

Total Cost of Ownership for Linux Web Servers in the Enterprise

September 2002

Table of Contents

1. EXECUTIVE SUMMARY	2
2. METHODOLOGY	2
3. COST BREAKDOWN	3
3.1. Software Purchase	3
3.2. Hardware Purchase and Maintenance	4
3.3. System Support and Administration	5
4. "SOFT" COSTS	6
4.1. Security	6
4.2. Availability	7
4.3. Scalability	7
4.4. Other Factors	7
5. CONCLUSION	8



22 Crescent Road
Westport, CT 06880
<http://www.rfgonline.com>

1. Executive Summary

Through numerous conversations with IT executives, RFG learned that Linux server deployments are well under way, in many cases with notable cost savings from those implementations. To further explore and quantify this anecdotal evidence, RFG performed a total cost of ownership (TCO) study to evaluate Linux deployments in the enterprise. Survey participants included mid- to large-sized companies, and of the companies polled, 14 yielded relevant data sets that were included in this study.

RFG found that Linux was the least expensive platform to deploy and operate. Although some initial costs were higher at points, the ability to massively scale the product horizontally without paying additional licensing fees can yield significant cost savings over the long term.

Case	Year 1	Year 2	Year 3
Linux	\$49,931	\$62,203	\$74,475
Solaris	\$421,718	\$491,619	\$561,520
Windows	\$91,724	\$141,193	\$190,662

2. Methodology

RFG surveyed IT executives from Global 2000 companies both inside and out of its client base and across a range of vertical markets to collect data on their Web server deployments for three target platforms – Linux, Sun's Solaris, and Microsoft's Windows. Although Linux can support a variety of functions, serving Web pages is a very visible and measurable function. More importantly, RFG found that companies polled were further along the migration path in terms of Web server deployments than for other applications.

The majority of Linux and Windows deployments were on Intel x86-architecture servers, with Microsoft's Internet Information Server (IIS) providing Web server functionality on Windows and Apache providing the same for Linux. The majority of Solaris deployments were on Sun SPARC-architecture systems, with Apache providing the Web server functionality.

However, the numbers of servers deployed in each case were not equivalent. SPARC-based deployments generally consisted of a few servers, each of which was configured with a large number of processors (vertically scaled). In contrast, x86-based deployments that handled similar workloads were generally deployed with a larger number of systems, each of which was configured with only one or two processors (horizontally scaled).

To make direct comparisons possible, RFG normalized the data collected using the concept of a "Processing Unit." To do this, RFG calculated the number of servers that would be required to process 100,000 hits per day as follows:

$$\text{servers_per_PU} = \frac{\text{servers_in_use} * (100,000)}{\text{daily_hits_for_site}}$$

Two data sets were excluded from the calculation of a Processing Unit, one of which was a green field deployment of Linux well in excess of 10,000 nodes, the other of which exceeded 300 nodes. This was done to make the comparison more fair, since deployments on this scale are unusual for Web server environments in the enterprise. IT executives should note that the greatest benefits of a Linux deployment may be realized by a deployment on this scale, and adjust their calculations accordingly if such an opportunity presents itself in their own departments.

Calculating the average number of servers in each Processing Unit yielded:

Environment	Servers per Processing Unit
Windows on x86-architecture hardware	7.6
Linux on x86-architecture hardware	7.4
Solaris on x86-architecture hardware	2.2

IT executives can use the figures above in conjunction with usage estimates or figures for their own Web server deployments to scale the figures presented in this document to their own environments. However, IT executives should note that the values above illustrate actual hit

counts experienced by the users polled, and should not be viewed as an indication of the maximum usage possible for each server. Further, precise server hardware configurations, software versions, application tuning, operating system tuning, and so forth are crucial determinants in a server's ability to handle a given workload.

Figures in this study are based on the following criteria:

- ? Deployments will be in use for a period of 3 years.
- ? The Web server for Windows systems is IIS, while the Web server for both Linux and Solaris is Apache.
- ? Shipping costs and other fees not specific to either the hardware or software products were excluded from consideration.
- ? "External" support hardware and software products, such as load balancing appliances, firewalls, and intrusion detection systems, were excluded from consideration.
- ? All prices are in U.S. Dollars (USD) and are based on retail pricing. Prices will be rounded to the nearest whole dollar *after* each calculation is performed.
- ? Software support options selected were those included in the purchase of the product.

Retail pricing was used to create a level playing field between all three deployment types. Because enterprise licensing contracts are highly negotiable, IT executives may be able to obtain more competitive pricing in very high volumes. This is true of all vendors considered in this study.

3. Cost Breakdown

This study compared four areas – software purchase prices, hardware purchase and maintenance prices, software maintenance and upgrade prices, and administrative costs. Administrative and hardware costs formed the largest chunk of the platform TCO, which for Linux is not surprising given that the product's low cost is the factor most often quoted by

companies considering deploying Linux in production environments.

3.1. Software Purchase

Software purchase costs per Processing Unit were as follows:

Case	Up Front	Year 2	Year 3	Total
Linux	\$400	\$0	\$0	\$400
Solaris	\$0	\$0	\$0	\$0
Windows	\$5,320	\$1,330	\$1,330	\$7,980

Software purchase costs were straightforward in this study, as the packages compared all ship bundled with Web server software, so other products are not required here.

Linux is available in a variety of distributions, which are packages that include the Linux kernel itself, as well as development tools, system libraries, utilities, the Web server software (Apache), and so forth. Red Hat was the preference among study participants by a slim margin. Most were experimenting with more than one option, and a few had even developed custom distributions for their own internal use. Despite its slim margin, Red Hat Linux 7.3 Professional, which retails for \$200, will be used for software pricing in this study.

Linux licensing models are more flexible than for Solaris and Windows. First, there are no client access licenses, per-processor charges, or other costs beyond the purchase of the software. Second, a company might purchase commercial versions for pilot projects and administrator training purposes, and then use a free, downloadable distribution for production deployments to take advantage of the economies of scale that make Linux a more compelling option.

Several participants had done exactly this, making costs for Linux software significantly lower across the servers deployed. In fact, only 27% of the Linux servers deployed were provisioned with purchased copies of their respective distributions.

Therefore, software purchase costs would be \$400 per Processing Unit.

According to Sun, Solaris licensing is typically included as part of the hardware purchase price for SPARC-based systems. Thus, the cost listed here is \$0, and the actual software cost is reflected in Section 3.2, Hardware Purchase and Maintenance. The majority of participants used Apache, an open source product, not Sun's iPlanet Web Server product. Thus, there are no further licensing costs for this case.

IT executives should note that Solaris 8 is available for the x86 architecture, which means it could be used to perform a comparison against Windows and Linux on the same hardware platform. However, survey participants had not done this, nor did they have any plans to do so. Further, earlier this year, Sun announced that it had withdrawn support for the x86 architecture beginning with Solaris version 9, and then recently issued a quiet retraction of that statement in response to customer complaints. Until the precise future of Solaris on x86-based systems is determined, RFG believes IT executives should avoid placing any bets on this combination.

Finally, Windows is licensed on a per-server basis, and is available for purchase without client access licenses (CALs) for \$700, or \$5,320 per Processing Unit. CALs are not required for anonymous Web users. However, authenticated users may require additional licensing. IT executives should check with Microsoft to be sure they are properly licensed for this type of deployment.

Microsoft's new Software Assurance program throws a wrinkle into product licensing, as it introduces the concept of a yearly maintenance fee for the ownership of the software. Microsoft's stated benefit is that this keeps the customer up to date with the latest version of the product. However, this occurs whether or not the customer finds the latest version compelling, and elects to perform the upgrade. This adds a 25% yearly maintenance fee to the purchase cost of the software. Yearly software maintenance costs are thus \$1,330.

3.2. Hardware Purchase and Maintenance

Hardware purchase and yearly maintenance costs were as follows:

Case	Purchase	Year 2 Maint.	Year 3 Maint.	Total
Linux	\$37,511	\$252	\$252	\$38,015
Solaris	\$345,400	\$21,083	\$21,083	\$387,566
Windows	\$38,524	\$259	\$259	\$39,042

RFG has found from numerous discussions with clients that other factors such as proper system administration, adequate performance monitoring, adequate system resources, and platform and application tuning play much more heavily than platform choice in a system's ability to handle user loads. This is especially true for horizontally scaled Web server clusters.

Therefore, in an effort to maintain a fair comparison for x86-based system costs, RFG configured identical systems from each of the four top x86 server hardware vendors in the market: Compaq, Dell, HP, and IBM.¹ Basic system specifications were a rack-mountable Pentium III 1.2 GHz server configured with 512MB RAM, and two 18GB SCSI hard disk drives in a hardware RAID-1 configuration, and 3-year same-business-day on-site warranty support. The average cost for this server from these vendors was \$5069, yielding a cost per Processing Unit of \$38,524 for Windows and \$37,511 for Linux.

In contrast, there was a good deal of similarity in Solaris hardware configurations. Although Solaris 8 is available for the x86 architecture, no survey participants operated with this configuration. The most common Sun servers in use were the SunFire 4800 and 6800. Among survey participants, the average cost per server was \$157,000, yielding a cost per Solaris Processing Unit of \$345,400.

Hardware maintenance costs were drawn directly from survey data. Per Processing Unit, customers

¹ Note: The initial survey that this TCO study was based on was performed prior to the acquisition of Compaq by HP. Since pricing data did not vary significantly after the product lines shifted, the original data set was kept for this case.

were paying \$252 per year for Linux deployments, and \$259 per year for Windows deployments. Maintenance for Sun systems was \$21,083 per Processing Unit per year on average.

3.3. System Support and Administration

Yearly externally purchased support and administrator salary costs per Processing Unit were as follows:

Case	Externally Purchased Support costs	Administrator Salary costs	3-Year Total for both
Linux	<\$10*	\$12,010	\$36,060
Solaris	\$19,309	\$29,509	\$146,454
Windows	\$1,520	\$46,360	\$143,640

* Although few survey participants did so, RFG believes IT executives should consider commercial support options to increase the success rates for their Linux deployments.

Support costs were those fees paid to consulting providers or product vendors for technical support, deployment planning, migration assistance, custom development not related to the Web sites themselves (such as the creation of custom Linux distributions or installation process), etc. Administrative costs were costs to employ system administrators to manage the deployment.

Support costs for Linux averaged less than \$10 per Processing Unit per year. Many administrators were taking advantage of free support resources, including mailing lists, news groups, Web site knowledge bases, and so on. However, RFG believes that IT executives will realize the highest success rates with Linux deployments when commercial support options are included as part of a deployment strategy.

The same costs for Windows averaged \$200 per server per year, or \$1,520 per Processing Unit. Finally, support costs for Solaris averaged \$8,777 per server per year, or \$19,309 per Processing Unit. Support costs for Solaris were high because most participants utilized Sun consulting services extensively for system support.

Participants using Windows generally licensed support options in the form of "incident packs"

that covered all servers in their organizations, reducing the cost per server to some degree. However, all Windows cases required some level of commercial support. Although Windows was designed for ease of use and administration, this design abstracts the administrative interfaces from the fundamental operating system layers. Thus, when serious problems occur on a Windows system, administrators are more likely to require external assistance, typically directly from Microsoft, in resolving them.

RFG wanted to determine administrative costs per Processing Unit, but unfortunately this made the calculations somewhat complex. To arrive at this figure, RFG began by collecting average administrator salary data from each participant, as well as the number of servers each administrator could manage. RFG then calculated the average administrative cost per server, and finally extended this value out to a Processing Unit for each deployment case. The raw data is as follows:

Case	Salary per Admin	Servers per Admin	Salary per Server	Salary per PU
Linux	\$71,400	44	\$1,623	\$12,010
Solaris	\$85,844	6.4*	\$13,413	\$29,509
Windows	\$61,000	10	\$6,100	\$46,360

* For the purposes of this survey, administrators were only counted as their duties related to support of the Web application deployed. Sun customers generally had very few, highly scaled systems, often managed by one or two dedicated administrators. Separate discussions with RFG clients indicate that this number is far higher – as much as 40 to 60 servers per administrator – across the entire enterprise.

It is interesting to note that despite the greater salary requirements for Linux and Solaris administrators, the greater numbers of servers they can manage yields a much lower cost per Processing Unit.

IT executives should note that administrator experience levels, system automation facilities, and network management tools could affect the numbers above. Several participants noted that their first-year administration costs for Linux were higher than they expected for future years, primarily due to one or more of the following factors:

1. Administrators of other Unix platforms, such as AIX, HP-UX, or Solaris, had been re-tasked, and required from 1-4 weeks to become familiar enough with Linux to manage the servers as effectively as they had on the prior platform.
2. System automation tasks, such as scheduled maintenance scripts, had not yet been written for the Linux systems.
3. Network management agents for products such as Openview, Tivoli, and Unicenter had not yet been deployed to monitor the Linux systems.

Finally, training and certification costs were included as part of the survey data, but differences between those costs for the various deployments were not large enough to draw a comparison.

4. "Soft" Costs

There are a number of hidden costs associated with each platform that are difficult to assign a monetary value to. For example, Windows owners are responsible for ensuring that they comply with Microsoft's licensing policies by maintaining inventories of software products installed. Microsoft has threatened audits and lawsuits where it believes companies have violated those policies. Partially out of a fear of such legal action and the penalties that they may carry, many companies have implemented costly software inventory products to perform these audits internally. These products may range from \$5 to \$50 per seat, depending on their functionality. Because Linux and Solaris are not licensed in this fashion, companies that use these products do not need to fear such penalties.

4.1. Security

Costs for system downtime due to breaches of security varied wildly across the study participants. Some quoted losses in the millions for each hour of system downtime. Others focused on the potential for lawsuits should private data, which may include credit card information or data controlled by privacy legislation such as the

Health Insurance Portability and Accountability Act (HIPAA), be obtained by hackers.

More importantly, participants were reluctant to disclose details about hacker penetrations or viral infections, despite the confidential nature of the survey. For many companies, it is standard policy to never disclose such issues to an outside party. Thus, security issues must be placed under the soft costs category. RFG believes IT executives can make their own assessments regarding the yearly cost of security issues by combining the following factors:

1. The amount of system downtime each server has historically averaged due to intrusions or viral infections.
2. The amount of system downtime each server has historically averaged due to the need to install patches and system updates.
3. The cost per hour of downtime for a given system.
4. The amount of time, and cost of that time, administrators have spent monitoring security bulletin services and software vendor security portals for patches, and installing and supporting those patches.
5. Amount of time, and cost of that time, administrators spend performing emergency repair activities after a disruptive incident occurs.

Historically, Microsoft has a poor reputation for product security, and although the vendor has been working to repair this image, RFG believes it will take quite some time to do so – perhaps as long as two to three years. Although there is some truth to the statement that all software products may contain undiscovered vulnerabilities, Microsoft is currently a political target for hackers seeking to prove a point, which makes them higher-risk systems for customers. Survey data showed that Windows installations require twice the number of administrator hours on average spent patching systems and dealing with other

security-related issues than either Solaris or Linux.

Survey participants using Windows commented more than once on the continuing requirement to reboot a server after applying a security patch. This affects the system's uptime values, and contributed to poorer performance of Windows in the area of availability as well as security.

RFG believes IT executives should identify security-related costs for their own environments in two areas. IT executives should start with administrator time spent identifying, applying, and supporting security patches, and determine the cost of this time to the company. IT executives should then consider costs for system downtime and/or public relations incidents related to intrusions, such as theft of confidential customer data like credit card information.

Looking down the road, RFG believes IT executives should also be extremely cautious when evaluating Microsoft's .NET platform products, which include game-changing elements to enable Web services, such as support for the Simple Object Access Protocol (SOAP). Although these may have beneficial long-term effects for the company, their security implications remain relatively unexplored. For example, firewall products designed specifically to filter and control SOAP-based communications are not yet on the market, and the exact security implications of SOAP have not yet been adequately determined.

4.2. Availability

System availability was treated as a soft cost, because the amount of data available regarding uptime figures would not support raw cost calculations. This is unfortunate, as participants quoted well over \$1Mil on average for expected revenue loss per hour of system downtime, making this a crucial issue.

However, operating system stability is only one aspect of availability. Along with other factors, the hardware platform plays a significant role here, and the Intel x86 architecture was not designed with high availability in mind. Although some hardware OEMs are working to bring

mainframe-class availability features to the platform, until they do, Windows will be limited by the architecture; it cannot run on other architectures.

Software design is another area where availability can be hampered. One participant noted that because Windows security patching requires a server reboot for the patch to take effect, this process creates most of their servers' measurable downtime. In contrast, most Linux patches can be handled by restarting only the affected service. This issue, coupled with the high rate of security patch release by Microsoft, has sharply cut that participant's availability figures.

4.3. Scalability

Survey data supported claims that Solaris is more vertically scalable than Windows. The same data was inconclusive with regards to the scalability of Linux vs. Solaris or Linux vs. Windows, as study participants had not performed the detailed benchmarking required to compare these systems. However, Linux cannot yet properly handle as many CPUs as Solaris. Until it does, users are likely to experience higher levels of vertical scalability with Solaris. However, although Linux is not yet as vertically scalable as Solaris, the results from the study show that this does not affect the platform cost.

4.4. Other Factors

IT executives should also consider the effect of becoming "married" to a given platform. Linux can run on a number of hardware architectures, including Intel, Power, and SPARC. IT executives that select Linux have the flexibility to change platforms should they decide that Intel is no longer the proper hardware platform choice for their companies. Java, Perl, and PHP scripts will move with the site and function properly on any platform. Linux thus provides multiple layers of flexibility, including hardware architecture, licensing model, administrator skill portability, and more. Further, Linux provides the freedom to choose distribution, support, and service providers that best suit the business needs of the company.

5. Conclusion

Three-year total costs of ownership for each deployment were as follows:

Case	Year 1	Year 2	Year 3
Linux	\$49,931	\$62,203	\$74,475
Solaris	\$394,218	\$464,119	\$534,020
Windows	\$91,724	\$141,193	\$190,662

Linux was the least expensive platform to deploy and operate. Although some initial costs were higher at points, the ability to massively scale the product horizontally without paying additional licensing fees can yield significant cost savings over the long term.

RFG believes IT executives should evaluate and pilot Linux server deployments, as the oft-quoted cost savings possible by deploying the platform were confirmed by this study. Further, to increase the cost savings possible through such a deployment, IT executives should explore educational programs for administrators as well as network and server management products, as these were the areas of greatest expense for Linux. Overall, given its low cost and flexible licensing requirements, lack of proprietary vendor goals, high level of security, and general stability and usability, Linux is worth considering for most types of server deployments.

