

AMD and Novell®: The Best-Engineered Linux Foundation for Enterprise Computing

EXECUTIVE SUMMARY

As the x86 industry's first native quad-core processor, Quad-Core AMD Opteron™ processors are designed to significantly improve performance in physical and virtual environments. From the beginning, SUSE® Linux Enterprise Server from Novell® has been optimized to take advantage of features unique to these processors, creating a stable, long-term solution that can meet critical enterprise demands for exceptional energy efficiency, advanced virtualization technology, and outstanding computing performance.

Quad-Core AMD Opteron processors deliver essential performance-enhancing and energy-saving elements like a shared 2-MB L3 cache, split planes for independently controlling processor and northbridge voltage and frequency, and physical memory addressing up to 256 TB. SUSE Linux Enterprise Server leverages AMD Virtualization™ (AMD-V™) technology with Rapid Virtualization Indexing—a two-dimensional approach to page tables that transfers memory addressing from the hypervisor to the hardware—to virtually run high-memory management workloads like databases and terminal services applications at near native performance. The hardware-controlled, independent clock speeds unique to Quad-Core AMD Opteron processors can also create significant performance improvements in physical and virtual SUSE Linux Enterprise Server environments. In addition, the updated GNU Compiler Collection (GCC) provided with SUSE Linux Enterprise Server takes advantage of several changes to the AMD processor microarchitecture, including advanced bit manipulation instructions and four new SSE4a instructions.

Future enhancements to SUSE Linux Enterprise Server are anticipated to unleash several other features already built in to Quad-Core AMD Opteron processors, including performance monitoring techniques that use instruction-based sampling to validate chips and gigabyte memory pages that can improve performance in heavy random access workloads like enterprise database applications.

CONTENTS

EXECUTIVE SUMMARY1

A LONG-TERM PARTNERSHIP3

 Why AMD and Novell? In a Nutshell.....3

IMPROVEMENTS FROM THE SILICON UP.....4

 Native Quad-Core Design.....4

 AMD Balanced Smart Cache.....4

 Increased Physical Memory Addressing.....4

OUT-OF-THE-BOX HARDWARE CERTIFICATIONS4

OPTIMIZING PERFORMANCE AND INCREASING POWER SAVINGS6

 Better Performance for Virtualized Environment.....6

 A Trip Down Memory Lane.....6

 Shadow Page Tables: High Overhead, Slow Performance.....6

 Rapid Virtualization Indexing: Blazing Fast Performance with SUSE Linux Enterprise Server7

 Industry-Leading Power Efficiency with Enhanced AMD PowerNow!™8

 Reduced Energy Consumption9

 Instruction-Level Performance Improvements9

 Advanced Bit Manipulation (ABM)9

 New SSE4a Instructions.....9

LOOKING TO THE FUTURE.....11

 Instruction-Based Sampling11

 Gigabyte Memory Pages.....11

 AMD and Novell = Investment Protection11

FINAL ANALYSIS13

A LONG-TERM PARTNERSHIP

AMD and Novell share a history of collaborating to bring breakthrough technology to your enterprise. SUSE Linux Enterprise Server is often the first Linux distribution to enable advanced features in AMD processors. SUSE Linux Enterprise engineers are among the first to develop code enhanced for new AMD processors: SUSE Linux Enterprise Server immediately supported AMD PowerNow!™ technology when it was introduced in 2005, and followed with strong support for AMD Virtualization™ (AMD-V™) technology and Xen virtualization in 2006. With the advent of Quad-Core AMD Opteron™ processors in 2007, Novell again stepped up by updating SUSE Linux Enterprise Server for advanced features like Enhanced AMD PowerNow! and AMD-V technology with Rapid Virtualization Indexing.

This strong technology partnership continues today: every year AMD and Novell collaborate on close to one hundred feature topics to fine tune how SUSE Linux Enterprise Server works with AMD Opteron processors. These combined engineering efforts give you the peace of mind that comes from building on a robust computing foundation that excels in performance and value.

WHY AMD AND NOVELL? IN A NUTSHELL...

Novell provides a robust, scalable, and flexible infrastructure that takes advantage of the most advanced features of AMD technology, so IT departments can fully realize the benefits of Quad-Core AMD Opteron processors. That's because AMD and Novell collaborated to fine tune the most advanced features of Quad-Core AMD Opteron processors for SUSE Linux Enterprise Server 10 environments, including:

- AMD Balanced Smart Cache—Helps increase multi-threaded application performance through three levels of dedicated and shared cache
- 48-bit physical addressing—Enables servers to have more physical memory
- AMD-V technology with Rapid Virtualization Indexing—Designed to significantly improve virtual machine performance through hardware-based memory management
- Enhanced AMD PowerNow! technology—Delivers performance-on-demand while reducing power consumption

- New advanced bit manipulation and SSE4a instructions—Improve instruction-level performance
- Reliability, availability, and serviceability (RAS) features—Help reduce hardware failure that could corrupt data or bring down servers

SUSE Linux Enterprise Server 10 SP 1 is the first release to fully support the features available in Quad-Core AMD Opteron processors. This release can help improve the performance of many guest operating systems in virtualized environments by including paravirtualized device drivers (guest operating system drivers designed to work with a virtual device rather than a real hardware device). The result includes:

- Unmodified SUSE Linux Enterprise Server 9 and 10 guest operating systems run with near native performance in virtual environments.
- Unmodified Windows® and Red Hat Linux guest operating systems experience significant performance improvements in virtual environments.

SUSE Linux Enterprise Server 11 is expected to bring additional optimization for features found in Quad-Core AMD Opteron processors:

- Support for the newest performance monitoring infrastructure, including better kernel support for AMD Opteron processor features
- Low power consumption as monitored against SpecPower2007 and other benchmarks
- More sophisticated scheduler strategies to enhance power savings through split plane technology from AMD

These upcoming enhancements demonstrate once again AMD and Novell's commitment to providing a robust, scalable, and flexible infrastructure for your IT solutions. In this paper, you'll get an in-depth look at how SUSE Linux Enterprise Server can help optimize the power-saving and performance-enhancing features found in Quad-Core AMD Opteron processors.

IMPROVEMENTS FROM THE SILICON UP

When you run Novell SUSE Linux Enterprise Server on Quad-Core AMD Opteron™ processors, expect excellent computing performance and exceptional energy efficiency. Fundamental design characteristics and advanced planning inherent in AMD processor architecture ensure that SUSE Linux Enterprise Server systems can experience significant performance benefits.

NATIVE QUAD-CORE DESIGN

Quad-Core AMD Opteron processors feature a native quad-core design. Instead of “gluing” two dual-core processors together and calling it quad-core, AMD designed its processors with four independent cores on a single processor die. AMD Direct Connect Architecture features multi-core technology on a single piece of silicon with the cores directly connected to the I/O subsystem, the memory controller, and other processors in the configuration. The result is highly efficient communication for improved performance and performance-per-watt with reduced latency.

Native multi-core technology from AMD enables virtual machines to efficiently share I/O and memory resources. Virtualization puts more stress on how memory is managed, but the AMD architecture is very efficient with memory management, particularly because the integrated memory controller resides on the processor itself. This design reduces latency and improves virtual machine performance.

AMD BALANCED SMART CACHE

AMD Balanced Smart Cache, found in Quad-Core AMD Opteron processors, includes a dedicated L1 and L2 cache on each core. It also includes a new, shared high-speed L3 cache, which is designed to improve core efficiency for better support of multi-threaded virtualization environments and to avoid thrashing. Because data enters the smaller, faster L1 cache and flows to the larger outer L2 and L3 caches, loads and stores happen more rapidly, and the L2 cache becomes an extension of the L1 cache (Figure 1).

In this cache structure, all data items fetched from memory are stored first in the L1 cache, which splits 128 KB equally between instructions and data. By leveraging a built-in hardware pre-fetch device, the processor can detect repeating data access patterns and load upcoming data items into L1 cache. If an item in L1 cache is not needed, it is evicted to that core's dedicated 512 KB L2 cache, which is

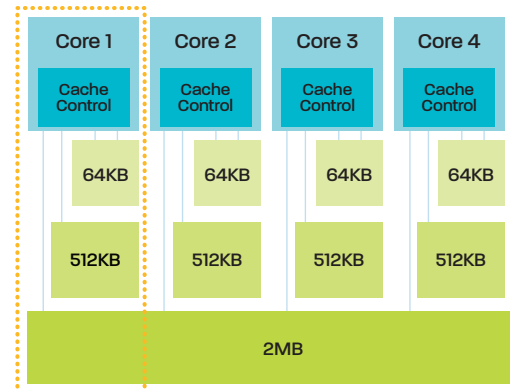


Figure 1: AMD Balanced Smart Cache

sized to accommodate the majority of today's working sets. Placing this “victim” data item in the L2 cache reduces access latency because it can be loaded back to the L1 cache very quickly.

Dedicated L1 and L2 cache per core helps performance of virtualized environments and large databases by reducing cache pollution associated with a shared L2 cache. The shared 2 MB L3 cache helps maximize the efficiency of the cache hierarchy by sharing data between cores while helping reduce latency to main memory. Quad-Core AMD Opteron processors have a sharing-aware replacement policy: if the data is still being used, leave it alone; if it's not, get rid of it. The result is a smart cache design that improves core efficiency for better support of multithreaded software environments like those found in SUSE Linux Enterprise Server installations.

INCREASED PHYSICAL MEMORY ADDRESSING

AMD always develops its processors with an eye on the future, and the increased memory addressing in Quad-Core AMD Opteron processors is no exception. Physical memory addressing increased from 40 bits to 48 bits, resulting in the ability to address up to 256 TB of physical memory. While the maximum memory available in a machine currently on the market is about 1 TB, SUSE Linux Enterprise Server systems running on these quad-core processors will be ready to take advantage of increased memory capacity as soon as it is available.

OUT-OF-THE-BOX HARDWARE CERTIFICATIONS

When you choose AMD and Novell, you can rest easy knowing that you're getting a server combination certified by both companies. AMD announced at LinuxWorld San

Francisco in August 2007 the availability of AMD Validated Server platforms certified for SUSE Linux Enterprise Server. This means system builders can build and ship systems with SUSE Linux Enterprise Server and carry the "YES Certified" sticker as long as they use endorsed parts.

AMD performed Novell YES Certified testing on barebones platforms from Supermicro Computer, Tyan Computer Corporation, and Uniwide Technologies, so systems builders can construct machines from those vendors and be certified to run SUSE Linux Enterprise Server. In all, 13 different platforms from these vendors have been certified.

OPTIMIZING PERFORMANCE AND INCREASING POWER SAVINGS

AMD and Novell rigorously tested Quad-Core AMD Opteron™ processors with SUSE Linux Enterprise Server to get the most out of key features customers want and need, including:

- A hardware solution to memory management in virtualized environments
- Enhanced power-saving techniques for multi-core processors
- Instruction-level performance improvements

SUSE Linux Enterprise Server is the first Linux distribution optimized to run hand-in-glove on Quad-Core AMD Opteron processors. This stable foundation provides long-term performance improvements for the workloads you want to run, be they virtualization, database, IT infrastructure, or compute-intensive.

BETTER PERFORMANCE FOR VIRTUALIZED ENVIRONMENT

Virtualization technology delivers great benefits like lowering power, hardware, and space requirements through consolidated workloads, running multiple operating systems on one physical server, and providing fast disaster recovery. However, virtualization environments have been plagued with performance challenges, especially in terms of memory management.

Linux-based virtualization environments use a software approach to memory management called shadow page tables. This approach requires many world switches, or transitions between the guest operating system and the hypervisor, which creates high overhead in the hypervisor and reduces the performance of virtualized workloads—especially memory-intensive ones like databases.

To address this challenge, Quad-Core AMD Opteron processors provide AMD Virtualization™ (AMD-V™) technology with Rapid Virtualization Indexing (RVI), a hardware-based solution to memory management in virtualized environments. AMD-V with RVI increases the performance of many virtualized applications by allowing virtual machines to directly manage memory with less hypervisor intervention and associated overhead.

A Trip Down Memory Lane

One of the greatest challenges associated with server

virtualization is how to virtualize the x86 virtual memory system. In a non-virtualized environment, the operating system manages memory by assigning virtual memory address spaces to applications. It uses page tables to map those virtual addresses to real system physical memory addresses.

A 48-bit virtual address in 64-bit mode in a 4-KB page table has four levels of address information. It can take up to four memory accesses to translate a virtual address to a physical address, plus one memory access to read the data. When the system needs to access physical memory, it “walks” each level of the page table until it finds the physical memory address that corresponds to the virtual memory address (Figure 2). Using virtual memory enhances security by isolating applications in their own memory spaces and presents more memory to applications than is physically present to avoid memory fragmentation.

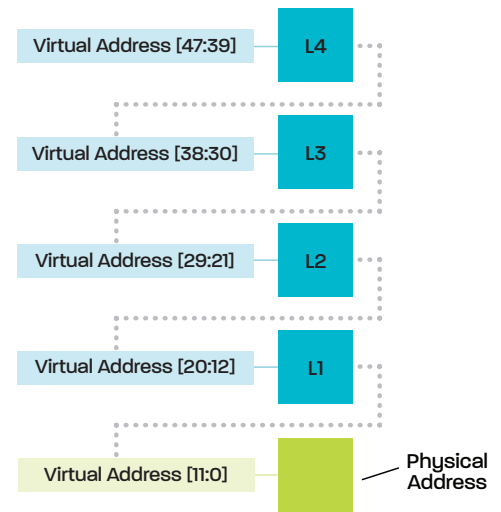


Figure 2: Standard x86 Page Table

Shadow Page Tables: High Overhead, Slow Performance

In a virtualized environment, the operating system becomes a guest on the physical, or host, machine. The hypervisor, or virtualization software, creates the illusion that the guest operating system is managing physical memory, when in fact the guest has no direct access to physical memory management on the host machine. Xen, the open source hypervisor used by SUSE Linux Enterprise Server to create virtualized environments, uses a software-based approach called shadow page tables to create this illusion.

Shadow page tables are duplicates of the page tables maintained by the guest operating system, and they contain the actual mapping of virtual to physical addresses. The hypervisor manages the shadow page tables and uses them to translate guest physical addresses to system physical addresses.

When a guest needs to modify a page table, the process involves a great deal of interaction between the guest and the hypervisor (Figure 3). The page tables that the guest operating system sees are located in write-protected memory, so when the guest tries to modify the page table, the system generates a page fault that is sent to the hypervisor. The hypervisor creates a copy of the table where the fault occurred, makes the original table writeable, and turns control of the table over to the guest. The guest modifies the page table and flushes the translation look-aside buffer (TLB), a buffer that holds cached translations of virtual to physical memory for faster access. The hypervisor traps the TLB flush, inspects the change, updates the corresponding shadow page table, and sets the guest page table as read-only again. It then allows the guest to resume operation.

Rapid Virtualization Indexing: Blazing Fast Performance with SUSE Linux Enterprise Server

Customers and hypervisor vendors alike have asked processor vendors to develop hardware support for virtualized memory management. With its Quad-Core AMD Opteron™ processors, AMD introduced AMD Virtualization™ (AMD-V™) technology with Rapid Virtualization Indexing (RVI), a highly efficient, hardware-based approach to virtualized memory management. RVI reduces the hypervisor overhead associated with software-based shadow paging algorithms, creating better performance for virtualized applications. Xen calls this technique hardware-assisted paging, and SUSE Linux Enterprise Server has been optimized to immediately take advantage of this improved memory management technology.

RVI uses a two-dimensional approach to page tables called nested page tables to enable the translation of virtual to physical addresses in hardware. This approach helps reduce hypervisor overhead by moving memory management back to the guest and allowing the guest to manipulate the page tables directly.

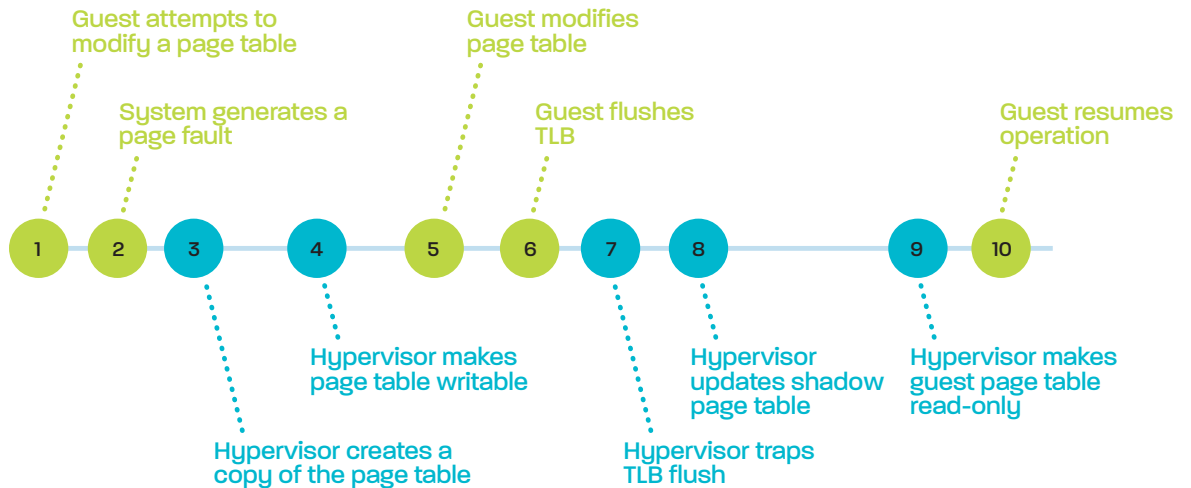


Figure 3: Software-Based Virtualized Memory Management with Shadow Page Tables

This process is complex and slow because of the number of world switches, or transitions between the guest operating system and the hypervisor. It creates an enormous amount of overhead in the hypervisor and limits the performance of virtualized workloads—especially memory-intensive ones like databases.

Nested page tables are really two page tables:

- The guest page table, which maps guest virtual addresses to guest physical addresses.
- The nested page table, which maps the guest physical address to the system physical address.

The nested page table is the mapping of the guest's view of physical memory to the system's actual physical memory. Since the guest physical address is **not** the actual system physical address, RVI uses the nested page table to convert guest physical addresses encountered during the page table walk to system physical addresses. This mapping can be created and held in memory, so the hypervisor does not need to intercept the guest's changes to its page tables.

RVI is designed to greatly increase performance of virtualized applications by allowing virtual machines to directly manage memory with less hypervisor intervention and associated overhead. RVI helps eliminate hypervisor cycles spent managing shadow pages, which can account for well over half of hypervisor time, and it reduces the number of world switches that a memory-intensive application must perform. The time it takes to perform a world switch has also been improved on Quad-Core AMD Opteron processors by 25% compared to Rev F for greater

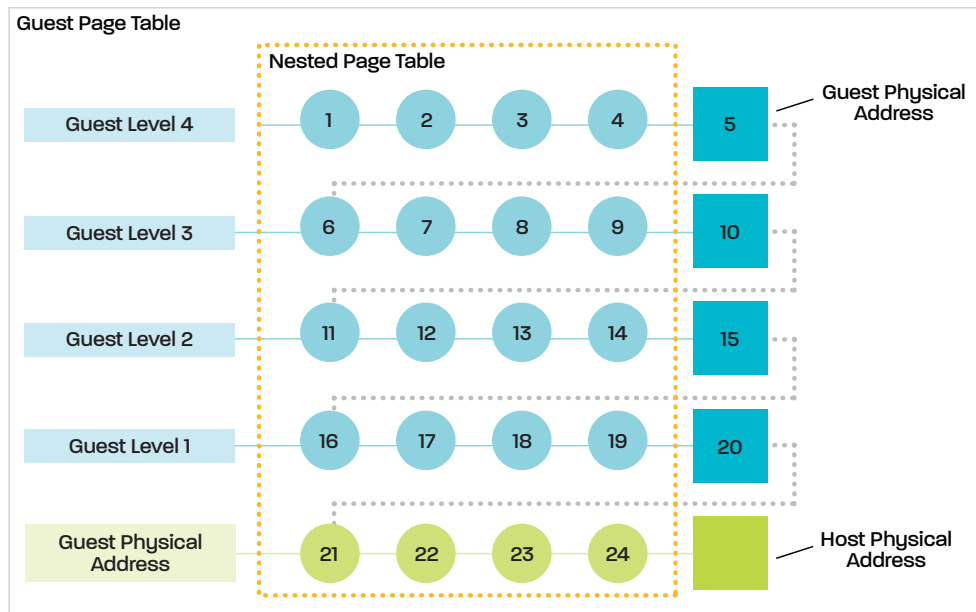


Figure 4: Nested Page Table

The nested page table adds a second layer, or dimension, to memory management. Each level of the guest page table has a corresponding level in the nested page table. Consequently, where a normal page walk could take up to five memory accesses to translate a virtual address to a physical address, a nested page walk could take up to 24 memory accesses (Figure 4).

AMD has developed a hardware solution called a page walk cache (PWC) to reduce the number of memory accesses required when the nested page table is walked. The PWC is a small, fast page entry cache that holds higher level entries from the page tables. The PWC can accelerate memory accesses because many page entry references are frequently reused. So, if a processor finds the page entry it needs in the PWC, it can avoid walking the entire page table and access the physical memory address faster.

application responsiveness. Virtualized workloads that run on SUSE Linux Enterprise Server will continue to perform at near-native levels because of Novell's optimization work with Quad-Core AMD Opteron™ processors and RVI.

INDUSTRY-LEADING POWER EFFICIENCY WITH ENHANCED AMD POWERNOW!™

Quad-Core AMD Opteron™ processors with Enhanced AMD PowerNow!™ technology are designed to help reduce your power and cooling costs by working with the operating system to deliver performance-on-demand capabilities for optimum performance-per-watt and power savings. SUSE Linux Enterprise Server has been optimized to dynamically change the clock speed of each core, to help reduce energy costs—especially for businesses that run their servers at below 60% utilization.

With Independent Dynamic Core technology from AMD, you can tweak each core to deal with the demands of a software workload because each CPU core is clocked independently. While the voltage of all four cores has to be equal, clock speed and the corresponding current draw can be reduced depending on load. This ability to independently manage the clock frequency of each core gives you more precise power management and associated power savings, as well as lowering total cost of ownership.

Reduced Energy Consumption

Quad-Core AMD Opteron processors can also create valuable power savings by employing independent voltage planes (called split planes) for the cores and the memory controller in the northbridge. Dual Dynamic Power Management™ (DDPM™) creates greater power savings through reduced energy consumption by allowing the processor to independently set the voltage for the cores and the northbridge. Servers with motherboards that have two separate voltage regulator modules can immediately take advantage of DDPM. For example, the processor can reduce voltage to the cores when they are not under load while the northbridge continues to run at an independent voltage. Additional voltage can also be sent to the northbridge to raise its frequency, resulting in a performance boost on power-constrained platforms.

SUSE Linux Enterprise Server 10 reduced the frequency of timer interrupts (ticks) to 250 Hz to take advantage of the C1e state in Quad-Core AMD Opteron processors. C1e is an enhanced processor sleeping state triggered by the BIOS if all CPU cores enter the C1 state. Reducing timer interrupts to 250 Hz can help significantly reduce energy consumption, resulting in additional power savings. And SUSE Linux Enterprise Server 11 is expected to have tickless timer interrupts and an optimized polling strategy for bundled applications to help maximize power savings from AMD Opteron processor C states.

INSTRUCTION-LEVEL PERFORMANCE IMPROVEMENTS

Besides power savings and enhanced performance in virtualized environments, AMD and Novell also focused on optimizing SUSE Linux Enterprise Server and Quad-Core AMD Opteron processors to improve instruction-level performance. Reducing the number of instructions that must be processed results in more efficient use of clock cycles, which contributes to better overall performance. GNU Compiler Collection (GCC) version 4.1, which ships with SUSE Linux Enterprise Server 10, has been updated to take advantage of these new features.

Advanced Bit Manipulation (ABM)

Advanced Bit Manipulation (ABM) offers two instructions that operate on general purpose registers: LZCNT (leading-zero count) and POPCNT (population count). By using ABM, programmers can create more efficient code because operations that previously required multiple instructions to complete can now be done with a single instruction.

LZCNT counts the number of leading zeros in a register or memory location. It starts at the most significant bit and works downward until a one is encountered. LZCNT is used to emulate floating point arithmetic for multi-precision numbers whose size is much larger than what a floating point unit can handle. It is also used for hashing in data mining, in compression, and for cryptographic purposes.

POPCNT identifies how many ones are set in a given value, and is particularly useful for communication, storage, and signal processing applications that use bit arrays for computations. Suppose that an application uses an array to store the results of a network transmit operation. Each element represents a true or a false, depending on whether the associated block transmitted correctly. This data can be used to calculate the amount of packet loss in the network transmit operation.

For example, if block numbers 16, 27, and 60 were not transmitted, the corresponding values in the array index would be set to one (and the rest would be zero). Using a bit array makes the array smaller and results in fewer memory accesses, but accessing an element in the array becomes more complex. By using POPCNT in this scenario, you can determine how many bits were set in the array with just one instruction. Besides this scenario, POPCNT can substantially increase performance in other applications including cryptography, encoding/decoding, databases (for quickly assessing information about data), and calculating Hamming distances.

New SSE4a Instructions

Streaming SIMD Extension (SSE) instructions are part of high-performance floating-point processing, which supports high-performance technical computing for financial analysis, biological sciences, and petroleum exploration. SSE instructions also play a critical role in media encoding and decoding for video compression, HD-DVD authoring, and other multimedia applications. Quad-Core AMD Opteron™ processors introduced four new SSE4a instructions:

- EXTRQ/INSERTQ: These instructions combine mask and shift operations into a single instruction. EXTRQ extracts bits (at a given offset/length) from the lower 64 bits of the XMM register and stores them into the least significant bits of the destination XMM register. INSERTQ inserts the least significant bits from the lower 64 bits of the XMM register into the lower 64 bits of the destination XMM register (providing offset and length).

- MOVNTSD/MOVNTSS: These instructions are streaming stores that can be done on scalar operands. MOVNTSD stores the DP XMM register into a 64-bit memory location, treating it as nontemporal to minimize cache pollution. This move is very useful when data is unlikely to be used again soon. MOVNTSS does the same, but for SP 32-bit values.

LOOKING TO THE FUTURE

Quad-Core AMD Opteron™ processors are designed with the future in mind, and Novell engineers expect to continue to collaborate with AMD to optimize SUSE Linux Enterprise Server. Future enhancements to SUSE Linux Enterprise Server are anticipated to unleash several features already built in to Quad-Core AMD Opteron processors, including new performance monitoring techniques and gigabyte memory pages.

INSTRUCTION-BASED SAMPLING

Software developers use performance monitoring techniques to validate new processors. Traditionally, developers use techniques like event-based sampling, where the system counts each time a given event (e.g., a floating point operation or cache miss) happens. Unfortunately, this method is imprecise and does not provide much insight into performance challenges at the instruction and source code levels. To provide developers with more precise information about processor performance, AMD created a new technique called Instruction-Based Sampling (IBS).

IBS gives insight into AMD64 instruction fetch behavior as well as the execution of operations that are issued from AMD64 instructions. Developers can use this information to analyze and improve the performance of programs that execute on Quad-Core AMD Opteron processors. IBS provides four advantages over traditional performance monitoring techniques:

1. Hardware events are attributed precisely to the instructions that cause the events.
2. A wide range of events are monitored and collected with each IBS sample. To obtain the same level of information using traditional methods requires multiple sampling runs or counter multiplexing.
3. IBS collects the virtual and physical addresses of load/store operands. Developers can use this information in profiling tools to associate specific data structures with x86 instructions that perform load/store operations.
4. IBS measures latency for key performance parameters, such as data cache miss latency.

Because IBS provides such precise information about instruction execution, developers could use it to create

automated optimization techniques (like profile-directed optimization) which require detailed, precise information about instruction-level program behavior. IBS support is unique to AMD processors and provides developers with detailed, useful information unavailable on processors from other vendors. Novell is expected to provide IBS support in the next major release of SUSE Linux Enterprise Server, marking another first in the AMD and Novell commitment to optimizing SUSE Linux Enterprise Server for AMD Opteron processors.

GIGABYTE MEMORY PAGES

Most memory management systems use 4 KB and 2 MB page sizes. The 4 KB pages provide fine granularity for memory access and mapping, but can create performance drains in heavy random access applications (like databases) because each memory access requires a complete page walk unless it is cached in the TLB. The 2 MB pages require less page walking, but applications whose random data access spans a range larger than 2 MB will still tend to run out of TLBs and have to do many page walks, which results in the same basic problem—slow performance.

AMD implemented support for 1 GB memory pages in its Quad-Core AMD Opteron processors, allowing the operating system to map a 1 GB section of memory at a time and configure how it is accessed. Because more data can be held in a contiguous section of physical memory, data access and performance can be greatly improved, especially for workloads like enterprise databases. Novell is expected to provide support for gigabyte memory pages in an upcoming release of SUSE Linux Enterprise Server, making it the first Linux distribution to optimize for AMD memory management techniques. The result: a foundation that dramatically improves performance in heavy random access workloads like enterprise database applications.

AMD AND NOVELL = INVESTMENT PROTECTION

Quad-Core AMD Opteron processors and Novell's flexible SUSE Linux Enterprise platform help you protect your IT investments. The Common Core Strategy from AMD enables seamless upgrades from dual- to quad-core processors and beyond, so you can minimize the cost of transitions and maximize your past investments in hardware, software, and personnel. Quad-Core AMD Opteron processors also feature Same Socket Technology that lets you maintain consistent thermal and power envelopes across multiple computing generations.

And with AMD's stable, long-term roadmap, your organization can leverage past investments in AMD Opteron™ processor-based platforms by upgrading to quad-core, since they are socket compatible with just a BIOS update, and will operate at the same power draw and thermal output as previous dual-core and single-core AMD Opteron processors. By choosing AMD and Novell, you can plan well-timed technology transitions that protect your IT investments and help decrease your IT total cost of ownership.

FINAL ANALYSIS

SUSE Linux Enterprise Server from Novell has always been enhanced to take advantage of key features in Quad-Core AMD Opteron™ processors. Because of this close collaboration, platforms based on AMD and Novell give data center managers and other server customers a robust, scalable, and flexible infrastructure that takes full advantage of the most advanced features of Quad-Core AMD Opteron processors and SUSE Linux Enterprise Server. When you choose AMD and Novell, you're choosing a stable, long-term solution that meets critical enterprise demands for exceptional energy efficiency, advanced virtualization technology, and outstanding computing performance.

To learn more about how Quad-Core AMD Opteron processors and SUSE Linux Enterprise Server can help improve workload performance in virtual environments and increase power savings, visit:

http://www.amd.com/us-en/Processors/ProductInformation/0,,30_118_8825,00.html
<http://www.novell.com/linux/>



Advanced Micro Devices
One AMD Place
Sunnyvale, CA 94088-3453
www.amd.com

¹ Data taken from "Barcelona: AMD's Next-Generation Quad-Core Microprocessor", a presentation given by Ben Sander, AMD Principal Member of Technical Staff, on March 28, 2007; slide 24. Filename: CART2007-Barcelona.pdf.

This White Paper may contain forward-looking statements, which are made pursuant to the safe harbor provisions of the U.S. Private Securities Litigation Reform Act of 1995. Forward-looking statements are generally preceded by words such as "plans," "expects," "believes," "anticipates" or "intends." Investors are cautioned that all forward-looking statements in this document involve risks and uncertainty that could cause actual results to differ materially from current expectations. Forward looking statements in this White Paper include the risk that Novell will not release SUSE Linux Enterprise Server 11 precisely as plans. We urge investors to review in detail the risks and uncertainties in the Company's filings with the United States Securities Exchange Commission.

AMD assumes no obligation to update or otherwise correct or revise this information. However, AMD reserves the right to revise this information and to make changes from time to time to the content hereof without obligation of AMD to notify any person of such revisions or changes.

© 2008 Advanced Micro Devices, Inc. All rights reserved. AMD, the AMD Arrow logo, AMD PowerNow!, DDP, Dual Dynamic Power Management, AMD-V, AMD Virtualization, AMD Opteron and combinations thereof are trademarks of Advanced Micro Devices, Inc. Other names are for informational purposes only and may be trademarks of their respective owners.