Planning Considerations for

Intel Extended Memory 64 Technology

on Servers and Workstations

Intel® Extended Memory 64 Technology (EM64T) is designed to expand the memory addressing capabilities of the 32-bit Intel architecture (IA-32) while the IA-32 systems continue to run the vast number of existing 32-bit x86-based applications. This article examines 32-bit, x86-based workstation and server applications whose performance may be improved by implementing EM64T. In addition, this article defines specific application scenarios in which the 64-bit Intel Itanium® architecture may help provide better performance than x86-based Intel systems running EM64T.

BY JOHN COOMBS AND JOHN FRUEHE

The recently announced Intel® Extended Memory 64 Technology (EM64T) is designed to increase the memory addressing capabilities of the 32-bit Intel architecture (IA-32) while the IA-32 systems continue to run existing 32-bit, x86-based applications. Intel Xeon™ and Pentium® processor–based systems that incorporate EM64T are expected to become available in the second half of 2004. EM64T-enabled systems are designed to fill the memory addressability gap between highly concurrent 64-bit Itanium® processor–based database and high-performance computing (HPC) servers and Pentium processor–based desktop systems and servers running 32-bit office productivity, Web, and file serving applications. EM64T can help enable Dell™ Precision™ workstations running Linux® operating systems with 64-bit addressing or Microsoft® Windows® XP to create larger and more detailed models and analyze or simulate more complex environments. Dell PowerEdge™ servers incorporating EM64T have the potential to help increase performance on memory-intensive applications as well as to help provide rich, layered commercial and information services to client and Web environments, which have become much more demanding and sophisticated over the last ten years.

The contrast between 32-bit and 64-bit architectures

The memory address ceiling of mainstream Intel processor–based architecture was last raised in 1985, when x86 memory addressing was increased from 16 bits to 32 bits. Before 1985, the 16-bit Intel 80286 processor allowed 2^{16} bytes (representing 65,536 memory locations) to be directly addressed—enough to catalog the population of a small city. In 1985, the 32-bit 80386 processor, enabled by the 32-bit Microsoft Windows NT® and Microsoft Windows 95 operating systems, was designed to address nearly 2^{32} bytes (representing roughly 4.3 billion memory locations)—enough to catalog every person in the world based on the reported global population in 1975.

Bits	Binary	Number of memory addresses (equivalence in bytes)	Relative scale
4	24	16	
9	28	256	
10	210	1,024 (1 KB)	
16	216	65,536 (64 KB)	Population of a small city
32	232	4,294,967,297 (4 GB)	Reported population of the world in 1975
64	264	18,446,744,073,709,600,000 (18 exabytes)	Estimated number of grains in a cubic mile of sand

Figure 1. Differences in scale between architectures

Understanding how addressability can scale may be helpful when considering what applications can best take advantage of 64-bit architecture. Although 64-bit architecture might initially appear to only double memory capacity, moving from 32 bits to 64 bits actually squares the available memory, providing a theoretical limit of 4.3 billion times 4.3 billion addressable memory locations. The resulting 18 exabytes[1] of memory addressability theoretically provides the capability to individually name each grain in a cubic mile of sand (see Figure 1).

Applications that may benefit from 64-bit memory addressing

In Intel Xeon, Itanium, and certain Pentium processors, 64-bit addressing is designed to provide a large, flat memory space that is easily addressed by several memory-intensive applications (see Figure 2). Such applications include databases; digital content creation (DCC) and mechanical computer-aided design (MCAD); Monte Carlo simulation; geophysical analysis and mechanical computer-aided engineering (MCAE); and directory, e-commerce, and messaging servers.

Databases. A very large memory address space can allow database administrators to cache all—or a significant part—of an Oracle® or Microsoft SQL Server™ database in buffer RAM. This approach can enable many more in-memory processes and help provide greater scalability using 64-bit memory addressability as compared to 32-bit addressability. Increased addressability and caching also may help provide improved online transaction processing (OLTP) performance on 64-bit Dell PowerEdge servers, because it helps enable both faster context switching and reduced I/O.

Digital content creation and mechanical computer-aided design. A very large memory address space can help enable designers in the areas of DCC and MCAD to create large 3-D models—in the 6 GB to 8 GB range—and virtual environments on platforms such as the Dell Precision workstation. Compared to 32-bit platforms, the increased memory addressability of 64-bit technology could also enable designers to assemble multiple components—for example, putting a virtual piston into a virtual engine block—quickly and without sacrificing detail simply to fit within the memory limits imposed by a 32-bit Windows process.

Monte Carlo simulation. Server and workstation applications—such as financial risk management and molecular modeling—that employ Monte Carlo techniques to simulate and test a broad range of very detailed outcomes are likely to benefit from the detail and range made possible by 64-bit addressing.

Geophysical analysis and mechanical computer-aided engineering. The 64-bit memory addressing on Intel Xeon and Itanium processor–based servers and HPC clusters can help geophysicists analyze complex environments in much greater detail than is possible when using 32-bit memory addressing. For example, such increased addressing capability may help to improve the likelihood of finding and recovering oil and natural gas.

Compared to 32-bit addressing, the 64-bit addressing capabilities of Itanium processors on Dell PowerEdge servers enable dense meshes for a detailed analysis of airflows around—and effects upon and between—structures in automotive, aerospace, and related industries that rely on MCAE applications. The 64-bit addressing also can help enhance visualization of server data on—and computational steering of server applications from—Intel Xeon and Pentium 4 processor–based Precision workstations, allowing greater interactivity and increased attention to detail than 32-bit variants.[2]

In Intel Xeon, Itanium, and selected Pentium processors, 64-bit addressing is designed to provide a large, flat memory space that is easily addressed by several memory-intensive applications.

Directory, e-commerce, and messaging servers. Directory servers on large corporate intranets can typically store and quickly provide access to more information about the user population when using a large, flat virtual address space. E-commerce servers may benefit from the faster encryption available with dedicated 64-bit registers as well as the larger memory and caches available with upcoming 64-bit Intel Xeon CPUs featuring EM64T, as compared with previous-generation 32-bit Intel Xeon processors. Messaging servers can usually be more easily consolidated and scaled when directories and indexes are enabled by 64-bit

[1] One exabyte = 2^{60} (1,152,921,504,606,846,976) bytes, or 1,024 petabytes.

[2] For more information, see "Windows XP 64-Bit Edition Version 2003 for Itanium-based Systems Technical Overview," Microsoft Corporation, http://download.microsoft.com/download/2/e/b/2eb773ec-944f-4000-a475-d8ee6834af62/WindowsXP64-BitforItanium.doc.

Memory used by example application and data set	64-bit workstation or HPC application example	Vertical market
3.2 GB	Finite element analysis: 5,000 elements with four degrees of freedom	MCAE
6 GB	Automated parking garage design: 500,000 components	MCAD
10 GB	Electronic design: 1,000,000-gate device	Electronic computer-aided design (ECAD)
26 GB	Video editing: typical two-hour movie	DCC
45 GB	Video rendering: one frame of special effects	DCC
56 GB	High-end engineering analysis: 1,000 time steps	MCAE
60 GB	Seismic analysis: multiple shots of one area	Geophysics
80 GB	Geographic survey: determining wireless paths through the city of Vienna with a 3-D geographic information system (GIS)	GIS

Figure 2. Examples of applications and data sets that can take advantage of 64-bit addressing

technology to scale well beyond the 32-bit addressability limit of roughly 4.3 billion messages.

The persistence of 32-bit applications

Many 32-bit applications are not pushing the upper limits of their 4 GB physical or virtual address space. Such applications can continue to run in 32-bit compatibility mode on a 64-bit platform, or in native mode on a 32-bit platform, for the foreseeable future without running out of address space and crashing. In addition, Web, file, and print servers may continue to run comfortably with a 32-bit operating system on 32-bit hardware, or in 32-bit mode with a 64-bit operating system and CPU, for many years to come. Similarly, very few, if any, of today's 32-bit consumer applications running on desktop or notebook PCs test the memory limits of a 2 GB to 3 GB virtual Windows process or push the 4 GB physical memory address limitation.

Some sophisticated users and independent software vendors (ISVs) use Intel and Microsoft extensions to the IA-32 platform such as Physical Address Extension (PAE) or Address Windowing Extension (AWE) to circumvent limitations of 32-bit addressing. These segmented memory extensions, which are similar to those that helped enable 1 MB memory addressing on the 16-bit Intel 80286 processor, work in some application environments, such as SQL Server, and not others—for instance, certain MCAD and MCAE applications. PAE and AWE are neither straightforward nor easy to use, and tend to be vendor- and application-specific. In addition, the processing power required to access and manage memory addresses beyond 4 GB can limit scalability because the processing cycles that PAE and AWE consume to address the additional memory can degrade application response as well as system performance. In short, PAE and AWE remain interim solutions on the path to industry-standard 64-bit workstations and servers.

Applications on EM64T versus Itanium servers

As EM64T becomes available in the second half of 2004 and beyond, distinctions may be drawn between the applications that are appropriate for EM64T and the applications that are appropriate for Itanium processor–based servers. Intel Xeon and Pentium 4 processor–based servers supporting EM64T are designed to run the spectrum of 32-bit applications. Next-generation Intel Xeon processors with EM64T, fast frontside buses, and fast clock cycles are designed to run 32-bit applications that require high native 32-bit performance. At the same time, they are intended to support the transition from 32-bit to 64-bit for Windows and Linux server applications by enabling 32-bit and 64-bit programs to run on the same hardware platform.

Cost- and performance-sensitive server applications that benefit from large data caches may also be appropriate for EM64T. Mixed application environments—for example, a 64-bit database running with a transaction monitor, virus checker, and other middleware not yet ported from the 32-bit environment—may also perform better on an EM64T-based server than on a similar 64-bit Itanium system.

For enterprises deploying applications on Itanium platforms, the application is often the first choice in the procurement cycle, followed by the operating system, and finally, the hardware. Organizations tend to buy an Itanium processor–based server to run a single application such as SAP or Microsoft SQL Server. In contrast, infrastructure servers, which tend to be Intel Xeon processor–based, are typically standardized at the organizational level and such servers are purchased as needed and deployed in various locations within the data center. For high-end and midrange servers running applications that are migrated from legacy 64-bit RISC systems, 32-bit x86 compatibility can be less critical.

IT managers deploying performance-oriented, 64-bit concurrent floating-point and database applications may want to consider Itanium-based systems like the two-processor Dell PowerEdge 3250 or the four-processor Dell PowerEdge 7250. Because many large e-commerce, enterprise resource planning (ERP), supply chain management (SCM), data warehousing, and business intelligence applications are largely concurrent, these applications are well suited to the Itanium architecture. Many Java™ Virtual Machines (JVMs) used in e-commerce, such as BEA® WebLogic JRockit™, can also benefit from just-in-time compilation, predication, and speculative loading—all features of the Itanium architecture. High-end, floating-point-intensive HPC applications such as reservoir simulation and computational fluid dynamics may also benefit from improved performance by taking advantage of the large address

space, large caches, and concurrent throughput of an Itanium processor–based server.

The migration process from 32-bit to 64-bit

Best practices suggest that IT managers considering a migration from 32-bit to 64-bit architecture—or from RISC systems to 64-bit industry-standard systems—first perform a thorough needs analysis. Administrators must also work with their operating system and application vendors to discuss the timing of 64-bit product and feature availability, so they can plan their 64-bit hardware acquisitions to coincide with important software releases as 64-bit Intel Xeon, Pentium 4, and Itanium processor–based environments rapidly mature.

Dell recommends that IT managers contact their operating system and application vendors early in the planning process when considering a 64-bit migration.

A thorough needs analysis can help determine which applications can take advantage of 64-bit architecture immediately, which applications can wait, and which applications will likely not migrate in the foreseeable future, if at all. For example, large database environments—which can benefit from the concurrent performance and large address space of a 64-bit Itanium processor–based platform—should target 64-bit platforms for additional performance and scalability. Application servers, which have less-stringent performance requirements, are more likely candidates for an EM64T migration. Meanwhile, client systems in even the largest and most complex database environments may have enough address space to perform adequately on the 32-bit platform.

Administrators may discover that large address spaces in themselves are not always advantageous. Benefits across various database environments are not uniform. Some data warehousing and decision-support applications perform full-table scans—bypassing data stored in buffer RAM—or invalidate tables frequently. In such cases, buffering data in memory may result in a performance penalty. Yet data warehousing and decision support may run considerably faster in a 64-bit Itanium environment because of the concurrency enabled by Itanium servers when they do not buffer a great quantity of data.

Dell recommends that IT managers contact their operating system and application vendors early in the planning process when considering a 64-bit migration. Many software vendors' 64-bit plans are just emerging, and others' plans are changing with the recent announcement of EM64T-based solutions. Compiled versions of either EM64T-based or Itanium processor–based 64-bit applications are required, so organizations may also continue to run 32-bit versions of applications while executing a planned migration to an EM64T-based or Itanium processor–based server.

Organizations must also be prepared for potential migration challenges, particularly with respect to the recently announced and yet-to-be-shipped EM64T-based systems. When 64-bit EM64T-based workstations and servers are released, both 32-bit and 64-bit versions of Linux and 32-bit versions of Windows will be available—but not 64-bit production versions of Windows. Although 32-bit operating systems will run on 64-bit hardware, they are designed to run only 32-bit applications. To run on 64-bit hardware with a 64-bit operating system, 32-bit and 64-bit applications will both require 64-bit device drivers. Drivers for the most pervasive devices may become available toward the end of 2004 and other drivers will arrive over time, in a manner analogous to that experienced with Windows NT in the first months after it shipped.

Important considerations for 64-bit migration

Intel Extended Memory 64 Technology is designed to greatly expand the memory addressing capabilities of x86-based workstations and servers while they run 32-bit x86-based applications. Enterprises running memory-intensive design, simulation, or analysis applications on 32-bit Intel processor–based workstations may want to consider migrating to 64-bit EM64T-based workstations beginning in the second half of 2004. IT managers with memory-intensive database, e-commerce, and messaging applications running on Intel Xeon or Pentium 4 processor–based servers may prefer to begin procuring EM64T-compatible hardware as soon as it is available and then plan a migration of operating systems and applications as these components become available from the software vendors. The robust support of Intel, Microsoft, and Dell for a mixed 32-bit and 64-bit environment can help smooth the migration to 64-bit architecture in both workstation and server environments. Managers of RISC-based UNIX® server environments, particularly servers dedicated to a single, large application, may consider migrating to Itanium processor–based servers because applications are available today.

John Coombs (john_coombs@dell.com) is a product planning senior consultant in the Dell Precision workstation business unit. John has an M.B.A. from the University of Chicago and a B.A. in Economics from Brown University, and has been in the technology field for 23 years.

John Fruehe (john_fruehe@dell.com) is a marketing strategist for the Dell Enterprise Product Group. He has worked at Dell for more than seven years; prior to that, he was at Compaq and Zenith Data Systems. John has a B.S. in Economics from Illinois State University and has been in the technology field for 13 years.

FOR MORE INFORMATION

Intel EM64T:
http://www.intel.com/technology/64bitextensions