White Paper

Technology Innovations for Storage: Overview and Benefits for Storage Professionals

Intel in Communications

run
Better, faster and further.

Contents

Executive Summary

The overall storage marketplace and data storage requirements have seen enormous growth in the last 6-7 years, even through the latest downturn beginning in the second half of 2000. This has created an environment in which storage solutions developers and IT managers are demanding scalable, high-performance, standards-based solutions which can be quickly and easily deployed and managed. The storage market segment relies on Intel to drive storage technology and initiatives as well as provide innovative features and solutions that solve the problems faced by today's storage professionals. Intel® processors with Hyper-Threading (HT) Technology and Intel® Extended Memory 64 Technology (Intel® EM64T) are two examples of technology innovations from Intel which deliver on this promise.

Threaded programming as well as processor and application parallelism is becoming more prevalent in today's storage environment. Intel processors with HT Technology add value to storage systems by doubling the number of logical processors available to applications and operating systems, enabling higher performance I/O, response time and overall number of transactions per second. Storage environments, especially large, high-performance systems have outgrown the memory and addressability limitations of traditional 32-bit architecture. Intel® Xeon™ processors with Intel EM64T provide a solution for developers of storage systems who need to address more memory in a flat address space beyond 4 GB. Intel EM64T provides the ability to address larger disk drives and/or decreased disk cluster sizes as well as more files, all of which contribute to more efficient use of disk drives and less wasted space.

The purpose of this paper is to provide a quick introduction to key storage technologies and challenges faced by storage professionals and discuss the benefits of HT Technology and Intel EM64T as they relate to storage. This document will serve only as a base level of understanding of these technologies. For more comprehensive papers written specifically for these technologies, refer to the bibliography at the end of this paper.

Introduction

Despite the recent market segment downturn, the worldwide storage marketplace has continued to grow every year since the late 1990s to nearly a $50B marketplace today (see Figure 1).

This enormous growth in the storage marketplace creates many challenges for storage equipment manufacturers and storage professionals in terms of cost, performance and the sheer magnitude in the amount of data to be stored. Intel is an industry leader in the development of key storage technologies and the delivery of standards-based technology and storage building blocks. Innovative microprocessor features such as HT Technology and Intel EM64T are two examples of technologies which add value to Intel architecture-based storage solutions.

Storage Overview

Storage systems are connected and typically classified as Direct Attached Storage (DAS), Network Attached Storage (NAS), or Storage Area Networks (SANs).

DAS connects the storage directly to the server typically over standard interfaces like SCSI. With DAS, the server directly owns the storage. The storage is usually located physically close to

Figure 1. Source: IDC

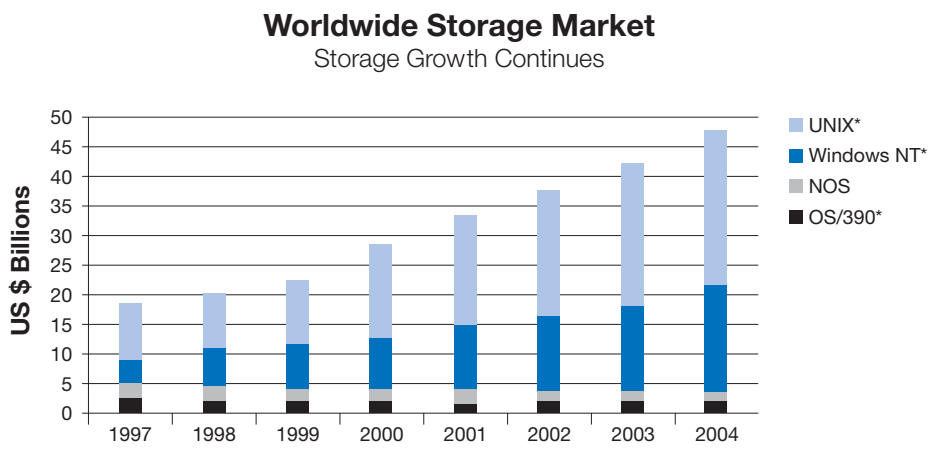

the server and is accessed at the block level. NAS connects the storage directly to a network such as a Local Area Network (LAN) running Gigabit Ethernet which provides file-level access to data using a standard protocol such as Network File System (NFS) or Common Internet File System (CIFS). Systems access data from a NAS device via a file "redirector." The file system on the NAS server determines if the location of the data requested by the application client is in its cache or on the storage.[1]

SANs are dedicated, high-speed networks of storage elements combining hardware, software, and networking components. They provide any-to-any connections of server and storage elements and access data at the block level. The connections or "fabric" between the storage elements are typically Fibre Channel. SANs have the advantage of creating a massive pool of storage, storage consolidation, centralized management, improved data access, scalability and security.

Three of the top challenges faced by storage professionals are cost, performance, and data growth. The growth of the Internet, communications and e-commerce have forced corporations and IT managers to continuously invest in, deploy and manage more storage hardware while maintaining or reducing management staff. This inconsistency creates a large gap in the amount of storage managed per staff member. In addition, the growth in storage requirements fuels the need for higher performance storage processing solutions, disk drives and storage interconnects.

Hyper-Threading Technology

Introduction to HT Technology

HT Technology is a form of Simultaneous Multi-Threading Technology (SMT), where multiple threads of software applications can be run simultaneously on one processor. This is achieved by duplicating the architectural state on each processor, while sharing one set of processor execution resources (see Figure 2). The architectural state tracks the flow of a program or thread, and the execution resources are the units on the processor that do the work: add, multiply, load, etc.[2]

Hyper-Threading Technology Benefits for Storage

In the high-performance storage category, most storage systems are traditional Dual-Processor (DP) designs with two physical processors on the main board. As processing requirements continue to increase, more simultaneous processing threads and more processors will be required. This creates a problem many storage equipment OEMs face every year. How to design a high-performance system while remaining in the same relative space and price point?

HT Technology provides a method to increase processing capability in the same relative footprint. Since Intel introduced the Intel Xeon processor with HT Technology in 2002, a DP system running two processors with HT Technology can provide a performance benefit of four processors or threads running simultaneously in the same relative space as a traditional DP system. In addition, storage OEMs now have the ability to take existing DP applications, re-package them in a UP system with processors enabled with HT Technology and provide solutions for their customers with lower cost and performance requirements, without having to change their application.

Figure 2. Source: Hyper-Threading Technology on the Intel® Xeon™ Processor Family for Servers WP[2]

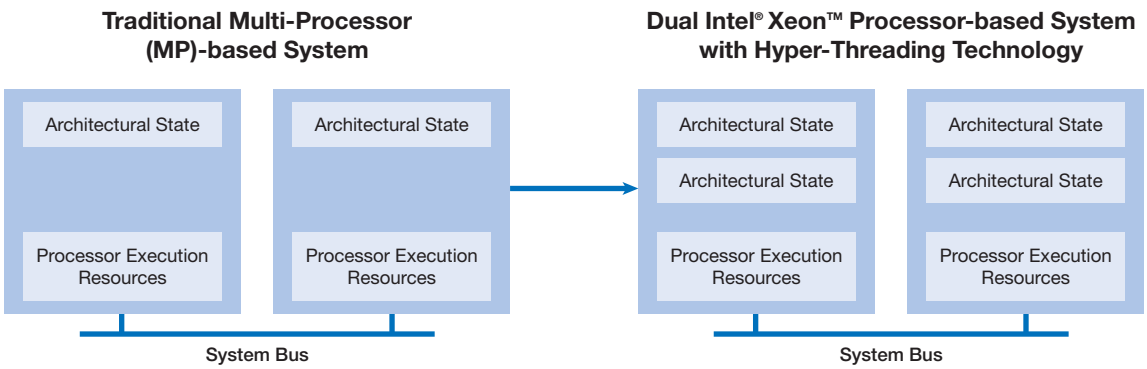

System scalability is another challenge faced by storage professionals from mid-range to the highest performance storage systems. To keep up with storage growth, IT managers demand standards-based, high-performance solutions that can be deployed quickly and meet the same space requirements. In June 2003, DataCore and Huge Systems both issued press releases that noted benefits of using processors with HT Technology in their products.[3,4]

Using Intel® solutions enabled with HT Technology effectively doubles the number of logical processors available for OS and applications software to increase performance. This performance enhancement has a significant, noticeable impact on the performance of server applications: it increases the number of transactions that can be processed, enables support for more users, and provides faster response times for many Internet and e-Business applications.

Intel® Extended Memory 64 Technology

Introduction to Intel® EM64T

In mid-2004, Intel Xeon processors enabled with Intel EM64T were introduced to the marketplace. Intel EM64T includes extended memory addressability (i.e., 64-bit pointers and registers), eight additional registers for Streaming Single Instruction, Multiple Data (SIMD) Extensions (SSE) and eight additional general-purpose registers, and double-precision integer support.

Intel EM64T Benefits for Storage

In many storage systems, one sure way to increase performance is to add a disk cache. Typically the strategy for the disk write and disk read caches are quite different. The write cache is positioned to absorb data that is written from the network in real time. The data is then written as fast as possible to the actual hard disk drives. The read caches, on the other hand, try to hold those files most often requested by the NAS or SAN client systems to hide normal disk latency in providing the data directly (see Figure 3).

Figure 3. Typical Disk Cache Symbolic Configuration

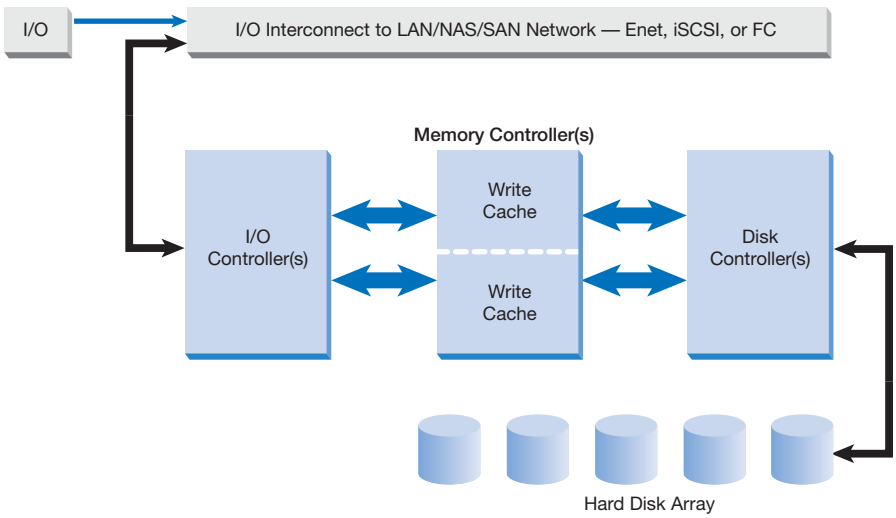

Table 1. Cluster Configurations

Disk size	Required cluster size for a FAT32 volume
16 GB	4 KB
64 GB	16 KB
128 GB	32 KB (the largest cluster size that Windows* 95/98 tools can handle)
256 GB	64 KB (not supported by Windows 95/98 or Windows Me)

Today, a mid-range storage system may be managing a total disk cache memory of 8 GB, while large systems may have cache sizes up to 256 GB. So far, 32-bit processors have been limited to total physical memory spaces of 4 GB. To address larger amounts of memory, systems were required to use either proprietary paging models or Intel's PAE (Physical Address Extension) model. Both of these models carry a performance penalty due to the extra page lookups required. The Intel EM64T modes in the new generation of Intel Xeon processors provide the ability to directly address 64 GB, enabling developers to break the 4 GB boundary and provide solutions with large disk caches, enough to service all but the largest storage systems.

Drive space and efficiency is another top concern for storage professionals. Address space limitation applies to disks just as it does for memory. This requires a change to a different file system, but a file system with more bits can use smaller data clusters, access larger disks, and handle more files.

As you'll recall, Windows* OS 95 changed the disk system from the 16-bit FAT16 file system that was limited to about 2 GB per partition to a 32-bit file system with a theoretical limit of 8 GB per partition, although that requires using 32 KB data clusters. Since a cluster is the smallest increment on a disk, any file smaller than 32 KB leaves wasted space, and any file that does not evenly fit into 32 KB increments leaves wasted space (see Table 1).[5]

The number of files is also an issue with large volumes because a 32-bit file system is limited to about 4.3 billion files. Although this may seem like a lot of files, this can be quite small when compared to a large SAN that can consist of multiple Terabytes (TB).

The extra address bits enabled with Intel EM64T help solve the drive efficiency problem as well as the limit on the number of files. Intel EM64T enables users to use smaller clusters and/or larger disks. Smaller clusters mean more files on the same size disk. The inverse also improves efficiency by using larger disks with the same cluster sizes.[6]

Finally, a challenge related to performance and memory is passing program parameters back and forth between the processor's internal registers and system memory. For some time, internal registers were used to pass variables because it was the fastest way to do so. As programming structures became more complex, programmers were forced to pass variables in main memory, which is much slower and results in a performance penalty.

Enabling Intel EM64T can add considerably to the on-chip registers visible to software. Essentially, eight additional 64-bit General Purpose Registers and eight Extended Multi-Media (XMM) 128-bit registers are added. Now that the on-chip register set has been enhanced, it may be possible to once again pass parameters register to register.

Summary

The storage marketplace and data storage requirements will continue to grow, prompting the need for high-performance, cost-effective solutions. Data management will continue to be an issue and fuel the need for standards-based solutions which can be deployed quickly. Processor and application parallelism are becoming more prevalent. Developers are increasingly turning to threading as a way of boosting overall system performance. Storage capacity requirements have reached the 4 GB limitations of traditional 32-bit architecture, creating a need for higher capacity disk caches and disk drives.

HT Technology and Intel EM64T are two technologies from Intel to address the challenges facing storage professionals. These technologies provide cost-effective methods to enhance overall system performance, increase drive efficiency and create the ability to break through existing 32-bit addressability barriers. They are only two of what will be many more processor technologies to be driven and delivered by Intel to benefit the storage marketplace.

Bibliography

1. White Paper, "An Overview of Network Attached Storage," The Evaluator Group, http://www.evaluatorgroup.com

2. Intel White Paper, "Hyper-Threading Technology on the Intel® Xeon™ Processor Family for Servers," http://www.intel.com/business/bss/products/hyperthreading/server/ht_server.pdf

3. DataCore Press Release, "DataCore Unveils Breakthrough storage Management Services," http://www.datacore.com/pressroom/pr_030604a.asp, June 2003

4. Huge Systems Press Release, "Huge Systems Announces Impressive Performance of Mediavault U-320RX Array with new Linux 2.6.6 kernel," http://www.hugesystems.com/PressRoom/Releases/Huge-U320-RX-Linux-Release.pdf, June 2003

5. Windows Hardware and Driver Central, http://www.microsoft.com/whdc/system/winpreinst/ntfs-preinstall.mspx

6. McPherson, James, "Why 64-bit is the 'new' catchword," *TechRepublic*, April 2003

For more information, visit the Intel Web site at: **developer.intel.com**

UNITED STATES AND CANADA
Intel Corporation
Robert Noyce Bldg.
2200 Mission College Blvd.
P.O. Box 58119
Santa Clara, CA 95052-8119
USA

EUROPE
Intel Corporation (UK) Ltd.
Pipers Way
Swindon
Wiltshire SN3 1RJ
UK

ASIA-PACIFIC
Intel Semiconductor Ltd.
32/F Two Pacific Place
88 Queensway, Central
Hong Kong, SAR

JAPAN
Intel Kabushiki Kaisha
P.O. Box 115 Tsukuba-gakuen
5-6 Tokodai, Tsukuba-shi
Ibaraki-ken 305
Japan

SOUTH AMERICA
Intel Semicondutores do Brazil
Rue Florida, 1703-2 and CJ22
CEP 04565-001 Sao Paulo-SP
Brazil

Printed in USA 0804/DLC/OCG/XX/PDF Please Recycle 303490-001US