



www.intel.com




Can a server give your business a competitive edge?
Yes. Innovative server technologies in the Intel® Xeon™ processor family deliver performance, versatility, value and dependability.



Powerful and cost-effective servers to run the broadest range of applications.

Intel® Xeon™ Processor Family

Dual and Multi-processing Servers with Hyper-Threading Technology

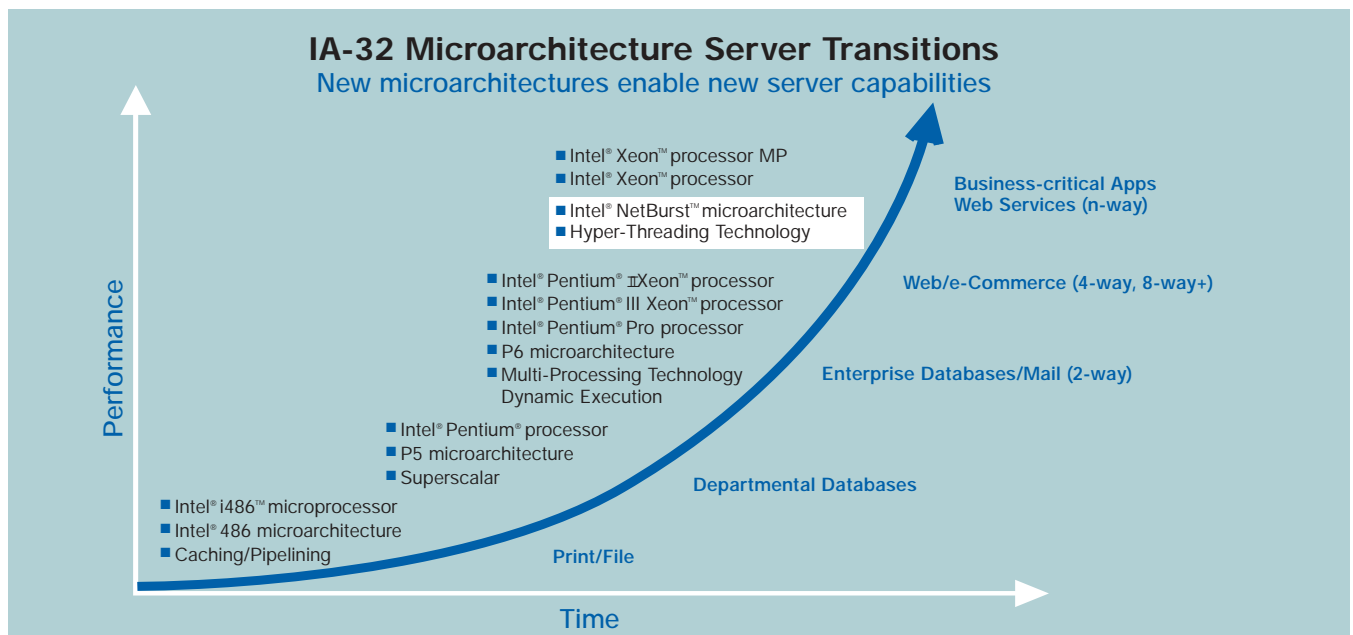


The Intel® Xeon™ processor family for servers features groundbreaking technology that increases computing power and throughput for today's Internet, e-Business, and enterprise server applications. The resulting increased transaction rates and reduced response times increase business productivity and enhance the overall customer enterprise computing experience. Innovative server technologies integrated into the Intel Xeon processor family give businesses a competitive edge.

Intel's commitment to providing businesses with industry-leading performance is long-standing, with each new processor introduction and microarchitecture shift representing significant performance advances. Starting with the Intel® i486™ microprocessor, Intel added advanced caching and pipelining to increase the number of instructions executed by the processor in a given clock cycle. Moving to the Intel® Pentium® processor, Intel developed superscalar techniques to process more instructions in parallel, further increasing execution throughput. The introduction of the

Intel® Pentium® Pro processor provided multi-processing technology – enabling multiple software threads to execute simultaneously across multiple processors. All these innovations from Intel have increased server capabilities along the way. The Intel Xeon processor family for servers represents the next leap forward in processor design and performance. The Intel® NetBurst™ microarchitecture is fully compatible with the Intel® P6¹ microarchitecture and runs all of today's IA-32 server software, but also adds several, advanced technologies. Along with Hyper-Threading Technology introduced in servers in Q1, 2002, and the Integrated Cache Subsystem, the Intel Xeon processor family for servers offers balanced platform improvements in memory and I/O throughput. Recent advances include continual frequency increases, iL3 cache of 1MB or 2MB in the Intel Xeon processor MP, and the introduction of a 533 MHz system bus for the Intel Xeon processor.

This powerful combination of innovative technologies results in a processor and server platform that is exceptionally well-suited and balanced for both the dual-processor (DP) and multi-processor (MP)-based server environments, and that is poised to bring greater levels of performance and price/performance benefits to a wide range of businesses.



1. Processors based on the P6 microarchitecture include the Intel® Pentium® Pro, Pentium® Ⅱ, Pentium® Ⅱ Xeon™, Pentium® III, and Pentium® III Xeon™.

The Intel Xeon processor family for servers with Intel NetBurst microarchitecture and Hyper-Threading Technology provides a strong foundation for further performance increases in DP and MP-based platforms, able to take on the most challenging workloads.

Intel® NetBurst™ Microarchitecture Brings Performance Headroom to Servers

Based on the Intel NetBurst microarchitecture, the Intel Xeon processor family fits seamlessly into today's industry-standard server infrastructure.

With the Intel NetBurst microarchitecture, Intel has significantly enhanced key P6 microarchitecture features, as well as introduced innovative features that help improve performance; such as higher frequency capabilities, a 400 MHz and introducing a 533 MHz system bus, Rapid Execution Engine, and an Execution Trace Cache. Other improved features include Streaming SIMD Extensions 2 (SSE2) and enhanced floating-point and multi-media units. These design features give server users the headroom they need to meet current and future performance needs.

System Bus – *(400 MHz for Intel Xeon Processor MP; 533 MHz for Intel Xeon Processor for DP-based servers)*

The system bus is designed to increase the throughput of multi-processing and multithreaded server applications and provide the necessary bandwidth for Hyper-Threading Technology when accessing system memory. It uses signaling and buffering schemes that allow for sustained data transfers. When one of the logical processors cannot find the data it needs in the Integrated Cache Subsystem, then the data must be transferred over the system bus from memory.

In systems with four or more processors, the system bus runs at 400 MHz and provides up to 3.2 GB/s bandwidth, which can be up to four times the previous-generation MP-based processor. For DP-based server applications, which typically are not so data intensive or cache hungry and derive a greater benefit from increased frequency and I/O, the system bus now runs at 533 MHz. This provides up to 4.27 GB/s bandwidth, over four times earlier generations of DP-based processors.

Higher Clock Frequencies – Relative to the P6 microarchitecture, the Intel NetBurst microarchitecture doubles the pipeline depth to a 20-stage pipeline. The longer pipeline allows the processor clock frequency to scale much higher than the P6 microarchitecture – offering nearly double the frequency at introduction. Higher frequencies lead to faster raw execution and increased performance for many server workloads.

Rapid Execution Engine – Runs Arithmetic Logic Units (ALUs) within the processor at twice the frequency of the processor core on certain instructions.

All integer-based instructions take advantage of this feature, with some of the operations done in half the time. Both Web and database servers can take advantage of this feature for improved user response times and increased transaction rates.

SSE2 – Extends the SIMD capabilities of Intel® MMX™ technology and the SSE instructions by adding 144 new instructions. New instructions enable software to execute tasks with fewer instructions in less time by taking advantage of data parallelism. The SSE2 feature can be used to provide significant performance improvements in front-end Internet server applications like media servers, secure transactions and next generation Web services.

Hyper-Threading Technology Introduces Thread-level Parallelism for Server Applications

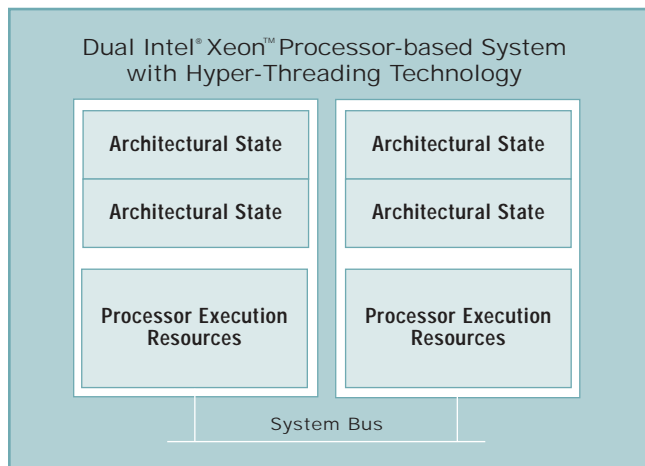
To gain a competitive edge, businesses must rapidly respond to increasingly demanding customer expectations. Businesses must also handle integration complexities more efficiently. The platform and applications that an enterprise deploys could make the difference between success and failure.

Traditionally, server applications are designed to take advantage of thread-level parallelism (TLP) by executing threads in parallel across multiple processors in a DP or MP-based server. The Intel Xeon processor family for servers adds an exciting innovation to TLP by supporting thread-level parallelism on a single processor, using Hyper-Threading Technology.

Hyper-Threading Technology is a form of simultaneous multithreading technology, where multiple threads from server applications can be run simultaneously on each processor within the server. This is made possible by duplicating the processor's architectural state, but continuing to provide only one set of shared processor execution resources. Each architectural state is called a "logical processor," and the operating system uses the two logical processors to track program execution for two different software threads, simultaneously.

The first logical processor tracks one thread, while the other logical processor tracks a second thread simultaneously. Because the two threads share one set of execution resources, the second thread can use execution resources that would otherwise be idle if only one thread was executing. The result is an increased utilization of the execution resources within each physical processor package.

This improvement in CPU resource utilization can yield higher processing throughput for multithreaded server applications. For example, one logical processor can execute a floating-point operation while the other logical processor executes an integer add and load operation.



Because most server applications are already multi-threaded, many of these applications can run "as is" on the Intel Xeon processor family for servers and provide performance improvements. Access to the second logical processor can be dependent on operating system version and software licensing configurations (consult your software vendor). Performance levels are highly dependent on the specific application, the quality of the application's multithreading, and other system dependencies. Some server applications have been shown to see significant improvements due to Hyper-Threading Technology.

Both DP and MP-based server environments can benefit from Hyper-Threading Technology. For DP-based server applications, such as front-end Web servers, user response times can increase. For MP-based server applications, such as those used for OLTP (Online Transaction Processing), Hyper-Threading Technology can offer increased transaction rates, and the ability to support a higher number of concurrent users.

Hyper-Threading Technology complements multi-processing systems by providing performance headroom. As more applications are optimized to take advantage of threading in the future, the Intel Xeon processor family can provide support for new e-Business and enterprise server capabilities.

While this Technology Overview focuses exclusively on servers, Hyper-Threading Technology is not restricted to servers. Intel has also introduced Hyper-Threading Technology into workstations, business desktops and consumer desktops. For a discussion of how Hyper-Threading Technology is implemented on workstations, and the benefits realized, please refer to the document "Intel's Hyper-Threading Technology: Delivering Performance Boosts for Workstations."

Integrated Cache Subsystem Improves Throughput for DP and MP-based Servers

The Integrated Cache Subsystem on the Intel Xeon processor family for servers is an innovative cache hierarchy designed specifically for the Intel NetBurst microarchitecture. It enables more efficient instruction execution and data access supporting greater processor throughput for server workloads.

Execution Trace Cache – This first-level instruction cache on the Intel Xeon processor family is very unique.

In the previous generation of processors, the instruction decoder was part of the execution loop. With the Intel NetBurst microarchitecture, the decoder is uncoupled from the execution loop by using the Execution Trace Cache. It stores decoded instructions (also known as micro-ops) and delivers them to the processor core at high speed. The Execution Trace Cache helps maximize throughput for the processor pipeline.

Level 1 Data Cache – The first-level data cache is designed to provide very fast access to frequently used data. Although previous processors offered an L1 data cache, the Intel Xeon processor's L1 data cache is much faster. This first-level data cache is tightly integrated with the Rapid Execution Engine, and just like the Rapid Execution Engine, is clocked at twice the processor's clock rate. This innovative design provides very high bandwidth to frequently accessed data needed by the Rapid Execution Engine.

Advanced Transfer Cache – (512KB)

The second-level cache delivers data and instructions to the processor core at high speed with larger cache lines.

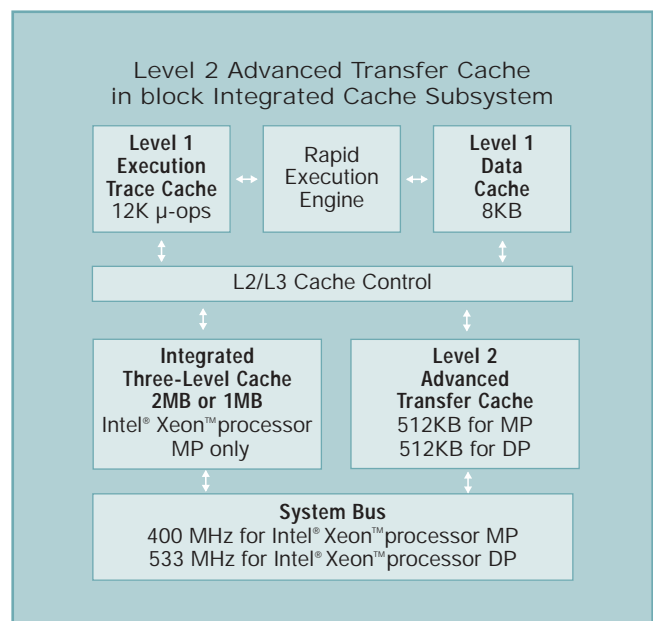
Because the Advanced Transfer Cache is clocked at the same rate as the processor core as faster processors are released, this cache speed increases correspondingly, providing high-speed access to key data. The larger cache line sizes also decrease average cache misses, thereby enhancing performance.

Integrated Three-Level Cache – (1MB or 2MB)

Available only on the Intel® Xeon™ processor MP, this large, integrated cache is specifically designed for mid-range and high-end server applications. The Intel Xeon processor MP represents Intel's first 32-bit processor to offer three levels of cache integrated into the processor package.

The Integrated Three-Level cache makes room for large instruction and data sets, providing a high-bandwidth path to memory, which reduces average memory latency and improves performance for mid-range and high-end server workloads.

For applications requiring very large data sets, the main memory in the system may be thought to be the cache for the local disk, which is the cache for the corporate database. ERP, SCM and CRM applications, which have planning modules, need access to large data sets for the analysis cycle. When the complete data set can be held in local memory, the analysis will execute much faster. Servers built for the Intel Xeon processor MP typically have large memory support, up to 32GB or 64GB, or up to 8x that supported in most DP-based platforms.



Cache memory on the processor chip is faster than main memory. It allows frequently accessed data to be available to the processors with less waiting. Some cache helps nearly all applications, but many applications are cache hungry because they require frequent data access cycles. Examples are database table searching or sorting which can be found in portions of CRM or SCM applications. In these cases, larger caches, and more levels of cache, improve performance by reducing the time to fetch the data. The Intel Xeon processor MP features three levels of integrated cache, with a large third-level cache of 1MB or 2MB.

Tightly integrated and balanced with the Intel NetBurst microarchitecture, the Integrated Cache Subsystem improves average cache latency and hit rates, and provides higher throughput when compared with P6 microarchitecture cache subsystems. This innovative, Three-Level cache architecture enhances the Intel Xeon processor family's data-handling capabilities for high-end server workloads.

Platform Enhancements on Intel Xeon Processor Family-based Servers

With Internet and enterprise server applications becoming larger and more complex, the increased bandwidth and throughput of the Intel Xeon processor family for servers would not be as effective without the accompanying platform improvements that create a truly balanced platform.

Double Data Rate (DDR) Memory – Is a technology derived from mature SDRAM technology. The secret to DDR memory's high performance is its ability to perform two data operations in one clock cycle – providing twice the throughput of SDRAM. Depending on server platform configurations, DDR memory can increase memory capacity up to four times greater than P6 microarchitecture-based MP servers.

DDR memory provides greater headroom and superior capabilities, increasing memory throughput and capacity for memory-intensive server workloads and in particular, enhances server capabilities for faster data mining.

PCI-X I/O – (*Peripheral Component Interconnect Extended*) Represents an extension to today's PCI bus technology that doubles the speed and amount of data exchanged between the processor and peripherals.

In MP-based platform configurations, PCI-X can increase I/O bandwidth from P6 microarchitecture-based servers by up to six times, and increase the number of I/O channels as well. PCI-X helps Intel Xeon processor family-based servers increase performance for high-bandwidth I/O devices, such as Gigabit Ethernet cards, Fiber Channel, Ultra3 SCSI, and processors that are interconnected.

The Intel® E7500 and Intel® E7501 Chipsets –

The Intel® E7500 chipset is optimized for dual-processing Intel Xeon processor-based servers at 400 MHz system bus and the Intel® E7501 chipset is optimized for dual-processing Intel Xeon processor-based servers at 533 MHz system bus. These chipset designs deliver maximized system bus, memory, and I/O bandwidth to enhance performance, scalability, and end-user productivity, while providing a smooth transition to next-generation server technologies.

Manageability and Reliability Enhancements –

Intel Xeon processor-based platforms providing increased availability, maximizing your server's uptime. One example is Hot-Plug capability, where a system I/O device or memory module can be changed "on the fly," without powering down the system. Other examples are features including error correction code, Intel® x4 Single Device Data Correction² memory and memory mirroring, where the platform can detect memory errors and correct them, disable the failing devices and switch to redundant modules, all while the server continues to operate. These systems can also offer enhancements in platform management software that allow the system administrator to more effectively manage the servers remotely. Intel Xeon processor-based servers are inherently more dependable offering platform-level manageability features once found only in mainframe servers.

2. In a x4 DDR memory device, the Intel® x4 Single Device Data Correction (x4 SDDC), provides error detection and correction for 1,2,3, or 4 data bits within that single device and provides error detection, up to 8 data bits, within two devices.

Intel Xeon Processor Family

A Winning Technology Combination for e-Business and the Enterprise

The Intel Xeon processor family for servers provides ample support for evolving your infrastructure to fit your business needs. With headroom for larger workloads and application enhancements, servers based on the Intel Xeon processor family with Intel NetBurst microarchitecture and Hyper-Threading Technology can scale right along with your growing business.

Multi-processor (MP) and dual-processor (DP)-based servers equipped with these processors offer industry-leading performance and offer the best value for businesses today. The Intel Xeon processor family for servers is fully compatible with existing server solutions. Enterprise and e-Businesses can increase their business-competitive advantage by moving to both the dual-processing and multi-processing versions of the Intel Xeon processor family.

Dual-processing servers based on the Intel Xeon processor deliver compute power at unparalleled value and versatility for departmental servers, Internet infrastructure (for example, Web-caching, streaming media, search engine and security), mail server and Small & Medium Business (SMB) server applications. The combination of the Intel NetBurst microarchitecture and Hyper-Threading Technology provides outstanding performance and headroom for peak server workloads, resulting in faster customer response times and support for more users.

The Intel Xeon processor MP has been designed exclusively for MP environments to support the demands of key enterprise applications. Features in the Intel Xeon processor MP match key enterprise application characteristics. Many applications execute in parallel operations, which makes great demands on the processor. Adding multiple processors improves solution performance. The Intel Xeon processor MP is specifically designed for multi-

processor operation in 4P, 8P and up to 32P servers. Multi-processor servers based on the Intel Xeon processor MP provide industry-leading compute power and throughput for e-Business and enterprise server applications; such as Customer Relationship Management, Supply Chain Management, Middleware and Application Servers, Collaboration/Mail, Media Servers, Site Management, Enterprise Databases, and Business Intelligence. The combination of the Intel NetBurst microarchitecture, Hyper-Threading Technology, and the Integrated Three-Level cache delivers outstanding throughput and headroom for large server workloads, resulting in increased transaction rates, faster response times, support for more users, as well as enhanced availability and application scalability.

The synergistic combination of the Intel NetBurst microarchitecture, Hyper-Threading Technology, and Integrated Cache Subsystem provided by the Intel Xeon processor family for servers offers the level of performance that rigorous DP and MP-based server environments require.

Servers equipped with Intel Xeon processors are dependable, versatile in being able to Scale Right (scale up or scale out) and deliver industry-leading performance.

When your business needs to maximize system throughput, increase transaction rates, and improve user response times – all at a business value that cannot be matched – ask your system and solutions providers today about servers using the Intel Xeon processor family.



Intel® Xeon™ Processor Family for Servers

- Performance
- Dependability
- Value
- Versatility



Information in this document is provided in connection with Intel products. No license, express or implied, by estoppel or otherwise, to any intellectual property rights is granted by this document. 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, or life sustaining applications. Intel may make changes to specifications and product descriptions at any time, without notice.

Designers must not rely on the absence or characteristics of any feature or instructions marked "reserved" or "undefined." Intel reserves these for future definition and shall have no responsibility whatsoever for conflicts or incompatibilities arising from future changes to them.

The Intel Xeon Processor and Intel Xeon Processor MP may contain design defects or errors known as errata which may cause the products to deviate from published specifications. Current characterized errata are available on request.

Copyright © 2002, Intel Corporation. All rights reserved. Intel, the Intel logo, Intel Inside, Intel NetBurst, Intel Xeon, i486, MMX and Pentium are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

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