

The Intel logo, consisting of the word "intel." in a lowercase, sans-serif font, positioned on a blue vertical bar on the left side of the page.

Migrating to Linux* on Intel[®] Architecture

The Intel Solution Services logo, featuring the word "Intel" in a small font above "solution" in a larger, bold font, with "services" in a smaller font below it, all in white text on a blue background.

Table of Contents

- Introduction 2
- The Decline of Proprietary Architectures 2
- The Rise of Linux and the Intel Architecture-based Platform..... 3
- Linux Enterprise Readiness 7
- Total Cost of Ownership 9
- Realizing the Benefits: Experiences within Intel 13
- Approach to Migration..... 13
- Case Study 18
- Conclusion 20
- Appendix A..... 21
- Appendix B..... 22
- Appendix C 26

Introduction

Businesses have the increasing challenge of improving the performance of their system infrastructure and delivering more, while balancing the pressures of reducing overall costs. At the same time, existing servers and equipment are aging such that they cannot support the growing capacity and performance demands. Businesses are increasingly seeing their legacy systems reach the end of life, either through expiring leases or vendors ceasing to support the systems. The ongoing support costs of existing machines and platforms must be analyzed to ensure capacity for potential future growth. The competitive marketplace leaves no room for falling behind or for accruing extraneous costs. In light of this, companies have to consider how to obtain a desirable return on investment (ROI), either through staying on the same platform or considering an alternative platform.

When evaluating what platform offers the best ROI, the following should be considered:

- How to reduce total cost of ownership (TCO)? Reduce manageability costs? Reduce license costs? Reduce support costs?
- How to improve performance and quality of service?
- How to allow for future growth and flexibility?

The challenge comes in evaluating and balancing all these demands to obtain the best ROI. Increasingly, businesses are looking to an Intel[®] Architecture-based platform running the Linux* operating system to reduce their TCO and boost performance, without compromising the availability and reliability that is needed for enterprise applications.

This document describes items to consider when migrating from a RISC-based environment to Linux* running on Intel architecture. This is not a how-to or integration guide; rather this paper documents some examples of other companies who have performed such a migration and some of the tools that have been used.

This paper analyzes some of growing trends in the marketplace in regards to migrating to Linux on Intel architecture, as well as the financial benefits, including TCO benefits. In addition, this paper summarizes some of the advantages of the Linux and Intel architecture-based platform, as well as proven methodologies in migrating to such a platform. Lastly, this paper provides key migration learnings and experiences from Intel and the industry.

The Decline of Proprietary Architectures

Many companies are realizing the risks and disadvantages of being on a proprietary, RISC-based platform. The market share for proprietary operating systems and for RISC-based servers has declined steadily over the past decade. Gartner Dataquest reports have repeatedly shown this trend of the declining RISC-UNIX market share. One example of this was seen in Western Europe during third quarter 2001 when Gartner Dataquest reported, "Shipments of RISC/UNIX-based servers fell by 36 percent and 33 percent in unit and revenue terms, respectively."¹

¹ "Gartner Dataquest Says Server Revenue in Western Europe Declined by 21 Percent in Third Quarter 2001." 16 November 2001. <http://www.gartner.com/5_about/press_releases/2001/pr20011218a.jsp>.

Companies using proprietary systems do not have the flexibility and selection of hardware original equipment manufacturers (OEMs) and software operating systems. The operating system is a key part of software development. New applications are written for specific operating system's application programming interfaces (API), which are the sets of commands that an operating system provides to give programs access to its functionality. System vendors have used the API as a vendor-specific standard that makes it difficult for customers to move off the platform after they have begun to use applications that are optimized for that particular API. Customers have found themselves locked into platforms once they have written and bought software to use on them.² In such situations, these companies are at the mercy of prices set by single vendors, and therefore, have limited options. With a Linux and Intel architecture-based platform, there are a greater number of hardware system vendors and Linux operating system vendors to select from, with potentially more competitive prices.

The reality of the cost of being "locked" into a proprietary platform is apparent, and more and more companies are realizing it. In 2002, Merrill Lynch decided to do a large-scale, company-wide Linux deployment in an effort to cut its costs and boost revenue. Rick Carey, the chief technology architect at Merrill Lynch who oversaw this effort, states, "We are telling all of our vendors that they need to have some kind of Linux strategy. We are hearing that consistently from everyone on Wall Street."³

Merrill Lynch was suffering from having proprietary hardware and many versions of UNIX to support. The developers wrote software for every version of UNIX, including the tools and patches. Carey said this approach was very time-consuming and expensive. If a UNIX project did not work for some reason, the technology was rarely transferable to another project. "When I have proprietary hardware and software, I have sunk costs into that project that I can't recover," says Carey.

The Rise of Linux and the Intel Architecture-based Platform

In contrast with the decline of the proprietary RISC-based platforms, the industry is seeing increasing growth for Linux on Intel architecture. The growth of Linux in the industry has made others take notice throughout the past couple of years.

"The Linux market segment was one of the more fortunate sectors during the unprecedented downturn of 2001," says IDC Analyst AI Gillen. "Unlike its close cousin the UNIX market (in which both shipments and revenue were down by a significant margin), Linux unit shipments were up in 2001 compared with 2000, and revenue did not experience a decline."⁴

The following diagram shows the Linux market trend that IDC observed in 2001. It provides historical data as well as revenue forecasts for the Linux operating environments for the period of 1999 through 2006.

² "Technology: Infrastructure Software." Linux Handbook. Goldman Sachs Global Equity Research. 2 January 2003. <http://h30046.www3.hp.com/uploads/whitepapers/GS_Linux_Handbook.pdf?pagesite=LINUX>.

³ DiCarlo, Lisa. "Wall Street Embraces Linux." Forbes. 27 March 2002. <<http://www.forbes.com/home/2002/03/27/0327linux.html>>.

⁴ Gillen, AI. "Worldwide Linux Operating Environments Forecast, 2002-2006: Client Shipments Pick Up the Pace." IDC Bulletin. February 200. Document # 26592. <<http://www.idc.com>>.

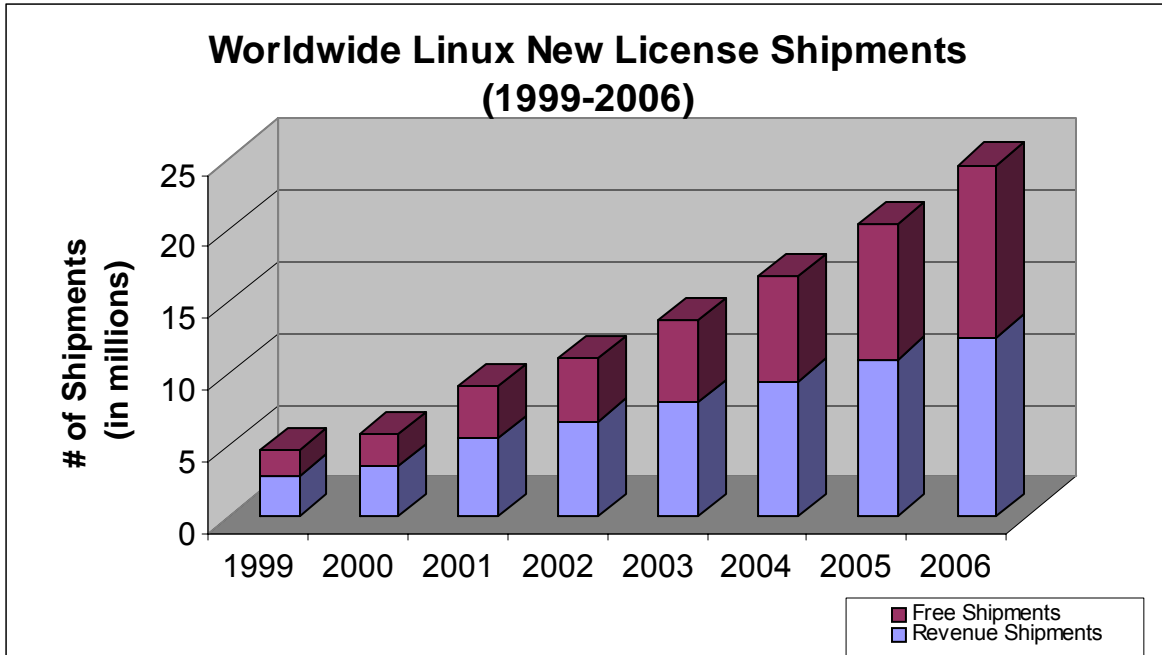


Figure 1 - Growth of Linux Shipments⁵

In a 2002 report, the IDC projected the worldwide server system shipments' compound annual growth rate (CAGR) would be 13.9 percent through 2006, from 4.3 million shipments in 2001 to 8.2 million shipments in 2006. The IDC estimates the shipments of servers running Linux would grow at a 33.6 percent CAGR (from 11.4 percent of total server shipments in 2001 to 25.2 percent in 2006).⁶

As the demands for Linux have increased, so have the demands for the Intel architecture-based platform. The Intel architecture-based platform running Linux provides the price and performance advantages that companies are looking for in comparison to proprietary, RISC-based platforms.

In comparing the number of server shipments from the first quarter of 2002 to the first quarter of 2003, the number of Intel® processor-based server shipments continues to be higher than the number of RISC-based server shipments. In addition, the number of Intel processor-based servers shipped showed an increase, while the number of RISC-based server shipments showed a decline. The following diagram shows some of the market trends.

⁵ Gillen, Al. "Worldwide Linux Operating Environments Forecast, 2002-2006: Client Shipments Pick Up the Pace." IDC Bulletin. February 2002. Document # 26592. <<http://www.idc.com>>.

⁶ Josselyn, Stephen, Mark Melenovsky, Jean Bozman, and Vernon Turner. "Worldwide Server Market Forecast Update, 2001-2006." IDC Report. November 2002. Document #28257. <<http://www.idc.com>>.

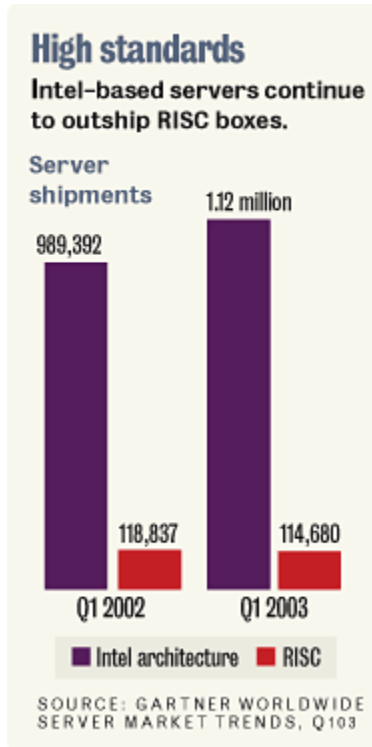


Figure 2 - Number of Server Shipments⁷

Much of the steady incline in Intel architecture-based server shipments can be attributed to the price and performance benefits of the platform. The following graph shows the highest performing and lowest cost per transaction per minute results for Intel architecture in each year, and general trending for the platform. The value improvement equates to a 47 percent compound annual rate of decline.

⁷ Gonzalez, Joseph and Jeffrey Hewitt. "Worldwide Server Market Trends, 1Q03." Gartner Report. 15 August 2003. <<http://www3.gartner.com/lnit>>.

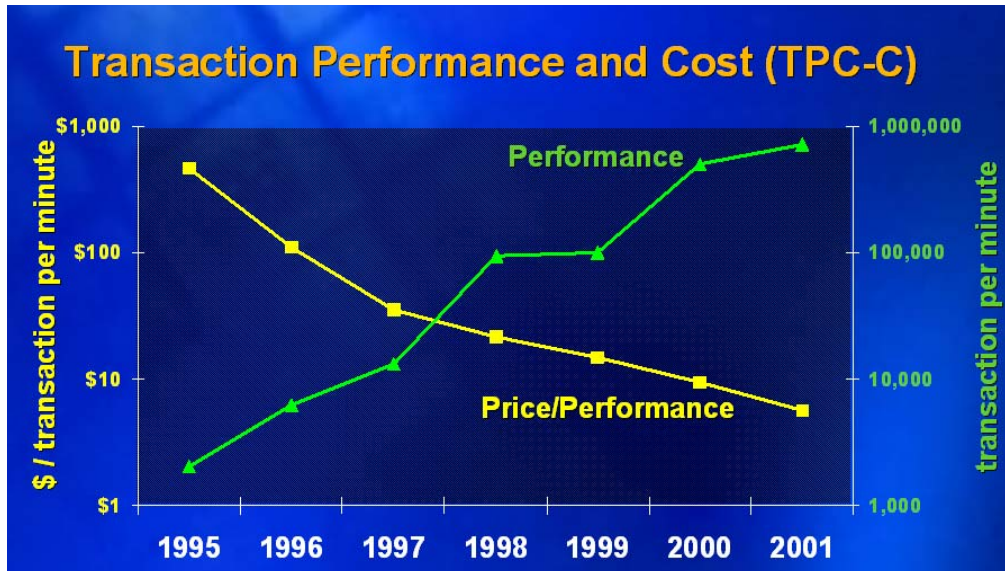


Figure 3 - Transaction Performance and Cost on Intel Architecture⁸

Intel architecture is compelling for various reasons. The increasing use of N-tier application architectures (e.g. back-end database management system servers, application servers handling business logic, front-end Web and presentation servers) enables a scale-out, or strength-in-numbers, approach to address both scalability and availability at the application server and front-end presentation server tiers. Application server and presentation server requirements can be met cost-effectively with Intel architecture-based servers. With the emergence of the server “blade” form factor, these servers are very dense and cost-effective solutions. Price-performance benefits also exist when consolidating applications and scaling up to larger servers. In addition, Intel architecture is gaining wide range acceptance and usage in the database back-end tier and is increasingly used as the high-end enterprise platform. All of these factors together will accelerate the growth of Intel architecture-based server unit volumes.

The combination of Linux and Intel architecture is compelling because the Linux operating system was developed and optimized for Intel architecture. Recognizing a need for a highly reliable and inexpensive operating system that would run on Intel architecture, Linus Torvalds wrote Linux in 1991. He made it publicly available with its open source code, and users have been enhancing Torvalds’ Linux ever since. It is accepted by the Intel architecture-based systems worldwide community and continues to be improved to serve the growing needs of its users.

E*TRADE is one example of a company that adopted Linux as a commoditized operating system to save money. “The important thing for us is to provide the best service for our customers,” says E*TRADE CTO Josh Levine. “The best way for E*TRADE to provide its customers with top-notch service is for the company itself to make money. This is why the company decided to use Linux.”⁹

Levine says the company replaced approximately 300 of its servers (possibly more), and implemented a commoditized operating system, which allowed the company to maximize cost efficiencies. Just four months after moving to Linux, Levine says the company has reduced technology costs by \$65 million. He

⁸ Transaction Processing Performance Council. <<http://www.tpc.org>>.

⁹ Olavsrud, Thor. “E*Trade Goes Linux with IBM.” *siliconvalley.internet.com*. January 30, 2002. <<http://siliconvalley.internet.com/news/article.php/965301>>.

adds the company expects to derive significant savings in the future as a result of moving to open source software.

As demonstrated by E*TRADE, the movement towards Linux running on the Intel platform is gaining momentum. The major drivers for the adoption of Linux on Intel architecture-based platforms are:

1. Linux is reaching enterprise readiness. Key infrastructure software vendors and major systems vendors support Linux. Enterprise versions of Linux have been released, and support and service models are in place.
2. Linux allows companies to lower total costs while maximizing high performance.
3. Migrating from UNIX to Linux is relatively easy. Because Linux is a UNIX-like operating system, there is an ease of migration, and the learning curve for staff to make the transition to Linux is not very high.
4. There is a trend towards server commoditization and server consolidation.

In November 2001, Merrill Lynch saw the advantages of improved performance, scalability, availability and lower cost with an Intel architecture and Linux solution. The company engaged Intel® Solution Services, Intel's professional services organization, to migrate its application from RISC-based systems to an Intel architecture-based system running Oracle9* Real Application Cluster (RAC) and Linux. In four months, the new platform showed an immediate two-fold performance improvement over the RISC-based platform and an anticipated 50 percent reduction in TCO.¹⁰

Linux Enterprise Readiness

The growth of Linux continues, and it is increasingly considered an enterprise-ready operating system. Previously, Linux was primarily used as a platform for file and print servers, email servers, Web and Internet servers, and firewall servers. Today, Linux is used for all tiers of enterprise applications, from application servers to highly available database clusters.

"In our view, Linux has evolved into an enterprise-class operating system that will have a significant and lasting presence in the IT landscape, and its continued emergence will cause considerable changes in the enterprise IT vendor ecosystem," reports Goldman Sachs Global Equity Research. "We believe its strongest effects will be seen in the corporate data center, where we see a shift occurring toward Linux-on-Intel servers away from the current paradigm of proprietary UNIX-on-RISC systems."¹¹

The enterprise environment is usually business-critical, and downtime and maintenance costs can be substantial. Businesses of every size need ways to minimize downtime without allocating huge amounts of computing resources, time and staff. Also, the design and deployment phases are more complex and the expected lifetime of the system is longer. The terms *reliability*, *availability*, *scalability* and *manageability* (RASM) are often used to describe the performance expectations for an enterprise-class environment.

¹⁰ "World Leading Finance House Migrates to Leading Platform." Intel® Solution Services. <http://www.intel.com/internetservices/intelsolutionservices/downloads/Merril_Lynch.pdf>.

¹¹ "Technology: Infrastructure Software." Linux Handbook. Goldman Sachs Global Equity Research. 2 January 2003. <http://h30046.www3.hp.com/uploads/whitepapers/GS_Linux_Handbook.pdf?pagesite=LINUX>.

Bloor Research Report

In January 2003, Bloor Research North America completed a study looking at the enterprise-readiness of Linux. Three years prior, Bloor Research had published a report concluding Linux was not yet ready to support large enterprise applications. After re-examining Linux scalability, availability, reliability, security, manageability, flexibility, as well as server consolidation characteristics in its 2003 report, Bloor Research now believes Linux is enterprise ready. Here is a summary of some of Bloor Research's findings.

Linux Scalability: Linux scales well vertically to 6-way symmetric multiprocessing (SMP) on Intel hardware. In the next three months, expect to find Linux 2.5 scaling to be back ported into Linux 2.4 kernel distributions, allowing for 8-way scaling. The next major revision, Linux 2.6 to be released in the coming year, will provide up to 16-way scaling. Linux also scales extremely well horizontally in distributed grid computing configurations.

Linux Availability: The provisions for availability have been added to the operating system and have been provided in the major Linux distributions. Failover extensions are now included and can be found in the base Linux kernel.

Linux Reliability: Almost 90 percent of Linux is installed on Intel platforms, and it is generally known to be reliable. The Linux operating environment has also proven to be reliable when used to run dedicated applications. When Linux systems have failed, the failures have largely been caused by incompatible applications contending for the same system resources, poorly written device drivers or limitations in the operating environment (e.g., early revisions of Linux were not written to exploit multi-processor environments). One way to avoid this dilemma is to buy "advanced server" or "enterprise server" pre-tested environments from reputable Linux suppliers.

Linux Security: When you think "Linux security," think "UNIX security." The security provided by both systems is similar; however, Bloor Research did find that Linux security and UNIX security vary greatly in one respect: openness. Because Linux is based on open source code, a huge community of developers closely scrutinizes Linux code, and they quickly reveal any code-related security issues, which can be addressed promptly. In addition, Linux developers can build their own layers of security directly on the Linux kernel. This is beneficial for enterprises and governments who want to invest in specialized security development.

Linux Manageability: Although many UNIX-based manageability tools, utilities and applications can be used to manage Linux environments, very sophisticated Linux management tools (including workload balancing, performance tuning and other management products) are available from grid vendors in the form of distributed resource management tools, utilities and applications. Also, some vendors with major commitments to systems management (such as Computer Associates, IBM and Sun) have completed UNIX to Linux systems management ports and currently offer rich suites of Linux management tools, utilities and applications.¹² For some examples of and references to management tools, please see the Appendix.

¹² Clabby, Joe. "Linux Enterprise Ready?: Scalability, Reliability, Security, Flexibility and Total Cost of Ownership Considerations." Bloor Research. January 2003.

Total Cost of Ownership

Given the enterprise readiness of Linux and the movement towards Linux and Intel architecture-based deployments, it is important to consider total cost when evaluating the Linux and Intel architecture-based platform. The use of Linux on Intel architecture-based servers is rising, but what is the cost of the solution and does it provide the best ROI? Does a Linux and Intel architecture-based environment provide lower TCO?

In considering this, it is important to first level-set on the use of the term “total cost of ownership”. The term “total cost of ownership” was created by the Gartner Group and is used to measure the real cost of deploying a new solution. It provides a balance for weighing the purchase price of a product with the deployment, maintenance and support costs associated with the product.

When calculating TCO, the following components are evaluated:

- Software – including operating system and applications
- Hardware
- Personnel
- Maintenance
- Facilities
- Migration
- Downtime
- Services

Robert Frances Group TCO Study

Robert Frances Group (RFG) is an advisory service for Global 2000 and mid-market executives concerned with managing the business of IT. RFG conducted numerous surveys with IT executives at Global 2000 companies and learned that Linux server deployments are well under way, with notable cost savings in many cases. To further explore and quantify this anecdotal evidence, RFG performed a TCO study to evaluate Linux deployments in the enterprise. The TCO results for Linux versus Solaris* calculated in the study are shown below.

The three-year TCO for each deployment was as follows:

Case	Year 1	Year 2	Year 3
Linux	\$49,931	\$62,203	\$74,475
Solaris	\$421,718	\$491,619	\$561,520

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 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.¹³

¹³ "Total Cost of Ownership for Linux Web Servers in the Enterprise." Robert Frances Group Study. July 2002. <<http://www.rfgonline.com>>.

IDC: Linux TCO Analysis

An IDC white paper focusing on the role of Linux in reducing the cost of enterprise computing compared the TCO of Linux on Intel architecture-based servers and UNIX on RISC-based servers. The associated costs running Linux on Intel architecture are not only dramatically lower for the hardware and software, but the costs are also comparable or lower for staffing.

Linux on the Intel platform has emerged as a viable alternative to proprietary solutions for enterprise computing. For enterprises with the right mix of requirements and skill, Linux offers tremendous potential to lower costs associated with supporting application workloads. The results of a study conducted by IDC on the cost of computing showed Linux running in an Intel architecture-based environment has a lower cost of ownership than that of the RISC and UNIX-based environment.

As the following figure demonstrates, Linux delivers a 1.8:1 cost advantage over the RISC/UNIX-based environment using Internet/intranet/extranet workloads and a 5.5:1 cost advantage for collaborative workloads.

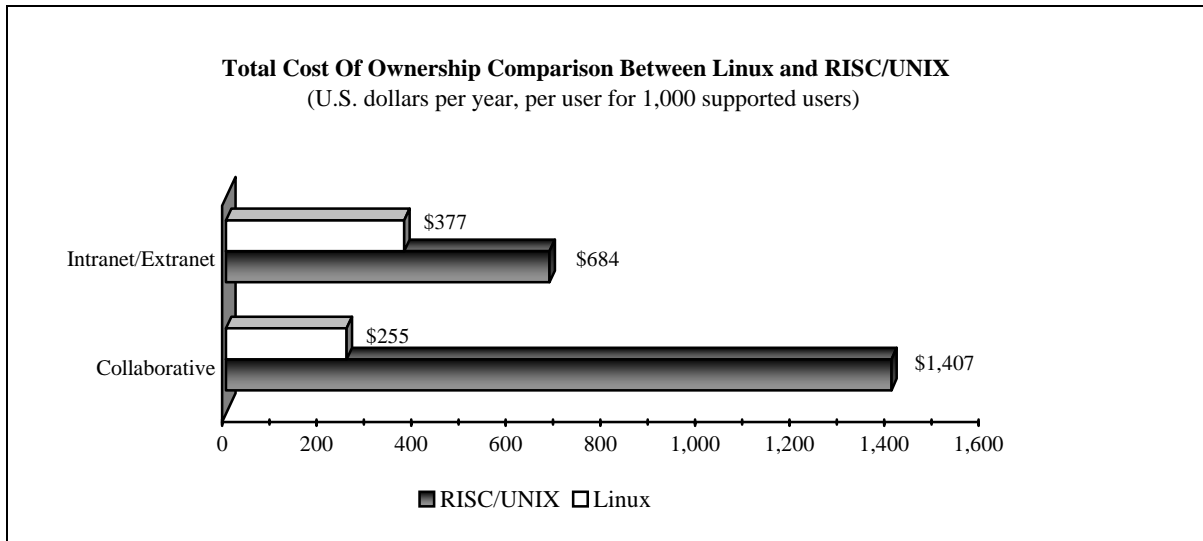


Figure 4 - TCO Comparison Between Linux and RISC/UNIX¹⁴

IDC TCO Process at a Glance

The process IDC uses to calculate TCO includes the following:

1. Determine the number of staff supporting the system
2. Determine the percentage of time spent supporting the system
3. Calculate weekly staffing hours
4. Convert to yearly staffing hours
5. Convert to cost dollars, including overhead allowance
6. Normalize to per user cost for 1,000 supported users

¹⁴ Gillen, Al, Dan Kusnetzky and Scott McLarmon. "The Role of Linux in Reducing the Cost of Enterprise Computing." IDC. November 2001. <<http://www.idc.com>>.

7. Determine hardware and software licensing costs; normalize to 1,000 supported users
8. Combine support cost and hardware, software and installation costs to determine TCO value

Internet/Intranet/Extranet Workloads

The following figure presents the TCO comparisons between Linux and RISC/UNIX workloads as determined by the IDC study. First year TCO comparisons show RISC/UNIX workload costs at \$684 per year per user for 1,000 users. By comparison, Linux workload costs are \$377. Figure 6 demonstrates the RISC/UNIX systems had a first year TCO at 80 percent higher than the Linux system.

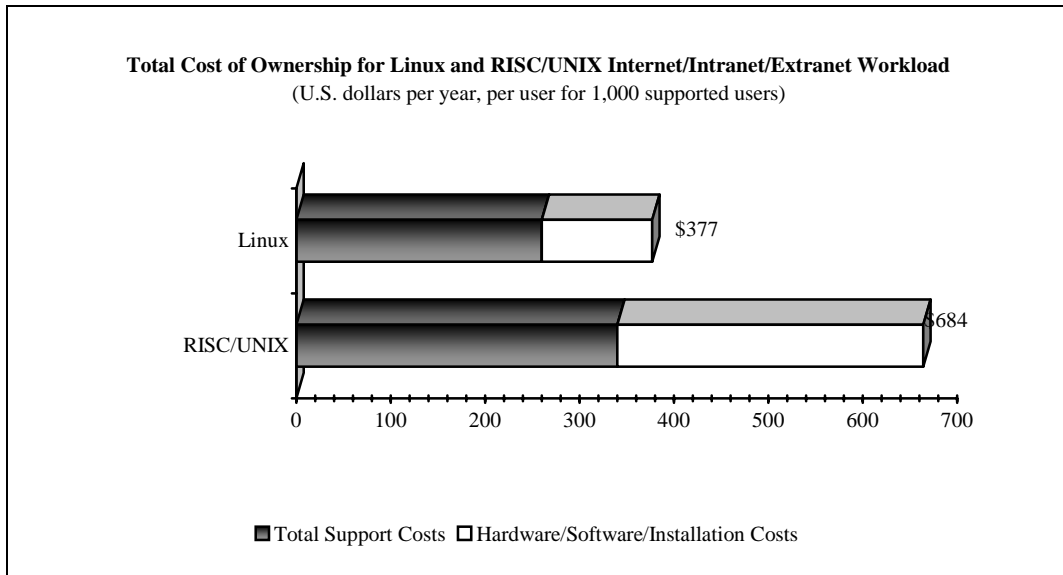


Figure 5 - TCO for Linux and RISC/UNIX Internet/Intranet/Extranet Workloads¹⁵

Support Cost Comparisons

It is interesting to compare areas with significant differences between these two platforms. Figure 7 and Figure 8 present summary TCO comparisons on a line-by-line basis for Internet/intranet/extranet workloads.

¹⁵ Gillen, Al, Dan Kusnetzky and Scott McLarmon. "The Role of Linux in Reducing the Cost of Enterprise Computing." IDC. November 2001. <<http://www.idc.com>>.




Figure 6 - Details of TCO of Linux and RISC/UNIX for Internet/Intranet/Extranet Workloads (U.S dollars per year, per user for 1,000 supported users)¹⁶

<i>(U.S. dollars per year, per user for 1,000 supported users)</i>		
Cost Area	Linux	RISC/UNIX
Deinstallation and disposal of desktop systems	4	30
Procurement	7	31
Administration	27	50
Web-site management	25	40
Asset management administration	6	16
System backup	9	16
Upgrades/moves/adds/changes	11	17
Network management	26	22
Planning/management	17	12
Database management	60	54
Operations	23	13
User support	51	41
Total support costs	266	341
Per-user server hardware/software/ installation costs	111	343
Total one-year cost of ownership	377	684

Source: IDC, 2001

Figure 7 - Details of TCO of Linux and RISC/UNIX for Internet/Intranet/Extranet Workloads¹⁷

¹⁶ Gillen, AI, Dan Kusnetzky and Scott McLarmon. "The Role of Linux in Reducing the Cost of Enterprise Computing." IDC. November 2001. <<http://www.idc.com>>.

Realizing the Benefits: Experiences within Intel

Engineering teams at Intel have implemented nearly 3,000 Intel architecture-based systems with the open code Linux operating system to provide high-performance, low-cost solutions for several of Intel's CPU projects. Linux on Intel architecture was chosen because of the stability and performance needed to deliver the required compute cycles within the project timeline, as well for the TCO cost advantages over that of a comparable RISC installation.

To build a production validation environment based on Intel architecture, it was necessary to select a UNIX operating system that could provide the required stability and performance. The team evaluated several variants of UNIX, including Linux. Each met most of the functionality and performance requirements, but Linux was chosen for the following long-term benefits:

- The Linux operating system was developed and optimized for Intel architecture.
- The open source software model ensures the quality of the code base.
- Industry momentum for Linux is growing rapidly. A variety of tools and utilities for design engineering are available.
- Worldwide support is available, which is an important requirement for global design teams.
- With Intel architecture's open standards, the best hardware platform can be selected. Unlike proprietary solutions, Intel architecture's open standard allows the same operating system to run on any vendor-supplied Intel architecture-based system.

Intel's implementation of Linux on Intel Architecture has proven to be stable and robust, providing significant cost reduction and improved performance of many applications in the development environment. On a per-unit basis for the Intel-based solution, Intel saved 88 percent of the cost of RISC-based hardware, operating system, and data center space.

For details about deploying an open source operating system and porting the design project environment, read the white papers entitled, "Linux on Intel[®] Architecture Works for Intel's Design Environment" (<http://www.intel.com/eBusiness/pdf/it/wp014705.pdf>) and "Migrating from RISC to Intel[®] Architecture: Total Cost of Ownership Analysis" (<http://www.intel.com/ebusiness/pdf/it/wp032102.pdf>).

Approach to Migration

Given the industry growth of Linux in the enterprise and its price/performance advantages, the question now is how to deploy the Linux operating system on Intel architecture? In addition, how can a company mitigate the risks that are involved and keep costs down? It is important to have a process and a plan in place for handling such a task.

The migration can be planned into phases, with each phase addressing a different level, from the operating system and tool verification to the porting of the code and testing and deployment.

¹⁷ Gillen, Al, Dan Kusnetzky and Scott McLarmon. "The Role of Linux in Reducing the Cost of Enterprise Computing." IDC. November 2001. <<http://www.idc.com>>.

This section outlines best practices in evaluating and implementing such a migration, and will conclude with an example of how Intel® Solution Services helped IP Dynamics, a secure communications software company, migrate its application to a Linux and Intel architecture-based platform.

Assessment

The first important step in the migration process is to perform an assessment. In this phase, business requirements, objectives and major drivers for the migration are evaluated. It is also important to identify the key issues and potential risk areas and efforts should be made to ensure these are addressed.

A major part of the assessment phase is to conduct the TCO and ROI analysis, similar to what was described previously. As part of this analysis, the TCO components that may influence the decision to proceed with the migration should be prioritized. It is important to understand all of the system and software components and the maintenance and support costs involved with the source platform. Based on that information, an enterprise company can assess the components of the target platform and evaluate the costs of the migration and the ongoing support costs of the new platform. The TCO and ROI financial analysis is an important step in determining the potential cost savings of the new platform and if migrating makes business sense.

Another important aspect of the assessment phase is the technical feasibility study to determine what the dependencies may be in the environment and what potentially may hinder a successful migration. Third party dependencies, as well as custom application code and components, need to be identified. It is important to determine which of the third party dependencies are available on the target platform and to determine what functional gaps may exist in the third party dependencies between the source platform and the target platform. Then, it is important to determine how to migrate the third party applications and also overcome the gaps that may exist. In addition, it is necessary to evaluate the efforts that must be taken to port the custom application code and components. It is key to identify these things in the beginning of the project to determine the work effort, identify potential risks and estimate the amount of time needed for the migration. In the assessment phase, the technical feasibility of migration is determined before investing in resources to do the migration.

In addition to the TCO analysis and the technical feasibility study, another aspect of the assessment phase is to determine test and validation plans for the newly migrated system to ensure the system is performing to expectations. It is important to collect information on the systems operations and the applications, including the system configurations and usage models. In addition, information should also be collected related to the performance metrics and the performance characteristics of the systems, applications and typical workloads. The performance metrics will be an important component in verifying the functionality and performance of the newly migrated system.

The end goal of the assessment is to understand each phase of the migration effort before committing resources. The assessment helps companies comprehend the true cost of the migration, identify technical dependencies and mitigate risk. The typical duration of a migration assessment is two days to one week for an onsite visit to gather the necessary information, followed by two to three weeks, offsite to process the information.

Use of Benchmarks

One aspect of assessing the TCO of a new platform is the use of public benchmark data. Benchmark information can be used in comparing the price and performance of different systems and making a best estimate or generalization as to sizing and the number of machines needed in the new environment to support the capacity of the original systems. Typically, when migrating from an older, RISC-based architecture to Intel architecture, the mapping of systems' performance per CPU from the old environment

to the new environment is not necessarily one-to-one. Generally, it is possible to reduce the number of new systems needed to sustain equivalent performance provided by the old systems through server consolidation. Fewer machines can translate into less support and hosting costs.

In choosing a benchmark to use for this comparison, a benchmark that is similar in functionality or a best fit in type to the application to be migrated should be selected. For example, for a Web server, the SpecWeb* benchmark can be used. Similarly, for a J2EE* application server, the SpecJAppServer* benchmark may be used. The diagram below shows an example of the relative performance of an older Sun Enterprise* server compared to a newer Dell PowerEdge* server using the SpecWeb99* benchmark. The results show the number of conforming simultaneous connections, and a higher results number indicates better performance results.

SpecWeb99 ¹⁸					
Vendor	System	Application	Processors	Results	Processor Type
Dell	PowerEdge* 6600	Red Hat Content Accelerator* 2	4	6700	Intel® Xeon™ processor MP, 2GHz
Sun Microsystems	Enterprise* 420R with SNCA	Zeus* 3.3.8.4	4	1400	Sun UltraSPARC* II, 450 MHz

Suppose the source platform is the Sun Microsystems' Enterprise* 420R, and the potential target platform is a Dell PowerEdge* 6600. Given the results ratio, it can be estimated that one Dell machine would provide equivalent to 4.78 (or 6700 divided by 1400) Sun machines. This indicates there is a definite potential to reduce of the number of machines with the migration. This information can then be used as a starting point to estimate the number of target servers and to evaluate the TCO for the migration.

However, even if there is a potential reduction in the number of servers needed to support the same capacity as the original systems, it is important to account for room for growth so the systems do not become saturated too soon. The performance metrics and acceptable usage information are important aspects in assessing what systems to use and how many systems to have for the new target environment.

Although benchmarks are a useful reference point for a type of solution, they do not necessarily map one-to-one with real life solutions and business needs. The functionality and behavior of an application may be different than the benchmark and have differing results. Because applications differ, it is important to determine the application sizing for the new platform. This is where a proof-of-concept of the target platform would be useful to gather performance and sizing data. Intel Solution Services can help estimate the TCO of migrating solutions to Intel architecture-based servers, as well as execute a proof-of-concept engagement to gather real performance and TCO data.

Architecture and Design

After the financial and technical feasibility assessments have been completed, the architecture and design phase can commence. In this phase, the main objective is to validate the feasibility of the migration and to create a migration plan. In the architecture and design phase, an analysis is performed on the hardware, software, database and infrastructure. The applications are also considered, with thought given to better candidates for a Linux migration. As part of this phase, a proof-of-concept can be

¹⁸ Standard Performance Evaluation Corporation. <<http://www.specbench.org>>.

completed to demonstrate and verify the migration plan. This is also helpful in proving the ease with which a migration can be completed and identifying hidden dependencies.

The assessment phase will provide a list of the third party applications and custom applications to be migrated. In the architecture and design phase, a detailed plan is created for the code porting and testing. Gaps from the source platform to the target platform are also identified and addressed. As part of a proof-of-concept, it is important to deploy and verify the base operating system and software tools to be used in the new target environment, before attempting to port the application itself.

Sizing Servers

Another important aspect of the architecture and design phase is to understand what size servers are needed. Addressing the size and type of the machines an enterprise needs is a task that can be simplified through the following means. Understanding the financial and business goals along with the current performance will help determine the server utilization requirements, the availability and size of the solution, and the extensibility for future growth.

When approaching a large migration to Intel architecture, companies typically choose to optimize those areas of business that are relatively modular and are large enough to realize a significant TCO reduction. That is to say, the best success comes from migrating less complex solutions that are common throughout the enterprise.

One of the key areas where financial and business goals guide the requirements of a solution is in availability. Depending on the volume of transactions and critical nature of a solution, even momentary downtime can impact business significantly. In order to achieve high availability with clustered software, most enterprises choose to run their applications on multiple systems. Certainly, there is no need to consolidate an application that runs well clustered with failover to a single server. This is where sizing and utilization of the solution comes into consideration.

Utilization

A major key to sizing is smart utilization and using resources effectively. If servers are underutilized, it is a waste of hardware, operational support resources, and ultimately, money. Overburdened servers can increase support costs and be a management nightmare, especially as a business grows. Utilization can be considered in a variety of ways based upon the type of application, the percent rule of availability and license utilization. First, it is important to consider the different types of applications and how they apply to utilization.

In the case of non-clustered software, there are two possibilities: stateless applications and monolithic applications. Stateless applications, such as those that typically run on Web servers, allow for virtually unlimited system scale out. These applications do not require communication between users, instances or servers, and therefore do not suffer from diminishing returns as servers are added. Choosing the right server for these applications has more to do with density and management considerations. The second possibility is with monolithic applications, those that can run only on one system. In this case, server sizing with respect to utilization implies smaller applications require smaller servers and larger applications require larger applications. Having numerous smaller servers may make management of the machines more complex. On the other hand, it may be possible to consolidate the smaller monolithic applications onto larger servers using software that creates multiple virtual machines and allows for multiple separate operating system environments. This reduces the number of physical servers that need to be managed.

In the case of clustered software, there are fundamental properties that help guide the system choice. For instance, if a solution is clustered on two servers and no degradation in service is a requirement, then each server would need to run at 50 percent (or less) utilization. As the number of servers in that cluster

increases, the maximum allowable utilization increases per server. This equates to 66 percent for three servers, 75 percent for four servers and so on. Because of this rule, it can easily be deduced that if an application scales well to three or more servers, equipment will be better utilized. There is a point of diminishing returns for clustered software because of the communication between the nodes. The more servers or nodes that are added in a clustered environment for applications, the more time the software spends communicating between those nodes and less time doing real work. In this case, network constraints may be an issue.

One key area of utilization often ignored is license utilization. Many enterprise software products have significant licensing costs. These costs can be per server, per user and per processor. One key opportunity for reducing TCO is in the area of per processor licensing. Many large application server and database vendors use this as a means of charging customers. The best ROI with this kind of software can be obtained by running it on the fastest processors available, thereby reducing the overall number of licenses required. This is especially apparent with software that has a very high cost relative to the system cost.

Solution Migration

The third step in the migration process is the solution migration phase. In this phase, the target infrastructure is built and software applications are ported to the new environment. This entails setting up and configuring the hardware environment and the software environment with all of its components. This includes the operating system, application packages, setup scripts and third party software. The custom software is ported and verified, and the data is also migrated. An important aspect of this phase is the testing and verification of the components, the ported applications and the code. It is important to conduct functional testing on the components, as well as end-to-end testing. The tests need to encompass the various usage models and meet the performance criteria specified in the earlier phases. The aspects of the migration need to be thoroughly documented, and operational information needs to be turned over to the proper parties.

Deployment Considerations and Strategies

For a multi-tiered, multi-server complex environment, it is important to determine the deployment strategies of the operating system and software packages for the various tiers of the infrastructure. Migration to a new platform provides opportunities to standardize the configurations to be deployed. It is important to consider configuration deployment management when dealing with a larger infrastructure of many servers. Typically, different usage groups of servers and software images need to be identified and evaluated before deployment. In this situation, automation and ease of management of the software become very important.

Deployment and management tools of a wide variety are available for the Linux platform. Red Hat Linux provides a technology called Kickstart*, in which the operating system and configuration specifications can be distributed to systems, installed in a short amount of time and be fully automated. In addition, Linux software packages are distributed in Red Hat Package Manager (RPM) packages that allow for easy installation and management. Examples of tools for software distribution and management are listed in Appendix B of this paper.

Optimization and Tuning

After the applications have been migrated to the new platform and tested and verified for correct functionality, the applications and systems need to be analyzed for optimization and tuning. Potentially, the configuration settings of the applications on a UNIX environment may not apply to a Linux and Intel architecture-based environment. The Linux and Intel architecture-based platform provides different opportunities for performance optimizations that differ from the RISC/UNIX environment because of the specific features of Intel architecture. Therefore, it is important to tune and optimize the systems to realize the full performance potential of the newly migrated systems.

Optimization and performance tuning requires analysis of bottlenecks in the applications and in the system. The system configurations can be tuned based on their functional usage. Application code can also be analyzed for potential optimization opportunities. In addition, Intel provides many tools and libraries that can significantly aid in increasing the performance of applications (see Appendix A).

Case Study

Applying the Methodology to IP Dynamics¹⁹

IP Dynamics, a California-based secure communications company, provides virtual private network (VPN) software with an emphasis on easy deployment and administration. IP Dynamics' Virtual Community Network (VCN) Manager* application uses domain name-style addressing, allowing network administrators and non-technical users to set up VPNs, easily assign access rights and security policies, and change user settings in minutes. A large telecom service provider and key customer of IP Dynamics needed an Intel architecture-based version of the VCN Manager application. The telecom service provider wanted a Linux and Intel architecture-based platform to save money and to leverage the existing skill set of its IT staff. With limited resources to port the RISC-based VCN application and database to Intel architecture, IP Dynamics needed outside resources to complete the job in a short timeframe, and they turned to Intel Solution Services for assistance.

TCO Savings

For IP Dynamics, having an Intel architecture-based version of its product would dramatically expand its available market by enabling its customers to lower their TCO. According to Huan Wang, director of IP Dynamics' product line management, the Intel architecture-based version of the VCN Manager requires one-third less startup costs and reduces operating costs compared to the RISC-based platform.

¹⁹ Intel® Solution Services' Final Report v1.0. 12 September 2002. Intel® Solution Services' ROI Analysis. 29 March 2003.




Figure 8 - Estimated Startup Cost Savings with VCN Manager Running on Intel Architecture (typical system supporting 30,000 to 40,000 users)²⁰

In the technical assessment and planning, IP Dynamics' source code was scanned using a porting assessment tool. This tool examined the source code for potential pitfalls in porting and highlighted potential issues, such as macros, data types and functions that were supported on the RISC-based platform, but not on Linux. By uncovering these issues, Intel Solution Services was able to address the potential problems or dependency risks at an early stage in the migration. It also aided the team in assessing the difficulty of the migration and in estimating the amount of time needed for the migration.

Architecture and Design Plan

The VCN application included multiple tiers and technologies, which made the migration complex. The specific goals of the project were to:

1. Migrate IP Dynamics' VCN suite from the RISC-based server to Red Hat Linux 8.0 running on Intel architecture-based servers.
2. Migrate the VCN database from the existing TimesTen* database to an Oracle9*i* database on Linux.
3. Create a common source-code tree that could be compiled for either the RISC-based or Intel architecture-based platform.

Application Migration and Regression Testing

During the migration phase, the team focused on resolving issues that surfaced during the assessment and design phases, which included missing header files and different functional data types, declarations and arguments. Resolving these issues was a joint effort between the Intel Solution Services' consultants and the IP Dynamics' developers. In addition, several third-party components of the VCN suite had to be analyzed to determine whether or not they supported Linux. The TimesTen database that IP Dynamics had been using did not support Linux, so the database was migrated to the Oracle9*i* database instead. Intel Solution Services' consultants and IP Dynamics' developers rewrote the VCN Java* and C++ code to support the Oracle database, which included modifying the SQL* statements and connection information for new schema and updating the APIs.

²⁰ Intel® Solution Services' Final Report v1.0. 12 September 2002. Intel® Solution Services' ROI Analysis. 29 March 2003.

Functionality testing was performed for three weeks in the IP Dynamics' quality assurance lab as a part of the migration process. Intel Solution Services employed a rigorous methodology to comprehensively test the systems. After the VCN suite passed the functionality tests, the application suite was stress tested and load tested using a load simulator.

Performance Tuning of VCN Suite

After the Linux version of the VCN suite was tested for functionality and load, the team tuned and optimized the application. The tuning and optimization phase focused on characterizing the performance of the VCN software and finding "hot spots" in the code using the Intel® VTune™ Performance Analyzer tool. Intel provided the hardware, tools and technical support to have the VCN Manager Suite tested on three Intel® Xeon™ processor-based servers, using 2-way, 4-way and 8-way configurations.

Results Delivered

In a three-month timeframe, the Intel Solution Services consultants successfully migrated the VCN application and its database from RISC-based servers to Linux running on Intel architecture-based servers, thereby accelerating the product's time-to-market. Huan Wang says, "Intel cut our development time in half." The migrated application not only met the immediate customer's needs, but also provided a platform to expand the market served by IP Dynamics, which now includes a large number of companies using Intel architecture-based systems.

"If we offer a more affordable product, it makes us more competitive," Wang explains. "But TCO also is an attraction for our customers. The Intel architecture-based version of the VCN Manager requires one-third less startup costs and reduces operating costs compared to the RISC-based platform."

Conclusion

The business challenges of today are unavoidable realities that demand thought and response. More and more businesses are realizing the dangers of being locked into a proprietary, RISC-based architecture. Linux, as an operating system, was initially developed and optimized for Intel architecture, so the combination of the two makes an easy stepping-stone for moving away from a RISC/UNIX-based architecture.

The Linux and Intel architecture-based platform provides reduced TCO for businesses in a variety of industry verticals, and the price performance benefits of the platform continue to grow. Linux has also matured in functionality, support and acceptance, such that it is now considered to be enterprise-ready by the industry. More and more companies are certifying their applications on the Linux and Intel architecture-based platform. Linux vendors have matured in their support of the operating system to provide versions that ensure the reliability, availability, scalability and manageability of the system.

Intel follows a proven migration methodology to help businesses capitalize on the value of a Linux and Intel architecture-based platform. Each phase of the methodology – assessment, design, implementation and optimization – helps mitigate risk and allows for a successful deployment of Linux on Intel architecture.

In light of the maturity of Linux in the enterprise and the price/performance advantages of the Linux and Intel architecture-based platform, businesses today should consider this new platform as a means of reducing TCO and maximizing ROI of their IT environments.

Appendix A

Linux Development Tools from Intel

Intel® Fortran and C/C++ Compiler for Linux

Intel has applied its expertise in software optimization for Intel® 32-bit and Intel® Itanium® processor-based systems to develop the next release of the Intel® Fortran and C/C++ Compiler for Linux. The Intel® Compiler delivers optimal performance on Intel 32-bit processors, including Intel® Pentium® 4 processors and Intel® Xeon™ processors, as well as the 64-bit Intel Itanium processor and the new Intel® Itanium® 2 processor. It offers improvements to its advanced optimization features, such as vectorization and Streaming SIMD Extensions 2 (SSE2) for Intel 32-bit processors, software pipelining for the Itanium processor family, and Interprocedural Optimization and Profile Guided Optimization. The Intel Compiler now offers specific optimization options – called tpp1 and tpp2 for Itanium and Itanium 2 processors, respectively – to help companies take advantage of the latest Intel 64-bit architecture.

The Intel® C++ Compiler for Linux has substantially improved its source and binary compatibility with GNU C/C++, providing stronger C++ ABI conformance, wider GNU Compiler Collection extensions support and the ability to build the Linux kernel with fewer modifications. The Intel Compiler has added extended language directives for software pipelining, loop unrolling and data prefetching that provides information to the compiler to enhance optimization of the application code. It also provides development and optimization capability for threaded applications through its new support of the OpenMP* 2.0 standard for C/C++ and the enhanced Auto-Parallization feature. For more information on the Intel® C/C++ Compiler for Linux, visit: <<http://www.intel.com/software/products/compilers/clin/whatsnew.htm>>.

According to Dr. Fons Rademakers, a senior scientist with CERN in Geneva, "The Intel compilers have performed excellently on our ROOT code. On average, the Intel C++ Compiler for Linux produces executables that run 30 percent faster than ones produced by gcc 3.2. Due to the excellent compatibility with the GNU compilers, the porting effort was reduced to a minimum. Furthermore, the strong technical support that is provided with the Intel Compilers enables very quick turnaround on fixes for customer issues."²¹

Intel® VTune™ Performance Analyzer for Linux

The Intel® VTune™ Performance Analyzer 1.0 for Linux provides a fully native Linux solution that allows higher levels of performance for software running on the latest Intel 32-bit processors, including Intel® Xeon™ processors and Intel® Pentium® 4 processors. This new product provides a command-line capability that allows performance data to be collected, analyzed and displayed for 32-bit Linux applications, kernels and drivers.

Intel® VTune™ Enterprise Analyzer

The Intel® VTune™ Enterprise Analyzer for Web applications, Java edition, is an enterprise Web performance analysis tool designed specifically for improving the performance of J2EE* applications. The VTune Enterprise Analyzer helps programmers gain insight into back-end Web infrastructure hotspots and bottlenecks as they develop distributed applications during stress testing. The tool also provides detailed information on inter-tier and object activities of Java-based Web sites.

²¹ Intel® Software Products – Customer Comments.
<<http://www.intel.com/software/products/global/custcomments.htm>>.

Appendix B

Linux Management Tools

Below are some examples of Linux management tools. This should not be considered an endorsement by Intel, but a short list provided to illustrate the variety of management tools available for the Linux platform.

Red Hat Command Center*

Red Hat is now offering its own system management software suite, called Red Hat Command Center*. Red Hat Command Center provides integrated application, network, system and transaction monitoring, ensuring the availability and performance of critical Web applications and infrastructures. Red Hat Command Center can be deployed in a matter of days to provide customers with a complete, browser-based view of their distributed computing infrastructure. For more information, visit: <http://www.redhat.com/rhn/commandcenter/>.

Ximian

Ximian Inc., the leading open source desktop company, announced the availability of Red Carpet Enterprise*, a powerful server-based solution that provides customers with centralized software maintenance and version management for servers and workstations running a broad range of leading Linux distributions. Deployed entirely behind the corporate firewall, Red Carpet Enterprise lets organizations automatically update groups of systems on demand or on preset schedules with custom company software standards, as well as software from Ximian, Linux distribution providers and third-party software vendors. As a result, companies can dramatically reduce the cost and complexity of updating their internal Linux servers and workstations with the security and performance of a behind-the-firewall solution.

BMC

BMC software and Red Hat are cooperating to provide end-to-end enterprise management for Linux customers.

Heorix eQ Management Suite* - Multiple Application Management Software

Heorix eQ Management Suite* - Multiple Application Management Software unifies management of the multiplatform enterprise by detecting and resolving complex application, system and IT infrastructure problems. The software suite supports UNIX, Linux, Windows* NT/2000/.NET*, Novell NetWare* and OpenVMS* platforms. Monitoring and management coverage includes packaged and custom applications, database systems, messaging platforms, Web servers, operating systems and infrastructure components.

Linux Image Management Tools

Platespin

Platespin provides remote deployment of ESX, which is fully integrated into product. Drag and drop deployment of ready-to-run VMware* operating environment images (OS, applications, etc.) accelerates uptake of VMware technology into production. It enables moves, back-up, ability to configure, re-configure and rapidly recover VM guest environments. It provides remote purposing and rapid deployment of servers, enables simultaneous and hands-free deployment, and reduces attended operator time by over 90 percent (drag & drop activation reduces operator intervention to two to three steps). Actual deployment time can be reduced by approximately 60 percent.

Open Country

Open Country enables nodes (connected devices), users, applications and content to all work harmoniously in a normalized environment. Software is simply and intelligently hosted, distributed, installed, upgraded and repaired. Nodes provision and configure themselves, and applications and content can be automatically upgraded, making those using or managing connected PCs, workstations, servers and embedded systems ideally suited to benefit from Open Country's solutions.

OC-Manager* is a complete node and application management solution, using Open Country's unique concept of normalization. Linux applications today need to be re-packaged to run on each separate flavor of Linux, and Open Country's normalized environment refreshingly enables an application to install on the Linux distributions of all major vendors, and to be launched and managed through a consistent user interface, regardless of distributions or desktop. Similarly, in an environment where deploying and managing a network of Linux systems is cumbersome and cryptic, Open Country's OC-Manager enables point and click deployment, management, duplication and recovery of Linux systems.

Hewlett-Packard Tools

Because manageability is critical to the adoption of Linux in IT environments, Hewlett-Packard (HP) has expanded its Linux management portfolio by adding administrative tools and a partner program. The tools are designed to ease the manageability of IT environments for a range of customer needs: from using Linux as Web servers or for application development to running vital applications on Linux where high availability is necessary. HP's Linux management partner program for independent software vendors (ISVs) uses HP Servicecontrol Manager* software as its manageability framework. Products from Linux software vendors, such as Aduva, Integrated Research Ltd., Network Shell, Symark Software and the TOLIS Group are integrated into HP Servicecontrol Manager software. Other Linux manageability and high availability products from HP include HP Tootools* 5.5, HP Servicecontrol Manager* 2.5, HP Process Resource Manager* and HP MC/Serviceguard* software.

IBM Tools

IBM Director* v4.1, available later this year, is the newest release of the industry leading client/server workgroup manager. The IBM Director tool provides customers with flexible capabilities to realize maximum system availability and lower IT costs. With IBM Director, IT administrators can view and track the hardware configuration of remote systems in detail and monitor the usage and performance of critical components, such as processors, disks and memory.

In addition to making available a new version of IBM Director, enhanced extensions to IBM Director will also be available for IBM customers who want advanced capabilities. IBM Director Server Plus Pack consists of five additional tools for advanced server managements to help optimize performance and maximize availability. These tools are tightly integrated with IBM Director for consistent management from a single console.

Linux Resource Management Tools

Aurema ARMTech*, Active Workload Manager

ARMTech* for Linux, Aurema's enterprise-class active resource management utility, provides a comprehensive solution for the allocation and control of major system resources for enterprise users deploying Linux on Intel 32-bit and Itanium® processor-based servers. The robust system management capabilities of ARMTech for Linux gives users an unprecedented level of control over critical system resources, enabling increased server utilization without affecting performance, as well as offering precise resource consumption accounting for billing and capacity planning. Aurema's Active Resource Management technology extends and simplifies the dynamic management and control of the resources required to run mission-critical applications at peak levels on Linux-based servers, providing results previously unavailable from traditional passive resource monitoring and workload management technologies.

While other system management tools offer the means to monitor system resources and workload usage, only Aurema's Application Workload Manager extends the capabilities of IBM Director by allowing control of how applications use valuable server resources. In consolidated server scenarios, Application Workload Manager dramatically increases the utilization of a machine's CPU, allowing multiple mission-critical applications to run effectively on a single IBM xSeries server, providing the predictability and performance required for successful server consolidation.

SteelEye's LifeKeeper*

SteelEye's LifeKeeper* is designed for continuous availability to meet stringent enterprise requirements. SteelEye's LifeKeeper for Linux high availability clustering is a software application that helps ensure the continuous availability of applications and maintains system and storage uptime. LifeKeeper helps maintain the high availability of clustered Linux systems by monitoring system and application health, working to maintain client connectivity and providing uninterrupted data access regardless of where clients reside — on the corporate Internet, intranet or extranet.

Linux Security and Performance Monitoring Tools

Hewlett-Packard

Hewlett-Packard has expanded its Linux-based offerings for the security and intelligent device markets, as well as a software developer's portal and a variety of other software tools to provide customers with Linux solutions. The range of offerings — including one of the industry's first security products for Linux-based systems — provides prevention, containment and detection with multiple layers of security. The software protects crucial server components, including the operating system and application layer, and it helps prevent unauthorized communication between programs, networks and files. To learn more about Hewlett-Packard's security tools, visit: <http://www.hp.com/linux>.

Tivoli

Tivoli has always embraced an open management philosophy by offering solutions that integrate across heterogeneous environments. Whether Linux is being integrated into an existing environment or an infrastructure has been built solely on Linux, Tivoli offers e-Business infrastructure management expertise. Tivoli systems management software equips companies to handle all aspects of systems management:

security, performance, availability, configuration and systems operation. With solutions built to be flexible, reliable and scalable, Tivoli can help keep e-Business applications and resources available and secure across the entire heterogeneous infrastructure.

Appendix C

Red Hat Enterprise Linux Support of RASM

In order to address the needs for an enterprise environment, Red Hat Linux Advanced Server* for the enterprise environment addresses RASM with the following features:

Reliability

Red Hat engineering has worked with OEMs and ISVs to ensure Red Hat Advanced Server meets reliability standards by subjecting it to stringent qualification and testing over an extended period of time. Red Hat has worked with major application vendors, including Oracle, IBM, SAP, VERITAS, Computer Associates and others, to ensure its products work seamlessly with Red Hat Enterprise Linux solutions. Acting on feedback from both enterprise customers and key OEM and ISV alliances, Red Hat has extended the life cycle of Red Hat Enterprise Linux releases; new versions of Red Hat Enterprise Linux products will be released on a 12- to 18-month schedule. Unlike consumer versions of Linux, the Red Hat Enterprise Linux release cycle ensures that applications certified on one release will typically not require any modifications or testing cycles before being certified on new releases. Also, Red Hat Enterprise Linux products are only available for deployment as a support and maintenance subscription, which includes access to Red Hat Network for system maintenance and management services.

Availability

From a user's viewpoint, the term *availability* is easily defined: the user's applications and data should be continuously accessible. This goal is reached through a hardware and software configuration that assures that in the event of the failure of any single hardware component, the overall system continues to operate. This is usually achieved by building in redundant components so there is no single point of failure, either through load balancing, clustering or a combination of the two. Red Hat Enterprise Linux Advanced Server provides two integrated high availability technologies. The first, IP Load Balancing, provides network load balancing for environments such as Web server farms. IP Load Balancing will distribute incoming network requests across a group of servers, which then service the request. Load balancing yields improved performance, and if a single server fails, incoming requests will continue to be redistributed across the remaining servers. From the remote users' viewpoint, the network requests can be serviced continuously, even in the face of a server failure.

The second technology, Red Hat Cluster Manager, provides high availability by using a technology widely used by other operating systems – application *failover*. Cluster Manager has been specifically developed for use with today's standard hardware products and does not require expensive, special-purpose hardware. A variety of application types can be clustered: generic, unmodified applications and custom, in-house applications; databases; file and server environments; enterprise commercial applications; Internet and open source applications; as well as messaging applications.

Scalability

Red Hat Enterprise Linux products are delivered with kernels that are optimized for SMP systems (up to 8-way). Red Hat has provided several significant kernel enhancements in Red Hat Enterprise Linux products to ensure the application can scale linearly to maximize Intel processor throughput and processing power.

Some features of Red Hat Linux designed to improve scalability for enterprise deployments include:

- Asynchronous I/O support, which allows a process to issue an I/O and immediately continue processing;

- Process scheduling, which allows data and instructions held in the CPU cache to be available for the process each time it executes, thus greatly improving performance; and
- Bounce buffer elimination, which drastically reduces buffer copying while continuing to allow device drivers to access the contents of an I/O buffer regardless of its location in memory.

Manageability

The manageability of a system can have a large effect on its cost-effectiveness and long-term TCO. Complex, error-prone or time-consuming administration tools can lead to extended downtimes and dissatisfied users. Consequently, all Red Hat Enterprise Linux products provide a comprehensive suite of management features, taking advantage of the very strong Linux command line environment and the many graphical user interface (GUI)-based administration utilities for use with GNOME and the K desktop environment (KDE).

Two manageability features in Red Hat Linux Advanced Server include a network console, which allows all kernel messages to be logged by another system on the network, and the Netcrash dump feature, which allows a Red Hat Enterprise Linux system to transmit a complete crash dump across the network to a sink node, making it possible to more quickly recover a crashed system to normal operation. The Linux and Intel architecture-based platform has been gaining credibility in production environments that require 24 X 7 X 365 availability and have stringent performance demands.

Legal Disclaimers

Information in this document is provided in connection with Intel® products. Except as provided in Intel's terms and conditions of sale for such products, INTEL ASSUMES NO LIABILITY WHATSOEVER AND INTEL DISCLAIMS ANY EXPRESS OR IMPLIED WARRANTY RELATING TO SALE AND/OR USE OF INTEL PRODUCTS, INCLUDING LIABILITY OR WARRANTIES RELATING TO FITNESS FOR A PARTICULAR PURPOSE, MERCHANTABILITY, OR INFRINGEMENT OF ANY PATENT, COPYRIGHT, OR OTHER INTELLECTUAL PROPERTY RIGHT. Intel products are not intended for use in medical, life-saving, life-sustaining, critical control or safety systems, or in nuclear facility applications.

Intel may make changes to specifications, product descriptions, and plans at any time, without notice.

Customer savings and results are specific to the client, their business goals and the configuration of their enterprise environment. Differences in business goals, software design, infrastructure or configuration may affect actual results.

Performance tests and ratings are measured using specific computer systems and/or components and reflect the approximate performance of Intel products as measured by those tests. Any difference in system hardware or software design or configuration may affect actual performance. Buyers should consult other sources of information to evaluate the performance of systems or components they are considering purchasing. For more information on performance tests and on the performance on Intel products, reference www.intel.com/performance/resources/limits.htm, or call (U.S.) 1-800-628-8686 or 1-916-356-3104.

Intel, the Intel logo, Intel Xeon, Itanium, Pentium and VTune are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

*Other names and brands may be claimed as the property of others.

Copyright © 2003, Intel Corporation. All rights reserved.

1103/MH

About Intel[®] Solution Services

Intel Corporation
2200 Mission College Blvd.
Santa Clara, CA 95054-1549

Intel[®] Solution Services is Intel Corporation's worldwide professional services organization, helping enterprise companies capitalize on the full value of Intel[®] architecture through consulting focused on architecture transitions.

Backed by the largest silicon manufacturing company and one of the largest e-Business corporations in the world, Intel Solution Services uses its foremost expertise in Intel architecture and next-generation technologies, as well as its relationships with key industry alliances, to design cost-effective, leading-edge solutions that help deliver superior business results.

Services are available through on-site consulting, as well as at Intel[®] Solution Centers located worldwide. The Intel Solution Centers are state-of-the-art environments for designing and testing high-performance solutions using Intel's best-known methods and technologies.

For more information, visit Intel Solution Services' Web site at: www.intel.com/internetservices/intelsolutionservices, or send an email to: solution-services-questions@intel.com.