

Best Practices

Microsoft® Exchange Server 2007 with Celeros SmartSAN™ Storage System Best Practices

© 2010-2011 Celeros Corporation
December, 2010



Table of Contents

| | |
|--|----|
| Introduction | 3 |
| Basics of Performance Tuning | 3 |
| Understanding the Context for Performance Tuning | 3 |
| Three Components That Influence Performance | 4 |
| Microsoft Exchange Server 2007 Considerations | 4 |
| Analyzing I/O Characteristics..... | 5 |
| Calculating Theoretical Usage | 5 |
| Storage Group and Database Design | 7 |
| Locating Databases within Storage Groups | 7 |
| Backup and Restore Factors | 7 |
| Determining the Best RAID Level | 7 |
| Choosing the Number of Drives to Put in a Volume Group | 8 |
| Clustering | 8 |
| Considering the Server Platform | 8 |
| Sharing Bandwidth with Multiple HBAs | 9 |
| Host Connectivity | 9 |
| Microsoft Windows Server Considerations | 10 |
| Aligning Host I/O with RAID Striping..... | 10 |
| NTFS Allocation Unit Size | 10 |
| Fibre Channel Switch Settings | 10 |
| Using Command Tag Queuing | 10 |
| Setting up the Storage System | 11 |
| Factors Influencing Storage Performance | 11 |
| Estimating Capacity Limits..... | 11 |
| Setting Global Parameters | 12 |
| Setting LUN-Specific Parameters | 13 |
| Setting the LUN-Specific Write Cache and Write Cache Mirroring | 13 |
| Setting the LUN-Specific Read Cache and Read Prefetch | 13 |
| Locating Volume Groups | 13 |
| Striping Volume Groups in Controller Enclosure..... | 13 |
| Using the SmartSAN Performance Monitor | 14 |
| Summary of Performance Tuning Best Practices | 14 |
| Conclusion | 15 |

Best Practices for Microsoft Exchange Server 2007 with the Celeros SmartSAN Storage System

Introduction

This document details the requirements necessary to implement a SmartSAN storage system in a Microsoft Exchange Server 2007 environment. This document is not intended to be an installation and configuration guide for Microsoft Exchange Server 2007, but rather a guide to tune and optimize the storage system within the Exchange environment.

For additional information about Exchange Server 2007, refer to the Microsoft Exchange server 2007 product page. Go to:

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

Disclaimer

Due to the highly customizable nature of an Exchange environment, you must take into consideration the specific environment and equipment to achieve optimal performance from a SmartSAN storage system. Also note that many procedures within this document should be attempted only by experienced Microsoft Exchange Server 2007 administrators and trained storage specialists with intimate knowledge of the working environment. This document results from the collaboration of numerous Celeros engineers and performance specialists.

Basics of Performance Tuning

Understanding the Context for Performance Tuning

The challenge of storage performance tuning is to understand and control interacting factors while accurately measuring Microsoft Exchange Server 2007 performance. Because the performance of the storage system accounts for only a portion of the overall Exchange performance, tuning must be done in context. The full context includes other contributing factors such as the I/O characteristics of Exchange and the components in the data path. These contributing factors include the following:

- Switches
- File system
- Operating system
- Server

With multiple parameters to consider, the task of performance tuning can seem difficult. To reduce the complexity of tuning, the SmartSAN storage system offers the performance monitoring features and flexible tuning controls in SmartSAN Storage Manager.

Three Components That Influence Performance

This document provides an overall approach to tuning I/O performance and also provides specific guidelines for using the storage system tuning controls. These recommendations begin with an overall analysis of the three elements that determine I/O performance:

- The Microsoft Exchange Server 2007 application
- Server platform (hardware, operating system, device drivers)
- SmartSAN storage system

Microsoft Exchange Server 2007 Considerations

When planning for Exchange 2007 storage, performance should be the primary focus, with space planning a secondary concern. A system that is first optimized for Exchange I/O performance will generally result in adequate space as well.

Measuring or predicting the performance of transactional I/O facilitates in determining whether the configuration meets the requirements for acceptable disk latency and disk throughput. For Microsoft Exchange Server 2007, the storage system is considered to be performing poorly if average read and write disk latencies are greater than 20ms for database volumes and 10ms for log volumes. Frequent or extended latency spikes above 50ms are also indicators of poor performance.

It is critical that you develop an understanding of the Exchange Server 2007 I/O requirements for your environment. For additional information, refer to the Microsoft document, *What Causes Exchange Disk I/O?* Go to:

<http://technet.microsoft.com/en-us/library/bb738154.aspx>

Analyzing I/O Characteristics

Microsoft Exchange Server 2007 accesses data either randomly or sequentially, depending on the data being accessed.

Database access is in random 8KB page sizes. However, I/O coalescing generates larger I/O sizes—up to 1MB. By increasing the page size to 8KB, an increase of 4KB over Microsoft Exchange Server 2003, disk I/O decreases by allowing larger amounts of data to be accessed in a single I/O. Exchange 2007 databases also take advantage of the 64-bit architecture allowing for a larger database cache, which results in a decrease in the read percentage of the total disk I/O. A system configured in accordance with Microsoft's recommended memory configuration guidelines results in a database read-to-write ratio of approximately 1:1.

Log access is sequential with writes ranging from 512 bytes up to the 1MB buffer size. All changes made to the database are written to the transaction logs first. Exchange 2007 Server transaction log files are 1MB in size, a decrease from 5MB in MS Exchange Server 2003. Due to a decrease in database reads and a smaller log file size, the log-to-database write ratio is approximately 3:4.

Calculating Theoretical Usage

When planning a new Microsoft Exchange Server 2007 implementation, it is critical that you determine the total number of users in your configuration and the approximate user mailbox size before you create storage volumes. Without these values, you cannot accurately determine the amount of storage to allocate. Remember to allow for future growth in your calculations.

After you have determined the total number of users, use the following procedure to predict the number of IOPS the Exchange configuration must support. Table 1 shows the estimated user IOPS profile as well as recommended user mailbox size.

Table 1 User Profiles

| User Type (Usage Profile) | Send/Receive Per Day (Approximately 50 KB Message Size) | Estimated IOPS Per Use (Cached Mode) | Mailbox Size |
|---------------------------|---|--------------------------------------|--------------|
| Light | 5 sent/20 received | 0.11 | 100 MB |
| Average | 10 sent/40 received | 0.18 | 250 MB |
| Heavy | 20 sent/80 received | 0.32 | 1024MB |
| Very Heavy | 30 sent/120 received | 0.48 | 2048 MB |

To determine the required database IOPS:

Number of users x (I/O profile of user + 20 percent for overhead) = required IOPS for database volumes

EXAMPLE

$4000 \text{ (users)} \times 0.384 \text{ (heavy user profile + 20 percent overhead}^1\text{)} = 1536 \text{ IOPS required for database volumes}$

To determine the required log IOPS:

Assume a database read-to-write ratio of 1:1. Continuing with the previous example, out of 1536 IOPS for the database volumes, there would be 768 IOPS for reads and 768 IOPS for writes.

Follow the log-to-database write ratio of 3:4 to determine the required IOPS for log volumes.

$768 \text{ DB writes/sec} \times (.75) = 576 \text{ log writes/sec}$

To determine the total required IOPS:

Total required IOPS = required IOPS for database volumes + required IOPS for log volumes

EXAMPLE

$1536 \text{ IOPS} + 576 \text{ IOPS} = 2112 \text{ IOPS}$

Divide the total IOPS by the number of Exchange storage groups to determine the IOPS each storage group will need to support. When configuring the storage volumes, this value should be the minimum IOPS that the storage group volume can support.

Total IOPS \div Number of storage groups = IOPS per storage group.

NOTE Prior to creating storage system volumes it is highly recommended that you use the Exchange 2007 Mailbox Server Role Storage Requirements Calculator developed by Microsoft. This calculator helps you determine the storage requirements, including I/O performance and capacity requirements. This calculator also suggests the optimal LUN layout based on the input parameters.

For additional information about the Mailbox Storage Calculator for Exchange 2007, including usage instructions, refer to the Microsoft document, *Mailbox Server Storage Design*. Go to:

<http://technet.microsoft.com/en-us/library/bb738147.aspx>

¹ Add 20 percent to the estimated IOPS to allow for additional Exchange activity during peak hours.

Storage Group and Database Design

Locating Databases within Storage Groups

In Exchange Server 2007, Microsoft recommends locating all databases in a storage group on the same physical LUN. It is also a best practice to place no more than one database in each storage group. Do not place multiple storage groups on the same physical LUN: use only one storage group per LUN. Remember to verify that your chosen LUN design can accommodate the required IOPS per storage group.

Backup and Restore Factors

If you plan on taking advantage of the Celeros VSS hardware provider to implement Exchange 2007 snapshot backups, design the LUNs to minimize impact to separate storage groups. Each Exchange 2007 storage group should contain a single transaction log LUN and a single database LUN that are independent of other storage groups. This design allows you to back up and restore a single Exchange storage group with no performance impact to other storage groups that might exist.

Determining the Best RAID Level

RAID1/10² and RAID5 are all viable options for Exchange 2007 database volumes because these RAID levels have similar characteristics for read environments. However, RAID1/10 offers the best performance while still providing data protection by mirroring each physical disk. In an Exchange 2007 environment, the transaction logs are the most important data set and low write latency is critical for server performance. With that in mind, logs should be placed on RAID1/10 volumes with protected write cache. Latency for the transaction log volumes should be less than 10ms.

Considering Exchange Server 2007's database read-to-write ratio of 1:1, RAID5 volumes are at a disadvantage due to the increased writes. RAID5 is challenged most by random writes, which can generate multiple drive I/Os for each host write. However, when designed properly, RAID5 provides the required IOPS and does so using fewer disks than RAID1/10. When you design your volumes, make sure to use enough disks in RAID5 volume groups to support the projected IOPS requirement calculated earlier in the "Calculating Theoretical Usage" section on page 3. Latency for the database volumes should be less than 20ms.

NOTE Different RAID levels can be tested by using the SmartSAN Dynamic RAID Migration feature, which allows the RAID level of a volume group to be changed while maintaining continuous access to the data.

² SmartSAN Storage Manager only shows the option RAID1. However, if you select four or more drives, RAID10 is automatically configured across the volume—two drives for user data and two for the mirrored data.

Choosing the Number of Drives to Put in a Volume Group

Microsoft Exchange Server 2007 is an I/O intensive application. Therefore, you must select a number of drives that matches the per volume group I/O rate needed to support the application. Make sure to account for the I/Os required to implement the data protection of the selected RAID level. You can add more drives to the configuration for a linear increase in performance, up to the point of controller saturation. More drives provide more spindles to service I/O. For example, in a RAID5 database volume, you can add more spindles to service I/O.

The higher the spin speed of the drive, the better. A 15K RPM drive has one-half less rotational latency than a 7.2K RPM drive.

NOTE The spindle count of an existing volume group can be increased using the Dynamic Capacity Expansion feature of SmartSAN Storage Manager.

Clustering

Microsoft Exchange Server 2007 offers multiple options for clustering:

- Local Continuous Replication
- Cluster Continuous Replication
- Single Copy Clusters
- Standby Continuous Replication

Here are some points to consider when implementing Exchange 2007 clustering.

- **Local continuous replication (LCR)** – A single-server solution that uses built-in technology to create and maintain a copy of a storage group on a second set of disks that are connected to the same server as the production storage group. LCR requires a second storage system to host the passive copy. The additional storage system needs similar capacity and performance as the active copy.
- **Cluster continuous replication (CCR)** – Combines the replication and replay features in Exchange 2007 with failover features in Windows Cluster services. CCR requires a second storage system to host the passive copy. The additional system needs similar capacity and performance as the active copy.
- **Single copy clusters (SCC)** – These clusters use Microsoft Windows Cluster services to manage access and node failover to a shared storage system.
- **Standby continuous replication (SCR)** – A new feature available in Exchange 2007 Service Pack 1. SCR is designed for locations that use standby recovery servers for high availability. SCR requires a second storage system to host the passive copy. The additional system needs similar capacity and performance as the active copy.

For more detailed information about clustering in an Exchange environment, refer to the Microsoft document, *High Availability*. Go to:

<http://technet.microsoft.com/en-us/library/bb124721.aspx>

Considering the Server platform

The server platform consists of the server hardware and the system software.

Considering the Server Hardware Architecture

Available bandwidth depends on the server hardware. The number of buses adds to the aggregate bandwidth, but the number of HBAs sharing a single bus can throttle the bandwidth.

Calculating Aggregate Bandwidth

An important limiting factor in I/O performance is the I/O capability of the Exchange server. The aggregate bandwidth of the server to the storage system is measured in MB/s and consists of the total capability of the buses to which the storage system is connected. For example, a 64-bit PCI bus clocked at 133MHz has a maximum bandwidth calculated by the following formula:

$$\text{PCI bus throughput (MB/s)} = \text{PCI Bus Width} \div 8 \times \text{Bus Speed}$$

$$64 \text{ bit} \div 8 \times 133\text{MHz} = 1062\text{MB/s} \approx 1\text{GB/s}$$

| MHz | PCI Bus Width | Throughput (MB/s) |
|-----|---------------|-------------------|
| 66 | 64 | 528 |
| 100 | 64 | 800 |
| 133 | 64 | 1064 |
| 266 | 64 | 2128 |
| 533 | 64 | 4264 |

Sharing Bandwidth with Multiple HBAs

Multiple HBAs on a bus share this single source of I/O bandwidth, and each HBA might have multiple ports, which can operate at various speeds. As a result, the ability to drive a storage system can be throttled³ by either the server bus or by the HBAs. Therefore, whenever you configure an Exchange server or whenever you analyze I/O performance, you must know how much server bandwidth is available and which devices are sharing that bandwidth.

Host Connectivity

The SmartSAN storage system was designed with high availability and performance in mind. The SmartSAN storage system features redundant power supplies, redundant controllers, redundant cache, and redundant internal architecture. To maximize this redundant technology and to prevent a single point of failure, redundant host connectivity and switching is recommended.

³ *Throttle* – To slow down I/O processing during low memory conditions, processing one sequence at a time in the order the request was received.

Microsoft Windows Server Considerations

Aligning Host I/O with RAID Striping

On physical disks that maintain 64 sectors per track, Microsoft Windows creates the partition starting at the 64th sector. Starting at the 64th sector causes misalignment with the underlying RAID striping and allows for the possibility of a single I/O operation to span multiple segments causing performance degradation.

Use the Microsoft Diskpart.exe utility to align the sectors of all Exchange 2007 storage volumes before formatting. Microsoft Exchange Server 2007 writes data in multiples of 8KB I/O operations, and I/O operation to a database can be from 8KB to 1MB. Therefore, make sure that the starting offset is a multiple of 8KB. Celeros recommends starting with a value of 64KB.

Microsoft provides the diskpart.exe utility as part of Windows 2003 Service Pack 1 support tools. For detailed steps about how to use Diskpart to align I/O with storage track boundaries, refer to the Microsoft document, *How to Align Exchange I/O with Storage Track Boundaries*. Go to:

<http://technet.microsoft.com/en-us/library/aa998219.aspx>

NTFS Allocation Unit Size

Configure the NTFS volumes that host the Exchange 2007 databases with an NTFS allocation unit size of 64KB. The recommendation of 64KB is based on the performance improvements seen with large sequential read operations.

Format the transaction log volumes with the default allocation unit size of 4KB. Testing has shown that changing to 64KB does not result in any increase in transaction log sequential throughput.

Fibre Channel Switch Settings

Use the Fibre Channel switch settings recommended by Celeros. For example, on Brocade switches, verify that the In-Order Delivery parameter is enabled. In a multi-switch SAN fabric, where I/O traverses inter-switch links, make sure to configure sufficient inter-switch link bandwidth.

Using Command Tag Queuing

Command Tag Queuing (CTQ) refers to the controller's ability to line up multiple SCSI commands for a single LUN and execute the commands in an optimized order that minimizes rotational latency and seek latency. The Celeros models vary in CTQ capability, generally up to 2048 per controller. CTQ is enabled by default on The SmartSAN storage system, but you must enable CTQ in Windows Server and on the HBA. Refer to the documentation from the HBA vendor for the necessary registry changes.

Setting up the Storage System

Factors Influencing Storage Performance

Drive I/O capacity is at the very heart of storage system performance. For The SmartSAN storage system, the number of drives in a configuration usually establishes the upper bound for storage system performance. Many various interacting factors determine how much of the raw performance of a group of drives that Microsoft Exchange Server 2007 can use. These factors include the following:

- Size of the cache
- Algorithms that manage the cache
- Number and type of host and drive-side channels
- How RAID parity calculations are performed
- Whether SCSI commands are queued for optimized execution by the controllers
- How controllers choose data paths

Estimating Capacity Limits

In Exchange Server 2007 environments, the number of drives in the drive group largely determines performance. The maximum IOPS from drive for a storage system is typically specified with a full complement of drives. Performance slows with fewer drives and can be approximated by a simple ratio.

The following formula yields an approximation only:

$$\text{Max I/O Rate from Max \# of Drives Supported} \times \text{Actual \# of Drives} \div \text{Max \# of Drives Supported} = \text{Max I/O Rate from Actual \# of Drives}$$

Many factors determine IOPS, including drive type, RPM, data layout, varying I/O sizes, volume group layout, controller architecture and workload.

Calculating Optimal Segment Size

The Celeros term segment size refers to the amount of data written to one drive in a volume group before writing to the next drive in the volume group. For example, in a RAID5

4+1 volume group with a segment size of 128KB, the first 128KB of the LUN storage capacity is written to the first drive, the next 128KB to the second drive, and so forth.

For a RAID1 2+2 volume group, 128KB of an I/O would be written to each of the two data drives and to the mirrors. If the I/O size is larger than the number of drives multiplied by 128KB, this pattern repeats until the entire I/O is completed.

The formula for optimal segment size is as follows:

$$\text{LUN segment size} = \text{LUN stripe width} \div \text{number of data drives}$$

For RAID5, the number of data drives is equal to the number of drives in the volume group minus 1. For example:

$$\text{RAID5, 4+1 with a 64KB segment size} \Rightarrow (5-1) \times 64\text{KB} = 256\text{KB stripe width}$$

For RAID1, the number of data drives is equal to the number of drives divided by 2. For example:

RAID1, 2+2 with a 64KB segment size => (2) x 64KB = 128KB stripe width

For small I/O requests, the segment size should be large enough to minimize the number of segments (drives in the LUN) that must be accessed to satisfy the I/O request to minimize segment boundary crossings. For Microsoft Exchange Server 2007 volumes, set the segment size to 64K or larger, so that the stripe width is at least as large as the median I/O size.

Setting Global Parameters

Setting the Caching Parameters

The cache block size is a global parameter for the storage system. Set the cache block size to 16KB. This large block size works for Microsoft Exchange Server 2007 for two reasons:

- The mix of random and sequential I/O used by Exchange Server 2007
- The use of I/O coalescing, which results in fewer but larger I/O

You easily can change the cache block size at any time to optimize for a particular workload during a specific time period.

Setting the Global Cache Flush

Two global parameters, Start Flushing and Stop Flushing, are provided to control the flushing of write data from the controller cache to the drives. Flushing begins when the percentage of unwritten data cache exceeds the Start Flushing level and stops when the percentage hits the Stop Flushing mark. Celeros recommends setting both parameters to the same value to cause a brief flushing operation to maintain a specified level of free space. Start with the default values and experiment. Values of 80 and 80, for Start Flushing and Stop Flushing respectively, are normally sufficient.

Setting the Global Media Scan

The impact of Media Scan is minimal, but the extra reads do represent a finite workload.

Therefore, consider the performance demands when setting Media Scan.

- In most cases, enable Media Scan and set the scan frequency to 15 days to enable periodic scans of the surface of all LUN drives.
- When absolute maximum performance is the objective, do not enable Media Scan.

You also can enable or disable Media Scan for each volume. See the following section on "Setting LUN-Specific Parameters."

Setting LUN-Specific Parameters

One way to limit the workload caused by Media Scan is to enable or disable Media Scan at the volume level, rather than globally.

- In most cases, enable Media Scan for each volume.
- If the goal is to maximize the performance of a LUN, disable Media Scan for a specific volume.

Setting the LUN-Specific Write Cache and Write Cache Mirroring

Write caching, or caching a drive segment to a memory buffer before writing to the drive, can increase I/O performance during data transfers. However, write caching also increases the risk of data loss if a controller or the controller's memory fails while unwritten data resides in cache memory. SmartSAN's write-cache mirroring protects data during a controller failure or a cache-memory failure. When you enable write cache mirroring, cached data is mirrored across two redundant controllers with the same cache size. The data written to the cache memory of one controller is also written to the cache memory of the alternate controller. Therefore, if one controller fails, the alternate can complete all outstanding write operations.

Enabling write caching with mirroring generally improves performance on Exchange 2007 volumes and is recommended.

Setting the LUN-Specific Read Cache and Read Prefetch

Always enable read cache. Enabling read cache allows the controllers to service reads from cache for any additional read requests to the data stored within that cache.

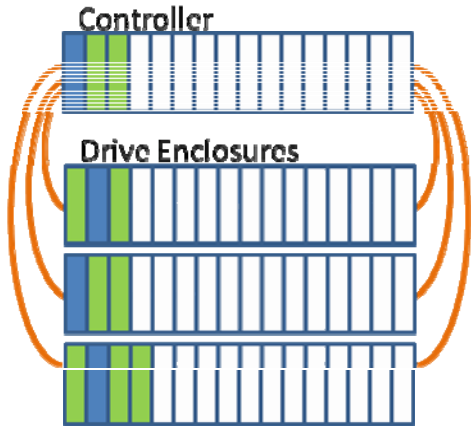
Cache read prefetch should be disabled on all Exchange 2007 volumes. This prevents the controllers from reading ahead and retrieving segments of data that will most likely not be used.

Locating Volume Groups

How you locate the volume groups depends on the controller model you are using.

Striping Volume Groups in Controller Enclosure

Stripe volume groups across all available drive enclosures to distribute I/O evenly across the drive-side "loops." A loop is an A/B pair of channels. Models might have one or two drive loops, but the approach is the same. Figure 1 shows the striping for two loops.



8+1 group striped across even and odd drives of 4 drive enclosures

Figure 1 Striping Volume Groups in Controller Enclosure Models

Because the green volume group (8+1) has more drives than there are drive enclosures, the green volume group uses multiple drives on all drive enclosures. If access to an entire drive enclosure is lost, the volume group data becomes inaccessible. However, the high component reliability and redundant design elements in The SmartSAN storage system makes losing access to an entire drive enclosure an extremely remote possibility. The choice for maximum tray protection is really targeted at human cabling errors during on-the-fly upgrades and drive enclosure reconfigurations.

Using the SmartSAN Performance Monitor

SmartSAN Storage Manager provides an integrated Performance Monitor that reports the following statistics for each volume in the storage system.

Table 3 Performance Monitoring Statistics

| Statistic | Description |
|----------------------|---|
| Total I/Os | Since Start of this monitoring session |
| Read Percentage | Percent of Read I/Os |
| Cache Hit Percentage | Percent of Reads satisfied from cache |
| Current KB/sec | Since last polling interval or requested Update |
| Max. KB/sec | Highest value since last Start |
| Current I/O/sec | Since last polling interval or requested Update |
| Max. I/O/sec | Highest value since last Start |

Use the Performance Monitor to guide the tuning process. Observe the cache hit percentage and the read/write mix for each LUN of interest while an application is running.

For detailed usage information about Performance Monitor, see the SmartSAN Storage Manager online help.

Summary of Performance Tuning Best Practices

Below is a summary of the recommended best practices to obtain optimal performance from a SmartSAN storage system in a Microsoft Exchange Server 2007 environment.

- Use all available host-side channels. Storage partitioning, switch zoning (where applicable), and multi-pathing HBA drivers are all useful tools to ensure that all channels are kept active.
- Balance I/O across the dual controllers of the storage system and strive to keep both controllers busy. Locate the database volume and its corresponding log volume on separate paths.
- Choose faster drives. A 15K RPM drive has one-half less rotational latency than a 7.2K RPM drive.
- Add more drives to the configuration for a linear increase in performance, up to the point of controller saturation. More drives provide more spindles for I/O.
- For the log LUNs, use RAID1/10 volumes and enable read caching and write caching with mirroring.
- For the database LUNs, use RAID1/10, RAID6, or RAID5, and enable read caching and write caching with mirroring.
- Configure the entire capacity of a volume group into a single volume. Multiple LUNs on one volume group typically increases seek time penalties.
- Isolate Exchange 2007 database and the log disk I/O on separate physical volume groups. Separate the transaction log files (sequential I/O) from databases (random I/O) to maximize I/O performance and increase fault tolerance. Storing sequentially accessed files separately keeps the disk heads in position for sequential I/O, reducing the amount of time required to locate data. For recoverability, separate a storage group's transaction logs and databases to ensure that a catastrophic failure of a specific set of disks does not cause a loss of both database and transaction logs.
- Choose a segment size of 64K or larger, so that the stripe width is at least as large as the median I/O size.
- Isolate Exchange 2007 volumes from other drive-intensive applications. Sharing the storage with other applications might have a negative impact on Exchange 2007 performance.
- If Auto Volume Transfer is not required for failover by any of the hosts that will use the storage system, disable Auto Volume Transfer in all host regions. Consult with a Celeros Customer and Technical Support representative for the procedure.

Conclusion

This document delivers important general guidelines for fine-tuning the Windows Server, Microsoft Exchange Server 2007, and the SmartSAN storage system. To continue to improve performance, learn as much as possible about the requirements for a Microsoft Exchange Server 2007 environment.

References

Microsoft [2007] *High Availability* [Clustering in an Exchange environment]

<http://technet.microsoft.com/en-us/library/bb124721.aspx>

Microsoft [2006] *How to Align Exchange I/O with Storage Track Boundaries* [Microsoft's usage details on diskpar]

<http://technet.microsoft.com/en-us/library/aa998219.aspx>

Microsoft [2007] *Mailbox Server Storage Design* [Using the Mailbox Storage Calculator for Exchange 2007 to predict storage needs]

<http://technet.microsoft.com/en-us/library/bb738147.aspx>

Microsoft [2006] *What Causes Exchange Disk I/O?* [Exchange Server 2007 I/O requirements]

<http://technet.microsoft.com/en-us/library/bb738154.aspx>

Legal Notices

Document Description

This document describes the best practices for using Microsoft Exchange 2007 with Celeros SmartSAN™ Storage Systems.

Ownership of Materials

The Document is provided as a courtesy to customers and potential customers of Celeros® Corporation ("Celeros"). Celeros assumes no obligation to correct any errors contained herein or to advise any user of liability for the accuracy or correctness of information provided herein to a user. Celeros makes no commitment to update the Document. Celeros reserves the right to change these legal terms and conditions from time to time at its sole discretion. In the case of any violation of these rules and regulations, Celeros reserves the right to seek all remedies available by law and in equity for such violations. Except as expressly provided herein, Celeros and its suppliers do not grant any express or implied right to you under any patents, copyrights, trademarks, or trade secret information. Other rights may be granted to you by Celeros in writing or incorporated elsewhere in the Document.

Trademark Acknowledgments

EzSAN, EzNAS, EzSANFiler, SmartSAN are trademarks or registered trademarks of Celeros Corporation. All other brand and product names may be trademarks of their respective companies.

Performance Information

Performance tests and ratings are measured using specific computer systems and/or components and reflect the approximate performance of Celeros products as measured by those tests. Any difference in system hardware or software design or configuration may affect actual performance. Buyers should consult other sources of information to evaluate the performance of systems or components they want to purchase.

Disclaimer

Celeros has provided this Document to enable a user to gain an understanding of SmartSAN . This document is designed to assist a user in making a general decision as to whether a SmartSAN storage system configuration is appropriate for such user's objectives. Neither this Document nor the Tool are designed or intended to guarantee that the configuration a user chooses will work in a specific manner. While the guidance provided by this Document can help a user to choose an appropriate configuration (or avoid a configuration that is not appropriate), there is no way Celeros can guarantee the exact performance and/or results of the Information contained in this Document. Accordingly, Celeros assumes no obligation whatsoever for the use of the Information provided in this Document, AND UNDER NO CIRCUMSTANCES WILL CELEROS OR ITS AFFILIATES BE LIABLE UNDER ANY CONTRACT, STRICT LIABILITY, NEGLIGENCE OR OTHER LEGAL OR EQUITABLE THEORY, FOR ANY SPECIAL, INDIRECT, INCIDENTAL, PUNITIVE OR CONSEQUENTIAL DAMAGES OR LOST PROFITS IN CONNECTION WITH THIS DOCUMENT.

THE INFORMATION AND MATERIALS PROVIDED IN THIS DOCUMENT IS PROVIDED "AS IS" AND CELEROS MAKES NO WARRANTIES EXPRESS, IMPLIED OR STATUTORY, INCLUDING, WITHOUT LIMITATION, THE IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE WITH RESPECT TO THE SAME. CELEROS EXPRESSLY DISCLAIMS ANY WARRANTY WITH RESPECT TO ANY TITLE OR NON-INFRINGEMENT OF ANY THIRD PARTY INTELLECTUAL PROPERTY RIGHTS, OR AS TO THE ABSENCE OF COMPETING CLAIMS, OR AS TO INTERFERENCE WITH USER'S QUIET ENJOYMENT.