The CMP Guide
to Intel® Architecture Server Deployment

Beyond Proprietary Computing

With highly scalable, enterprise server building blocks, businesses can now move beyond isolated, proprietary computing systems to realize a high-performance, cost-effective processing infrastructure that enhances both operational stability and business agility.

INSIDE

Enterprise server deployment scenarios for:

- XML-based e-commerce
- Real-time business intelligence
- Enterprise customer relationship management
- Extended supply chain management

SKU#250437-001

Executive Overview

The convergence of highly scalable server technologies and advanced multi-tier application software has set the stage for new computing architectures that greatly enhance the competitive performance of IT-driven enterprises. Leading companies are now moving away from isolated, vertical islands of proprietary computing and toward a more open and broadly deployable computing model. The open, multivendor approach is based on a high-performance server infrastructure that scales and adapts to every area of the end-to-end enterprise: front office, back office and wherever business goals demand the highest possible levels of application performance.

The open computing model discussed in this guide is built with servers based on Intel® processors and thousands of widely available enterprise IT products and services that enhance the Intel Architecture (IA). With a global vendor community now delivering IA-based solutions for all areas of the enterprise, IT architects can rapidly deploy mission-critical applications with unprecedented levels of scalability, performance and cost effectiveness. Open computing also facilitates strong interoperability between new distributed multi-tier applications and existing RISC and mainframe systems.

Accelerated innovation and reliable operations

Until recently, proprietary RISC and mainframe systems have been the traditional option for certain core enterprise applications. Proprietary systems draw on a limited pool of industry resources and are consequently costly, slow to advance and capable of addressing only a limited range of enterprise computing needs. Many enterprises depend on multiple proprietary systems, a situation that has led to fragmented, inefficient computing architectures with limited interoperability and limited cross-functional integration of key business processes.

The open computing approach that began with Intel-based PCs on the desktop in the 1980s has been extended to front-end Web servers, mid-tier application servers and recently to the most mission-critical transaction-processing systems in the back office. Today, advanced IA processors and supporting server hardware and software products give IT architects a very powerful and cost-effective set of multipurpose building blocks that apply to the widest possible range of transactional, analytical and content-serving applications.

Open computing has the unique ability to support reliable operational systems while at the same time accelerating corporate innovation through flexible and rapid time-to-benefit deployments. Intel uses the term "macroprocessing" to describe the benefits of this approach. Macroprocessing delivers:

- High levels of end-to-end scalability with full interoperability for hardware and software building blocks and no proprietary barriers.
- Balanced-performance computing that eliminates bottlenecks throughout the processor and system architectures.
- Data center-class reliability, availability and serviceability (RAS) built into processors and server platforms.
- Cost-effective server solutions driven by open and competitive market conditions.
- The Intel® e-Business Network, one of the world's largest vendor communities, delivers market-tested best practices, fast-paced technology innovation and best-of-class solutions.

Building the optimum hardware architecture

New enterprise applications based on extensive Internet connectivity and highly interactive e-business services demand ever-increasing performance and scalability from servers. Meanwhile, corporate executives are calling for accelerated application deployment schedules and no compromise on system availability or reliability. In this environment, IT architects need a server infrastructure that supports fluid vertical and horizontal scaling of processing resources *(see **Figure A**)*. In many cases, it's better to deploy applications in a highly modular, distributed manner with numerous identical servers—a "scale-out" strategy. In a scale-out architecture, each application or software module can have its own dedicated server and processing

Figure A: **Can your enterprise IT architecture do this?**
The open computing model works equally well with many servers in large server farms, or a massive consolidated server with powerful centralized processing capabilities—or both.

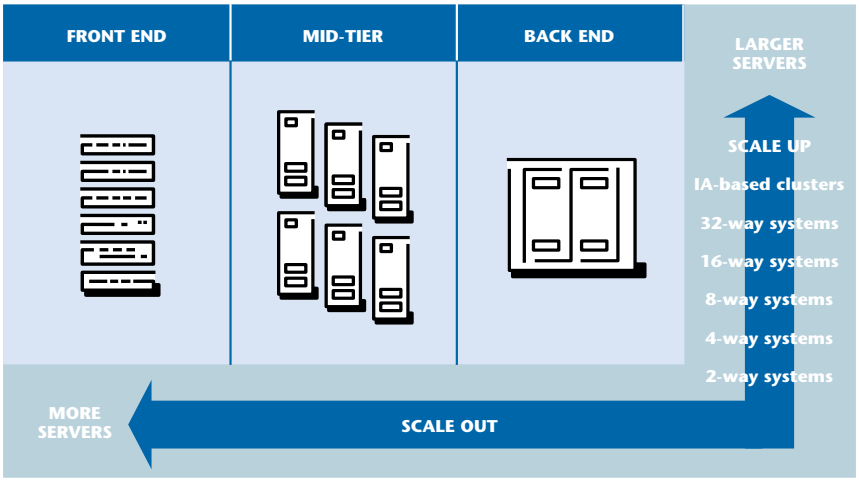

resources. Scale out is an important strategy in many leading enterprises, but in some cases, there is a need to consolidate key applications and large data stores on a single powerful centralized server that contains dozens of symmetric multiprocessing (SMP) processors—this is a "scale-up" strategy.

When IT architects can scale out and scale up reliable, high-performance processing power on demand, it's called "scale right." Using scale-right methods, enterprises can intelligently deploy computing resources exactly where they are needed throughout the enterprise. If modularity, redundancy and fault tolerance are crucial, scaled-out clusters provide the most suitable choice. If hardware consolidation and centralized memory/storage/performance scaling issues apply, scaling up with many processors in a single system is a logical alternative. The Intel Architecture is the only computing model available today that supports the full range of scale-up and scale-out capabilities for enterprise servers in an open computing infrastructure.

Intel Processors for Servers: Industry-leading performance, enterprise-wide coverage

The Intel Xeon™ processor family scales from slim profile dual-processor front-end Web servers up to massive 32-way or greater data center servers running high-volume transactional and computational applications. Server clusters integrate hundreds of Xeon processors in fault-tolerant and high-availability configurations. Complementing Xeon processors on the high end of the processing spectrum, Intel Itanium™ processors support multi-terabyte datasets, enterprise transaction processing, large in-memory databases and intensive floating-point processing for scientific, engineering, analytical and security applications. Between them, the Xeon and Itanium processor families cover the entire range of enterprise computing requirements.

At the level of operating and database software, recent advancements mean that mainstream enterprise applications can now take full advantage of extensive symmetric multiprocessing resources within a single server or across multiple SMP nodes in tightly coupled server clusters. Enterprise-grade database and clustering products based on cost-effective IA-based servers are now widely available from such companies as HP, IBM, Microsoft and Oracle.

Along with the requirement for end-to-end scalability, the need for improved IT economics at all levels of the enterprise is encouraging the push for open, multi-vendor computing. Intel processors are manufactured in great volume and used by virtually all major server manufacturers and the majority of integrators. With a worldwide user base and over 80 percent of the enterprise server market segment share (source: IDC), IA is extending desktop "volume economics" into the high-end enterprise server environment. This results in acquisition and ownership costs that are typically several times less than those of proprietary systems—and a cost/performance ratio that can be several times greater.

Beyond volume economics, the operational stability enabled by IA processors is another important driver, as open computing takes its place in core enterprise systems. Exclusive diagnostic and error-checking features in IA processors are now supported by major operating software products that provide data center-class reliability, availability and serviceability that was previously available only in expensive proprietary RISC and mainframe systems.

A global development effort

Whether IT architects have a mission-critical back-office perspective or a highly focused, front-office mandate, an open, IA-based enterprise server infrastructure delivers the broadest possible range of hardware and software solutions from over 55,000 worldwide members of the Intel® e-Business Network. As a result, enterprises can select the optimal solution for their computing infrastructure without being locked into a single vendor.

According to Mike Fister, Vice President and General Manager of Intel's Enterprise Platforms Group: "No single company in the enterprise server market is capable of solely sustaining the investment required for consistently high levels of innovation, performance and value. Intel's view is that it takes an entire industry to deliver these characteristics with the pace, flexibility and choice required by today's challenging IT environment."

This viewpoint is shared by many different types of enterprises that have built their most strategic, mission-critical applications on IA-based servers and the open computing model, including: Black & Decker, CERN, Home Shopping Network, Merrill Lynch, Nasdaq, Philips Lighting, San Diego Supercomputing Center, Yahoo! and many more. In all types of enterprises open computing lets IT architects reap the benefits of new, distributed architectures and leading-edge e-business applications—without leaving core RISC and mainframe systems behind.

The application scenarios on the following pages all rely on the same essential set of macroprocessing principles: end-to-end scalability, balanced processor performance, major cost economies, data center-class system availability and a broad selection of interoperable hardware and software building blocks delivered by the global community of IA developers.

Open Computing Server Deployment Scenarios

The ultimate enterprise building blocks

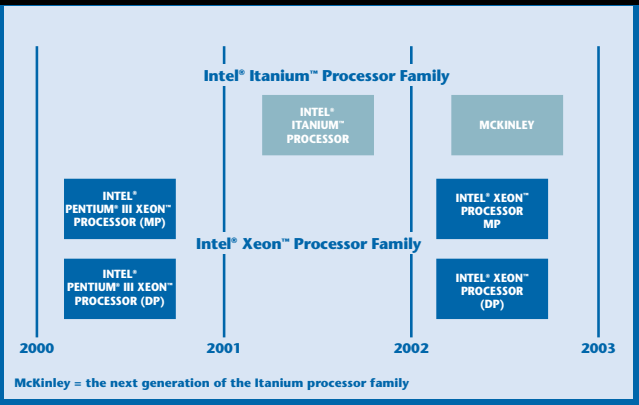

Figure B: **Intel® processors for servers—Evolution in action**

Servers based on Intel® Pentium® III Xeon™ processors are widely available from a large number of system vendors. To assist IT planning, this guide focuses on the latest dual and multiprocessor servers in the Xeon and Itanium processor families (*see* ***Figure B***). Numerous hardware makers are offering enterprise servers based on Intel Xeon processors starting in early 2002. Servers based on the Intel Itanium processor family are supplied by more than 25 system vendors.

Intel® Xeon™ Processors—Price/performance perfection

Servers based on the Intel Xeon processor family are the primary multi-purpose building blocks of the open computing approach. The latest dual and multiprocessor (MP) Xeon processors surpass previous IA 32-bit incarnations with a balanced performance model that eliminates scalability restrictions and bottlenecks throughout the memory, cache, I/O and execution aspects of the server architecture. MP servers based on the Xeon processor family deliver some of the best performance benchmarks in the industry—at a fraction of the cost of RISC and mainframe technologies.

The Xeon processor features Intel's exclusive NetBurst™ microarchitecture, which is based on hyper-pipelined technology that enables speeds well in excess of 2 GHz. To complement its high frequencies, NetBurst deploys a 400 MHz front-side bus and a design that minimizes latencies at every stage of the processing path, tripling bandwidth for faster data transfer and improved performance. Xeon-based servers contain up to 32 or more processors; Xeon-based clusters support over 100 processors. Dual processor Xeon-based servers are ideally suited for front-end Internet infrastructure, delivering speeds up to 2.20 GHz or more; with Level 2 cache up to 512KB. For the demanding application and database workloads in the mid-tier and backend, multiprocessor servers based on the Xeon Processor MP run at up to 1.60 GHz and feature an advanced three-level cache subsystem, with up to 1MB of Level 3 cache in addition to 256KB of Level 2 cache.

Hyper-Threading technology, an exclusive aspect of the Xeon processor family microarchitecture, lets a single physical Xeon processor appear as two logical processors to multi-threaded software. This allows multiple threads to be executed simultaneously, improving resource utilization by over 40 percent for certain multi-threaded applications. The Xeon processor's Hyper-Threading technology represents the latest innovation in a long history of technology leadership from Intel.

Intel® Itanium™ Processor Family—Beyond 64-bit

The Intel Itanium processor has a groundbreaking parallel architecture, multi-terabyte memory address space, a massive set of internal execution registers (328 registers), and a 64-bit wide data bus. The Itanium processor family is particularly well-suited for applications with very large datasets (e.g., many gigabytes or terabytes), complex floating point computations or intensive encryption/decryption (e.g., secure socket layer-SSL) requirements. The Itanium processor's exclusive Explicitly Parallel Instruction Computing (EPIC) design combines state-of-the-art speculation, predication and extensive instruction-level parallelism. EPIC executes up to 20 operations per clock cycle, permitting more work in less time. Other Itanium processor highlights:

- The industry's broadest data center operating system support.
- An industry-wide effort to port 64-bit databases, enterprise applications and infrastructure applications.
- Advanced built-in system monitoring and diagnostic features for data center-class reliability, availability and serviceability.
- The next generation Itanium processor—project name McKinley—is further improved by faster processor speeds, a 400 MHz memory bus and a number of microarchitectural enhancements.

It should be noted that the Itanium processor family complements the Xeon processor family but does not replace it. The two architectures coexist with each fulfilling its own role, creating a high-performance, end-to-end enterprise computing solution that will continue to receive technical refinements from Intel's architecture labs going forward.

OPEN COMPUTING SCENARIO 1

XML-based e-commerce

Figure C: **A scalable computing architecture for high-volume e-commerce**
Revenue-generating e-commerce applications are stable and scalable when deployed with XML and an IA-based server infrastructure.

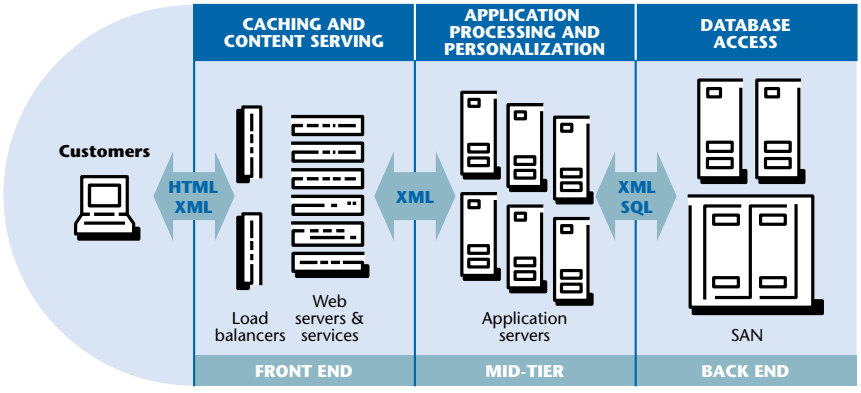

The move to XML-based Web services is changing the computing landscape for a great range of applications inside and between enterprises. This is largely due to XML's inherent platform independence and its ability to "self-define" itself for each new environment. Leading IT vendors Ariba, BEA, Commerce One, HP, IBM, Microsoft, Oracle, SAP, SAS and many others are actively engaged in building XML-based solutions for e-commerce and other mission-critical enterprise applications. Setting precedents for industry cooperation, Microsoft.NET and the J2EE development environments are both relying heavily on XML for data exchange and interoperability in a full range of e-commerce applications with hierarchical and modular software designs.

XML e-commerce applications are particularly effective when deployed in a multi-tiered ("n-tier") architecture. In this environment, the internal modules of applications (front-end edge services, mid-tier application logic and back-end database services) are decoupled and distributed onto dedicated servers as necessary, to improve performance, availability and scalability. XML is the glue used to integrate the various tiers of the enterprise e-commerce application. Scalable, high-performance servers are the key enablers of this approach.

New front-end performance requirements

In the front end of the present e-commerce scenario, traffic is routed to a front-end server farm and serviced by a number of Web servers, load balancers and caching appliances. Traditionally, corporate Web sites have served static HTML pages that put only modest loads on front-end servers. In contrast, XML e-commerce and inter-application Web services involve more data-rich interactions and processor-intensive security encryption/decryption. Because XML is more complex and flexible than HTML, its markup is more verbose and more computationally intensive to process. The combination of data-rich client requests and business-to-business (B2B) application interactions is creating a new yardstick for server performance requirements in the front end.

Dual Xeon-based servers are a good match for the front end of the present scenario. In cases with higher levels of transaction processing or XML complexity, it may be advisable to scale up to a 4-way server based on the Xeon Processor MP. As application interactions become more complex and dynamic, Xeon-based servers reduce front-end response times by virtue of their balanced-performance NetBurst™ microarchitecture. NetBurst combines high processor frequency, a high-speed 400 MHz memory bus, optimized caching and a number of microarchitecture enhancements that result in superior 32-bit program execution and reduced processor latencies. This translates into significant performance increases across a broad range of applications. The Xeon processor's support for the 64-bit PCI-X bus standard allows numerous high-speed I/O devices and network connections without I/O becoming a choke point—even during extremely heavy periods of customer activity.

Multi-threading the middle tier

Once an incoming request is received by the front end, if it requires further processing, it is forwarded to mid-tier application servers. Most XML and application logic processing takes place in the mid-tier. If a database query is necessary, the middle tier forwards it to the back end via SQL or XML. In high-volume shopping cart and e-procurement applications, mid-tier servers must handle hundreds or thousands of XML messages concurrently.

The mid-tier is typically an architecturally volatile area of the enterprise architecture. It has the high availability requirements of the data center, but it is also much more subject to rapid "quick win" application deployments. Modularity is one way that IT departments can ensure that existing applications are not disrupted by new deployments. Four-way or greater servers are used to prototype, test and deploy new applications on their own platforms, which insulates the existing infrastructure from risk.

The scenario example deploys the Intel Xeon Processor MP, which meets mid-tier XML processing requirements with its advanced three-level cache subsystem. Advanced caching systems reduce contention for main memory in multiprocessor configurations, as does the 400 MHz front-side bus. The Xeon processor's Hyper-Threading technology executes two threads simultaneously on each processor, increasing the number of XML transactions that can be executed concurrently. Since incoming XML requests create a large number of application threads that all contend for a finite set of processor resources,

Hyper-Threading technology enables smoothly scalable performance.

Increasingly, software vendors are porting key enterprise applications to the Itanium processor family. Very large caches of up to 4MB, large memory addressability and parallel execution will enable very fast execution of instructions, making Itanium-based mid-tier servers a good fit for the most computationally intensive environments.

Back-end scalability

The high performance and cost-effectiveness of the Intel Xeon Processor MP is unlocking new possibilities in the database tier of high-volume e-commerce applications. It's now possible to build very high-performance databases with clusters of 4-way, 8-way or greater servers that run a variety of high-end clustering products. For example:

- **IBM.** High-performance DB2 database implementations are now optimized for Windows running on IA-based servers and clustered systems. IBM deploys servers based on the Intel Xeon and Itanium processor families in both Microsoft and Linux operating environments.
- **Microsoft.** Windows 2000 Datacenter Edition now adds failover clustering services to its already very high availability (99.9% to 99.999%). This enables the highest levels of reliability and virtually non-stop service availability on IA-based servers in the data center.
- **Oracle.** Oracle9*i* Real Application Clusters (RAC) create a single database image across multiple IA-based servers for high-end enterprise transaction processing and back-office applications.

Back-end clusters based on Xeon processors support SAN/NAS storage solutions and dynamic application load balancing. These clusters are now being deployed in mission-critical data-center environments in many types of enterprises—with significant performance, scalability and flexibility advantages, compared to proprietary RISC and mainframe systems.

If e-commerce applications call for massive memory addressing, intensive floating-point performance and leading-edge security algorithms, the Itanium processor family is the best choice. SMP and clustered servers based on the Itanium processor extend the Intel Architecture to the top of the high-end server space.

OPEN COMPUTING SCENARIO 2

Real-time business intelligence

As Business Intelligence (BI) systems strive to capture data from a multitude of customer touch points and operational areas throughout the enterprise, very large datasets are generated. In some cases terabytes of data are accumulated at a rate of tens or hundreds of gigabytes per month over the course of years.

With the current emphasis on interactive and ad hoc analytical queries, BI applications now exceed the capabilities of traditional back-office systems and batch processes that take hours or overnight to complete. To support new real-time demands on data warehouses and data marts, servers must provide industrial strength processing power and memory addressing, with plenty of headroom for future growth.

Real-time performance is what BI users want, but at what cost? Data warehouse projects can exceed even the most generous budgets. A related challenge is the development schedule. Large, monolithic data warehouse projects can take 18 to 24 months to develop—if everything goes well.

There is also a strong requirement for front-office/back-office integration that embraces both best-of-breed and legacy applications. This is particularly challenging, and very few enterprises today have a single view of the customer across all key departments and application areas.

Advanced computing infrastructure for data warehouses

Due to huge data stores and complex processing requirements, data warehousing is a good candidate for a combined scale-out and scale-up approach—scale right. Depending on specific application needs, the computing infrastructure for data warehouses and data marts can be scaled up with very large central servers, or scaled out with many 4-way and 8-way devices. In BI applications where

Figure D: **Enterprise-wide business intelligence—with data center-class performance**
Very large and high performance data warehouse applications are now supported on cost-effective Xeon processor-based clusters, running clustering software from IBM, Microsoft, Oracle and others.

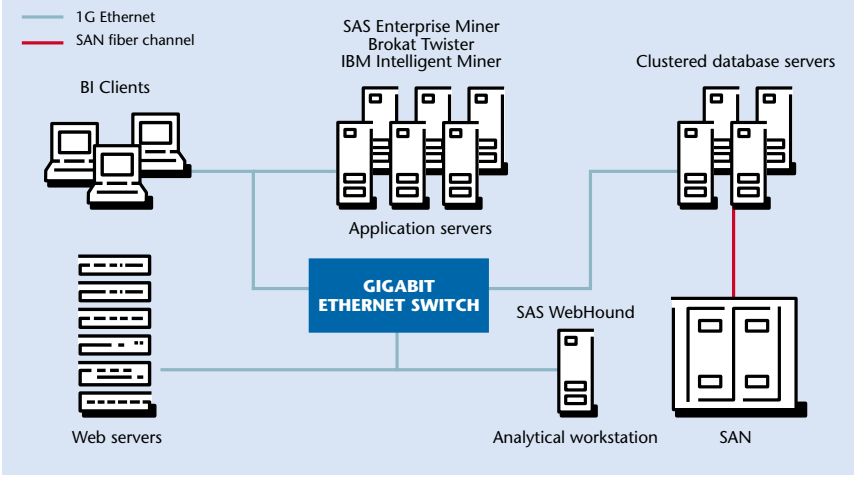

very high levels of performance are required from a single scaled-up system, Xeon-based servers with up to 32 or more processors in a single machine may be deployed.

In the back end of the example configuration, four 8-way servers based on the Xeon Processor MP work together in an automatic failover cluster that ensures data warehouse services will run with maximum uptime. IBM DB2, Microsoft SQL Server and Oracle9*i* RAC are good candidates for cluster-aware back-end software. With the widest operating system support of any data center processor, Xeon-based servers have the option of Unix, Linux or Microsoft for O/S software. This choice of operating environments and the resulting platform and application interoperability is not available from proprietary RISC and mainframe systems.

In the mid-tier, 4-way or greater servers based on the Xeon Processor MP are deployed to run analytical programs, as well as business intelligence, decision support and various e-business applications. As servers and storage devices are added, performance scales smoothly, and database operations are balanced across additional nodes within the cluster.

Parallel Business Intelligence

For the most rigorous analytical applications and the largest datasets, the Itanium processor family should be considered for both application and back-end processing requirements. Itanium-based servers can seamlessly access many gigabytes of data in main memory, which is more efficient than "demand paged" virtual memory addressing. With vast amounts of directly addressed main memory, Itanium processors minimize high-latency disk accesses that seriously degrade BI processing during real-time decision support and business analytics. With Itanium processors, entire data marts can be kept permanently in memory, using in-memory database products.

With the ability to execute up to 20 operations per clock cycle, Itanium-based servers can manipulate gigantic tabular row/column arrays within gigabyte or terabyte datasets, without becoming processor-bound. This enhanced throughput is necessary when enterprise BI applications must execute rapid sequential table and index reads, multi-table joins and extensive sorting of live data. BI is also facilitated by the Itanium processor's industry-leading internal register resources (128 integer, 128 floating-point, 72 specialized), which far surpass the generally limited number of execution registers found in proprietary 64-bit RISC architectures. The Itanium processor's internal resources allow it to handle unprecedented amounts of intermediate execution data, which is required for optimum parallelism, predication and speculation. Predication is the ability to execute both sides of a program branch simultaneously. This speeds branch execution and thus overall performance as it eliminates delays associated with branching decisions and penalties (should the wrong branch be executed). Speculation lets the Itanium processor execute certain instructions before they are required—delivering even greater performance advantages. The Itanium processor's advanced predication and speculation services optimize and "parallelize" program branches that occur frequently in ad hoc analytic and data-mining environments.

Leading database and enterprise application programs are being rapidly ported to the Itanium processor environment by such companies as Ansys, BEA, BMC, Computer Associates, Covalent, IBM, i2, J.D. Edwards, McAfee, Microsoft, MSC.software, Oracle, RSA Security, SAP, SAS, TimesTen and others. HP-UX operating software and related system-management programs are now available on Itanium-based platforms. With its datacenter-class reliability, availability and serviceability, and built-in cluster file system, HP-UX for the Itanium processor family provides an ideal platform for enterprise database management systems such as Oracle9*i* and 9*i* RAC.

Servers based on the Itanium processor family are available in 4-way and up to 32-way systems—for scale-up possibilities without the huge costs associated with back-office RISC and mainframe solutions. With clustering, hundreds or thousands of processors are supported in a massively parallel unified system *(see **NCSA builds a "TeraGrid" with the Intel Itanium processor family**)*.

NCSA builds a "TeraGrid" with Intel® Itanium™ processors

When faced with the very expensive proposition of upgrading its aging proprietary supercomputer network, the National Center for Supercomputing Applications (NCSA) found it could do a lot more with clusters based on the Itanium processor family—at a fraction of the cost. The planned TeraGrid system will use a network of over 3,300 Intel processors to create a distributed supercomputing system linking four sites capable of 13.6 teraflops (13.6 trillion calculations per second) and more than 450 terabytes of storage.

The TeraGrid will be accessible to researchers across the United States so that they can more quickly analyze and simulate complex scientific problems. Examples of research areas include: molecular modeling for disease cures, research on alternative energy sources, as well as climate and atmospheric simulations.

Initial massively clustered system implementations such as TeraGrid are found primarily in the scientific and technical computing arena. But expect to see these powerful methods focused on e-business and core enterprise applications in the near future.

Vendors for Open Computing

To help end-user enterprises be more competitive and market-responsive, these leading IT vendors provide a wealth of powerful computing solutions based on the Intel® Xeon™ or Intel® Itanium™ processor families.

Server Hardware

Acer
Caliber
Compaq
Dell
Egenera
Fujitsu Siemens Computers
HP
IBM
Marathon Technologies
MicronPC
NEC Technologies
Network Appliance
Premio
Stratus Technologies
Supermicro
Unisys

Strategic Alliances

BEA Systems and Intel for advanced e-business solutions: Together, BEA and Intel offer a Java application server platform that combines flexibility, scalability and best-in-class price/performance for today's complex enterprise environment. J2EE-based BEA WebLogic, optimized on Intel® Architecture, empowers enterprises to deploy Intel-based servers to run their mission-critical e-business applications and Web services.

i2 and Intel for dynamic, scalable supply chain automation: The i2 and Intel alliance combines i2's expertise in dynamic value chain management solutions with Intel's expertise in high-performance e-business solutions that enable open standards-based computing architectures.

IBM and Intel for scalable e-business and database solutions: As the leading provider of middleware and enterprise IT infrastructure built on open standards, IBM provides e-business and database software across multiple platforms and on the latest Intel Architecture servers. Building on a 20-year relationship, IBM and Intel today work jointly on major technology initiatives, marketing programs and software solutions for Intel-based architectures, including the Xeon™ and Itanium™ processor families.

Microsoft and Intel for a whole universe of software choices: With the new Microsoft.NET Enterprise Servers and Web services, powered by Intel server solutions, enterprises get proven levels of reliability, scalability, availability and manageability. Key Enterprise Servers for the Intel Architecture include: SQL 2000, BizTalk 2000, Commerce 2000, Application Center, Exchange 2000 and Host Integration Server.

Oracle and Intel for enterprise solutions: Intel and Oracle have worked together for years to produce powerful products for enterprise computing. Oracle's database, tools and applications, combined with current and new Intel technology such as the Itanium™ processor, give customers all the power, scalability, availability and reliability they need for today's and tomorrow's most demanding e-business solutions.

SAP AG and Intel for enterprise software with a competitive edge: Intel has been working closely with SAP AG since the 1994 launch of SAP R/3 on Intel Architecture. Today, all mySAP.com solutions are available on the Intel Architecture-based platform with superior performance and compelling cost while meeting the reliability, availability and scalability needs of enterprise customers.

Server Software

Acuo Technologies
BEA Systems
BroadVision
Caldera
Cognos
Commerce One
Computer Associates International
EMC
Entrust Technologies
Genesys Telecommunications Laboratories
Groove Networks
HP
i2
IBM
Interwoven
Javelin Technologies
J.D. Edwards
Lutris Technologies
Microsoft
Network Associates
Openwave Systems
Oracle
PeopleSoft
Peregrine Systems
Pivotal
Plumtree Software
Portal Software
PTC
Rational Software
Red Hat
Relativity Technologies
Reuters
RiskMetrics
Rogue Wave Software
RSA Security
SAP AG
SAS Institute
Siebel Systems
SpeechWorks
SteelEye Technology
SunGard
Symantec
TIBCO Software
TimesTen Performance Software
VERITAS Software
webMethods

Solution Providers

Accenture
Arthur Andersen
Avanade
Cambridge Technology Partners
Cap Gemini Ernst &Young
Computer Sciences Corporation
Deloitte & Touche
DiamondCluster
Digitas
Dimension Data
EDS
Inforte
Keane
KPMG
Luminent
Plural
PricewaterhouseCoopers
Questra
Scient
Silverline Technologies

Enterprise Servers At Work

Whether they use servers based on Intel® Xeon™ or Intel® Itanium™ processors or both, many of the most exciting and powerful companies in the world rely on Intel's scalable, open computing model for enterprise servers that definitely get the job done.

Best Buy. When the number one retailer of consumer electronics in the United States launched its BestBuy.com e-commerce site in mid-2000, it became one of the top 50 most popular sites on the Web in less than two months. BestBuy.com standardized on Intel processor-based servers from Compaq to drive the site. To start with, BestBuy.com deployed over 100 Intel-based servers and housed them in an Exodus Communications data center in Chicago. Crucial to the choice of Intel-based servers was the hardware's track record for scalability and availability—not optional considerations with such a massive customer base.

Black & Decker. When the Black & Decker Power Tools and Accessories Group undertook major e-business initiatives last year, it chose servers based on Intel® Pentium® III and Intel® Pentium® III Xeon™ processors to provide a scalable platform on which it could quickly deploy best-of-class solutions. Thanks to the flexibility of the Intel Architecture and the strength of the Intel® e-Business Network of over 55,000 IT solution providers, it took Black & Decker just a few short months to roll out robust, new information portals.

Commerce One. Commerce One provides industry-leading online e-marketplaces for B2B e-commerce, bringing together buyers and suppliers from around the world. Using XML, the company is able to connect thousands of buyers and suppliers who conduct mission-critical business over Commerce One.net's servers every day. Commerce One chose Intel-based Compaq servers to power its e-marketplaces because they provide total reliability and scalability for mission-critical applications that need to be up and running 24x7. And Compaq's consulting services provide world-class expertise in using the Intel platform.

Nasdaq. Since its inception in 1971, as the world's first electronic stock market, Nasdaq has set a precedent for technological trading innovation. Today, it is the fastest-growing stock market in the United States, trading over a billion shares a day. Nasdaq relies on over a thousand Intel-based servers for its corporate infrastructure and external Web sites. A typical Nasdaq server is a Dell PowerEdge configured with two or four Intel® Pentium® III Xeon™ processors. Intel-based servers also power Nasdaq's MarketWatch and other specialized business-critical OLTP applications, meeting their requirements for extremely high OLTP transaction rates and industry-leading hardware availability.

Philips. As a large European manufacturer, Philips Lighting Poland makes products that are sold through distribution channels across the world. The challenge of managing suppliers, customers, deliveries and order handling prompted Philips to seek an industrial-grade ERP solution that would streamline its production process, greatly reduce overhead and accelerate responsiveness. Driving the system: servers based on Intel® Pentium® III Xeon™ processors running Oracle databases that are accessed by Java servlets and a Java Database Connectivity (JDBC) application programming interface.

Yahoo! As the leading Web destination in the world, Yahoo! all but invented the Web portal concept and then went on to innovate around a wide range of online consumer services, free e-mail and Internet message chat, personalized information delivery, real-time stock quotes and a rich diversity of multimedia and news content. Yahoo! maintains over 1,500 servers at co-location facilities around the globe. Each successive generation of Intel processors has enabled Yahoo! to scale up its operation. Yahoo! is now working on integrating Intel Itanium processors into its architecture.

Liz Claiborne. To keep pace in the volatile and ever-changing retail market, Liz Claiborne relies heavily on its IT infrastructure to automate its business processes, from sales and marketing to inventory and distribution-center management. At the same time, it must keep IT costs down to succeed in the low-margin retail industry. Liz Claiborne is testing the Intel® Itanium™ processor's binary compatibility capabilities by running applications originally targeted for HP PA-RISC-based platforms on servers based on the Itanium processor. Without incurring great additional development costs, companies such as Liz Claiborne can leverage existing software, taking advantage of the improved performance and scalability of Itanium™ processor-based systems.

Preussag Group. As a leading European travel services provider with 80,000 employees worldwide, Preussag clients count on it for exceptional customer service, data warehousing/data mining and Web-based services. Preussag relies on mySAP.com and SAP R/3 applications for its software infrastructure. Due to the ever-increasing memory and performance requirements, Preussag is evaluating cost-effective, Itanium-based Dell PowerEdge servers, which feature a rack-optimized server design that accommodates up to four processors, 64 gigabytes of memory and hot-swap hard drives/power supplies.

Reuters. As a leading provider of financial, technology and news services, Reuters constantly seeks to maintain up-to-the-minute, real-time information delivery capability. The Reuters CTO Innovation Lab tested the Intel Itanium-based server from Compaq to determine its capability for achieving cost savings, improved performance and consolidation of systems for operational efficiency. When Reuters tested Microsoft SQL Server 64-bit and an internal application on the Intel® Itanium™ processor family, it found significant performance advantages in running both applications.

OPEN COMPUTING SCENARIO 3

Enterprise Customer Relationship Management

With the economy putting customer retention and customer lifetime value at center stage, CRM applications have an increasingly mission-critical nature. CRM is becoming part of the operational workflow of the enterprise, and hence it must be highly reliable and available to a degree that was previously associated only with back-office online transaction processing. Going forward, a fully realized enterprise CRM function must manage data to and from an ever-expanding universe of customer touch points: direct sales, call centers, direct mail, e-commerce, e-mail, voice-over IP, mobile commerce, interactive TV and so on. The goal is to use all this data to personalize and enhance customer interactions in real time.

A flexible scale-out/scale-up architecture gives IT designers a viable platform for real-time customization and personalization of dynamic e-commerce content on an enterprise scale. Pools of identical servers are configured in each tier of the architecture; availability and performance scale incrementally as the size of the server farm in each tier grows. Storage is decoupled from servers so it can scale on its own, without requiring the upgrading or replacement of servers. Processing loads are distributed between servers in each tier by load-balancing software or specialized load-balancing appliances.

In the current scenario *(see **Figure E**)*, Xeon-based servers were chosen as universal building blocks that will adapt readily to all the various tiers of the eCRM architecture. In the front end, 2-way servers are configured in a load-balanced server farm. If one server fails, its traffic is automatically routed to another working server within the farm by front-end load directors. All front-end servers have access to content that is concentrated in a network-attached storage (NAS) device with fault-tolerant RAID disks. Typical front-end dual-processor server configuration: 2.0 GHz or greater Xeon processors with 512KB cache, 512MB-1GB main memory and dual network interface cards.

Figure E: **Processors for the customer-driven enterprise**
eCRM and content personalization applications benefit greatly when high-performance, cost-effective processing power can be deployed throughout the n-tier architecture on a scalable, as-needed basis.

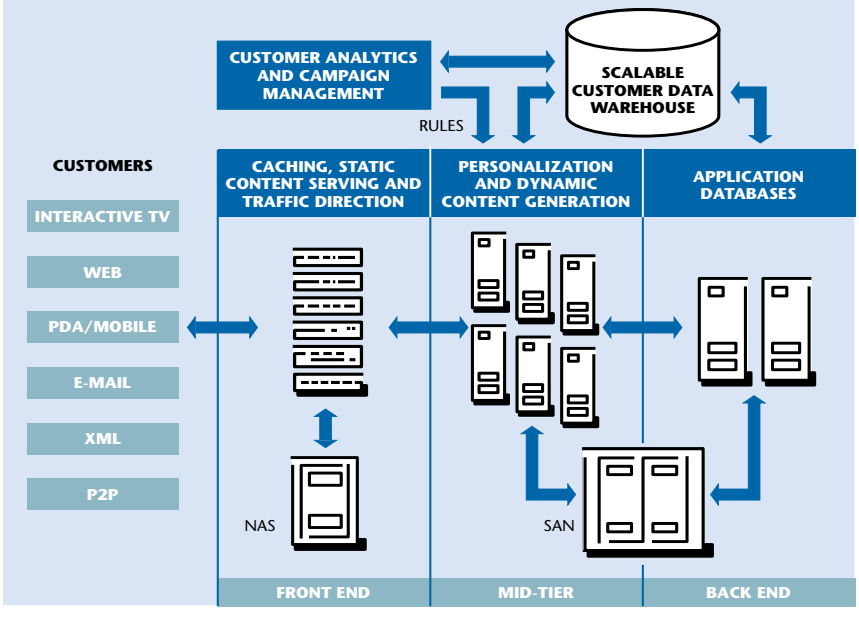

High-performance business logic

The online personalization requirements of eCRM range from dynamically customized Web pages to e-business applications that process XML-based forms in real time. Personalization can be determined by the characteristics of a customer's browsing session or by customer profiles that are stored in the back-end system (or cached toward the front of the n-tier architecture for maximum responsiveness). XML-based applications that customize customer transactions increasingly rely on mid-tier component-based business logic. Mid-tier components are driving higher utilization levels on the application servers due to increased computational and I/O requirements. Leading application server vendors that address the needs of this environment include Allaire, ATG, BEA Systems, IBM, Iona Technologies, Microsoft, Oracle and SilverStream.

Mid-tier eCRM applications can be memory-bound, due to their need for rapid access to large amounts of non-sequential data generated by real-time customer activities. The example scenario configures servers based on the Xeon Processor MP for the mid-tier. The processor's microarchitecture provides up to 64GB of high-speed, high-reliability memory (DDR and ECC), so processors will not idle excessively while waiting for memory access. The Xeon Processor MP's advanced three-level cache design further reduces latencies with up to 1MB or more of integrated, very high-speed cache memory. The effects of the advanced caching and performance-balanced microarchitecture allow mid-tier eCRM applications to maintain very high transaction levels during peak workloads—hundreds of thousands of transactions per minute. This means that more customers can be supported per server without degradation.

Back-end resources for eCRM

For back-end eCRM data management and data-intensive analytical programs, 4- and 8-way servers based on Xeon Processor MP or clustered server configurations of

up to 32 or more processors are deployed in the database tier. NAS latencies are generally too great for this tier, so SAN storage is used for high-bandwidth disk services that will not bottleneck Xeon processors during high-speed block I/O transfers. Using the latest distributed database software from IBM, Microsoft, Oracle and others, Xeon-based servers can cost-effectively scale datasets up to hundreds of gigabytes, with automatic failover for virtually nonstop availability. It's always a challenge to integrate customer systems that run on different closed, proprietary platforms. In contrast, servers with Xeon processors have the broadest possible O/S and application software support, which greatly facilitates integration of best-of-breed eCRM programs with existing and legacy systems.

As the amount of data in enterprise CRM continues to grow, there will be an increasing need for Itanium processor-based servers deployed for back-end databases. Itanium-based servers can quickly and efficiently feed data into the application servers.

OPEN COMPUTING SCENARIO 4

Extended supply chain management

Some businesses are working toward efficient internal procurement systems, while others are focused on full-blown supply chain management (SCM) networks that unite suppliers, customers and partners. In either case, interoperability is key to SCM and extended enterprise integration efforts, with the best tools being the ones flexible enough to accommodate both legacy and best-of-breed software resources.

Consequently, many supply chain applications (procurement, forecasting, order management, etc.) are using XML and n-tier solutions for integration. XML is used inside advanced component-based applications, and it also fills the ongoing need to talk to legacy systems in the back office and elsewhere. As shown in Figure F, this can be accomplished with an XML interface to an enterprise application integration (EAI) middleware system, or in some cases a direct XML interface to the ERP program itself.

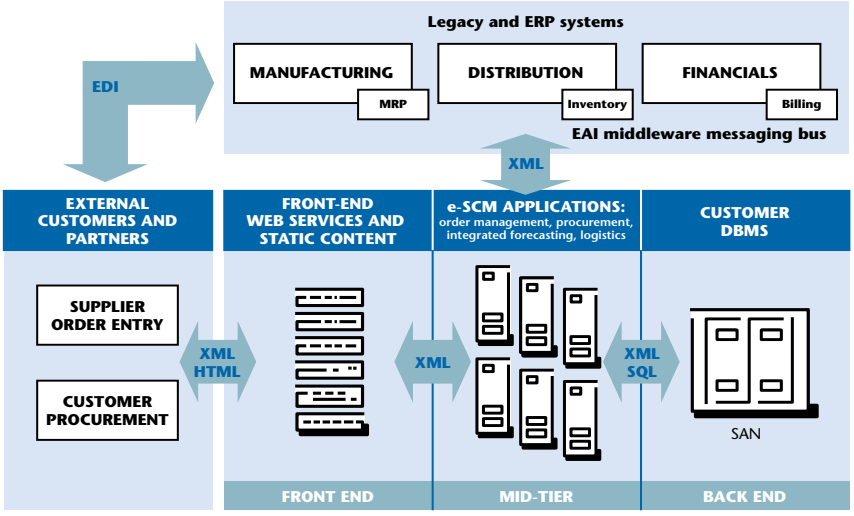

Figure F: **Seamless supply chain automation**
When SCM services are built with loosely coupled n-tier structures and high-performance IA-based servers, seamless transactional interactions with both customers and suppliers are possible.

In highly distributed SCM architectures, any weak points in the processing infrastructure can degrade user experience, jeopardizing supply chain efficiency. Consequently, the front end in this scenario is serviced by rack-mounted, dual-processor Xeon-based servers with 512MB or more of high-speed DDR/ECC memory. The Xeon processor's balanced microarchitecture allows the front end to accommodate large peaks and bursts of incoming supply chain messages. To ensure availability, Xeon-based servers are given redundant hot-swap disks, power supplies and fans. In high availability environments, the Xeon processor's built-in thermal sensing and unique integrated system-management features enable proactive fault diagnosis and avoidance for maximum front-end up time.

Mid-tier application server performance scaling

The modularity and loose coupling of the n-tier SCM model greatly facilitates asynchronous communications with external customers and suppliers via XML, SOAP, UDDI and similar data exchange standards. Asynchronous messaging allows traffic to adapt flexibly to the variability of IP network connections throughout the n-tier architecture and its Internet extensions. The distributed application environment meets enterprise needs for business-process extensibility, but the combined processing demands from external and internal users can load mid-tier application servers heavily. Performance scalability here is a key requirement.

In proprietary architectures, processing or I/O bottlenecks can lead to major "fork-lift" upgrades that replace entire processing systems at great expense. To avoid this, the sample scenario deploys a highly modular framework of 4- and 8-way servers based on the Intel® Xeon™ Processor MP. Each server supports four gigabytes of memory, to cope with escalating computational and I/O requirements in the mid-tier. Scalable clusters are also a cost-effective option here, when a very large number of transactions or SKUs are involved. Each mid-tier server can have its own database, or it can share SAN

database services in the back end.

Another characteristic of the mid-tier is the need to keep "state" information about user sessions. Stateful sessions are aware of customer histories, real-time inventory status, etc. When designing mission-critical SCM applications, it helps to limit session-management chores to a single tier (e.g., mid-tier), which avoids the need for state-management services at multiple layers of the architecture. Dynamic and "stateful" content that is typically served from the middle tier is a potential performance bottleneck. Consequently, it helps to keep customer profile information and other session data in the memory as much as possible. This is facilitated through the use of large server memories and in-memory databases. In this scenario, servers based on the Xeon processor ensure that the processing infrastructure can meet the challenges of the mid-tier SCM application processes as performance demands grow and evolve.

The SCM back end: Scale it right

Each enterprise has a unique need for processing power in the back end of the supply chain system. Ideally, IT architects should have the option of scaling processing power and storage, either by adding additional servers that scale out the back end or by using a few powerful scaled-up servers. Both of these approaches have their advantages.

In the sample SCM scenario, the back end is "scaled right" with a cluster of load-balanced Xeon-based servers. Servers in the cluster are configured for automatic failover to a hot back-up server when a cluster node fails or is taken off line for maintenance. Cascading cluster failover gives each machine more than one failover partner, ensuring extremely high levels of availability and multiple layers of redundancy. With support for up to 64 gigabytes of high-speed memory in each server and highly optimized transactional I/O, Xeon Processor MP-based clusters are suitable for many back-end supply chain requirements.

In cases where large cluster configurations of 8-way or greater Xeon Processor MP-based servers are not enough, Intel Itanium-based servers are deployed to fill out the high end of the datacenter processing continuum. The Itanium processor is used when supply chain automation requires multi-terabyte datasets, a huge flat address space and the highest levels of floating-point computation. As supply chain systems demand enterprise-wide data warehouses, real-time data mining, large in-memory databases and massive amounts of analytical processing, Itanium-based servers are configured as powerful centralized units or in scalable clusters. To support all this hardware firepower, key SCM in-memory database software such as i2 LiveCache and SAP-APO is being ported to the Itanium-based platforms.

Open computing architectures drive new business opportunities

Open computing and related IA macroprocessing principles are a reality today, thanks to Intel's technology leadership, worldwide industry support and market-tested best practices. Open computing lets enterprises choose from a whole universe of IT innovations, enabling the rapid pursuit of new opportunities that were often too expensive or too time-consuming with proprietary systems. To succeed in the open computing environment, the key is to make the right choices. First, choose open, not proprietary. Then, choose the right IA-based products and services by looking for solutions that deliver on the promise of Intel's macroprocessing vision: highly scalable and interoperable building blocks, advanced performance, major cost economies and data center-class reliability. The ultimate result is business agility that doesn't reduce operational stability and an open architecture that is widely deployable across the enterprise today.

MACROPROCESSING: KEY PRINCIPLES FOR IA-BASED COMPUTING

Highly scalable and interoperable building blocks	High-performance, cost-effective systems	Available and reliable server platforms	One of the world's largest solutions communities
Interoperable building blocks of versatile computing power. These are deployed via a "scale-right" strategy that allows both "scale-out" and "scale-up" techniques across the front end, mid-tier and back end of the enterprise.	Powerful and cost-effective enterprise computing with price/performance ratios typically several times greater than those of proprietary systems—and acquisition and ownership costs that are typically several times less.	Data center-class reliability, availability, serviceability and manageability built in IA processors from the ground up, and incorporated in leading data center operating and application software.	The Intel® e-Business Network (IeBN), with 55,000+ hardware, software and service providers, delivers the most innovative, interoperable and market-tested enterprise solutions possible.

For more information, please visit www.intel.com/ebusiness

GROUP PUBLISHING DIRECTOR: Joseph Braue
ASSOCIATE PUBLISHER, WEST: Pamala McGlinchey
DIRECTOR OF PROJECT MANAGEMENT: Karen White
PROJECT MANAGER: Lisa Broscritto
WRITER: Steve King
DESIGN TEAM: CMP Creative Services East
FOR MORE INFORMATION: jbraue@cmp.com, or 212-592-8214

CMP Custom Media Solutions