IT@Intel White Paper
Intel IT
IT Best Practices
Mission-Critical Business Solutions
May 2011

Migrating Mission-Critical Environments to Intel® Architecture

By using a phased approach and implementing best practices derived from our experiences, we minimized or eliminated most risks and successfully completed each migration with little or no interruption to core business activities.

David Aires
General Manager, Factory Automation

Elwood Coslett
Data Center Engineering Director, Intel IT

Gerald T. Seaman
Senior Integration ERP Architect, Intel IT

Executive Overview

As part of a multiyear strategy, in 1992 Intel IT began migrating Intel's most mission-critical computing environments from proprietary systems to Intel® architecture. These environments include silicon design, manufacturing, and global enterprise resource planning (ERP).

Intel's business results rely on these environments, so we approached the migrations with caution. Ultimately, however, the migrations were less difficult—and the benefits greater—than we originally anticipated. By using a phased approach and implementing best practices derived from our experiences, we minimized or eliminated most risks and successfully completed each migration with little or no interruption to core business activities.

Success depended upon breaking the migration for each environment into manageable phases, timing migration to coincide with other upgrades to business-critical systems, and developing a center of technical expertise to provide leadership and guidance in defining and executing the overall strategy.

Although each environment required a different strategy to meet specific business requirements, applying best practices enabled us to deploy the new Intel®-based platforms with confidence, while achieving a variety of benefits and business value:

- Improved system performance to increase productivity and support faster time to market for Intel® products.
- Increased return on investment and decreased operational costs due to industry-standard servers and clients.
- Ability to choose the best hardware for the workload because Intel's open standard enables the same OS on any Intel-based platform.
- Decreased data center footprint due to space- and energy-efficient servers.
- Savings of more than USD 1.4 billion over five years.[1]

Our successful migration of mission-critical environments to Intel architecture demonstrates that this platform provides the necessary reliability, availability, and scalability, and that migration can be accomplished with minimal downtime and disruption to key business functions.

Contents

IT@INTEL

The IT@Intel program connects IT professionals around the world with their peers inside our organization – sharing lessons learned, methods and strategies. Our goal is simple: Share Intel IT best practices that create business value and make IT a competitive advantage. Visit us today at www.intel.com/IT or contact your local Intel representative if you'd like to learn more.

BACKGROUND

When Intel IT chooses a hardware platform, the choice is based on business need, such as computing performance, reliability, and scalability. For example, we originally chose proprietary RISC-based platforms for our silicon design and manufacturing environments because RISC was the only computing architecture available that could provide the performance and stability required by our aggressive product development cycles.

However, by the early 1990s, Intel®-based systems could support enterprise workloads and had reached performance parity with proprietary hardware. Intel® architecture therefore became a platform option for our mission-critical environments, including silicon design, manufacturing, and enterprise resource planning (ERP). Using Intel-based systems would also enable us to also take advantage of industry-standard and open-source OSs such as Microsoft Windows* and Linux*.

We determined that running our mission-critical environments on Intel architecture could also provide two compelling business benefits:

- **Improved performance to support faster time to market.** Intel's mission-critical environments require consistent improvements in processor performance to keep pace with product design cycles and the increasing complexity of Intel® products. The improvement cadence of proprietary hardware has been erratic and slow compared to that of Intel architecture—which delivers consistent performance improvements every 12 to 18 months.
- **Reduced costs for platforms and software.** Due to the cost and performance advantages of Intel-based systems, migrating from proprietary hardware could significantly reduce our total cost of ownership (TCO) and increase productivity. In addition, using Microsoft Windows and Linux would support our enterprise strategy of using standard building blocks to maximize the value and efficiency of existing tools and support staff.

SOLUTION

Because migrating mission-critical environments such as silicon design, manufacturing, and ERP is a significant endeavor, we developed a phased, multiyear strategy to transition to Intel architecture. As we made progress toward our goal, we formalized best practices that enabled us to maximize the business value of the migrations—particularly by minimizing downtime.

We began replacing proprietary mainframes in our ERP environment in 1992 and eliminating RISC-based servers in our silicon design environment in 1998. By the end of 2012, we will have replaced almost all RISC-based servers with Intel-based servers in our mainstream manufacturing environment.[2]

We used slightly different strategies for migrating the different environments, each optimized to address specific business needs. Each migration to Intel architecture has resulted in significant cost reductions and increased compute performance.

Table 1 summarizes our work. We are currently finalizing migration in our manufacturing environment.

Silicon Design Environment

Silicon design engineering teams need considerable computing power to create each new generation of Intel® processors. They must have rapid and reliable access to application binaries and project data. Because downtime and sluggish performance diminish their productivity, they need stable and highly available systems that scale to meet growing compute demands.

In 1998, the silicon design environment encompassed a community of 16,000 engineers at more than 45 design sites worldwide. At that time, Intel's silicon design environment faced several challenges:

- Accelerating hardware and software complexity necessitated a 10-fold increase in computing power for the development of each new generation of Intel architecture.

Table 1. Migrating Mission-Critical Environments to Intel® Architecture

Environment	Reasons for Migration	Pre-Migration Platform	Post-Migration Platform	Time Frame
Silicon Design	• Compute power requirements growing exponentially • RISC-based platforms becoming prohibitively expensive	• Proprietary version of UNIX* running on RISC-based servers	• Linux* running on a mix of Intel® Itanium® processor-based servers and Intel® Xeon® processor-based servers	1998 to 2004
Manufacturing	• Supplier ending support for existing hardware • Required faster transaction speeds	• Open VMS* and UNIX running on RISC-based servers	• Open VMS or UNIX running on Intel® Itanium® processor-based servers • Microsoft Windows* running on Intel® Xeon® processor-based servers	Wafer fabrication factories: 2004 to 2008 Assembly and test factories: 2010 to 2012
Enterprise Resource Planning (ERP)	• Supplier no longer supported existing infrastructure • Required larger amounts of in-memory processing and parallel processing to speed transactions	• Proprietary OS running on mainframes	• Proprietary OS running on Intel® Pentium® processor-based servers • Microsoft Windows* running on high-performance Intel® processor-based servers	Stage 1: 1992 to 1996 Stage 2: 2004 to 2008

- As the number of simultaneous designs increased, product development timelines shrank.
- The need to invest in improving design team productivity conflicted with the mandate to contain costs.

OPEN-SOURCE OS AND PLATFORM FOR SILICON DESIGN

After thoroughly examining compute requirements for the silicon design environment, we decided to investigate using Linux* on Intel architecture.

We chose Linux because it was developed and optimized for Intel architecture and because it would be compatible with legacy UNIX* OSs and existing tools. We also hoped to boost industry momentum—specifically, the use of Linux in electronic design automation (EDA) environments—by working closely with suppliers to identify gaps in required applications and to improve software availability and compatibility.

Unlike proprietary RISC-based platforms, this open standard would enable us to run the same OS on Intel architecture-based platforms obtained from any supplier. This allowed us to select the best hardware over the long term. The price/performance benefits of using this platform were significant and provided a number of opportunities to reduce overall costs while improving the performance of many applications due to the speed of the Intel processors.

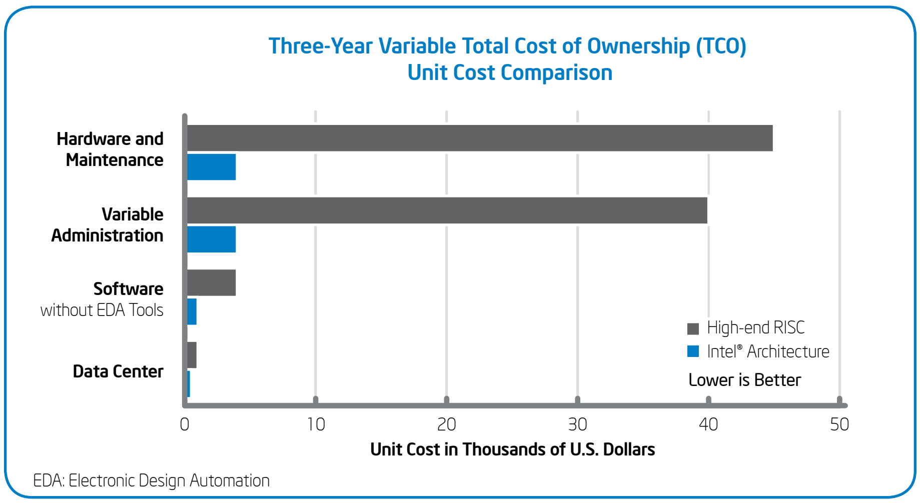

Figure 1. When we compared total cost of ownership (TCO) for systems based on Intel® architecture and RISC, we found that Intel architecture could provide savings in the hundreds of millions of dollars over a three-year period.

PROOF OF CONCEPT AND TCO STUDY

To demonstrate the benefits of running Linux on Intel architecture, we conducted a proof of concept (PoC) in 1998, which took about 50 days to complete. During the PoC, we successfully ported a CPU-intensive application, batch tools, and public domain scripts and tools to Linux; integrated Linux into the silicon design environment; and tested all necessary scripts. Linux on Intel architecture provided design teams with the required capacity, and teams also found that their tools ran faster on Intel-based platforms than on RISC-based platforms.

We designed a TCO study to compare the price/performance value of Intel architecture and RISC-based platforms for EDA. We discovered that the Intel-based platform running Linux provided a 90-percent performance advantage over the fastest RISC-based platform we tested. These results allowed us to reduce the number of servers we purchased for silicon design engineering by 47 percent.

TCO for Intel-based systems was about six percent of a comparable RISC-based installation, which provided hundreds of millions of dollars in savings over a three-year period. On a per-unit basis, we saved 88 percent of the cost of RISC hardware, OS, and data center space. These results are illustrated in Figure 1.[3]

As a result of this analysis, we were compelled to move our design teams to Intel architecture.

THREE-PHASE MIGRATION PLAN FOR SILICON DESIGN

We used a three-phase approach to migration in our silicon design environment, focusing first on the highest-priority applications and design steps. In defining and prioritizing each phase, we considered those tools that were most critical to design engineers, those that were easiest to migrate, and those that offered the best return on investment (ROI).

Phase 1. As shown in the high-level flowchart of our silicon design process in Figure 2, our highest priority was the architecture and logic design steps of product development. These steps represent 70 percent of our computational need and are batch-oriented for the most part.

Phase 2. We migrated the full chip validation and tapeout steps. Although these activities do not involve many applications, the applications are compute- and memory-intensive, and Intel architecture could support these demands.

Phase 3. We migrated the circuit design and physical layout steps. Applications in these steps use both interactive and batch processing, and are highly graphical.

Linux Implementation

To roll out the new OS on a global scale, we used several methods that contributed to a smooth migration:

- Small team of experts. We formed a team responsible for managing the OS distribution, including developing the OS roadmap, patching, hardware compatibility testing, and the overall release methodology.
- Pilot users. At the time we started the migrations, roughly 20,000 users needed to move onto Intel-based platforms. To make this process easier and to overcome natural resistance to change, we enlisted small groups of volunteers to run workloads on the new platform and provide feedback about their experiences. We used this feedback to improve the overall user experience.
- Intel Architecture Lab Partnership. We worked with various teams to copy their applications and test cases onto systems in our Intel Architecture Lab. We learned how to run and benchmark these applications in the lab, and then provided results to users. If users found system or application anomalies, we addressed them so that they were not faced with these errors when they went to run the tests in their local environments. The Intel Architecture Lab became a strategic asset. Using hundreds of hours of workloads, we can determine how a new platform will perform prior to its release into the production environment, relative to the previous generation platform.
- Partnering with users. Soliciting user feedback resolved a number of issues and also helped to build a level of trust and credibility between the migration team and users. Users also appreciated that IT was responsible for some migration overhead.

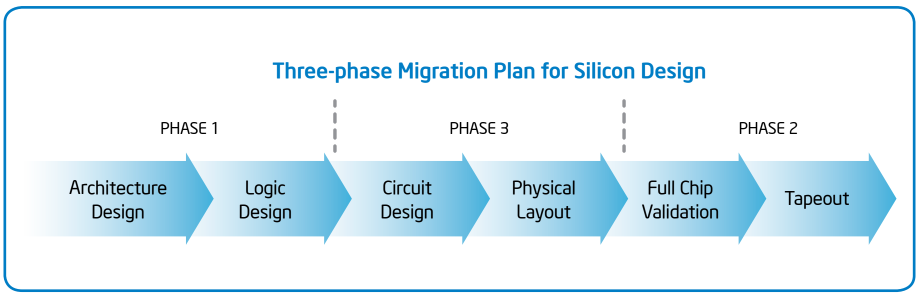

Figure 2. We prioritized our silicon design environment migration, focusing first on those applications that represented the highest need and therefore the highest return on investment.

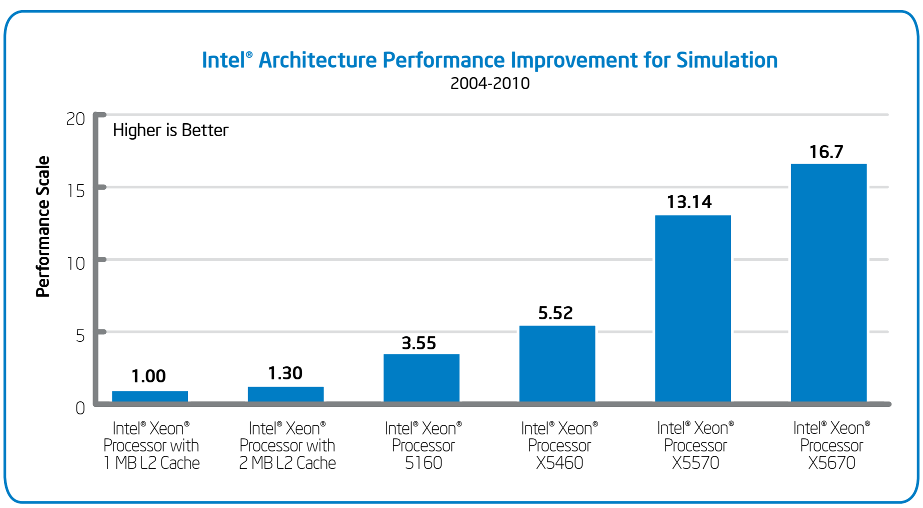

Figure 3. Each new generation of Intel® processor increases performance. Between 2004 and 2010, the performance of Intel® Xeon® processors increased almost 16x when running our most intensive simulation application. Intel internal measurements May 2007, November 2007, February 2009, and March 2010.

RESULTS OF SILICON DESIGN MIGRATION

We completed our RISC migration in the silicon design environment by 2004. The price/performance advantages have continued to deliver additional benefits post-migration, due to significant architectural and manufacturing process improvements in each generation of architecture. Figure 3 shows performance increases in workloads for the most frequently used batch application employed in the simulation of Intel chip designs.

Migrating Intel's silicon design environment to Intel architecture also offered significant benefits in the areas of cost efficiency and server management efficiency.

Cost Efficiency

Because Intel-based platforms are far more affordable than similar RISC-based platforms, we significantly reduced capital expenditures. In the six-year period between 2000 and 2005, we achieved savings of USD 1.4 billion. As shown in Figure 4, this figure included software, data center, and hardware savings. For example, migration reduced data center space requirements by 20,000 square feet.

Server Management and Data Center Efficiency

Standardizing on Intel architecture enabled us to manage servers more efficiently. With an open-source platform, we automated many of our management processes. Each time a new processor or version of Linux was released, we could quickly and easily upgrade, port the tools, and take immediate advantage of performance improvements. Previously, we couldn't create these efficiencies because we were using proprietary RISC-based platforms from multiple suppliers.

As shown in Figure 5, standardization enabled us to increase the number of servers each IT support engineer can effectively manage, from only 20 servers in 1998 to 300 servers today.

Manufacturing Environment

By early 2004, two pressing business needs became apparent in our manufacturing environment:

- Our legacy RISC-based platforms were quickly approaching end-of-life status, and RISC suppliers were no longer supporting their products. This situation would seriously jeopardize the high availability and reliability of our manufacturing environment.
- We needed to improve performance to handle larger amounts of data and increasing numbers of transactions.

After clearly demonstrating the benefits of migrating to Intel architecture in our silicon design environment—including increased performance, reduced TCO, and improved server management—we felt we could address these problems by migrating from RISC-based platforms to Intel-based platforms in our manufacturing environment.

However, a major hardware migration in Intel's manufacturing environment could potentially threaten factory production schedules. Therefore, we had to avoid unscheduled downtime. The measure of our success was migrating the entire manufacturing environment to a new architecture while maintaining full manufacturing output—with limited scheduled downtime for the conversions.

OVERVIEW OF THE MANUFACTURING ENVIRONMENT MIGRATION

Intel has two types of factories:

- Wafer fabrication factories, which produce silicon wafers in high volumes. Many integrated circuit products can be made from a single wafer.
- Assembly and test factories, which slice the wafers into separate integrated circuit products and put them into final packages that are tested for quality and functionality.

In both our wafer fabrication and assembly and test factories, every aspect of the manufacturing process is tracked by a highly complex manufacturing execution system (MES). The MES consists of hundreds of applications and scripts, and monitors all materials, routes, and production steps. Each application and script is critical to the successful completion of the manufacturing process. Our migration strategy would have to make sure that all these tools and scripts would work, without interruption, on the new platform.

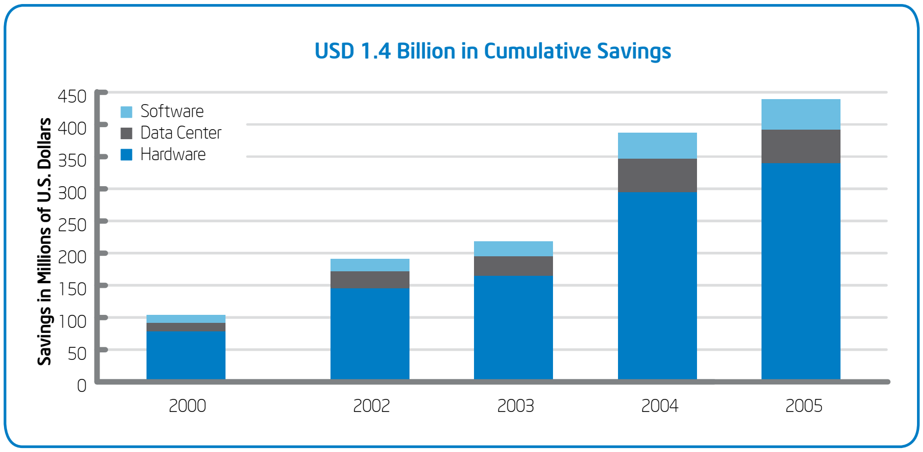

Figure 4. Migrating our silicon design environment from RISC to Intel® architecture saved Intel more than USD 1.4 billion over a period of six years. Source: "Intel Takes the RISC Out of Computing." *Premier IT Magazine*. Winter 2007.

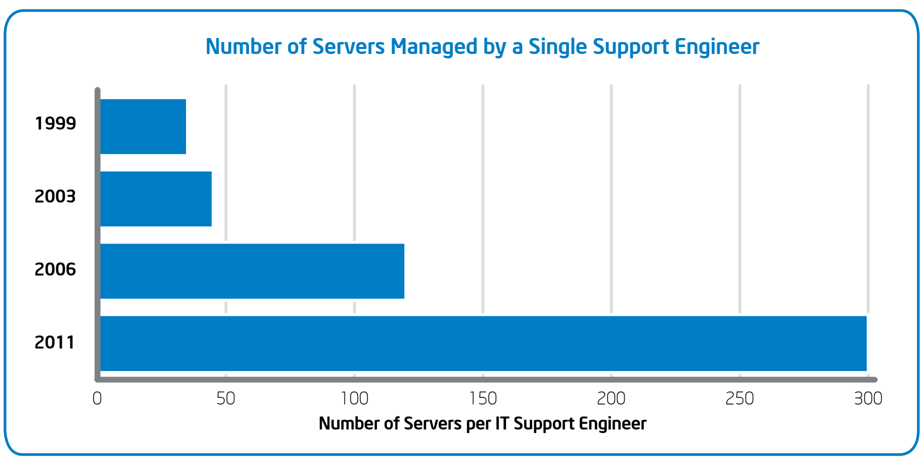

Figure 5. Standardizing on Intel architecture has significantly increased our server management efficiency, measured by how many servers a single support engineer can manage.

We devised two strategies to meet the different needs of each type of factory. We converted wafer fabrication plants first to align with a new MES platform and the implementation of a new 300 mm wafer size. To minimize risk, we ran both old and new platforms in parallel, or started new technology on the new platform as part of a new factory start-up process. In assembly and test, we are incrementally inserting Intel architecture in the production MES environment.

WAFER FABRICATION MIGRATION

Our strategy for migrating Intel's wafer fabrication factories to Intel architecture consisted of two important elements:

- **Timing.** We aligned the introduction of Intel architecture into our wafer fabrication factories with the introduction of 65 nanometer (nm) technology in 2004. By conducting the migration and the conversion of the manufacturing process at the same time, we avoided disturbing each wafer fabrication factory twice—thereby maintaining the best possible level of business continuity and productivity.
- **Tight control.** We further minimized any negative effects of migration by initially developing a new MES with a new platform based on Intel processors in our development factory. Once we completed and validated the new platform—which took about two years—we then used our highly systematic and methodical "Intel Copy Exactly" technology to transfer the new platform into high-volume manufacturing. Intel Copy Exactly is a methodology we use to transfer a known methodology of manufacturing from one site to another so we can maintain repeatability, efficiency, and reliability for manufacturing silicon chips.

  The Intel Copy Exactly process took about six months and minimized the risk of introducing errors and problems into high-volume manufacturing by replicating every detail—including hardware and software components—which might affect the manufacturing process.

With the exception of some legacy factories, for which upgrades would not be cost-effective, all of Intel's wafer fabrication factories now run on Intel architecture.

ASSEMBLY AND TEST MIGRATION

In our assembly and test factories, we have begun replacing RISC-based platforms while the factories are in full production. In order to enable this migration strategy, the Technology Development group upgraded and validated the existing legacy MES and related software components—including the OS—to work on both the old RISC platform and the new Intel® Itanium® processor-based platform.

Once we had a software stack that was qualified on both RISC and Intel architecture, we began replacing the RISC servers with Intel-based servers on a server-by-server basis in a multi-node configuration. This is an ongoing process; the first production assembly and test factory will be completed in the first half of 2011, and we expect to complete the conversions of Intel's worldwide assembly and test factories to Intel architecture by the end of 2012.

RESULTS OF THE MANUFACTURING MIGRATION

Migrating to Intel architecture in our manufacturing environment has afforded us several benefits:

- Wafer fabrication factories can handle larger amounts of data and more transactions in less time, using industry-standard, highly reliable hardware.
- Assembly and test factories no longer rely on outdated, unsupported hardware. This improves business continuity and productivity. In addition, based on internal measurements, we have seen a 10 percent to 50 percent performance improvement, depending on the application and transaction. We have also experienced a greater than 50 percent decrease in resource utilization.

Enterprise Resource Planning Environment

Intel's ERP system is crucial to business operations, automating the integration of internal and external management information across the company. Because it facilitates the flow of information between all business functions, including order management and human resources, a disruption to the ERP system can potentially have a negative effect on business continuity.

Originally, Intel's ERP system ran on mainframes using a proprietary processor architecture. From 1992 to 1996, we migrated our ERP environment from the mainframes to a proprietary OS running on Intel® Pentium® processor-based servers.

In 2004, we faced a two-fold dilemma:

- Our server supplier was acquired, and the new company decided to discontinue the line of servers we were using. This threatened the business continuity of our mission-critical ERP environment, because soon no upgrades or support would be available.
- Our supply chain in-memory database applications required more memory than was available on the platform we were using.

We therefore decided to upgrade the OS, database, and software solution stack simultaneously with migration to a newer Intel architecture.

Since completing Stage 2 of the ERP migration in 2008, we have continued to upgrade the ERP environment on a regular cadence to take advantage of compute performance and technology improvements such as multi-threading and multi-core processors in new generations of Intel architecture.

ERP MIGRATION STRATEGY

Because the true cost of a platform migration does not lie in the hardware, but rather in the downtime cost to the business, upgrading the entire software and hardware platform all at once meant we could avoid bringing the system down multiple times. It also meant we would need to monitor many changes occurring simultaneously.

To help the entire process run smoothly, we developed a project plan that defined our goals and a method for accomplishing them. Similar to the silicon design migration, we used a focused, streamlined approach. We also gathered a team of internal and external experts to guide and optimize the migration process.

We divided the migration into four phases: explore, plan and reference design, testing, and deployment. These separate phases enabled us to evaluate and test where we might encounter problems with the solution stack. By the time we reached the final step, deployment, we had worked through potential problems, validated the software stack, and could deploy with confidence.

RESULTS OF STAGE 2 OF THE ERP MIGRATION

By forming a team of technical experts, including supplier consultants, and a methodical multiphase approach, we successfully migrated our ERP environment.

- We completed the entire migration in 12 months—including testing, validation, and deployment—with only 18 hours of total planned downtime.
- We increased system reliability by migrating to a platform and OS that is fully supported.
- By migrating to newer Intel architecture, we have been able to introduce multilanguage capabilities into the ERP database and user interface, take advantage of parallel processing, and implement very large memory models for the database.
- Moving to industry-standard servers allowed us to consolidate our database and application tiers for a significant reduction in TCO.
- We improved system performance, reducing transaction response time and batch runtimes.

Selecting the Right Hardware

Intel-based platforms offer a broad choice of hardware, which allowed us to carefully choose which generation of Intel architecture we used for various system components as we migrated each of our mission-critical environments. This approach enabled us to use servers that are optimized for specific types of workloads.

For example, in situations where a scale-up approach makes sense, such as our manufacturing database, we chose Intel Itanium processor-based solutions. In other situations, such as our ERP environment, scale-out is a better solution, and therefore we chose Intel® Xeon® processor-based solutions.

We also used a combination of two-socket and four-socket servers for our mission-critical systems, depending on whether scale-up or scale-out was the best approach. Scale-up made sense for our ERP database because it is capable of taking advantage of multi-threading. Therefore, deploying it on four-socket servers enabled us to increase transaction speeds due to parallel processing. In other cases, such as silicon design servers, scale-out better met our computing needs; we simply added more two-socket servers to the computing clusters as necessary to meet performance requirements.

BEST PRACTICES FOR MIGRATION

Although we developed different migration strategies to align with the specific needs of each environment, we formalized a set of best practices that apply across mission-critical platforms:

- **Align with business goals.** Our migrations addressed a variety of specific business challenges, such as improving performance, reducing capital expenditures, enhancing productivity, and transitioning to a fully supported platform when hardware or software products approached end-of-life status.
- **Improve continuity through methodical implementation.** To reduce the complexity of a platform migration, we broke it into manageable phases with rapid and clear success defined at the end of each phase.
- **Select appropriate hardware and software.** We evaluated and chose a combination of hardware and software that best met the business needs of the environment.
- **Coordinate the solution stack.** If several changes were necessary, such as to both the hardware and the OS or application stack, we made them all at once to minimize disrupting mission-critical business activities.
- **Establish centers of technical competency.** We invested in developing a center of technical competency that could provide leadership and guidance in defining and executing the migration strategy.
- **Engage other stakeholders.** We worked with suppliers early in the process to make them aware of gaps in the application suites necessary to make the migration complete from an end-to-end perspective, and enabled them to contribute to the overall success of our migration objectives. For some migrations, such as silicon design, we also engaged users in the migration process, involving them in pilot projects and asking for user feedback
- **Quantify and track business value.** We measured and tracked the business value of the migrations on an ongoing basis in order to quantify benefits.
- **Test and validate, then copy exact.** We thoroughly tested the new platform, then froze the environment and deployed it across all affected environments with no modifications.

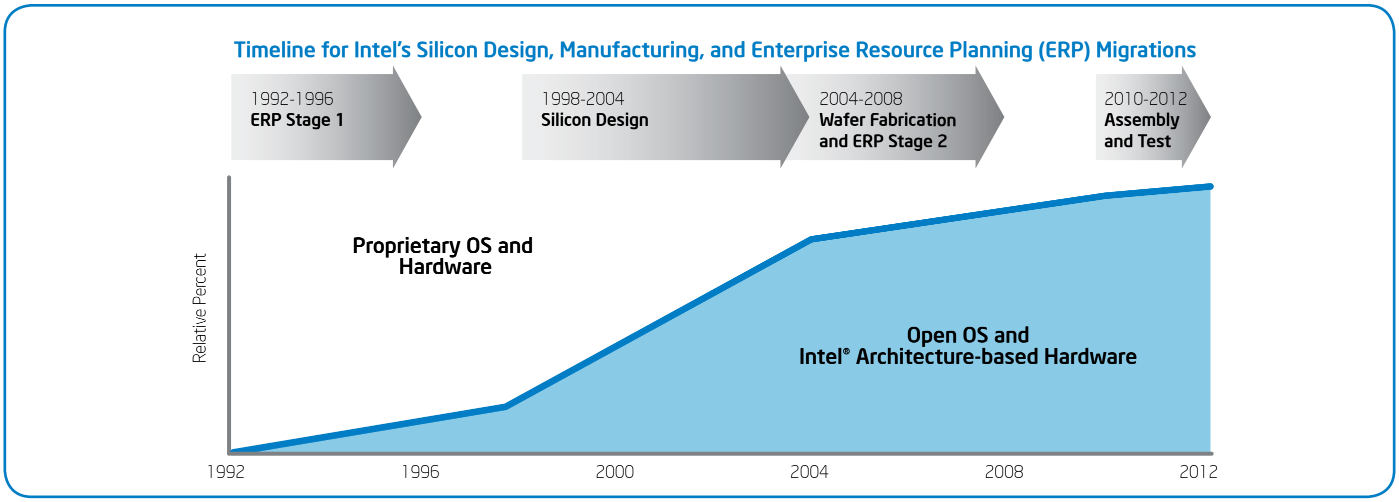

Figure 6. A multiyear strategy of migrating mission-critical systems to Intel® architecture enabled us to improve performance and reduce total cost of ownership without interrupting vital business processes.

CONCLUSION

We anticipate that Intel's most mission-critical environments—silicon design, manufacturing except for a few legacy facilities, and global ERP—will run entirely on Intel architecture by the end of 2012, as shown in Figure 6. We embarked on these migrations with caution, given Intel's dependence on established systems. Our experience has shown that implementing best practices can aid in a more successful and smooth conversion—thereby maximizing both business continuity and the business value of the migration.

Using best practices, we minimized or eliminated most risks while benefiting from the migration to Intel architecture in our mission-critical systems:

- Improved system performance to increase productivity and support faster time to market for our products.
- Increased ROI and decreased operational costs due to industry-standard servers and clients.
- Gained the ability to choose the best hardware for the workload because Intel's open standard enables the same OS on any Intel architecture-based platform.
- Decreased data center footprint due to space- and energy-efficient servers.
- Achieved savings of more than USD 1.4 billion over five years.

With Intel's silicon design, most manufacturing, and ERP environments running on Intel-based platforms, we have demonstrated that Intel architecture provides the performance and reliability needed by mission-critical applications, and that migration can be accomplished with minimal downtime and disruption to key business processes.

For more information on Intel IT best practices, visit www.intel.com/it.

CONTRIBUTORS

Ajay Chandramouly
Industry Engagement Manager, Intel IT

ACRONYMS

EDA	electronic design automation
ERP	enterprise resource planning
MES	manufacturing execution system
nm	nanometer
PoC	proof of concept
ROI	return on investment
TCO	total cost of ownership

1 "Intel Takes the RISC Out of Computing." *Premier IT Magazine*. Winter 2007.

2 We have legacy RISC-based platforms in older factories that would not be cost effective to replace and will likely run on the same hardware and software until those facilities close.

3 "Migrating from RISC to Intel® Architecture: Total cost of ownership (TCO) analysis," May 2003.

Printed in USA
0511/AC/KC/PDF

Please Recycle
325335-001US