



Veritas Storage Foundation™ *for Windows®*

Enhancing Application
Performance with VxCache

In Veritas Storage Foundation *for Windows*

Enhancing Application Performance with VxCache

Contents

Introduction	4
Caching	4
Before cache	5
After cache	5
Memory Addressing	6
Physical Address Extension x86	6
VxCache	7
Determine memory available for VxCache	7
Determine the volumes on which to enable VxCache	7
Monitoring VxCache	8
VxCache performance	9
Naming conventions	10
Best practices	13
Conclusion	14
References	14

Introduction

Performance is a critical aspect of any server deployment and remains the focus of Windows system administrators as well as Microsoft. In order to maintain competitive advantage and stay within increasingly restricted budgets, businesses need a way to increase performance while maintaining the bottom line. Although disk drive performance continues to rise, the fastest way to access data is from memory. The process of storing data in memory for quick access is called caching. The Veritas Storage Foundation *for Windows* VxCache feature allows volume-level caching, resulting in enhanced read performance for the volumes on which it has been enabled. Figure 1 shows the benefits of enabling VxCache on a volume containing Microsoft® Exchange databases. Response times are progressively lower, that is, quicker, as more memory is configured for VxCache. See the section on “VxCache performance” for additional details.

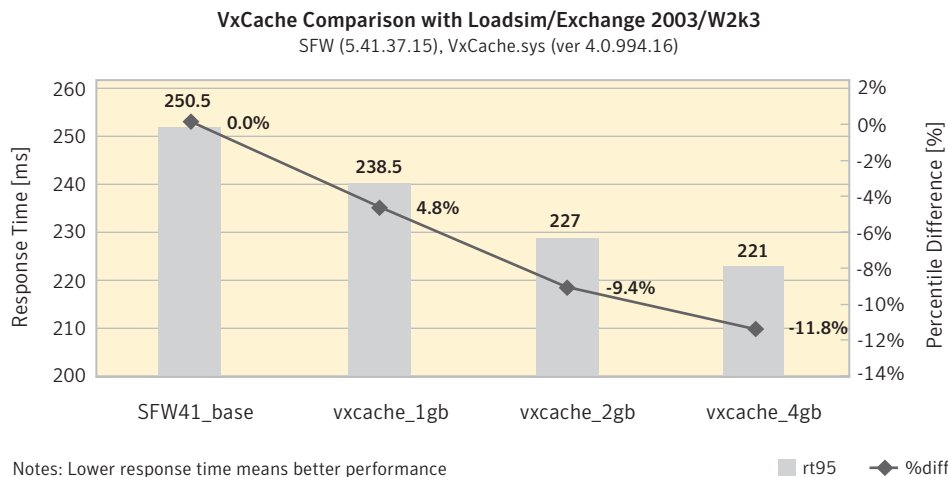


Figure 1. VxCache enabled on Exchange database volumes

Caching

Among the various types of caching are memory caching, hardware and software disk caching, and page caching. Caching appears on every computer in a variety of forms and allows a computer's tasks to be done more rapidly. Data which is used frequently is duplicated in a temporary location which has faster access times than the original data location. This data caching speeds data access times and streamline application response times.

Veritas Storage Foundation *for Windows*: Enhancing Application Performance with VxCache

Before cache

A Microsoft Exchange Server 2003 user uses the Microsoft Outlook application to access an unread email. The request is made to the Exchange Mailbox Server which does a read access from the data volumes which contain the email message and the message is sent to Outlook for access by the user. Later, another user who was sent the same email message wants to read the same message. Without caching, the Exchange Server must complete another read access to retrieve the email message from the data volumes. Since disk storage has much slower access than cache memory, the access times for the first user is delayed and when the second user requests the same email message there is the same delay for this user as well.

After cache

The same Microsoft Exchange Server 2003 user uses the Microsoft Outlook application to access an unread email. The request is made to the Exchange Mailbox Server which does a read access from the data volumes which contain the email message and the message is sent to Outlook for access by the user. When the email message is read from the data volume, caching is used to write a copy of the email message to computer memory which has much faster access times than disk volumes. Later, another user who was sent the same email message wants to read the same message. With caching, the Exchange Server obtains the email message from the disk cache which provides nearly instantaneous access to the email message.

Based on this example, some important facts to note are:

- Cache technology is the use of a faster but smaller memory type to accelerate a slower but larger memory type.
- When using a cache, you must check the cache to see if an item is there. If it is there, it's a cache hit. If not, it's a cache miss, and the computer must wait for a round-trip from the larger, slower memory area.
- A cache has a maximum size that is much smaller than the larger storage area.
- It is possible to have multiple layers of cache.

Memory addressing

Windows operating system (OS) 32-bit editions support only 32-bit memory addressing and therefore can access only 4 GB (2^{32}) of physical memory. Today, hardware is available that offers much more physical memory than 4 GB. Physical Address Extension (PAE) x86 and 4-gigabyte tuning (4GT) are two technologies available to the Windows OS that allow increased access to physical or virtual memory by user mode applications. By default, on systems with 4 GB of memory, the address space is split in two, with 2 GB of virtual address space directly accessible to applications and the other 2 GB accessible only to the kernel. 4GT uses the /3GB switch in the boot.ini to change this relationship so that applications have increased access to 3 GB of virtual address space, with the remaining 1 GB going to the kernel. This, however, does not provide access to memory above 4 GB.

Physical Address Extension x86

PAE x86 allows 32-bit Windows servers¹ to access memory above 4 GB by changing the addressing mode to support 36-bit memory addressing. This makes it possible to address up to 64 GB (2^{36}) of memory and thereby increase OS performance, especially when the server hosts multiple applications. Applications can also directly access memory above 4 GB if they support the Address Windowing Extensions (AWE) API set.

Note: Win32 Address Windowing Extensions provide user applications with 32-bit virtual addressing to use more than 32-bit regions of physical memory. AWE allows applications to use physical nonpaged memory (memory that cannot be paged to disk) and map sections of that memory to specific, equal-sized address ranges within the applications' 32-bit virtual address space, allowing them to quickly manipulate large amounts of physical memory. AWE is suitable for state-of-the-art database management systems and other memory-intensive applications.

PAE allows the OS and applications to treat all of the physical memory in a system as general-purpose memory and use it for caching and virtual memory management without major changes. More applications can be running before paging occurs. Paging is the process of moving infrequently used parts of a program's memory from RAM to hard disk, specifically, the page file. More detailed information on what PAE is and how it works with Windows is available on Microsoft's Web site.

¹ PAE is not necessary on 64-bit versions of the Windows Server family.

Veritas Storage Foundation *for Windows*: Enhancing Application Performance with VxCache

VxCache

Veritas Storage Foundation *for Windows* includes a feature called VxCache that can be enabled on a volume-by-volume basis to improve I/O performance by dedicating a portion of a system's physical memory to caching. A memory pool of up to 4 GB can be set aside at the system level for VxCache, and volumes can be assigned to utilize this pool of memory for block-level caching by specifically enabling VxCache on them. The VxCache driver can address either a portion of extended memory or PAE memory on systems that are PAE enabled. VxCache uses write-through caching in which writes go directly to disk, eliminating the potential for data loss on VxCache-enabled dynamic volumes in the event of a system crash. Performance can be monitored via the VxCache Volume Monitoring window.

Implementing VxCache is essentially a two-step process. The first step is to enable it at the server level. A prerequisite to this step is to determine the amount of physical memory available for VxCache. VxCache cannot be enabled on systems with less than 512 MB of physical memory. But note that while it can be enabled on systems containing 512 MB of cache, 500 MB of memory is reserved for the OS and applications. The system should contain additional memory for VxCache. The next step is to enable VxCache at the volume level. A preliminary analysis should be done to determine whether a volume is a candidate for VxCache or not.

Determine memory available for VxCache

To determine the amount of physical memory available on a system for VxCache, use Perfmon, Task Manager, or a similar tool to determine the amount of memory in use on the system. If it is less than 500 MB, use 500 MB as the base amount, as VxCache assumes this amount as the minimum that will be used by the OS. Determine the total physical memory in the system (via System Properties). VxCache cannot be enabled on systems that have less than 512 MB of physical memory. The difference between the memory being used and the total memory on the system will give the amount of memory available for VxCache allocation.

Determine the volumes on which to enable VxCache

After determining the amount of memory available for VxCache, identify the volumes for which VxCache should be enabled. Not all volumes are suitable for VxCache. Target volumes should be read intensive, allowing quick retrieval of data from memory compared with disk during read operations. Enabling VxCache on a volume containing Exchange databases, for example, would make accessing an email attachment sent to several recipients faster, as each recipient accesses

it from memory rather than the hard disk each time it is accessed. Write-intensive volumes, such as a volume containing the transaction logs for Exchange, would not be suitable for VxCache, as they would not benefit from caching. Also, volumes containing the databases for AWE-aware applications, such as Microsoft SQL, would not be candidates for VxCache.

Monitoring VxCache

VxCache utilization can be monitored at the volume level using VxCache Volume Monitoring. Enabling monitoring on a particular VxCache-enabled volume will launch the VxCache Volume Monitoring window, which displays various statistics, including the volume name, total VxCache memory, and more (see Figure 2).

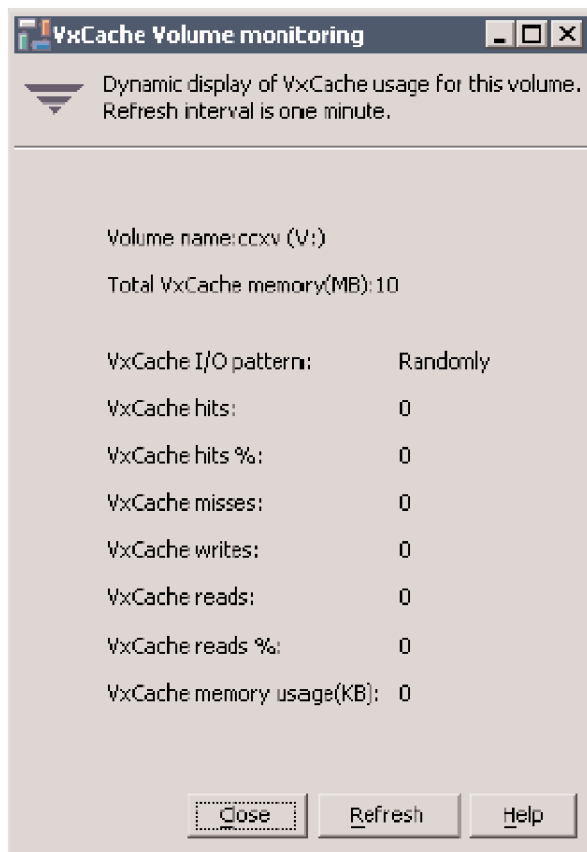


Figure 2. VxCache Volume Monitoring window

Veritas Storage Foundation *for Windows*: Enhancing Application Performance with VxCache

Monitoring these statistics over a period of time can help determine whether or not VxCache should be enabled on a volume. When VxCache reads outnumber VxCache writes, performance has been improved and VxCache should be enabled on that volume. If over time, however, the number of writes is more than the number of reads, there may be a degradation of performance with VxCache, as each write is being written to the volume and to memory. In this case, it would probably be best to disable VxCache on that volume (no need for the overhead of checking the backpack if you're going back to the storeroom anyway). Table 1 gives details of the statistics monitored.

Table 1. VxCache Volume Monitoring statistics

Term	Description
Volume name	Name of the volume that has VxCache enabled
Total VxCache memory	Total amount of VxCache memory available
VxCache I/O pattern	I/O pattern that VxCache uses to improve performance
VxCache hits	Number of requests (reads and writes) satisfied by VxCache
VxCache hits %	Percentage of hits of the total number of requests (satisfied and not satisfied) made to VxCache
VxCache misses	Number of reads and writes not satisfied by VxCache
VxCache writes	Number of writes to VxCache
VxCache reads	Number of reads satisfied by VxCache
VxCache reads %	Percentage of reads of the total number of all requests (reads and writes) satisfied by VxCache
VxCache memory usage	Amount of VxCache memory used by this volume

VxCache performance

Performance varies based on the amount of available VxCache memory, the size and number of the volumes using VxCache, and how data is accessed. As stated previously, VxCache should be enabled on volumes that experience a high proportion of reads, such as those containing Exchange databases, and not on volumes that are more write intensive, such as the volume containing the Exchange transaction logs.

Using Loadsim, an Exchange benchmarking tool, VxCache was evaluated on a Windows Server 2003 system running Exchange Server 2003. A baseline was established with Veritas Storage Foundation *for Windows* running on a PAE-enabled system. Differences in performance and hence the baseline between systems running with and without Veritas Storage Foundation *for Windows* (SFW) were negligible. Testing also included enabling 4GT, where the /3GB switch is added to the boot.ini file.

Veritas Storage Foundation *for Windows*: Enhancing Application Performance with VxCache

Baseline performance showed the Loadsim 95th percentile response time (RT95) as 250.5 ms. VxCache with 1 GB of cache showed a response time of 238.5 ms, or 4.8 percent faster than baseline. VxCache with 2 GB of cache showed a response time of 227 ms, or 9.4 percent faster than baseline. With 4 GB of memory, VxCache showed a response time of 221 ms, or 11.8 percent faster than baseline. In the following tables, lower response times indicate better (faster) performance. Table 2 compares performance vertically, with VxCache at various memory levels, against the baseline (as previously defined). Table 3 introduces the /3GB switch into the equation and compares performance horizontally (with /3GB versus without /3GB) at baseline and different VxCache levels. Table 4 summarizes the performance of all tested configurations, making a vertical comparison against the baseline. The volume layouts used in this testing were a 120 GB RAID-0 volume for the databases and a 36 GB simple volume for the transaction logs.

Naming conventions

- Except for SFW41_noPAE_/3GB, all test cases used the PAE switch.
- SFW41_base = SFW 4.1 with PAE enabled as baseline
- SFW41_noPAE_/3GB = SFW 4.1 with 4GT (/3GB switch) and no PAE
- SFW41_base_/3GB = SFW 4.1 with 4GT (/3GB switch) and PAE enabled
- VxCache_1gb = VxCache with 1 GB of cache and PAE enabled
- VxCache_1gb_/3gb = VxCache with 1 GB of cache and PAE and 4GT (/3GB switch) enabled
- VxCache_2gb = VxCache with 2 GB of cache and PAE enabled
- VxCache_1gb_/3gb = VxCache with 2 GB of cache and PAE and 4GT (/3GB switch) enabled
- VxCache_4gb = VxCache with 4 GB of cache and PAE enabled
- RT95 = Loadsim 2003 response time in milliseconds, based on the 95th percentile
- Difference (%) = Percentage difference in response time compared with the baseline. Increasingly negative numbers indicate a faster response time. Positive numbers indicate a slower response time.
- Performance Impact (%) = A percentage representation of how performance is affected, for example, 4% indicates a 4 percent increase in performance; -1% indicates a 1 percent decrease in performance.
- Up arrows indicate an increase in performance; down arrows indicate a decrease in performance.

Veritas Storage Foundation *for Windows*:
Enhancing Application Performance with VxCache

Table 2. VxCache versus the baseline

Configuration	RT95	Difference (%)	Performance Impact (%)
SFW41_base	250.5	baseline	baseline
VxCache_1gb	238.5	-4.8	4.8 (↑)
VxCache_2gb	227	-9.4	9.4 (↑)
VxCache_4gb	221	-11.8	11.8 (↑)

Table 3. 4-gigabyte tuning (/3GB switch) comparison

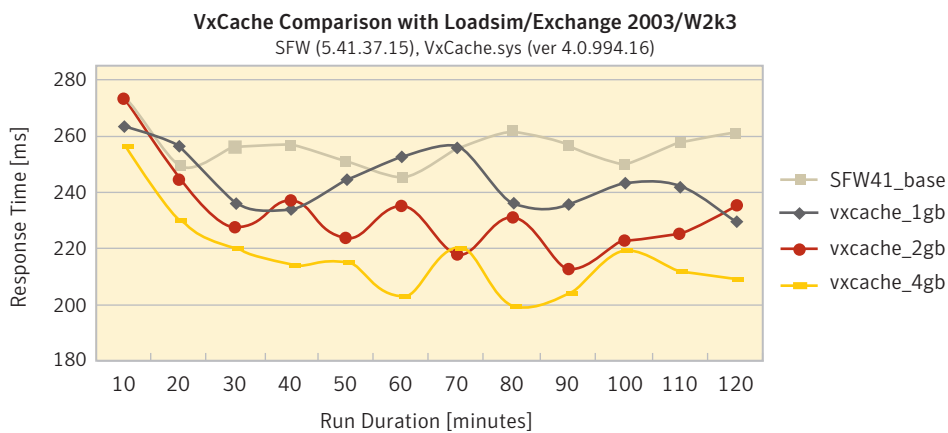
Configuration	RT95	RT95_3GB	Difference (%)	Performance Impact (%)
SFW41_base	250.5	243.5	-2.8	2.8 (↑)
VxCache_1gb	238.5	238	-0.2	0.2 (↑)
VxCache_2gb	227	231	1.8	-1.8 (↓)

Table 4. Performance summary

Configuration	RT95	Difference (%)	Performance Impact (%)
SFW41_base	250.5	baseline	0
SFW41_noPAE_/3GB	249.5	-0.4	0.4 (↑)
SFW41_base_/3GB	243.5	-2.8	2.8 (↑)
VxCache_1gb	238.5	-4.8	4.8 (↑)
VxCache_1gb_/3GB	238	-5.0	5.0 (↑)
VxCache_2gb	227	-9.4	9.4 (↑)
VxCache_2gb_/3GB	231	-7.8	7.8 (↑)
VxCache_4gb	221	-11.8	11.8 (↑)

Veritas Storage Foundation *for Windows*: Enhancing Application Performance with VxCache

Figure 1 showed a direct relationship between increased performance and the amount of memory available, specifically, the more VxCache memory, the shorter the response time and consequently the better the performance. Figure 3 shows how performance varies for different amounts of cache as run duration increases. On average, performance increases (response times are lower) as cache increases. Note that 4GT (/3GB switch) testing suggested that performance was slightly improved when using Veritas Storage Foundation *for Windows* 4.1 without VxCache, while it was slightly lower when VxCache was enabled. Table 2 shows that enabling the /3GB switch on the baseline configuration (no VxCache) improved performance by 2.8 percent. When VxCache was enabled in addition to the /3GB switch, performance remained essentially unchanged with 1 GB VxCache memory, and it actually dropped at the 2 GB VxCache level. This can probably be attributed to the limited amount of nonpaged pool memory available for memory management. However, drawing a final conclusion on using 4 gigabyte tuning (/3GB switch) is difficult, as results fluctuate within a 2 percent margin, which is within the margin of error attributable to the testing environment.



Notes: Lower response time means better performance

Figure 3. Cache performance response time versus run duration

Best practices

- VxCache requires a minimum of 512 MB of memory on a system before it can be enabled. While the Administrator's Guide recommends a minimum of 1 GB to allow 500 MB for the OS and applications and 500 MB for VxCache, the actual memory (above 512 MB) required for VxCache would depend on the types of volumes on which it will be enabled and the way data is accessed. For volumes that "qualify" for VxCache, the more memory available for caching (up to 4 GB), the better the performance.
- Determine the amount of memory available on the system for VxCache and the volumes that would benefit from VxCache before enabling it.
- Enable VxCache on read-intensive volumes such as the volumes containing Exchange databases. Write-intensive volumes would not benefit from VxCache.
- Applications that are AWE aware, such as Microsoft SQL, would not benefit from VxCache, as they have direct access to memory above 4 GB.
- Using the /3GB switch with VxCache should be done carefully, as it reduces the amount of nonpaged pool memory available for managing memory made available by VxCache and can result in diminishing returns as more memory is assigned for VxCache.
- Use VxCache Volume Monitoring to monitor the read/write ratio of a volume over time to determine whether or not performance is being increased and whether or not VxCache should be enabled on that volume.

Conclusion

Reading data from memory is always faster than from disk. VxCache can be enabled at the volume level, making it possible to enable it only on the volumes that would benefit from it. Performance enhancements of up to 40 percent have been observed for Microsoft Exchange when VxCache is enabled. And with its volume-level granularity, you can choose the volumes on which to enable VxCache, optimizing performance.

References

1. PAE X86 and 4GT technical references on Resource and Memory Management Technologies page of Microsoft's Web site (www.microsoft.com/resources/documentation/windowsServ/2003/all/techref/en-us/W2K3TR_scale_memgt_over.asp)
2. HowStuffWorks Web site (www.howstuffworks.com)

About Symantec

Symantec is a global leader in infrastructure software, enabling businesses and consumers to have confidence in a connected world.

The company helps customers protect their infrastructure, information, and interactions by delivering software and services that address risks to security, availability, compliance, and performance. Headquartered in Cupertino, Calif., Symantec has operations in 40 countries.

More information is available at www.symantec.com.

For specific country offices and contact numbers, please visit our Web site. For product information in the U.S., call toll-free 1 (800) 745 6054.

Symantec Corporation
World Headquarters
20330 Stevens Creek Boulevard
Cupertino, CA 95014 USA
+1 (408) 517 8000
1 (800) 721 3934
www.symantec.com

Copyright © 2007 Symantec Corporation. All rights reserved. Symantec, the Symantec Logo, Veritas, and Veritas Storage Foundation are trademarks or registered trademarks of Symantec Corporation or its affiliates in the U.S. and other countries. Microsoft and Windows are either registered trademarks or trademarks of Microsoft Corporation in the United States and/or other countries. Other names may be trademarks of their respective owners. Printed in the U.S.A.
03/07 12151277