

The Effect of L3 Cache Size

on MMB2 Workloads



The Intel® Xeon™ processor MP is designed for powerful servers and can be used in a wide range of deployments. It supports either a 512 KB or 1 MB level 3 (L3) cache. Dell tested the impact of L3 cache sizes on the MAPI (Messaging Application Programming Interface) Messaging Benchmark 2 (MMB2) workload using a Dell™ PowerEdge™ 6600 server. These tests show how cache size affects the performance of the system memory, the Microsoft® Exchange Information Store, and the system processor.

BY SCOTT STANFORD

The Intel® Xeon™ processor MP uses Hyper-Threading technology, which increases utilization of processor resources, to provide responsive performance. Another factor affecting this processor's performance is the size of its level 3 (L3) cache. This article analyzes the Intel Xeon processor MP L3 cache under the MAPI (Messaging Application Programming Interface) Messaging Benchmark 2 (MMB2) workload. MMB2 mimics enterprise-level messaging workloads in the Microsoft® Exchange 2000 Server environment and places demand on processor, memory, and disk I/O resources.

Establishing the test environment

The Dell™ team tested a Dell PowerEdge™ 6600 server with four Intel Xeon processors MP¹ at 1.5 GHz. Previous studies² have shown that Hyper-Threading improves the use of available processing power in an Exchange 2000 Enterprise Server environment, and therefore Dell engineers enabled Hyper-Threading in these tests.

The Dell engineers connected the PowerEdge 6600 server to two Dell PowerVault™ 220S SCSI enclosures, which held the logs and databases for Exchange. Each PowerVault 220S contained one RAID-1+0 array for the Exchange transaction logs and another RAID-1+0 array for the Exchange Information Store databases. The PowerEdge 6600 server was configured as follows:

- Four Intel Xeon processors MP at 1.5 GHz with 512 KB L3 cache or four Intel Xeon processors MP at 1.5 GHz with 1 MB L3 cache
- Microsoft Windows® 2000 Advanced Server operating system with Service Pack 3 (SP3) and Microsoft Exchange 2000 Enterprise Server with SP3
- Four 1 GB, 200 MHz error-correcting code (ECC) double data rate (DDR) memory modules
- One PowerEdge Expandable RAID Controller 3/ Dual Channel (PERC 3/DC) controlling the

¹Processors were of the Foster class, not the Gallatin class.

²For more information on Hyper-Threading and Exchange 2000 Server performance, see "Impact of Hyper-Threading Technology on Exchange 2000 Performance" by Ramesh Radhakrishnan and Scott Stanford in Dell *Power Solutions*, August 2002.




Figure 1. Lazy Write Pages per second at 512 KB and 1 MB cache sizes




Figure 2. Page Faults per second at 512 KB and 1 MB cache sizes

operating system on channel A and the PowerVault 220S on channel B

- One PERC 3/DC controlling a second PowerVault 220S on channel A
- One single-port Intel PRO/1000 MT Peripheral Component Interconnect Extended (PCI-X) network adapter card

The PowerEdge 6600 internal 1×8 backplane contained eight 15,000 rpm U160 disk drives in a RAID-1+0 configuration supporting four partitions for the operating system, Exchange executables, the Microsoft Active Directory® directory service, and the paging file housed in the internal 1×8 bay.³

Examining subsystem performance monitor counters

To characterize how the L3 cache size affects system performance under a 4,000-user MMB2 workload, Dell engineers examined several performance monitor counters for system memory, Exchange Information Store, and the microprocessor.

Lazy Write Pages

The file system cache is an area of system memory where applications store frequently accessed information to avoid accessing the data

from disk. The Lazy Write Pages per second performance counter provides one method for measuring how frequently information stored inside the file system cache is written to a physical disk medium. As Figure 1 shows, at a 4,000-user MMB2 workload, doubling the L3 cache size from 512 KB to 1 MB reduced the number of Lazy Write Pages per second by almost five, or 33 percent. The larger cache reduced the need to access relatively slower system memory, and this in turn reduced access to the even slower disk I/O subsystem.

Page Faults

For memory requests that cannot be satisfied by the contents in the processor cache, additional transactions involving reads or writes to main memory addresses must occur. Requests for data that is present in physical memory blocks not immediately available to the processor cache can cause a memory access fault. The faulted page is referred to as a soft fault when the requested data is located elsewhere in physical memory.⁴ When the data does not reside in physical memory, hard faults occur. In contrast to soft faults, which typically cause only minuscule delays, hard faults require reads to disk that can result in additional slowdowns to satisfy a pending I/O transaction.

Just as the larger L3 cache helped reduce the number of Lazy Write Pages per second, it also dropped the number of Page Faults per second—a performance counter that measures the number of hard faults—by more than 7 percent when compared to the 512 KB L3 cache processors. Figure 2 shows Page Faults per second data for the processor cache size configurations at a 4,000-user MMB2 workload.

Exchange Information Store

As the test results demonstrate, larger cache size can improve file system efficiency and reduce the number of requests from main

The send queue lengths of the Exchange Information Store can indicate how efficiently pending MMB2 transactions are committed to mail databases.

memory to slower disk I/O subsystems. The send queue lengths of the Exchange Information Store can indicate how efficiently pending MMB2 transactions are committed to mail databases. Writes or reads made to the mail databases while messages are committed depend on—and can be negatively or positively affected by—an efficient file system. Like the number of Lazy Write Pages and Page Faults, message queue size is

³PERC 3/DC policies set to default and 64 KB stripe size. NT file system (NTFS) allocation unit set to default.

⁴Windows 2000 Resource Kit, Performance Counter Reference Help Files.

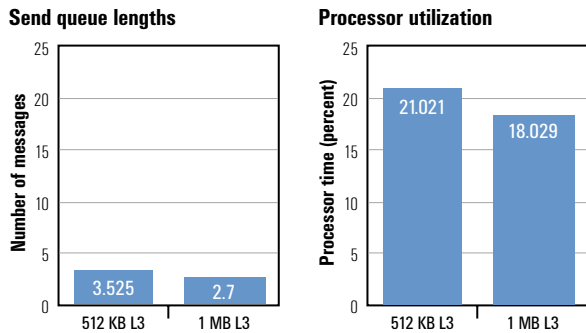


Figure 3. Send queue lengths of the Exchange Information Store mailbox and processor utilization at 512 KB and 1 MB cache sizes




Figure 4. System Processor Queue Length at 512 KB and 1 MB cache sizes

directly related to file system and disk I/O efficiencies, and a larger L3 cache size can reduce inefficiencies.

At a 4,000-user MMB2 workload, the send queue lengths of the Information Store mailbox were reduced by nearly one message for the duration of the test. Although this reduced queue length may not seem substantial, the amount of processor time dedicated to handling the same pending message queues with the 1 MB L3 cache was only 18 percent. As shown in Figure 3, the processor with the 512 KB L3 cache worked harder to maintain a slower message delivery capability at more than 21 percent processor utilization.

Processor Queue Length


Processor Queue Length is a processor utilization performance counter that tracks the number of waiting, or ready, threads in a queue. Fewer threads in a waiting state result in lower workload queue lengths. Figure 4 shows that the 1 MB L3 cache processor configuration had half, or 50 percent, of the number of waiting threads in queue than the smaller 512 KB L3 cache implementation.

Improving processor performance through larger cache

Many more performance monitor counters can be used to further explore the impact of Intel Xeon processor MP L3 cache sizes on MMB2, or any other workload. The test cases in this article used

As the test results demonstrate, larger cache size can improve file system efficiency and reduce the number of requests from main memory to slower disk I/O subsystems.

some key memory, Microsoft Exchange, and system counters to provide a snapshot of server and application health. Dell benchmarking tests showed that the PowerEdge 6600 server containing four Intel Xeon processors MP at 1.5 GHz scales well as the L3 cache size increases.

Configurations with the larger L3 cache utilize physical memory more effectively, so more data requests can be handled by system memory rather than slower disk I/O subsystems. Processors with larger L3 cache provide more effective file system cache activity and shorter message and processor queue lengths. The larger L3 cache size improves overall system efficiency by allowing the same frequency processor to do more work with less effort. 

Scott Stanford (scott_stanford@dell.com) is a systems engineer in the Server Performance and Analysis Lab at Dell. His current work focuses on Exchange Server 2000 benchmarking. Scott served in the U.S. Peace Corps in Nepal and with the U.S. Army, 24th Infantry Division. Before Dell, he worked in the public sector as an information services manager. He has an M.S. in Community and Regional Planning from The University of Texas at Austin, and a B.S. from Texas A&M University. Scott is A+ and N+ certified and a Microsoft Certified Systems Engineer (MCSE).

FOR MORE INFORMATION

Dell PowerEdge server performance benchmarks:

http://www.dell.com/us/en/esg/topics/products_benchmark_pedge_exchange.htm

Dell and Microsoft Exchange:

<http://www.dell.com/exchange>

Microsoft Exchange:

<http://www.microsoft.com/exchange>

Intel Xeon processors:

<http://www.intel.com/xeon>

