

Return on Investment in Linux Implementations Survey Results and Conclusions

January, 2004

Table of Contents

OVERVIEW	2
ROI versus TCO	2
Methodology and Population.....	2
RESULTS.....	3
Hardware	3
Operating System Licensing and Support	4
Layered and Application Software	4
Personnel / Human resources (HR)	6
Facilities.....	6
Incremental Revenues.....	7
ADDITIONAL FINDINGS	7
CONCLUSIONS.....	9

By **Evan Bauer and Chad Robinson**
The Robert Frances Group



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

Overview

The purpose of this study was to look beyond the cost issues that have dominated media discussions on the deployment of Linux-based systems and examine what longer-term benefits might have been achieved as measured through return on investment (ROI). RFG performed an in-depth survey of fifteen enterprises with significant Linux initiatives more than one year old.

We found a wide range of critical economic factors amongst the survey population; however all of the enterprises surveyed reported economic impacts resulting in positive ROI within one year. Those enterprises in low-margin industries tended to focus on incremental cost savings in hardware, software, and manpower in calculating their ROI; major financial services firms with more aggressive implementations that included application servers and high-performance computing (HPC) applications reported the highest returns, with several projects returning over 500 percent of the initial implementation investment.

ROI versus TCO

Above and beyond the reduction in Total Cost of Ownership (TCO) for Linux-based implementations (see the RFG study "Total Cost of Ownership for Linux Web Servers in the Enterprise"), a more interesting and general question is "what kind of Return on Investment are firms getting from Linux deployments?" For most initiatives today, businesses are using expected return on investment (expected ROI) as the basis for investment decisions and measured ROI as a means for evaluating the success of projects.

While TCO measures both direct and indirect costs of alternative implementations, it makes no measure of the financial benefits gained from that implementation. Return on Investment, in comparison, includes cost factors but also

includes the value of the resulting benefits to the enterprise attributable to the investment. In examining ROI, one looks at the total return (positive or negative) over a given period as well as the time required to achieve positive ROI -- or the "time to return". In order to shed light on the ROI for Linux-based projects, RFG conducted in-depth interviews with fifteen enterprises with more than a year's experience in deploying Linux-based systems and applications.

Methodology and Population

This survey was based on in-depth interviews with individual companies; it was not a mass-broadcast survey of a statistically significant sample for the global enterprise population. All participants answered positively to the question "Does your firm have one or more production Linux implementations that you would be willing to discuss?" There was no additional filtering of the participants in the study. These results do provide strong indications of both the growing general utility of the Linux platform and the very short time-to-return on project investments. Our survey population included the following industries:

- Higher Education
- Financial Services
- Food Services
- Mass Media
- High-Tech Manufacturing
- Retail

We were surprised to find that enterprises -- either in or across industries -- were not gathering similar cost or return metrics on their technology investments, even though all members of the sample reported positive ROI from their Linux investments. Therefore we are reporting results based on the metrics of each of the participating enterprises anecdotally and not trying to normalize them to provide any statistical summarization of dollars returned.

Our study made use of a three-section, 35-question survey instrument that examined corporate policy, infrastructure implementations, and application specifics (the application sections were repeated for each Linux-based application deployed by the responder). The formulation of the questions allowed the participants the opportunity to provide additional information regarding the specifics of and critical factors in their implementations, as well as how they measured both costs and returns. All interviews were conducted by the authors of the study, either in person or by telephone.

Results

In analyzing the results of the surveys, RFG classified the elements of ROI for platform investments in the following categories:

- Hardware
- System Software
- Layered and Application Software
- Manpower / Human resources (HR)
- Facilities
- Incremental Revenues

We will look at each category in terms of its effects on the Return on Investment to the Linux adopters.

Hardware

Notably, the universally cited area of Linux-driven initial cost savings was in the purchase of server hardware, with particular emphasis on the use of high-performance, low-cost, [Intel Corp.](#)-based hardware as a substitute for either RISC (in 12 cases) or mainframe (in one case) planned server purchases (two cases were green field deployments). All Intel (Linux on Intel) deployments were on two-way or four-way servers, both in rack and blade server configurations. These lower-end configurations were attributable to two factors:

1. Participants chose to scale-out horizontally in their Linux deployments (using clustering, load balancing, and grids) versus the scale-up approach (larger SMP configurations) used with traditional servers.
2. Participants cited the performance of the Intel Xeon IA-32 processors as providing comparable performance to at least twice as many RISC chips in their production implementations.

One firm replacing RISC servers with computational grids (built on a combination of open source and internally developed software) purchased 200 [Hewlett-Packard \(HP\)/Compaq Corp.](#) ProLiant Intel servers for roughly \$1MM, compared to a proposed \$12MM for a Sun SPARC configuration. Three firms with High-Performance Computing (HPC) workloads reported Intel price-performance versus alternative RISC/UNIX architectures of more than 10:1.

Three hardware vendors dominated the Linux implementations: HP, [IBM Corp.](#), and [Egenera, Inc.](#) HP ProLiant hardware was the most frequently cited for individual servers, with IBM xSeries a close second. Egenera dominated the blade server configurations (in terms of numbers of CPUs deployed), with IBM second, and only one responder indicating deployment of the HP blade server offering.

Note that firms making use of Linux on IBM's zSeries mainframe systems did not significantly count the cost of hardware in their ROI calculations. These firms were making use of spare capacity on existing mainframe hardware and took advantage of special pricing from IBM to make use of the resources.

Only three firms in the survey indicated any official deployments of Linux on desktops

(though most believed they had unofficial installations on the desktops of developers, systems administrators, and security specialists). One indicated some savings for engineering desktops using Linux on high-end PC workstations versus RISC workstation hardware.

Operating System Licensing and Support

Operating system (OS) licensing cost savings were rarely cited by participants as a major contributor to their ROI; it was Linux's efficiency, flexibility, and manageability that contributed substantially to the cost savings and new opportunity enablement. The one exception to this was in the more apples-to-apples comparison with [Microsoft Corp.](#) Windows Server environments. Users with major distributed deployments (branch offices, cash registers, and ATM machines) cited savings on each device in the hundreds of dollars. Savings of \$300 per ATM seemed inconsequential until the responders described a national branch-banking network with tens of thousands of devices deployed.

Most enterprises selected either the [Red Hat, Inc.](#) Advanced Server (first among our responders), followed by the [SuSE, Inc.](#) Enterprise Server as their official Linux distributions. Two responders had engineered their own builds from sources available on the Internet. Following their practices with proprietary UNIX and Windows operating system deployments, the major financial institutions had their own engineered "builds" derived from the standard vendor OS distributions.

The real system software cost savings found by our survey was in the support costs associated with major Linux deployments. Linux adopters have been taking advantage of the additional support flexibility that an open source software product provides. The first area of savings was

due to competition in the area of support services. Although Red Hat was the most frequently selected distribution, IBM had the two largest enterprise operating system support contracts issued by the enterprises surveyed. More firms relied on their hardware vendors for support than on the Linux distribution vendor.

Other organizations reported that they were purchasing support for a smaller percentage of their servers, relying on internal staff skills and the open source community to resolve most issues -- augmenting with per incident support packages. Two large enterprises had chosen to do all of their support using in-house skills and Internet Linux-community resources.

It is worth noting at this point that none of the survey respondents believed that various legal actions by Caldera International Inc. (The SCO Group) were of sufficient concern to slow or delay any of their Linux deployment plans. Most are tracking the issues and have discussed them with their legal departments, but see no reason for end-user organizations not to take advantage of the Linux opportunity. One firm described SCO's actions as "the last flailing shot of a dying bird."

Layered and Application Software

This was the area of greatest diversity among the study's participants. We will look at responses in this space both for their specific contributions to achieving ROI and for what they can tell us about the state and maturity of the Linux platform today.

Ports of existing applications to Linux dominated in the initial deployments at all firms surveyed, with new application rollouts to Linux coming into play 6-12 months later for about half the firms. The initial application ports were, in all but two cases, from proprietary RISC/UNIX systems, with one additional from Windows and a second from a mainframe

platform. Enterprises chose to do these ports both to explore "the next promising platform" and to lower hardware purchase costs needed for upgrades, replacements, or consolidations. With the single exception of the Windows port (where the savings were seen from reduced administration and support costs), firms reported hardware purchase savings of 50-90 percent.

Applications ported were written in C, C++, and Java. In all surveyed cases, firms using application servers (WebLogic or WebSphere) did not make a change to an open source J2EE server environment (e.g. JBOSS or Tomcat), though several indicated that they expected to evaluate those environments in the future. One responder has implemented a Linux and blade server-based utility for hosting Web and Web services applications using BEA WebLogic clusters.

They found superior per-chip performance to their previous Sun UltraSparc environment when using the Sun Java Virtual Machine (JVM) for the Intel Linux platform. Testing the use of both the BEA/Intel JRocket JVM and the IBM JVM has shown substantial additional application performance increases and one of the alternative JVMs is planned for production deployment in the near future. Another responder has been working with Oracle to build a 9i and 9i AS (application server) combined implementation on Linux, production rollout is expected for early 2004.

The highest ROI applications reported were all in the high performance computing (HPC) category. No two firms surveyed reported using the same software packages to support workload distribution and correlation. Commercial products mentioned were the [DataSynapse GridServer](#) and [Platform Computing LSF](#). Open source packages used in HPC clusters and grids were ACE, the Globus Toolkit, and MPI. Two firms reported using these in combination with

in-house developed workload distribution software.

Utility server deployments were made at 70 percent of the respondents. In these cases, additional open source packages (some seen as "just part of the operating system") were used. Leading those cited were Apache for Web serving, Samba for file and print services, and Squid for Web proxy services. These were the largest server counts of the individual deployments

Three firms identified database servers as the next area for high ROI deployments. Interest seems to be split between Oracle and Sybase (both the IQ and ASE products), with one firm interested in porting their DB2 databases. Two of the largest financial services firms surveyed were testing in preparation for volume deployments of Oracle 9i RAC clusters on Egenera BladeFrame hardware.

All firms surveyed responded that (with the exception of new grid deployments) their Linux-based systems were to take part in their existing middleware architectures. IBM WebSphere MQ (formerly MQ Series) dominated the enterprise middleware selection for Linux deployments, with TIBCO and Sybase products also mentioned. One firm with a proprietary intra-system messaging product has ported from Windows to Linux. The port allowed them to increase the numbers of connections per server from 5,000 (a Windows limit) to 15,000; since the system had not previously been CPU-constrained, two or more years of planned additional hardware purchases were no longer required.

Availability of independent software vendor (ISV) products (both applications and libraries) was cited by 33 percent of the responders as their number one constraint on deployment, although all but one indicated that the situation

had substantially improved over the last 12-18 months.

Personnel / Human resources (HR)

Personnel costs (employees, contractors, consultants, and out-sourced labor) remain the single largest item in most enterprise IT budgets. Our survey showed that firms are finding substantial ROI contributions through the personnel elements of their Linux deployments.

Linux systems yield increased server manageability over Windows. For enterprises with large utility server deployments or computational grids containing hundreds of servers these cost saving were substantial. The combination of an open environment, strong scripting and GUI administrative tools, and the ability of systems administrators and systems programmers to get "much closer" to the OS itself provided administrator to server ratios "approaching an order of magnitude" better for Linux deployments.

Linux servers, surprisingly, allowed a higher ratio of servers to systems administrators (SA) than did firms' legacy RISC UNIX environments. One respondent reported that they had "50-60 servers per administrator for Solaris and 90-100 for Linux." RFG believes that some of this may be attributable to Linux being the newer environment -- obsolete and inefficient practices may not be baked in to standard operations. Linux deployments tend to be going in as a platform, not just as incremental servers, meaning that firms are considering management issues from the beginning of the deployment. At the same time, the Linux community puts an ongoing emphasis on manageability and tools for software distribution (including patching) and system maintenance -- this seems to be paying off for our respondents.

In most enterprises, training costs associated with deploying Linux were negligible. The large financial institutions all had training programs associated with their deployments, usually of one to two weeks of instruction per individual. This was, however, not the norm. More than half of our sample had no formal training program, nor did they feel that they needed one. As one respondent told us "Linux is just another UNIX", existing skills (especially for developers) remain applicable, new features and facilities are open and accessible. More than one enterprise observed that their best technical people had already been working with Linux before it became an organizational initiative. "We did our pilot entirely with volunteers, most had good Linux knowledge and experience before the project began."

Firms doing recruiting of new technologists from universities found that the opportunity to work with Linux (versus Windows or UNIX) environments was a draw to job candidates. One firm indicated that the new graduates were bringing Linux skills with them from school.

None of the large UNIX shops in our survey reported adding any additional headcount when adding Linux to their operating system portfolio. Whether adding dozens or thousands of servers, existing headcount was sufficient to get the work done. Linux work (including the creation of derived or custom distributions) was part time work (reported between 1 percent and 25 percent) for engineers and systems administrators.

Facilities

Facilities is another category where contribution to ROI is made on the savings, not revenue side of the equation. In this case, Linux by itself made no direct contribution to the population's bottom-line; but as an element in server consolidation strategies the contribution was measurable at about 40 percent of the sites

surveyed. Server count went down amongst our respondents for three reported reasons:

1. Faster Intel chips required fewer server instances for the same application or utility workload.
2. Use of virtualization (both VMware on Intel and VM on zSeries) allowed a single, more powerful, physical server instance to provide multiple, necessarily independent, operating system instances. This was cited at multiple locations for development and testing, with production use of virtualized servers planned for 2004.
3. Users of blade server technologies reported that the ability to do N+1 failover strategies (rather than have a dedicated spare box for each application server) reduced the need for unused spare capacity by "30-60 percent".

When combined with the footprint and wiring savings provided by blade servers, substantial reductions in facilities costs were achieved. One firm noted, however, that blade server densities require careful planning for power, cooling, and raised floor loading "you can't just drop these things in".

Incremental Revenues

The firms that reported remarkable ROIs were those whose Linux investments were part of new business initiatives that helped them drive new businesses or increases in incremental revenues. In the study sample, these all came from High Performance Computing (HPC) applications where horizontal scaling of processing power enabled new applications. Although the applications themselves were not dependent on the choice of operating systems, the economics of large-scale Linux deployments are lowering both the threshold and risks associated with their implementation.

One enterprise reported constructing a compute cluster for under \$1 million on a Linux blade server as an alternative to a \$5 million RISC-based supercomputer. The trading system deployed has paid for itself more than ten times over during its first year of operation; although a RISC configuration would also have produced positive returns, its total cost had been seen as raising the risk of the project to an unacceptable level. So without the Intel price-point, the business opportunity would never have been exploited.

The deployment of Linux-based grid and cluster configurations showed remarkable benefits in making very computationally intensive operations available "without a NASA-scale budget." We found the trend among the leading firms was to first parallelize existing applications (using libraries like MPI and ACE) for quick initial returns, then to start the construction of standard computational grid facilities to provide utility or backbone capability for a range of new applications. Even higher ROI can be expected for the applications built on these systems, as engineering and infrastructure costs will be spread over a wider range of application systems. Additional benefits will be gained through development teams working on common toolkits -- resulting both code and experience.

Additional Findings

There were other data points and trends -- in addition to the elements that contributed to return on investment for their Linux-based projects.

- Use of Linux was not an automatic indication that enterprises would select other open source packages over commercial products. At least one of database, middleware, backup/restore, and other packages were purchased for

their Linux implementations by all of our respondents.

- Desktop deployments are still immature, though 70 percent of those surveyed indicated possible future interest in Linux desktops for at least some of their population. It is also worth noting that application development for Linux servers was done in most firms on a mix of Linux and Windows desktops.
- ISV support for the Linux platform has progressed dramatically since RFG's TCO survey of just one year ago. In many cases second-tier vendors of vertical application packages are holding back deployments in some industries, but our survey respondents indicated that they are putting pressure on those firms to port when the end-user economics warrant it.
- The dominant deployment platform remains 2-4 CPU Intel-based systems, either rack mounted or blades. Few participants had any interest in scaling this number vertically within an SMP -- though two financial institutions were interested in higher performance clustering incorporating distributed shared memory.
- There was very little interest in 64-bit Linux today, though some considering large future database deployments are tracking the situation. Most respondents saw more near-term promise in the [Advanced Micro Devices \(AMD\)](#) 64-bit Opteron chipset versus Intel's IA64.

Conclusions

RFG believes, based on this study and our ongoing research on the applications and opportunities for Linux and Open Source that Linux adopters should continue to see consistent strong returns on investment for the foreseeable future. In particular we see strengthening returns from:

- Linux-based deployments of Java and PHP application servers, scaling horizontally for availability while providing very low effort ports from UNIX and Windows server platforms.
- Continued revenue growth from both compute and data grid implementations. Enterprises that invest in these facilities should be able to create new revenue generating opportunities from high performance computing and data mining without the cost of creating platforms and capacity for each new application.
- New support and services models -- encouraged and enabled by competition and the web-based Open Source community -- will allow enterprises to buy more flexible support offerings for their Linux-based systems, lowering costs without necessarily adding support risk.
- Linux use will encourage further adoption of a variety of strong Open Source software packages, further lowering costs and encouraging new projects that will combine these tools and platforms into application systems. Strong 2004 candidates include Apache, Tomcat, JBoss, MySQL/MaxDB, Squid, PHP, among others. Like Linux, all of these packages have commercial support available.
- Hardware selection flexibility without requiring adoption of multiple operating systems. With developer and SA skill-sets common across Linux platforms, firms will lower costs and improve time-to-market by leveraging Linux expertise across blade-based Intel compute grids, zLinux-based data engines, and a variety of configurations in-between.
- Desktop Linux opportunities, though continuing to lag server deployments, should provide compelling ROI opportunities for engineering, technical, and IT workstations at the high-end as well as for "mail and web" users at the low-end within both large and small enterprises. The results seen by very early adopters in our survey will, as the word spreads, encourage other organizations looking for opportunities and options on the desktop.
- The release of the long-awaited 2.6 version of the Linux kernel, which enhances SMP performance, enables 64-bit implementations, improves TCP/IP performance, enhances security, and improves desktop capabilities. Most enterprises won't start deploying 2.6-based systems until mid-2004, but the technical strengths of the Linux platform remain on the rise.

As RFG's 2003 study indicates, senior IT executives can look improve their platform ROI by implementing Linux-based solutions that take advantage of reduced hardware and support costs while enabling innovation by taking the sticker-shock out of large-scale data and compute grid deployments.

This report was sponsored by IBM.