HOME
```

Unix

The commands **`env`**, **`set`**, and **`printenv`** display all environment variables and their values. `env` and `set` are also used to set environment variables and are often incorporated directly into the shell. `printenv` can also be used to print a single variable by giving that variable name as the sole argument to the command.

In Unix, the following commands can also be used, but are often dependent on a certain shell.

```
export VARIABLE=value  # for Bourne, bash, and related shells
setenv VARIABLE value  # for csh and related shells
```

Working principles of environment variables

A few simple principles govern how environment variables achieve their effect.

Local to process

Environment variables are local to the process in which they were set. That means that if we spawn two shell processes and change the value of an environment variable in one, that change will not be seen by the other.

Inheritance

When a child process is created, it inherits all the environment variables and their values from the parent process. Usually, when a program calls another program, it first creates a child process by forking, then the child adjusts the environment as needed and lastly the child replaces itself with the program to be called. This procedure gives the calling program control over the environment of the called program.

Case-sensitive

In Unix and Unix-like systems the names of environment variables are case-sensitive.

Persistence

Environment variables persistence can be session-wide or system-wide.

DOS, OS/2 and Windows (Command Prompt)

In DOS, OS/2 and Windows, the **`set`** command without any arguments displays all environment variables along with their values.

To set a variable to a particular value, use:

```
set VARIABLE=value
```

However, this is temporary. Permanent change to the environment variable can be achieved through editing the registry (not recommended for novices) and using the Windows Resource Kit application `setx.exe`. With the introduction of Windows Vista, the `setx` command became part of Windows.

Users of the Windows GUI can manipulate variables via <Control Panel:System:Advanced:Environment Variables>; through the Windows Registry this is done changing the values under HKCU\Environment (for user specific variables) and HKLM\SYSTEM\CurrentControlSet\Control\Session Manager\Environment (for System variables).

To set a variable whose name is in another variable, you can do:

```
set %VARNAME%=value
```

This feature allows certain interesting applications. For example, you may create a uni-dimensional array of elements (vector) this way:

```
set VECTOR[%I%]=value of element subscript %I%

:MkVec
set VECNAME=%1
set i=0
:loop
    shift
    if "%1" == "" goto exitloop
    set /a i+=1
    set %VECNAME%[%i%]=%1
```

```
    goto loop
:exitloop
exit /B %i%

call :MkVec DOWNAME=Monday Tuesday Wednesday Thursday Friday Saturday Sunday
```

To see the current value of a particular variable, use:

```
echo %VARIABLE%
```

or

```
set VARIABLE
```

Note: Please take note that doing so will print out all variables beginning with 'VARIABLE'. Another example is:

```
C:\> set p
Path=c:\.. ..
PATHEXT=.COM;.EXE;.BAT;
PROCESSOR_ARCHITECTURE=.. ..
PROCESSOR_IDENTIFIER=x8..
PROCESSOR_LEVEL=6..
PROCESSOR_REVISION=1706..
ProgramFiles=C:\Program.. .
PROMPT=$P$G
```

To see the value of an array element a double expansion is required: one for the subscript value and an additional expansion for the array element. This may be achieved via Delayed !VARIABLE! Expansion this way:

```
set DOW=value of Day of Week (1..7)
echo !DOWNAME[%DOW%]!
```

To delete a variable, the following command is used:

```
set VARIABLE=
```

Windows PowerShell

To set a system variable:

```
Set-Content -Path Env:VARIABLE -value value
$Env:VARIABLE = value # alternate form
```

Case-insensitivity

In Windows the names of environment variables are case-insensitive.

Unexported variables

In Unix shells, variables may be assigned without the **`export`** keyword. Variables defined in this way are displayed by the **`set`** command, but are **not** true environment variables, as they are stored only by the shell and not recognized by the kernel. The `printenv` command will not display them, and child processes do not inherit them.

```
VARIABLE=value
```

However, if used in front of a program to run, the variables will be exported to the environment and thus appear as real environment variables to the program:

```
VARIABLE=value program_name [arguments]
```

The tool that gives closest parallel in Windows is the SETLOCAL/ENDLOCAL commands that prevent variables from being set globally.

Security

On Unix, a setuid program is given an environment chosen by its caller, but it runs with different authority from its caller. The dynamic linker will usually load code from locations specified by the environment variables LD_LIBRARY_PATH and LD_PRELOAD and run it with the process's authority. If a setuid program did this, it would be insecure, because its caller could get it to run arbitrary code and hence misuse its authority. For this reason, libc unsets these environment variables at startup in a setuid process. setuid programs usually unset unknown environment variables and check others or set them to reasonable values.

Common environment variables

Examples of Unix environment variables

$PATH

Contains a colon-separated list of directories that the shell searches for commands that do not contain a slash in their name (commands with slashes are interpreted as file names to execute, and the shell attempts to execute the files directly). It is equivalent to the Windows `%PATH%` variable. See: Path (computing).

$HOME

Contains the location of the user's home directory. Although the current user's home directory can also be found out through the C functions `getpwuid` and `getuid`, `$HOME` is often used for convenience in various shell scripts (and other contexts). Using the environment variable also gives the user the possibility to point to another directory.

$PWD

This variable points to the current directory. Equivalent to the output of the command pwd when called without arguments.

$DISPLAY

Contains the identifier for the display that X11 programs should use by default.

$LD_LIBRARY_PATH

On many Unix systems with a dynamic linker, contains a colon-separated list of directories that the dynamic linker should search for shared objects when building a process image after `exec`, before searching in any other directories.

$LANG, $LC_ALL, $LC_...

LANG is used to set to the default locale. For example, if the locale values are `pt_BR`, then the language is set to (Brazilian) Portuguese and Brazilian practice is used where relevant. Different aspects of localization are controlled by individual LC_-variables (LC_CTYPE, LC_COLLATE, LC_DATE etc.). LC_ALL can be used to force the same locale for all aspects.

Examples of DOS environment variables

%COMSPEC%

This variable contains the full path to the command processor, command.com.

%PATH%

This variable contains a semicolon-delimited list of directories in which the command interpreter will search for executable files. Equivalent to the Unix $PATH variable (although note that PATH on Windows additionally performs the same task as LD_LIBRARY_PATH on Unix-like systems). Note that %PATH% can also be set like this PATH=c:\dos; where SET isn't required.

%TEMP% and %TMP%

These variables contain the path to the directory where temporary files should be stored.

Examples from Microsoft Windows

Discrete value variables

These variables generally expand to discrete values, such as the current working directory, the current date, or a random number. Some of these are true environment variables and will be expanded by all functions that handle environment variables. Others, like %CD% simply look like environment variables and will only be expanded by some functions and shells. They are not case sensitive.

%CD%

This variable points to the current directory. Equivalent to the output of the command `cd` when called without arguments.

%DATE%

This variable expands to the current date. The date is displayed according to the current user's date format preferences.

The following is a way of reformatting the date and time for use in file copies. The example assumes UK format of day month year and the time is set for a 24 hour clock.

```
@echo off
echo %DATE% %TIME%
for /F "tokens=1-3 delims=/" %%a in ("%DATE%") do set MTH=%%a& set DAY=%%b& set YR=%%c
for /F "tokens=1-3 delims=:." %%a in ("%TIME%") do set HR=%%a& set MIN=%%b& set SEC=%%c
if "%HR:~0,1%"==" " set HR=0%HR:~1,1%
set MYDATE=%YR%%MTH%%DAY%-%HR%%MIN%%SEC%
echo %MYDATE%
```

%ERRORLEVEL%

This variable points to the current error level. If there was an error in the previous command, this is what you need to check against to find out about that.

%RANDOM%

This variable returns a random number between 0 and 32767.

%TIME%

This variable points to the current time. The time is displayed according to the current user's time format preferences.

System path variables

These variables refer to locations of critical operating system resources, and as such generally are not user-dependent.

%AppData%

Contains the full path to the Application Data folder of the logged-in user. Does not work on Windows NT 4.0 SP6 UK.

%LOCALAPPDATA%

This variable is the temporary files of Applications. Its uses include storing of Desktop Themes, Windows Error Reporting, Caching and profiles of web browsers.

%ComSpec%

This variable contains the full path to the command processor; on Windows NT based operating systems this is `cmd.exe`, while on Windows 9x and ME it is the DOS command processor, `COMMAND.COM`.

%PATH%

This variable contains a semicolon-delimited (do not put spaces in between) list of directories in which the command interpreter will search for an executable file that matches the given command. Equivalent to the Unix `$PATH` variable.

%ProgramFiles%

This variable points to Program Files directory, which stores all the installed program of Windows and others. The default on English-language systems is `C:\Program Files`. In 64-bit editions of Windows (XP, 2003, Vista), there are also **%ProgramFiles(x86)%** which defaults to `C:\Program Files (x86)` and **%ProgramW6432%** which defaults to `C:\Program Files`. The %ProgramFiles% itself depends on whether the process requesting the environment variable is itself 32-bit or 64-bit (this is caused by Windows-on-Windows 64-bit redirection).

%CommonProgramFiles%

This variable points to Common Files directory. The default is `C:\Program Files\Common Files`.

%SystemDrive%

The %SystemDrive% variable is a special system-wide environment variable found on Microsoft Windows NT and its derivatives. Its value is the drive upon which the system folder was placed. Also see next item.

The value of %SystemDrive% is in most cases `C:`.

%SystemRoot%

The %SystemRoot% variable is a special system-wide environment variable found on Microsoft Windows NT and its derivatives. Its value is the location of the system folder, including the drive and path.

The drive is the same as %SystemDrive% and the default path on a clean installation depends upon the version of the operating system. By default, on a clean installation:

- Windows NT 5.1 (Windows XP) and newer versions use `\WINDOWS`
- Windows NT 5.0 (Windows 2000), Windows NT 4.0 and Windows NT 3.1 use `\WINNT`
- Windows NT 3.5x uses `\WINNT35`
- Windows NT 4.0 Terminal Server use `\WTSRV`

%WinDir%

This variable points to the Windows directory (on Windows NT-based operating systems it is identical to the %SystemRoot% variable, above). If the System is on drive C: then the default values are:

- `C:\WINDOWS` on Windows 95, Windows 98, Windows Me, Windows XP, Windows Server 2003, Windows Vista, Windows Server 2008 and Windows 7

- `C:\WINNT` for Windows NT 4, and Windows 2000

Note that Windows NT 4 Terminal Server Edition by default installs to `C:\WTSRV`.

User management variables

These variables store information related to resources and settings owned by various user profiles within the system. As a general rule, these variables do not refer to critical system resources or locations that are necessary for the OS to run.

%AllUsersProfile% (%PROGRAMDATA% for Windows Vista, Windows 7)

The %AllUsersProfile%(%PROGRAMDATA%) variable expands to the full path to the All Users profile directory. This profile contains resources and settings that are used by all system accounts. Shortcut links copied to the All Users' Start menu or Desktop folders will appear in every user's Start menu or Desktop, respectively.

%UserDomain%

The variable holds the name of the Workgroup or Windows Domain to which the current user belongs. The related variable, %LOGONSERVER%, holds the hostname of the server that authenticated the current user's logon credentials (name and password). For Home PCs, and PCs in a Workgroup, the authenticating server is usually the PC itself. For PCs in a Windows Domain, the authenticating server is a domain controller (a primary domain controller, or PDC, in Windows NT 4-based domains).

%UserProfile%

The %UserProfile% variable is a special system-wide environment variable found on Microsoft Windows NT and its derivatives. Its value is the location of the current user's profile directory, in which is found that user's HKCU registry hive (`NTUSER`).

Users can also use the **%USERNAME%** variable to determine the active users login identification.

Windows GUI forced variable expansion

In certain cases it is not possible to create file paths containing environment variables using the Windows GUI, and it is necessary to fight with the user interface to make things work as intended.

- In Windows 7, a shortcut may not contain the variable %USERNAME% in unexpanded form. Trying to create shortcut to *\\server\share\accounts\%USERNAME%* or *C:\users\%USERNAME%* will be silently changed to replace %USERNAME% with the account name of the currently logged-in user, when the OK button is pressed on the shortcut properties.
  - This can only be overridden if the %USERNAME% variable is part of a parameter to some other program in the shortcut. For example, *%SYSTEMROOT%\Explorer.exe C:\Users\%USERNAME%* is not expanded when OK is clicked, but this shortcut is treated as unsafe and displays a warning when opened.
- In Group Policy Management on Server 2008 R2, a profile folder can not be redirected to a custom folder hierarchy. For example, the desktop can not be redirected to *\\server\share\accounts\%USERNAME%\custom\path\desktop*. Upon pressing OK, this is silently changed to "Create a folder for each user in the root path" with the path *\\server\share\accounts* pointing to "\username\desktop".
  - This behavior can only be overridden if the path contains a variable or drive letter that is not currently resolvable at the time of editing the GPO. For example if a mapping for drive O: does not exist on the server, then the path *O:\folder\%username%\CustomTarget* is not expanded when OK is clicked.
- A domain user account may not contain a profile path or home folder path containing an unexpanded %USERNAME% variable. Upon clicking OK, this is silently replaced with the user's account name.

- This causes problems for new user creation that is performed by copying an existing user account, if there are additional folders listed after the username in the path. For a pre-existing account with a profile path of *\server\share\accounts\DomainUser\profile* the Microsoft Management Console doesn't know which part of the path contains the previous user's name and doesn't change the path during the copy, resulting in the new account pointing to the other account's profile/home paths. The profile/home paths must be manually re-edited to point to the correct location.

Default Values on Microsoft Windows

Variable	Windows XP	Windows Vista/7
%ALLUSERSPROFILE% and %PROGRAMDATA%	C:\Documents and Settings\All Users	C:\ProgramData
%APPDATA%	C:\Documents and Settings\{username}\Application Data	C:\Users\{username}\AppData\Roaming
%COMPUTERNAME%	{computername}	{computername}
%COMMONPROGRAMFILES%	C:\Program Files\Common Files	C:\Program Files\Common Files
%COMMONPROGRAMFILES(x86)%	C:\Program Files (x86)\Common Files	C:\Program Files (x86)\Common Files
%COMSPEC%	C:\Windows\System32\cmd.exe	C:\Windows\System32\cmd.exe
%HOMEDRIVE%	C:	C:
%HOMEPATH%	\Documents and Settings\{username}	\Users\{username}
%LOCALAPPDATA%		C:\Users\{username}\AppData\Local
%LOGONSERVER%	\\{domain_logon_server}	\\{domain_logon_server}
%PATH%	C:\Windows\system32;C:\Windows;C:\Windows\System32\Wbem;{plus program paths}	C:\Windows\system32;C:\Windows;C:\Windows\System32\Wbem;{plus program paths}
%PATHEXT%	.COM;.EXE;.BAT;.CMD;.VBS;.VBE;.JS;.WSF;.WSH	.com;.exe;.bat;.cmd;.vbs;.vbe;.js;.jse;.wsf;.wsh;.msc
%PROGRAMDATA%		%SystemDrive%\ProgramData
%PROGRAMFILES%	%SystemDrive%\Program Files	%SystemDrive%\Program Files
%PROGRAMFILES(X86)%	%SystemDrive%\Program Files (x86) (only in 64-bit version)	%SystemDrive%\Program Files (x86) (only in 64-bit version)
%PROMPT%	Code for current command prompt format. Code is usually PG	Code for current command prompt format. Code is usually PG
%PSModulePath%		%SystemRoot%\system32\WindowsPowerShell\v1.0\Modules\
%PUBLIC%		%SystemDrive%\Users\Public
{Drive}:\$Recycle.Bin	C:\Recycle.Bin	C:\$Recycle.Bin
%SystemDrive%	C:	C:
%SystemRoot%	The Windows directory, usually C:\Windows, formerly C:\WINNT	%SystemDrive%\Windows
%TEMP% and %TMP%	%SystemDrive%\Documents and Settings\{username}\Local Settings\Temp	%SystemDrive%\Users\{username}\AppData\Local\Temp
%USERDOMAIN%	{userdomain}	{userdomain}
%USERDATA%	%SystemDrive%\Documents and Settings\{username}	%SystemDrive%\Users\{username}
%USERNAME%	{username}	{username}
%USERPROFILE%	%SystemDrive%\Documents and Settings\{username}	%SystemDrive%\Users\{username}
%WINDIR%	C:\Windows	C:\Windows

In this list, there is no environment variable that refers to the location of the user's My Documents folder, so there is no standard method for setting a program's home directory to be the My Documents folder.

External links

Unix

- `environ(7)` [1]: user environment – Linux Conventions and Miscellany Manual

Windows

- Environment Variables Wiki [2] — Wiki of Windows XP, Vista & 7 environment variables
- Eveditor [3] — Modern environment variables editor for Windows
- User Environment Variables [4] — Microsoft article on User Environment Variables and User Profiles
- Environment Variable Reference [5] — *Has a list showing which environment variables are for 9x WinNTx etc*
- Windows XP Command Shell Overview [6] with a list of environment variables — *Microsoft.com*
- How To Manage Environment Variables in Windows XP [7] — *Microsoft.com*
- Path Manager (pathman.exe) [8] — *Command line tool from Microsoft for editing PATH environment variable on Windows*
- Environment Variables in Windows XP [9] — *Computer Education*
- RapidEE (Rapid Environment Editor) [10] — *Windows environment variables editor*
- (EnvMan) Windows Environment Variables Manager [11] — *Environment Variables Editor for Windows*
- Getting Environment Variables in .NET [12] — *How to get environment variables in .NET*

References

[1] http://linux.die.net/man/7/environ
[2] http://environmentvariables.org
[3] http://eveditor.com
[4] http://msdn.microsoft.com/en-us/library/windows/desktop/bb776899%28v=vs.85%29.aspx
[5] http://www.scriptlogic.com/support/CustomScripts/environmentVariableReference.html
[6] http://technet.microsoft.com/en-gb/library/bb490954.aspx
[7] http://support.microsoft.com/default.aspx?scid=kb;en-us;310519
[8] http://www.microsoft.com/downloads/details.aspx?FamilyID=33543464-e0fb-4b36-80d0-4e0aba57d47f&DisplayLang=en
[9] http://vlaurie.com/computers2/Articles/environment.htm
[10] http://www.rapidee.com
[11] http://env-man.blogspot.com/2007/04/envman-user-guide.html
[12] http://msdn.microsoft.com/en-us/library/system.environment.getenvironmentvariable.aspx

Configuration file

In computing, **configuration files**, or **config files** configure the initial settings for some computer programs. They are used for user applications, server processes and operating system settings. The files are often written in ASCII (rarely UTF-8) and line-oriented, with lines terminated by a newline or carriage return/line feed pair, depending on the operating system. They may be considered a simple database.

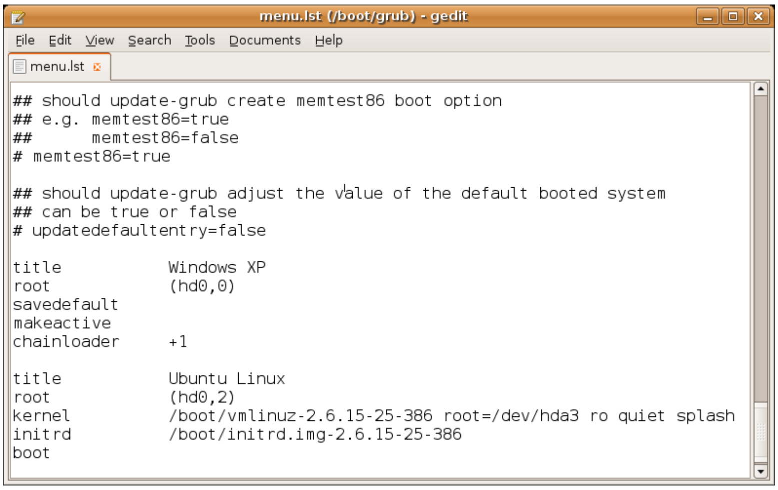

A configuration file for GNU GRUB being modified with gedit. This file contains a list of Operating Systems, which GNU GRUB reads and presents to the user as a menu.

Some applications provide tools to create, modify, and verify the syntax of their configuration files; these sometimes have graphical interfaces. For other programs, system administrators may be expected to create and modify files by hand using a text editor. For server processes and operating-system settings, there is often no standard tool, but operating systems may provide their own graphical interfaces such as YaST or debconf.

Some computer programs only read their configuration files at startup. Others periodically check the configuration files for changes. Users can instruct some programs to re-read the configuration files and apply the changes to the current process, or indeed to read arbitrary files as a configuration file. There are no definitive standards or strong conventions.

UNIX/Linux

Across the Unix variants hundreds of configuration-file formats exist. Each application or service may have a unique format. Historically, Unix operating system settings were often modified only by editing configuration files. Almost all formats allow entries to be disabled by prepending a special comment character, turning that entry into a comment.

The configuration files on Unix-type operating systems are traditionally documented using manpages, though other forms of online help are also used. In many cases the default configuration files distributed with a program contain extensive internal documentation in the form of comments. It is rare for a file to be completely undocumented, except in cases where a graphical configuration tool is the preferred method of configuring a program.

Unix user applications often create a file or directory in the home directory of the user upon startup. To hide the file or directory from casual listing of the contents of the home directory, the name of the file or directory is prepended with a period, giving rise to the nickname "dotfile" or "dot file". Server processes often use configuration files stored in `/etc`, but they may also use their installation directory or a location defined by the system administrator.

Configuration files also do more than just modify settings, they often (in the form of an "rc file") run a set of commands upon startup (for example, the "rc file" for a shell might instruct the shell to change directories, run certain programs, delete or create files —- many things which do not involve modifying variables in the shell itself and so were not in the shell's dotfiles); according to the Jargon File, this convention is borrowed from "runcom files" on the CTSS operating system;[1] see run commands for details. This functionality can and has been extended for programs written in interpreted languages such that the configuration file is actually another program rewriting or extending or customizing the original program; Emacs is the most prominent such example. The "rc" naming convention of "rc files" was inspired by the "runcom" facility mentioned above and does not stand for "resource configuration" or "runtime configuration" as is often wrongly guessed.[1]

"rc" files are traditionally files which end in the ".rc" suffix and which contain data and information that is used as configuration information for the associated program. Typically the name of that program is the first part of the rc file's name, with the ".rc" suffix being used to indicate the file's purpose.

On UNIX variants dot files remain "hidden" from listing by default. On Mac OS X these files are sometimes called "hidden files" although other mechanisms exist on Mac OS X to hide a file from view in various tools. The Explorer interface of Microsoft Windows XP does not allow the user to rename a file with an initial '.' though it does allow access to such files, and Windows' Notepad program does allow files to be saved with such names. Where Unix programs that use dotfiles are ported to Windows, they are sometimes modified to accept some other naming convention; for example, GNU Emacs permits its configuration file to be named `_emacs` instead of `.emacs`.[2]

IBM's AIX uses an Object Data Manager (ODM) database to store some system settings, some of which need to be available at boot time.

Microsoft DOS

DOS primarily relies on two files called `CONFIG.SYS` and `AUTOEXEC.BAT`. These were retained up to Windows 98SE, but were not strictly required to run Windows applications.

Microsoft Windows

The Microsoft Windows family of operating systems and their attendant applications utilize a similar system of configuration files. Windows 3.0 had an API for INI files (from "initialization"). Many Windows programs abandoned configuration files to use the Windows Registry to store information.

IBM OS/2

IBM's OS/2 uses a binary format, also with a .INI suffix, but this differs from the Windows versions. It contains a list of lists of untyped key-value pairs.[3] Two files control system-wide settings: OS2.INI and OS2SYS.INI. Application developers can choose whether to use them or create a specific file for their applications.

Configuration languages

Many language specifications have been created specifically to describe and retain configurations. These are frequently not Turing complete (nor need to be, by definition). A notable exception is Lua, which started out specifically as a configuration language for use in other programs. It evolved into a complete programming language, but retains a phrasing that allows configuration descriptions to be read directly into a native, stateful, tabulated set of variable-key pairings accessible to other programs (via a library), as well as allowing (self or external) invocation of commands to augment configuration activities.

The class includes all markup languages. The trend in the increase of XML and YAML (among other formats) for use as configuration-file formats is at least partially attributable to the increase in popularity of open source and platform neutral software applications and libraries. Moreover, the specifications describing these formats are routinely made available to the public, thus increasing the availability of parsers and emitters across programming languages.

References

[1] "rc file" (http://www.catb.org/jargon/html/R/rc-file.html). Catb.org. . Retrieved 2012-02-29.
[2] "Installing Emacs - GNU Emacs FAQ For MS Windows" (http://www.gnu.org/software/emacs/windows/Installing-Emacs.html). Gnu.org. . Retrieved 2012-02-29.
[3] *The OS/2 INI Files* (http://www.os2voice.org/VNL/past_issues/VNL0904H/vnewsf4.htm) by James J. Weinkam.

External links

- libprf1 (http://libprf1.tigris.org/) - alpha cross-platform, multi-language support for accessing the Preferences Registry Format (PRF) 1.0 configuration files

Windows Registry

The **Windows Registry** is a hierarchical database that stores configuration settings and options on Microsoft Windows operating systems. It contains settings for low-level operating system components as well as the applications running on the platform. The kernel, device drivers, services, SAM, user interface and third party applications all make use of the registry. The registry also provides a means to access counters for profiling system performance.

When first introduced with Windows 3.1, the Windows registry's primary purpose was to store configuration information for COM-based components. With the introduction of Windows 95 and Windows NT, its use was extended to tidy up the profusion of per-program INI files that had previously been used to store configuration settings for Windows programs.[1][2] It is not a requirement for a Windows application to use the Windows Registry—for example, .Net Framework applications use XML files for configuration, while portable applications usually keep their configuration data within files in the directory/folder where the application executable resides.

Rationale

Prior to the Registry, .INI files stored each program's settings into a text file, often located in a shared location that did not allow for user-specific settings in a multi-user scenario. By contrast, the Windows registry stores all application settings in one logical repository (but a number of discrete files) and in a standardized form. According to Microsoft this offers several advantages over .INI files.[2][3] Since file parsing is done much more efficiently with a binary format, it may be read from or written to more quickly than an INI file. As well, strongly typed data can be stored in the registry, as opposed to the text information stored in .INI files. This is a benefit when editing keys manually using `regedit.exe`, the built-in Windows registry editor. Because user-based registry settings are loaded from a user-specific path rather than from a read-only system location, the registry allows multiple users to share the same machine, and also allows programs to work for less privileged users. Backup and restoration is also simplified as the registry can be accessed over a network connection for remote management/support, including from scripts, using the standard set of APIs, as long as the Remote Registry service is running and firewall rules permit this.

As the registry is constructed as a database, it offers improved system integrity with features such as atomic updates. If two processes attempt to update the same registry value at the same time, one process's change will precede the others and the overall consistency of the data will be maintained. Where changes are made to .INI files, such race conditions can result in inconsistent data which doesn't match either attempted update. Windows Vista and later operating systems provide transactional updates to the registry by means of the Kernel Transaction Manager, extending the atomicity guarantees across multiple key and/or value changes, with traditional commit-abort semantics. (Note however that NTFS provides such support for the file system as well, so the same guarantees could, in theory, be obtained with traditional configuration files.)

Structure

Keys and values

The registry contains two basic elements: keys and values.

Registry *keys* are similar to folders—in addition to values, each key can contain subkeys, which may contain further subkeys, and so on. Keys are referenced with a syntax similar to Windows' path names, using backslashes to indicate levels of hierarchy. Each subkey has a mandatory name, which is a non-empty string that cannot contain any backslash, and whose letter case is insignificant.

The hierarchy of registry keys can only be accessed from a known root key handle (which is anonymous but whose effective value is a constant numeric handle) that is mapped to the content of a registry key preloaded by the kernel from a stored "hive", or to the content of a subkey within another root key, or mapped to a registered service or DLL that provides access to its contained subkeys and values.

E.g. HKEY_LOCAL_MACHINE\Software\Microsoft\Windows refers to the subkey "Windows" of the subkey "Microsoft" of the subkey "Software" of the HKEY_LOCAL_MACHINE root key.

There are seven predefined root keys, traditionally named according to their constant handles defined in the Win32 API, or by synonymous abbreviations (depending on applications):

- HKEY_LOCAL_MACHINE or HKLM
- HKEY_CURRENT_CONFIG or HKCC (only in Windows 9x/Me and NT-based versions of Windows)
- HKEY_CLASSES_ROOT or HKCR
- HKEY_CURRENT_USER or HKCU
- HKEY_USERS or HKU
- HKEY_PERFORMANCE_DATA (only in NT-based versions of Windows, but invisible in the Windows Registry Editor)
- HKEY_DYN_DATA (only in Windows 9x/Me, and visible in the Windows Registry Editor)

Like other files and services in Windows, all registry keys may be restricted by access control lists (ACLs), depending on user privileges, or on security tokens acquired by applications, or on system security policies enforced by the system (these restrictions may be predefined by the system itself, and configured by local system administrators or by domain administrators). Different users, programs, services or remote systems may only see some parts of the hierarchy or distinct hierarchies from the same root keys.

Registry *values* are name/data pairs stored within keys. Registry values are referenced separately from registry keys. Each registry value stored in a registry key has a unique name whose letter case is not significant. The Windows API functions that query and manipulate registry values take value names separately from the key path and/or handle that identifies the parent key. Registry values may contain backslashes in their names, but doing so makes them difficult to distinguish from their key paths when using some legacy Windows Registry API functions (whose usage is deprecated in Win32).

The terminology is somewhat misleading, as each registry key is similar to an associative array, where standard terminology would refer to the name part of each registry value as a "key". The terms are a holdout from the 16-bit registry in Windows 3, in which registry keys could not contain arbitrary name/data pairs, but rather contained only one unnamed value (which had to be a string). In this sense, the entire registry was like a single associative array where the registry keys (in both the registry sense and dictionary sense) formed a hierarchy, and the registry values were all strings. When the 32-bit registry was created, so was the additional capability of creating multiple named values per key, and the meanings of the names were somewhat distorted.[4] For compatibility with the previous behavior, each registry key may have a "default" value, whose name is the empty string.

Each value can store arbitrary data with variable length and encoding, but which is associated with a symbolic type (defined as a numeric constant) defining how to parse this data. The standard types are:

List of standard registry value types

Type ID	Symbolic type name	Meaning and encoding of the data stored in the registry value
0	REG_NONE	No type (the stored value, if any)
1	REG_SZ	A string value, normally stored and exposed in UTF-16LE (when using the Unicode version of Win32 API functions), usually terminated by a null character
2	REG_EXPAND_SZ	An "expandable" string value that can contain environment variables, normally stored and exposed in UTF-16LE, usually terminated by a null character
3	REG_BINARY	Binary data (any arbitrary data)
4	REG_DWORD / REG_DWORD_LITTLE_ENDIAN	A DWORD value, a 32-bit unsigned integer (numbers between 0 and 4,294,967,295 [$2^{32} - 1$]) (little-endian)
5	REG_DWORD_BIG_ENDIAN	A DWORD value, a 32-bit unsigned integer (numbers between 0 and 4,294,967,295 [$2^{32} - 1$]) (big-endian)
6	REG_LINK	A symbolic link (UNICODE) to another registry key, specifying a root key and the path to the target key
7	REG_MULTI_SZ	A multi-string value, which is an ordered list of non-empty strings, normally stored and exposed in UTF-16LE, each one terminated by a null character, the list being normally terminated by a second null character.
8	REG_RESOURCE_LIST	A resource list (used by the *Plug-n-Play* hardware enumeration and configuration)
9	REG_FULL_RESOURCE_DESCRIPTOR	A resource descriptor (used by the *Plug-n-Play* hardware enumeration and configuration)
10	REG_RESOURCE_REQUIREMENTS_LIST	A resource requirements list (used by the *Plug-n-Play* hardware enumeration and configuration)
11	REG_QWORD / REG_QWORD_LITTLE_ENDIAN	A QWORD value, a 64-bit integer (either big- or little-endian, or unspecified) (Introduced in Windows XP)

Hives

The Registry comprises a number of logical sections, or "hives"[5] (the word hive constitutes an in-joke).[6] Hives are generally named by their Windows API definitions, which all begin "HKEY". They are frequently abbreviated to a three- or four-letter short name starting with "HK" (e.g. HKCU and HKLM). Technically, they are predefined handles (with known constant values) to specific keys that are either maintained in memory, or stored in hive files stored in the local filesystem and loaded by the system kernel at boot time and then shared (with various access rights) between all processes running on the local system, or loaded and mapped in all processes started in a user session when the user logs on the system.

The HKEY_LOCAL_MACHINE (local machine-specific configuration data) and HKEY_CURRENT_USER (user-specific configuration data) nodes have a similar structure to each other; user applications typically look up their settings by first checking for them in "HKEY_CURRENT_USER\Software\Vendor's name\Application's name\Version\Setting name", and if the setting is not found, look instead in the same location under the HKEY_LOCAL_MACHINE key. However, the converse may apply for administrator-enforced policy settings where HKLM may take precedence over HKCU. The Windows Logo Program has specific requirements for where different types of user data may be stored, and that the concept of least privilege be followed so that administrator-level access is not required to use an application.[7][8]

HKEY_LOCAL_MACHINE (HKLM)

Abbreviated HKLM, HKEY_LOCAL_MACHINE stores settings that are specific to the local computer.[9]

The key located by HKLM is actually not stored on disk, but maintained in memory by the system kernel in order to map there all other subkeys. Applications cannot create any additional subkeys. On NT-based versions of Windows, this key contains four subkeys, "SAM", "SECURITY", "SYSTEM", and "SOFTWARE", that are loaded at boot time within their respective files located in the %SystemRoot%\System32\config folder. A fifth subkey, "HARDWARE", is volatile and is created dynamically, and as such is not stored in a file (it exposes a view of all the currently detected Plug-and-Play devices). On Windows Vista and above, a sixth subkey is mapped in memory by the kernel and populated from boot configuration data (BCD).

- The "HKLM\SAM" key usually appears as empty for most users (unless they are granted access by administrators of the local system or administrators of domains managing the local system). It is used to reference all "Security Accounts Manager" (SAM) databases for all domains into which the local system has been administratively authorized or configured (including the local domain of the running system, whose SAM database is stored a subkey also named "SAM": other subkeys will be created as needed, one for each supplementary domain). Each SAM database contains all builtin accounts (mostly group aliases) and configured accounts (users, groups and their aliases, including guest accounts and administrator accounts) created and configured on the respective domain, for each account in that domain, it notably contains the user name which can be used to log on that domain, the internal unique user identifier in the domain, a cryptographic hash of each user's password for each enabled authentication protocol, the location of storage of their user registry hive, various status flags (for example if the account can be enumerated and be visible in the logon prompt screen), and the list of domains (including the local domain) into which the account was configured.
- The "HKLM\SECURITY" key usually appears empty for most users (unless they are granted access by users with administrative privileges) and is linked to the Security database of the domain into which the current user is logged on (if the user is logged on the local system domain, this key will be linked to the registry hive stored by the local machine and managed by local system administrators or by the builtin "System" account and Windows installers). The kernel will access it to read and enforce the security policy applicable to the current user and all applications or operations executed by this user. It also contains a "SAM" subkey which is dynamically linked to the SAM database of the domain onto which the current user is logged on.
- The "HKLM\SYSTEM" key is normally only writable by users with administrative privileges on the local system. It contains information about the Windows system setup, data for the secure random number generator (RNG), the list of currently mounted devices containing a filesystem, several numbered "HKLM\SYSTEM\Control Sets" containing alternative configurations for system hardware drivers and services running on the local system (including the currently used one and a backup), a "HKLM\SYSTEM\Select" subkey containing the status of these Control Sets, and a "HKLM\SYSTEM\CurrentControlSet" which is dynamically linked at boot time to the Control Set which is currently used on the local system. Each configured Control Set contains:
  - an "Enum" subkey enumerating all known Plug-and-Play devices and associating them with installed system drivers (and storing the device-specific configurations of these drivers),
  - a "Services" subkey listing all installed system drivers (with non device-specific configuration, and the enumeration of devices for which they are instantiated) and all programs running as services (how and when they can be automatically started),
  - a "Control" subkey organizing the various hardware drivers and programs running as services and all other system-wide configuration,
  - a "Hardware Profiles" subkey enumerating the various profiles that have been tuned (each one with "System" or "Software" settings used to modify the default profile, either in system drivers and services or in the applications) as well as the "Hardware Profiles\Current" subkey which is dynamically linked to one of these profiles.

- The "HKLM\SOFTWARE" subkey contains software and Windows settings (in the default hardware profile). It is mostly modified by application and system installers. It is organized by software vendor (with a subkey for each), but also contains a "Windows" subkey for some settings of the Windows user interface, a "Classes" subkey containing all registered associations from file extensions, MIME types, Object Classes IDs and interfaces IDs (for OLE, COM/DCOM and ActiveX), to the installed applications or DLLs that may be handling these types on the local machine (however these associations are configurable for each user, see below), and a "Policies" subkey (also organized by vendor) for enforcing general usage policies on applications and system services (including the central certificates store used for authentifying, authorizing or disallowing remote systems or services running outside of the local network domain).
- The "HKLM\SOFTWARE\Wow6432Node" key is used by 32-bit applications on a 64-bit Windows OS, and is equivalent but separate to "HKLM\SOFTWARE". The key path is transparently presented to 32-bit applications by WoW64 as HKLM\SOFTWARE[10] (in a similar way that 32-bit applications see %SystemRoot%\Syswow64 as %SystemRoot%\System32)

HKEY_CURRENT_CONFIG (HKCC)

Abbreviated HKCC, HKEY_CURRENT_CONFIG contains information gathered at runtime; information stored in this key is not permanently stored on disk, but rather regenerated at boot time. It is a handle to the key "HKEY_LOCAL_MACHINE\System\CurrentControlSet\Hardware Profiles\Current", which is initially empty but populated at boot time by loading one of the other subkeys stored in "HKEY_LOCAL_MACHINE\System\CurrentControlSet\Hardware Profiles".

HKEY_CLASSES_ROOT (HKCR)

Abbreviated HKCR, HKEY_CLASSES_ROOT contains information about registered applications, such as file associations and OLE Object Class IDs, tying them to the applications used to handle these items. On Windows 2000 and above, HKCR is a compilation of user-based HKCU\Software\Classes and machine-based HKLM\Software\Classes. If a given value exists in both of the subkeys above, the one in HKCU\Software\Classes takes precedence.[11] The design allows for either machine- or user-specific registration of COM objects. The user-specific classes hive, unlike the HKCU hive, does not form part of a roaming user profile.

HKEY_USERS (HKU)

Abbreviated HKU, HKEY_USERS contains subkeys corresponding to the HKEY_CURRENT_USER keys for each user profile actively loaded on the machine, though user hives are usually only loaded for currently logged-in users.

HKEY_CURRENT_USER (HKCU)

Abbreviated HKCU, HKEY_CURRENT_USER stores settings that are specific to the currently logged-in user.[12] The HKEY_CURRENT_USER key is a link to the subkey of HKEY_USERS that corresponds to the user; the same information is accessible in both locations. On Windows NT-based systems, each user's settings are stored in their own files called NTUSER.DAT and USRCLASS.DAT inside their own Documents and Settings subfolder (or their own Users sub folder in Windows Vista and above). Settings in this hive follow users with a roaming profile from machine to machine.

HKEY_PERFORMANCE_DATA

This key provides runtime information into performance data provided by either the NT kernel itself, or running system drivers, programs and services that provide performance data. This key is not stored in any hive and not displayed in the Registry Editor, but it is visible through the registry functions in the Windows API, or in a simplified view via the Performance tab of the Task Manager (only for a few performance data on the local system) or via more advanced control panels (such as the Performances Monitor or the Performances Analyzer which allows collecting and logging these data, including from remote systems).

HKEY_DYN_DATA

This key is used only on Windows 95, Windows 98 and Windows Me.[13] It contains information about hardware devices, including Plug and Play and network performance statistics. The information in this hive is also not stored on the hard drive. The Plug and Play information is gathered and configured at startup and is stored in memory.[14]

Editing

Manual editing

The Windows registry can be edited manually using programs such as regedit.exe and on older versions of Windows, regedt32.exe, although these tools do not expose some of registry's metadata such as the last modified date. They also implement workarounds in code that allow Registry keys to be renamed, as the underlying APIs do not support this capability.

As a careless change could cause irreversible damage, a backup of the registry before editing is recommended by Microsoft. Incorrect changes made to the registry with the Registry Editor can cause system-wide problems in Windows, which may require re-installing Windows to fix.[15] Editing the registry is sometimes necessary when working around Windows-specific issues e.g. problems when logging onto a domain can be resolved by editing the registry.[16]

A simple implementation of the current registry tool appeared in Windows 3.x, called the "Registration Info Editor" or "Registration Editor". This was basically just a database of applications used to edit embedded OLE objects in documents.

Windows 9x operating systems included REGEDIT.EXE which could be used in Windows and also in real mode MS-DOS.[17] Windows NT introduced permissions for Registry editing. Windows NT 4.0 and Windows 2000 were distributed with both the Windows 9x REGEDIT.EXE program and Windows NT 3.x's REGEDT32.EXE program. There were several differences between the two editors on these platforms:

- REGEDIT.EXE had a left-side tree view that begins at "My Computer" and lists all loaded hives. REGEDT32.EXE had a left-side tree view, but each hive had its own window, so the tree displays only keys.
- REGEDIT.EXE represented the three components of a value (its name, type, and data) as separate columns of a table. REGEDT32.EXE represented them as a list of strings.
- REGEDIT.EXE supported right-clicking of entries in a tree view to adjust properties and other settings. REGEDT32.EXE required all actions to be performed from the top menu bar.
- REGEDIT.EXE supported searching for key names, values, or data throughout the entire registry, whereas REGEDT32.EXE only supported searching for key names in one hive at a time.
- Earlier versions of REGEDIT.EXE did not support editing permissions. Therefore, on those early versions, only REGEDT32.EXE could access the full functionality of an NT registry. REGEDIT.EXE in Windows XP and above supports editing permissions.
- REGEDIT.EXE only supported string (REG_SZ), binary (REG_BINARY), and DWORD (REG_DWORD) values. REGEDT32.EXE supported those, plus expandable string (REG_EXPAND_SZ) and multi-string (REG_MULTI_SZ). Attempting to edit unsupported key types with REGEDIT.EXE on Windows 2000 or Windows NT 4.0 would result in irreversible conversion to a supported type.[18]

Windows XP was the first system to integrate these two programs into one, adopting the old REGEDIT.EXE interface and adding the REGEDT32.EXE functionality. The differences listed above are not applicable on Windows XP and newer systems; REGEDIT.EXE is the improved editor, and REGEDT32.EXE is deprecated. On Windows XP and above, the Registry Editor also supports multiple instances if the -m switch is specified.

The Registry Editor allows users to perform the following functions:

- Creating, manipulating, renaming and deleting registry keys, subkeys, values and value data
- Importing and exporting `.REG` files, exporting data in the binary hive format
- Loading, manipulating and unloading registry hive format files (Windows NT-based systems only)
- Setting permissions based on ACLs (Windows NT-based systems only)
- Bookmarking user-selected registry keys as Favorites
- Finding particular strings in key names, value names and value data
- Remotely editing the registry on another networked computer

It is also possible to edit the registry under Linux using the open source Offline NT Password & Registry Editor to edit the files.[19]

`.REG` files

`.REG` files (also known as Registration entries) are text-based human-readable files for exporting and importing portions of the registry. On Windows 2000 and later NT-based operating systems, they contain the string *Windows Registry Editor Version 5.00* at the beginning and are Unicode-based. On Windows 9x and NT 4.0 systems, they contain the string *REGEDIT4* and are ANSI-based.[20] Windows 9x format `.REG` files are compatible with Windows 2000 and later NT-based systems. The Registry Editor on Windows on these systems also supports exporting `.REG` files in Windows 9x/NT format. Data is stored in `.REG` files in the following syntax:[20]

```
[<Hive Name>\<Key Name>\<Subkey Name>]
"Value Name"=<Value type>:<Value data>
```

The Default Value of a key can be edited by using @ instead of "Value Name":

```
[<Hive Name>\<Key Name>\<Subkey Name>]
@=<Value type>:<Value data>
```

String values do not require a <Value type> (see example), but backslashes ("\") need to be written as a double-backslash ("\\").

For example, to add the values "Value A", "Value B", "Value C", "Value D", "Value E", "Value F", "Value G", "Value H", "Value I", "Value J", and "Value K" to the HKLM\SOFTWARE\Microsoft key,

```
Windows Registry Editor Version 5.00
[HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft]
"Value A"="<String value data>"
"Value B"=hex:<Binary data (as comma-delimited list of hexadecimal values)>
"Value C"=dword:<DWORD value integer>
"Value D"=hex(7):<Multi-string value data (as comma-delimited list of hexadecimal values)>
"Value E"=hex(2):<Expandable string value data (as comma-delimited list of hexadecimal values)>
"Value F"=hex(b):<QWORD value (as comma-delimited list of 8 hexadecimal values, in little endian byte order)>
"Value G"=hex(4):<DWORD value (as comma-delimited list of 4 hexadecimal values, in little endian byte order)>
"Value H"=hex(5):<DWORD value (as comma-delimited list of 4 hexadecimal values, in big endian byte order)>
"Value I"=hex(8):<REG_RESOURCE_LIST (as comma-delimited list of hexadecimal values)>
"Value J"=hex(a):<REG_RESOURCE_REQUIREMENTS_LIST (as comma-delimited list of hexadecimal values)>
"Value K"=hex(0):
```

Data from `.REG` files can be added/merged with the registry by double-clicking these files or using the /s switch in the command line. `.REG` files can also be used to remove registry data.

To remove a key (and all subkeys, values and data), the key name must be preceded by a minus sign ("-").[20]

For example, to remove the HKLM\SOFTWARE\Microsoft key (and all subkeys, values and data),

```
[-HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft]
```

To remove a value (and its data), the values to be removed must have a minus sign ("-") after the equal sign ("=").[20]

For example, to remove only the "Value A" and "Value B" values (and their data) from the HKLM\SOFTWARE\Microsoft key,

```
[HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft]
"Value A"=-
"Value B"=-
```

To remove only the (Default) value of the key HKLM\SOFTWARE\Microsoft (and its data),

```
[HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft]
@=-
```

Lines beginning with a semicolon are considered as comments:

```
; This is a comment. This can be placed in any part of a .reg file
[HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft]
"Value"="Example string"
```

Command line editing

The registry can be manipulated in a number of ways from the command line. The `Reg.exe` and `RegIni.exe` utility tools are included in Windows XP and later versions of Windows. Alternative locations for legacy versions of Windows include the Resource Kit CDs or the original Installation CD of Windows.

Also, a `.REG` file can be imported from the command line with the following command:

RegEdit.exe /s *file*

The /s means the file will be *silent merged* to the Registry. If the `/s` parameter is omitted the user will be asked to confirm the operation. In Windows 98, Windows 95 and at least some configurations of Windows XP the `/s` switch also causes `RegEdit.exe` to ignore the setting in the registry that allows administrators to disable it. When using the `/s` switch `RegEdit.exe` does not return an appropriate return code if the operation fails, unlike `Reg.exe` which does.

```
RegEdit.exe /e file
```

exports the whole registry to a `.REG` file, while any of

```
RegEdit.exe /e file HKEY_CLASSES_ROOT[\<key>]
RegEdit.exe /e file HKEY_CURRENT_CONFIG[\<key>]
RegEdit.exe /e file HKEY_CURRENT_USER[\<key>]
RegEdit.exe /e file HKEY_LOCAL_MACHINE[\<key>]
RegEdit.exe /e file HKEY_USERS[\<key>]
```

export the specified (sub)key only.

The default association for `.REG` files in many versions of Microsoft Windows.

We can use also `Reg.exe`. Here is a sample to display the value of the registry value Version:

```
Reg.exe QUERY HKLM\Software\Microsoft\ResKit /v Version
```

Other command line options include a VBScript or JScript together with CScript, WMI or `WMIC.exe` and Windows PowerShell.

Registry permissions can be manipulated through the command line using `RegIni.exe` and the `SubInACL.exe` [21] tool. For example, the permissions on the HKEY_LOCAL_MACHINE\SOFTWARE key can be displayed using:

```
SubInACL.exe /keyreg HKEY_LOCAL_MACHINE\SOFTWARE /display
```

Programs or scripts

The registry can be edited through the APIs of the Advanced Windows 32 Base API Library (advapi32.dll).[22]

List of Registry API functions			
RegCloseKey	RegOpenKey	RegConnectRegistry	RegOpenKeyEx
RegCreateKey	RegQueryInfoKey	RegCreateKeyEx	RegQueryMultipleValues
RegDeleteKey	RegQueryValue	RegDeleteValue	RegQueryValueEx
RegEnumKey	RegReplaceKey	RegEnumKeyEx	RegRestoreKey
RegEnumValue	RegSaveKey	RegFlushKey	RegSetKeySecurity
RegGetKeySecurity	RegSetValue	RegLoadKey	RegSetValueEx
	RegNotifyChangeKeyValue	RegUnLoadKey	

Many programming languages offer built-in runtime library functions or classes that wrap the underlying Windows APIs and thereby enable programs to store settings in the registry (e.g. `Microsoft.Win32.Registry` in VB.NET and C#, or `TRegistry` in Delphi and Free Pascal). COM-enabled applications like Visual Basic 6 can use the WSH `WScript.Shell` object. Another way is to use the Windows Resource Kit Tool, `Reg.exe` by executing it from code,[23] although this is considered poor programming practice.

Similarly, scripting languages such as Perl (with `Win32::TieRegistry` [24]), Windows Powershell and Windows Scripting Host also enable registry editing from scripts.

COM self-registration

Prior to the introduction of Registration-Free COM, developers were encouraged to add initialization code to in-process and out-of-process binaries to perform the Registry configuration required for that object to work. For in-process binaries such as .DLL and .OCX files, the modules typically exported a function called DllInstall()[25] that could be called by installation programs or invoked manually with utilities like Regsvr32.exe;[26] out-of-process binaries typically support the commandline arguments /Regserver and /Unregserver that created or deleted the required Registry settings.[27] COM applications that break because of DLL Hell issues can commonly be repaired with RegSvr32.exe or the /RegServer switch without having to re-invoke installation programs.[28]

Advanced functionality

Windows exposes APIs that allows user-mode applications to register to receive a notification event if a particular registry key is changed.[29] APIs are also available to allow kernel-mode applications to filter and modify registry calls made by other applications.[30]

Windows also supports remote access to the registry of another computer via the `RegConnectRegistry` function[31] if the Remote Registry service is running, correctly configured and its network traffic is not firewalled.[32]

Locations

The Registry is physically stored in several files, which are generally obfuscated from the user-mode APIs used to manipulate the data inside the Registry. Depending upon the version of Windows, there will be different files and different locations for these files, but they are all on the local machine. The location for system Registry files in Windows NT is \Windows\System32\Config; the user-specific HKEY_CURRENT_USER user registry hive is stored in `Ntuser.dat` inside the user profile. There is one of these per user; if a user has a roaming profile, then this file will be copied to and from a server at logout and login respectively. A second user-specific Registry file named UsrClass.dat contains COM registry entries and does not roam by default.

Windows NT-based operating systems

Windows NT-based systems store the registry in a binary hive format which can be exported, loaded and unloaded by the Registry Editor in these operating systems. The following system Registry files are stored in `%SystemRoot%\System32\Config\`:

- `Sam` – HKEY_LOCAL_MACHINE\SAM
- `Security` – HKEY_LOCAL_MACHINE\SECURITY
- `Software` – HKEY_LOCAL_MACHINE\SOFTWARE
- `System` – HKEY_LOCAL_MACHINE\SYSTEM
- `Default` – HKEY_USERS\.DEFAULT
- `Userdiff` – Not associated with a hive. Used only when upgrading operating systems.[33]

The following file is stored in each user's profile folder:

- `%UserProfile%\Ntuser.dat` – HKEY_USERS\<User SID> (linked to by HKEY_CURRENT_USER)

For Windows 2000, Server 2003 and Windows XP, the following additional user-specific file is used for file associations and COM information:

- `%UserProfile%\Local Settings\Application Data\Microsoft\Windows\Usrclass.dat` (path is localized) – HKEY_USERS\<User SID>_Classes (HKEY_CURRENT_USER\Software\Classes)

For Windows Vista and later, the path was changed to:

- `%UserProfile%\AppData\Local\Microsoft\Windows\Usrclass.dat` (path is not localized) alias `%LocalAppData%\Microsoft\Windows\Usrclass.dat` – HKEY_USERS\<User SID>_Classes (HKEY_CURRENT_USER\Software\Classes)

Windows 2000 kept an alternate copy of the registry hives (.ALT) and attempts to switch to it when corruption is detected.[34] Windows XP and Windows Server 2003 do not maintain a `System.alt` hive because NTLDR on those versions of Windows can process the `System.log` file to bring up to date a System hive that has become inconsistent during a shutdown or crash. In addition, the `%SystemRoot%\Repair` folder contains a copy of the system's registry hives that were created after installation and the first successful startup of Windows.

Windows 95, 98, and Me

The registry files are stored in the `%WINDIR%` directory under the names `USER.DAT` and `SYSTEM.DAT` with the addition of `CLASSES.DAT` in Windows Me. Also, each user profile (if profiles are enabled) has its own `USER.DAT` file which is located in the user's profile directory in `%WINDIR%\Profiles\<username>\`.

Windows 3.11

The only registry file is called `REG.DAT` and it is stored in the `%WINDIR%` directory.

Backups and recovery

Different editions of Windows have supported a number of different methods to back up and restore the registry of the years, some of which are now deprecated:

- System Restore can back up the registry and restore it as long as Windows is bootable, or from the Windows Recovery Environment starting with Windows Vista.
- NTBackup can back up the registry as part of the *System State* and restore it. Automated System Recovery in Windows XP can also restore the registry.
- On Windows NT-based systems, the *Last Known Good Configuration* option in startup menu relinks the `HKLM\SYSTEM\CurrentControlSet` key, which stores hardware and device driver information.
- Windows 98 and Windows Me include command line (Scanreg.exe) and GUI (Scanregw.exe) registry checker tools to check and fix the integrity of the registry, create up to five automatic regular backups by default and restore them manually or whenever corruption is detected. The registry checker tool backs up the registry, by default, to `%Windir%\Sysbckup` Scanreg.exe can also run from MS-DOS.
- The Windows 95 CD-ROM included an Emergency Recovery Utility (ERU.exe) and a Configuration Backup Tool (Cfgback.exe) to back up and restore the registry. Additionally Windows 95 backs up the registry to the files system.da0 and user.da0 on every successful boot.
- Windows NT 4.0 included `RDISK.EXE`, a utility to back up and restore the entire registry.[35]
- The Windows 2000 Resource Kit contained an unsupported pair of utilities called Regback.exe and RegRest.exe for backup and recovery of the Registry.[36]

Policy

Group policy

Windows 2000 and later versions of Windows use Group Policy to enforce Registry settings. Policy may be applied locally to a single computer using `gpedit.msc`, or to multiple users and/or computers in a domain using `gpmc.msc`.

Legacy systems

With Windows 95, Windows 98, Windows Me and Windows NT, administrators can use a special file to be merged into the registry, called a policy file (`POLICY.POL`). The policy file allows administrators to prevent non-administrator users from changing registry settings like, for instance, the security level of Internet Explorer and the desktop background wallpaper. The policy file is primarily used in a business with a large number of computers where the business needs to be protected from rogue or careless users.

The default extension for the policy file is `.POL`. The policy file filters the settings it enforces by user and by group (a "group" is a defined set of users). To do that the policy file merges into the registry, preventing users from circumventing it by simply changing back the settings. The policy file is usually distributed through a LAN, but can be placed on the local computer.

The policy file is created by a free tool by Microsoft that goes by the filename `poledit.exe` for Windows 95/Windows 98 and with a computer management module for NT-based systems. The editor requires administrative permissions to be run on systems that uses permissions. The editor can also directly change the current registry settings of the local computer and if the remote registry service is installed and started on another computer it can also change the registry on that computer. The policy editor loads the settings it can change from `.ADM` files, of which one is included, that contains the settings the Windows shell provides. The `.ADM` file is plain text and supports easy localisation by allowing all the strings to be stored in one place.

.INI file virtualization

Windows NT kernels support redirection of INI file-related APIs into a virtual file in a Registry location such as HKEY_CURRENT_USER using a feature called "InifileMapping".[37] This functionality was introduced to allow legacy applications written for 16-bit versions of Windows to be able to run under Windows NT platforms on which the System folder is no longer considered an appropriate location for user-specific data or configuration. Non-compliant 32-bit applications can also be redirected in this manner, even though the feature was originally intended for 16-bit applications.

Registry virtualization

Windows Vista has introduced limited Registry virtualization, whereby poorly written applications that do not respect the principle of least privilege and instead try to write user data to a read-only system location (such as the HKEY_LOCAL_MACHINE hive), can be redirected to a more appropriate location, without changing the application itself. The operation is transparent to the application, as it does not know that its Registry operations have been directed elsewhere.

Similarly, application virtualization redirects all of an application's Registry operations to a non-Registry backed location, such as a file. Used together with file virtualization, this approach allows applications to run without being installed on the location machine.

Low integrity processes may also use registry virtualization. For example as Internet Explorer 7 or 8 running in "Protected Mode" on Windows Vista and above will automatically redirect registry writes by ActiveX controls to a sandboxed location in order to frustrate some classes of security exploits.

Lastly, the Application Compatibility Toolkit[38] provides shims that can transparently redirect HKEY_LOCAL_MACHINE or HKEY_CLASSES_ROOT Registry operations to HKEY_CURRENT_USER to address "LUA" bugs that cause applications not to work for limited users.

Disadvantages

The Windows Registry has several disadvantages. Although it is relatively easy to back up and restore the database to the same Windows installation, moving per-program or user-specific settings to a different machine is difficult because so many programs use settings that are dependent on the local machine.[39]

Critics have also labeled the registry in Windows 95 as a single point of failure, because re-installation of the operating system is required if the Registry is corrupted. Windows NT introduced a transaction log to protect the database during updates.[40]

Its centralized and hierarchical structure, while providing the many benefits discussed above, also makes the registry a weak link in the system and an easy target for both malicious and unintentional damage by users.[40]

Equivalents in other operating systems

In contrast to the Windows registry's binary-based database model, some other operating systems use separate plain-text files for daemon and application configuration, but group these configurations together for ease of management.

- Under Unix-like operating systems e.g. Linux that follow the Filesystem Hierarchy Standard, system-wide configuration files (information similar to what would appear in HKEY_LOCAL_MACHINE on Windows) are traditionally stored in files in `/etc/` and its subdirectories, or sometimes in `/usr/local/etc`. Per-user information (information that would be roughly equivalent to that in HKEY_CURRENT_USER) is stored in hidden directories and files (that start with a period/full stop) within the user's home directory. However XDG-compliant applications should refer to the environment variables defined in the Base Directory specification.[41]
- Applications running on Apple Inc.'s Mac OS X operating system typically store settings in property list files which are usually stored in each user's Library folder.
- RISC OS uses directories for configuration data, which allows applications to be copied into application directories, as opposed to the separate installation process that typifies Windows applications; this approach is also used on the ROX Desktop for Linux.[42] This directory-based configuration also makes it possible to use different versions of the same application, since the configuration is done "on the fly".[43] If one wishes to remove the application, it is possible to simply delete the folder belonging to the application.[44][45] This will often not remove configuration settings which are stored independently from the application, usually within the computer's !Boot structure, in !Boot.Choices or potentially anywhere on a network fileserver. It is possible to copy programs between computers running RISC OS by copying the application directories belonging to the programs, however some programs may require re-installing, e.g. when shared files are placed outside of an application directory.[43]
- IBM AIX (a Unix variant) uses a registry component called Object Data Manager (ODM). The ODM is used to store information about system and device configuration. An extensive set of tools and utilities provides users with means of extending, checking, correcting the ODM database. The ODM stores its information in several files, default location is /etc/objrepos.
- The GNOME desktop environment uses a registry-like interface called GConf for storing configuration settings for the desktop and applications. However, in GConf, all application settings are stored in separate files, thereby partially eliminating a single point of failure.
- The Elektra Initiative provides an alternative back-end for text configuration files for the Linux operating system, similar to the registry.
- While not an operating system, the Wine compatibility layer, which allows Windows software to run on a Unix-like system, also employs a Windows-like registry as text files in the WINEPREFIX folder: system.reg (HKEY_LOCAL_MACHINE), user.reg (HKEY_CURRENT_USER) and userdef.reg.[46]

Notes

[1] "Windows 2000 Registry: Latest Features and APIs Provide the Power to Customize and Extend Your Apps" (http://msdn.microsoft.com/msdnmag/issues/1100/Registry/). . Retrieved 2007-07-19.

[2] "The System Registry" (http://msdn.microsoft.com/en-us/library/ms970651.aspx). .

[3] "Windows 95 Architecture Components". Microsoft. http://www.microsoft.com/technet/archive/win95/rk31_arc.mspx?mfr=true. Retrieved 2008-04-29. "The following table shows other difficulties or limitations caused by using .INI files that are overcome by using the Registry."

[4] Raymond Chen, "Why do registry keys have a default value?" (http://blogs.msdn.com/oldnewthing/archive/2008/01/18/7145021.aspx)

[5] "Registry hives" (http://msdn2.microsoft.com/en-us/library/ms724877.aspx). . Retrieved 2007-07-19.

[6] Chen, Raymond (2011-08-08). "Why is a registry file called a "hive"?" (http://blogs.msdn.com/b/oldnewthing/archive/2003/08/08/54618.aspx). The Old New Thing. . Retrieved 2011-07-29.

[7] When applications fail to execute because they request more privileges than they require (and are denied those privileges), this is known as a limited user application (LUA) bug.

[8] "Designed for Windows XP Application Specification" (http://www.microsoft.com/downloads/details.aspx?displaylang=en&FamilyID=209e3d65-f0be-4eef-8602-73bb9bc29d54). Microsoft. 2002-08-20. . Retrieved 2009-04-08.

[9] "HKEY_LOCAL_MACHINE" (http://technet.microsoft.com/en-us/library/cc959046.aspx). Microsoft. 2009. . Retrieved 2009-04-08.

[10] http://msdn.microsoft.com/en-us/library/aa384253(v=vs.85).aspx

[11] "Description of the Microsoft Windows registry" (http://support.microsoft.com/kb/256986). . Retrieved 2008-09-25.

[12] "HKEY_CURRENT_USER" (http://technet.microsoft.com/en-us/library/cc976337.aspx). Microsoft. 2009. . Retrieved 2009-04-08.

[13] Description of the HKEY_DYN_DATA Registry Key in Windows 95, Windows 98, and Windows 98 SE (http://support.microsoft.com/kb/250568)

[14] A Closer Look at HKEY_DYN_DATA (http://docs.rinet.ru/Registratura/htm/apf.htm)

[15] "Incorrect Icons Displayed for .ico Files" (http://support.microsoft.com/kb/171325). November 15, 2006. . Retrieved 31 March, 2012.

[16] "You may receive a "STOP 0x00000035 NO_MORE_IRP_STACK_LOCATIONS" error message when you try to log on to a domain" (http://support.microsoft.com/kb/906866). October 9, 2011. . Retrieved 31 March, 2012. This page tells the user to edit the registry when resolving the issue.

[17] Using Registry Editor in Real Mode (http://support.microsoft.com/kb/131352)

[18] Microsoft's *Windows 2000 Security Hardening Guide* (http://www.microsoft.com/technet/security/prodtech/windows2000/win2khg/05sconfg.mspx) version 1.3, published May 15, 2003

[19] Offline Registry Editor user manual/doc (http://pogostick.net/~pnh/ntpasswd/regedit.txt)

[20] How to add, modify, or delete registry subkeys and values by using a registration entries (.reg) file (http://support.microsoft.com/kb/310516/en-us)

[21] http://www.microsoft.com/downloads/details.aspx?FamilyID=e8ba3e56-d8fe-4a91-93cf-ed6985e3927b&DisplayLang=en

[22] "Reading and Writing Registry Values with Visual Basic" (http://www.windowsdevcenter.com/lpt/a/5016). . Retrieved 2007-07-19.

[23] "REG command in Windows XP" (http://www.petri.co.il/reg_command_in_windows_xp.htm). . Retrieved 2007-07-19.

[24] http://search.cpan.org/~adamk/Win32-TieRegistry-0.25/TieRegistry.pm

[25] "DllInstall Function" (http://msdn.microsoft.com/en-us/library/windows/desktop/bb759846(v=vs.85).aspx). Microsoft. 2012-03-07. . Retrieved 2012-03-22.

[26] "Regsvr32" (http://msdn.microsoft.com/en-us/library/ms859484.aspx). Microsoft. . Retrieved 2012-03-22.

[27] "How to: Register Automation Servers" (http://msdn.microsoft.com/en-us/library/05c9yccf(v=vs.80).aspx). Microsoft. . Retrieved 2012-03-22.

[28] "How to re-register PowerPoint 2000, PowerPoint 2003, PowerPoint 2007 and PowerPoint 2010" (http://support.microsoft.com/kb/291893). Microsoft. January 2012. . Retrieved 2012-03-22.

[29] "RegNotifyChangeKeyValue function" (http://msdn.microsoft.com/en-us/library/windows/desktop/ms724892(v=vs.85).aspx). Microsoft. .

[30] "Registering for Notifications" (http://msdn.microsoft.com/en-us/library/windows/hardware/ff560895(v=vs.85).aspx). Microsoft. .

[31] "RegConnectRegistry function" (http://msdn.microsoft.com/en-us/library/windows/desktop/ms724840(v=vs.85).aspx). Microsoft. .

[32] "How to Manage Remote Access to the Registry" (http://support.microsoft.com/kb/314837). Microsoft. .

[33] "Overview of the Windows NT Registry" (http://technet.microsoft.com/en-us/library/cc722562.aspx). . Retrieved 2011-12-02.

[34] "Inside the Registry" (http://technet.microsoft.com/en-us/library/cc750583.aspx). . Retrieved 2007-12-28.

[35] How To Backup, Edit, and Restore the Registry in Windows NT 4.0 (http://support.microsoft.com/kb/323170)

[36] "Technical Reference to the Registry: Related Resources" (http://technet.microsoft.com/en-us/library/cc939136.aspx). Microsoft. . Retrieved 2011-09-09.

[37] "Chapter 26 - Initialization Files and the Registry" (http://www.microsoft.com/technet/archive/ntwrkstn/reskit/26_ini.mspx?mfr=true). Microsoft. . Retrieved 2008-03-03.

[38] "Microsoft Application Compatibility Toolkit 5.0" (http://technet.microsoft.com/en-us/appcompat/aa905102.aspx). Microsoft. . Retrieved 2008-07-26.

[39] The Windows Registry (http://www.tech-pro.net/intro_reg.html)

[40] Forensic Analysis of the Windows Registry (http://www.docstoc.com/docs/2381418/Forensic-Analysis-of-the-Windows-Registry-Computer-Forensics)
[41] "XDG Base Directory Specification" (http://standards.freedesktop.org/basedir-spec/latest/index.html). .
[42] "Application directories" (http://roscidus.com/desktop/AppDirs). . Retrieved 17 May 2012.
[43] "Case Studies Of The Top 132 Annoyances With Operating Systems Other Than RISC OS" (http://productsdb.riscos.com/admin/ros_test.htm). . Retrieved 3 April, 2012. Page from the riscos.com website. Mentioned in points 82 and 104.
[44] "RISC OS tour" (http://www.riscos.org/tour/index.html). . Retrieved 2007-07-19.
[45] "The RISC OS Products Directory" (http://productsdb.riscos.com/admin/riscos.htm). 02 Nov 2006. . Retrieved 1 April, 2012.
[46] 3.2. Using the Registry and Regedit (http://www.winehq.org/docs/wineusr-guide/using-regedit) (Wine User Guide)

Footnotes

References

• Russinovich, Mark E.; Solomon, David A. (2005). *Microsoft Windows Internals* (Fourth Edition ed.). Microsoft Press. pp. 183–236. ISBN 978-0-7356-1917-3.

External links

• Windows Registry info & reference (http://msdn2.microsoft.com/en-us/library/ms724871.aspx) in the MSDN Library

Windows Management Instrumentation

Windows Management Instrumentation (**WMI**) is a set of extensions to the Windows Driver Model that provides an operating system interface through which instrumented components provide information and notification. WMI is Microsoft's implementation of the Web-Based Enterprise Management (WBEM) and Common Information Model (CIM) standards from the Distributed Management Task Force (DMTF).

WMI allows scripting languages like VBScript or Windows PowerShell to manage Microsoft Windows personal computers and servers, both locally and remotely. WMI is preinstalled in Windows 2000 and newer OSs. It is available as a download for Windows NT,[1] Windows 95 and Windows 98.[2]

Microsoft also provides a command line interface to WMI called **Windows Management Instrumentation Command-line** (**WMIC**).[3]

Purpose of WMI

The purpose of WMI is to define a proprietary set of environment-independent specifications which allow management information to be shared between management applications. WMI prescribes enterprise management standards and related technologies for Windows that work with existing management standards, such as Desktop Management Interface (DMI) and SNMP. WMI complements these other standards by providing a uniform model. This model represents the managed environment through which management data from any source can be accessed in a common way.

Development process

Because WMI abstracts the manageable entities with CIM and a collection of providers, the development of a provider implies several steps. The major steps can be summarized as follows:

Step 1 – Create the manageable entity model

- Define a model
- Implement the model

Step 2 – Create the WMI Provider

- Determine the provider type to implement
- Determine the hosting model of the provider
- Create the provider template with the ATL wizard
- Implement the code logic in the provider
- Register the provider with WMI and the system

Step 3 – Test the provider

Step 4 – Create consumer sample code

Importance of WMI providers

Since the release of the first WMI implementation during the Windows NT 4.0 SP4 era (as an out-of-band download), Microsoft has consistently added WMI providers to Windows. Under Windows NT 4.0, Microsoft had roughly 15 WMI providers available once WMI was installed. When Windows 2000 was released, there were 29 WMI providers as part of the operating system installation. With the release of Windows Server 2003, Microsoft included in the platform more than 80 WMI providers. Windows Vista includes 13 new WMI providers,[4] taking the number close to around 100 in all, and Windows Server 2008 includes some more including providers for IIS 7, PowerShell and virtualization. This has been a sign for many customers that WMI became at Microsoft, the “ubiquitous” management layer of Windows, even if this commitment has never been explicit from Microsoft.

During these last years, due to a constant increasing exposure of management data through WMI in Windows, more and more people in the IT systems management field started to develop scripts and automation procedures based on WMI. Beyond the scripting needs, most leading management software in the world, such as MOM, SMS, ADS, HP OpenView for Windows (HPOV), BMC Software or CA, Inc. are WMI-enabled and capable to consume and provide WMI information through various *User Interfaces*. This enables administrators and operators not capable of scripting or programming on top of WMI to enjoy the benefits of WMI without even learning about it. However, if they want to, because WMI is scriptable, it gives them the opportunity to consume WMI information from scripts or from any Enterprise Management software that is WMI-aware.

Features

For someone willing to develop one or many WMI providers, WMI offers many features out of the box. Here are the most important advantages:

1. *Automation interfaces:* Because WMI comes with a set of automation interfaces ready to use, all management features supported by a WMI provider and its set of classes get the scripting support for free out-of-the box. Beyond the WMI class design and the provider development, the Microsoft development and test teams are not required to create, validate and test a scripting model as it is already available from WMI.
2. *.NET Management interfaces:* Because the System.Management namespace [5] relies on the existing COM/DCOM plumbing, the created WMI provider and its set of WMI classes becomes automatically available to all .NET applications independently of the language used (e.g. C#, VB.NET). Beyond the WMI class design and the provider development, like for scripting, the Microsoft development and test teams are not required to create,

validate and test new assemblies to support a new namespace in the .NET Framework as this support is already available from WMI for free.

3. *C/C++ COM/DCOM programming interfaces:* Like most components in Windows, COM/DCOM programmers can leverage the features of the provider they develop at the COM/DCOM interfaces level. Like in previous environments (scripting and .NET Framework), a COM/DCOM consumer just needs to interact with the standard set of WMI COM interfaces to leverage the WMI provider capabilities and its set of supported WMI classes. To make all management information available from the native APIs, the WMI provider developer just needs to interact with a set of pre-defined WMI COM interfaces. This will make the management information available at the WMI COM level automatically. Moreover, the scripting COM interface object model is very similar to the COM/DCOM interface object model, which makes it easy for developers to be familiar with the scripting experience.
4. *Remoting capabilities over DCOM and SOAP:* More than simply offering local COM capabilities, as management is all about remoting, WMI offers the DCOM transport. In addition, SOAP transport will be available in Windows Server 2003 **R2** through the WS-Management initiative led by Microsoft, Intel, Sun Microsystems and Dell. This initiative allows to run any scripts remotely or to consume WMI data through a specific set of interfaces handling SOAP requests/responses. The advantage for the WMI provider developer is that when he exposes all his features through WMI, *Windows Remote Management*/WS-Management can in turn consume that information as well (embedded objects in WMI instances are not supported in Windows Server 2003 R2. It is however a target for Vista). All the layering to WS-Management and the mapping of the CIM data model to SOAP comes for free out of the WMI/WS-Management solution. In the event DCOM must be used, implementing DCOM requires the presence of a proxy DLL deployed on each client machine. As WMI is available in the Windows operating system since Windows 2000, these issues are eliminated.
5. *Support for Queries:* WMI offers support for WQL[6] queries out of the box. This means that if a provider is not designed to support queries, WMI supports it by using an enumeration technique out of the provider.
6. *Eventing capabilities:* WMI offers the capability to notify a subscriber for any event it is interested in. WMI uses the WMI Query Language (WQL) to submit WQL event queries and defines the type of events to be returned. The eventing mechanism, with all related callbacks, is part of the WMI COM/DCOM and automation interfaces. Anyone writing a WMI provider can have the benefit of this functionality at no cost for his customers. It will be up to the consumer to decide how it wants to consume the management information exposed by the WMI provider and its related set of WMI classes.
7. *Code template generator:* To speed up the process of writing a WMI provider including all COM/DCOM interfaces and related definitions, the WMI team developed the *WMI ATL Wizard* to generate the code template implementing a provider. The code generated is based on the WMI class model initially designed by the developer. The WMI provider developer will be able to interface the pre-defined COM/DCOM interfaces for the WMI provider with its set of native APIs retrieving the management information to expose. The exercise consists in filling the “gaps” in the provider code to create the desired interfacing logic.
8. *Predictability:* Predictability is an important concern for IT professionals because it defines the capability of someone having an experience with a set of interfaces managing a Windows component to apply this knowledge right away, intuitively, to any other manageable Windows component without having relearn everything from ground up. Predictability for a customer is a real gain as it increases the Return of Investment (ROI). A person facing such a situation simply expects things to work the same way based on his previous experience. The constant increase of COM programming/scriptable interfaces has a huge impact on the predictability, as this makes it difficult for customers to automate, manage Windows and leverage their existing knowledge. WMI with CIM address this problem by always exposing the same programming object model (COM/DCOM, Automation, .NET) whatever the manageable entity is.
9. *Protect existing customer investments:* Protecting customers and partners investment motivates customers to invest in technologies. As Microsoft did invest a lot these past years in writing WMI providers, customers and

partners invested in tools leveraging the WMI capabilities of Windows. Therefore, they naturally continue to exploit these capabilities instead of having to use a new set of specific interfaces for each Windows manageable component. A specific set of interfaces means having a specific set of agents or in-house developed software based on a new model or set of interfaces especially dedicated to a component or technology. By leveraging the capabilities of WMI today, customers and partners can leverage the work investment made in the past while minimizing their costs in developments, learning curves and new discoveries. This will also have a great impact on the stability and reliability of their infrastructure as they continue to leverage an existing implementation with an improved technology.

10. *Provide a logical and unified administration model:* As briefly described before in the introduction, this model is based on an industry standard called CIM defined by the DMTF (http://www.dmtf.org).The CIM class-based schema is defined by a consortium of constructors and software developers that meets the requirements of the industry. This implies that not only Microsoft leverages the WMI capabilities, but also any other third party constructors or developers write their own code to fit into the model. For instance, Intel is doing this for some their network driver adapters and software. HP is leveraging existing WMI providers and implementing their own WMI providers in their HP Open View Enterprise Management software. IBM consumes WMI from the Tivoli management suite, MOM and SMS are also consuming and providing WMI information. Lastly, Windows XP SP2 leverages WMI to get information status from anti-virus software and firewalls.

WMI tools

Some WMI tools can also be useful during the design and development phases. These tools are:

- *The MOF compiler (MOFComp.exe):* The Managed Object Format (MOF) compiler parses a file containing Managed Object Format statements and adds the classes and class instances defined in the file to the CIM repository. The MOF format is a specific syntax to define CIM class representation in an ASCII file (e.g. MIB are to SNMP what MOF files are to CIM). MOFComp.exe is included in every WMI installation. Every definition existing in the CIM repository is initially defined in an MOF file. MOF files are located in %SystemRoot%\System32\WBEM. During the WMI setup, they are loaded in the CIM repository.
- *The WMI Administrative Tools:* The WMI Administrative Tools are made of four tools: WMI CIM Studio, WMI Object Browser, WMI Event Registration and WMI Event Viewer. WMI Administrative Tools can be downloaded here [7]. The most important tool for a WMI provider developer is WMI CIM Studio as it helps in the initial WMI class creation in the CIM repository. It uses a web interface to display information and relies on a collection of ActiveX components installed on the system when it runs for the first time. WMI CIM Studio provides the ability to:
    - Connect to a chosen system and browse the CIM repository in any namespace available.
    - Search for classes by their name, by their descriptions or by property names.
    - Review the properties, methods and associations related to a given class.
    - See the instances available for a given class of the examined system.
    - Perform Queries in the WQL language.
    - Generate an MOF file based on selected classes.
    - Compile an MOF file to load it in the CIM repository.
- *WinMgmt.exe:* WinMgmt.exe is not a tool; it is the executable that implements the WMI Core service. Under the Windows NT family of operating systems, WMI runs as a service. On computers running Windows 98, Windows 95 or Windows Me, WMI runs as an application. Under the Windows NT family of operating systems, it is also possible to run this executable as an application, in which case, the executable runs in the current user context. For this, the WMI service must be stopped first. The executable supports some switches that can be useful when starting WMI as a service or as an application. WMI provider developers who may want to debug their providers essentially need to run the WMI service as an application.[8]

- *WBEMTest.exe:* WBEMTest.exe is a WMI tester tool, which is delivered with WMI. This tool allows an administrator or a developer to perform most of the tasks from a graphical interface that WMI provides at the API level. Although available under all Windows NT-based operating systems, this tool is not officially supported by Microsoft. WBEMTest provides the ability to:
  - Enumerate, open, create and delete classes.
  - Enumerate, open, create and delete instances of classes.
  - Select a namespace.
  - Perform data and event queries.
  - Execute methods associated to classes or instances.
  - Execute every WMI operation asynchronously, synchronously or semi-asynchronously.
- The WMI command line tool (WMIC): WMIC is a command-line tool designed to ease WMI information retrieval about a system by using some simple keywords (aliases). WMIC.exe is only available under Windows XP Professional, Windows Server 2003, Windows Vista, Windows 7 and Windows Server 2008. By typing "WMIC /?" from the command-line, a complete list of the switches and reserved keywords is available.
  - There is a Linux port of WMI command line tool, written in Python, based on Samba4 called 'wmi-client' [9]
- *WBEMDump.exe:* WBEMDump is a tool delivered with the Platform SDK. This command line tool comes with its own Visual C++ project. The tool can show the CIM repository classes, instances, or both. It is possible to retrieve the same information as that retrieved with WMIC. WBEMDump.exe requires more specific knowledge about WMI, as it doesn't abstract WMI as WMIC. However, it runs under Windows NT 4.0 and Windows 2000. It is also possible to execute methods exposed by classes or instances. Even if it is not a standard WMI tool delivered with the system installation, this tool can be quite useful for exploring the CIM repository and WMI features.

Wireless networking example

In the .NET framework, the ManagementClass class represents a Common Information Model (CIM) management class. A WMI class can be a Win32_LogicalDisk in the case of a disk drive, or a Win32_Process, such as a running program like Notepad.exe.

This example shows how "MSNdis_80211_ServiceSetIdentifier" WMI class is used to find the SSID of the Wi-Fi network that the system is currently connected to in the language C#:

```
ManagementClass mc = new ManagementClass("root\\WMI",
"MSNdis_80211_ServiceSetIdentifier", null);
ManagementObjectCollection moc = mc. GetInstances();

foreach (ManagementObject mo in moc)
{
   string wlanCard = (string)mo["InstanceName"];
   bool active;
   if (!bool.TryParse((string)mo["Active"], out active))
   {
      active = false;
   }
   byte[] ssid = (byte[])mo["Ndis80211SsId"];
}
```

The "MSNdis_80211_ServiceSetIdentifier" WMI class is only supported on Windows XP and Windows Server 2003.

WMI driver extensions

The WMI extensions to WDM provide kernel-level instrumentation such as publishing information, configuring device settings, supplying event notification from device drivers and allowing administrators to set data security through a WMI provider known as the *WDM provider*. The extensions are part of the WDM architecture; however, they have broad utility and can be used with other types of drivers as well (such as SCSI and NDIS). The WMI Driver Extensions service monitors all drivers and event trace providers that are configured to publish WMI or event trace information. Instrumented hardware data is provided by way of drivers instrumented for WMI extensions for WDM. WMI extensions for WDM provide a set of Windows device driver interfaces for instrumenting data within the driver models native to Windows, so OEMs and IHVs can easily extend the instrumented data set and add value to a hardware/software solution. The WMI Driver Extensions, however, are not supported by Windows Vista and later operating systems.[10]

References

[1] WMI Redistributable for Windows NT (http://www.microsoft.com/downloads/details.aspx?familyid=C174CFB1-EF67-471D-9277-4C2B1014A31E&displaylang=en)
[2] WMI Redistributable for Windows 95 and Windows 98 (http://www.microsoft.com/downloads/details.aspx?familyid=98a4c5ba-337b-4e92-8c18-a63847760ea5&displaylang=en)
[3] Description of WMIC (http://support.microsoft.com/kb/290216)
[4] Windows Vista Client Manageability (http://download.microsoft.com/download/b/e/3/be37cbce-425e-45c2-a9f5-378026b5be81/04-d-WinMgmtTech-v03-TOURB-FINAL.ppt)
[5] System.Management namespace (http://msdn2.microsoft.com/en-us/library/system.management.aspx)
[6] WMI Query Language via PowerShell (http://www.ravichaganti.com/blog/?page_id=2134)
[7] http://www.microsoft.com/downloads/details.aspx?FamilyID=6430f853-1120-48db-8cc5-f2abdc3ed314&DisplayLang=en
[8] WMI Tasks: Computer Software (http://msdn2.microsoft.com/en-us/library/aa394588.aspx)
[9] http://packages.ubuntu.com/hardy/python/wmi-client
[10] Deprecated components in Windows Vista and Windows Server 2008 (http://msdn2.microsoft.com/en-us/library/aa480152.aspx#appcomp_topic16)

External links

- Microsoft
  - WMI at the Windows Hardware Developer Central (http://www.microsoft.com/whdc/system/pnppwr/wmi/default.mspx)
  - CIM terminology (http://msdn2.microsoft.com/en-us/library/ms811530.aspx)
  - WMI Overview and Background (http://msdn2.microsoft.com/en-us/library/ms811553.aspx)
  - WMI and CIM overview (http://msdn2.microsoft.com/en-us/library/ms811552.aspx)
  - How improved support for WMI makes PowerShell the best environment to use and script WMI (http://blogs.msdn.com/powershell/archive/2006/06/26/647038.aspx)
  - Microsoft WMI Webcast (http://msevents.microsoft.com/CUI/WebCastEventDetails.aspx?culture=en-US&EventID=1032290321&CountryCode=US)
- Others
  - Use Java to work with WMI from any platform (http://www.j-interop.org)
  - Working with WMI providers to PowerShell (http://searchwincomputing.techtarget.com/generic/0,295582,sid68_gci1250536,00.html)
  - Working with WMI using Python (includes cookbook examples) (http://timgolden.me.uk/python/wmi.html)
  - Working with WMI using VC++ (Examples) (http://www.vedivi.com/support/blog/71-how-to-enable-remote-desktop-programmatically.html)
  - (http://srimedia.wetpaint.com)

- WMI (and SMART) Components for Delphi (http://www.magsys.co.uk/delphi/magwmi.asp)
- WMIC Snippets - many example queries with output shown (http://quux.wiki.zoho.com/WMIC-Snippets.html)
- WMIC - A Hidden Gem: simple introduction to a great tool (http://codeslammer.wordpress.com/2008/02/21/wmic-a-hidden-gem/)
- WMI Client for Linux (http://packages.ubuntu.com/wmi-client)
- List of WMI Win32 classes divided in sections with source examples in different scripting languages (http://activexperts.com/admin/wmi/)

Daemon (computing)

In Unix and other multitasking computer operating systems, a **daemon** (/ˈdeɪmən/ or /ˈdiːmən/)[1] is a computer program that runs as a background process, rather than being under the direct control of an interactive user. Typically daemon names end with the letter *d*: for example, `syslogd` is the daemon that implements the system logging facility and `sshd` is a daemon that services incoming SSH connections.

In a Unix environment, the parent process of a daemon is often, but not always, the init process. A daemon is usually created by a process forking a child process and then immediately exiting, thus causing init to adopt the child process. In addition, a daemon or the operating system typically must perform other operations, such as dissociating the process from any controlling terminal (tty). Such procedures are often implemented in various convenience routines such as *daemon(3)* in Unix.

Systems often start daemons at boot time: they often serve the function of responding to network requests, hardware activity, or other programs by performing some task. Daemons can also configure hardware (like udevd on some GNU/Linux systems), run scheduled tasks (like cron), and perform a variety of other tasks.

Terminology

The term was coined by the programmers of MIT's Project MAC. They took the name from Maxwell's demon, an imaginary being from a famous thought experiment that constantly works in the background, sorting molecules.[2] Unix systems inherited this terminology. Daemons are also characters in Greek mythology, some of whom handled tasks that the gods could not be bothered with. BSD and some of its derivatives have adopted a daemon as its mascot, although this mascot is a variation of the demons that appear in Christian artwork.

The word *daemon* is an alternative spelling of *demon,*[3] and is pronounced /ˈdiːmən/ ***DEE**-mən*. In the context of computer software, the original pronunciation /ˈdiːmən/ has drifted to /ˈdeɪmən/ ***DAY**-mən* for some speakers.[1]

Alternate terms for *daemon* are *service* (Microsoft Windows NT), *subsystem* (IBM z/OS), *server virtual machine* (IBM VM), *ghost job* (XDS UTS).

Creation

In a strictly technical sense, a Unix-like system process is a daemon when its parent process terminates and the daemon is assigned the init process (process number 1) as its parent process and has no controlling terminal. However, more commonly, a daemon may be any background process, whether a child of `init` or not.

The common method for a process to become a daemon involves:

- Dissociating from the controlling tty
- Becoming a session leader
- Becoming a process group leader
- Executing as a background task by forking and exiting (once or twice). This is required sometimes for the process to become a session leader. It also allows the parent process to continue its normal execution.
- Setting the root directory ("/") as the current working directory so that the process does not keep any directory in use that may be on a mounted file system (allowing it to be unmounted).
- Changing the umask to 0 to allow open(), creat(), et al. operating system calls to provide their own permission masks and not to depend on the umask of the caller
- Closing all inherited files at the time of execution that are left open by the parent process, including file descriptors 0, 1 and 2 (stdin, stdout, stderr). Required files will be opened later.
- Using a logfile, the console, or /dev/null as stdin, stdout, and stderr

Implementation in MS-DOS and Microsoft Windows

In the Microsoft DOS environment, daemon-like programs were implemented as Terminate and Stay Resident (TSR) software. On Microsoft Windows NT systems, programs called Windows services perform the functions of daemons. They run as processes, usually do not interact with the monitor, keyboard, and mouse, and may be launched by the operating system at boot time. In Windows 2000 and later versions, Windows services are configured and manually started and stopped using the Control Panel, a dedicated control/configuration program, the Service Controller component of the Service Control Manager (`sc` command), or the `net start` and `net stop` commands.

However, any Windows application can perform the role of a daemon, not just a service, and some daemons for Windows have the option of running as a normal process.

Implementation in Mac OS

On the original Mac OS, optional features and services were provided by files loaded at startup time that patched the operating system; these were known as system extensions and control panels. Later versions of classic Mac OS augmented these with fully fledged faceless background applications: regular applications that ran in the background. To the user, these were still described as regular system extensions.

Mac OS X, which is a Unix-like system, has daemons. There is a category of software called *services* as well, but these are different in concept from Windows services.

Etymology

According to Fernando J. Corbato who worked on Project MAC in 1963 his team is the first to use the term daemon. The use of the term daemon was inspired by Maxwell's demon, an imaginary agent in physics and thermodynamics that helped to sort molecules.[4]

> We fancifully began to use the word daemon to describe background processes which worked tirelessly to perform system chores.

In the general sense, daemon is an older form of the word demon, from the Greek δαίμων. In the Unix System Administration Handbook, Evi Nemeth states the following about daemons:[5]

> Many people equate the word "daemon" with the word "demon", implying some kind of satanic connection between UNIX and the underworld. This is an egregious misunderstanding. "Daemon" is actually a much older form of "demon"; daemons have no particular bias towards good or evil, but rather serve to help define a person's character or personality. The ancient Greeks' concept of a "personal daemon" was similar to the modern concept of a "guardian angel"—*eudaemonia* is the state of being helped or protected by a kindly spirit. As a rule, UNIX systems seem to be infested with both daemons and demons. (p.403)

A further characterization of the mythological symbolism is that a daemon is something which is not visible yet is always present and working its will. Plato's Socrates describes his own personal daemon to be something like the modern concept of a moral conscience:

> "The favour of the gods has given me a marvelous gift, which has never left me since my childhood. It is a voice which, when it makes itself heard, deters me from what I am about to do and never urges me on."
>
> —Character of Socrates in "Theages", Plato[6]

Notable service daemons in Unix-like systems

- amd: Berkeley Automounter
- anacron: Executed delayed cron tasks at boot time
- apmd: Advanced Power Management Daemon
- arpwatch: watches for Ethernet IP address pairings that are resolved using the ARP protocol
- atd: Runs jobs queued using the at tool
- biod: Cooperates with a remote nfsd to handle client Network File System requests
- blued: Mac OS X Bluetooth daemon
- bootparmd: Internet Bootstrap Protocol server daemon
- configd: A daemon that maintains dynamic configuration information about the computer and its environment (mainly the network)
- crond: Task scheduler daemon
- cupsd: CUPS printer daemon
- devfsd: Device file system (devfs)
- dhcpd: Dynamic Host Configuration Protocol and Internet Bootstrap Protocol server
- fetchmail: daemon to retrieve mail from servers at regular intervals
- fingerd: Finger protocol server
- ftpd: File Transfer Protocol (FTP) server
- gated: routing daemon that handles multiple routing protocols and replaces routed and egpup
- gpm: General Purpose Mouse Daemon
- httpd: Hypertext Transfer Protocol (HTTP) daemon (web server)
- identd: Provides the identity of a user of a particular TCP connection
- inetd: Internet Superserver Daemon

- init: Initial process that spawns other processes
- imapd: Internet Message Access Protocol (IMAP) server daemon
- innd: InterNetNews server daemon
- ipchains: A deprecated packet forwarding and firewall daemon
- isdn: Integrated Services Digital Network (ISDN) interfacing server daemon
- kerneld: Automatically loads and unloads kernel modules
- klogd: Kernel log daemon (intercepts and logs Linux kernel messages)[7]
- kswapd: Kernel page swapping daemon
- kudzu: Detects and configures new or changed hardware during boot
- launchd: Init and operating system service management implementation on Mac OS X
- lpd: Line Printer Daemon protocol
- memcached: In-memory distributed memory caching daemon
- mpd: Music Player Daemon
- mountd: Part of typical Network File System implementation
- mysql: Database server daemon
- named: A Domain Name System (DNS) server daemon
- nfsd: Network File System (NFS) daemon
- nmbd: Network Message Block Daemon; part of Samba
- ntpd: Network Time Protocol (NTP) service daemon
- pcmcia: Provides generic PCMCIA services
- portmap
- postfix: A mail transfer agent used as a replacement for sendmail
- postgresql: Database server daemon
- retchmail: A mail retrieval agent
- rlprd: Remote line printer proxy daemon
- routed: Manages routing tables
- rpcbind: Remote Procedure Call Bind Daemon
- rquotad: Remote quota reporting, associated with Network File System
- rwalld: Allows users to write messages on remote terminals using rwall
- rwhod: Maintains the database used by the rwho and ruptime tools
- sched, swapper: Copies process regions to swap space to reclaim physical pages of memory
- sendmail: A mail transfer agent daemon
- smbd: Samba, an Server Message Block (SMB) daemon
- smtpd: Simple Mail Transfer Protocol Daemon
- snmpd: Simple Network Management Protocol Daemon
- sound: A sound server daemon
- squid: A web page caching proxy server daemon
- sshd: Secure Shell Server Daemon
- syncd: Keeps the file systems synchronized with system memory
- syslogd: Syslog daemon
- tcpd: Service wrapper restricts access to inetd based services through hosts.allow and hosts.deny
- telnetd: Telnet server daemon
- vhand: The "page stealing daemon" releases pages of memory for use by other processes
- vsftpd: "Very Secure FTP Daemon"
- webmin: Web based administration server daemon
- xfsd: X font server daemon
- xinetd: Enhanced Internet Superserver Daemon

- ypbind: A bind server for Network Information Service ("Yellow Pages")

References

[1] Eric S. Raymond. "daemon" (http://catb.org/~esr/jargon/html/D/daemon.html). *The Jargon File*. . Retrieved 2008-10-22.
[2] Fernando J. Corbató (2002-01-23). "Take Our Word for It" (http://www.takeourword.com/TOW146/page4.html). . Retrieved 2006-08-20.
[3] "Merriam-Webster definition of daemon" (http://www.merriam-webster.com/dictionary/daemon). *Merriam-Webster Online*. . Retrieved 2009-08-05.
[4] "The Origin of the word Daemon" (http://ei.cs.vt.edu/~history/Daemon.html). .
[5] "The BSD Daemon" (http://www.freebsd.org/copyright/daemon.html). Freebsd.org. . Retrieved 2008-11-15.
[6] There is some doubt as to whether "Theages" was genuinely authored by Plato but this passage accurately articulates the Greek conception of a daemon.
[7] http://unixhelp.ed.ac.uk/CGI/man-cgi?klogd+8

External links

- start-stop-daemon man page (http://manpages.ubuntu.com/manpages/lucid/en/man8/start-stop-daemon.8.html)
- daemon man page (http://manpages.ubuntu.com/manpages/lucid/en/man1/daemon.1.html)
- Tutorial: Unix Daemons in Perl (http://www.webreference.com/perl/tutorial/9/)
- How to daemonize in Linux (http://www-theorie.physik.unizh.ch/~dpotter/howto/daemonize)
- Unix Daemon Server Programming (http://www.enderunix.org/docs/eng/daemon.php)
- Linux Daemon Writing HOWTO (http://www.netzmafia.de/skripten/unix/linux-daemon-howto.html)

Runlevel

The term ***runlevel*** refers to a mode of operation in one of the computer operating systems that implement Unix System V-style initialization. Conventionally, seven runlevels exist, numbered from zero to six; though up to ten, from zero to nine, may be used. **S** is sometimes used as a synonym for one of the levels. Only one "runlevel" is executed on bootup - run levels are ***not*** executed sequentially, i.e. either runlevel 2 OR 3 OR 4 is executed, ***not*** 2 then 3 then 4.

"Runlevel" defines the state of the machine after boot. Different runlevels are typically assigned to:

- single-user mode
- multi-user mode without network services started
- multi-user mode with network services started
- system shutdown
- system reboot

The exact setup of these configurations will vary from OS to OS, and from one Linux distribution to another. For example, runlevel 4 might be multi-user, GUI, no-server on one distribution, and nothing on another. Note the difference in the Red Hat and Slackware distributions charted in this article. However, "runlevels" do commonly follow patterns described in this article. When installing Linux, it is best to consult that distribution's available user guides.

In standard practice, when a computer enters runlevel zero, it halts, and when it enters runlevel six, it reboots. The intermediate runlevels (1-5) differ in terms of which drives are mounted, and which network services are started. Default runlevels are typically 3, 4, or 5. Lower run levels are useful for maintenance or emergency repairs, since they usually don't offer any network services at all. The particular details of runlevel configuration differ widely among operating systems, and also among system administrators.

The runlevel system replaced the traditional `/etc/rc` script used in Version 7 Unix.

Standard runlevels

Standard runlevels

ID	Name	Description
0	Halt	Shuts down the system.
S	Single-User Mode	Does not configure network interfaces or start daemons.[runlevel_S_as_1]
6	Reboot	Reboots the system.

= Almost all systems use runlevel 1 for this purpose. This mode is intended to provide a safe environment to perform system maintenance. Originally this runlevel provided a single terminal (console) interface running a root login shell. The increasing trend towards physical access to the computer during the boot process has led to changes in this area.

Linux

The Linux operating system can make use of runlevels through the programs of the sysvinit project. After the Linux kernel has booted, the `init` program reads the `/etc/inittab` file to determine the behavior for each runlevel. Unless the user specifies another value as a kernel boot parameter, the system will attempt to enter (start) the default runlevel.

Linux Standard Base specification

Conforming implementations are not required to provide these exact run levels or give them the meanings described here, and may map any level described here to a different level which provides the equivalent functionality.[1]

LSB 4.1.0

ID	Name	Description
0	Halt	Shuts down the system.
1	Single-User Mode	Mode for administrative tasks.[2][behavour_of_runlevel_1]
2	Multi-User Mode	Does not configure network interfaces and does not export networks services.[behavour_of_runlevel_23]
3	Multi-User Mode with Networking	Starts the system normally.[behavior_of_runlevel_23]
4	Not used/User-definable	For special purposes.
5	Start the system normally with appropriate display manager. (with GUI)	As runlevel 3 + display manager.
6	Reboot	Reboots the system.

= The additional behavior of this runlevel varies greatly. All distributions provide at least one virtual terminal. Some distributions start a login shell as the superuser; some require correctly entering the superuser's password first; others provide a login prompt, allowing any user access.

= In some cases, runlevels 2 and 3 function identically; offering a Multi-User Mode with Networking.

Debian Linux

Debian, as well as most of the distributions based on it, like early Ubuntu, does not make any distinction between runlevels 2 to 5.

Debian Linux runlevels

ID	Description
S	Only run on boot (replaces `/etc/rc.boot`)
0	Halt
1	Single-User mode
2-5	Full Multi-User with console logins and display manager if installed
6	Reboot

Ubuntu

Ubuntu 6.10 (Edgy Eft) and later contain Upstart as a replacement for the traditional init-process, but they still use the traditional init scripts and Upstart's SysV-rc compatibility tools to start most services and emulate runlevels.

Red Hat Linux and Fedora

Red Hat as well as most of its derivatives (such as CentOS) uses runlevels like this [3] :

Red Hat Linux/Fedora runlevels

Code	Information
0	Halt
1	Single-User mode
2	Multi-user mode console logins only (without networking)
3	Multi-User mode, console logins only
4	Not used/User-definable
5	Multi-User mode, with display manager as well as console logins (X11)
6	Reboot

SUSE Linux

SUSE uses a similar setup to Red Hat:

SUSE Linux runlevels

ID	Description
0	Halt
1 or **S**	Single-User mode
2	Multi-User mode without networking
3	Multi-User mode, console logins only
4	Not used/User-definable
5	Multi-User mode with display manager
6	Reboot

The services that run under a specific runlevel can be modified with YaST | System Services (runlevel), **insserv** command or with **chkconfig** command like the Red Hat based distributions.

Slackware Linux

Slackware Linux uses runlevel 1 for maintenance, as on other Linux distributions; runlevels 2, 3 and 5 identically configured for a console (with all services active); and runlevel 4 adds the X Window System.

Slackware Linux runlevels

ID	Description
0	Halt
1	Single-User mode
2	Unused but configured the same as runlevel 3
3	Multi-User mode without display manager
4	Multi-User mode with display manager
5	Unused but configured the same as runlevel 3
6	Reboot

Arch Linux

Arch Linux runlevels[4]

ID	Description
0	Halt
1	Single-User (Maintenance Mode)
2	Not used
3	Multi-User
4	Not used
5	Multi-User with X11
6	Reboot

Gentoo Linux

Gentoo Linux runlevels

ID	Description
0	Halt
1 or **S**	Single-User mode
2	Multi-User mode without networking
3	Multi-User mode
4	Aliased for runlevel **3**
5	Aliased for runlevel **3**
6	Reboot

Unix

System V Releases 3 and 4

System V runlevels

ID	Description
0	Shut down system, power-off if hardware supports it (only available from the console)
1	Single-User mode, all filesystems unmounted but root, all processes except console processes killed
2	Multi-User mode
3	Multi-User mode with RFS (and NFS in Release 4) filesystems exported
4	Multi-User, User-definable
5	Halt the operating system, go to firmware
6	Reboot
s, **S**	Identical to **1**, except current terminal acts as the system console

Solaris[5]

Solaris runlevels

ID	Description
0	Operating system halted; (SPARC only) drop to OpenBoot prompt
S	Single-User mode with only root filesystem mounted (as read-only)
1	Single-User mode with all local filesystems mounted (read-write)
2	Multi-User mode with most daemons started
3	Multi-User mode; identical to **2** (runlevel **3** runs both `/sbin/rc2` and `/sbin/rc3`), with filesystems exported, plus some other network services started.
4	Alternative Multi-User mode, User-definable
5	Shut down, power-off if hardware supports it

6	Reboot

HP-UX

HP-UX runlevels

ID	Description
0	System halted
S	Single-User mode, booted to system console only, with only root filesystem mounted (as read-only)
s	Single-User mode, identical to **S** except the current terminal acts as the system console
1	Single-User mode with local filesystems mounted (read-write)
2	Multi-User mode with most daemons started and Common Desktop Environment launched
3	Identical to runlevel **2** with NFS exported
4	Multi-User mode with VUE started instead of CDE
5, **6**	Not used/User-definable

FreeBSD, OpenBSD and NetBSD

The BSD variants don't use the concept of run levels, although on some versions init(8) provides an emulation of some of the common run levels.

AIX

AIX does not follow the System V R4 (SVR4) run level specification, with run levels from 0 to 9 available, as well as from a to c. 0 and 1 are reserved, 2 is the default normal multi-user mode and run levels from 3 to 9 are free to be defined by the administrator. Run levels from a to c allow the execution of processes in that run level without killing processes started in another.

AIX runlevels

ID	Name	Description
0		reserved
1		reserved
2	Normal Multi-User mode	default mode

External links

- Runlevel Definition [6] - by The Linux Information Project (LINFO)
- What are run levels? [7] - LinuxQuestions.org

References

[1] "Chapter 20. System Initialization 20.5. Run Levels" (http://refspecs.freestandards.org/LSB_4.1.0/LSB-Core-generic/LSB-Core-generic/runlevels.html). *Linux Standard Base Core Specification 4.1*. 2011. . Retrieved 2011-04-21.

[2] "Chapter 15. Commands and Utilities 15.2. Command Behavior" (http://refspecs.freestandards.org/LSB_4.1.0/LSB-Core-generic/LSB-Core-generic/shutdown.html). *Linux Standard Base Core Specification 4.1*. 2011. . Retrieved 2011-04-21.

[3] (http://docs.redhat.com/docs/en-US/Red_Hat_Enterprise_Linux/3/html/Reference_Guide/s1-boot-init-shutdown-sysv.html) Redhat runlevels

[4] "Arch Linux Runlevels" (https://wiki.archlinux.org/index.php/Runlevels). . Retrieved 2010-12-31.

[5] "Solaris Boot Process" (http://blogs.sun.com/boot/entry/amrita_sadhukhan). Amrita Sadhukhan (Sun). . Retrieved 2009-11-04.

[6] http://www.linfo.org/runlevel_def.html

[7] http://wiki.linuxquestions.org/wiki/Run_Levels

Windows service

On Microsoft Windows operating systems, a **Windows service** is a long-running executable that performs specific functions and it is designed not to require user intervention. Windows services can be configured to start when the operating system is booted and run in the background as long as Windows is running, or they can be started manually when required. They are similar in concept to a Unix daemon. Many appear in the processes list in the Windows Task Manager, most often with a username of `SYSTEM`, `LOCAL SERVICE` or `NETWORK SERVICE`, though not all processes with the `SYSTEM` username are services. The remaining services run through svchost.exe as DLLs loaded into memory.

Managing services

Once a service is installed, it can be managed by launching "Services" from the Windows Control Panel → *Administrative Tools* or typing "Services.msc" in the *Run* command on Start menu. In Windows Vista and later, it can be started or stopped in the Services tab in Windows Task Manager, where its process can also be found. The "Services" management console provides a brief description of the service functions and displays the path to the service executable, its current status, startup type, dependencies and the account under which the service is running. It enables users to:

- Start, stop, pause or restart services.
- Specify service parameters.
- Change the startup type which includes *Automatic*, *Manual* and *Disabled*:
  - *Automatic* starts the services at system logon,
  - *Manual* starts a service as required or when called from an application (according to definition, but only some of the time in practice, depending on the service),
  - *Disabled* completely disables the service and prevents it and its dependencies from running.
  - *Automatic (Delayed)* is a new startup type introduced in Windows Vista, that starts the service a short while after the system has finished its booting and initial busy operations, so that the system boots up faster.
- Change the account under which the service logs on.
- Configure recovery options upon service failure.
- Export the list of services as a text file or a CSV file.

In Windows XP Service Pack 3 and Windows Vista, besides the *Services* management console, users can manipulate services using MSConfig. The use of MSConfig to manage services, however, causes a prompt on the next startup. MSConfig can hide all operating system services for troubleshooting. Users can also use the *SC* command in the command prompt.

Developing a Windows service

A *Windows Service* is created using development tools such as Microsoft Visual Studio or Embarcadero Delphi. Windows provides an interface called the Service Control Manager that manages the starting and stopping of services. An application that wants to be a service needs to first be written in such a way that it can handle start, stop, and pause messages from the Service Control Manager. Then, in one or more API calls, the name of the service and other attributes such as its description are registered with the Service Control Manager. Although typically services do not have a user interface, developers can add forms and other UI components. In this case, the "Allow service to interact with desktop" should be checked on the *Logon* tab in the *Service properties* dialog (though care should be taken with this approach as this can cause a security risk since any logged in user would be able to interact with the service).

External links

- Microsoft Developer Network - Services [1]
- INSTSRV.EXE [2] - a Windows NT resource kit program that allows you to install an arbitrary application as a service
- Registry values [3] for Start, Type & ErrorControl

References

[1] http://msdn2.microsoft.com/en-us/library/ms685141.aspx
[2] http://support.microsoft.com/default.aspx?scid=kb;en-us;137890
[3] http://windowsitpro.com/article/articleid/15169/what-are-the-errorcontrol-start-and-type-values-under-the-----services-subkeys.html

Syslog

Syslog is a standard for computer data logging. It separates the software that generates messages from the system that stores them and the software that reports and analyzes them.

Syslog can be used for computer system management and security auditing as well as generalized informational, analysis, and debugging messages. It is supported by a wide variety of devices (like printers and routers) and receivers across multiple platforms. Because of this, syslog can be used to integrate log data from many different types of systems into a central repository.

Messages refer to a facility (`auth, authpriv, daemon, cron, ftp, lpr, kern, mail, news, syslog, user, uucp, local0, ... , local7`) and are assigned a severity (`Emergency, Alert, Critical, Error, Warning, Notice, Info` or `Debug`) by the sender of the message.

Configuration allows directing messages to various local devices (console), files (/var/log/) or remote syslog daemons. Care must be taken when updating the configuration as omitting or misdirecting message facilities or severities can cause important messages to be ignored by syslog or overlooked by the administrator.

logger is a command line utility that can send messages to the syslog.

Some implementations permit the filtering and display of syslog messages.

Syslog is now standardized within the Syslog working group of the IETF.

History

Syslog was developed in the 1980s by Eric Allman as part of the Sendmail project, and was initially used solely for Sendmail. It proved so valuable that other applications began using it as well. Syslog has since become the standard logging solution on Unix and Unix-like systems; there have also been a variety of syslog implementations on other operating systems and is commonly found in network devices such as routers.

Syslog functioned as a *de facto* standard, without any authoritative published specification, and many implementations existed, some of which were incompatible. The Internet Engineering Task Force documented the status quo in RFC 3164. Since then, additions to syslog have been worked on. RFC 3164 was made obsolete by RFC 5424[1]

At different points in time, various companies have attempted patent claims on syslog.[2][3] This had little effect on the use and standardization of the protocol.

Outlook

Various groups are working on draft standards detailing the use of syslog for more than just network and security event logging, such as its proposed application within the health care environment.

Regulations, such as SOX, PCI DSS, HIPAA, and many others are requiring organizations to implement comprehensive security measures, which often include collecting and analyzing logs from many different sources. Syslog has proven to be an effective format to consolidate logs, as there are many open source and proprietary tools for reporting and analysis. Converters exist from Windows Event Log as well as other log formats to syslog.

An emerging area of managed security services is the collection and analysis of syslog records for organizations. Companies calling themselves Managed Security Service Providers attempt to apply analytics techniques (and sometimes artificial intelligence algorithms) to detect patterns and alert customers of problems.

Facility Levels

The list of Facilities available:[4]

Facility Number	Facility Description
0	kernel messages
1	user-level messages
2	mail system
3	system daemons
4	security/authorization messages
5	messages generated internally by syslogd
6	line printer subsystem
7	network news subsystem
8	UUCP subsystem
9	clock daemon
10	security/authorization messages
11	FTP daemon
12	NTP subsystem
13	log audit
14	log alert
15	clock daemon
16	local use 0 (local0)
17	local use 1 (local1)
18	local use 2 (local2)
19	local use 3 (local3)
20	local use 4 (local4)
21	local use 5 (local5)
22	local use 6 (local6)
23	local use 7 (local7)

Severity levels

RFC 5424 [5] defines eight severity levels:

Code	Severity	Description	General Description
0	Emergency	System is unusable.	A "panic" condition usually affecting multiple apps/servers/sites. At this level it would usually notify all tech staff on call.
1	Alert	Action must be taken immediately.	Should be corrected immediately, therefore notify staff who can fix the problem. An example would be the loss of a backup ISP connection.
2	Critical	Critical conditions.	Should be corrected immediately, but indicates failure in a primary system, an example is a loss of primary ISP connection.
3	Error	Error conditions.	Non-urgent failures, these should be relayed to developers or admins; each item must be resolved within a given time.
4	Warning	Warning conditions.	Warning messages, not an error, but indication that an error will occur if action is not taken, e.g. file system 85% full - each item must be resolved within a given time.
5	Notice	Normal but significant condition.	Events that are unusual but not error conditions - might be summarized in an email to developers or admins to spot potential problems - no immediate action required.
6	Informational	Informational messages.	Normal operational messages - may be harvested for reporting, measuring throughput, etc. - no action required.
7	Debug	Debug-level messages.	Info useful to developers for debugging the application, not useful during operations.

Internet standards

The Domain Name System is defined by Request for Comments (RFC) documents published by the Internet Engineering Task Force (Internet standards). The following is a list of RFCs that define the Syslog protocol:[6]

- RFC 3195 *Reliable Delivery for syslog*
- RFC 5424 *The Syslog Protocol*
- RFC 5425 *TLS Transport Mapping for Syslog*
- RFC 5426 *Transmission of Syslog Messages over UDP*
- RFC 5427 *Textual Conventions for Syslog Management*
- RFC 5848 *Signed Syslog Messages*
- RFC 6012 *Datagram Transport Layer Security (DTLS) Transport Mapping for Syslog*

External links

- IETF syslog working group [7]
- SANS Paper [8] The Ins and Outs of System Logging Using Syslog
- NIST SP 800-92 Guide to Computer Security Log Management (PDF) [9]
- NetLogger [10] methodology and tools for debugging and analysis of complex distributed applications

References

[1] R. Gerhards, The Syslog Protocol, RFC 5424 (http://tools.ietf.org/html/rfc5424)
[2] "LXer: Patent jeopardizes IETF syslog standard" (http://lxer.com/module/newswire/view/64026/index.html). .
[3] "IETF IPR disclosure on HUAWEI's patent claims" (http://datatracker.ietf.org/public/ipr_detail_show.cgi?ipr_id=724). .
[4] "Syslog Facilities" (http://www.kiwisyslog.com/help/syslog/protocol_facilities.htm). . Retrieved 15 February 2012.
[5] http://tools.ietf.org/html/rfc5424
[6] "Security Issues in Network Event Logging (syslog)" (http://datatracker.ietf.org/wg/syslog/). IETF. .
[7] http://www.ietf.org/html.charters/syslog-charter.html
[8] http://www.sans.org/reading_room/whitepapers/logging/ins-outs-system-logging-syslog_1168
[9] http://csrc.nist.gov/publications/nistpubs/800-92/SP800-92.pdf
[10] https://sites.google.com/a/lbl.gov/netlogger/

Event Viewer

Event Viewer A component of Microsoft Windows	
Event Viewer in Windows Vista	
Details	
Type	Utility software
Included with	Windows NT and all its successors
Service name	Windows Event log (`eventlog`)
Description	This service manages events and event logs. It supports logging events, querying events, subscribing to events, archiving event logs, and managing event metadata.

Event Viewer in Windows XP

Event Viewer, a component of Microsoft's Windows NT line of operating systems, lets administrators and users view the event logs on a local or remote machine. In Windows Vista, Microsoft overhauled the event system.[1]

Overview

Windows NT has featured event logs since its release in 1993. Applications and operating system components can make use of this centralized log service to report events that have taken place, such as a failure to start a component or complete an action. The system defines three log sources: System, Application and Security.

Microsoft intends the System and Application log sources for use by the Windows operating system and Windows applications respectively. Only the Local Security Authority Subsystem Service (lsass.exe) can directly write to the Security log.

The Event Viewer uses event IDs to define the uniquely identifiable events that a Windows computer can encounter. For example, when a user's authentication fails, the system may generate Event ID 672.

Windows NT 4.0 added support for defining "event sources" (i.e. the application which created the event) and performing backups of logs.

Windows 2000 added the capability for applications to create their own log sources in addition to the three system-defined "System", "Application", and "Security" log files. Windows 2000 also replaced NT4's Event Viewer with a Microsoft Management Console (MMC) snap-in.

Windows Server 2003 added the `AuthzInstallSecurityEventSource()` API calls so that applications could register with the security event logs, and write security audit entries.[2]

Versions of Windows based on the Windows NT 6.0 kernel (Windows Vista and Windows Server 2008) no longer have a 300-megabyte limit to their total size. Prior to NT 6.0, the on-disk files were opened as memory-mapped files in kernel memory space, which used the same memory pools as other kernel components.

Windows XP (commandline)

Windows XP provides a set of three commandline tools, useful to task automation:

- eventquery.vbs - Official script to query, filter and output results based on the event logs. Discontinued after XP.
- eventcreate - a command (continued in Vista and Seven) to put custom events in the logs.
- eventtriggers - a command to create event driven tasks. Discontinued after XP, replaced by the "Attach task to this event" feature.

Windows Vista

Event Viewer consists of a rewritten event tracing and logging architecture on Windows Vista.[1] It has been rewritten around a structured XML log-format and a designated log type to allow applications to more precisely log events and to help make it easier for support technicians and developers to interpret the events. The XML representation of the event can be viewed on the *Details* tab in an event's properties. It is also possible to view all potential events, their structures, registered *event publishers* and their configuration using the *wevtutil* utility, even before the events are fired. There are a large number of different types of event logs including Administrative, Operational, Analytic, and Debug log types. Selecting the *Application Logs* node in the *Scope* pane reveals numerous new subcategorized event logs, including many labeled as diagnostic logs. Analytic and Debug events which are high frequency are directly saved into a trace file while Admin and Operational events are infrequent enough to allow additional processing without affecting system performance, so they are delivered to the Event Log service. Events are published asynchronously to reduce the performance impact on the *event publishing* application. Event attributes are also much more detailed and show EventID, Level, Task, Opcode, and Keywords properties.

Users can filter event logs by one or more criteria or by a limited XPath 1.0 expression, and custom views can be created for one or more events. Using XPath as the query language allows viewing logs related only to a certain subsystem or an issue with only a certain component, archiving select events and sending traces on the fly to support technicians.

Filtering using XPath 1.0

1. Open Windows Event Log
2. Expand out *Windows Logs*
3. Select the log file that is of interest to you (In the example below, we use the *Security* event log)
4. Right-click on the Event Log and select *Filter Current Log...*
5. Change the selected tab from *Filter* to *XML*
6. Check the box to *Edit query manually'*
7. Paste your query into the text box. You will find sample queries below.

Here are examples of simple custom filters for the new Window Event Log:

1. Select all events in the Security Event Log where the account name involved (TargetUserName) is "JUser"

   ```
   <QueryList><Query Id="0" Path="Security"><Select
   Path="Security">*[EventData[Data[@Name="TargetUserName"]="JUser"]]</Select></Query></
   ```

2. Select all events in the Security Event Log where any Data node of the EventData section is the string "JUser"

   ```
   <QueryList><Query Id="0" Path="Security"><Select
   Path="Security">*[EventData[Data="JUser"]]</Select></Query></QueryList>
   ```

3. Select all events in the Security Event Log where any Data node of the EventData section is "JUser" or "JDoe"

   ```
   <QueryList><Query Id="0" Path="Security"><Select
   Path="Security">*[EventData[Data="JUser" or
   Data="JDoe"]]</Select></Query></QueryList>
   ```

4. Select all events in the Security Event Log where any Data node of the EventData section is "JUser" and the Event ID is "4471"

```
<QueryList><Query Id="0" Path="Security"><Select
Path="Security">*[System[EventID="4771"]] and
*[EventData[Data="JUser"]]</Select></Query></QueryList>
```

5. Real world example for a package called Goldmine which has two @Names

```
<QueryList><Query Id="0" Path="Application"><Select
Path="Application">*[System[Provider[@Name='GoldMine' or
@Name='GMService']]]</Select></Query></QueryList>
```

Caveats:

- There are limitations [3] to Microsoft's implementation of Xpath[4]
- Queries using Xpath string functions will result in error[5]

Event subscribers

Major *event subscribers* include the Event Collector service and Task Scheduler 2.0. The Event Collector service can automatically forward event logs to other remote systems, running Windows Vista, Windows Server 2008 or Windows Server 2003 R2 on a configurable schedule. Event logs can also be remotely viewed from other computers or multiple event logs can be centrally logged and monitored agentlessly and managed from a single computer. Events can also be directly associated with tasks, which run in the redesigned Task Scheduler and trigger automated actions when particular events take place.

References

[1] New tools for Event Management in Windows Vista (http://www.microsoft.com/technet/technetmag/issues/2006/11/EventManagement/)

[2] "AuthzInstallSecurityEventSource Function" (http://msdn2.microsoft.com/en-us/library/Aa376314.aspx). . Retrieved 2007-10-05.

[3] http://msdn.microsoft.com/en-us/library/dd996910%28VS.85%29.aspx#limitations

[4] "Microsoft's Implementation and Limitations of XPath 1.0 in Windows Event Log" (http://msdn.microsoft.com/en-us/library/dd996910(VS.85).aspx#limitations). . Retrieved 2009-08-07.

[5] "filter events using an Xpath query" (http://www.open-a-socket.com/index.php/2009/04/). . Retrieved 2011-09-20.

External links

- Official sources:
  - Developer documentation for event logging (NT 3.1 through XP) (http://msdn2.microsoft.com/en-us/library/aa363652.aspx), (Windows Vista) (http://msdn2.microsoft.com/en-us/library/aa385780.aspx)
  - Windows 2000 Security Event Descriptions (Part 1 of 2) (http://support.microsoft.com/?kbid=299475), (Part 2 of 2) (http://support.microsoft.com/?kbid=301677)
  - Windows Server 2003 Security – Threats and Countermeasures – Chapter 6: Event Log (http://www.microsoft.com/technet/security/guidance/serversecurity/tcg/tcgch06n.mspx) from Microsoft TechNet
  - Events and Errors (http://technet.microsoft.com/en-us/library/cc754424.aspx) (Windows Server 2008) on Microsoft TechNet
- Other:
  - Windows Eventlog Viewer (http://www.tzworks.net/prototype_page.php?proto_id=4) Free tool that can be run on Windows, Linux or Mac OS-X
  - eventid.net (http://www.eventid.net/search.asp) – Contains several thousand Windows event log entries along with troubleshooting suggestions for each of them

- Trap to event log with eventcreate (http://www.loriotpro.com/ServiceAndSupport/How_to/Create_Windows_Event_On_Trap_Receipt_EN.php)

Hardening (computing)

In computing, **hardening** is usually the process of securing a system by reducing its surface of vulnerability. A system has a larger vulnerability surface the more that it does; in principle a single-function system is more secure than a multipurpose one. Reducing available vectors of attack typically includes the removal of unnecessary software, unnecessary usernames or logins and the disabling or removal of unnecessary services.

There are various methods of hardening Unix and Linux systems. This may involve, among other measures, applying a patch to the kernel such as Exec Shield or PaX; closing open network ports; and setting up intrusion-detection systems, firewalls and intrusion-prevention systems. There are also hardening scripts and tools like Bastille Linux, JASS [1] for Solaris systems and Apache/PHP Hardener [2] that can, for example, deactivate unneeded features in configuration files or perform various other protective measures.

External links

- IT Security Topic — Hardening [3] - www.colorado.edu
- Hardening Your Computing Assets [4]PDF - www.globalsecurity.org

References

[1] http://sun.com/software/security/jass/
[2] http://www.syhunt.com/hardener/
[3] http://www.colorado.edu/cns/security/awareness/hardening/
[4] http://www.globalsecurity.org/military/library/report/1997/harden.pdf

Microsoft Baseline Security Analyzer

Microsoft Baseline Security Analyzer

Screenshot of Microsoft Baseline Security Analyzer analysis result	
Developer(s)	Microsoft
Initial release	16 August 2004[1]
Stable release	2.2 / 10 August 2010[2]
Operating system	Windows Server 2008 R2, Windows 7, Windows Server 2008, Windows Vista, Windows Server 2003, Windows XP and Windows 2000[2]
Platform	IA-32 and x86-64[2]
Size	1.5 ~ 1.7 MB[2]
Translation available	English, German, French and Japanese[2]
Type	Computer security
License	Freeware
Website	www.microsoft.com/mbsa [3]

Microsoft Baseline Security Analyzer (**MBSA**) is a software tool released by Microsoft to determine security state by assessing missing security updates and less-secure security settings within Microsoft Windows, Windows components such as Internet Explorer, IIS web server, and products Microsoft SQL Server, and Microsoft Office macro settings. Security updates are determined by the current version of MBSA using the Windows Update Agent present on Windows computers since Windows 2000 Service Pack 3. The less-secure settings, often called Vulnerability Assessment (VA) checks, are assessed based on a hard-coded set of registry and file checks. An example of a VA might be that permissions for one of the directories in the /www/root folder of IIS could be set at too low a level, allowing unwanted modification of files from outsiders.

Versions 1.2.1 and below run on NT4, Windows 2000, Windows XP, and Windows Server 2003, provide support for IIS versions 5 through 6, SQL Server 7 and 2000, Internet Explorer 5.01 and 6.0 only, and Microsoft Office 2000 through 2003. Security update assessment is provided by an integrated version of Shavlik's HFNetChk 3.8 scan tool. MBSA 1.2.1 was localized into English, German, French and Japanese versions and supported security assessment for any locale.

Version 2.0 retained the hard-coded VA checks, but replaced the Shavlik security assessment engine with Microsoft Update technologies which adds dynamic support for all Microsoft products supported by Microsoft Update. MBSA 2.0.1 was released to support the revised Windows Update (WU) offline scan file (WSUSSCN2.CAB). MBSA 2.1 added Vista and Windows Server 2008 support, a new Vista-styled GUI interface, support for the latest Windows Update Agent (3.0), a new Remote Directory (/rd) feature and extended the VA checks to x64 platforms.

See Also

- Belarc Advisor
- RSBAC

References

[1] "Download Details: Microsoft Baseline Security Analyzer v1.2.1 (for IT Professionals)" (http://www.microsoft.com/downloads/details.aspx?FamilyID=b13ebd6b-e258-4625-b0a3-64a4879f7798&displaylang=en). *Microsoft Download Center*. Microsoft Corporation. . Retrieved 13 October 2009.
[2] "Download Details: Microsoft Baseline Security Analyzer 2.2 (for IT Professionals)" (http://www.microsoft.com/downloads/en/details.aspx?displaylang=en&FamilyID=02be8aee-a3b6-4d94-b1c9-4b1989e0900c). *Microsfot Download Center*. Microsoft Corporation. 6 August 2010. . Retrieved 21 November 2009.
[3] http://www.microsoft.com/mbsa

External links

- Official website (http://technet.microsoft.com/en-us/security/cc184924.aspx)
  - Forum (http://social.technet.microsoft.com/Forums/en-US/MBSA)
  - Newsgroup (http://www.microsoft.com/communities/newsgroups/en-us/default.aspx?dg=microsoft.public.security.baseline_analyzer)
- Microsoft Office Visio 2007 Connector for the Microsoft Baseline Security Analyzer (MBSA) 2.1 (http://technet.microsoft.com/en-us/security/cc184925.aspx)

Troubleshooting

Troubleshooting is a form of problem solving, often applied to repair failed products or processes. It is a logical, systematic search for the source of a problem so that it can be solved, and so the product or process can be made operational again. Troubleshooting is needed to develop and maintain complex systems where the symptoms of a problem can have many possible causes. Troubleshooting is used in many fields such as engineering, system administration, electronics, automotive repair, and diagnostic medicine. Troubleshooting requires identification of the malfunction(s) or symptoms within a system. Then, experience is commonly used to generate possible causes of the symptoms. Determining which cause is most likely is often a process of elimination - eliminating potential causes of a problem. Finally, troubleshooting requires confirmation that the solution restores the product or process to its working state.

In general, troubleshooting is the identification of, or diagnosis of "trouble" in the management flow of a corporation or a system caused by a failure of some kind. The problem is initially described as symptoms of malfunction, and troubleshooting is the process of determining and remedying to the causes of these symptoms.

A system can be described in terms of its expected, desired or intended (usually, for artificial systems, its purpose). Events or inputs to the system are expected to generate specific results or outputs. (For example selecting the "print" option from various computer applications is intended to result in a hardcopy emerging from some specific device). Any unexpected or undesirable behavior is a symptom. Troubleshooting is the process of isolating the specific cause or causes of the symptom. Frequently the symptom is a failure of the product or process to produce any results. (Nothing was printed, for example).

The methods of forensic engineering are especially useful in tracing problems in products or processes, and a wide range of analytical techniques are available to determine the cause or causes of specific failures. Corrective action can then be taken to prevent further failures of a similar kind. Preventative action is possible using failure mode and effects analysis (FMEA) and fault tree analysis (FTA) before full scale production, and these methods can also be used for failure analysis.

Aspects

Most discussion of troubleshooting, and especially training in formal troubleshooting procedures, tends to be domain specific, even though the basic principles are universally applicable.

Usually troubleshooting is applied to something that has suddenly stopped working, since its previously working state forms the expectations about its continued behavior. So the initial focus is often on recent changes to the system or to the environment in which it exists. (For example a printer that "was working when it was plugged in over there"). However, there is a well known principle that correlation does not imply causality. (For example the failure of a device shortly after it's been plugged into a different outlet doesn't necessarily mean that the events were related. The failure could have been a matter of coincidence.) Therefore troubleshooting demands critical thinking rather than magical thinking.

It's useful to consider the common experiences we have with light bulbs. Light bulbs "burn out" more or less at random; eventually the repeated heating and cooling of its filament, and fluctuations in the power supplied to it cause the filament to crack or vaporize. The same principle applies to most other electronic devices and similar principles apply to mechanical devices. Some failures are part of the normal wear-and-tear of components in a system.

A basic principle in troubleshooting is to start from the simplest and most probable possible problems first. This is illustrated by the old saying "When you see hoof prints, look for horses, not zebras", or to use another maxim, use the KISS principle. This principle results in the common complaint about help desks or manuals, that they sometimes first ask: "Is it plugged in and does that receptacle have power?", but this should not be taken as an affront, rather it should serve as a reminder or conditioning to always check the simple things first before calling for help.

A troubleshooter could check each component in a system one by one, substituting known good components for each potentially suspect one. However, this process of "serial substitution" can be considered degenerate when components are substituted without regards to a hypothesis concerning how their failure could result in the symptoms being diagnosed.

Simple and intermediate systems are characterized by lists or trees of dependencies among their components or subsystems. More complex systems contain cyclical dependencies or interactions (feedback loops). Such systems are less amenable to "bisection" troubleshooting techniques.

It also helps to start from a known good state, the best example being a computer reboot. A cognitive walkthrough is also a good thing to try. Comprehensive documentation produced by proficient technical writers is very helpful, especially if it provides a theory of operation for the subject device or system.

A common cause of problems is bad design, for example bad human factors design, where a device could be inserted backward or upside down due to the lack of an appropriate forcing function (behavior-shaping constraint), or a lack of error-tolerant design. This is especially bad if accompanied by habituation, where the user just doesn't notice the incorrect usage, for instance if two parts have different functions but share a common case so that it isn't apparent on a casual inspection which part is being used.

Troubleshooting can also take the form of a systematic checklist, troubleshooting procedure, flowchart or table that is made before a problem occurs. Developing troubleshooting procedures in advance allows sufficient thought about the steps to take in troubleshooting and organizing the troubleshooting into the most efficient troubleshooting process. Troubleshooting tables can be computerized to make them more efficient for users.

Some computerized troubleshooting services (such as Primefax, later renamed Maxserve), immediately show the top 10 solutions with the highest probability of fixing the underlying problem. The technician can either answer additional questions to advance through the troubleshooting procedure, each step narrowing the list of solutions, or immediately implement the solution he feels will fix the problem. These services give a rebate if the technician takes an additional step after the problem is solved: report back the solution that actually fixed the problem. The computer uses these reports to update its estimates of which solutions have the highest probability of fixing that particular set of symptoms.[1]

Half-splitting

Efficient methodical troubleshooting starts with a clear understanding of the expected behavior of the system and the symptoms being observed. From there the troubleshooter forms hypotheses on potential causes, and devises (or perhaps references a standardized checklist of) tests to eliminate these prospective causes. This approach is often called "Divide and Conquer".

Two common strategies used by troubleshooters are to check for frequently encountered or easily tested conditions first (for example, checking to ensure that a printer's light is on and that its cable is firmly seated at both ends). This is often referred to as "milking the front panel."[2]

Then, "bisect" the system (for example in a network printing system, checking to see if the job reached the server to determine whether a problem exists in the subsystems "towards" the user's end or "towards" the device).

This latter technique can be particularly efficient in systems with long chains of serialized dependencies or interactions among its components. It's simply the application of a binary search across the range of dependencies and is often referred to as "half-splitting".[3]

Reproducing symptoms

One of the core principles of troubleshooting is that reproducible problems can be reliably isolated and resolved. Often considerable effort and emphasis in troubleshooting is placed on reproducibility ... on finding a procedure to reliably induce the symptom to occur.

Once this is done then systematic strategies can be employed to isolate the cause or causes of a problem; and the resolution generally involves repairing or replacing those components which are at fault.

Intermittent symptoms

Some of the most difficult troubleshooting issues relate to symptoms that are only intermittent. In electronics this often is the result of components that are thermally sensitive (since resistance of a circuit varies with the temperature of the conductors in it). Compressed air can be used to cool specific spots on a circuit board and a heat gun can be used to raise the temperatures; thus troubleshooting of electronics systems frequently entails applying these tools in order to reproduce a problem.

In computer programming race conditions often lead to intermittent symptoms which are extremely difficult to reproduce; various techniques can be used to force the particular function or module to be called more rapidly than it would be in normal operation (analogous to "heating up" a component in a hardware circuit) while other techniques can be used to introduce greater delays in, or force synchronization among, other modules or interacting processes.

Intermittent issues can be thus defined:

> An intermittent is a problem for which there is no known procedure to consistently reproduce its symptom.
>
> —Steven Litt, [4]

In particular he asserts that there is a distinction between frequency of occurrence and a "known procedure to consistently reproduce" an issue. For example knowing that an intermittent problem occurs **"within"** an hour of a particular stimulus or event ... but that sometimes it happens in five minutes and other times it takes almost an hour ... does not constitute a "known procedure" even if the stimulus does increase the frequency of observable exhibitions of the symptom.

Nevertheless, sometimes troubleshooters must resort to statistical methods ... and can only find procedures to increase the symptom's occurrence to a point at which serial substitution or some other technique is feasible. In such cases, even when the symptom seems to disappear for significantly longer periods, there is a low confidence that the root cause has been found and that the problem is truly solved.

Also, tests may be run to stress certain components to determine if those components have failed. [5]

Multiple problems

Isolating single component failures which cause reproducible symptoms is relatively straightforward.

However, many problems only occur as a result of multiple failures or errors. This is particularly true of fault tolerant systems, or those with built-in redundancy. Features which add redundancy, fault detection and failover to a system may also be subject to failure, and enough different component failures in any system will "take it down."

Even in simple systems the troubleshooter must always consider the possibility that there is more than one fault. (Replacing each component, using serial substitution, and then swapping each new component back out for the old one when the symptom is found to persist, can fail to resolve such cases. More importantly the replacement of any component with a defective one can actually increase the number of problems rather than eliminating them).

Note that, while we talk about "replacing components" the resolution of many problems involves adjustments or tuning rather than "replacement." For example, intermittent breaks in conductors --- or "dirty or loose contacts" might simply need to be cleaned and/or tightened. All discussion of "replacement" should be taken to mean "replacement or adjustment or other maintenance."

References

[1] "Troubleshooting at your fingertips" by Nils Conrad Persson. "Electronics Servicing and Technology" magazine 1982 June.

[2] "Hewlett Packard Bench Briefs" (http://www.hparchive.com/Bench_Briefs/HP-Bench-Briefs-1982-01-02.pdf). Hewlett Packard. . Retrieved 14 October 2011.

[3] Sullivan, Mike (Nov 15, 2000). "Secrets of a super geek: Use half splitting to solve difficult problems" (http://articles.techrepublic.com.com/5100-10878_11-5029507.html). *TechRepublic*. . Retrieved 22 October 2010.

[4] http://www.troubleshooters.com/tpromag/9812.htm#DefinitionofanIntermittent

[5] http://www.ocf.berkeley.edu/~joyoung/trouble/page1.shtml

Software Package Management

Installation (computer programs)

Installation (or **setup**) of a computer program (including device drivers and plugins), is the act of making the program ready for execution. Because the process varies for each program and each computer, programs (including operating systems) often come with an *installer*, a specialized program responsible for doing whatever is needed for their installation.

Software installation can occur on one machine at a time or can be automated through software distribution software to update multiple computers at once. This is especially useful for large organizations where there may be hundreds or thousands of computers needing to be updated continuously in order to operate most effectively.[1] Managing installations and updates is a part of Software Asset Management, which is generally a subset of a company's corporate risk management strategy.

Overview

Some computer programs can be executed by simply copying them into a folder stored on a computer and executing them. Other programs are supplied in a form unsuitable for immediate execution and therefore need an installation procedure. Once installed, the program can be executed again and again, without the need to reinstall before each execution.

Common operations performed during software installations include:

- Making sure that required system requirements are present
- Checking for existing versions of the software
- Creating or updating program files and folders
- Adding configuration data such as configuration files, Windows registry entries or environment variables
- Making the software accessible to user, for instance by creating links, shortcuts or bookmarks
- Configuring components that run automatically, such as daemons or Windows services
- Performing product activation

Necessity

As mentioned earlier, some computer program need no installation. This was once usual for many programs which run on DOS, Mac OS, Atari TOS and AmigaOS. As the computing environment grew more complex however, the need for tangible installation presented itself.

Nowadays, a class of modern applications that do not need installation are known as *portable applications*, as they may be roamed around onto different computers and run. Similarly, there are *live operating systems*, which do not need installation and can be run directly from a bootable CD, DVD, or USB flash drive. Examples are AmigaOS 4.0, various Linux distributions, MorphOS or Mac OS versions 1.0 through 9.0. (See live CD and live USB.) Finally, web applications, which run inside a web browser, do not need installation.

Types

Attended installation

This is the most common form of installation. An installation process usually needs a user who attends it to make choices, such as accepting or declining an end-user license agreement (EULA), specifying preferences such as the installation location, supplying passwords or assisting in product activation. In graphical environments, installers that offer a wizard-based interface are common. Attended installers may ask users to help mitigate the errors. For instance, if the disk in which the computer program is be installed was full, the installer may ask the user to specify another target path.

Silent installation

Installation that does not display messages or windows during its progress. "Silent installation" is not the same as "unattended installation" (see below): All silent installations are unattended but not all unattended installations are silent. The reason behind a silent installation may be convenience or subterfuge. Malware is always installed silently.

Unattended installation

Installation that is performed without user interaction during its progress or with no user present at all. An unattended installation either does not require the user to supply anything or has received all necessary input prior to the start of installation. Such input may be in the form of command line switches or an *answer file*, a file that contains all the necessary parameters. Windows XP is an example of an operating system that can be installed with an answer file. In unattended installation, it is assumed that there is no user to help mitigate errors. For instance, if the installation medium was faulty, the installer should fail the installation, as there is no user to fix the fault or replace the medium. Unattended installers may record errors in a computer log for later review.

Headless installation

Installation performed without using a computer monitor connected. In attended forms of headless installation, another machine connects to the target machine (for instance, via a local area network) and takes over the display output. Since a headless installation does not need a user at the location of the target computer, unattended headless installers may be used to install a computer software on multiple machines at the same time.

Scheduled or automated installation

An installation process that runs on a preset time or when a predefined condition transpires, as opposed to an installation process that starts explicitly on a user's command. For instance, a system administrator willing to install a later version of a computer program that is being used can schedule that installation to occur when that program is not running. An operating system may automatically install a device driver for a device that the user connects. (See plug and play.) Malware may also be installed automatically. For example, the infamous Conficker was installed when the user plugged an infected device to his computer.

Clean installation

A clean installation is one that is done in the absence of any interfering elements such as old versions of the computer program being installed or leftovers from a previous installation. In particular, the clean installation of an operating system is an installation in which the target disk partition is erased before installation. Since the interfering elements are absent, a clean installation may succeed where an unclean installation may fail or may take significantly longer.

Network installation

An installation of a program from a shared network resource. This may simply be a copy of the original media but software publishers which offer site licenses for institutional customers may provide a version intended for installation over a network.

Installer

An *installation program* or *installer* is a computer program that installs files, such as applications, drivers, or other software, onto a computer. Some installers are specifically made to install the files they contain; other installers are general-purpose and work by reading the contents of the software package to be installed.

The differences between a package management system and an installer are:

Package management system	Installer
Usually part of an operating system.	Each product comes bundled with its own installer.
Uses one installation database.	Performs its own installation, sometimes recording information about that installation in a registry.
Can verify and manage all packages on the system.	Works only with its bundled product.
One package management system vendor.	Multiple installer vendors.
One package format.	Multiple installation formats.

Bootstrapper

During the installation of computer programs it is sometimes necessary to update the installer or package manager itself. To make this possible, a technique called bootstrapping is used. The common pattern for this is to use a small executable file (e.g. setup.exe) which updates the installer and starts the real installation after the update. This small executable is called bootstrapper. Sometimes the bootstrapper installs other prerequisites for the software during the bootstrapping process too.

Common types

Cross platform installer builders that produce installers for Windows, Mac OS X and Linux include InstallAnywhere (Flexera Software), JExpress (DeNova),[2] and InstallBuilder (BitRock Inc.).

Installers for Microsoft Windows include Windows Installer, a software installation component. Additional third party commercial tools for creating installers for Windows include InstallShield (Flexera Software), Advanced Installer (Caphyon Ltd),[3] InstallAware (InstallAware Software),[4] Wise Installation Studio (Wise Solutions, Inc.), SetupBuilder (Lindersoft, Inc.),[5] Installer VISE (MindVision Software), MSI Studio (ScriptLogic Corporation), Actual Installer (Softeza Development),[6] Smart Install Maker (InstallBuilders Company),[7] MSI Factory and Setup Factory (Indigo Rose Software), Centurion Setup (Gammadyne Corporation).[8] Free installer-authoring tools include NSIS, IzPack, Clickteam, InnoSetup, InstallSimple and WiX.

Mac OS X includes Installer, a native Package Manager software. Mac OS X also includes a separate software updating application, Software Update but only supports Apple and system software. Included in the dock as of 10.6.6, the Mac App Store shares many attributes with the successful App Store for iOS devices, such as a similar app approval process, the use of Apple ID for purchases, and automatic installation and updating. Although this is Apple's preferred delivery method for Mac OS X,[9] previously purchased licenses can not be transferred to the Mac App Store for downloading or automatic updating. Commercial applications for Mac OS X may also use a third-party installer, such as Mac version of Installer VISE (MindVision Software) or InstallerMaker (StuffIt).

References

[1] "Software Distribution" (http://www.kace.com/products/systems-management-appliance/features/software-distribution.php). Dell KACE. . Retrieved 2012-07-06.

[2] Denova.com (http://denova.com)

[3] "Caphyon setup authoring solutions" (http://www.advancedinstaller.com/). Advanced Installer. .

[4] "Installation Software for Windows Programs" (http://www.installaware.com/). InstallAware. .

[5] "Lindersoft: Software Installation Solutions - High-quality Software Installations" (http://www.setupbuilder.com/). Setupbuilder.com. . Retrieved 2010-05-02.

[6] "Installation Software for Windows Programs" (http://www.actualinstaller.com/). Actual Installer. . Retrieved 2010-05-02.

[7] "Smart Install Maker - Custom setup files made easy" (http://www.sminstall.com/). Sminstall.com. . Retrieved 2010-05-02.

[8] "Centurion Setup" (http://www.gammadyne.com/centurion.htm). Gammadyne Corporation. .

[9] http://www.apple.com/mac/app-store/great-mac-apps.html

External links

- The Application Deployment Information Center (http://www.appdeploy.com/)

Package management system

In software, a **package management system**, also called **package manager**, is a collection of software tools to automate the process of installing, upgrading, configuring, and removing software packages for a computer's operating system in a consistent manner. It typically maintains a database of software dependencies and version information to prevent software mismatches and missing prerequisites.

Packages are distributions of software, applications and data. Packages also contain metadata, such as the software's name, description of its purpose, version number, vendor, checksum, and a list of dependencies necessary for the software to run properly. Upon installation, metadata is stored in a local package database.

Package management systems are designed to save organizations time and money through remote administration and software distribution technology that eliminate the need for manual installs and updates. This can be particularly useful for large enterprises whose operating systems are based on Linux and other Unix-like systems, typically consisting of hundreds or even thousands of distinct software packages; in the former case, a package management system is a convenience, in the latter case it becomes essential.[1]

Impact

Ian Murdock has commented that package management is "the single biggest advancement Linux has brought to the industry", that it blurs the boundaries between operating system and applications, and that it makes it "easier to push new innovations [...] into the marketplace and [...] evolve the OS".[2]

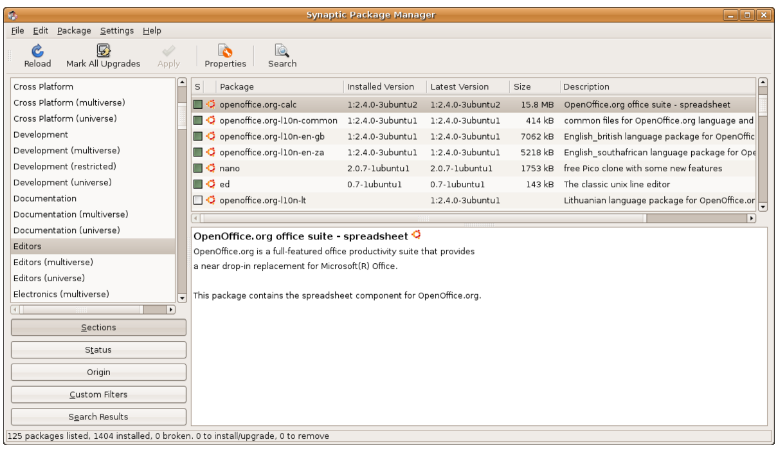

Synaptic, a popular graphical frontend for package management in Linux

Terminology

A package management system is often called an "install manager". This can lead to confusion between a package management system and an installer. The differences include:

Package management system	Installer
Usually part of an operating system.	Each product comes bundled with its own installer.
Uses one installation database.	Performs its own installation, sometimes recording information about that installation in a registry.
Can verify and manage all packages on the system.	Works only with its bundled product.
One package management system vendor.	Multiple installer vendors.
One package format.	Multiple installation formats.

A package, for package managers, denotes a specific set of files bundled with the appropriate metadata for use by a package manager. This can be confusing, as some programming languages often use the word "package" as a specific form of software library. Furthermore, that software library can be distributed in a package of files bundled for a package manager.

Functions

Package management systems are charged with the task of organizing all of the packages installed on a system. Typical functions of a package management system include:

- Verifying file checksums to ensure correct and complete packages;
- Verifying digital signatures to authenticate the origin of packages;
- Applying file archivers to manage encapsulated files;
- Upgrading software with latest versions, typically from a software repository;
- Grouping of packages by function to reduce user confusion;
- Managing dependencies to ensure a package is installed with all packages it requires. This resolved the problem known as Dependency Hell.

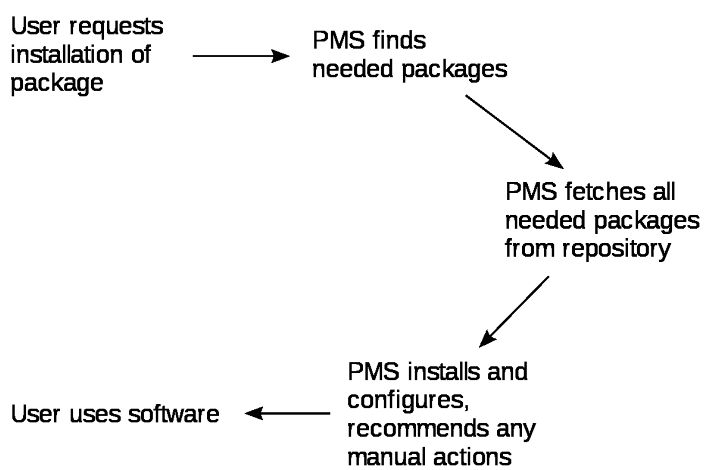

Illustration of a package management system being used to download new software. Manual actions can include accepting a license agreement or selecting some package-specific configuration options.

Some additional challenges are met by only a few package management systems.

Challenges with shared libraries

Computer systems which rely on dynamic library linking, instead of static library linking, share executable libraries of machine instructions across packages and applications. In these systems, complex relationships between different packages requiring different versions of libraries results in a challenge colloquially known as "dependency hell". On Microsoft Windows systems, this is also called "DLL hell" when working with dynamically linked libraries. Good package management systems become vital on these systems.[3]

Front-ends for locally compiled packages

System administrators may install and maintain software using tools other than package management software. For example, a local administrator may download unpackaged source code, compile it, and install it. This may cause the state of the local system to fall out of synchronization with the state of the package manager's database. The local administrator will be required to take additional measures, such as manually managing some dependencies or integrating the changes into the package manager.

There are tools available to ensure that locally compiled packages are integrated with the package management. For distributions based on .deb and .rpm files as well as Slackware Linux, there is CheckInstall, and for recipe-based systems such as Gentoo Linux and hybrid systems such as Arch Linux, it is possible to write a recipe first, which then ensures that the package fits into the local package database.

Maintenance of configuration

Particularly troublesome with software upgrades are upgrades of configuration files. Since package management systems, at least on Unix systems, originated as extensions of file archiving utilities, they can usually only either overwrite or retain configuration files, rather than applying rules to them. There are exceptions to this that usually apply to kernel configuration (which, if broken, will render the computer unusable after a restart). Problems can be caused if the format of configuration files changes. For instance, if the old configuration file does not explicitly disable new options that should be disabled. Some package management systems, such as Debian's dpkg, allow configuration during installation. In other situations, it is desirable to install packages with the default configuration and then overwrite this configuration, for instance, in headless installations to a large number of computers. (This kind of pre-configured installation is also supported by dpkg.)

Repositories

To give users more control over the kinds of software that they are allowing to be installed on their system (and sometimes due to legal or convenience reasons on the distributors' side), software is often downloaded from a number of software repositories.[4]

Upgrade suppression

When a user interacts with the package management software to bring about an upgrade, it is customary to present the user with the list of things to be done (usually the list of packages to be upgraded, and possibly giving the old and new version numbers), and allow the user to either accept the upgrade in bulk, or select individual packages for upgrades. Many package management systems can be configured to never upgrade certain packages, or to upgrade them only when critical vulnerabilities or instabilities are found in the previous version, as defined by the packager of the software. This process is sometimes called *version pinning*.

For instance:

- yum supports this with the syntax *exclude=openoffice**,[5]
- pacman with *IgnorePkg = openoffice*[6] (to suppress upgrading openoffice in both cases)
- dpkg and dselect support this partially through the *hold* flag in package selections
- APT extends the *hold* flag through the complex "pinning" mechanism[7]
    - Users can also blacklist a package[8]
- aptitude has "hold" and "forbid" flags
- portage supports this through the package.mask configuration file
- DotPet from .PET packages in Puppy Linux, but is efficient in console and lightweight

Cascading package removal

Some of the more advanced package management features offer "cascading package removal",[6] in which all packages that depend on the target package and all packages that only the target package depends on, are also removed.

Common package management systems and formats

Package formats

Each package manager relies on the format and metadata of the packages it can manage. That is, package managers need groups of files to be bundled for the specific package manager along with appropriate metadata, such as dependencies. Often, a core set of utilities manages the basic installation from these packages and multiple package managers use these utilities to provide additional functionality.

For example, yum relies on rpm as a backend. Yum extends the functionality of the backend by adding features such as simple configuration for maintaining a network of systems. As another example, the Synaptic Package Manager provides a graphical user interface by using the Advanced Packaging Tool (apt) library, which, in turn, relies on dpkg for core functionality.

Alien is a program that converts between different Linux package formats. It supports conversion between Linux Standard Base conform RPM, deb, Stampede (.slp) and Slackware (tgz) packages.

Free and open source software systems

By the nature of free and open source software, packages under similar and compatible licenses are available for use on a number of operating systems. These packages can be combined and distributed using configurable and internally complex packaging systems to handle many permutations of software and manage version-specific dependencies and conflicts. Some packaging systems of free and open source software are also themselves released as free and open source software. One typical difference between package management in proprietary operating systems, such as Mac OS X and Windows, and those in free and open source software, such as Linux, is that free and open source software systems permit third-party packages to also be installed and upgraded through the same mechanism, whereas the package management systems of Mac OS X and Windows will only upgrade software provided by Apple and Microsoft, respectively (with the exception of some third party drivers in Windows). The ability to continuously upgrade third party software is typically added by adding the URL of the corresponding repository to the package management's configuration file.

Application-level package managers

Besides the systems-level application managers, there are some add-on package managers for operating systems with limited capabilities and for programming languages where developers need the latest libraries.

In contrast to systems-level application managers, application-level package managers focus on a small part of the software system. They typically reside within a directory tree that is not maintained by the systems-level package manager (like c:\cygwin or /usr/local/fink). However, this is not the case for the package managers that deal with programming libraries. This leads to a conflict as both package managers claim to "own" a file and might break upgrades.

References

[1] "Software Distribution" (http://www.kace.com/products/systems-management-appliance/features/software-distribution.php). Dell KACE. . Retrieved 2012-07-11.

[2] "How package management changed everything" (http://ianmurdock.com/2007/07/21/how-package-management-changed-everything/). ianmurdock.com. . Retrieved 2008-03-01.

[3] Chris, Tucker (2007–03–15). *Optimal Package Install/Uninstall Manager* (http://cseweb.ucsd.edu/~lerner/papers/opium.pdf). UC San Diego. p. 1. . Retrieved 2011–09–14.

[4] "Linux repository classification schemes" (http://braintickle.blogspot.com/2006/01/linux-repository-classification.html). braintickle.blogspot.com. . Retrieved 2008-03-01.

[5] "CentOS yum pinning rpms" (http://web.archive.org/web/20071102203232/http://lists.centos.org/pipermail/centos/2005-May/046320.html). centos.org. Archived from the original (http://lists.centos.org/pipermail/centos/2005-May/046320.html) on 2007-11-02. . Retrieved 2008-03-01.

[6] "pacman(8) Manual Page" (http://www.archlinux.org/pacman/pacman.8.html). archlinux.org. . Retrieved 2008-03-01.

[7] "How to keep specific versions of packages installed (complex)" (http://www.debian.org/doc/manuals/apt-howto/ch-apt-get.en.html#s-pin). debian.org. . Retrieved 2008-03-01.

[8] "Apt pinning to blacklist a package" (http://linux.derkeiler.com/Mailing-Lists/Debian/2009-07/msg00476.html). . Retrieved 2010-08-19.

External links

- Distrowatch Comparison of Package Management Systems (http://distrowatch.com/dwres.php?resource=package-management)

Advanced Packaging Tool

Advanced Packaging Tools

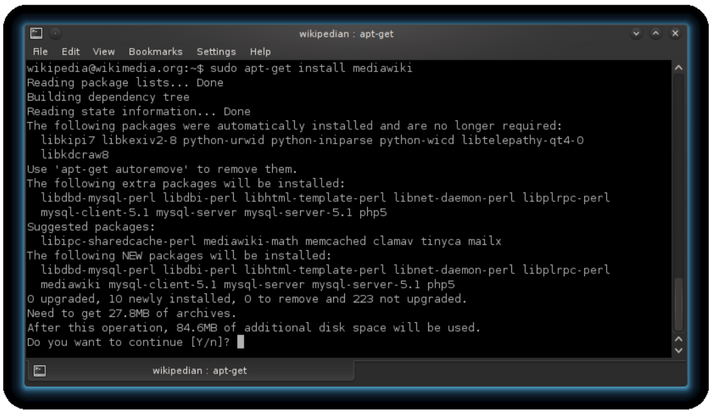

`apt-get` requesting confirmation before an installation

Initial release	16 August 1998[1]
Stable release	0.8.10.3[2] / April 15, 2011
Preview release	0.8.16~exp13 / March 13, 2012
Programming language used	C++
Operating system	Cross-platform
Platform	Unix and Unix-like operating systems
Type	Package management system
License	GNU General Public License
Website	wiki.debian.org/Apt [3] wiki.debian.org/Teams/Apt [4]

The **Advanced Packaging Tool**, or **APT**, is a free user interface that works with core libraries to handle the installation and removal of software on the Debian GNU/Linux distribution and its variants.[5] APT simplifies the process of managing software on Unix-like computer systems by automating the retrieval, configuration and installation of software packages, either from binary files or by compiling source code.[5]

APT was originally designed as a front-end for dpkg to work with Debian's `.deb` packages, but it has since been modified to also work with the RPM Package Manager system via apt-rpm.[6] The Fink project has ported APT to Mac OS X for some of its own package management tasks, and APT is also available in OpenSolaris (included in the Nexenta OS distribution).[7] Cydia is a package manager for jailbroken iOS partially based on APT (ported to iOS as part of the related Telesphoreo project).[8][9]

Usage

There is no single "apt" program; apt is itself the package name containing the set of tools (and requiring the libraries) that support its functionality. A significant part of apt is a C++ library of functions (another package known as *libapt*) which are used by these related front-end programs for dealing with packages, such as `apt-get` and `apt-cache`. They are commonly used in examples due to their simplicity and ubiquity; apt-get and apt-cache are of "*important*" priority in all current Debian releases, and are therefore installed in a default Debian installation. Apt can be functionally considered to be a front-end to `dpkg`, and a friendlier front end to this than `dselect`. While dpkg performs actions on individual packages, apt tools manage relations (especially dependencies) between them, as well as sourcing and management of higher-level versioning decisions (release tracking and version pinning).

APT is often hailed as one of Debian's best features.[10][11][12][13] It is remarked that this quality comes from the strict quality controls of Debian policy.[14][15]

A major feature in APT is the way it calls dpkg — it does topological sorting of the list of packages to be installed or removed and calls dpkg in the best possible sequence. In some cases, it utilizes the `--force` options in dpkg. However, it only does this when it is unable to calculate how to avoid the reason dpkg requires the action to be forced.

Installation of software

An install directive is followed by the name of one or more packages desired for installation. Each package name is phrased as just the name portion of the package, not a fully qualified filename (for instance, in a Debian GNU/Linux system, libc6 would be the argument provided, not `libc6_1.9.6-2.deb`). Notably, all packages containing dependencies required by the package(s) specified for installation will also be automatically retrieved and installed. This was an original distinguishing characteristic of apt-based package management systems whereby software installation failure due to missing dependencies, a type of dependency hell, was specifically avoided.

Another such distinction is remote repository retrieval of packages. A location configuration file (`/etc/apt/sources.list`) is used to locate the desired packages and retrieve them, and also obtain information about available (but uninstalled) packages.

Other command option features (switches) may be used to override decisions made by apt-get's conflict resolution system. If a hyphen is appended to the package name (with no intervening space), the identified package will be removed if it is installed. Similarly, a plus sign can be used to designate a package to install. A specific version of a package can be selected for installation by following the package name with an equals and the version of the package to select. This will cause that version to be located and selected for install. Alternatively a specific distribution can be selected by following the package name with a slash and the version of the distribution or the archive name (stable, testing, unstable).

Both of the version selection mechanisms can downgrade packages and must be used with care.

Finally, the `apt_preferences` mechanism allows creating an alternative installation policy for individual packages.

If no package matches the given expression and the expression contains one of '`.`', '`?`' or '`*`', it is assumed to be a POSIX regular expression and it is applied to all package names in the database. Any matches are then installed (or removed). Note that matching is done by substring, so "`lo.*`" matches "how-lo" and "lowest". If this is undesired, the regular expression can be anchored with a '`^`' or '`$`' character, or a more specific regular expression can be created.

Update, upgrade and dist-upgrade

- **update** is used to resynchronize the **package index** files from their sources. The lists of available packages are fetched from the location(s) specified in `/etc/apt/sources.list`. For example, when using a Debian archive, this command retrieves and scans the `Packages.gz` files, so that information about new and updated packages is available. An update should always be performed before a safe-upgrade or dist-upgrade. Be aware that the overall progress meter will not always be correct as the size of the package files cannot be known in advance.
- **upgrade** is used to install the newest versions of all packages currently installed on the system from the sources enumerated in `/etc/apt/sources.list`. Packages currently installed with new versions available are retrieved and upgraded; under no circumstances are currently installed packages removed, or packages not

already installed retrieved and installed. New versions of currently installed packages that cannot be upgraded without changing the install status of another package will be left at their current version. An update must be performed first, so that apt-get knows that new versions of packages are available.

- **dist-upgrade**, in addition to performing the function of upgrade, also intelligently handles changing dependencies with new versions of packages; apt-get has a "smart" conflict resolution system, and it will attempt to upgrade the most important packages at the expense of less important ones if necessary. The `/etc/apt/sources.list` file contains a list of locations from which to retrieve desired package files.[16] aptitude has a smarter dist-upgrade feature called full-upgrade.[17]

Configuration and files

`/etc/apt` has the apt configuration folders and files.

apt-config is the APT Configuration Query program.[18] `apt-config dump` shows the configuration.[19]

Files

- `/etc/apt/sources.list`: Locations to fetch packages from.
- `/etc/apt/sources.list.d/`: Additional source list fragments.
- `/etc/apt/apt.conf`: APT configuration file.
- `/etc/apt/apt.conf.d/`: APT configuration file fragments.
- `/etc/apt/preferences`: version preferences file. This is where you would specify "pinning", i.e. a preference to get certain packages from a separate source or from a different version of a distribution.
- `/var/cache/apt/archives/`: storage area for retrieved package files.
- `/var/cache/apt/archives/partial/`: storage area for package files in transit.
- `/var/lib/apt/lists/`: storage area for state information for each package resource specified in `sources.list`
- `/var/lib/apt/lists/partial/`: storage area for state information in transit.

Sources

APT relies on the concept of repositories in order to find software and resolve dependencies. For apt, a repository is a directory containing packages along with an index file. This can be specified as a networked or CDROM location. The Debian project keeps a central repository of over 25,000 software packages ready for download and installation.

For extra packages, any number of additional repositories can be added to APT's *sources.list* configuration file (`/etc/apt/sources.list`) and then be queried by APT. Graphical front-ends often allow modifying `sources.list` more simply (apt-setup). Once a package repository has been specified (like during the system installation), packages in that repository can be installed without specifying a source.

In addition to network repositories, compact discs and other storage media (USB keydrive, hard disks...) can be used as well, using apt-cdrom [20] or adding `file:/`[21] to the source list file. Apt-cdrom can specify a different folder than a cd-rom, using the `-d` option (i.e. a hard disk or a USB keydrive). The Debian CDs available for download contain Debian repositories. This allows non-networked machines to be upgraded. Also one can use apt-zip.

Problems may appear when several sources offer the same package(s). Systems that have such possibly conflicting sources can use APT pinning to control which sources should be preferred.

APT pinning

The *APT pinning* feature allows administrators to force APT to choose particular versions of packages which may be available in different versions from different repositories. This allows administrators to ensure that packages are not upgraded to versions which may conflict with other packages on the system, or that have not been sufficiently tested for unwelcome changes.

In order to do this, the *pins* in APT's *preferences* file (`/etc/apt/preferences`) must be modified,[22] although graphical front-ends often make pinning simpler.

Front-ends

Several other front-ends to APT exist, which provide more advanced installation functions and more intuitive interfaces. These include:

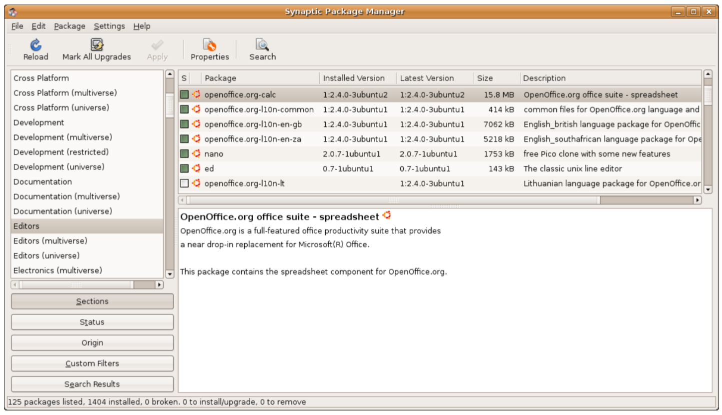

Synaptic Package Manager is one of the frontends available for APT

- *Synaptic Package Manager*, a GTK+ graphical user interface
- *Ubuntu Software Center*, a GTK+ graphical user interface replacement for Synaptic
- *aptitude*, a versatile alternative to dselect
- *KPackage*, part of KDE
- *Adept Package Manager*, a graphical user interface for KDE (deb, rpm, bsd)
- *PackageKit*, a freedesktop.org frontend.
- *GDebi*, a GTK-based tool sponsored for Ubuntu. (There is also a Qt version, available in the Ubuntu repositories a gdebi-kde.)
- *apt-cdrom*, a way to add a new CDROM to APT's list of available sources.lists (list of available repositories). It is necessary to use apt-cdrom to add CDs to the APT system, it cannot be done by hand.
- *apt-zip*, a way to use apt with removable media, specifically USB flash drives.
- *apt:foo*, an interface for downloading and installing software by clicking on a web-link (experimental) [23]
- *gnome-apt*, a gtk/GNOME-widget-based graphical front-end.
- *Cydia*, a package manager for jailbroken iOS partially based on APT (ported to iOS as part of the Telesphoreo project).[24][25]
- *Hildon Application Manager (Maemo Application)*, a Maemo front-end
- *APT Daemon*, a front-end that runs as a service to allow standard users to install software through PolicyKit and is in turn the framework that the Ubuntu Software Center (along with the Linux Mint Software Manager) uses to not be root and still run.

APT front-ends can:

- Search for new packages.
- Upgrade packages.
- Install or remove packages.
- Upgrade the whole system to a new release.

APT front-ends can list the dependencies of packages being installed or upgraded, ask the administrator if packages recommended or suggested by newly installed packages should be installed too, automatically install dependencies and perform other operations on the system such as removing obsolete files and packages.

History

The original effort that led to the apt-get program was the dselect replacement project known by its codename *deity*.[26] This project was commissioned by Brian White, the Debian Release Manager at the time. The very first functional version of apt-get was called dpkg-get and was only intended to be a test program for the core library functions that would underpin the new UI.[27]

Much of the original development of APT was done on IRC, so records have been lost. The 'Deity Creation Team' mailing list archives include only the major highlights.

The Deity name was abandoned as the official name for the project due to concerns over the religious nature of the name. The APT name was eventually decided after considerable internal and public discussion. Ultimately the name was proposed on IRC, accepted and then finalized on the mailing lists.[28] As originally used, APT is not an acronym, but a proper name. The name gained mindshare during IRC discussions due to the variety of possible acronym expansions and it was ultimately decided that the official use of APT would be as a proper name and no official expansion would ever be presented by the team.

APT was introduced in 1998 and original test builds were circulated on IRC. The first Debian version that included it was Debian 2.1, released on 9 March 1999.[29]

In the end the original goal of the Deity project of replacing the dselect UI was a failure. Work on the user interface (UI) portion of the project was abandoned (the UI directories were removed from the CVS system) after the first public release of apt-get. The response to APT as a dselect method and a command line utility was so great and positive that all development efforts focused on maintaining and improving the tool. It was not until much later that several independent people built UIs on top of the capable libapt-pkg.

The final push of the first APT era was to build a complete dpkg replacement (libapt-inst). This project was also a failure, however the partial code found a use as part of the secretive 'Project Betsy' program, which resulted in the highly efficient apt-ftparchive and libapt python bindings. After this, the original author faded away and maintainership of APT languished.

Eventually, a new team picked up the project, began to build new features and released version 0.6 of APT which introduced the Secure APT feature, using strong cryptographic signing to authenticate the package repositories.[30]

References

[1] "Advanced Packaging Tool" (http://dictionary.sensagent.com/advanced+packaging+tool/en-en/#cite_note-22). sensagent. . Retrieved 2012-02-22.

[2] "Debian -- Details of package apt" (http://packages.qa.debian.org/a/apt.html). packages.qa.debian.org. . Retrieved 2012-04-02.

[3] http://wiki.debian.org/Apt

[4] http://wiki.debian.org/Teams/Apt

[5] "apt-get(8) - Linux man page" (http://linux.die.net/man/8/apt-get). linux.die.net. Archived (http://web.archive.org/web/20080516191248/http://linux.die.net/man/8/apt-get) from the original on 16 May 2008. . Retrieved 2008-05-12.

[6] "APT-RPM" (http://apt-rpm.org/about.shtml). apt-rpm.org. Archived (http://web.archive.org/web/20080421014529/http://apt-rpm.org/about.shtml) from the original on 21 April 2008. . Retrieved 2008-05-12.

[7] "Fink - About" (http://www.finkproject.org/about.php). www.finkproject.org. Archived (http://web.archive.org/web/20080511210043/http://www.finkproject.org/about.php) from the original on 11 May 2008. . Retrieved 2008-05-12.

[8] "Bringing Debian APT to the iPhone" by Jay Freeman (saurik (http://www.saurik.com/id/1)

[9] Telesphoreo Tangelo documentation (http://www.telesphoreo.org/)

[10] Byfield, Bruce (2004-12-09). "An apt-get primer" (http://www.webcitation.org/5p6i2iySC). Archived from the original (http://www.linux.com/archive/feed/40745) on 2010-04-19. .

[11] "From the archives: the best distros of 2000" (http://www.tuxradar.com/content/archives-best-distros-2000). Tux Radar. .

[12] Dorgan, David (2004-01-19). "Migrating to Debian" (http://www.webcitation.org/5p6ikqQPs). linux.ie. Archived from the original (http://www.linux.ie/lists/pipermail/cork/2004-January/003744.html) on 2010-04-19. .

[13] "Mobile Linux development with Familiar and a minimal Debian" (http://www.vieka.com/mobiletux/debian.htm). Mobile Tux. .

[14] Why Debian (http://people.debian.org/~srivasta/talks/why_debian/talk.html)

[15] Debian policy manual (http://www.debian.org/doc/debian-policy/)

[16] die.net Linux Man pages (http://linux.die.net/man/8/apt-get)
[17] Discussion on dist-upgrade vs. full-upgrade (http://forums.debian.net/viewtopic.php?f=10&t=39379)
[18] Apt-Config (http://www.debianadmin.com/manpages/aptconfigmanpage.htm)
[19] Query APT Configuration Using apt-config - Debian Admin (http://www.debianadmin.com/query-apt-configuration-using-apt-config.html)
[20] apt-cdrom (http://ccrma.stanford.edu/planetccrma/man/man8/apt-cdrom.8.html)
[21] i.e. Deb file:/mnt/install stable main contrib. non-free
[22] "Apt Pinning" (http://wiki.debian.org/AptPreferences). Debian Wiki. Archived (http://web.archive.org/web/20061004194955/http://wiki.debian.org/AptPreferences) from the original on 4 October 2006. . Retrieved 2006-09-19.
[23] OzOS | apt:foo (http://cafelinux.org/OzOs/node/3)
[24] "Bringing Debian APT to the iPhone" by Jay Freeman (saurik (http://www.saurik.com/id/1)
[25] Telesphoreo Tangelo documentation (http://www.telesphoreo.org/)
[26] Deity Mailing List (http://lists.debian.org/debian-user/1997/04/msg00786.html), 1997-04.
[27] Deity Mailing List (http://lists.debian.org/deity/1998/03/msg00013.html), 1998-03.
[28] Deity Mailing List (http://lists.debian.org/deity/1998/03/msg00106.html), 1998-03.
[29] "A Brief History of Debian" (http://debian.org/doc/manuals/project-history/ch-detailed.en.html#s4.1). Debian.org. .
[30] "Secure APT" (http://wiki.debian.org/SecureApt). Debian Wiki. Archived (http://web.archive.org/web/20060901104920/http://wiki.debian.org/SecureApt#head-2ff8d9725145887b6c9f83b002a41d1433002309) from the original on 1 September 2006. . Retrieved 2006-09-05.

External links

- apt-get Manual page (http://linux.die.net/man/8/apt-get)
- APT HOWTO (http://www.debian.org/doc/manuals/apt-howto/)
- apt-cache (http://linux.die.net/man/8/apt-cache)
- apt.conf (http://linux.die.net/man/5/apt.conf)
- apt_preferences (http://linux.die.net/man/5/apt_preferences)
- Python-apt (https://launchpad.net/python-apt)
- Using apt offline (http://www.fifi.org/doc/apt/offline.html/index.html)
- apt-zip (http://wiki.debian.org/AptZip):
  - Offline functionality (http://lists.debian.org/deity/2006/03/msg00090.html).
- apt-offline (http://apt-offline.alioth.debian.org)
- Adding source lists using apt-setup (http://www.debianadmin.com/adding-source-lists-using-apt-setup.html)
- apt-cdrom (http://linux.die.net/man/8/apt-cdrom)
- apturl (https://launchpad.net/apturl/)

dpkg

dpkg

Original author(s)	Ian Murdock
Developer(s)	The dpkg team
Stable release	1.15.8.11
Preview release	1.16.2
Programming language used	C, C++, Perl[1]
Operating system	Unix-like, POSIX
Type	Package management system
License	GNU General Public License
Website	www.debian.org/doc/FAQ/ch-pkgtools.en.html [2]

dpkg is the software at the base of the Debian package management system. `dpkg` is used to install, remove, and provide information about .deb packages.

`dpkg` itself is a low level tool; higher level tools, such as APT, are used to fetch packages from remote locations or deal with complex package relations. Tools like aptitude or synaptic are more commonly used than `dpkg` on its own, as they have a more sophisticated way of dealing with package relationships and a friendlier interface.

The Debian package "dpkg" provides the `dpkg` program, as well as several other programs necessary for run-time functioning of the packaging system, including `dpkg-statoverride`, `dpkg-divert` and `update-alternatives`. It also includes the programs such as start-stop-daemon and install-info, and the latter is kept mostly due to backwards compatibility (it is developed and distributed separately nowadays). The Debian package "dpkg-dev" includes the numerous build tools described below.

History

dpkg was originally created by Matt Welsh, Carl Streeter and Ian Murdock, first as a Perl program, and then later the main part was rewritten in C by Ian Jackson in 1993. The name *dpkg* was originally the short for "Debian package", but the meaning of that phrase has evolved significantly, as dpkg the software is orthogonal to the deb package format as well as the Debian Policy which defines how Debian packages behave.

Example use

To install a .deb package:

```
dpkg -i debFileName
```

where *debFileName* is the name of the Debian software package.

The list of installed packages can be obtained with:

```
dpkg -l [optional pattern]
```

To remove an installed package:

```
dpkg -r packagename
```

Development tools

dpkg-dev contains a series of development tools required to unpack, build and upload Debian source packages. These include:

- **dpkg-source** packs and unpacks the source files of a Debian package.
- **dpkg-gencontrol** reads the information from an unpacked Debian tree source and generates a binary package control package, creating an entry for this in Debian/files.
- **dpkg-shlibdeps** calculates the dependencies of runs with respect to libraries.
- **dpkg-genchanges** reads the information from an unpacked Debian tree source that once constructed creates a control file (.changes).
- **dpkg-buildpackage** is a control script that can be used to construct the package automatically.
- **dpkg-distaddfile** adds a file input to debian/files.
- **dpkg-parsechangelog** reads the changes file (changelog) of an unpacked Debian tree source and creates a conveniently prepared output with the information for those changes.

References

[1] "dpkg on alioth" (http://alioth.debian.org/projects/dpkg). Alioth (Debian). 2012.04.11. . Retrieved 2012.03.20.
[2] http://www.debian.org/doc/FAQ/ch-pkgtools.en.html

External links

- Debian's dpkg package (http://wiki.debian.org/Teams/Dpkg)
- Debian dpkg mailing list (http://lists.debian.org/debian-dpkg/)
- dpkg(8) (http://manpages.debian.net/cgi-bin/man.cgi?query=dpkg&apropos=0&sektion=0&manpath=Debian+Sid&format=html&locale=en) manual page
- General Origin handling (http://www.hadrons.org/~guillem/debian/docs/origin.proposal).

Windows Installer

Windows Installer

Help screen of Windows Installer 5.0 running on Windows 7	
Original author(s)	Microsoft
Stable release	5.0
Operating system	Microsoft Windows
Type	Installer
License	Proprietary

Windows Installer Package[1]

Filename extension	`.msi, .msp`
Internet media type	`application/x-ole-storage, text/mspg-legacyinfo`
Developed by	Microsoft
Extended from	Cabinet (file format)
Open format?	No

The **Windows Installer** (previously known as **Microsoft Installer**[2]) is a software component used for the installation, maintenance, and removal of software on modern Microsoft Windows systems. The installation information, and often the files themselves, are packaged in *installation packages*, loosely relational databases structured as COM Structured Storages and commonly known as "MSI files", from their default file extension. Windows Installer contains significant changes from its predecessor, Setup API. New features include a GUI framework and automatic generation of the uninstallation sequence. Windows Installer is positioned as an alternative to stand-alone executable installer frameworks such as older versions of InstallShield and NSIS.

Microsoft encourages third parties to use Windows Installer as the basis for installation frameworks, so that they synchronize correctly with other installers and keep the internal database of installed products consistent. Important features such as *rollback* and *versioning* depend on a consistent internal database for reliable operation. Furthermore, Windows Installer facilitates the principle of least privilege by performing software installations by proxy for unprivileged users.

Logical structure of packages

A package describes the installation of one or more full *products* (Windows Installer does not handle dependencies between products) and is universally identified by a GUID (the PackageCode property). A product is made up of *components*, grouped into *features*.

Products

A single, installed, working program (or set of programs) is a product. A product is identified by a unique GUID (the ProductCode property) providing an authoritative identity throughout the world. The GUID, in combination with the version number (ProductVersion property) allow for release management of the software's files and registry keys. A product is not the same as a package: a package is identified by a unique GUID (stored in the Summary Information Stream) and provides identity and release management for all the contents of the .MSI file. Release management of a

Product pertains only to changes in the files and registry keys that make up a software application. Release management of a package includes the package logic and other meta data that relates to how the package executes when running. For example, changing an EXE in the software application may require the ProductCode and/or ProductVersion to be changed for release management of the software application. Only adding a launch condition (with the software application remaining exactly the same as the previous version) would still require the PackageCode to change for release management of the .MSI file itself.

Features

A *feature* is a hierarchical group of components—a feature can contain any number of components and other features (a feature contained in another feature is called a "subfeature"). Many software packages only involve one feature. More complex installation programs usually display a "custom setup" dialog box at run time, from which the end user can select which features to install or remove.

The package author defines the product features. A word-processing program, for example, might provide features for the main program executable, the program's help files, and optional spelling checker and stationery modules.

Components

A *component* is the basic unit of a product—each component is treated by Windows Installer as a unit: the install developer cannot, for example, use a condition to specify to install just part of a component. Components can contain files, directories, COM components, registry keys, shortcuts, and other data. The end user does not directly interact with components.

Components are identified globally by GUIDs, thus the same component can be shared among several features of the same package or multiple packages, ideally through the use of Merge Modules (although, for this to work correctly, different components should not share any sub-components).

Key paths

A *key path* is a specific file, registry key, or ODBC data source that the package author specifies as critical for a given component. Because a file is the most common type of key path, the term *key file* is commonly used. A component can contain at most one key path; if a component has no explicit key path, the component's destination directory is taken to be the key path. When an MSI-based application is launched, Windows Installer checks the existence of these critical files or registry keys (that is, the key paths). If there is a mismatch between the current system state and the value specified in the MSI package (e.g., a key file is missing), then the related feature is re-installed. This process is also known as *self-healing* or *self-repair*. No two components should use the same key path.

Setup phases

User interface

The user interface phase typically queries the target system and displays an installation wizard and enables the user to change various options that will affect the installation.

However, the user interface sequence should not make any changes to the system. Three reasons for this are as follows.

1. A user can install an MSI package in quiet mode, bypassing this phase entirely, by running the msiexec.exe command-line utility with the /qn (or /qb or /qr) option and specifying on the command line all the information that the wizard would normally gather. Therefore, any actions that occur in the user interface sequence will not be performed during a silent installation.

2. Similarly, clicking the *Remove* button in the Add or Remove Programs panel runs a product's uninstaller with a basic user interface, again with the result that any actions that occur in the user interface sequence will not be performed.
3. Actions that make system changes should not be scheduled in the user interface sequence as the user interface sequence runs with user privileges, and not with elevated privileges, as described in the following section.

Actions in the user interface sequence of a normal installation are defined in the InstallUISequence table. Similarly, there is an AdminUISequence in which can be placed dialog boxes and actions to display and perform from within an administrative installation wizard.

Execute

When the user clicks the *Finish* or *Install* button in a typical MSI installation wizard, installation proceeds to the Execute phase, in which software components are actually installed. The Execute phase makes system changes, but does not display any user-interface elements.

Execute phase happens in two steps:

Immediate mode. In this phase, Windows Installer receives instructions, either from a user or an application, to install or uninstall features of a product. The requests cause the execution of *sequences* of *actions*, which query the installation database to build an internal *script* describing the execution phase in detail.

Deferred mode. In this phase, the script built in immediate mode is executed in the context of the privileged Windows Installer service (specifically, the LocalSystem account). The script must be executed by a privileged account because of the heterogeneity of the scenarios in which a setup operation is initiated—for example, elevated privileges are necessary to serve on-demand installation requests from non-privileged users. (In order to run with elevated privileges, however, the package must be deployed by a local administrator or advertised by a system administrator using Group Policy.)

Execute sequence actions for a normal installation are stored in the InstallExecuteSequence table. An MSI database can also contain AdminExecuteSequence and AdvtExecuteSequence tables to define actions to perform for administrative and advertised installations.

Rollback

All installation operations are transactional.[3] For each operation that Windows Installer performs, it generates an equivalent undo operation that would undo the change made to the system. In case any script action fails during deferred execution, or the operation is cancelled by the user, all the actions performed until that point are rolled back, restoring the system to its original state. Standard Windows Installer actions automatically write information into a rollback script; package authors who create custom actions that change the target system should also create corresponding rollback actions (as well as uninstallation actions and uninstallation-rollback actions). As a design feature, if applied correctly this mechanism will also rollback a failed uninstall of an application to a good working state.

Other features

Advertisement

Windows Installer can *advertise* a product rather than actually installing it.[4] The product will appear installed to the user, but it will not actually be installed until it is run for the first time by triggering an entry point (by means of a Start menu shortcut, by opening a document that the product is configured to handle, or by invoking an advertised COM class). A package can be advertised by an administrator using Group Policy or other deployment mechanism, or by running the msiexec executable with the /jm (for per-machine advertisement) or /ju (for per-user advertisement) switch. It should also be noted that some MSI packages authored in Installshield may prevent the use of these and other Native MSI features.

The user must have administrator privileges to complete the advertised installation; in most workplaces, end users are not administrators and this method of distribution will fail. Microsoft created a workaround via Group Policies to "Elevate user privileges" during MSI installations. This is often seen by system administrators as compromising security since any MSI would automatically gain administrator privileges.

Installation on demand

Similar to advertisement, it consists in the installation of *features* as soon as the user tries to use them.[5]

Administrative installation

An administrative installation creates an uncompressed source image for a product, typically to be used for installing or running an application from a network location.[6] An administrative installation is not a typical installation, in that it does not create any shortcuts, register COM servers, create an Add or Remove Programs entry, and so on. Often an administrative installation enables a user to install the product in such a way that its features run from the uncompressed installation source.

Administrative installations are also useful when creating a Windows Installer patch, which requires uncompressed images of the earlier and current versions of a product in order to compute binary file differences. An administrative installation is performed by running the msiexec executable with the /a switch.

Custom actions

The developer of an installer package may write code to serve their own purpose, delivered in a DLL. This can be executed during the installation sequences, including when the user clicks a button in the user interface, or during the InstallExecuteSequence. Custom Actions typically validate product license keys, or initialise more complex services. Developers should normally provide inverse custom actions for use during uninstallation...

Msiexec provides a way to break after loading a specified custom action DLL but before invoking the action.[7]

Merge modules and nested executables

A Windows Installer package may contain another package to be installed at the same time. These are ideally provided as a .msm file component, but may also be a separate executable program which will be unpacked from the installer package during the InstallExecuteSequence and can be run immediately. The file can then optionally be deleted before the end of the InstallExecuteSequence, and so is ideal for using with older installers.

Miscellaneous

Windows Installer allows applications to run directly from a network share, without the need for a local copy (*run from source*); it can repair broken installations by restoring damaged or deleted files, registry entries and application shortcuts; it supports per-user installation of applications; it can resolve component identifiers into paths, allowing applications to avoid hard-coded file paths; and it natively supports *patches* (.msp files made out of patch creation properties) and other customizations of packages through manipulations (*transforms* or .mst files) of a package's relational database. Version 2.0 onwards, it supports digital signatures and version 3.0 onwards, delta compression for patches.

It is also unique among installation software frameworks for Windows in that it is highly transparent. The full API and all command-line options are documented; packages are freely viewable and editable, both with free tools and programmatically (as opposed to the proprietary and even weakly encrypted packages of InstallShield); and the format for file archives is the well documented cabinet file format.

File extraction only

Individual files can be extracted from a .MSI file without carrying out a full installation, either by using the Installer with an appropriate command line[8] [9] or by certain third-party software such as 7-Zip. **NOTE:** Not all files can be extracted using these methods such as different version of a file getting installed based on option will fail to be properly extracted

Windows Vista

Windows Installer 4.0, which was shipped with Windows Vista, incorporates new capabilities to take advantage of Vista's User Account Control architecture. MSI packages can be marked as not requiring elevated privileges to install, thus allowing a package to install without prompting the user for Administrator credentials (provided that the installation does not write to any areas that a regular user does not have access to, including Program Files). Windows Installer also works in conjunction with the Restart Manager; when installing or updating an application or system component with "full" user interface mode, the user will be displayed a list of affected applications that can be shut down, and then restarted after files have been updated. Installer actions running in silent mode perform these application restarts automatically. System services and tray applications can also be restarted in this manner.

Developing installer packages

Creating an installer package for a new application is non-trivial. It is necessary to specify which files must be installed, to where, with what registry keys. Any non-standard operations can be done using Custom Actions, which are typically developed in DLLs. There are a number of commercial and freeware products to assist in creating MSI packages, including Visual Studio, InstallShield and WiX. To varying degrees, the user interface and behavior in less common situations such as unattended installation, may be configured. Once prepared, an installer package is "compiled" by reading the instructions and files from the developer's local machine, and creating the .msi file.

The user interface (dialog boxes) presented at the start of installation can be changed or configured by the setup engineer developing a new installer. There is a limited language of buttons, textfields and labels which can be arranged in a sequence of dialogue boxes. An installer package should be capable of running without any UI, for what is called "unattended installation".

ICE validation

Microsoft provides a set of Internal Consistency Evaluators, or ICEs, that can be used to detect potential problems with an MSI database.[10] The ICE rules are combined into CUB files, which are stripped-down MSI files containing custom actions that test the target MSI database's contents for validation warnings and errors. ICE validation can be performed with the Platform SDK tools Orca and msival2, or with validation tools that ship with the various authoring environments.

For example, some of the ICE rules are:

- ICE09: Validates that any component destined for the System folder is marked as being permanent.
- ICE24: Validates that the product code, product version, and product language have appropriate formats.
- ICE33: Validates that the Registry table is not used for data better suited for another table (Class, Extension, Verb, and so on).

Addressing ICE validation warnings and errors is an important step in the release process.

Versions

Version	Included with[11]	Also available for
1.0	Office 2000	
1.1	Windows 2000 RTM, SP1, SP2	Windows 95/98 Windows NT 4.0 SP6
1.2	Windows Me	
2.0	Windows XP RTM, SP1 Windows 2000 SP3, SP4 Windows Server 2003 RTM	Windows 95/98/Me Windows NT 4.0 SP6 Windows 2000 RTM, SP1, SP2
3.0	Windows XP SP2	Windows 2000 SP3, SP4 Windows XP RTM, SP1 Windows Server 2003 RTM
3.1	Windows XP SP3 Windows Server 2003 SP1, SP2 Windows XP Professional x64 Edition RTM, SP2	Windows 2000 SP3, SP4 Windows XP RTM, SP1, SP2 Windows Server 2003 RTM
4.0	Windows Vista RTM, SP1 Windows Server 2008 RTM	
4.5[12]	Windows Vista SP2 Windows Server 2008 SP2	Windows XP SP2, SP3 Windows Server 2003 SP1, SP2 Windows XP Professional x64 Edition RTM, SP2 Windows Vista RTM, SP1 Windows Server 2008 RTM[13]
5.0	Windows 7 RTM Windows Server 2008 R2 RTM	

Version may be checked by running `msiexec.exe`. (See screenshot at the top.)

Tools

Name	Publisher	Description	License	Website
Advanced Installer	Caphyon Ltd.	An installer that features one freeware edition and four commercial editions. Can create, validate and edit .msi packages.	Freemium	www.advancedinstaller.com [14]
InstallAware	InstallAware Software	A commercial setup creator that comes in four commercial editions	Shareware	[www.installaware.com www.installaware.com]
InstallShield	Flexera Software	InstallShield is a software tool for creating installers or software packages primarily used for installing software for Microsoft Windows desktop and server platforms.	Shareware	[www.installshield.com www.installshield.com]
InstEdit	Google Code	InstEd is a free .msi editor. Paid version with more features available also.	Freeware	[www.instedit.com www.instedit.com]
IsWiX	CodePlex	Industrial Strength Windows Installer XML (IsWiX) is a document editor based on the Fireworks Application Framework. IsWiX enables non-setup developers to collaborate with setup developers using WiX projects.	Microsoft Public License	iswix.codeplex.com [15]
Microsoft Visual Studio	Microsoft	Microsoft Visual Studio is capable of building Windows Installer Deployment projects that can create installer packages.[16]	Shareware	www.microsoft.com/vstudio [17]

MSI Studio	ScriptLogic	An .msi editor aim towards system administrators who wish to repackage installation packages or optimize the installation for their own environment/needs.	Shareware	www.scriptlogic.com [18]
WiX	CodePlex	WiX (Windows Installer XML) is a free and open-source set of tools that helps build a Windows Installer packages from an XML document. It can be either used from command-line or integrated into Microsoft Visual Studio. SharpDevelop, a free and open-source alternative to Visual Studio has adopted WiX.[19]	Common Public License	[wix.codeplex.com wix.codeplex.com]
7-Zip	SourceForge	7-Zip is an open source file archiver utility, and can extract the contents of MSI files.[20]	GNU Lesser General Public License	[www.7-zip.org www.7-zip.org]

References

[1] http://filext.com/file-extension/msi
[2] when setup isn't just xcopy (http://blogs.msdn.com/robmen/archive/2003/10/11/56487.aspx)
[3] Rollback Installation (http://msdn.microsoft.com/en-us/library/aa371370(VS.85).aspx)
[4] Advertisement (http://msdn.microsoft.com/en-us/library/aa367548(VS.85).aspx)
[5] Installation-On-Demand (http://msdn.microsoft.com/en-us/library/aa369293(VS.85).aspx)
[6] Administrative Installation (http://msdn.microsoft.com/en-us/library/aa367541(VS.85).aspx)
[7] http://msdn.microsoft.com/en-us/library/aa368264%28VS.85%29.aspx
[8] Extracting files from a .MSI file from the command line (http://thebackroomtech.com/2007/08/23/howto-extract-files-from-a-msi-file-using-the-windows-command-line/)
[9] How to extract MSI files from the command line/prompt in Windows XP/Vista/7 (http://cyberst0rm.blogspot.com/2011/07/how-to-extract-contentdata-from-msi.html)
[10] Internal Consistency Evaluators - ICEs (http://msdn.microsoft.com/en-us/library/aa369554(VS.85).aspx)
[11] Released Versions of Windows Installer (http://msdn.microsoft.com/en-us/library/aa371185(VS.85).aspx)
[12] What's New in Windows Installer 4.5 (http://msdn.microsoft.com/en-us/library/bb736325(VS.85).aspx)
[13] (http://msdn.microsoft.com/en-us/library/aa371185(VS.85).aspx)
[14] http://www.advancedinstaller.com/
[15] http://iswix.codeplex.com/
[16] "Windows Installer Deployment" (http://msdn.microsoft.com/en-us/library/2kt85ked.aspx). *Microsoft Developers Network (MSDN) Library*. Microsoft Corporation. August 2008. . Retrieved 1 October 2009. "Windows Installer deployment enables you to create installer packages to be distributed to users; the user runs the setup file and steps through a wizard to install the application."
[17] http://www.microsoft.com/vstudio
[18] http://www.scriptlogic.com/
[19] "SharpDevelop Features" (http://www.icsharpcode.net/OpenSource/SD/Features.aspx). IC#Code. . Retrieved 31 March 2012.
[20] "Windows Installer Tools & Tips" (http://installsite.org/pages/en/msi/tips.htm#7zip). InstallSite Stefan Krueger. . Retrieved 19 November 2009.

External links

- Windows installer start page (http://msdn.microsoft.com/en-us/library/cc185688(VS.85).aspx)
- Windows Installer Team Blog (http://blogs.msdn.com/windows_installer_team/) at MSDN.
  - Windows Installer 4.5 expected around Q2-2008 (blog) (http://blogs.msdn.com/windows_installer_team/archive/2007/12/10/release-date-for-msi-4-5-now-looks-like-q2-ish.aspx)
- InstallSite.org (http://installsite.org/) "Resources for Setup Developers": This site publishes a useful released MSI version matrix (http://installsite.org/pages/en/msifaq/a/1001.htm#matrix).
- Windows Installer 4.5 Software Development Kit (http://download.microsoft.com/download/7/c/4/7c426dfc-46e2-4ded-bab4-3b33600ad7d1/msi45sdk.msi) standalone version, XP & Vista
- Windows Installer 5.0 Software Development Kit (http://blogs.msdn.com/windows_installer_team/archive/2009/08/12/latest-sdk-for-windows-installer-5-0.aspx) standalone for 5.0 not available, only the version

integrated with Windows 7 SDK
- MSI frequently asked questions (http://sourceforge.net/projects/mayadeploy/files/FAQ/)

Windows Update

Windows Update A component of Microsoft Windows	
Windows Update running on Windows 7	
Details	
Type	Network service
Included with	Windows 98 and later
Description	Windows Update
Related components	
BITS, Windows Installer, Internet Explorer, Windows Genuine Advantage	

Windows Update is a service provided by Microsoft that provides updates for the Microsoft Windows operating system and its installed components, including Internet Explorer. An optional feature disables access to Windows Update, enabling instead access to **Microsoft Update**, an expanded version of the service which provides updates not just for the operating system and Internet Explorer, but also for other Microsoft software running under Windows, such as Microsoft Office, Windows Live applications, and Microsoft Expression Studio. Updates are normally provided over an Internet connection, although there is provision for updates to be installed on computers without an Internet connection.

There are different kinds of updates. *Security updates* or *critical updates* protect against vulnerabilities to malware and security exploits. Other updates correct errors that aren't related to security, or enhance functionality.

Security updates are routinely provided on the second Tuesday of each month, Patch Tuesday, but can be provided whenever a new update is urgently required to prevent a newly discovered or prevalent exploit targeting Windows users. Windows Update can be configured to install critical updates automatically so long as the computer is connected to the Internet, without the user needing to install them manually, or even be aware that an update is required.

Windows Vista, Windows Server 2008, Windows 7, Windows Server 2008 R2, Windows 8, and Windows Server 2012 provide a Control Panel to configure update settings and check for updates. The Windows Update Control Panel is also the means to download Windows Ultimate Extras, optional software for Windows Vista Ultimate Edition. For previous versions of Microsoft Windows, updates can be downloaded from the Windows Update website, using Internet Explorer.

History

Windows Update web site

Windows Update was introduced as an Internet web site with the launch of Windows 95. A link to Windows Update on the Start Menu gave access to additional downloads for the operating system. At the time of Windows 98's release Windows Update offered additional desktop themes, games, device driver updates, and optional components such as NetMeeting.[1] Windows 95 and Windows NT 4 were retroactively given the ability to access the Windows Update website, and download updates designed for those operating systems, starting with the release of versions of Internet

Explorer 4 for those operating systems. The initial focus of Windows Update was on free add-ons and new technologies for Windows; security fixes for Outlook Express, Internet Explorer and other applications appeared later, as did access to beta versions of upcoming Microsoft software, most notably Internet Explorer 5. Fixes to Windows 98 to resolve the Year 2000 problem were distributed using Windows Update in December 1998. Microsoft attributed the sales success of Windows 98 in part to Windows Update.[2]

Windows Update requires Internet Explorer or a third-party web browser that uses Microsoft's MSHTML layout engine, as it must support the use of an ActiveX control to house the software that is executed on the user's computer. While details have changed from version to version, it has always scanned the computer to find what operating system components and software are installed, and compared the versions of those components with the latest available versions. The ActiveX component then interfaces with Windows Installer to install or update those components, and to report the success or failure of those installations back to Microsoft's servers.

The first version of the Windows Update web site (usually referred to as "v3") did not require any personally-identifiable information to be sent to Microsoft. In order for the v3 ActiveX control to determine what updates were needed, the entire list of available software on Windows Update was downloaded to the user's computer when they visited the Windows Update web site. As the number of updates offered by Windows Update grew, this resulted in performance concerns. Arie Slob, writing for the Windows-help.net newsletter in March 2003, noted that the size of the update list had exceeded 400KB, which caused delays of more than a minute for dial-up users.[3]

Windows Update v4, released in conjunction with Windows XP in 2001, changed this by having the ActiveX control submit a list of the hardware components to Microsoft's servers, which then returns a list of only those device drivers available for that machine. It also narrowed down the list of available updates for the operating system and related components by sending details of what operating system version, service pack, and locale are installed. German technology web site tecchannel.de published an analysis of the Windows Update communication protocol in February 2003, which received wide attention on technology web sites. The report, which was the first to contain extensive details of how the Windows Update communication protocol worked, also discovered that the make and model of the computer, the amount of free disk space, and the Windows product key, were sent.[4]

Critical Update Notification Tool/Utility

Shortly after the release of Windows 98, Microsoft released a **Critical Update Notification Tool** (later called **Critical Update Notification Utility**) through Windows Update, which installed a background tool on the user's computer that checked the Windows Update web site on a regular schedule for new updates that have been marked as "Critical". By default, this check occurred every five minutes, and when Internet Explorer was started, though the user could configure the next check to occur only at certain times of the day or on certain days of the week. The check was performed by querying the server for a file, "cucif.cab", which contains a list of all the critical updates released for the user's operating system. The Critical Update Notification Tool then compared this list with the list of installed updates on the user's machine, and displayed a message to the user informing them of new critical updates if they were available. Once the check executed, any custom schedule defined by the user was reverted to the default; Microsoft stated that this was by design in order to ensure that users received notification of critical updates in a timely manner.[5]

An analysis done by security researcher H D Moore in early 1999 was critical of this approach, describing it as "horribly inefficient" and susceptible to attacks. In a posting to BugTraq, he explained that, "every single Windows 98 computer that wishes to get an update has to rely on a single host for the security. If that one server got compromised one day, or an attacker cracks the MS DNS server again, there could be millions of users installing trojans every hour. The scope of this attack is big enough to attract crackers who actually know what they are doing..."[6]

The Critical Update Notification tool continued to be promoted by Microsoft through 1999 and the first half of 2000. Initial releases of Windows 2000 shipped with the tool, but Windows 95 and Windows NT 4.0 were not supported. It was superseded by **Automatic Updates** in Windows Me and Windows 2000 SP4.

Automatic Updates

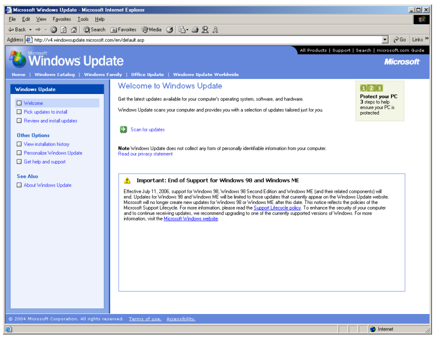

Windows Update v4 in Windows Me.

With the release of Windows Me in 2000, Microsoft introduced **Automatic Updates** as a replacement for the Critical Update Notification tool. Unlike its predecessor, Automatic Updates includes the ability to download and install updates without using a web browser. Instead of the five minute schedule used by its predecessor, the Automatic Updates client checks the Windows Update servers once a day. The user is given the option to download available updates then prompt the user to install them, or to notify the user prior to downloading any available updates. After Windows Me is installed, the user is prompted via a notification balloon to configure the Automatic Updates client.

The Windows Update web site itself was significantly updated to match the visual style of Windows XP.

Windows XP and Windows 2000 Service Pack 3 include Background Intelligent Transfer Service, a protocol for transferring files in the background without user interaction. As a system component, it is capable of monitoring the user's Internet usage, and throttling its own bandwidth usage in order to prioritize user-initiated activities. The Automatic Updates client for these operating systems was updated to use this system service.

Microsoft Update

At the February 2005 RSA Conference, Microsoft announced the first beta of **Microsoft Update**, an optional replacement for Windows Update that provides security patches, service packs and other updates for both Windows and other Microsoft software.[7] The initial release in June 2005 provided support for Microsoft Office 2003, Exchange 2003, and SQL Server 2000, running on Windows 2000, XP, and Server 2003. Over time, the list has expanded to include other Microsoft products, such as Windows Live, Windows Defender, Visual Studio, runtimes and redistributables, Zune Software, Virtual PC and Virtual Server, CAPICOM, Microsoft Lync, and other server products. It also offers Silverlight and Windows Media Player as optional downloads if applicable to the operating system. A persistent bug in Microsoft Update affecting XP computers with limited internal memory is that it allows the update programs wuauclt.exe and svchost.exe to claim 100% of the computers' memory for extended periods of time (up to hours) making affected computers unusable.

MS Office Update

Microsoft Office Update was a free online service that allowed users to detect and install updates for certain Microsoft Office products. This update service supported Office 2000, Office XP, Office 2003, and Office 2007. On 1 August 2009, Microsoft decommissioned the service.[8] Users are now required to use Microsoft Update. However, as Microsoft Update does not work with Office 2000, Office 2000 users no longer have any method of automatically detecting and installing updates. This is not a limitation for existing installations of Office 2000, because the product is no longer supported and so no new updates are being produced. However, it is a serious limitation for anyone re-installing MS Office 2000.

Windows Vista, Windows Server 2008, and Windows 7

In Windows Vista, Windows Server 2008 and later, the web site is no longer used to provide a user interface for selecting and downloading updates. In its place, the Automatic Updates control panel has been expanded to provide similar functionality. Support for Microsoft Update is also built into the operating system, but is turned off by default. The revised Windows Update can also be set to automatically download and install both *Important* and *Recommended* updates. In prior versions of Windows, such updates were only available through the Windows Update web site.

In versions of Windows prior to Vista, updates requiring a reboot would pop up a dialog box every number of specified minutes requesting that users reboot their machines.[9] This dialog box was changed to allow the user to select a longer period of time (up to 4 hours) before being prompted again. The revised dialog box also displays under other applications, instead of on top of them.

In Windows 7 and Vista[10] once automatic updates have finished, the computer will be shut down after a countdown, sometimes causing the countdown to finish and the system to reboot while the user is in the middle of using the computer (or away from the computer and not wanting it to reboot for various reasons), possibly losing data, gameplay advancement, etc.

Windows Update makes use of Transactional NTFS, a file system feature introduced with Windows Vista, when performing updates to Windows system files. This feature helps Windows recover cleanly in the event of an unexpected shut-down during an update, as the transactioning system will ensure that changes are committed to the file system (in particular, to the persistent files of the registry) in an atomic fashion.[11]

Statistics

At the beginning of 2005, Windows Update was being accessed by about 150 million people,[12] with about 112 million of those using Automatic Updates.[13]

As of 2008, Windows Update had about 500 million clients, processed about 350 million unique scans per day, and maintained an average of 1.5 million simultaneous connections to client machines. On Patch Tuesday, the day Microsoft typically releases new software updates, outbound traffic could exceed 500 gigabits per second.[14] Approximately 90% of all clients used automatic updates to initiate software updates, with the remaining 10% using the Windows Update web site. The web site is built using ASP.NET, and processes an average of 90,000 page requests per second.

References

[1] Gartner, John (August 24, 1995). "Taking Windows 98 For A Test-Drive" (http://www.techweb.com/wire/story/win98/TWB19980625S0017). TechWeb. . Retrieved 2008-07-29.

[2] "Strong Holiday Sales Make Windows 98 Best-Selling Software of 1998" (http://www.microsoft.com/presspass/press/1999/feb99/holislpr.mspx) (Press release). Microsoft. February 9, 1999. . Retrieved 2008-07-29.

[3] Slob, Arie (March 22, 2003). "Windows Update is Spying on You!" (http://www.windowsnewsletter.com/html-archive/2003/22mar2003.html). windows-help.net. . Retrieved 2008-07-30.

[4] Leyden, John (February 28, 2003). "Windows Update keeps tabs on all system software" (http://www.theregister.co.uk/2003/02/28/windows_update_keeps_tabs/). The Register. . Retrieved 2008-07-30.

[5] "Description of the Windows Critical Update Notification Utility (MSKB244420)" (http://support.microsoft.com/kb/q224420/). *Knowledge Base*. Microsoft. December 5, 2007. . Retrieved 2008-07-29.

[6] Moore, H D (January 29, 1999). "How the MS Critical Update Notification works..." (http://seclists.org/bugtraq/1999/Jan/0400.html). . Retrieved 2008-07-29.

[7] "Microsoft Update Site Launched" (http://www.helpwithwindows.com/microsoft-update.html). helpwithwindows.com. June 10, 2005. . Retrieved 2008-07-30.

[8] "Microsoft Decommissioning the Office Update Service - FAQ" (http://office.microsoft.com/en-us/downloads/FX010402221033.aspx). Office.microsoft.com. . Retrieved 2010-04-23.

[9] Atwood, Jeff (May 13, 2005). "XP Automatic Update Nagging" (http://www.codinghorror.com/blog/archives/000294.html). Coding Horror: .NET And Human Factors. . Retrieved 2006-09-22.

[10] http://www.microsoft.com/windows/downloads/windowsupdate/faq.mspx
[11] "NTFS Beta Chat Transcript" (https://blogs.technet.com/filecab/articles/457811.aspx). *The Filing Cabinet*. TechNet Blogs. July 12, 2006. . Retrieved 22 September 2006.
[12] "RSA Conference 2005: "Security: Raising the Bar" (speech transcript)" (http://www.microsoft.com/presspass/exec/billg/speeches/2005/02-15RSA05.aspx). *PressPass*. Microsoft. February 15, 2005. . Retrieved 2008-07-30.
[13] "Microsoft Announces Availability of New Solutions to Help Protect Customers Against Spyware and Viruses" (http://www.microsoft.com/presspass/press/2005/jan05/01-06newsolutionspr.mspx). *PressPass*. Microsoft. January 6, 2005. . Retrieved 2008-07-30.
[14] "Introducing the Microsoft.com Engineering Operations Team" (http://technet.microsoft.com/en-us/library/cc627316.aspx). *Microsoft TechNet*. Microsoft. 2008. . Retrieved 31 May 2012.

External links

- Microsoft Update service (http://update.microsoft.com/microsoftupdate)
- Latest Computer Security Updates: Summary - Microsoft Security (http://www.microsoft.com/protect/computer/updates/bulletins/default.mspx) / RSS (http://www.microsoft.com/protect/rss/default.mspx)
- Microsoft Technical Security Notifications (http://www.microsoft.com/technet/security/bulletin/notify.mspx)
- Windows Update Explained (http://download.microsoft.com/download/a/9/4/a94af289-a798-4143-a3f8-77004f7c2fd3/Windows Update Explained.docx)

Windows Server Update Services

Windows Server Update Services (WSUS)

WSUS can display precise information about which updates each client needs	
Developer(s)	Microsoft Corporation
Stable release	3.0 SP2 / August 25, 2009[]
Operating system	Windows Server 2008 R2 Windows Server 2008 SP1 or later Windows Server 2003 SP2 or later Windows Small Business Server 2011 Windows Small Business Server 2008 Windows Small Business Server 2003[]
Platform	**Hardware:** IA-32 and x86-64 **Software:** .NET Framework 2.0 and Microsoft Management Console
Size	80 MB (almost)[]
Translation available	English[]
Type	Remote administration
License	Freeware
Website	www.microsoft.com/wsus [1]

Windows Server Update Services (**WSUS**), previously known as **Software Update Services** (**SUS**), is a computer program developed by Microsoft Corporation that enables administrators to manage the distribution of updates and hotfixes released for Microsoft products to computers in a corporate environment. WSUS downloads these updates from the Microsoft Update website and then distributes them to computers on a network. WSUS runs on Windows Server and is free to licensed Microsoft customers.

History

The first version of WSUS was known as Software Update Services (SUS).[2] It only delivered hotfixes and patches for Microsoft operating systems. SUS ran on a Windows Server operating system and downloaded updates for the specified versions of Windows from the remote Windows Update site, operated by Microsoft. Clients could then download updates from this internal server, rather than connecting directly to Windows Update.[3] Support for SUS by Microsoft was originally planned to end on December 6, 2006, but based on user feedback, the date was extended to July 10, 2007.[4]

WSUS builds on SUS by expanding the range of software it can update. The WSUS infrastructure allows automatic downloads of updates, hotfixes, service packs, device drivers and feature packs to clients in an organization from a central server(s).

Operation

Windows Server Update Services 2.0 and above comprise a repository of update packages from Microsoft. It allows administrators to approve or decline updates before release, to force updates to install by a given date, and to obtain extensive reports on what updates each machine requires. System administrators can also configure WSUS to approve certain classes of updates automatically (critical updates, security updates, service packs, drivers, etc.). One can also approve updates for "detection" only, allowing an administrator to see what machines will require a given update without also installing that update.

Administrators can use WSUS with Group Policy for client-side configuration of the Automatic Updates client, ensuring that end-users can't disable or circumvent corporate update policies. WSUS does not require the use of Active Directory; client configuration can also be applied by local group policy or by modifying the Windows registry.

WSUS uses .NET Framework, Microsoft Management Console and Internet Information Services. WSUS 3.0 uses either SQL Server Express or Windows Internal Database as its database engine. WSUS 2.0 used MSDE. System Center Configuration Manager (SCCM) interoperates with WSUS, and is able to import third party security updates into the product.[5]

Version history

Version	Date	Comments
2.0 Release Candidate	March 22, 2005	
2.0 RTW	June 6, 2005	
2.0 Service Pack 1	May 31, 2006	Adds support for Windows Vista clients, additional client languages, and using Microsoft SQL Server 2005 as a database backend, as well as performance improvements with the web-based user interface
3.0 beta 2	August 14, 2006	MMC based UI and loads of new features
3.0 Release Candidate	February 12, 2007	
3.0 RTW	April 30, 2007	WSUS 3.0 and WSUS Client 3.0 were made available via WSUS on 22 May 2007[6]
3.0 Service Pack 1 Release Candidate	November 1, 2007	
3.0 Service Pack 1 RTW[7]	February 7, 2008	
3.0 Service Pack 2 RTW	August 25, 2009	Included as a role within Server 2008 R2

References

[1] http://www.microsoft.com/wsus
[2] Foust, Mark; Chellis, James; Sheltz, Matthew; Sage London, Suzan (2006). "Chapter 7: Planning Server-Level Security". *MCSE Windows Server 2003 network infrastructure planning and maintenance study guide*. Hoboken, New Jersey: John Wiley and Sons. p. 532. ISBN 978-0-7821-4450-5.
[3] "Software Update Services" (http://technet.microsoft.com/en-us/windowsserver/bb466186.aspx). *Microsoft TechNet*. Microsoft Corporation. . Retrieved 4 May 2011.
[4] Keizer, Gregg (16 Nov 2006). "Microsoft Keeps Software Update Services Alive Until July" (http://www.informationweek.com/news/194400595). *InformationWeek* (UBM TechWeb). . Retrieved 4 May 2011.
[5] "About System Center Updates Publisher" (http://technet.microsoft.com/en-us/library/bb632895.aspx). Microsoft. . Retrieved 11 August 2011.
[6] Harder, Bobbie (22 May 2007). "Updates for WSUS available today" (http://blogs.technet.com/b/wsus/archive/2007/05/22/updates-for-wsus-available-today.aspx). *WSUS Product Team Blog*. Microsoft Corporation. . Retrieved 4 May 2011.
[7] Cole, Cecilia (7 Feb 2008). "WSUS 3.0 SP1 is now RTM" (http://blogs.technet.com/b/wsus/archive/2008/02/07/wsus-3-0-sp1-is-now-rtm.aspx). *WSUS Product Team Blog*. Microsoft Corporation. . Retrieved 4 May 2011.

External links

- Official website (http://technet.microsoft.com/windowsserver/bb332157) on Microsoft TechNet
- Download: Windows Server Update Services 3.0 SP2 (http://www.microsoft.com/download/en/details.aspx?displaylang=en&id=5216)
- WSUS Product Team Blog (http://blogs.technet.com/wsus/default.aspx)
- WSUS Wiki (http://wsus.editme.com/) – contains many detailed documents on WSUS operation, known issues, and troubleshooting
- WSUS.DE (http://www.wsus.de) - German WSUS-Community (Microsoft CLIP Community)
- WuInstall (http://www.wuinstall.com) - Control installation of WSUS updates from command line

WiX

Windows Installer XML toolset

Original author(s)	Rob Mensching
Developer(s)	Microsoft
Stable release	3.5.2519 / Jan 31, 2011
Preview release	3.6 (updated weekly)
Programming language used	C++, C#
Operating system	Windows
Type	Software development tools
License	Common Public License
Website	wixtoolset.org [1]

The **Windows Installer XML** (**WiX**, pronounced "wicks"), is a free software toolset that builds Windows Installer (MSI) packages from an XML document. It supports a command-line environment that developers may integrate into their build processes to build MSI and MSM setup packages. WiX was the first software released by Microsoft under an open-source license called Common Public License.

The WiX distribution includes Votive, a Visual Studio add-in that allows creating and building WiX setup projects using the Visual Studio IDE. Votive supports syntax highlighting and IntelliSense for `.WXS` source files and adds a WiX setup project type to Visual Studio.

Internal structure

WiX is composed of components, whose names derive from a play on words on the "wick(s)" of a candle.

Candle

Candle, the compiler, is used to compile the XML documents to object files that contain symbols and references to symbols.

Light

Light, the linker, takes one or more object files and links the references in the object files to the appropriate symbols in other object files. Light is also responsible for collecting all of the binaries, packaging them appropriately, and generating the final MSI or MSM file.

Lit

Lit, the library tool, is an optional tool that can be used to combine multiple object files into libraries that can be parsed by Light.

Dark

Dark, the decompiler, takes existing MSI and MSM files and generates XML documents that represent the package.

Tallow/Heat

Tallow is a tool for generating WiX filelist code by traversing directory trees. It produces a WiX "fragment" which can be incorporated with other WiX source files at compile time. It is replaced in WiX 3.0 by a more general "harvesting" tool known as **Heat**. There is also an unofficial version of Tallow known as **Mallow**,[2] which adds synchronization capabilities and improved component id generation.

There is also a new tool called **Paraffin**,[3] which provides support for initial creation of a fragment and synchronization.

Pyro

Pyro is the WiX tool to create Patch files (msp) without the Windows Installer SDK.

Burn

As of version 3.5, **Burn**, a prerequisite bootstrapper and installer chainer tool is under development.[4] Planned features include small size, proper elevation support for Vista/7, a very customizable UI and progress indicators, and automatic downloads of required. The old Burn source code was abandoned in July 2009 and reimplemented based on NETFX4 bootstrapper. [5] The new code is released in WiX 3.6 on September 2010. [6] [7]

History

On April 5, 2004, WiX was the first Microsoft project to be released under an externally created Open Source license, the Common Public License. It was also the first Microsoft Shared Source project to be hosted externally (on SourceForge).

Rob Mensching, the original author and lead developer of WiX, works on WiX in his spare time. At the time of release he said, "I did not feel that many people inside Microsoft understood what the Open Source community was really about and I wanted to improve that understanding by providing an example."

As of 2006, several other Microsoft employees from various product divisions of the company work on WiX with Mensching, meeting after business hours once a week to coordinate development efforts and write code. WiX has proven to be so popular with Microsoft development teams that many of Microsoft's software products, such as SQL Server 2005, Office 2007, and Microsoft Codename Oslo are packaged using WiX.

As of 2007, WiX version 2.0 is considered stable and production quality, and is no longer being developed.

As of July 4, 2009, WiX version 3.0 is considered release quality.

WiX version 3.5 was released January 31, 2011.

WiX version 3.6 is the version currently being developed, which concentrates on Burn.

References

[1] http://wixtoolset.org
[2] WiX Tips and Tricks: Part 2 - Tallow, Component GUIDs and Uninstallation (http://osdir.com/ml/windows.devel.wix.user/2005-02/msg00115.html)
[3] John Robbins' Blog : WiX: A Better TALLOW – PARAFFIN (Part 3 of 3) (http://www.wintellect.com/cs/blogs/jrobbins/archive/2007/10/21/wix-a-better-tallow-paraffin.aspx)
[4] Rob Mensching's blog: Let's talk about Burn. (http://robmensching.com/blog/posts/2009/7/14/Lets-talk-about-Burn)
[5] Burn moves to a new foundation (http://robmensching.com/blog/posts/2010/1/15/Burn-moves-to-a-new-foundation)
[6] Change of plans for WiX (http://robmensching.com/blog/posts/2010/4/28/Change-of-plans-for-WiX-v3.5)
[7] The install is ugly and buggy but it's functional (http://robmensching.com/blog/posts/2010/9/6/Burn-baby.-Burn)

- Mallett, Steve. "WiX Project Lead Interviewed On CPL Licensing" (http://developers.slashdot.org/article.pl?sid=04/04/23/2229232). Slashdot.org. Retrieved 2004-04-24.
- Mensching, Rob. "Windows Installer XML (WiX) toolset has released as Open Source on SourceForge.net" (http://blogs.msdn.com/robmen/archive/2004/04/05/107709.aspx). Microsoft Developer Network. Retrieved 2004-04-24.

External links

- Official website (http://wixtoolset.org)
- wix on codeplex.com (http://wix.codeplex.com/)
- Project Summary (http://sourceforge.net/projects/wix/) on SourceForge.net
- Interview with Rob Mensching of Microsoft's WiX Project (http://osdir.com/Article380.phtml)

Automatic Installations

Preboot Execution Environment

The **Preboot eXecution Environment** (**PXE**, also known as **Pre-Execution Environment**; sometimes pronounced "pixie") is an environment to boot computers using a network interface independently of data storage devices (like hard disks) or installed operating systems.

PXE was introduced as part of the Wired for Management framework by Intel and is described in the specification (version 2.1) published by Intel and Systemsoft on September 20, 1999.[1] It makes use of several network protocols like Internet Protocol (IPv4), User Datagram Protocol (UDP), Dynamic Host Configuration Protocol (DHCP) and Trivial File Transfer Protocol (TFTP) and of concepts like Globally Unique Identifier (GUID), Universally Unique Identifier (UUID) and Universal Network Device Interface and extends the firmware of the PXE client (the computer to be bootstrapped via PXE) with a set of predefined Application Programming Interfaces (APIs).

Chain

The firmware on the client tries to locate a PXE redirection service on the network (Proxy DHCP) in order to receive information about available PXE boot servers. After parsing the answer, the firmware will ask an appropriate boot server for the file path of a network bootstrap program (NBP), download it into the computer's random-access memory (RAM) using TFTP, possibly verify it, and finally execute it. If only one NBP is used among all PXE clients it could be specified using BOOTP without any need of a proxy DHCP, but a TFTP boot server is still required.

Availability

PXE was designed to be applicable to many system architectures. The 2.1 version of the specification assigns architecture identifiers to six system types, including IA-64 and DEC Alpha. However, the specification only completely covers IA-32. Intel included PXE in the EFI for IA-64, creating a de-facto standard with the implementation.

Protocol

The PXE protocol is approximately a combination of DHCP and TFTP, albeit with subtle modifications to both. DHCP is used to locate the appropriate boot server or servers, with TFTP used to download the initial bootstrap program and additional files.

To initiate a PXE bootstrap session the PXE firmware broadcasts a DHCPDISCOVER packet extended with PXE-specific options (*extended DHCPDISCOVER*) to port 67/UDP (DHCP server port). The PXE options identify the firmware as capable of PXE, but they will be ignored by standard DHCP servers. If the firmware receives DHCPOFFERs from such servers, it may configure itself by requesting one of the offered configurations.

Proxy DHCP

If a PXE redirection service (Proxy DHCP) receives an *extended DHCPDISCOVER*, it replies with an *extended DHCPOFFER* to the client's port 68/UDP (DHCP client port).

An *extended DHCPOFFER* contains mainly:

- a PXE Discovery Control field to recommend multicasting, broadcasting, or unicasting to contact PXE boot servers
- a list of IP addresses of each available PXE Boot Server Type
- a PXE Boot Menu with each entry representing a PXE Boot Server Type
- a PXE Boot Prompt telling the user to press a certain key to see the boot menu
- a timeout to launch the first boot menu entry if it expires

The Proxy DHCP service may also run on the same host as the standard DHCP service. Since two services cannot share port 67/UDP, the Proxy DHCP runs on port 4011/UDP and expects the *extended DHCPDISCOVER packets* from PXE Clients to be DHCPREQUESTs. The standard DHCP service has to send a special combination of PXE options in its DHCPOFFER, so the PXE client knows to look for a Proxy DHCP on the same host, port 4011/UDP.

Boot server contact

To contact a PXE Boot Server the booting system must have an IP address (perhaps from a DHCP server).

It multicasts or unicasts a *DHCPREQUEST* packet extended with PXE-specific options (*extended DHCPREQUEST*) to port 4011/UDP or broadcasts it to port 67/UDP. This packet contains the PXE Boot Server **type** and the PXE Boot Layer, allowing multiple boot server types to run from one daemon. The *extended DHCPREQUEST* may be a *DHCPINFORM*.

A PXE Boot Server receiving an *extended DHCPREQUEST* configured for the requested **type** and client architecture responds with an *extended DHCPACK* including:

- the complete file path to download the NBP via TFTP.
- PXE Boot Server **type** and PXE Boot Layer it answered
- the multicast TFTP configuration, if MTFTP as described in the PXE specification should be used.

The booting system accepts information from only one *extended DHCPOFFER*.

A 2.1 version PXE Boot Server supports "Boot Integrity Services" ([2]) allowing the Client to verify downloaded NBPs using a checksum file which is downloaded from the same boot server as the NBP.

To get the file path of this *credentials* file another exchange of *extended DHCPREQUEST* and *extended DHCPACK* is required.

Network bootstrap program

After receiving the requested *extended DHCPACK*, the *Network Bootstrap Program* is uploaded into RAM and after it is verified or if verification is not required, the NBP will be executed. It has access to the APIs of the PXE firmware extension (Pre-boot, UDP, TFTP, Universal Network Device Interface (UNDI)). Its functions or tasks are not described in the PXE specification.

Integration

The *PXE Client/Server Protocol* was designed so:

- it can be used in the same network as an existing DHCP environment without interference
- it can be integrated completely into standard DHCP services
- it can be easily extended at the most important points without a call for papers
- every service (DHCP, Proxy DHCP, Boot Server) can be implemented standalone or in any combination of them.

Additionally the PXE firmware extension was designed as an Option ROM for the IA-32 BIOS so you can make a personal computer (PC) PXE-capable by installing a NIC that provides a PXE Option ROM. Note, this procedure also applies to the newer AMD64 processor standard for PC.

The design goal of utilizing existing DHCP and TFTP servers cannot be achieved in a strictly conforming implementation. Some aspects of the PXE protocol require that the DHCP and TFTP servers be modified and communicate. One specific example is using multicast, where DHCP packets provide the multicast group information rather than an opening RFC-2090 multicast TFTP exchange. The impact of this is minimal as the most common PXE client implementation (written by Intel and provided at no cost as a linkable IA32 binary module) interoperates with a combination of isolated DHCP and unicast TFTP servers.

Resources

Specifications, RFCs and other documents about PXE:

- PXE specification [3] - The Preboot Execution Environment specification v2.1 published by Intel & Systemsoft.
- BIS specification [4] - The Boot Integrity Services specification v1.0 published by Intel.
- Remote Boot Protocol Draft [5] - draft of the PXE Client/Server Protocol included in the PXE specification.
- LTSP - Linux Terminal Server Project
- Lan Core [6] - Open Source Thin Client solution.

References

[1] "Preboot Execution Environment (PXE) Specification" (http://download.intel.com/design/archives/wfm/downloads/pxespec.pdf) (PDF). . Retrieved 2009-02-18.

[2] "Boot Integrity Services Application Programming Interface" (http://download.intel.com/design/archives/wfm/downloads/bisspec.pdf) (PDF). . Retrieved 2009-02-18.

[3] ftp://download.intel.com/design/archives/wfm/downloads/pxespec.pdf

[4] ftp://download.intel.com/design/archives/wfm/downloads/bisspec.pdf

[5] http://quimby.gnus.org/internet-drafts/draft-henry-remote-boot-protocol-00.txt

[6] http://lancore.sourceforge.net/

External links

- PXE error codes (http://h18013.www1.hp.com/products/servers/management/rdp/knowledgebase/00000138.html) - A catalogue of PXE error codes
- PXE, aka Pre-Execution Environment (http://www.itstuff.ca/2007/09/pxe-aka-pre-execution-environment-part.html) - How to boot from Network - Part 1
- PXE, aka Pre-Execution Environment (http://www.itstuff.ca/2007/12/pxe-aka-pre-execution-environment-and.html) - How to boot from Network - Part 2
- How To: Windows PXE Install using Serva (http://www.vercot.com/~serva/howto/WindowsPXE1.html)
- How To: Debian PXE Install using Serva (http://www.vercot.com/~serva/howto/DebianPXE1.html)
- How To: Ubuntu PXE Install using Serva (http://www.vercot.com/~serva/howto/UbuntuPXE1.html)
- PXE setup instructions from the centos wiki (http://wiki.centos.org/HowTos/PXE/PXE_Setup)
- Set up a Debian network boot server (http://wiki.debian.org/PXEBootInstall)

Dynamic Host Configuration Protocol

The **Dynamic Host Configuration Protocol** (**DHCP**) is a network protocol that is used to configure network devices so that they can communicate on an IP network. A DHCP client uses the DHCP protocol to acquire configuration information, such as an IP address, a default route and one or more DNS server addresses from a DHCP server. The DHCP client then uses this information to configure its host. Once the configuration process is complete, the host is able to communicate on the internet.

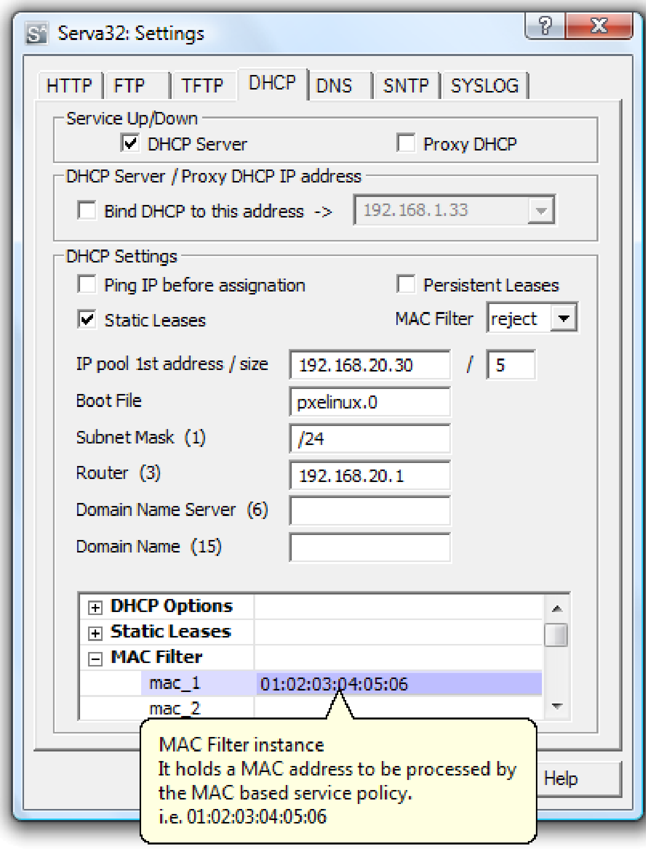

A DHCP Server settings tab

The DHCP server maintains a database of available IP addresses and configuration information. When it receives a request from a client, the DHCP server determines the network to which the DHCP client is connected, and then allocates an IP address or prefix that is appropriate for the client, and sends configuration information appropriate for that client.

Because the DHCP protocol must work correctly even before DHCP clients have been configured, the DHCP server and DHCP client must be connected to the same network link. In larger networks, this is not practical. On such networks, each network link contains one or more DHCP relay agents. These DHCP relay agents receive messages from DHCP clients and forward them to DHCP servers. DHCP servers send responses back to the relay agent, and the relay agent then sends these responses to the DHCP client on the local network link.

DHCP servers typically grant IP addresses to clients only for a limited interval. DHCP clients are responsible for renewing their IP address before that interval has expired, and must stop using the address once the interval has expired, if they have not been able to renew it.

DHCP is used for IPv4 and IPv6. While both versions serve much the same purpose, the details of the protocol for IPv4 and IPv6 are sufficiently different that they may be considered separate protocols.[1]

Hosts that do not use DHCP for address configuration may still use it to obtain other configuration information. Alternatively, IPv6 hosts may use stateless address autoconfiguration. IPv4 hosts may use link-local addressing to achieve limited local connectivity.

History

DHCP was first defined as a standards track protocol in RFC 1531 in October 1993, as an extension to the Bootstrap Protocol (BOOTP). The motivation for extending BOOTP was that BOOTP required manual intervention to add configuration information for each client, and did not provide a mechanism for reclaiming disused IP addresses.

Many worked to clarify the protocol as it gained popularity, and in 1997 RFC 2131 was released, and remains as of 2011 the standard for IPv4 networks. DHCPv6 is documented in RFC 3315. RFC 3633 added a DHCPv6 mechanism for prefix delegation. DHCPv6 was further extended to provide configuration information to clients configured using stateless address autoconfiguration in RFC 3736.

The BOOTP protocol itself was first defined in RFC 951 as a replacement for the Reverse Address Resolution Protocol RARP. The primary motivation for replacing RARP with BOOTP was that RARP was a data link layer protocol. This made implementation difficult on many server platforms, and required that a server be present on each individual network link. BOOTP introduced the innovation of a *relay agent*, which allowed the forwarding of BOOTP packets off the local network using standard IP routing, thus one central BOOTP server could serve hosts on many IP subnets.[2]

Technical overview

Dynamic Host Configuration Protocol automates network-parameter assignment to network devices from one or more DHCP servers. Even in small networks, DHCP is useful because it makes it easy to add new machines to the network.

When a DHCP-configured client (a computer or any other network-aware device) connects to a network, the DHCP client sends a broadcast query requesting necessary information to a DHCP server. The DHCP server manages a pool of IP addresses and information about client configuration parameters such as default gateway, domain name, the name servers, other servers such as time servers, and so forth. On receiving a valid request, the server assigns the computer an IP address, a lease (length of time the allocation is valid), and other IP configuration parameters, such as the subnet mask and the default gateway. The query is typically initiated immediately after booting, and must complete before the client can initiate IP-based communication with other hosts. Upon disconnecting, the IP address is returned to the pool for use by another computer. This way, many other computers can use the same IP address within minutes of each other.

Depending on implementation, the DHCP server may have three methods of allocating IP-addresses:

- *dynamic allocation*: A network administrator assigns a range of IP addresses to DHCP, and each client computer on the LAN is configured to request an IP address from the DHCP server during network initialization. The request-and-grant process uses a lease concept with a controllable time period, allowing the DHCP server to reclaim (and then reallocate) IP addresses that are not renewed.
- *automatic allocation*: The DHCP server permanently assigns a free IP address to a requesting client from the range defined by the administrator. This is like dynamic allocation, but the DHCP server keeps a table of past IP address assignments, so that it can preferentially assign to a client the same IP address that the client previously had.
- *static allocation*: The DHCP server allocates an IP address based on a table with MAC address/IP address pairs, which are manually filled in (perhaps by a network administrator). [Only requesting clients with a MAC address listed in this table will be allocated an IP address]. This feature (which is not supported by all DHCP servers) is variously called *Static DHCP Asignment* (by DD-WRT), *fixed-address* (by the dhcpd documentation), *Address Reservation* (by Netgear), *DHCP reservation* or *Static DHCP* (by Cisco/Linksys), and *IP reservation* or *MAC/IP binding* (by various other router manufacturers).

Technical details

DHCP uses the same two ports assigned by IANA for BOOTP: destination UDP port 67 for sending data to the server, and UDP port 68 for data to the client. DHCP communications are connectionless in nature.

DHCP operations fall into four basic phases: IP discovery, IP lease offer, IP request, and IP lease acknowledgement. These points are often abbreviated as DORA (Discovery, Offer, Request, Acknowledgement).

DHCP clients and servers on the same subnet communicate via UDP broadcasts, initially. If the client and server are on different subnets, a DHCP Helper or DHCP Relay Agent may be used. Clients requesting renewal of an existing lease may communicate directly via UDP unicast, since the client already has an established IP address at that point.

DHCP discovery

The client broadcasts messages on the physical subnet to discover available DHCP servers. Network administrators can configure a local router to forward DHCP packets to a DHCP server from a different subnet. This client-implementation creates a User Datagram Protocol (UDP) packet with the broadcast destination of 255.255.255.255 or the specific subnet broadcast address.

A DHCP client can also request its last-known IP address (in the example below, 192.168.1.100). If the client remains connected to a network for which this IP is valid, the server may grant the request. Otherwise, it depends whether the server is set up as authoritative or not. An authoritative server will deny the request, making the client ask for a new IP address immediately. A non-authoritative server simply ignores the request, leading to an implementation-dependent timeout for the client to give up on the request and ask for a new IP address.

DHCPDISCOVER

UDP Src=0.0.0.0 sPort=68 Dest=255.255.255.255 dPort=67			
OP	**HTYPE**	**HLEN**	**HOPS**
0x01	0x01	0x06	0x00
XID			
0x3903F326			
SECS		**FLAGS**	
0x0000		0x0000	
CIADDR (Client IP Address)			
0x00000000			
YIADDR (Your IP Address)			
0x00000000			
SIADDR (Server IP Address)			
0x00000000			
GIADDR (Gateway IP Address)			
0x00000000			
CHADDR (Client Hardware Address)			
0x00053C04			
0x8D590000			
0x00000000			
0x00000000			

192 octets of 0s, or overflow space for additional options. BOOTP legacy
Magic Cookie
0x63825363
DHCP Options
DHCP option 53: DHCP Discover
DHCP option 50: 192.168.1.100 requested
DHCP option 55: Parameter Request List: Request Subnet Mask (1), Router (3), Domain Name (15), Domain Name Server (6)

DHCP offer

When a DHCP server receives an IP lease request from a client, it reserves an IP address for the client and extends an IP lease offer by sending a DHCPOFFER message to the client. This message contains the client's MAC address, the IP address that the server is offering, the subnet mask, the lease duration, and the IP address of the DHCP server making the offer.

The server determines the configuration based on the client's hardware address as specified in the CHADDR (Client Hardware Address) field. Here the server, 192.168.1.1, specifies the IP address in the YIADDR (Your IP Address) field.

DHCPOFFER

UDP Src=192.168.1.1 sPort=67 Dest=255.255.255.255 dPort=68			
OP	**HTYPE**	**HLEN**	**HOPS**
0x02	0x01	0x06	0x00
0x00000000			
YIADDR (Your IP Address)			
0xC0A80164			
SIADDR (Server IP Address)			
0xC0A80101			
GIADDR (Gateway IP Address)			
0x00000000			
CHADDR (Client Hardware Address)			
0x00053C04			
0x8D590000			
0x00000000			
0x00000000			
192 octets of 0s. BOOTP legacy			
Magic Cookie			
0x63825363			
DHCP Options			
DHCP option 53: DHCP Offer			
DHCP option 1: 255.255.255.0 subnet mask			

DHCP option 3: 192.168.1.1 router
DHCP option 51: 86400s (1 day) IP lease time
DHCP option 54: 192.168.1.1 DHCP server
DHCP option 6: DNS servers 9.7.10.15, 9.7.10.16, 9.7.10.18

DHCP request

In response to the offer Client requests the server. The client replies DHCPRequest, unicast to the server, requesting the offered address. A client can receive DHCP offers from multiple servers, but it will accept only one DHCP offer. Based on the Transaction ID field in the request, servers are informed whose offer the client has accepted. When other DHCP servers receive this message, they withdraw any offers that they might have made to the client and return the offered address to the pool of available addresses. In some cases DHCP request message is broadcast, instead of being unicast to a particular DHCP server, because the DHCP client has still not received an IP address. Also, this way one message can let all other DHCP servers know that another server will be supplying the IP address without missing any of the servers with a series of unicast messages.

DHCPREQUEST

UDP Src=0.0.0.0 sPort=68 Dest=255.255.255.255 dPort=67			
OP	**HTYPE**	**HLEN**	**HOPS**
0x01	0x01	0x06	0x00
XID			
0x3903F326			
SECS		**FLAGS**	
0x0000		0x0000	
CIADDR (Client IP Address)			
0x00000000			
YIADDR (Your IP Address)			
0x00000000			
SIADDR (Server IP Address)			
0xC0A80101			
GIADDR (Gateway IP Address)			
0x00000000			
CHADDR (Client Hardware Address)			
0x00053C04			
0x8D590000			
0x00000000			
0x00000000			
192 octets of 0s. BOOTP legacy			
Magic Cookie			
0x63825363			
DHCP Options			

DHCP option 53: DHCP Request
DHCP option 50: 192.168.1.100 requested
DHCP option 54: 192.168.1.1 DHCP server.

DHCP acknowledgement

When the DHCP server receives the DHCPREQUEST message from the client, the configuration process enters its final phase. The acknowledgement phase involves sending a DHCPACK packet to the client. This packet includes the lease duration and any other configuration information that the client might have requested. At this point, the IP configuration process is completed.

The protocol expects the DHCP client to configure its network interface with the negotiated parameters.

DHCPACK

UDP Src=192.168.1.1 sPort=67 Dest=255.255.255.255 dPort=68			
OP	**HTYPE**	**HLEN**	**HOPS**
0x02	0x01	0x06	0x00
XID			
0x3903F326			
SECS		**FLAGS**	
0x0000		0x0000	
CIADDR (Client IP Address)			
0x00000000			
YIADDR (Your IP Address)			
0xC0A80164			
SIADDR (Server IP Address)			
0xC0A80101			
GIADDR (Gateway IP Address switched by relay)			
0x00000000			
CHADDR (Client Hardware Address)			
0x00053C04			
0x8D590000			
0x00000000			
0x00000000			
192 octets of 0s. BOOTP legacy			
Magic Cookie			
0x63825363			
DHCP Options			
DHCP option 53: DHCP ACK			
DHCP option 1: 255.255.255.0 subnet mask			
DHCP option 3: 192.168.1.1 router			

DHCP option 51: 86400s (1 day) IP lease time	
DHCP option 54: 192.168.1.1 DHCP server	
DHCP option 6: DNS servers 9.7.10.15, 9.7.10.16, 9.7.10.18	

After the client obtains an IP address, the client may use the Address Resolution Protocol (ARP) to prevent IP conflicts caused by overlapping address pools of DHCP servers.

DHCP information

A DHCP client may request more information than the server sent with the original DHCPOFFER. The client may also request repeat data for a particular application. For example, browsers use *DHCP Inform* to obtain web proxy settings via WPAD. Such queries do not cause the DHCP server to refresh the IP expiry time in its database.

DHCP releasing

The client sends a request to the DHCP server to release the DHCP information and the client deactivates its IP address. As client devices usually do not know when users may unplug them from the network, the protocol does not mandate the sending of *DHCP Release*.

Client configuration parameters in DHCP

A DHCP server can provide optional configuration parameters to the client. RFC 2132 describes the available DHCP options defined by Internet Assigned Numbers Authority (IANA) - DHCP and BOOTP PARAMETERS [3].
A DHCP client can select, manipulate and overwrite parameters provided by a DHCP server.[4]

DHCP options

The following tables list the available DHCP options, as stated in RFC2132.[5]

RFC1497 vendor extensions[6]

Code	Name	Length	Notes
0	Pad[7]	1 octet	Can be used to pad other options so that they are aligned to the word boundary
1	Subnet Mask[8]	4 octets	Must be sent after the router option (option 3) if both are included
2	Time Offset[9]	4 octets	
3	Router	multiples of 4 octets	Available routers, should be listed in order of preference
4	Time Server	multiples of 4 octets	Available time servers to synchronise with, should be listed in order of preference
5	Name Server	multiples of 4 octets	Available IEN116 name servers, should be listed in order of preference
6	Domain Name Server	multiples of 4 octets	Available DNS servers, should be listed in order of preference
7	Log Server	multiples of 4 octets	Available log servers, should be listed in order of preference.
8	Cookie Server	multiples of 4 octets	
9	LPR Server	multiples of 4 octets	
10	Impress Server	multiples of 4 octets	
11	Resource Location Server	multiples of 4 octets	
12	Host Name	minimum of 1 octet	
13	Boot File Size	2 octets	Length of the boot image in 4KiB blocks

14	Merit Dump File	minimum of 1 octet	Path where crash dumps should be stored
15	Domain Name	minimum of 1 octet	
16	Swap Server	4 octets	
17	Root Path	minimum of 1 octet	
18	Extensions Path	minimum of 1 octet	
255	End	1 octet	Used to mark the end of the vendor option field

IP Layer Parameters per Host[10]

Code	Name	Length	Notes
19	IP Forwarding Enable/Disable	1 octet	
20	Non-Local Source Routing Enable/Disable	1 octet	
21	Policy Filter	multiples of 8 octets	
22	Maximum Datagram Reassembly Size	2 octets	
23	Default IP Time-to-live	1 octet	
24	Path MTU Aging Timeout	4 octets	
25	Path MTU Plateau Table	multiples of 2 octets	

IP Layer Parameters per Interface[11]

Code	Name	Length	Notes
26	Interface MTU	2 octets	
27	All Subnets are Local	1 octet	
28	Broadcast Address	4 octets	
29	Perform Mask Discovery	1 octet	
30	Mask Supplier	1 octet	
31	Perform Router Discovery	1 octet	
32	Router Solicitation Address	4 octets	
33	Static Route	multiples of 8 octets	A list of destination/router pairs

Link Layer Parameters per Interface[12]

Code	Name	Length	Notes
34	Trailer Encapsulation Option	1 octet	
35	ARP Cache Timeout	4 octets	
36	Ethernet Encapsulation	1 octet	

TCP Parameters[13]

Code	Name	Length	Notes
37	TCP Default TTL	1 octet	
38	TCP Keepalive Interval	4 octets	
39	TCP Keepalive Garbage	1 octet	

Application and Service Parameters[14]

Code	Name	Length	Notes
40	Network Information Service Domain	minimum of 1 octet	
41	Network Information Servers	multiples of 4 octets	
42	Network Time Protocol Servers	multiples of 4 octets	
43	Vendor Specific Information	minimum of 1 octets	
44	NetBIOS over TCP/IP Name Server	multiples of 4 octets	
45	NetBIOS over TCP/IP Datagram Distribution Server	multiples of 4 octets	
46	NetBIOS over TCP/IP Node Type	1 octet	
47	NetBIOS over TCP/IP Scope	minimum of 1 octet	
48	X Window System Font Server	multiples of 4 octets	
49	X Window System Display Manager	multiples of 4 octets	
64	Network Information Service+ Domain	minimum of 1 octet	
65	Network Information Service+ Servers	multiples of 4 octets	
68	Mobile IP Home Agent	multiples of 4 octets	
69	Simple Mail Transport Protocol (SMTP) Server	multiples of 4 octets	
70	Post Office Protocol (POP3) Server	multiples of 4 octets	
71	Network News Transport Protocol (NNTP) Server	multiples of 4 octets	
72	Default World Wide Web (WWW) Server)	multiples of 4 octets	
73	Default Finger Server	multiples of 4 octets	
74	Default Internet Relay Chat (IRC) Server	multiples of 4 octets	
75	StreetTalk Server	multiples of 4 octets	
76	StreetTalk Directory Assistance (STDA) Server	multiples of 4 octets	

DHCP Extensions[15]

Code	Name	Length	Notes
50	Requested IP Address	4 octets	
51	IP Address Lease Time	4 octets	
52	Option Overload	1 octet	
66	TFTP server name	minimum of 1 octet	
67	Bootfile name	minimum of 1 octet	
53	DHCP Message Type	1 octet	
54	Server Identifier	4 octets	
55	Parameter Request List	minimum of 1 octet	
56	Message	minimum of 1 octet	
57	Maximum DHCP Message Size	2 octets	
58	Renewal (T1) Time Value	4 octets	
59	Rebinding (T2) Time Value	4 octets	
60	Vendor class identifier	minimum of 1 octet	
61	Client-identifier	minimum of 2 octets	

Vendor identification

An option exists to identify the vendor and functionality of a DHCP client. The information is a variable-length string of characters or octets which has a meaning specified by the vendor of the DHCP client. One method that a DHCP client can utilize to communicate to the server that it is using a certain type of hardware or firmware is to set a value in its DHCP requests called the Vendor Class Identifier (VCI) (Option 60). This method allows a DHCP server to differentiate between the two kinds of client machines and process the requests from the two types of modems appropriately. Some types of set-top boxes also set the VCI (Option 60) to inform the DHCP server about the hardware type and functionality of the device. The value that this option is set to give the DHCP server a hint about any required extra information that this client needs in a DHCP response.

DHCP relaying

In small networks, where only one IP subnet is being managed, DHCP clients communicate directly with DHCP servers. However, DHCP servers can also provide IP addresses for multiple subnets. In this case, a DHCP client that has not yet acquired an IP address cannot communicate directly with the DHCP server using IP routing, because it doesn't have a routable IP address, nor does it know the IP address of a router. In order to allow DHCP clients on subnets not directly served by DHCP servers to communicate with DHCP servers, DHCP relay agents can be installed on these subnets. The DHCP client broadcasts on the local link; the relay agent receives the broadcast and transmits it to one or more DHCP servers using unicast. The relay agent stores its own IP address in the GIADDR field of the DHCP packet. The DHCP server uses the GIADDR to determine the subnet on which the relay agent received the broadcast, and allocates an IP address on that subnet. When the DHCP server replies to the client, it sends the reply to the GIADDR address, again using unicast. The relay agent then retransmits the response on the local network.

Reliability

The DHCP protocol provides reliability in several ways: periodic renewal, rebinding, and failover. DHCP clients are allocated leases that last for some period of time. Clients begin to attempt to renew their leases once half the lease interval has expired. They do this by sending a unicast DHCPREQUEST message to the DHCP server that granted the original lease. If that server is down or unreachable, it will fail to respond to the DHCPREQUEST. However, the DHCPREQUEST will be repeated by the client from time to time, so when the DHCP server comes back up or becomes reachable again, the DHCP client will succeed in contacting it, and renew its lease.

If the DHCP server is unreachable for an extended period of time, the DHCP client will attempt to rebind, by broadcasting its DHCPREQUEST rather than unicasting it. Because it is broadcast, the DHCPREQUEST message will reach all available DHCP servers. If some other DHCP server is able to renew the lease, it will do so at this time.

In order for rebinding to work, when the client successfully contacts a backup DHCP server, that server must have accurate information about the client's binding. Maintaining accurate binding information between two servers is a complicated problem; if both servers are able to update the same lease database, there must be a mechanism to avoid conflicts between updates on the independent servers. A standard for implementing fault-tolerant DHCP servers was developed at the Internet Engineering Task Force.[16][17]

If rebinding fails, the lease will eventually expire. When the lease expires, the client must stop using the IP address granted to it in its lease. At that time, it will restart the DHCP process from the beginning by broadcasting a DHCPDISCOVER message. Since its lease has expired, it will accept any IP address offered to it. Once it has a new IP address, presumably from a different DHCP server, it will once again be able to use the network. However, since its IP address has changed, any ongoing connections will be broken.

Security

The base DHCP protocol does not include any mechanism for authentication.[18] Because of this, it is vulnerable to a variety of attacks. These attacks fall into three main categories:

- Unauthorized DHCP servers providing false information to clients.[19]
- Unauthorized clients gaining access to resources.[19]
- Resource exhaustion attacks from malicious DHCP clients.[19]

Because the client has no way to validate the identity of a DHCP server, unauthorized DHCP servers can be operated on networks, providing incorrect information to DHCP clients. This can serve either as a denial-of-service attack, preventing the client from gaining access to network connectivity, or as a man-in-the-middle attack. Because the DHCP server provides the DHCP client with server IP addresses, such as the IP address of one or more DNS servers,[19] an attacker can convince a DHCP client to do its DNS lookups through its own DNS server, and can therefore provide its own answers to DNS queries from the client.[20] This in turn allows the attacker to redirect network traffic through itself, allowing it to eavesdrop on connections between the client and network servers it contacts, or to simply replace those network servers with its own.[20]

Because the DHCP server has no secure mechanism for authenticating the client, clients can gain unauthorized access to IP addresses by presenting credentials, such as client identifiers, that belong to other DHCP clients. This also allows DHCP clients to exhaust the DHCP server's store of IP addresses—by presenting new credentials each time it asks for an address, the client can consume all the available IP addresses on a particular network link, preventing other DHCP clients from getting service.

DHCP does provide some mechanisms for mitigating these problems. The Relay Agent Information Option protocol extension (RFC 3046) allows network operators to attach tags to DHCP messages as these messages arrive on the network operator's trusted network. This tag is then used as an authorization token to control the client's access to network resources. Because the client has no access to the network upstream of the relay agent, the lack of authentication does not prevent the DHCP server operator from relying on the authorization token.[18]

Another extension, Authentication for DHCP Messages (RFC 3118), provides a mechanism for authenticating DHCP messages. Unfortunately RFC 3118 has not seen widespread adoption because of the problems of managing keys for large numbers of DHCP clients.[21]

Notes

[1] Ralph Droms; Ted Lemon (2003). *The DHCP Handbook*. SAMS Publishing. p. 436. ISBN 0-672-32327-3.

[2] Bill Croft; John Gilmore (September 1985). "RFC 951 - Bootstrap Protocol" (http://tools.ietf.org/html/rfc951#section-6). *Network Working Group*. .

[3] http://www.iana.org/assignments/bootp-dhcp-parameters

[4] In Unix-like systems this client-level refinement typically takes place according to the values in a `/etc/dhclient.conf` configuration file.

[5] Alexander, Steve; Droms, Ralph (March 1997). *DHCP Options and BOOTP Vendor Extensions* (https://tools.ietf.org/html/rfc2132). IETF. RFC 2132. . Retrieved June 10, 2012.

[6] Alexander, Steve; Droms, Ralph (March 1997). "RFC 2132: DHCP Options and BOOTP Vendor Extensions" (http://tools.ietf.org/html/rfc2132#section-3). IETF. Section 3: RFC 1497 vendor extensions. . Retrieved 2012-07-26.

[7] Alexander, Steve; Droms, Ralph (March 1997). "RFC 2132: DHCP Options and BOOTP Vendor Extensions" (http://tools.ietf.org/html/rfc2132#section-3.1). IETF. Section 3.1: Pad Option. . Retrieved 2012-07-26.

[8] Alexander, Steve; Droms, Ralph (March 1997). "RFC 2132: DHCP Options and BOOTP Vendor Extensions" (http://tools.ietf.org/html/rfc2132#section-3.3). IETF. Section 3.3: Subnet Mask. . Retrieved 2012-07-26.

[9] Alexander, Steve; Droms, Ralph (March 1997). "RFC 2132: DHCP Options and BOOTP Vendor Extensions" (http://tools.ietf.org/html/rfc2132#section-3.4). IETF. Section 3.4: Time Offset. . Retrieved 2012-07-26.

[10] Alexander, Steve; Droms, Ralph (March 1997). "RFC 2132: DHCP Options and BOOTP Vendor Extensions" (http://tools.ietf.org/html/rfc2132#section-4). IETF. Section 4: IP Layer Parameters per Host. . Retrieved 2012-07-26.

[11] Alexander, Steve; Droms, Ralph (March 1997). "RFC 2132: DHCP Options and BOOTP Vendor Extensions" (http://tools.ietf.org/html/rfc2132#section-5). IETF. Section 5: IP Layer Parameters per Interface. . Retrieved 2012-07-26.

[12] Alexander, Steve; Droms, Ralph (March 1997). "RFC 2132: DHCP Options and BOOTP Vendor Extensions" (http://tools.ietf.org/html/rfc2132#section-6). IETF. Section 6: Link Layer Parameters per Interface. . Retrieved 2012-07-26.

[13] Alexander, Steve; Droms, Ralph (March 1997). "RFC 2132: DHCP Options and BOOTP Vendor Extensions" (http://tools.ietf.org/html/rfc2132#section-7). IETF. Section 7: TCP Parameters. . Retrieved 2012-07-26.

[14] Alexander, Steve; Droms, Ralph (March 1997). "RFC 2132: DHCP Options and BOOTP Vendor Extensions" (http://tools.ietf.org/html/rfc2132#section-8). IETF. Section 8: Application and Service Parameters. . Retrieved 2012-07-26.

[15] Alexander, Steve; Droms, Ralph (March 1997). "RFC 2132: DHCP Options and BOOTP Vendor Extensions" (http://tools.ietf.org/html/rfc2132#section-9). IETF. Section 9: DHCP Extensions. . Retrieved 2012-07-26.

[16] Droms, Ralph; Kinnear, Kim; Stapp, Mark; Volz, Bernie; Gonczi, Steve; Rabil, Greg; Dooley, Michael; Kapur, Arun (March 2003). *DHCP Failover Protocol* (https://tools.ietf.org/html/draft-ietf-dhc-failover-12). IETF. I-D draft-ietf-dhc-failover-12. . Retrieved May 09, 2010.

[17] The IETF proposal provided a mechanism whereby two servers could remain loosely in sync with each other in such a way that even in the event of a total failure of one server, the other server could recover the lease database and continue operating. Due to the length and complexity of the specification, it was never published as a standard; however, the techniques described in the specification are in wide use, with one open source implementation in the ISC DHCP server as well as several commercial implementations.

[18] Michael Patrick (January 2001). "RFC 3046 - DHCP Relay Agent Information Option" (http://tools.ietf.org/html/rfc3046#section-7). *Network Working Group*. .

[19] Ralph Droms (March 1997). "RFC 2131 - Dynamic Host Configuration Protocol" (http://tools.ietf.org/html/rfc2131#section-7). *Network Working Group*. .

[20] Sergey Golovanov (Kaspersky Labs) (June 2011). "TDSS loader now got "legs"" (http://www.securelist.com/en/blog/208188095/TDSS_loader_now_got_legs). .

[21] Ted Lemon (April 2002). "Implementation of RFC 3118" (http://www.ietf.org/mail-archive/web/dhcwg/current/msg00876.html). .

References

External links

- RFC 2131 - Dynamic Host Configuration Protocol
- RFC 2132 - DHCP Options and BOOTP Vendor Extensions
- RFC 3046 - DHCP Relay Agent Information Option
- RFC 3942 - Reclassifying Dynamic Host Configuration Protocol Version Four (DHCPv4) Options
- RFC 4242 - Information Refresh Time Option for Dynamic Host Configuration Protocol for IPv6
- RFC 4361 - Node-specific Client Identifiers for Dynamic Host Configuration Protocol Version Four (DHCPv4)
- RFC 4436 - Detecting Network Attachment in IPv4 (DNAv4)

Trivial File Transfer Protocol

Trivial File Transfer Protocol (**TFTP**) is a file transfer protocol notable for its simplicity. It is generally used for automated transfer of configuration or boot files between machines in a local environment. Compared to FTP, TFTP is extremely limited, providing no authentication, and is rarely used interactively by a user.

Due to its simple design, TFTP could be implemented using a very small amount of memory. It is therefore useful for booting computers such as routers which may not have any data storage devices. It is an element of the Preboot Execution Environment (PXE) network boot protocol, where it is implemented in the firmware ROM / NVRAM of the host's network card.

It is also used to transfer small amounts of data between hosts on a network, such as IP phone firmware or operating system images when a remote X Window System terminal or any other thin client boots from a network host or server. The initial stages of some network based installation systems (such as Solaris Jumpstart, Red Hat Kickstart, Symantec Ghost and Windows NT's Remote Installation Services) use TFTP to load a basic kernel that performs the actual installation. It was used for saving router configurations on Cisco routers, but was later augmented by other protocols.[1]

TFTP was first defined in 1980 by IEN 133.[2] It is currently defined by RFC 1350. There have been some extensions to the TFTP protocol documented in later RFC's (see the section on Extensions, below). TFTP is based in part on the earlier protocol EFTP, which was part of the PUP protocol suite. TFTP support appeared first as part of 4.3 BSD.

Due to the lack of security, it is dangerous to use it over the Internet. Thus, TFTP is generally only used on private, local networks.

Overview

Trivial File Transfer Protocol (TFTP) is a simple protocol to transfer files. It has been implemented on top of the User Datagram Protocol (UDP) using port number 69. TFTP is designed to be small and easy to implement, and therefore it lacks most of the features of a regular FTP. TFTP only reads and writes files (or mail) from/to a remote server. It cannot list directories, and currently has no provisions for user authentication.

In **TFTP**, any transfer begins with a request to read or write a file, which also serves to request a connection. If the server grants the request, the connection is opened and the file is sent in fixed length blocks of 512 bytes. Each data packet contains one block of data, and must be acknowledged by an acknowledgment packet before the next packet can be sent. A data packet of less than 512 bytes signals termination of a transfer. If a packet gets lost in the network, the intended recipient will timeout and may retransmit his last packet (which may be data or an acknowledgment), thus causing the sender of the lost packet to retransmit that lost packet. The sender has to keep just one packet on hand for retransmission, since the lock step acknowledgment guarantees that all older packets have been received. Notice that both machines involved in a transfer are considered senders and receivers. One sends data and receives acknowledgments, the other sends acknowledgments and receives data.

TFTP typically uses UDP as its transport protocol, but it is not a requirement. Data transfer is initiated on port 69, but the data transfer ports are chosen independently by the sender and receiver during initialization of the connection. The ports are chosen at random according to the parameters of the networking stack, typically from the range of Ephemeral ports.[3]

TFTP defines three modes of transfer: netascii, octet, and mail. Netascii is a modified form of ASCII, defined in RFC 764. It consists of an 8-bit extension of the 7-bit ASCII character space from 0x20 to 0x7F (the printable characters and the space) and eight of the control characters. The allowed control characters include the null (0x00), the line feed (LF, 0x0A), and the carriage return (CR, 0x0D). Netascii also requires that the end of line marker on a host be translated to the character pair CR LF for transmission, and that any CR must be followed by either a LF or the null.

Octet allows for the transfer of arbitrary 8-bit bytes, with the received file identical to the sent file. More correctly, if a host receives an octet file and then returns it, the returned file must be identical to the original.[4] The Mail transfer mode uses Netascii transfer, but the file is sent to an email recipient by specifying that recipient's email address as the file name. RFC 1350 declared this mode of transfer obsolete.

No security or authentication is provided by the protocol specification. Unix implementations often restrict file transfers to a single configured directory, and only to read from files with world readability, and only write to already existing files that have world writeability.

Protocol walkthrough

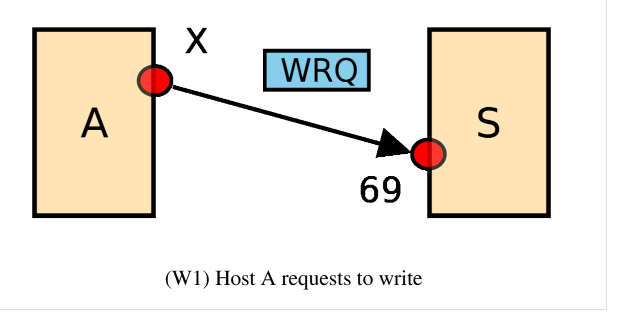

(W1) Host A requests to write

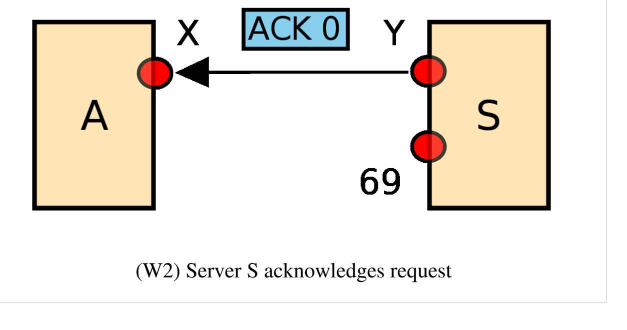

(W2) Server S acknowledges request

1. The initiating host A sends an RRQ (read request) or WRQ (write request) packet to host S at port number 69, containing the filename and transfer mode.
2. S replies with an ACK (acknowledgement) packet to WRQ and directly with a DATA packet to RRQ. Packet is sent from a freshly allocated ephemeral port, and all future packets to host S should be to this port.
3. The source host sends numbered DATA packets to the destination host, all but the last containing a full-sized block of data (512 bytes). The destination host replies with numbered ACK packets for all DATA packets.
4. The final DATA packet must contain less than a full-sized block of data to signal that it is the last. If the size of the transferred file is an exact multiple of the block-size, the source sends a final DATA packet containing 0 bytes of data.
5. Receiver responds to each DATA with associated numbered ACK. Sender responds to the first received ACK of a block with DATA of the next block.

6. If an ACK is not eventually received, a retransmit timer resends DATA packet.

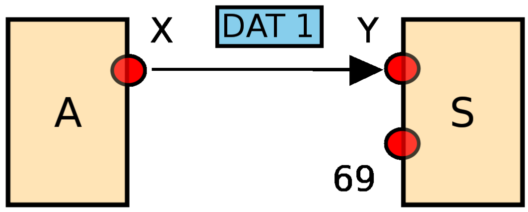

(W3) Host A sends numbered data packets

Additional details

- The original versions of TFTP, prior to RFC 1350, displayed a particularly bad protocol flaw which was named Sorcerer's Apprentice Syndrome (after "The Sorcerer's Apprentice" segment of *Fantasia*) when it was discovered.
- In the early days of work on the TCP/IP protocol suite, TFTP was often the first protocol implemented on a new host type, because it was so simple.

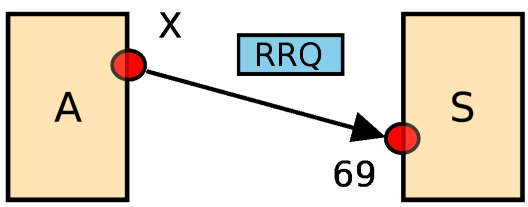

(R1) Host A requests to read

Extensions

- The original protocol has a file size limit of 32 MB. In 1998 this limit was extended to 4 GB by RFC 2347 which introduced option negotiation and RFC 2348 which introduced block-size negotiation. If the server and client support block number wraparound, file size is essentially unlimited.
- Since TFTP utilizes UDP, it has to supply its own transport and session support. Each file transferred via TFTP constitutes an independent exchange. Classically, this transfer is performed in lock-step, with only one packet (either a block of data, or an 'acknowledgement') ever in flight on the network at any time. Due to this lack of windowing, TFTP provides low throughput over high latency links. Note that Windows 2008 introduced pipelined TFTP as part of Windows Deployment Services (WDS) and uses an 8 packet window by default. This dramatically improves performance for things like PXE booting.

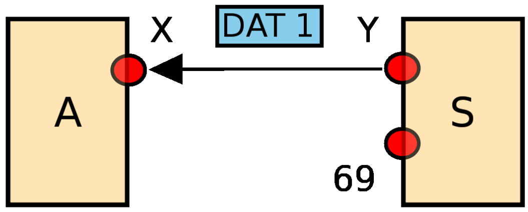

(R2) Server S sends data packet 1

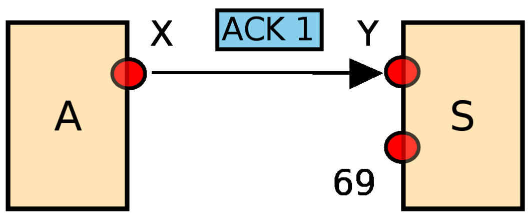

(R3) Host A acknowledges data packet 1

Known TFTP implementations

GNU inetutils [5]
: the GNU project network suite includes a TFTP client/server implementation

tftp-hpa
: An opensource TFTP host published under BSD license.

atftp [6]
: a GPL client/server implementation of the TFTP protocol for Linux

tftp-server
: a GPL, multi-threaded TFTP server for Linux

TFTP Server [7]
: a free mobile TFTP server for Apple iPhone and iPad

TFTP Server [8]
: TFTP server for OS X

Tftpd32 [9]
: an opensource (EUPL) IPv6 ready TFTP and DHCP server/service for Windows

haneWIN DHCP [10]
: a shareware TFTP and DHCP server/service for Windows

References

[1] http://www.cisco.com/en/US/docs/ios/11_3/configfun/command/reference/frcfgfil.html#wp14279
[2] Karen R. Sollins (1980-01-29). *The TFTP Protocol* (https://tools.ietf.org/rfcmarkup?url=ftp://ftp.rfc-editor.org/in-notes/ien/ien133.txt). IETF. IEN 133. . Retrieved 2010-05-01.
[3] Karen R. Sollins (July 1992). *The TFTP Protocol (Revision 2)* (https://tools.ietf.org/html/rfc1350). IETF. RFC 1350. . Retrieved 2010-05-01.
[4] RFC 1350, page 5.
[5] http://www.gnu.org/software/inetutils/
[6] http://freshmeat.net/projects/atftp/
[7] http://www.mobiletftpserver.com
[8] http://ww2.unime.it/flr/tftpserver/
[9] http://tftpd32.jounin.net/
[10] http://www.hanewin.net/dhcp-e.htm

Further reading

- RFC 906 – Bootstrap loading using TFTP, R. Finlayson, June 1984.
- RFC 1350 – TFTP Protocol (revision 2), K. R. Sollins, July 1992. (This superseded the preceding, RFC 783 and earlier FTP RFCs back to the original IEN 133)
- RFC 1785 – TFTP Option Negotiation Analysis, G. Malkin, A. Harkin, March 1995.
- RFC 2090 – TFTP Multicast Option, A. Emberson, February 1997. (*Status: Experimental*)
- RFC 2347 – TFTP Option Extension, G. Malkin, A. Harkin, May 1998. (This superseded the preceding, RFC 1782)
- RFC 2348 – TFTP Blocksize Option, G. Malkin, A. Harkin, May 1998. (This superseded the preceding, RFC 1783)
- RFC 2349 – TFTP Timeout Interval and Transfer Size Options, G. Malkin, A. Harkin, May 1998 (This superseded the preceding, RFC 1784).
- RFC 3617 – Uniform Resource Identifier (URI) Scheme and Applicability Statement for the Trivial File Transfer Protocol (TFTP), E. Lear, October 2003.

Disk cloning

Disk cloning is the process of copying the contents of one computer hard disk to another disk or to an "image" file. This may be done straight from one disk to another, but more often, the contents of the first disk are written to an image file as an intermediate step, then the second disk is loaded with the contents of the image. Typically, this is done for archiving purposes, to restore lost or damaged data, or to move wanted data into a new disk, though other reasons also exist.

Unlike standard copying functions, disk cloning involves copying hidden and in-use files, and thus presents special challenges, as those types of files are typically not available for copying. Additional complications arise when the process is used for networked computers, as the network must be able to distinguish between different computers. Post-cloning operations may be necessary to address these and other issues.

Common Uses

There are a number of notable uses for disk cloning software. These include:

- **Reboot and restore** – a technique in which the disk of a computer is automatically wiped and restored from a "clean", master image, which should be in full working order and should have been swept for viruses. This is used by some cybercafes and some training and educational institutes, and helps ensure that even if a user does misconfigure something, downloads inappropriate content or programs, or infects a computer with a virus, the computer will be restored to a clean, working state. The reboot and restore process can either take place irregularly when a computer shows signs of malfunctioning, on a regular basis (e.g., nightly) or even, in some cases, every time a user logs off, which is the safest approach (although that does involve some downtime).
- **Provisioning new computers** – Provisioning with a standard set of software so that a new user is ready to go straight away with a complete application suite and does not have to waste time installing individual applications. This is often done by original equipment manufacturers and larger companies.
- **Hard drive upgrade** – An individual user may use disk copying (cloning) to upgrade to a new, usually larger, hard disk.
- **Full system backup** – A user may create a comprehensive backup of their operating system and installed software.
- **System recovery** – An OEM can provide media that can restore a computer to its original factory software configuration.
- **Transfer to another user** – A system sold or given to another person may be reset by reloading a known, previously-saved image that contains no personal files or information.

How it works

This article is specific to disk cloning on the x86 platform; specific details may not apply to other platforms.

To provision the hard disk of a computer without using disk cloning software, the following steps are generally required for each computer:

1. Create one or more partitions on the disk
2. Format each partition to create a file system on it
3. Install the operating system
4. Install device drivers for the particular hardware
5. Install application software

With disk cloning, this is simplified to:

1. Install the first computer, as above.

2. Create an image of the hard disk (optional)
3. Clone the first disk, or its image, to the remaining computers.

This can be referred to simply as a recovery disc.

History

Before Windows 95, some computer manufacturers used hardware disk copying machines to copy software. This had the disadvantages of copying not just the used data on the disk, but also unused sectors, as the hardware used was not aware of the structures on the disks. A larger hard disk could not be copied to a smaller one, and copying a smaller one to a larger left the remaining space on the new disk unused. The two disks required identical geometries.

Other manufacturers and companies partitioned and formatted disks manually, then used file copy utilities or archiving utilities, such as tar or zip to copy files. It is not sufficient simply to copy all files from one disk to another, because there are special boot files or boot tracks which must be specifically placed for an operating system to run, so additional manual steps were required.

Windows 95 compounded the problems because it was larger than earlier popular operating systems, and thus took more time to install. The long filenames added to the FAT filesystem by Microsoft in Windows 95 were not supported by most copy programs, and the introduction of the FAT32 filesystem in 1997 caused problems for others. The growth of the personal computer market at this time also made a more efficient solution desirable.

Ghost was introduced in 1996 by Binary Research. It initially supported only FAT filesystems directly, but it could copy but not resize other filesystems by performing a sector copy on them. Ghost added support for the NTFS filesystem later that year, and also provided a program to change the Security Identifier (SID) which made Windows NT systems distinguishable from each other. Support for the ext2 filesystem was added in 1999.

Competitors to Ghost soon arose, and a features war has carried on to the present day. Many disk cloning programs now offer features which go beyond simple disk cloning, such as asset management and user settings migration.

On UNIX based computer systems, dd was more commonplace due to the lack of filesystem support in Ghost.

Post-cloning operations

Two machines with identical names are said not to be allowed on the same network, and, for Windows NT and its successors, two machines with identical security IDs (SIDs, aka Security Identifier) are said not to be allowed on the same Active Directory domain.[1][2][3] A disk cloning program should change these as part of copying the disk or restoring the image. Some operating systems are also not well suited to changes in hardware, so that a clone of Windows XP for example may object to being booted on a machine with a different motherboard, graphics card and network card, especially if non-generic drivers are used. Microsoft's solution to this is Sysprep, a utility which runs hardware detection scans and sets the SID and computer name freshly when the machine boots. Microsoft recommends that Sysprep be set up on all machines before cloning, rather than allow third party programs to configure them. Similarly, Linux systems simply require the necessary kernel modules to be available (or compiled directly into the kernel), in order to support new hardware when the machine boots. However there are ways to help make images for cloning with Windows more portable. One such example would be a product called Universal Imaging Utility [4] from Binary Research (original developers of Symantec's Ghost) which incorporates a large number of hardware device drivers into the sysprep routine.

When it comes to "Domain SID", the Domain SID is recomputed each time a computer enters a domain. Thus, all the "post-cloning operations" that are based on "leave the domain and then rejoin the domain" will actually cause a re-creation of the Domain SID for the computer that joins the domain.

In other words, duplicated SIDs are usually not a problem with Microsoft Windows systems

There are files in some Microsoft operating systems (called BOOTSECT.*) which are copies of the Boot Partition Block (BPB) used by alternate operating systems that Microsoft Windows loader (NTLDR) can load. BOOTSECT.* files may have to be altered if partition sizes or layouts are changed during the clone.

Linux systems usually boot using either the LILO or GRUB bootloaders. These contain lists of absolute disk sectors in their MBR, which must be altered by the cloning program as the files they refer to are likely not to be in the same location of the destination disk. For example, if the original boot loader script points to the system being on a disk on channel 0 and the system being of the second partition, the target computer will need to have the same configuration.

Operating environment

A disk cloning program needs to be able to read even protected operating system files on the source disk, and must guarantee that the system is in a consistent state at the time of reading. It must also overwrite any operating system already present on the destination disk. To simplify these tasks, most disk cloning programs can run under an operating system different from the native operating system of the host computer, for example, MS-DOS or an equivalent such as PC-DOS or DR-DOS, or Linux. The computer is booted from this operating system, the cloning program is loaded and copies the Windows file system. Many programs (e.g. Acronis True Image) can clone a disk, or make an image, from within Windows, with special provision for copying open files; but an image cannot be restored onto the Windows System Drive under Windows.

A disc cloning program running under non-Windows operating systems must have device drivers or equivalent for all devices used. The manufacturers of some devices do not provide suitable drivers, so the manufacturers of disk cloning software must write their own drivers, or include device access functionality in some other way. This applies to tape drives, CD and DVD readers and writers, and USB and FireWire drives. Cloning software contains its own TCP/IP stack for multicast transfer of data where required.

Image transfer

The simplest method of cloning a disk is to have both the source and destination disks present in the same machine, but this is often not possible. Disk cloning programs can link two computers by a parallel cable, or save and load images to an external USB drive or network drive. As disk images tend to be very large (usually a minimum of several hundred MB), performing several clones at a time puts excessive stress on a network. The solution is to use multicast technology. This allows a single image to be sent simultaneously to many machines without putting greater stress on the network than sending an image to a single machine.

Image manipulation

Although disk cloning programs are not primarily backup programs, they are sometimes used as such. A key feature of a backup program is to allow the retrieval of individual files without needing to restore the entire backup. Disk cloning programs either provide a Windows Explorer-like program to browse image files and extract individual files from them, or allow an image file to be mounted as a read-only filesystem within Windows Explorer.

Some such programs allow deletion of files from images, and addition of new files.

Notes

[1] What are the problems with workstations having the same SID? (http://www.windowsitpro.com/article/articleid/14919/what-are-the-problems-with-workstations-having-the-same-sid.html)

[2] *Problems With Duplicate SIDs* (http://web.archive.org/web/20070212020819/http://appdeploy.com/articles/sids.shtml), archived from the original (http://www.appdeploy.com/articles/sids.shtml) on 20070212,

[3] The problem with duplicated SIDs in a Workgroup of computers running Windows NT/2K/XP is only related to different user accounts having the same SID. This could lead to unexpected access to shared files or files stored on a removable storage: If some ACLs (Access control lists) are set on a file, the actual permissions can be associated with a user SID. If this user SID is duplicated on a cloned computer (because the computer SID is duplicated and because the user SIDs are built based on the computer SID + a sequential number), a user of a second computer (cloned from the first one) could have access to the files that the user of a first computer has protected.

[4] http://www.binaryresearch.net/products/the_universal_imaging_utility

References

Preseed

Preseeding is a method for automating the installation of the Debian GNU/Linux operating system and its derivatives. Answers to installation questions, which would normally be answered interactively by an operator, are predetermined and supplied via a configuration file (and sometimes boot parameters). This is similar to unattended installations of Windows operating systems using an answer file (see Installation (computer programs)).

Many Debian-based operating systems support preseed, because it is a feature of the Debian-Installer (also known as "d-i"). For instance, although Ubuntu is commonly installed via the user-friendly Ubiquity installer, preseeding the d-i is the recommended method for automating Ubuntu installations [1] and for customizing install CDs [2].

Note that preseeding automates the operating system installation, but it does not necessarily continue to detailed configuration or application installation in the same way as Fully Automatic Installation

References

[1] "Automating the installation using preseeding" (https://help.ubuntu.com/8.04/installation-guide/powerpc/appendix-preseed.html). Ubuntu documentation team. . Retrieved 2009-11-20.

[2] "InstallCDCustomization" (https://help.ubuntu.com/community/InstallCDCustomization). Ubuntu documentation team. . Retrieved 2009-11-20.

External links

- Preseeding the d-i (http://wiki.debian.org/DebianInstaller/Preseed)

Windows Deployment Services

Windows Deployment Services is a technology from Microsoft for network-based installation of Windows operating systems. It is the successor to Remote Installation Services.[1] WDS is intended to be used for remotely deploying Windows Vista, Windows 7 and Windows Server 2008, but also supports other operating systems because unlike its predecessor RIS, which was a method of automating the installation process, WDS uses disk imaging, in particular the Windows Imaging Format (WIM). WDS is included as a *Server Role* in all 32-bit and 64-bit versions of Windows Server 2008, and is included as an optionally installable component with Windows Server 2003 Service Pack 2.

Overview

The Windows Deployment Service is the combined updated and redesigned versions of Remote Installation Service (RIS) and Automated Deployment Services (ADS). The deployment of Windows 7, Windows Vista, Windows Server 2008, Windows Server 2003, and Windows XP can be fully automated and customized through the use of unattended installation scripting files. Tasks that can be made automatic include naming the machine, having the machine join a domain, adding or removing programs and features, and installing server roles (in the case of Windows Server 2008). Windows Vista and Windows Server 2008 are installed from a set of source files on the server, often copied from the product's installation media.

WDS expands upon simple scripted installation by giving the technician the ability to capture, store, and deploy image-based installation packages. A major new feature available in the Windows Server 2008 versions of WDS is that it supports IP Multicast deployments. Multicasting allows new clients to join an existing multicast deployment that has already started; the WDS server will wrap the multicast so that any client who joined the deployment after it started can receive data it is missing. WDS's multicast uses the standard internet protocol IGMP. WDS also supports x64-based computers with Extensible Firmware Interface (EFI).

WDS contains the ability to deploy other operating systems such as Windows PE, Windows XP, and Windows 2000, but the installation of these operating systems cannot be performed with source files or controlled with unattended scripts. The unsupported operating system to be deployed must first be installed and configured on a workstation; an image of the finished operating system configuration is then captured with the Windows Automated Installation Kit, and this captured image can be deployed through WDS.

There are also several types of Windows 7 Migration Software that conduct an in-depth analysis of what is present on the machine. These software back up existing software, then asess and prepare for new software by conducting hardware and software inventories, deploy the software along with pre and post installation tasks, and distribute maintenance updates as part of a patch management strategy.[2]

Automated image capture and apply

WDS functions in conjunction with the Preboot Execution Environment (PXE) to load a miniature version of Windows known as Windows PE for installation and maintenance tasks. WDS functions as both a storage repository for the PXE network boot images as well as a repository for the actual operating system images to be installed on the target computer.

When multiple boot images are available, PXE booting via WDS will present the end-user with a boot menu to select the image to load.

Windows PE automation using WAIK

To simplify the tasks of capturing and applying images, two special scripted Windows PE boot images can be created which automate these tasks. These scripted Windows PE boot images are created using the Windows Automated Installation Kit, in combination with Windows 7 installation media containing the source WIM images, and then added to the WDS server's boot image repository. The Windows PE boot images may be either 32- or 64-bit, but 32-bit tends to be more universally compatible across all potential hardware types.

A difficulty of Windows PE booting is that it needs to include network drivers and disk controller drivers intended to work with the target hardware to be imaged. The process of adding drivers to the Windows PE boot image can be automated using the WDS server console:

1. Select the source WIM image, which may be either a new one created from original Windows 7 installation DVDs (32- or 64-bit), or a previously configured WIM.
2. Select the drivers to install into the WIM
3. WDS mounts the WIM to a virtual path, adds drivers to the virtual path, and generates a new WIM
4. The updated WIM image is added to the boot image section of the WDS repository

This process can be repeated at a later time when a new system type needs to be captured but the current Windows PE Capture boot image does not include network drivers for it. The boot image is updated with the additional drivers using the WDS interface and automatically re-added to the WDS boot image collection to replace the original.

For specialty one-off systems this WIM driver update process is not necessary if the hard drive of the target system to be captured is removed from the source system after sysprepping, and is either installed in a computer with supported network drivers, or attached to the supported system using an external "USB to hard drive" adapter.

Automated capture process

Using the Windows PE Capture boot image, the general imaging process is:

1. Build the original source OS, which may be Windows XP, Vista, or Windows 7.
2. Run Sysprep on the source OS, and reboot. Sysprep can be run by itself, or can use an auto-installation script to automate the image's first-run setup
3. As the computer starts, enter the boot menu and select network booting
4. Run the Windows PE Capture boot image
5. Select the sysprepped drive to be imaged (the automated Windows PE capture tool will not allow capture of systems that have not been sysprepped)
6. Enter a name and description for the image, where to store the WIM locally and what to call the image file.
7. If desired, upload the image to a WDS server after imaging completes. Enter the server name, and log on.
8. Start the capture

The image creation process involves first creating a compressed WIM locally before it is uploaded to the WDS repository. This requires enough local free space on the source system to fully compress the operating system and all programs before the upload occurs. For laptops and portables with limited storage, an external USB drive can be used for temporary WIM storage. As designed, Microsoft does not allow the WIM to be written to a temporary network storage location before upload to WDS.

The automated capture process uses a mechanism called data deduplication to significantly reduce the amount of space needed on a file server for captured images. Two WIMs are used to store each uploaded WDS image:

- a special "Resource.WIM" shared across all images
- a second WIM containing the specific file names, dates, and filesystem structure for each individual system

The upload process involves comparing the single captured WIM with the data already stored in the Resource.WIM on the WDS server, and generating the secondary WIM containing the specific differences between the captured WIM and the data already in the Resource.WIM. Disk storage for all images may be reduced by 50% to 95%

depending on the amount of operating system and program data duplicated across the images.

Automated apply process

Applying a captured image involves running a second Windows PE "Apply" boot image on the target system to receive the image. This boot image also needs the appropriate network and disk controller drivers as with the Windows PE Capture boot image.

1. The system is booted using PXE network booting and the Windows PE Apply image is loaded.
2. The operator logs on to the domain, and selects the boot image to apply.
3. A disk partitioning screen appears and the location for the target image is selected. If the target storage is unformatted, a default partition set is created. For Vista and Windows 7, a small 100 megabyte boot partition is created for storing bootloader data separate from the rest of the system partition. This boot partition is normally hidden from the Windows Vista/7 user.
4. The image data is applied to the selected partition, and the system reboots, either running the Sysprep manual mini-setup process or following the script created during the initial Sysprepping.

The WDS image creator may optionally select a separate WAIK / Sysprep installation script to be applied to the image during the first boot. This alternate script is selected within WDS by viewing the properties of each uploaded system image.

WDS automation and dual-boot systems

WDS automated capture and apply do not directly support dual-boot operating systems. Only one operating system at a time can be captured and deployed, and the automated capture process only deals with single partitions containing a sysprepped Windows OS installation. It will not capture data-only partitions.

However, WDS automated capture and apply can be used to duplicate and create dual-boot systems, if each OS is captured and applied separately. Generally, Windows XP and older operating systems need to be applied to the target system before Vista or Windows 7, due to the different boot loaders used by the newer operating systems.

Manual image capture and deploy

It is technically possible to create scripts that manually perform the imaging, capture, and apply processes, using command line tools provided by Microsoft. However, the methods for doing this are complex and difficult.

In general, the tools involved are:

- **dism** - Deployment Image Servicing and Management, used to add drivers to Windows PE boot images.
- **imagex** - used to capture and apply images. Creates either a single WIM structure, or can deduplicate data using a second shared resource WIM. Does not require a Windows Deployment Server to capture or apply images, and can work solely with a logged-on network share or mapped drive letter.
- **wdsutil** - used to manage the WDS server without the graphical user interface, and to add captured images to the repository.

Using imagex to manually create a WIM does not require the source operating system to be sysprepped or for the source partition to contain a Windows operating system. Any type of Windows-accessible file system can be imaged, including MSDOS, but the source system either needs to be able to run Windows PE or the source system's hard drive is moved into a newer system that supports Windows PE.

Microsoft generally requires Windows 2000, XP, Vista, and Windows 7 to be sysprepped before imaging, due to certain security-related disk data that Microsoft requires to be unique across duplicated system images. Sysprep randomizes this data when the image is applied to a new system.

Imagex does not have any disk formatting and partitioning capabilities. Separate Windows command line tools such as **diskpart** are needed to define partitions on the target system for imagex to use.

References

[1] "Windows Deployment Services" (http://msdn2.microsoft.com/en-us/library/aa967394.aspx). . Retrieved 2007-07-21.
[2] "Windows 7 Deployment Services and Windows 7 Migration Needs" (http://www.kace.com/solutions/windows-7-migration.php). Dell KACE. .

Configuration Management: One Host

CFEngine

CFEngine

Developer(s)	Mark Burgess, CFEngine AS [1]
Stable release	3.3.4 / June 21, 2012
Operating system	Cross-platform
Platform	Unix, Linux, Windows
Type	Configuration management System administration Network management
License	GNU General Public License
Website	www.cfengine.com [1]

CFEngine is a popular open source configuration management system, written by Mark Burgess. Its primary function is to provide automated configuration and maintenance of large-scale computer systems, including the unified management of servers, desktops, embedded networked devices, mobile smartphones, and tablet computers.

History: CFEngine 1 & CFEngine 2

The CFEngine project began in 1993 as a way for author Mark Burgess (then a post-doctoral fellow of the Royal Society at Oslo University, Norway) to get his work done by automating the management of a small group of workstations in the Department of Theoretical Physics. Like many post-docs and PhD students, Burgess ended up with the task of managing Unix workstations, scripting and fixing problems for users manually. Scripting took too much time, the flavours of Unix were significantly different, and scripts had to be maintained for multiple platforms, drowning in exception logic.

After discussing the problems with a colleague, Burgess wrote the first version of CFEngine (*the configuration engine*) which was published as an internal report [2] and presented at the CERN computing conference. It gained significant attention from a wider community because it was able to hide platform differences using a domain-specific language.

A year later, Burgess finished his post-doc but decided to stay in Oslo and took at job lecturing at Oslo University College. Here he realized that there was little or no research being done into configuration management, and he set about applying the principles of scientific modelling to understanding computer systems. In a short space of time, he developed the notion of convergent operators, which remains a core of CFEngine.

In 1998, dissatisfied with the level of understanding in the area and the ad hoc discussions of computer security at the time, Burgess wrote "*Computer Immunology*" [3], a paper at the USENIX/LISA08 conference. It laid out a manifesto for creating self-healing systems, reiterated a few years later by at IBM in their form of Autonomic Computing. This started a research effort which led to a major re-write, **CFEngine 2**, which added features for machine learning, anomaly detection and secure communications.

CFEngine 3: Promise Theory

Between 1998 and 2004, CFEngine grew in adoption along with the popularity of Linux as a computing platform.

During this time, Mark Burgess developed Promise Theory, a model of distributed cooperation for self-healing automation.[4]

In 2008, after more than five years of research, **CFEngine 3** was introduced, which incorporated Promise Theory as "a way to make CFEngine both simpler and more powerful at the same time", according to Burgess. The most significant re-write of the project to date, CFEngine 3 also integrated knowledge management and discovery mechanisms—allowing configuration management to scale to automate enterprise-class infrastructure.

Portability

CFEngine provides an operating system-independent interface to Unix-like host configuration. It requires some expert knowledge to deal with peculiarities of different operating systems, but has the power to perform maintenance actions across multiple Unix-like hosts. CFEngine can be used on Windows hosts, and is widely used for managing large numbers of Unix hosts that run heterogeneous operating systems e.g. Solaris, Linux, AIX, and HPUX. Statistics collected by the supporting commercial company CFEngine AS [5] indicate hundreds of thousands of hosts running cfengine, with the largest sites recorded at 50,000.

Research-based

Shortly after its inception, CFEngine inspired a field of research into automated configuration management. The CFEngine project claims to attempt to place the problem of configuration management in a scientific framework. Its author Mark Burgess has developed a range of theoretical tools and results to talk about the problem, and has written several text books and monographs explaining them.

Commercialization

In June 2008 the company CFEngine AS [1] was formed as a collaboration between author Mark Burgess, Oslo University College and the Oslo Innovation Centre in order to support users of CFEngine. In April 2009, the company launched the first commercial version of CFEngine - CFEngine Nova [6]. Current version of CFEngine Nova is 3. February 2011, the company received its first round of funding, from FERD Capital [7]. The company has offices in Oslo, Norway and Palo Alto, California, United States of America. Haavard Nord, one of the founders of Qt, is the chairman, Thomas Ryd the CEO, and Mark Burgess acts as the CTO.

Convergence

One of the main ideas in CFEngine is that changes in computer configuration should be carried out in a *convergent* manner.[8][9] This means that each change operation made by the agent should have the character of a fixed point. Rather than describing the steps needed to make a change, CFEngine language describes the final state in which one wants to end up. The agent then ensures that the necessary steps are taken to end up in this "policy compliant state". Thus, CFEngine can be run again and again, whatever the initial state of a system, and it will end up with a predictable result. CFEngine supports the item of statistical compliance with policy, meaning that a system can never guarantee to be exactly in an ideal or desired state, rather one approaches (converges) towards the desired state by best-effort, at a rate that is determined by the ratio of the frequency of environmental change to the rate of CFEngine execution.[10]

User base

CFEngine is used in both large and small companies, as well as in many universities and governmental institutions. Sites as large as 50,000 machines are reported, while sites of several thousand hosts running under cfengine are common. According to statistics from the Cfengine AS [1], probably several million computers run CFEngine around the world, and users from more than 100 countries have been registered. Facebook is also a user of CFEngine.

References

[1] http://www.cfengine.com
[2] http://www.iu.hio.no/~mark/papers/cfengine_history.pdf
[3] http://www.usenix.org/publications/library/proceedings/lisa98/full_papers/burgess/burgess.pdf
[4] M. Burgess and A. Couch, Autonomic Computing Approximated by Fixed-Point Promises, Proceedings of First IEEE International Workshop on Modelling Autonomic Communication Environments (MACE2006). p197-222. pdf (http://research.iu.hio.no/papers/MACE-v2.pdf)
[5] http://www.cfengine.com/
[6] http://www.cfengine.com/nova
[7] http://www.ferd.no/lang/en/show.do?page=278;619
[8] M. Burgess, Cfengine: a site configuration engine, USENIX Computing systems, Vol8, No. 3 1995 pdf (http://www.iu.hio.no/~mark/papers/paper1.pdf)
[9] M. Burgess, Configurable immunity for evolving human-computer systems, Science of Computer Programming 51 2004, p197-213 pdf (http://research.iu.hio.no/papers/immunity.pdf)
[10] Burgess M (2003). "On the theory of system administration" (http://research.iu.hio.no/papers/sysadmtheory3.pdf) (PDF). *Science of Computer Programming* **49**: 1-46. .

External links

- Official website (http://www.cfengine.com/)
- Mark Burgess's talk about CFEngine @ Google Tech Talks (https://www.youtube.com/watch?v=4CCXs4Om5pY) on YouTube
- CFEngine Introduction (Spanish) (http://woop.es/2011/06/introduccion-cfengine/)

Infrastructure Basics: DHCP and DNS

Network Time Protocol

Network Time Protocol (**NTP**) is a networking protocol for clock synchronization between computer systems over packet-switched, variable-latency data networks.

In operation since before 1985, NTP is one of the oldest Internet protocols in use. NTP was originally designed by David L. Mills of the University of Delaware, who still develops and maintains it with a team of volunteers.

NTP uses the User Datagram Protocol (UDP) on port number 123.

Overview

NTP provides Coordinated Universal Time (UTC) including scheduled leap second adjustments. No information about time zones or daylight saving time is transmitted; this information is outside its scope and must be obtained separately.

NTP uses Marzullo's algorithm and is designed to resist the effects of variable latency. NTP can usually maintain time to within tens of milliseconds over the public Internet,[1] and can achieve 1 millisecond accuracy in local area networks under ideal conditions.[2]

As of June 2010, the current reference implementation is version 4 (NTPv4), which is a proposed standard as documented in RFC 5905. It succeeds version 3, specified in RFC 1305.

The protocol uses the User Datagram Protocol (UDP) on port number 123.[3]

A less complex implementation of NTP, using the same protocol but without requiring the storage of state over extended periods of time, is known as the **Simple Network Time Protocol** (SNTP). It is used in some embedded devices and in applications where high accuracy timing is not required (RFC 1361, RFC 1769, RFC 2030, RFC 4330 and RFC 5905).

NTP software implementations

Unix

For modern Unix systems, the NTP client is implemented as a daemon process that runs continuously in user space (ntpd). Because of sensitivity to timing, however, it is important to have the standard NTP clock phase-locked loop implemented in kernel space. All recent versions of Linux, BSD, Mac OS X, Solaris and AIX are implemented in this manner.

The NTP packet is a UDP datagram, carried on port 123.[4]

Microsoft Windows

Microsoft Windows NT 4.0 did not come with an NTP implementation. The reference implementation of NTP can be used on NT4 systems.[5]

All Microsoft Windows versions since Windows 2000 and Windows XP include the Windows Time Service ("w32time"),[6] which has the ability to sync the computer clock to an NTP server. The version in Windows 2000 and Windows XP only implements Simple NTP, and violates several aspects of the NTP version 3 standard.[7] Beginning with Windows Server 2003 and Windows Vista, a compliant implementation of full NTP is included.[8] "

However, Microsoft does not guarantee that the Windows Time Service will be particularly accurate and will not support even 1-second accuracy:

> The W32Time service is not a full-featured NTP solution that meets time-sensitive application needs. The W32Time service is primarily designed to do the following:

- Make the Kerberos version 5 authentication protocol work.
- Provide loose sync time for client computers.

> The W32Time service cannot reliably maintain sync time to the range of 1 to 2 seconds. Such tolerances are outside the design specification of the W32Time service.[9]

Clock strata

NTP uses a hierarchical, semi-layered system of levels of clock sources. Each level of this hierarchy is termed a *stratum* and is assigned a layer number starting with 0 (zero) at the top. The stratum level defines its distance from the reference clock and exists to prevent cyclical dependencies in the hierarchy. It is important to note that the stratum is not an indication of quality or reliability, it is common to find *stratum 3* time sources that are higher quality than other *stratum 2* time sources. This definition of *stratum* is also different from the notion of clock strata used in telecommunication systems.

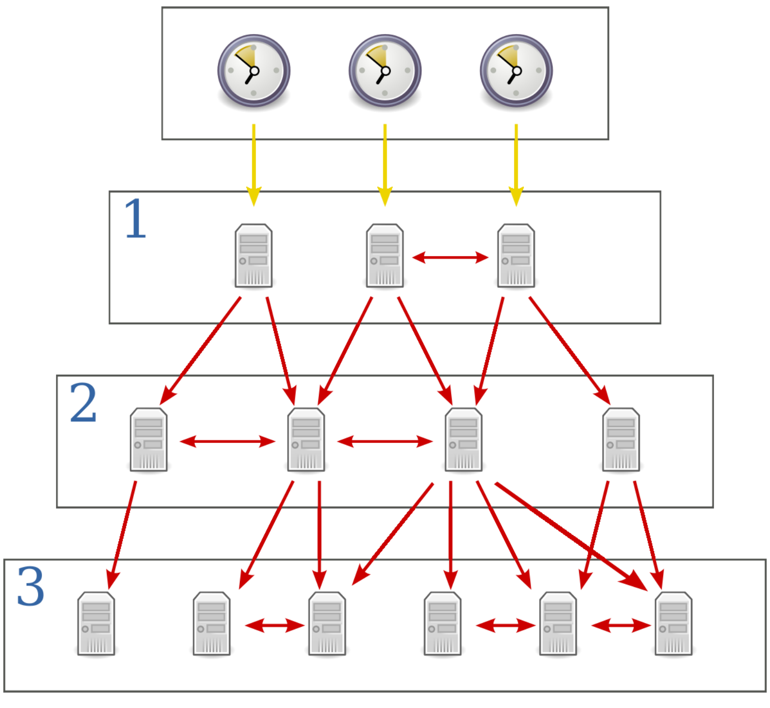

Yellow arrows indicate a direct connection; red arrows indicate a network connection.

Stratum 0

> These are devices such as atomic (caesium, rubidium) clocks, GPS clocks or other radio clocks. Stratum-0 devices are traditionally not attached to the network; instead they are locally connected to computers (e.g., via an RS-232 connection using a pulse per second signal).

Stratum 1

These are computers attached to Stratum 0 devices. Normally they act as servers for timing requests from Stratum 2 servers via NTP. These computers are also referred to as time servers.

The U.S. Naval Observatory Alternate Master Clock at Schriever AFB (Colorado) is a Stratum 0 source for NTP

Stratum 2

These are computers that send NTP requests to Stratum 1 servers. Normally a Stratum 2 computer will reference a number of Stratum 1 servers and use the NTP algorithm to gather the best data sample, dropping any Stratum 1 servers that seem obviously wrong. Stratum 2 computers will peer with other Stratum 2 computers to provide more stable and robust time for all devices in the peer group. Stratum 2 computers normally act as servers for Stratum 3 NTP requests.

Stratum 3

These computers employ exactly the same NTP functions of peering and data sampling as Stratum 2, and can themselves act as servers for lower strata.

While NTP (depending on what version of NTP protocol in use) supports up to 256 strata, only the first 16 are employed and any device at Stratum 16 is considered to be unsynchronized.

NTP timestamps

The 64-bit timestamps used by NTP consist of a 32-bit part for seconds and a 32-bit part for fractional second, giving NTP a time scale that rolls over every 2^{32} seconds (136 years) and a theoretical resolution of 2^{-32} seconds (233 picoseconds). NTP uses an epoch of January 1, 1900. The first rollover occurs in 2036, prior to the UNIX year 2038 problem.

Implementations should disambiguate NTP time using a knowledge of the approximate time from other sources. Since NTP only works with the differences between timestamps and never their absolute values, the wraparound is invisible as long as the timestamps are within 68 years of each other. This means that the rollover will be invisible for most running systems, since they will have the correct time to within a very small tolerance. However, systems that are starting up need to know the date within no more than 68 years. Given the large allowed error, it is not expected that this is too onerous a requirement. One suggested method is to set the clock to no earlier than the system build date. Many systems use a battery powered hardware clock to avoid this problem.

Even so, future versions of NTP may extend the time representation to 128 bits: 64 bits for the second and 64 bits for the fractional-second. The current NTP4 format has support for *Era Number* and *Era Offset*, that when used properly should aid fixing date rollover issues. According to Mills, "The 64 bit value for the fraction is enough to resolve the amount of time it takes a photon to pass an electron at the speed of light. The 64 bit second value is enough to provide unambiguous time representation until the universe goes dim."[10][11]

Clock synchronization algorithm

To synchronize its clock with a remote server, the NTP client must compute the round-trip delay time and the offset. The round-trip delay is computed as $\delta = (t_3 - t_0) - (t_2 - t_1)$, where t_0 is the time of the request packet transmission, t_1 is the time of the request packet reception, t_2 is the time of the response packet transmission and t_3 is the time of the response packet reception. $t_3 - t_0$ is the time elapsed on the client side between the emission of the request packet and the reception of the response packet, while $t_2 - t_1$ is the time the server waited before sending the answer. The offset is given by $\theta = \frac{(t_1 - t_0) + (t_2 - t_3)}{2}$.[12]

The NTP synchronization is correct when both the incoming and outgoing routes between the client and the server have symmetrical nominal delay. If the routes do not have a common nominal delay, the synchronization has a systematic bias of half the difference between the forward and backward travel times.[13]

Leap seconds

NTP delivers UTC time. UTC is subject to scheduled leap seconds to synchronize the timescale to the rotation of the earth. When a leap second is added, NTP is suspended for 1 second. Because NTP has no mechanism for remembering the history of leap seconds, leap seconds cause the entire NTP timescale to shift by 1 second.[14][15]

Security concerns

Only a few security problems have been identified in the reference implementation of the NTP codebase in its 25+ year history.[16][17] The protocol has been undergoing revision and review over its entire history. As of January 2011, there are no security revisions in the NTP specification and no reports at CERT.[18] The current codebase for the reference implementation has been undergoing security audits from several sources for several years now, and there are no known high-risk vulnerabilities in the current released software.[19]

Notes

[1] "Executive Summary: Computer Network Time Synchronization" (http://www.eecis.udel.edu/~mills/exec.html). . Retrieved 2011-11-21.

[2] "NTP FAQ" (http://www.ntp.org/ntpfaq/NTP-s-algo.htm#Q-ACCURATE-CLOCK). The NTP Project. . Retrieved 2011-08-27.

[3] "Port Numbers" (http://www.iana.org/assignments/port-numbers). The Internet Assigned Numbers Authority (IANA). .

[4] Page 16 (http://tools.ietf.org/html/rfc5905)

[5] "NTP 4.x for Windows NT" (http://doc.ntp.org/4.2.6/hints/winnt.html). NTP Documentation Archive. . Retrieved 2011-01-12.

[6] "Windows Time Service Technical Reference" (http://technet.microsoft.com/en-us/library/cc773061(WS.10).aspx). technet.microsoft.com. 2011-08-17. . Retrieved 2011-09-19.

[7] "Windows Time Service page at NTP.org" (http://support.ntp.org/bin/view/Support/WindowsTimeService). Support.ntp.org. 2008-02-25. . Retrieved 2011-01-12.

[8] "How the Windows Time Service Works" (http://technet.microsoft.com/en-us/library/cc773013(WS.10).aspx). technet.microsoft.com. 2010-03-12. . Retrieved 2011-09-19.

[9] Support boundary to configure the Windows Time service for high accuracy environments (http://support.microsoft.com/kb/939322)

[10] University of Delaware Digital Systems Seminar presentation by David Mills, 2006-04-26

[11] 2^{-64} seconds is about 54 zeptoseconds (light would travel 16.26 picometres, or approximately 0.31 × Bohr radius), and 2^{64} seconds is about 585 billion years.

[12] Mills, David L. (2005), *NTP Clock Discipline Modelling and Analysis* (http://www.eecis.udel.edu/~mills/database/brief/algor/algor.pdf), p. 10, , retrieved 2011-07-13

[13] Gotoh, T; Imamura, K; Kaneko, A (2002). "Improvement of NTP time offset under the asymmetric network with double packets method" (http://ieeexplore.ieee.org/stamp/stamp.jsp?tp=&arnumber=1034915&userType=inst). Conference on Precision Electromagnetic Measurements. pp. 448–449. doi:10.1109/CPEM.2002.1034915.

[14] David Mills. "The NTP Timescale and Leap Seconds" (http://www.eecis.udel.edu/~mills/leap.html). . Retrieved 2011-09-15.

[15] David Mills. "NTP Timescale and Leap Seconds" (http://doc.ntp.org/4.1.2/leap.htm). . Retrieved 2011-09-15.

[16] "Security Notice" (http://support.ntp.org/security). Support.ntp.org. 2009-12-10. . Retrieved 2011-01-12.

[17] results returned by a search on "Network Time Protocol" (http://search.cert.org/search?q="Network+Time+Protocol"&btnG.x=0&btnG.y=0&btnG=Go&entqr=0&ud=1&sort=date:D:L:d1&output=xml_no_dtd&oe=UTF-8&ie=UTF-8&client=default_frontend&

proxystylesheet=default_frontend&site=default_collection) at http://cert.org]
[18] in RFC 778, RFC 891, RFC 956, RFC 958, RFC 1305, RFC 5905, and the NTPv4 specification (http://www.eecis.udel.edu/~mills/database/reports/ntp4/ntp4.pdf), and no results returned by a search on "Network Time Protocol" at http://cert.org]
[19] "Code Audit" (http://support.ntp.org/Main/CodeAudit). Support.ntp.org. 2009-06-13. . Retrieved 2011-01-12.

References

Relevant RFCs

- RFC 958, (Obsoleted by: 1059, 1119, 1305) Network Time Protocol (NTP)
- RFC 1059, (Obsoleted by: 1119, 1305) Network Time Protocol (Version 1) Specification and Implementation
- RFC 1119, (Obsoleted by: 1305) Network Time Protocol (Version 2) Specification and Implementation
- RFC 1305, (Obsoleted by: 5905) Network Time Protocol (Version 3) Specification, Implementation and Analysis
- RFC 1361, (Obsoleted by: 1769) Simple Network Time Protocol (SNTP)
- RFC 1769, (Obsoleted by: 2030, 4330) Simple Network Time Protocol (SNTP)
- RFC 2030, (Obsoleted by: 4330) Simple Network Time Protocol (SNTP) Version 4 for IPv4, IPv6 and OSI
- RFC 4330, (Obsoleted by: 5905) Simple Network Time Protocol (SNTP) Version 4 for IPv4, IPv6 and OSI
- RFC 5905, Network Time Protocol Version 4: Protocol and Algorithms Specification

Further reading

- Mills, David L.. *Computer Network Time Synchronization: The Network Time Protocol*. Taylor & Francis / CRC Press. ISBN 0-8493-5805-1.

External links

- Official site (http://ntp.org/)
  - NTP Public Services Project (http://support.ntp.org/) This is the home for the NTP (Network Time Protocol) Public Services Project that provides public support for the NTP Project and the IETF NTP Working Group. Also a Wiki with lists of NTP servers.
  - Pool of public NTP Time Servers (http://www.pool.ntp.org/)
  - For a comprehensive list of NTP Servers (http://support.ntp.org/bin/view/Servers/WebHome)
- IETF NTP working group (http://www.ietf.org/html.charters/ntp-charter.html)
- David Mills' NTP project page (http://www.eecis.udel.edu/~mills/ntp.html)
  - NTP version 4 Release Notes (http://www.eecis.udel.edu/~mills/ntp/html/release.html)
- Time and NTP paper (http://www.ijs.si/time/)
- NTP Survey 2005 (http://www.ntpsurvey.arauc.br/)

Domain Name System

The **Domain Name System** (**DNS**) is a hierarchical distributed naming system for computers, services, or any resource connected to the Internet or a private network. It associates various information with domain names assigned to each of the participating entities. A **Domain Name Service** resolves queries for these names into IP addresses for the purpose of locating computer services and devices worldwide. By providing a worldwide, distributed keyword-based redirection service, the Domain Name System is an essential component of the functionality of the Internet.

An often-used analogy to explain the Domain Name System is that it serves as the phone book for the Internet by translating human-friendly computer hostnames into IP addresses. For example, the domain name www.example.com translates to the addresses `192.0.43.10` (IPv4) and `2620:0:2d0:200::10` (IPv6). Unlike a phone book, however, DNS can be quickly updated and these updates distributed, allowing a service's location on the network to change without affecting the end users, who continue to use the same hostname. Users take advantage of this when they recite meaningful Uniform Resource Locators (URLs) and e-mail addresses without having to know how the computer actually locates the services.

The Domain Name System distributes the responsibility of assigning domain names and mapping those names to IP addresses by designating authoritative name servers for each domain. Authoritative name servers are assigned to be responsible for their particular domains, and in turn can assign other authoritative name servers for their sub-domains. This mechanism has made the DNS distributed and fault tolerant and has helped avoid the need for a single central register to be continually consulted and updated. Additionally, the responsibility for maintaining and updating the master record for the domains is spread among many domain name registrars, who compete for the end-user's, domain-owner's, business. Domains can be moved from registrar to registrar at any time.

The Domain Name System also specifies the technical functionality of this database service. It defines the DNS protocol, a detailed specification of the data structures and communication exchanges used in DNS, as part of the Internet Protocol Suite.

Overview

The Internet maintains two principal namespaces, the domain name hierarchy[1] and the Internet Protocol (IP) address spaces.[2] The Domain Name System maintains the domain name hierarchy and provides translation services between it and the address spaces. Internet name servers and a communication protocol implement the Domain Name System.[3] A DNS name server is a server that stores the DNS records for a domain name, such as address (A) records, name server (NS) records, and mail exchanger (MX) records (see also list of DNS record types); a DNS name server responds with answers to queries against its database.

History

The practice of using a name as a simpler, more memorable abstraction of a host's numerical address on a network dates back to the ARPANET era. Before the DNS was invented in 1982, each computer on the network retrieved a file called *HOSTS.TXT* from a computer at SRI (now SRI International).[4][5] The HOSTS.TXT file mapped names to numerical addresses. A hosts file still exists on most modern operating systems by default and generally contains a mapping of "localhost" to the IP address 127.0.0.1. Many operating systems use name resolution logic that allows the administrator to configure selection priorities for available name resolution methods.

The rapid growth of the network made a centrally maintained, hand-crafted HOSTS.TXT file unsustainable; it became necessary to implement a more scalable system capable of automatically disseminating the requisite information.

At the request of Jon Postel, Paul Mockapetris invented the Domain Name System in 1983 and wrote the first implementation. The original specifications were published by the Internet Engineering Task Force in RFC 882 and RFC 883, which were superseded in November 1987 by RFC 1034[1] and RFC 1035.[3] Several additional Request for Comments have proposed various extensions to the core DNS protocols.

In 1984, four Berkeley students—Douglas Terry, Mark Painter, David Riggle, and Songnian Zhou—wrote the first Unix implementation, called The Berkeley Internet Name Domain (BIND) Server.[6] In 1985, Kevin Dunlap of DEC significantly re-wrote the DNS implementation. Mike Karels, Phil Almquist, and Paul Vixie have maintained BIND since then. BIND was ported to the Windows NT platform in the early 1990s.

BIND was widely distributed, especially on Unix systems, and is the dominant DNS software in use on the Internet.[7] With the heavy use and resulting scrutiny of its open-source code, as well as increasingly more sophisticated attack methods, many security flaws were discovered in BIND. This contributed to the development of a number of alternative name server and resolver programs. BIND version 9 was written from scratch and now has a security record comparable to other modern DNS software.

Structure

Domain name space

The domain name space consists of a tree of domain names. Each node or leaf in the tree has zero or more *resource records*, which hold information associated with the domain name. The tree sub-divides into *zones* beginning at the root zone. A DNS zone may consist of only one domain, or may consist of many domains and sub-domains, depending on the administrative authority delegated to the manager.

Administrative responsibility over any zone may be divided by creating additional zones. Authority is said to be *delegated* for a portion of the old space, usually in the form of sub-domains, to another nameserver and administrative entity. The old zone ceases to be authoritative for the new zone.

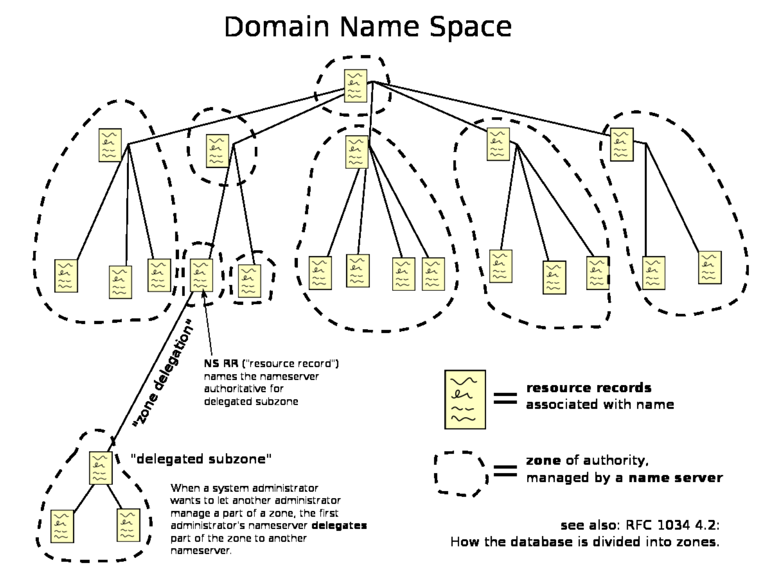

The hierarchical Domain Name System, organized into zones, each served by a name server

Domain name syntax

The definitive descriptions of the rules for forming domain names appear in RFC 1035, RFC 1123, and RFC 2181. A domain name consists of one or more parts, technically called *labels*, that are conventionally concatenated, and delimited by dots, such as `example.com`.

- The right-most label conveys the top-level domain; for example, the domain name `www.example.com` belongs to the top-level domain `com`.
- The hierarchy of domains descends from right to left; each label to the left specifies a subdivision, or subdomain of the domain to the right. For example: the label `example` specifies a subdomain of the `com` domain, and `www` is a sub domain of `example.com`. This tree of subdivisions may have up to 127 levels.

- Each label may contain up to 63 characters. The full domain name may not exceed a total length of 253 characters in its external dotted-label specification.[8] In the internal binary representation of the DNS the maximum length requires 255 octets of storage.[1] In practice, some domain registries may have shorter limits.
- DNS names may technically consist of any character representable in an octet. However, the allowed formulation of domain names in the DNS root zone, and most other sub domains, uses a preferred format and character set. The characters allowed in a label are a subset of the ASCII character set, and includes the characters *a* through *z*, *A* through *Z*, digits *0* through *9*, and the hyphen. This rule is known as the *LDH rule* (letters, digits, hyphen). Domain names are interpreted in case-independent manner.[9] Labels may not start or end with a hyphen.[10]
- A hostname is a domain name that has at least one IP address associated. For example, the domain names `www.example.com` and `example.com` are also hostnames, whereas the `com` domain is not.

Internationalized domain names

The permitted character set of the DNS prevented the representation of names and words of many languages in their native alphabets or scripts. ICANN has approved the Internationalizing Domain Names in Applications (IDNA) system, which maps Unicode strings into the valid DNS character set using Punycode. In 2009 ICANN approved the installation of IDN country code top-level domains. In addition, many registries of the existing top level domain names (TLD)s have adopted IDNA.

Name servers

The Domain Name System is maintained by a distributed database system, which uses the client-server model. The nodes of this database are the name servers. Each domain has at least one authoritative DNS server that publishes information about that domain and the name servers of any domains subordinate to it. The top of the hierarchy is served by the root nameservers, the servers to query when looking up (*resolving*) a TLD.

Authoritative name server

An *authoritative* name server is a name server that gives answers that have been configured by an original source, for example, the domain administrator or by dynamic DNS methods, in contrast to answers that were obtained via a regular DNS query to another name server. An authoritative-only name server only returns answers to queries about domain names that have been specifically configured by the administrator.

An authoritative name server can either be a *master* server or a *slave* server. A master server is a server that stores the original (*master*) copies of all zone records. A slave server uses an automatic updating mechanism of the DNS protocol in communication with its master to maintain an identical copy of the master records.

Every DNS zone must be assigned a set of authoritative name servers that are installed in NS records in the parent zone, and should be installed (to be authoritative records) as self-referential NS records on the authoritative name servers.

When domain names are registered with a domain name registrar, their installation at the domain registry of a top level domain requires the assignment of a *primary* name server and at least one *secondary* name server. The requirement of multiple name servers aims to make the domain still functional even if one name server becomes inaccessible or inoperable.[11] The designation of a primary name server is solely determined by the priority given to the domain name registrar. For this purpose, generally only the fully qualified domain name of the name server is required, unless the servers are contained in the registered domain, in which case the corresponding IP address is needed as well.

Primary name servers are often master name servers, while secondary name server may be implemented as slave servers.

An authoritative server indicates its status of supplying definitive answers, deemed *authoritative*, by setting a software flag (a protocol structure bit), called the *Authoritative Answer* (*AA*) bit in its responses.[3] This flag is usually reproduced prominently in the output of DNS administration query tools (such as dig) to indicate *that the responding name server is an authority for the domain name in question.*[3]

Operation

Address resolution mechanism

Domain name resolvers determine the appropriate domain name servers responsible for the domain name in question by a sequence of queries starting with the right-most (top-level) domain label.

The process entails:

1. A network host is configured with an initial cache (so called *hints*) of the known addresses of the root nameservers. Such a *hint file* is updated periodically by an administrator from a reliable source.
2. A query to one of the root servers to find the server authoritative for the top-level domain.
3. A query to the obtained TLD server for the address of a DNS server authoritative for the second-level domain.
4. Repetition of the previous step to process each domain name label in sequence, until the final step which returns the IP address of the host sought.

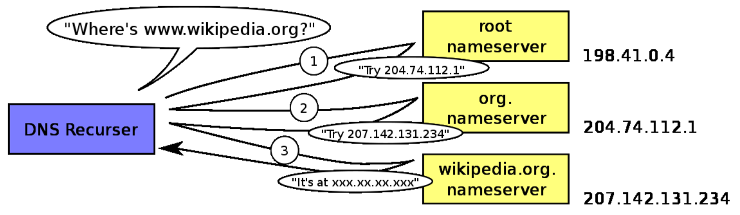

A DNS recursor consults three nameservers to resolve the address www.wikipedia.org.

The diagram illustrates this process for the host www.wikipedia.org.

The mechanism in this simple form would place a large operating burden on the root servers, with every search for an address starting by querying one of them. Being as critical as they are to the overall function of the system, such heavy use would create an insurmountable bottleneck for trillions of queries placed every day. In practice caching is used in DNS servers to overcome this problem, and as a result, root nameservers actually are involved with very little of the total traffic.

Recursive and caching name server

In principle, authoritative name servers are sufficient for the operation of the Internet. However, with only authoritative name servers operating, every DNS query must start with recursive queries at the root zone of the Domain Name System and each user system must implement resolver software capable of recursive operation.

To improve efficiency, reduce DNS traffic across the Internet, and increase performance in end-user applications, the Domain Name System supports DNS cache servers which store DNS query results for a period of time determined in the configuration (time-to-live) of the domain name record in question. Typically, such *caching* DNS servers, also called *DNS caches*, also implement the recursive algorithm necessary to resolve a given name starting with the DNS root through to the authoritative name servers of the queried domain. With this function implemented in the name server, user applications gain efficiency in design and operation.

The combination of DNS caching and recursive functions in a name server is not mandatory; the functions can be implemented independently in servers for special purposes.

Internet service providers typically provide recursive and caching name servers for their customers. In addition, many home networking routers implement DNS caches and recursors to improve efficiency in the local network.

DNS resolvers

The client-side of the DNS is called a DNS resolver. It is responsible for initiating and sequencing the queries that ultimately lead to a full resolution (translation) of the resource sought, e.g., translation of a domain name into an IP address.

A DNS query may be either a non-recursive query or a recursive query:

- A *non-recursive query* is one in which the DNS server provides a record for a domain for which it is authoritative itself, or it provides a partial result without querying other servers.
- A *recursive query* is one for which the DNS server will fully answer the query (or give an error) by querying other name servers as needed. DNS servers are not required to support recursive queries.

The resolver, or another DNS server acting recursively on behalf of the resolver, negotiates use of recursive service using bits in the query headers.

Resolving usually entails iterating through several name servers to find the needed information. However, some resolvers function more simply by communicating only with a single name server. These simple resolvers (called "stub resolvers") rely on a recursive name server to perform the work of finding information for them.

Circular dependencies and glue records

Name servers in delegations are identified by name, rather than by IP address. This means that a resolving name server must issue another DNS request to find out the IP address of the server to which it has been referred. If the name given in the delegation is a subdomain of the domain for which the delegation is being provided, there is a circular dependency. In this case the nameserver providing the delegation must also provide one or more IP addresses for the authoritative nameserver mentioned in the delegation. This information is called *glue*. The delegating name server provides this glue in the form of records in the *additional section* of the DNS response, and provides the delegation in the *answer section* of the response.

For example, if the authoritative name server for `example.org` is `ns1.example.org`, a computer trying to resolve `www.example.org` first resolves `ns1.example.org`. Since `ns1` is contained in `example.org`, this requires resolving `example.org` first, which presents a circular dependency. To break the dependency, the nameserver for the `org` top level domain includes glue along with the delegation for `example.org`. The glue records are address records that provide IP addresses for `ns1.example.org`. The resolver uses one or more of these IP addresses to query one of domain's authoritative servers, which allows it to complete the DNS query.

Record caching

The DNS Resolution Process reduces the load on individual servers by *caching* DNS request records for a period of time after a response. This entails the local recording and subsequent consultation of the copy instead of initiating a new request upstream. The time for which a resolver caches a DNS response is determined by a value called the time to live (TTL) associated with every record. The TTL is set by the administrator of the DNS server handing out the authoritative response. The period of validity may vary from just seconds to days or even weeks.

As a noteworthy consequence of this distributed and caching architecture, changes to DNS records do not propagate throughout the network immediately, but require all caches to expire and refresh after the TTL. RFC 1912 conveys basic rules for determining appropriate TTL values.

Some resolvers may override TTL values, as the protocol supports caching for up to 68 years or no caching at all. Negative caching, i.e. the caching of the fact of non-existence of a record, is determined by name servers authoritative for a zone which must include the Start of Authority (SOA) record when reporting no data of the requested type exists. The value of the *MINIMUM* field of the SOA record and the TTL of the SOA itself is used to establish the TTL for the negative answer.

Reverse lookup

A reverse lookup is a query of the DNS for domain names when the IP address is known. Multiple domain names may be associated with an IP address. The DNS stores IP addresses in the form of domain names as specially formatted names in pointer (`PTR`) records within the infrastructure top-level domain arpa. For IPv4, the domain is `in-addr.arpa`. For IPv6, the reverse lookup domain is `ip6.arpa`. The IP address is represented as a name in reverse-ordered octet representation for IPv4, and reverse-ordered nibble representation for IPv6.

When performing a reverse lookup, the DNS client converts the address into these formats, and then queries the name for a PTR record following the delegation chain as for any DNS query. For example, assume the IPv4 address `208.80.152.2` is assigned to Wikimedia. It is represented as a DNS name in reverse order like this: `2.152.80.208.in-addr.arpa`. When the DNS resolver gets a PTR (reverse-lookup) request, it begins by querying the root servers (which point to ARIN's servers for the `208.in-addr.arpa` zone). On ARIN's servers, `152.80.208.in-addr.arpa` is assigned to Wikimedia, so the resolver sends another query to the Wikimedia nameserver for `2.152.80.208.in-addr.arpa`, which results in an authoritative response.

Client lookup

Users generally do not communicate directly with a DNS resolver. Instead DNS resolution takes place transparently in applications such as web browsers, e-mail clients, and other Internet applications. When an application makes a request that requires a domain name lookup, such programs send a resolution request to the DNS resolver in the local operating system, which in turn handles the communications required.

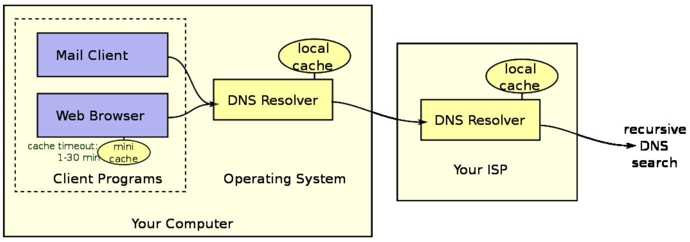

DNS resolution sequence

The DNS resolver will almost invariably have a cache (see above) containing recent lookups. If the cache can provide the answer to the request, the resolver will return the value in the cache to the program that made the request. If the cache does not contain the answer, the resolver will send the request to one or more designated DNS servers. In the case of most home users, the Internet service provider to which the machine connects will usually supply this DNS server: such a user will either have configured that server's address manually or allowed DHCP to set it; however, where systems administrators have configured systems to use their own DNS servers, their DNS resolvers point to separately maintained nameservers of the organization. In any event, the name server thus queried will follow the process outlined above, until it either successfully finds a result or does not. It then returns its results to the DNS resolver; assuming it has found a result, the resolver duly caches that result for future use, and hands the result back to the software which initiated the request.

Broken resolvers

An additional level of complexity emerges when resolvers violate the rules of the DNS protocol. A number of large ISPs have configured their DNS servers to violate rules (presumably to allow them to run on less-expensive hardware than a fully compliant resolver), such as by disobeying TTLs, or by indicating that a domain name does not exist just because one of its name servers does not respond.[12]

As a final level of complexity, some applications (such as web-browsers) also have their own DNS cache, in order to reduce the use of the DNS resolver library itself. This practice can add extra difficulty when debugging DNS issues, as it obscures the freshness of data, and/or what data comes from which cache. These caches typically use very short caching times—on the order of one minute.[13]

Internet Explorer represents a notable exception: versions up to IE 3.x cache DNS records for 24 hours by default. Internet Explorer 4.x and later versions (up to IE 8) decrease the default time out value to half an hour, which may be changed in corresponding registry keys.[14]

Other applications

The system outlined above provides a somewhat simplified scenario. The Domain Name System includes several other functions:

- Hostnames and IP addresses do not necessarily match on a one-to-one basis. Multiple hostnames may correspond to a single IP address: combined with virtual hosting, this allows a single machine to serve many web sites. Alternatively, a single hostname may correspond to many IP addresses: this can facilitate fault tolerance and load distribution, and also allows a site to move physical locations seamlessly.
- There are many uses of DNS besides translating names to IP addresses. For instance, Mail transfer agents use DNS to find out where to deliver e-mail for a particular address. The domain to mail exchanger mapping provided by MX records accommodates another layer of fault tolerance and load distribution on top of the name to IP address mapping.
- E-mail Blacklists: The DNS system is used for efficient storage and distribution of IP addresses of blacklisted e-mail hosts. The usual method is putting the IP address of the subject host into the sub-domain of a higher level domain name, and resolve that name to different records to indicate a positive or a negative. Here is a hypothetical example blacklist:
  - 102.3.4.5 is blacklisted => Creates 5.4.3.102.blacklist.example and resolves to 127.0.0.1
  - 102.3.4.6 is not => 6.4.3.102.blacklist.example is not found, or default to 127.0.0.2
  - E-mail servers can then query blacklist.example through the DNS mechanism to find out if a specific host connecting to them is in the blacklist. Today many of such blacklists, either free or subscription-based, are available mainly for use by email administrators and anti-spam software.
- Sender Policy Framework and DomainKeys, instead of creating their own record types, were designed to take advantage of another DNS record type, the TXT record.
- To provide resilience in the event of computer failure, multiple DNS servers are usually provided for coverage of each domain, and at the top level, thirteen very powerful root servers exist, with additional "copies" of several of them distributed worldwide via Anycast.
- Dynamic DNS (sometimes called DDNS) allows clients to update their DNS entry as their IP address changes, as it does, for example, when moving between ISPs or mobile hot spots.

Protocol details

DNS primarily uses User Datagram Protocol (UDP) on port number 53 to serve requests.[3] DNS queries consist of a single UDP request from the client followed by a single UDP reply from the server. The Transmission Control Protocol (TCP) is used when the response data size exceeds 512 bytes, or for tasks such as zone transfers. Some resolver implementations use TCP for all queries.

DNS resource records

A Resource Record (RR) is the basic data element in the domain name system. Each record has a type (A, MX, etc.), an expiration time limit, a class, and some type-specific data. Resource records of the same type define a resource record set (RRset). The order of resource records in a set, returned by a resolver to an application, is undefined, but often servers implement round-robin ordering to achieve Global Server Load Balancing. DNSSEC, however, works on complete resource record sets in a canonical order.

When sent over an IP network, all records use the common format specified in RFC 1035:[15]

RR (Resource record) fields

Field	Description	Length (octets)
NAME	Name of the node to which this record pertains	(variable)
TYPE	Type of RR in numeric form (e.g. 15 for MX RRs)	2
CLASS	Class code	2
TTL	Count of seconds that the RR stays valid (The maximum is 2^{31}-1, which is about 68 years.)	4
RDLENGTH	Length of RDATA field	2
RDATA	Additional RR-specific data	(variable)

NAME is the fully qualified domain name of the node in the tree. On the wire, the name may be shortened using label compression where ends of domain names mentioned earlier in the packet can be substituted for the end of the current domain name.

TYPE is the record type. It indicates the format of the data and it gives a hint of its intended use. For example, the *A* record is used to translate from a domain name to an IPv4 address, the *NS* record lists which name servers can answer lookups on a DNS zone, and the *MX* record specifies the mail server used to handle mail for a domain specified in an e-mail address (see also List of DNS record types).

RDATA is data of type-specific relevance, such as the IP address for address records, or the priority and hostname for MX records. Well known record types may use label compression in the RDATA field, but "unknown" record types must not (RFC 3597).

The *CLASS* of a record is set to `IN` (for *Internet*) for common DNS records involving Internet hostnames, servers, or IP addresses. In addition, the classes Chaos (`CH`) and Hesiod (`HS`) exist.[16] Each class is an independent name space with potentially different delegations of DNS zones.

In addition to resource records defined in a zone file, the domain name system also defines several request types that are used only in communication with other DNS nodes (*on the wire*), such as when performing zone transfers (AXFR/IXFR) or for EDNS (OPT).

Wildcard DNS records

The domain name system supports *wildcard domain names* which are names that start with the *asterisk label*, '*', e.g., `*.example`.[1][17] DNS records belonging to wildcard domain names specify rules for generating resource records within a single DNS zone by substituting whole labels with matching components of the query name, including any specified descendants. For example, in the DNS zone *x.example*, the following configuration specifies that all subdomains (including subdomains of subdomains) of *x.example* use the mail exchanger *a.x.example*. The records for *a.x.example* are needed to specify the mail exchanger. As this has the result of excluding this domain name and its subdomains from the wildcard matches, all subdomains of *a.x.example* must be defined in a separate wildcard statement.

The role of wildcard records was refined in RFC 4592, because the original definition in RFC 1034 was incomplete and resulted in misinterpretations by implementers.[17]

Protocol extensions

The original DNS protocol had limited provisions for extension with new features. In 1999, Paul Vixie published in RFC 2671 an extension mechanism, called Extension mechanisms for DNS (EDNS) that introduced optional protocol elements without increasing overhead when not in use. This was accomplished through the `OPT` pseudo-resource record that only exists in wire transmissions of the protocol, but not in any zone files. Initial extensions were also suggested (EDNS0), such as increasing the DNS message size in UDP datagrams.

Dynamic zone updates

Dynamic DNS updates use the `UPDATE` DNS opcode to add or remove resource records dynamically from a zone data base maintained on an authoritative DNS server. The feature is described in RFC 2136. This facility is useful to register network clients into the DNS when they boot or become otherwise available on the network. Since a booting client may be assigned a different IP address each time from a DHCP server, it is not possible to provide static DNS assignments for such clients.

Security issues

Originally, security concerns were not major design considerations for DNS software or any software for deployment on the early Internet, as the network was not open for participation by the general public. However, the expansion of the Internet into the commercial sector in the 1990s changed the requirements for security measures to protect data integrity and user authentication.

Several vulnerability issues were discovered and exploited by malicious users. One such issue is DNS cache poisoning, in which data is distributed to caching resolvers under the pretense of being an authoritative origin server, thereby polluting the data store with potentially false information and long expiration times (time-to-live). Subsequently, legitimate application requests may be redirected to network hosts operated with malicious intent.

DNS responses are traditionally not cryptographically signed, leading to many attack possibilities; the Domain Name System Security Extensions (DNSSEC) modify DNS to add support for cryptographically signed responses. Several extensions have been devised to secure zone transfers as well.

Some domain names may be used to achieve spoofing effects. For example, paypal.com and paypa1.com are different names, yet users may be unable to distinguish them in a graphical user interface depending on the user's chosen typeface. In many fonts the letter *l* and the numeral *1* look very similar or even identical. This problem is acute in systems that support internationalized domain names, since many character codes in ISO 10646, may appear identical on typical computer screens. This vulnerability is occasionally exploited in phishing.[18]

Techniques such as forward-confirmed reverse DNS can also be used to help validate DNS results.

Domain name registration

The right to use a domain name is delegated by domain name registrars which are accredited by the Internet Corporation for Assigned Names and Numbers (ICANN), the organization charged with overseeing the name and number systems of the Internet. In addition to ICANN, each top-level domain (TLD) is maintained and serviced technically by an administrative organization, operating a registry. A registry is responsible for maintaining the database of names registered within the TLD it administers. The registry receives registration information from each domain name registrar authorized to assign names in the corresponding TLD and publishes the information using a special service, the whois protocol.

ICANN publishes the complete list of TLD registries and domain name registrars. Registrant information associated with domain names is maintained in an online database accessible with the WHOIS service. For most of the more than 240 country code top-level domains (ccTLDs), the domain registries maintain the WHOIS (Registrant, name servers, expiration dates, etc.) information. For instance, DENIC, Germany NIC, holds the DE domain data. Since about 2001, most gTLD registries have adopted this so-called *thick* registry approach, i.e. keeping the WHOIS data in central registries instead of registrar databases.

For `COM` and `NET` domain names, a *thin* registry model is used: the domain registry (e.g. VeriSign) holds basic WHOIS (registrar and name servers, etc.) data. One can find the detailed WHOIS (registrant, name servers, expiry dates, etc.) at the registrars.

Some domain name registries, often called *network information centers* (NIC), also function as registrars to end-users. The major generic top-level domain registries, such as for the `COM`, `NET`, `ORG`, `INFO` domains, use a registry-registrar model consisting of many domain name registrars.[19][20] In this method of management, the registry only manages the domain name database and the relationship with the registrars. The *registrants* (users of a domain name) are customers of the registrar, in some cases through additional layers of resellers.

Internet standards

The Domain Name System is defined by Request for Comments (RFC) documents published by the Internet Engineering Task Force (Internet standards). The following is a list of RFCs that define the DNS protocol.

- RFC 920, *Domain Requirements* – Specified original top-level domains
- RFC 1032, *Domain Administrators Guide*
- RFC 1033, *Domain Administrators Operations Guide*
- RFC 1034, *Domain Names - Concepts and Facilities*
- RFC 1035, *Domain Names - Implementation and Specification*
- RFC 1101, *DNS Encodings of Network Names and Other Types*
- RFC 1123, *Requirements for Internet Hosts—Application and Support*
- RFC 1178, *Choosing a Name for Your Computer* (FYI 5)
- RFC 1183, *New DNS RR Definitions*
- RFC 1591, *Domain Name System Structure and Delegation* (Informational)
- RFC 1912, *Common DNS Operational and Configuration Errors*
- RFC 1995, *Incremental Zone Transfer in DNS*
- RFC 1996, *A Mechanism for Prompt Notification of Zone Changes (DNS NOTIFY)*
- RFC 2100, *The Naming of Hosts* (Informational)
- RFC 2136, *Dynamic Updates in the domain name system (DNS UPDATE)*
- RFC 2181, *Clarifications to the DNS Specification*
- RFC 2182, *Selection and Operation of Secondary DNS Servers*
- RFC 2308, *Negative Caching of DNS Queries (DNS NCACHE)*
- RFC 2317, *Classless IN-ADDR.ARPA delegation* (BCP 20)

- RFC 2671, *Extension Mechanisms for DNS (EDNS0)*
- RFC 2672, *Non-Terminal DNS Name Redirection*
- RFC 2845, *Secret Key Transaction Authentication for DNS (TSIG)*
- RFC 3225, *Indicating Resolver Support of DNSSEC*
- RFC 3226, *DNSSEC and IPv6 A6 aware server/resolver message size requirements*
- RFC 3597, *Handling of Unknown DNS Resource Record (RR) Types*
- RFC 3696, *Application Techniques for Checking and Transformation of Names* (Informational)
- RFC 4343, *Domain Name System (DNS) Case Insensitivity Clarification*
- RFC 4592, *The Role of Wildcards in the Domain Name System*
- RFC 4635, *HMAC SHA TSIG Algorithm Identifiers*
- RFC 4892, *Requirements for a Mechanism Identifying a Name Server Instance* (Informational)
- RFC 5001, *DNS Name Server Identifier (NSID) Option*
- RFC 5452, *Measures for Making DNS More Resilient against Forged Answers*
- RFC 5625, *DNS Proxy Implementation Guidelines* (BCP 152)
- RFC 5890, *Internationalized Domain Names for Applications (IDNA):Definitions and Document Framework*
- RFC 5891, *Internationalized Domain Names in Applications (IDNA): Protocol*
- RFC 5892, *The Unicode Code Points and Internationalized Domain Names for Applications (IDNA)*
- RFC 5893, *Right-to-Left Scripts for Internationalized Domain Names for Applications (IDNA)*
- RFC 5894, *Internationalized Domain Names for Applications (IDNA):Background, Explanation, and Rationale* (Informational)
- RFC 5895, *Mapping Characters for Internationalized Domain Names in Applications (IDNA) 2008* (Informational)
- RFC 5966, *DNS Transport over TCP - Implementation Requirements*
- RFC 6195, *Domain Name System (DNS) IANA Considerations* (BCP 42)

Security

- RFC 4033, *DNS Security Introduction and Requirements*
- RFC 4034, *Resource Records for the DNS Security Extensions*
- RFC 4035, *Protocol Modifications for the DNS Security Extensions*
- RFC 4509, *Use of SHA-256 in DNSSEC Delegation Signer (DS) Resource Records*
- RFC 4470, *Minimally Covering NSEC Records and DNSSEC On-line Signing*
- RFC 5011, *Automated Updates of DNS Security (DNSSEC) Trust Anchors*
- RFC 5155, *DNS Security (DNSSEC) Hashed Authenticated Denial of Existence*
- RFC 5702, *Use of SHA-2 Algorithms with RSA in DNSKEY and RRSIG Resource Records for DNSSEC*
- RFC 5910, *Domain Name System (DNS) Security Extensions Mapping for the Extensible Provisioning Protocol (EPP)*
- RFC 5933, *Use of GOST Signature Algorithms in DNSKEY and RRSIG Resource Records for DNSSEC*

References

[1] RFC 1034, *Domain Names - Concepts and Facilities*, P. Mockapetris, The Internet Society (November 1987)
[2] RFC 781, *Internet Protocol - DARPA Internet Program Protocol Specification*, Information Sciences Institute, J. Postel (Ed.), The Internet Society (September 1981)
[3] RFC 1035, *Domain Names - Implementation and Specification*, P. Mockapetris, The Internet Society (November 1987)
[4] RFC 3467, *Role of the Domain Name System (DNS)*, J.C. Klensin, J. Klensin (February 2003)
[5] Cricket Liu, Paul Albitz (2006). *DNS and BIND* (http://oreilly.com/catalog/9780596100575) (5th ed.). O'Reilly. p. 3. .
[6] Douglas Brian Terry, Mark Painter, David W. Riggle and Songnian Zhou, *The Berkeley Internet Name Domain Server* (http://www.eecs.berkeley.edu/Pubs/TechRpts/1984/5957.html), Proceedings USENIX Summer Conference, Salt Lake City, Utah, June 1984, pages 23–31.
[7] "DNS Server Survey" (http://mydns.bboy.net/survey/). .
[8] RFC 2181, *Clarifications to the DNS Specification*, R. Elz, R. Bush (July 1997)
[9] Network Working Group of the IETF, January 2006, RFC 4343: Domain Name System (DNS) Case Insensitivity Clarification
[10] RFC 3696, *Application Techniques for Checking and Transformation of Names*, J.C. Klensin, J. Klensin
[11] "Name Server definition at techterms.com" (http://www.techterms.com/definition/nameserver). .
[12] "Providers ignoring DNS TTL?" (http://ask.slashdot.org/story/05/04/18/198259/providers-ignoring-dns-ttl). Slashdot. 2005. . Retrieved 2012-04-07.
[13] Ben Anderson: Why Web Browser DNS Caching Can Be A Bad Thing (http://dyn.com/web-browser-dns-caching-bad-thing/)
[14] "How Internet Explorer uses the cache for DNS host entries" (http://support.microsoft.com/default.aspx?scid=KB;en-us;263558). Microsoft Corporation. 2004. . Retrieved 2010-07-25.
[15] RFC 5395, *Domain Name System (DNS) IANA Considerations*, D. Eastlake 3rd (November 2008), Section 3
[16] RFC 5395, *Domain Name System (DNS) IANA Considerations*, D. Eastlake 3rd (November 2008), p. 11
[17] RFC 4592, *The Role of Wildcards in the Domain Name System*, E. Lewis (July 2006)
[18] APWG. "Global Phishing Survey: Domain Name Use and Trends in 1H2010." 10/15/2010 apwg.org (http://www.apwg.org/reports/APWG_GlobalPhishingSurvey_1H2010.pdf.)
[19] ICANN accredited registrars (http://www.icann.org/registrars/accredited-list.html)
[20] VeriSign COM and NET registry (http://www.verisign.com/information-services/naming-services/com-net-registry/page_002166.html)

External links

- Vixie, Paul (2007-05-04). "DNS Complexity – Although it contains just a few simple rules, DNS has grown into an enormously complex system." (http://www.acmqueue.com/modules.php?name=Content&pa=showpage&pid=481). ACM Queue.
- Zytrax.com (http://www.zytrax.com/books/dns/), Open Source Guide – DNS for Rocket Scientists, an on-line technical.
- Create a zone file (http://www.zonefile.org/?lang=en) on zonefile.org
- Domain Name System (http://www.microsoft.com/dns) on Microsoft TechNet

Root name server

A **root name server** is a name server for the Domain Name System's root zone. It directly answers requests for records in the root zone and answers other requests returning a list of the designated authoritative name servers for the appropriate top-level domain (TLD). The root name servers are a critical part of the Internet because they are the first step in translating (resolving) human readable host names into IP addresses that are used in communication between Internet hosts.

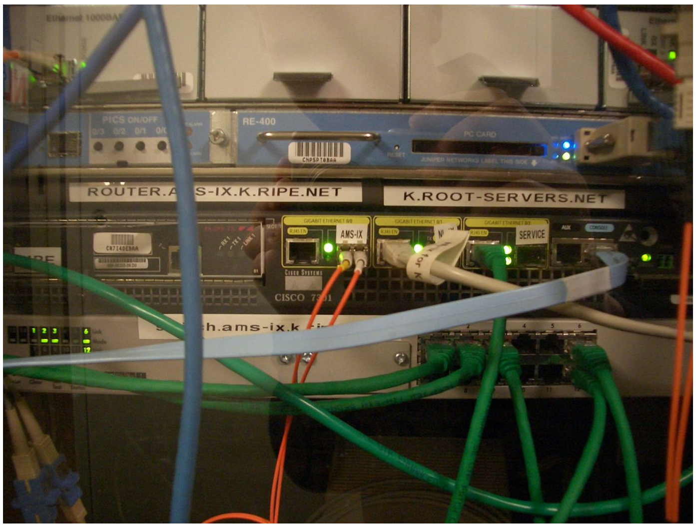

A Cisco 7301 router, part of the AMS-IX mirror of the K root-server.

A combination of limits in the DNS and certain protocols, namely the practical size of unfragmented User Datagram Protocol (UDP) packets, resulted in a limited number of root server addresses that can be accommodated in DNS name query responses. This limit has determined the number of name server installations at (currently) 13 clusters, serving the needs of the entire public Internet worldwide.

Root domain

The Domain Name System is a hierarchical naming system for computers, services, or any resource participating in the Internet. The top of that hierarchy is the root domain. The root domain does not have a formal name and its label in the DNS hierarchy is an empty string. All fully qualified domain names (FQDNs) on the Internet can be regarded as ending with this empty string for the root domain, and therefore ending in a full stop character (the label delimiter), e.g., `www.example.com.`. This is generally implied rather than explicit, as modern DNS software does not actually require that the terminating dot be included when attempting to translate a domain name to an IP address.

The root domain contains all top-level domains of the Internet. As of June 2009, there are 20 generic top-level domains (gTLDs) and 248 country code top-level domains (ccTLDs) in the root domain.[1] In addition, the `ARPA` domain is used for technical name spaces in the management of Internet addressing and other resources. A `TEST` domain is used for testing internationalized domain names.

Resolver operation

When a computer on the Internet needs to resolve a domain name, it uses resolver software to perform the lookup. A resolver breaks the name up into its labels from right to left. The first component (TLD) is queried using a root server to obtain the responsible authoritative server. Queries for each label return more specific name servers until a name server returns the answer of the original query.

In practice, most of this information does not change very often over a period of hours and therefore it is cached by intermediate name servers or by a name cache built into the user's application. DNS lookups to the root nameservers may therefore be relatively infrequent. A survey in 2003 [2] reports that only 2% of all queries to the root servers were legitimate. Incorrect or non-existent caching was responsible for 75% of the queries, 12.5% were for unknown TLDs, 7% were for lookups using IP addresses as if they were domain names, etc. Some misconfigured desktop computers even tried to update the root server records for the TLDs. A similar list of observed problems and recommended fixes has been published in RFC 4697.

Although any local implementation of DNS can implement its own private root name servers, the term "root name server" is generally used to describe the thirteen well-known root name servers that implement the root name space domain for the Internet's official global implementation of the Domain Name System.

Root server addresses

There are currently 13 root name servers specified, with names in the form `letter.root-servers.net`, where `letter` ranges from A to M. This does not mean there are 13 physical servers; each operator uses redundant computer equipment to provide reliable service even if failure of hardware or software occur. Additionally, nine of the servers operate in multiple geographical locations using a routing technique called anycast, providing increased performance and even more fault tolerance.

Ten servers were originally in the United States; some are now operated via anycast. Three servers were originally located in Stockholm (I), Amsterdam (K), and Tokyo (M).

Letter	IPv4 address	IPv6 address	AS-number[3]	Old name	Operator	Location #sites (global/local)[4]	Software
A	198.41.0.4	2001:503:ba3e::2:30	AS26415	ns.internic.net	Verisign	Distributed using anycast 6/0	BIND
B	192.228.79.201 (since January 2004; originally was 128.9.0.107)[5]	2001:478:65::53 (not in root zone yet)	AS4	ns1.isi.edu	USC-ISI	Marina Del Rey, California, U.S. 0/1	BIND
C [6]	192.33.4.12	2001:500:2::c (not in root zone yet)	AS2149	c.psi.net	Cogent Communications	Distributed using anycast 6/0	BIND
D [7]	128.8.10.90	2001:500:2d::d	AS27	terp.umd.edu	University of Maryland	College Park, Maryland, U.S. 1/0	BIND
E	192.203.230.10	N/A	AS297 [8]	ns.nasa.gov	NASA	Mountain View, California, U.S. 1/0	BIND
F [9]	192.5.5.241	2001:500:2f::f	AS3557	ns.isc.org	Internet Systems Consortium	Distributed using anycast 2/47	BIND 9[10]
G [11]	192.112.36.4	N/A	AS5927	ns.nic.ddn.mil	Defense Information Systems Agency	Distributed using anycast 6/0	BIND
H [12]	128.63.2.53	2001:500:1::803f:235	AS13	aos.arl.army.mil	U.S. Army Research Lab	Aberdeen Proving Ground, Maryland, U.S. 2/0	NSD
I [13]	192.36.148.17	2001:7fe::53	AS29216	nic.nordu.net	Autonomica	Distributed using anycast 38	BIND

J	192.58.128.30 (since November 2002; originally was 198.41.0.10)	2001:503:c27::2:30	AS26415		Verisign	Distributed using anycast 63/7	BIND
K [14]	193.0.14.129	2001:7fd::1	AS25152 [15]		RIPE NCC	Distributed using anycast 5/13	NSD[16]
L [17]	199.7.83.42 (since November 2007; originally was 198.32.64.12)[18]	2001:500:3::42	AS20144 [19]		ICANN	Distributed using anycast 107	NSD[20]
M [21]	202.12.27.33	2001:dc3::35	AS7500 [22]		WIDE Project	distributed using anycast 5/1	BIND

Older servers had their own name before the policy of using similar names was established.

The choice of 13 nameservers was made because of limitations in the original DNS specification, which specifies a maximum packet size of 512 bytes when using the User Datagram Protocol (UDP).[23] The addition of IPv6 addresses for the root nameservers requires more than 512 bytes, which is facilitated by the EDNS0 extension to the DNS standard.[24] While only 13 names are used for the root nameservers, there are many more physical servers; A, C, F, G, I, J, K, L and M servers now exist in multiple locations on different continents, using anycast address announcements to provide decentralized service. As a result most of the physical root servers are now outside the United States, allowing for high performance worldwide.

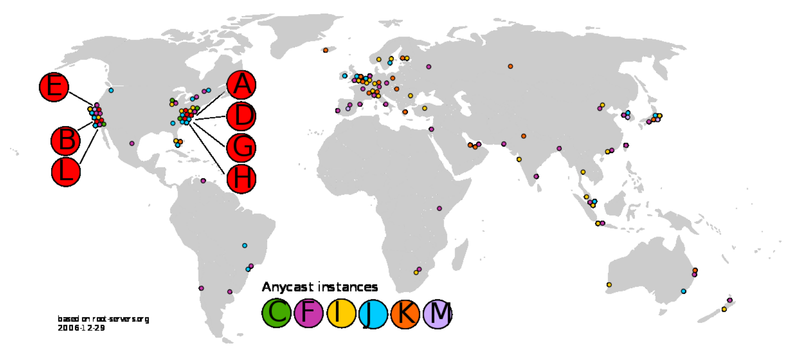

At the end of 2006 there were a total of 13 root nameservers, including Anycast servers.

There are also several alternative namespace systems with an alternative DNS root using their own set of root nameservers that exist in parallel to the mainstream nameservers. The first, AlterNIC, generated a substantial amount of press.

The function of a root name server may also be implemented locally, or on a provider network. Such servers are synchronized with the official root zone file as published [25] by ICANN, and do not constitute an alternate root.

As the root nameservers are an important part of the Internet, they have come under attack several times, although none of the attacks have ever been serious enough to severely affect the performance of the Internet.

Root server supervision

The DNS Root Server System Advisory Committee is an ICANN committee. However, the root zone is controlled by the United States Department of Commerce who must approve all changes to the root zone file requested by ICANN. ICANN's bylaws[26] assign authority over the operation of the root nameservers of the Domain Name System to the DNS Root Server System Advisory Committee.

Root zone file

The root zone file is a small (about 200 kB) data set[27] whose publication is the primary purpose of root nameservers.

The root zone file is at the apex of a hierarchical distributed database called the Domain Name System (DNS). This database is used by almost all Internet applications to translate worldwide unique names like www.wikipedia.org into other identifiers such as IP addresses.

The contents of the root zone file is a list of names and numeric IP addresses of the authoritative DNS servers for all top-level domains (TLDs) such as com, org, edu, or the country code top-level domains. On 12 December 2004, there were 258 TLDs and 773 different authoritative servers for those TLDs listed. Other name servers forward queries for which they do not have any information about authoritative servers to a root name server. The root name server, using its root zone file, answers with a referral to the authoritative servers for the appropriate TLD or with an indication that no such TLD exists.[28]

References

[1] "List of top-level domains" (http://data.iana.org/TLD/tlds-alpha-by-domain.txt). ICANN. .

[2] "Wow, That’s a Lot of Packets" (http://dns.measurement-factory.com/writings/wessels-pam2003-paper.pdf) (PDF). 2003. . Retrieved 2008-02-05.

[3] AS-numbers and IP-addresses from Root-servers.org homepage (http://root-servers.org/) checked 15 November 2010

[4] Location and sites from Root-servers.org homepage (http://root-servers.org/) checked 15 November 2010

[5] "New IPv4 address for b.root-servers.net" (http://www.root-servers.org/news/new-ip-b.html). .

[6] http://c.root-servers.org/

[7] http://d.root-servers.org/

[8] http://www.peeringdb.com/view.php?asn=297

[9] http://f.root-servers.org/

[10] F-root | Internet Systems Consortium (https://www.isc.org/community/f-root/)

[11] http://www.nic.mil/

[12] http://h.root-servers.org/

[13] http://i.root-servers.org/

[14] http://k.root-servers.org/

[15] http://www.peeringdb.com/view.php?asn=25152

[16] K-root Homepage (http://k.root-servers.org/)

[17] http://l.root-servers.org/

[18] "Advisory — "L Root" changing IP address on 1st November" (http://blog.icann.org/?p=227). ICANN. .

[19] http://www.peeringdb.com/view.php?asn=20144

[20] l.root-servers.net (http://l.root-servers.org/)

[21] http://m.root-servers.org/

[22] http://www.peeringdb.com/view.php?asn=7500

[23] RFC 1035 Domain names - implementation and specification

[24] ICANN: Accommodating IP Version 6 Address Resource Records for the Root of the Domain Name System (http://www.icann.org/committees/security/sac018.pdf)

[25] http://www.internic.net/zones/root.zone

[26] ICANN Bylaws XI-2.3 (http://www.icann.org/general/bylaws.htm#XI-2.3)

[27] IANA: Root Files (http://www.iana.org/domains/root/files)

[28] ISOC, *DNS Root Name Servers explained for the non-expert*, (Available online (http://www.isoc.org/briefings/019/), accessed 19 March 2010.)

Notes

- Root Server Technical Operations Association (http://www.root-servers.org/)
- Root Servers' Geographical Locations on Google Maps (http://www.circleid.com/posts/dns_root_servers_google_maps/)
- DNS Root Server System Advisory Committee (http://www.rssac.org/)
- DNS Root Name Servers Explained For Non-Experts (http://www.isoc.org/briefings/019/)
- DNS Root Name Servers Frequently Asked Questions (http://www.isoc.org/briefings/020/)
- Location of Root servers in Asia-Pacific (http://www.apnic.net/services/rootserver/)
- Bogus Queries received at the Root Servers (http://www.bind9.net/dnshealth/)
- ORSN, Open Root Server Network with IPv6 support in europe (http://www.orsn.org/)
- RFC 2826 - IAB Technical Comment on the Unique DNS Root
- RFC 2870 - Root Name Server Operational Requirements
- RFC 4697 - Observed DNS Resolution Misbehavior (from observations on the Root Servers)

External links

- Root Server Technical Operations Association (http://www.root-servers.org/)
- ftp://ftp.internic.net/domain/
- http://private.dnsstuff.com/info/roottimes.htm Root Server response times
- DNS root nameservers explained for non-experts (http://www.isoc.org/briefings/019/)

Dynamic DNS

Dynamic DNS or **DDNS** is a method of updating, in real time, a Domain Name System (DNS) to point to a changing IP address on the Internet. This is used to provide a persistent domain name for a resource that may change location on the network.

There are two very different mechanisms the term is used to describe. At the administration levels of the internet, "dynamic DNS updating" refers to systems that are used to update traditional DNS records without manual editing. These mechanisms are explained in RFC 2136 and use the TSIG mechanism to provide security. The second sort of DDNS is a particular type of DNS server that allows lightweight and immediate updates to its local database, often using a web-based form. These are used by individuals and small systems

Background

The Internet uses an IP address to refer to all resources on the Internet in a fashion similar to phone numbers. Like phone numbers, most IP addresses cannot be easily remembered, and are looked up in the internet equivalent of a phone book. This is the Domain Name System, or DNS. DNS automatically translates human-readable hostnames like "wikipedia.org" into the corresponding IP address, 208.80.152.201.

DNS is based on a distributed database that takes some time to update globally. When DNS was first introduced, the database was small and could be easily maintained by hand. As the system grew this task became difficult for any one site to handle, and a new management structure was introduced to spread out the updates among many domain name registrars. Registrars today offer end-user updating to their account information, typically using a web-based form, and the registrar then pushes out update information to other DNS servers.

Due to the distributed nature of the DNS systems and its registrars, updates to the global DNS system may take hours to distribute. Thus DNS is only suitable for services that do not change their IP address very often, as is the case for most large services like the Wikipedia or Google. Smaller services, however, are generally much more likely to move from host to host over shorter periods of time. Servers being run on certain types of internet service provider, cable modems in particular, are likely to change their IP address over very short periods of time, on the order of days or hours.

DDNS

Dynamic DNS is a system that addresses the problem of rapid updates. The term is used in two contexts which, while technically similar, have very different purposes and user populations.

Administrators who maintain the internet, and those who are associated with domains, use devices such as routers, gateways and computer systems using the Internet Protocol Suite, to keep DNS server information up to date. This process is **dynamic DNS updating**, and is often controlled directly by the domain owner by functionality presented by their registrar. This is the basic updating methodology of the modern DNS system, and has the slow-update problems noted above.

Standard users of the internet who connect to it via an Internet Service Provider (ISP) will be allocated a numeric IP address by the ISP. The address may either be constant ("static"), or may change from one session on the internet to the next ("dynamic"). If it is necessary to be able to access the computer from another location, a numeric address is inconvenient to remember, and an address which changes unpredictably makes connection next to impossible.

For these users there are a number of providers who use the same protocols as administrators to provide a **Dynamic DNS service** for end users. These are generally implemented in the user's router or computer, which notices changes to its IP address and automatically sends an update message to the DDNS service provider. The communication between the user's computer and the DDNS provider is not standardised, varying from one provider to another, although a few standard web-based methods of updating have emerged over time.

RFC 2136 dynamic DNS update

The standardized method of dynamically updating nameserver records is prescribed by RFC 2136, commonly known as 'Dynamic DNS Update' or 'DDNS'. Unlike updates through a DDNS provider, RFC 2136 is a protocol in its own right, with its own security mechanisms, and for use with managed nameservers. While RFC 2136 supports all DNS record types (including zone and user), it is most commonly used for dynamic hosts. In this form it is used primarily as an extension of the DHCP system, and in which the authorized DHCP servers register the clients' records with the nameserver(s) (Windows servers are an exception: by default, Windows servers only register 'A' records and the DHCP clients are expected to register the reverse pointers). This form of support for RFC 2136 is provided by a plethora of client and server software, including those that are components of most current operating systems. Support for RFC 2136 is also an integral part of many directory services, including LDAP and Windows' Active Directory domains.

History

In the initial stages of the Internet (ARPANET) addressing of hosts on the network was achieved by static translation tables that mapped hostnames to IP addresses. The tables were maintained manually in form of the hosts file. The Domain Name System brought a method of distributing the same address information automatically online through recursive queries to remote databases configured for each network, or domain. Even this DNS facility still used static lookup tables at each participating node. IP addresses, once assigned to a particular host, rarely changed and the mechanism was initially sufficient. However, the rapid growth of the Internet and the proliferation of personal computers in the workplace and in homes created the substantial burden for administrators of keeping track of assigned IP addresses and managing their address space. The Dynamic Host Configuration Protocol (DHCP) allowed enterprises and Internet service providers (ISPs) to assign addresses to computers automatically as they powered up. In addition, this helped conserve the address space available, since not all devices might be actively used at all times and addresses could be assigned as needed. This feature required that DNS servers be kept current automatically as well. The first implementations of *dynamic DNS* fulfilled this purpose: Host computers gained the feature to notify their respective DNS server of the address they had received from a DHCP server or through self-configuration. This protocol-based DNS update method was documented and standardized in IETF publication RFC 2136 in 1997 and has become a standard part of the DNS protocol (see also nsupdate program).

The explosive growth and proliferation of the Internet into people's homes brought a growing shortage of available IP addresses. DHCP became an important tool for ISPs as well to manage their address spaces for connecting home and small-business end-users with a single IP address each by implementing network address translation (NAT) at the customer premise router. The private network behind these routers uses address space set aside for these purposes (RFC 1918), masqueraded by the NAT device. This, however, broke the end-to-end principle of Internet architecture and methods were required to allow private networks, with frequently changing external IP addresses, to discover their public address and insert it into the Domain Name System in order to participate in Internet communications more fully. Today, numerous providers, called *Dynamic DNS* service providers, offer such technology and services on the Internet.

Function

In Microsoft Windows networks, dynamic DNS is an integral part of Active Directory, because domain controllers register their network service types in DNS so that other computers in the Domain (or Forest) can access them.

Increasing efforts to secure Internet communications today involve encryption of all dynamic updates via the public Internet, as these public dynamic DNS services have been abused increasingly to design security breaches. Standards-based methods within the DNSSEC protocol suite, such as TSIG, have been developed to secure DNS updates, but are not widely in use. Microsoft developed alternative technology (GSS-TSIG) based on Kerberos authentication.

Some freeware DNS server software supports a different dynamic update procedure. They have a built-in DHCP server which automatically updates or adds the DNS data internally with entries about addresses dynamically allotted by the DHCP server, without the user needing to configure dynamic updates. One such server is Dual DHCP DNS [1].

DDNS for ISP users

Dynamic DNS providers offer a software client program that automates the discovery and registration of the client system's public IP addresses. The client program is executed on a computer or device in the private network. It connects to the DDNS provider's systems with a unique login name; the provider uses the name to link the discovered public IP address of the home network with a hostname in the domain name system. Depending on the provider, the hostname is registered within a domain owned by the provider, or within the customer's own domain name. These services can function by a number of mechanisms. Often they use an HTTP service request since even restrictive environments usually allow HTTP service. The provider might use RFC 2136 to update the DNS servers.

Many home networking modem/routers have clients for several DDNS providers built into their firmware.

DDNS for Security Appliance Manufacturers

Dynamic DNS is an expected feature or even requirement for IP based security appliances like DVRs and IP cameras. Many options are available for today's manufacturer, and these include the use of existing DDNS services like dyn.com or no-ip.com or the use of custom services hosted by the manufacturer themselves. A manufacturer may choose to develop their own DDNS software or use a readymade software product like MintDNS or GNUDIP.

In almost all cases, a simple http based update API is used as it allows for easy integration of a DDNS client into a device's firmware. There are many pre-made tools that exist that can help to ease the burden of client development even further like cURL or Inadyn. Most web-based DDNS services use a standard user name and password security schema. This requires that a user first create an account at the DDNS server website and then configure their device to send updates to the DDNS server whenever an IP address change is detected. This is the method used by almost all web based DDNS services.

Some device manufacturers go a step further by only allowing their DDNS Service to be used by the devices they manufacture, and also eliminate the need for user names and passwords altogether. Generally this is accomplished by encrypting the device's MAC address using an encryption algorithm kept secret on both the DDNS server and within the device's firmware. The resulting decryption or decryption failure is used to secure or deny updates. Resources for the development of custom DDNS services are generally limited and involve a full software development cycle to design and field a secure and robust DDNS server. Some pre-made software products are available, though, and a few companies that specialize in the development of secure custom DDNS services do exist.

External links

- Dynamic DNS Services [2] at the Open Directory Project
- List of Dynamic DNS Providers [3]

References

[1] http://sourceforge.net/projects/dhcp-dns-server/
[2] http://www.dmoz.org/Computers/Software/Internet/Servers/Address_Management/Dynamic_DNS_Services/
[3] http://dnslookup.me/dynamic-dns/

Microsoft DNS

Microsoft DNS is the name given to the implementation of domain name system services provided in Microsoft Windows operating systems.

Overview

The Domain Name System support in Microsoft Windows NT, and thus its derivatives Windows 2000, Windows XP, and Windows Server 2003, comprises two clients and a server. Every Microsoft Windows machine has a DNS lookup client, to perform ordinary DNS lookups. Some machines have a Dynamic DNS client, to perform Dynamic DNS Update transactions, registering the machines' names and IP addresses. Some machines run a DNS server, to publish DNS data, to service DNS lookup requests from DNS lookup clients, and to service DNS update requests from DNS update clients.

The server software is only supplied with the server versions of Windows.

DNS lookup client

Applications perform DNS lookups with the aid of a DLL. They call library functions in the DLL, which in turn handle all communications with DNS servers (over UDP or TCP) and return the final results of the lookup back to the applications.

Microsoft's DNS client also has optional support for local caching, in the form of a *DNS Client* service (also known as *DNSCACHE*). Before they attempt to directly communicate with DNS servers, the library routines first attempt to make a local IPC connection to the DNS Client service on the machine. If there is one, and if such a connection can be made, they hand the actual work of dealing with the lookup over to the DNS Client service. The DNS Client service itself communicates with DNS servers, and caches the results that it receives.

Microsoft's DNS client is capable of talking to multiple DNS servers. The exact algorithm varies according to the version, and service pack level, of the operating system; but in general all communication is with a *preferred* DNS server until it fails to answer, whereupon communication switches to one of several *alternative* DNS servers.

The effects of running the DNS Client service

There are several minor differences in system behavior depending on whether the DNS Client service is started:

- **Parsing of the "hosts" file**: The lookup functions read only the hosts file if they cannot off-load their task onto the DNS Client service and have to fall back to communicating with DNS servers themselves. In turn, the DNS Client service reads the "hosts" file once, at startup, and only re-reads it if it notices that the last modification timestamp of the file has changed since it last read it. Thus:
  - With the DNS Client service running: The "hosts" file is read and parsed only a few times, once at service startup, and thereafter whenever the DNS Client service notices that it has been modified.
  - Without the DNS Client service running: The "hosts" file is read and parsed repeatedly, by each individual application program as it makes a DNS lookup.
- **The effect of multiple answers in the "hosts" file**: The DNS Client service does not use the "hosts" file directly when performing lookups. Instead, it (initially) populates its cache from it, and then performs lookups using the data in its cache. When the lookup functions fall back to doing the work themselves, however, they scan the "hosts" file directly and sequentially, stopping when the first answer is found. Thus:
  - With the DNS Client service running: If the "hosts" file contains multiple lines denoting multiple answers for a given lookup, all of the answers in the cache will be returned.

  - Without the DNS Client service running: If the "hosts" file contains multiple lines denoting multiple answers for a given lookup, only the first answer found will be returned.
- **Fallback from preferred to alternative DNS servers**: The fallback from the preferred DNS server to the alternative DNS servers is done by whatever entity, the DNS Client service or the library functions themselves, is actually performing the communication with them. Thus:
  - With the DNS Client service running: Fallback to the alternative DNS servers happens globally. If the preferred DNS server fails to answer, all subsequent communication is with the alternative DNS servers.
  - Without the DNS Client service running: Any fallback to the alternative DNS servers happen locally, within each individual process that is making DNS queries. Different processes may be in different states, some talking to the preferred DNS server and some talking to alternative DNS servers.

Differences from other systems

Linux distributions and various versions of Unix have a generalized name resolver layer. The resolver can be controlled to use a **hosts** file or Network Information Service (NIS), by configuring the Name Service Switch.

Dynamic DNS Update client

Whilst DNS lookups read DNS data, DNS updates *write* them. Both workstations and servers running Windows attempt to send Dynamic DNS update requests to DNS servers.

Workstations running Windows attempt to register their names and their IP addresses with DNS servers, so that other machines may locate them by name. Prior to Windows Vista (and Windows Server 2008) this registration is performed by the *DHCP Client* service. It is thus necessary to run the DHCP Client service on pre-Vista machines, even if DHCP isn't being used to configure the machine in order to dynamically register a machine's name and address for DNS lookup. The DHCP Client service registers name and address data whenever they are changed (either manually by an administrator or automatically by the granting or revocation of a DHCP lease). In Windows Vista (and Windows Server 2008) Microsoft moved the registration functionallity from the *DHCP Client* service to the *DNS Client* service.

Servers running Microsoft Windows also attempt to register other information, in addition to their names and IP addresses, such as the locations of the LDAP and Kerberos services that they provide.

DNS server

Microsoft Windows server operating systems can run the *DNS Server* service. This is a monolithic DNS server that provides many types of DNS service, including caching, Dynamic DNS update, zone transfer, and DNS notification. DNS notification implements a push mechanism for notifying a select set of secondary servers for a zone when it is updated.

Microsoft's "DNS Server" service was first introduced in Windows NT 3.51 as an add-on with Microsoft's collection of BackOffice services (at the time was marked to be used for testing purposes only). Some sources claim that the DNS server implementation in Windows NT 3.51 was a fork of ISC's BIND version 4.3, but this is not true. The DNS server implementation in Windows NT 3.51 was written by Microsoft. The DNS server component in all subsequent releases of Windows Server have built upon that initial implementation and do not use BIND source code. However, Microsoft has taken care to ensure good interoperability with BIND and other implementations in terms of zone file format, zone transfer, and other DNS protocol details.

As of 2004, it was the fourth most popular DNS server (counting BIND version 9 separately from versions 8 and 4) for the publication of DNS data.[1]

Like various other DNS servers, Microsoft's DNS server supports different database *back ends*. Microsoft's DNS server supports two such back ends. DNS data can be stored either in *master files* (also known as *zone files*) or in the

Active Directory database itself. In the latter case, since Active Directory (rather than the DNS server) handles the actual replication of the database across multiple machines, the database can be modified on any server ("multiple-master replication"), and the addition or removal of a *zone* will be immediately propagated to all other DNS servers within the appropriate Active Directory "replication scope". (Contrast this with BIND, where when such changes are made, the list of *zones*, in the `/etc/named.conf` file, has to be explicitly updated on each individual server.)

Microsoft's DNS server can be administered using either a graphical user interface, the "DNS Management Console", or a command line interface, the **dnscmd** utility. New to Windows Server 2012 is a fully featured PowerShell provider for DNS server management.

Common issues

Prior to Windows Server 2003 and Microsoft Windows 2000 Service Pack 3, the most common problem encountered with Microsoft's DNS server was cache pollution. Although Microsoft's DNS Server had a mechanism for properly dealing with cache pollution, the mechanism was turned off by default.[2]

In 2004, a common problem involved the feature of the Windows Server 2003 version of Microsoft's DNS server to use EDNS0, which a large number of firewalls could not cope with.[3]

References

[1] Moore, Don (2004). "DNS server survey" (http://mydns.bboy.net./survey/). . Retrieved 2005-01-06.
[2] How to prevent DNS cache pollution (http://support.microsoft.com/kb/241352)
[3] Microsoft (April 6, 2006). "An external DNS query may cause an error message in Windows Server 2003" (http://support.microsoft.com/kb/828731/). Microsoft. . Retrieved 2006-05-08.

External links

- Microsoft Domain Name System (DNS) (http://www.microsoft.com/dns)
- Jonathan de Boyne Pollard (2004). "Your firewall is preventing you from using EDNS0." (http://homepage.ntlworld.com./jonathan.deboynepollard/FGA/dns-edns0-and-firewalls.html). *Frequently Given Answers*. — the problems with EDNS0 and firewalls, and how to fix them
- Domain Name System Client Behavior in Windows Vista (http://www.microsoft.com/technet/network/ipv6/vista_dns.mspx)
- What's new in DNS in Windows Server 2008 (http://technet.microsoft.com/en-us/library/cc753143.aspx)
- What's new in DNS in Windows Server 2008 R2 (http://technet.microsoft.com/en-us/library/dd378952.aspx)

Active Directory

Active Directory (AD) is a directory service created by Microsoft for Windows domain networks. It is included in most Windows Server operating systems.

Active Directory provides a central location for network administration and security. Server computers that run Active Directory are called domain controllers. An AD domain controller authenticates and authorizes all users and computers in a Windows domain type network—assigning and enforcing security policies for all computers and installing or updating software. For example, when a user logs into a computer that is part of a Windows domain, Active Directory checks the submitted password and determines whether the user is a system administrator or normal user.[1]

Active Directory makes use of Lightweight Directory Access Protocol (LDAP) versions 2 and 3, Kerberos and DNS.

History

Active Directory was previewed in 1999, released first with Windows NT server,Then Windows 2000 Server edition, and revised to extend functionality and improve administration in Windows Server 2003. Additional improvements were made in Windows Server 2003 R2, Windows Server 2008 and Windows Server 2008 R2, and with the release of the later the domain controller role was renamed **Active Directory Domain Services**.

Structure

Objects

An Active Directory structure is a hierarchical arrangement of information about objects. The objects fall into two broad categories: resources (e.g., printers) and security principals (user or computer accounts and groups). Security principals are assigned unique security identifiers (SIDs).

Each object represents a single entity—whether a user, a computer, a printer, or a group—and its attributes. Certain objects can contain other objects. An object is uniquely identified by its name and has a set of attributes—the characteristics and information that the object represents— defined by a schema, which also determines the kinds of objects that can be stored in Active Directory.

The schema object lets administrators extend or modify the schema when necessary. However, because each schema object is integral to the definition of Active Directory objects, deactivating or changing these objects can fundamentally change or disrupt a deployment. Schema changes automatically propagate throughout the system. Once created, an object can only be deactivated—not deleted. Changing the schema usually requires planning.[2]

Sites

A **Site** object in Active Directory represents a geographic location that hosts networks.[3]

Forests, trees, and domains

The Active Directory framework that holds the objects can be viewed at a number of levels. The forest, tree, and domain are the logical divisions in an Active Directory network.

Within a deployment, objects are grouped into domains. The objects for a single domain are stored in a single database (which can be replicated). Domains are identified by their DNS name structure, the namespace.

A tree is a collection of one or more domains and domain trees in a contiguous namespace, linked in a transitive trust hierarchy.

At the top of the structure is the *forest.* A forest is a collection of trees that share a common global catalog, directory schema, logical structure, and directory configuration. The forest represents the security boundary within which users, computers, groups, and other objects are accessible.

Forest-Widgets Corp		
	Tree-Eastern	
		Domain-Boston
		Domain-New York
		Domain-Philly
	Tree-Southern	
		Domain-Atlanta
		Domain-Dallas

Domain-Dallas		
	OU-Marketing	
		Hewitt
		Aon
		Steve
	OU-Sales	
		Bill
		Ralph

Example of the geographical organizing of zones of interest within trees and domains.

Organizational units

The objects held within a domain can be grouped into Organizational Units (OUs).[4] OUs can provide hierarchy to a domain, ease its administration, and can resemble the organization's structure in managerial or geographical terms. OUs can contain other OUs—domains are containers in this sense. Microsoft recommends using OUs rather than domains for structure and to simplify the implementation of policies and administration. The OU is the recommended level at which to apply group policies, which are Active Directory objects formally named Group Policy Objects (GPOs), although policies can also be applied to domains or sites (see below). The OU is the level at which administrative powers are commonly delegated, but delegation can be performed on individual objects or attributes as well.

Organizational Units are an abstraction for the administrator and do not function as containers; the underlying domain is the true container. It is not possible, for example, to create user accounts with an identical username (sAMAccountName) in separate OUs, such as "fred.staff-ou.domain" and "fred.student-ou.domain", where "staff-ou" and "student-ou" are the OUs. This is so because sAMAccountName, a user object attribute, must be unique within the domain. However, two users in different OUs can have the same Common Name (CN), the first component of the Distinguished Name (DN) of the user. Thus from the point of view of the DN, OUs do function as containers.

As the number of users in a domain increases, conventions such as "first initial, middle initial, last name" (Western order) or the reverse (Eastern order) fail for common family names like *Li (李)*, *Smith* or *Garcia*. Workarounds include adding a digit to the end of the username. Alternatives include creating a separate ID system of unique employee/student id numbers to use as account names in place of actual user's names, and allowing users to nominate their preferred word sequence within an acceptable use policy.

Because duplicate usernames cannot exist within a domain, account name generation poses a significant challenge for large organizations that cannot be easily subdivided into separate domains, such as students in a public school system or university who must be able to use any computer across the network.

Shadow groups

In Microsoft's Active Directory, OUs do not confer access permissions, and objects placed within OUs are not automatically assigned access privileges based on their containing OU. This is a design limitation specific to Active Directory. Other competing directories such as Novell NDS are able to assign access privileges through object placement within an OU.

Active Directory requires a separate step for an administrator to assign an object in an OU as a member of a group also within that OU. Relying on OU location alone to determine access permissions is unreliable, because the object may not have been assigned to the group object for that OU.

A common workaround for an Active Directory administrator is to write a custom PowerShell or Visual Basic script to automatically create and maintain a *user group* for each OU in their directory. The scripts are run periodically to update the group to match the OU's account membership, but are unable to instantly update the security groups anytime the directory changes, as occurs in competing directories where security is directly implemented into the directory itself. Such groups are known as *Shadow Groups*. Once created, these shadow groups are selectable in place of the OU in the administrative tools.

Microsoft refers to shadow groups in the Server 2008 Reference documentation, but does not explain how to create them. There are no built-in server methods or console snap-ins for managing shadow groups.[5]

The division of an organization's information infrastructure into a hierarchy of one or more domains and top-level OUs is a key decision. Common models are by business unit, by geographical location, by IT Service, or by object type and hybrids of these. OUs should be structured primarily to facilitate administrative delegation, and secondarily, to facilitate group policy application. Although OUs form an administrative boundary, the only true security boundary is the forest itself and an administrator of any domain in the forest must be trusted across all domains in the forest.[6]

Physical matters

Sites are physical (rather than logical) groupings defined by one or more IP subnets.[7] AD also holds the definitions of connections, distinguishing low-speed (e.g., WAN, VPN) from high-speed (e.g., LAN) links. Site definitions are independent of the domain and OU structure and are common across the forest. Sites are used to control network traffic generated by replication and also to refer clients to the nearest domain controllers. Microsoft Exchange Server 2007 uses the site topology for mail routing. Policies can also be defined at the site level.

Physically the Active Directory information is held on one or more peer domain controllers (DCs), replacing the NT PDC/BDC model. Each DC has a copy of the Active Directory. Servers joined to Active Directory that are not domain controllers are called Member Servers.[8]

The Active Directory database is organized in *partitions*, each holding specific object types and following a specific replication pattern. AD synchronizes changes using *multi-master replication*.[9] Microsoft often refers to these partitions as 'naming contexts'.[10] The 'Schema' partition contains the definition of object classes and attributes within the Forest. The 'Configuration' partition contains information on the physical structure and configuration of the forest (such as the site topology). Both replicate to all domain controllers in the Forest. The 'Domain' partition holds all objects created in that domain and replicates only to Domain Controllers within its domain. So, for example, a user created in Domain X would be listed only in Domain X's domain controllers. A subset of objects in the domain partition replicate to domain controllers that are configured as global catalogs. Global catalog (GC) servers provide a global listing of all objects in the Forest.[11] Global Catalog servers replicate to themselves all objects from all domains and hence, provide a global listing of objects in the forest. However, to minimize replication traffic and keep the GC's database small, only selected attributes of each object are replicated. This is called the partial attribute set (PAS). The PAS can be modified by modifying the schema and marking attributes for replication to the GC.[12] Earlier versions of Windows used NetBIOS to communicate. Active Directory is fully integrated with DNS and requires TCP/IP—DNS. To be fully functional, the DNS server must support SRV resource records or service records.

Replication

Active Directory replication is 'pull' rather than 'push', meaning that replicas pull changes from the server where the change was effected.[13] The *Knowledge Consistency Checker* (KCC) creates a replication topology of *site links* using the defined *sites* to manage traffic. Intrasite replication is frequent and automatic as a result of change notification, which triggers peers to begin a pull replication cycle. Intersite replication intervals are typically less frequent and do not use change notification by default, although this is configurable and can be made identical to intrasite replication.

Each link can have a 'cost' (e.g., DS3, T1, ISDN etc.) and the KCC alters the site link topology accordingly. Replication may occur transitively through several site links on same-protocol *site link bridges*, if the cost is low, although KCC automatically costs a direct site-to-site link lower than transitive connections. Site-to-site replication can be configured to occur between a *bridgehead server* in each site, which then replicates the changes to other DCs within the site. Replication for Active Directory zones is automatically configured when DNS is activated in the domain based by site.

Replication of Active Directory uses Remote Procedure Calls (RPC) over IP (RPC/IP). Between Sites SMTP can be used for replication, but only for changes in the Schema, Configuration, or Partial Attribute Set (Global Catalog) NCs. SMTP cannot be used for replicating the default Domain partition.[14]

Database

The Active Directory database, the *directory store*, in Windows 2000 Server uses the JET Blue-based Extensible Storage Engine (ESE98) and is limited to 16 terabytes and 2 billion objects (but only 1 billion security principals) in each domain controller's database. Microsoft has created NTDS databases with more than 2 billion objects.[15] (NT4's Security Account Manager could support no more than 40,000 objects). Called NTDS.DIT, it has two main tables: the *data table* and the *link table*. In Windows Server 2003 a third main table was added for security descriptor single instancing.<ref name="blogs.technet.com"/

Programmatic interface

The features of Active Directory may be accessed programmatically via the COM interfaces provided by *Active Directory Service Interfaces*.[16]

Single server operations

Flexible Single Master Operations Roles (**FSMO**, sometimes pronounced "fizz-mo") operations are also known as operations master roles. Although domain controllers allow simultaneous updates in multiple places, certain operations are supported only on a single server. These operations are performed using the roles listed below:

Role Name	Scope	Description
Schema Master	1 per forest	Schema modifications
Domain Naming Master	1 per forest	Addition and removal of domains if present in root domain
PDC Emulator	1 per domain	Provides backwards compatibility for NT4 clients for PDC operations (like password changes). The PDC runs domain specific processes such as the Security Descriptor Propagator (SDPROP), and is the master time server within the domain. It also handles external trusts, the DFS consistency check, holds current passwords and manages all GPOs as default server.
RID Master	1 per domain	Allocates pools of unique identifiers to domain controllers for use when creating objects
Infrastructure Master	1 per domain/partition	Synchronizes cross-domain group membership changes. The infrastructure master cannot run on a global catalog server (GCS) (unless all DCs are also GCs, or environment consists of a single domain).

Trusting

To allow users in one domain to access resources in another, Active Directory uses trusts.[17]

Trusts inside a forest are automatically created when domains are created. The forest sets the default boundaries of trust, and implicit, transitive trust is automatic for all domains within a forest.

Terminology

One-way trust

: One domain allows access to users on another domain, but the other domain does not allow access to users on the first domain.

Two-way trust

: Two domains allow access to users on both domains.

Trusting domain

: The domain that allows access to users from a trusted domain.

Trusted domain

: The domain that is trusted; whose users have access to the trusting domain.

Transitive trust

: A trust that can extend beyond two domains to other trusted domains in the forest.

Intransitive trust

: A one way trust that does not extend beyond two domains.

Explicit trust

: A trust that an admin creates. It is not transitive and is one way only.

Cross-link trust

: An explicit trust between domains in different trees or in the same tree when a descendant/ancestor (child/parent) relationship does not exist between the two domains.

Shortcut

: Joins two domains in different trees, transitive, one- or two-way.

Forest

: Applies to the entire forest. Transitive, one- or two-way

Realm

Can be transitive or nontransitive, one- or two-way

External

Connect to other forests or non-AD domains. Nontransitive, one- or two-way.[18]

Windows Server 2003 introduced the *forest root trust*. This trust can be used to connect Windows Server 2003 forests if they are operating at the 2003 forest functional level. Authentication across this type of trust is Kerberos based (as opposed to NTLM). Forest trusts are transitive for all the domains in the trusted forests. Forest trusts, however, are not transitive at a forest level. So where domains inside two trusting forests trust each other, forests A that trusts Forest B will not automatically trust Forest C because it is trusted by forest B. In that sense forest A will not automatically (transatively) trust forest C.

Lightweight Directory Service

Active Directory Lightweight Directory Service (*AD LDS*), formerly known as *Active Directory Application Mode* (ADAM),[19] is a light-weight implementation of Active Directory. AD LDS is capable of running as a service on computers running Microsoft Windows Server. AD LDS shares the code base with Active Directory and provides the same functionality as Active Directory, including an identical API, but does not require the creation of domains or domain controllers.

Like Active Directory, AD LDS provides a *Data Store* for storage of directory data and a *Directory Service* with an LDAP *Directory Service Interface*. Unlike Active Directory, however, multiple AD LDS instances can be run on the same server.

Unix integration

Varying levels of interoperability with Active Directory can be achieved on most Unix-like operating systems through standards-compliant LDAP clients, but these systems usually do not interpret many attributes associated with Windows components, such as Group Policy and support for one-way trusts.

Third-parties offer Active Directory integration for Unix platforms (including UNIX, Linux, Mac OS X, and a number of Java- and UNIX-based applications), including:

- *Centrify DirectControl* (Centrify Corporation) – Active Directory-compatible centralized authentication and access control[20]
- *Centrify Express* (Centrify Corporation) – A suite of free Active Directory-compliant services for centralized authentication, monitoring, file-sharing and remote access
- *UNAB* (Computer Associates)
- *TrustBroker* (CyberSafe Limited) – An implementation of Kerberos
- *PowerBroker Identity Services*, formerly *Likewise* (BeyondTrust, formerly Likewise Software) – Allows a non-Windows client to join Active Directory[20]
- *Authentication Services* (Quest Software)
- *ADmitMac* (Thursby Software Systems)[20]
- *Samba* – Can act as a domain controller[21][22]

The schema additions shipped with Windows Server 2003 R2 include attributes that map closely enough to RFC 2307 to be generally usable. The reference implementation of RFC 2307, nss_ldap and pam_ldap provided by PADL.com, support these attributes directly. The default schema for group membership complies with RFC 2307bis (proposed).[23] Windows Server 2003 R2 includes a Microsoft Management Console snap-in that creates and edits the attributes.

An alternate option is to use another directory service such as 389 Directory Server (formerly Fedora Directory Server, FDS), eB2Bcom ViewDS v7.1 XML Enabled Directory or Sun Microsystems Sun Java System Directory Server, which can perform two-way synchronization with AD and thus provide a "deflected" integration, as Unix and

Linux clients authenticate to this while Windows Clients authenticate to AD. Another option is to use OpenLDAP with its *translucent* overlay, which can extend entries in any remote LDAP server with additional attributes stored in a local database. Clients pointed at the local database see entries containing both the remote and local attributes, while the remote database remains completely untouched.

Administration (querying, modifying, and monitoring) of Active Directory can be achieved via many scripting languages, including PowerShell, VBScript, JScript/JavaScript, Perl, Python, and Ruby.[24][25][26][27]

References

[1] "Active Directory on a Windows Server 2003 Network" (http://technet.microsoft.com/en-us/library/cc780036(WS.10).aspx#w2k3tr_ad_over_qbjd). *Active Directory Collection*. Microsoft. 13 March 2003. . Retrieved 25 December 2010.

[2] *Windows Server 2003: Active Directory Infrastructure*. Microsoft Press. 2003. pp. 1–8–1–9. ISBN 0-7356-1438-5.

[3] "Managing Sites" (http://technet.microsoft.com/en-us/library/bb727051.aspx). Microsoft CorporationA. . "An Active Directory site object represents a collection of Internet Protocol (IP) subnets, usually constituting a physical Local Area Network (LAN)."

[4] "Organizational Units". Microsoft Corporation. 2011. "An organizational unit in Active Directory is analogous to a directory in the file system"

[5] Microsoft Server 2008 Reference, discussing shadow groups used for fine-grained password policies: http://technet.microsoft.com/en-us/library/cc770394%28WS.10%29.aspx

[6] "Specifying Security and Administrative Boundaries" (http://technet.microsoft.com/en-us/library/cc755979(WS.10).aspx). Microsoft Corporation. 23 January 2005. . "However, service administrators have abilities that cross domain boundaries. For this reason, the forest is the ultimate security boundary, not the domain."

[7] "Sites overview" (http://technet.microsoft.com/en-us/library/cc782048(WS.10).aspx). Microsoft Corporation. 21 January 2005. . "A site is a set of well-connected subnets."

[8] "Planning for domain controllers and member servers" (http://technet.microsoft.com/en-us/library/cc737059(WS.10).aspx). Microsoft Corporation. 21 January 2005. . "[...] member servers, [...] belong to a domain but do not contain a copy of the Active Directory data."

[9] "Directory data store" (http://technet.microsoft.com/en-us/library/cc736627(WS.10).aspx). Microsoft Corporation. 21 January 2005. . "Active Directory uses four distinct directory partition types to store [...] data. Directory partitions contain domain, configuration, schema, and application data."

[10] Andreas Luther. "Active Directory Replication Traffic" (http://technet.microsoft.com/en-us/library/bb742457.aspx). Microsoft Corporation. . Retrieved 26 May 2010. "The Active Directory is made up of one or more naming contexts or partitions."

[11] "What Is the Global Catalog?" (http://technet.microsoft.com/en-us/library/cc728188(WS.10).aspx). Microsoft Corporation. 10 December 2009. . "[...] a domain controller can locate only the objects in its domain. [...] The global catalog provides the ability to locate objects from any domain [...]"

[12] "Attributes Included in the Global Catalog" (http://msdn.microsoft.com/en-us/library/ms675160(VS.85).aspx). Microsoft Corporation. 26 August 2010. . "The isMemberOfPartialAttributeSet attribute of an attributeSchema object is set to TRUE if the attribute is replicated to the global catalog. [...] When deciding whether or not to place an attribute in the global catalog remember that you are trading increased replication and increased disk storage on global catalog servers for, potentially, faster query performance."

[13] "What Is the Active Directory Replication Model?" (http://technet.microsoft.com/en-us/library/cc737314(WS.10).aspx). Microsoft Corporation. 28 March 2003. . "Domain controllers request (pull) changes rather than send (push) changes that might not be needed."

[14] "What Is Active Directory Replication Topology?" (http://technet.microsoft.com/en-us/library/cc775549(WS.10).aspx). Microsoft Corporation. 28 March 2003. . "SMTP can be used to transport nondomain replication [...]"

[15] efleis (8 June 2006). "Large AD database? Probably not this large" (http://blogs.technet.com/efleis/archive/2006/06/08/434255.aspx). Blogs.technet.com. . Retrieved 20 November 2011.

[16] Active Directory Service Interfaces (http://msdn.microsoft.com/en-us/library/aa772170(VS.85).aspx), Microsoft

[17] "Domain and Forest Trusts Technical Reference" (http://technet.microsoft.com/en-us/library/cc738955(WS.10).aspx). Microsoft Corporation. 28 March 2003. . "Trusts enable [...] authentication and [...] sharing resources across domains or forests"

[18] "How Domain and Forest Trusts Work" (http://technet.microsoft.com/en-us/library/cc773178(WS.10).aspx). Microsoft Corporation. 3 June 2006. . Retrieved 1 June 2010. "Defines several kinds of trusts. (automatic, shortcut, forest, realm, external)"

[19] "AD LDS" (http://msdn.microsoft.com/en-us/library/aa705886(VS.85).aspx). Microsoft. . Retrieved 28 April 2009.

[20] Edge, Charles S., Jr; Smith, Zack; Hunter, Beau (2009). "ch. 3". *Enterprise Mac Administrator's Guide*. New York City: Apress. ISBN 978-1-4302-2443-3.

[21] "Samba4/Releases/4.0.0alpha13" (http://wiki.samba.org/index.php/Samba4/Releases/4.0.0alpha13). *SambaPeople*. SAMBA Project. Archived (http://web.archive.org/web/20101115160233/http://wiki.samba.org/index.php/Samba4/Releases/4.0.0alpha13) from the original on 15 November 2010. . Retrieved 29 November 2010.

[22] "The great DRS success!" (http://people.samba.org/people/2009/10/05#drs-success). *SambaPeople*. SAMBA Project. 5 October 2009. Archived (http://web.archive.org/web/20091013094528/http://people.samba.org/people/2009/10/05) from the original on 13 October 2009. . Retrieved 2 November 2009.

[23] "RFC 2307bis" (http://www.padl.com/~lukeh/rfc2307bis.txt). . Retrieved 20 November 2011.
[24] "Active Directory Administration with Windows PowerShell" (http://technet.microsoft.com/en-us/library/dd378937(WS.10).aspx). Microsoft. . Retrieved 7 June 2011.
[25] "Using Scripts to Search Active Directory" (http://technet.microsoft.com/library/ee692830.aspx). Microsoft. . Retrieved 22 May 2012.
[26] "ITAdminTools Perl Scripts Repository" (http://www.itadmintools.com/2011/09/itadmintools-perl-script-repository.html). ITAdminTools.com. . Retrieved 22 May 2012.
[27] "Win32::OLE" (http://search.cpan.org/~jdb/Win32-OLE-0.1709/lib/Win32/OLE.pm). Perl_Open-Source_Community. . Retrieved 22 May 2012.

External links

- Microsoft's Active Directory Page (http://www.microsoft.com/windowsserver2008/en/us/active-directory.aspx)
- Microsoft Technet: White paper: Active Directory Architecture (http://technet.microsoft.com/en-us/library/bb727030.aspx) (Single technical document that gives an overview about Active Directory.)
- Microsoft Technet: Detailed description of Active Directory on Windows Server 2003 (http://technet.microsoft.com/en-us/library/cc782657(WS.10).aspx)
- Microsoft MSDN Library: [MS-ADTS]: Active Directory Technical Specification (http://msdn.microsoft.com/en-us/library/cc223122.aspx) (part of the Microsoft Open Specification Promise)
- Active Directory Application Mode (ADAM) (http://technet2.microsoft.com/windowsserver/en/library/6a4ae393-17e9-425d-b619-969dab507a6d1033.mspx?mfr=true)
- Powershell scripts to search Active Directory (http://www.happysysadm.com/2010/07/powershell-script-to-search-active.html)

Cloud computing

Cloud computing is the delivery of computing and storage capacity [1] as a service [2] to a community of end-recipients. The name comes from the use of a cloud-shaped symbol[3] as an abstraction for the complex infrastructure it contains in system diagrams[4]. Cloud computing entrusts services with a user's data, software and computation over a network.

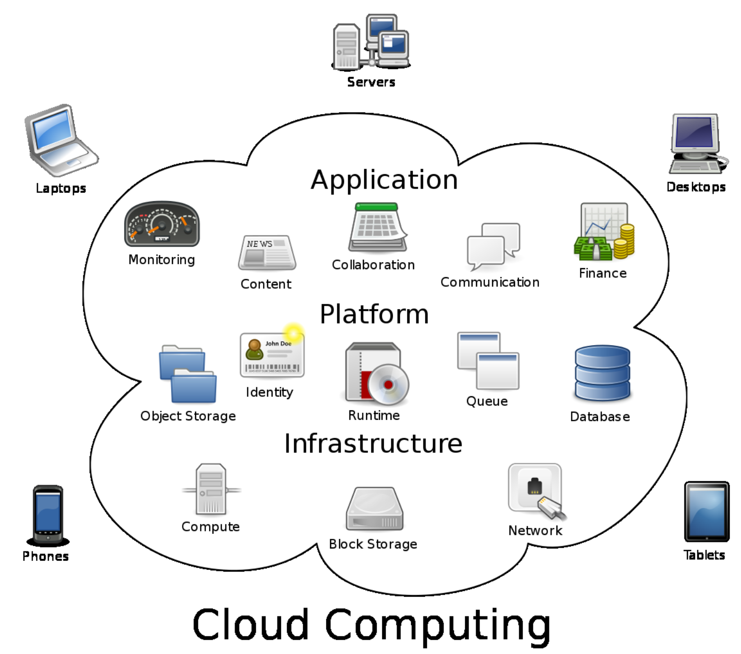

Cloud computing logical diagram

There are three types of cloud computing:[5]

- Infrastructure as a Service (IaaS),
- Platform as a Service (PaaS), and
- Software as a Service (SaaS).

Using Software as a Service, users also rent application software and databases. The cloud providers manage the infrastructure and platforms on which the applications run.

End users access cloud-based applications through a web browser or a light-weight desktop or mobile app while the business software and user's data are stored on servers at a remote location. Proponents claim that cloud computing allows enterprises to get their applications up and running faster, with improved manageability and less maintenance, and enables IT to more rapidly adjust resources to meet fluctuating and unpredictable business demand.[6][7]

Cloud computing relies on sharing of resources to achieve coherence and economies of scale similar to a utility (like the electricity grid) over a network (typically the Internet).[8] At the foundation of cloud computing is the broader concept of converged infrastructure and shared services.[9]

History

The origin of the term *cloud computing* is obscure, but it appears to derive from the practice of using drawings of stylized clouds to denote networks in diagrams of computing and communications systems. The word *cloud* is used as a metaphor for the Internet, based on the standardized use of a cloud-like shape to denote a network on telephony schematics and later to depict the Internet in computer network diagrams as an abstraction of the underlying infrastructure it represents. In the 1990s, telecommunications companies who previously offered primarily dedicated point-to-point data circuits, began offering virtual private network (VPN) services with comparable quality of service but at a much lower cost. By switching traffic to balance utilization as they saw fit, they were able to utilize their overall network bandwidth more effectively. The cloud symbol was used to denote the demarcation point between that which was the responsibility of the provider and that which was the responsibility of the users. Cloud computing extends this boundary to cover servers as well as the network infrastructure.[10]

The underlying concept of cloud computing dates back to the 1950s; when large-scale mainframe became available in academia and corporations, accessible via thin clients / terminal computers. Because it was costly to buy a mainframe, it became important to find ways to get the greatest return on the investment in them, allowing multiple users to share both the physical access to the computer from multiple terminals as well as to share the CPU time, eliminating periods of inactivity, which became known in the industry as time-sharing.[11].

As computers became more prevalent, scientists and technologists explored ways to make large-scale computing power available to more users through time sharing, experimenting with algorithms to provide the optimal use of the infrastructure, platform and applications with prioritized access to the CPU and efficiency for the end users.[12].

John McCarthy opined in the 1960s that "computation may someday be organized as a public utility." Almost all the modern-day characteristics of cloud computing (elastic provision, provided as a utility, online, illusion of infinite supply), the comparison to the electricity industry and the use of public, private, government, and community forms, were thoroughly explored in Douglas Parkhill's 1966 book, *The Challenge of the Computer Utility*. Other scholars have shown that cloud computing's roots go all the way back to the 1950s when scientist Herb Grosch (the author of Grosch's law) postulated that the entire world would operate on dumb terminals powered by about 15 large data centers.[13] Due to the expense of these powerful computers, many corporations and other entities could avail themselves of computing capability through time sharing and several organizations, such as GE's GEISCO, IBM subsidiary The Service Bureau Corporation, Tymshare (founded in 1966), National CSS (founded in 1967 and bought by Dun & Bradstreet in 1979), Dial Data (bought by Tymshare in 1968), and Bolt, Beranek and Newman marketed time sharing as a commercial venture.

The ubiquitous availability of high capacity networks, low cost computers and storage devices as well as the widespread adoption of hardware virtualization, service-oriented architecture, autonomic, and utility computing have led to a tremendous growth in cloud computing.[14][15][16]

After the dot-com bubble, Amazon played a key role in the development of cloud computing by modernizing their data centers, which, like most computer networks, were using as little as 10% of their capacity at any one time, just to leave room for occasional spikes. Having found that the new cloud architecture resulted in significant internal efficiency improvements whereby small, fast-moving "two-pizza teams" could add new features faster and more easily, Amazon initiated a new product development effort to provide cloud computing to external customers, and launched Amazon Web Service (AWS) on a utility computing basis in 2006.[17][18]

In early 2008, Eucalyptus became the first open-source, AWS API-compatible platform for deploying private clouds. In early 2008, OpenNebula, enhanced in the RESERVOIR European Commission-funded project, became the first open-source software for deploying private and hybrid clouds, and for the federation of clouds.[19] In the same year, efforts were focused on providing quality of service guarantees (as required by real-time interactive applications) to cloud-based infrastructures, in the framework of the IRMOS European Commission-funded project, resulting to a **real-time cloud environment**.[20] By mid-2008, Gartner saw an opportunity for cloud computing "to shape the relationship among consumers of IT services, those who use IT services and those who sell them"[21] and observed that "organisations are switching from company-owned hardware and software assets to per-use service-based models" so that the "projected shift to computing... will result in dramatic growth in IT products in some areas and significant reductions in other areas."[22]

On March 1, 2011, IBM announced the Smarter Computing [23] framework to support Smarter Planet.[24] Among the various components of the Smarter Computing foundation, cloud computing is a critical piece.

In 2012, Dr. Biju John and Dr. Souheil Khaddaj incorporated the semantic term into the cloud "Cloud computing is a universal collection of data which extends over the internet in the form of resources (such as information hardware, various platforms, services etc.) and forms individual units within the virtualization environment. Held together by infrastructure providers, service providers and the consumer, then it is semantically accessed by various users." (CLUSE 2012), Bangalore, April 2012[25]

Similar systems and concepts

Cloud computing shares characteristics with:

- Autonomic computing — Computer systems capable of self-management.[26]
- Client–server model — *Client–server computing* refers broadly to any distributed application that distinguishes between service providers (servers) and service requesters (clients).[27]
- Grid computing — "A form of distributed and parallel computing, whereby a 'super and virtual computer' is composed of a cluster of networked, loosely coupled computers acting in concert to perform very large tasks."
- Mainframe computer — Powerful computers used mainly by large organizations for critical applications, typically bulk data processing such as census, industry and consumer statistics, police and secret intelligence services, enterprise resource planning, and financial transaction processing.[28]
- Utility computing — The "packaging of computing resources, such as computation and storage, as a metered service similar to a traditional public utility, such as electricity."[29][30]
- Peer-to-peer — Distributed architecture without the need for central coordination, with participants being at the same time both suppliers and consumers of resources (in contrast to the traditional client–server model).
- Cloud gaming - Also called On-demand gaming is a way of delivering to games to computers. The gaming data will be stored in the provider's server, so that gaming will be independent of client computers used to play the game.

Characteristics

Cloud computing exhibits the following key characteristics:

- **Agility** improves with users' ability to re-provision technological infrastructure resources.
- **Application programming interface** (API) accessibility to software that enables machines to interact with cloud software in the same way the user interface facilitates interaction between humans and computers. Cloud computing systems typically use REST-based APIs.
- **Cost** is claimed to be reduced and in a public cloud delivery model capital expenditure is converted to operational expenditure.[31] This is purported to lower barriers to entry, as infrastructure is typically provided by a third-party

and does not need to be purchased for one-time or infrequent intensive computing tasks. Pricing on a utility computing basis is fine-grained with usage-based options and fewer IT skills are required for implementation (in-house).[32] The e-FISCAL project's state of the art repository[33] contains several articles looking into cost aspects in more detail, most of them concluding that costs savings depend on the type of activities supported and the type of infrastructure available in-house.

- **Device and location independence**[34] enable users to access systems using a web browser regardless of their location or what device they are using (e.g., PC, mobile phone). As infrastructure is off-site (typically provided by a third-party) and accessed via the Internet, users can connect from anywhere.[32]
- **Virtualization** technology allows servers and storage devices to be shared and utilization be increased. Applications can be easily migrated from one physical server to another.
- **Multitenancy** enables sharing of resources and costs across a large pool of users thus allowing for:
  - **Centralization** of infrastructure in locations with lower costs (such as real estate, electricity, etc.)
  - **Peak-load capacity** increases (users need not engineer for highest possible load-levels)
  - **Utilisation and efficiency** improvements for systems that are often only 10–20% utilised.[17]
- **Reliability** is improved if multiple redundant sites are used, which makes well-designed cloud computing suitable for business continuity and disaster recovery.[35]
- **Scalability and Elasticity** via dynamic ("on-demand") provisioning of resources on a fine-grained, self-service basis near real-time, without users having to engineer for peak loads.[36][37]
- **Performance** is monitored, and consistent and loosely coupled architectures are constructed using web services as the system interface.[32]
- **Security** could improve due to centralization of data, increased security-focused resources, etc., but concerns can persist about loss of control over certain sensitive data, and the lack of security for stored kernels.[38] Security is often as good as or better than other traditional systems, in part because providers are able to devote resources to solving security issues that many customers cannot afford.[39] However, the complexity of security is greatly increased when data is distributed over a wider area or greater number of devices and in multi-tenant systems that are being shared by unrelated users. In addition, user access to security audit logs may be difficult or impossible. Private cloud installations are in part motivated by users' desire to retain control over the infrastructure and avoid losing control of information security.
- **Maintenance** of cloud computing applications is easier, because they do not need to be installed on each user's computer and can be accessed from different places.

Service models

Cloud computing providers offer their services according to three fundamental models:[8][40] Infrastructure as a service (IaaS), platform as a service (PaaS), and software as a service (SaaS) where IaaS is the most basic and each higher model abstracts from the details of the lower models.

Infrastructure as a service (IaaS)

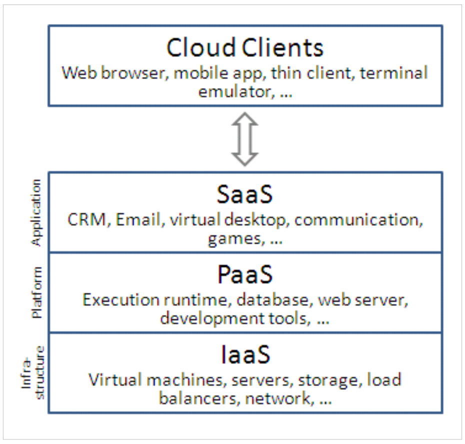

In this most basic cloud service model, cloud providers offer computers, as physical or more often as virtual machines, and other resources. The virtual machines are run as guests by a hypervisor, such as Xen or KVM. Management of pools of hypervisors by the cloud operational support system leads to the ability to scale to support a large numbers of virtual machines. Other resources in IaaS clouds include images in a virtual machine image library, raw (block) and file-based storage, firewalls, load balancers, IP addresses, virtual local area networks (VLANs), and software bundles[41]. IaaS cloud providers supply these resources on demand from their large pools installed in data centers. For wide area connectivity, the Internet can be used or -- in carrier clouds -- dedicated virtual private networks can be configured.

To deploy their applications, cloud users then install operating system images on the machines as well as their application software. In this model, it is the cloud user who is responsible for patching and maintaining the operating systems and application software. Cloud providers typically bill IaaS services on a utility computing basis, that is, cost will reflect the amount of resources allocated and consumed.

IaaS refers not to a machine that does all the work, but simply to a facility given to businesses that offers users the leverage of extra storage space in servers and data centers.

Examples of IaaS include: Amazon CloudFormation (and underlying services such as Amazon EC2), Rackspace Cloud, Google Compute Engine, and RightScale.

Platform as a service (PaaS)

In the PaaS model, cloud providers deliver a computing platform typically including operating system, programming language execution environment, database, and web server. Application developers can develop and run their software solutions on a cloud platform without the cost and complexity of buying and managing the underlying hardware and software layers. With some PaaS offers, the underlying computer and storage resources scale automatically to match application demand such that cloud user does not have to allocate resources manually.

Examples of PaaS include: Amazon Elastic Beanstalk, Heroku, EngineYard, Google App Engine, and Microsoft Azure.

Software as a service (SaaS)

In this model, cloud providers install and operate application software in the cloud and cloud users access the software from cloud clients. The cloud users do not manage the cloud infrastructure and platform on which the application is running. This eliminates the need to install and run the application on the cloud user's own computers simplifying maintenance and support. What makes a cloud application different from other applications is its elasticity. This can be achieved by cloning tasks onto multiple virtual machines at run-time to meet the changing work demand.[42] Load balancers distribute the work over the set of virtual machines. This process is inconspicuous to the cloud user who sees only a single access point. To accommodate a large number of cloud users, cloud applications can be multitenant, that is, any machine serves more than one cloud user organization. It is common to refer to special types of cloud based application software with a similar naming convention: desktop as a service, business process as a service, Test Environment as a Service, communication as a service.

The pricing model for SaaS applications is typically a monthly or yearly flat fee per user.[43]

Examples of SaaS include: Google Apps, Quickbooks Online and Salesforce.com.

Cloud clients

Users access cloud computing using networked client devices, such as desktop computers, laptops, tablets and smartphones. Some of these devices - *cloud clients* - rely on cloud computing for all or a majority of their applications so as to be essentially useless without it. Examples are thin clients and the browser-based Chromebook. Many cloud applications do not require specific software on the client and instead use a web browser to interact with the cloud application. With Ajax and HTML5 these Web user interfaces can achieve a similar or even better look and feel as native applications. Some cloud applications, however, support specific client software dedicated to these applications (e.g., virtual desktop clients and most email clients). Some legacy applications (line of business applications that until now have been prevalent in thin client Windows computing) are delivered via a screen-sharing technology.

Deployment models

Public cloud

Public cloud applications, storage, and other resources are made available to the general public by a service provider. These services are free or offered on a pay-per-use model. Generally, public cloud service providers like Amazon AWS, Microsoft and Google own and operate the infrastructure and offer access only via Internet (direct connectivity is not offered).[32]

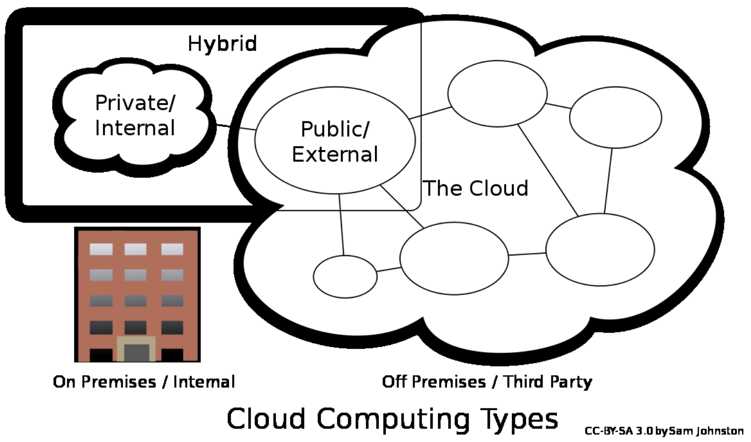

Cloud computing types

Community cloud

Community cloud shares infrastructure between several organizations from a specific community with common concerns (security, compliance, jurisdiction, etc.), whether managed internally or by a third-party and hosted internally or externally. The costs are spread over fewer users than a public cloud (but more than a private cloud), so only some of the cost savings potential of cloud computing are realized.[8]

Hybrid cloud

Hybrid cloud is a composition of two or more clouds (private, community or public) that remain unique entities but are bound together, offering the benefits of multiple deployment models.[8]

By utilizing "hybrid cloud" architecture, companies and individuals are able to obtain degrees of fault tolerance combined with locally immediate usability without dependency on internet connectivity. Hybrid Cloud architecture requires both on-premises resources and off-site (remote) server based cloud infrastructure.

Hybrid clouds lack the flexibility, security and certainty of in-house applications.[44] Hybrid cloud provides the flexibility of in house applications with the fault tolerance and scalability of cloud based services.

Private cloud

Private cloud is cloud infrastructure operated solely for a single organization, whether managed internally or by a third-party and hosted internally or externally.[8] Undertaking a private cloud project requires a significant level and degree of engagement to virtualize the business environment, and it will require the organization to reevaluate decisions about existing resources. When it is done right, it can have a positive impact on a business, but every one of the steps in the project raises security issues that must be addressed in order to avoid serious vulnerabilities.[45]

They have attracted criticism because users "still have to buy, build, and manage them" and thus do not benefit from less hands-on management,[46] essentially "[lacking] the economic model that makes cloud computing such an intriguing concept".[47][48]

Architecture

Cloud architecture,[49] the systems architecture of the software systems involved in the delivery of cloud computing, typically involves multiple *cloud components* communicating with each other over a loose coupling mechanism such as a messaging queue. Elastic provision implies intelligence in the use of tight or loose coupling as applied to mechanisms such as these and others.

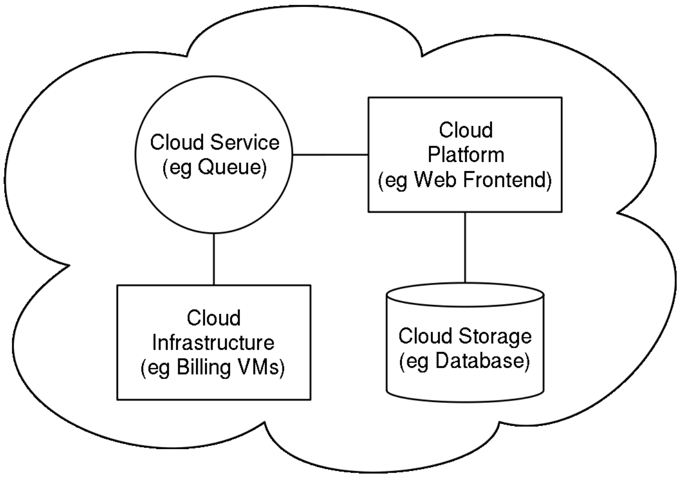

Cloud computing sample architecture

The Intercloud

The Intercloud[50] is an interconnected global "cloud of clouds"[51][52] and an extension of the Internet "network of networks" on which it is based.[53][54][55]

Cloud engineering

Cloud engineering is the application of engineering disciplines to cloud computing. It brings a systematic approach to the high-level concerns of commercialisation, standardisation, and governance in conceiving, developing, operating and maintaining cloud computing systems. It is a multidisciplinary method encompassing contributions from diverse areas such as systems, software, web, performance, information, security, platform, risk, and quality engineering.

Issues

Privacy

The cloud model has been criticised by privacy advocates for the greater ease in which the companies hosting the cloud services control, thus, can monitor at will, lawfully or unlawfully, the communication and data stored between the user and the host company. Instances such as the secret NSA program, working with AT&T, and Verizon, which recorded over 10 million phone calls between American citizens, causes uncertainty among privacy advocates, and the greater powers it gives to telecommunication companies to monitor user activity.[56] Using a cloud service provider (CSP) can complicate privacy of data because of the extent to which virtualization for cloud processing (virtual machines) and cloud storage are used to implement cloud service.[57] The point is that CSP operations, customer or tenant data may not remain on the same system, or in the same data center or even within the same provider's cloud. This can lead to legal concerns over jurisdiction. While there have been efforts (such as US-EU Safe Harbor) to "harmonise" the legal environment, providers such as Amazon still cater to major markets (typically the United States and the European Union) by deploying local infrastructure and allowing customers to select "availability zones."[58] Cloud computing poses privacy concerns because the service provider may access the data that is on the cloud at any point in time. They could accidentally or deliberately alter or even delete information.[59]

Compliance

In order to obtain compliance with regulations including FISMA, HIPAA, and SOX in the United States, the Data Protection Directive in the EU and the credit card industry's PCI DSS, users may have to adopt *community* or *hybrid* deployment modes that are typically more expensive and may offer restricted benefits. This is how Google is able to "manage and meet additional government policy requirements beyond FISMA"[60][61] and Rackspace Cloud or QubeSpace are able to claim PCI compliance.[62]

Many providers also obtain a SAS 70 Type II audit, but this has been criticised on the grounds that the hand-picked set of goals and standards determined by the auditor and the auditee are often not disclosed and can vary widely.[63] Providers typically make this information available on request, under non-disclosure agreement.[64][65]

Customers in the EU contracting with cloud providers outside the EU/EEA have to adhere to the EU regulations on export of personal data.[66]

U.S. Federal Agencies have been directed by the Office of Management and Budget to use a process called FedRAMP (Federal Risk and Authorization Management Program) to assess and authorize cloud products and services. Federal CIO Steven VanRoekel issued a memorandum to federal agency Chief Information Officers on December 8, 2011 defining how federal agencies should use FedRAMP. FedRAMP consists of a subset of NIST Special Publication 800-53 security controls specifically selected to provide protection in cloud environments. A subset has been defined for the FIPS 199 low categorization and the FIPS 199 moderate categorization. The FedRAMP program has also established a Joint Acceditation Board (JAB) consisting of Chief Information Officers from DoD, DHS and GSA. The JAB is responsible for establishing accreditation standards for 3rd party organizations who will perform the assessments of cloud solutions. The JAB will also review authorization packages and may grant provisional authorization (to operate). The federal agency consuming the service will still have the final responsibility for final authority to operate.[67]

Legal

As can be expected with any revolutionary change in the landscape of global computing, certain legal issues arise; everything from trademark infringement, security concerns to the sharing of propriety data resources .

Open source

Open-source software has provided the foundation for many cloud computing implementations, one prominent example being the Hadoop framework.[68] In November 2007, the Free Software Foundation released the Affero General Public License, a version of GPLv3 intended to close a perceived legal loophole associated with free software designed to be run over a network.[69]

Open standards

Most cloud providers expose APIs that are typically well-documented (often under a Creative Commons license[70]) but also unique to their implementation and thus not interoperable. Some vendors have adopted others' APIs and there are a number of open standards under development, with a view to delivering interoperability and portability.[71]

Security

As cloud computing is achieving increased popularity, concerns are being voiced about the security issues introduced through adoption of this new model. The effectiveness and efficiency of traditional protection mechanisms are being reconsidered as the characteristics of this innovative deployment model can differ widely from those of traditional architectures.[72] An alternative perspective on the topic of cloud security is that this is but another, although quite broad, case of "applied security" and that similar security principles that apply in shared multi-user mainframe security models apply with cloud security.[73]

The relative security of cloud computing services is a contentious issue that may be delaying its adoption.[74] Physical control of the Private Cloud equipment is more secure than having the equipment off site and under someone else’s control. Physical control and the ability to visually inspect the data links and access ports is required in order to ensure data links are not compromised. Issues barring the adoption of cloud computing are due in large part to the private and public sectors' unease surrounding the external management of security-based services. It is the very nature of cloud computing-based services, private or public, that promote external management of provided services. This delivers great incentive to cloud computing service providers to prioritize building and maintaining strong management of secure services.[75] Security issues have been categorised into sensitive data access, data segregation, privacy, bug exploitation, recovery, accountability, malicious insiders, management console security, account control, and multi-tenancy issues. Solutions to various cloud security issues vary, from cryptography, particularly public key infrastructure (PKI), to use of multiple cloud providers, standardisation of APIs, and improving virtual machine support and legal support.[72][76][77]

Cloud computing offers many benefits, but it also is vulnerable to threats. As the uses of cloud computing increase, it is highly likely that more criminals will try to find new ways to exploit vulnerabilities in the system. There are many underlying challenges and risks in cloud computing that increase the threat of data being compromised. To help mitigate the threat, cloud computing stakeholders should invest heavily in risk assessment to ensure that the system encrypts to protect data; establishes trusted foundation to secure the platform and infrastructure; and builds higher assurance into auditing to strengthen compliance. Security concerns must be addressed in order to establish trust in cloud computing technology.

Sustainability

Although cloud computing is often assumed to be a form of "green computing", there is no published study to substantiate this assumption.[78] Siting the servers affects the environmental effects of cloud computing. In areas where climate favors natural cooling and renewable electricity is readily available, the environmental effects will be more moderate. (The same holds true for "traditional" data centers.) Thus countries with favorable conditions, such as Finland,[79] Sweden and Switzerland,[80] are trying to attract cloud computing data centers. Energy efficiency in cloud computing can result from energy-aware scheduling and server consolidation.[81] However, in the case of distributed clouds over data centers with different source of energies including renewable source of energies, a small compromise on energy consumption reduction could result in high carbon footprint reduction.[82]

Abuse

As with privately purchased hardware, customers can purchase the services of cloud computing for nefarious purposes. This includes password cracking and launching attacks using the purchased services.[83] In 2009, a banking trojan illegally used the popular Amazon service as a command and control channel that issued software updates and malicious instructions to PCs that were infected by the malware.[84]

Research

Many universities, vendors and government organisations are investing in research around the topic of cloud computing:[85][86]

- In October 2007, the Academic Cloud Computing Initiative (ACCI) was announced as a multi-university project designed to enhance students' technical knowledge to address the challenges of cloud computing.[87]
- In April 2009, UC Santa Barbara released the first open source platform-as-a-service, AppScale, which is capable of running Google App Engine applications at scale on a multitude of infrastructures.
- In April 2009, the St Andrews Cloud Computing Co-laboratory was launched, focusing on research in the important new area of cloud computing. Unique in the UK, StACC aims to become an international centre of excellence for research and teaching in cloud computing and will provide advice and information to businesses interested in using cloud-based services.[88]
- In October 2010, the TClouds (Trustworthy Clouds) project was started, funded by the European Commission's 7th Framework Programme. The project's goal is to research and inspect the legal foundation and architectural design to build a resilient and trustworthy cloud-of-cloud infrastructure on top of that. The project also develops a prototype to demonstrate its results.[89]
- In December 2010, the TrustCloud research project [90][91] was started by HP Labs Singapore to address transparency and accountability of cloud computing via detective, data-centric approaches[92] encapsulated in a five-layer TrustCloud Framework. The team identified the need for monitoring data life cycles and transfers in the cloud,[90] leading to the tackling of key cloud computing security issues such as cloud data leakages, cloud accountability and cross-national data transfers in transnational clouds.
- In July 2011, the High Performance Computing Cloud (HPCCLoud) project was kicked-off aiming at finding out the possibilities of enhancing performance on cloud environments while running the scientific applications - development of HPCCLoud Performance Analysis Toolkit which was funded by CIM-Returning Experts Programme - under the coordination of Prof. Dr. Shajulin Benedict.
- In June 2011, the Telecommunications Industry Association developed a Cloud Computing White Paper, to analyze the integration challenges and opportunities between cloud services and traditional U.S. telecommunications standards.[93]
- In 2011, FEMhub launched NCLab, a free SaaS application for science, technology, engineering and mathematics (STEM). NCLab has more than 10,000 users as of July 2012.

References

[1] "Cloud Computing" (http://www.academicroom.com/topics/cloud-computing). Academic Room. . Retrieved 2012-06-16.

[2] "Cloud Computing" (http://www.academicroom.com/topics/cloud-computing). Academic Room. . Retrieved 2012-06-16.

[3] Figure 8, "A network 70 is shown schematically as a cloud", US Patent 5,485,455, column 17, line 22, filed Jan 28, 1994

[4] Cloud symbol 49 in Figure 1, "Internet cloud 106", US Patent 6,069,890, column 7, line 64

[5] 'The Institute IEEE', "A View Inside the Cloud", by Ania Monaco, June 2012, page 8

[6] "Baburajan, Rajani, "The Rising Cloud Storage Market Opportunity Strengthens Vendors," infoTECH, August 24, 2011" (http://it.tmcnet.com/channels/cloud-storage/articles/211183-rising-cloud-storage-market-opportunity-strengthens-vendors.htm). It.tmcnet.com. 2011-08-24. . Retrieved 2011-12-02.

[7] "Oestreich, Ken, "Converged Infrastructure," CTO Forum, November 15, 2010" (http://www.thectoforum.com/content/converged-infrastructure-0). Thectoforum.com. 2010-11-15. . Retrieved 2011-12-02.

[8] "The NIST Definition of Cloud Computing" (http://csrc.nist.gov/publications/nistpubs/800-145/SP800-145.pdf). National Institute of Science and Technology. . Retrieved 24 July 2011.

[9] "Kerravala, Zeus, Yankee Group, "Migrating to the cloud is dependent infrastructure," Tech Target" (http://www.convergedinfrastructure.com/Path-to-the-Cloud/). Convergedinfrastructure.com. . Retrieved 2011-12-02.

[10] "July, 1993 meeting report from the IP over ATM working group of the IETF" (http://mirror.switch.ch/ftp/doc/ietf/ipatm/atm-minutes-93jul.txt). CH: Switch. . Retrieved 2010-08-22.

[11] Strachey, Christopher (June 1959). "Time Sharing in Large Fast Computers". *Proceedings of the International Conference on Information processing, UNESCO*. paper B.2.19: 336-341.

[12] Corbató, Fernando J.. "An Experimental Time-Sharing System" (http://larch-www.lcs.mit.edu:8001/~corbato/sjcc62/). *SJCC Proceedings*. MIT. . Retrieved 3 July 2012.

[13] Ryan; Falvey; Merchant (October 2011), "Regulation of the Cloud in India" (http://ssrn.com/abstract=1941494), *Journal of Internet Law* **15** (4),

[14] "Cloud Computing: Clash of the clouds" (http://www.economist.com/displaystory.cfm?story_id=14637206). The Economist. 2009-10-15. . Retrieved 2009-11-03.

[15] "Gartner Says Cloud Computing Will Be As Influential As E-business" (http://www.gartner.com/it/page.jsp?id=707508). Gartner. . Retrieved 2010-08-22.

[16] Gruman, Galen (2008-04-07). "What cloud computing really means" (http://www.infoworld.com/d/cloud-computing/what-cloud-computing-really-means-031). *InfoWorld*. . Retrieved 2009-06-02.

[17] "Jeff Bezos' Risky Bet" (http://www.businessweek.com/magazine/content/06_46/b4009001.htm), *Business Week*,

[18] "Amazon's early efforts at cloud computing partly accidental" (http://itknowledgeexchange.techtarget.com/cloud-computing/2010/06/17/amazons-early-efforts-at-cloud-computing-partly-accidental/), *IT Knowledge Exchange*, Tech Target, 2010-6-17,

[19] B Rochwerger, J Caceres, RS Montero, D Breitgand, E Elmroth, A Galis, E Levy, IM Llorente, K Nagin, Y Wolfsthal, E Elmroth, J Caceres, M Ben-Yehuda, W Emmerich, F Galan. "The RESERVOIR Model and Architecture for Open Federated Cloud Computing", IBM Journal of Research and Development, Vol. 53, No. 4. (2009)

[20] D Kyriazis, A Menychtas, G Kousiouris, K Oberle, T Voith, M Boniface, E Oliveros, T Cucinotta, S Berger, "A Real-time Service Oriented Infrastructure", International Conference on Real-Time and Embedded Systems (RTES 2010), Singapore, November 2010

[21] Keep an eye on cloud computing (http://www.networkworld.com/newsletters/itlead/2008/070708itlead1.html), Amy Schurr, Network World, 2008-07-08, citing the Gartner report, "Cloud Computing Confusion Leads to Opportunity". Retrieved 2009-09-11.

[22] Gartner Says Worldwide IT Spending On Pace to Surpass Trillion in 2008 (http://www.gartner.com/it/page.jsp?id=742913), Gartner, 2008-08-18. Retrieved 2009-09-11.

[23] http://www-03.ibm.com/systems/data/flash/smartercomputing/

[24] "Launch of IBM Smarter Computing" (https://www-304.ibm.com/connections/blogs/IBMSmarterSystems/date/201102?lang=en_us). . Retrieved 1 March 2011.

[25] (http://www.iseec.org.uk/cluse/index.html) 12.

[26] "What's In A Name? Utility vs. Cloud vs Grid" (http://www.datacenterknowledge.com/archives/2008/Mar/25/whats_in_a_name_utility_vs_cloud_vs_grid.html). Datacenterknowledge.com. . Retrieved 2010-08-22.

[27] "Distributed Application Architecture" (http://java.sun.com/developer/Books/jdbc/ch07.pdf). Sun Microsystem. . Retrieved 2009-06-16.

[28] "Sun CTO: Cloud computing is like the mainframe" (http://itknowledgeexchange.techtarget.com/mainframe-blog/sun-cto-cloud-computing-is-like-the-mainframe/). Itknowledgeexchange.techtarget.com. 2009-03-11. . Retrieved 2010-08-22.

[29] "It's probable that you've misunderstood 'Cloud Computing' until now" (http://portal.acm.org/citation.cfm?id=1496091.1496100&coll=&dl=ACM&CFID=21518680&CFTOKEN=18800807). TechPluto. . Retrieved 2010-09-14.

[30] Danielson, Krissi (2008-03-26). "Distinguishing Cloud Computing from Utility Computing" (http://www.ebizq.net/blogs/saasweek/2008/03/distinguishing_cloud_computing/). Ebizq.net. . Retrieved 2010-08-22.

[31] "Recession Is Good For Cloud Computing – Microsoft Agrees" (http://www.cloudave.com/link/recession-is-good-for-cloud-computing-microsoft-agrees). CloudAve. . Retrieved 2010-08-22.

[32] "Defining "Cloud Services" and "[http://www.cloudreviews.com/blog/what-is-hot-in-cloud-computing Cloud Computing (http://blogs.idc.com/ie/?p=190)""]. IDC. 2008-09-23. . Retrieved 2010-08-22.
[33] "e-FISCAL project state of the art repository" (http://www.efiscal.eu/state-of-the-art). .
[34] Farber, Dan (2008-06-25). "The new geek chic: Data centers" (http://news.cnet.com/8301-13953_3-9977049-80.html). CNET News. . Retrieved 2010-08-22.
[35] King, Rachael (2008-08-04). "Cloud Computing: Small Companies Take Flight" (http://www.businessweek.com/technology/content/aug2008/tc2008083_619516.htm). Businessweek. . Retrieved 2010-08-22.
[36] "Defining and Measuring Cloud Elasticity" (http://digbib.ubka.uni-karlsruhe.de/volltexte/1000023476). KIT Software Quality Departement. . Retrieved 13 August 2011.
[37] "Economies of Cloud Scale Infrastructure" (http://www.youtube.com/watch?v=nfDsY3f4nVI). Cloud Slam 2011. . Retrieved 13 May 2011.
[38] "Encrypted Storage and Key Management for the cloud" (http://www.cryptoclarity.com/CryptoClarityLLC/Welcome/Entries/2009/7/23_Encrypted_Storage_and_Key_Management_for_the_cloud.html). Cryptoclarity.com. 2009-07-30. . Retrieved 2010-08-22.
[39] Mills, Elinor (2009-01-27). "Cloud computing security forecast: Clear skies" (http://news.cnet.com/8301-1009_3-10150569-83.html). CNET News. . Retrieved 2010-08-22.
[40] Voorsluys, William; Broberg, James; Buyya, Rajkumar (February 2011). "Introduction to Cloud Computing" (http://media.johnwiley.com.au/product_data/excerpt/90/04708879/0470887990-180.pdf). In R. Buyya, J. Broberg, A.Goscinski. *Cloud Computing: Principles and Paradigms*. New York, USA: Wiley Press. pp. 1–44. ISBN 978-0-470-88799-8. .
[41] Amies, Alex; Sluiman, Harm; Liu, Guo Ning (July 2012). "Infrastructure as a Service Cloud Concepts" (http://www.ibmpressbooks.com/bookstore/product.asp?isbn=9780133066845). *Developing and Hosting Applications on the Cloud*. IBM Press. ISBN 978-0-13-306684-5. .
[42] Hamdaqa, Mohammad. *A Reference Model for Developing Cloud Applications* (http://www.stargroup.uwaterloo.ca/~mhamdaqa/publications/A REFERENCEMODELFORDEVELOPINGCLOUD APPLICATIONS.pdf). .
[43] Chou, Timothy. *Introduction to Cloud Computing: Business & Technology* (http://www.scribd.com/doc/64699897/Introduction-to-Cloud-Computing-Business-and-Technology). .
[44] Stevens, Alan (June 29, 2011). "When hybrid clouds are a mixed blessing" (http://www.theregister.co.uk/2011/06/29/hybrid_cloud/). *The Register* (http://www.theregister.co.uk). . Retrieved March 28, 2012.
[45] "Is a Private Cloud Really More Secure?" (http://content.dell.com/us/en/enterprise/d/large-business/private-cloud-more-secure.aspx). Dell.com. . Retrieved 07-11-12.
[46] Foley, John. "Private Clouds Take Shape" (http://www.informationweek.com/news/services/business/showArticle.jhtml?articleID=209904474). InformationWeek. . Retrieved 2010-08-22.
[47] Haff, Gordon (2009-01-27). "Just don't call them private clouds" (http://news.cnet.com/8301-13556_3-10150841-61.html). CNET News. . Retrieved 2010-08-22.
[48] "There's No Such Thing As A Private Cloud" (http://www.informationweek.com/cloud-computing/blog/archives/2009/01/theres_no_such.html). InformationWeek. 2010-06-30. . Retrieved 2010-08-22.
[49] "Building GrepTheWeb in the Cloud, Part 1: Cloud Architectures" (http://developer.amazonwebservices.com/connect/entry.jspa?externalID=1632&categoryID=100). Developer.amazonwebservices.com. . Retrieved 2010-08-22.
[50] Bernstein, David; Ludvigson, Erik; Sankar, Krishna; Diamond, Steve; Morrow, Monique (2009-05-24). *Blueprint for the Intercloud – Protocols and Formats for Cloud Computing Interoperability* (http://www2.computer.org/portal/web/csdl/doi/10.1109/ICIW.2009.55). IEEE Computer Society. pp. 328–336. doi:10.1109/ICIW.2009.55. .
[51] "Kevin Kelly: A Cloudbook for the Cloud" (http://www.kk.org/thetechnium/archives/2007/11/a_cloudbook_for.php). Kk.org. . Retrieved 2010-08-22.
[52] "Intercloud is a global cloud of clouds" (http://samj.net/2009/06/intercloud-is-global-cloud-of-clouds.html). Samj.net. 2009-06-22. . Retrieved 2010-08-22.
[53] "Vint Cerf: Despite Its Age, The Internet is Still Filled with Problems" (http://www.readwriteweb.com/archives/vint_cerf_despite_its_age_the.php?mtcCampaign=2765). Readwriteweb.com. . Retrieved 2010-08-22.
[54] "SP360: Service Provider: From India to Intercloud" (http://blogs.cisco.com/sp/comments/from_india_to_intercloud/). Blogs.cisco.com. . Retrieved 2010-08-22.
[55] Canada (2007-11-29). "Head in the clouds? Welcome to the future" (http://www.theglobeandmail.com/servlet/story/LAC.20071129.TWLINKS29/TPStory/Business). Toronto: Theglobeandmail.com. . Retrieved 2010-08-22.
[56] Cauley, Leslie (2006-05-11). "NSA has massive database of Americans' phone calls" (http://www.usatoday.com/news/washington/2006-05-10-nsa_x.htm). USATODAY.com. . Retrieved 2010-08-22.
[57] Winkler, Vic (2011). *Securing the Cloud: Cloud Computer Security Techniques and Tactics* (http://www.elsevier.com/wps/find/bookdescription.cws_home/723529/description#description). Waltham, MA USA: Elsevier. p. 60. ISBN 978-1-59749-592-9. .
[58] "Feature Guide: Amazon EC2 Availability Zones" (http://developer.amazonwebservices.com/connect/entry.jspa?externalID=1347&categoryID=112). Amazon Web Services. . Retrieved 2010-08-22.
[59] "Cloud Computing Privacy Concerns on Our Doorstep" (http://cacm.acm.org/magazines/2011/1/103200-cloud-computing-privacy-concerns-on-our-doorstep/fulltext). .
[60] "FISMA compliance for federal cloud computing on the horizon in 2010" (http://searchcompliance.techtarget.com/news/article/0,289142,sid195_gci1377298,00.html). SearchCompliance.com. . Retrieved 2010-08-22.

[61] "Google Apps and Government" (http://googleenterprise.blogspot.com/2009/09/google-apps-and-government.html). Official Google Enterprise Blog. 2009-09-15. . Retrieved 2010-08-22.

[62] "Cloud Hosting is Secure for Take-off: Mosso Enables The Spreadsheet Store, an Online Merchant, to become PCI Compliant" (http://www.rackspace.com/cloud/blog/2009/03/05/cloud-hosting-is-secure-for-take-off-mosso-enables-the-spreadsheet-store-an-online-merchant-to-become-pci-compliant/). Rackspace. 2009-03-14. . Retrieved 2010-08-22.

[63] "Amazon gets SAS 70 Type II audit stamp, but analysts not satisfied" (http://searchcloudcomputing.techtarget.com/news/article/0,289142,sid201_gci1374629,00.html). SearchCloudComputing.com. 2009-11-17. . Retrieved 2010-08-22.

[64] "Assessing Cloud Computing Agreements and Controls" (http://wistechnology.com/articles/6954/). WTN News. . Retrieved 2010-08-22.

[65] "Cloud Certification From Compliance Mandate to Competitive Differentiator" (http://www.youtube.com/watch?v=wYiFdnZAlNQ). Cloudcor. . Retrieved 2011-09-20.

[66] "How the New EU Rules on Data Export Affect Companies in and Outside the EU | Dr. Thomas Helbing – Kanzlei für Datenschutz-, Online- und IT-Recht" (http://www.thomashelbing.com/en/how-new-eu-rules-data-export-affect-companies-and-outside-eu). Dr. Thomas Helbing. . Retrieved 2010-08-22.

[67] "FedRAMP" (http://www.gsa.gov/portal/category/102371). *U.S. General Services Administration*. 2012-06-13. . Retrieved 2012-06-17.

[68] "Open source fuels growth of cloud computing, software-as-a-service" (http://www.networkworld.com/news/2008/072808-open-source-cloud-computing.html). Network World. . Retrieved 2010-08-22.

[69] "AGPL: Open Source Licensing in a Networked Age" (http://redmonk.com/sogrady/2009/04/15/open-source-licensing-in-a-networked-age/). Redmonk.com. 2009-04-15. . Retrieved 2010-08-22.

[70] GoGrid Moves API Specification to Creative Commons (http://www.gogrid.com/company/press-releases/gogrid-moves-api-specification-to-creativecommons.php)

[71] "Eucalyptus Completes Amazon Web Services Specs with Latest Release" (http://ostatic.com/blog/eucalyptus-completes-amazon-web-services-specs-with-latest-release). Ostatic.com. . Retrieved 2010-08-22.

[72] Zissis, Dimitrios; Lekkas (2010). "Addressing cloud computing security issues" (http://www.sciencedirect.com/science/article/pii/S0167739X10002554). *Future Generation Computer Systems*. doi:10.1016/j.future.2010.12.006. .

[73] Winkler, Vic (2011). *Securing the Cloud: Cloud Computer Security Techniques and Tactics* (http://www.elsevier.com/wps/find/bookdescription.cws_home/723529/description#description). Waltham, MA USA: Syngress. pp. 187, 189. ISBN 978-1-59749-592-9. .

[74] "Are security issues delaying adoption of cloud computing?" (http://www.networkworld.com/news/2009/042709-burning-security-cloud-computing.html). Network World. . Retrieved 2010-08-22.

[75] "Security of virtualization, cloud computing divides IT and security pros" (http://www.networkworld.com/news/2010/022210-virtualization-cloud-security-debate.html). Network World. 2010-02-22. . Retrieved 2010-08-22.

[76] Armbrust, M; Fox, A., Griffith, R., Joseph, A., Katz, R., Konwinski, A., Lee, G., Patterson, D., Rabkin, A., Zaharia, (2010). "A view of cloud computing.". *Communication of the ACM* **53** (4): 50–58. doi:10.1145/1721654.1721672.

[77] Anthens, G. "Security in the cloud". *Communications of the ACM* **53** (11). doi:10.1145/1839676.1839683.

[78] James Urquhart (January 7, 2010). "Cloud computing's green paradox" (http://news.cnet.com/8301-19413_3-10428065-240.html). CNET News. . Retrieved March 12, 2010. "...there is some significant evidence that the cloud is encouraging more compute consumption"

[79] Finland – First Choice for Siting Your Cloud Computing Data Center. (http://www.fincloud.freehostingcloud.com/). Retrieved 4 August 2010.

[80] Swiss Carbon-Neutral Servers Hit the Cloud. (http://www.greenbiz.com/news/2010/06/30/swiss-carbon-neutral-servers-hit-cloud). Retrieved 4 August 2010.

[81] Berl, Andreas, et al., Energy-Efficient Cloud Computing (http://comjnl.oxfordjournals.org/content/53/7/1045.short), The Computer Journal, 2010.

[82] Farrahi Moghaddam, Fereydoun, et al., Low Carbon Virtual Private Clouds (http://ieeexplore.ieee.org/search/srchabstract.jsp?tp=&arnumber=6008718), IEEE Cloud 2011.

[83] Alpeyev, Pavel (2011-05-14). "Amazon.com Server Said to Have Been Used in Sony Attack" (http://www.bloomberg.com/news/2011-05-13/sony-network-said-to-have-been-invaded-by-hackers-using-amazon-com-server.html). Bloomberg. . Retrieved 2011-08-20.

[84] Goodin, Dan (2011-05-14). "PlayStation Network hack launched from Amazon EC2" (http://www.theregister.co.uk/2011/05/14/playstation_network_attack_from_amazon/). The Register. . Retrieved 2012-05-18.

[85] "Cloud Net Directory. Retrieved 2010-03-01" (http://www.cloudbook.net/directories/research-clouds). Cloudbook.net. . Retrieved 2010-08-22.

[86] "– National Science Foundation (NSF) News – National Science Foundation Awards Millions to Fourteen Universities for Cloud Computing Research – US National Science Foun" (http://www.nsf.gov/news/news_summ.jsp?cntn_id=114686). Nsf.gov. . Retrieved 2011-08-20.

[87] Rich Miller (2008-05-02). "IBM, Google Team on an Enterprise Cloud" (http://www.datacenterknowledge.com/archives/2008/05/02/ibm-google-team-on-an-enterprise-cloud/). DataCenterKnowledge.com. . Retrieved 2010-08-22.

[88] "StACC - Collaborative Research in Cloud Computing" (http://www.cs.st-andrews.ac.uk/stacc). University of St Andrews department of Computer Science. . Retrieved 2012-06-17.

[89] "Trustworthy Clouds: Privacy and Resilience for Internet-scale Critical Infrastructure" (http://www.tclouds-project.eu). . Retrieved 2012-06-17.

[90] Ko, Ryan K. L.; Jagadpramana, Peter; Lee, Bu Sung (2011). "Flogger: A File-centric Logger for Monitoring File Access and Transfers within Cloud Computing Environments" (http://www.hpl.hp.com/techreports/2011/HPL-2011-119.pdf). *Proceedings of the 10th IEEE International Conference on Trust, Security and Privacy of Computing and Communications (TrustCom-11)*. .

[91] Ko, Ryan K. L.; Jagadpramana, Peter; Mowbray, Miranda; Pearson, Siani; Kirchberg, Markus; Liang, Qianhui; Lee, Bu Sung (2011). "TrustCloud: A Framework for Accountability and Trust in Cloud Computing" (http://www.hpl.hp.com/techreports/2011/HPL-2011-38.pdf). *Proceedings of the 2nd IEEE Cloud Forum for Practitioners (IEEE ICFP 2011), Washington DC, USA, July 7–8, 2011*. .

[92] Ko, Ryan K. L. Ko; Kirchberg, Markus; Lee, Bu Sung (2011). "From System-Centric Logging to Data-Centric Logging - Accountability, Trust and Security in Cloud Computing" (http://www.hpl.hp.com/people/ryan_ko/RKo-DSR2011-Data_Centric_Logging.pdf). *Proceedings of the 1st Defence, Science and Research Conference 2011 - Symposium on Cyber Terrorism, IEEE Computer Society, 3–4 August 2011, Singapore*. .

[93] "Publication Download" (http://www.tiaonline.org/market_intelligence/publication_download.cfm?file=TIA_Cloud_Computing_White_Paper). Tiaonline.org. . Retrieved 2011-12-02.

cloud reviews (http://www.cloudreviews.com/)

External links

- Cloud Computing Portal and Scholarly Resources on Academic Room (http://www.academicroom.com/physical-sciences/information-technology/cloud-computing)
- The NIST Definition of Cloud Computing (http://csrc.nist.gov/publications/nistpubs/800-145/SP800-145.pdf). Peter Mell and Timothy Grance, NIST Special Publication 800-145 (September 2011). National Institute of Standards and Technology, U.S. Department of Commerce.
- Guidelines on Security and Privacy in Public Cloud Computing (http://nvlpubs.nist.gov/nistpubs/sp/2011/sp800-144.pdf). Wayne Jansen and Timothy Grance, NIST Special Publication 800-144 (December 2011). National Institute of Standards and Technology, U.S. Department of Commerce.
- Cloud Computing - Benefits, risks and recommendation for information security (http://www.enisa.europa.eu/activities/risk-management/files/deliverables/cloud-computing-risk-assessment). Daniele Cattedu and Giles Hobben, European Network and Information Security Agency 2009.

Configuration Management: Many Hosts

Secure Shell

Secure Shell (**SSH**) is a network protocol for secure data communication, remote shell services or command execution and other secure network services between two networked computers that it connects via a secure channel over an insecure network: a server and a client (running SSH server and SSH client programs, respectively).[1] The protocol specification distinguishes two major versions that are referred to as SSH-1 and SSH-2.

The best-known application of the protocol is for access to shell accounts on Unix-like operating systems. It was designed as a replacement for Telnet and other insecure remote shell protocols such as the Berkeley rsh and rexec protocols, which send information, notably passwords, in plaintext, rendering them susceptible to interception and disclosure using packet analysis.[2] The encryption used by SSH is intended to provide confidentiality and integrity of data over an unsecured network, such as the Internet.

Definition

SSH uses public-key cryptography to authenticate the remote computer and allow it to authenticate the user, if necessary.[1] Anyone can produce a matching pair of different keys (public and private). The public key is placed on all computers that must allow access to the owner of the matching private key (the owner keeps the private key secret). While authentication is based on the private key, the key itself is never transferred through the network during authentication.

SSH only verifies if the same person offering the public key also owns the matching private key. Hence, in all versions of SSH it is important to verify unknown public keys, i.e. associate the public keys with identities, before accepting them as valid. Accepting an attacker's public key without validation will authorize an unauthorized attacker as a valid user.

Key management

On Unix-like systems, the list of authorized keys is stored in the home folder of the user that is allowed to log in remotely, in the file ~/.ssh/authorized_keys.[3] This file is only respected by ssh if it is not writable by anything apart from the owner and root. When the public key is present on one side and the matching private key is present on another side, typing in the password is no longer required (some software like MPI stack may need this password-less access to run properly). However, for additional security the private key itself can be locked with a passphrase.

The private key can also be looked for in standard places, but its full path can also be specified as a command line setting (the switch *-i* for ssh). The ssh-keygen utility produces the public and private keys, always in pairs.

SSH also supports password-based authentication that is encrypted by automatically generated keys. In this case the attacker could imitate the legitimate side, ask for the password and obtain it (man-in-the-middle attack). However this is only possible if the two sides have never authenticated before, as SSH remembers the key that the remote side once used. Password authentication can be disabled.

Usage

SSH is typically used to log into a remote machine and execute commands, but it also supports tunneling, forwarding TCP ports and X11 connections; it can transfer files using the associated SSH file transfer (SFTP) or secure copy (SCP) protocols.[1] SSH uses the client-server model.

The standard TCP port 22 has been assigned for contacting SSH servers,[4] though administrators frequently change it to a non-standard port, believing it adds an additional security measure.

An SSH client program is typically used for establishing connections to an SSH daemon accepting remote connections. Both are commonly present on most modern operating systems, including Mac OS X, most distributions of GNU/Linux, OpenBSD, FreeBSD, Solaris and OpenVMS. Notably, Windows is one of the few modern desktop/server OSs that does not include SSH by default. Proprietary, freeware and open source versions of various levels of complexity and completeness exist.

SSH is important in cloud computing to solve connectivity problems, avoiding the security issues of exposing a cloud-based virtual machine directly on the Internet. An SSH tunnel can provide a secure path over the Internet, through a firewall to a virtual machine[5].

History and development

Version 1.x

In 1995, Tatu Ylönen, a researcher at Helsinki University of Technology, Finland, designed the first version of the protocol (now called **SSH-1**) prompted by a password-sniffing attack at his university network. The goal of SSH was to replace the earlier rlogin, TELNET and rsh protocols, which did not provide strong authentication nor guarantee confidentiality. Ylönen released his implementation as freeware in July 1995, and the tool quickly gained in popularity. Towards the end of 1995, the SSH user base had grown to 20,000 users in fifty countries.

In December 1995, Ylönen founded SSH Communications Security to market and develop SSH. The original version of the SSH software used various pieces of free software, such as GNU libgmp, but later versions released by SSH Secure Communications evolved into increasingly proprietary software.

It is estimated that, as of 2000, there were 2 million users of SSH.[6]

Notable vulnerabilities

In 1998 a vulnerability was described in SSH 1.5 which allowed the unauthorized insertion of content into an encrypted SSH stream due to insufficient data integrity protection from CRC-32 used in this version of the protocol.[7][8] A fix known as SSH Compensation Attack Detector[9] was introduced into most implementations. Many of these updated implementations contained a new integer overflow vulnerability[10] that allowed attackers to execute arbitrary code with the privileges of the SSH daemon, typically root.

In January 2001 a vulnerability was discovered that allows attackers to modify the last block of an IDEA-encrypted session.[11] The same month, another vulnerability was discovered that allowed a malicious server to forward a client authentication to another server.[12]

Since SSH-1 has inherent design flaws which make it vulnerable, it is now generally considered obsolete and should be avoided by explicitly disabling fallback to SSH-1. Most modern servers and clients support SSH-2.

Version 1.99

In January 2006, well after version 2.1 was established, RFC 4253 specified that an SSH server which supports both 2.0 and prior versions of SSH should identify its protoversion as 1.99.[13] This is not an actual version but a method to identify backward compatibility.

OpenSSH and OSSH

In 1999, developers wanting a free software version to be available went back to the older 1.2.12 release of the original SSH program, which was the last released under an open source license. Björn Grönvall's OSSH was subsequently developed from this codebase. Shortly thereafter, OpenBSD developers forked Grönvall's code and did extensive work on it, creating OpenSSH, which shipped with the 2.6 release of OpenBSD. From this version, a "portability" branch was formed to port OpenSSH to other operating systems. As of 2005, OpenSSH was the single most popular SSH implementation, coming by default in a large number of operating systems. OSSH meanwhile has become obsolete.[14] OpenSSH continues to be maintained and now supports both 1.x and 2.0 versions.

Version 2.x

"Secsh" was the official Internet Engineering Task Force's (IETF) name for the IETF working group responsible for version 2 of the SSH protocol.[15] In 2006, a revised version of the protocol, **SSH-2**, was adopted as a standard. This version is incompatible with SSH-1. SSH-2 features both security and feature improvements over SSH-1. Better security, for example, comes through Diffie-Hellman key exchange and strong integrity checking via message authentication codes. New features of SSH-2 include the ability to run any number of shell sessions over a single SSH connection.[16] Due to SSH-2's superiority and popularity over SSH-1, some implementations such as Lsh[17] and Dropbear[18], only support SSH-2 protocol.

Vulnerabilities

In November 2008, a vulnerability was discovered for all versions of SSH which allowed recovery of up to 32 bits of plaintext from a block of ciphertext that was encrypted using what was then the standard default encryption mode, CBC.[19]

Internet standard documentation

The following RFC publications by the IETF "secsh" working group document SSH-2 as a proposed Internet standard.

- RFC 4250, The Secure Shell (SSH) Protocol Assigned Numbers
- RFC 4251, The Secure Shell (SSH) Protocol Architecture
- RFC 4252, The Secure Shell (SSH) Authentication Protocol
- RFC 4253, The Secure Shell (SSH) Transport Layer Protocol
- RFC 4254, The Secure Shell (SSH) Connection Protocol
- RFC 4255, Using DNS to Securely Publish Secure Shell (SSH) Key Fingerprints
- RFC 4256, Generic Message Exchange Authentication for the Secure Shell Protocol (SSH)
- RFC 4335, The Secure Shell (SSH) Session Channel Break Extension
- RFC 4344, The Secure Shell (SSH) Transport Layer Encryption Modes
- RFC 4345, Improved Arcfour Modes for the Secure Shell (SSH) Transport Layer Protocol

It was later modified and expanded by the following publications.

- RFC 4419, Diffie-Hellman Group Exchange for the Secure Shell (SSH) Transport Layer Protocol (March 2006)
- RFC 4432, RSA Key Exchange for the Secure Shell (SSH) Transport Layer Protocol (March 2006)

- RFC 4462, Generic Security Service Application Program Interface (GSS-API) Authentication and Key Exchange for the Secure Shell (SSH) Protocol (May 2006)
- RFC 4716, The Secure Shell (SSH) Public Key File Format (November 2006)
- RFC 5656, Elliptic Curve Algorithm Integration in the Secure Shell Transport Layer (December 2009)

Uses

SSH is a protocol that can be used for many applications across many platforms including most Unix variants (Linux, the BSD's including Apple's OS X, & Solaris), as well as Microsoft Windows. Some of the applications below may require features that are only available or compatible with specific SSH clients or servers. For example, using the SSH protocol to implement a VPN is possible, but presently only with the OpenSSH server and client implementation.

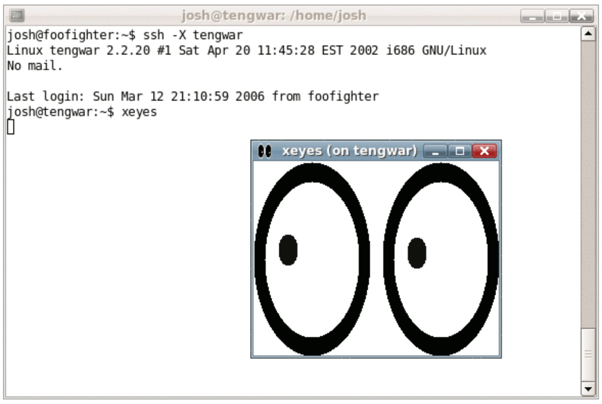

Example of tunneling an X11 application over SSH: the user 'josh' has SSHed from the local machine 'foofighter' to the remote machine 'tengwar' to run xeyes.

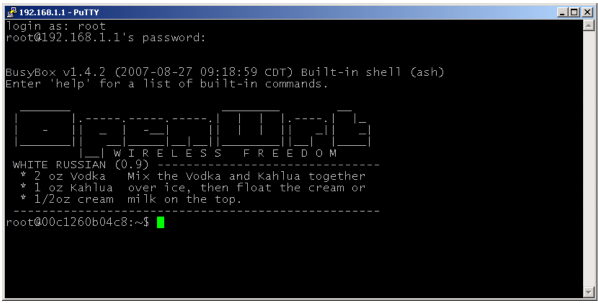

Logging into OpenWrt via SSH using PuTTY running on Windows.

- For login to a shell on a remote host (replacing Telnet and rlogin)
- For executing a single command on a remote host (replacing rsh)
- Secure file transfer
- In combination with rsync to back up, copy and mirror files efficiently and securely
- For forwarding or tunneling a port (not to be confused with a VPN, which routes packets between different networks, or bridges two broadcast domains into one).
- For using as a full-fledged encrypted VPN. Note that only OpenSSH server and client supports this feature.
- For forwarding X from a remote host (possible through multiple intermediate hosts)
- For browsing the web through an encrypted proxy connection with SSH clients that support the SOCKS protocol.
- For securely mounting a directory on a remote server as a filesystem on a local computer using SSHFS.
- For automated remote monitoring and management of servers through one or more of the mechanisms discussed above.
- For development on a mobile or embedded device that supports SSH.

File transfer protocols using SSH

There are multiple mechanisms for transferring files using the Secure Shell protocols.

- Secure copy (SCP), which evolved from RCP protocol over SSH
- rsync, intended to be more efficient than SCP
- SSH File Transfer Protocol (SFTP), a secure alternative to FTP (not to be confused with FTP over SSH)
- Files transferred over shell protocol (a.k.a. FISH), released in 1998, which evolved from Unix shell commands over SSH

Architecture

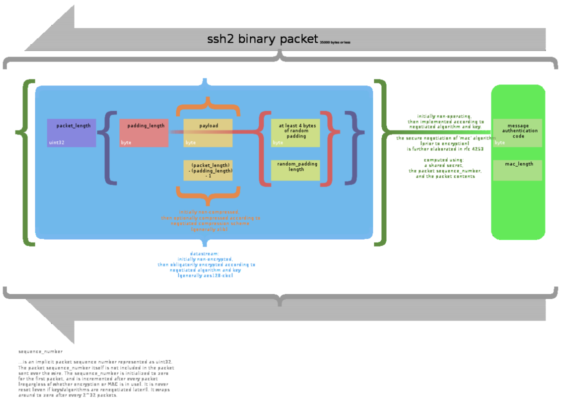

Diagram of the SSH-2 binary packet.

The SSH-2 protocol has an internal architecture (defined in RFC 4251) with well-separated layers. These are:

- The *transport* layer (RFC 4253). This layer handles initial key exchange as well as server authentication, and sets up encryption, compression and integrity verification. It exposes to the upper layer an interface for sending and receiving plaintext packets with sizes of up to 32,768 bytes each (more can be allowed by the implementation). The transport layer also arranges for key re-exchange, usually after 1 GB of data has been transferred or after 1 hour has passed, whichever is sooner.
- The *user authentication* layer (RFC 4252). This layer handles client authentication and provides a number of authentication methods. Authentication is *client-driven*: when one is prompted for a password, it may be the SSH client prompting, not the server. The server merely responds to the client's authentication requests. Widely used user authentication methods include the following:
  - *password*: a method for straightforward password authentication, including a facility allowing a password to be changed. This method is not implemented by all programs.
  - *publickey*: a method for public key-based authentication, usually supporting at least DSA or RSA keypairs, with other implementations also supporting X.509 certificates.
  - *keyboard-interactive* (RFC 4256): a versatile method where the server sends one or more prompts to enter information and the client displays them and sends back responses keyed-in by the user. Used to provide one-time password authentication such as S/Key or SecurID. Used by some OpenSSH configurations when PAM is the underlying host authentication provider to effectively provide password authentication, sometimes leading to inability to log in with a client that supports just the plain *password* authentication method.
  - GSSAPI authentication methods which provide an extensible scheme to perform SSH authentication using external mechanisms such as Kerberos 5 or NTLM, providing single sign on capability to SSH sessions. These methods are usually implemented by commercial SSH implementations for use in organizations, though OpenSSH does have a working GSSAPI implementation.
- The *connection* layer (RFC 4254). This layer defines the concept of channels, channel requests and global requests using which SSH services are provided. A single SSH connection can host multiple channels simultaneously, each transferring data in both directions. Channel requests are used to relay out-of-band channel specific data, such as the changed size of a terminal window or the exit code of a server-side process. The SSH client requests a server-side port to be forwarded using a global request. Standard channel types include:

  - *shell* for terminal shells, SFTP and exec requests (including SCP transfers)
  - *direct-tcpip* for client-to-server forwarded connections
  - *forwarded-tcpip* for server-to-client forwarded connections
- The SSHFP DNS record (RFC 4255) provides the public host key fingerprints in order to aid in verifying the authenticity of the host.

This open architecture provides considerable flexibility, allowing SSH to be used for a variety of purposes beyond a secure shell. The functionality of the transport layer alone is comparable to Transport Layer Security (TLS); the user authentication layer is highly extensible with custom authentication methods; and the connection layer provides the ability to multiplex many secondary sessions into a single SSH connection, a feature comparable to BEEP and not available in TLS.

Enhancements

These are intended for performance enhancements of SSH products:

- SSH-over-SCTP: support for SCTP rather than TCP as the connection oriented transport layer protocol.
- ECDSA: support for elliptic curve DSA rather than DSA or RSA for signing.
- ECDH: support for elliptic curve Diffie-Hellman rather than plain Diffie-Hellman for encryption key exchange.
- UMAC: support for UMAC rather than HMAC for MAC/integrity.

References

[1] Network Working Group of the IETF, January 2006, RFC 4252, The Secure Shell (SSH) Authentication Protocol
[2] SSH Hardens the Secure Shell (http://www.serverwatch.com/news/print.php/3551081), Serverwatch.com
[3] SSH setup manual (http://wiki.qnap.com/wiki/How_To_Set_Up_Authorized_Keys)
[4] port-numbers assignments (http://www.iana.org/assignments/port-numbers) at iana.org
[5] Amies A, Wu C F, Wang G C, Criveti M (2012). Networking on the cloud (http://www.ibm.com/developerworks/cloud/library/cl-networkingtools/index.html) *IBM developerWorks*, June 21.
[6] Nicholas Rosasco and David Larochelle. "How and Why More Secure Technologies Succeed in Legacy Markets: Lessons from the Success of SSH" (http://www.cs.virginia.edu/~drl7x/sshVsTelnetWeb3.pdf). *Quoting Barrett and Silverman,* SSH, the Secure Shell: The Definitive Guide, *O'Reilly & Associates (2001)*. Dept. of Computer Science, Univ. of Virginia. . Retrieved 2006-05-19.
[7] SSH Insertion Attack (http://www.coresecurity.com/content/ssh-insertion-attack)
[8] Weak CRC allows packet injection into SSH sessions encrypted with block ciphers (http://www.kb.cert.org/vuls/id/13877), US-CERT
[9] SSH CRC-32 Compensation Attack Detector Vulnerability (http://www.securityfocus.com/bid/2347/discuss), SecurityFocus
[10] SSH CRC32 attack detection code contains remote integer overflow (http://www.kb.cert.org/vuls/id/945216), US-CERT
[11] Weak CRC allows last block of IDEA-encrypted SSH packet to be changed without notice (http://www.kb.cert.org/vuls/id/315308), US-CERT
[12] SSH-1 allows client authentication to be forwarded by a malicious server to another server (http://www.kb.cert.org/vuls/id/684820), US-CERT
[13] RFC 4253, section 5. Compatibility With Old SSH Versions (http://tools.ietf.org/html/rfc4253#section-5.1), IETF
[14] OSSH Information for VU#419241 (https://www.kb.cert.org/vuls/id/MIMG-6L4LBL)
[15] Secsh Protocol Documents (http://www.vandyke.com/technology/drafts.html), VanDyke Software, Inc.
[16] SSH Frequently Asked Questions (http://www.snailbook.com/faq/ssh-1-vs-2.auto.html)
[17] Official website of Lsh (http://www.lysator.liu.se/~nisse/lsh/)
[18] Official website of Dropbear (https://matt.ucc.asn.au/dropbear/dropbear.html)
[19] SSH CBC vulnerability (http://www.kb.cert.org/vuls/id/958563), US-CERT

Further reading

- Daniel J. Barrett, Richard E. Silverman, and Robert G. Byrnes, *SSH: The Secure Shell (The Definitive Guide)*, O'Reilly 2005 (2nd edition). ISBN 0-596-00895-3
- Michael Stahnke, *Pro OpenSSH*, Apress 2005 ISBN 1-59059-476-2
- Tatu Ylönen (12 July 1995). "Announcement: Ssh (Secure Shell) Remote Login Program" (http://groups.google.com/group/comp.security.unix/msg/67079d812a19f499?dmode=source&hl=en). comp.security.unix. Original announcement of Ssh
- Himanshu Dwivedi; *Implementing SSH*, Wiley 2003. ISBN 978-0-471-45880-7
- This article is based on material taken from the Free On-line Dictionary of Computing prior to 1 November 2008 and incorporated under the "relicensing" terms of the GFDL, version 1.3 or later.

External links

- Old homepage for IETF 'secsh' working group, which has concluded (http://www.ietf.org/html.charters/OLD/secsh-charter.html) (for SSH-2)
- SSH Protocols (http://www.snailbook.com/protocols.html)

WS-Management

Web Services-Management (**WS-Management**) is a DMTF open standard defining a SOAP-based protocol for the management of servers, devices, applications and various Web services. The DMTF has published the standards document DSP0226 with version v1.1.0 of 2010-03-03[1].

The specification is based on DMTF open standards and Internet standards for Web services. WS-Management was originally developed by a coalition of vendors. The coalition started with AMD, Dell, Intel, Microsoft, Sun Microsystems and expanded to a total of 13 members before being subjugated to the DMTF in 2005.

WS-Management provides a common way for systems to access and exchange management information across the IT infrastructure. The specification is quite rich, supporting much more than get/set of simple variables, and in that it is closer to WBEM or Netconf than to SNMP. A mapping of the DMTF-originated Common Information Model into WS-Management was also defined.

Implementations and application support

- Microsoft has implemented the WS-Management standard in Windows Remote Management 1.1 (WinRM),[2] available for Windows XP, Windows Server 2003, Windows Vista and Windows Server 2008.
- Using WS-Management (WinRM 2.0), Windows PowerShell 2.0 allows scripts and cmdlets to be invoked on a remote machine or a large set of remote machines.[3]
- Novell is working to develop an open source implementation of the WS-Management specification for SUSE Linux Enterprise.[4]
- WinRM 2.0 for Windows XP and Windows Server 2003 was released on Oct 26, 2009.[5]
- A European research project (ITEA 2 programme, a strategic pan-European programme for advanced pre-competitive R&D in Software-intensive Systems and Services), named SODA [6] (Service Oriented Device and Delivery Architecture) developed several implementations of WS-Managament in ANSI C, Java, and for OSGi. These implementations are specifically targeted to be used with an open web service protocol stack named DPWS (Devices Profile for Web Services), and were optimized to be integrated in micro-devices with only 100kB of memory. These implementations are open source, and source code is available on the SOA4D forge [7].

References

[1] http://www.dmtf.org/standards/published_documents/DSP0226_1.1.pdf
[2] "Windows Vista Management Features" (http://technet.microsoft.com/de-at/windowsvista/aa905069(en-us).aspx). . Retrieved 2007-12-12.
[3] "What's New in CTP of PowerShell 2.0" (http://blogs.msdn.com/powershell/archive/2007/11/06/what-s-new-in-ctp-of-powershell-2-0.aspx). . Retrieved 2007-12-12.
[4] "Microsoft and Novell Announce Technical Collaboration for Customers" (http://www.novell.com/news/press/microsoft_and_novell_announce_technical_collaboration_for_customers). . Retrieved 2007-12-12.
[5] "Windows Management Framework Core for Windows XP" (http://www.microsoft.com/downloads/details.aspx?FamilyID=60cb5b6c-6532-45e0-ab0f-a94ae9ababf5). . Retrieved 2009-10-28.
[6] http://www.soda-itea.org
[7] http://www.soa4d.org

External links

- WS-Management specifications (http://www.dmtf.org/standards/wsman)
- Openwsman: Open-source implementation of WS-Management (http://www.openwsman.org/)
- Wiseman: Open-source java implementation of WS-Management (https://wiseman.dev.java.net/)
- SOA4D (Service Oriented Architecture for Devices): Open-source C and Java implementation of DPWS stack and WS-Management (http://www.soa4d.org)
- A live WS-Management example for experimentation (http://wsman.msft.net)
- WinRM (Windows Remote Management): a Microsoft Windows implementation of WS-Management Protocol based on SOAP (Simple Object Access Protocol) (http://msdn.microsoft.com/en-us/library/aa384426.aspx)

Group Policy

Group Policy is a feature of the Microsoft Windows NT family of operating systems that control the working environment of user accounts and computer accounts. Group Policy provides the centralized management and configuration of operating systems, applications, and users' settings in an *Active Directory* environment.

Operation

Group Policy in part controls what users can and cannot do on a computer system, for example: to enforce a password complexity policy that prevents users from choosing an overly simple password, to allow or prevent unidentified users from remote computers to connect to a network share, to block access to the Windows Task Manager or to restrict access to certain folders. A set of such configurations is called a Group Policy Object (GPO).

As part of Microsoft's *IntelliMirror* technologies, Group Policy aims to reduce the cost of supporting users. IntelliMirror technologies relate to the management of disconnected machines or roaming users and include *roaming user profiles*, *folder redirection*, and *offline files*.

Although Group Policy is more often seen in use for enterprise environments, it is also common in schools, smaller businesses, and other kinds of smaller organizations.

Enforcement

To accomplish the goal of central management of a group of computers, machines should receive and enforce GPOs. A GPO that resides on a single machine only applies to that computer. To apply a GPO to a group of computers, Group Policy relies on Active Directory (or third party products like ZENworks Desktop Management) for distribution. Active Directory can distribute GPOs to computers that are part of a Windows domain.

By default, Microsoft Windows refreshes its policy settings every 90 minutes with a random 30 minutes offset. On Domain controllers, Microsoft Windows does so every five minutes. During the refresh, it discovers, fetches and applies all GPOs that apply to the machine and logged on user. Some settings, such as automated software installation, drive mappings, startup scripts or logon scripts are only applied during startup or user logon. Since Windows XP, a refresh of the group policy can be manually initiated by the user using the `gpupdate` command from a command prompt.[1]

Group Policy Objects are processed in the following order (from top to bottom):[2]

1. **Local** - Any settings in the computer's local policy. Prior to Windows Vista, there was only one local group policy stored per computer. Windows Vista and later Windows versions allow individual group policies per user accounts.[3]
2. **Site** - Any Group Policies associated with the *Active Directory site* in which the computer resides. (An Active Directory site is a logical grouping of computers that is meant to facilitate management of computers based on their physical proximity.) If multiple policies are linked to a site, they are processed in the order set by the administrator.
3. **Domain** - Any Group Policies associated with the Windows domain in which the computer resides. If multiple policies are linked to a domain, they are processed in the order set by the administrator.
4. **Organizational Unit** - Group policies assigned to the *Active Directory organizational unit (OU)* in which the computer or user are placed. (OUs are logical units that help organizing and managing a group of users, computers or other Active Directory objects.) If multiple policies are linked to an OU, they are processed in the order set by the administrator.

Inheritance

A policy setting inside a hierarchical structure is ordinarily passed from parent to children, and from children to grandchildren, and so forth. This is termed inheritance. It can be blocked or enforced to control what policies are applied at each level. If a higher level administrator (enterprise administrator) creates a policy that has inheritance blocked by a lower level administrator (domain administrator), this policy will still be processed.

Where a Group Policy Preference Settings is configured and there is also an equivalent Group Policy Setting configured, then the value of the Group Policy Setting will take precedence.

Filtering

WMI filtering is the process of customizing the scope of the GPO by choosing a Windows Management Instrumentation (WMI) filter to apply.

Local Group Policy

Local Group Policy (LGP) is a more basic version of the Group Policy used by Active Directory. Prior to Windows Vista, LGP could enforce a GPO for a single local computer, but could not make policies for individual users or groups. Windows Vista allows setting local Group Policy for individual users.[4] LGP can be applied to a computer on a domain, and it can be used on *Windows XP Home Edition*.

Group Policy Preferences

There is a set of group policy setting extensions that were previously known as PolicyMaker. Microsoft bought PolicyMaker and then integrated them with Windows Server 2008. Microsoft has since released a migration tool that allows users to migrate PolicyMaker items to Group Policy Preferences.[5]

Group Policy Preferences adds a number of new configuration items. These items also have a number of additional targeting options that can be used to granularly control the application of these setting items.

Group Policy Preferences are compatible with x86 and x64 versions of Windows XP, Windows Server 2003, and Windows Vista with the addition of the Client Side Extensions (also known as CSE).[6][7][8][9][10][11]

Client Side Extensions are now included in Windows Server 2008, Windows 7, and Windows Server 2008 R2.

Group Policy Management Console

Originally, Group Polices were modified using the Group Policy Edit tool that was integrated with Active Directory Users and Computers Microsoft Management Console (MMC) snap-in, but it was later split into a separate MMC snap-in called the Group Policy Management Console (GPMC). The GPMC is now a user component in Windows Server 2008 and Windows Server 2008 R2 and is provided as a download as part of the Remote Server Administration Tools for Windows Vista and Windows 7.[12][13][14][15]

Advanced Group Policy Management

Microsoft has also released a tool to make changes to Group Policy called Advanced Group Policy Management[16] (a.k.a. AGPM). This tool available for any organisation that has licensed the Microsoft Desktop Optimization Pack (a.k.a. MDOP). This advanced tool allows administrators to have a check in/out process for modification Group Policy Objects, track changes to Group Policy Objects, and implement approval workflows for changes to Group Policy Objects.

To use this software you must license all of your Windows Active Directory clients for MDOP.

AGPM consists of two parts - server and client. The server is a Windows Service that stores its Group Policy Objects in an archive located on the same computer or a network share. The client is a snap-in to the Group Policy Management Console, and connects to the AGPM server. Configuration of the client is performed via Group Policy.

Security

Group Policy settings are enforced voluntarily by the targeted applications. In many cases, this merely consists of disabling the user interface for a particular function without disabling lower-level means of accessing it.[17]

Alternatively, a malevolent user can modify or interfere with the application so that it cannot successfully read its Group Policy settings, thus enforcing potentially lower security defaults or even returning arbitrary values.[18]

Windows 8 Enhancements

Windows 8 introduced a new feature called Group Policy Update. This feature allows an administrator to force a group policy update on all computers with accounts in a particular Organizational Unit. This update creates a scheduled task on the computer and it will run the GPUPDATE command within 10 minutes with a random offset to not over load the domain controller.

Group Policy Infrastructure Status which can report when any Group Policy Objects are not replicated correct amongst domain controllers.[19]

Group Policy Results Report also has a new feature that times the execution of individual components when doing a Group Policy Update.[20]

References

[1] Gpupdate (http://technet.microsoft.com/en-us/library/bb490983.aspx)
[2] "Group Policy processing and precedence" (http://technet.microsoft.com/en-us/library/cc785665.aspx). Microsoft Corporation. 22 April 2012. .
[3] Group Policy - Apply to a Specific User or Group - Windows 7 Forums (http://www.sevenforums.com/tutorials/151415-group-policy-apply-specific-user-group.html)
[4] Step-by-Step Guide to Managing Multiple Local Group Policy Objects (http://go.microsoft.com/fwlink/?LinkID=73434)
[5] Group Policy Preference Migration Tool (GPPMIG) (http://www.microsoft.com/downloads/details.aspx?FamilyID=35791cb6-710b-48c4-aaa1-90db170bcf2a&displaylang=en)
[6] Group Policy Preference Client Side Extensions for Windows XP (KB943729) (http://www.microsoft.com/downloads/details.aspx?displaylang=en&FamilyID=e60b5c8f-d7dc-4b27-a261-247ce3f6c4f8)
[7] Group Policy Preference Client Side Extensions for Windows XP x64 Edition (KB943729) (http://www.microsoft.com/downloads/details.aspx?displaylang=en&FamilyID=249c1aed-c1f1-4a0b-872e-ef0a32170625)
[8] Group Policy Preference Client Side Extensions for Windows Vista (KB943729) (http://www.microsoft.com/downloads/details.aspx?familyid=AB60DC87-884C-46D5-82CD-F3C299DAC7CC&displaylang=en)
[9] Group Policy Preference Client Side Extensions for Windows Vista x64 Edition (KB943729) (http://www.microsoft.com/downloads/details.aspx?familyid=B10A7AF4-8BEE-4ADC-8BBE-9949DF77A3CF&displaylang=en)
[10] Group Policy Preference Client Side Extensions for Windows Server 2003 (KB943729) (http://www.microsoft.com/downloads/details.aspx?familyid=BFE775F9-5C34-44D0-8A94-44E47DB35ADD&displaylang=en)
[11] Group Policy Preference Client Side Extensions for Windows Server 2003 x64 Edition (KB943729) (http://www.microsoft.com/downloads/details.aspx?familyid=29E83503-7686-49F3-B42D-8E5ED23D5D79&displaylang=en)
[12] Microsoft Group Policy Team (2009-12-23). "How to Install GPMC on Server 2008, 2008 R2, and Windows 7 (via RSAT)" (http://blogs.technet.com/grouppolicy/archive/2009/12/23/how-to-install-rsat.aspx). .
[13] Microsoft Remote Server Administration Tools for Windows Vista (http://www.microsoft.com/downloads/details.aspx?displaylang=en&FamilyID=9ff6e897-23ce-4a36-b7fc-d52065de9960)
[14] Microsoft Remote Server Administration Tools for Windows Vista for x64-based Systems (http://www.microsoft.com/downloads/details.aspx?displaylang=en&FamilyID=d647a60b-63fd-4ac5-9243-bd3c497d2bc5)
[15] Remote Server Administration Tools for Windows 7 (http://www.microsoft.com/downloads/details.aspx?FamilyID=7d2f6ad7-656b-4313-a005-4e344e43997d&displaylang=en)
[16] Microsoft Windows Enterprise | Enhancing Group Policy (http://www.microsoft.com/windows/enterprise/products/mdop/agpm.aspx)
[17] Raymond Chen, "Shell policy is not the same as security" (http://blogs.msdn.com/oldnewthing/archive/2008/01/31/7337160.aspx)
[18] Mark Russinovich, "Circumventing Group Policy as a Limited User (http://blogs.technet.com/b/markrussinovich/archive/2005/12/12/circumventing-group-policy-as-a-limited-user.aspx)
[19] Updated: What's new with Group Policy in Windows 8 (http://www.grouppolicy.biz/2011/10/whats-new-with-group-policy-in-windows-8/)
[20] Windows 8 Group Policy Performance Troubleshooting Feature (http://www.grouppolicy.biz/2012/01/windows-8-group-policy-performance-troubleshooting-feature/)

Further reading

1. "Group Policy for Beginners" (http://technet.microsoft.com/en-us/library/hh147307.aspx). *Windows 7 Technical Library*. Microsoft. 27 April 2011. Retrieved 22 April 2012.
2. "Group Policy Management Console" (http://msdn.microsoft.com/en-us/library/windows/desktop/aa814316(v=vs.85).aspx). *Dev Center - Desktop*. Microsoft. 3 February 2012. Retrieved 22 April 2012.
3. "Step-by-Step Guide to Managing Multiple Local Group Policy Objects" (http://technet.microsoft.com/en-us/library/cc766291.aspx). *Windows Vista Technical Library*. Microsoft. Retrieved 22 April 2012.
4. "Group Policy processing and precedence" (http://technet.microsoft.com/en-us/library/cc785665.aspx). *Windows Server 2003 Product Help*. Microsoft. 21 January 2005. Retrieved 22 April 2012.

External links

- Official website (http://www.microsoft.com/grouppolicy)
- Group Policy Team Blog (http://blogs.technet.com/grouppolicy)
- Group Policy Settings Reference for Windows and Windows Server (http://www.microsoft.com/download/en/details.aspx?displaylang=en&id=25250)

System Center Configuration Manager

System Center Configuration Manager

Developer(s)	Microsoft Corporation
Stable release	Configuration Manager 2012 RTM / March 2012
Development status	Active
Operating system	Microsoft Windows
Platform	x86 x64
Translation available	Multilingual
Type	Systems management
License	MS-EULA
Website	Microsoft System Center Configuration Manager [1]

System Center Configuration Manager (CM12 *or* CM07 *or* ConfigMgr *or* Configuration Manager), formerly **Systems Management Server** (SMS), is a systems management software product by Microsoft for managing large groups of Windows-based computer systems. Configuration Manager provides remote control, patch management, software distribution, operating system deployment, network access protection, and hardware and software inventory.

There have been three major iterations of SMS. The 1.x versions of the product defined the scope of control of the management server (the site) in terms of the NT domain that was being managed. Since the 2.x versions, that site paradigm has switched to a group of subnets that will be managed together. Since SMS 2003, the site could also be defined as one or more Active Directory sites. The most frequently used feature is inventory management, which provides both hardware and software inventory across a business enterprise.

The major difference between the 2.x product and SMS 2003 is the introduction of the Advanced Client. The Advanced Client communicates with a more scalable management infrastructure, namely the Management Point. A Management Point (MP) can manage up to 25000 Advanced Clients.

The Advanced Client was introduced to provide a solution to the problem that a managed laptop might connect to a corporate network from multiple locations and should not always download content from the same place within the enterprise (though it should always receive policy from its own site). When an Advanced Client is within another location (SMS Site), it may use a local distribution point to download or run a program which can conserve bandwidth across a WAN.

The current generation of the product, System Center Configuration Manager 2012, was initially released in March 2012.[2]

Version history

Product	Revision	Released	Service Pack	Feature Pack	Version/Build	Notes
Systems Management Server (SMS)	1.0	1994				
Systems Management Server (SMS)	1.1	1995				
Systems Management Server (SMS)	1.2	1996				
Systems Management Server (SMS)	2.0	1999				
Systems Management Server (SMS)	2003	2003				
Systems Management Server (SMS)	2003 R2	2006				
System Center Configuration Manager (ConfigMgr)	2007	2007	Beta 1		4.00.5135.0000	
System Center Configuration Manager (ConfigMgr)	2007	2007	RTM		4.00.5931.0000	
System Center Configuration Manager (ConfigMgr)	2007	2008 (May)	SP1		4.00.6221.1000	
System Center Configuration Manager (ConfigMgr)	2007	2010	SP1	KB977203	4.00.6221.1193[3]	This update (KB 977203) may be run against SP1 or SP2 clients.
System Center Configuration Manager (ConfigMgr)	2007	2008		R2	no change	The R2 feature add-on requires at least SP1, and can be installed after SP2.
System Center Configuration Manager (ConfigMgr)	2007	2009	SP2		4.00.6487.2000	
System Center Configuration Manager (ConfigMgr)	2007	2010	SP2	KB977203	4.00.6487.2111[4]	This update (KB 977203) may be run against SP1 or SP2 clients.
System Center Configuration Manager (ConfigMgr)	2007	2010[5]		R3	4.00.6487.2157	The R3 update requires SP2 Client.
System Center Configuration Manager (ConfigMgr)	2012 (Formerly "v.Next")[6]	2010/05/26	Beta 1			
System Center Configuration Manager (ConfigMgr)	2012	2011/03/23	Beta 2		5.00.7561.0000	
System Center Configuration Manager (ConfigMgr)	2012	2011/10/28	Release Candidate 1		5.00.7678.0000	

System Center Configuration Manager (ConfigMgr)	2012	2012/01/17	Release Candidate 2[7]		5.00.7703.0000	
System Center Configuration Manager (ConfigMgr)	2012	2012/03/31	RTM		5.00.7711.0000	

External links

- System Center Configuration Manager homepage [1]
- Configuration Manager TechCenter on TechNet [8]
- myITforum.com [9]
- Windows Management User Group [10]
- FAQShop (FAQ specifically for SMS and for ConfigMgr) [11]
- SMSUG.ca User Group for SMS in Canada [12]
- The Configuration Manager Support Team Blog [13]
- Nexus SC: The System Center Team Blog [14]
- SCCM - System Center Configuration Manager Blog [15]
- Configuration Manager OSD Feature Team Blog [16]
- Managing Privileged Accounts from within System Center Configuration Manager (Product Datasheet from Lieberman Software) [17]
- QMX natively extends the capabilities of SCCM to Non-Windows domain [18]
- Community of SCCM Smart Agents for: Applications, Hardware, Infrastructure, Network, Operating Systems, Power, Printers, Security, Storage, Telecom, Virtualization and more. [19]

References

[1] http://www.microsoft.com/systemcenter/configurationmanager/en/us/default.aspx
[2] http://www.microsoft.com/en-us/server-cloud/system-center/configuration-manager-2012.aspx
[3] http://support.microsoft.com/kb/977203
[4] http://support.microsoft.com/kb/977203
[5] http://www.zdnet.com/blog/microsoft/microsoft-messenger-connect-configuration-manager-2007-r3-no-longer-in-beta/7679
[6] http://blogs.technet.com/b/systemcenter/archive/2010/11/09/announcing-configuration-manager-v-next-official-name.aspx
[7] http://www.microsoft.com/systemcenter/en/us/configuration-manager/cm-vnext-beta.aspx
[8] http://technet.microsoft.com/configmgr/default.aspx
[9] http://www.myITforum.com
[10] http://wmug.co.uk
[11] http://www.faqshop.com/configmgr2007/default.htm
[12] http://www.SMSUG.ca
[13] http://blogs.technet.com/configurationmgr/default.aspx
[14] http://blogs.technet.com/systemcenter/default.aspx
[15] http://scug.be/blogs/sccm/default.aspx
[16] http://blogs.technet.com/b/inside_osd
[17] http://www.liebsoft.com/uploadedFiles/wwwliebsoftcom/MARCOM/Solutions/SCCM_Datasheet.pdf
[18] http://www.quest.com/management-xtensions-device-management-cm
[19] http://communities.quest.com/community/management-extensions/sccm

User Accounts: Identities and Authentication

passwd (file)

In Unix-like operating systems the **`/etc/passwd`** file is a text-based database of information about users that may login to the system or other operating system user identities that own running processes.

In many operating systems this file is just one of many possible back-ends for the more general passwd name service.

The file's name originates from one of its initial functions as it contained the data used to verify passwords of user accounts. However, on modern Unix systems the security-sensitive password information is instead often stored in a different file using shadow passwords, or other database implementations.

The /etc/passwd file typically has file system permissions that allow it to be readable by all users of the system (*world-readable*), although it may only be modified by the superuser or by using a few special purpose privileged commands.

File format

The /etc/passwd file is a text file with one record per line, each describing a user account. Each record consists of seven fields separated by colons.[1] The ordering of the records within the file is generally unimportant.

An example record may be:

```
jsmith:x:1001:1000:Joe Smith,Room 1007,(234)555-8910,(234)555-0044,email:/home/jsmith:/bin/sh
```

The fields, in order from left to right, are[2]:

1. The first field is the user name, i.e. the string a user would type in when logging into the operating system: the logname. Each record in the file must have a unique user name field.
2. The second field stores information used to validate a user's password; however in most modern uses this field is usually set to "x" (or some other indicator) with the actual password information being stored in a separate shadow password file. Setting this field to an asterisk "*" is the typical way to deactivate an account to prevent it being used.
3. The third field is the user identifier, the number that the operating system uses for internal purposes. It does not have to be unique.
4. The fourth field is the group identifier. This number identifies the primary group of the user; all files that are created by this user may initially be accessible to this group.
5. The fifth field, called the Gecos field, is commentary that describes the person or account. Typically, this is a set of comma-separated values including the user's full name and contact details.
6. The sixth field is the path to the user's home directory.
7. The seventh field is the program that is started every time the user logs into the system. For an interactive user, this is usually one of the system's command line interpreters (shells).

References

[1] Introduction to Processes and the Unix Shell (http://www.rwc.uc.edu/thomas/Intro_Unix_Text/Process.html)
[2] Understanding /etc/passwd File Format (http://www.cyberciti.biz/faq/understanding-etcpasswd-file-format/)

External links

- Manual page from Unix First Edition describing /etc/passwd (http://man.cat-v.org/unix-1st/5/passwd)

Shadow password

In computing, Unix-like operating systems use the **shadow password** database mechanism to increase the security level of passwords by restricting all but highly privileged users' access to encrypted password data. Typically, that data is kept in files owned by and accessible only by the super user (i.e., on Unix-like systems, the *root user*, and on many others, the *administrator account*).

Design

Systems administrators can reduce the likelihood of brute force attacks by making the list of hashed passwords unreadable by unprivileged users. The obvious way to do this is to make the `passwd` database itself readable only by the root user. However, this would restrict access to other data in the file such as username-to-userid mappings, which would break many existing utilities and provisions. One solution is a "shadow" password file to hold the password hashes separate from the other data in the world-readable *passwd* file. For local files, this is usually `/etc/shadow` on Linux and Unix systems, or `/etc/master.passwd` on BSD systems; each is readable only by *root*. (Root access to the data is considered acceptable since on systems with the traditional "all-powerful root" security model, the root user would be able to obtain the information in other ways in any case). Virtually all recent Unix-like operating systems use shadowed passwords.

The shadow password file does not entirely solve the problem of attacker access to hashed passwords, as some network authentication schemes operate by transmitting the encrypted password over the network (sometimes in cleartext, eg: Telnet [1]), making it vulnerable to interception. Copies of system data, such as system backups written to tape or optical media, can also become a means for illicitly obtaining hashed passwords. In addition, the functions used by legitimate password-checking programs need to be written in such a way that malicious programs cannot make large numbers of authentication checks at high rates of speed.

Usage

On a system without shadowed passwords (typically older Unix systems dating from before 1990 or so), the *passwd* file holds the following user information for each user account:

- Username
- Salt combined with the current hash of the user's password (usually produced from a cryptographic hash function)
- Password expiration information
- User ID (UID)
- Default group ID (GID)
- Full name
- Home directory path
- Login shell

The passwd file is readable by all users so that name service switch can work (e.g., to ensure that user names are shown when the user lists the contents of a folder), but only the root user can write to it. This means that an attacker with unprivileged access to the system can obtain the hashed form of every user's password. Those values can be used to mount a brute force attack offline, testing possible passwords against the hashed passwords relatively quickly without alerting system security arrangements designed to detect an abnormal number of failed login attempts. Users often select passwords vulnerable to such password cracking techniques.[2]

With a shadowed password scheme in use, the `/etc/passwd` file typically shows a character such as '`*`', or '`x`' in the password field for each user instead of the hashed password, and `/etc/shadow` usually contains the following user information:

- User login name
- salt and hashed password OR a status exception value e.g.:
  - "idsalt$encrypted", where "$id" is the hashing algorithm used (On GNU/Linux, "`$1$`" stands for MD5, "`$2$`" is Blowfish, "`$5$`" is SHA-256 and "`$6$`" is SHA-512, crypt(3) manpage [3], other Unix may have different values, like NetBSD [4]).
  - "NP" or "!" or null - No password, the account has no password.
  - "LK" or "*" - the account is Locked, user will be unable to log-in
  - "!!" - the password has expired
- Days since epoch of last password change
- Days until change allowed
- Days before change required
- Days warning for expiration
- Days before account inactive
- Days since Epoch when account expires
- Reserved

The format of the shadow file is simple, and basically identical to that of the password file, to wit, one line per user, ordered fields on each line, and fields separated by colons. Many systems require the order of user lines in the shadow file be identical to the order of the corresponding users in the password file.

To modify the contents of the shadow file on most systems, users generally invoke the `passwd` program, which in turn largely depends on PAM. For example, the type of hash used is dictated by the configuration of the `pam_unix.so` module. By default, the MD5 hash has been used, while current modules are also capable of stronger hashes such as blowfish, SHA256 and SHA512.

History

Password shadowing first appeared in UNIX systems with the development of System V Release 3.2 in 1988 and BSD4.3 Reno in 1990. But, vendors who had performed ports from earlier UNIX releases did not always include the new password shadowing features in their releases, leaving users of those systems exposed to password file attacks.

System administrators may also arrange for the storage of passwords in distributed databases such as NIS and LDAP, rather than in files on each connected system. In the case of NIS, the shadow password mechanism is often still used on the NIS servers; in other distributed mechanisms the problem of access to the various user authentication components is handled by the security mechanisms of the underlying data repository.

In 1987 the author of the original *Shadow Password Suite*, Julie Haugh, experienced a computer break-in and wrote the initial release of the Shadow Suite containing the `login`, `passwd` and `su` commands. The original release, written for the SCO Xenix operating system, quickly got ported to other platforms. The Shadow Suite was ported to Linux in 1992 one year after the original announcement of the Linux project, and was included in many early distributions, and continues to be included in many current Linux distributions.

References

[1] http://tools.ietf.org/html/rfc2877
[2] Rob Lemos (2002-05-22). "Passwords: the weakest link?" (http://www.news.com/2009-1001-916719.html). *CNET News.com*. . Retrieved 2008-02-19.
[3] http://www.kernel.org/doc/man-pages/online/pages/man3/crypt.3.html
[4] http://netbsd.gw.com/cgi-bin/man-cgi?crypt+3+NetBSD-current

External links

- authconfig (http://linux.die.net/man/8/authconfig), a command-line tool for controlling the use of shadow passwords
- An example shadow file (http://configuration.logfish.net/index.php/etc/shadow), showing the general layout of the file

Security Accounts Manager

The **Security Accounts Manager** (SAM) file in Windows XP, Windows Vista and Windows 7. It stores users' passwords in a hashed format (in LM hash and NTLM hash). Since a hash function is one-way, this provides some measure of security for the storage of the passwords.

In an attempt to improve the security of the SAM database against offline software cracking, Microsoft introduced the SYSKEY function in Windows NT 4.0. When SYSKEY is enabled, the on-disk copy of the SAM file is partially encrypted, so that the password hash values for all local accounts stored in the SAM are encrypted with a key (usually also referred to as the "SYSKEY").

In the case of online attacks, it is not possible to simply copy the SAM file to another location. The SAM file cannot be moved or copied while Windows is running, since the Windows kernel obtains and keeps an exclusive filesystem lock on the SAM file, and will not release that lock until the operating system has shut down or a "Blue Screen of Death" exception has been thrown. However, the in-memory copy of the contents of the SAM can be dumped using various techniques, making the password hashes available for offline brute-force attack.

Removing LM hash

Most versions of Windows can be configured to disable the creation and storage of valid LM hashes when the user changes their password. This is the default setting in Windows Vista, but was disabled by default in previous versions of Windows. Note: enabling this setting does not immediately clear the LM hash values from the SAM, but rather enables an additional check during password change operations that will instead store a "dummy" value in the location in the SAM database where the LM hash is otherwise stored. (This dummy value has no relationship to the user's password — it is the same value used for all user accounts.)

As well, LM hashes cannot be calculated when the user chooses a password of over fourteen characters in length. Thus, when a user (or administrator) sets a password of fifteen characters or longer, the LM hash value is set to a "dummy" value, which is not valid for authentication purposes.

Related attacks

In Windows NT 3.51, NT 4.0 and 2000, an attack was devised to bypass the local authentication system. If the SAM file is deleted from the hard drive (e.g. mounting the Windows OS volume into an alternate operating system), the attacker could log in as any account with no password. This flaw was corrected with Windows XP, which shows an error message and shuts down the computer. However there exist recently developed software utilities which, by the aforementioned methodology of using either an emulated virtual drive, or boot disk (usually Unix/Linux) based environment to mount the local drive housing the active NTFS partition, and using programmed software routines and function calls from within assigned memory stacks to isolate the SAM file from the Windows NT system installation directory structure (default: `c:\windows\system32\config`) and, depending on the particular software utility being used, remove the password hashes stored for user accounts in their entirety, or in some cases, modify the user account passwords directly from this environment.

This software has both a highly pragmatic and beneficial use as a password clearing or account recovering utility for individuals who have lost or forgotten their windows account passwords, as well as a possible use as a malicious software security bypassing utility. Essentially granting a user with enough ability, experience, and familiarity with both the cracking utility software and the security routines of the Windows NT kernel (as well as offline and immediate local access to the target computer) the capability to entirely bypass/remove the windows account passwords from a potential target computer. Only recently, Microsoft released a utility called LockSmith, which is part of MSDart. MSDart is not freely available to end-users, however.

Home directory

A **Home directory** is a file system directory on a multi-user operating system containing files for a given user of the system. The specifics of the home directory (such as its name and location) is defined by the operating system involved; for example, Windows systems between 2000 and 2003 keep home directories in a folder called *Documents and Settings*.

Description

A user's home directory is intended to contain that user's files; including text documents, music, pictures or videos, etc. It may also include their configuration files of preferred settings for any software they have used there and might have tailored to their liking: web browser bookmarks, favorite desktop wallpaper and themes, passwords to any external services accessed via a given software, etc. The user can install executable software in this directory, but it will only be available to users with permission to this directory. The home directory can be organized further with the use of sub-directories.

The content of a user's home directory is protected by file system permissions, and by default is only accessible to that user and administrators. Any other user that has been granted administrator privileges has authority to access any protected location on the filesystem including other users home directories.

Benefits of Home Directories

Separating user data from system-wide data avoids redundancy and makes backups of important files relatively simple. Furthermore, Trojan horses, viruses and worms running under the user's name and with their privileges will in most cases only be able to alter the files in the user's home directory, and perhaps some files belonging to workgroups the user is a part of, but not actual system files.

Default Home Directory per Operating System

System	Path	Variable
Microsoft Windows NT	<root>\WINNT\Profiles\<username>	%UserProfile%
Microsoft Windows 2000, XP and 2003	<root>\Documents and Settings\<username>	
Microsoft Windows Vista and 7	<root>\Users\<username>	
Unix-Based [1]	<root>/home/<username>	$HOME and ~/
Unix-Derived	/var/users/<username> /u01/<username> /usr/<username> /user/<username> /users/<username>	
SunOS / Solaris	/export/home/<username>	
Linux (FHS)	/home/<username>	
AT&T Unix (original version)	<root>/usr/<username>	$HOME
Mac OS X	/Users/<username>	$HOME and ~/, and path to home folder (in AppleScript)
OpenVMS	<device>:[<username>]	SYS$LOGIN

Other Features per Operating System

Unix

In Unix, a user will be automatically placed into their home directory upon login. The `~user` shorthand variable refers to a user's home directory (allowing the user to navigate to it from anywhere else in the filesystem, or use it in other Unix commands). The ~ (tilde character) shorthand command refers to that particular user's home directory.

The Unix superuser has access to all directories on the filesystem, and hence can access home directories of all users. The superuser's home directory on older systems was `/`, but on many newer systems it is located at `/root` (Linux, BSD), or `/var/root` (Mac OS X).

VMS

In the OpenVMS operating system, a user's home directory is called the "root directory", and the equivalent of a Unix/DOS/Windows/AmigaOS "root directory" is referred to as the "Master File Directory".

Contrast with Single-user Operating Systems

Single-user operating systems simply have a single directory or partitions for all users files, there is no individual directory setup per user (though users can still setup and maintain directories inside this main working directory manually).

- AmigaOS versions 2 and up have "System" and "Work" partitions on hard disks by default.

- BeOS (and its successors) have a /home directory which contain the files belonging to the single user of the system.
- Versions of Windows prior Windows 95 OEM Service Release 2 did not have a user folder, but since that release, `\My Documents` became in effect the single user's home directory.
- NeXTSTEP and OPENSTEP in a single user, non-networked setup, `/me` is used, as well as `/root` when logged in as superuser.

References

[1] "Home Directory Definition" (http://www.linfo.org/home_directory.html). Accessed on July 23, 2009

Microsoft Windows profile

Microsoft Windows profile refers to the user profile that is used by the Microsoft Windows operating system to represent the characteristics of the user.

Windows XP

Profile creation

Establishing a user account on the computer (or on its parent domain) *does not*, in fact, create a profile for that user. The profile is created the first time the user interactively logs on at the computer. Note that logging on across a network to access shared folders does not create a profile.

At first logon, a folder will typically be created under "Documents and Settings" (standard folder on English version of Windows 2000, XP and Windows Server 2003) matching the logon name of the user. Should a folder of that name already exist, the profile-creation process will create a new one, typically named username.computername, on workgroup computers, or username.domainname on Active Directory member computers.

Once a profile folder has been created, Windows will never automatically rename that folder. Thus if the username itself is subsequently changed, the profile folder will remain as is, and the profile will no longer match the username, which could lead to confusion. For this reason the administrator might want to avoid renaming user accounts if at all possible, or rename the folder manually and edit the registry to reflect the changes.

The new profile is in fact created by making a copy of a special profile, named Default User. It is permissible to modify this Default User profile (within certain guidelines) so as to provide a customised working environment for each new user. Modification of the Default User profile should ideally be done from new, prior to any users logging-on to the computer; if a user has already logged on once or more, the Default Profile has no effect whatsoever for that user.

Profile contents

NTUSER.DAT - Within the root of the profile, a file named NTUSER.DAT contains the user's personalized settings for the majority of software installed on the computer, including Windows itself. When the user logs on, NTUSER.DAT becomes merged with the computer's registry, such that it appears as the HKEY_CURRENT_USER branch of the registry tree. NTUSER.DAT is held open for writing (i.e., *locked*) whenever the user is logged on. Paradoxically, that lock means that this file cannot be easily copied by its owner.

My Documents - This folder is intended to contain the user's work, and in Windows XP-aware programs, dialog boxes will typically prompt the user to store documents here. 'My Documents' as a shortcut also appears on the desktop, and in My Computer. It is here that these shortcuts point.

Favorites, Cookies and History - These folders are used by Microsoft's Internet Explorer web browser to store surfing data. They are not used by alternative browsers such as Firefox or Opera, which typically store their data under "Application Data."

Nethood, Printhood - These folders contain the network shares and printers 'discovered' by the user with the My Network Places applet, in the form of shortcuts.

Start Menu - This folder contains the shortcuts present on the same-named Desktop feature.

Desktop - Similarly, this folder contains files and shortcuts present on the user's desktop.

Application Data is provided mainly for the use of programmers, as a place to store data that is related to specific software, but which does not fall into the category of documents that a user might open directly. This folder was made necessary by Windows' 'good practice' programming guidelines, which now prohibit the storage of temporary data of any kind in the Program Files folder.

Local Settings is functionally similar to *Application Data*, and in fact it contains a second subfolder of that name. It also contains the temporary files generated by Windows programs themselves, and as a result of Internet Explorer's online activities. For standalone computers the two folders are functionally similar, but on networks employing Roaming profiles, the *Local Settings* folder is *not* included in the profile-synchronization process. Thus, data in the *Local Settings* folder will not be copied between computers when the user roams.

Note: Some of these subfolders are hidden from the user's view in Explorer. To see them you must uncheck *Hide System Folders* in the folder options.

Special profiles

Default User - Plays a role in the profile-creation process, see above.

All Users - This profile is present mainly to answer an issue related to software installation. It provides a way for setup-programs to create desktop or start-menu shortcuts which will be visible to all users of the computer, not just the user running the setup program. The Application Data section may also contain program-data common to all users. *All Users* acts purely as an information-store, it is never loaded as an active profile.

Administrator - All versions of NT-based Windows have an administrator account and corresponding profile, although on XP this account may only be visible on the log-on screen if the computer is started in safe mode. In Windows Vista, it is disabled by default.

History and origin

Historically, the Windows 95/98 product line did not employ user-profiling as standard, with all users sharing the same settings, although that feature could be activated in Control Panel.

The user-profiling scheme in force today owes its origins to Windows NT, which stored its profiles within the system folder itself, typically under C:\WINNT\Profiles\. Windows 2000 saw the change to a separate "Documents and Settings" folder for profiles, and in this respect is virtually identical to Windows XP and Windows Server 2003.

Recent developments

Windows Vista's profiles are functionally similar to those of Windows XP, but with some differences. Perhaps the key difference is that they are stored in a "C:\Users" folder, instead of "C:\Documents and Settings." Therefore, to cater for programs which are not Vista-ready, a symbolic link is also provided under the name of "Documents and Settings" which invisibly redirects any attempted access of the latter to "C:\Users." The fact that the profile-root folder appears twice in any folder-listing has adverse implications for any backup program. Backup software needs to be aware of this Vista idiosyncrasy, or else it is possible to double the size of the backup by copying what appear to be two separate folders, both potentially containing the bulk of the data on the computer.

A second change in Vista is that the media-specific "My Pictures" and "My Music" folders are now outside of the "My Documents" folder, instead of being subfolders (as well as removing the "My" Prefix - so My Documents becomes Documents, for example).

References

- Vista Profile Guide [1]
- Microsoft on User Profiles [2]
- Changing the Windows User Profiles Directory via the registry [3]
- User Profile Structure [4]

References

[1] http://smallbusiness.itworld.com/4374/nls_windows_vistaprint060531/page_1.html
[2] http://www.microsoft.com/resources/documentation/windows/xp/all/proddocs/en-us/userprofile_overview.mspx?mfr=true
[3] http://www.neotitans.com/resources/changing-user-profiles-directory.html
[4] http://technet.microsoft.com/en-us/library/cc775560.aspx

Roaming user profile

	C:\Documents and Settings\{username}		
		Application Data	
		Cookies	
		Desktop	
		Favorites	
		Local Settings	
			Application Data
			History
			Temp
			Temporary Internet Files
		My Documents	
			My Music
			My Pictures
			My Videos
		Recent	
		NetHood	
		PrintHood	
		SendTo	
		Start Menu	
		Templates	
		NTUSER.DAT	
		ntuser.dat.LOG	
		ntuser.ini	

Folder layout of typical Windows 2000/XP user profile. Normally everything except the items within "Local Settings" are stored on the file server as part of a roaming profile.

A **roaming user profile** is a concept in the Microsoft Windows NT family of operating systems (and now adopted by Linux operating systems such as Ubuntu) that allows users with a computer joined to a Windows Server domain to log on to any computer on the same network and access their documents and have a consistent desktop experience, such as applications remembering toolbar positions and preferences, or the desktop appearance staying the same.

Method of operation

All Windows operating systems since Windows 2000 are designed from the start to be able to support roaming profiles. Normally, a standalone computer stores the user's documents, desktop items, application preferences, and desktop appearance on the local computer in two divided sections, consisting of the portion that could roam plus an additional temporary portion containing items such as the web browser cache. The Windows registry is similarly divided to support roaming; there are System and Local Machine hives that stay on the local computer, plus a separate User hive (HKEY_CURRENT_USER) designed to be able to roam with the user profile.

When a roaming user is created, the user's profile information is instead stored on a centralized file server accessible from any network-joined desktop computer. The login prompt on the local computer checks to see if the user exists in the domain rather than on the local computer; no pre-existing account is required on the local computer. If the domain login is successful, the roaming profile is copied from the central file server to the desktop computer, and a local account is created for the user.

When the user logs off from the desktop computer, the user's roaming profile is copied from the local computer back to the central file server, not including the temporary local profile items. Because this is a copy and not a move/delete, the user's profile information remains on the local computer in addition to being copied to the network.

When the user logs in on a second desktop computer, this process repeats, copying the roaming profile from the server to the second desktop computer, and then copying back from the desktop to the server when the user logs off.

When the user returns to the first desktop computer and logs in, the roaming profile is copied over the previous profile information, replacing it. If profile caching is enabled, the server is capable of only copying the newest files to the local computer, reusing the existing local files that have not changed since the last login, and thereby speeding up the login process.

Profile copying limitations

Roaming reduces network performance

Due to the profile copying at login and logout, a roaming profile set up using the default configuration can be extremely slow and waste considerable amounts of time for users with large amounts of data in their account.

When Microsoft designed Internet Explorer, the programmers made an explicit decision to store cookies and favorites as tiny individual files less than a kilobyte each, rather than storing this data as a single large consolidated file, such as the *bookmark.html* file used by Mozilla Firefox. Microsoft also stores shortcut files in the *Recent* profile folder, linking to recently opened files and folders.

File servers tend to only transfer large files several megabytes in size at the fastest possible network speed. Hundreds of very small files only a kilobyte per file can reduce network performance by 90%. As a profile ages and accumulates hundreds to thousands of cookies, favorites, and Recent items, the login and logout times become progressively slower, even though these files occupy only a few megabytes of profile data.

Local caching of the user profile on a desktop computer hard drive can reduce and improve login and logout times, but at the penalty of cluttering up the hard drive with profile data from every cached user who logs in. Local caching is more suitable where people tend to use the same computer every day. Local profile caching is not useful where hundreds to thousands of students need to be able to use any computer across a school or university campus—the cumulative cached data from so many different profiles can consume all available lab computer disk space.

Roaming profiles and WAN links

Users with a roaming profile can encounter crippling logon delays when logging in over a WAN. If connected to the domain from a remote site, after authentication, Windows will attempt to pull the user's profile from the location specified in Active Directory. If the location happens to be across a WAN link it can potentially slow the WAN down to a crawl and cause the logon to fail (after a very lengthy delay).

Users with a roaming profile working from a remote site should login to the machine *before* connecting to the network, (so that the machine uses its cached local copy) and connect to the network after logon has completed. Another option is to remove the roaming profile path from Active Directory prior to their departure. This must be done in enough time that the change is replicated to the relevant Domain Controller at the remote site.

Not compatible with gigabytes of user files

DV video editing in a roaming profile generally results in unacceptably slow login times because the video file segments are also copied back and forth from server to desktop. A one-hour 15 gigabyte DV file takes 20 minutes to copy over a 100 megabit LAN connection. If this were present in a roaming profile it would take at least 20 minutes for the user to login and 20 minutes to logout.

Even for small DV editing projects consisting of several short source clips, each clip still uses 250 megabytes per minute, and typically the source clips are retained when creating the final DV movie project. A small project consisting of four 5 minute clips to generate one 10 minute movie, totals 7.5 gigabytes of DV data, and requires at least 10 minutes to transfer over a 100 megabit network connection before the user's desktop appears and they can begin to do any work.

In a school environment where such editing projects are not mission-critical and do not absolutely need to be backed up on an expensive tape archive system every night, the applications requiring such excessively large amounts of user data are instead usually run on a stand-alone local account that does not roam, to bypass these network storage and retrieval problems.

Third-party companies don't use it correctly

Many third-party companies do not understand the difference between the roaming part of the profile and the non-roaming temporary section, and so will store their temporary files in the roaming portion, adding megabytes of unnecessary data to the profile and increasing the roaming login and logout times. Some go as far as intentionally storing data in the wrong location, as a method of encouraging the purchase of more expensive "enterprise" versions of their software.[1]

Mass-user logins/logouts cause congestion

In a school environment, roaming can result in severe network congestion and slowness when an entire lab of students log off of computers at the same time, and then within minutes are attempting to login in somewhere else. Account data inconsistency problems can result if the students begin to login in the second location before the profile uploading and logout from the first location has finished.

Misbehaving third-party programs don't exit

Some programs installed on desktop computers do not properly release control of the User registry during logoff, and can result in corrupted profiles because the User registry copying never successfully completes. To deal with this, Microsoft created a utility known as the *User Profile Hive Cleanup Service* which will forcefully remap the file handles for these misbehaving programs so that the profile copying can finish successfully and the account logoff is successful.[2] However, the hung program may remain on the local computer still holding the local cached copy of the User registry in a busy state, until the computer is rebooted.

Roaming accounts don't backup until logoff

The most recent version of a file in a roaming profile without redirection is stored only on the local computer, and stays there until the user logs off, whereupon it transfers to the server. If nightly server backups are done, and a roaming user does not log off for days at a time, their roaming account documents are not being included in the nightly backup.

Further, if a roaming user uses standby or hibernation to turn off the computer at night, their profile is still not copied to the network. In this manner it is possible for a roaming account's documents to not be backed up for days to weeks at a time, and there is the potential for considerable data loss if the local hard drive suffers a catastrophic failure during these long periods of not logging off the roaming account from the local computer.

Multiple logins can overwrite each other

Due to the underlying file copying mechanism from server to desktop, roaming assumes the user account is logged on to only a single computer at a time. Documents in a roaming profile copied down to the local machine have no network awareness of each other, and it is not possible to use file locking to alert the user that the file is already open.

Logging onto multiple computers with one account, and opening the same document multiple times on each computer can result in inconsistencies and loss of saved changes if the file is modified on two different computers at the same time:

- When the first computer with the modified document logs off, the changes are written to the network copy of the profile.
- When the second computer logs off, the different document version overwrites the previously saved changes during profile logout.

Folder redirection to improve performance

To deal with these profile copying problems, it is possible to override the default operation of roaming, and set up user accounts so that certain parts of the profile are accessed by the local computer directly on a central file server rather than copying to the local computer first.[3] This requires that the central server and network be extremely reliable and always available, because if the server is down, users can not access their files from a local cached copy.

To the end-user, folder redirection generally does not appear to function any differently from using a normal standalone computer. Redirecting the user's My Documents and Desktop to be accessed directly on a file server are the first two big steps for speeding up roaming profiles. However, as 3rd party software have begun to store more and more data in the *Application Data* portion of the roaming profile, it has also become useful to redirect that to also be accessed directly on the server.

The question may be raised as to why the entire roaming profile can not be accessed directly on the server, and no copying needs to be done at all. The reasoning for this appears to be that certain Microsoft programs running all the time on the client computer can not tolerate the sudden loss of their data folders if the server goes down or the network is disconnected. Some portions must still be copied back and forth before the desktop appears so that these folders are available if the network-redirected folders go down.

Limitations of redirection

Roaming folder redirection is not compatible with laptop users who take their computer home with them. Normally the entire local cached copy of the profile would be available when the laptop is used away from the building network, but with redirection the most important portions are not stored locally and so are unavailable when used away from the network. Microsoft Direct Access system allows users to connect to the building network over a WAN link transparently.

Redirecting *Application Data* and logging on multiple times running the same program on each computer can result in corruption of the application settings, due to multiple programs trying to access the same settings files on the server. Some programs such as OpenOffice.org will detect if the preferences are in use somewhere else and alert the user.

Redirection limitations of UNC paths

Some programs do not work properly with redirected profile folders that refer to a UNC file path on a server share: *\\server\share\username\Application Data*

- Windows Command prompt cannot have a UNC working directory, so batch files usually fail.
- It is not possible to install Microsoft Office VSTO add ins on a UNC path. (ApData can be a natural place for users to install addins without administration privilidges.)
- Adobe Reader has been incompatible with Application Data located on a UNC file path since at least version 9.0, which would crash with a runtime error[4]. Adobe Reader X (10.0) is partially compatible but will not run in document protection mode on a UNC path.
- Open Office 3.3 is similarly incompatible with Application Data on a UNC path, and the software crashes on startup.[5] A fix has been developed and will be available in an upcoming release.

These problems with UNC paths can usually be fixed by having the folders redirected to a drive mapping for the UNC share:

- Drive N: (say) is mapped to \\server\share\userhomedir
- AppDir folder redirection to user home directory:N:\Application Data

However, use of drive mappings is generally deprecated by Microsoft, and UNC-only redirection paths are the preferred implementation.

Mandatory profiles

Folder redirection with mandatory profiles

Folder redirection may be used with mandatory profiles, and is useful in situations where it is desirable to "lock down" the general desktop appearance but still allow users to save documents to the network. For example, this can be used as a generic account for anyone to use without a password for temporary use.

Redirecting My *Documents* and the *Desktop* in a mandatory profile will allow documents to be saved, but at logoff, any changes to the desktop appearance such as the desktop picture, Internet Explorer cookies, Favorites, and the Recent documents opened list are reverted to the original state.[6]

Setup methods

Active Directory

A roaming user profile must first be set up on the domain controller to which client computers are joined. In Windows 2000 and later versions, this is set using the *Active Directory Users and Computers* snap-in. Windows NT 4.0 and earlier used the *User Manager for Domains* program. A user profile location is set on the server and can be customized, as required. When a user logs onto a domain, the roaming user profile is downloaded from the server onto the local computer and applied. When the user logs off, the changes made to the roaming profile are transferred back to the domain controller.

Although a roaming user profile may be stored in any shared folder of a computer available inside a local Microsoft Windows network, using the domain controller is recommended because the profile data should be available at any workstation the user tries to log on to. Should the server not be available, the user will still be able to log on using a cached copy of the profile on his workstation, unless the profile is super-mandatory.

Enabling roaming profiles for a workstation running Windows NT 4.0, Windows 2000, Windows XP Professional, Windows Vista Business or Ultimate is done by specifying a location on the server where the users' profiles are located; this is done under *User Manager for Domains* in Windows NT 4.0 Server and *Active Directory Users and Computers* in Windows 2000 and later. Workstations running Windows 95, 98 or Me can also have roaming profiles, roaming profiles become available in Windows 9x when a home directory on the network is specified for the user and multiple desktop settings have been enabled under the *Passwords* box in the Windows Control Panel.

Roaming profiles on Windows 95, 98 and Me are all compatible with each other so if a network has mixture of Windows 95 and Windows 98 workstations the same user profile may be used for each workstation. This is also the case with Roaming profiles between Windows NT 4.0, Windows 2000, Windows XP but there may be some compatibility issues due to differences in each version of Windows. Roaming profiles in Windows Vista and Windows 7 are compatible with each other but these versions are not compatible with earlier versions of Windows. A separate profile folder with the extension .V2 will be created when using Roaming profiles with Windows Vista or 7. The easiest solution is to have all workstations running the same version of Windows. (see Compatibility section)

Novell eDirectory

For roaming to work with Novell servers, the Novell product "ZENworks Desktop Management" needs to be installed on the server, and its associated workstation management package installed on each of the client computers. Within the directory, a *User Package* object is created, which enables roaming, specifies where the roaming profile is stored, and also stores any associated group policies for each version of Windows where users will login. The User Package also enables *Dynamic Local User*, which functions similar to Active Directory, allowing an account created in eDirectory to login on any desktop computer even if no local account exists in advance, and assigns local account privileges such as User, Power User, or Administrator to the newly created local user account.

The User Package can be associated with a specific user account in the directory, or is associated with an organizational unit that then applies to all user accounts within that OU. The User Package also enables additional ZENworks Desktop Management functions, such as remote view and remote control of the desktop computer, network printers that follow the user from one desktop to the next, and the scheduling of events that are to be run wherever the user is logged in.

Windows 3.x

While Windows 3.x does not contain user profiles it was possible for users to have their own personalised desktop in a business environment. Windows 3.x had an administrative setup option which network administrators could use by typing setup.exe /a Windows could then be installed to a network share. Windows setup was then run from each local machine to install a few local files making Windows 3.1 capable of being run over a network. The local files could be saved to a user's home directory on a Novell or Windows NT Domain network allowing the user to have his or her settings roam between machines, the local machine in this scenario did not require a hard drive and could have been booted from a floppy or network card.

Advantages of roaming user profiles

- Enforcement of administrative control by using mandatory user profiles which helps to protect the user's environment from being damaged by the user himself/herself.
- Users can access their data anywhere in the network with more reliability
- Easier backup as most data is in one location on the server

Disadvantages of roaming user profiles

Each time a user logs into a workstation, all of the files and settings are transferred over the network; the result is that the login process takes longer than if the user were to use a local profile. This is particularly the case if the profile is large in size. The login time may be reduced if the profile is cached as some files can be loaded from the local workstation and by using folder redirection to redirect folders that can grow to a large size, like My Documents, to a network share.

However, this limitation has been addressed in Windows Server 2008 Active Directory by allowing folder redirection of virtually all folders that were previously stored in a user's profile (including My Music, Favorites, and others) to a centralized and secured network share. This means that a user's roaming profile can easily be reduced to size smaller than 20MB, thus eliminating the long login times that were experienced with previous versions of AD. When using folder redirection and automatic caching of offline files, all of a user's files and preferences are available offline and synced in a much more efficient manner than previously possible when the computer is reconnected to the network using Remote Differential Compression (RDC).

Another problem is related to different set of applications installed on machines, applications stores information into Local Settings and some into the registry, but only the registry is transferred across. It can corrupt application functionality under roaming profile.

2000/XP and Vista/Win7 compatibility

While Windows XP and Windows 2000 profiles are basically similar, Windows Vista and its successor Windows 7 use an entirely different profile structure. Thus, a user who switches-desk between the two classes of OS cannot have personal data transferred automatically, as would normally happen with roaming profiles. Instead, two distinct server-side profiles are created for this user.

This is an important consideration for any site intending to introduce Vista or Windows 7 computers into an existing Windows 2000/XP roaming-profile network. If possible it should be planned that users will not have to migrate regularly between the two classes of OS.

Windows Vista and 7 will get their profile stored on the server with .V2 added (example: \\server\profiles\username.V2)

Redirected folder sharing

Redirected network folders are able to override the separation between 2000/XP and Vista/Win7. For example, both types of profiles can be redirected to use a single Documents folder, and a single Desktop folder, so that the user's account documents are consistent between the two profiles, even if all other account settings will be different.

Redirected sharing of folders such as Application Data may lead to data corruption, since Microsoft did not intend this for their application data to be shared between the different OS versions.

Alternatives

Mylogon - Allows any-user logon to a client computer whilst maintaining the same local settings. May be preferable to Roaming Profiles for small-site networks, data-entry stations and the like.

[[AppSense [7]] User Virtualization] - Solution Description from official AppSense product overview document: "User Virtualization is recognized as the most effective and scalable approach to managing the user component of the desktop.

User Virtualization is an infrastructure technology solution that virtualizes, centralizes, manages and applies the user environment on to a desktop as required. User Virtualization spans all desktops across multiple OS platforms, desktop and application delivery mechanisms, devices and locations.

By separating the user from the desktop and managing it as a separate component, organizations can adopt multi-platform/delivery mechanism desktop environments. AppSense simplifies desktop management overhead, reduces operational costs, improves end user experience and ensure that user settings and corporate policy is applied to set up, configure and personalize a desktop, no matter how that desktop is delivered or where it is hosted."

Liquidware Labs' User Virtualization Management Solution, ProfileUnity [8] - Solution description from official document - "ProfileUnity is a lightweight yet powerful user virtualization management solution. ProfileUnity runs without any software to install on end-point PCs or virtual machines. At logon a managed profile is synched with the local user session in Windows own native format. No databases or additional servers are relied upon, thus keeping cost and complexity low. The solution is scalable to hundreds of thousands of users in a single organization due to its lightweight approach of compressing the user profile, registry, and settings while in storage yet fully uncompressing it in native format on the end-point. Because ProfileUnity uses Windows native format, organizations can use the solution for migration only or leave it installed for heterogeneous Windows desktops to exist."

RES Workspace Manager [9] - Excerpt from the official product description: "With RES Software, your IT team can offer users a work environment that is location and time independent. You can even make the desktop user independent—or better yet, make the user desktop-independent. By separating users from their physical desktop, you can manage changes more easily and give them an optimal experience anytime, anywhere.

RES Workspace Manager lets you pick the level of management and control you want for your organization, today. You can always upgrade as your organization grows or needs new features."

References

[1] http://forums.counterpath.com/viewtopic.php?f=6&t=6744&start=0
[2] User Profile Hive Cleanup Service (http://www.microsoft.com/downloads/details.aspx?FamilyId=1B286E6D-8912-4E18-B570-42470E2F3582)
[3] Microsoft TechNet, Windows Server 2008, Group Policy Management, User Folder Redirection, Folder Redirection Overview (http://technet.microsoft.com/en-us/library/cc732275(WS.10).aspx)
[4] Adobe knowledge-base: Runtime error | Roaming Profile workflows | Acrobat, Reader 9 http://kb2.adobe.com/cps/404/kb404597.html
[5] bug report, Open Office 3.3 incompatible with redirected Application Data http://openoffice.org/bugzilla/show_bug.cgi?id=115778
[6] Microsoft MSDN, 2012-06-06, "Mandatory User Profiles", Quote: "[...] a user can modify his or her desktop, but the changes are not saved when the user logs off.", http://msdn.microsoft.com/en-us/library/windows/desktop/bb776895(v=vs.85).aspx
[7] http://www.appsense.com/
[8] http://www.LiquidwareLabs.com/
[9] http://www.ressoftware.com/

- Windows Help

External links

- Microsoft MSDN Library: Reference on User Profiles (http://msdn2.microsoft.com/en-us/library/aa372123.aspx)
- Microsoft Technet: Windows Server 2003 Product Help: User Profiles best practices (http://technet.microsoft.com/en-us/library/cc738803(v=ws.10))
- Microsoft Technet: Windows Server 2003: Operations Whitepaper: Best Practices for User Profiles (http://technet.microsoft.com/de-de/library/cc784484(v=ws.10).aspx)

Directory service

A **directory service** is the software system that stores, organizes and provides access to information in a directory. In software engineering, a directory is a map between names and values. It allows the lookup of values given a name, similar to a dictionary. As a word in a dictionary may have multiple definitions, in a directory, a name may be associated with multiple, different pieces of information. Likewise, as a word may have different parts of speech and different definitions, a name in a directory may have many different types of data.

Directories may be very narrow in scope, supporting only a small set of node types and data types, or they may be very broad, supporting an arbitrary or extensible set of types. In a telephone directory, the nodes are names and the data items are telephone numbers. In the DNS the nodes are domain names and the data items are IP addresses (and alias, mail server names, etc.). In a directory used by a network operating system, the nodes represent resources that are managed by the OS, including users, computers, printers and other shared resources. Many different directory services have been used since the advent of the Internet but this article focuses mainly on those that have descended from the X.500 directory service.

Introduction

A simple directory service called a naming service, maps the names of network resources to their respective network addresses. With the name service type of directory, a user doesn't have to remember the physical address of a network resource; providing a name will locate the resource. Each resource on the network is considered an object on the directory server. Information about a particular resource is stored as attributes of that object. Information within objects can be made secure so that only users with the available permissions are able to access it. More sophisticated directories are designed with namespaces as Subscribers, Services, Devices, Entitlements, Preferences, Content and so on. This design process is highly related to Identity management.

A directory service defines the namespace for the network. A namespace in this context is the term that is used to hold one or more objects as named entries. The directory design process normally has a set of rules that determine how network resources are named and identified. The rules specify that the names be unique and unambiguous. In X.500 (the directory service standards) and LDAP the name is called the **Distinguished name** (DN) and is used to refer to a collection of attributes (relative distinguished names) which make up the name of a directory entry.

A directory service is a shared information infrastructure for locating, managing, administering, and organizing common items and network resources, which can include volumes, folders, files, printers, users, groups, devices, telephone numbers and other objects. A directory service is an important component of a NOS (Network Operating System). In the more complex cases a directory service is the central information repository for a Service Delivery Platform. For example, looking up "computers" using a directory service might yield a list of available computers and information for accessing them.

Replication and Distribution have very distinct meanings in the design and management of a directory service. The term replication is used to indicate that the same directory namespace (the same objects) are copied to another directory server for redundancy and throughput reasons. The replicated namespace is governed by the same authority. The term distribution is used to indicate that multiple directory servers, that hold different namespaces, are interconnected to form a distributed directory service. Each distinct namespace can be governed by different authorities.

Comparison with relational databases

There are a number of things that distinguish a traditional directory service from a typical relational database. Of course there are exceptions, but in general:

- directory information is read more often than it is written; this makes features related to transactions and rollback less important.
- data can be redundant if it helps performance.

Directory schemas are defined as object classes, attributes, name bindings and knowledge (namespaces), where an object class has:

- Must - attributes that each of its instances must have
- May - attributes that can be defined for an instance, but can be omitted with the absence treated somewhat like NULL in a relational database
- Attributes are sometimes multi-valued allowing multiple naming attributes at one level such as machine type and serial number concatenated or multiple phone numbers for "work phone".
- Attributes and object classes are standardized throughout the industry and formally registered with the IANA for their object ID. Therefore directory applications seek to reuse much of the standard classes and attributes to maximize the benefit of existing directory server software.
- Object instances are slotted into namespaces. That is, each object class inherits from its parent object class (and ultimately from the root of the hierarchy) adding attributes to the must/may list.
- Directory services are often a central component in the security design of an IT system and have a correspondingly fine granularity regarding access control: who may operate in which manner on what information. Also see: ACLs

Implementations of directory services

Directory services were part of an Open Systems Interconnection (OSI) initiative to get everyone in the industry to agree to common network standards to provide multi-vendor interoperability. In the 1980s, the ITU and ISO came up with a set of standards - X.500, for directory services, initially to support the requirements of inter-carrier electronic messaging and network name lookup. The Lightweight Directory Access Protocol, LDAP, is based on the directory information services of X.500, but uses the TCP/IP stack and a string encoding scheme of the X.500 protocol DAP, giving it more relevance on the Internet.

There have been numerous forms of directory service implementations from different vendors. Systems developed before the advent of X.500 include:

- **Domain Name System:** (DNS), the first directory service on the Internet, which is still used everywhere today.
- **Hesiod:** was based on DNS and used at MIT's Project Athena.
- **Network Information Service:** (NIS), originally named Yellow Pages (YP), was Sun Microsystems' implementation of a directory service for Unix network environments. It served a similar role as Hesiod.
- **NetInfo:** was developed by NeXT in the late 1980s for NEXTSTEP. After being acquired by Apple, it was released as open source and used as the directory service for Mac OS X before being deprecated in favor of the LDAP-based Open Directory. Support for NetInfo was completely removed with the release of 10.5 Leopard.
- **Banyan VINES:** was the first scalable directory services offering.
- **NT Domains:** was developed by Microsoft to provide directory services for Windows machines prior to the release the LDAP-based Active Directory in Windows 2000. Windows Vista continues to support NT Domains, but only after relaxing the minimum authentication protocols it supports.

Among the LDAP/X.500 based implementations are:

- **Active Directory:** Microsoft's modern directory service for Windows, originating from the X.500 directory it created for use in Exchange Server, first shipped with Windows 2000 Server and is supported by successive versions of Windows.

- **eDirectory:** This is Novell's implementation of directory services. It supports multiple architectures including Windows, NetWare, Linux and several flavours of Unix and has long been used for user administration, configuration management, and software management. eDirectory has evolved into a central component in a broader range of Identity management products. It was previously known as Novell Directory Services.
- **eNitiatives ViewDS Directory Server:** ViewDS[1][2][3] was originally developed by Telstra Research Laboratories in Clayton Victoria Australia (previously Telecom Australia) as an X.500 Directory server known as View500 to run on line White & Yellow Pages services. ViewDS was acquired by eNitiatives in 2000. It differs from other X.500 Directory products in that it has an in built indexing engine capable of indexing all attributes and also supports a range of different types of searching and matching on entries, such as word matching, stem matching, synonym matching, acronym matching, component matching, misspelling matching, and sounds like matching. This matching is available on multiple languages including Pinyin and Traditional Mandarin. ViewDS is LDAPv3 compliant and is also the world's first Directory to support the XACML[4] standard for Policy Based Access control onto all attributes stored in the directory, with an inbuilt combined Policy Decision Point (PDP) and Policy Information Point (PIP) as well as two Policy Administration Tools (PAP). ViewDS is also CCEB ACP133EdD (Military)[5] and IATA ATN-AMHS standards compliant, and supports the storage of XML objects, data and schema in the Directory using the draft IETF XML Enabled Directory standard. It also supports SPMLv2.0, DSMLv2 and SCIM 1.0. ViewDS is widely used in the Government,[6][7][8]Aviation, Health & Defence sectors.
- **Red Hat Directory Server:** Red Hat released a directory service, that it acquired from AOL's **Netscape Security Solutions** unit,[9] as a commercial product running on top of Red Hat Enterprise Linux called Red Hat Directory Server and as the community supported 389 Directory Server project.
- **Open Directory:** Apple's Mac OS X Server uses a directory service named Open Directory, which implements LDAP using a customized build of OpenLDAP and integrates support for both SASL and Kerberos authentication. It uses a plugins architecture to work with other LDAPv3 directories, including proprietary solutions like Active Directory and eDirectory.
- **Apache Directory Server:** Apache Software Foundation offers a directory service called ApacheDS.
- **Oracle Internet Directory:** (OID) is Oracle Corporation's directory service, which is compatible with LDAP version 3.
- **CA Directory:** CA Directory contains pre-caching engine which can index all attributes that are used in LDAP search filters, and caching all attributes returned in search results.
- **Alcatel-Lucent Directory Server:** CTIA 2009 - 4G Service Creation & Development Award Winner offering enhanced performance, high availability and proven efficiencies[10]
- **Sun Java System Directory Server:** Sun Microsystems' current directory service offering[11]
- **OpenDS:** An open source directory service implementation from scratch in Java, backed by Sun Microsystems[12]
- **IBM Tivoli Directory Server** It is a customized build of an old release of OpenLDAP.
- **DirX Directory** Server from Atos (ex-Siemens software)
- **Windows NT Directory Services (NTDS)**, later renamed Active Directory, replaces the former NT Domain system.
- **Critical Path Directory Server**
- **OpenLDAP** Derived from the original University of Michigan reference LDAP implementation (as are the Netscape/Red Hat/Fedora/Sun JSDS servers) but significantly evolved. It supports all current computer architectures, including Unix and Unix derivatives, Linux, Windows, z/OS, and a variety of embedded/realtime systems.
- **Isode Limited:** High performance and high availability LDAP and X.500 servers.

- UnboundID Directory Server [13]: **A commercial high-performance Directory Server product produced by the UnboundID Corporation.**
- Lotus Domino

There are also plenty of open-source tools to create directory services, including OpenLDAP and the Kerberos protocol, and Samba software which can act as a Windows Domain Controller with Kerberos and LDAP backends. Administration is done using GOsa or Samba provided SWAT.

Using name services

Unix OSs

Name services on Unix systems are typically configured through nsswitch.conf. Information from name services can be retrieved using getent.

Notes

[1] (http://www.viewds.com)
[2] http://ats.business.gov.au/companies-and-technologies/information-technology/enterprise-software-and-systems/eb2bcom
[3] http://blogs.kuppingercole.com/kuppinger/category/directory-services/
[4] http://www.networkworld.com/newsletters/dir/2011/082911id2.html
[5] (http://www.businesswire.com/news/home/20110111005269/en/Telos-Corporation-Partners-eNitiatives-Feature-ViewDS-AMHS)
[6] http://www.directory.gov.au.com
[7] http://www.vic.gov.au.com/contactsandservices/directory
[8] http://www.directory.act.gov.au
[9] (http://www.informationweek.com/story/showArticle.jhtml?articleID=48800390)
[10] "Alcatel-Lucent 8661 Directory Server" (http://www.alcatel-lucent.com/wps/portal/products/detail?LMSG_CABINET=Solution_Product_Catalog&LMSG_CONTENT_FILE=Products/Product_Detail_000558.xml). Alcatel-lucent.com. . Retrieved 2012-01-09.
[11] "Oracle and Sun" (http://www.sun.com/software/products/directory_srvr_ee/). Sun.com. 2010-09-07. . Retrieved 2012-01-09.
[12] "Java.net" (http://opends.dev.java.net/). Opends.dev.java.net. . Retrieved 2012-01-09.
[13] http://www.unboundid.com

References

- Carter, Gerald (2003). *LDAP System Administration*. O'Reilly Media. ISBN 978-1-56592-491-8.

Lightweight Directory Access Protocol

The **Lightweight Directory Access Protocol** (**LDAP**; /ˈɛldæp/) is an application protocol for accessing and maintaining distributed directory information services over an Internet Protocol (IP) network.[1] LDAP is defined in terms of ASN.1 and transmitted using BER.

Directory services may provide any organized set of records, often with a hierarchical structure, such as a corporate electronic mail directory. Similarly, a telephone directory is a list of subscribers with an address and a phone number.

LDAP is specified in a series of Internet Engineering Task Force (IETF) Standard Track Requests for comments (RFCs). The latest version is Version 3, published as RFC 4510.

Origin and influences

Telecommunication companies' understanding of directory requirements was well developed after some 70 years of producing and managing telephone directories. These companies introduced the concept of directory services to information technology and computer networking, their input culminating in the comprehensive X.500 specification,[2] a suite of protocols produced by the International Telecommunication Union (ITU) in the 1980s.

X.500 directory services were traditionally accessed via the X.500 Directory Access Protocol (DAP), which required the Open Systems Interconnection (OSI) protocol stack. LDAP was originally intended to be a lightweight alternative protocol for accessing X.500 directory services through the simpler (and now widespread) TCP/IP protocol stack. This model of directory access was borrowed from the DIXIE and Directory Assistance Service protocols.

Standalone LDAP directory servers soon followed, as did directory servers supporting both DAP and LDAP. The latter has become popular in enterprises, as LDAP removed any need to deploy an OSI network. Today, X.500 directory protocols including DAP can also be used directly over TCP/IP.

The protocol was originally created by Tim Howes of the University of Michigan, Steve Kille of Isode Limited, and Wengyik Yeong of Performance Systems International, circa 1993. Mark Wahl of Critical Angle Inc., Tim Howes, and Steve Kille started work in 1996 on a new version of LDAP, LDAPv3, under the aegis of the Internet Engineering Task Force (IETF). LDAPv3, first published in 1997, superseded LDAPv2 and added support for extensibility, integrated the Simple Authentication and Security Layer, and better aligned the protocol to the 1993 edition of X.500. Further development of the LDAPv3 specifications themselves and of numerous extensions adding features to LDAPv3 has come through the IETF.

In the early engineering stages of LDAP, it was known as *Lightweight Directory Browsing Protocol*, or *LDBP*. It was renamed with the expansion of the scope of the protocol beyond directory browsing and searching, to include directory update functions. It was given its *Lightweight* name because it was not as network intensive as its DAP predecessor and thus was more easily implemented over the internet due to its relatively modest bandwidth usage.

LDAP has influenced subsequent Internet protocols, including later versions of X.500, XML Enabled Directory (XED), Directory Service Markup Language (DSML), Service Provisioning Markup Language (SPML), and the Service Location Protocol (SLP).

Protocol overview

A client starts an LDAP session by connecting to an LDAP server, called a Directory System Agent (DSA), by default on TCP port 389. The client then sends an operation request to the server, and the server sends responses in return. With some exceptions, the client does not need to wait for a response before sending the next request, and the server may send the responses in any order.

The client may request the following operations:

- StartTLS — use the LDAPv3 Transport Layer Security (TLS) extension for a secure connection
- Bind — authenticate and specify LDAP protocol version
- Search — search for and/or retrieve directory entries
- Compare — test if a named entry contains a given attribute value
- Add a new entry
- Delete an entry
- Modify an entry
- Modify Distinguished Name (DN) — move or rename an entry
- Abandon — abort a previous request
- Extended Operation — generic operation used to define other operations
- Unbind — close the connection (not the inverse of Bind)

In addition the server may send "Unsolicited Notifications" that are not responses to any request, e.g. before it times out a connection.

A common alternative method of securing LDAP communication is using an SSL tunnel. This is denoted in LDAP URLs by using the URL scheme "ldaps". The default port for LDAP over SSL is 636. The use of LDAP over SSL was common in LDAP Version 2 (LDAPv2) but it was never standardized in any formal specification. This usage has been deprecated along with LDAPv2, which was officially retired in 2003.[3]

Directory structure

The protocol accesses LDAP directories, which follow the 1993 edition of the X.500 model:

- An entry consists of a set of attributes.
- An attribute has a name (an *attribute type* or *attribute description*) and one or more values. The attributes are defined in a *schema* (see below).
- Each entry has a unique identifier: its *Distinguished Name* (DN). This consists of its *Relative Distinguished Name* (RDN), constructed from some attribute(s) in the entry, followed by the parent entry's DN. Think of the DN as the full file path and the RDN as its relative filename in its parent folder (e.g. if */foo/bar/myfile.txt* were the DN, then *myfile.txt* would be the RDN).

Be aware that a DN may change over the lifetime of the entry, for instance, when entries are moved within a tree. To reliably and unambiguously identify entries, a UUID might be provided in the set of the entry's *operational attributes*.

An entry can look like this when represented in LDAP Data Interchange Format (LDIF) (LDAP itself is a binary protocol):

```
dn: cn=John Doe,dc=example,dc=com
cn: John Doe
givenName: John
sn: Doe
telephoneNumber: +1 888 555 6789
telephoneNumber: +1 888 555 1232
```

```
mail: john@example.com
manager: cn=Barbara Doe,dc=example,dc=com
objectClass: inetOrgPerson
objectClass: organizationalPerson
objectClass: person
objectClass: top
```

"dn" is the distinguished name of the entry; it's neither an attribute nor a part of the entry. "cn=John Doe" is the entry's RDN (Relative Distinguished Name), and "dc=example,dc=com" is the DN of the parent entry, where "dc" denotes 'Domain Component'. The other lines show the attributes in the entry. Attribute names are typically mnemonic strings, like "cn" for common name, "dc" for domain component, "mail" for e-mail address and "sn" for surname.

A server holds a subtree starting from a specific entry, e.g. "dc=example,dc=com" and its children. Servers may also hold references to other servers, so an attempt to access "ou=department,dc=example,dc=com" could return a *referral* or *continuation reference* to a server that holds that part of the directory tree. The client can then contact the other server. Some servers also support *chaining*, which means the server contacts the other server and returns the results to the client.

LDAP rarely defines any ordering: The server may return the values of an attribute, the attributes in an entry, and the entries found by a search operation in any order. This follows from the formal definitions - an entry is defined as a set of attributes, and an attribute is a set of values, and sets need not be ordered.

Operations

Expand discussion of referral responses to various operations, especially modify, for example where all modifies must be directed from replicas to a master directory.

Add

The ADD operation is used to insert a new entry into the directory server database.[4] If the distinguished name in the add request already exists in the directory, then the server will not add a duplicate entry but will set the result code in the add result to decimal 68, "entryAlreadyExists".[5]

- LDAP-compliant servers will never dereference the distinguished name transmitted in the add request when attempting to locate the entry, that is, distinguished names are never de-aliased.
- LDAP-compliant servers will ensure that the distinguished name and all attributes conform to naming standards
- The entry to be added must not exist, and the immediate superior must exist.

```
dn: uid=user,ou=people,dc=example,dc=com
changetype: add
objectClass: top
objectClass: person
uid: user
sn: last-name
cn: common-name
userPassword: password
```

In the above example, `uid=user,ou=people,dc=example,dc=com` must not exist, and `ou=people,dc=example,dc=com` must exist.

Bind (authenticate)

When an LDAP session is created, that is, when an LDAP client connects to the server, the **authentication state** of the session is set to anonymous. The BIND operation establishes the authentication state for a session.

Simple BIND and SASL PLAIN can send the user's DN and password in plaintext, so the connections utilizing either Simple or SASL PLAIN should be encrypted using Transport Layer Security (TLS). The server typically checks the password against the `userPassword` attribute in the named entry. Anonymous BIND (with empty DN and password) resets the connection to anonymous state.

SASL (Simple Authentication and Security Layer) BIND provides authentication services through a wide range of mechanisms, e.g. Kerberos or the client certificate sent with TLS.[6]

BIND also sets the LDAP protocol version. The version is an integer and at present must be either 2 (two) or 3 (three), although the standard supports integers between 1 and 127 (inclusive) in the protocol. If the client requests a version that the server does not support, the server must set the result code in the BIND response to the code for a protocol error. Normally clients should use LDAPv3, which is the default in the protocol but not always in LDAP libraries.

BIND had to be the first operation in a session in LDAPv2, but is not required in LDAPv3 (the current LDAP version). In LDAPv3, each successful BIND request changes the authentication state of the session and each unsuccessful BIND request resets the authentication state of the session.

Delete

To delete an entry, an LDAP client transmits a properly formed delete request to the server.[7]

- A delete request must contain the distinguished name of the entry to be deleted
- Request controls may also be attached to the delete request
- Servers do not dereference aliases when processing a delete request
- Only leaf nodes (entries with no subordinates) may be deleted by a delete request. Some servers support an operational attribute `hasSubordinates` whose value indicates whether an entry has any subordinate entries, and some servers support an operational attribute `numSubordinates`[8] indicating the number of entries subordinate to the entry containing the `numSubordinates` attribute.

Delete requests are subject to access controls, that is, whether a connection with a given authentication state will be permitted to delete a given entry is governed by server-specific access control mechanisms.

Search and Compare

The Search operation is used to both search for and read entries. Its parameters are:

baseObject

: The name of the base object entry (or possibly the root) relative to which the search is to be performed.

scope

: What elements below the baseObject to search. This can be `BaseObject` (search just the named entry, typically used to read one entry), `singleLevel` (entries immediately below the base DN), or `wholeSubtree` (the entire subtree starting at the base DN).

filter

: Criteria to use in selecting elements within scope. For example, the filter `(&(objectClass=person)(|(givenName=John)(mail=john*)))` will select "persons" (elements of objectClass `person`) where the matching rules for `givenName` and `mail` determine whether the values for those attributes match the filter assertion. Note that a common misconception is that LDAP data is case-insensitive, whereas in fact matching rules and ordering rules determine matching, comparisons, and relative value relationships. If the example filters were required to match the case of the attribute value, an *extensible match filter* must be used, for example, `(&(objectClass=person)(|(givenName:caseExactMatch:=John)(mail:caseExactSubstringsMatc`

derefAliases

Whether and how to follow alias entries (entries that refer to other entries),

attributes

Which attributes to return in result entries.

sizeLimit, timeLimit

Maximum number of entries to return, and maximum time to allow search to run. These values, however, cannot override any restrictions the server places on size limit and time limit.

typesOnly

Return attribute types only, not attribute values.

The server returns the matching entries and potentially continuation references. These may be returned in any order. The final result will include the result code.

The Compare operation takes a DN, an attribute name and an attribute value, and checks if the named entry contains that attribute with that value.

Modify

The MODIFY operation is used by LDAP clients to request that the LDAP make changes to an existing entries.[9] Attempts to modify entries that do not exist will fail. MODIFY requests are subject to access controls as implemented by the server.

The MODIFY operation requires that the distinguished name (DN) of the entry be specified, and a sequence of changes. Each change in the sequence must be one of:

- add (add a new value, which must not already exist in the entry)
- delete (delete an existing value)
- replace (replace an existing value with a new value)

LDIF example of adding a value to an attribute:

```
dn: dc=example,dc=com
changetype: modify
add: cn
cn: the-new-cn-value-to-be-added
```

To replace the value of an existing attribute, Use the `replace` keyword. If the attribute is multi-valued, the client must specify the value of the attribute to delete.

To delete an attribute from an entry, use the keyword `delete` and the changetype designator `modify`. If the attribute is multi-valued, the client must specify the value of the attribute to delete.

There is also a modify-increment extension which allows an incrementable attribute value to be incremented by a specified amount. The modify-increment extension uses object identifier `1.3.6.1.1.14`. The following example using LDIF increments `employeeNumber` by `5`:

```
dn: uid=user.0,ou=people,dc=example,dc=com
changetype: modify
increment: employeeNumber
employeeNumber: 5
```

When LDAP servers are in a replicated topology, LDAP clients should consider using the post-read control to verify updates instead of a search after an update.[10] The post-read control is designed so that applications need not issue a search request after an update – it is bad form to retrieve an entry for the sole purpose of checking that an update worked because of the replication eventual consistency model. An LDAP client should not assume that it connects to the same directory server for each request because architects may have placed load-balancers or LDAP proxies or both between LDAP clients and servers.

Modify DN

Modify DN (move/rename entry) takes the new RDN (Relative Distinguished Name), optionally the new parent's DN, and a flag that says whether to delete the value(s) in the entry that match the old RDN. The server may support renaming of entire directory subtrees.

An update operation is atomic: Other operations will see either the new entry or the old one. On the other hand, LDAP does not define transactions of multiple operations: If you read an entry and then modify it, another client may have updated the entry in the meantime. Servers may implement extensions[11] that support this, though.

Extended operations

The Extended Operation is a generic LDAP operation that can define new operations that were not part of the original protocol specification. StartTLS is one of the most significant extensions. Other examples include the Cancel and Password Modify.

StartTLS

The StartTLS operation establishes Transport Layer Security (the descendant of SSL) on the connection. It can provide data confidentiality (to protect data from being observed by third parties) and/or data integrity protection (which protects the data from tampering). During TLS negotiation the server sends its X.509 certificate to prove its identity. The client may also send a certificate to prove its identity. After doing so, the client may then use SASL/EXTERNAL. By using the SASL/EXTERNAL, the client requests the server derive its identity from credentials provided at a lower level (such as TLS). Though technically the server may use any identity information established at any lower level, typically the server will use the identity information established by TLS.

Servers also often support the non-standard "LDAPS" ("Secure LDAP", commonly known as "LDAP over SSL") protocol on a separate port, by default 636. LDAPS differs from LDAP in two ways: 1) upon connect, the client and server establish TLS before any LDAP messages are transferred (without a StartTLS operation) and 2) the LDAPS connection must be closed upon TLS closure.

It should be noted that some "LDAPS" client libraries only encrypt communication, they do not check the host name against the name in the supplied certificate[12]

LDAPS was used with LDAPv2, because the StartTLS operation had not yet been defined. The use of LDAPS is deprecated, and modern software should only use StartTLS.

Abandon

The Abandon operation requests that the server abort an operation named by a message ID. The server need not honor the request. Unfortunately, neither Abandon nor a successfully abandoned operation send a response. A similar Cancel extended operation does send responses, but not all implementations support this.

Unbind

The Unbind operation abandons any outstanding operations and closes the connection. It has no response. The name is of historical origin, and is *not* the opposite of the Bind operation.[13]

Clients can abort a session by simply closing the connection, but they should use Unbind.[14] Unbind allows the server to gracefully close the connection and free resources that it would otherwise keep for some time until discovering the client had abandoned the connection. It also instructs the server to cancel operations that can be canceled, and to not send responses for operations that cannot be canceled.[15]

LDAP URLs

An LDAP URL format exists, which clients support in varying degrees, and servers return in referrals and continuation references (see RFC 4516):

```
ldap://host:port/DN?attributes?scope?filter?extensions
```

Most of the components described below are optional.

- *host* is the FQDN or IP address of the LDAP server to search.
- *port* is the network port (default port 389) of the LDAP server.
- *DN* is the distinguished name to use as the search base.
- *attributes* is a comma-separated list of attributes to retrieve.
- *scope* specifies the search scope and can be "base" (the default), "one" or "sub".
- *filter* is a search filter. For example `(objectClass=*)` as defined in RFC 4515.
- *extensions* are extensions to the LDAP URL format.

For example, "`ldap://ldap.example.com/cn=John%20Doe,dc=example,dc=com`" refers to all user attributes in John Doe's entry in `ldap.example.com`, while "`ldap:///dc=example,dc=com??sub?(givenName=John)`" searches for the entry in the default server (note the triple slash, omitting the host, and the double question mark, omitting the attributes). As in other URLs, special characters must be percent-encoded.

There is a similar non-standard `ldaps:` URL scheme for LDAP over SSL. This should not be confused with LDAP with TLS, which is achieved using the StartTLS operation using the standard `ldap:` scheme.

Schema

The contents of the entries in a subtree are governed by a schema known as a directory information tree (DIT).

The schema of a Directory Server defines a set of rules that govern the kinds of information that the server can hold. It has a number of elements, including:

- Attribute Syntaxes—Provide information about the kind of information that can be stored in an attribute.
- Matching Rules—Provide information about how to make comparisons against attribute values.
- Matching Rule Uses—Indicate which attribute types may be used in conjunction with a particular matching rule.
- Attribute Types—Define an object identifier (OID) and a set of names that may be used to refer to a given attribute, and associates that attribute with a syntax and set of matching rules.

- Object Classes—Define named collections of attributes and classify them into sets of required and optional attributes.
- Name Forms—Define rules for the set of attributes that should be included in the RDN for an entry.
- Content Rules—Define additional constraints about the object classes and attributes that may be used in conjunction with an entry.
- Structure Rule—Define rules that govern the kinds of subordinate entries that a given entry may have.

Attributes are the elements responsible for storing information in a directory, and the schema defines the rules for which attributes may be used in an entry, the kinds of values that those attributes may have, and how clients may interact with those values.

Clients may learn about the schema elements that the server supports by retrieving an appropriate subschema subentry.

The schema defines *object classes*. Each entry must have an objectClass attribute, containing named classes defined in the schema. The schema definition of the classes of an entry defines what kind of object the entry may represent - e.g. a person, organization or domain. The object class definitions also define the list of attributes that must contain values and the list of attributes which may contain values.

For example, an entry representing a person might belong to the classes "top" and "person". Membership in the "person" class would require the entry to contain the "sn" and "cn" attributes, and allow the entry also to contain "userPassword", "telephoneNumber", and other attributes. Since entries may have multiple ObjectClasses values, each entry has a complex of optional and mandatory attribute sets formed from the union of the object classes it represents. ObjectClasses can be inherited, and a single entry can have multiple ObjectClasses values that define the available and required attributes of the entry itself. A parallel to the schema of an objectClass is a class definition and an instance in Object-oriented programming, representing LDAP objectClass and LDAP entry, respectively.

Directory servers may publish the directory schema controlling an entry at a base DN given by the entry's subschemaSubentry operational attribute. (An *operational attribute* describes operation of the directory rather than user information and is only returned from a search when it is explicitly requested.)

Server administrators can add additional schema entries in addition to the provided schema elements. A schema for representing individual people within organizations is termed a white pages schema.

Variations

A lot of the server operation is left to the implementor or administrator to decide. Accordingly, servers may be set up to support a wide variety of scenarios.

For example, data storage in the server is not specified - the server may use flat files, databases, or just be a gateway to some other server. Access control is not standardized, though there has been work on it and there are commonly used models. Users' passwords may be stored in their entries or elsewhere. The server may refuse to perform operations when it wishes, and impose various limits.

Most parts of LDAP are extensible. Examples: One can define new operations. *Controls* may modify requests and responses, e.g. to request sorted search results. New search scopes and Bind methods can be defined. Attributes can have *options* that may modify their semantics.

Other data models

As LDAP has gained momentum, vendors have provided it as an access protocol to other services. The implementation then recasts the data to mimic the LDAP/X.500 model, but how closely this model is followed varies. For example, there is software to access SQL databases through LDAP, even though LDAP does not readily lend itself to this.[16] X.500 servers may support LDAP as well.

Similarly, data previously held in other types of data stores are sometimes moved to LDAP directories. For example, Unix user and group information can be stored in LDAP and accessed via PAM and NSS modules. LDAP is often used by other services for authentication.

An example of such data model is the GLUE Schema,[17] which is used in a distributed information system based on LDAP that enable users, applications and services to discover which services exist in a Grid infrastructure and further information about their structure and state.

Usage

An LDAP server may return referrals to other servers for requests that it cannot fulfill itself. This requires a naming structure for LDAP entries so one can find a server holding a given DN or distinguished name, a concept defined in the X.500 Directory and also used in LDAP. Another way of locating LDAP servers for an organization is a DNS server resource record (SRV).

An organization with the domain example.org may use the top level LDAP DN `dc=example,dc=org` (where *dc* means domain component). If the LDAP server is also named ldap.example.org, the organization's top level LDAP URL becomes `ldap://ldap.example.org/dc=example,dc=org`.

Primarily two common styles of naming are used in both X.500 [2008] and LDAPv3. These are documented in the ITU specifications and IETF RFCs. The original form takes the top level object as the country object, such as c=US, c=FR. The domain component model uses the model described above. An example of country based naming could be c=FR, o=Some Organization, ou=Some Organizational Unit L=Locality, or in the US: c=US, st=CA, o=Some Organization ou=Organizational Unit, L=Locality, and CN=Common Name.

References

[1] LDAP: Framework, Practices, and Trends (http://www2.computer.org/portal/web/csdl/doi/10.1109/MIC.2004.44)
[2] The X.500 series - ITU-T Rec. X.500 to X.521
[3] Tools.ietf.org (http://tools.ietf.org/html/draft-zeilenga-ldapv2-04)
[4] Add section of RFC4511 (http://tools.ietf.org/html/rfc4511#section-4.7)
[5] LDAP result codes (http://tools.ietf.org/html/rfc4511#appendix-A)
[6] SASL Mechanisms at IANA (http://www.iana.org/assignments/sasl-mechanisms/sasl-mechanisms.xml)
[7] RFC4511: delete request (http://tools.ietf.org/html/rfc4511#section-4.8)
[8] Boreham Draft (numSubordinates) (http://tools.ietf.org/html/draft-boreham-numsubordinates-01)
[9] Modify Section of RFC4511 (http://tools.ietf.org/html/rfc4511#section-4.6)
[10] read-entry controls (http://tools.ietf.org/html/rfc4527)
[11] INTERNET-DRAFT LDAP Transactions draft-zeilenga-ldap-txn-15.txt (http://www.rfc-editor.org/internet-drafts/draft-zeilenga-ldap-txn-15.txt)
[12] Shibboleth Security alert 20120227 (http://shibboleth.internet2.edu/secadv/secadv_20120227.txt)
[13] Tools.ietf.org (http://tools.ietf.org/html/rfc4511#section-4.3)
[14] Tools.ietf.org (http://tools.ietf.org/html/rfc4511#section-5.3)
[15] Tools.ietf.org (http://tools.ietf.org/html/rfc4511#section-3.1)
[16] Openldap.org (http://www.openldap.org/doc/admin24/backends.html#SQL)
[17] SourceForge : Project Home (http://forge.gridforum.org/sf/projects/glue-wg)

- ITU-T Rec. X.680, "Abstract Syntax Notation One (ASN.1) - Specification of Basic Notation", 1994
- Basic encoding rules (BER) - ITU-T Rec. X.690, "Specification of ASN.1 encoding rules: Basic, Canonical, and Distinguished Encoding Rules", 1994

- RFC 3641 - Generic String Encoding Rules (GSER) for ASN.1 Types
- RFC 4346 - The TLS Protocol Version 1.1
- RFC 4422 - Simple Authentication and Security Layer (SASL)
- SASL mechanisms (http://www.iana.org/assignments/sasl-mechanisms) registered at IANA
- This article is based on material taken from the Free On-line Dictionary of Computing prior to 1 November 2008 and incorporated under the "relicensing" terms of the GFDL, version 1.3 or later.

Further reading

- Arkills, B (2003). *LDAP Directories Explained: An Introduction and Analysis* (http://www.informit.com/store/product.aspx?isbn=0-201-78792-X). Addison-Wesley Professional. ISBN 0-201-78792-X.
- Carter, G (2003). *LDAP System Administration* (http://oreilly.com/catalog/9781565924918). O'Reilly Media. ISBN 1-56592-491-6.
- Donley, C (2002). *LDAP Programming, Management, and Integration*. Manning Publications. ISBN 1-930110-40-5.
- Howes, T; Smith, M; Good, G (2003). *Understanding and Deploying LDAP Directory Services* (http://www.informit.com/store/product.aspx?isbn=0-672-32316-8). Addison-Wesley Professional. ISBN 0-672-32316-8.
- Rhoton, J (1999). *Programmer's Guide to Internet Mail: SMTP, POP, IMAP, and LDAP*. Elsevier. ISBN 1-55558-212-5.
- Voglmaier, R (2003). *The ABCs of LDAP: How to Install, Run, and Administer LDAP Services*. Auerbach Publications. ISBN 0-8493-1346-5.

External links

- Devshed.com (http://www.devshed.com/c/a/Administration/Understanding-LDAP-part-1/), Understanding LDAP, A simple, light introductory tutorial for LDAP.
- Skills-1st.co.uk (http://www.skills-1st.co.uk/papers/ldap-schema-design-feb-2005/index.html), LDAP schema design
- Capitalhead.com (http://capitalhead.com/articles/troubleshooting-ldap-ssl-connection-issues-between-microsoft-ilmmiis--novell-edirectory-873.aspx), Troubleshooting LDAP SSL connection issues between Microsoft ILM/MIIS & Novell eDirectory 8.7.3
- Prasannatech.com (http://web.archive.org/web/20080713040421/http://www.prasannatech.com/ldapdesign.html), LDAP schema design - A Case Study

RFCs

LDAP is specified in a series of Request for Comments documents:

- RFC 4510 - LDAP: Technical Specification Road Map (Obsoletes: RFC 2251, RFC 2252, RFC 2253, RFC 2254, RFC 2255, RFC 2256, RFC 2829, RFC 2830, RFC 3377, RFC 3771)
- RFC 4511 - LDAP: The Protocol (Obsoletes RFC 2251, RFC 2830 & RFC 3771)
- RFC 4512 - LDAP: Directory Information Models (Obsoletes RFC 2251, RFC 2252, RFC 2256 & RFC 3674)
- RFC 4513 - LDAP: Authentication Methods and Security Mechanisms (Obsoletes RFC 2251, RFC 2829 & RFC 2830)
- RFC 4514 - LDAP: String Representation of Distinguished Names (Obsoletes RFC 2253)
- RFC 4515 - LDAP: String Representation of Search Filters (Obsoletes RFC 2254)
- RFC 4516 - LDAP: Uniform Resource Locator (Obsoletes RFC 2255)
- RFC 4517 - LDAP: Syntaxes and Matching Rules (Obsoletes RFC 2252 & RFC 2256, Updates RFC 3698)
- RFC 4518 - LDAP: Internationalized String Preparation

- RFC 4519 - LDAP: Schema for User Applications (Obsoletes RFC 2256, Updates RFC 2247, RFC 2798 & RFC 2377)

The following RFCs detail LDAP-specific Best Current Practices:

- RFC 4520 (also BCP 64) - Internet Assigned Numbers Authority (IANA) Considerations for the Lightweight Directory Access Protocol (LDAP) (replaced RFC 3383)
- RFC 4521 (also BCP 118) - Considerations for Lightweight Directory Access Protocol (LDAP) Extensions

The following is a partial list of RFCs specifying LDAPv3 extensions:

- RFC 2247 - Use of DNS domains in distinguished names (Updated by RFC 4519 & RFC 4524)
- RFC 2307 - Using LDAP as a Network Information Service
- RFC 2589 - LDAPv3: Dynamic Directory Services Extensions
- RFC 2649 - LDAPv3 Operational Signatures
- RFC 2696 - LDAP Simple Paged Result Control
- RFC 2798 - inetOrgPerson LDAP Object Class (Updated by RFC 3698, RFC 4519 & RFC 4524)
- RFC 2830 - LDAPv3: Extension for Transport Layer Security
- RFC 2849 - The LDAP Data Interchange Format (LDIF)
- RFC 2891 - Server Side Sorting of Search Results
- RFC 3045 - Storing Vendor Information in the LDAP root DSE
- RFC 3062 - LDAP Password Modify Extended Operation
- RFC 3296 - Named Subordinate References in LDAP Directories
- RFC 3671 - Collective Attributes in LDAP
- RFC 3672 - Subentries in LDAP
- RFC 3673 - LDAPv3: All Operational Attributes
- RFC 3687 - LDAP Component Matching Rules
- RFC 3698 - LDAP: Additional Matching Rules
- RFC 3829 - LDAP Authorization Identity Controls
- RFC 3866 - Language Tags and Ranges in LDAP
- RFC 3909 - LDAP Cancel Operation
- RFC 3928 - LDAP Client Update Protocol
- RFC 4370 - LDAP Proxied Authorization Control
- RFC 4373 - LBURP
- RFC 4403 - LDAP Schema for UDDI
- RFC 4522 - LDAP: Binary Encoding Option
- RFC 4523 - LDAP: X.509 Certificate Schema
- RFC 4524 - LDAP: COSINE Schema (replaces RFC 1274)
- RFC 4525 - LDAP: Modify-Increment Extension
- RFC 4526 - LDAP: Absolute True and False Filters
- RFC 4527 - LDAP: Read Entry Controls
- RFC 4528 - LDAP: Assertion Control
- RFC 4529 - LDAP: Requesting Attributes by Object Class
- RFC 4530 - LDAP: entryUUID
- RFC 4531 - LDAP Turn Operation
- RFC 4532 - LDAP Who am I? Operation
- RFC 4533 - LDAP Content Sync Operation
- RFC 4876 - Configuration Profile Schema for LDAP-Based Agents
- RFC 5020 - LDAP entryDN Operational Attribute

LDAPv2 was specified in the following RFCs:

- RFC 1777 - Lightweight Directory Access Protocol (replaced RFC 1487)
- RFC 1778 - The String Representation of Standard Attribute Syntaxes (replaced RFC 1488)
- RFC 1779 - A String Representation of Distinguished Names (replaced RFC 1485)

LDAPv2 was moved to historic status by the following RFC:

- RFC 3494 - Lightweight Directory Access Protocol version 2 (LDAPv2) to Historic Status

LDAP Data Interchange Format

LDIF

Filename extension	`.ldif`
Type of format	Data interchange
Standard(s)	RFC 2849

The **LDAP Data Interchange Format** (**LDIF**) is a standard plain text data interchange format for representing LDAP (Lightweight Directory Access Protocol) directory content and update requests. LDIF conveys directory content as a set of records, one record for each object (or entry). It represents update requests, such as Add, Modify, Delete, and Rename, as a set of records, one record for each update request.

LDIF was designed in the early 1990s by Tim Howes, Mark C Smith, and Gordon Good while at the University of Michigan. LDIF was updated and extended in the late 1990s for use with Version 3 of LDAP. This later version of LDIF is called version 1 and is formally specified in RFC 2849, an IETF Standard Track RFC. RFC 2849, authored by Gordon Good, was published in June 2000 and is currently a Proposed Standard.

A number of extensions to LDIF have been proposed over the years. One extension has been formally specified by the IETF and published. RFC 4525, authored by Kurt Zeilenga, extended LDIF to support the LDAP Modify-Increment extension. It is expected that additional extensions will be published by the IETF in the future.

Content Record Format

Each content record is represented as a group of attributes, with records separated from one another by blank lines. The individual attributes of a record are represented as single logical lines (represented as one or more multiple physical lines via a line-folding mechanism), comprising "name: value" pairs. Value data that do not fit within a portable subset of ASCII characters are marked with '::' after the attribute name and encoded into ASCII using base64 encoding. The content record format is a subset of the Internet Directory Information type.RFC 2425 [1]

Tools that employ LDIF

The OpenLDAP utilities include tools for exporting data from LDAP servers to LDIF content records (`ldapsearch`), importing data from LDIF content records to LDAP servers (`ldapadd`), and applying LDIF change records to LDAP servers (`ldapmodify`).

LDIF is one of the formats for importing and exporting address book data that the address books in Netscape Communicator and in the Mozilla Application Suite support. Yahoo! Mail does not encode certain characters properly when one exports their Yahoo! address book in LDIF format. For example, ampersand (&) is encoded as an HTML Extended Character (&) instead of the ampersand character. As a result, when the LDIF file is imported into Thunderbird, for example, a text phrase like "John & Jane Doe" comes out in one's address book as "John & Jane Doe". The only corrective means at the moment is manually editing the address book after an Import.

Microsoft Windows 2000 Server and Windows Server 2003 include an LDIF based command line tool named LDIFDE for importing and exporting information in Active Directory.

JXplorer is a cross platform open source java application that can browse and do basic editing of LDIF files.

LDIF fields

```
dn: distinguished name
```

This refers to the name that uniquely identifies an entry in the directory.

```
   dc: domain component
```

This refers to each component of the domain. For example www.google.com would be written as DC=www,DC=google,DC=com

```
   ou: organizational unit
```

This refers to the organizational unit (or sometimes the user group) that the user is part of. If the user is part of more than one group, you may specify as such, e.g., OU= Lawyer,OU= Judge.

```
   cn: common name
```

This refers to the individual object (person's name; meeting room; recipe name; job title; etc.) for whom/which you are querying.

Examples of LDIF

This is an example of a simple directory entry with several attributes, represented as a record in LDIF:

```
dn: cn=The Postmaster,dc=example,dc=com
objectClass: organizationalRole
cn: The Postmaster
```

This is an example of an LDIF record that modifies multiple single-valued attributes for two different directory entries (this format is used by Microsoft's LDIFDE tool):

```
dn: CN=John Smith,OU=Legal,DC=example,DC=com
changetype: modify
replace:employeeID
employeeID: 1234
-
replace:employeeNumber
employeeNumber: 98722
-
replace: extensionAttribute6
extensionAttribute6: JSmith98
-

dn: CN=Jane Smith,OU=Accounting,DC=example,DC=com
changetype: modify
replace:employeeID
employeeID: 5678
-
replace:employeeNumber
employeeNumber: 76543
-
replace: extensionAttribute6
```

```
extensionAttribute6: JSmith14
-
```

Note: the "-" character between each attribute change is required. Also note that each directory entry ends with a "-" followed by a blank line. The final "-" is required.

This is an example of an LDIF file that adds a telephone number to an existing user:

```
dn: cn=Peter Michaels, ou=Artists, l=San Francisco, c=US
changetype: modify
add: telephonenumber
telephonenumber: +1 415 555 0002
```

An example of LDIF containing a control:

```
version: 1
dn: o=testing,dc=example,dc=com
control: 1.3.6.1.1.13.1 false cn
changetype: add
objectClass: top
objectClass: organization
o: testing
```

RFCs

- RFC 2849 — The LDAP Data Interchange Format (LDIF) - Technical Specification
- RFC 4510 — Lightweight Directory Access Protocol (LDAP): Technical Specification Road Map
- RFC 4525 — LDAP Modify-Increment Extension

External links

- MSDN : Windows 2003 : Using the LDIFDE Tool [2]
- MSDN : Active Directory : LDIF Scripts [3]

References

[1] http://www.ietf.org/rfc/rfc2425.txt
[2] http://msdn2.microsoft.com/en-us/library/ms870068.aspx
[3] http://msdn2.microsoft.com/en-us/library/ms677268.aspx

Kerberos (protocol)

Stable release	krb5-1.10.1 / March 8, 2012
Website	web.mit.edu/kerberos/ [1]

Kerberos (/ˈkɛərbərəs/) is a computer network authentication protocol which works on the basis of "tickets" to allow nodes communicating over a non-secure network to prove their identity to one another in a secure manner. Its designers aimed primarily at a client–server model, and it provides mutual authentication—both the user and the server verify each other's identity. Kerberos protocol messages are protected against eavesdropping and replay attacks. Kerberos builds on symmetric key cryptography and requires a trusted third party, and optionally may use public-key cryptography by utilizing asymmetric key cryptography during certain phases of authentication.[2] Kerberos uses port 88 by default.

History and development

MIT developed Kerberos to protect network services provided by Project Athena. The protocol was named after the character *Kerberos* (or *Cerberus*) from Greek mythology which was a monstrous three-headed guard dog of Hades. Several versions of the protocol exist; versions 1–3 occurred only internally at MIT.

Steve Miller and Clifford Neuman, the primary designers of Kerberos version 4, published that version in the late 1980s, although they had targeted it primarily for Project Athena.

Version 5, designed by John Kohl and Clifford Neuman, appeared as RFC 1510 in 1993 (made obsolete by RFC 4120 in 2005), with the intention of overcoming the limitations and security problems of version 4.

MIT makes an implementation of Kerberos freely available, under copyright permissions similar to those used for BSD. In 2007, MIT formed the Kerberos Consortium to foster continued development. Founding sponsors include vendors such as Oracle, Apple Inc., Google, Microsoft, Centrify Corporation and TeamF1 Inc., and academic institutions such as KTH-Royal Institute of Technology, Stanford University, MIT and vendors such as CyberSafe offering commercially supported versions.

Authorities in the United States classified Kerberos as auxiliary military technology and banned its export because it used the DES encryption algorithm (with 56-bit keys). A non-US Kerberos 4 implementation, KTH-KRB developed at the Royal Institute of Technology in Sweden, made the system available outside the US before the US changed its cryptography export regulations (*circa* 2000). The Swedish implementation was based on a limited version called eBones. eBones was based on the exported MIT Bones release (stripped of both the encryption functions and the calls to them) based on version Kerberos 4 patch-level 9.

Windows 2000 and later use Kerberos as their default authentication method. Some Microsoft additions to the Kerberos suite of protocols are documented in RFC 3244 "Microsoft Windows 2000 Kerberos Change Password and Set Password Protocols". RFC 4757 documents Microsoft's use of the RC4 cipher. While Microsoft uses the Kerberos protocol, it does not use the MIT software.

Many UNIX and UNIX-like operating systems, including FreeBSD, Apple's Mac OS X, Red Hat Enterprise Linux, Oracle's Solaris, IBM's AIX, HP's OpenVMS, Univention's Univention Corporate Server and others, include software for Kerberos authentication of users or services. Embedded implementation of the Kerberos V authentication protocol for client agents and network services running on embedded platforms is also available from companies such as TeamF1, Inc.

As of 2005, the IETF Kerberos working group is updating the specifications. Recent updates include:

- Encryption and Checksum Specifications" (RFC 3961).
- Advanced Encryption Standard (AES) Encryption for Kerberos 5 (RFC 3962).

- A new edition of the Kerberos V5 specification "The Kerberos Network Authentication Service (V5)" (RFC 4120). This version obsoletes RFC 1510, clarifies aspects of the protocol and intended use in a more detailed and clearer explanation.
- A new edition of the GSS-API specification "The Kerberos Version 5 Generic Security Service Application Program Interface (GSS-API) Mechanism: Version 2." (RFC 4121).

Protocol

Description

The client authenticates itself to the Authentication Server (AS) which forwards the username to a Key Distribution Center (KDC). The KDC issues a Ticket Granting Ticket (TGT), which is time stamped, encrypts it using the user's password and returns the encrypted result to the user's workstation. If successful, this gives the user desktop access.

When the client needs to communicate with another node ("principal" in Kerberos parlance) it sends the TGT to the Ticket Granting Service (TGS), which shares the same host as the KDC. After verifying the TGT is valid and the user is permitted to access the requested service, the TGS issues a Ticket and session keys, which are returned to the client.

The client then sends the Ticket and keys to the service server (SS).

Here is another description.

The client authenticates to the AS *once* using a long-term *shared secret* (e.g. a password) and receives a Ticket Granting Ticket (TGT) from the AS. Later, when the client wants to contact some SS, it can (re)use this ticket to get additional tickets from TGS, for SS, without resorting to using the shared secret. The latter tickets can be used to prove authentication to the SS.

The phases are detailed below.

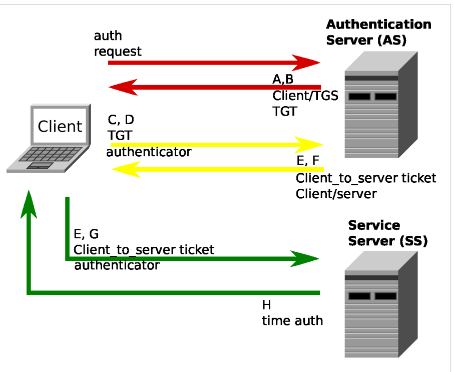

Kerberos negotiations

User Client-based Logon

1. A user enters a username and password on the client machines.
2. The client performs a one-way function (hash usually) on the entered password, and this becomes the secret key of the client/user.

Client Authentication

1. The client sends a cleartext message of the user ID to the AS requesting services on behalf of the user. (Note: Neither the secret key nor the password is sent to the AS.) The AS generates the secret key by hashing the password of the user found at the database (e.g. Active Directory in Windows Server).
2. The AS checks to see if the client is in its database. If it is, the AS sends back the following two messages to the client:
    - Message A: *Client/TGS Session Key* encrypted using the secret key of the client/user.
    - Message B: *Ticket-Granting-Ticket* (which includes the client ID, client network address, ticket validity period, and the *client/TGS session key*) encrypted using the secret key of the TGS.
3. Once the client receives messages A and B, it attempts to decrypt message A with the secret key generated from the password entered by the user. If the user entered password does not match the password in the AS database, the client's secret key will be different and thus unable to decrypt message A. With a valid password and secret key the client decrypts message A to obtain the *Client/TGS Session Key*. This session key is used for further communications with the TGS. (Note: The client cannot decrypt Message B, as it is encrypted using TGS's secret

key.) At this point, the client has enough information to authenticate itself to the TGS.

Client Service Authorization

1. When requesting services, the client sends the following two messages to the TGS:
   - Message C: Composed of the TGT from message B and the ID of the requested service.
   - Message D: Authenticator (which is composed of the client ID and the timestamp), encrypted using the *Client/TGS Session Key*.
2. Upon receiving messages C and D, the TGS retrieves message B out of message C. It decrypts message B using the TGS secret key. This gives it the "client/TGS session key". Using this key, the TGS decrypts message D (Authenticator) and sends the following two messages to the client:
   - Message E: *Client-to-server ticket* (which includes the client ID, client network address, validity period and *Client/Server Session Key*) encrypted using the service's secret key.
   - Message F: *Client/Server Session Key* encrypted with the *Client/TGS Session Key*.

Client Service Request

1. Upon receiving messages E and F from TGS, the client has enough information to authenticate itself to the SS. The client connects to the SS and sends the following two messages:
   - Message E from the previous step (the *client-to-server ticket*, encrypted using service's secret key).
   - Message G: a new Authenticator, which includes the client ID, timestamp and is encrypted using *Client/Server Session Key*.
2. The SS decrypts the ticket using its own secret key to retrieve the *Client/Server Session Key*. Using the sessions key, SS decrypts the Authenticator and sends the following message to the client to confirm its true identity and willingness to serve the client:
   - Message H: the timestamp found in client's Authenticator plus 1, encrypted using the *Client/Server Session Key*.
3. The client decrypts the confirmation using the *Client/Server Session Key* and checks whether the timestamp is correctly updated. If so, then the client can trust the server and can start issuing service requests to the server.
4. The server provides the requested services to the client.

Drawbacks and Limitations

- Single point of failure: It requires continuous availability of a central server. When the Kerberos server is down, no one can log in. This can be mitigated by using multiple Kerberos servers and fallback authentication mechanisms.
- Kerberos has strict time requirements, which means the clocks of the involved hosts must be synchronized within configured limits. The tickets have a time availability period and if the host clock is not synchronized with the Kerberos server clock, the authentication will fail. The default configuration per MIT [3] requires that clock times are no more than five minutes apart. In practice Network Time Protocol daemons are usually used to keep the host clocks synchronized.
- The administration protocol is not standardized and differs between server implementations. Password changes are described in RFC 3244 [4].
- Since all authentication is controlled by a centralized KDC, compromise of this authentication infrastructure will allow an attacker to impersonate any user.

Related Requests For Comments

- RFC 2712 — Addition of Kerberos Cipher Suites to Transport Layer Security (TLS)
- RFC 4120 — The Kerberos Network Authentication Service (V5)
- RFC 4537 — Kerberos Cryptosystem Negotiation Extension
- RFC 4556 — Public Key Cryptography for Initial Authentication in Kerberos (PKINIT)
- RFC 4752 — The Kerberos V5 (GSSAPI) Simple Authentication and Security Layer (SASL) Mechanism
- RFC 6111 — Additional Kerberos Naming Constraints
- RFC 6112 — Anonymity Support for Kerberos
- RFC 6113 — A Generalized Framework for Kerberos Pre-Authentication
- RFC 6251 — Using Kerberos Version 5 over the Transport Layer Security (TLS) Protocol

References

[1] http://web.mit.edu/kerberos/
[2] RFC 4556, abstract
[3] http://web.mit.edu/Kerberos/krb5-1.5/krb5-1.5.4/doc/krb5-admin/Clock-Skew.html
[4] http://www.ietf.org/rfc/rfc3244.txt

Notes

- SDK Team. "Microsoft Kerberos (Windows)" (http://msdn.microsoft.com/en-us/library/aa378747(VS.85).aspx). *MSDN Library*.
- B. Clifford Neuman and Theodore Ts'o (September 1994). "Kerberos: An Authentication Service for Computer Networks" (http://gost.isi.edu/publications/kerberos-neuman-tso.html). *IEEE Communications* **32** (9): 33–8. doi:10.1109/35.312841.
- John T. Kohl, B. Clifford Neuman, and Theodore Y. T'so (1994). "The Evolution of the Kerberos Authentication System" (ftp://athena-dist.mit.edu/pub/kerberos/doc/krb_evol.PS). In Johansen, D.; Brazier, F. M. T. (Postscript). *Distributed open systems*. Washington: IEEE Computer Society Press. pp. 78–94. ISBN 0-8186-4292-0.
- Cisco Systems Kerberos Overview- An Authentication Service for Open Network Systems (http://www.cisco.com/en/US/tech/tk59/technologies_white_paper09186a00800941b2.shtml)

External links

- How Kerberos Authentication Works (http://learn-networking.com/network-security/how-kerberos-authentication-works)
- Kerberos page (http://web.mit.edu/kerberos/) at MIT
- Kerberos Working Group at IETF (http://www.ietf.org/html.charters/krb-wg-charter.html)
- Kerberos Consortium (http://www.kerberos.org/) at MIT
- White Papers (http://www.kerberos.org/software/whitepapers.html) at MIT
- Vendor Documentation and Specifications (http://www.kerberos.org/docs/links.html) at MIT
- Kerberos How-to (http://www.kerberos.org/software/adminkerberos.pdf)
- The Kerberos FAQ (http://www.faqs.org/faqs/kerberos-faq/general/) (last modified 8/18/2000)
- Shishi, a free Kerberos implementation for the GNU system (http://josefsson.org/shishi/)
- Designing an Authentication System: A Dialogue in Four Scenes. Humorous play concerning how the design of Kerberos evolved. (http://web.mit.edu/kerberos/www/dialogue.html)
- Description of Kerberos 5 in the SPORE library (http://www.lsv.ens-cachan.fr/spore/kerberos.html)
- Kerberos Authentication in Windows Server 2003 (http://www.microsoft.com/windowsserver2003/technologies/security/kerberos/default.mspx)
- Kerberos Tutorial (http://www.kerberos.org/software/tutorial.html)

- Novell Inc's Comment to the Proposed Settlement between Microsoft and the Department of Justice - Microsoft purposefully breaks Kerberos interoperability (http://www.usdoj.gov/atr/cases/ms_tuncom/major/mtc-00029523.htm)
- Kerberos in FreeBSD (http://www.freebsd.org/doc/en/books/handbook/kerberos5.html)
- Embedded Kerberos Implementation (http://teamf1.com/home/product/authagent-kerberos/) by TeamF1
- Heimdal, an implementation of Kerberos 5 (http://www.h5l.org/)

User account policy

User Account Policy is a document which outlines the requirements for requesting and maintaining an account on computer systems or networks, typically within an organization. It is very important for large sites where users typically have accounts on many systems. Some sites have users read and sign an Account Policy as part of the account request process.

Policy Content

- Should state who has the authority to approve account requests.
- Should state who is allowed to use the resources (e.g., employees or students only)
- Should state any citizenship/resident requirements.
- Should state if users are allowed to share accounts or if users are allowed to have multiple accounts on a single host.
- Should state the users' rights and responsibilities.
- Should state when the account should be disabled and archived.
- Should state how long the account can remain inactive before it is disabled.
- Should state password construction and aging rules.

Example

Some example wording: "Employees shall only request/receive accounts on systems they have a true business need to access. Employees may only have one official account per system and the account ID and login name must follow the established standards. Employees must read and sign the acceptable use policy prior to requesting an account."

External links

- National Institute for Standards and Technology [1]

References

[1] http://csrc.nist.gov/

Monitoring

Nagios

Nagios

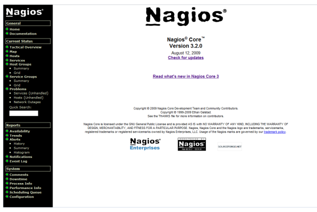

Screenshot of the Nagios web interface

Original author(s)	Ethan Galstad and others [1]
Initial release	March 14, 1999[2]
Stable release	3.4.1 / May 15, 2012[3]
Operating system	Unix-like
Type	Network monitoring
License	GNU General Public License
Website	www.nagios.org [4]

Nagios (/ˈnɑːɡiːoʊs/) is a popular open source computer system monitor, network monitoring and infrastructure monitoring software application. Nagios offers complete monitoring and alerting for servers, switches, applications, and services and is considered the defacto industry standard in IT infrastructure monitoring. It watches hosts and services. It alerts users when things go wrong and alerts them again when those wrong things get better/resolved.

Nagios, originally created under the name *NetSaint*, was written and is currently maintained by Ethan Galstad, along with a group of developers actively maintaining both official and unofficial plugins. N.A.G.I.O.S. is a recursive acronym: "Nagios Ain't Gonna Insist On Sainthood",[5] "Sainthood" being a reference to the original name *NetSaint*, which was changed in response to a legal challenge by owners of a similar trademark.[6] "Agios" is also a transliteration of the Greek word *άγιος* which means "saint".

Nagios was originally designed to run under Linux, but also runs well on other Unix variants. It is free software, licensed under the terms of the GNU General Public License version 2 as published by the Free Software Foundation.

In a 2006 survey among the nmap-hackers mailing list, 3243 people responded when asked for their favorite network security tools. Nagios came in 67th overall and 5th among traffic monitoring tools. Nmap itself was excluded from the list. Update: The 2011 list has been released and Nagios is listed at #69. [7]

Overview

Nagios is Open Source Software licensed under the GNU GPL V2.

- Monitoring of network services (SMTP, POP3, HTTP, NNTP, ICMP, SNMP, FTP, SSH)
- Monitoring of host resources (processor load, disk usage, system logs) on a majority of network operating systems, including Microsoft Windows with the NSClient++ [8] plugin or Check MK.
- Monitoring of anything else like probes (temperature, alarms...) which have the ability to send collected data via a network to specifically written plugins
- Monitoring via remotely-run scripts via Nagios Remote Plugin Executor
- Remote monitoring supported through SSH or SSL encrypted tunnels.
- Simple plugin design that allows users to easily develop their own service checks depending on needs, by using the tools of choice (shell scripts, C++, Perl, Ruby, Python, PHP, C#, etc.)
- Plugins available for graphing of data (Nagiosgraph [9], PNP4Nagios [10], Splunk for Nagios [11], and others available)
- Parallelized service checks available
- Ability to define network host hierarchy using "parent" hosts, allowing detection of and distinction between hosts that are down and those that are unreachable
- Contact notifications when service or host problems occur and get resolved (via e-mail, pager, SMS, or any user-defined method through plugin system)
- Ability to define event handlers to be run during service or host events for proactive problem resolution
- Automatic log file rotation
- Support for implementing redundant monitoring hosts
- Optional web-interface for viewing current network status, notifications, problem history, log files, etc.
- Data storage is done in text files rather than database

Nagios Agents

NRPE

Nagios Remote Plugin Executor (NRPE) is a Nagios agent that allows remote systems monitoring using scripts that are hosted on the remote systems. It allows for monitoring resources such as disk usage, system load or number of users currently logged in. Nagios periodically polls the agent on the remote system using the `check_nrpe` plugin.

NRDP

Nagios Remote Data Processor (NDRP) is a Nagios agent with a flexible data transport mechanism and processor. It is designed with an architecture that allows it to be easily extended and customized. NRDP uses standard ports and protocols (HTTP(S) and XML) and can be implemented as a replacement for NSCA.

NSClient++

This program is mainly used to monitor Windows machines. Being installed on a remote system NSClient++ listens to port TCP 1248. Nagios plugin that is used to collect information from this addon is called `check_nt`. As NRPE, NSClient++ allows to monitor the so called "private services" (memory usage, CPU load, disk usage, running processes, etc.)

References

[1] http://www.nagios.org/about/team

[2] first release of NetSaint from the changelog at http://web.archive.org/web/20060501150621/http://www.netsaint.org/changelog.php

[3] Nagios 3.x Version History (http://www.nagios.org/development/history/nagios-3x.php)

[4] http://www.nagios.org

[5] Galstad, Ethan (2003-05-03). offi "Nagios: FAQs : What does Nagios mean?" (http://support.nagios.com/knowledgebase/faqs/index.php?option=com_content&view=article&id=52&catid=35&faq_id=2&expand=false&showdesc=true). *Nagios: Frequently Asked Questions*. Nagios Enterprises, LLC. offi. Retrieved 2009-03-06. "The official meaning is that N.A.G.I.O.S. is a recursive acronym which stands for "Nagios Ain't Gonna Insist On Sainthood"."

[6] "2005-02-22 - Ethan Galstad" (http://archive.fosdem.org/2005/index/interviews/interviews_galstad.html). *Fosdem 2005*. 2005-02-22. . Retrieved 2009-03-06. "Although we were able to eventually reach an amicable agreement on my future use of the name "NetSaint", I felt it was prudent to change the name in order to prevent any future mishaps."

[7] Top 6 Traffic Monitoring Tools http://sectools.org/traffic-monitors.html

[8] http://nsclient.org/

[9] http://sourceforge.net/projects/nagiosgraph/

[10] http://sourceforge.net/projects/pnp4nagios

[11] http://exchange.nagios.org/directory/Addons/Log-File-Management/Splunk-For-Nagios/details

Further reading

- Barth, Wolfgang; (2006) *Nagios: System And Network Monitoring* (http://web.archive.org/web/20080531191735/http://www.nostarch.com/frameset.php?startat=nagios) - No Starch Press ISBN 1-59327-070-4
- Barth, Wolfgang; (2008) " Nagios: System And Network Monitoring, 2nd edition (http://www.nostarch.com/nagios_2e.htm) *- No Starch Press ISBN 1-59327-179-4*
- Turnbull, James; (2006) *Pro Nagios 2.0* (http://www.apress.com/book/bookDisplay.html?bID=10096) - San Francisco: Apress ISBN 1-59059-609-9
- Josephsen, David; (2007) *Building a Monitoring Infrastructure with Nagios* (http://www.pearson.ch/Informatik/PrenticeHall/1471/9780132236935/Building_a_Monitoring_Infrastructure.aspx) - Prentice Hall ISBN 0-13-223693-1
- Dondich, Taylor; (2006) *Network Monitoring with Nagios* (http://www.oreilly.com/catalog/networknagios/index.html) - O'Reilly ISBN 0-596-52819-1
- Schubert, Max et al.; (2008) *Nagios 3 Enterprise Network Monitoring* (http://www.nagios3book.com/) - Syngress ISBN 978-1-59749-267-6

External links

- Project website (http://www.nagios.org)
- Company website (http://www.nagios.com)
- Nagios Exchange (http://exchange.nagios.org) - thousands of Nagios plugins, addons, and utilities
- Nagwin - Nagios for Windows (http://www.itefix.no/i2/nagwin)

Comparison of network monitoring systems

This is a comparison of some notable network monitoring systems.

Name	IP SLA Reports	Logical Grouping	Trending	Trend Prediction	Auto Discovery	Agent	SNMP	Syslog	Plugins	Triggers / Alerts	WebApp	Distributed Monitoring	Inventory	Data Storage Method	License	Maps	Access Control	IPv6
AdRem NetCrunch	No	Yes	Yes	No	Yes	No	Yes	Yes	Yes	Yes	Viewing, Acknowledging	No	Yes	SQL	Commercial	Yes	Yes	Unknown
AccelOps	Yes	Yes	Yes	Yes	Yes	Supported	Yes	Yes	Yes	Yes	Full Control	Yes	Yes	PostgreSQL	Commercial	Yes	Yes	Unknown
AggreGate Network Manager	Yes	Yes	Yes	Yes	Yes	Supported	Yes	Yes	Yes	Yes	Full Control	Yes	Yes	MySQL, MS SQL, PostgreSQL, Oracle, Firebird, HSQLDB	Limited free, Commercial	Yes	Yes	Unknown
Argus	Yes	Yes	No	No	No	Yes	Yes	Yes	Yes	Yes	Viewing, Acknowledging	Yes	Unknown	Berkeley DB	Artistic License	No	Yes	Yes
CA Spectrum	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	MySQL	Commercial	Yes	Yes	Yes
Avaya VPFM	Yes	Yes	Yes	No	Yes	Yes	Yes	Yes	Yes	Yes	Full Control	Yes	Yes	MySQL	Commercial	Yes	Yes	Yes
Cacti	Yes	Yes	Yes	Yes	Via plugin	No	Yes	Yes	Yes	Yes	Full Control	Yes	Yes	RRDtool, MySQL	GPL	Plugin	Yes	Yes
collectd	No	No	No	No	Push model; multicast possible	Supported	Yes	Yes	Yes	Yes	Viewing	Yes	No	RRDtool	GPLv2	No	Apache ACL	Yes
Dhyan Network management System	Yes	Yes	Yes	Unknown	Yes	Supported	Yes	Yes	Yes	Yes	Full Control	Yes	Yes	MySQL, Oracle, Derby	Commercial	Yes	Yes	Yes
dopplerVUE	Yes	Yes	Yes	No	Yes	Yes	Yes	Yes	Yes	Yes	Full Control	Yes	Yes	MS SQL	Commercial	Yes	Yes	Yes
ExtraHop	Yes	Yes	Yes	No	Yes	N/A	Yes	No	No	Yes	Full Control	With ECM	Yes	Proprietary	Commercial	Yes	Yes	Yes
Name	**IP SLA Reports**	**Logical Grouping**	**Trending**	**Trend Prediction**	**Auto Discovery**	**Agent**	**SNMP**	**Syslog**	**Plugins**	**Triggers / Alerts**	**WebApp**	**Distributed Monitoring**	**Inventory**	**Data Storage Method**	**License**	**Maps**	**Access Control**	**IPv6**
FreeNATS	Yes	Yes	No	No	Yes	Yes	No	Via plugin	Yes	In PHP Code	Full Control	No	No	MySQL	GPL	No	Yes	Unknown
Ganglia	No	Yes	Yes	No	Via gmond check in	Yes	Via plugin	No	Yes	No	Viewing	Yes	Unknown	RRDtool	BSD	Yes	No	Unknown

GroundWork Monitor	Yes	Yes	Yes	Yes	Yes	Supported	Yes	Yes	Supported	Yes	Viewing, Acknowledging, Reporting, Most maintenance	Yes	Via NEDI	SQL, RRDTool	Limited free, Commercial	Yes	Yes	No
HP Network Node Manager (NNMi)	Yes	Yes	Yes	Yes	Yes	No	Yes	Via integration	Yes	Yes	Full Control	Yes	Yes	PostgreSQL, Oracle Database	Commercial	Yes	Yes	Yes
Kaseya Network Monitor	Yes	Yes	Yes	No	Yes	No	Yes	Yes	Yes	Yes	Full Control	Yes	Yes	FirebirdSQL	Commercial	Yes	Yes	Unknown
IBM Tivoli Network Manager	Possible via configuration	Yes	Yes	Yes	Yes	No	Yes	Yes	Yes	Yes	Yes	Yes	Yes	MySQL, Oracle Database, DB2	Commercial	Yes	Yes	Yes
Icinga	Via plugin	Yes	Yes	No	Via plugin	Supported	Via plugin	Via plugin	Yes	Yes	Full Control	Yes	Via plugin	MySQL, PostgreSQL, Oracle Database	GPL	Yes	Yes	Yes
InterMapper	Yes	Yes	Yes	No	Yes	Yes	Yes	Yes	Yes	Yes	Viewing	Yes	Yes	PostgreSQL	Limited free, Commercial	Yes	Yes	Yes
IPHost Network Monitor	Yes	Yes	Yes	No	Yes	Yes	Yes	No	Yes	Yes	Viewing, Acknowledging, Reporting	Yes	No	FirebirdSQL	Commercial	No	No	Unknown
isyVmon	Yes	Yes	Yes	No	Via plugin	Yes	Yes	Yes	Yes	Yes	Full Control	Yes	Via plugin	RRDtool, MySQL	Limited free, Commercial	Via plugin	Yes	Yes
LiveAction	Yes	Yes	Yes	No	Yes	No	Yes	No	Yes	Yes	Viewing, Reporting	Yes	Yes	Yes	Commercial	Yes	Yes	Yes
Name	**IP SLA Reports**	**Logical Grouping**	**Trending**	**Trend Prediction**	**Auto Discovery**	**Agent**	**SNMP**	**Syslog**	**Plugins**	**Triggers / Alerts**	**WebApp**	**Distributed Monitoring**	**Inventory**	**Data Storage Method**	**License**	**Maps**	**Access Control**	**IPv6**
Munin	No	No	Yes	Yes	No	Yes	Yes	No	Yes	Partial	Viewing	Yesvia nodes	Unknown	RRDtool	GPL	Unknown	Unknown	Yes
Nagios	Via plugin	Yes	Yes	No	Via plugin	Supported	Via plugin	Via plugin	Yes	Yes	Yes	Yes	Via plugin	Flat file, SQL	GPL	Yes	Yes	Yes
NagiosXI	Via plugin	Yes	Yes	No	Via plugin	Supported	Via plugin	Via plugin	Yes	Yes	Full Control	Yes	Via plugin	MySQL, PostgreSQL	Commercial	Yes	Yes	Yes
NetQoS Performance Center	Yes	Yes	Yes	Yes	Yes	No	Yes	Yes	Yes	Yes	Viewing, Acknowledging, Reporting	Yes	Yes	Yes	Commercial	Yes	Yes	Unknown
Network Instruments Observer Infrastructure	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Full Control	Yes	Yes	Proprietary Database	Commercial	Yes	Yes	Yes

NetXMS	No	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Full Control	Yes	No	MySQL, MS SQL, Oracle, PostgreSQL, SQLite	GPL	Yes	Yes	No
Nimsoft Monitoring Solution	Yes	Yes	Yes	Yes	Yes	Supported	Yes	Yes	Yes	Yes	Viewing, Acknowledging, Reporting	Yes	Yes	SQLServer, Oracle	Commercial	Yes	Yes	Yes
Observium	Yes	No	No	No	Yes	No	Yes	Yes	No	Yes	Full Control	No	Yes	RRDtool, MySQL	Limited free	Yes	Yes	Yes
Name	**IP SLA Reports**	**Logical Grouping**	**Trending**	**Trend Prediction**	**Auto Discovery**	**Agent**	**SNMP**	**Syslog**	**Plugins**	**Triggers / Alerts**	**WebApp**	**Distributed Monitoring**	**Inventory**	**Data Storage Method**	**License**	**Maps**	**Access Control**	**IPv6**
OpenKBM	Yes	Yes	Yes	Yes	Yes	Supported	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Proprietary with JDBC support	Commercial	Yes	Yes	Yes
OpenNMS	Yes	Yes	Yes	Unknown	Yes	Supported	Yes	Yes	Yes	Yes	Full Control	Yes	Yes	JRobin, PostgreSQL	GPLv3	Yes	Yes	Yes
OPNET's AppResponse Xpert	Yes	Yes	Yes	Yes	Yes	No	Yes	No	Yes	Yes	Viewing, Acknowledging, Reporting	Yes	No	Yes	Commercial	Yes	Yes	Unknown
Opsview	Yes	Yes	Yes	No	Yes	Yes	Yes	Yes	Yes	Yes	Full Control	Yes	No	SQL	GPL	Yes	Yes	Yes
op5 Monitor	Yes	Yes	Yes	No	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Via plugin	Flat file, SQL	Commercial	Yes	Yes	Yes
Owl	No	Yes	Yes	Yes	Yes	Yes	No	No	Yes	Yes	Viewing, Acknowledging, Reporting	Yes	Yes	Oracle, MySQL, Sybase	Commercial	Yes	Yes	No
PacketTrap	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Viewing, Reporting	Yes	Yes	SQL	Commercial	Yes	Yes	Yes
Pandora FMS	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Full Control	Yes	Yes	MySQL, PostgreSQL, Oracle	GPLv2; (Enterprise edition available)	Yes	Yes	Yes
Performance Co-Pilot	No	Yes	Yes	No	No	Yes	No	No	Yes	Yes	No	Yes	No	Flat file	GPL, LGPL	No	No	Unknown
PRTG Network Monitor	Yes	Yes	Yes	Yes	Yes	Supported	Yes	Yes	Yes	Yes	Full Control	Yes	Yes	Proprietary	Freeware and Commercial	Yes	Yes	Yes
Scrutinizer	Yes	Yes	Yes	No	No	No	Yes	Yes	Yes	Yes	Viewing, Acknowledging, Reporting	Yes	Yes	MySQL	Limited free, Commercial	Yes	Yes	Yes
Name	**IP SLA Reports**	**Logical Grouping**	**Trending**	**Trend Prediction**	**Auto Discovery**	**Agent**	**SNMP**	**Syslog**	**Plugins**	**Triggers / Alerts**	**WebApp**	**Distributed Monitoring**	**Inventory**	**Data Storage Method**	**License**	**Maps**	**Access Control**	**IPv6**
ServersCheck	Yes	Yes	Yes	No	Yes	Supported	Yes	Yes	Yes	Yes	Full Control	Yes	No	Flat file, ODBC	Commercial	Yes	Yes	Unknown
SevOne	Yes	Yes	Yes	Yes	Yes	Supported	Yes	No	Yes	Yes	Full Control	Yes	Yes	MySQL	Commercial	Yes	Yes	Yes

Shinken	Via plugin	Yes	Yes	No	Via plugin	Yes	Via plugin	Via plugin	Yes	Yes	Viewing, Acknowledging, Reporting	Yes	Via plugin	Flat file, MySQL, Oracle, CouchDB, Sqlite, MongoDB	AGPL	Yes	Yes	Yes
Solarwinds	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Full Control	Yes	Yes	SQL	Commercial	Yes	Yes	Yes
Spiceworks	No	Yes	Yes	No	Yes	Supported	Yes	No	Yes	Yes	Full Control	Yes	Yes	Sqlite	Commercial (Free)	Yes	Yes	No
TclMon	Yes	Yes	Yes	No	Yes	Supported	Yes	Yes	Yes	Yes	Viewing	No	No	RRDTool	BSD	Yes	Yes	No
Verax NMS	Yes	Yes	Yes	Yes	Yes	No	Yes	Yes	Yes	Yes	Full Control	Yes	No	Oracle, MySQL	Commercial	Yes	Yes	Unknown
WhatsUpGold	Yes	Yes	Yes	No	Yes	Yes	Yes	Yes	Yes	Yes	Full Control	Yes	Yes	SQL	Commercial	Yes	Yes	Yes
Xymon/Hobbit	Yes	Yes	Yes	No	No	Yes	Via Plugin	No	Yes	Yes	Viewing, Acknowledging	Yes	No	Flat file	GPL	No	Apache ACL	No
Name	**IP SLA Reports**	**Logical Grouping**	**Trending**	**Trend Prediction**	**Auto Discovery**	**Agent**	**SNMP**	**Syslog**	**Plugins**	**Triggers / Alerts**	**WebApp**	**Distributed Monitoring**	**Inventory**	**Data Storage Method**	**License**	**Maps**	**Access Control**	**IPv6**
Zabbix	Yes	Yes	Yes	No	Yes	Supported	Yes	Yes	Yes	Yes	Full Control	Yes	Yes	Oracle, MySQL, PostgreSQL, IBM DB2, SQLite	GPL	Yes	Yes	Yes
Zenoss	Yes	Yes	Yes	Yes	Yes	Supported	Yes	Yes	Yes	Yes	Full Control	Yes	Yes	ZODB, MySQL, RRDtool	GPL	Yes	Yes	Yes
Zyrion Traverse	Yes	Yes	Yes	Yes	Yes	Supported	Yes	Yes	Yes	Yes	Full Control	Yes	Yes	SQL	Commercial	Yes	Yes	Unknown
Name	**IP SLA Reports**	**Logical Grouping**	**Trending**	**Trend Prediction**	**Auto Discovery**	**Agent**	**SNMP**	**Syslog**	**Plugins**	**Triggers / Alerts**	**WebApp**	**Distributed Monitoring**	**Inventory**	**Data Storage Method**	**License**	**Maps**	**Access Control**	**IPv6**

Legend

Product Name

> The name of the software, linked to its Wikipedia article.

IP SLAs Reports

> Features reports on IP SLAs.

Logical Grouping

> Supports arranging the hosts or devices it monitors into user-defined groups.

Trending

> Provides trending of network data over time.

Trend Prediction

> The software features algorithms designed to predict future network statistics.

Auto Discovery

> The software automatically discovers hosts or network devices it is connected to.

Agent

The product relies on a software agent that must run on hosts it is monitoring, so that data can be pushed back to a central server. "Supported" means that an agent may be used, but is not mandatory. An SNMP daemon does not count as an agent.

SNMP

Able to retrieve and report on SNMP statistics.

Syslog

Able to receive and report on Syslogs.

Plugins

Architecture of the software based on a number of 'plugins' that provide additional functionality.

Triggers/Alerts

Capable of detecting threshold violations in network data, and alerting the administrator in some form.

WebApp

Runs as a web-based application.

- No: There is no web-based frontend for this software.
- Viewing: Network data can be viewed in a graphical web-based frontend.
- Acknowledging: Users can interact with the software through the web-based frontend to acknowledge alarms or manipulate other notifications.
- Reporting: Specific reports on network data can be configured by the user and executed through the web-based frontend.
- Full Control: ALL aspects of the product can be controlled through the web-based frontend, including low-level maintenance tasks such as software configuration and upgrades.

Distributed Monitoring

Able to leverage more than one server to distribute the load of network monitoring.

Inventory

Keeps a record of hardware and/or software inventory for the hosts and devices it monitors.

Data Storage Method

Main method used to store the network data it monitors.

License

License released under (e.g. GPL, BSD license, etc.).

Maps

Features graphical network maps that represent the hosts and devices it monitors, and the links between them.

Access Control

Features user-level security, allowing an administrator to prevent access to certain parts of the product on a per-user or per-role basis.

IPv6

Supports monitoring IPv6 hosts and/or devices, receiving IPv6 data, and running on an IPv6-enabled server. Supports communication using IPv6 to the SNMP agent via an IPv6 address.

Network File Systems

Network File System

Network File System (**NFS**) is a distributed file system protocol originally developed by Sun Microsystems in 1984,[1] allowing a user on a client computer to access files over a network in a manner similar to how local storage is accessed. NFS, like many other protocols, builds on the Open Network Computing Remote Procedure Call (ONC RPC) system. The Network File System is an open standard defined in RFCs, allowing anyone to implement the protocol.

Versions and variations

Original NFS version

The implementation details are defined in RFC 1094. Sun used version 1 only for in-house experimental purposes. When the development team added substantial changes to NFS version 1 and released it outside of Sun, they decided to release the new version as v2, so that version interoperation and RPC version fallback could be tested.[2]

NFSv2

Version 2 of the protocol (defined in RFC 1094, March 1989) originally operated entirely over UDP. Its designers meant to keep the protocol stateless, with locking (for example) implemented outside of the core protocol. People involved in the creation of NFS version 2 include Rusty Sandberg, Bob Lyon, Bill Joy, and Steve Kleiman.

NFSv2 only allowed the first 2 GB of a file to be read.

NFSv3

Version 3 (RFC 1813, June 1995) added:

- support for 64-bit file sizes and offsets, to handle files larger than 2 gigabytes (GB);
- support for asynchronous writes on the server, to improve write performance;
- additional file attributes in many replies, to avoid the need to re-fetch them;
- a READDIRPLUS operation, to get file handles and attributes along with file names when scanning a directory;
- assorted other improvements.

At the time of introduction of Version 3, vendor support for TCP as a transport-layer protocol began increasing. While several vendors had already added support for NFS Version 2 with TCP as a transport, Sun Microsystems added support for TCP as a transport for NFS at the same time it added support for Version 3. Using TCP as a transport made using NFS over a WAN more feasible.

NFSv4

Version 4 (RFC 3010, December 2000; revised in RFC 3530, April 2003), influenced by AFS and CIFS, includes performance improvements, mandates strong security, and introduces a stateful protocol.[3] Version 4 became the first version developed with the Internet Engineering Task Force (IETF) after Sun Microsystems handed over the development of the NFS protocols.

NFS version 4.1 (RFC 5661, January 2010) aims to provide protocol support to take advantage of clustered server deployments including the ability to provide scalable parallel access to files distributed among multiple servers (pNFS extension).

Other extensions

WebNFS, an extension to Version 2 and Version 3, allows NFS to integrate more easily into Web-browsers and to enable operation through firewalls. In 2007, Sun Microsystems open-sourced their client-side WebNFS implementation.[4]

Various side-band protocols have become associated with NFS, including:

- The byte-range advisory Network Lock Manager (NLM) protocol (added to support UNIX System V file locking APIs).
- The remote quota reporting (RQUOTAD) protocol (to allow NFS users to view their data-storage quotas on NFS servers).

NFS over RDMA is an adaptation of NFS that uses RDMA as a transport.[5][6]

Platforms

NFS is often used with Unix operating systems (such as Solaris, AIX and HP-UX) and Unix-like operating systems (such as Linux and FreeBSD). It is also available to operating systems such as the classic Mac OS, OpenVMS, Microsoft Windows, Novell NetWare, and IBM AS/400. Alternative remote file access protocols include the Server Message Block (SMB, also known as CIFS), Apple Filing Protocol (AFP), NetWare Core Protocol (NCP), and OS/400 File Server file system (QFileSvr.400). SMB and NetWare Core Protocol (NCP) occur more commonly than NFS on systems running Microsoft Windows; AFP occurs more commonly than NFS in Macintosh systems; and QFileSvr.400 occurs more commonly in AS/400 systems.

Typical implementation

Assuming a Unix-style scenario in which one machine (the client) requires access to data stored on another machine (the NFS server):

1. The server implements NFS daemon processes (running by default as `nfsd`) in order to make its data generically available to clients.
2. The server administrator determines what to make available, exporting the names and parameters of directories (typically using the `/etc/exports` configuration file and the `exportfs` command).
3. The server security-administration ensures that it can recognize and approve validated clients.
4. The server network configuration ensures that appropriate clients can negotiate with it through any firewall system.
5. The client machine requests access to exported data, typically by issuing a `mount` command. (The client asks the server (rpcbind) which port the NFS server is using, the client connects to the NFS server (nfsd), nfsd passes the request to mountd)
6. If all goes well, users on the client machine can then view and interact with mounted filesystems on the server within the parameters permitted.

Note that automation of the NFS mounting process may take place — perhaps using `/etc/fstab` and/or automounting facilities.

Protocol development versus competing protocols

1980s

NFS and ONC figured prominently in the network-computing war between Sun Microsystems and Apollo Computer, and later the UNIX wars (ca 1987-1996) between AT&T and Sun on one side, and Digital Equipment, HP, and IBM on the other.

During the development of the ONC protocol (called SunRPC at the time), only Apollo's Network Computing System (NCS) offered comparable functionality. Two competing groups developed over fundamental differences in the two remote procedure call systems. Arguments focused on the method for data-encoding — ONC's External Data Representation (XDR) always rendered integers in big-endian order, even if both peers of the connection had little-endian machine-architectures, whereas NCS's method attempted to avoid byte-swap whenever two peers shared a common endianness in their machine-architectures. An industry-group called the Network Computing Forum formed (March 1987) in an (ultimately unsuccessful) attempt to reconcile the two network-computing environments.

Later, Sun and AT&T announced that the two firms would jointly develop AT&T's next version of UNIX: System V Release 4. This caused many of AT&T's other licensees of UNIX System V to become concerned that this would put Sun in an advantaged position, and it ultimately led to Digital Equipment, HP, IBM, and others forming the Open Software Foundation (OSF) in 1988. Ironically, Sun and AT&T had previously competed over Sun's NFS versus AT&T's Remote File System (RFS), and the quick adoption of NFS over RFS by Digital Equipment, HP, IBM, and many other computer vendors tipped the majority of users in favor of NFS.

OSF solicited the proposals for various technologies, including the remote procedure call (RPC) system and the remote file access protocol. In the end, proposals for these two requirements, called respectively, the Distributed Computing Environment (DCE), and the Distributed File System (DFS) won over Sun's proposed ONC and NFS. DCE derived from a suite of technologies, including NCS and Kerberos. DFS used DCE as the RPC and derived from the Andrew File System (AFS).

1990s

Sun Microsystems and the Internet Society (ISOC) reached an agreement to cede "change control" of ONC RPC so that ISOC's engineering-standards body, the Internet Engineering Task Force (IETF), could publish standards documents (RFCs) documenting the ONC RPC protocols and could extend ONC RPC. OSF attempted to make DCE RPC an IETF standard, but ultimately proved unwilling to give up change control. Later, the IETF chose to extend ONC RPC by adding a new authentication flavor based on GSSAPI, RPCSEC GSS, in order to meet IETF's requirements that protocol standards have adequate security.

Later, Sun and ISOC reached a similar agreement to give ISOC change control over NFS, although writing the contract carefully to exclude NFS version 2 and version 3. Instead, ISOC gained the right to add new versions to the NFS protocol, which resulted in IETF specifying NFS version 4 in 2003.

2000s

By the 21st century, neither DFS nor AFS had achieved any major commercial success as compared to CIFS or NFS. IBM, which had previously acquired the primary commercial vendor of DFS and AFS, Transarc, donated most of the AFS source code to the free software community in 2000. The OpenAFS project lives on. In early 2005, IBM announced end of sales for AFS and DFS.

Present

NFSv4.1 adds the Parallel NFS pNFS [7] capability, which enables data access parallelism. The NFSv4.1 protocol defines a method of separating the filesystem meta-data from the location of the file data; it goes beyond the simple name/data separation by striping the data amongst a set of data servers. This is different from the traditional NFS server which holds the names of files and their data under the single umbrella of the server. There exist products which are multi-node NFS servers, but the participation of the client in separation of meta-data and data is limited. The NFSv4.1 client can be enabled to be a direct participant in the exact location of file data and avoid solitary interaction with the single NFS server when moving data.

The NFSv4.1 pNFS server is a collection of server resources or components; these are assumed to be controlled by the meta-data server.

The pNFS client still accesses a single meta-data server for traversal or interaction with the namespace; when the client moves data to and from the server it may be directly interacting with the set of data servers belonging to the pNFS server collection.

In addition to pNFS, NFSv4.1 provides Sessions, Directory Delegation and Notifications, Multi-server Namespace, ACL/SACL/DACL, Retention Attributions, and SECINFO_NO_NAME.

References

[1] "Design and Implementation of the Sun Network Filesystem" (http://citeseerx.ist.psu.edu/viewdoc/summary?doi=10.1.1.14.473). USENIX. 1985. .

[2] * *NFS Illustrated* (2000) by Brent Callaghan - ISBN 0-201-32570-5

[3] "NFS Version 4" (http://www.usenix.org/events/usenix05/tech/italks.html#nFSv4). USENIX. 2005-04-14. .

[4] yanfs.dev.java.net (https://yanfs.dev.java.net/)

[5] Tom Talpey (February 28, 2006). "NFS/RDMA Implementation(s) Update" (http://www.connectathon.org/talks06/talpey-cthon06-nfs-rdma.pdf). Network Appliance, Inc.. .

[6] Brent Callaghan (January 28, 2002). "NFS over RDMA" (http://www.usenix.org/events/fast02/wips/callaghan.pdf). Sun Microsystems. .

[7] http://www.pnfs.com

External links

- RFCs
  - RFC 5661 - Network File System (NFS) Version 4 Minor Version 1 Protocol
  - RFC 3530 - NFS Version 4 Protocol Specification
  - RFC 2054 - WebNFS Specification
  - RFC 2339 - Sun/ISOC NFS Change Control Agreement
  - RFC 2203 - RPCSEC_GSS Specification
  - RFC 1813 - NFS Version 3 Protocol Specification
  - RFC 1790 - Sun/ISOC ONC RPC Change Control Agreement
  - RFC 1094 - NFS Version 2 Protocol Specification
- IETF: Network File System Version 4 (nfsv4) Charter (http://www.ietf.org/html.charters/nfsv4-charter.html)
- Linux NFS Overview, FAQ and HOWTO Documents (http://nfs.sourceforge.net/)

- Christopher Smith (2006-05-02). "Linux NFS-HOWTO" (http://nfs.sourceforge.net/nfs-howto/index.html). Retrieved 2010-12-16.

- IBM: NFSv4 delivers seamless network access (http://www-128.ibm.com/developerworks/linux/library/l-nfsv4.html?ca=dgr-lnxw06NFSv4SeamlessNetAccess)
- NFS operation explained with sequence diagrams (http://www.eventhelix.com/RealtimeMantra/Networking/NFS_Protocol_Sequence_Diagram.pdf)

Server Message Block

In computer networking, **Server Message Block** (**SMB**), also known as **Common Internet File System** (**CIFS**, /ˈsɪfs/) operates as an application-layer network protocol[1] mainly used for providing shared access to files, printers, serial ports, and miscellaneous communications between nodes on a network. It also provides an authenticated inter-process communication mechanism. Most usage of SMB involves computers running Microsoft Windows, where it was known as "Microsoft Windows Network" before the subsequent introduction of Active Directory. Corresponding Windows services are the "Server Service" (for the server component) and "Workstation Service" (for the client component).

The Server Message Block protocol can run atop the Session (and lower) network layers in several ways:

- directly over TCP, port 445;[2]
- via the NetBIOS API, which in turn can run on several transports:[3]
  - on UDP ports 137, 138 & TCP ports 137, 139 — see NetBIOS over TCP/IP;
  - on several legacy protocols such as NBF (incorrectly referred to as *NetBEUI*).

History

Barry Feigenbaum originally designed SMB at IBM with the aim of turning DOS "Interrupt 33" (21h) local file access into a networked file system[4]. Microsoft has made considerable modifications to the most commonly used version. Microsoft merged the SMB protocol with the LAN Manager product which it had started developing for OS/2 with 3Com circa 1990, and continued to add features to the protocol in Windows for Workgroups (circa 1992) and in later versions of Windows.

SMB was originally designed to run on top of the NetBIOS/NetBEUI API (typically implemented with NBF, NetBIOS over IPX/SPX, or NBT). Since Windows 2000, SMB runs, by default, with a thin layer, similar to the Session Message packet of NBT's Session Service, on top of TCP, using TCP port 445 rather than TCP port 139 — a feature known as "direct host SMB".[2]

At around the time when Sun Microsystems announced WebNFS,[5] Microsoft launched an initiative in 1996 to rename SMB to **Common Internet File System** (**CIFS**), and added more features, including support for symbolic links, hard links, larger file sizes, and an initial attempt at supporting direct connections over TCP port 445 without requiring NetBIOS as a transport (a largely experimental effort that required further refinement). Microsoft submitted some partial specifications as Internet-Drafts to the IETF,[6] though these submissions have expired.

The Samba project originated with the aim of reverse engineering the SMB protocol and implementing an SMB server to allow MS-DOS clients to use SMB to access files on Sun Microsystems machines.[7] Because of the importance of the SMB protocol in interacting with the widespread Microsoft Windows platform, Samba became a popular free implementation of a compatible SMB client and server for interoperating with non-Microsoft operating systems.

Microsoft introduced SMB2 with Windows Vista in 2006, and later improved on it in Windows 7.

Implementation

Client-server approach

SMB works through a client-server approach, where a client makes specific requests and the server responds accordingly. One section of the SMB protocol specifically deals with access to filesystems, such that clients may make requests to a file server; but some other sections of the SMB protocol specialize in inter-process communication (IPC). The Inter-Process Communication (IPC) share or ipc$ is a network share on computers running Microsoft Windows. This virtual share is used to facilitate communication between processes and computers over SMB, often to exchange data between computers that have been authenticated.

Developers have optimized the SMB protocol for local subnet usage, but users have also put SMB to work to access different subnets across the Internet—exploits involving file-sharing or print-sharing in MS Windows environments usually focus on such usage.

SMB servers make their file systems and other resources available to clients on the network. Client computers may want access to the shared file systems and printers on the server, and in this primary functionality SMB has become best-known and most heavily used. However, the SMB file-server aspect would count for little without the NT domains suite of protocols, which provide NT-style domain-based authentication at the very least. Almost all implementations of SMB servers use NT Domain authentication to validate user-access to resources.

Samba

Samba is a free software re-implementation of the SMB/CIFS networking protocol, originally developed by Andrew Tridgell. As of version 3, Samba provides file and print services for Microsoft Windows clients and can integrate with a Windows NT 4.0 server domain, either as a Primary Domain Controller (PDC) or as a domain member. Samba4 installations can act as an Active Directory domain controller or member server, at Windows 2008 domain and forest functional levels.[8]

Performance issues

NetBIOS

The use of the SMB protocol has often correlated with a significant increase in broadcast traffic on a network. However the SMB itself does not use broadcasts—the broadcast problems commonly associated with SMB actually originate with the NetBIOS service location protocol. By default, a Microsoft Windows NT 4.0 server used NetBIOS to advertise and locate services. NetBIOS functions by broadcasting services available on a particular host at regular intervals. While this usually makes for an acceptable default in a network with a smaller number hosts, increased broadcast traffic can cause problems as the size of the network increases. The implementation of name resolution infrastructure in the form of Windows Internet Naming Service (WINS) or Domain Name System (DNS) resolves this problem. WINS was a proprietary implementation used with Windows NT 4.0 networks, but brought about its own issues and complexities in the design and maintenance of a Microsoft network.

Since the release of Windows 2000, the use of WINS for name resolution has been deprecated by Microsoft, with hierarchical Dynamic DNS now configured as the default name resolution protocol for all Windows operating systems. Resolution of (short) NETBIOS names by DNS requires that a DNS client expand short names, usually by appending a connection-specific DNS suffix to its DNS lookup queries. WINS can still be configured on clients as a secondary name resolution protocol for interoperability with legacy Windows environments and applications. Further, Microsoft DNS servers can forward name resolution requests to legacy WINS servers in order to support name resolution integration with legacy (pre-Windows 2000) environments that do not support DNS.

WAN performance issues

Network designers have found that latency has a significant impact on the performance of the SMB 1.0 protocol, that it performs more poorly than other protocols like FTP. Monitoring reveals a high degree of "chattiness" and a disregard of network latency between hosts.[9] For example, a VPN connection over the Internet will often introduce network latency. Microsoft has explained that performance issues come about primarily because SMB 1.0 is a block-level rather than a streaming protocol, that was originally designed for small LANs; it has a block size that is limited to 64K, SMB signing creates an additional overhead and the TCP window size is not optimized for WAN links.[10] Solutions to this problem include the updated SMB 2.0 protocol, Offline Files, TCP window scaling and WAN acceleration devices from various network vendors that cache and optimize SMB 1.0.[11]

Microsoft's modifications

Microsoft added several extensions to its own SMB implementation. For example, it added NTLM, then NTLMv2 authentication protocols in order to address security weakness in the original LanMan authentication. LanMan authentication derived from the original legacy SMB specification's requirement to use IBM "LanManager" passwords, but implemented DES in a flawed manner that allowed passwords to be cracked.[12] Later, Kerberos authentication was also added. The NT 4.0 Domain logon protocols initially used 40-bit encryption outside of the United States of America, because of export restrictions on stronger 128-bit encryption[13] (subsequently lifted in 1996 when President Bill Clinton signed Executive Order 13026[14]). Opportunistic locking support has changed with each server release.

Opportunistic locking

In the SMB protocol, **opportunistic locking** is a file locking mechanism designed to improve performance by controlling caching of network files by the client. Contrary to the traditional locks, **OpLocks** are not used in order to provide mutual exclusion. The main goal of OpLocks is to provide synchronization for caching. There are three types of opportunistic locks:

Batch Locks

Batch OpLocks were created originally to support a particular behavior of MS-DOS batch file execution operation in which the file is opened and closed many times in a short period, which is a performance problem. To solve this, a client may ask for a OpLock of type "batch". In this case, the client delays sending the close request and if a subsequent open request is given, the two requests cancel each other.

Exclusive Locks

When an application opens in "shared mode" a file hosted on an SMB server which is not opened by any other process (or other clients) the client receives an **exclusive OpLock** from the server. This means that the client may now assume that it is the only process with access to this particular file, and the client may now cache all changes to the file before committing it to the server. This is a performance improvement, since fewer round-trips are required in order to read and write to the file. If another client/process tries to open the same file, the server sends a message to the client (called a *break* or *revocation*) which invalidates the exclusive lock previously given to the client. The client then flushes all changes to the file.

Level 2 OpLocks

If an exclusive OpLock is held by a client and a locked file is opened by a third party, the client has to relinquish its exclusive OpLock to allow the other client's write/read access. A client may then receive a "Level 2 OpLock" from the server. A Level 2 OpLock allows the caching of read requests, but excludes write caching.

Breaks

In contrast with the SMB protocol's "standard" behavior, a break request may be sent *from* server *to* client. It informs the client that an OpLock is no longer valid. This happens, for example, when another client wishes to open a file in a way that invalidates the OpLock. The first client is then sent an OpLock break and required to send all its local changes (in case of batch or exclusive OpLocks), if any, and acknowledge the OpLock break. Upon this acknowledgment the server can reply to the second client in a consistent manner.

SMB2

Microsoft introduced a new version of the Server Message Block (SMB) protocol (SMB 2.0 or SMB2) with Windows Vista in 2006.[15] Although the protocol is proprietary, its specification has been published to allow other systems to interoperate with Microsoft operating systems that use the new protocol.[16]

SMB2 reduces the 'chattiness' of the SMB 1.0 protocol by reducing the number of commands and subcommands from over a hundred to just nineteen.[] It has mechanisms for pipelining, that is, sending additional requests before the response to a previous request arrives, thereby improving performance over high latency links. It adds the ability to compound multiple actions into a single request, which significantly reduces the number of round-trips the client needs to make to the server, improving performance as a result.[9] SMB1 also has a compounding mechanism — known as AndX — to compound multiple actions, but Microsoft clients rarely use AndX. It also introduces the notion of "durable file handles": these allow a connection to an SMB server to survive brief network outages, as are typical in a wireless network, without having to incur the overhead of re-negotiating a new session.

SMB2 includes support for symbolic links. Other improvements include caching of file properties, improved message signing with HMAC SHA-256 hashing algorithm and better scalability by increasing the number of users, shares and open files per server among others.[9] The SMB1 protocol uses 16-bit data sizes, which amongst other things, limits the maximum block size to 64K. SMB2 uses 32 or 64-bit wide storage fields, and 128 bits in the case of file-handles, thereby removing previous constraints on block sizes, which improves performance with large file transfers over fast networks.[9]

Windows Vista and later operating systems use SMB2 when communicating with other machines running Windows Vista or later. SMB1 continues in use for connections with older versions of Windows, as well as systems like Samba and various vendors' NAS solutions. Samba 3.5 also includes experimental support for SMB2.[17] Samba 3.6 fully supports SMB2, except the modification of user quotas using the Windows quota management tools.[18]

SMB2 brings a number of benefits to third party implementers of SMB protocols. SMB1, originally designed by IBM, was reverse engineered, and later became part of a wide variety of non-Windows operating systems such as Xenix, OS/2 and VMS (Pathworks). X/Open standardised it partially; it also had draft IETF standards which lapsed. (See http://ubiqx.org/cifs/Intro.html for historical detail.) SMB2 is also a relatively clean break with the past. Microsoft's SMB1 code has to work with a large variety of SMB clients and servers. SMB1 features many versions of information for commands (selecting what structure to return for a particular request) because features such as Unicode support were retro-fitted at a later date. SMB2 involves significantly reduced compatibility-testing for implementers of the protocol. SMB2 code has considerably less complexity since far less variability exists (for example, non-Unicode code paths become redundant as SMB2 requires Unicode support).

SMB 2.1

SMB 2.1, introduced with Windows 7 and Server 2008 R2, introduced minor performance enhancements with a new opportunistic locking mechanism.[19]

SMB 3.0

SMB 3.0 (previously named SMB 2.2)[20] will be introduced with Windows 8[20] and Windows Server 2012.[20] It will bring several significant changes that are aimed to add functionality and improve SMB2 performance, notably in virtualized data centers like SMB2 RDMA Transport Protocol and multichannel.[21]

Features

The SMB "Inter-Process Communication" (IPC) system provides named pipes and was one of the first inter-process mechanisms commonly available to programmers that provides a means for services to inherit the authentication carried out when a client first connected to an SMB server.

Some services that operate over named pipes, such as those which use Microsoft's own implementation of DCE/RPC over SMB, known as MSRPC over SMB, also allow MSRPC client programs to perform authentication, which over-rides the authorization provided by the SMB server, but only in the context of the MSRPC client program that successfully makes the additional authentication.

Since Windows domain controllers use SMB to transmit policies at login, they have packet-signing enabled by default to prevent man-in-the-middle attacks; the feature can also be turned on for any server running Windows NT 4.0 Service Pack 3 or later.[22] The design of Server Message Block version 2 (SMB2) aims to mitigate this performance-limitation by coalescing SMB signals into single packets.

SMB supports opportunistic locking — a special type of locking-mechanism — on files in order to improve performance.

SMB serves as the basis for Microsoft's Distributed File System implementation.

Security

Over the years, there have been many security vulnerabilities in Microsoft's implementation of the protocol or components that it directly relies on,[23][24][25] with the most recent vulnerability (at time of writing) involving the SMB2 implementation.[26]

Specifications for SMB and SMB2 Protocols

The specifications for the SMB are proprietary and were originally closed, thereby forcing other vendors and projects to reverse-engineer the protocol in order to interoperate with it. The SMB 1.0 protocol was eventually published some time after it was reverse engineered, whereas the SMB 2.0 procotol was made available from Microsoft's MSDN Open Specifications Developer Center from the outset.[27] There are a number of specifications that are relevant to the SMB protocol:

- MS-CIFS [28] MS-CIFS is a recent replacement (2007) for the draft-leach-cifs-v1-spec-02.txt a document widely used to implement SMB clients, but also known to have errors of omission and commission.
- MS-SMB [29] Specification for Microsoft Extensions to MS-CIFS
- MS-SMB2 [30] Specification for the SMB 2 protocol
- MS-FSSO [31] Describes the intended functionality of the Windows File Access Services System, how it interacts with systems and applications that need file services, and how it interacts with administrative clients to configure and manage the system.

References

[1] "Microsoft SMB Protocol and CIFS Protocol Overview" (http://msdn.microsoft.com/en-us/library/aa365233(VS.85).aspx). Microsoft. 2009-10-22. . Retrieved 2009-11-01.

[2] "Direct hosting of SMB over TCP/IP" (http://support.microsoft.com/kb/204279). Microsoft. 2007-10-11. . Retrieved 2009-11-01.

[3] Richard Sharpe (8 October 2002). "Just what is SMB?" (http://samba.anu.edu.au/cifs/docs/what-is-smb.html). . Retrieved 18 July 2011.

[4] WhichNAS "NAS Network Protocols Knowledge Base - NAS Basics" (http://www.whichnas.net/nasbasics/nas-network-protocols)

[5] YANFS yet another nfs (http://www.sun.com/software/webnfs/overview.xml)

[6] * Common Internet File System Protocol (CIFS/1.0) (http://www.tools.ietf.org/html/draft-heizer-cifs-v1-spec)

- CIFS Logon and Pass Through Authentication (http://www.tools.ietf.org/html/draft-leach-cifs-logon-spec)
- CIFS/E Browser Protocol (http://www.tools.ietf.org/html/draft-leach-cifs-browser-spec)
- CIFS Printing Specification (http://www.tools.ietf.org/html/draft-leach-cifs-print-spec)
- CIFS Remote Administration Protocol (http://www.tools.ietf.org/html/draft-leach-cifs-rap-spec)
- A Common Internet File System (CIFS/1.0) Protocol (http://www.tools.ietf.org/html/draft-leach-cifs-v1-spec)

[7] Tridgell, Andrew (June 27, 1997). "A bit of history and a bit of fun" (http://www.rxn.com/services/faq/smb/samba.history.txt). . Retrieved 2011-07-26.

[8] http://samba.2283325.n4.nabble.com/Samba-4-functional-levels-td3322760.html

[9] Jose Barreto (2008-12-09). "SMB2, a Complete Redesign of the Main Remote File Protocol for Windows" (http://blogs.technet.com/josebda/archive/2008/12/05/smb2-a-complete-redesign-of-the-main-remote-file-protocol-for-windows.aspx). Microsoft. . Retrieved 2009-11-01.

[10] Neil Carpenter (2004-10-26). "SMB/CIFS Performance Over WAN Links" (http://blogs.technet.com/neilcar/pages/247903.aspx). Microsoft. . Retrieved 2009-11-01.

[11] Mark Rabinovich, Igor Gokhman. "CIFS Acceleration Techniques" (http://www.snia.org/events/storage-developer2009/presentations/monday/MarkRabinovich-IgorGokhman-CIFS_Acceleration_Techniques.pdf). Storage Developer Conference, SNIA, Santa Clara 2009. .

[12] Christopher Hertel (1999). "SMB: The Server Message Block Protocol" (http://ubiqx.org/cifs/SMB.html). . Retrieved 2009-11-01.

[13] "Description of Microsoft Windows Encryption Pack 1" (http://support.microsoft.com/kb/159709). Microsoft. 2006-11-01. . Retrieved 2009-11-01.

[14] "US Executive Order 13026" (http://www.gpo.gov/fdsys/pkg/WCPD-1996-11-18/pdf/WCPD-1996-11-18-Pg2399.pdf). United States Government. 1996. . Retrieved 2009-11-01.

[15] Navjot Virk and Prashanth Prahalad (March 10, 2006). "What's new in SMB in Windows Vista" (http://blogs.msdn.com/chkdsk/archive/2006/03/10/548787.aspx). *Chk Your Dsks*. MSDN. . Retrieved 2006-05-01.

[16] "(MS-SMB2): Server Message Block (SMB) Version 2 Protocol Specification" (http://msdn.microsoft.com/en-us/library/cc246482(PROT.13).aspx). Microsoft. 2009-09-25. . Retrieved 2009-11-01.

[17] Samba 3.5 - Release Notes Archive (http://www.samba.org/samba/history/samba-3.5.0.html)

[18] Samba 3.6 - Release Notes Archive (http://samba.org/samba/history/samba-3.6.0.html)

[19] "Implementing an End-User Data Centralization Solution" (http://www.microsoft.com/downloads/details.aspx?displaylang=en&FamilyID=d8541618-5c63-4c4d-a0fd-d942cd3d2ec6). Microsoft. 2009-10-21. pp. 10–11. . Retrieved 2009-11-02.

[20] Jeffrey Snover (19 April 2012). "Windows Server Blog: SMB 2.2 is now SMB 3.0" (http://blogs.technet.com/b/windowsserver/archive/2012/04/19/smb-2-2-is-now-smb-3-0.aspx). Microsoft. . Retrieved 14 June 2012.

[21] "" (http://www.snia.org/sites/default/files2/SDC2011/presentations/keynote/ThomasPfenning_The_Future_of_File_Protocols-final.pdf). .

[22] "Overview of Server Message Block signing" (http://support.microsoft.com/kb/887429). Microsoft. 2007-11-30. . Retrieved 2009-11-01.

[23] "MS02-070: Flaw in SMB Signing May Permit Group Policy to Be Modified" (http://support.microsoft.com/kb/329170). Microsoft. 2007-12-01. . Retrieved 2009-11-01.

[24] "MS09-001: Vulnerabilities in SMB could allow remote code execution" (http://support.microsoft.com/kb/958687). Microsoft. 2009-01-13. . Retrieved 2009-11-01.

[25] "MS08-068: Vulnerability in SMB could allow remote code execution" (http://support.microsoft.com/kb/957097). Microsoft. 2009-02-26. . Retrieved 2009-11-01.

[26] "MS09-050: Vulnerabilities in SMB could allow remote code execution" (http://support.microsoft.com/kb/975517). Microsoft. 2009-10-13. . Retrieved 2009-02-26.

[27] Windows Protocols (http://msdn.microsoft.com/en-us/library/cc216517(PROT.10).aspx)

[28] http://msdn.microsoft.com/en-us/library/ee442092%28PROT.10%29.aspx

[29] http://msdn.microsoft.com/en-us/library/cc246231%28PROT.10%29.aspx

[30] http://msdn.microsoft.com/en-us/library/cc246482%28PROT.10%29.aspx

[31] http://msdn.microsoft.com/en-us/library/ee392367%28PROT.10%29.aspx

External links

- Hertel, Christopher (2003). *Implementing CIFS — The Common Internet FileSystem* (http://www.ubiqx.org/cifs/Book.html). Prentice Hall. ISBN 0-13-047116-X. (Text licensed under the Open Publication License, v1.0 or later, available from the link above.)
- Technical details about SMB/CIFS (http://ubiqx.org/cifs/)
- Common Internet File System (CIFS) File Access Protocol (http://www.microsoft.com/downloads/details.aspx?FamilyID=c4adb584-7ff0-4acf-bd91-5f7708adb23c&displaylang=en) - Technical details from Microsoft Corporation
- the NT LM 0.12 dialect of SMB (http://www.samba.org/samba/ftp/specs/smb-nt01.doc). In Microsoft Word format
- Samba development information (http://devel.samba.org/)
- Introduction to the Common Internet File System (CIFS): Leverage the Power of this Popular Network File Sharing Protocol (http://www.embeddedcomponents.com/marketplace/makers/visualitynq/intro/) Online introduction to CIFS: Lecture/blog by Ron Fredericks
- Zechner, Anton (2007). "Source-code of a free SMB server for small embedded systems" (http://members.inode.at/anton.zechner/az/AzSmb.en.htm)
- Description of the update that implements Extended Protection for Authentication in the Server service (http://support.microsoft.com/kb/2345886)
- Post about SMB3 in the Windows Server Blog (http://blogs.technet.com/b/windowsserver/archive/2012/04/19/smb-2-2-is-now-smb-3-0.aspx)

Samba (software)

Samba

Initial release	1992 [1]
Stable release	3.6.5 / April 30, 2012
Development status	active
Operating system	Multiplatform
Type	Network file system
License	GNU General Public License version 3
Website	www.samba.org [2]

Samba is a free software re-implementation of the SMB/CIFS networking protocol, originally developed by Andrew Tridgell. As of version 3, Samba provides file and print services for various Microsoft Windows clients and can integrate with a Windows Server domain, either as a Primary Domain Controller (PDC) or as a domain member. It can also be part of an Active Directory domain.

Samba runs on most Unix and Unix-like systems, such as GNU/Linux, Solaris, AIX and the BSD variants, including Apple's Mac OS X Server (which was added to the Mac OS X client in version 10.2). Samba is standard on nearly all distributions of Linux and is commonly included as a basic system service on other Unix-based operating systems as well. Samba is released under the GNU General Public License. The name *Samba* comes from SMB (Server

Message Block), the name of the standard protocol used by the Microsoft Windows network file system.

Early history

Andrew Tridgell developed the first version of Samba Unix in December 1991 and January 1992, as a PhD student at the Australian National University, using a packet sniffer to do network analysis of the protocol used by DEC Pathworks server software. At the time of the first releases, versions 0.1, 0.5 and 1.0, all from the first half of January 1992, it did not have a proper name, and Tridgell just called it "a Unix file server for Dos Pathworks". At the time of version 1.0, he realized that he "had in fact implemented the netbios protocol" and that "this software could be used with other PC clients".

With a focus on interoperability with Microsoft's LAN Manager, Tridgell released "netbios for unix", nbserver, version 1.5 in December 1993. This release was the first to include client-software as well as a server. Also, at this time GPL2 was chosen as license.

Midway through the 1.5-series, the name was changed to *smbserver*. However, Tridgell got a trademark notice from the company "Syntax", who sold a product named *TotalNet Advanced Server* and owned the trademark for "SMBserver". The name "Samba" was derived by running the Unix command grep through the system dictionary looking for words that contained the letters S, M, and B, in that order (i.e. `grep -i '^s.*m.*b' /usr/share/dict/words`).[3]

Versions 1.6, 1.7, 1.8, and 1.9 followed relatively quickly, with the latter being released in January 1995. Tridgell considers the adoption of CVS in May 1996 to mark the birth of the Samba Team, though there had been contributions from other people, especially Jeremy Allison, previously.[4]

Version 2.0.0 was released in January 1999, and version 2.2.0 in April 2001.

Version History

Note: *** All users of versions 3.6.3 and previous releases should look to upgrading due to a **serious** issues with alert CVE-2012-1182 [5] (Contributors please provide updates, literally millions of Samba hardware devices implemented in the last decade could be rendered obsolete) ***

Version 3.0.0, released on 23 September 2003, was a major upgrade. Samba gained the ability to join Active Directory as a member, though not as a domain controller.[6] Subsequent point-releases to 3.0 have added minor new features. Currently, the latest release in this series is 3.0.37, released 1 October 2009, and shipped on a voluntary basis.[7] The 3.0.x series officially reached end-of-life on 5 August 2009.[7]

Version 3.1 was used only for development.

With version 3.2, the project decided to move to time-based releases. New major releases, such as 3.3, 3.4, etc. will appear every 6 months. New features will only be added when a major release is done, point-releases will be only for bug fixes.[8] Also, 3.2 marked a change of license from GPL2 to GPL3, with some parts released under LGPL3.[9] The main technical change in version 3.2 was to autogenerate much of the DCE/RPC-code that used to be handcrafted. Version 3.2.0 was released on 1 July 2008.[10] It will be updated on an as-needed basis for security issues only[11] and its current release is 3.2.15 from 1 October 2009. The 3.2.x series officially reached end-of-life on 1 March 2010.[11]

Version 3.3.0 was released 27 January 2009 and is now at version 3.3.16 in this branch.[12]

Version 3.4.0 was released 3 July 2009. This was the first release to include both Samba 3 and Samba 4 source code.[13]

Version 3.4.14 was released 26 July 2011. It is the latest stable release of the Samba 3.4 series.[14]

Version 3.5.0 was released 1 March 2010. This was the first release to include experimental support for SMB2.[15]

Version 3.5.11 was released on 4 August 2011. It is the latest stable release of the Samba 3.5 series.[16][17]

Version 3.6.0 was released on 9 August 2011. This is the first branch which includes full support for SMB2.[18]

Version 4.0 is planned as a major rewrite that will enable Samba to be an Active Directory domain controller. After three years of development, the first technical preview (4.0.0TP1) was released in January 2006.[19] Subsequently, new previews and then alphas have followed regularly. Beta 3 was released on 2nd July 2012 .[20]

Features

Samba allows file and print sharing between computers running Windows and computers running Unix. It is an implementation of dozens of services and a dozen protocols, including:

- The NetBIOS over TCP/IP (NBT)
- SMB
- CIFS (an enhanced version of SMB)
- DCE/RPC or more specifically, MSRPC, the Network Neighborhood suite of protocols
- A WINS server also known as a NetBIOS Name Server (NBNS)
- The NT Domain suite of protocols which includes NT Domain Logons
- Secure Accounts Manager (SAM) database
- Local Security Authority (LSA) service
- NT-style printing service (SPOOLSS), NTLM and more recently Active Directory Logon which involves a modified version of Kerberos and a modified version of LDAP.

All these services and protocols are frequently incorrectly referred to as just NetBIOS or SMB. The NetBIOS and WINS protocols are deprecated on Windows.

Samba sets up network shares for chosen Unix directories (including all contained subdirectories). These appear to Microsoft Windows users as normal Windows folders accessible via the network. Unix users can either mount the shares directly as part of their file structure using the smbmount command or, alternatively, can use a utility, smbclient (libsmb) installed with Samba to read the shares with a similar interface to a standard command line FTP program. Each directory can have different access privileges overlaid on top of the normal Unix file protections. For example: home directories would have read/write access for all known users, allowing each to access their own files. However they would still not have access to the files of others unless that permission would normally exist. Note that the netlogon share, typically distributed as a read only share from `/etc/samba/netlogon`, is the logon directory for user logon scripts.

Samba services are implemented as two daemons:

- smbd, which provides the file and printer sharing services, and
- nmbd, which provides the NetBIOS-to-IP-address name service. NetBIOS over TCP/IP requires some method for mapping NetBIOS computer names to the IP addresses of a TCP/IP network.

Samba configuration is achieved by editing a single file (typically installed as `/etc/smb.conf` or `/etc/samba/smb.conf`). Samba can also provide user logon scripts and group policy implementation through `poledit`.

Samba is included in most Linux distributions and is started during the boot process. On Red Hat, for instance, the `/etc/rc.d/init.d/smb` script runs at boot time, and starts both daemons. Samba is not included in Solaris 8, but a Solaris 8-compatible version is available from the Samba website.

Samba includes a web administration tool called *Samba Web Administration Tool* (SWAT).[21][22][23]

Samba TNG

Samba TNG

Developer(s)	Samba TNG team
Stable release	0.5-rc1 / December 3, 2009
Development status	Unmaintained
Operating system	Cross-platform
Type	Microsoft networking
License	GNU General Public License
Website	www.samba-tng.org [24]

Samba TNG (The Next Generation) was forked in late 1999, after disagreements between the Samba Team leaders and Luke Leighton about the directions of the Samba project. They failed to come to an agreement on a development transition path which allowed the research version of Samba he was developing (known at the time as Samba-NTDOM) to slowly be integrated into Samba.[25]

Since the project started, development has been minimal, due to a lack of developers. As such the Samba TNG team frequently recommends to people who are unsure of which program to use to try Samba instead, as they have more developers and are able to support more platforms and situations.[26]

One of the key goals of the Samba TNG project is to rewrite all of the NT Domains services as FreeDCE projects.[27] Making this rewriting goal difficult is the fact that services were all developed manually through network reverse-engineering, with limited or no reference to DCE/RPC documentation.

The key differences between the two programs are in the implementation of the NT Domains suite of protocols and MSRPC services. Samba makes all the NT Domains services available from a single place, whereas Samba TNG has separated each service into its own program.

ReactOS has started using Samba TNG services for its SMB implementation. The developers of both projects were interested in seeing the Samba TNG design used to help get ReactOS talking to Windows networks. They have been working together to adapt the network code and build system. The multi-layered and modular approach made it easy to port each service to ReactOS.[28]

References

[1] "10 years of Samba" (http://www.samba.org/samba/docs/10years.html). . Retrieved 2011-08-12.
[2] http://www.samba.org/
[3] "A bit of history and a bit of fun" (http://www.rxn.com/services/faq/smb/samba.history.txt). 1997-06-27. . Retrieved 2009-05-28.
[4] 10 years of Samba! (http://samba.org/samba/docs/10years.html)
[5] https://www.samba.org/samba/security/CVE-2012-1182
[6] The first stable release of Samba 3.0 is available (http://lists.samba.org/archive/samba-announce/2003/000021.html)
[7] Release Planning for Samba 3.0 - SambaWiki (http://wiki.samba.org/index.php/Release_Planning_for_Samba_3.0)
[8] Monday, April 28 - Samba Mashup Report (http://lists.samba.org/archive/samba-announce/2008/000136.html)
[9] Samba Adopts GPLv3 for Future Releases (http://lists.samba.org/archive/samba-announce/2007/000122.html)
[10] [ANNOUNCE] Samba 3.2.0 Available for Download (http://lists.samba.org/archive/samba-announce/2008/000145.html)
[11] Release Planning for Samba 3.2 - SambaWiki (http://wiki.samba.org/index.php/Release_Planning_for_Samba_3.2)
[12] Samba 3.3.16 Available for Download (http://samba.org/samba/history/samba-3.3.16.html)
[13] Samba - Release Notes Archive (http://www.samba.org/samba/history/samba-3.4.0.html)
[14] Samba 3.4.14 Available for Download (http://samba.org/samba/history/samba-3.4.14.html)
[15] Samba - Release Notes Archive (http://www.samba.org/samba/history/samba-3.5.0.html)
[16] Seeger, Karolin (2011-06-14). "Samba 3.5.9 Available for Download" (http://lists.samba.org/archive/samba-announce/2011/000232.html). *samba-announce mailing list*. .

[17] Samba 3.5.11 Available for Download (http://samba.org/samba/history/samba-3.5.11.html)
[18] Samba 3.6.0 Available for Download (http://samba.org/samba/history/samba-3.6.0.html)
[19] Samba 4.0.0TP1 Available for Download (http://news.samba.org/releases/4.0.0tp1/)
[20] Samba4 Releases 4.0.0beta2 (https://wiki.samba.org/index.php/Samba4/Releases/4.0.0beta2)
[21] Chapter 37. SWAT: The Samba Web Administration Tool (http://us3.samba.org/samba/docs/man/Samba-HOWTO-Collection/SWAT.html)
[22] Configure IT Quick: Simplify Samba configuration with SWAT (http://articles.techrepublic.com.com/5100-10878_11-1035709.html)
[23] Linux.com :: SWAT your Samba problems (http://www.linux.com/feature/125452)
[24] http://www.samba-tng.org/
[25] "Project FAQ - What is the relationship between Samba and Samba TNG?" (http://wiki.samba-tng.org/doku.php/pages/faq#what_is_the_relationship_between_samba_and_samba-tng). . Retrieved 2008-02-19.
[26] "Project FAQ - Which should I use - Samba or Samba TNG?" (http://wiki.samba-tng.org/doku.php/pages/faq#which_should_i_use_-_samba_or_samba-tng). . Retrieved 2008-02-19.
[27] "Project FAQ - What's all this about FreeDCE?" (http://wiki.samba-tng.org/doku.php/pages/faq#what_s_all_this_about_freedce). . Retrieved 2008-02-19.
[28] Vincent, Brian. "Interview with Steven Edwards" (http://www.winehq.org/?interview=14). Wine HQ. . Retrieved 2008-02-19.

External links

- Official websites:
  - Samba Official Website (http://www.samba.org/)
  - How Samba was written (http://samba.org/ftp/tridge/misc/french_cafe.txt)
  - http://us1.samba.org/samba/tng.html Article (http://us1.samba.org/samba/tng.html) on Samba-TNG from Andrew Tridgell (Oct. 2000)
  - " Using Samba (http://samba.org/samba/docs/man/using_samba/toc.html)"-E-book (Published by O'Reilly as ISBN 0-596-00256-4)
- Other:
  - Samba-3 by Example (http://www.informit.com/content/images/0131472216/downloads/0131472216_book.pdf), book licensed under the OPL
  - A history of Samba, written in 1994 (http://www.rxn.com/services/faq/smb/samba.history.txt), by Andrew Tridgell
  - Samba-TNG project homepage (http://www.samba-tng.org/)
  - Coverage on the Samba code fork (http://news.zdnet.com/2100-9595_22-524722.html?legacy=zdnn) from ZDNet

Distributed file system

In computing, a **distributed file system** or **network file system** is any file system that allows access to files from multiple hosts sharing via a computer network.[1] This makes it possible for multiple users on multiple machines to share files and storage resources.

The client nodes do not have direct access to the underlying block storage but interact over the network using a protocol. This makes it possible to restrict access to the file system depending on access lists or capabilities on both the servers and the clients, depending on how the protocol is designed.

In contrast, in a shared disk file system all nodes have equal access to the block storage where the file system is located. On these systems the access control must reside on the client.

Distributed file systems may include facilities for transparent replication and fault tolerance. That is, when a limited number of nodes in a file system go offline, the system continues to work without any data loss.

The difference between a distributed file system and a distributed data store can be vague, but DFSes are generally geared towards use on local area networks.

History and examples

The first file servers were developed in the 1970s. In 1976 Digital Equipment Corporation created the File Access Listener (FAL), an implementation of the Data Access Protocol as part of DECnet Phase II which became the first widely used network file system. In 1985 Sun Microsystems created the file system called "Network File System" (NFS) which became the first widely used Internet Protocol based network file system. Other notable network file systems are Andrew File System (AFS), Apple Filing Protocol (AFP), NetWare Core Protocol (NCP), and Server Message Block (SMB) which is also known as Common Internet File System (CIFS).

Transparency

Transparency is usually built into distributed file systems, so that files accessed over the network can be treated the same as files on local disk by programs and users. The multiplicity and dispersion of servers and storage devices are thus made invisible. It is up to the network file system to locate the files and to arrange for the transport of the data.

Performance

A common performance measurement of a network file system is the amount of time needed to satisfy service requests. In conventional systems, this time consists of a disk-access time and a small amount of CPU-processing time. But in a network file system, a remote access has additional overhead due to the distributed structure. This includes the time to deliver the request to a server, the time to deliver the response to the client, and for each direction, a CPU overhead of running the communication protocol software. The performance of a network file system can be viewed as one dimension of its transparency; to be fully equivalent, it would need to be comparable to that of a local disk.

Concurrent file updates

Concurrency control becomes an issue when more than one person or client is accessing the same file and want to update it. Hence updates to the file from one client should not interfere with access and updates from other clients. Concurrency control or locking may either be built into the file system or provided by an add-on protocol.

References

[1] Silberschatz, Galvin (1994). *Operating System concepts*, chapter 17 *Distributed file systems*. Addison-Wesley Publishing Company. ISBN 0-201-59292-4.

External links

- *A distributed file system for distributed conferencing system* (http://purl.fcla.edu/fcla/etd/UFE0001123) ("A *DFS* for the *DCS*") by Philip S Yeager, Thesis, University of Florida, 2003. (pdf)

Email Servers

Simple Mail Transfer Protocol

Simple Mail Transfer Protocol (**SMTP**) is an Internet standard for electronic mail (e-mail) transmission across Internet Protocol (IP) networks. SMTP was first defined by RFC 821 (1982, eventually declared STD 10),[1] and last updated by RFC 5321 (2008)[2] which includes the extended SMTP (ESMTP) additions, and is the protocol in widespread use today. SMTP uses TCP port 25. The protocol for new submissions (MSA) is effectively the same as SMTP, but it uses port 587 instead. SMTP connections secured by SSL are known by the shorthand SMTPS, though SMTPS is not a protocol in its own right.

While electronic mail servers and other mail transfer agents use SMTP to send and receive mail messages, user-level client mail applications typically only use SMTP for sending messages to a mail server for relaying. For receiving messages, client applications usually use either the Post Office Protocol (POP) or the Internet Message Access Protocol (IMAP) or a proprietary system (such as Microsoft Exchange or Lotus Notes/Domino) to access their mail box accounts on a mail server.

History

Various forms of one-to-one electronic messaging were used in the 1960s. People communicated with one another using systems developed for specific mainframe computers. As more computers were interconnected, especially in the US Government's ARPANET, standards were developed to allow users of different systems to e-mail one another. SMTP grew out of these standards developed during the 1970s.

SMTP can trace its roots to two implementations described in 1971, the Mail Box Protocol, which has been disputed to actually have been implemented,[3] but is discussed in RFC 196 and other RFCs, and the SNDMSG program, which, according to RFC 2235, Ray Tomlinson of BBN "invents" for TENEX computers the sending of mail across the ARPANET.[4][5][6] Fewer than 50 hosts were connected to the ARPANET at this time.[7]

Further implementations include FTP Mail [8] and Mail Protocol, both from 1973.[9] Development work continued throughout the 1970s, until the ARPANET converted into the modern Internet around 1980. Jon Postel then proposed a Mail Transfer Protocol in 1980 that began to remove the mail's reliance on FTP.[10] SMTP was published as RFC 788 in November 1981, also by Postel.

The SMTP standard was developed around the same time as Usenet, a one-to-many communication network with some similarities.

SMTP became widely used in the early 1980s. At the time, it was a complement to Unix to Unix Copy Program (UUCP) mail, which was better suited for handling e-mail transfers between machines that were intermittently connected. SMTP, on the other hand, works best when both the sending and receiving machines are connected to the network all the time. Both use a store and forward mechanism and are examples of push technology. Though Usenet's newsgroups are still propagated with UUCP between servers,[11] UUCP mail has virtually disappeared[12] along with the "bang paths" it used as message routing headers.

The article about sender rewriting contains technical background info about the early SMTP history and source routing before RFC 1123.

Released with 4.1cBSD, right after RFC 788, Sendmail was one of the first (if not the first) mail transfer agents to implement SMTP.[13] Over time, as BSD Unix became the most popular operating system on the Internet, sendmail became the most common MTA.[14] Some other popular SMTP server programs include Postfix, qmail, Novell

GroupWise, Exim, Novell NetMail, Microsoft Exchange Server, Sun Java System Messaging Server.

Message submission (RFC 2476) and SMTP-AUTH (RFC 2554) were introduced in 1998 and 1999, both describing new trends in e-mail delivery. Originally, SMTP servers were typically internal to an organization, receiving mail for the organization *from the outside*, and relaying messages from the organization *to the outside*. But as time went on, SMTP servers (Mail transfer agents), in practice, were expanding their roles to become message submission agents for Mail user agents, some of which were now relaying mail *from the outside* of an organization. (e.g. a company executive wishes to send e-mail while on a trip using the corporate SMTP server.) This issue, a consequence of the rapid expansion and popularity of the World Wide Web, meant that SMTP had to include specific rules and methods for relaying mail and authenticating users to prevent abuses such as relaying of unsolicited e-mail (spam).

As this protocol started out purely ASCII text-based, it did not deal well with binary files, or characters in many non-English languages. Standards such as Multipurpose Internet Mail Extensions (MIME) were developed to encode binary files for transfer through SMTP. Mail transfer agents (MTAs) developed after Sendmail also tended to be implemented 8-bit-clean, so that the alternate "just send eight" strategy could be used to transmit arbitrary text data (in any 8-bit ASCII-like character encoding) via SMTP. Mojibake was still a problem due to differing character set mappings between vendors, although the email addresses themselves still allowed only ASCII. 8-bit-clean MTAs today tend to support the 8BITMIME extension, permitting binary files to be transmitted almost as easily as plain text. Recently the SMTPUTF8 extension was created to support UTF-8 text, allowing international content and addresses in non-Latin scripts like Cyrillic or Chinese.

Many people contributed to the core SMTP specifications, among them Jon Postel, Eric Allman, Dave Crocker, Ned Freed, Randall Gellens, John Klensin, and Keith Moore.

Mail processing model

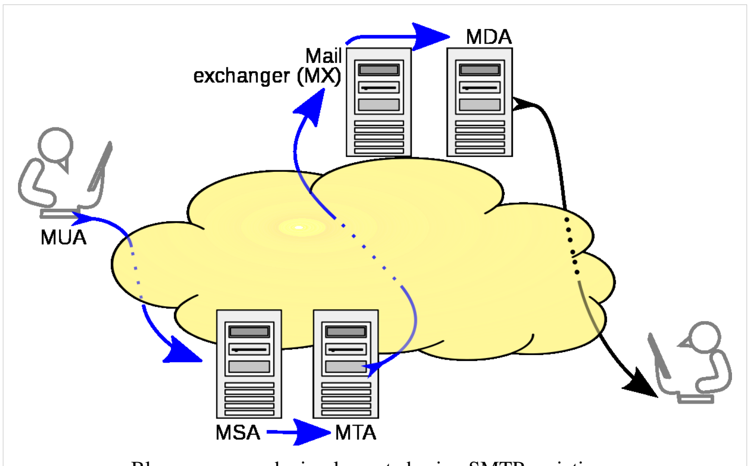

Blue arrows can be implemented using SMTP variations.

Email is submitted by a mail client (MUA, mail user agent) to a mail server (MSA, mail submission agent) using SMTP on TCP port 587. Most mailbox providers still allow submission on traditional port 25. From there, the MSA delivers the mail to its mail transfer agent (MTA, mail transfer agent). Often, these two agents are just different instances of the same software launched with different options on the same machine. Local processing can be done either on a single machine, or split among various appliances; in the former case, involved processes can share files; in the latter case, SMTP is used to transfer the message internally, with each host configured to use the next appliance as a smart host. Each process is an MTA in its own right; that is, an SMTP server.

The boundary MTA has to locate the target host. It uses the Domain name system (DNS) to look up the mail exchanger record (MX record) for the recipient's domain (the part of the address on the right of @). The returned MX record contains the name of the target host. The MTA next connects to the exchange server as an SMTP client. (The article on MX record discusses many factors in determining which server the sending MTA connects to.)

Once the MX target accepts the incoming message, it hands it to a mail delivery agent (MDA) for local mail delivery. An MDA is able to save messages in the relevant mailbox format. Again, mail reception can be done using many computers or just one —the picture displays two nearby boxes in either case. An MDA may deliver messages directly to storage, or forward them over a network using SMTP, or any other means, including the Local Mail Transfer Protocol (LMTP), a derivative of SMTP designed for this purpose.

Once delivered to the local mail server, the mail is stored for batch retrieval by authenticated mail clients (MUAs). Mail is retrieved by end-user applications, called email clients, using Internet Message Access Protocol (IMAP), a protocol that both facilitates access to mail and manages stored mail, or the Post Office Protocol (POP) which typically uses the traditional mbox mail file format or a proprietary system such as Microsoft Exchange/Outlook or Lotus Notes/Domino. Webmail clients may use either method, but the retrieval protocol is often not a formal standard.

SMTP defines message *transport*, not the message *content*. Thus, it defines the mail *envelope* and its parameters, such as the envelope sender, but not the header or the body of the message itself. STD 10 and RFC 5321 define SMTP (the envelope), while STD 11 and RFC 5322 define the message (header and body), formally referred to as the Internet Message Format.

Protocol overview

SMTP is a connection-oriented, text-based protocol in which a mail sender communicates with a mail receiver by issuing command strings and supplying necessary data over a reliable ordered data stream channel, typically a Transmission Control Protocol (TCP) connection. An *SMTP session* consists of commands originated by an SMTP client (the initiating agent, sender, or transmitter) and corresponding responses from the SMTP server (the listening agent, or receiver) so that the session is opened, and session parameters are exchanged. A session may include zero or more SMTP transactions. An *SMTP transaction* consists of three command/reply sequences (see example below.) They are:

1. **MAIL** command, to establish the return address, a.k.a. Return-Path, 5321.From, mfrom, or envelope sender. This is the address for bounce messages.
2. **RCPT** command, to establish a recipient of this message. This command can be issued multiple times, one for each recipient. These addresses are also part of the envelope.
3. **DATA** to send the *message text*. This is the content of the message, as opposed to its envelope. It consists of a *message header* and a *message body* separated by an empty line. DATA is actually a group of commands, and the server replies twice: once to the *DATA command* proper, to acknowledge that it is ready to receive the text, and the second time after the end-of-data sequence, to either accept or reject the entire message.

Besides the intermediate reply for DATA, each server's reply can be either positive (2xx reply codes) or negative. Negative replies can be permanent (5xx codes) or transient (4xx codes). A **reject** is a permanent failure by an SMTP server; in this case the SMTP client should send a bounce message. A **drop** is a positive response followed by message discard rather than delivery.

The initiating host, the SMTP client, can be either an end-user's email client, functionally identified as a mail user agent (MUA), or a relay server's mail transfer agent (MTA), that is an SMTP server acting as an SMTP client, in the relevant session, in order to relay mail. Fully capable SMTP servers maintain queues of messages for retrying message transmissions that resulted in transient failures.

A MUA knows the *outgoing mail* SMTP server from its configuration. An SMTP server acting as client, i.e. *relaying*, typically determines which SMTP server to connect to by looking up the MX (Mail eXchange) DNS resource record for each recipient's domain name. Conformant MTAs (not all) fall back to a simple A record in case no MX record can be found. Relaying servers can also be configured to use a smart host.

An SMTP server acting as client initiates a TCP connection to the server on the "well-known port" designated for SMTP: port 25. MUAs should use port 587 to connect to an MSA. The main difference between an MTA and an MSA is that SMTP Authentication is mandatory for the latter only.

SMTP vs mail retrieval

SMTP is a delivery protocol only. It cannot *pull* messages from a remote server on demand. Other protocols, such as the Post Office Protocol (POP) and the Internet Message Access Protocol (IMAP) are specifically designed for retrieving messages and managing mail boxes. However, SMTP has a feature to initiate mail queue processing on a remote server so that the requesting system may receive any messages destined for it (see Remote Message Queue Starting below). POP and IMAP are preferred protocols when a user's personal computer is only intermittently powered up, or Internet connectivity is only transient and hosts cannot receive message during off-line periods.

Remote Message Queue Starting

Remote Message Queue Starting is a feature of SMTP that permits a remote host to start processing of the mail queue on a server so it may receive messages destined to it by sending the `TURN` command. This feature however was deemed insecure[15] and was extended in RFC 1985 with the `ETRN` command which operates more securely using an authentication method based on Domain Name System information.

On-Demand Mail Relay

On-Demand Mail Relay (**ODMR**) is an SMTP extension standardized in RFC 2645 that allows e-mail to be relayed to an authenticated recipient.

Internationalization

Many users whose native script is not Latin based have had difficulty with the Latin email address requirement. Often this leads to meaningless, but easy to type, locale addresses.

RFC 6531 was created to solve that problem, providing internationalization features for SMTP, the SMTPUTF8 extension. RFC 6531 provides support for multi-byte and non-ASCII characters in email addresses, such as Pelé@live.com (simple diacritic), δοκιμή@παράδειγμα.δοκιμή, and 测 试@测 试.测 试. Current support is limited, but there is strong interest in broad adoption of RFC 6531 and the related RFCs in countries like China that have a large user base where Latin (ASCII) is a foreign script.

Outgoing mail SMTP server

An e-mail client needs to know the IP address of an SMTP server and this has to be given as part of its configuration (usually given as a DNS name). The server will deliver outgoing messages on behalf of the user.

Outgoing mail server access restrictions

Server administrators need to impose some control on which clients can use the server. This enables them to deal with abuse, for example spam. Two solutions have been in common use:

- In the past, many systems imposed usage restrictions by the *location* of the client, only permitting usage by clients whose IP address is one that the server administrators control. Usage from any other client IP address is disallowed.
- Modern SMTP servers typically offer an alternative system that requires authentication of clients by credentials before allowing access.

Restricting access by location

Under this system, an *ISP's* SMTP server will not allow access by users who are 'outside the ISP's network'. More precisely, the server may only allow access to users with an IP address provided by the ISP, which is equivalent to requiring that they are connected to the Internet using that same ISP. A mobile user may often be on a network other than that of their normal ISP, and will then find that sending email fails because the configured SMTP server choice is no longer accessible.

This system has several variations. For example, an organisation's SMTP server may only provide service to users on the same network, enforcing this by firewalling to block access by users on the wider Internet. Or the server may perform range checks on the client's IP address. These methods were typically used by corporations and institutions such as universities which provided an SMTP server for outbound mail only for use internally within the organisation. However, most of these bodies now use client authentication methods, as described below.

By restricting access to certain IP addresses, server administrators can readily recognise the IP address of any abuser. As it will be a meaningful address to them, the administrators can deal with the rogue machine or user.

Where a user is mobile, and may use different ISPs to connect to the internet, this kind of usage restriction is onerous, and altering the configured outbound email SMTP server address is impractical. It is highly desirable to be able to use email client configuration information that does not need to change.

Client authentication

Modern SMTP servers typically require authentication of clients by credentials before allowing access, rather than restricting access by location as described earlier. This more flexible system is friendly to mobile users and allows them to have a fixed choice of configured outbound SMTP server.

Open relay

A server that is accessible on the wider Internet and does not enforce these kinds of access restrictions is known as an open relay. This is now generally considered a bad practice worthy of blacklisting.

Ports

Server administrators choose whether clients use TCP port 25 (SMTP) or port 587 (Submission), as formalized in RFC 6409 (previously RFC 2476), for relaying outbound mail to a mail server. The specifications and many servers support both. Although some servers support port 465 for legacy *secure SMTP* in violation of the specifications, it is preferable to use standard ports and standard ESMTP commands[16] according to RFC 3207 if a secure session needs to be used between the client and the server.

Some servers are set up to reject all relaying on port 25, but valid users authenticating on port 587 are allowed to relay mail to any valid address.

Some Internet service providers intercept port 25, redirecting traffic to their own SMTP server regardless of the destination address. This means that it is not possible for their users to access an SMTP server outside the ISP's network using port 25.

Some SMTP servers support authenticated access on an additional port other than 25 to allow users to connect to them even if port 25 is blocked.

SMTP transport example

A typical example of sending a message via SMTP to two mailboxes (*alice* and *theboss*) located in the same mail domain (*example.com* or *localhost.com*) is reproduced in the following session exchange. (In this example, the conversation parts are prefixed with *S:* and *C:*, for *server* and *client*, respectively; these labels are not part of the exchange.)

After the message sender (SMTP client) establishes a reliable communications channel to the message receiver (SMTP server), the session is opened with a greeting by the server, usually containing its fully qualified domain name (FQDN), in this case *smtp.example.com*. The client initiates its dialog by responding with a `HELO` command identifying itself in the command's parameter with its FQDN (or an address literal if none is available).[2]

```
S: 220 smtp.example.com ESMTP Postfix
C: HELO relay.example.org
S: 250 Hello relay.example.org, I am glad to meet you
C: MAIL FROM:<bob@example.org>
S: 250 Ok
C: RCPT TO:<alice@example.com>
S: 250 Ok
C: RCPT TO:<theboss@example.com>
S: 250 Ok
C: DATA
S: 354 End data with <CR><LF>.<CR><LF>
C: From: "Bob Example" <bob@example.org>
C: To: "Alice Example" <alice@example.com>
C: Cc: theboss@example.com
C: Date: Tue, 15 January 2008 16:02:43 -0500
C: Subject: Test message
C:
C: Hello Alice.
C: This is a test message with 5 header fields and 4 lines in the message body.
C: Your friend,
C: Bob
C: .
S: 250 Ok: queued as 12345
C: QUIT
S: 221 Bye
```

`{The server closes the connection}`

The client notifies the receiver of the originating email address of the message in a `MAIL FROM` command. In this example, the email message is sent to two mailboxes on the same SMTP server: one each for each recipient listed in the `To` and `Cc` header fields. The corresponding SMTP command is `RCPT TO`. Each successful reception and execution of a command is acknowledged by the server with a result code and response message (e.g., 250 Ok).

The transmission of the body of the mail message is initiated with a `DATA` command after which it is transmitted verbatim line by line and is terminated with an end-of-data sequence. This sequence consists of a new-line (<CR><LF>), a single full stop (period), followed by another new-line. Since a message body can contain a line with just a period as part of the text, the client sends *two* periods every time a line starts with a period; correspondingly, the server replaces every sequence of two periods at the beginning of a line with a single one. Such escaping method is called *dot-stuffing*.

The server's positive reply to the end-of-data, as exemplified, implies that the server has taken the responsibility of delivering the message. A message can be doubled if there is a communication failure at this time, e.g. due to a power shortage: Until the sender has received that `250` reply, it must assume the message was not delivered. On the other hand, after the receiver has decided to accept the message, it must assume the message has been delivered to it. Thus, during this time span, both agents have active copies of the message that they will try to deliver.[17] The probability that a communication failure occurs exactly at this step is directly proportional to the amount of filtering that the server performs on the message body, most often for anti-spam purposes. The limiting timeout is specified to be 10 minutes.[18]

The `QUIT` command ends the session. If the second recipient were located elsewhere, the client would `QUIT` and connect to the appropriate SMTP server after the first message had been queued. The information that the client sends in the `HELO` and `MAIL FROM` commands are added (not seen in example code) as additional header fields to the message by the receiving server. It adds a `Received` and `Return-Path` header field, respectively.

Optional extensions

Although optional and not shown in this example, many clients ask the server for the SMTP extensions that the server supports, by using the `EHLO` greeting of the extended SMTP specification (RFC 1870). Clients fall back to `HELO` only if the server does not respond to `EHLO`.

Modern clients may use the ESMTP extension keyword `SIZE` to query the server for the maximum message size that will be accepted. Older clients and servers may try to transfer excessively sized messages that will be rejected after consuming network resources, including connect time to network links that is paid by the minute.

Users can manually determine in advance the maximum size accepted by ESMTP servers. The client replaces the `HELO` command with the `EHLO` command.

```
S: 220 smtp2.example.com ESMTP Postfix
C: EHLO bob.example.org
S: 250-smtp2.example.com Hello bob.example.org [192.0.2.201]
S: 250-SIZE 14680064
S: 250-PIPELINING
S: 250 HELP
```

Thus *smtp2.example.com* declares that it will accept a fixed maximum message size no larger than 14,680,064 octets (8-bit bytes). Depending on the server's actual resource usage, it may be currently unable to accept a message this large. In the simplest case, an ESMTP server will declare a maximum SIZE with only the EHLO user interaction.

Security and spamming

The original SMTP specification did not include a facility for authentication of senders. Subsequently, the SMTP-AUTH extension was defined by RFC 2554.[19] The SMTP extension (ESMTP) provides a mechanism for email clients to specify a security mechanism to a mail server, authenticate the exchange, and negotiate a security profile (Simple Authentication and Security Layer, SASL) for subsequent message transfers.

Microsoft products implement the proprietary Secure Password Authentication (SPA) protocol through the use of the SMTP-AUTH extension.

However, the impracticality of widespread SMTP-AUTH implementation and management means that E-mail spamming is not and cannot be addressed by it.

Modifying SMTP extensively, or replacing it completely, is not believed to be practical, due to the network effects of the huge installed base of SMTP. Internet Mail 2000 was one such proposal for replacement.

Spam is enabled by several factors, including vendors implementing MTAs that are not standards-compliant, which make it difficult for other MTAs to enforce standards, security vulnerabilities within the operating system (often exacerbated by always-on broadband connections) that allow spammers to remotely control end-user PCs and cause them to send spam, and a lack of "intelligence" in many MTAs.

There are a number of proposals for sideband protocols that will assist SMTP operation. The Anti-Spam Research Group (ASRG) of the Internet Research Task Force (IRTF) is working on a number of E-mail authentication and other proposals for providing simple source authentication that is flexible, lightweight, and scalable. Recent Internet Engineering Task Force (IETF) activities include MARID (2004) leading to two approved IETF experiments in 2005, and DomainKeys Identified Mail in 2006.

Related Requests For Comments

- RFC 1123 – Requirements for Internet Hosts—Application and Support (STD 3)
- RFC 1870 – SMTP Service Extension for Message Size Declaration (obsoletes: RFC 1653)
- RFC 2505 – Anti-Spam Recommendations for SMTP MTAs (BCP 30)
- RFC 2920 – SMTP Service Extension for Command Pipelining (STD 60)
- RFC 3030 – SMTP Service Extensions for Transmission of Large and Binary MIME Messages
- RFC 3207 – SMTP Service Extension for Secure SMTP over Transport Layer Security (obsoletes RFC 2487)
- RFC 3461 – SMTP Service Extension for Delivery Status Notifications (obsoletes RFC 1891)
- RFC 3463 – Enhanced Status Codes for SMTP (obsoletes RFC 1893)
- RFC 3464 – An Extensible Message Format for Delivery Status Notifications (obsoletes RFC 1894)
- RFC 3798 - Message Disposition Notification
- RFC 3834 – Recommendations for Automatic Responses to Electronic Mail
- RFC 4952 – Overview and Framework for Internationalized E-mail
- RFC 4954 – SMTP Service Extension for Authentication (obsoletes RFC 2554)
- RFC 5068 – E-mail Submission Operations: Access and Accountability Requirements (BCP 134)
- RFC 5321 – The Simple Mail Transfer Protocol (obsoletes RFC 821 aka STD 10, RFC 974, RFC 1869, RFC 2821)
- RFC 5322 – Internet Message Format (obsoletes RFC 822 aka STD 11, and RFC 2822)
- RFC 5336 - SMTP Extension for Internationalized Email Addresses (updates RFC 2821, RFC 2822, and RFC 4952)
- RFC 5504 - Downgrading Mechanism for Email Address Internationalization
- RFC 6409 – Message Submission for Mail (obsoletes RFC 4409, RFC 2476)
- RFC 6522 – The Multipart/Report Content Type for the Reporting of Mail System Administrative Messages (obsoletes RFC 3462, and in turn RFC 1892)

References

[1] RFC 821, *Simple Mail Transfer Protocol*, J.B. Postel, The Internet Society (August 1982)
[2] RFC 5321, *Simple Mail Transfer Protocol*, J. Klensin, The Internet Society (October 2008)
[3] The History of Electronic Mail (http://www.multicians.org/thvv/mail-history.html), *Tom Van Vleck: "*It is not clear this protocol was ever implemented*"*
[4] *The First Network Email* (http://openmap.bbn.com/~tomlinso/ray/firstemailframe.html), Ray Tomlinson, BBN
[5] Picture of " The First Email Computer (http://openmap.bbn.com/~tomlinso/ray/ka10.html)" by Dan Murphy, a PDP-10
[6] Dan Murphy's TENEX and TOPS-20 Papers (http://www.opost.com/dlm/tenex/)
[7] RFC 2235
[8] RFC 469 - Network Mail Meeting Summary
[9] RFC 524 - A Proposed Mail Protocol
[10] RFC 772 - Mail Transfer Protocol
[11] Tldp.org (http://tldp.org/HOWTO/Usenet-News-HOWTO/x64.html)

[12] draft-barber-uucp-project-conclusion-05 - The Conclusion of the UUCP Mapping Project (http://tools.ietf.org/html/draft-barber-uucp-project-conclusion-05)

[13] Eric Allman (1983), *Sendmail - An Internetwork Mail Router* (http://docs.freebsd.org/44doc/smm/09.sendmail/paper.pdf), BSD UNIX documentation set, Berkeley: University of California, , retrieved 29 June 2012

[14] Craig Partridge (2008), *The Technical Development of Internet Email* (http://www.ir.bbn.com/~craig/email.pdf), IEEE Annals of the History of Computing, IEEE Computer Society, doi:10.1109/MAHC.2008.32,

[15] RFC 1985, *SMTP Service Extension for Remote Message Queue Starting*, J. De Winter, The Internet Society (August 1996)

[16] RFC 3207 specifies only the well-known port 25 and the "Submission port," which is TCP port 587, for the STARTTLS command, the precursor for an encrypted SMTP session using TLS. It makes no mention of the unofficial port 465.

[17] RFC 1047

[18] rfc5321#section-4.5.3.2.6 (http://tools.ietf.org/html/rfc5321#section-4.5.3.2.6)

[19] RFC 2554, *SMTP Service Extension for Authentication*, J. Myers (March 1999)

Further reading

- Hughes, L (1998). *Internet e-mail Protocols, Standards and Implementation*. Artech House Publishers. ISBN 0-89006-939-5.
- Hunt, C (2003). *sendmail Cookbook*. O'Reilly Media. ISBN 0-596-00471-0.
- Johnson, K (2000). *Internet Email Protocols: A Developer's Guide*. Addison-Wesley Professional. ISBN 0-201-43288-9.
- Loshin, P (1999). *Essential Email Standards: RFCs and Protocols Made Practical*. John Wiley & Sons. ISBN 0-471-34597-0.
- Rhoton, J (1999). *Programmer's Guide to Internet Mail: SMTP, POP, IMAP, and LDAP*. Elsevier. ISBN 1-55558-212-5.
- Wood, D (1999). *Programming Internet Mail*. O'Reilly. ISBN 1-56592-479-7.

External links

- Essential Internet Protocols - SMTP (http://www.vanemery.com/Protocols/SMTP/smtp.html)
- SMTP Sequence Diagram (http://www.eventhelix.com/RealtimeMantra/Networking/SMTP_Sequence_Diagram.pdf) (PDF)
- The Case For E-mail Security (http://luxsci.com/extranet/articles/email-security.html) - Security and Insecurity in SMTP, POP and IMAP.
- Picture of the first computers to send and receive a network email, 2 PDP-10s (http://history-computer.com/Internet/Maturing/Tomlinson.html)
- Email Address Internationalization IETF Working Group (http://www.ietf.org/html.charters/eai-charter.html)
- SMTP TLS Transport real-time test (http://www.CheckTLS.com/TestReceiver?LEVEL=3) - Live version of above example, mostly for TLS (i.e. secure) email but useful for non-TLS too

Internet Message Access Protocol

Internet message access protocol (**IMAP**) is one of the two most prevalent Internet standard protocols for e-mail retrieval, the other being the Post Office Protocol (POP).[1] Virtually all modern e-mail clients and mail servers support both protocols as a means of transferring e-mail messages from a server.

E-mail protocols

The Internet Message Access Protocol (commonly known as IMAP) is an Application Layer Internet protocol that allows an e-mail client to access e-mail on a remote mail server. The current version, IMAP version 4 revision 1 (IMAP4rev1), is defined by RFC 3501 [2]. An IMAP server typically listens on well-known port 143. IMAP over SSL **(IMAPS)** is assigned well-known port number 993.

IMAP supports both on-line and off-line modes of operation. E-mail clients using IMAP generally leave messages on the server until the user explicitly deletes them. This and other characteristics of IMAP operation allow multiple clients to manage the same mailbox. Most e-mail *clients* support IMAP in addition to POP to retrieve messages; however, fewer e-mail *services* support IMAP.[3] IMAP offers access to the mail storage. Clients may store local copies of the messages, but these are considered to be a temporary cache.

Incoming e-mail messages are sent to an e-mail server that stores messages in the recipient's e-mail box. The user retrieves the messages with an e-mail client that uses one of a number of e-mail retrieval protocols. Some clients and servers preferentially use vendor-specific, proprietary protocols, but most support the Internet standard protocols, SMTP for sending e-mail and POP and IMAP for retrieving e-mail, allowing interoperability with other servers and clients. For example, Microsoft's Outlook client uses a proprietary protocol to communicate with a Microsoft Exchange Server server as does IBM's Notes client when communicating with a Domino server, but all of these products also support POP, IMAP, and outgoing SMTP. Support for the Internet standard protocols allows many e-mail clients such as Pegasus Mail or Mozilla Thunderbird (see comparison of e-mail clients) to access these servers, and allows the clients to be used with other servers (see list of mail servers).

History

IMAP was designed by Mark Crispin in 1986 as a remote mailbox protocol, in contrast to the widely used POP, a protocol for retrieving the contents of a mailbox.[4]

IMAP was previously known as **Internet Mail Access Protocol**, **Interactive Mail Access Protocol** (RFC 1064), and **Interim Mail Access Protocol**.[5]

Original IMAP

The original *Interim Mail Access Protocol* was implemented as a Xerox Lisp machine client and a TOPS-20 server.

No copies of the original interim protocol specification or its software exist. Although some of its commands and responses were similar to IMAP2, the interim protocol lacked command/response tagging and thus its syntax was incompatible with all other versions of IMAP.

IMAP2

The interim protocol was quickly replaced by the *Interactive Mail Access Protocol* (IMAP2), defined in RFC 1064 (in 1988) and later updated by RFC 1176 (in 1990). IMAP2 introduced command/response tagging and was the first publicly distributed version.

IMAP3

IMAP3 is an extinct and extremely rare variant of IMAP.[6] It was published as RFC 1203 in 1991. It was written specifically as a counter proposal to RFC 1176, which itself proposed modifications to IMAP2.[7] IMAP3 was never accepted by the marketplace.[8][9] The IESG reclassified RFC1203 "Interactive Mail Access Protocol - Version 3" as a Historic protocol in 1993. The IMAP Working Group used RFC1176 (IMAP2) rather than RFC1203 (IMAP3) as its starting point.[10][11]

IMAP2bis

With the advent of MIME, IMAP2 was extended to support MIME body structures and add mailbox management functionality (create, delete, rename, message upload) that was absent in IMAP2. This experimental revision was called IMAP2bis; its specification was never published in non-draft form. An internet draft of IMAP2bis was published by the IETF IMAP Working Group in October 1993. This draft was based upon the following earlier specifications: unpublished *IMAP2bis.TXT* document, RFC1176, and RFC1064 (IMAP2).[12] The *IMAP2bis.TXT* draft documented the state of extensions to IMAP2 as of December 1992.[13] Early versions of Pine were widely distributed with IMAP2bis support[6] (Pine 4.00 and later supports IMAP4rev1).

IMAP4

An IMAP Working Group formed in the IETF in the early 1990s took over responsibility for the IMAP2bis design. The IMAP WG decided to rename IMAP2bis to IMAP4 to avoid confusion with a competing IMAP3 proposal from another group that never got off the ground. The expansion of the IMAP acronym also changed to the *Internet Message Access Protocol*

Advantages over POP

Connected and disconnected modes of operation

When using POP, clients typically connect to the e-mail server briefly, only as long as it takes to download new messages. When using IMAP4, clients often stay connected as long as the user interface is active and download message content on demand. For users with many or large messages, this IMAP4 usage pattern can result in faster response times.

Multiple clients simultaneously connected to the same mailbox

The POP protocol requires the currently connected client to be the only client connected to the mailbox. In contrast, the IMAP protocol specifically allows simultaneous access by multiple clients and provides mechanisms for clients to detect changes made to the mailbox by other, concurrently connected, clients. See for example RFC3501 section 5.2 which specifically cites "simultaneous access to the same mailbox by multiple agents" as an example.

Access to MIME message parts and partial fetch

Usually all Internet e-mail is transmitted in MIME format, allowing messages to have a tree structure where the leaf nodes are any of a variety of single part content types and the non-leaf nodes are any of a variety of multipart types. The IMAP4 protocol allows clients to separately retrieve any of the individual MIME parts and also to retrieve portions of either individual parts or the entire message. These mechanisms allow clients to retrieve the text portion of a message without retrieving attached files or to stream content as it is being fetched.

Message state information

Through the use of flags defined in the IMAP4 protocol, clients can keep track of message state; for example, whether or not the message has been read, replied to, or deleted. These flags are stored on the server, so different clients accessing the same mailbox at different times can detect state changes made by other clients. POP provides no mechanism for clients to store such state information on the server so if a single user accesses a mailbox with two different POP clients (at different times), state information—such as whether a message has been accessed—cannot be synchronized between the clients. The IMAP4 protocol supports both pre-defined system flags and client defined keywords. System flags indicate state information such as whether a message has been read. Keywords, which are not supported by all IMAP servers, allow messages to be given one or more tags whose meaning is up to the client. Adding user created tags to messages is an operation supported by some web-based e-mail services, such as Gmail.

Multiple mailboxes on the server

IMAP4 clients can create, rename, and/or delete mailboxes (usually presented to the user as folders) on the server, and copy messages between mailboxes. Multiple mailbox support also allows servers to provide access to shared and public folders. The *IMAP4 Access Control List (ACL) Extension* (RFC 4314) may be used to regulate access rights.

Server-side searches

IMAP4 provides a mechanism for a client to ask the server to search for messages meeting a variety of criteria. This mechanism avoids requiring clients to download every message in the mailbox in order to perform these searches.

Built-in extension mechanism

Reflecting the experience of earlier Internet protocols, IMAP4 defines an explicit mechanism by which it may be extended. Many extensions to the base protocol have been proposed and are in common use. IMAP2bis did not have an extension mechanism, and POP now has one defined by RFC 2449.

Disadvantages

While IMAP remedies many of the shortcomings of POP, this inherently introduces additional complexity. Much of this complexity (e.g., multiple clients accessing the same mailbox at the same time) is compensated for by server-side workarounds such as Maildir or database backends.

The IMAP specification has been criticised for being insufficiently strict and allowing behaviours that effectively negate its usefulness. For instance, the specification states that each message stored on the server has a "unique id" to allow the clients to identify the messages they have already seen between sessions. However, the specification also allows these UIDs to be invalidated with no restrictions, practically defeating their purpose.[14]

Unless the mail storage and searching algorithms on the server are carefully implemented, a client can potentially consume large amounts of server resources when searching massive mailboxes.

IMAP4 clients need to maintain a TCP/IP connection to the IMAP server in order to be notified of the arrival of new mail. Notification of mail arrival is done through in-band signaling, which contributes to the complexity of client-side IMAP protocol handling somewhat.[15] A private proposal, push IMAP, would extend IMAP to implement push e-mail by sending the entire message instead of just a notification. However, push IMAP has not been generally accepted and current IETF work has addressed the problem in other ways (see the Lemonade Profile for more information).

Unlike some proprietary protocols which combine sending and retrieval operations, sending a message and saving a copy in a server-side folder with a base-level IMAP client requires transmitting the message content twice, once to SMTP for delivery and a second time to IMAP to store in a sent mail folder. This is remedied by a set of extensions defined by the IETF LEMONADE Working Group for mobile devices: URLAUTH (RFC 4467) and CATENATE (RFC 4469) in IMAP and BURL (RFC 4468) in SMTP-SUBMISSION. POP servers don't support server-side folders so clients have no choice but to store sent items on the client. Many IMAP clients can be configured to store sent mail in a client-side folder, or to BCC oneself and then filter the incoming mail instead of saving a copy in a folder directly. In addition to the LEMONADE "trio", Courier Mail Server offers a non-standard method of sending using IMAP by copying an outgoing message to a dedicated outbox folder.

Like POP, IMAP is an e-mail only protocol. As a result, items such as contacts, appointments or tasks cannot be managed or accessed using IMAP.

References

[1] Pegoraro, Rob (2004-03-24). "Internet Providers Should Find Their Way to IMAP" (http://www.washingtonpost.com/wp-dyn/articles/A10089-2004Mar20.html). Washington Post. . Retrieved 2008-06-25.

[2] http://tools.ietf.org/html/rfc3501

[3] Mullet, Diana (2000). *Managing IMAP*. O'Reilly. p. 25. ISBN 0-596-00012-X.

[4] The IMAP Connection - IMAP Status and History (http://www.imap.org/about/history.status.html)

[5] http://www.iana.org/assignments/service-names

[6] "RFC 2061 - IMAP4 COMPATIBILITY WITH IMAP2BIS" (http://tools.ietf.org/html/rfc2061). IETF. 1996. . Retrieved 2010-08-21.

[7] "INTERACTIVE MAIL ACCESS PROTOCOL - VERSION 3" (http://tools.ietf.org/html/rfc1203). IETF. 1991. . Retrieved 2010-08-21.

[8] "IMAP2, IMAP2bis, IMAP3, IMAP4, IMAP4rev1 (LAN Mail Protocols)" (http://stason.org/TULARC/networking/lans-mail-protocols/03-IMAP2-IMAP2bis-IMAP3-IMAP4-IMAP4rev1-LAN-Mail-Protoc.html). . Retrieved 2010-08-21.

[9] "IMAP Overview, History, Versions and Standards" (http://www.tcpipguide.com/free/t_IMAPOverviewHistoryVersionsandStandards-3.htm). . Retrieved 2010-08-21.

[10] "Protocol Action: Interactive Mail Access Protocol - Version 3 to Historic (IETF mail archive)" (http://www.ietf.org/mail-archive/web/ietf/current/msg01656.html). 1993. . Retrieved 2010-08-21.

[11] "Innosoft and POP/IMAP protocols? (mail archive)" (http://www.pmdf.process.com/ftp/info-pmdf/aug.1993?httpd=content&type=text/plain; charset=ISO-8859-1). 1993. . Retrieved 2010-08-21.

[12] "INTERACTIVE MAIL ACCESS PROTOCOL - VERSION 2bis (Internet Draft)" (http://tools.ietf.org/html/draft-ietf-imap-imap2bis-02). IETF. 1993. . Retrieved 2010-08-21.

[13] "IMAP2BIS -- EXTENSIONS TO THE IMAP2 PROTOCOL (DRAFT)" (http://ftp.zcu.cz/pub/network/imap/old/IMAP2bis.TXT). 1992. . Retrieved 2010-08-21.

[14] "IMAP implementation in Sup, an e-mail client written in Ruby" (http://sup.rubyforge.org/svn/trunk/lib/sup/imap.rb). rubyforge.com. . Retrieved 2011-02-22.

[15] "IMAP IDLE: The best approach for 'push' e-mail" (http://www.isode.com/whitepapers/imap-idle.html). Isode.com. . Retrieved 2009-07-30.

Further reading

- Heinlein, P; Hartleben, P (2008). *The Book of IMAP: Building a Mail Server with Courier and Cyrus*. No Starch Press. ISBN 1-59327-177-8.
- Hughes, L (1998). *Internet e-mail Protocols, Standards and Implementation*. Artech House Publishers. ISBN 0-89006-939-5.
- Johnson, K (2000). *Internet E-mail Protocols: A Developer's Guide*. Addison-Wesley Professional. ISBN 0-201-43288-9.
- Loshin, P (1999). *Essential E-mail Standards: RFCs and Protocols Made Practical*. John Wiley & Sons. ISBN 0-471-34597-0.
- Mullet, K (2000). *Managing IMAP*. O'Reilly Media. ISBN 0-596-00012-X.
- Rhoton, J (1999). *Programmer's Guide to Internet Mail: SMTP, POP, IMAP, and LDAP*. Elsevier. ISBN 1-55558-212-5.
- Wood, D (1999). *Programming Internet Mail*. O'Reilly. ISBN 1-56592-479-7.

External links

- "IMAP Protocol Mailing List" (http://www.imapwiki.org/ImapProtocolList).
- RFC 3501 - specification of IMAP version 4 revision 1
- RFC 2683 - IMAP Implementation Suggestions RFC
- RFC 2177 - IMAP4 IDLE command

Electronic mailing list

An **electronic mailing list** is a special usage of email that allows for widespread distribution of information to many Internet users. It is similar to a traditional mailing list — a list of names and addresses — as might be kept by an organization for sending publications to its members or customers, but typically refers to four things — a list of email addresses, the people ("subscribers") receiving mail at those addresses, the publications (email messages) sent to those addresses, and a *reflector*, which is a single email address that, when designated as the recipient of a message, will send a copy of that message to all of the subscribers.

How automated electronic mailing lists work

Electronic mailing lists are usually fully or partially automated through the use of special mailing list software and a reflector address that are set up on a server capable of receiving email. Incoming messages sent to the reflector address are processed by the software, and, depending on their content, are acted upon internally (in the case of messages containing commands directed at the software itself) or are distributed to all email addresses subscribed to the mailing list. Depending on the software, additional addresses may be set up for the purpose of sending commands.

Many electronic mailing list servers have a special email address in which subscribers (or those that want to be subscribers) can send commands to the server to perform such tasks as subscribing and unsubscribing, temporarily halting the sending of messages to them, or changing available preferences. The common format for sending these commands is to send an email that contains simply the command followed by the name of the electronic mailing list the command pertains to. Examples: *subscribe anylist* or *subscribe anylist John Doe*. Some list servers also allow people to subscribe, unsubscribe, change preferences, etc. via a website.

Electronic mailing list servers can be set to forward messages to subscribers of a particular mailing list either individually as they are received by the list server or in digest form in which all messages received on a particular day by the list server are combined into one email that is sent once per day to subscribers. Some mailing lists allow individual subscribers to decide how they prefer to receive messages from the list server (individual or digest).

Types

Announcement list

One type of electronic mailing list is an *announcement list*, which is used primarily as a one-way conduit of information and can only be "posted to" by selected people. This may also be referred to by the term *newsletter*. Newsletter and promotional emailing lists are employed in various sectors as parts of direct marketing campaigns.

Discussion list

Another type of electronic mailing list is a *discussion list*, in which any subscriber may post. On a discussion list, a subscriber uses the mailing list to send messages to all the other subscribers, who may answer in similar fashion. Thus, actual discussion and information exchanges can happen. Mailing lists of this type are usually topic-oriented (for example, politics, scientific discussion, joke contests), and the topic can range from extremely narrow to "whatever you think could interest us". In this they are similar to Usenet newsgroups, and share the same aversion to off-topic messages. The term discussion group encompasses both these types of lists and newsgroups.

List security

On both discussion lists and newsletter lists precautions are taken to avoid spamming.

Discussion lists often require every message to be approved by a moderator before being sent to the rest of the subscribers, although higher traffic lists typically only moderate messages from new subscribers, and only then for a time. Companies sending out promotional newsletters have the option of working with whitelist mail distributors, which agree to standards and high fines from ISPs should any of the opt-in subscribers complain. In exchange for their compliance and agreement to prohibitive fines, the emails sent by whitelisted companies are not blocked by spam filters, which can often reroute these legitimate, non-spam emails.[1]

Subscription

Some mailing lists are open to anyone who wants to join them, while others require an approval from the list owner before one can join. Joining a mailing list is called "subscribing" and leaving a list is called "unsubscribing".

Archives

A mailing list archive is a collection of past messages from one or more electronic mailing lists. Such archives often include searching and indexing functionality. Many archives are directly associated with the mailing list, but some organizations like Gmane collect archives from multiple mailing lists hosted at different organizations; thus, one message sent to one popular mailing list can end up in many different archives. Gmane had archives of over 9000 mailing lists as of 16 January 2007. Some popular free software programs for collecting mailing list archives are Hypermail, MHonArc and FUDforum.

References

[1] "What is a "Whitelist" and why do I want to work with a "Whitelisted" Mail Distributor?" (http://www.bethesda-list.com/lists/what-is-a-whitelist). .

External links

- Discussion_groups article at LISWiki, a Library science wiki
- Mailing List Management Packages (http://seaotter.berkeley.edu/calmail/maillists/maillist-managers-matrix.html)

Comparison of mail servers

This is a **comparison of mail servers**: mail transfer agents, mail delivery agents, and other computer software which provide e-mail services.

Feature comparison

	Server OS support			Features											Storage			
Mail server	**Linux/Unix**	**Windows**	**Mac OS**	**SMTP**	**POP3**	**IMAP**	**IMAP IDLE**	**SMTP over TLS**	**POP over TLS**	**IPv6**	**NNTP**	**SSL**	**Webmail**	**ActiveSync**	**Database**	**Filesystem**	**Other**	**License**
agorum core	Yes	Yes	No	Yes	No	Yes	?	No	No	?	No	Yes	Yes	No	Yes	No	No	Open Source/GPLv2
Apache James	Yes	Yes	Yes	Yes	Yes	Yes	?	Yes	Yes	?	Yes	Yes	No	No	Yes	Yes	No	Open Source/ASLv2
Atmail	Yes	No	Yes	Yes	Yes	Yes	Yes	?	Yes	?	No	Yes	Yes	Yes	Yes	Yes	No	Proprietary (Free 5 user)
Axigen	Yes	Yes	No	Yes	Yes	Yes	?	Yes	Yes	No	No	Yes	Ajax & Basic	Yes	No	No	Yes	Proprietary (free version: email 100 users / calendar 5 users)
Citadel	Yes	No	Yes	Yes	Yes	Yes	?	Yes	Yes	?	?	Yes	Yes	No	Yes	No	No	Open Source/GPLv3
CommuniGate Pro	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	No	Yes	Yes	Yes	Yes	Yes	Yes	Proprietary
Courier Mail Server	Yes	No	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	?	Yes	Yes	No	No	maildir	No	Open Source
Cyrus IMAP	Yes	No	Yes	No	Yes	Yes	Yes	No	Yes	Yes	Yes	Yes	No	No	Yes (indexes)	Yes	Yes	Open Source/BSD
DBMail	Yes	via Cygwin	Yes	Postfix, Exim, Sendmail, qmail	Yes	Yes	Yes	Postfix, Exim, Sendmail, qmail	Yes	Yes	?	Yes	No	No	PostgreSQL, MySQL, SQLite	No	No	GPL
Digital Integration iMail5	Yes	Yes	No	Yes	Yes	Yes	?	Yes	Yes	?	No	Yes	Yes	No	Yes	Yes	Yes	Proprietary
Dovecot	Yes[1]	No[1]	Yes[1]	No	Yes	Yes	Yes	No	Yes	Yes	No	Yes	No	No	Yes	maildir, mbox, dbox	No	Open Source/Mixed: MIT and LGPL 2.1[2]
Eudora Internet Mail Server	No	No	Yes	Yes	Yes	Yes	?	No	No	?	No	No	No	No	No	No	Yes	Proprietary

Exim	Yes	via Cygwin	Yes	Yes	No	No	No	Yes	No	Yes	No	Yes	No	No	Yes	Yes	No	Open Source/GPL
FirstClass	Yes	Yes	Yes	Yes	Yes	Yes	?	Yes	Yes	?	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Proprietary
Gordano Messaging Suite	Yes	Yes	No	Yes	Yes	Yes	?	Yes	Yes	?	No	Yes	Yes	No	Yes	Yes	Yes	Proprietary
GroupWise	Yes	Yes	No	Yes	Yes	Yes	?	Yes	Yes	?	?	Yes	Yes	Yes	Yes	No	No	Proprietary
Hexamail Server	Yes	Yes	No	Yes	Yes	Yes	?	Yes	Yes	?	No	Yes	Yes	No	No	Yes	No	Proprietary
hMailServer	No	Yes	No	Yes	Yes	Yes	Yes	V 5.0	V 5.0	Yes	No	V 5.0	Yes	No	Yes	No	No	GNU GPL (V5.0 Proprietary)
IBM Lotus Domino	Yes	Yes	No	Yes	Yes	Yes	?	Yes	Yes	?	Yes	Yes	Yes	Yes	Yes	No	No	Proprietary
IndiMail	Yes	No	Yes	Yes	Yes	Yes	?	Yes	Yes	Yes	No	Yes	No	No	Yes	Yes	No	Open Source/GPLv3
Ipswitch IMail Server	No	Yes	No	Yes	Yes	Yes	?	Yes	Yes	?	No	Yes	Yes	Yes	No	Yes	No	Proprietary
JMail	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	No	Yes	Yes	Yes	Yes	Yes	Yes	Proprietary
Kerio Connect	Yes	Yes	Yes	Yes	Yes	Yes	?	Yes	Yes	?	Yes	Yes	Yes	Yes	Yes	Yes	No	Proprietary
Kolab	Yes	No	No	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Free Software/GPLv3
MagicMail	Yes	No	No	Yes	Yes	Yes	?	Yes	Yes	?	No	Yes	Yes	No	Yes	Yes	Yes	Proprietary
MailSite Fusion	No	Yes	No	Yes	Yes	Yes	?	Yes	Yes	?	No	Yes	Yes	Yes	Yes	Yes	Yes	Proprietary
Mailtraq	No	Yes	No	Yes	Yes	Yes	Yes	Yes	Yes	?	Yes	Yes	Yes	No	Yes	Yes	Yes	Proprietary
MDaemon Mail Server	No	Yes	No	Yes	Yes	Yes	Yes	Yes	Yes	No	No	Yes	Yes	Yes [Requires one time fee] [3]	Yes [Requires one time fee]	Yes	No	Proprietary
Meldware Mail Server	Yes	Yes	Yes	Yes	Yes	Yes	?	Yes	Yes	?	?	Yes	Yes	No	Yes	Yes	No	Open source/LGPL
Mercury Mail Transport System	No	Yes	No	Yes	Yes	Yes	?	Yes	Yes	?	No	Yes	Yes	No	No	Yes	No	Proprietary/donationware
Microsoft Exchange Server	No	Yes	No	Yes	Yes	Yes	Yes [4][5]	Yes	Yes	Yes (2007 sp1 onwards) [6]	Yes	Yes	Yes	Yes	ESE only	Yes (up to 2003 only) [7]	Yes	Proprietary EULA
Mirapoint	Yes	No	No	Yes	Yes	Yes	?	Yes	Yes	?	No	Yes	Yes	No	Yes	Yes	Yes	Proprietary
NetMail	Yes	Yes	No	Yes	Yes	Yes	?	Yes	Yes	?	?	Yes	Yes	No	Yes	No	No	Proprietary
Open-Xchange	Yes	No	No	Yes	Yes	Yes	?	Yes	Yes	?	?	Yes	Yes	Yes	Yes	No	No	Dual license (GPL)
Oracle Communications Messaging Server	Yes	No	No	Yes	Yes	Yes	?	Yes	Yes	?	No	Yes	Yes	No	Yes	Yes	No	Proprietary
Postfix	Yes	No	Yes	Yes	No	No	No	Yes	No	Yes	No	Yes	No	No	Yes	Yes	Yes	Open source/IBM Public License
qmail	Yes	No	Yes	Yes	Yes	No	No	No	No	?	?	No	No	No	No	Yes	No	Public domain

Qpopper	Yes	via Cygwin	Yes	No	Yes	No	No	No	Yes	?	No	Yes	No	No	No	Yes	Yes	BSD-style
Sendmail	Yes	No	Yes	Yes	No	No	No	Yes	No	Yes	No	Yes	No	No	?	Yes	?	Open source/Sendmail License
SparkEngine	Yes	Yes	Yes	Yes	No	No	No	Yes	No	?	No	Yes	No	No	Yes	Yes	Yes	Proprietary
SurgeMail	Yes	Yes	Yes	Yes	Yes	Yes	?	Yes	Yes	?	No	Yes	Yes	No	Yes	Yes	Yes	Proprietary (Free 5 user)
Synovel Collabsuite	Yes	No	No	Yes	Yes	Yes	?	Yes	Yes	?	?	Yes	Yes	No	Yes	Yes	?	Open Source/Proprietary
UW IMAP	Yes	Yes	Yes	No	Yes	Yes	Yes	No	Yes	?	No	Yes	WebPine	No	Yes (Indexes)	Yes	No	Open Source/Apache license
WinGate	No	Yes	No	Yes	Yes	Yes	Yes	Yes	Yes	No	No	Yes	Yes	No	No	Yes	No	Proprietary
Xeams	Yes	Yes	Yes	Yes	Yes	Yes	?	Yes	Yes	?	No	Yes	No	No	No	Yes	No	Proprietary/Free with support option
Zarafa	Yes	No	No	External [8]	Yes	Yes	Yes	External [8]	Yes	No	No	Yes	Yes	Yes	Yes	Yes	No	Open Source/Proprietary
Zimbra	Yes	No	Yes	Yes	Yes	Yes	Yes	Yes	Yes	No	No	Yes	Yes	Yes	Yes	Yes	No	Open Source/Proprietary
Mail server	**Linux/Unix**	**Windows**	**Mac OS X**	**SMTP**	**POP3**	**IMAP**	**IMAP IDLE**	**SMTP over TLS**	**POP over TLS**	**IPv6**	**NNTP**	**SSL**	**Webmail**	**ActiveSync**	**Database**	**Filesystem**	**Other**	**License**
	Server OS support			**Features**										**Storage**				

Authentication

Mail server	SMTP AUTH	POP before SMTP	APOP	Filesystem	Database	LDAP	Other
agorum core	Yes	No	No	No	Yes	Yes	PLAIN, LOGIN, CRAM-MD5, NTLM
Apache James	Yes	?	?	?	Yes	Yes	PLAIN, LOGIN
Atmail	Yes	Yes	Yes	Yes	Yes	Yes	PLAIN, LOGIN, CRAM-MD5, DIGEST-MD5, Active Directory, any LDAP-compatible source
Axigen	Yes	No	Yes	No	No	Yes	GSSAPI, CRAM-MD5, DIGEST-MD5, APOP, CLRTXT, LOGIN, PLAIN, LDAP, Internal
CommuniGate Pro	Yes	Yes	Yes	Yes	Yes	Yes	CLRTXT, CRAM-MD5, DIGEST-MD5, APOP, GSSAPI, NTLM, MSN, SESSIONID, External Authentication Plug-in
Courier Mail Server	Yes	Yes	Yes	Yes	Yes	Yes	Managed by Courier authentication library [9] which can use PAM, Userdb, PostgreSQL(beta), MySQL, LDAP, and EXTERNAL.
Cyrus IMAP	No	Yes	Yes	Yes	Yes	Yes	Supports all Cyrus SASL authentication methods. Also supports X.509 PKI auth via STARTTLS and EXTERNAL.

Cyrus SASL Library	No	Yes	Yes	Yes	Yes	Yes	Berkeley DB, GDBM, or NDBM (sasldb), PAM, MySQL, PostgreSQL, SQLite, LDAP, Active Directory(LDAP), DCE, Kerberos 4 and 5, proxied IMAP auth, getpwent, shadow, SIA, Courier Authdaemon, httpform, APOP and SASL mechanisms: ANONYMOUS, CRAM-MD5, DIGEST-MD5, EXTERNAL, GSSAPI, LOGIN, NTLM, OTP, PASSDSS, PLAIN, SRP
Digital Integration iMail5	Yes	No	Yes	No	No	No	
Dovecot	No	Yes	Yes	Yes	Yes	Yes	Not an SMTP server, but offers SASL to MTAs. [10] [11] POP-before-SMTP via DRAC plugin. PAM, MySQL, PostgreSQL, SQLite, LDAP, Active Directory(LDAP), Kerberos 5, proxied IMAP auth, getpwent, shadow, SIA, BSDauth, Vpopmail. APOP and SASL mechanisms: ANONYMOUS, PLAIN, LOGIN, CRAM-MD5, DIGEST-MD5, SCRAM-SHA1, EXTERNAL, GSSAPI, NTLM, OTP, SKEY, RPA.
Eudora Internet Mail Server	Yes	No	Yes	Yes	No	No	CRAM-MD5, DIGEST-MD5, PLAIN, LOGIN
Exim	Yes	Yes	Yes	Yes	Yes	Yes	Cyrus SASL, Dovecot SASL, CRAM-MD5, PLAIN, LOGIN, SPA
FirstClass	Yes	?	?	Yes	Yes	Yes	
Gordano Messaging Suite	Yes	Yes	Yes	No	Yes	Yes	PLAIN, LOGIN, CRAM-MD5, DIGEST-MD5, APOP, ODBC, Active Directory, NT Domain
GroupWise	Yes	Yes	Yes	Yes	Yes	Yes	eDirectory, Any LDAPv3-compliant source
Hexamail Server	Yes	Yes	Yes	Yes	No	Yes	CLRTXT, CRAM-MD5, DIGEST-MD5, APOP, GSSAPI, NTLM, MSN, SESSIONID
IndiMail	Yes	Yes	No	Yes	Yes	Yes	CLRTXT, PLAIN, LOGIN, CRAM-MD5, CRAM-SHA1, CRAM-RIPEMD, DIGEST-MD5, APOP, Checkpassword, authindi, authldap, Pluggable Authentication Modules, Name Service Switch
Ipswitch IMail Server	Yes	No	Yes	Yes	Yes	Yes	Active Directory, Windows authentication, CRAM-MD5, PLAIN, LOGIN
JMail	Yes	No	Yes	No	No	No	CRAM-MD5, PLAIN, LOGIN, APOP, HTTP DIGEST. Authentication through a JID system
Kolab	Yes	Yes	Yes	Yes	Yes	Yes	Kolab employs Cyrus and Postfix and thus supports all Cyrus SASL authentication methods and X.509 PKI auth via STARTTLS and EXTERNAL.
Kerio Connect	Yes	Yes	Yes	Yes	No	Yes	Active Directory, Apple Open Directory, ActiveSync, Pluggable Authentication Modules
MagicMail	Yes	No	No	Yes	Yes	No	?
MailSite Fusion	Yes	No	Yes	Yes	Yes	Yes	NT Domain, Active Directory, PLAIN, LOGIN, NTLM, SCRAM-MD5, CRAM-MD5
Mailtraq	Yes	Yes	Yes	Yes	Yes	Yes	NT Domain, Active Directory, Local Directory
MDaemon Mail Server	Yes	Yes	Yes	No	Yes	Yes	Active Directory, PLAIN, CRAM-MD5, LOGIN
Meldware Mail Server	Yes	No	Yes	Yes	Yes	Yes	?

Mercury Mail Transport System	Yes	Yes	Yes	Yes	No	?	Internal, LOGIN, PLAIN, CRAM-MD5, Netware, LDAP?, partial NT Domain
Microsoft Exchange Server	Yes	?	Yes	No	No	Yes	Active Directory, PLAIN, NTLM
Mirapoint	Yes	Yes	Yes	Yes	Yes	Yes	Plaintext, LDAP, NTLM, APOP, Kerberos_v4, Kerberos_v5, Radius, Sessionid
NetMail	Yes	Yes	Yes	Yes	Yes	Yes	eDirectory, LDAP
Oracle Communications Messaging Server	Yes	Yes	Yes	Yes	Yes	Yes	Plain, Login, CRAM-MD5, Digest-MD5, Certificate
Postfix	Yes	No	No	Yes	Yes	Yes	Supports all Cyrus SASL authentication methods except for APOP.
Qpopper	No	Yes	Yes	Yes	Yes	Yes	Pluggable Authentication Modules,GSSAPI, CRAM-MD5, DIGEST-MD5, APOP, PLAIN, LDAP, CMU SASL library [12], Kerberos
qmail	?	?	?	?	?	?	?
Sendmail	Yes	Yes	No	Yes	Yes	Yes	Supports all Cyrus SASL authentication methods except for APOP. Also supports X.509 PKI auth via STARTTLS and EXTERNAL.
SurgeMail	Yes	Yes	Yes	Yes	Yes	Yes	Internal, LDAP, PAM, MYSQL, ODBC, Unix Password, Windows Domain or external protocol
WinGate	Yes	Yes	Yes	Yes	Yes	No	NT Domain, Active Directory, CRAM-MD5, SASL PLAIN, SASL LOGIN
Zarafa	Yes	Yes	?	Yes	Yes	Yes	Internal, Active Directory, any LDAP-compatible source, Unix
Zimbra	Yes	No	Yes	Yes	Yes	Yes	Internal, LDAP, Open LDAP, Active Directory
Mail server	**SMTP AUTH**	**POP before SMTP**	**APOP**	**Filesystem**	**Database**	**LDAP**	**Other**

Antispam Features

Mail Server	DNSBL	SURBL	Spamtraps	Greylisting	SPF	Tarpit	Bayesian filters	Regular expressions	Embedded Antivirus	Embedded Antispam
agorum core	No	No	No	No	No	No	No	Yes	No	No
Apache James	?	?	?	?	?	?	Yes	?	?	?
Atmail	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes ClamAV	Yes SpamAssassin and Exim ACLs
Axigen	Yes	No	No	Yes	Yes	No	Yes, in SpamAssassin	Yes	Yes Commtouch	Yes SpamAssassin
Citadel	Yes	Yes	No	No	Optional (via SpamAssassin)	No	Optional (via SpamAssassin)	No	Optional (ClamAV)	Optional (SpamAssassin)
CommuniGate Pro	Yes	Yes	Yes	No	Yes	Yes	Optional	Yes	Optional	Optional

Courier Mail Server	Yes	?	Yes	Yes[13]	Yes	Yes	Yes[13]	Yes[13]	Yes Clamav	Yes[13]
Digital Integration iMail5	Yes	No	No	Yes	No	No	No	Optional (GWAVA)	Optional (GWAVA)	Proprietary, Optional (GWAVA)
Dovecot	?	?	?	?	?	?	?	?	?	?
Eudora Internet Mail Server	Yes	?	Yes	Yes	?	?	?	?	?	?
Exim	Yes	Yes	Yes	Yes	Yes	Yes	Optional with spamassassin etc.	Yes	Optional any	Optional any
FirstClass	Yes	Yes	No	?	Yes	Yes	Yes	Yes	Yes Symantec	?
Gordano Messaging Suite	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes Commtouch, Authentium	Yes Commtouch
GroupWise	Yes	?	?	?	?	?	?	?	?	?
Hexamail Server	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes BitDefender	Yes
hMailServer	Yes	Yes	No	Yes	Yes	No	No	Yes	Yes ClamAV	Yes SpamAssassin
IBM Lotus Domino	Yes[14]	?	?	?	No[15]	?	?	?	?	?
IndiMail	Yes	Yes	No	Yes	Yes	Yes	Yes	Yes	Yes ClamAV	Yes Bogofilter
Ipswitch IMail Server	Yes	Yes	No	No	Yes	No	Yes	Yes	Yes Symantec, BitDefender	Yes Commtouch
JMail	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes ClamAV and/or optional any	Yes and/or built-in
Kolab	Yes	?	?	Yes	Yes	?	Optional with spamassassin etc.	Yes	Optional (via content filter[16] interface)	Optional (via content filter[16] interface)
Kerio Connect	Yes	Yes	?	No	Yes	Yes	Yes	?	Optional Sophos + optional secondary (dual)	Yes
MagicMail	Yes	No	Yes	No	?	?	No	Yes	Yes	Yes
MailSite Fusion	Yes	Yes	Yes	Yes	Yes	No	Yes	Yes	Yes	Yes
Mailtraq	?	?	?	?	?	?	?	?	?	?
MDaemon Mail Server	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Optional Kaspersky Labs	Yes

Meldware Mail Server	?	?	?	?	?	?	?	?	?	?
Mercury Mail Transport System	Yes	?	?	Yes (GreyWall extension)	?	?	Yes (SpamHalter extension)	?	Yes (ClamWall extension)	?
Microsoft Exchange Server	Yes (2003 & later)	?	?	Yes	Yes	Yes (2003 & later)	?	?	?	?
Mirapoint	Yes	No	Yes	Yes	Yes	Yes	Yes	Yes	Yes Commtouch, Sophos, F-Secure	Yes Commtouch, Sophos, SPamAssassin
NetMail	?	?	?	?	?	?	?	?	?	?
Oracle Communications Messaging Server	Yes	Yes	Yes	Yes	Yes	Yes	Optional with spamassassin etc.	Yes	Optional any	Optional any
Postfix	Yes	?	?	Yes	Yes	?	Optional with spamassassin etc.	Yes	Optional (via content filter[16] interface)	Optional (via content filter[16] interface)
qmail	?	?	?	Yes	?	?	No	?	No	No
Sendmail	?	?	?	Yes	Optional	?	?	?	Yes ClamAV	Yes SpamAssassin
SparkEngine	?	?	?	?	?	?	?	?	?	?
SurgeMail	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Synovel Collabsuite	Yes	Yes	Yes	Yes	Yes	?	Yes	Yes	YesClamAV	Yes SpamAssassin
UW IMAP	?	?	?	?	?	?	?	?	?	?
WinGate	?	?	?	?	?	?	?	?	?	?
Zarafa	No[8]	No[8]	No[8]	No[8]	No[8]	No[8]	No[8]	No[8]	No	No
Zimbra	Yes	?	?	?	?	?	?	?	Yes ClamAV	Yes SpamAssassin
Mail Server	**DNSBL**	**SURBL**	**Spamtraps**	**Greylisting**	**SPF**	**Tarpit**	**Bayesian filters**	**Regular expressions**	**Embedded Antivirus**	**Embedded Antispam**

Notes and references

[1] Dovecot Operating System Compatibility (http://wiki.dovecot.org/OSCompatibility). Retrieved on 2008-09-03.
[2] Dovecot licensing information: COPYING (http://www.dovecot.org/doc/COPYING). Retrieved on 2008-09-03.
[3] ActiveSync for MDaemon (http://www.altn.com/Purchase/ActiveSync/)
[4] http://technet.microsoft.com/en-us/library/cc540463(EXCHG.80).aspx
[5] http://technet.microsoft.com/en-us/library/ff848256.aspx
[6] http://technet.microsoft.com/en-us/library/bb629624(EXCHG.80).aspx
[7] Exchange Installable File System (ExIFS) (http://msdn.microsoft.com/en-us/library/ms876378(EXCHG.65).aspx)
[8] Zarafa Collaboration Platform architecture diagram (http://doc.zarafa.com/trunk/Administrator_Manual/en-US/html/_architecture.html). Retrieved on 2011-12-10.
[9] http://www.courier-mta.org/authlib/
[10] http://wiki.dovecot.org/DovecotFeatures
[11] http://wiki.dovecot.org/Authentication.
[12] http://asg.web.cmu.edu/sasl/
[13] with Courier-pythonfilter filtering framework (http://swik.net/courier-pythonfilter)
[14] http://www-01.ibm.com/support/docview.wss?uid=swg21102611
[15] http://www-01.ibm.com/support/docview.wss?rs=463&context=SSKTMJ&dc=DB520&uid=swg21227751&loc=en_US&cs=UTF-8&lang=en&rss=ct463lotus
[16] http://www.postfix.org/FILTER_README.html

Microsoft Exchange Server

Microsoft Exchange Server

Developer(s)	Microsoft Corporation
Initial release	April 11, 1993
Stable release	Microsoft Exchange Server 2010 SP2[1] / December 4, 2011
Development status	Active
Programming language used	C and C++
Operating system	Microsoft Windows
Platform	x86-64 (64-bit)
Translation available	Multilingual
Type	Collaborative software
License	Proprietary (MS-EULA)
Website	www.microsoft.com/exchange [2]

Microsoft Exchange Server is the server side of a client–server, collaborative application product developed by Microsoft. It is part of the Microsoft Servers line of server products and is used by enterprises using Microsoft infrastructure products. Exchange's major features consist of electronic mail, calendaring, contacts and tasks; that work with Microsoft Outlook on PC and Mac, wireless synchronization of email, calendar, contacts with major mobile devices and browser-based access to information; and support for data storage.

History

Planning the migration from Microsoft's internal "legacy XENIX-based messaging system" to the Exchange Server environment began in April 1993,[3] and the process was completed in the late 1996 when the last XENIX server on the MS corporate backbone had been removed.[4]

Exchange 1.0

Microsoft Exchange 1.0

Developer(s)	Microsoft
Stable release	4.00.835.1374 (version 5.0) / October 14, 1996
Operating system	Microsoft Windows
Type	E-mail client
License	Proprietary EULA
Website	Exchange update for Windows 95 [5]

Windows Messaging, initially called Microsoft Exchange, is an e-mail client that was included with Windows 95 (beginning with OSR2), Windows 98, and Windows NT 4.0. (In Windows 98, it is not installed by default, but available as a separate program in the setup CD.) Microsoft Exchange gained wider usage with the release of Windows 95, as this was the only e-mail client that came bundled with it. Exchange was included throughout later releases of Windows up until the initial release of Windows 98, which by then also included Outlook Express 4.0.

- The original version lacked support of Internet mail (SMTP and POP3). They are only available with the separate Microsoft Plus! pack.
- HTML e-mail was shown in such a way that the message contained an *.ATT or *.htm attachment, which had to be saved and then viewed in a browser, as MS Exchange did not have support for HTML-formatted messages. Similarly, e-mail that did not use traditional message formatting was delivered in the form of text attachments with the *.ATT extension, which could be opened through Notepad. These files were in turn saved in the active Temp directory and some sensitive e-mail could therefore have been made available for other users to see.
- International characters were unsupported. Some e-mail that was sent with a non-ASCII or non-7/8-bit character set, was shown in the form of text attachments, which had to be saved and then read in a web browser, with the browser's text encoding set for a specified code page.
- Microsoft Fax, also called Microsoft at Work Fax (AWF), was the fax component to provide Send-and-Receive Fax capability; sent and received faxes were stored in the same .pst file as other messages, a first attempt at unified messaging by Microsoft. It also provided the ability to act as a fax server,[6] a capacity not available in later versions of Windows until Windows Vista.

In 1996, Microsoft Exchange was renamed to Windows Messaging, because of Microsoft's release of another Exchange product which was meant for servers. Windows Messaging had two branches of successors:

- In software bundled with Windows itself, these were Internet Mail and News in Windows 95 (and bundled with Internet Explorer 3), which was succeeded by Outlook Express 4.0 in Windows 98 (bundled with Internet Explorer 4.0 in Windows 95) and throughout newer Windows systems. These did not use the .pst file type.[7]
- Microsoft Outlook became the professional-grade and more direct successor of MS Exchange Client, which still uses the .pst file type. Because Microsoft Outlook used the same basic Windows Messaging profile, account, and e-mail settings (MAPI), Microsoft Exchange users not familiar with it may have thought that Outlook duplicated those settings and made copies of all their mail while they were trying out the new Microsoft Outlook 97. Thus, some MS Exchange users could have unknowingly deleted all their e-mail, thinking it was a copy, as Microsoft

Outlook did not have any front-end feature to notify users that it was actually using the same MS Exchange / Windows Messaging account.

Exchange Server 4.0

Microsoft began a preliminary planning of the Exchange 4.0 migration in April 1993.[4] In January 1995, 500 users were running on Exchange Beta 1, 5,000+ users running on Exchange Beta 2A in September 1995, and finally all 32,000 Microsoft mailboxes successfully migrated to Exchange and Microsoft Exchange shipped in April 1996.[8] Microsoft IT Group actually migrated all Microsoft employees to the Exchange platform before the product had the official Release status.[4]

Exchange Server 4.0, released on April, 1996,[9] was the original version of Exchange Server sold to the public, positioned as an upgrade to Microsoft Mail 3.5. The original version of Microsoft Mail (written by Microsoft) had been replaced, several weeks after Lotus acquired cc:Mail, by a package called *Network Courier*, acquired during the purchase of Consumer Software Inc. in April 1991.[10] Exchange Server was however an entirely new X.400-based client–server mail system with a single database store that also supported X.500 directory services. The directory used by Exchange Server eventually became Microsoft's Active Directory service, an LDAP-compliant directory server. Active Directory was integrated into Windows 2000 as the foundation of Windows Server domains.

Exchange Server 5.0

On May 23, 1997, **Exchange Server 5.0** was released, which introduced the new Exchange Administrator console, as well as opening up "integrated" access to SMTP-based networks for the first time. Unlike Microsoft Mail (which required a standalone SMTP relay), Exchange Server 5.0 could, with the help of an add-in called the Internet Mail Connector, communicate directly with servers using SMTP. Version 5.0 also introduced a new Web-based e-mail interface called Exchange Web Access, which was rebranded as Outlook Web Access in a later Service pack. Along with Exchange Server version 5.0, Microsoft released version 8.01 of Microsoft Outlook, version 5.0 of the Microsoft Exchange Client and version 7.5 of Microsoft Schedule+ to support the new features in the new version of Exchange Server.

Exchange Server 5.5

Introduced November 1997, was sold in two editions, Standard and Enterprise. They differ in database store size, mail transport connectors and clustering capabilities. The Standard Edition had the same 16 GB database size limitation as earlier versions of Exchange Server, while the Enterprise Edition had an increased limit of 16 TB (although Microsoft's best practices documentation recommends that the message store not exceed 100 GB). The Standard Edition includes the Site Connector, MS Mail Connector, Internet Mail Service (previously "Internet Mail Connector"), and Internet News Service (previously "Internet News Connector"), as well as software to interoperate with cc:Mail, Lotus Notes and Novell GroupWise. The Enterprise Edition adds an X.400 connector, and interoperability software with SNADS and PROFS. The Enterprise Edition also introduced two node clustering capability. Exchange Server 5.5 introduced a number of other new features including a new version of Outlook Web Access with Calendar support, support for IMAP4 and LDAP v3 clients and the Deleted Item Recovery feature. Exchange Server 5.5 was the last version of Exchange Server to have separate directory, SMTP and NNTP services. There was no new version of Exchange Client and Schedule+ for version 5.5, instead version 8.03 of Microsoft Outlook was released to support the new features of Exchange Server 5.5.

Exchange 2000 Server

Exchange 2000 Server (v6.0, code name Platinum), released on November 29, 2000, overcame many of the limitations of its predecessors. For example, it raised the maximum sizes of databases and increased the number of servers in a cluster from two to four. However, many customers were deterred from upgrading by the requirement for a full Microsoft Active Directory infrastructure to be in place, as unlike Exchange Server 5.5, Exchange 2000 Server had no built-in Directory Service, and had a dependency upon Active Directory. The migration process from Exchange Server 5.5 did not have any in-place upgrade path, and necessitated having the two systems online at the same time, with user-to-mailbox mapping and a temporary translation process between the two directories. Exchange 2000 Server also added support for instant messaging, but that capability was later spun off to Microsoft Office Live Communications Server.

Exchange Server 2003

Microsoft Exchange Server 2003

Developer(s)	Microsoft
Stable release	Microsoft Exchange Server 2003 SP2 (1.0) / October 19, 2005
Operating system	Microsoft Windows
Type	Server Client

Exchange Server 2003 (v6.5, code name Titanium) debuted on September 28, 2003. Exchange Server 2003 (currently at Service Pack 2) can be run on Windows 2000 Server (only if Service Pack 4 is first installed) and 32-bit Windows Server 2003, although some new features only work with the latter. Like Windows Server 2003, Exchange Server 2003 has many compatibility modes to allow users to slowly migrate to the new system. This is useful in large companies with distributed Exchange Server environments who cannot afford the downtime and expense that comes with a complete migration.

The June 2, 2003, release of **Exchange Server 2003** made the migration from pre-2000 versions of Exchange significantly easier (although still involved the same basic steps), and many users of Exchange Server 5.5 waited for the release of Exchange Server 2003 to upgrade. The upgrade process also required upgrading a company's servers to Windows 2000. Some customers opted to stay on a combination of Exchange Server 5.5 and Windows NT 4.0, both of which are no longer supported by Microsoft.

One of the new features in Exchange Server 2003 is enhanced disaster recovery [11] which allows administrators to bring the server online more quickly. This is done by allowing the server to send and receive mail while the message stores are being recovered from backup. Some features previously available in the Microsoft Mobile Information Server 2001/2002 products have been added to the core Exchange Server product, like Outlook Mobile Access and server-side Exchange ActiveSync, while the Mobile Information Server product itself has been dropped. Better anti-virus and anti-spam protection have also been added, both by providing built-in APIs that facilitate filtering software and built-in support for the basic methods of originating IP address, SPF ("Sender ID"), and DNSBL filtering which were standard on other open source and *nix-based mail servers. Also new is the ability to drop inbound e-mail before being fully processed, thus preventing delays in the message routing system. There are also improved message and mailbox management tools, which allow administrators to execute common chores more quickly. Others, such as Instant Messaging and Exchange Conferencing Server have been extracted completely in order to form separate products. Microsoft now appears to be positioning a combination of Microsoft Office, Microsoft Office Live Communications Server, Live Meeting and Sharepoint as its collaboration software of choice. Exchange Server is now to be simply e-mail and calendaring.

Exchange Server 2003 added several basic filtering methods to Exchange Server. They are not sophisticated enough to eliminate spam, but they can protect against DoS and mailbox flooding attacks. Exchange Server 2000 supported the ability to block a sender's address, or e-mail domain by adding '*@domain.com', which is still supported in Exchange Server 2003. Added filtering methods in Exchange Server 2003 are:

Connection filtering

Messages are blocked from DNS RBL lists[12] or from manually specified IP addresses/ranges

Recipient filtering

Messages blocked when sent to manually specified recipients on the server (for intranet-only addresses) or to any recipients not on the server (stopping spammers from guessing addresses)

Sender ID filtering

Sender ID, a form of Sender Policy Framework (SPF)

Intelligent Message Filter

A free Microsoft add-on that uses heuristic message analysis to block messages or direct them to the "Junk E-Mail" folder in Microsoft Outlook clients.[13]

Exchange 2003 mainstream support ended on April 14, 2009.[14]

Editions

Exchange Server 2003 is available in two versions, Standard Edition and Enterprise Edition. Standard Edition supports up to two storage groups (with one of the storage groups, called the recovery storage group, being reserved for database recovery operations) and a maximum of 2 databases per storage group. Each database is limited to a maximum size of 16GB.[15] Beginning with the release of Service Pack 2, Standard Edition allows a maximum database size of 75 GB, but only supports 16 GB by default; larger sized databases have to be updated-in with a registry change.[16] Enterprise Edition allows a 16 TB maximum database size, and supports up to 4 storage groups with 5 databases per storage group for a total of 20 databases per server.[17]

Exchange Server 2003 is included with both Microsoft Small Business Server 2003 Standard and Premium editions and is 32-bit only, and will not install on the various 64-bit versions of Windows Server 2003.

Exchange Server 2007

Microsoft Exchange Server 2007

Developer(s)	Microsoft
Stable release	Microsoft Exchange Server 2007 SP3 (8.03.0083.006) / June 20, 2010
Operating system	Microsoft Windows
Type	Server Client

Exchange Server 2007 (currently at Service Pack 3) was released on November 30, 2006, to business customers as part of Microsoft's roll-out wave of new products. It includes new clustering options, 64-bit support for greater scalability, voice mail integration, better search and support for Web services, better filtering options, and a new Outlook Web Access interface. Exchange 2007 also dropped support for Exchange 5.50 migrations, routing groups, admin groups, Outlook Mobile Access, X.400, and some API interfaces, amongst other features.[18]

Exchange Server 2007 (v8, code name E12, or with SP1 v8.1) runs only on 64-bit x86-64 versions of Windows Server. This requirement applies to supported production environments only; a 32-bit trial version is available for download and testing. Hence, companies currently running Exchange Server on 32-bit hardware will be required to replace or migrate hardware if they wish to upgrade to the new version. Companies that are currently running

Exchange Server on 64-bit capable hardware are still required to migrate from their existing Exchange 2000/2003 servers to a new 2007 server since in-place upgrades are not supported in 2007.

The first beta of Exchange Server 2007 (then named "Exchange 12" or E12) was released in December 2005 to a very limited number of beta testers. A wider beta was made available via TechNet Plus and MSDN subscriptions in March 2006 according to the Microsoft Exchange team blog.[19] On April 25, 2006, Microsoft announced that the next version of Exchange Server would be called **Exchange Server 2007**.

Exchange server 2007 comes in two editions, Exchange Server 2007 Standard edition and Exchange Server 2007 Enterprise edition. Standard edition can have 5 databases in up to 5 storage groups, while in Enterprise edition this is extended to 50 databases in up to 50 storage groups.

SCC and CCR are not supported in standard edition but LCR and SCR is supported. While in Exchange 2007 Enterprise Edition SCC, LCR, CCR and SCR are supported.

Exchange Server 2007 is an integrated part of the Innovative Communications Alliance products.[20]

New features

The principal enhancements, as outlined by Microsoft, are:[21]

- Protection: anti-spam, antivirus, compliance, clustering with data replication, improved security and encryption
- Improved Information Worker Access: improved calendaring, unified messaging, improved mobility, improved web access
- Improved IT Experience: 64-bit performance & scalability, command-line shell & simplified GUI, improved deployment, role separation, simplified routing
- Exchange Management Shell: a new command-line shell and scripting language for system administration (based on Windows PowerShell). Shell users can perform every task that can be performed in the Exchange Server graphical user interface plus additional tasks, and can program often-used or complex tasks into scripts that can be saved, shared, and re-used. The Exchange Management Shell has over 375 unique commands to manage features of Microsoft Exchange Server 2007.[22]
- "Unified Messaging" that lets users receive voice mail, e-mail, and faxes in their mailboxes, and lets them access their mailboxes from cell phones and other wireless devices. Voice commands can be given to control and listen to e-mail over the phone (and also send some basic messages, like "I'll be late")
- Increased the database maximum size limit. Database size is now limited to 16TB per database[23]
- Increased the maximum number of storage groups and mail databases per server, to 5 each for Standard Edition (from 1 each in Exchange Server 2003 Standard), and to 50 each for Enterprise Edition (from 4 groups and 20 databases in Exchange Server 2003 Enterprise).
- You can configure Outlook Anywhere (formerly known as RPC over HTTP) to provide external access to Microsoft Exchange Server 2007 for your clients. If you want Microsoft Office Outlook 2007 user profiles to be automatically configured to connect to Exchange 2007, configure the Autodiscover service. This also provides external URLs for Exchange services such as the Availability service and offline address book.

Exchange Server 2010

Microsoft Exchange 2010

Developer(s)	Microsoft
Stable release	Microsoft Exchange Server 2010 SP2 Update Rollup 3 (14.2.309.2) / May 29, 2012
Operating system	Microsoft Windows
Type	Server Client
License	Proprietary EULA
Website	[24]

Microsoft reached the RTM (Release To Manufacturing) milestone for Exchange Server 2010 on May, 2009, and it was officially launched on November 9, 2009.[25] A 120 day trial is downloadable from Microsoft.[24] Exchange Server 2010 (currently at Service Pack 2) is available in two server editions; Standard edition and Enterprise edition.

Major changes from previous versions of Exchange Server include:

- The high availability options for Mailbox Databases (SCC: Single Copy Clustering, CCR: Clustered Continuous Replication and LCR: Local Continuous Replication) and site resiliency functionality (SCR: Standby Continuous Replication) have been replaced by Database Availability Groups (DAGs) in Exchange Server 2010. Major DAG benefits include providing database level high availability (as opposed to server level), support for up to sixteen (16) copies of each database, and flexible configuration (databases copies may be added / removed at will without requiring major server reconfiguration). Each server that runs the Enterprise edition of Exchange Server 2010 can host up to 100 database copies.
- High availability for the Client Access Server role in Exchange Server 2010 is provided by using Client Access Server (CAS) arrays. A CAS array can contain multiple Client Access Servers in an Active Directory site and provide a single name endpoint for client connections. CAS arrays cannot span multiple Active Directory sites.
- In Exchange Server 2007, a clustered mailbox server could not be combined with any other roles. In Exchange Server 2010, the Mailbox Server Role may be combined with the Client Access Server and/or Hub Transport roles, regardless of whether or not the mailbox server participates in a Database Availability Group. (However, since Database Availability Groups use Windows Failover Clustering, and Microsoft does not support the combination of Windows Failover Clustering and Windows Network Load Balancing on the same server, a multi-role deployment will require the use of a 3rd party load balancer to provide load balancing and fault tolerance for the Client Access Server role).
- With the introduction of the RPC Client Access service, all Outlook clients access their mailbox database through the Client Access Server role. This abstraction layer allows for improved load balancing and redundancy and minimal client impact in the event of a database level *-over ("switchover" or "failover") event.
- Exchange Server 2010 provides cost savings in required hardware. Storage performance requirements (measured in IOPS: Input/Output operations Per Second) have been reduced by approximately 70% over Exchange Server 2007, and by approximately 90% over Exchange Server 2003. According to a case study, Microsoft IT was able to reduce hardware costs by 75% during the migration from Exchange Server 2007 to Exchange Server 2010.
- Exchange Server 2010 extends the large mailbox support introduced in Exchange Server 2007, and also introduces a Personal Archive feature to allow messages to be retained longer without the need for a 3rd party archival system. The Personal Archive is implemented as a secondary mailbox for archive-enabled users, and in Exchange Server 2010 Service Pack 1, the Personal Archive may be located on a different database than the primary mailbox, which may reside on a different disk if desired.

- The compliance and legal search features have been enhanced. What was formerly known as the "Dumpster" in previous versions of Exchange (a special storage area for messages which have been deleted from the Deleted Items folder or "permanently deleted" from a regular folder, such as the Inbox) has been evolved into the Recoverable Items folder in Exchange Server 2010. If configured appropriately, the Recoverable Items folder allows for a "tamper proof" storage area (users cannot circumvent the Recoverable Items folder to bypass legal discovery), which also provides a revision history of any modified items.
- Administration delegation can now be performed at a granular level due to Exchange Server 2010's implementation of Role Based Access Control (RBAC). Users and administrators can be given extremely fine grained abilities for functions provided both within the Exchange Management Console or Exchange Management Shell and in Outlook Web App. For example, a compliance officer may be given the ability to perform cross mailbox discovery searches within Outlook Web App; a help desk technician may be granted the ability to set an Out Of Office message for other employees within the company, or a branch administrator in a remote office may be granted the permission to perform specific Exchange Management Shell commands that pertain only to the Exchange server in their branch office.
- Outlook Web App includes improvements (including, for example, the ability for users to track their sent messages and printable calendar views) and the "Premium" experience is now available across multiple browsers (including Safari and Firefox).
- Distribution groups can now be "moderated", meaning that distribution groups can now be configured to allow users to join at will or only with a group moderator's permission, and individual messages sent to distribution groups can now be approved or denied by a moderator.
- Exchange Server 2010 introduces a transport concept called "Shadow Redundancy" which protects e-mail messages while they are in transit. If a Hub Transport server or an Edge Transport server fails after it has received a message for processing, but before it was able to deliver it to the next "hop" server, the server which sent the message to that transport server is now able to detect the failure and redeliver the message to a different Hub Transport or Edge Transport server for processing.

In January 2011, Microsoft Exchange Server 2010 won InfoWorld's 2011 Technology of the Year Award for Best Mail Server.

Clustering and high availability

Exchange Server Enterprise Edition supports clustering of up to 4 nodes when using Windows 2000 Server, and up to 8 nodes with Windows Server 2003. Exchange Server 2003 also introduced active-active clustering, but for two-node clusters only. In this setup, both servers in the cluster are allowed to be active simultaneously. This is opposed to Exchange's more common active-passive mode in which the failover servers in any cluster node cannot be used at all while their corresponding home servers are active. They must wait, inactive, for the home servers in the node to fail. Subsequent performance issues with active-active mode have led Microsoft to recommend that it should no longer be used.[26] In fact, support for active-active mode clustering has been discontinued with Exchange Server 2007.

Exchange's clustering (active-active or active-passive mode) has been criticized because of its requirement for servers in the cluster nodes to share the same physical data. The clustering in Exchange Server provides redundancy for Exchange Server as an *application*, but not for Exchange *data*.[27] In this scenario, the data can be regarded as a single point of failure, despite Microsoft's description of this set up as a "Shared Nothing" model.[28] This void has however been filled by ISV's and storage manufacturers, through "site resilience" solutions, such as geo-clustering and asynchronous data replication.[29] Exchange Server 2007 introduces new cluster terminology and configurations that address the shortcomings of the previous "shared data model".[30]

Exchange Server 2007 provides built-in support for asynchronous replication modeled on SQL Server's "Log shipping"[31] in CCR (Cluster Continuous Replication)[32] clusters, which are built on MSCS MNS (Microsoft

Cluster Service—Majority Node Set) clusters, which do not require shared storage. This type of cluster can be inexpensive and deployed in one, or "stretched" across two datacenters for protection against site-wide failures such as natural disasters. The limitation of CCR clusters is the ability to have only two nodes and the third node known as "voter node" or file share witness[33] that prevents "split brain"[33] scenarios, generally hosted as a file share on a Hub Transport Server.[34] The second type of cluster is the traditional clustering that was available in previous versions, and is now being referred to as SCC (Single Copy Cluster). In Exchange Server 2007 deployment of both CCR and SCC clusters has been simplified and improved; the entire cluster install process takes place during Exchange Server installation. LCR or Local Continuous Replication[34] has been referred to as the "poor man's cluster". It is designed to allow for data replication to an alternative drive attached to the same system and is intended to provide protection against local storage failures. It does not protect against the case where the server itself fails.

In November 2007, Microsoft released SP1 for Exchange Server 2007. This service pack includes an additional high-availability feature called SCR (Standby Continuous Replication). Unlike CCR which requires that both servers belong to a Windows cluster, typically residing in the same datacenter, SCR can replicate data to a non-clustered server, located in a separate datacenter.

With Exchange Server 2010, Microsoft introduced the concept of the Database Availability Group (DAG). A DAG contains Mailbox servers that become members of the DAG. Once a Mailbox server is a member of a DAG, the Mailbox Databases on that server can be copied to other members of the DAG. When you add a Mailbox server to a DAG, the Failover Clustering Windows role is installed on the server and all required clustering resources are created.

Licensing

Like Windows Server products, Exchange Server requires Client Access Licenses, which are different from Windows CALs. Corporate license agreements, such as the Enterprise Agreement, or EA, include Exchange Server CALs. It also comes as part of the Core CAL. Just like Windows Server and other server products from Microsoft, you can choose to use User or Device CALs. Device CALs are assigned to a device (workstation, laptop or PDA). User CALs, are assigned to a user or employee (not a mailbox). User CALs allow a user to access Exchange e-mail from any device. User and Device CALs are the same price, however cannot be used interchangeably. For Service Providers looking to host Microsoft Exchange, there is an SPLA (Service Provider License Agreement) available whereby Microsoft receives a monthly service fee in the place of the traditional Client Access Licenses. Two types of Exchange CAL are available: Exchange CAL Standard and Exchange CAL Enterprise. The Enterprise CAL is an add-on license to the Standard CAL.

Exchange Hosting

Microsoft Exchange Server can also be purchased as a hosted service from a number of providers.[35] Though Exchange Hosting has been around for more than 10 years, it is only recently that many providers have been marketing the service as "Cloud Computing" or Software-as-a-Service. Exchange Hosting allows for Microsoft Exchange Server to be running in the Internet also called the Cloud and managed by a "Hosted Exchange Server provider" instead of building and deploying the system in-house.

Exchange Online

Microsoft Exchange Online is an email, calendar and contacts solution delivered as a cloud service, hosted by Microsoft. It is essentially the same service offered by hosted Exchange providers and it is built on the same technologies as Microsoft Exchange Server. Exchange Online provides end users with a familiar email experience across PCs, the Web and mobile devices, while giving IT administrators or small businesses and professionals web-based tools for managing their online deployment.[36]

Microsoft Exchange is available both as on-premises software and as a hosted service with Exchange Online. Customers can also choose to combine both on-premises and online options in a hybrid deployment.

History

Exchange Online was first provided as a hosted service in dedicated customer environments in 2005 to select pilot customers.[37] Microsoft launched a multi-tenant version of Exchange Online as part of the Business Productivity Online Standard Suite in November 2008.[38] In June 2011, as part of the commercial release of Microsoft Office 365, Exchange Online was updated with the capabilities of Exchange Server 2010.

Exchange Server 2010 was developed concurrently as a server product and for the Exchange Online service.

Clients

Microsoft Exchange Server uses a proprietary RPC protocol, MAPI/RPC,[39] that was designed to be used by the Microsoft Outlook client. Clients capable of using the proprietary features of Exchange Server include Microsoft Outlook and Novell Evolution. Exchange Web Services (EWS), an alternative to the MAPI protocol, is a documented SOAP based protocol introduced with Exchange Server 2007 which significantly reduces synchronization time between the server vs. WebDAV, which is used by Exchange Server 2003. Exchange Web Services is used by the latest version of Microsoft Entourage for Mac and Microsoft Outlook for Mac. Also, since the release of Mac OS X v10.6 (also known as Mac OS X Snow Leopard), Mac computers running OS X include some support for this technology via Apple's Mail application. Built-in support with Mac OS X 10.6 requires the Exchange organization to be running Exchange Server 2007 SP1/SP2 or Exchange Server 2010.

Mac users wishing to access Exchange e-mail running on Exchange Server 2000 or 2003 must use Microsoft's Entourage client versions X, 2004 or 2008. Alternatively a limited version of Outlook Web Access is available to Mac users using a web browser. Entourage X, 2004 and 2008 do not support synchronizing tasks and notes with Exchange Servers 2000, 2003, 2007 or 2010. However Entourage 2008 "Web Services Edition", which is a free download from Microsoft for users of Office 2008, does support synchronizing tasks and notes with Exchange Server 2007 SP1 rollup update 4 or later (including Exchange 2010). Microsoft Outlook for Mac 2011 has replaced Entourage "Web Services Edition" but also requires Exchange Server 2007 or later.

E-mail hosted on an Exchange Server can also be accessed using SMTP, POP3 and IMAP4 protocols, using clients such as Outlook Express, Mozilla Thunderbird, and Lotus Notes. (These protocols must be enabled on the server. Recent versions of Exchange Server turn them off by default.)

Exchange Server mailboxes can also be accessed through a web browser, using Outlook Web Access (OWA), called Outlook Web App in Exchange Server 2010. Exchange Server 2003 also featured a version of OWA for mobile devices, called Outlook Mobile Access (OMA).

DavMail Gateway [40] allows any email client to connect to a Microsoft Outlook server with Outlook Web Access (OWA).

GNOME Evolution project can be used to Connect to MS-Exchange (in OWA mode for Exchange 2000/2003, native mode for Exchange 2007).[41] Evolution is now also available for Windows.[42][43]

ActiveSync

ActiveSync A component of Microsoft Windows	
ActiveSync 4.5 on Windows XP	
Details	
Replaced by	Windows Mobile Device Center

Microsoft ActiveSync

Initial release	1.0 / September 10, 1996
Stable release	4.5 / February 13, 2007
License	EULA
Website	ActiveSync 4.5 download [44] MSDN: ActiveSync [45]

Support for Exchange ActiveSync was added to Microsoft Exchange Server 2003. Exchange ActiveSync, in the context of Exchange Server, allows a compliant device such as a Windows Mobile device to securely synchronize mail, contacts and other data directly with an Exchange server. Since its inception, ActiveSync has become a popular mobile access standard for businesses due to cross-platform support from companies like Nokia and Apple Inc.[46] as well as its advanced device security and compliance features.

Support for Push E-mail was added to Exchange Server 2003 with Service Pack 2. Windows Mobile 5.0 requires the "Messaging and Security Feature Pack (MSFP)", later versions of the mobile operating system, such as Windows Phone 7, have the capability built in.[47] Many other devices now support ActiveSync push e-mail, such as the iPhone and Android Phones.[48]

Exchange Server 2007 and Exchange Server 2010 support the use of Exchange ActiveSync Policies. By using Exchange ActiveSync Policies, administrators can secure the devices that connect to the organization or remotely deactivate features on the devices. Administrators or users can also remotely wipe a lost mobile device.

References

[1] "Microsoft Exchange Server 2010 Service Pack 2 (SP2)" (http://www.microsoft.com/download/en/details.aspx?id=28190). . Retrieved December 4, 2011.

[2] http://www.microsoft.com/exchange

[3] Microsoft's Migration to Microsoft Exchange Server (http://web.archive.org/web/20050507010335/http://www.microsoft.com/technet/archive/itsolutions/intranet/build/exchgdep.mspx)

[4] XENIX – Microsoft Short-lived Love Affair with Unix (http://www.softpanorama.org/People/Torvalds/Finland_period/xenix_microsoft_shortlived_love_affair_with_unix.shtml)

[5] http://microsoft.com/windows95/downloads/contents/WURecommended/S_WUCommunications/W95ExchUpdate

[6] How to Set Up a Fax Server or Fax Client with Microsoft Fax (http://support.microsoft.com/kb/141292/EN-US/)

[7] The Microsoft Exchange Server 4.0 team invented .pst format for personal archives (http://support.microsoft.com/kb/297019/)

[8] Exchange Server 2007 Design and Architecture at Microsoft (http://technet.microsoft.com/en-us/library/bb894728.aspx)

[9] Build numbers and release dates for Exchange Server (http://support.microsoft.com/kb/158530/)

[10] Paul Korzeniowski (1992). "E-mail becoming foundation for networked applications—electronic mail; local area network" (http://findarticles.com/p/articles/mi_m0SMG/is_n1_v12/ai_11725946). Software Magazine. . Retrieved 2007-07-02.

[11] "Exchange Server backup and Disaster Recovery" (http://blog.vembu.com/ms-exchange-2003-2007-2010-backup-restore-using-storegrid-2/). . Retrieved 2012-05-23.

[12] "Implementing and Configuring Blacklist Support in Exchange Server 2003" (http://www.msexchange.org/tutorials/Blacklist_Support_Exchange_2003.html). . Retrieved 2007-07-02.

[13] "Exchange Intelligent Message Filter" (http://www.microsoft.com/technet/prodtechnol/exchange/downloads/2003/imf/default.mspx). . Retrieved 2007-07-02.
[14] "Microsoft Support Lifecycle" (http://support.microsoft.com/lifecycle/?p1=1773). .
[15] Exchange Server 2003 Administrator's Pocket Consultant
[16] "Registry tweak to set a 75gb store limit on Exchange 2003 Standard Sp2" (http://www.mrtweak.com/exchange-server/registry-tweak-to-set-a-75gb-store-limit-on-exchange-2003-standard-sp2/). . Retrieved 2007-07-02.
[17] "Exchange 2003 editions" (http://www.microsoft.com/exchange/evaluation/editions.mspx). . Retrieved 2007-07-02.
[18] Microsoft TechNet (http://technet.microsoft.com/en-us/library/79df6a3a-29a8-4935-b143-8a66c8d082d4.aspx)
[19] Terry Myerson (March 1, 2006). "Exchange 12 Beta 1 Community Technology Preview" (http://msexchangeteam.com/archive/2006/03/01/420941.aspx). *The Microsoft Exchange Team Blog*. . Retrieved 2009-07-16.
[20] Microsoft, Nortel unveil ICA's first products (http://www.infoworld.com/article/07/01/17/HNicafirstproducts_1.html)
[21] Microsoft Exchange Server Website (http://www.microsoft.com/exchange/)
[22] Exchange 2007 Cmdlet List (http://technet.microsoft.com/en-us/library/bb123703.aspx)
[23] Rodney Buike. "Understanding the Exchange Information Store" (http://www.msexchange.org/articles/Understanding-Exchange-Information-Store.html). . Retrieved 2008-12-19.
[24] Microsoft Exchange Server 2010 (http://microsoft.com/exchange/2010/)
[25] Microsoft releases Exchange 2010, acquires Teamprise | Beyond Binary - CNET News (http://news.cnet.com/8301-13860_3-10393208-56.html?tag=mncol;title)
[26] "Considerations when deploying Exchange on an Active/Active cluster" (http://support.microsoft.com/default.aspx?scid=kb;en-us;815180&product=exch2003). . Retrieved 2007-07-02. (Logging in required)
[27] "The benefits of Windows 2003 clustering with Exchange 2003" (http://blogs.technet.com/exchange/archive/2004/06/09/152186.aspx). . Retrieved 2007-07-02.
[28] "Exchange Clustering Concepts" (http://www.microsoft.com/technet/prodtechnol/exchange/2003/insider/clustering.mspx). . Retrieved 2007-07-02.
[29] "Storage Glossary: Basic Storage Terms" (http://www.microsoft.com/windowsserversystem/storage/storgloss.mspx). . Retrieved 2007-07-02.
[30] "High Availability" (http://web.archive.org/web/20071024040430/http://microsoft.com/technet/prodtechnol/exchange/E2k7Help/d2efb6f9-f70a-4f96-9f8d-f7aad6ae83d7.mspx?mfr=true). Archived from the original (http://www.microsoft.com/technet/prodtechnol/exchange/E2k7Help/d2efb6f9-f70a-4f96-9f8d-f7aad6ae83d7.mspx?mfr=true) on 2007-10-24. . Retrieved 2007-07-02.
[31] "Frequently asked questions—SQL Server 2000—Log shipping" (http://support.microsoft.com/kb/314515/en-us). . Retrieved 2007-07-02.
[32] "High availability" (http://www.microsoft.com/technet/prodtechnol/exchange/E2k7Help/d2efb6f9-f70a-4f96-9f8d-f7aad6ae83d7.mspx). . Retrieved 2007-07-02.
[33] "An update is available that adds a file share witness feature and a configurable cluster heartbeats feature to Windows Server 2003 Service Pack 1-based server clusters" (http://support.microsoft.com/kb/921181/en-us). . Retrieved 2007-07-02.
[34] "High Availability" (http://www.microsoft.com/technet/prodtechnol/exchange/E2k7Help/d2efb6f9-f70a-4f96-9f8d-f7aad6ae83d7.mspx). . Retrieved 2007-07-02.
[35] "Hosted Exchange Partner Directory" (http://www.microsoft.com/serviceproviders/solutions/catalog.aspx). . Retrieved 2007-07-02.
[36] "Microsoft Exchange Online for Enterprises Service Description" (http://www.microsoft.com/download/en/details.aspx?id=13602). *Office 365 for Enterprise Service Descriptions*. Microsoft Corporation. .
[37] Fried, Ina. "Microsoft hops into managed PC business" (http://news.cnet.com/Microsoft-hops-into-managed-PC-business/2100-1011_3-5609320.html). *CNET News*. CNET. . Retrieved 7 July 2011.
[38] "Exchange Online and SharePoint Online Out of Beta and Ready for Purchase" (http://www.microsoft.com/presspass/press/2008/nov08/11-17ExchangeSharePointOnlinePR.mspx). *Microsoft News Center*. Microsoft Corporation. . Retrieved 7 July 2011.
[39] Exchange Server Protocols (http://msdn.microsoft.com/en-us/library/cc307725(EXCHG.80).aspx)
[40] http://davmail.sourceforge.net/
[41] Evolution/FAQ - GNOME Live! (http://live.gnome.org/Evolution/FAQ#Evolution_Exchange_.28formerly_known_as_Connector.29)
[42] Evolution (software)
[43] Windows - Evolution (http://www.go-evolution.org/Windows#Evolution_Installers_for_Windows)
[44] http://www.microsoft.com/downloads/en/details.aspx?FamilyID=9e641c34-6f7f-404d-a04b-dc09f8141141
[45] http://msdn.microsoft.com/en-us/library/aa913903.aspx
[46] "Microsoft Exchange ActiveSync Licensees" (http://www.microsoft.com/exchange/2007/evaluation/features/owa_mobile.mspx). Microsoft. . Retrieved 2009-09-22.
[47] "Exchange ActiveSync: Frequently Asked Questions" (http://technet.microsoft.com/en-us/exchange/bb288524.aspx). TechNet. . Retrieved 2009-09-22.
[48] "Apple - iPhone in Business" (http://www.apple.com/iphone/business/). Apple Computer. . Retrieved 2009-09-22.

Further reading

- Walther, Henrik. *How to Cheat at Configuring Exchange Server 2007*. ISBN 978-1-59749-137-2.
- Morimoto, Rand; Michael Noel, Chris Amaris, Andrew Abbate, Mark Weinhardt. *Exchange Server 2007 Unleashed*. ISBN 978-0-672-32920-3.
- Morimoto, Rand; Michael Noel, Chris Amaris, Andrew Abbate, Mark Weinhardt. *Exchange Server 2010 Unleashed*. ISBN 978-0-672-33046-9.
- McBee, Jim; Barry Gerber. *Mastering Microsoft Exchange Server 2007*. ISBN 978-0-470-04289-2.
- Cavalancia, MCSE, MCT, MCNE, MCNI, Nick. *Microsoft Exchange Server 2007: A Beginner's Guide*. ISBN 978-0-07-148639-2.
- Luckett, Richard; Bharat Suneja, William Lefkovics. *Microsoft Exchange Server 2007: The Complete Reference*. ISBN 978-0-07-149084-9.
- Jagott, Siegfried; Joel Stidley. *Microsoft Exchange Server 2010 Best Practices*. ISBN 978-0-7356-2719-2.
- Redmond, Tony. *Microsoft Exchange Server 2010 Inside Out (including SP1)*. ISBN 978-0-7356-4061-0.

External links

- "Microsoft Exchange" (http://www.microsoft.com/exchange/)
- "Tips and Tricks for manual Exchange Server recovery in case of damage" (http://www.repair-file.com/exchange-server-recovery.htm)
- "EHLO - The Exchange Team Blog" (http://msexchangeteam.com/)
- "A brief history of time - Exchange Server way" (http://msexchangeteam.com/archive/2008/01/02/447806.aspx) Written by the Exchange Team
- "Microsoft TechNet - Exchange Server Home" (http://technet.microsoft.com/en-us/exchange/default.aspx)
- "Exchange Server Supportability Matrix" (http://technet.microsoft.com/en-us/library/ee338574.aspx)
- "OpenChange: a portable Open Source implementation of Microsoft Exchange Server and Exchange protocols." (http://www.openchange.org/)
- "Features no longer supported by Exchange Server 2007" (http://searchexchange.techtarget.com/tip/0,289483,sid43_gci1223025,00.html)
- "Active Directory LDAP Compliance" (http://www.microsoft.com/windowsserver2003/techinfo/overview/ldapcomp.mspx) Microsoft Corporation, December 2, 2003. Retrieved November 2, 2005.
- "Evaluation of Microsoft Exchange 2010 Beta at IT Reviews July 27 2009" (http://www.itreviews.co.uk/software/s777.htm)
- "iPhone Exchange Setup Guide + Common Activesync / Exchange myths" (http://computerrepairservice.net/blog/iphone-exchange-setup-guide)
- "Windows Phone 7 - Exchange Setup Guide" (http://www.cloudwire.co/setup/windows-phone7)
- "Evolution connecting to Exchange" (http://live.gnome.org/Evolution/FAQ#Evolution_Exchange_.28formerly_known_as_Connector.29)
- "Evolution Windows NSIS Installer" (http://www.go-evolution.org/Windows#Evolution_Installers_for_Windows)
- Exchange Web Forms (http://msdn.microsoft.com/en-us/library/aa142651(EXCHG.65).aspx)
- Passionate team working on Microsoft Exchange server (http://www.MSExchangeGuru.com)
- Exchange Server Tutorials (http://exchangeserverpro.com)
- Exchange Server Consultancy Support (http://exchangeserver.biz/)
- MS Exchange Backup and Disaster Recovery Softwares (http://storegrid.vembu.com/online-backup/ms-exchange-backup.php)

Article Sources and Contributors

System administrator *Source*: http://en.wikipedia.org/w/index.php?oldid=504196051 *Contributors*: 16x9, ABF, ADobkin, Aamodp, Abune, Adalie, Aditya, Akhil 0950, Al Lemos, Andkore, Andy Dingley, Aneeshjoseph, Aurick, Autarch, B137, BD2412, Banej, BarretB, Beland, Bob2000, Bovineone, Btmiller, BuddhaInside, CBAwiki, Callmejosh, Canadian-Bacon, Chris Roy, Christian75, CompREM, Coren, Crazytales, CyntWorkStuff, DBigXray, DMGCR, DMahalko, DanPhilpott, David Gerard, Davidshq, Daytona2, Dcstamm, DerHexer, Desalvionjr, Discospinster, Djamund, Doctorgage, DrMrScience, Drgrussell, Druiloor, Dveeden, EagleOne, Egsan Bacon, Einar414, Ejl, Elsendero, Entens, Epbr123, Exelsiar, Finalcoz, Frap, Fruckert, Fubar Obfusco, Gafaddict, Gardar Rurak, Gee totes, Gogo Dodo, Goliathlibrary, GoneAwayNowAndRetired, Grafen, Halbared, Harland1, Haseo9999, Heimstern, Hella, Hm2k, HobbesPDX, Howmuchi, Hudi124, Hux, Hydriz, Hydrogen Iodide, Ib80, Ilir pz, Immunize, Invisible decoder, Iridescent, Ixfd64, JZCL, Jahiegel, JamesAM, Jamoche, Javert, Jayden54, Jbond00747, Jeff3000, Jesse.kelly, John254, Jokestress, Jolomo, Josh Parris, Joy, Jsbillings, Jtan163, Justinfr, Karlski2010, Khym Chanur, Kiran.thunga, Kkm010, Knight Samar, KnightRider, Krellis, Kris DKET, Kuszi, Kvedulv, Kyorosuke, La goutte de pluie, Lance Albin, Leafyplant, Lear's Fool, Lewik, LittleWink, MBisanz, MER-C, Manop, McSly, Mdd, Mdd4696, Michael Hardy, MikeDogma, Mintleaf, Mmortal03, Mo5h, Moideenshakir, Moverton, Mudklipz6969, Muhandes, NaSPA1, Ned Scotland, Netoholic, Nickjo, NightMonkey, NillaGoon, Nneonneo, Not hylas( ), Nthomas, Octahedron80, Oktanouc, Olasunmbo Olayinka, Ostevens, Ouc, Owenebrady, Paapasolo, Packet2, Philip Trueman, Pinethicket, Pit, Poor Yorick, Public Menace, Punanimal, RDates, RainbowOfLight, Rcutty, Redeyed Treefrog, Rfsmit, Rhobite, Ringbang, Rlevse, Robert Brockway, RoyBoy, Russcelt, Sade, Saint Aardvark, Sasquatch, ScienceWick2011, Seowens, Sephiroth storm, Seven of Nine, Shanemodonnell, Shanes, Shashidahal, SimonP, SirBruce, SirLeopold, SlimVirgin, Smihael, Snowolf, SpacemanSpiff, Spellcast, Spick And Span, Spug, Stackleschwien, Stevenplunkett, Strykerhorse, Sudozero, Surv1v4l1st, Taibah U, Thaurisil, Thingg, Thumperward, Tobias Bergemann, TreyHarris, Ufu4u, UncleBubba, Unixguy, Vegthura, Vipinhari, Vishnava, Vivio Testarossa, Weylinp, World10, Xiahou, Xiong Chiamiov, Zondor, 419 anonymous edits

Usage share of operating systems *Source*: http://en.wikipedia.org/w/index.php?oldid=503947988 *Contributors*: 1exec1, Abhkum, Althepal, Antigin, Arichnad, Arnos78, Athaba, Barry McGuiness, Bender235, BerlinSight, Bkil, Bomazi, Bonadea, Cengelsm, Charonn0, Cimmind, Cityscape4, Cntras, CommonsDelinker, Czarkoff, DMTagatac, Daniel.Cardenas, DanielWaterworth, Darklilac, Darklock, Emmebi76, Faagel, Gary King, Get knowledge spread knowledge, Ghettoblaster, GoingBatty, Gona.eu, Gregory Heffley, HAl, Harumphy, HeWhoMowedTheLawn, Heroman1234, History2007, Hullaballoo Wolfowitz, I384403, Idler921, J.delanoy, Jamison Lofthouse, Jasper Deng, Jdm64, Jerryobject, Jimthing, Jman8088, Joaopaulo1511, Kev19, Khr0n0s, Kkm010, Koavf, LPfi, Lakeyboy, LarrHoff, Lineplus, M gol, Marcos Marado, Mclaud2000, Mdwh, Melizg, MillenniumB, Mo ainm, Morrison.a.n, Mr.Dust, Mranderson2008, NTox, NeonMerlin, Nicolas Love, Nigelj, Nono64, Nwusr123log, Oscailt, OspreyPL, Ozalperen, Ozob, Perspectoff, Pogson, Pstavroulis, Quantumelfmage, Rasmasyean, Raysonho, Realflash, Rfvuhbtg, Rich Farmbrough, Rodrigue, Rwwww, Ryan Norton, SF007, Shalineth, Shikamarusam, Sibtain85, Slasher-fun, Smalljim, Starkiller88, SterlingNorth, Strcat, Synchronism, TheGeneralUser, TheGrg, Thecheesykid, This, that and the other, ToolmakerSteve, Unflavoured, Universalcosmos, UrbanTerrorist, Useerup, Wageslave, Wikante, Wikiolap, WilliamH, Xpclient, Yworo, 270 anonymous edits

Linux distribution *Source*: http://en.wikipedia.org/w/index.php?oldid=504904597 *Contributors*: 192.35.241.xxx, 199.183.24.xxx, 1exec1, 1wheel, 6.27, A.M., A3li, A876, Aav, Acerperi, Activedreamer, AdrianTM, Ahoerstemeier, Ahunt, Akanemoto, Alex at kms, Alfem, Alicekey, AlistairMcMillan, Alpha Omicron, Alphax, Am088, Amazing Blair, Andre Engels, Andrewthorne, Andthorn, Angelsh, Anja98, Antaeus Feldspar, Anthony, Ap, Archanamiya, Arianit Kukaj, Ashawley, Avouac, Bassem, Battulafamily, Bdcomp, Bearcat, Beefball, Beland, BenBurch, Bernopedia, Bobblehead, Bowmanjj, BrokenSegue, Bryan Derksen, Burgundavia, Butko, C45207, CFeyecare, Calebbell, CanisRufus, Capt-google, CaribDigita, Cellstorm, Centrx, CesarB, Cfailde, Chadloder, Chargh, Charles Gaudette, Chealer, Cheese Sandwich, Chris Pickett, Chris55, ChrisHodgesUK, Christian List, Cidtux, Cjfsyntropy, Cokoli, CommonsDelinker, Connor Behan, Conversion script, Coolcanine, CoolingGibbon, Crt, Curps, Cwolfsheep, Cyber rigger, D. F. Schmidt, DIG, Damian Yerrick, Danakil, Dannyb13, David Gerard, DavidCary, Dawnseeker2000, Dcljr, Dejvid, DevastatorIIC, Dispenser, Dmerrill, DocWatson42, Dondilly, Dori, Dp, Dylan Lake, Edgarde, Edward, Egil, ElBenevolente, Elfguy, Elisonniven, Ellmist, Eloquence, Emc2, Emperorbma, ErkinBatu, Eugrus, Everyking, Evice, FT2, Fady saber, Falcon8765, Farizluqman, Fayazbhai, FleetCommand, Fluteflute, Fo0bar, Freakofnurture, Freed42, Fubar Obfusco, Fuzzie, GL9!, Gadfium, Galoubet, Garo, Gbeeker, Gene93k, Geraldshields11, Germ, Ghettoblaster, Giftlite, Giggy, Gilgamesh, Giraffedata, Glimz, Gmaxwell, Gobonobo, Gofurix, Gogo Dodo, GraemeL, Grandscribe, GrapeSteinbeck, Greyweather, Gronky, Gruepig, Gurch, Gwaka Lumpa, Gwicke, Gzyeah, Haakon, HaeB, Haham hanuka, Hairy Dude, Ham Pastrami, Hannes Hirzel, Harvester, Hayabusa future, Hfastedge, Hoary, Hulsie, Hurricane111, Hux, Hvn0413, IEEE, IRWolfie-, Icechen1, Ida Shaw, Imrathor, Imtiaz Rahi, Intgr, Ipsign, Irishbitte, Isilanes, Itjrscmanaus, Iune, Ivan Pozdeev, JB82, JGXenite, JLaTondre, Jaalto, Jaan513, Jagginess, Jamesontai, Jannex, Jarble, Jdm64, Jedimike, Jeffq, Jerebin, Jerome Charles Potts, Jitinarora, Jmundo, Joakim, JoeSmack, John J. Bulten, Jojoyee, Jordi Burguet Castell, Joshua-the-grey, Joy, Jpk, Julien, Kaal, Kaare, KamasamaK, Kaobear, Karen Johnson, Karol Langner, Kaveh, Kbolino, Kdammers, Ke4roh, Keithonearth, Kenny sh, Kentynet, Kgrindal, Kim Bruning, KimDabelsteinPetersen, Kimiko, Kingmundi, KitchM, Kjb, Kkmurray, Kl4m-AWB, Knownot, Kozuch, Ksn, Kumarat9pm, Kupojsin, KuriosD, Kwertii, LPfi, Ladislav, Lakefall, Laudaka, Lentower, Lesserangel, LiDaobing, Ligart, Lightedbulb, Linton, Linuxerist, Llywrch, Logicwiki, Lordmetroid, Lotje, Lucid, MER-C, Maerk, Man with one red shoe, Manop, MarSch, Markus Krötzsch, Martani, Massysett, Materialscientist, Matonen, Matthew2414, Max rspct, Maximus Rex, Meiskam, Member, MementoVivere, MethodOne, Mikenolte, Minesweeper, Minghong, MiniStephan, Mintguy, Miquonranger03, Mlk, Modemac, Morte, Mortense, Motor, MrMarmite, Mrzaius, Mulad, Mushroom, Naddy, Nandhp, Nanshu, NapoliRoma, Neonz, NerdyNSK, NeroArcNume, Neustradamus, NicM, NickBush24, Nigelj, Nikai, Nomorecorruptcops, Noneexistent, Norm, NuclearWarfare, Nyat, ObfuscatePenguin, Odyssomay, Olki, PL Przemek, Paddu, Pakaran, Paradoxsociety, Paul Richter, Pazuzu156, Pcmacdon, Peak, Peng, Pengo, Perkster, Peterfa, PhilKnight, Pichote, Pictureuploader, PierreAbbat, Pimlottc, Pinkville, Plugwash, Pmlineditor, Pravinvet, Prolog, Proteus71, Psb777, Pschemp, Pudgenet, Pwilt328, RAMChYLD, Rahul.quara20, Razimantv, Rbrwr, Rdash, RedWolf, ReddishShift, Reisio, RexNL, Rich Farmbrough, Rjackson72, Rmstallman, RobLa, Robert Brockway, Robomaeyhem, Robust Physique, Ross Uber, Ruud Koot, Rwalker, Rwxrwxrwx, RxS, S. Ugarte, SF007, Salgueiro, Samsara, Samw, Sietse Snel, SimonP, Sir Lewk, Sj, SkonesMickLoud, Skylerb, Smyth, Snoyes, SpaceFlight89, Spyhawk, Stan Shebs, Stanleypng, Stbalbach, Stevenj, Stevertigo, Stewartadcock, StormyJerry, Strcat, Struthious Bandersnatch, StuartBrady, SudoGhost, TakuyaMurata, Takuzinis, Tarquin, Tb33, Tcascardo, Tcms2330, Technobadger, Teggles, The Anome, The Firewall, Thev00d00, This, that and the other, Thue, Thumperward, Tim Graff, Tlroche, Tobias Bergemann, Tomgibbons, Tonybuaa, Toussaint, Trontonic, Troykitch, Trufetes en, Turbodin, Tverbeek, Two halves, Ubupc, Unixguy, Vaganyik, Val42, Valdisvi, Valent, Vanished User 03, Vespristiano, Victorgehrke, ViperSnake151, Viranch, Vlyalcin, Walshga, Walter.bender, Wernher, Wesley, Wik, WikiSolved, Wmpatricio, Woohookitty, Writtenonsand, Ww, Wwwolf2, Wwwolf3, XP1, XTaran, Xinit, Yama, Yngwin, Yonkie, Yono, Youssefsan, Ysangkok, Yworo, Zaiger, Zen-master, Zenupassio, Zoeblackmore, Zoicon5, Սահակ, පසිඳු කාවින්ද, 616 anonymous edits

Debian *Source*: http://en.wikipedia.org/w/index.php?oldid=504498051 *Contributors*: (, (14$H, *drew, -mushroom, 16@r, 4th-otaku, 62.2.17.xxx, A.M., A7Bo, ABCD, ALK, AThing, Acerperi, AdrianTM, Ahunt, Ajdlinux, Akanemoto, Alexr, Alinutza b, Alisha.4m, AlistairMcMillan, AlloG, Altonbr, Amxitsa, Andareed, Andorin Kato, Andreas T. Hagli, Andrel, Andrew Rodland, Andrewferrier, Android Mouse, Angela, Angr, Anomie, Antaeus Feldspar, Antidrugue, Ap, Arch o median, Archanamiya, Ardonik, Aripollak, Arite, ArnoldReinhold, Aronzak, Asav, Ashley Pomeroy, Ashley Y, Atjesse, Attys, Auno, Avbidder, AzertyFab, Aztak, Azul, BW, Bacon-chan, Baszoetekouw, Bdesham, Bdsatish, Belorn, Ben-Zin, Bender235, Benjamin Mako Hill, Beno1000, Benuski, Beofluff, Betterworld, Bevo, Bhaak, BiT, Blehfu, Boborok, Bongwarrior, Bookofjude, Borislavsabev, Brandon, Brianski, Bucephalus, Burgundavia, Byap, CFeyecare, CGA, CIreland, Cajunman4life, Canaima, CanisRufus, Capricorn42, Cassowary, Catapult, Cbpares it3a, Centrx, CesarB, Champaign IL, CharlesC, Chealer, Chevelle572, Chi And H, Cillophako, Ciphergoth, Cjcollier, Classical geographer, Clbt, ClementSeveillac, Colin Marquardt, CommonsDelinker, Computermatt, Conscious, Conversion script, Cookiecaper, CoolingGibbon, D64, DFulbright, DHN, DJSnuggles, DVD R W, Daf, Daff555, DalekClock, Damian Yerrick, Dan100, Daniel15, DannyKitty, Dante.regis, Danusch, Darac, Darana, Dataphile, David Gerard, Davide89v, Dawnseeker2000, Dawynn, Dbenbenn, Dbmoodb, Deeahbz, Deerslayer, Deflective, Dejvid, Deleteme42, Delirium, Delta04, Demitsu, Denis C., Denny, Dfred, Dieman, Digerati1338, Dinomite, Dkrogers, Dondelelcaro, Doors5678, Dori, Doulos Christos, Dragon 280, Drange net, Drt1245, Dtrebbien, Duckwitz, Dwandelt, Dwheeler, Dylan Lake, Dysprosia, EVula, Eanschuessler, Easyas12c, EdH, Ederkbula, Edward, Eeera, Ehamberg, Ehn, El Cubano, Elyk53, EmersonSmalltalk, Emperorbma, Eriberto mota, ErikStewart, ErkinBatu, Eurleif, Evice, Fatmangoz, Feral2k, Fibonacci, Fishcakes, Fluteflute, Fo0bar, Foolishgrunt, Fractal3, FranklinPiat, Frap, Fratrep, Fredrik, Freebiekr, Freed42, Fritz Saalfeld, Froth, Frozenbacon, Fsateler, Fsolda, Fubar Obfusco, Fuzzie, Gaius Cornelius, Gamera2, Gardrek, GastonFiore, Gaurav, Gb-uk, Gerfried Fuchs, Geronimooo, Gilliam, Glycerinester, Gmaxwell, Gnp, Gnu.bhai, Gobonobo, Gogo Dodo, Gosub, GrAfFiT, Graham87, GrahamHardy, Grandscribe, Grendelkhan, Gronky, Guaka, Gurugantum, Guy Peters, Gwalla, Gwern, Gwicke, Gyro Copter, Gürkan Sengün, H2g2bob, HTML2011, HVS, Hadal, Hard2Win, Hellmark, Henning Makholm, Henry Flower, Hertz1888, Hew, Hobart, Honeyman, Htaccess, Hugo 87, HugoEckener, HuntingStorm, Hvn0413, Hyad, Hydrargyrum, Hydrox, IGeMiNix, Ibroadfo, Iggymwangi, Indefatigable, InfernalPanda, Insaneinside, Intersofia, Isbilia, Isidore, Isilanes, Ismouton, Itai, Iune, Izogi, JVz, JWCurtis2003, JWilk, Jaalto, Jacco, Jacoplane, Jae, Jakob5675a, Jason108, JasonWoof, Jdm64, Jdorwin, Jdowland, Jdthood, Jeffrey Smith, Jerebin, Jeremy Visser, Jerryobject, Jfmantis, Jgrahn, Jleedev, Jm51, Joe Wreschnig, Joe routt, JoeSmack, Joeblakesley, Johannes121, John Nixon, John of Reading, Johnleemk, JonathanRidgway, Jonkerz, Jordi, Jordi Burguet Castell, JorgePeixoto, Joshk, Joy, Jredmond, Jsnx, Julesd, Justcorrected, Jxgreat, KAMiKAZOW, KDesk, Kaare, Kabads, Kaobear, Kaol, Karol Langner, Kedadi, Keithonearth, Kelch, Kesla, KimDabelsteinPetersen, Kimiko, Kingpin13, Kisama, Kjoonlee, Kl4m-AWB, KneeLess, Kohlrabi, Koolabsol, Kouign-amann, Koyaanis Qatsi, Kozuch, Kreutzm, Kvaks, Kwamikagami, Kyng, LFaraone, LOL, LX, Lahgoon, Lan3y, Leandrod, LeeHunter, Legoktm, Lentower, Leranedo, Lezek, Lightedbulb, Lightmouse, LilHelpa, LimoWreck, LinuxDemos, Linuxbeak, Locos epraix, Lokpest, Lopifalko, LordKenTheGreat, Lordmetroid, Lotje, Lucifer Spam, Lukejamesoconnor, Lysander89, Mac, MadSurgeon, Madduck, Man with one red shoe, Marco Krohn, Marcusfriedman, Margamanterola, Mark alfred, Markhobley, Marko knoebl, Markus Krötzsch, Marky1124, Martin451, Marudubshinki, Masatran, Matt Crypto, Matthew Woodcraft, Matusz, Maulkin, Mav, Mdosisza, Mdwh, Meandtheshell, Meelar, Melancholie, Melesse, Melnakeeb, Mentifisto, Merphant, Metz2000, Mgius, Michael Hardy, MichaelBillington, Mike92591, Mikecron, Mikeyo, Mineralè, Miranda, Mmj, Mmmready, Mms, Momet, Monz, Mrfebruary, Mrmrbeaniepiece, Msikma, Mst14, Mwtoews, NSK Nikolaos S. Karastathis, Naive cynic, Najoj, Nanshu, Nate Silva, Nazrani, NeoChaosX, Neonz, Neuro, NicoNet, Nihiltres, Nikai, Noahspurrier, Nooby god, Nopetro, Nurg, OKtosiTe, Olexandr Kravchuk, OllieWilliamson, Oneiros, Orgads, Oscarnet75, OverlordQ, Owlbuster, PL Przemek, Palfrey, Palosirkka, Paradoxian, Passargea, Paul Stansifer, Paxcoder, Pboi20, Pbryan, Pejman47, Petchboo, Pgan002, Phil Boswell, PhilHibbs, Philipp Kern, Phleg, Phobie, Phoenix-forgotten, Phoenix2275, Pictureuploader, PieterDeBruijn, Plasticup, Plugwash, Pmc, Pmsyyz, Pmyteh, Pne, Poccil, Prairie Dad, Praka123, Praveen Illa, Principalkec, Prohlep, Prolog, Przemub, PseudoSudo, Ptomes, Python eggs, Qji, Quadell, Quebec99, R'n'B, Rainkid, Raphael Frey, Ravemaker, Ray Dassen, RazvanLNX, Rbuj, Reedy, Reisio, Rfl, Rich Farmbrough, RichiH, Richwales, Ringbang, Ripero, Rjwilmsi, Rlmorgan, Robert impey, Robertmh, Robust Physique, Roman V. Odaisky, Rossumcapek, Royanee, Rsaxvc, Rsrikanth05, Rudiger Junior, Rwxrwxrwx, Ryan spencer, Ryguasu, Rythie, SF007, SVG, Sam Hocevar, Samsara, Samuelsen, Santa naz, Sbharris, Scheater5, SciFi-er, Scientus, Sesse, Sfan00 IMG, Sfmammamia, Shadowjams, Shaka0, Shooke, Sigondronggondrong, Silas S. Brown, Skizzik, Skx, SlackerMom, Slashme, SoLowRockerMan, Solipsist, Somebody in the WWW, Speculatrix, Sperling, Spiritllama, Spoon!, Squash, Static623, Stefan Kögl, SteveJothen, SteveSims, Stevertigo, Str4nd, StuffOfInterest, SudoGhost, Sundar2000, Suruena, Svetovid, Svick, Swineinfluenza, Synchrite, Tbjablin, Tedickey, Telofy, TerminalPreppie, TexasAndroid, The Anome, TheSameGuy, Tholme, Thomasmallen, Thue, Thumperward, Thuresson, Thv, Tim Retout, Timo Honkasalo, Timrichardson, Timwi, Tkinias, Tkinkhorst, Tlroche, Tohd8BohaithuGh1, Tom harrison, Tommy Kronkvist, TommyG, Tothwolf, Toussaint, Tr606, Tragek, Traroth, Treekids, Trusilver, Turnstep, Tzartzam, Untrammelled, Unyoyega, VARGUX, VShaka, Vedant lath, Vigilius, VodkaJazz, Voidxor, VoluntarySlave, WJetChao, Waves00, Wdrev, Whkoh, Wikidrone, Will Beback Auto, William Avery, Winston Chuen-Shih Yang, Wireless friend, Wwwolf3, XTaran, Ximian99, Xypron, Yekrats, Yoe, Yoghurt, Yono, Youssefsan, Yworo, ZOrK, Zoicon5, Zundark, Zvonsully, Ævar Arnfjörð Bjarmason, Александър, Սահակ, 941 anonymous edits

Ubuntu (operating system) *Source*: http://en.wikipedia.org/w/index.php?oldid=504586967 *Contributors*: (chubbstar), ++Martin++, -mushroom, 0imagination, 122589423KM, 16@r, 1812ahill, 1o8, 2fargon, 2p0rk, 41523, 4th-otaku, 7severn7, A.M., A2-computist, A20 Gate, AGiorgio08, ALUOPline2, AThing, Aanhorn, Aapo Laitinen, Aasami, Abu badali, Acerperi, Achilles2.0, Actam, Adagio Cantabile, Admrboltz, Adrian J. Hunter, AdrianTM, Aeon17x, Aerographer1981, Aezram, Aeæ, Afccarl1994, Afiler, AgadaUrbanit, Ahunt, Ahy1, Airplaneman, Ajithrshenoy, Aka042, Akanemoto, Alacard, Alansohn, Alerante, Alex Perrier, Alex2610, AlexanderHaas, Alexd, Alexdragon, Alexkon, Alinutza b, Alisha.4m, AlistairMcMillan, Aljjam, Aljoriz, Alkivar, Allens, Almejo, Alpha Beta Epsilon, AlphaAqua, Alphamouse, Altonbr, Alukasz, Amuletxheart, Andareed, Andre Engels, Andrewpmk, Andromeda, Andy16666, Andypandy.UK, Andypham3000, Anon user, Anonymous Dissident, Antaeus Feldspar, Antandrus, Antidrugue, Antilived, Anttipng, Apoc2400, Apollo2011, Arcandam, Archanamiya, Arckeda, Ardnivar, ArglebargleIV, ArielGold, Arite, Arjayay, Arkrishna, Arpabone, Artagnon, Artem Karimov, Arundhati bakshi, Arvindn, Asdirk, Aseld, Ash211, Ashkbsdhkb, Ashley Y, Ashwinrajeev, Asparagus, Athertek, Atlantic 252, Atropos, Audacious Hope, Audin, Audriusa, Aurora Studios, Aus jack, AussieScribe, Austin512, Austriacus, Avatar of Justice, Avenue, Avono, Awefweuaifh, AwesomeMachine, Awickert, Axlrosen, Az1568, Azrael Nightwalker, B, B Fizz, B694kp8d, B9 hummingbird hovering, Backslash Forwardslash, Bakken, Balaam's Miracle, Balls23, Balrog-kun, Balubino, Banazir, Basawala, Battoe19, BazookaJoe, Beano, Beao, Bearcat, Beefball, Beetstra, Beland, BenBen1234, Benandorsqueaks, Bender235, Bender2k14, Benixau, Benjamin Mako Hill, Benjaminevans82, Beno1000, Bensin, Bento00, Berean Hunter, Berland, Bestadvocate, Bevo, BhaiSaab, BiT, BigDunc, BigglesPiP, Bilalakhtar96, Biscuittin, Bjcubsfan, Bk109, Blainster, Blake-, Blaxxter, Blehfu, Blooper4912, BlueCanary9999, Bobbo, Bobkevin713, Bobnorwal, Bobo 178, Bodsda, Bojan, Bomac, Bombadil1986, Bonadea, Bongwarrior, Bonius, Bootstoots, Boricuaspacepr, Boudhayan Gupta, Boxalox, BradBeattie, Bradkoch2007, Brax365, Brazil4Linux, BreadBasket, Brianmarx, Brianreading, Brianrusso, Brianski, Brighterorange, Brinerustle, Brisvegas, BrokenSegue, Bruce89, Bruce89palm, Brucevdk, Btate, Bubba73, Bubzie5, Bungalowbill, Burgundavia, BurnDownBabylon, Bvc123q, Bvidinli, Bz22-NJITWILL, C'est moi, C.Fred, C628, CFeyecare, CGA, CMBJ, CUSENZA Mario, CWii, Cadsuane Melaidhrin, Cakez, Caltas, Cam493, Can't sleep, clown will eat me, CanManJam, Canaima, Candorz, Capricorn42, Careyo, Carlos Ferrer, Catalaalatac, Cbm, Cedrus-Libani, Celestianpower, Cenk Manisalı, Centrx, Ceros, CesarB, Ceyockey, Chainz, Chameleon, ChaosData, Chappy84, Chargh, Charles Gaudette, CharlesC, Chavash, Chealer, Chenzw, Cherry Cotton, Chicagoman, Chickenofbristol, Chillum, Chimin 07, Chiok, Chris the speller, Chris83, Chrisjniles, ChristTrekker, Christian75, Christofer C. Bell, Christopher.murch, Christthi, Chun-hian, Ciao 90, Cindamuse, Ciphergoth, Ck135667, CloudNine, Cnidus.eudoxus, Cobi, CodeOne, Cogiati, ColdFusion650, Colejohnson66, Colin Watson, Colonies Chris, Commander, CommonsDelinker, Computerdude33, ConMan, Conalhamill, Connormah, Conti, Coolcom, CoolingGibbon, Coralmizu, Corevette, Coughinink, CountZer0, Cowplopmorris, Cozzerob, CrackDragon, Craigular, Creidieki, Crispmuncher, CristiS, Cueax, Cujocat22, Cutter, Cyan.aqua, Cybercat, Cybercobra, Cyclonius, Cyrus XIII, DJSnuggles, DalekClock, Dan Atkinson, Dancraggs, Daniel 1992, Daniel.Cardenas, Danieljackson, Danielmiester, Danno uk, Danski14, Danutzdobrescu, Dariustriplet, Darkodelta, Darkwind6000, Darrelljon, Darxus, David Fuchs, David Latapie, David Rolek, DavidBailey, Davidweiner23, Dawnseeker2000, Dayofswords, Daytonlowell, Dbolton, Dcljr, Ddxc, Debresser, Decltype, Deeahbz, Deflective, Defza, Delan Azabani, Delirium of disorder, Delldot, Deltabeignet, Demrit, Denis C., DerHexer, Destin, Dfrg.msc, Dharit-NJITWILL, Dhughes, DiThi, Digitize, Dirk Gently, Discospinster, Diti, Djibouti,Djibouti, Djsubterrain, Djyang, Dkatten, Dollarwallet, DonovanE, DotMacintosh, Doulos Christos, Dovi, Download, Dr. Weazel, Dr. Weazel, Dr. Zed, Dr.Koljan, Dragon 280, DreamsReign, Drjsc, Drmies, Drumond.douglas, Dsavi, Dsym, Dueyfinster, Dungodung, Dusty777, Dvdkid919, Dwo, Dylnuge, Dzubint, Déjà Vu, Eanschuessler, Easyas12c, EdgeOfEpsilon, Eeblefish, Eekster, Egret Of Regret, EhJJ, Ekhcsub, Eleventytwelve1112, Eli Loewen, EliasAlucard, Elipongo, Elizium23, Eloquence, Eltouristo, Emerson1234567890, Emurphy42, Emx, Energizing mouth spray, Enk, Enochlau, Enquire, Epbr123, EpicFlame, Epixstudios, Eraserhead1, Erdal Ronahi, Ergosteur, Erianna, Eric, Eric B. and Rakim, ErkinBatu, Estel, Estemi, Estherschindler, Eug, Euphrosyne, EvanProdromou, Evanh2008, Evanreyes, EvelinaB, Evice, Evil Monkey, Ewald, Exercisephys, Exeva, EyeRmonkey, FF2010, Fakewcfrog, Falcorian, Favonian, Fcendejas, Feedmecereal, Felos, Fhfhfhfhf, Fholson, Fixman, Fl, FleetCommand, Fletcher, Flik, Fluteflute, FogMasher, Forrest Ivie, Fotoni, Foxj, FrF, Frafa, Frak, Frap, Fraxoj, Frecklefoot, Freedom to share, Freeworld5, Freqk, Freshgavin, Friendly Ed, Friviere, FrozenPurpleCube, Frymaster, Fs, Fsolda, Fubar Obfusco, Fuhnny, Fullerene, Funkruf, Furrykef, Furste, Future Perfect at Sunrise, Fuzheado, Fxer, GNosis, GSlicer, Gab86, Gaelle Desbordes, Gaius Cornelius, Gala.martin, Gamaliel, Ganesh.rao, Garazy, Gareth Aus, Garglebutt, Gargleflump, Gary King, Gdavidp, Gdupont, Geoff918, Geopgeop, Georgeryp, Geronimooo, Ghettoblaster, Ghost Freeman, Giftlite, Gigs, Giltramirez, Gingerdan, Gioto, Girialam, Giuliopp, Glider87, Glolrk2007, Gogo Dodo, Googleplex12, Googolit, Gordmoo, Gowtham vmj, GozzoMan, Gr0ff, Gr33ndata, Grandscribe, Greenrd, GregAsche, GregorB, Grendelkhan, Greyweather, Gronky, Grosscha, Group29, Groupmsl, Gruepig, Gth759k, GuelphGryphon98, Guilherme Paula, Gunny01, Gurch, Gustavb, Gutsygibbon, Guyjohnston, Gürkan Sengün, HAL(Old), Hadlock, Haha169, Haham hanuka, Hairy Dude, Hal7000, Hamish2k, Hanagata San, Hankwang, Hansjorn, HappyInGeneral, Happyrabbit, Hapsiainen, Haris.tv, Harizotoh9, Harryboyles, Harumphy, Haseo9999, Hawkian, Hawtspinz, Hdtvinfoeu, Heirpixel, Hello. I'm new here, but I'm sure I can help out., Herakleitoszefesu, Heron, Highpeakster, Highpriority, Hiplibrarianship, Historyofmadness, Hn, HoCkEy PUCK, Hobophobe, Hoo man, Hopugop, Hounddog32, Hp notebook, Htonl, Hu12, Huseyx2, Hvn0413, Hyad, Hydrargyrum, Hydriz, I need a name, ICEAGE, IDontBelieveYou, IMSoP, IWantFroyo, Iamorgassvdhjl, Ian Lancaster, Ianjm, Ianml, Icebuster7, Id1337x, Ida Shaw, Iggy Koopa, Ike9898, Iketsi, Ikonoko123, Ikrshnan, Il Palazzo-sama, Ilya, Ilya Voyager, ImMrDrake, Imaginos28, ImpliedFibre, InMyHumbleOpinion, Index2, Infocidal, Int21h, InverseHypercube, Inzy, Ironic editor, Isaac Dupree, Isilanes, Itsmine, Izanbardprince, J.delanoy, JAGFin1, JCDenton2052, JGMacon, JGXenite, JHunterJ, JJC1138, JLaTondre, JPat, Jackmatty, Jackol, JackyR, Jacobolus, Jaduncan, Jaffa phil, Jagoly, Jake7401, Jakeyboyzang, James Foster, James90064, Jamie Kitson, Jan Hofmann, JanCeuleers, Jancikotuc, Jango4, Jarble, JasonDClinton, Jasper Deng, Jatos, Javitomad, Jaycee55, Jeandré du Toit, Jebethe, Jeffq, Jeffwang, Jeliekbigbusanas, Jeltz, Jeni, Jerebin, Jeremy Visser, Jerryobject, Jeshan, JetBlast, Jhertel, Jhmac77, Jim1138, Jiminy Krikkitt, Jmedina399, Jnw222, JoanneB, JoeSmack, Joeblakesley, Joelr31, Joewski, Jogloran, John, John Lucas, John Nevard, John Nixon, Johnh123, Jojit fb, Jon Harald Søby, Jonas Jørgensen, Jongleur100, Jonik, Jonjesbuzz, Jonne, Jonny2BeGood, Jonnyclavin1616, Jons2006, Joo-hyun1, Jopo, JorgePeixoto, Josejuan05, Josh the Nerd, Joshsmith65536, Joshuaes, Jouva, Joy, Jpalanca, Jpkotta, Jpvinall, Jrg dnn, Jrtayloriv, Julesd, Julianbrelsford, Justcorrected, Justthai, K;;m5m k;;m5m, KAMiKAZOW, KD5TVI, KDesk, Kaobear, Kareeser, Karl Dickman, Karlwick, Katalaveno, Katanzag, Kayedj, Kbabeatbox, Kbdank71, Kc8ukw, Keenan Pepper, Kenny Strawn, Kernkernlee, Kesla, Kev, Kevin S., Kevin908, KevinCole, Khoikhoi, Kiand, KimDabelsteinPetersen, Kimchi.sg, Kinema, King Of Cake, Kinghajj, Kingoftherings, Kingpin13, Kjbyumwbvf, Kjoonlee, Kkm010, Kl4m, Kl4m-AWB, Klungel, KnowledgeOfSelf, Koavf, Koertefa, Koman90, Koolabsol, Korg, Koumed, Kozuch, Kragnerac, Krtek2125, Krukouski, KrzysztofKlimonda, Kslays, Ktec776, Kulshrax, Kungfuadam, Kuru, Kurykh, Kwamikagami, Kwertii, Kyorosuke, L3iggs, LFaraone, LGiunruux, LadyPhi, Landroni, Lantay77, Laos, Lar, Largon, Laserlight, Latiligence, Laurens-af, Laurentrepond, LauriO, Laurynas126, Lavenderbunny, Law, Le Enfente Orange, LeedsKing, Lengau, Lengyeltom, Leo dubiam, Leolaursen, Leonard^Bloom, Lester, Letdorf, Lewys93, Liftarn, Lightedbulb, Ligulem, LilHelpa, Limulus, Linfocito B, Liorkaplan, Lisamh, Lishy Guy, LizardJr8, Lkt1126, Llewelyn MT, Localzuk, Locos epraix, Lopifalko, Lord Darth Vader, Lorg, Lorian, Lost*in*color, Lotje, Lovysinghal, Lproven, Ltmboy, Luca Ghio, Lucasoutloud, Lucid, Lucifer Spam, Lucohami, Luigii249, Luisgarciadiaz, Luna Santin, Lurker, Luxdormiens, Lwdjaymac, Lyo, Lysander89, M gol, M5, MER-C, Maaparty, Mabeenot, Mac, Mac Lover, Mac v, Macewan, Macrowiz, Mad Pierrot, Madduck, MadeOfEyelashes, MadnessInside, Madskillzlol, Magister Mathematicae, Mahanga, Mahjongg, Makima, MalTsilna, Malomeat, Man with two legs, Mandavi, Mann jess, Manu.valentin, Manzabar, Marc-André Aßbrock, Marcelivan, Marcelomp3, Marcorosas1991, Marginoferror, MarienZ, Mark Bergsma, Mark Renier, Mark waugh, Marko Knoebl, Marko knoebl, Marminnetje, Marozols, Martarius, Martin451, Martinor, Martinultima, Marty-109, Marudubshinki, Masatran, Materialscientist, Mateusza, Matpe815, Matrixtom, Matt Crypto, Matt Crypto on Carts, Matt me, Matt.Hoy, MattGiuca, Mattb90, Matthewrbowker, MattieTK, Mattl, Mattva01, MaxSmith777, Maxheadway, Mazin07, Mbarne, McCart42, Mcornils, Md towhid, Mdd4696, Mdy66, Mdz, Meathunch, Mecanismo, Medovina, Meeples, Melancholie, Mentifisto, Merbenz, MetaEntropy, Metal.lunchbox, Metroidpal45, Mewtu, Mhoenig, Michael A. White, Michael963, Michaelzeng7, MichealH, Michelleem, Mike Rosoft, Miko3k, MilFlyboy, Miles Monroe, Mini-Geek, Miniwonder, Minna Sora no Shita, Mipadi, Miranda, Mirv, MisfitToys, MisterSheik, Mitchumch, Mitsuhirato, Miwunderlich, Mkpeeters, Mm40, Mnemo, Momet, Momposi, MontelEdwards, Mordomo, Mortense, Moxfyre, Mpt, Mr.Aks, MrMarmite, MrNate, MrOllie, MrVibrating, Mrand, Mrzaius, Ms2ger, Msikma, Mssnlayam, Mualif02, Muntuwandi, Muro de Aguas, MuzikJunky, MyNameIsNotBob, Myredroom, N.Nahber, N5iln, NCLI, NSLKFNBKLFSNBFSLKBNKLSFNBKLSF, Nabbia, Nachoman-au, Nadiratharanga, NapoliRoma, Napsterbater, Nate Silva, Nathanhillinbl, Nattgew, Naught101, Ncmvocalist, Ndenison, Nealmcb, Neil9999, NeilFraser, NeoChaosX, Nerdpill, NerdyNSK, Nerval, Nessup, Netalarm, Neurolysis, Neustradamus, NeutralPoint, Neutralaccounting, Nextil, Nguyen Thanh Quang, Nick, Nickrud, Nicolas Love, Niduzzi, Nigelj, Nikai, Nileshbandhiya, Nilocia, Ninnnu, Nixeagle, NoaHall, Noelitonoelito, Noldoaran, Nonagonal Spider, Noneexistent, Nopetro, Norf star, Norm, Norm mit, Northamerica1000, Nrbelex, Ntse, Ntsimp, Numbo3, Nuwewsco, Nvdokh, Oberiko, Obueno, Ojan, Okungnyo, Oldag07, Olenick, Olivaro, OliverTwisted, Om26er, Omaryak, Omegatron, OnesimusUnbound, Oni Lukos, Oobermicro7, Opensourcefuture, OranL, Orbital fox, Organ123, OsamaK, Oscarnet75, Oscarthecat, Ossguy, OsvaldoGago, Otternox, OuroborosCobra, Owen, Owlbuster, Oxymoron83, Ozguy89, Ozzie The Owl, PET, PGScooter, Pablo Alcayaga, Pablo-flores, Pamri, Panfakes, Pankkake, Panosdotk, ParticleMan, Pasado, Passeportout, Patato, Patlex, PatriciaJH, Paul-loves-cake, Paulyche, Paxsimius, Pbryan, Pcu123456789, Peachey88, Pelago, Penguinjuice, Percy train, Personalinformation, Perspectoff, Peter S., Peteturtle, Peyre, Phalacee, Phatsphere, Pheneviox, Phible, Philip Trueman, PhnomPencil, Physadvoc, Pile0nades, Pinethicket, Pinkmanlovemouth, Pion, Pitel, Pixel ;-), PizzaMan, Pladask, Pmedema, Pne, Pnm, Pointblankstare, Poppafuze, Power piglet, Pragmatik, Praka123, Prara, Prcr, Prfe FCE CPE cambridge, PriceChild, Prodego, Propaniac, Prosallus, Prowikipedians, Pschemp, Pseudo daoist, PseudoOne, Psubhashish, Ptcamn, Ptomes, Qrk, QuantumEleven, Quantumelfmage, Quebec99, Quercusrobur, Quiddity, Quispiam, QuiteUnusual, Qwertyus, R'n'B, R000t, R3m0t, RCM2309, RCalixte, Racecar56, Radu.cotescu, Ragnarok Addict, Rahulthewall3000, RainbowOfLight, Ramvi, Randomengine, Raoulharris, Raul654, Ravemaker, Raven in Orbit, Raysonho, Razorflame, Rchamberlain, Rdwallis, Reach Out to the Truth, Rebroad, Red and Green and Brown and Blue..., RedxelaSinnak, Reedy, Regancy42, Reijimitsurgi, Reinthal, Remember the dot, Remixz, Reparator, Repetition, Respawn2003, Restecp, Retired username, RevDan, RexNL, Rfl, Rgaugyan, Rhe br, Rholton, Rich Farmbrough, Richard Maxwell, Richard0612, Richardcavell, Rick Browser, Rickjpelleg, Riksweeney, Rjanag, Rjshade, Rjwilmsi, Rlmorgan, Rmcfanatic, Rmky87, Robert Brockway, Robert Moyse, RobertG, Robust Physique, Rodlitenstuga, RolandH, RolandKwee, Roly Perera, Romeu, Romperomperompe, Rosiestep, Rotring, Roujo, Rowan187, Rpgdude, Rprpr, Rsg, Rtc, Rubinstein, Rugby471, Ruiz, Rythie, Ryulong, S Carpenter, SCARECROW, SF007, Sagaci, Salvio giuliano, Sam Barsoom, Sam Burke, SamJohnston, Sampritb, Samsara, Samuel Grant, Sanderd17, Sango123, Saoshyant, SarekOfVulcan, Sasank, Sashin, Savidan, Sceptre, SchuminWeb, SciCorrector, Scientus, Scjessey, Scorpus57, Scott Ritchie, Scott3, Scrantoncity, Scrdcow, Scytale on, Sdbinwiiexe, Sdornan, Sdrtirs, SeanCollins, SeanLegassick, Seans Potato Business, Seantrac, Secleinteer, Secretlondon, Servant Saber, Shadow1, Shadowjams, Shae-bichette, Shagywashere, Shardz, ShaunL, Shervinafshar, Shivray, Shizhao, Shreevatsa, SickTwist, Siddhanathan, Sikon, SilverWolf7, Silverhand, SilvervWolf7, Simesa, Simple Bob, Simxp, Sin-man, Sjasty, Skarebo, Skybon, Skysmith, Sladen, Slashme, Slavenweb, Sleeping123, Slicky, Slusk, Smitty, Smyth, Snap monkey, Snarius, Snigbrook, Snori, SnowPrak, Soeb, Soliloquial, Some Person, Something1235678, Sommers, Sonicology, Sonicsuns, Southwest, Spalding, Spazturtle, Spellmaster, Spencer, Sperling, Spike the Dingo, Spitzl, Spoon!, Spoondish, SpuriousQ, SpyMagician, Srborlongan, Sreyan, Srippon, Srjsignal, Ss73666, Stausifr, Steel, Stephantom, Stephanvaningen, Stephen MUFC, StephenDennis, StephenH, Stephenchou0722, Stephenmac7, Stevage, SteveSims, Steven Walling, StevenT1, Stikonas, Stormtalon, Stormwatch, Strangnet, Subiet, Subsonix, SudoGhost, Sujit900, Sundar2000, SuperGerbil, SuperMoonMan, Superkirbyartist, Superm401, Surfaz, Suruena, Sushiflinger, Sv1xv, Svaksha, Sverdrup, SwineFlew?, Swordkeyblade, Syherp, Sykil, Syndicate, Synergy, T. Moitie, T34, TFOWR, Tacoboy42, Taejo, TakuyaMurata, Takuzinis, Tami.hearn, Tangotango, Tarehart, Tavilis, Tavrian, Tazpa, Tcalight, Tcrow777, Tdfjytfjukckjhg, Team4Technologies, Techman224, Teddks, Tedickey, Tedwardo2, TempestSA, TerminalPreppie, Terrifictriffid, Tetrahedron93, Thanhbv, The Monster, The Signal, The Thing That Should Not Be, The bellman, The rooker, TheAMmollusc, TheBilly, TheIntersect, TheSameGuy, TheSmuel, TheTechFan, Thebestofall007, Theda, Thedjatclubrock, Themagicmanfromtrent, Thenickdude, Theo148, Thephotoman, Theproc64, Theunixgeek, Thialfihar, Thingg, Tholme, Thomasmallen, Thorpe, Thryduulf, Thumperward, Thunderbird2, Thunderpenguin, Tiddly Tom, Tide rolls, Tim1988, TimingClock, Timothylord, Tiptoety, Tjwagner, Tm T, Tnchris, Tobias Conradi, Tobias K., Toehead2001, Tohd8BohaithuGh1, Tom harrison, TommyKiwi, Tomwsulcer, Tony1, TonyYarusso, Toosage4u, Torrg, Toussaint, Towsonu2003, Traxs7, Trent the god, TrevorLSciAct, Trilobite, Trmnshwrx, Troy 07, Trusilver, TrygveFlathen, Tufriast, Tulkas, TurboForce, Tuxstar, Tuxthepenguin933, Twermund, Twinxor, Twipley, Txuspe, Typhoon, Ubernostrum, Ubunterooster, Ubuntutorial, Ugriffin, Ukigison, Ulf Abrahamsson, Ullas13, Umapathy, Unamed102, Undeference, Uni4dfx, Unohustr, Upapilot, Urashimataro, Urhixidur, Ushimitsudoki, Utcursch, Uzusan, VARGUX, VBGFscJUn3, VShaka, Val42, Valaqil, Ventolin, Versus22, VertetNoir, Victoria99, VinayNeekhra, Vincentt, ViperSnake151, Vivekvc, Voidvector, VoluntarySlave, W.F.Galway, WJetChao, WPIsFlawed, Walter Görlitz, WalterGR, WanderingHermit, Waqas1987, Warfvinge, Wassersturm, WatchAndObserve, Wayne Slam, WazzaMan, Wellivea1, Wellparp, Werson, Wfox, WhatisFeelings?, Widefox, WikUWerHere, Wiki alf, WikiEdit, Wikicleaner2, Wikieditor96, Wikien2009, Wikievil666, Wikijeff, WikipedianMarlith, Wikitanvir, Wikitester, Willbeonekenobi, Willv624, Willyarnold, Wilstrup, WinterSpw, Winterst, Wiretse, Wissons, Wizardus13, Wo.luren, Woohookitty, Wouterstomp, Writtenonsand, Wvoutlaw2002, Wōdenhelm, X-Fi6, XDanielx, Xaminmo, Xcentaur, Xeku, Xela, Xelgen, Xenon54, Xequence, Ximian99, Xiphiidae, Xoxjoy, Xpclient, Xpertise, Xyzzyva, Yearofthedragon, Yk4ever, Ynnad73, Yohisha, YorkBW, Yoshi:3, Ysaimanoj, Yurik, Yworo, ZZyXx, ZachPruckowski, Zagadka77, Zappa-king 1337, ZeWrestler, Zen Clark, ZerglingChamp, ZeroUm, Zeus, Zib redlektab, Zidonuke, Zig973, Zombifier, Zondor, ZooFari, Zundark, Zvn, Zwolf101, Ævar Arnfjörð Bjarmason, පසිඳු කාවින්ද, ⚗, 図, 2551 anonymous edits

Microsoft Windows *Source*: http://en.wikipedia.org/w/index.php?oldid=504220163 *Contributors*: -Majestic-, 03jmgibbens, 1(), 16@r, 1nt2, 2mcm, 2toy mora, 62.253.64.xxx, 67773732TYU, 68DANNY2, 9ms, A Raider Like Indiana, A gnome, A-giau, A7x, ACCOM2222, Aanhorn, Abhishikt, AbsoluteFlatness, Academic Challenger, Acroterion, Adam Mirowski, Adamacious, Adamodell, Admin@pcrevs.com, Ae-a, Aeæ, Afro Article, Ageha Winds, Ahodacsek, Ahoerstemeier, Aidan W, Aido2002, Aihtdikh, AimalCool, Airplaneman, Ajm81, Akamad, Akhristov, Akira-otomo, Aksi great, Albert0057, Aldie, Alegoo92, Alemily, Alfio, AlistairMcMillan, Allstarecho, Almafeta, Alpha 4615, Alphathon, Alphius, Alshaheen15, Althepal, Am088, Amakuru, Amcfreely, Amrykid, Andre Engels, Andrevan, Andrewtechhelp, Andros 1337, Andy16666, Andyh2, Angeljon121, Ann Stouter, Anog, Anonymous Dissident, Anonymous56789, Antandrus, Anthall1991, Antimatter15, Antique Rose, Anville, Anþony, Aomarks, Aqair, Aranel, Arch dude, Archer3, Archivist, Aria1561, Arnoldkul, Arthena, Ashdurbat, AstroNomer, Astroview120mm, Aude, Audrius u, AussieLegend, AvantgardeMVC, Aveilleux, Avenue, Avihut, Avocado27, Ayjay1545, Azrael Nightwalker, B, BNSF Man, BUzTeD, Badwolf2212, Bakery2k, Banes, Barek, Barrera marquez, Bartosz, Bbq man, Bbriggs1, Bdoserror, Bdshort1, Beao, Beinsane, Ben-Zin, BenBen1234, Benandorsqueaks, Benc, Berek, Bergsten, Betelgeuse, Bevo, Bhadani, Bibliomaniac15, Big Bird, Big Booger, BigCow, Biker Biker, BillG, BillWSmithJr, Billyswong, Bimmerosx, Binsurf, Binzisimpsons, Bissinger, Blackanddarkness, Blackcap95, Blackcats, Blakkandekka, Blaxthos, Bleakcomb, Blobglob, Blowdart, Blubberboy92, BlueCaper, Blueforce4116, Bluestriker, Bo98, Bobo192, Bobwrit, Bogods, Bollyjeff, Boogie314, Boris Allen, BostonMA, Boylett, Brandizzi, Brandon Brown, Brian0918, BrianGo28, BrianRecchia, Brianski, Brisvegas, BrokenSegue, Bubba hotep, Buchanan-Hermit, Buddha24, Bunnyhop11, Burntsauce, CRFWNY, CWenger, Cacophony, Cadiomals, Caffelice, Caltas, Camembert, Can't sleep, clown will eat me, Canadian-Bacon, Candamir, Candorwien, CanisRufus, Caper13, CaptainVindaloo, Casper10, Casper2k3, Cayindra, Cbrown1023, Cburnett, Ceeon, Celestra, CesarB, Ceyockey, Cff12345, Cgnabod, Cgs, Charles Gaudette, Charles dye, Chazz, Cherkash, Chille, Chmpoure, Chocolatemilk94, Choppie3000, Chowbok, Chris 73, Chris Pickett, Chris the speller, Christian List, Chriswoz, Ciao 90, CitronManden, Cityscape4, Cjcamilla, Cjordan93, ClamDip, ClockworkTroll, Clyde1998, Cncccer, Cncxbox, Codificate, Codyblevins, Coffee, Colejohnson66, Colin Hill, Cometstyles, Commander Keane, Compfreak7, Computer boy2, Computerdan000, ConCompS, Coniosis, Conman23456, Conversion script, Coolcaesar, CoolingGibbon, Corporal clegg48, Courcelles, Cp111, Cpiral, Crazycomputers, Crazyromo, Crem23, Cremepuff222, Cristan, Crpietschmann, Cryptic, Cstanners, Ctbolt, Cuvtixo, Cvinoth, Cwolfsheep, Cyktsui, Cynical, Cyrius, D, D thadd, DHN, DOHC Holiday, DaDrumBum, DalekClock, Damian Yerrick, Dan D. Ric, Daniel5127, DanielRigal, Danny Beaudoin, DarkFalls, DarkHorizon, Darranc, Darremon, Darthnader37, Dasani, DataMatrix, Databases, Davelane, David Biddulph, David Gerard, David1409, Davidjk, Dcandeto, Dcolvin, Dcshank, Deanhowell123, Debackerl, Deeahbz, Dehumanizer, Delirium, Delldot, Deltabeignet, Demmy, Denniss, Derek Parnell, Deryck Chan, Dethomas, Deus2, Dieboybun, Diegogrez, DigbyDalton, Digita, DigitalLife, Dimre01, Dina, Dinhtuydzao, Dinjired, Dinnerface, DinosaursLoveExistence, Dipankan001, Disorganisation Man, Dj789, Djegan, Djhybrid117, Djmasala, Dmerrill, DmitryKo, Dog1818, Dojarca, Dolphinn, Downloaddude1258, Dp462090, Dragar Gt, Dragon 280, DragonflySixtyseven, Dragonhelmuk, Dstln, Dtech, Dudesleeper, Dungodung, Dust Filter, Dustin gayler, Dwheeler, Dylan Flaherty, Dysprosia, Dzubint, EH74DK, Eagleal, Ed g2s, Edgar181, Edgarde, EdgeOfEpsilon, Edward, Eelamstylez77, Eggsacute, Egil, Eivindsol, El Dominio, Elassint, Eliotwiki, ElliotAdderton, Ellmist, Elm-39, Eloquence, Elwood00, Emacsuser, Enco1984, Endofskull, Enigmaaaaa, Enno, Enochlau, Ente75, Eraserhead1, Ergosteur, Esanchez7587, Escape Orbit, Esebi95, Evice, Evil Monkey, Evil saltine, Evosoho, Ewlyahoocom, Exodite, F80, FF2010, Fabsss, Faisal.akeel, Faithlessthewonderboy, Falco McCloud, Fantasy, Faradayplank, FatalError, Fatjoe151, Favonian, FayssalF, Fbv65edel, Fdp, Felipe Aira, Fennec, Ferkelparade, Ffx, Fiskegalen92, Fitzebwoy, Fiver2552, Flamurai, Flanakin, FleetCommand, Flying Bishop, FlyingPenguins, Foo1995, Fowl2, FrancoGG, Frap, Fraslet, Freakofnurture, Frecklefoot, Fred Bradstadt, Frederik.Questier, FreeKresge, Freeeekyyy, Frenchman113, FreplySpang, FrozenPurpleCube, Frozenport, Fuck You, Funandtrvl, Furrykef, Fuzheado, GDonato, GNUtoo, GSK, Gabrielpokemon, Gadfium, Gaius Cornelius, Galactor213, Galoubet, Galwhaa, GamerXp, Gamerzworld, Gaming&Computing, Ganeshotaku, Ganfon, Gang65, Gaodifan, Gardar Rurak, Gary King, Gary Kirk, Gavin.perch, Gazpacho, Ged UK, Geek45, Geoffspear, George Adam Horváth, Georgia guy, Geraki, Gerbrant, Ghettoblaster, Gilbertogm, Gilliam, Glen, Gnepets, Goatasaur, Gobonobo, Gogo Dodo, Goldom, Googler459, Googlesucks56789, Gosub, Graham87, GrandPoohBah, Grayshi, Gregory Heffley, Grin, Grm wnr, Groggy Dice, Gronky, Grunt, Gryllida, Gscshoyru, Gtdp, Gunnar Guðvarðarson, Gustyfalcon, Guy Harris, Guyjohnston, HJ Mitchell, Haakon, Haeleth, Hahafatpeople, Hairchrm, Hall Monitor, Halsteadk, Hammersoft, Harryboyles, Haseo9999, Hbomb phd mom, Hdt83, Heapchk, Hebrides, Helixblue, Hello32020, Hendrixski, Henry W. Schmitt, HereToHelp, Hirosho, Holbred, Howardjp, Hungupbg, Hurtstotalktoyou, Husky, Hydr, Hypnoticcyst, I have wood, I8189720, IGod, ILOVELOL32, IMSoP, IXella007, Iain99, Ialsoagree, Ian Pitchford, Ianjones50, Ichangefacts, Iggymwangi, Ilaiho, Illegal Operation, Illyria05, Ilya, IlyaHaykinson, Ilyanep, Indefatigable, Indon, Injust, Interframe, Ioprwe, Ipodsocool, Iridescent, IsUsername, ItsProgrammable, Ivan Pozdeev, Ixfd64, J Di, J36miles, JCarriker, JIP, JLaTondre, JWSchmidt, JYOuyang, JYolkowski, Jacob Hnri 5, JacquesStrap, Jaericho, Jake Wasdin, Jamieostrich, Jan Hofmann, Jasper Deng, Jauerback, Jaxl, Jcbparry, Jcurtin, Jdlowery, Jdm64, Jeffwang, Jeh, Jennie--x, Jeremy Visser, JeremyA, Jerryseinfeld, Jerrysmp, Jesant13, Jesse V., Jesse Viviano, Jesus5555, Jesus764, Jesuss, JettaMann, Jevel66, Jiffy, Jigs41793, Jimmi Hugh, Jimothytrotter, Jimthing, Jiy, Jketola, Jkonline, Jmath666, Jmjglick, Jmoynihan08hm77, JoJan, JoThousand, Joanjoc, JoanneB, Joe2832, Joejoejo, Joffeloff, Johann Wolfgang, John, John Ericson, John Goettle, John Quincy Adding Machine, JohnJamesWIlson, JohnOwens, Johnclow13, Johnleemk, Johnmc, Johnuniq, Jojit fb, Jon vs, JonasL, Jonatan Swift, Jondel, Jonghyunchung, Jopxton, Jordan015, Jordanyoung17, Jose Concepcion, Josh the Nerd, Joshua Issac, JoshuaArgent, Joyous!, Jpoke89, Jrcure, Jsmethers, Jtc, Ju66l3r, JuJube, JuanC08, Julian Mendez, Jumbo Snails, Justanother, Jyoz, Kakomu, Kar.ma, Karimarie, Karlo2002, Karnesky, Katarighe, Kaycubs, Kaysov, Kbh3rd, Kbolino, Kelly Martin, Kenny Strawn, Kesla, Kevinkor2, Keyser Söze, Khym Chanur, Kierenj, Kigali1, KingpinE7, Kinneyboy90, Kipholbeck, Kissmeplease, Kittyhawk2, Kjhonsa, Kkm010, Klingoncowboy4, KnowledgeOfSelf, Knutux, Kohlmalo, Kolonuk, Koman90, Kop, Korossyl, Kozuch, KrakatoaKatie, Krawi, KrisBogdanov, KumardipSarkar, Kungfuadam, Kungming2, Kurieeto, Kushared, Kwen, Kypr8, L Kensington, La goutte de pluie, Lacrimosus, Lahiru k, Lantay77, Larry laptop, Lbs6380, Lcarsdata, Lead$peaker, Leandrod, Lectonar, LeeHunter, Leif, Leithp, LeoNomis, Letdorf, Leuko, Liberty Miller, Lightmouse, Linear88, Linuxbeak, Linuxerist, LionKimbro, Llalala, Localzuk, Loismum, Lonaowna, LonelyBeacon, Looksliketrent, Lord British, Lordalpha1, Lotje, Lou.weird, Lowellian, Luce007, Luk, Lumpbucket, Luna Santin, Lupin, M Johnson, MADCastro2012, MCC, MER-C, MERC, MJ94, MJGR, MK8, MMMEEE, MZMcBride, Ma8thew, Mac, MacGeekGuy, Macurry, Madcow 93, Maester mensch, Magister Mathematicae, Mailer diablo, MainFrame, Malo, Man123123, Marceki111, Marcok, Mark, Mark85296341, MarkGallagher, MarkSpearmint, Markhurd, Markie2, Markustwofour, Martarius, Martyx, Marudubshinki, Marx Gomes, Marysunshine, Massysett, MasterCole, Mastershake86, Mathwizxp, Matt Crypto, Matt.forestpath, MattGiuca, Mattarata, Mav, Mawcowboybillsbrick7, Max Schwarz, Maxim Masiutin, Maximus Rex, McLovin34, Mcwatson, Meelar, Mehran, Mephiles602, Metalim, Methanegas, Metz2000, Mguy77, Mgw854, Michael Angelkovich, Michael Ray, Michaeldadmum, Michaelpremsrirat, MichealH, Midgrid, Midkay, Midnight Green, Mike4ty4, MikeZuniga, Mild Bill Hiccup, Milfman, Minesweeper, MinorItem, Miros 0571, Mithent, Mjpieters, Mkbashier, Mkdw, Mlscdi, Mmoll21m, Mmoller, Mms, Mo0, Modster, Mokk123, Monedula, Monothleft, MontyB, Moonknightus, Moonridersaregay, Moonwolf14, Mr Bartels, Mr. Lefty, Mr. XYZ, Mrdelayer, Mrleewilliams, Ms2ger, Mugunth Kumar, Mulad, Mushroom, Mvmarier, Mwtoews, Mxn, Myanw, Nachmore, Naddy, Naltrexone, Nanshu, NapoliRoma, Nathan nfm, Nathanlilienthal13, Neckername, Neelchauhan, Neo The User, Neo2256, NewGuy4, Newfraferz87, Nick2253, Nickenzi, Nicolas Love, Nimc, Nintendog master 54, Nishkid64, Nixdorf, Njan, Nmnmnm, NotinREALITY, Noypi380, Nukleartoaster, Nurg, O, Odd bloke, Off!, Okwestern, OlEnglish, Olioster, Ollie the Magic Skater, Omicronpersei8, Ondrejsv, Oneupthextraman, Opagecrtr, Opelio, OregonD00d, Oscarthecat, Ost316, OtherPerson, Otter Mii-kun, OwenX, Oxymoron83, Ozherb, PPGMD, Packard Bell Legend, Padsquad43, Pakaran, Patrick, Patrickweeks007, Patrolman89, Paul Stansifer, Paulfp, Pavel Vozenilek, Pcuser42, Persian Poet Gal, Person1936, Petaluma Paranormal, Peter Grey, Peteturtle, Petrwiki, Peyman4u, Peyre, Pgiii, Pgk, Phanton, PhelperRF, Phil Sandifer, Philip Trueman, Phoogenb, Pigman, Pile0nades, Pinikas, Piratesmack, Plugwash, PoWwOw123, Pokyrek, Poor Yorick, Poweroid, Prjrv787, Prolog, Propound, Pschulz01, PsyMar, Puffin, Puramyun31, Pursey, QUAZWRATH, Qantasplanes, Qudder, Quinsareth, Quoladdie, Qwerty124gg, Qwitchibo, Qxl32, Qxz, R'n'B, RB972, RTC, RadioActive, RadioFan, Radon210, Raja99, Ramanpotential, RandomP, Rasmasyean, Rbuj, Rchamberlain, Rcmouse1010, Rd syringe, Rderijcke, Rdsmith4, Reallikeunreal, Rebroad, Red Dalek, Red Director, RedKlonoa, RedXII, Redvers, Reilly, Reisio, Remember the dot, Repetition, Resplendent, RevolverOcelotX, Rgbc2000, Rhesusmonkeyboy, Rhobite, Rhtc, Rich Farmbrough, Richard Lotspard, Richard626, Ricjl, RickK, Riluve, Rincewind32, Rjecina, Rjwilmsi, Rl, Rlevse, Robartin, Robbie940, Robert Buzink, Robert H, Robert Xia, RobertG, Robertd, Robinhw, Rocastelo, Rocket71048576, Rockymountains, Rodeo90, Roketjack, RolandH, Rolandhelper, Ronark, Rsantmann, Rsm99833, Running, Rursus, Rutherfordjigsaw, Ruud Koot, Rwwww, Ryan Norton, Ryan Postlethwaite, Ryan t moua, S0aasdf2sf, SF007, SHARD, SHeumann, SNIyer12, ST47, SYSS Mouse, Sae1962, Sam Hocevar, Samuel, Samuel Blanning, Samvscat, Sango123, Sasquatch, Sauronjim, Scepia, SchmuckyTheCat, SchuminWeb, Science4sail, Sciencegeek10188, Scientus, Scifiintel, Sciurinæ, Scorpiona, Scottymoze, Scriberius, Scuiqui fox, SeanMack, Sebmathews, Sebrat, Secretlondon, Selfdiscipline, Selivanow, September 11 terrorist, Sethoeph, Sfan00 IMG, Shadowfyren, Shadowjams, Shakumafu, Shandris, Shanes, Shashank Shekhar, Shelmac, Sherylanne, Shibboleth, Shinjiman, Shnout, Sigma 7, Sigmundpetersen, Sikon, Silvergoat, SimonEast, SimonP, SimonTrew, Simple Bob, Simxp, SirGrant, Siroxo, Sjakkalle, Sjorford, Skirks, Skiwi, Skyeap, Slashuer, Slathering, Sleepeeg3, Slipknotmetal, SmackEater, Smokizzy, SmoothNikola, Smurfy, Smyth, Snaxe920, Snoyes, SoSaysChappy, Solphusion, Somaditya, Somebody in the WWW, Soumyasch, SouthernNights, Southpark20, Spe88, Spear of fire, SpecMode, Specter01010, SpeedyGonsales, Spencerperry, Sperling, Spliced, Spookfish, Spug, Sri555666, Srnelson, Ssd, Starionwolf, Steel, Steezly, StephenH, Stephenchou0722, Steve2011, Stevenh123, Stevenmitchell, Stevietheman, Stilroc, Stino v, Stonda, StuIsCool, StuThomas, SudoGhost, Supercooljs2, Supers, Supremeknowledge, Suruena, SusanLesch, Susvolans, Swabjob, Sweetback, Swollib, Swotboy2000, SyntaxError55, Syrthiss, Szhang21, TAG.Odessa, TJ Spyke, Ta bu shi da yu, Tacvek, Tad Lincoln, Tadas12, TakuyaMurata, Tangotango, Tannin, Tarashav, Tawker, Tbhotch, Teamcritical, Tech2blog, Techman224, Template namespace initialisation script, Tepidpond, Terence, Tfgbd, Thaek, Thalakan, Thavian, The Colclough, The Disco Times, The Epopt, The Fish, The Geneticist, The Rambling Man, The wub, TheChrisD, TheDoober, TheGeneralUser, TheKMan, TheNewPhobia, ThePointblank, Thealexweb, Theazman1, Thegreenj, Themodernizer, Thetehror, Thewallowmaker, Thewikipedian, Thingg, Think outside the box, This, that and the other, Thisisntfake1, Thomas Larsen, Thompson.matthew, Thorpe, Throup, Thu, Thumperward, TiCPU, Tiggerjay, TimR, Timberframe, Timmeh, Timwi, Tinton5, ToasterOS, Tom Morris, Tomcage9, Tommy Irianto, Tompagenet, Tomy9510, Tonsofpcs, Tony1, Tothwolf, Towel401, Towsonu2003, Toyotaboy95, Tracer9999, Train2104, Treekids, Tregoweth, TrekMaster, Trevor MacInnis, Tripacer99, Trivialist, Trusilver, Tsunaminoai, Ttwaring, TubularWorld, TuneyLoon, Turionaltec, TuukkaH, Twernt, Twickline, TyA, Tyomitch, Tyraios, Tyrol5, UMC2, Uli, Ultimus, Umofomia, UncleBubba, Unimaginative Username, Unixman83, Unknownperson1234, Uriel8, Utcursch, Vahid83, Valenciano, Vanished user 5zariu3jisj0j4irj, Vbrtrmn, Vedek Dukat, Vegaswikian, Vincent.premysler, Vontafeijos, Vrenator, W1 m2, Wackymacs, Walter Görlitz, WalterGR, Wanderson9, Wangmike, Wapcaplet, Warlordwolf, Warren, Weararedjacket, Wengier, Werdan7, Werideatdusk33, Westley Turner, Whiteford8, Whkoh, Widefox, Wiki alf, Wiki fanatic, WikiBone, WikiFew, WikiMan44, Wikidexel2, Wikinger, WikipedianMarlith, Wikiwaka101, Windowsknowitall@msn.com, WindowzRULZlolZ, Windsok, Wknight94, Wlklpedla is meant to be vanda1ized, Wmahan, Wmplayer, Woohookitty, Wow1000, Wragge, Wsoder, Wykis, X1987x, XJamRastafire, Xation, Xelgen, Xeysz, Xgmx, Ximian99, Xlation, Xmachina, Xp54321, Xpclient, XrXeJoeXaXpXeXr, Xtreme racer, Xxglennxx, YUL89YYZ, Yama, Yamamoto Ichiro, Yamla, Yasirniazkhan, Yeeshenhao, Yehiahassan, Yergizmo, Ynhockey, Yousifnet, Yuckfoo, Yuhong, Z.E.R.O., Z98, ZFU738, Zakfleming, Zapvet, Zeeboid, Zepheus, Zer0431, Zntrip, Zondor, Zoney, Zundark, , 乌尔奇奥拉, 每日飞龙, 1542 anonymous edits

Windows Server 2008 *Source*: http://en.wikipedia.org/w/index.php?oldid=501545261 *Contributors*: 1exec1, 28bytes, A Raider Like Indiana, ALEF7, Aaron1a12, Abune, Achilles2.0, Adi4094, Aealtrus, Aero Flame, Ageha Winds, Ahunt, Ajh16, Akhristov, Alex Bakharev, AlistairMcMillan, Amadeus93, Andonic, Andy Dingley, Anteaus, Anthonyslive, Ar-wiki, AubreyEllenShomo, AussieLegend, Avocado27, Ayleuss, Baldursgate, Banes, Barvinok, Bebo2good1, Beest, Benwing, Bevo, Blowdart, Blue9637, BlueNovember, Bobby122, Bonadea, Bovineone, Brandon, Brianski, Bugtrio, Camanda, Cameron Scott, Catfish777, Cgwaters, ChazBeckett, Christopher G Lewis, Closedmouth, Compgeek123, CoolingGibbon, Cosmicdreams, Courcelles, CyberSkull, DARTH SIDIOUS 2, DBigXray, Damian Yerrick, Damienotis, Danger, Daniel Lawrence, Daniel.Cardenas, DeadEyeArrow, Demerzel7, DerekMorr, Dfxdeimos, Diamond145, Dorsal Axe, Doughboy58, Download, Drahgo, Drwebster93, EagleOne, Easter r0x, Eh.yeoh, ElinorD, Enochlau, Erfa, Errorlines, Evil Merlin, Evildeathmath, Eyas, Firthr, Fjehoel, Fjurden, FleetCommand, Fluteflute, Fowl2, Fran z, General Jazza, Georgia guy, Gfoley4, Ghettoblaster, Gioto, GoodwinC, Granburguesa, Gronky, Guinness, Gunnar Guðvarðarson, Guy Harris, HR4, Haemo, Hamburglar2011, Hede2000, Hidbaty223, Hoo man, Hza100, IMSoP, Igoldste, Illegal Operation, IllusiveMonkey, Indefatigable, InfotechGuyz, Ixfd64, Jabencarsey, Jamieostrich, Jasper Deng, Jeh, Jeremy Visser, Jgw, Jhdezjr, Jim.Liu, Jim1138, Jiy, Jjupiter100, Jolson88, Josh the Nerd, Jpeeling, Jrenier, Jusdafax, Karl2620, Karnesky, Katarighe, Kaustubhg123, Kkk, Koman90, Kozuch, KrisBogdanov, Kubanczyk, Letdorf, Leuko, Linda Editor, Lonaowna, LtNOWIS, MalibuTheCat, MaratIk, Markhoney, Masgatotkaca, Master of Puppets, Materialscientist, MehrdadAfshari, Melnakeeb, Men72194, Mephiles602, MikeRS, Misi91, MithrandirAgain, Mix Bouda-Lycaon, Mmoller, Mo ainm, Mononomic, Mortense, MrN9000, Msikma, MushroomCloud, Musp, My76Strat, Mário e Dário, NJA, Neolandes, Nicolas Love, Njan, Northgrove, Noypi380, Ntotten, OKtosiTe, OckhamRazor77, OnePt618, Panarchy, Paul kenneth mora, Paul1337, Pgilman, Phlegat, Pmsyyz, Poli, Prari, Qixil, R'n'B, Racconish, Raysonho, Rchandra, Reedy, Remember the dot, RenniePet, Requestion, Rich Farmbrough, Riki, Rjwilmsi, Robenel, RobertL, Ronark, Rovastar, Rsocol, Ryan Norton, Scott.somohano, Senator Palpatine, Sherkhon, Shshme, SkyF, Slwri, Smbrannon, Smjg, Socrates2008, Soumyasch, SpigotMap, Ssssstu, Stenis, Stephenchou0722, Steven Zhang, Stonkie, Stualden, Suggestednickname, SunCreator, Super-Magician, Swalot, Syp, Ta bu shi da yu, TerraFrost, Tfl, The FORCE One, TheGrimReaper NS, Themodernizer, Tide rolls, TigerGeek, Tim Boothby, Timbatron, Tomwchow, TonyW, Trusilver, Tslocum, Tygrrr, UU, UncleDouggie, Uusitunnus, VMS Mosaic, Vegaswikian, Vlad, Wangowiki, Warren, Werideatdusk33, WikiCle, WikipedianYknOK, Winterst, Wiretse, Woohookitty, XMog, Xadith12, Xavier114fch, Xcountry99, Xerces8, Xionbox, Xpclient, Zantolak, Ziggitysplit, Zr40, පසිඳු කාවින්ද, 633 anonymous edits

Windows 7 *Source*: http://en.wikipedia.org/w/index.php?oldid=504694952 *Contributors*: (Turnip), (chubbstar), -Majestic-, .:Alex:., 00iddy, 10metreh, 118.9O.9O.116, 13josh31, 1990alan, 2001:4BA0:FFF3:EF:0:0:0:2, 24dot, 2605:8900:5000:1001:6:0:12C:2, 28421u2232nfenfcenc, 2toy mora, 3dwarehouse, 465456645 boyaah 5464996, 5 albert square, 512upload, 66 is back, 66seals, 68DANNY2, 77.7.53.40o, A Raider Like Indiana, A. S. Castanza, ABF, AGiorgio08, AMK1211, Aateyya, Abani79, Abhijitsathe, Abhkum, Abours, Adacosta, Adamhauner, Adondai, Adoniscik, Aduckandaduck, Adw2000, Aeroavi, After Midnight, AgentB1, Ahkilinc, Ahmedunbreakabletato, Ahunt, AimalCool, Airplaneman, Airstrike, Ajnauron, Akhristov, Akldawgs, Aladdin Sane, Alansohn, AlbionSection, Ale jrb, Aleenf1, Alex43223, Alexanderino, Alexf, Alexius08, AlexiusHoratius, Alfie66, Aliasd, Alib, Alibaba, Alice in the wonder yupi world, Alig salty, Alinor, AlistairMcMillan, Allanbeaudry, Allstarecho, Alshain01, Althepal, Aluvus, Alwaystech, Ameliorate!, Amindzx, Amplitude101, Amuletxheart, AnOddName, Andrew c, Andrew35133, Andrew8326, AndrewJD, Andrewrp, Andy, Andy Johnston, Andy Marchbanks, Andylong, Andyso, Angr, Anivaros, Anna Lincoln, Ant roedr der, Anthm76, Anthony Appleyard, Anthonyhcole, Aperson1234567, Apparition11, Arbero, Archie Paulson, ArglebargleIV, Arjun G. Menon, Art LaPella, Artichoker, Aruffell, Arvindn, Ashadeofgrey, Aspensti, Asvs13, Attilios, AuburnPilot, Aurista25, AussieLegend, Austin512, Avenged Eightfold, Avicennasis, AxG, Axelriet, Axl, Ayleuss, Ayushvijay, Azekeal, BB RULZ07, BCube, BNSF1995, BOFH2005, BWCNY, Barek, Barneca, Barts1a, Battamer, Bbatsell, Beano, Beast of traal, Bebo2good1, Beetstra, Bellemonde, Ben b, Bender235, Bevo, Bewildebeast, Biasoli, Biblevistaheresy, Bieniu rick, Bigdaddyrob2g, Bigjimr, Bihco, Biker Biker, BillyTFried, Billykwooten, Bilz0r, Blake-, Blankness63, Blast73, Blehfu, BlindEagle, Blob5824, Blomskoe, Blound, Blowdart, Blue9637, Blurpeace, Bob the Wikipedian, Bobblewik, Bobo192, Bobrengo, Bodnotbod, Boky, Bomberboysk, Bonadea, Bongwarrior, Booticon, Bop It Guy20, BorgQueen, Bourgeoisdude, Bovineboy2008, Brancamj, Brandon, BrandonR, Breawycker, Brentdax, Brettrayfiel, Brian Kendig, Bridgeplayer, Broken Fruit, Bruceprn, Bryan.burgers, Bsutcliffe, Bswenso2, Bubba73, Bubzie5, Bumcheekcity, C.dotson94, CHONGOR, CHouseh, CJMiller, CL, CWY2190, CWenger, Cablehorn, Calu2000, CalumCook234, CalusReyma, Calvinturbo, Calwiki, Campoftheamericas, Caper13, CapitalR, Capricorn42, CaptainVindaloo, CardinalDan, Carewser, Carnildo, Carobonos1, Cars4me, Causa sui, Cdaughertyjr, Ceeon, Celarnor, Centurionguy, Chad01, Charles Gaudette, Chasingsol, Chatfecter, ChazBeckett, Chenglou, Chief SIRQ, Chilicheez, Chluk2425, Chocobogamer, Chowbok, Chris122990, Chrisch, ChrisfromHouston, Chrismgallen, Christian Swenson, Christopher Kraus, Chuunen Baka, Cireshoe, Cirt, CiudadanoGlobal, Ckatz, Clam0p, Claritas, Cleduc, Closedmouth, Coasterlover1994, CobraBK, Cobraman156, Codenamepenryn, Cody Cooper, Cody.feilding.nz, Coffee, Cold Phoenix, Coma28, Cometstyles, CommonsDelinker, Contents rest..., Control-alt-delete, CoolingGibbon, Corvus cornix, Coviti, Crash Underride, Crazlunatic, Crazyboy899, Csigabi, Cst17, Ctmt, Cullend11, Cumbiagermen, Curlz06, CyberSkull, Cybercat, Cyclonenim, Cyzor, Czechnmymail, D, D3, DARTH SIDIOUS 2, DBigXray, DJBoddington, DRGrim, DRJHMUFC, Dadycupcakes, Daec.wz.cz, Daikiki, Dajavax, Dale.90, DaltinWentsworth, Damaster98, Daminstant, DanShihisAsian, Dando120, Daniel, Daniel Hen, Daniel.Cardenas, DanielDeibler, DanielPharos, DannyDaWriter, Dansiman, Darbao, Darkfate, Darkquest21, Darremon, Darth NormaN, Darth Panda, Dasani, David3718, Dbeachy1, Dcandeto, Debresser, Deemon, Defusiondex, Deggle, Dehm46, Dekisugi, DenisYurkin, Dennylin93, DerHexer, Derangedfirewire, Deviantomega, Dffgd, DiaNoCHe, Diaa abdelmoneim, DiaphragmMonkey, Diego Grez, Discospinster, Dittyh39, DizzyITTech, Djlollyb, Dmillman, DmitryKo, Dngnta, DoYourBusiness, DoirB, Dominar Rygel XVI, DominatorMatrix, Donate to this project create knowledge, Donfbreed, Dorsal Axe, Doshindude, DoubleBlue, Dougofborg, Dr.Franich, Dr.K., Dr.Koljan, Drewhiggins, Dripdk, Droopy345, Dspradau, Duncan, Dunmanhigh, Dvferret, Dwyatt 101, Dynaflow, E.R.UT, EBAY RULEZ and stuff, ERK, ESkog, Earringguyz, Eastlaw, Ecbfilms001, Ed the Penguin, Eeekster, Eegorr, Eelamstylez77, Efficiency54WS, Ejg930, Elektrik Shoos, Elite9001, Elockid, Emarsee, Emilhem, EmpMac, Emry, Enfestid, Ensign Q, Ente75, Entegy, Enviroboy, Epbr123, Eraserhead1, Eric-Wester, EryZ, Estoy Aquí, Estudiarme, Evice, Evil Eccentric, EwokiWiki, Excirial, Extransit, F1racer328, F44control2, FChurca, FF2010, FT2, Faalagorn, Fabrictramp, Fahadinc, Fahadsadah, Falcon8765, Falcon9x5, Fantasy Game Productions, Farry, Feezo, Felanox, Fernvale, Ferrari895, Fetchcomms, Ffgamera, Fieldday-sunday, File 891 Sheet 17., Fin Archangel, Finalnight, Fishnet37222, Fld300b, FleetCommand, Flewis, Flippytoon123, Fluteflute, Fnlayson, FokkerTISM, FormerUser1, Fourpole, Fowl2, Fran z, FrankTobia, Frap, Freakgeek, Freyr, Fry1989, Fryed-peach, Fshy, Fumitol, Funny little guy, Fvw, Fæ, GB fan, GILLY710, GSK, GTR45766, Gabriellundmark, Gabrielpokemon, Gail, Gaius Cornelius, Gaius Octavius, Gallonigher, Gamersedge, Gammondog, Garas, Garde, Gary King, Gary63, Gateman1997, Gatemansgc, Gboyers, Ged UK, General Heed, General Rommel, Geobijes, Geoffrey, Georgia guy, Georgy90, Gfoley4, Gggh, Ghettoblaster, Giggett, Gioto, Godfollower4ever, Gogo Dodo, Goodone121, Gordon Ecker, Gpia7r, Gppande, Gpvos, Gr1bb, Gracz54, Grafen, Graham87, GrandDrake, Granpire Viking Man, Grayshi, Green caterpillar, Greenkangaroox, GregorB, Grim23, Gromlakh, Gu1dry, Gurch, Guy Harris, Guy M, Gwabell, Gwernol, Gyrferret, HJ Mitchell, HR 40000, HR4, Habanero-tan, Hadley falls, Haipa Doragon, HairyWombat, Halaster, Halsteadk, Ham Pastrami, HamburgerRadio, Hamzic, Hankyeol, Happy Dude, Happyandyk, Havelbaude, Havermayer, Hdtvinfoeu, Healy6991, Helge Skjeveland, Helios solaris, Helixweb, Hello5959us, Helvetica, Henrygb, Hergio, Heroman1234, Heron, Heymid, Hidbaty223, Hilary Gage, Hkpolitik, Hmich176, Hmwith, Hodgson, Homestereo, Honeymanl, Hoo man, Hooriaj, Hp notebook, Hugo 87, Hullaballoo Wolfowitz, Hunocsi, Hunter Kahn, Hydrargyrum, Hza100, I Love Pi, I42, I5bala, IBoy2G, IGAT, IMSoP, IMaMac, IW.HG, Ian2894, Icairns, Ice Truck Killer, Igob8a, Igoldste, Illegal Operation, Imaninjapirate, Imperator3733, ImperatorExercitus, Indefatigable, Indrek, Intelati, Inter, Iridescent, IsUsername, Island Monkey, Ismail, It Is Me Here, Ithizar, Ixfd64, J matula, J.delanoy, JBsupreme, JCDenton2052, JForget, JLaTondre, Jac16888, Jack Phoenix, Jackfork, Jackrodgers1, JacobJetto, Jacono, Jake Wartenberg, Jakezing, JamesWeb, Jamie-forsyth, JamieS93, Jamietw, Jamsta, Jan Hofmann, Jandalhandler, JaredJenkins2, JasonJD48, Jasper Deng, Jawadreventon, Jay2020, Jazzyrhythms, Jdcompguy, Jeepday, Jeh, Jellypuzzle, Jenrzzz, Jeremiahbarrar, Jeremy Visser, Jeroen888, Jerryobject, Jesant13, Jesse V., JetBlast, Jgunaratne, Jholman, Jhsounds, Jigesh, Jim10701, Jimmy stallman 35, Jimthing, Jj98, Jjupiter100, Jman8088, Jmcollier, Jnc, Jock Boy, Joecool1885, Joey-das-WBF, Joffeloff, John, John Goettle, John the gardener, John254, JohnCD, Johnfernandogentile, Johnleemk, Johnny Jupiter, JonMcLoone, Jonathanmjefferies, Jonhoo, Jonjames1986, Jopleka, Jorgenev, Jose Concepcion, Josephegorman, Josh the Nerd, Jrugordon, Jsalims80, Jtalledo, JuJube, JuanC08, Juhachi, Juliancolton, Julzz, Jungholee90, Junkcops, Jusdafax, Jvd897, Jwoodger, Jwsolutions, Језка, KGasso, Kakarot and goku are the same, KamasamaK, Kanags, Kaobear, Katana111, Katarighe, Kazzo111, Kbolino, Keilana, KennethMPennington, Keraunoscopia, Kesla, Keyang tay, Khr0n0s, KieferSkunk, Kiersere, King Mustard, King of Hearts, Kingofpunial, Kingoftherings, Kingpin13, Kipholbeck, Kirachinmoku, Kleb t, Kmhkmh, KnightRider, KnowledgeOfSelf, Kollision, Koman90, Koratnia, Kozuch, Kralizec!, KrisBogdanov, Ktr101, Kubek15, Kubigula, Kurze, Kvasi, Kyoungil, Kypr8, Kyteto, L Kensington, LOL, LP-mn, La lozono lolita, Lamename3000, Lampman, Larry V, Lawrencema, Lazera, Lazerlevi, Ldmerriam, Lego castel mthrfckr, Leigh-Kennedy, Leptonsoup337, Letdorf, Leviel, Lewys93, Lifung, LightSpeed, LightSpeed3, Lights, Lightsup55, Ligulem, Linefeed, Linuxbeak, Linuxlove8088, Lion golden mayer, LisKiller, Lisss001, LittleBenW, Livewings, LizardJr8, Llamabr, Lnatan25, Localzuk, Longzheng, Lookatmygrlyfont, LoopZilla, Lopifalko, Lordalpha1, Lorian, Lotje, LovesMacs, Luigiacruz, Lumidek, Luna Santin, Lysander89, Lyverbe, M.O.X, M.nelson, M3talsyr3n, M4dRefluX, MBisanz, MER-C, MET920, MLVD, MZMcBride, Macarron, Macedoniarulez, MadGuy7023, Madlobster, Madman10101, MadmanNova, Magister Mathematicae, MainFrame, Malik Shabazz, Manamanam, Mandarax, Marcinw2, Mardus, Marek69, Marilyn.hanson, Mark.P.Bartlett, MarkKB, Marksurge, Martarius, Martin steven david, Martyhol, Mason092, MassKnowledgeLearner, Master martin, MasterAshZero, MasterOfTheXP, Mathias-S, MattTait, Matturn, Max Naylor, Maxis ftw, Mayhew132, Mazevedo, Mb1000, Mcmatter, Mdebets, Mdikici, Medessec, Meewam, MementoVivere, Mennonot, Menziep, Mephiles602, Merlion444, Metarhyme, Mf393, Mghostsoft, Micahblitz, Michael Zimmermann, Microsofkid, Midkay, Midway, Mikaida, Mike Payne, Mike411, MikeAllen, MikeZuniga, Mikebar, Mikecintel, Mikepelley, Milad Mosapoor, Mild Bill Hiccup, Minghong, Minna Sora no Shita, MinorContributor, Mipadi, Miss Madeline, MithrandirAgain, Mkat09, Mktsay123, Mmmaaahhh, Mmxx, MocniSkopac, MoltenBoron, Monaarora84, Monger, Moofant, Moogleboi, Mordien, Morqueozwald, Mortense, Moyce, Mr Lman22, Mr Mo, MrTechNews123, Mrcomputerman7, Mscuthbert, Mtp318, Mudkiller, Mugunth Kumar, Muminovs, Mutchy126, MuzikJunky, Mvjs, Mwarren us, Mygerardromance, MynesFOFF, Myscrnnm, Mythdon, N419BH, NJA, NaBUru38, Naruto kqod, NawlinWiki, Ncmvocalist, Ndrasen, Ndrfx, Neelchauhan, Neilksos, Nemanjanoob, Neon, Newsoxy, Ngz303, Nhellviktor, Nichalp, Nicholas761, Nick2253, Nicolas Love, Nicoli nicolivich, Nightstallion, NikSRB, Nikpapag, Nil Einne, Nio32, NitrogenTSRH, Njuuton, Nn123645, Noddog, Noescom, Norch123, Northgrove, NorwegianBlue, NotIdleChatter, Nottheking, Nrbelex, Nsaa, Nuck Chorris 0, NuclearWarfare, Numen, Nuujinn, Nyat, Nyr56, Odie5533, Officer781, Ohconfucius, Ohnoitsjamie, Ohyeahuhhuh, Okiefromokc, Okungnyo, OlavN, Oldhamlet, Oldiesmann, Omar060997, Omphaloscope, On this lost highway, One-eyed pirate, Oneiros, Onesimos, Ooddiittyy, Organman91, OriginalGamer, Oscarthecat, Oxymoron83, Patriarch, Pauerbach08, Paullloydjohnson, Pcuser42, Pencilcase123, Peripitus, Perspectoff, Peteforsyth, Peter Campbell, Peyman4u, Pgerrity, Phantomsteve, Pharaoh of the Wizards, PhatKangaroo, Philip

Trueman, Phoogenb, Phynicen, Piano non troppo, Pigsonthewing, Pikachu 976 976, Pinethicket, Pinkadelica, PirateSmackK, Pixelsnader, Piz d'Es-Cha, Pizzahut2, Pk1223, Plau, Playclever, Pmlineditor, Pmsyyz, PoisonedQuill, Pojet beau, PokeMan999, Pokiron, Pontiacsunfire08, Possum, PowerComputer, Prcjac, PrinceGaz, Proofreader77, PsY.cHo, PseudoOne, Pt, Public Menace, Purplesplat, Pwt898, Pyfan, Qasdfdsaq, QueenCake, QuiescentWonder, Quilokos, Quincy Calunga, R'n'B, R2y2a2n, RG2, RJaguar3, RLE64, Rahulbhagchandani, Rakeshonsou, Random contributor, RandomAct, Rasmasyean, RaviC, Rayqwmon46, Raysonho, Rebroad, Reconsider the static, Red thirty, Redekopmark, Redvers, RegularBreaker, Rehman, Reixo, Remurmur, Renesis, Resplendent, Ribeirowaldo, Rich Farmbrough, Rich Janis, Richardguk, Richtom80, Rigby27, Rjd0060, Rjwilmsi, Rlbarton, Rmosler2100, Rmsuperstar99, Roadbloc, Robert Bond, Robert K S, RobyWayne, Rockfang, Rolandhelper, Ronark, Ronhjones, RoyBoy, Rrburke, Rrose Selavy, Rsocol, Rtyq2, Ryan Norton, Ryan-alexander, Ryanquek95, S0aasdf2sf, SF007, SNIyer12, Sabrown100, Sadara, SadaraX, Saiarcot895, Saladbar1, Salamurai, SalopianJames, Sam Korn, Samlaptop, Samlaptop85213, San Diego Ca, Sandgem Addict, Sandstig, Sarahannedavid, Sarv007, Saumyzardasch, Sayden, SchfiftyThree, Schlegel, SchmuckyTheCat, SchuminWeb, Scientus, Scj2315, Scooter20, Sdorand, Sdornan, SeLfkiLL, Seabhcan, Seegoon, Servant Saber, Sexyparty, Sfm 7, Shadiac, Shadowjams, Shanew2, Shjacks45, Shshme, Shubinator, Sigma 7, Silvergoat, Simeon, Similar a lo que, Simon.hess, Simonalexander2005, Simxp, Sinistrum, SirMetal, Sirgregmac, Sjö, Sk8rSoda, Skier Dude, Skies, Skizzik, Skorp, SkyBullet, SkyWalker, Skycaptain95, Sliqer, Sloppy85, Slvr hwk tony tgr, SmartRice, Smorgan06, SnDPhoenix, Snider1, Snnw, Snowballsn, SoWhy, Sockatume, Socrates2008, Solidsnake1211, Solmonburk li, SomeCallMe 1, SonicAD, Sonicdude558, Soopersonic25, Soosern55, Sotcr, Soumyasch, Sourav bhattacharyya, South Bay, Soviet689, SpaceFlight89, Specs112, Spencerk96, Spotty11222, Sri555777, Srw985, Ssnautilus, Ssssstu, Stagyar Zil Doggo, Starfleet bop, Starionwolf, Starkiller88, Steel, Stefan Majewsky, Stentie, Stephen85442, StephenHodges, Stephenb, Stephenchou0722, Stephenw32768, Stepshep, SteveWiilliams, Steveprutz, Stevertigo, Stevie works, Stifle, Stj6, Stroppolo, StuffOfInterest, Stukov, Sturm55, Suggestednickname, Sumj, Super Dragonite, Super wiki editor, Super-Magician, Supercooljs2, Superm401, Supermanasfsd, Supertrouperdc, Surv1v4l1st, Svarya, Syahirsyazwan, Symode09, SynergyBlades, Sysrpl, T.Saito, THEN WHO WAS PHONE?, TKhaldi, TPK, TSwiftfan88, Ta bu shi da yu, Tabletop, Tagus, Tainted Conformity, Tamiera, Tampa Gator, Tamrood, TangLab, Tanhueiming, Tapir Terrific, Taylor Karras, Tbhotch, Tcncv, Tedder, Tedickey, Tedzdog, Teiladnam, Texas™, Th1rt3en, ThaSawh, Thalter, The 888th Avatar, The Giant Puffin, The Gow, The High Fin Sperm Whale, The Merciful, The Missing Piece, The Thing That Should Not Be, The Ultimate Koopa, The zepellin from queen, TheGeekHead, TheGeneralUser, TheNewPhobia, TheSameGuy, TheSpeedster, TheTechFan, TheWeirdMouse, Thealexweb, Thecurran91, Thefalcon2k, Themcman1, Themoby1, Themodernizer, Thewildwolf, Thingg, Thisiswe, Thljcl, Thlump, Thor netherlands, Thumperward, Thunderpenguin, Tide rolls, Tigerghost, Tikiwont, Tillmo, TimUsher01, Tima518, Timmeh, Timotheus Canens, Timwi, Tintenfischlein, Tintero, TinyClanger, Tiptoety, Tjdwowh, Tkgd2007, Tkynerd, Tmwerty, TnT3, Tobias Schmidbauer, Tobybuk, Toehead2001, Tohd8BohaithuGh1, Tom743, Tommy2010, Tony1, TonyW, Tow, Tpbradbury, Tracer9999, Tresiden, Tri400, Tricore5, Triplight, Tripter, Troy 07, Trusilver, Tstrobaugh, Ttat1108, Tuanese, Turbolinux999, Tyler-willard, Tyler9xp, UAhmad7, UNV, Ubergeekguy, Ubuntu vista3, Ukexpat, Ultimus, Unclebert11, Undertradeover, Unionhawk, United153, Untoldpromise, Urvabara, Useight, Utcursch, Utkarsh apple, Uturnaroun, Valar, Vanished user 5zariu3jisj0j4irj, Vatech09, Versus22, Vikizh, Vinh1000, ViperSnake151, Vistalover, Vivek825, Vlad, WJBscribe, Wacky WeLsH LaD, Wadertader, Wakari07, Waqas1987, Warren, Washburnmav, Wattlebird, Wavelength, Wayne Slam, Webmaster.legend, Webmaster961, Weebee777, Wenli, Wereon, White.matthew.09, WhosAsking, Whytehorse1413, Widefox, WikiLaurent, Wikidexel2, Wikimedium7, Wikipedian06, Wikipelli, WildCowboy, Will2k, Williameis, Wiltabone, Windows s3cks, Windows72106, Windowsvistafan, Wingedbunny1, Winston, Winwolf3d, Wjoakes4, Wmplayer, Woohookitty, Worldprofile, Wpuser0, WrcF1, Wreiad, Wteoh1, X-Fi6, X201, XRDoDRX, XX55XX, XXenternamehereXx, Xaban, Xdenizen, Xena princess plk, Xenfreak, Xetxo, Ximian99, Xinpeijin, Xionbox, Xompanthy, Xpclient, Xsparkoflife, Xtreme racer, Xvincentx77, Yakiv Gluck, Yamla, Yamomwa3, Yeeshenhao, Yerpo, Yofootball127, Yottamol, Your1727, Ysangkok, Yuckhil, Yulu, Zabby1982, Zayani, Zgillis, Ziiike, ZirixZero, Zntrip, Zollerriia, Zoltanator, ZooFari, Zrs 12, Zundark, Zvn, Étienne Trudeau, 675בר, පසිඳු කාවින්ද, 2453 anonymous edits

Command-line interface *Source*: http://en.wikipedia.org/w/index.php?oldid=504692196 *Contributors*: 4C, 4lex, AAA!, ACupOfCoffee, AVRS, Abs Like Jesus, Adam78, AeonicOmega, After Midnight, Akhristov, Al Lemos, Aladdin Sane, Alan Millar, Alanfeld, Aldie, Algorerhythms, AllenDowney, Ancheta Wis, Andre Engels, AndyBQ, Arch o median, ArchonMagnus, Arnero, ArnoldReinhold, Artagnon, Arvindn, Asmpgmr, B33R, Bart Fox, BartlebyScrivener, Bdesham, Bevo, BirdValiant, Bongwarrior, BurntSky, Burzmali, CQ, Cadsuane Melaidhrin, Cander0000, CapitalSasha, ChangChienFu, Chauncey27, Chealer, Chridd, Christopherlin, Cometstyles, Conversion script, Crazycomputers, Cybjit, Czarkoff, DARTH SIDIOUS 2, DMacks, DRKD, DXBari, Dancter, Daniel5127, DanielRigal, Dgutson, Dicklyon, Diego Moya, Dmbrown00, Dragice, Duffjohnson, Dugwiki, Dusik, Dzubint, ENeville, Edward, Elfguy, Elideb, Eman2129, Ericd, Erik Sandberg, Eruantalon, Evil Monkey, Exlex, FSII, Fabartus, FabioBatista, Fastily, FayssalF, Fibonacci, Fieldday-sunday, FleetCommand, Fourchannel, Frietjes, Gate28, Gator1, Getonyourfeet, Gforce20, Ghettoblaster, Gpearson2, Graham87, Grendelkhan, Guy Harris, Gwernol, HendrixEesti, Hirzel, Hvn0413, Hydrargyrum, IMSoP, Ideoplex, Incnis Mrsi, Intgr, Isnow, IvanLanin, Ixfd64, JIP, JLaTondre, JaGa, Jakuzem, Javawizard, Jerome Charles Potts, Jesin, Jgrahn, Jncraton, Justinstein, Kate, Kbdank71, Kbolino, Kejoxen, Khalid hassani, Killiondude, Kim Bruning, Koman90, Kpacquer, Kurtis1991, Kyle Barbour, Lexein, Lightmouse, Linnell, Liontooth, Lloydpick, LordRM, Lradrama, Mac, Majorly, Makkadiamond, Massysett, Matthiaspaul, Matusz, Meandtheshell, Merphant, Michael Hardy, Michaeldsuarez, Milan Kerslager, Millstream3, Minghong, Minimice, Minkee.s, Modster, MrOllie, Mrh30, Myanw, Myth723, Nasukaren, Naturalui, Nayuki, Neelix, Nick Warren, Nofxjunkee, O, Obradovic Goran, Obuli, Omegatron, Ortolan88, Panarchy, Pctechbytes, Peter S., PhS, Philip Trueman, Piotrus, Pnm, Possum, Ppntori, Prolog, Pt, Pwt-wsu-sp, QuixoticKate, R'n'B, Radagast83, Radiojon, Radomir, Random832, Rctdude2, RealWorldExperience, Recury, RedWolf, Rich Farmbrough, Rob Hooft, Ronz, Rsciaccio, Rwwww, SF007, SJK, SchreiberBike, SchuminWeb, Simon1220, SimonP, Sjhan81, Slady, Smalljim, Snoyes, Soumyasch, Splintercellguy, Stephen Gilbert, Svetovid, Svick, Tacoman359, Tannin, Tarquin, TedPavlic, Tedickey, Teza1532, Tghe-retford, The Thing That Should Not Be, Thuresson, TimShell, Timc, Timneu22, Tizio, Tobias Bergemann, Tobias Hoevekamp, Tom Morris, Tomaxer, Tommy, Trevj, Troy.hester, Trusilver, Ts1and2fanatic, TuukkaH, Uchaschiysya, Ucucha, Unixguy, Unknown-xyz, Uriyan, Urparticle, Vadmium, VampWillow, Virtlink, Wa3frp, Wandering Ghost, Warinthepocket, Warren, Wendy.krieger, Wikiold1, Wonderstruck, Wtshymanski, Ww, XJamRastafire, Xandi, Yath, Ysangkok, Zeno Gantner, Zeth, Zundark, Zwilson, ZxxZxxZ, 계정명뢸로하지, 436 anonymous edits

Bash (Unix shell) *Source*: http://en.wikipedia.org/w/index.php?oldid=502107509 *Contributors*: 194.237.150.xxx, A. B., A:-)Brunuś, ALavelle, Abhinavpriyadarshi8, Abhishek77in, AbstractClass, Ahunt, Aivalli, Akanemoto, Alismayilov, AlistairMcMillan, AllyUnion, Ameliorate!, Android Mouse, Anthony Appleyard, Apokrif, Arrenlex, Arvindn, Aythun, Ayucat, Azimout, B4hand, Bgough, BiT, BradBeattie, BrianFennell, Brianreading, Brighterorange, Bs27975, CFeyecare, CYD, CalumH93, Calvin 1998, Canaima, Cfajohnson, Charles Gaudette, Chatter, Chealer, Chester Markel, Chris 73, Chris Pickett, ChristTrekker, Conversion script, Cpiral, Cprompt, CyberSkull, DStoykov, Dalphi, Danakil, Daniel.Cardenas, David Siegwart, Deineka, Demonkoryu, DenisKrivosheev, DixonD, Djmckee1, Dmbrown00, Dmulter, Doctor J, Donkdonk, Dricherby, Drj, Ds13, Dysprosia, Echimu, EdJohnston, Edward, Eliashc, Elsendero, Emperorbma, Epastore, Eyu100, Falcon9x5, Falcor84, Feezo, Fennec, Finlay McWalter, Franl, Frap, Fred Bradstadt, Free Software Knight, Gdr, Gennaro Prota, Geo Swan, Gesslein, Ghettoblaster, Gigglesworth, Gilgamesh, Glenn, Gnufs, Graham87, Grandscribe, GregorB, Grim23, Gronky, Grover cleveland, Guaka, Guanaco, Gwern, H2g2bob, Hanru, Harvester, He's a very naughty boy, Henry Flower, Hhielscher, Hjsb, Hu, Husky, Hvn0413, ICEAGE, Ionutzmovie, Isnow, J.delanoy, JackPotte, Jacobolus, James.m.henderson, Jeff Carr, Jerome Charles Potts, Jerryobject, Jic, JimD, Jnc, JoeSmack, JonHarder, Jonik, Jouster, Joy, Julesd, KJS77, Kbdank71, Kbolino, Kenguest, Kernel.package, KimDabelsteinPetersen, Kl4m, Kl4m-AWB, Krellis, Krischik, Kusunose, Leolaursen, Lerdsuwa, Linuxguy2003, Lupin, Luxana, M7, MER-C, Mac, Mac c, Mad101dan, Mairi, Man with one red shoe, Marasmusine, Marchash, Mark Renier, Markhobley, Martarius, Marudubshinki, Masaru.wota, Matt B., Mcaruso, Meandtheshell, Mefistofele, Mentifisto, Merphant, Miaow Miaow, MightyWarrior, Mipadi, Morte, Mortense, Mpontes, Msnicki, MushroomCloud, Mysid, Nasukaren, Netoholic, Nicolais, Nikai, Niomiedc, Nysin, Oneiros, OsamaK, Ospalh, Packetspike, Paul Richter, Pbroks13, PhilippWeissenbacher, Pilaftank, Pixel ;-), Pmsyyz, Pnorcks, Popup, Prolog, Qwertylurker, Qwertyus, RJHall, Radiojon, Raffaele Megabyte, RainbowOfLight, Rangek, Reach Out to the Truth, Reisio, Renku, Rich Farmbrough, Rogerx, Roman V. Odaisky, Ruakh, SF007, Sango123, Sean Kelly, Sharkb, Shaydon, ShellCoder, Shlomif, SimonP, Siodhe, Skiidoo, Skrio, Smyth, Sontofle, Speck-Made, Spsneo, Starfarer, Stassats, Stevenj, Stevenwagner, Strait, Stubblyhead, SuperJr, Sygrous, Synchrite, Taras, Tassedethe, Tcsetattr, Tedickey, Tgeller, Thumperward, TinaSDCE, Tjmoel, Tomaxer, TotoBaggins, Toussaint, Traroth, Turbo95, Tzartzam, UnDeRTaKeR, Unixguy, Unyoyega, Uriyan, Utcursch, V.R.S., VQuakr, Vadmium, Vlad, Wikante, Winston Chuen-Shih Yang, Winterst, Wm, Xaje, Youngcada, Yurik, Zachlipton, Zearin, Zundark, 265 anonymous edits

Batch file *Source*: http://en.wikipedia.org/w/index.php?oldid=502742674 *Contributors*: 0sudoman, 1230049-0012394-C, 21.abhijit, ABF, Aacini, Absolutelyamazingalexander, Adam1213, Ae-a, Aitias, Akadruid, Akhristov, AlanUS, Alansohn, AlexKarpman, AlistairMcMillan, Andres, Andy Dingley, Anish9807, Anne226, Anomie, Another-anomaly, Arj, Armslurp, BIL, BananaFiend, Benhoyt, Benzen, Berserkerus, Bevo, Bobbo, Bryan Derksen, Caceo, CanOfWorms, Canthusus, Cassidy981, Chaheel Riens, Charles dye, ChemicalBit, Chibibrain, Chief SIRQ, Chris the speller, Chrisahn, Christian75, Christopherwoods, Claunia, Creol, Cribrad, Cuzimwhiterite, Cwolfsheep, DJMatus23, DMahalko, Damian Yerrick, Damicatz, David Haslam, Dawnseeker2000, Dcoetzee, Dougall, Dr. F.C. Turner, Dugas, Eclipse97, Elsendero, Emperorbma, Ethd, Evice, FF2010, Far55589, Fejj the ritual, Feureau, ForemanDomai, Frap, Fred Gandt, Frymaster, Fsdrsf, Ghettoblaster, Giraffedata, Glane23, Glorfindel83, GrahamColm, Grawity, Gwernol, Gzorg, Gzornenplatz, Hansfn, Harmil, Helpsloose, Holdenmcrutch, Husond, II MusLiM HyBRiD II, IRP, Iball, Insanity Incarnate, Iohannes Animosus, Ismellaprep, Jack Phoenix, Jasper Deng, Jcamilli, Jd027, Jensgb, Jesse Viviano, JesseHogan, Jh51681, Joedf, Joey-das-WBF, Jogar2, Josh Parris, Josh the Nerd, Justinstein, KD5TVI, Kbolino, Kithira, Kizzmequik 74, Krashlandon, Kruckenberg.1, Ksnow, Kungfuadam, Ld100, LiDaobing, Liftarn, Lihanrui, Lo2u, MSTCrow, Markhurd, Martarius, Mastergyan singh, MattieTK, Mattnelsen, Maxis ftw, Mazca, Mike Rosoft, Milan Kerslager, Mister Nighttime, Mr.Z-man, MrZombax, N-Man, Needlenose, NightKev, Ninly, Nono64, Notkeel, P1415926535, Pdn, PeteX, Peter Karlsen, Pforret, Pixnion, Pol098, Protious, Psychonaut, RCX, Rathgemz, Rebroad, Redquark, Retama, Rhubbarb, Ringbang, S3819, Screenshotguy, Seth ze, Sherip99, Shmuel, Shuipzv3, Silas S. Brown, Silentmacro, Sir Nicholas de Mimsy-Porpington, Skew-t, Snori, Sofian Rahmani, Splibubay, Stefan2, Stephen Poppitt, Stickee, Sunny bacon, SwiftEagle, Tbhotch, Tedickey, The guardians27, Thedarxide, Tide rolls, Timharwoodx, Tobias Bergemann, TomBrown16, Tomato86, Troy.hester, Tuntable, Tutterz, U52983, Uncle G, Urparticle, Veltas, Verbal, Versageek, Vrmlguy, W.F.Galway, Wadethefred, Wendy.krieger, Wermlandsdata, Wernher, Whiskey Pete, Wikiji, Winterst, Wtshymanski, Xpclient, Yath, Yobmod, Yonkie, Zzuuzz, 430 anonymous edits

Windows PowerShell *Source*: http://en.wikipedia.org/w/index.php?oldid=502858384 *Contributors*: -Majestic-, Aaron Brenneman, Abune, Adetaylor, AdrianLozano, Akhristov, Alex Sims, AlistairMcMillan, Anetode, Anna Frodesiak, Anomicene, AnonMoos, Antifumo, Apetrovskiy, Ardonik, Arvindn, Ascendingquickly, Asharism, Asmoore82, Astronautics, Autopilot, Avicennasis, Bevo, Black Falcon, Bleakgadfly, Blue Slime, Boingbop, Bovineone, Brighterorange, Btx40, CRGreathouse, CWenger, Cambridge, ma, Candace Gillhoolley, Capricorn42, Cardinal Direction, Cdc, CecilWard, Chowbok, Christoph hausner, Chromjuwel, Clear issue, Contentasaurus, Corbin Benton Davenport, Coutcin, Cwolfsheep, Cybercobra, Cyberherbalist, DBMwS, DHR, DStoykov, DadaNeem, Dalesc, Danim, Davemck, Dawnseeker2000, DerEikopf, Derek Ross, Derek farn, Dewritech, Dimitri.Koens, Disavian, Dogbertwp, DragonHawk, Dreftymac, Drphilharmonic, Dspn, Earthsound, Ecephd, ElectronicsEnthusiast, Evice, ExZsoldier, Faganp, Fbthibault, FleetCommand, Fran McCrory, Frap, Frietjes, Furrykef, Fæ, Galoubet, GavinOB, Gene Thomas, Gentleben8282, Georgeryp, Getcrunk, Ghettoblaster, Gil-Estel, Glider51, GregorB, Grouse, Haeleth, Halr9000, Harryboyles, Helvitica Bold, Herbee, Hergio, Heywoodg, HodracirK, Honkster, Hu12, Huntscorpio, IMSoP, ITurtle, Ian Pitchford, Ihavenolife, Intersofia, Issuesixty soulsgreat, JLaTondre, JVottero, James keeler, Jeffq, Jeffthejiff, Jhicks61, Jiddisch, Jkliff, Joey-das-WBF, Josh the Nerd, JoshuaGrainger, Jsaylor3, Jsnover, Julesd, Juneblender, KJRehberg, Kaini, Kinger, Koavf, Kurt, Kushboy, Leotohill, Little Professor, Llywrch, Lofote, Logan, Logical Cowboy, M gol, M4gnum0n, M7, Mahanga, Marcika, Mario23, Mark Renier, MarkPeters45, Markbuxbaum, Marudubshinki, Matěj Grabovský, Mechcozmo, Mephiles602, Mernen, Miami33139, Mishac, Mistermoe, Mjb, Mnruxter, Modster, MoraSique, Mortense, Mosquitoeater, Mrbitter, Mrivera1, Msnicki, Mwtoews, Nikthestoned, Nimbex, Northgrove, Ocalvo, Offby1, Ohnoitsjamie, Oising, Open4D, OsamaSajid, PTSE, Paul, Paul A, Peppage, Pgp688, Piet Delport, Pol098, PowerChelle, PowerWF, R keith hill, R'n'B, R3m0t, Red Thrush, Reedy, Resing, RickSheikh, Ringbang, Rjwilmsi, Rl, Rob Russell, Robertbowerman, RsRado, Rsciaccio, SURIV, Scott Paeth, Shinjiman, Silvergoat, Snori, Socrates2008, Soifranc, Soumyasch, SpaceFalcon2001, Stormie, Susfele, Svick, Tabledhote, Tagus, Tariqabjotu, Tedickey, Tennisman1432, Tfl, Themfromspace, Thoward37, Thumperward, Tim1357, Timothyleerussell, TinJack, Tonydesweet, Torc2, Toussaint, Tulcod, Tyomitch, Unfree, Uzume, VX, Varlaam, Vasiliy Faronov, Vbaculum, Warren, Wavelength, Wdanwatts, Web-Crawling Stickler, Willy888, XP1, Xpclient, Yath, ZZyXx, ZacBowling, Zbpop, Zelphar, ZeroOne, Zerofire0, Zeus, Zickzack, Zsinj, 503 anonymous edits

Environment variable *Source*: http://en.wikipedia.org/w/index.php?oldid=504179142 *Contributors*: 16@r, Aacini, Abdull, Abu ari, Aledeniz, AlistairMcMillan, Andreas Kaufmann, AndrewHowse, Ante Aikio, Apokrif, Arda Xi, Avulas, Batpox, Ben Ben, Born2cycle, Buttertea, CanisRufus, Cburley, Celique, Charles dye, Chrstphr, Ciotog, ComSpex, Corner-face-jacks, DMahalko, DVD R W, Diego Queiroz, Dillard421, Djmitche, Dmitri Maximov, Dolda2000, EagleOne, Episteme-jp, Everyking, Evice, Frap, Gaius Cornelius, Gary King, Ghettoblaster, Glenn, Hyad, Ignacioerrico, Ilyaroz, IraChesterfield, Ixanezis, Jaho, Jasonpearce, Jeff3000, Jerome Charles Potts, Joey-das-WBF, John Bokma, JohnMashey, Josh the Nerd, Kaiserdaemon, Khym Chanur, Kri, LPfi, Lambent Ametrine, Limited Atonement, LittleBenW, Llywrch, Madmardigan53, MarkMcClelland, Martinkunev, Marudubshinki, Matthiaspaul, Michaeldadmum, Miko3k, Mion, Mozillaman, MrJones, Mrzaius, Mulad, Nikai, Omveer15681, PGibbons, Paraphrased, Phantom xxiii, Pjn, Pmsyyz, Pol098, Polluks, Progably, R'n'B, Redline, Reisio, Requestion, Riedl, Rjwilmsi, Romanc19s, Romeu, Roybristow, Seashorewiki, SebastianHelm, SlimVirgin, Snori, Ssd, Stepa, Sunman600, Tannin, That Guy, From That Show!, The locster, Theta4, Thnidu, Timharwoodx, Tizio, Victor.spain, Vlad.setchin, Wermlandsdata, Wikid77, WinFORs, Winterst, Wittyameta, Wknight94, Wlievens, Wtpollard, Xerces8, Xzaranos, YordanGeorgiev, 188 anonymous edits

Configuration file *Source*: http://en.wikipedia.org/w/index.php?oldid=501178689 *Contributors*: 16@r, AKA MBG, Abdull, Alansohn, AlistairMcMillan, Bagwill, Benjaminoakes, CanisRufus, Dreftymac, EdgeOfEpsilon, Eeekster, Error, Figure, Fragglet, Fratrep, Friedo, Furrykef, Gadfium, Gardar Rurak, Gwern, Icey, Isilanes, Jamelan, JavaTenor, Jgrahn, Jjk, JonHarder, Joy, KP-Adhikari, Kusunose, Kutabar, Lampak, Leolaursen, Lifeformnoho, Longhair, Massysett, Mecanismo, Nandhp, Nbarth, Nixeagle, Omegatron, Perey, Phatom87, Poppafuze, Psychonaut, RedWolf, Rjwilmsi, Rwwww, Scott5114, Shanekboyce, SteveLoughran, The undertow, The weaver, Themfromspace, Timharwoodx, Walter Görlitz, Wereon, Winterst, روخو, 59 anonymous edits

Windows Registry *Source*: http://en.wikipedia.org/w/index.php?oldid=504575623 *Contributors*: 1812ahill, 94knux, A:-)Brunuś, AThing, Abb615, Abce2, Ahmad ghare, Ajraddatz, Akhristov, Albertsoft, Aldaden, Alenaross07, Alerante, Alexius08, AlistairMcMillan, Amin20, Anbu121, Andyjsmith, Angel leo ef, AngelOfSadness, AnonMoos, Antandrus, Antonio Jose Torres, Anupam, Apachegila, Arosa, Artie p, AussieLegend, BasBloemsaat, Billyoneal, Bleakgadfly, Blitterbug, Blowdart, Bobbydown, Bobo192, Bodragon, Boeing720, Bonadea, Borgx, Brack1969, Btx40, Bungalowbill, CJMiller, CUTKD, Cab88, Can't sleep, clown will eat me, Canaima, Capricorn42, Cedear, Cenarium, Cheng, Chris the speller, Chrislk02, Christian Storm, Christopher G Lewis, Chromancer, CliffC, Closedmouth, Cmdrjameson, Cncccer, Cognitivelydissonant, Collard, Comrade009, Corpx, Cosmo0, Courcelles, Cpl Syx, CrazyTerabyte, CyberSkull, DStoykov, Dah31, DanielPharos, Daveb, Davey77, David Gerard, David spector, Davigoli, Debog, Dee Earley, Dekingsey, Delan Azabani, DerEikopf, Derekrogerson, Djegan, Dkostic, Dlrohrer2003, DoggyDude96a, Doniago, Doulos Christos, DrYak, Dreftymac, Długosz, Eagleal, Echodelta, Edward Z. Yang, Eghanvat, El Hoy, Entbark, EoGuy, Er Komandante, EvanED, Everyking, Evice, Fairce, Faust6, Ferret, Figgie123, Fishnet37222, FleetCommand, Foant, ForthOK, Fram, Fran z, Frap, FreeRangeFrog, Fripon, Furrykef, Gerbrant, Ghettoblaster, Glennp, Goodone121, Goplat, GraemeL, GregorB, Guycalledryan, Gzkn, Haakon, Happyandyk, HartzR, Hello5959us, Heron, HobgoblinOfLittleMinds, Hooperbloob, Howard McCay, Huon, Hyad, I already forgot, Iainelder, IceHunter, Igoldste, Ike9898, InverseHypercube, Iridescent, Isheden, J. Finkelstein, JakobVoss, Jarble, Jasens, Jasper Deng, Jeltz, Jeremy Visser, Jesant13, Jesus Carp, Jetkins, Jglass7268, Jlshoem, Joaopaulo1511, John Hay, JonHarder, Joseph MB, Jossi, Jrutski, JulesH, Jwoodger, K Dominik, Kant.shashi.s, Kar.ma, Karl Meier, Kb, Kbdank71, Killerj, Kotiwalo, Krótki, Kuzaar, L Kensington, Lacharge, Ldo, Light current, Lilpiggia12, LindsayH, LittleOldMe, Localzuk, Lofote, Loren.wilton, Lyml, M7, MER-C, Magioladitis, Mani1, MaraNeo127, Markhobley, Martyn Lovell, Masiano, Materialscientist, Mbssbs, MegaBrutal, Mel Etitis, Mephiles602, Metabolix, Michael Daly, Michaelas10, Mild Bill Hiccup, Mipadi, Mjb, Mkoyle, Mlist, Monkeyman, Morriske, Mortense, Motyka, MrArt, MrESaulved!, MrFirewall, Mteorman, Mwtoews, Myanw, Mystic8277, NeF, Nicd, Nightowl12, Noclador, Nopetro, Notopia, Ohnoitsjamie, Omarabid, Omegatron, Omicronpersei8, Opalatnik, Orange Suede Sofa, OwenX, Oxymoron83, PAMPPN, Pascal.kotte, Pb30, Pelagic, Pesterjack, PeterReid, Peterl, Phantomsteve, Phatom87, PhiLho, Phiarc, Philip Trueman, Piano non troppo, Piotrus, Pogogunner, Praveentech, PseudoSudo, Public Menace, Pyabo, Qst, Quietbritishjim, RJaguar3, Red Thrush, Redleaf, Registrymaster, Reisio, Reswobslc, Richard416282, RingtailedFox, Rl, Rmsuperstar99, Robocoder, Ronark, Ronhjones, Roninbk, Rory096, Ross Fraser, Runtime, Rwxrwxrwx, SNIa, SchfiftyThree, Scientus, SealedSun, Secarrie, Senori, Setubuet, Shaddim, Shanesnh, Shirik, Simetrical, Simoncraythorn, Simxp, Slimscreen, Smyth, Snek01, SoWhy, Socrates2008, Soumyasch, SpuriousQ, SteveLoughran, Stevenmitchell, Success sathiya, SunCreator, Svarya, Ta bu shi da yu, Tagus, Tanvir Ahmmed, Tech Nerd, TechnoFaye, Technobadger, Tenox, The Anome, The Thing That Should Not Be, TheMightyOrb, Theda, Themodernizer, Thespian, Thingg, Think outside the box, Thumperward, Tide rolls, Tiibiidii, Titoxd, Tizio, Tobias Bergemann, Tognopop, Tom Pippens, Travis Evans, Trevj, Trevor Andersen, Tuntable, TurboForce, TuringMachine17, Twas Now, Twp, Ulric1313, Useerup, Valleyvalleyd, Valvicus, VanessaJoy, Veekay, Vegaswikian, Velella, Verdy p, Vidhatanand, Viggio, Vilerage, Vivekk, Vusys, Warren, Wasbeer, Wayfarer, Wiki alf, Wikievil666, Will Beback, Winterst, Wmahan, WolfmanSF, Xdenizen, Xgenei, Xpclient, Xzaranos, Yamamoto Ichiro, Yeung29hk, Yoenit, Yuhong, Yuletide, Ywong137, Zeth, Zowie, Zundark, Zzuuzz, 783 anonymous edits

Windows Management Instrumentation *Source*: http://en.wikipedia.org/w/index.php?oldid=501739112 *Contributors*: 2D, Abrahami, Abune, Adrianmoss, Ahbond, Akhristov, Alikhalidkhan, AlistairMcMillan, Arakunem, Ashok srinath, BWCNY, Brad101, Candace Gillhoolley, Circuitsgate, Cwolfsheep, Darkwind, EagleOne, Enviroboy, Fiach6383, Frap, Fuzzy, Gandalf44, Gardar Rurak, Ghettoblaster, Gorpkaiser, Haakon, Hoo man, IvanLanin, Jamelan, Jazzycat, John of Reading, Jpbowen, Leotohill, Lousyd, Martarius, Maximus Rex, Mwtoews, No1lakersfan, Olathe, Pearle, Pinkeltje, Pozcircuitboy, Rchaganti, Ringbang, SGBailey, Seifried, SkyMachine, Soumyasch, Srig83819, Tedickey, Tim Ivorson, Torc2, UtherSRG, Vikramrc, VirtualTycoon, Warren, WorldDownInFire, Xpclient, Yaronf, Yewtaah, Yurik, 126 anonymous edits

Daemon (computing) *Source*: http://en.wikipedia.org/w/index.php?oldid=504030342 *Contributors*: A plague of rainbows, AThing, Aardvark92, Absolwent, Ajfweb, Alan012, Albevier@helixit.net, Aleksa Lukic, Alfred Centauri, Alkim2, Ankur garg21, Anomie, ArinArin, Atlant, B Fizz, Beinsane, Benjamin Mako Hill, Borgx, Bovineone, Brian Kendig, Btilm, C. A. Russell, CWenger, Capslock0733, Captain panda, Chris the speller, Cleared as filed, Clq, Coachs, ComradeSlice, Copros, CyberSkull, Cybercobra, Dan100, Deboerjo, Deepugn, Deflective, Demonseth17, Diego898, DoctorODDS, E-pen, EEMIV, EWS23, Edward, EoGuy, ErrantX, Evilgohan2, Exidor, Fluzwup, Frap, Furrykef, Futurix, Ghiraddje, Giacomo Ritucci, Gscshoyru, Guy Harris, Gwern, H2g2bob, Hdante, Hmains, Intersofia, IvarTJ, Ivolucien, Jargon777, Jay, Jayjg, Jeh, Jerome Charles Potts, Jetxee, Jftsang, Jjk, Jnc, Joanna Ingles, John Vandenberg, JonHarder, Jowa fan, Julienlecomte, Kaptanteo, Karimarie, Kbdank71, Kbrose, Kerfuffler, King of Hearts, Knutux, Kop, Kubx, Kwamikagami, Labalius, Lapsus Linguae, Lasse Havelund, Legoktm, Letdorf, Lhasapso, LifeOnJinx, Loadmaster, MCBastos, MCWNT, Mac, MagnusW, Maimai009, Mallow40, Mark Richards, MattGiuca, Maury Markowitz, Maxim Razin, Michael Zimmermann, Mindmatrix, Mipadi, Mmtmmt, Mofoburrell, Mosca, Mykolas OK, Nad, Nickshanks, Nonrational, Nsupathy, Odpadek, Owen, Paceo, Pakaran, Patrick, Peter Flass, Pmj, Pondrumm, Progdev, Puttyschool, QcRef87, Quantumobserver, R'n'B, Ramu50, Ravi.varshney, Refael Alats, RekishiEJ, Revolus, Rich Janis, Rjwilmsi, Robert Brockway, Roboshed, Rubicon, Sargdub, Sdfisher, Shandris, Shinmawa, SimonP, Singingwolfboy, Skaizun, Slady, Slp1, Smably, Snowolf, Softtest123, Solitonwave, SpeedyGonsales, SpuriousQ, Sss367, Stephenb, Strait, Struthious Bandersnatch, Sungrazer, Synetech, Techiepedia, Tesi1700, The Monster, Threh, Thumperward, Tigga en, TimBentley, Tobias Bergemann, Toddintr, Unimaginative Username, Vadmium, Varlaam, Vbs, Velvetsmog, Wbeaumo, Wereon, Weyes, WhisperToMe, Wurblzap, X34l0t, Xoneca, Xrblsnggt, ZX81, Zoef1234, 232 anonymous edits

Runlevel *Source*: http://en.wikipedia.org/w/index.php?oldid=498332300 *Contributors*: 525252a, Abdull, Agarvin, Aidin Sabetian, AlanH, Alison9, Andareed, AnthonyA7, Arnocs, Banej, Beno1000, Burgundavia, CanisRufus, CesarB, Chachra, Clorox, Conan, CyberNeon, DavidAyers, Dawnseeker2000, Dcoetzee, Demian12358, Drevicko, EdwinGroothuis, Eeekster, Emperorbma, Epbr123, Exaton, Frap, Freedomlinux, Gamebouille, Henryhuang, Hm2k, Imroy, Jez9999, John Nevard, Kbolino, Kpirkle, Letdorf, Lkesteloot, Mbuell72, MiG, Mwtoews, Palosirkka, Pedant17, Radagast83, Ravenpi, Rbenech, Rdsmith4, Reedy, Rholt, Rickan, SF007, Shirulashem, Silvioq, Snori, Spoon!, Strcat, TheParanoidOne, Thumperward, Tjameson, Trizt, Tyler, Unixguy, Vssun, Wknight94, روخو, 111 anonymous edits

Windows service *Source*: http://en.wikipedia.org/w/index.php?oldid=502105605 *Contributors*: Anphanax, Birdsusing nnn, Bobo192, Bobscola, Borgx, Brainhell, Campoftheamericas, Chamath Mc, Cocjh1, Concisor, Dancter, DanielPharos, Ency, Femto, Frap, Furrykef, H2g2bob, Jamelan, Jarble, Jasper Deng, Jdmwood, Jim Henry, John Vandenberg, Mephiles602, Mh7kJ, Mrzaius, NickelShoe, Ohnoitsjamie, Piet Delport, RattleMan, Robenel, Rsargent, Seema2000, Soumyasch, Spike0xFF, Stepanz, Tyomitch, Utcursch, Veinor, Warren, Xpclient, Xrblsnggt, 96 anonymous edits

Syslog *Source*: http://en.wikipedia.org/w/index.php?oldid=490814502 *Contributors*: Abune, Adams kevin, Adrianmsmith, Alexisfisher, AntonChuvakin, Asfarer, Azakhark, Berford, Blaxthos, Brolin Empey, Bschoeni, Chris the speller, ChrisGualtieri, Claytondukes, Clemente, DGerman, Dnakla, DokReggar, Eguroos, Eldar40k, Felipe1982, Frap, Fudoreaper, Hilitec, Hu12, Husky, Jeromebell, Julesd, KLuwak, Kbrose, Klaber, Kwi, Manadar, Martinkop, Mhenriday, Michael B. Trausch, MrOllie, Mykhal, NoExec, Novasource, Ohnoitsjamie, Pjahr, Rgerhards, Ryanaxp, Scientus, Secret, Sirex98, Slaniel, Squash1978, Starsoftware, Stephen Turner, Superp, Svick, Tedickey, Themfromspace, Thomasknowles, Thumperward, Trusilver, Veekay, Warren, Woohookitty, YordanGeorgiev, 197 anonymous edits

Event Viewer *Source*: http://en.wikipedia.org/w/index.php?oldid=502257444 *Contributors*: Agrigorof, Akhristov, Art LaPella, Blurback, Capslock99, CliffC, Dfantom, Donkeykong333, Edd Porter, Eidab, Espoo, Fleasoft, FleetCommand, Hfrancois, Ishida639, Jeromebell, Julesd, Jusses2, Kurt.hudson, Lionel Elie Mamane, Mortense, MrOllie, Niceguyedc, Nurg, Rjwilmsi, RoySmith, Sampart, Sceptre, Soumyasch, Srehorn, Stayman Apple, The Evil Overlord, Tide rolls, Trusilver, Vegaswikian, Veinor, ViperSnake151, Warren, Wasisnt, Wikit2007, Winterst, Xpclient, Zulu1001, 51 anonymous edits

Hardening (computing) *Source*: http://en.wikipedia.org/w/index.php?oldid=498663992 *Contributors*: Antti29, Atama, CanisRufus, DanielPharos, Equendil, Frap, Interiot, Jarble, JonHarder, K kokkinos, Kozuch, Msbalensiefer, NapoliRoma, Nikai, RHaworth, Sverdrup, Tordek ar, Tremaster, Victor Stinner, Wmahan, Www-data, Yonkie, 21 anonymous edits

Microsoft Baseline Security Analyzer *Source*: http://en.wikipedia.org/w/index.php?oldid=492862194 *Contributors*: BWCNY, BladeRunner99, Carlo87, Cwolfsheep, DanielPharos, Dugn, FleetCommand, Instinct, Lkt1126, Ohconfucius, Pmsyyz, Random User 937494, Rich Farmbrough, Ryan Norton, Silvergoat, Spikey, Tuanese, Warren, Xpclient, 16 anonymous edits

Troubleshooting *Source*: http://en.wikipedia.org/w/index.php?oldid=495654732 *Contributors*: AbsolutDan, Aki009, Atlant, Audionaut, Blakegripling ph, Bobster999, Borgx, Btrewern, Camw, Chamal N, Chuckhoffmann, CliffC, Credible58, DVdm, Darrell Greenwood, Debresser, Delcoder, Denisarona, DerHexer, Duckbill, EagleOne, Eternalbizlinks, EuTuga, Evan G, Gadget850, General announcement, Georgesidhom, Glenn.isaac, Greyskinnedboy, Hbent, Intelligentsium, Jaberwockynmt, Jamelan, Jeremy Visser, JimD, Jimwilliams84, Julie Dancer, JustAnotherJoe, Karengpve, Kuru, LuisOntanon, Lwalt, Magister Mathematicae, Malcolma, Marschv, Mcfly85, Mermaid from the Baltic Sea, MichaelBillington, Mightyafro, Mikiemike, Paulscho, Peterlewis, Pinethicket, Pol098, R'n'B, Reedy, Rich Farmbrough, Rjwilmsi, Ronelex, Ronhjones, Ronz, S.Örvarr.S, SJP, Skysmith, SmokeySteve, Spalding, Taelus, TangentCube, Teratornis, The Thing That Should Not Be, TheDJ, TheParanoidOne, Tlesher, Trusilver, Zhenqinli, Zzlukes, 162 anonymous edits

Installation (computer programs) *Source*: http://en.wikipedia.org/w/index.php?oldid=501019631 *Contributors*: -Majestic-, 16@r, A.M.962, AVRS, Albeetle, Altenmann, BabbaQ, Betacommand, Blakegripling ph, Blanchardb, Bovineone, Courcelles, Crimenr1, Csc300c0, DVdm, Dames01, David Latapie, Den fjättrade ankan, Diamondland, ERcheck, Ebrescia, Elkman, Emre D., FT2, Fe'our, FleetCommand, Fraggle81, Furrykef, Gennaro Prota, Gil mo, Grozo, Hanyhussain, Happyisenough, Harriv, Harry the Dirty Dog, Harryboyles, Hebrides, Invitamia, Jemiller226, Jerryobject, JonHarder, Joy, Jpbowen, Julien.ponge, Jzylstra, Kalazadoo, Kazu89, La comadreja, LastChanceToBe, Lhasapso, Logan1337, Lupo, M4gnum0n, Mac, Mahanga, Mark Arsten, Mdsam2, Mike, Miles, Mintleaf, Moc03, Mr. Shoeless, Mr.muscolo, MrJones, NerdyNSK, Pearle, Pol098, Porthugh, R'n'B, Radagast83, Raffaele Megabyte, RedWolf, Rfc1394, Robenel, Robertd, Rsm99833, Salix alba, Sergei, Smack, Socrates2008, SpaceFlight89, Stevenrasnick, Struway, Swarm, Tgeairn, The Nut, Thomas Maierhofer, Thumperward, Vinhtantran, Vlad, Vladislav Biryukov, VladislavBiryukov, Walk&check, Wbm1058, Wikiuser2007ex, Wrldwzrd89, Zalnas, Zondor, Zumbo, 111 anonymous edits

Package management system *Source*: http://en.wikipedia.org/w/index.php?oldid=501797783 *Contributors*: 16@r, 1exec1, A.Abdel-Rahim, Aacri, AhmadSherif, Alasdairmacdonald, Alerante, AlistairMcMillan, Alt-sysrq, Andrew Rodland, Anon lynx, Antonielly, Apyule, Ash-Fox, Ausjak09, Bahadır, Bchilll, Bcornec, Beinsane, Blizinsk, Borgx, Burpelson AFB, CQ, CanisRufus, Cbayly, Charles Merriam, Chealer, Chesemonkyloma, Chowbok, Ckamps, Damian Yerrick, Dawnseeker2000, DeXXus, Derek Ross, Dielson Sales, Dori, Dreftymac, ENeville, Edward, Epbr123, Eraserhead1, Erik Sandberg, ErkinBatu, Feezo, Ffangs, Fiftyquid, Fleminra, Frap, Gaeddal, Gareth.randall, Gert4gt, Gopher23, Greenrd, Gronky, Guy Harris, Gwern, Hhielscher, HowardBGolden, Isilanes, JLaTondre, Janizary, Jernejovc, Jesdisciple, Jim Kring, Jj137, Jjk, Jklamo, Jnc, Joy, Jstrater, Kbolino, Kbrose, KitchM, Kl4m-AWB, Kosebamse, Kreutzm, Letdorf, LittleWink, Luc4, Lucid, Mbiven, Mikemaccana, Miles, Mindmatrix, Nabber00, NapoliRoma, NicM, OMouse, Oogbus, OrgasGirl, Peak, Pegua, PeteVerdon, Phok, Plaes, Psz, PuerExMachina, Pxma, R'n'B, RJHall, Reedy, Rhwawn, Ricgal, Rnsanchez, Robert Merkel, Roland Richter, Ron Ritzman, RyanTMulligan, SF007, Sal.farina, Salahx, SamB, Samsara, Samuella, Separa, Simon04, Skreyola, Sligocki, Sman789, Smithj, Spamsara, SpookyET, Stannered, StenSoft, Stephen Gilbert, SteveLoughran, Tedder, Tero, Thijs en, ThijsN, Thumperward, Tjyang, Toddyt1, Toutoune25, Trevj, Vaganyik, ViperSnake151, Wapcaplet, Wvdschel, Xyb, Yefi, Yworo, 159 anonymous edits

Advanced Packaging Tool *Source*: http://en.wikipedia.org/w/index.php?oldid=496825114 *Contributors*: A.M., Adam Zivner, Adsims2001, Aigarius, Ajdlinux, Ajunne, Aldie, Alerante, AlisonW, AntiRush, Applepeeler, Astronautics, Athalis, Bacchiad, Bender235, BiT, Binarybits, Bisqwit, Bluemoose, Brianreading, Bruce89, Calamari, Can't sleep, clown will eat me, Chealer, Clorox, Cowplopmorris, Czarkoff, DMacks, Daev, Danakil, DanielPharos, DavidDouthitt, Dawnseeker2000, Deemo, Delan Ahmad, Demitsu, DevAnubis, Dragon 280, Dreamyshade, Dromadar, Dumpendebat, Edward, Emesee, Emperorbma, Enigmasoldier, Epbr123, ErkinBatu, Ernstp, Faisal.akeel, Fleminra, Frankipedia, Frap, Giant toaster, Gioto, Glenn, Gothmog.es, Grandscribe, Greenrd, Gronky, Guanaco, H2g2bob, Hamish2k, Haneluya, Imz, JLaTondre, Jandalhandler, Jaycee55, Jeffrey Smith, Jeltz, Jengelh, Jeremy Visser, Jnc, JorgePeixoto, Joy, Jóna Þórunn, KDesk, Kenny Strawn, Kjoonlee, Kl4m, Kl4m-AWB, Kvng, LFaraone, Leandrod, Lengau, Lightmouse, Linghow, LittleDan, Mac, Magnus.de, Mahjongg, ManoloKosh, Markhobley, Master shepherd, Maximus Rex, Mclaud2000, Medovina, Meoblast001, Merphant, Mkjo, Mmmready, Monk127, Morqueozwald, MrKIPS, Musiphil, Nickaubert, Nk, Nneonneo, Nopetro, Notwist, Olivier, Omegatron, Orangemike, Orderud, PatchesTheCaveman, PatriceNeff, Paul Pogonyshev, PhilKnight, Phoenix-forgotten, PichuUmbreon, Piet Delport, Pluke, Pmmollet, Pnm, Poobah, Poppafuze, Psychonaut, Raph, Rchamberlain, Rchandra, Redlazer, Reisio, Resrever, Rpeh, Rwh86, SF007, Samsara, SeanTater, Secular mind, SickTwist, Sirowain, Snori, Somercet, St33med, SudoGhost, Tdangkhoa, Tempmj, Tero, Terrible Tim, Thephotoman, ThijsN, Thumperward, Timothy Gu, Unara, Unforgettableid, Venemo, Where, Widefox, Wwwwolf, Yonatan, Yworo, Zidane2k1, Пика Пика, 167 anonymous edits

dpkg *Source*: http://en.wikipedia.org/w/index.php?oldid=499743837 *Contributors*: (:Julien:), 16@r, Alfasst, Astronautics, Bachinchi, Beathovn, Bowmanjj, Calvinps, Chealer, Coldacid, CoolingGibbon, Danakil, Demitsu, Dmd, Doors5678, Edward, Emperorbma, Fleminra, Fopkins, Frap, Grandscribe, Gronky, Jleedev, Jnc, Joy, Jredmond, KDesk, KP-Adhikari, Kl4m-AWB, Locos epraix, Lotje, Mac, Magnus.de, Matthaeus123, Medovina, Meeuw, Merphant, Michael B. Trausch, Mulad, Nasir8891, Ohura, Ootachi, PradeepWarade, Praveen Illa, Proger, Rich Farmbrough, Sanxiyn, Skyhook19, Sme, Tobias Bergemann, Totakeke423, Where, ZZyXx, 43 anonymous edits

Windows Installer *Source*: http://en.wikipedia.org/w/index.php?oldid=504794093 *Contributors*: Abdull, Akhristov, Alasdairking, Anadem, Anna Lincoln, Archgriffin, BWCNY, Biruitorul, Bovineone, Burgercat, Caltas, CanadianNine, CesarB, ChorizoLasagna, Christopher G Lewis, Chrpai, Cjrcl, Ckatz, Codename Lisa, Concavelenz, CyberStormZA, DanielPharos, Dannyxyz22, Dayal Ramachandran, DennisIsMe, DragonHawk, EagleOne, Emre D., Evrik, Faganp, Fishnet37222, Flcelloguy, FleetCommand, Fran z, Frap, FreeRangeFrog, GSMR, General Rommel, Ggasp, Ghettoblaster, Gogo Dodo, Goodone121, Greba, Harryboyles, Heaths, Ivan Pozdeev, JFQueralt, Jacob Poon, Jaericho, James Arthur Reed, Jeh, Jerome Charles Potts, John Vandenberg, John of Reading, Jorend, Julesd, Kiamlaluno, Killdevil, Legalize, Leikfaz, Linux731, Little Professor, MER-C, Ma7moud05, Martijn Hoekstra, Maxwell's Demon, Mayankjohri, Mcorazao, Mephiles602, Mhowkins, MichaelR., Mikepelley, Mipadi, Misterds, Mortense, MrJones, Msihelp, Ngyikp, Nil Einne, Nol888, O, OLLY2K9, OhanaUnited, Pilaftank, Piyushnasa, Pol098, Porphyrula, RandomUser1234, RekishiEJ, RenniePet, Requestion, Rich Farmbrough, Rlinhartpdx, Robertd, Ronz, Russell Pack, S.Örvarr.S, SF007, Sam Hocevar, Scot72, Scottwilleke, Sevvie, Skizzik, Socrates2008, Sopoforic, Soumyasch, SquarePeg, Stephenchou0722, Stevage, Swati.agrawal, Synchronism, Thumperward, ToddGreer, Toussaint, TreyVanRiper, Txomin, Umar1996, Uzume, Vegaswikian, Warren, Wiikipedian, WikHead, Wikifane12, William James Croft, Wing gundam, Woohookitty, Xmm0, Xpclient, Xrblsnggt, Yagibear, Yms, Yukoba, Zabadab, Zadirion, Zenohockey, පසිඳු කාවින්ද, 233 anonymous edits

Windows Update *Source*: http://en.wikipedia.org/w/index.php?oldid=500364155 *Contributors*: -Majestic-, 16@r, A:-)Brunuś, Abanima, Abune, AdamJacobs, Adraeus, Akhristov, AlistairMcMillan, Andrevan, Arnoutf, Astroview120mm, AussieLegend, Bebo2good1, Bevo, Boffy b, Bonusballs, Bsfmig, Bubba73, CMG, Calmer Waters, Capricorn42, Ch egor, Chadlupkes, Ckatz, Cyberlemming, Cyrius, Czeror, Danny, Darianthomson, Ddehghan, Decat, DevastatorIIC, Djsonik, Dysprosia, EagleOne, Edward, Elkman, Eric567, EwokiWiki, FT2, Fernvale, Fishnet37222, FivePointPalmExplodingHeart, FleetCommand, Fran z, Furrykef, GPHemsley, Gardar Rurak, Gargaj, Gary King, Gary63, GoddersUK, Graham87, Grutness, Gvf, Hooperbloob, I Love Pi, Im-Dark, Insolectual, Itsmine, Iuhkjhk87y678, J.delanoy, Jasper Deng, Jesant13, Jhdezjr, JonHarder, Jorgenpfhartogs, Justin Eiler, Kaini, Katarighe, Kickass331, Kimchi.sg, Kipholbeck, Kozuch, Krackedude, La goutte de pluie, Leeyc0, Lemonade51, Lowellian, MMiladin, Mahanga, Manfromthemoon, Maniago, Marsilies, Mephiles602, Metarhyme, Minghong, Mjb, Moralis, Mschumacher69, Mvgfootball, Naryathegreat, Navy Blue, Nbarbettini, Nealmcb, Necessary Evil, Network master, Nick Number, Nyttend, Off!, Panser Born, Peace and Passion, Phoenix-forgotten, Pol098, Poor Yorick, Poweroid, R'n'B, Rchamberlain, Reb0rn087, Rebroad, RenniePet, Reptaboy, Rich Farmbrough, Ronark, Rubybarett, SNIyer12, Schmiteye, SchuminWeb, Sebastian Jaenicke, ShineyMetal, Shinjiman, Silvergoat, Sir Baka Jones, Soulkeeper, Soumyasch, Spe88, Splittorff, SquareWheelGames, Stephenchou0722, SufferWell1396, SuperSonicTH, Swotboy2000, T-1, Talrias, Tbhotch, Tedickey, ThePlaz, Themodernizer, Thompson.matthew, Thumperward, Thunderpenguin, Tosha, Tuanese, UAhmad7, Vegas Bleeds Neon, Veyronb, Vlad, Voidvector, Warren, Wiki Raja, Wikinger, Wjemather, Xpclient, Yamamoto Ichiro, Yeeshenhao, YuMaNuMa, Zackaback, Zalgo, Zirix, Δ, පසිඳු කාවින්ද, 158 anonymous edits

Windows Server Update Services *Source*: http://en.wikipedia.org/w/index.php?oldid=500386860 *Contributors*: Akhristov, AlanK, Batareikin, Bhami, Chinmayparekh, Colfer2, CosineKitty, DARTH SIDIOUS 2, DerekMorr, FleetCommand, Gary63, Genius101, Henare, J04n, Janipewter, Jeffmartel, JonHarder, La goutte de pluie, MCs88, Mocky3497, Muramas, Mushroom, NevilleDNZ, Nityanandparab, Notbyworks, Officer781, Paul Comfort, Pedant17, Pmsyyz, Requestion, Rich Farmbrough, Romanc19s, Scott.somohano, Snori, Socrates2008, Soumyasch, Syp, Tassedethe, Tico285, Tinucherian, Vegaswikian, Vindictive Warrior, WLU, Warren, Wpadron, Ziggitysplit, Zoombody, Zootm, 82 anonymous edits

WiX *Source*: http://en.wikipedia.org/w/index.php?oldid=498701996 *Contributors*: AmirCohen, BCube, Chrisg123, Christopher G Lewis, Ciphergoth, Cmdrjameson, Den fjättrade ankan, DmitTrix, Dysprosia, EagleOne, Firsfron, Francs2000, Frap, Fredrik, Free Software Knight, Gadfium, GregorB, Herakleitoszefesu, Ian Vaughan, Ijcammarata, Jessemonroy650, Jim1138, Joelholdsworth, John Vandenberg, Julesd, K1Bond007, Kl4m-AWB, Ksaha.pune, Leotohill, M4gnum0n, Mad burcher, Mcld, Miles, Mms, Neustradamus, Niteowlneils, Nzd, Penartur, RenniePet, Rich Farmbrough, Ronark, Runtime, S.Örvarr.S, Securiger, Socrates2008, Soumyasch, Stevage, ThreeBlindMice, Thumperward, Tobewan, VladV, WalterGR, Warren, Widefox, Zootm, Ævar Arnfjörð Bjarmason, 62 anonymous edits

Preboot Execution Environment *Source*: http://en.wikipedia.org/w/index.php?oldid=503831517 *Contributors*: Ajayvb, Ary29, Baa, Betterworld, Bluemike, Bovineone, Canaima, Chnv, Cmdrjameson, Cwolfsheep, DARTH SIDIOUS 2, DGerman, David Legrand, Dillo, Dreadstar, Electron9, Frap, Frencheigh, Fumitol, GermanX, Greenrd, Guaka, Hadal, Hpa, IOLJeff, J.smith, James.d.carlson, Jdthood, Jeh, Jenglish02, Jobbon, Joel7687, Jolomo, JonHarder, Jrdioko, Komet, KumarDSukhani, Lonwolve, Lxr1234, Mayuresh314159, Mirror Vax, Mozzerati, Nirkus, Nsda, OwenBlacker, Pedant17, Rbowman, Rchandra, Ringbang, RobinMiller, Ruby.red.roses, Sam Hocevar, ScytheBlade1, SimonEast, SudoGhost, Thumperward, Tobias Bergemann, Transfinite, Voomoo, Welshie, WikHead, WinTakeAll, Winterst, Wkpxed, 144 anonymous edits

Dynamic Host Configuration Protocol *Source*: http://en.wikipedia.org/w/index.php?oldid=504946722 *Contributors*: 78.26, 909078L, Abhayakara, Abhi3385, Abisys, Abune, Acroterion, AdamJudd, Adoniscik, Afiler, Aitias, Alaniaris, Alansohn, Aldie, Alex rosenberg35, Alfio, Alison22, AlistairMcMillan, Allstarecho, Almightylinuxgod, Alphachimp, Alvin-cs, Amaccuish, Ampsarus, Andareed, Andre Engels, Andrei Stroe, Animum, Anuragkrgupta, Apokrif, Arichnad, Arkrishna, Ashakiran, Ashwin, Atomician, BL, BW, Bardo, BartRoos, BenAveling, Billdorr, Bobby D. DS., Bogdangiusca, Bookbrad, BrianOfRugby, Brucefulton, CALR, Can't sleep, clown will eat me, Celarnor, Celinama, Centrx, CharlieEchoTango, Chealer, Chris 73, Chris the speller, Chrumps, ClanCC, Cleared as filed, Cliffb, Closedmouth, Coastalsteve984, Conversion script, Coutcin, Ctmt, DARTH SIDIOUS 2, DKEdwards, DStoykov, DVdm, Danger, Daniel.Cardenas, Davee, Davidkazuhiro, Davidoff, Davidstrauss, Dbu, DeadEyeArrow, Dee Jay Randall, Denelson83, Desiremore, Dgies, Dgw, DigitalSorceress, Dols, Dreadstar, Dremora, Dylan Lake, EagleOne, Earendur, Echosmoke, Edward, Egjose, Ehn, Electrified mocha chinchilla, Elwell, Enjoi4586, ErikWarmelink, Evil Monkey, Ewlyahoocom, Feureau, Flash.killer, Fleung, Friday, FrostytheSnownoob, Fudoreaper, Furrykef, G Sisson, Gaius Cornelius, Gamera2, Gareth Owen, Gary King, Giftlite, Gimboid13, Grafen, Grenavitar, Guru cool, Gwalker nz, Hackingyourdna, Hadal, Hanjifi, Hannes Hirzel, HarisM, Harshbhai, Hazphi, Hazzamon, Hcberkowitz, Heathbar5477, Helix84, Hephaestos, Hide1713, Hughcharlesparker, Hughey, Infofarmer, Ironywrit, IsUsername, JTN, JaGa, Jasoltape, JasonWoof, Jasper Deng, Jimsve, Joeinwap, Joerg Reiher, Josh Parris, Joy, Jrvz, Jude Rosario, Karora, Kartik Agaram, Kawanzaleroy, Kbrose, Kentyman, Kevinmon, Kgfleischmann, Kinema, King Toadsworth, Kingpin13, Kinu, Krellis, Ksn, Kubigula, Kurykh, Kvng, LAAFan, Labongo, Lacen, Lahiru k, Lightmouse, Lkinkade, Looxix, Loren.wilton, Lowellian, M4gnum0n, M7, MER-C, Mahewa, Makelelecba, Mam711, Mandarax, Mani1, Materialscientist, MattGiuca, Matthuxtable, Maximbo, Mcgonrya, Mechanical digger, Memming, Michael Devore, Mike Dillon, Mike Rosoft, Mindmatrix, Ministry of Truth, Mking30, Moeron, Moondyne, Mr Stephen, MrExplosive, Msabramo, Mwtoews, NCdave, Nachmore, Nafmosaved, Nastajus, Naved 12, Netsnipe, Ngriffeth, Nichtich, Nikai, Nimiew, Nixdorf, Njroberts, Nnp, Nubiatech, Nutster, Nv8200p, Ocker3, Ocram, Ojw, Omniplex, One, Ootachi, PCHS-NJROTC, PRBryson, PV=nRT, Pablicosta, Pagingmrherman, Pandemic, Parag2010, Parbatyogesh, Paris.butterfield, PeaceNT, Pedant17, Peng, Phatom87, Philcha, Piet Delport, Pinkadelica, Pion, Pokemonmegaman, Prohlep, Pseudomonas, QRX, Quarl, RJaguar3, Raanoo, Radiant chains, Raghav, RattleMan, Razimantv, Rchandra, RedWolf, Reub2000, Rfc1394, Rich Farmbrough, Richwales, Rikboven, Rjstott, Rjwilmsi, Rl201, Rob Hooft, Rodrigo.haus, Roris.solaris, Rrburke, Rursus, SJP, Sam Korn, Sameer.sujeet, Savantas83, Schaea, Scott.somohano, ScottSteiner, Sdrtirs, Sesshomaru, Shady69, Shields020, Shiftoften66, Shiryaev, Skyfiler, Slazenger, Snarius, SpK, Spasemunki, SpeedyGonsales, SportWagon, Srbanator, Srharriott, Stephan Leeds, Stevenm.ict, Suburbanslice, Suffusion of Yellow, Sunray, Superborsuk, Superm401, Suruena, Swatithorve, TYelliot, Taochen, Taoqirkhosa, Tcb, Teemuk, Tellyaddict, Template namespace initialisation script, Teraya, The Thing That Should Not Be, The.valiant.paladin, Theking17825, Thomas d stewart, Thumperward, Tide rolls, Tim.sylvester, Tothwolf, Totty3478, Tremilux, True Pagan Warrior, Trusilver, Tyler.szabo, Ummit, UncleBubba, Utcursch, VernoWhitney, Vijeshchandran, Vocaro, Waggers, Wasisnt, Wdrazo, Wereon, Wiki104, Wikisierracharlie, Wilson.canadian, Winterst, Wisco, Wisher, Wizzard2k, Wk muriithi, Wrs1864, XLiquidIceX, Xhienne, Xionbox, Xoxon1kk1, Yesteraeon, YordanGeorgiev, Yuckfoo, Zac439, Zedla, ^demon, 909 anonymous edits

Trivial File Transfer Protocol *Source*: http://en.wikipedia.org/w/index.php?oldid=501966955 *Contributors*: 16@r, Anir1uph, Arichnad, B Fizz, Belamp, Beno1000, Biasoli, Blahma, Blowdart, Bonobo-ape, Ceyockey, Chealer, Courcelles, Crasshopper, Cupselus, Cwolfsheep, Danallen46, DavidHalko, DerekMorr, Dminkovsky, Dysprosia, EagleOne, EdwinGroothuis, Elthe, FatalError, Fcorthay, FranklinPiat, Frap, Gardar Rurak, Gerrit, GoingBatty, Guy Harris, Gzornenplatz, Haakon, Hairy Dude, HamidRezaBazoo, Hpa, Insanity Incarnate, Intgr, Ivansanchez, J.delanoy, JTN, JamesBurns, Jason Creighton, Jaymie94, Jim.henderson, Jnc, JonHarder, JustAGal, Kate, Kenyon, Kumarat9pm, Kwi, Liftarn, Lzur, Maarten Hermans, Macfac12, Mange01, Maxim, MinorContributor, Mormegil, Mstrobel, Nk, Norm, Ozga, Paxswill, Popup, Psychonaut, Radiojon, RasberryMoon, RedWolf, Rfportilla, Riki, Robneild, Rojomoke, Ryan Roos, Sam Hocevar, Shanes, SimonEast, SlackerMom, SpaceFlight89, Spatz, Strait, Superm401, Suruena, Tdmalone, Tim.sylvester, Timwi, Tobias Bergemann, Tylerdave, Unchqua, Vaxquis, Vertium, Vipinhari, Visualkev, Vocaro, Warren, Woohookitty, XqRG, Zanimum, Zhukun, Zundark, Δ, 165 anonymous edits

Disk cloning *Source*: http://en.wikipedia.org/w/index.php?oldid=495930751 *Contributors*: A.t.bruland, AlistairMcMillan, Anoopjohn, Austin.adams, Badger151, Bcnstony, Bluebusy, Bluemoose, Briansp, Brilliant Pebble, Ceyockey, CliffC, CyberSkull, Dan9186, Donvinzk, Drderail, Ehudshapira, Ellisbjohns, Fairycherry, Fortissimoplayer, Gadfium, Gaius Cornelius, Geoffr, GermanX, GraemeL, Greenrd, GregorB, Gregory.dulin, GsmithEIDW, Hohum, Hux, Iacap, J.delanoy, Jabar, JeffHyman, Jerome Charles Potts, Jhansonxi, Jjcf89, JonHarder, Jsaxton86, Lutskovp, Mac, Maltrap, Marasmusine, Maxime.Debosschere, Minghong, Minna Sora no Shita, Mongair, Neil916, NickSills, Omegatron, Pebblig, Philip Trueman, Pnm, Pol098, Quantumor, Rajah, Requestion, SF007, Santiago Roza (Kq), Sasha1024, Scot72, Scott D, Scott McNay, Scrap104, Semiclear, Some guy, Spellbinder, Stephenchou0722, Suaven, Svick, SymlynX, Techtonik, Thumperward, TingChong Ma, TomRobinson, Tomchiukc, ViperSnake151, W.F.Galway, X-filler2, Xmager, 162 anonymous edits

Preseed *Source*: http://en.wikipedia.org/w/index.php?oldid=459083154 *Contributors*: Artemgy, Athaenara, Dawynn, Frap, GuillaumeDelacour, INGfertility, Kdkeller, SF007, YUL89YYZ, 1 anonymous edits

Windows Deployment Services *Source*: http://en.wikipedia.org/w/index.php?oldid=501936685 *Contributors*: Crohnie, DMahalko, Ghettoblaster, InfotechGuyz, Materialscientist, Mlaffs, Scortiaus, Socrates2008, Soumyasch, TEMHoward, The wub, Xcvzxcvzxcvxcvzxcv, Xpclient, 27 anonymous edits

CFEngine *Source*: http://en.wikipedia.org/w/index.php?oldid=501161741 *Contributors*: Asav, Ayaham, Carlp101, CeciliaPang, CesarB, Csabo, DerekMorr, Djbclark, Druiloor, Emsearcy, Fiftyquid, Frap, Glenn, Great Cthulhu, HopeSeekr of xMule, Joshr, Joy, Kb, Kenyon, Ketiltrout, Megateuf, Midinastasurazz, MikhailGusarov, Minitrue, Moocha, NapoliRoma, NorwegianWikiMan, Ouc, Plasticity, RAFPeterM, Raffen, Rich Farmbrough, Shd, Tabletop, Thorwald, 98 anonymous edits

Network Time Protocol *Source*: http://en.wikipedia.org/w/index.php?oldid=503741886 *Contributors*: AThing, AbsoluteFlatness, Aeluwas, Alchemistmatt, Alexcat, Anasule, Arkrishna, Army1987, Aulis Eskola, Banej, Bco2003, Bdesham, Bear475, Bellhead, Blu28, Bluesky2000, Breno, Bryan Derksen, Bryan.burgers, Btm792, Calimo, Callidior, Carnildo, Coenen, ColderPalace1925, Coneslayer, Cxbrx, DanGunn, Daniel exb, DavidDouthitt, Dragon 280, Dulciana, ECMarshall, EagleOne, Echoray, Eidab, Eisnel, ElBenevolente, Electron9, Enjoi4586, Equendil, ErikStewart, Eyreland, Flums, Frap, Fred Bradstadt, Frodet, Fudoreaper, Galentx, Galoubet, Geronime, Giftlite, Gilesmorant, Glenn burdett, GrahamBignell, Hashar, Herbee, Hgerstung, Hu12, Ineuw, Isaac Dupree, JTN, Jaho, Janizary, Jean-Claude DeMars, Jec, Jeh, Jeremy Visser, Jim no.6, JimQ, Jnc, Joncnunn, Joonasl, Joy, Jushi, Kaare, Kaihsu, Karn, Kaszeta, Kbrose, KenrickKoo123, Kenyon, Kim Meyrick, Kinema, Kotasik, Kvng, Kyng, Lee Carre, Leeyc0, Lkesteloot, Lokeshmandvekar, Luspari, Lytel, Mannafredo, Marchioa, MarkMLl, Mds6901, Mdwyer, Meand, Michael Hardy, Mindmatrix, Mmccalpin, Mormegil, Mrdavidlaing, Mulad, Mushroom, N0YKG, N3ddy, Naddy, Naive cynic, NapoliRoma, Nixdorf, Ojigiri, Omniplex, PeterJeremy, Phaedrav, PhilipMW, PierreAbbat, R1c0ch3t, Raul654, RedWolf, Redrose64, Rees11, Rmallins, Rrandall, Sae1962, Saimhe, Sbauman, Serglogin, Sgeiger, Shadow1, Simbad82, Sir sigurd, Sleske, Smyth, Snodnipper, Spalding, Spiritia, Stenn, SteveKostecke, Sue D. Nymme, Suruena, T0pem0, Tcp-ip, Template namespace initialisation script, The Anome, Thumperward, Thunderbritches, Tobias Bergemann, Ummit, UncleBubba, Velella, Vleo64, Widefox, Widomski, Wikisierracharlie, Wrs1864, Xmm0, Xpclient, Yamara, Zahakiel, Åkebråke, 220 anonymous edits

Domain Name System *Source*: http://en.wikipedia.org/w/index.php?oldid=504801115 *Contributors*: (, 194.17.59.xxx, 28421u2232nfenfcenc, 2D, 4th-otaku, 777sms, A-moll9, ANONYMOUS COWARD0xC0DE, Abhayakara, Abhinav paulite, Adamantios, Adambiswanger1, Ahoerstemeier, Ahu, Aldie, Ale2006, AlexJP, Alfio, AlistairMcMillan, Allstarecho, Alphachimp, Alphax, Although, Alvestrand, AnK, Anastrophe, Andareed, Andrejj, Andrew Hampe, AndrewN, Andrewpmk, Anjaliiiiiiiiii, Anna Lincoln, Anon lynx, Anthony Appleyard, Anwar saadat, Apacheguru, Arcade, Ardenn, Arichnad, Arifin banjar, Arjayay, Art LaPella, Arvindn, Avoided, Axonizer, Ayla, Azaziel, Azraell, Bahaanet, Balajisarathi, Banej, Barek, Bart133, Bayerischermann, Beachy, Beard0, Bedekar.nilesh, Ben-Zin, BenM, Beno1000, Bevo, Bhadani, Bhairab.online, BigMoneyJim, Bigflavor, Binaryspiral, Blanchardb, Bloggeruk, Blue Dot, Bluezy, Boba5fett, Bobblewik, Bobo192, Bobrayner, Bocafan76, Borislav, Bovineone, Brandon, Breno, BrianTung, Brianga, Brichcja, Brucevdk, BryDye, Bryan Derksen, Btharper1221, Byronknoll, Cain Mosni, Calvin 1998, Calvinkrishy, Can't sleep, clown will eat me, CapitalR, Capricorn42, CardinalDan, Cate, Cbuckley, Centrx, Ceyockey, Chairman S., Chasindream, Chealer, Checco, Chenzw, Chipp, Chocolateboy, Chris fluppy, ChrisCork, Chrisdab, Christian List, Cjkporter, ClaretAsh, Cleared as filed, CobraA1, Codes02, Cohesion, Colfer2, Cometstyles, Conversion script, Coredesat, Corti, Cremepuff222, Cst17, Cyanoa Crylate, Cybjit, Cynric, DGMDGM, DNSstuff, Damian Yerrick, David Haslam, David.Monniaux, DavidCary, Davidshq, Davipo, Dbo789, Dcsutherland, Dee.tee, Delivery:435, Demitsu, Dennis Brown, DerHexer, Derat, Diannaa, Discospinster, Dithridge, Djkrajnik, Dman727, Dolda2000, DoomBringer, Doomdayx, Dougalf, Dr.queso, Draicone, Dreg mogul, Dripdk, Drj, Drjan, Drowne, Dtobias, Dwheeler, Dysprosia, EdJohnston, Editor plus, Edward K, Edward321, Egmontaz, Ehn, Elinruby, Eloquence, Elvey, Emdx, Epastore, Epbr123, Ephillips, Equendil, Eras-mus, Eric Kvaalen, ErikWarmelink, Evil Monkey, Explicit, Feezo, Fenitharbour, Feydey, Fifthnail, FisherQueen, FleetCommand, Fleminra, Flubbit, Forseti, Fortunecookie289, Frap, Fred Gandt, Fudoreaper, Funkysh, Funnyfarmofdoom, FvdP, Fæ, GFellows, Garas, Gary King, Gdr, Geoffspear, Gerbrant, Gerixau, Gilesmorant, Glane23, Glenn, Glrx, Gobbleswoggler, Gogo Dodo, GoingBatty, Graciella, GraemeL, Graham87, Grayshi, Greg.smith.tor.ca, Groogle, Guanxi, Gurch, Gutzalpus, Haakon, Hadal, HaeB, Hairy Dude, HalJor, Hamtechperson, Hardaker, Harris7, Hayabusa future, Hbatra, Hcberkowitz, Hdt83, Helix84, Helixblue, Henriwatson, Hilaryam89, Honta, Horsten, Htaccess, Hu12, IJMacD, Ibuzzo, Idont Havaname, Insanity Incarnate, Inter, InverseHypercube, Iohannes Animosus, Ispabierto, Iulianu, J.delanoy, JForget, JH-man, JSGould, JTN, Jackbrownwan, Jake Wartenberg, Jamyskis, Jarsyl, Jasynnash2, Jauerback, Javawizard, Jawsper, Jcw69, Jdthood, Jec, Jesant13, Jfpierce, Jgeer, Jh51681, Jhasai, Jhawkinson, Jiddisch, Jinlye, Jivecat, Jj137, Jlaidman, Jmcw37, Jnc, JoanneB, Joe Sewell, Jog1973, Johan Elisson, John Broughton, John Vandenberg, JohnSmith777, Johnarett, Jonah S., Joonas Govenius, Joseph Dwayne, Joseph Solis in Australia, Jpatokal, JuJube, Jy, Jyoshimi, Jyoujjawal, KGasso, KLLvr283, Kaconk666, Karada, Karengpve, Kateshortforbob, Katieh5584, Kbrose, Kelly Martin, Kenguest, Kinema, Kingpin13, Kinneyboy90, Kitplane01, Kizor, Kjdunlap, Kjetil r, KnightRider, Kop, Kraftlos, Krawi, Krellis, Krooga, Kungfuadam, Kurumban, Kvng, Kwamikagami, Kwertii, Labongo, Lambtron, Leandrod, Lee Daniel Crocker, Lemming, Lentoyip, LeoNomis, Liamdaly87, Liftarn, Lightdarkness, Lightmouse, Limajean34983, LimoWreck, LionKimbro, Lod, Lotje, Ludde23, Luna Santin, Lupin, Lysdexia, MC MasterChef, MER-C, MWD27, Maestro25, Mahanga, Mam711, Mange01, Marc Mongenet, Marcel van b, Marcod'Itri, Marechal Ney, Mark Arsten, Mark Bergsma, Mark7-2, Materialscientist, Matilda, Matt Darby, Matt.smart, MattGiuca, Matthewvelie, Mattias.Campe, Mattrobinson78, Maury Markowitz, Max Cheung, Maximillion Pegasus, Mblumber, Mendel, Merlion444, Merovingian, MetsFan76, Mfwitten, Mgechert, Michael Hardy, Michel Jansen, Mike Rosoft, Mikeborschow, Mindmatrix, Minesweeper, Minghong, MonRa, Moondyne, Moonriddengirl, Moskvax, MrJones, MrKris, Mrguyguy226, Mrzaius, Msantcroos, Msjadex, Mwtoews, Mysekurity, N0YKG, NTox, Nabber00, Nakon, Name99, Nanshu, Naraht, Narayanacharan, Natkeeran, NawlinWiki, Nburden, Nealcardwell, Neil916, Netesq, Netkinetic, NevilleDNZ, Nevlino, NewEnglandYankee, Niceguyedc, NickCT, NickW557, Nickmaine, Nilesh2293, Nimiew, Nishanth.fairfield, Niteowlneils, Nivix, Nneonneo, NotAnonymous0, Notheruser, Nubiatech, O, Oddity-, Olddocks, Oldwes, Oliver341, Omegagod, Omicronpersei8, Omiksemaj, Omniplex, Openhazel, Optimist on the run, Oshah, Oskar Sigvardsson, Ott2, Ottawahitech, Oxymoron83, Ozruta, Pagingmrherman, Pak21, Parsecboy, Patrick, Paul Stansifer, Paulson74, Pcb21, Pedant17, Peter Karlsen, Peterbrown1, Pgan002, Pharaoh of the Wizards, Phatom87, Philarete, Philip Trueman, PhilipMW, Piano non troppo, Pierre Monteux, Pilom, Pinethicket, Piotras, Pjrich, Pointer1, Poweroid, Pranav.ota, Puffin, Qasimalikhawaja, QmunkE, Quarl, Quincywalltech, Quintote, Qwertyca, Qwertyus, R4f, RAM, RachaelJetkins, Ragib, RainbowOfLight, Rameshbabu.itian, RandomStringOfCharacters, Rapomon, Rasbelin, Rasmus Faber, Ravedave, Raybellis, Reach Out to the Truth, RedHillian, Reeser8, Regibox, ReiDx, RetiredUser2, Rfc1394, Riana, Rich Farmbrough, Rick Block, Rick Sidwell, Rikkus, Rjd0060, Rjgodoy, Rob Hooft, Robertb-dc, Roger Davies, Ron Ritzman, Ronz, Rottenblackberry, Roughana, RoyBoy, Rrius, Rror, Rshin, RussBlau, Ruwanindika, Rwashi, S3000, SDC, SJP, SMC89, ST47, SWAdair, Salokha, Samboy, Sanbio, Sanjiv swarup, Scarian, Sceptre, Schmloof, Schnolle, Science4sail, Scolobb, Scott.somohano, Scratchy, Seaphoto, Sen86, Serezniy, Serie, Shadowhillway, Shadowjams, Shalom Yechiel, Shams9598, Shenme, Shentino, Shii, Simigh, Simon Dodd, Simonjonzie, Sionus, SirPavlova, Skarebo, Sleigh, Smalljim, Smallpond, Smappy, Smashville, Smichae, Snowolf, Snoyes, Socialservice, Sokool, SolarisBigot, Solveforce, Some jerk on the Internet, Sonomadiver@gmail.com, Sonuahluwalia, SpaceFlight89, Sparkie82, Spellmaster, SperoSpes, Spinboy, Spishco, Sstrader, StaticGull, StealthFox, Stephan Leeds, Stephen Gilbert, Stib, Stwalkerster, SuperHamster, Suruena, SusanLarson, Sweet473, SyntaxError55, Syp, TAG.Odessa, Takis, Tardis, Taw, TedPavlic, Tedder, Teppic, The Anome, The King's Peach, The Thing That

Should Not Be, The locster, Thefunkygibbon, Thumperward, Tide rolls, Tim1988, Timwi, Tirerim, Tobias Bergemann, Tobias Conradi, Tonsofpcs, Tony1, Torzsmokus, Totty3478, Towel401, Trantjd, Tree Biting Conspiracy, Tregoweth, Trusilver, Ttony21, Tualha, Turnstep, Twajn, Twinxor, Tylerdmace, Tyomitch, Ufirst, Unara, Uncle G, UncleBubba, Updatebjarni, UrsaFoot, Use Bombs, Valenciano, Vanished user 39948282, Vanisheduser12345, Vary, Vera Cruz, Versageek, VictorianMutant, Vivio Testarossa, Vldhcp, Vocaro, Vrenator, Vssun, W2bh, WRJ, Wasisnt, Wavelength, Wayward, Webmaster4india, Webzero, Weike9009, Wesley, Weyes, Wigren, Wiki alf, Wiki104, WikiReviewer.de, Wikibase, Wikibob, Wikibofh, Wikiborg, Wikijeff, Wikipedia.org@parkerpress.com, WikipediaEdit, Wikipelli, Will Beback Auto, Wk muriithi, Wknight94, Wmahan, Wrs1864, Xardion, XaroPoo, Xionbox, Yamamoto Ichiro, Yes-Chef?, Yogi de, Yurik, Yyy, ZSBoS, ZaffaNE, Zelphar, Zfr, Zidane2k1, Zro, Zundark, Zzuuzz, Пика Пика, 1764 anonymous edits

Root name server *Source*: http://en.wikipedia.org/w/index.php?oldid=501671287 *Contributors*: Aaronw, Abdull, AdamJacobMuller, Aeon1006, Alterego, Alvestrand, Amcl, Andareed, AndrewWTaylor, Ardonik, Bsmntbombdood, Buckshot06, C J B Scholten, C'est moi, Camembert, Cantus, Cate, CesarB, CheeToS, Chrisahn, Chupon, Codywohlers, CrypticBacon, D.i.l., Dan Atkinson, Dazzla, Dkarrenberg, Dmb, Dolda2000, Electriccatfish2, Ellery, Epohs, Flibble, Flyerman85, Fudoreaper, Gamersedge, GermanX, Gpmidi, Gracefool, Hannigan, Herve661, Iceshark7, Ilse@, Iorek85, Isarra, JDowning, JTN, Jake Wartenberg, Jasper Deng, John Hubbard, John of Reading, JohnSmith777, Karada, Katpatuka, Kbrose, Keinstein, Kizor, Kwertii, Liastnir, Lindberg G Williams Jr, Lukeaw, M7, MMuzammils, Maarten Hermans, Macshune, Madpoet, Magnus.de, Manchurian candidate, Marcod'Itri, Matthew Woodcraft, Mauls, Mbethke, Michael Hardy, Mindmatrix, Mlibby, Mro, Nakon, Niteowlneils, Nkagrawal, Nubiatech, Nv8200p, ORSN, Pace212, Pascal666, Pietbarber, Portstorm, Prunk, PseudoSudo, RHaworth, Rich Farmbrough, Robchurch, Roman Rakitov, Ruud Koot, Shii, Shiryaev, Sloverlord, Stovetopcookies, Stu, Sword, Syp, TakuyaMurata, Teknic, The Anome, Tonkie67, Txuspe, UsrBin, Vipinhari, Vishu kapur, Vkojouharov, WibblyLeMoende, Wikicrazier2011, Will-h, Winterheart, Wk muriithi, Wrs1864, Wwallacee, Xeryus, Xyphor, Yob kejor, Yonatan, ZeroUm, Zumbo, Zundark, 151 anonymous edits

Dynamic DNS *Source*: http://en.wikipedia.org/w/index.php?oldid=499514723 *Contributors*: (, A13xu, AbsolutDan, Aellis, Alerante, Amcl, Andy.melnikov, BellsTheorem, Bjfs, Breno, Byronknoll, Caliston, Cananian, CecilWard, Chealer, Chenchaoyi, Chuunen Baka, Davidshq, Davidu, Ddnsmogul, Edupedro, Etienne.navarro, EugeneKay, Folajimi, Frap, Fredrik, Gardar Rurak, Ghat, Gifer17, Gothicx, Haakon, HalJor, Icseaturtles, Imjustmatthew, Imroy, Imroy9, Intgr, InverseHypercube, Iridescence, J.delanoy, JTN, JaGa, Jcdericco, Jgarzik, Josh Parris, Karada, Kbrose, Keydaimon, Krellis, Krkwrgr, Kusma, Maury Markowitz, Mboverload, Mdwyer, Medende, Msebast, Nick in syd, Paraphelion, Pgr94, Philip Trueman, PierreAbbat, Pol098, Quarl, Requestion, Rookkey, Schzmo, Scott.somohano, Scottprive, Secondewijzer, Smallpond, Smithd, Smyth, Splendour, Square87, Thparkth, Tillmo, Tomchiukc, Tschild, Versageek, Wikimavin, Woohookitty, Wrs1864, Xpclient, Zondor, 179 anonymous edits

Microsoft DNS *Source*: http://en.wikipedia.org/w/index.php?oldid=498846705 *Contributors*: AJR, AlistairMcMillan, CanisRufus, CecilWard, CharlesClarkson, Chris the speller, Dasarp.mail, Dawnseeker2000, EagleOne, FastHandSteve, Fleminra, Frap, Jeff Westhead, Jonathan de Boyne Pollard, Kbrose, Kusma, Ligulem, Madman, Nasnema, Nrqm, Peyre, Pnm, Psychonaut, QuiteUnusual, Rjwilmsi, Scott.somohano, Shanes, Sietse Snel, Smallpond, SolarisBigot, Southpaw018, The Black Void, Tmh, Warren, Xpclient, ZacBowling, 38 anonymous edits

Active Directory *Source*: http://en.wikipedia.org/w/index.php?oldid=504895263 *Contributors*: 1exec1, 2mcm, 5 albert square, 7, AThing, Acdx, Achilees09, Adrian, Alereon, AlistairMcMillan, Alok.vij, Alro, Amniarix, Anaxial, Andre Engels, Andygreeny11, Anteaus, Arjun01, Art LaPella, Atomician, Av99, Avidmosh, Badgernet, Bairamkhan, Banano818, Barrylb, Ben Ben, Bishtmech, Bobprime, Bookgrrl, Bookofjude, Borgx, BrianWren, C45207, Calmer Waters, Canadian-Bacon, Canderson7, CanisRufus, Capricorn42, Carlosayam, Cbustapeck, Chasingsol, Chopeen, Chuunen Baka, Chzwhl, Chzz, Conversion script, Coremayo, Courcelles, Crkey, DMahalko, DRE, Dan East, DanShearer, DeepNorth, Defkkon, Defufna, Demiurge1000, Demonkoryu, Denisarona, DerHexer, Dhana82, Difu Wu, Disavian, Djoko, Dlrohrer2003, Dogcow, DougsTech, Dpv, Dqd, DragonHawk, Drmies, Dysprosia, E. Ripley, E. Sn0 =31337=, EEMIV, EagleOne, Eagleridge, Ed!, EdBever, Edupedro, Edward, Emendation, EnamelWildcat, Extra999, Falcon8765, Favonian, Ff1959, Finlay McWalter, FleetCommand, Frap, Frencheigh, Gaff, GandalfDaGraay, Gareth Griffith-Jones, Gary King, GcSwRhIc, Gdarin, Gkf101z, Gladmax, Gohiking, Groink, Gökhan, Hephaestos, Here, Highlandsun, Huffers, Igoldste, Imamusic, Init, JEETENDRAKAPSE, James086, Jann.poppinga, Jared Preston, Jasper Deng, Jay, Jdthood, Jeronim, Jesse Viviano, JesseGarrett, Jim1138, Jlundell, Joe.kraska, Joe407, Joelster, John Vandenberg, JohnWhitlock, Jonabbey, Jonathonblocker, Julesd, Justinm99, Jwreigh, Jxan3000, KD5TVI, Kasperd, KathrynLybarger, Kevin B12, Kingpin13, Kitamozihr, Kl4m, Kozuch, Kramsallad, Kungming2, Kwamikagami, Kycook, Kyle1278, Lankiveil, Leeyc0, Lfstevens, Ligulem, Lilac Soul, Limbo socrates, Little Mountain 5, Lotje, Luvcraft, Luís Felipe Braga, Magnus.de, Marc Kupper, MarkWahl, MarsRover, Matlou, MaverickSolutions, Merlion444, Michael B. Trausch, Michael Hardy, Michael McDonnell, Michaelsoft, MidgleyDJ, Mikemaccana, Millerj90, Minna Sora no Shita, Mkdw, Moondyne, Moskvax, Mr little irish, MrOllie, Mrichards0, Muhandes, Muhgcee, Mullachv, Natalie Erin, Nigelm, Nikhilgulhane, Niklarin, Njan, Noahspurrier, Northamerica1000, Nullie, Ohconfucius, Omicronpersei8, Opes, Orfen, PaulGarner, Paulprout, Peter bertok, Pgan002, Phatom87, Philip Trueman, Pill, Pinkadelica, Plugwash, Pnm, Poccil, PoliticalJunkie, Prolog, Psy guy, Public Menace, PuerExMachina, QTCaptain, Qvestchun, Qwirty, Rcsheets, Rdeluca, Redvers, Reedy, Requestion, Retodon8, RexNL, Rich Farmbrough, Ricky81682, Rillian, Rjwilmsi, Rlian, Robinvasan, Roblu, Rosiehausler, Rshilwant, Rupert Clayton, Rusty Russell, RxS, Ryan Norton, RyanG, SJP, Sahil 11186, Salmon92, Sam Hocevar, SamBBaran, Saruman7, SaulPerdomo, Scot72, Script1, Senator Palpatine, Sendahmed, Shadowjams, Shane grimm, Sherly255, Sietse Snel, Sillygates, Singhsro, SixSix, Smalljim, Sniffnoy, Soifranc, Soumyasch, SpaceFlight89, Spearhead, Spidern, SpuriousQ, Srice13, StopBob, Sujay.shinde, Superm401, Suprax, Susfele, Susieber, Sveno, Syp, Ta bu shi da yu, Tactical, Tawker, Tbhotch, TehRabbitt, Tellyaddict, Terence, TestPilot, Tfl, The Thing That Should Not Be, Thecheesykid, Thedarxide, Thenoblenew, ThreeBlindMice, Tide rolls, Tim.alsop, Tobias Bergemann, Tom k&e, TomRawlinson, Totty3478, Tpk5010, Tristan Horn, Turb0chrg, TuxyQ, Twimoki, Twsx, Tyler.szabo, UltraViolet-64, Uncle G, Unknown entity, Vegaswikian, Vladislav Artukov, Volt42, W00tfest99, Warren, Welsh, Whisky drinker, Wiki alf, Wikijeff, Wikipelli, Windofkeltia, Wk muriithi, Wrs1864, Xenophonf, Xpclient, Yggdrasil42, Ykhwong, Yooden, Zachlipton, Zaxxon, Zidonuke, Zophar1, 1014 anonymous edits

Cloud computing *Source*: http://en.wikipedia.org/w/index.php?oldid=504740507 *Contributors*: 1337H4XX0R, 1exec1, 21hattongardens, 28bytes, 623des, 842U, 9Engine, A. B., A.K.Karthikeyan, A.Ward, A520, A8UDI, A930913, Aaditya025, Aavindraa, Abb615, Abligh, Abrahami, Aburreson, Acandus, Acather96, Achimew, Acrooney, AcuteSys, Adam Hauner, Aditya.smn, AdityaTandon, Adjectivore, Adlawlor, Aflagg99, Agemoi, Ahunt, Ajraddatz, Akropp, Alankc, Alanmonteith, Alansohn, Alex5678, Alexamies, Alexlange, Allenmwnc, Allens, Amartya71, Ameliorate!, American Eagle, Amoebat, Amore proprio, AnOddName, Anamika.search, Anandcv, Andersoniooi, Andromani, Andy Dingley, Andyg511, Aneesh1981, Angrywikiuser, Anikingos, Anna Frodesiak, Anne.naimoli, Anonymous Dissident, Ansumang, AoV2, ArglebargleIV, Ashakeri3596, Ashley Pomeroy, Ashleymcneff, Asieo, AstralWiki, Auntof6, Aurifier, Avermapub, Avillaba, Avlnet, Avneralgom, Avoided, Axorc, B Fizz, BD2412, BDavis27, BaldPark, Bamtelim, Barmijo, Beanilkumar, Beer2beer, Beetstra, Bejnar, Belfry, Beltman R., Bencey, Bentogoa, Bernd in Japan, Berny68, Betswiki, Bidgee, Bigfella1123, Biglovinb, Bijoysr, Bikeborg, Bill.albing, Biscuittin, Bkengland, Blaxthos, Blvrao, Bob bobato, Bob.gourley@comcast.net, BobGourley, Bodhran, Bonadea, Bongwarrior, BostonRed, Bovineone, Bped1985, Bpgriner, BrennaElise, BrianWren, Brianwc, Bruce404, Bshende, Bubba73, Bulldog73, Bumm13, Burrows01, Bus stop, Businesstecho, ByM4k5, C. A. Russell, CPieras, Cajetan da kid, Calabe1992, CanadianLinuxUser, Candace Gillhoolley, Capricorn42, Carl presscott, Casieg, Casliber, CastAStone, Cbdorsett, Cemgurkok, Certitude1, Chadastrophic, Chanakal, ChangChienFu, Chapmagr, Chintan4u, Chishtigharana, Chmyr, Chris the speller, Chris55, Christian75, Chriswriting, Chuahcsy, CitizenB, Ckatz, ClamDip, CliffC, CloudComputing, Cloudcoder, Cloudcto, Cloudfest, Cloudnitc, Cloudreviewer, Cloudxtech, Clovis Sangrail, Clutterpad, Cmartell, Colleenhaskett, Colonies Chris, ConcernedVancouverite, Cornelius383, Courcelles, Cpl Syx, Crackerspeanut12, Craig.Coward, Cronos4d, Crysb, Csabo, Cst17, Curtbeckmann, Cybercool10, Czarkoff, D rousslan, DARTH SIDIOUS 2, DBigXray, DJ Clayworth, DSP-user, DVdm, Dadomusic, Dan6hell66, DanStern, Dancter, Danielchalef, Dannymjohnson, Darktemplar, Darkyeffectt, DarrenOP, DaudSharif, Davewild, David H Braun (1964), DavidLam, Davidogm, Dbake, Dbrisinda, Declan Clam, Deepalja, Deicool, Deineka, DeleoB, DellTechWebGuy, Denisarona, Derekvicente, Dethomas, Devinek22, DexDor, Dfxdeimos, Diego Moya, Digestor81, Dimos2k, Dinferno, Dinhthaihoang, Dlohcierekim, Dlohcierekim's sock, Dlowenberger, Dlrohrer2003, Dmcelroy1, Dmichaud, Dogbertwp, Dogsbody, Dominik92, DoriSmith, Dougher, Dpsd 02011, Dr.K., DrFurey, DragonflyReloaded, Drphilharmonic, DuineSidhe, Dyepoy05, E Wing, E0steven, EDC5370, EEldridge1, ESkog, EagerToddler39, Earthlyreason, Ebhakt, Eco30, Edchi, Eddiev11, Edeskonline, Editorfry, Eeekster, Ego White Tray, Ehheh, Einasmadi, ElKevbo, Elassint, Elasticprovisioner, Elauminri, Elbarcino, Eleven even, EliyahuStern, Elminster Aumar, EmilyEgnyte, Emrekenci, Enc1234, Eniagrom, Epbr123, Er.madnet, Eric, EricEnfermero, EricTheRed20, Estahl, Esthon, Ethicalhackerin2010, Ethoslight, Eusbls, Evagarfer, Everything counts, Evil saltine, EwokiWiki, Explorer09, Falcon8765, FatalError, Favonian, FayssalF, Fb2ts, Fcalculators, Feinoha, Felipebm, Felixchu, Fergus Cloughley, Fieldday-sunday, Figaronline, Figureskatingfan, Finngall, Firsfron, Fishtron, Five-toed-sloth, ForthOK, Francis.w.usher, Frap, Frdfm, Freddymay, Frozen Wind, Fvw, Fylbecatulous, Fzk chc, Fæ, GRevolution824, GVnayR, Gamber34, Gandalf44, Garima.rai30, Gary King, Gawali.jitesh, GcSwRhIc, Ggafraser, GhostModern, Giftlite, Ginsengbomb, Gkorland, Glane23, Glass Sword, Gmm24, GorillaWarfare, Gotocloud, Gozzilli, Grafen, Gram123, Gratridge, Greensburger, GregorB, Grim Littlez, Guillaume2303, Guliolopez, Gxela, HJ Mitchell, Haakon, Hallows AG, HamburgerRadio, HandThatFeeds, HangingCurve, Happyinmaine, Hardduck, Haroldpolo, Harryboyles, Haseo9999, Hatfields, Helotrydotorg, Hemapi, HenryLarsen, Heracles31, Heroeswithmetaphors, Heron, Heyitscory, Hfoxwell, History2007, Hogne, Hu12, Hughesjp, Hutch8700, Huygens 25, IANYL, IDrivebackup, IHSscj, Iamdavinci, Iamthenewno2, Ianlovesgolf, Ianmolynz, Identity20, Imeson, Iminvinciblekris, Imllorente, Imnotminkus, Imtiyaz 81, Imtiyazali4all, Indigokk, Info20072009, IronGargoyle, Ispabierto, Iste Praetor, Itangalo, Itzkishor, IvanLanin, J.delanoy, J04n, JDC321, JForget, JLRedperson, JLaTondre, JMJeditor, JaGa, Jack Bauer00, Jack007, Jackoboss, Jacob Poon, Jakeburns99, Jakemoilanen, Jamalystic, JamesBWatson, JamesLWilliams2010, JanaGanesan, Jandalhandler, Janniscui, Jarble, Jarrad Lewis, Jasapir, Javaeu, Jay-Sebastos, Jayvd, Jbadger88, Jbekuykendall, Jbucket, Jburns2009, Jdlh, Jean.julius, Jef4444, Jeff G., Jeh, Jerichochang97, Jerryobject, Jesse.gibbs.elastra, Jhodge88, Jht4060, Jim1138, JimDelRossi, Jimmi Hugh, Jimsimwiki, Jineshvaria, Jinlye, Jkelleyy, Jldupont, Jleedev, Jmc, Jmillerfamily, JoeBulsak, Johananl, John254, JohnCD, JohnJamesWIlson, JohnnyJohnny, Johnpltsui, Johnuniq, Jojalozzo, Jon.ssss, Jonathan.robie, JonathonReinhart, Jorunn, Jose Icaza, JoshDuffMan, JoshuaZ, Jpgordon, Jtschoonhoven, Julesd, Juliano, Juliashapiro, Jusdafax, Justinc, Jwt1801, JzG, Kawana, Keilana, Kenny Strawn, Kessler, Keynoteworld, Kgfleischmann, Kibbled bits, Kieransimkin, Kimdino, Kkh03, Klilidiplomus, Kmsmgill, Knodir, Knoxi171, Kompere, KotetsuKat, Kozuch, Kpalsson, Krenair, Kridjsss, Kristiewells, Kronostar, Kubanczyk, Kunitakako, Kuru, Kvng, Kyerussell, L Kensington, L200817s, Lainagier, Lairdp, LandoSr, LarryEFast, Larrymcp, Laser813, Latha as, Latticelattuce, Lavanyalava, Lbecque, Learn2009, Lee.kinser, Leifnsn, Leinsterboy, LemonairePaides, Lerichard, Letdorf, Li Yue Depaul, Linuxbabu, Lionheartf, Lipun4u, LittleBenW, Littlewild, Liujiang, Liyf, LizardJr8, Lmbhull, Lmp90, Loadbang, Logan, Lone boatman, Lord Roem, Lorddunvegan, Lordfaust, Lotje, Luigi Baldassari, Luk, Lysy, M.O.X, MJ94, MPSTOR, Maasmiley, Mabdul, Macecanuck, Machilde, Macholl, Mahjongg, MainFrame, Makru, Malke 2010, Malleus Fatuorum, Manasprakash79, Mandarax, Manishekhawat, Manop, Manusnake, Mariolina, Mark Arsten, Mark Renier, MarkCPhinn, Markoo3, Markqu, Marokwitz, MarshallWilensky, Martin.ashcroft, Martinwguy, Materialscientist, Mathematrucker, Mathglot, Mathonius, Mattg82, Maurice Carbonaro, MaxedoutnjNJITWILL, Mayur, Mbblake, Mckaysalisbury, Mdd, Meitar, Metapsyche, Metatinara, Mheikkurinen, Mhodapp, Michael Hardy, Michaelmas1957, MichealH, Miguelcaldas, Mike Rosoft, Mild Bill Hiccup, MillinKilli, Mindmatrix, Miracle Pen, Miym, Mkdonqui, Mlavannis, Mmanie, Mobilecloud, Mohamed Magdy, Momoricks, Monkey Bounce, Monkeyfunforidiots, Moonriddengirl, Mortense, Moswento, Motyka, Mpcirba, Mr little irish, MrKG, MrMarmite, MrOllie, Msfitzgibbonsaz, Msileinad, Muhandes, Mwarren us, Mwmaxey, Mychalmccabe, Nabeez, Nadeemamin, Nakakapagpabagabag, Nasnema, Nasserjune, Natishalom, NatterJames, Navywings, NawlinWiki, Nbarth, NeD80, Neko-chan, Nelliejellynoonaa, Nesjo, Neustradamus, NewEnglandYankee, Nfr-Maat, Nick Number, Nightscream, Nikoschance, Nimbus prof, Niri.M, Njcaus, Njkool, Nnemo, No1can, NocturneNoir, Nogburt, NonNobisSolum, Noookncranny, Nopetro, North wiki, Northamerica1000, Now3d, Nurnware, Ny156uk, Nyllo, OSborn, Obdurodon, Ohnoitsjamie, Ohspite, Oleg Alexandrov, Olegwiki, OllieFury, Oluropo, Om23, OmbudsTech, Onmytoes4eva, Optatus, Orange Suede Sofa, OsamaBinLogin, OsamaK, Oskilian, Otto ter Haar, Ourhistory153, PCHS-NJROTC, PJH4-NJITWILL, PMstdnt, Padmaja cool, Paj mccarthy, Palapa, Panoramak, Parveson, Patrickwooldridge, PaulQSalisbury, Paulmmn, PavelSolin, PcCoffee, Pea.hamilton, Pegordon, Person1988, Pete maloney, Peterduffell, Pgan002, Pgj1997, Phazoni, Phil.gagner, PhilJackson, Philg88, Philopappos86, Pinethicket, Piperh, Pisapatis, Pleft, Pmell, Pmj, Pmorinreliablenh, Pointillist, Poliverach, Pontificalibus, Popnose, Porqin, Pottersson, Ppachwadkar, Priyo123, Pumpmeup, Qarakesek, Qoncept, Quaeler, Quantling, QuentinUK, Qusayfadhel, R39132, RHaworth, RLSnow, RS-HKG, Rakesh india, Ramnathkc, Ramu50, Randhirreddy, Rangoon11, Ras67, Raushan Shahi, RayAYang, Razorflame, Ready, RealWorldExperience, Realtimer, Reaper Eternal, Rebuker, Recognizance, Reconsider the static, Reservoirhill, Reshadipoor, RetiredWikipedian789, Rich Farmbrough, Rich Janis, Richard Harvey, Richard.McGuire88, Richramos, Rickyphyllis, Rjwilmsi, Rms77, Rnb, Robert.Harker, Robert.ohaver, Robofish, Robscheele, Rolf-Peter Wille, Ronen.hamias, Ronz, Rosiestep, Rossumcapek, Rotiro, Rsiddharth, Rursus, Rutger72, Ruud Koot, Rw2, Ryanrs, SBunce, SC, Sakaal, Salix alba, SamJohnston, Samuraiguy, San2011, Sandipk singh, Sanjiv swarup, Sanmurugesan, Sanspeur, Sapenov, Sarathc, SarekOfVulcan, Sariman, SchreyP, Scientus, Scott A Herbert, ScottSteiner, Scottonsocks, Scwlong, Sean R Fox, Seaphoto, Sebastiangarth, Sfoskett, Shadowjams, Shajulin, Shakeelrashed, Sharktopus, Shawnconaway, Shaysom09, Shierro, Shinerunner, Shirt58, Shoaibnz, ShortBus, Shri ram r, Siddii, Silver seren, Simonmartin74, Sissyneck, Sivanesh, Sjames1, Skarebo, Skizzik, Skovigo, Skrewler, Skumarvimal, Skyrail, Slady, Slwri, Smileyranger, Smjg, Snaxe920, Snideology, Snowolf, Snydeq, SoWhy, SocialRadiusOly, Softdevusa, Songjin, Sp33dyphil, SpecMode, Spiritia, Spojrzenie, Sponsion, Sql er2, SqueakBox, Srastogis, Ssabihahmed, Staceyeschneider, Stbrodie1, StefanB sv, Stephan Leeds, Stephenb, SteveLoughran, Steven Walling, Stevenro, Steveozone, Stevevictor134, Stickee, Stickyboi, StrategicBlue, Stuartyeates, Subbumv, SudoGhost, Sumdeus, SunnySideOfStreet, Superbeecat, Superm401, Surajpandey10, SuzanneIAM, Sweerek, Swolfsg, Syp, TJK4114, Taicl13, Tangurena, TastyPoutine, Tatatemc, Tbhotch, Tdp610, Techman224, Technobadger, TechyOne, Teles, Terrillja, Tgeairn, ThaddeusB, The High Fin Sperm Whale, The Thing That Should Not Be, The wub, TheNightFly, Thecheesykid, Themfromspace, Thesurfpup, Thewebartists013, Thexman20, ThomasTrappler, Thorwald, Thumperward, Tiburondude, Tide rolls, TimFreeman701, Timwayne, Tinton5, Tisane, Tkdcoach, Tmguru, Tobias Bergemann, Tobych, Tom Morris, Tom6, Tomhubbard, Tony1, TonyHagale, Tonyshan, Torturedgenius, Tpbradbury, Tramen12, Tree Hugger, Trusilver, Trödel, Tsmitty31, Tungsten, TutterMouse, Tuwa, TylerFarell, Tylerskf, Ugarit, Ummelgroup, Uncle Dick, UncleBubba, UncleDouggie, Understated1, Undiscovered778, Uozef, Utcursch, ValStepanova, Valenciano, Valentina Ochoa, Vasiliy Faronov, Vasq0123, Veerapatr.k, Vegaswikian, Velella, Vermtt, VernoWhitney, Vertium, VicoSystems1, Victor falk, Vicweast, Vigilius, VijayKrishnaPV, Vipinhari, Vishwaraj.anand00, Vmops, Vondruska, Vpfaiz, Vpolavia, Vranak, Vrenator, Wackywace, Walshga, Wasami007, Washburnmav, Watal, Wavelength, WeisheitSuchen, Weternuni, Whitehatpeople, Whitehouseseo, WikHead, Wiki13, WikiLaurent, WikiPuppies, WikiScrubber, Wikipelli, Willit63, Windchaser, Winterst, Wizzard, Wlouth, WookieInHeat, Wooptoo, WootFTW, Wretman, Wtmitchell, Wtsao, Wtshymanski, Xe7al, Xiler, Xinconnu, Xxyt65r, Y2l2, YUL89YYZ, Yaraman, Yaris678, Ycagen, Yerpo, Yinchunxiang, Yonidebest, YoungManBlues, Yuchan0803, Yworo, ZX81, Zachnh, Zevnik, Zfish118, Zhanghaisu, Zundark, Zzuuzz, Çalıştay, Σ, 2514, דוד ש. anonymous edits

Secure Shell *Source*: http://en.wikipedia.org/w/index.php?oldid=504839013 *Contributors*: 194.237.150.xxx, 4lex, 613crazydude, 72Dino, 90 Auto, @modi, AThing, Adi4094, AeonicOmega, AgentPeppermint, Aillema, Alan J Shea, Alerante, Alexamies, Alexav8, Alismayilov, AlistairMcMillan, Allens, Alynna Kasmira, Amakuha, Analogue73, Andre Engels, AndriuZ, Anon lynx, Apokrif, Armando, Arthena, Artorius, Astronouth7303, Audriusa, Awalden0808, B4hand, BD2412, Ben Ben, Benjaminevans82, Beno1000, Bentogoa, Bernium, Bill.D Nguyen, Biot, Bomac, Borgx, Bovineone, Bpence, C.Fred, CCFreak2K, CWenger, Caltas, Captainjackfromgermany, Captainspalding, Cdc, Cflm001, Chealer, Chowbok, Cisguy, CommonsDelinker, Controloye, Conversion script, Corpx, Crazycomputers, CyrilB, D.j.potts@bcs.org.uk, DStoykov, Dan Austin, Dan Forward, Daniel.Cardenas, Davidsxls, Dawnseeker2000, DeadEyeArrow, Demonburrito, Denis bider, DerHexer, Dionyziz, Dissolve, Dlonceveau, Dmarquard, Doniago, Drinking.coffee, Drrngrvy, Drt1245, Dtwong, Dwlegg, EEMIV, EEPROM Eagle, Earthpigg, Eddie Nixon, Eeekster, Efa, Ehheh, Ehn, EmilSit, Emperorbma, Ency, Enjoi4586, Equendil, Ernstdehaan, Faradayplank, Finavon, Fleminra, Flo422, Folajimi, Frap, Galactic Dominator, Galoubet, Gardar Rurak, Gary63, Gbeeker, Gilesmorant, Giygas73, Glenn, Gosub, GotaForce, Grawity, Gronky, Guy Harris, Gwking, Hankwang, Harani66, Hawaiian717, Helix84, Hfastedge, Hobart, Honta, Hu12, Hugo 87, Hyad, Icarus4586, Inkling, Interiot, InverseHypercube, IvanLanin, JLaTondre, JTN, Jackal242, Janizary, Jasper Deng, Jdthood, Jemocri, Jerome Charles Potts, Jesse Viviano, JesseHogan, Jfromcanada, Jgrahn, JidGom, Jinlei, John Vandenberg, Jonkerz, KJRehberg, Karnesky, Kbrose, Kenguest, Kgfleischmann, Knarrff, Kotepho, Kravietz, Krellis, Krischik, Krithin, Kukini, Kwi, Kyng, L Kensington, L33th4x0rguy, LFaraone, LOL, Larry V, Laukster, Laurusnobilis, Lbecque, Ledow, Lerdthenerd, Lightdarkness, Ligulem, LittleBenW, LittleDan, Lklundin, Lord Bob, Lord Heinrich, Lotje, Lysdexia, MCBastos, Mabdul, Manuel Anastácio, Markpeak, Markskilbeck, Matt Crypto, Mayevski, Mcaruso, Mchirico, Mdupont, MeekMark, MegaHasher, Mentifisto, Merriam, Michael Hardy, MichaelBillington, Midnightcomm, Mikeblas, Mileswu, Mindmatrix, Mmernex, Mobystar, Mr.sanjibdhar, MrOllie, Mwtoews, Mykhal, NONCENSORED Popeye, Nandhpm, Nanshu, Nbauman, Nealmcb, Nicolas1981, NilsB, Nixdorf, Nk, Noahspurrier, Noya Watan, Nurg, Ohconfucius, Olaf Davis, Olivier, Omniplex, Only2sea, Ortzinator, Pascalv, Pearle, Peppergrower, Perspective, PhilipMW, Phoe6, Piepie, Pierremb, Plugwash, Prikryl, PrisonerOfIce, Profoss, Pseudometric, Qji, QueenCake, Quiddity, Random seed, Reconsider the static, ReiDx, Rich Farmbrough, Rifleman 82, Rjwilmsi, Robert Merkel, Rock drum, Romal, Romanm, Root2, Rootless, Rpresser, Schneelocke, SecurityBulletins.com, Seifried, Shd, Shellreef, Shibboleth, Shikaga, Siddhant, SimonEast, Sleske, Sohailms, Solvedo, Stannered, StefanPett, Stephan Leeds, Steven Luo, Superm401, Suruena, Suso, Sverdrup, Syp, TakuyaMurata, Tarquin, TastyPoutine, TerraFrost, The Thing That Should Not Be, The Zig, Thue, Tianjiao, TinaSDCE, Tirwhan, Tlroche, Tobias Bergemann, Tristanb, Tyrel, UncleBubba, Unixguy, Unixtastic, Urhixidur, Varnav, Verbose, WakiMiko, Welsh, Weregeek, Wiehenkm, WikiFlier, Winryder, Wjmallard, WojPob, Wojtekmejor, Wronkiew, Wrs1864, Ww, Youssefsan, Yueni, ZachPruckowski, ZeroOne, 442 anonymous edits

WS-Management *Source*: http://en.wikipedia.org/w/index.php?oldid=456225334 *Contributors*: AlistairMcMillan, Anrie Nord, Deville, Dkoper, EagleOne, Emersoni, Ghettoblaster, Joshrcohen, Maiquel, Mandarax, Modulum, Nashif, Nv8200p, Padlia abhay, Pnm, Toutoune25, Warren, Winterst, Wireless friend, Yaronf, Yumeyao, 25 anonymous edits

Group Policy *Source*: http://en.wikipedia.org/w/index.php?oldid=499052050 *Contributors*: Abdsufa, Aegicen, Ahoerstemeier, Alanburchill, Alcazar84, AlistairMcMillan, Amarshrestha, Andymease, Anetode, Anon lynx, Bd84, Borgx, CecilWard, Cheeni, Cleared as filed, Closedmouth, CobraWiki, D4mo, DanielPharos, DeadEyeArrow, Dmp1ce, Ebjers, Egamma, EoGuy, Epbr123, FT2, Fabricationary, FleetCommand, GGTM, Gary King, Gary Kirk, Ggugvrunt, Ghettoblaster, GoingBatty, Hadal, Hamitr, Harryeds, Ike9898, Irish Souffle, Jamelan, Jamesofur, Jeni, Jeronim, Jftuga, Kubieziel, Lightner, LilHelpa, MNGoldenEagle, Malvineous, Markwill, Marokwitz, Mephiles602, Mild Bill Hiccup, Moskvax, MrOllie, Nelsonheber, PatchesTheCaveman, PatrickDunfordNZ, Peyre, Pfkninenines, ProhibitOnions, RHaworth, Rajeshthennan, Rascalchops, Rich Farmbrough, Roblu, RonGrevink, RyanG, Sam Burke, Saqib, Scortiaus, Script1, Shandris, Smyth, Socrates2008, Soumyasch, Strike Eagle, Success sathiya, Syp, Ta bu shi da yu, Tech Nerd, Technophant, Thegalaxy, Turb0chrg, Utkarshashish, Vdm, Vegaswikian, W3bbo, Warren, Xpclient, Ykhwong, YolanCh, Zero sharp, Ziggitysplit, 207 anonymous edits

System Center Configuration Manager *Source*: http://en.wikipedia.org/w/index.php?oldid=503868244 *Contributors*: Abune, Akwoolf, AlistairMcMillan, Anoopmannur, Anrie Nord, Antti29, Biker Biker, Casimmons, Chamal N, ChangChienFu, CyberSkull, DancingPenguin, Danrawlings, Dfletter, EagleOne, Ehn, Engiorasta, Etrigan, Etsouris, Feinoha, FirestarXL, Frap, Gary King, Gavsto2006, Ghettoblaster, Hamitr, Humanoc, Jamelan, Jonik, JonnyH, Jsaylor3, JustinRossi, Kenyon, Kku, Krytenofsmeg, MarkDev88, Mcompton69, Mcorr, MeltBanana, Mhowkins, Misterds, Myleslong, Perturbance, Pmsyyz, PolarYukon, Protokon, Psychorob, Raystewart1, Rich Farmbrough, Richardt1980, Rodtrent, Rotshild, Signalhead, SiliconCerebrate, Smsarmad, Soumyasch, SpasticGorilla, Ta bu shi da yu, Tawker, TheParanoidOne, Tjkpay, Tonkie, Tweek, Warren, 140 anonymous edits

passwd (file) *Source*: http://en.wikipedia.org/w/index.php?oldid=494635148 *Contributors*: Alecv, Alevchuk, Apokrif, Cesium 133, Dmeranda, Download, Glenn, IMSoP, Kbrose, Lost.goblin, Marios.agathangelou, Mikemaccana, Miko3k, Mild Bill Hiccup, Paulderson, Pimlottc, Pnm, Stephen Morley, Thorwald, Vi2, WhiteDragon, 14 anonymous edits

Shadow password *Source*: http://en.wikipedia.org/w/index.php?oldid=494431237 *Contributors*: Alecv, Alerante, Alpha 4615, AndrewNeo, Andrewpmk, ArnoldReinhold, Asymmetric, Bender235, Bneely, CanisRufus, Chamal N, Chzz, CraSH, Dmeranda, FranklinPiat, Frap, Gbeeker, Ghepeu, Hmwith, JECompton, Jean Sreng, Jerome Charles Potts, Jfhaugh, JoeAshcraft, Joecamelman, Johnfn, KnightRider, Kotepho, Mikemaccana, MrJosiahT, Muerte, NapoliRoma, NoobX, Pingveno, Politepunk, RHaworth, Scientus, Scovetta, Sdfisher, Stereo, Taejo, Tholme, Thumperward, Tzadikv, Unixguy, Wiesel 99, Ww, Zaius, 48 anonymous edits

Security Accounts Manager *Source*: http://en.wikipedia.org/w/index.php?oldid=502934834 *Contributors*: Akhristov, AlistairMcMillan, Avocado27, Bb1netusaf2004, Bryan Derksen, Cwolfsheep, Dake, Edupedro, Frap, Gerold Broser, Grawity, Jasper Deng, JonHarder, Mauls, Mlondeen, ParanoidMike, Pnm, RedWolf, RichardGatinho, SF007, Spoon!, Template namespace initialisation script, Thugzclub, Tom Pippens, Urhixidur, Vitriol, Warren, Xpclient, Zobh, 34 anonymous edits

Home directory *Source*: http://en.wikipedia.org/w/index.php?oldid=492080776 *Contributors*: 16@r, Abdull, Adam Nohejl, Alerante, AlistairMcMillan, Appleseed, Beno1000, Bgruber, Blue Em, Bradv, CapitalSasha, Cschutijser, Darrenmoffat, Downcoder, Dysprosia, Eloquence, Emre D., Emx, Frap, Fredrik, Ghettoblaster, IMSoP, Jerome Charles Potts, Josh the Nerd, Jwoodger, Kawai kiwi, Kolyma, Korg, Leeannedy, Lotje, Lunarglow, MarkMascolino, Mernen, MichaelBillington, NapoliRoma, O1ive, Piccolomomo, Pilaftank, Pride the Arrogant, Psychonaut, Pxma, Raul654, Rbrwr, Spitzak, Thumperward, Tideflat, Twenex, Unixguy, Victor.zamanian, Warren, WilliamTheaker, Yath, Zzuuzz, روخو, 42 anonymous edits

Microsoft Windows profile *Source*: http://en.wikipedia.org/w/index.php?oldid=491279819 *Contributors*: A:-)Brunuś, JamesBWatson, KrisBogdanov, L736E, LinguistAtLarge, Lofote, Nabeth, Vrenator, Wikipelli, 14 anonymous edits

Roaming user profile *Source*: http://en.wikipedia.org/w/index.php?oldid=503844785 *Contributors*: A:-)Brunuś, AlistairMcMillan, Anteaus, BW95, Bhowden, Bmamlin, Caiaffa, CameronHarris, Camsipe, DMahalko, Doctorfluffy, Dosman, Edderso, FleetCommand, Frap, GoldenGlory84, JC Chu, JLaTondre, John of Reading, Josh the Nerd, Jæs, Kozuch, Marek69, PR123123, Pugglewuggle, Rjwilmsi, Rprpr, Shanes, SwirlBoy39, TomRES, Tuntable, Varinder k, VernoWhitney, Warren, WikHead, Xpclient, Zidane2k1, 85 anonymous edits

Directory service *Source*: http://en.wikipedia.org/w/index.php?oldid=501716004 *Contributors*: A bit iffy, Ahkilinc, Alan lloyd, AlistairMcMillan, Anetode, Angela, Anrie Nord, Asciberras, Bdamm, Beefyt, Biblbroks, Biznet, CanisRufus, Captain-tucker, Charivari, Chase me ladies, I'm the Cavalry, Chealer, Chrebbie, Closedmouth, Crediblehulk, Crkey, DStoykov, Danakil, David Shay, Davnor, Drake Redcrest, EagleOne, Edupedro, Edward, Elf-friend, Esprit15d, EvanProdromou, Falcon8765, Ff1959, Fmkrm, GerardM, Gerritdewitt, Hanswaarle, Hellisp, Hgfernan, Highlandsun, Isode, J.delanoy, JLaTondre, JakobVoss, Jay, Jiddisch, Jklowden, John Vandenberg, Josh the Nerd, Keebler, Kenyon, Khalid hassani, Kjell André, Kozuch, Kuru, L Kensington, Ligulem, Limbo socrates, Lullabud, Mahanchian, Mark Renier, MarkWahl, Mcskoufis, Mewsoft, Michael%Sappir, Mikemaccana, Mild Bill Hiccup, Minesweeper, Mo ainm, Motea, MrOllie, Ncmvocalist, Nickcarr, Nixdorf, Norberthegyaljai, Npmccallum, Ohnoitsjamie, Pnm, Poc0123, Ppp, Premil, R'n'B, Rajesh301190, Ranisquibbs, RedWolf, Reedy, Riaanvn, Richard.e.morton, Ringbang, Robvamos, Ronz, Ros Power, Rubix, SF007, Samrolken, Seomad2010, Shadowjams, Shai-kun, Spliffy, SuperHamster, Sutch, Ta bu shi da yu, The Anome, This lousy T-shirt, ToOb, Tobias Bergemann, Tomato86, Tps777, Vijayvkarthik, ViperSnake151, Wcs9000, WikHead, Wikidrone, Xezbeth, Zundark, 209 anonymous edits

Lightweight Directory Access Protocol *Source*: http://en.wikipedia.org/w/index.php?oldid=504677400 *Contributors*: 0, 62.202.117.xxx, 62.253.64.xxx, A8UDI, AJackl, ANGELUS, Abeld, Ahoerstemeier, Alexbatko, Almightylinuxgod, Amillar, Amux, Andareed, Andrew Findlay, Anna Frodesiak, Anna Lincoln, Anon lynx, Aragorn2, Arcade, Arknascar44, Armando, Asciberras, Atstarr, Augustan, Avijit.ghosal, BabsiSnoecks, Bahar101, Bcwhite, Beland, Benjamin Barenblat, Bevo, BigSmoke, Brandon, Breno, Brevantes, CKlunck, Cavrdg, Charles Brooking, Chealer, Chris Q, Christian75, Chuunen Baka, Combatentropy, Cometstyles, Crkey, DGG, Dalesc, Danpritts, Darrien, DavidHallett, Deflective, Denisarona, Dennette, DerHexer, Dewet, Diningphil, Dmsar, Doczilla, Dogcow, Dolda2000, DoubleBlue, DrRogla, Dustimagic, ESkog, EZio, EagleOne, Ed Poor, Edward, EmirA, Emmanuel JARRI, Emperorbma, Equendil, Exien, Ff1959, Flewis, Folajimi, Fred Bradstadt, Gamma, Georgesawyer, GerardM, Ghettoblaster, Giraffedata, Gman124, Gpallis, Great Cthulhu, Grovejary, Grs1969, Gurch, HFuruseth, Helon, Hhielscher, Highlandsun, Hotlorp, Hu12, IWhisky, IanManka, Iridescence, IronJohnSr, Isnow, J-stan, J.delanoy, JTN, Jaanis2010, JakobVoss, Jarsyl, Javawizard, Jay, Jbm212, Jdthood, Jeffschuler, Jiddisch, Joe Jarvis, JoeHenzi, Jonathan de Boyne Pollard, Joonga, Jordan Brown, Jordav, Juliank, Jwilleke, Jxn, Kbrose, Kdz, Khendon, Konman72, Kozuch, Kuru, Kvng, Kwamikagami, LCE1506, Lazki, Lee Carre, Lesterrivera, Levin, Looxix, Lost on belmont, Lotje, Lukegm, Lysdexia, MMSequeira, Magioladitis, Mallamace, Manavmohanty, Manishearth, Manster, Marc Esnouf, Marek69, MarkWahl, Marty.offe, Matt Crypto, Maximaximax, MilkMiruku, Mindmatrix, Monkbel, MooingLemur, Mortein, Mr stupid, Mrbill, Nabber00, Neelkanthtriptathi, Nick, Nippoo, OlEnglish, Omniplex, Oneiros, Oxymoron83, Pabix, Parpauta, Pascal.Tesson, PatrikR, Pedant17, Peterblaise, Pgan002, Phatom87, Pieffe, Piet Delport, Pit@0xff.at, Pnm, Polluks, Poor Yorick, Premil, Pxma, RFightmaster, RW Dutton, Rageear, Randy Johnston, Rasmus Faber, RedWolf, Redjar, Remi0o, Renaissancee, Rhubbarb, Rick Block, Rnapier, RodWiddowson, Ronabop, RossPatterson, Ruakh, Rwoodsmall, S3000, SJK, Sadads, Saligron, Sams sun, Scratchy, SebastianBreier, Shadow1, ShelfSkewed, Shshme, Smaines, SmartIcon, Snupha, Southen, Srpnor, StaticGull, Stephan Leeds, Stevage, Stickee, Subversive.sound, TRS-80, Taw, Taxman, Tbird1965, Tedp, Tero, The Master of Mayhem, The undertow, Thingg, Thumperward, Tide rolls, Timotab, Tlesher, Tobias Bergemann, Tobias Hoevekamp, Tothwolf, Towo, Treydrake, Turnstep, Tyler.szabo, UncleBubba, Uselesswarrior, Vanished user 39948282, Varlaam, VictorAnyakin, Vitund, WRJ, Waddey, Wageslave, Wbenton, Wereon, Wiki alf, Wilsonhughes, Wk muriithi, Wojtekpia, Woohookitty, Wsheward, Yingyang, ZZyXx, 638 anonymous edits

LDAP Data Interchange Format *Source*: http://en.wikipedia.org/w/index.php?oldid=461673926 *Contributors*: Bovineone, Cosort, Cxw, Dkrms, EagleOne, Ettrig, Ff1959, Fys, Gharman, GordonGood, Grawity, HFuruseth, Harro5, Highlandsun, Hu12, Jennifer Brooks, JimScott, Joejsullivan, Jonathan de Boyne Pollard, Juliank, Kdz, MartinDK, Navur, Pegacat, Smaines, Veledan, Wlly n Whlls, 44 anonymous edits

Kerberos (protocol) *Source*: http://en.wikipedia.org/w/index.php?oldid=503926631 *Contributors*: Abune, Agou, Aldie, AlexHajnal, AlistairMcMillan, Amux, Andareed, Andrejj, Angela, Ankurcha, Anon lynx, ArnoldReinhold, Aron Håkanson, Arvindn, AySz88, Bergsten, Blonkm, Bmusician, Bovineone, Bubba nuts, Bunnyhop11, Camaro889, Camelium, Cdinesh, Chealer, Choas, Christopher G Lewis, Chuunen Baka, Clemente, Cmglee, Colonies Chris, Conversion script, Crakkpot, Crono logical, Ctrlsoft, Danakil, David McBride, Dawnseeker2000, Deflective, Doug A. Hole, Dsonck92, Duhang, Eranb, Ericmc783, FleetCommand, Flexdream, Frits.a.m.storms, Fudoreaper, Fyyer, Gareth Jones, Gingekerr, Gioto, Gkf101z, Glenn, Grafor, Grawity, GregorB, Gronky, Guoguo12, Hadal, Haikz, Hede2000, Heybeppe, IanHarvey, Ignacioerrico, Int21h, Intgr, InverseHypercube, Ishi Gustaedr, Isilanes, IvanLanin, JLaTondre, JTN, Jdthood, Jeffblank, Jerome Charles Potts, Jjplaya209, Jldugger, Jmorgan, Joegreen, JonHarder, Jordan Brown, Julesd, Karlosian, Ke4roh, Kernel.package, Kl4m-AWB, Knguyeniii, Kowh, Koyaanis Qatsi, Kwamikagami, Liorma, Lovemachine, LucienBoland, MITalum, Makaristos, Marevoula, Mariehuynh, Matt Crypto, Meetabu, Mezzaluna, Michael Hardy, Mikemaccana, Mindmatrix, Miremare, Mirror Vax, Mitsuhirato, Modify, MrBlueSky, Mrwojo, Mwtoews, Nageh, Nikai, Noloader, Nono64, Ntrval, OS2Warp, Omicronpersei8, OrangeDog, P40tomahawk, Pedant17, Petko, Petrb, Peyna, Pgan002, PhilipMW, Piet Delport, Piperx, Pmsyyz, Pne, Pnm, Polacopr0, PseudoSudo, RDBrown, RPage, RanchoRosco, Rebroad, RedWolf, Redhanker, Reelrt, Reuteras, Rfc1394, Richi, Rick.G, Romal, Ronark, Rsukhija, Sajjen, SamHartman, Saqib, Schultz.Ryan, Scotto-san, Seeker68, Sergey Sergeevich Bondarenko, Shepard, Shreyasjoshis, Simon80, Stephan Leeds, StewartNetAddict, Strait, Tedp, The Storm Surfer, Thecheesykid, Tin, Tina scr, Tipiac, Trasz, Treekids, Vectro, Verrai, Vladislav Artukov, Wandie, We64, Werson, West Brom 4ever, Wiglaf, WikHead, Wildzd01, Winryder, Wknight94, Ww, Wwwwolf, Xinconnu, Yaronf, Yoe, Yonkie, Zac439, Zeroshell, Zntrip, 319 anonymous edits

User account policy *Source*: http://en.wikipedia.org/w/index.php?oldid=218609139 *Contributors*: Edward, Frank Kai Fat Chow, Leolaursen, N5iln, Samohyl Jan, 4 anonymous edits

Nagios *Source*: http://en.wikipedia.org/w/index.php?oldid=504487552 *Contributors*: ADobkin, AbelChiaro, Al3ph, AlistairMcMillan, Alpha 4615, Asl, AutomaticStrikeout, Beefyt, Bellenion, Bento00, Bliksim, Bluemoose, Blurpeace, Boinger, Brandon, Cain Mosni, CanisRufus, Chairman S., Ckoniecny, Corpx, Corti, Costas Skarlatos, Craig.trader, Cuddlyable3, D6, Daniel Quinlan, Dawnseeker2000, Deeahbz, Deflective, DeltaQuad, Demi, Dguenault, Dropnagios, Elwell, Emerime, Evert, Fheigl, Filik, Frap, Frederik Holden, Glasnt, Glenn, Graham87, Hm2k, JGXenite, Japo, JayMGoldberg, Jcalcote, JesseW, Jhein, JohnChrysostom, Joy, Jrtayloriv, Julio123, KagamiNoMiko, KingJ, Kirilln, Kl4m-AWB, Klykken, KmiZole, Kop, KuduIO, Kuru, Kwamikagami, Lausser, LimoWreck, Lucious, Lystrata, Maebmij, Mark Renier, Mayz, Mikeshk, MrOllie, MrOllie2, Mro, Mrzaius, Mstarr4, Nachoman-au, Nagios, Nagiosinc, Naparuba, NerdyNSK, NetVicious, Netdiva, NightMonkey, Noisy, Nurg, Orderud, Pearle, Pengo, Pereant antiburchius, Plaes, Pol098, Polluks, Pshankland, R'n'B, Rahgu, Raysonho, RedWolf, Reedy, Reswobslc, Rick Block, Rkarlsba, Robofish, Ronmore, SET, SQB, Samboy, Scenebad, Sietse Snel, Skittels, Sloanr, Socketubs, Srice13, Steavor, Stompersly, T sasahara, TJRC, TMLutas, Tbhotch, Tevkar, Thumperward, Tom Morris, Tookr, Treekids, Two Bananas, Uncopy, Upholder, Vernetto, Vssun, Wereon, Woudenberg, Yaronf, Ysangkok, 255 anonymous edits

Comparison of network monitoring systems *Source*: http://en.wikipedia.org/w/index.php?oldid=503959754 *Contributors*: 3vi1Jay, Accelops, Adamantios, Adamathefrog, Agujero Negro, Akp982, Akshaygs, Al3ph, Annajonna, Anton555a, Assaraf, Astonlam, Bernderk, Bhagatsingh.i, Bkosbab, Bostonvaulter, Bradni1234, BrewskE, Bunnyhop11, Claytondukes, CliffC, Codesmythe, Cvandeplas, DGerman, Dandv, Danhash, Dariorode, Darodewiki, Debresser, Dhyankarthik, Dkhenry, Dmatsunaga, Dominiksf, Donzerus, Doppler2, Drginm, Duffbeerforme, FedericoCeratto, Fibbsanton, Fowartehlluz, Frtoso, Funderburg, G Allegre, GandGTarr, Ghettoblaster, Gmcookie, Gnewf, Goldfndr, Grasmick, Gwmetz, HDReport, Harryboyles, Havoyan, Hfarris1, Hm2k, Hu12, Ikkewel, Innate ideas, Iqlas, JGehlbach, JLaTondre, Jasonfward, Jaw000, Jevansen, Jinxynix, Jkister, Jlmarina, Jmturner, Jonathanmlsw, Joseanes, Jsmith1108, Junruh5, Jwalling, Katicub, Kingpin13, Klerik, Krzysin, Ksalex, Kthompsobass, Kubanczyk, Kuru, Lahnfeear, Lbsbaw, Leewalker, M13l13, MBisanz, MSPCentral, MSPguy, Mabudsxx, ManagementMan, Matach, Matthewhray, Mattskid, Mattyevans, Mayz, Mclynd, Michael93555, Mike.lifeguard, MikeD123, Misslis08, Mmiracle, Monkeygiblet, MonsterSRM, MuffledThud, Naparuba, Natz.nano, Ndownes, NeilN, Netgt, NightMonkey, Normis99, Nsb83, Nww mag, Nyttend, Octotron, OlegD, Pachristensen, Pat01752, Peecee, PigFlu Oink, Plui31, Pozcircuitboy, Pravinspa, Pydubreucq, Qkowlew, RKings, Rayaraddi, Richlv, RobinH99, SDCA91902, SalopianJames, Sammael 42, Scottgwikip, Sd58619, ShakataGaNai, Shpritikin, SiD3WiNDR, Sintrex, Skandha101, Snookmz, Sonic1980, Ssd, StephenBuxton, Steppres, Steverfrancis, Stompersly, Tbockius, Terrillja, Tetryon, The doctor23, TheRealNightRider, Thekaleb, Theredia, Timeport101, Twistedwikibeats, Varnav, Woohookitty, Yhabibzai, 319 anonymous edits

Network File System *Source*: http://en.wikipedia.org/w/index.php?oldid=500484016 *Contributors*: Alexbatko, Andareed, Atomota, Audriusa, Beland, Brianski, Bsroiaadn, Burschik, Byrial, CanisRufus, Chealer, Cokoli, Cplatz, Craphound, Daniel.Cardenas, David Levy, DavidHalko, Dgw, Dittaeva, Dogcow, Dontraub, Edward, Eike Welk, Emperorbma, Etienne.navarro, Fatespeaks, Fnielsen, Fuzheado, GarethEvans, Geek66, Ghakko, Ghettoblaster, Glenn, Gpvos, Gronky, Guy Harris, Haakon, Hairy Dude, Happywaffle, Harryboyles, Hashar, HedgeHog, Helix84, Hella, Hellisp, Herakleitoszefesu, Hhielscher, Hippietrail, Hossein.ir, Ilikeplums, Ilyathemuromets, Intgr, Ivanl, JesseHogan, Jimj wpg, Jordan Brown, Joy, JzG, KJRehberg, Karl-Henner, Khendon, Kirankumar9949, Liam987, Libkeiser, M4gnum0n, Magnus.de, Mange01, Mark Foskey, Mattpickman, Mcld, Mennis, Merlin83b, Mindmatrix, Minesweeper, Mipadi, Mre5765, Mwtoews, NTK, Nfs user, Nixdorf, Nono64, Nubiatech, Nuno Tavares, Ohiostandard, Oli Filth, Omicronpersei8, Orderud, OrgasGirl, Pedant17, Peter S., Peterklevy, PlaysWithLife, Quarsaw, Qwe, Qwertyus, Raanoo, Raysonho, Razimantv, Rees11, Richard W.M. Jones, Robneild, Romx, RuM, S.Örvarr.S, Seaphoto, Secfan, Sfoskett, Shooke, SingingSenator, Smyth, SteinbDJ, Stormie, Stuart Morrow, Tanner-Christopher, Template namespace initialisation script, TenPoundHammer, Thu, Thue, Thumperward, Tony1958, Toussaint, Tyir, UkPaolo, Vedranf, Vegetator, Viajero, Whitelabrat, Wknight94, WojPob, Wrs1864, Ævar Arnfjörð Bjarmason, 129 anonymous edits

Server Message Block *Source*: http://en.wikipedia.org/w/index.php?oldid=504300300 *Contributors*: 4crickj, A bit iffy, Aaron Brenneman, Abune, Aitrus56, Akhristov, Albe, Alex French, AlexPlank, Alexius08, AlistairMcMillan, Allter, Amux, Analoguedragon, Android Mouse, Angusmca, Barek, Bellfoundry, BenBreen2003, Bishopolis, Blaxthos, Bletch, Borgx, Bovineone, Brad.budokan, Bumm13, Car031, Cburnett, Cobone, CosineKitty, Courcelles, Crackrjackr, Ctrlsoft, Cwolfsheep, Darrien, David Gerard, Deflective, Dgranja, Dinjiin, Dkronst, Dogcow, Drange net, ENeville, Eaefremov, Emperorbma, Emurphy42, Etz Haim, Exos2222, FireyFly, Franciosi, Frap, G-force, Gaius Cornelius, Ghakko, Gistca, Gogo Dodo, Grawity, Groink, Guanxi, Guy Harris, Gz33, Haakon, Hairy Dude, Haseo9999, Heron, Hippietrail, Hotdogger, Hqduong, Huiguo01, Iambk, Ian Pitchford, Imeson, Incnis Mrsi, Jamelan, Jannex, Jbattersby, Jd2157, JesseW, Jimmi Hugh, Julesd, Kadin2048, Karada, Kbolino, Kbrose, Kr2006, Kvng, Kwamikagami, Laithjazi, Leandrod, Lkcl, Loftenter, Mabdul, Mange01, Mark.taneo, Marokwitz, Mboverload, Memodude, Mentifisto, MetsFan76, Michael Hardy, Miraceti, Mitch feaster, Mre5765, Mwtoews, Myersj, Nakon, Neilbmartin, Nicolas Melay, Nicolas1981, Nixdorf, NotinREALITY, Novasource, Nubiatech, Numa, Oznt, Paul Weaver, Pauly04, Pearle, Pedant17, Peyre, Piperh, Plutor, Pnm, Project2501a, Public Menace, Qwertyus, Raanoo, Raysonho, Rchandra, Rebrane, Reedy, Reinderien, ReverendDarkness, Richard W.M. Jones, Rjwilmsi, Robmv, Rockear, Romanc19s, RonFredericks, Rprpr, SeL, Sfoskett, Shell Kinney, Sietse Snel, Simon the Likable, Slashme, SmilesALot, Socrates2008, Spellmaster, Sprmw7, Stephan Leeds, SteveLoughran, Stevenrasnick, Stewartadcock, Strait, Superm401, Swiftbrand, Ta bu shi da yu, Thayts, The Anome, Thursbysoftware, Trailrunner13, Trasz, Travis B., UncleBubba, Visualitynq, Warren, WegianWarrior, Winspool, Wk muriithi, Wknight94, Woohookitty, Wrs1864, Xpclient, ماني, 264 anonymous edits

Samba (software) *Source*: http://en.wikipedia.org/w/index.php?oldid=503898440 *Contributors*: 777sms, Ajdlinux, Alamaison, AlistairMcMillan, Am088, Ancheta Wis, Andrewferrier, AnonMoos, ArlindoPereira, Arruah, Austin512, Balderdash707, BenRG, Benjaminevans82, Bevo, Bezapt, Big Elite Feet, Bioprogrammer, Bogdangiusca, Bomazi, Budgee, CWY2190, CanisRufus, Capricorn42, Catapult, Charles Gaudette, CharlesC, Chiaki Minoru, Chpfeiffer, Cintari, Conversion script, Coviti, Ctrlsoft, Cutter, DStoykov, Damian Yerrick, Daniel.Cardenas, Darkonc, Darkride, Darrien, David Gerard, Davidprior, Deeahbz, Deeb, Deineka, Direvus, Dittaeva, Dlrohrer2003, DownloadMASTER, Druckus, E rulez, EagleOne, Ellmist, Elvenlord Elrond, Elwikipedista, Emre D., Espetkov, FCYTravis, Fibonacci, Fligh, Frankie0607, Frap, Gamester17, Gioto, Glenn, Grendelkhan, Gronky, Haikupoet, Halibutt, Hippietrail, Horst F JENS, IMSoP, Iambk, Ianozsvald, Ilion2, Inonit, Isilanes, JECompton, JavaTenor, Jeremy Visser, Jerryobject, JesseHogan, Jleedev, Joconnor, Jonsafari, Jshadias, Julesd, K.lee, K1Bond007, Kalpase, Khym Chanur, Kinema, Kl4m-AWB, Kozuch, KrzysztofKlimonda, Ktdreyer, LaosLos, Lkcl, Locos epraix, LoneFox, Malcolmxl5, Marokwitz, Martarius, Maxim, Melnakeeb, Michael Hardy, MikeMass, MilkMiruku, Millosh, Mindmatrix, Mohearn, Mortense, Mozzerati, Mwtoews, Nandesuka, Neiljt, Nevit, Ngreising, Nikidimi, Nisselua, Nixdorf, NotinREALITY, Numa, Paddles, Philippe.Petrinko, Phongn, PieterDeBruijn, Pitel, Plasticboob, Pne, Ppetru, Prius 2, Quebec99, R-Joe, Rchamberlain, Reedy, Regeseane, Reisio, Renegade54, Rich Farmbrough, Ron Ritzman, Roodscreen, Rune Kock, Rusty, SKopp, Sam916, Samboy, Scanos, Scott A Herbert, Secretlondon, Series8217, Sharcho, SheeEttin, Siafu, SimonD, Snoofer, Snowolf, Socrates2008, Stewartadcock, Stickee, TXiKi, Tedp, Tedwardo2, Telso, The Thing That Should Not Be, Thumperward, Timotheus Canens, Timothylord, Tmassey, Tobybuk, Topdeck, TotoCZ, Turnstep, Tverbeek, Twinxor, Unixguy, Urxn, Verbose, Vindictive Warrior, Vjcyber, Vybr8, Wanted, Weevil, WhalesPsuchsdichs, Where, Whollabilla, Wikijeff, Wknight94, Yaberger, Ykhwong, Zeuscho, روخو, 267 anonymous edits

Distributed file system *Source*: http://en.wikipedia.org/w/index.php?oldid=492024801 *Contributors*: AlistairMcMillan, Beland, Boryna, Cloud101tw, David Levy, Fishnet37222, History2007, Intgr, Isilanes, JerkerNyberg, Katieh5584, Kvng, Mbloore, Pok148, Radugramescu, Raysonho, Ringbang, Robofish, Storio, TheAMmollusc, Welsh, Wknight94, 20 anonymous edits

Simple Mail Transfer Protocol *Source*: http://en.wikipedia.org/w/index.php?oldid=504393718 *Contributors*: 164.58.10.xxx, 1exec1, 21655, Aapo Laitinen, Abdull, Adrian.benko, Aelfgifu, AhmedHan, Aitias, Alansohn, Aldie, Ale2006, AlistairMcMillan, Amatulic, AnK, Andareed, Andrew.cudzilo, Ans, Anwar saadat, Archwyrm, Arvindn, Astrps, Augurr, Babbage, BabsiSnoecks, Bakulev, Barefootguru, Beland, Benefros, Beno1000, Bigdumbdinosaur, Bluerasberry, Bonadea, Breno, Brianhf, Brownb2, Bshurtz, Bunnyhop11, C777, Cafeduke, CecilWard, Cedders, Cedric dlb, Chbarts, Chris the speller, Chrisahn, Christian List, ChristopherTStone, Cinedictum, Coneslayer, Conversion script, Corti, Cryptic C62, Cwolfsheep, DStoykov, Dave-ros, Dawnseeker2000, DerHexer, Devrishighosh, Diego Grez, DirkTheDaring, Dirkbb, EQuail, Ebradsha, Edgarde, ElBenevolente, Elmindreda, Elvey, Enjoi4586, Epbr123, Equendil, Erdal Ronahi, ErikWarmelink, Excirial, FIL123456789, Felixkasza, Fences and windows, Fleminra, Fredhoysted, Fredrik, Free Bear, Frnknstn, GBL, Gargaj, Ghettoblaster, Gianfranco, GoingBatty, Graylorde, Grin, Gscshoyru, Hac13, Hai Dang Quang, Hairy Dude, Hariva, HenryLi, Hirzel, Htaccess, Htonl, Hymek, IRbaboon, Ianeiloart, Incnis Mrsi, Intgr, Ivolucien, JTN, JVz, Jake Wartenberg, Jasper Deng, Jhawk2k, Jim1138, Jjatria, Jleedev, Jmbox80, JohnJSmith, Johnuniq, Jonik, Joy, Julie Deanna, Kbrose, Kelly Martin, Kentucho, Kerberos976, Kinema, Krellis, Kuru, L Kensington, Lalala2005, Lapsus Linguae, Laudaka, Ldorfman, Leif, LeoNomis, Liangent, Lotje, MER-C, MKoltnow, MRLC94, Mabdul, Macar, Macquigg, Mange01, Marcus Qwertyus, Mathiastck, Maxí, Mdmkolbe, Mentifisto, Message From Xenu, Michael Leeman, Milliped, Milomedes, Mindmatrix, Mordomo, Ms2ger, Mwtoews, NTox, Nabber00, Nacarlson, Nanshu, Napalm Llama, Nbarth, Neillucas, Nepenthes, Nixdorf, Nrawat, Nubiatech, Ohnoitsjamie, Omegatron, Omniplex, Osias, Pagingmrherman, Parasitical, Phatom87, Planetary Chaos Redux, Plugwash, Pokemonblackds, Pol098, Pratyeka, Quadro, Rajeev1486, Reach Out to the Truth, Recognizance, Rich Farmbrough, Richwales, Rick Block, Rory O'Kane, Rwxrwxrwx, SJFriedl, Samitalways4u, Samuel, Sceptre, Sciplus, Scurless, Sdfisher, Sdrtirs, Seffer, Shep68, Shlomif, Simsvc, Sjmsteffann, Southpaw018, Speight, Squire55, Sshoe, Staffwaterboy, StaticGull, Stephan Leeds, Stormie, Suruena, TGDA, Tedp, Template namespace initialisation script, The Anome, TheKMan, Theopolisme, Thumperward, TkGy, Trusilver, Twocs, UBJ 43X, UncleBubba, Uruiamme, Usien6, Varantes, Versageek, VictorianMutant, Vlhsrp, Washburnmav, Wilinckx, Will Beback Auto, WintersChild, Woohookitty, Wrs1864, Xbt, Xoder, Yahya Abdal-Aziz, Yaronf, Yiannisl, Zac67, ZeroOne, Zondor, Zophelia, Тивернополник, 429 anonymous edits

Internet Message Access Protocol *Source*: http://en.wikipedia.org/w/index.php?oldid=504834559 *Contributors*: *drew, 13 of Diamonds, 217.126.156.xxx, ASk, Adam Marsh, Addshore, Ajh16, Alaniaris, Ale2006, AlexJTaylor, Alexcat, AlistairMcMillan, Allen Moore, Anclation, Android Mouse, Anirvan, Anton Mravcek, Apollo2011, Areteichi, Arjun024, Armando, Avisp, AzaToth, BabsiSnoecks, Barefootguru, Bdesham, Beetstra, Beno1000, Billgordon1099, BobJohannessen, Brandonsturgeon, Caiaffa, Chadlupkes, Charlesrh, Chealer, Chrisahn, Christian List, Cmh294, Conversion script, Cosmoincarlow, Cwolfsheep, Cybercobra, DaleKiefling, Danejasper, Danitaz, Defrenrokorit, Deli nk, Doodledoo, Drable, Draney13, Duplicity, EagleOne, El Cubano, Eliashc, Enjoi4586, Equendil, Eric-Wester, Erpingham, Ersalan, Espetkov, Estephan500, Excirial, Falcon Kirtaran, Faradayplank, FisherQueen, Fleminra, Frap, Fuzzygenius, Gardar Rurak, Gary King, Giftlite, Gkanch, Glenn Anderson, Golem, Grafen, Guinness, Gwydion1, Gzbr, Hadrianheugh, Hairy Dude, HalifaxRage, Haseo9999, Hgd, Hhielscher, Hi878, Honta, Hyad, Imz, Intrigued-user, Ivan Pozdeev, J.delanoy, JTN, Jlang, Joanjoc, Joconnor, Johayek, Johnjosephbachir, JonHarder, Jonabbey, Jonik, Joseph Spiros, Julesd, Juliancolton, Karnesky, Kaszeta, Kbrose, Kedadi, Keegan, KneeLess, Koweja, Koyaanis Qatsi, Krellis, Ktinga, L.Petrlik, Lathspell, LeaveSleaves, LiDaobing, Liftarn, LobStoR, Mabdul, Mairi, Martin Kealey, Mboverload, Mcsee, Mdcougar, Mloughran, Moloch09, Mouse Nightshirt, Mr.Z-man, Mwtoews, Nagy, Natkeeran, NaturalBornKiller, Ned Scott, Nemec, NerdyNSK, Nestea Zen, Nixdorf, Nomeata, Nubiatech, Oceans and oceans, Oliphaunt, Omniplex, Omphaloscope, P00r, PeconicSky, Peyre, Pgk, Pgr94, Pitr, Pmsyyz, Pol098, Popnose, R27182818, RG, Red Slash, RedWolf, Rich Farmbrough, Rick Block, Rjwilmsi, Ron Ritzman, Ronark, Rsforward, RussNelson, Rwxrwxrwx, Ryochiji, SJP, Sakurina, Sciencewatcher, Sdfisher, Sdrtirs, SidP, Sinan, Snori, Speedoflight, Squadri, StevenMcCoy, Suruena, SusanLesch, THEN WHO WAS PHONE?, Thoric, Thumperward, Tim Ivorson, TimoSirainen, Timosa, Tktravis, TravisHein, Trou, Ts4z, Tyler9xp, UncleBubba, Uthbrian, V.s.mousavi, Voidxor, WJetChao, Walling, Wcoole, Wik, WintersChild, Wrs1864, Wsheward, Yaronf, Yosri, Z4ns4tsu, Zac67, ZeroOne, Александр Цамутали, Петър Петров, 387 anonymous edits

Electronic mailing list *Source*: http://en.wikipedia.org/w/index.php?oldid=503827480 *Contributors*: 16@r, Adam D. Audette, Adam audette, Ale2006, Alireza-Hassani, Alvestrand, Angoca, Arjuno3, Atlant, Atulsnischal, AvicAWB, Barbara937, Barefootguru, Bean2thousand, Beevvy, Benfitts, BenoyGeorge, Bizexpatiate, Bonadea, Brandorr, Brooks27, Brossow, Buffs, Businesslists, Cab88, Caknuck, CalumH93, Can't sleep, clown will eat me, Carados, Cchcap, Chealer, Chmod007, Ciampix, Codyrank, Commander, Cybersquid, D.Shaw 78, DMacks, Danrandow, Danwestern, Db173, Dcljr, Deed89, Deli nk, Delicious carbuncle, Drahreg01, Dxcuser, Eloquence, Ency, Epbr123, ErikWarmelink, Etz Haim, FF2010, Fanyavizuri, Finlay McWalter, Frecklefoot, From-cary, Gardar Rurak, GraemeL, Haakon, Hairy Dude, Harryzilber, Hildanknight, Hkhurana2000, Hydrargyrum, IMSoP, Icedog, Inarius, Ipridian, Iridescent, Ishmaelville, Jeroen, Jerome Charles Potts, Jim Henry, John Hubbard, Jokes Free4Me, JonHarder, Joy, Jpyeron, KarlDubost, Kbrose, Latchkey, Leandrod, Leif, Lexor, Lhoriman, Liao, Longhair, Lotje, Mahjongg, Mairi, Marcus Qwertyus, Marenco, Mark91, Mav, Mdirolf, Michael Hardy, Michaelbusch, Michaelprather, Mikeloi, Mintleaf, Mjb, Neilc, Neilpw1, Nfds, Nickpenamiami, Noam55, Ohnoitsjamie, Omegatron, Open2universe, Oral24, Orzelf, Oxymoron83, Pak21, Paul Salber, Peter, Pleather, Pmc, Pnm, Pwforaker, Quasar2112, Querencia, Radagast83, Rajesh.lalwani, RedWolf, Reindra, Renrenren, Rhobite, Rick Norwood, Rillian, RockMFR, RockyRoad29, Rom1sub, Seraphim, Sheskar, Simon Grabowski, Skagedal, SlubGlub, Soph, SpaceFlight89, Srtxg, Stephen Gilbert, Stev0, Stvfetterly, Swoovy420, Teri Ives, The Transhumanist, The undertow, Thv, Vevo90, VincentV, Vorpal Suds, Vsmith, Wangi, Wickethewok, WikHead, Wizardist, Wourriezer, Wrs1864, XDanielx, Xelgen, Ykhwong, Zoe, Zondor, Zoney, Zzuuzz, 130 anonymous edits

Comparison of mail servers *Source*: http://en.wikipedia.org/w/index.php?oldid=503812853 *Contributors*: 2001:470:E9B4:0:9D38:261A:CEFE:9079, Agateller, Akrus, Alchaemist, Ale2006, Alien life form, Anabolism, Anastrophe, Arfon, Arokha, Athaba, Atmail, Ausxan, Barbacana, Belenus, Bielle, ChristTrekker, Coaxial, Colonies Chris, Cpugmire, Crimoid, Cristina Zainea, D V, DanielRigal, Dk white, Dozhd, Draney13, Drbreznjev, Drcarver, Duckbill, EagleOne, Enfcyco, Eth000, FFTravisM, Fajo66, Ftiercel, Futurepower(R), GBuilder, Gashok, Geekdog, GoingBatty, Golem, HLachman, HangingCurve, Hexamail, Ianeiloart, Icewarp, Icewarp ltd, Imailds, Imdonatello, Ivolucien, Jallison, Jhartmann, Jhawk2k, Jigs0027, Jjk, Jlaidman, Jonabbey, Jorgen919, KitchM, Kjwu, Laurusnobilis, Ldsandon, Lennartgoosens, LukeShu, Mailenable, Mailsite, Marcostas, Massimobolzoni, MaxMVP, Mbhangui, MeaninglessName, MezzyMez, Mikeblas, Milo-oostergo, Mmurphy2k, Mr.Z-man, Oscarthecat, PanderMusubi, R'n'B, Reboot, Reisio, Requiem18th, Reukiodo, Rmdashrf, Rossnixon, Rum132, Rwxrwxrwx, Sideris, Smiteri, Syscomet, Themask123, TimoSirainen, Txt.file, Voidxor, WolfWings, Zac67, Ze.hfox, Zicko1, Важнов Алексей Геннадьевич, 178 anonymous edits

Microsoft Exchange Server *Source*: http://en.wikipedia.org/w/index.php?oldid=504080547 *Contributors*: 95j, A876, Aapo Laitinen, Aardila, Abelson, Abrahami, Alanfeld, Alex Libman, Alexf, AlistairMcMillan, Amplitude101, Andy1979s, Annvu, Apyule, Backdoor2world, Beland, Bender235, BfMGH, Big Bird, Bobo192, Bonadea, Bsoo, Bsuneja, Btornado, Bunnyhop11, Bwmcq, Bwrs, Canaima, Carbonite, Cflee, Chilly Penguin, Chrissawer, Chunt006, Cleared as filed, CommonsDelinker, Craigcarry1, DStoykov, Danieleran, Darklilac, Dekaptein, Delfassy, Deli nk, Diagonalfish, Diamondland, Digemedi, Disdero, Dmarquard, DocWatson42, Dogears, Dperegudov, Draicone, EEPROM Eagle, EJDyksen, EagleOne, Eastmain, Edemaine, Edomenig, Edward, EhJJ, Euchiasmus, Fdiv bug, Fragglet, Fredrik, Frozen Wind, Gardar Rurak, Gary King, Ghettoblaster, Gilesmorant, Gogo Dodo, Gohan71, Guoguo12, Haakon, Heracles31, Hmrox, Hutcher, I dream of horses, Iliank, Indefatigable, Infostellar, Ixfd64, IzKen, J.delanoy, JLaTondre, Jasper Deng, Jeremy Visser, Jesster79, Jhwinder, Jim.Liu, Jmcw37, Jmundo, Joel8r, JohnExchange, Jonathonblocker, Jonbryce, Joy, Keandobus, Kennedye, Keramos, Khalid hassani, Khazar2, Kkm010, Kozuch, Kuru, Kypr8, L337h1um, LMB, Lakorai, Lenin1234, Leuqarte, LionKimbro, Lofote, Logan, Lord Eru, Lwc, MER-C, Maginks, MainFrame, Majid ahmadi, Marcika, Mardus, MaxMVP, Mdlueck, Meeples, Meirpolaris, Melnakeeb, Mendel129, Mikeblas, Mindmatrix, Mishra78, Modster, Mondebleu, Mtlk, NSR, Natalie Erin, Nate Silva, Netalarm, NevilleDNZ, Newtronic, Nigelj, Nixdorf, Njan, Nort, Northgrove, OceanO, Od Mishehu, Odie5533, Ohura, Omarmalali, OnePt618, Orbital, PPGMD, Pah861, Panarchy, Panyd, Pascal666, Pelekan, Perhelion, Peterl, Picassob, Pmsyyz, Pol098, Possum, Pyeman73, R.O.C, Raitchison, Rebroad, Redrose64, Reedy, Retired user 0001, Richard Slater, Richi, Riki, Rjwilmsi, RobertL, Royan, Sambellajoe, SaxicolousOne, Scott Paeth, Senator Palpatine, Shouran, Shshme, Simontbates, Smiteri, Soumyasch, Sourceholder, Stephan Leeds, Stevenbdjr, Sunnybanthiya, SuperHamster, SuperTycoon, Superiaklubbene, SusanLarson, Svenster, Syp, Ta bu shi da yu, Tbhotch, Tdada, TechSoEasy, Tedickey, Teguhth, Th.held, TheCuriousGnome, Thewyzewun, TimBray, Tmcw, Trusilver, Tumbleweed05, Twojenkins, U, Valodzka, Varlaam, Veinor, Vertium, Vin1317, Vinh1000, Vlad, Vryl, Vt100, Waggers, Warren, Wdavidpence, West.andrew.g, Widefox, WikHead, Wikidear, Wikiolap, WinCEB, XantiSS, Xp fun, Xpclient, Yuhong, Zachary, Zack Howes, Zbnet, Zigger, Δ, 笨鸟不想飞, 590 anonymous edits

Image Sources, Licenses and Contributors

File:Wikimedia OS share pie chart.png *Source*: http://en.wikipedia.org/w/index.php?title=File:Wikimedia_OS_share_pie_chart.png *License*: Creative Commons Attribution-Sharealike 3.0 *Contributors*: User:I384403

File:Rack001.jpg *Source*: http://en.wikipedia.org/w/index.php?title=File:Rack001.jpg *License*: Creative Commons Attribution-Sharealike 3.0,2.5,2.0,1.0 *Contributors*: Jfreyre

File:IBM System z10.jpg *Source*: http://en.wikipedia.org/w/index.php?title=File:IBM_System_z10.jpg *License*: Creative Commons Attribution-Sharealike 2.0 *Contributors*: Roberto Berlim from Brazil

Image:Operating systems used on top 500 supercomputers.svg *Source*: http://en.wikipedia.org/w/index.php?title=File:Operating_systems_used_on_top_500_supercomputers.svg *License*: Public Domain *Contributors*: Benedikt.Seidl

Image:Gldt.svg *Source*: http://en.wikipedia.org/w/index.php?title=File:Gldt.svg *License*: unknown *Contributors*: Andreas Lundqvist

File:Linux-Distribution en.svg *Source*: http://en.wikipedia.org/w/index.php?title=File:Linux-Distribution_en.svg *License*: Creative Commons Attribution-ShareAlike 3.0 Unported *Contributors*: Linux-Distribution.svg: Gissi, Phrood, Erik Streb derivative work: Mewtu (talk)

Image:Debian.png *Source*: http://en.wikipedia.org/w/index.php?title=File:Debian.png *License*: GNU General Public License *Contributors*: Debian Project

Image:Fedora 16 default Desktop.png *Source*: http://en.wikipedia.org/w/index.php?title=File:Fedora_16_default_Desktop.png *License*: unknown *Contributors*: Elisonniven, Martin H., 1 anonymous edits

Image:GentooTen.png *Source*: http://en.wikipedia.org/w/index.php?title=File:GentooTen.png *License*: unknown *Contributors*: Gentoo 10.0 LiveDVD Team

Image:Mandriva2010eng.png *Source*: http://en.wikipedia.org/w/index.php?title=File:Mandriva2010eng.png *License*: GNU General Public License *Contributors*: Mandriva

File:OpenSUSE_12.1_KDE_desktop.png *Source*: http://en.wikipedia.org/w/index.php?title=File:OpenSUSE_12.1_KDE_desktop.png *License*: Creative Commons Attribution-Sharealike 3.0 *Contributors*: Wiedemann

Image:Slackware.png *Source*: http://en.wikipedia.org/w/index.php?title=File:Slackware.png *License*: GNU General Public License *Contributors*: Uploader

Image:Ubuntu_11.10.png *Source*: http://en.wikipedia.org/w/index.php?title=File:Ubuntu_11.10.png *License*: GNU General Public License *Contributors*: Funkruf / Ubuntu Foundation

Image:PuppyLinux525.png *Source*: http://en.wikipedia.org/w/index.php?title=File:PuppyLinux525.png *License*: GNU General Public License *Contributors*: Mick Amadio, chief developer, Larry Short, coordinator et al

Image:SLAX.png *Source*: http://en.wikipedia.org/w/index.php?title=File:SLAX.png *License*: GNU General Public License *Contributors*: http://hacktolive.org/

Image:Linux Mint Lisa.png *Source*: http://en.wikipedia.org/w/index.php?title=File:Linux_Mint_Lisa.png *License*: GNU General Public License *Contributors*: Linux Mint

Image:Pclos20092.jpg *Source*: http://en.wikipedia.org/w/index.php?title=File:Pclos20092.jpg *License*: GNU General Public License *Contributors*: PCLinuxOS Developer Team

File:Debian-OpenLogo.svg *Source*: http://en.wikipedia.org/w/index.php?title=File:Debian-OpenLogo.svg *License*: Public Domain *Contributors*: Debian

File:Debian.png *Source*: http://en.wikipedia.org/w/index.php?title=File:Debian.png *License*: GNU General Public License *Contributors*: Debian Project

File:Loudspeaker.svg *Source*: http://en.wikipedia.org/w/index.php?title=File:Loudspeaker.svg *License*: Public Domain *Contributors*: Bayo, Gmaxwell, Husky, Iamunknown, Mirithing, Myself488, Nethac DIU, Omegatron, Rocket000, The Evil IP address, Wouterhagens, 19 anonymous edits

File:Debian-aptitude.png *Source*: http://en.wikipedia.org/w/index.php?title=File:Debian-aptitude.png *License*: GNU General Public License *Contributors*: VShaka (talk) 10:00, 18 February 2009 (UTC). Original uploader was VShaka at en.wikipedia

File:Debian-organigram.svg *Source*: http://en.wikipedia.org/w/index.php?title=File:Debian-organigram.svg *License*: Creative Commons Attribution-Sharealike 3.0 Germany *Contributors*: Martin F. Krafft

File:Debian-package-cycl.svg *Source*: http://en.wikipedia.org/w/index.php?title=File:Debian-package-cycl.svg *License*: Creative Commons Attribution 2.5 *Contributors*: The image is copyright © Martin F. Krafft, based on the work by Kevin Mark. It is released under the terms of the CC by-sa 2.5 licence

File:Debian Installer graphical etch.png *Source*: http://en.wikipedia.org/w/index.php?title=File:Debian_Installer_graphical_etch.png *License*: GNU General Public License *Contributors*: User:IngerAlHaosului

File:Debian-cd-cover1.png *Source*: http://en.wikipedia.org/w/index.php?title=File:Debian-cd-cover1.png *License*: GNU General Public License *Contributors*: Tobias Wagner

image:Debian ports timeline.png *Source*: http://en.wikipedia.org/w/index.php?title=File:Debian_ports_timeline.png *License*: Creative Commons Attribution-Sharealike 3.0 *Contributors*: Scientus

File:Debian family tree 11-06.png *Source*: http://en.wikipedia.org/w/index.php?title=File:Debian_family_tree_11-06.png *License*: GNU Free Documentation License *Contributors*: 4th-otaku, Dratini0

File:Logo-ubuntu no(r)-black orange-hex.svg *Source*: http://en.wikipedia.org/w/index.php?title=File:Logo-ubuntu_no(r)-black_orange-hex.svg *License*: Public Domain *Contributors*: Colejohnson66, ColinCameron, Cyan.aqua

File:Ubuntu 12.04 Final Live CD Screenshot.png *Source*: http://en.wikipedia.org/w/index.php?title=File:Ubuntu_12.04_Final_Live_CD_Screenshot.png *License*: GNU General Public License *Contributors*: Ahunt, 1 anonymous edits

File:Nexus S with Ubuntu.jpg *Source*: http://en.wikipedia.org/w/index.php?title=File:Nexus_S_with_Ubuntu.jpg *License*: Creative Commons Attribution 2.0 *Contributors*: bakaotaku

File:Ubuntu family tree 11-06.png *Source*: http://en.wikipedia.org/w/index.php?title=File:Ubuntu_family_tree_11-06.png *License*: GNU Free Documentation License *Contributors*: 4th-otaku, Rudloff (usurped)

File:Ubuntu TV screenshot.png *Source*: http://en.wikipedia.org/w/index.php?title=File:Ubuntu_TV_screenshot.png *License*: unknown *Contributors*: Canonical Ltd.

File:Kubuntu_12_04.png *Source*: http://en.wikipedia.org/w/index.php?title=File:Kubuntu_12_04.png *License*: GNU General Public License *Contributors*: Funkruf / Ubuntu Fondation

File:Ubuntu 12-04 Server Installation.png *Source*: http://en.wikipedia.org/w/index.php?title=File:Ubuntu_12-04_Server_Installation.png *License*: unknown *Contributors*: Ubuntu Developers

File:Eucalyptus full.jpg *Source*: http://en.wikipedia.org/w/index.php?title=File:Eucalyptus_full.jpg *License*: Copyrighted free use *Contributors*: Yegulalp Serdar

File:20111102-145438-UDSPrecise-XL.jpg *Source*: http://en.wikipedia.org/w/index.php?title=File:20111102-145438-UDSPrecise-XL.jpg *License*: Creative Commons Attribution-Sharealike 3.0 *Contributors*: Exeva, LFaraone, Michael Barera

File:Windows Family Tree.svg *Source*: http://en.wikipedia.org/w/index.php?title=File:Windows_Family_Tree.svg *License*: Creative Commons Attribution-ShareAlike 3.0 Unported *Contributors*: J.int, Linfocito B, NOKIA 3120 classic, 1 anonymous edits

File:Viridian Architecture.svg *Source*: http://en.wikipedia.org/w/index.php?title=File:Viridian_Architecture.svg *License*: GNU Free Documentation License *Contributors*: Soumyasch

File:Linux command-line. Bash. GNOME Terminal. screenshot.png *Source*: http://en.wikipedia.org/w/index.php?title=File:Linux_command-line._Bash._GNOME_Terminal._screenshot.png *License*: unknown *Contributors*: User:ZxxZxxZ

File:bash screenshot.png *Source*: http://en.wikipedia.org/w/index.php?title=File:Bash_screenshot.png *License*: unknown *Contributors*: Akinom, Andrew pmk, Berland, Emx, Insuranze, Kahlil88, Leileilol, MarSoft, Pixel ;-), Shooke, Túrelio, 5 anonymous edits

File:Windows PowerShell 1.0 PD.png *Source*: http://en.wikipedia.org/w/index.php?title=File:Windows_PowerShell_1.0_PD.png *License*: Public Domain *Contributors*: Microsoft Corporation

Image:Bash demo.png *Source*: http://en.wikipedia.org/w/index.php?title=File:Bash_demo.png *License*: unknown *Contributors*: Chealer

File:Windows PowerShell icon.png *Source*: http://en.wikipedia.org/w/index.php?title=File:Windows_PowerShell_icon.png *License*: unknown *Contributors*: Ankit Maity, Shinjiman, Skier Dude, Skybon

Image:Gnu grub config file.png *Source*: http://en.wikipedia.org/w/index.php?title=File:Gnu_grub_config_file.png *License*: GNU General Public License *Contributors*: Original uploader was Icey at en.wikipedia

Image:Synaptic Package Manager.png *Source*: http://en.wikipedia.org/w/index.php?title=File:Synaptic_Package_Manager.png *License*: GNU General Public License *Contributors*: http://hacktolive.org/

Image:Pms.svg *Source*: http://en.wikipedia.org/w/index.php?title=File:Pms.svg *License*: Public Domain *Contributors*: Traced by User:Stannered

Image:Apt-get_install_mediawiki.png *Source*: http://en.wikipedia.org/w/index.php?title=File:Apt-get_install_mediawiki.png *License*: GNU General Public License *Contributors*: Debian project (Konsole terminal emulator from the KDE project)

File:Win Update ME.png *Source*: http://en.wikipedia.org/w/index.php?title=File:Win_Update_ME.png *License*: unknown *Contributors*: User:SufferWell1396

File:DHCP Server.png *Source*: http://en.wikipedia.org/w/index.php?title=File:DHCP_Server.png *License*: Creative Commons Attribution-Sharealike 3.0 *Contributors*: User:Patpat

File:Tftp-wrq.svg *Source*: http://en.wikipedia.org/w/index.php?title=File:Tftp-wrq.svg *License*: Public Domain *Contributors*: Ozga

File:Tftp-ack0.svg *Source*: http://en.wikipedia.org/w/index.php?title=File:Tftp-ack0.svg *License*: Public Domain *Contributors*: Ozga

File:Tftp-dat1-up.svg *Source*: http://en.wikipedia.org/w/index.php?title=File:Tftp-dat1-up.svg *License*: Public Domain *Contributors*: Ozga

File:Tftp-rrq.svg *Source*: http://en.wikipedia.org/w/index.php?title=File:Tftp-rrq.svg *License*: Public Domain *Contributors*: Ozga, 1 anonymous edits

File:Tftp-dat1-dwn.svg *Source*: http://en.wikipedia.org/w/index.php?title=File:Tftp-dat1-dwn.svg *License*: Public Domain *Contributors*: Ozga

File:Tftp-ack1.svg *Source*: http://en.wikipedia.org/w/index.php?title=File:Tftp-ack1.svg *License*: Public Domain *Contributors*: Ozga

Image:Network Time Protocol servers and clients.svg *Source*: http://en.wikipedia.org/w/index.php?title=File:Network_Time_Protocol_servers_and_clients.svg *License*: unknown *Contributors*: User:Bdesham

Image:Usno-amc.jpg *Source*: http://en.wikipedia.org/w/index.php?title=File:Usno-amc.jpg *License*: Public Domain *Contributors*: US Air Force

Image:Domain name space.svg *Source*: http://en.wikipedia.org/w/index.php?title=File:Domain_name_space.svg *License*: Public Domain *Contributors*: Bjankuloski06en, GeorgHH, Ggia, Herbythyme, LionKimbro, Ma-Lik, Mdd, Waldir, 12 anonymous edits

Image:An example of theoretical DNS recursion.svg *Source*: http://en.wikipedia.org/w/index.php?title=File:An_example_of_theoretical_DNS_recursion.svg *License*: Public Domain *Contributors*: LionKimbro, Syp, The Anome, Waldir, 4 anonymous edits

Image:DNS in the real world.svg *Source*: http://en.wikipedia.org/w/index.php?title=File:DNS_in_the_real_world.svg *License*: Public Domain *Contributors*: LionKimbro, Syp, Waldir, 2 anonymous edits

Image:ams-ix.k.root-servers.net.jpg *Source*: http://en.wikipedia.org/w/index.php?title=File:Ams-ix.k.root-servers.net.jpg *License*: Creative Commons Attribution-Sharealike 3.0 *Contributors*: D0ktorz, Ilse@, Kgyt, R. Koot, Yonatanh

Image:Root-current.svg *Source*: http://en.wikipedia.org/w/index.php?title=File:Root-current.svg *License*: Creative Commons Attribution-ShareAlike 3.0 Unported *Contributors*: Matthäus Wander, Waldir, 1 anonymous edits

File:A help portal.gif *Source*: http://en.wikipedia.org/w/index.php?title=File:A_help_portal.gif *License*: Copyrighted free use *Contributors*: Zureks

File:Icons-mini-page tree.gif *Source*: http://en.wikipedia.org/w/index.php?title=File:Icons-mini-page_tree.gif *License*: Copyrighted free use *Contributors*: Avatar, Bayo, Rocket000

File:Icons-mini-page url.gif *Source*: http://en.wikipedia.org/w/index.php?title=File:Icons-mini-page_url.gif *License*: Copyrighted free use *Contributors*: Avatar

File:Icons-mini-folder.gif *Source*: http://en.wikipedia.org/w/index.php?title=File:Icons-mini-folder.gif *License*: Copyrighted free use *Contributors*: Avatar

File:Icons-mini-icon user.gif *Source*: http://en.wikipedia.org/w/index.php?title=File:Icons-mini-icon_user.gif *License*: Copyrighted free use *Contributors*: Avatar, Rocket000

File:Cloud computing.svg *Source*: http://en.wikipedia.org/w/index.php?title=File:Cloud_computing.svg *License*: Creative Commons Attribution-Sharealike 3.0 *Contributors*: Sam Johnston

File:Cloud computing layers.png *Source*: http://en.wikipedia.org/w/index.php?title=File:Cloud_computing_layers.png *License*: unknown *Contributors*: User:CommonsHelper2 Bot, User:Jan Luca, User:Magnus Manske

File:Cloud computing types.svg *Source*: http://en.wikipedia.org/w/index.php?title=File:Cloud_computing_types.svg *License*: Creative Commons Attribution-Sharealike 3.0 *Contributors*: Sam Joton

File:CloudComputingSampleArchitecture.svg *Source*: http://en.wikipedia.org/w/index.php?title=File:CloudComputingSampleArchitecture.svg *License*: Creative Commons Attribution-Share Alike *Contributors*: Sam Johnston, Australian Online Solutions Pty Ltd

File:X11 ssh tunnelling.png *Source*: http://en.wikipedia.org/w/index.php?title=File:X11_ssh_tunnelling.png *License*: Creative Commons Attribution-ShareAlike 3.0 Unported *Contributors*: Grön, Jleedev, Suwa, Sven, 2 anonymous edits

File:OpenWrtPuTTY.png *Source*: http://en.wikipedia.org/w/index.php?title=File:OpenWrtPuTTY.png *License*: unknown *Contributors*: Casablanca

File:Ssh binary packet alt.svg *Source*: http://en.wikipedia.org/w/index.php?title=File:Ssh_binary_packet_alt.svg *License*: Creative Commons Attribution 2.5 *Contributors*: Traced by User:Stannered, original by en:User:Verbose using RFC4253 as source.

File:Icons-mini-page text.gif *Source*: http://en.wikipedia.org/w/index.php?title=File:Icons-mini-page_text.gif *License*: Copyrighted free use *Contributors*: Avatar

File:Kerberos.svg *Source*: http://en.wikipedia.org/w/index.php?title=File:Kerberos.svg *License*: Creative Commons Attribution-Sharealike 3.0 *Contributors*: User:Dsonck92

File:Nagios main screen.png *Source*: http://en.wikipedia.org/w/index.php?title=File:Nagios_main_screen.png *License*: GNU General Public License *Contributors*: Nagios Enterprises, LLC

File:Samba_Logo.png *Source*: http://en.wikipedia.org/w/index.php?title=File:Samba_Logo.png *License*: Public Domain *Contributors*: User:SreeBot

Image:SMTP-transfer-model.svg *Source*: http://en.wikipedia.org/w/index.php?title=File:SMTP-transfer-model.svg *License*: GNU Free Documentation License *Contributors*: Ale2006-from-en

File:Exchange-2010logo.svg *Source*: http://en.wikipedia.org/w/index.php?title=File:Exchange-2010logo.svg *License*: Public Domain *Contributors*: Microsoft Corporation Derivative work: Perhelion (talk)

License