

EMC CLARiiON CX4 Enterprise Flash Drives with Microsoft Exchange

Applied Technology

Abstract

This white paper examines the use of Microsoft Exchange with enterprise Flash drives, including use cases, performance characteristics, and general guidelines for database file placement.

June 2009

Copyright © 2008, 2009 EMC Corporation. All rights reserved.

EMC believes the information in this publication is accurate as of its publication date. The information is subject to change without notice.

THE INFORMATION IN THIS PUBLICATION IS PROVIDED “AS IS.” EMC CORPORATION MAKES NO REPRESENTATIONS OR WARRANTIES OF ANY KIND WITH RESPECT TO THE INFORMATION IN THIS PUBLICATION, AND SPECIFICALLY DISCLAIMS IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Use, copying, and distribution of any EMC software described in this publication requires an applicable software license.

For the most up-to-date listing of EMC product names, see EMC Corporation Trademarks on EMC.com

All other trademarks used herein are the property of their respective owners.

Part Number h5966.1

Table of Contents

Executive summary	4
Introduction	4
Audience	4
Technology overview	4
EMC CLARiiON CX4 series	4
Enterprise Flash drives	5
Microsoft Exchange on enterprise Flash drives	6
Placing “Heavy” Exchange users onto Flash drives	7
Exchange environments best suited for enterprise Flash drives	7
Exchange on Flash recommendations	9
Exchange log and database file placement	9
RAID protection	9
Cache setting.....	9
Exchange virtualization	9
Enterprise Flash drive vs. hard disk drive performance with Exchange.....	9
Extending the value of enterprise Flash drives.....	10
EMC SourceOne for Information Lifecycle Management.....	10
CLARiiON Virtual Provisioning	12
Conclusion	13
References	13
Appendix: Jetstress results	14
Microsoft Exchange Server Jetstress: Results for TEST1	14
28 HDD (24 drives in four RAID 1/0 3+3 for DB and one RAID 1/0 2+2 for all logs)	14
Microsoft Exchange Server Jetstress: Results for TEST2.....	16
5 EFDs (RAID 5 4+1) for the Exchange DBs and 4 fibre drives (RAID 1/0 2+2) for the Exchange logs (Write cache disabled on the EFD LUNs)	16
Microsoft Exchange Server Jetstress: Results for TEST3.....	18
5 EFDs (RAID 5 4+1) for the Exchange DBs and 4 fibre drives (RAID 1/0 2+2) for the Exchange logs (Write cache enabled on the LUNs)	18

Executive summary

The EMC CLARiiON® FLARE® operating system integrates enterprise-class Flash drives into the CLARiiON CX4 storage array. With Flash drives, EMC has created a new ultra-performance storage tier that removes previous performance limitations imposed by magnetic disk drives.

For years, the most demanding enterprise applications have been limited by the performance of magnetic disk media. Also known as EFDs, enterprise Flash drives deliver unprecedented performance and response times, which are benefits well suited for latency-sensitive applications like Microsoft Exchange

Enterprise Flash drives, also known as solid state drives (SSD), contain no moving parts, which removes much of the storage latency associated with traditional magnetic disk drives. A CLARiiON CX4 with enterprise Flash drives can deliver application response times of just a few milliseconds and up to 30 times more I/O operations per second (IOPS) than traditional Fibre Channel hard disk drives (HDD). Additionally, because there are no mechanical components, Flash drives consume significantly less energy than hard disk drives. When replacing a larger number of HDDs with a lesser number of enterprise Flash drives, energy consumption can be reduced up to 90 percent or more for a given IOPS workload.

The high-performance characteristics of enterprise Flash drives eliminate the need for organizations to purchase large numbers of traditional hard disk drives, while only utilizing a small portion of their capacity to satisfy the IOPS requirements of Microsoft Exchange. The practice of underutilizing a hard disk drive capacity for increased performance is commonly referred to as short-stroking. Enterprise Flash drives can increase Microsoft Exchange application performance and eliminate the need to short-stroke drives, thus keeping storage footprint and power consumption to a minimum and reducing total cost of ownership (TCO).

Exchange environments that are good candidates to benefit from EFD include those where:

- There is a set of users with high I/O requirements and/or exceptional response time requirements
- User mailboxes are relatively small, or they can be augmented by email archiving
- Power, cooling, and floor space are particularly important factors – either because of increasing ongoing costs or limited availability

Introduction

This white paper addresses the advantages of deploying enterprise Flash drives in a Microsoft Exchange environment, including the performance, power savings, and consolidation benefits. Use cases and best practices will be examined, and a performance example comparing Flash drives to HDDs is also discussed.

Audience

This white paper is intended for Microsoft Exchange database administrators and storage architects who want to understand how enterprise Flash drives can be deployed with Microsoft Exchange.

Technology overview

EMC CLARiiON CX4 series

EMC CLARiiON leads the midrange storage market in providing customers with cost-effective solutions that deliver the highest levels of performance, functionality, and reliability. The unique combination of a breakthrough hardware design and advanced software capabilities enables EMC CLARiiON CX4 series systems to meet the growing IT challenges of today's midsize and large enterprises—scaling system capacity and performance, simplifying management in complex environments, and delivering increasing levels of information availability and protection for critical applications and data.



Figure 1. CLARiiON CX4 midtier array family

The powerful CX4 series—CX4-120, CX4-240, CX4-480, and CX4-960—include offerings that are the ideal fit whether your requirements are for departmental applications or support for data-center-class business-critical systems. For companies looking for better ways to address their IT challenges, the CLARiiON CX4 series delivers advanced storage capabilities not found in other midrange storage systems. These capabilities include the following:

- Safeguards to ensure continuous data availability and integrity
- Modular hot-pluggable I/O
- Advanced array-based information replication and virtual LUN technology facilitating data mobility within and between systems
- Large mirrored write cache
- Support for four tiers of storage ranging from ultra-fast solid state Flash disks, to 15k and 10k rpm HDDs, to low power consumption SATA disks
- Easy-to-administer virtual provisioning

Enterprise Flash drives

Enterprise-class EMC Flash drives are constructed with nonvolatile semiconductor NAND Flash memory. They are packaged in a standard 3.5-inch disk drive form factor and employ a 4 Gb/s Fibre Channel interface compatible with existing CLARiiON CX4 drive-array enclosures. (Note that EFDs are not supported on CX3 UltraScale™ or older CLARiiON systems.) The enterprise Flash drives used in the CX4 differ significantly from the solid state technology used in consumer electronics, particularly in their performance and reliability characteristics. To satisfy enterprise-level drive requirements NAND single level cell Flash technology was made more robust with dynamic wear leveling functions, which ensures all cells in the Flash memory are used evenly, to minimize the risk of “wear-out” common to less advanced Flash devices. Additionally, enterprise Flash drives include bad block remapping and multi-bit error correction. Because of these reliability enhancements and the fact that the drive has no moving parts the life expectancy of the Flash drive exceeds that of traditional hard disk drives.

The lack of moving parts also offers additional benefits. As previously discussed, enterprise Flash drives can deliver unprecedented performance, especially well suited for applications that require consistently low read/write response times. Additionally enterprise Flash drives consume considerably less power when compared to traditional rotating media.

The following figures offer an introductory overview to the IOPS and power benefits offered by enterprise Flash drives. Figure 2 shows the relative IOPS rates for various drive technologies that can be expected to deliver suitable response times for Microsoft Exchange. It can be seen that while short-stroking conventional hard disk drives can deliver slightly more IOPS over a fully stroked drive, the performance attained is only a fraction of what is possible with enterprise Flash drive technology.

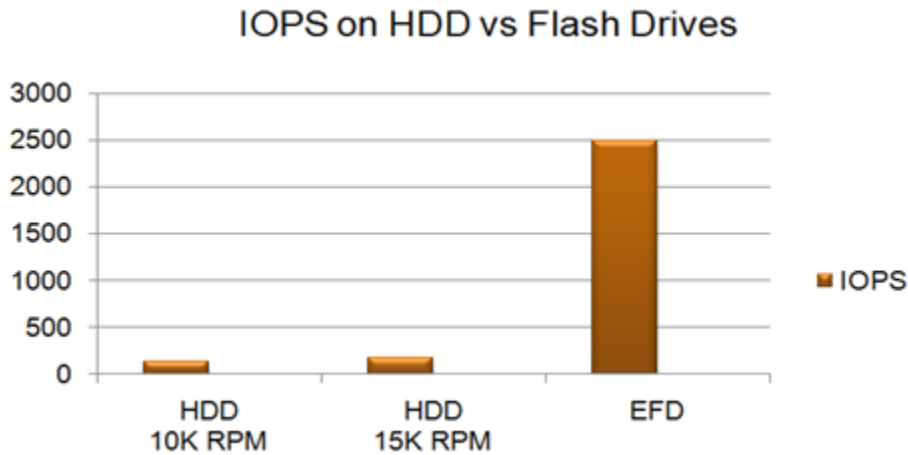


Figure 2. IOPS for different drive types

Figure 3 outlines, in addition to the performance benefits, the power savings that can be achieved when comparing enterprise Flash drives against traditional hard disk drives. The replacement of a larger number of short-stroked HDDs with a reduced number of enterprise Flash drives can deliver a significant reduction in power and cooling requirements – by 90 percent or more. The savings in power will clearly depend on the exact configuration, and to what extent the traditional HDDs have been short-stroked, as well as the total number of HDDs being replaced. For more detailed information, consult with an EMC representative for specific power usage derived from the EMC Power Calculator.




Figure 3. Flash drive vs. HDD power consumption

Microsoft Exchange on enterprise Flash drives

Under normal operating conditions, write requests are written directly to CLARiiON cache resulting in a host I/O response time that is not dependent on physical disk access. For read requests, if the requested data is in cache (either because of a recent read or write request or a sequential prefetch), the request is also immediately serviced without physical disk access. A read serviced from cache without disk access is referred to as a read cache hit. If the requested data is not in cache the data must be retrieved from disk. This is referred to as a read cache miss. A read request that results in a cache miss has an increased I/O response time due to the innate mechanical delays of hard disk drives. Therefore, from a performance perspective enterprise Flash drives are most advantageous for workloads that exhibit low cache hit rates.

Microsoft Exchange environments will generally exhibit low cache hit rates because the application I/O profile is typically small random 4K (Exchange 2003) or 8K (Exchange 2007) blocks, making them prime candidates for enterprise Flash drives.

Placing “Heavy” Exchange users onto Flash drives

A common practice for Exchange administrators is to categorize users and their related mailbox activities into profiles. Such categorization is important when considering user activity and has a direct impact on storage sizing, including IOPS, as previously discussed. Many standard guidelines and sizing calculators break users down into three profiles: Heavy, Medium, or Light. Additionally, some environments may need to include “Very Heavy” or “BlackBerry” profiles to represent even higher user activity.

Depending on the number of users and their respective profiles, it may not be practical or possible to utilize enterprise Flash drives for all Exchange databases in an environment. In such an environment, a subset of Heavy, Very Heavy, or BlackBerry users can be placed onto specific storage groups and databases that reside on Flash drives. Segregating the most active users onto enterprise Flash drives has several benefits. The most obvious benefit is the performance gain and enhanced experience those specific users will see by having their mailboxes on enterprise Flash drives (especially those running Outlook in online mode). Additionally, non-heavy users will benefit when considering the removal of the heavier workload from the non-Flash storage. By removing the heavy user workload from HDDs, the remaining users on HDDs will see better performance and subsequently an enhanced user experience.

The method of placing or moving heavy users specifically onto enterprise Flash drives is a manual process for creating storage tiers in the environment. Another way to create storage tiers is to utilize an archiving product such as EMC SourceOne™. EMC SourceOne is discussed in the section “EMC SourceOne for Information Lifecycle Management” on page 10.

Exchange environments best suited for enterprise Flash drives

Exchange Server administrators, in certain environments, have found it necessary to short-stroke traditional hard disk drives to achieve the required number of IOPS at the desired latencies. The downside of short-stroking hard disk drives is the total number of drives as well as the total number of storage arrays can increase significantly. This increased number of components causes added initial expense and ongoing higher demand and costs for power, cooling, and floor space. The introduction of enterprise Flash drives to such an environment can lead to both increased performance and also a consolidation of storage units, with resulting infrastructure savings.

Consider the read/write workload in the Exchange environment. With Exchange 2003, most customers will see a predominately (70 percent) random read workload, which will benefit greatly from enterprise Flash drives. Exchange 2007 introduces several database engine improvements including a larger database page size, as well as support for larger server memory configurations. Due to these changes the read/write ratio changes to a 50/50 ratio in most Exchange 2007 environments. Exchange 2007 workloads or write heavy workloads in general, however, can also benefit from enterprise Flash drives.

Microsoft Exchange environments, from a storage performance perspective, are generally sized based on a forecasted number of IOPS per user. For example a very heavy Exchange 2007 user is usually rated at around .5 to .7 IOPS. However the IOPS metric can vary dramatically depending on user profiles and mailbox clients (cached mode Outlook, BlackBerry, and others), Microsoft Exchange version (Exchange 2003 vs. Exchange 2007), and other general I/O activity on the server (collaborative applications, replication, anti-virus). In addition, the size of each user mailbox limits the number of users able to generate a workload against a given number of drives.

For example, a fully utilized 300 GB 15k rpm drive will yield roughly 266 GB of usable space and sustain roughly 180 IOPS at a reasonable latency when fully stroked. A 2 GB mailbox quota would allow storage capacity for approximately 100 users (with reasonable free space), while a 200 MB mailbox quota would support 1,000 users for the same amount of space. If the assumed profile is 1 IOPS per user, the 100 2 GB mailbox count (100 IOPS) would perform well, while the 1,000 100 MB mailboxes (1,000 IOPS) would

easily overwhelm a hard disk drive. To extend this example, if the assumed profile was 3 IOPS per user due to heavy BlackBerry usage in the environment, then even the 100 2 GB mailbox count, when taking RAID protection overhead into account, would exceed the processing capacity of the mirrored pair. The calculations for this example follow:

Assumptions are as follows:

- A single fully stroked 15k HDD can sustain 180 random 4K IOPS at a desirable latency (8K IOPS for Exchange 2007)
- Exchange 2003 workload is 70 percent read and 30 percent write
- Exchange 2007 workload is 50 percent read and 50 percent write
- The formula to calculate required number of physical drives is:
 $(\text{IOPS} \times \%R) + \text{WP} (\text{IOPS} \times \%W) / \text{Drive IOPS} = \text{Required Drives}$
 - IOPS = expected I/O operations per second
 - %R = percentage of IOPS that are reads
 - %W = percentage of IOPS that are writes
 - WP = RAID Write Penalty (RAID 1 or RAID 10 = 2, RAID 5 = 4)

Example 1: 2 GB mailbox profile with 3 IOPS per user

200 GB database space / 2 GB usable space per user = 100 users

100 users x 3 IOPS per user = 300 IOPS

Exchange 2003 = $(300 \times .7) + 2 (300 \times .3) / 180 = 3$ required drives = 4 drives for RAID 1/0

Exchange 2007 = $(300 \times .5) + 2 (300 \times .5) / 180 = 3$ required drives = 4 required drives

Example 2: 200 MB mailbox profile with 1 IOPS per user

200 GB database space / .2 GB usable space per user = 1,000 users

1,000 users x 1 IOPS per user = 1,000 IOPS

Exchange 2003 = $(1,000 \times .7) + 2 (1,000 \times .3) / 180 = 8$ required drives

Exchange 2007 = $(1,000 \times .5) + 2 (1,000 \times .5) / 180 = 10$ required drives

Example 3: 200 MB mailbox profile with 3 IOPS per user

200 GB database space / .2 GB usable space per user = 1,000 users

1,000 users x 3 IOPS per user = 3,000 IOPS

Exchange 2003 = $(3,000 \times .7) + 2 (3,000 \times .3) / 180 = 22$ required drives

Exchange 2007 = $(3,000 \times .5) + 2 (3,000 \times .5) / 180 = 26$ required drives

Table 1. RAID 10 drives required to meet Exchange I/O requirements

Exchange version	Mailbox size per user	IOPS per user	Drives required for I/O (300 GB / 15k rpm)
2003	2 GB	3	4
2007	2 GB	3	4
2003	200 MB	1	8
2007	200 MB	1	10
2003	200 MB	3	22
2007	200 MB	3	26

In the third example, with a small mailbox size and very high I/O requirement, it takes the I/O capacity of 26 15k rpm fibre drives to meet the demand of a set of users whose combined Exchange 2007 storage requirements fit on a single drive. As mailbox density in combination with IOPS requirements pushes conventional hard disk drives further past their performance capabilities, enterprise Flash drives become increasingly well suited.

Exchange on Flash recommendations

Exchange log and database file placement

The Exchange log workload is sequential write activity, under normal mailbox processing. Since writes are serviced from CLARiiON cache, the writes to the Exchange log will not generally benefit from Flash drives. In an effort to optimally utilize the enterprise Flash drive capacity with the most appropriate workload, it is reasonable to place the Exchange logs onto HDDs. Placing the logs on hard disk drives allows more capacity to be used by the Exchange database files, the best candidates for use on enterprise Flash drives.

RAID protection

It is recommended to use RAID 5 as the protection mechanism for the database files when deployed on enterprise Flash drives. RAID 5 is the optimal trade-off of protection, cost, and performance for enterprise Flash drives. While RAID 5 volumes have additional overhead of recalculating parity on write operations, this penalty is negligible in comparison to the high performance characteristics of the enterprise Flash drives. On the other hand, with the Exchange log workload on HDDs, it is recommended for the log LUNs to be in a RAID 10 configuration.

Cache setting

The CX4 LUN cache setting for enterprise Flash drives is disabled by default. Because EFDs perform so well with Exchange, the CLARiiON cache may be better reserved for other drive types, especially when other workloads are shared on the storage subsystem. However, for maximum performance you should enable the write cache on EFD-based LUNs.

Exchange virtualization

Enterprise Flash drives offer the same benefits to an Exchange environment deployed on a virtual server platform such as VMware Virtual Infrastructure or Microsoft Hyper-V. Thus far, testing has indicated that the best benefit will be with placement of the most active database LUNs of the Exchange mailbox virtual machine on enterprise Flash drives.

Enterprise Flash drive vs. hard disk drive performance with Exchange

Testing was performed to show the benefit of enterprise Flash drives compared to traditional hard disk drives. Three configurations on a CX4-480 were compared with the requirement of staying within the Microsoft Exchange recommended latency limits for Exchange database and log LUNs:

- Average DB Read Latency < 20 ms
- Average Log Write Latency <10 ms

The workload for the tests was generated by Jetstress version 08.02.0060 x64 bit. Along with this version of Jetstress, the Exchange 2007 SP1 ESE libraries were used.

- **Test 1** = 24 Fibre Channel 15k rpm HDDs (divided into four RAID 1/0 3+3 groups used for Exchange databases) and four fibre drives (RAID 1/0 2+2) for the Exchange logs.
- **Test 2** = 5 EFDs configured in RAID 5 4+1 for the Exchange databases and the same four Fibre Channel drives (RAID 1/0 2+2) for the Exchange logs (EFD LUN write cache OFF).
- **Test 3** = 5 EFDs configured in RAID 5 4+1 for the Exchange databases and four Fibre Channel drives for the Exchange logs (EFD LUN write cache ON).

Figure 4 shows the total IOPS performance of each test achieved along with the average database latencies. All tests were with Exchange logs on Fibre Channel drives and had excellent 1 millisecond latencies.




Figure 4. IOPS and latency comparison between HDDs and enterprise Flash drives

The 24-disk hard drive configuration achieved 4861 IOPS and performed well. However, in addition to achieving a per-drive I/O throughput of almost 10 times that of HDDs, the tests on enterprise Flash drives exhibited read response times to the host that were very low at an average of 7 ms. While the database latency rose on the EFD tests as a result of the write cache being off, results were still well under recommended limits. Well-performing database writes did not affect user response times at all. In all cases write latency on the logs was 1 ms or less.

Note that the HDD configuration used in this comparison used RAID 1/0 (the preferred RAID type for Exchange LUNs for best performance), compared to the RAID 5 enterprise Flash configuration. Had the HDDs been configured as RAID 5, the performance difference (response time as well as IOPS) would have been even more dramatic.

Additional information about the configuration and performance test results can be found within the Jetstress performance test reports in the appendix.

Extending the value of enterprise Flash drives

The deployment of enterprise Flash drives, as seen from the previous example, delivers unprecedented performance for Microsoft Exchange. The following sections discuss how EMC SourceOne and CLARiiON Virtual Provisioning™ can be used to make the most efficient use of available storage capacity on EFD with Microsoft Exchange.

EMC SourceOne for Information Lifecycle Management

Generally speaking, newly added mailbox items are accessed more frequently and generate the highest workload for enterprise Flash storage. Efficient space utilization of the enterprise Flash drives therefore benefits from the movement of less frequently accessed, that is, older, mailbox items on to a more cost-effective tier of storage.

EMC SourceOne is an e-mail archiving solution designed to help companies automate the management of their e-mail environment. SourceOne can assist with storage management, and reduce the cost and risk of legal discovery, as well as assist in creating e-mail retention and disposal policies to meet regulatory requirements. In the context of enterprise Flash drives, SourceOne can provide an Information Lifecycle Management strategy (ILM) to help make the most efficient use of enterprise Flash drive resources.

Using administrator-set policies, SourceOne copies unique messages and attachments from Exchange mailbox servers and retains them within the compressed SourceOne repository. The repository can reside on denser, more cost-effective storage, providing an additional storage tier within the Exchange environment. With messages stored independently, e-mails and associated attachments can be truncated or “shortcut” within the Exchange databases to reference the copy stored within the repository. When this shortcutting ability is applied to older messages and attachments that are not frequently accessed, space is freed within the Exchange database, functionally increasing user mailbox quota. By extension, the capacity available on the enterprise Flash drives will be increased and allow for storage of newer messages and additional user mailboxes.

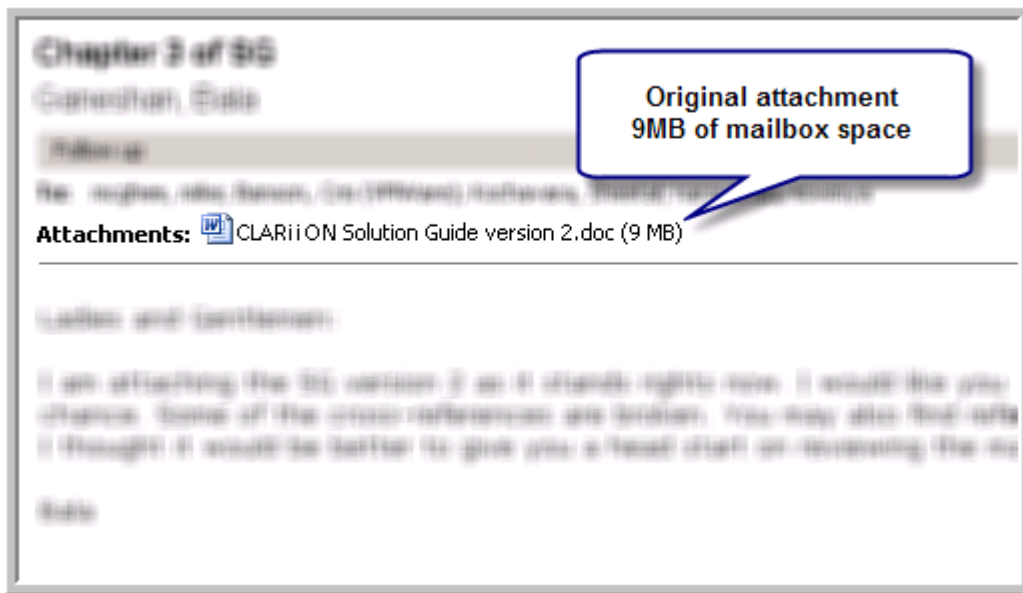


Figure 5. Original attachment size

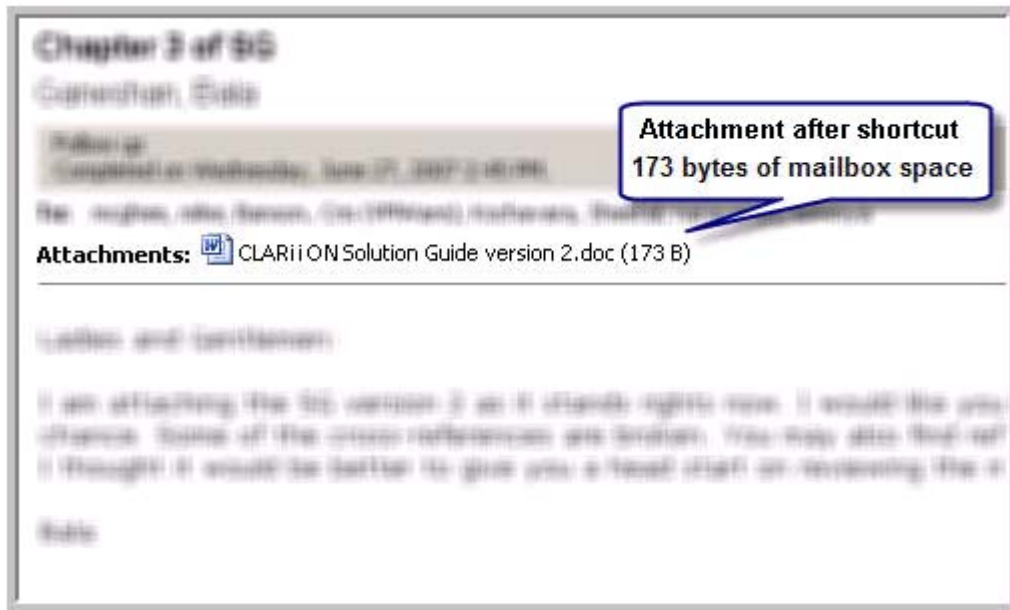


Figure 6. Attachment size after shortcut with SourceOne

It is easy to see from Figure 5 and Figure 6 that a significant amount of space savings can be achieved within the enterprise Flash tier of storage by implementing SourceOne. By combining the use of SourceOne along with enterprise Flash drives, an ILM strategy emerges that allows for the most efficient use of an ultra-performance Tier 0, while allowing for a more cost-effective tier of storage for long-term retention.

CLARiiON Virtual Provisioning

CLARiiON Virtual (thin) Provisioning is a feature that improves storage utilization and simplifies storage administration. CLARiiON thin LUNs can present more storage to Exchange than may be physically available. Space is physically allocated only as it is actually used. For example, you can configure a 2 TB pool of storage on EFD and assign LUNs to this space that present a total of 4 TB to Exchange. You do not need to add more physical capacity to the pool until the original physical 2 TB capacity is almost filled. This maximizes storage efficiency on the Flash drives.

Conclusion

Incorporation of enterprise Flash drives into CLARiiON CX4 provides a new Tier 0 storage layer that is capable of delivering very high I/O performance at a very low latency, which can dramatically improve storage performance for Microsoft Exchange. Traditional magnetic disk drive technology no longer defines the performance boundaries for mission-critical storage environments. The costly approach of spreading workloads over dozens or hundreds of underutilized disk drives is no longer necessary. Depending on the Exchange environment, enterprise Flash drives can lead to increased performance and savings in power, cooling, and data center floor space requirements.

EMC CLARiiON now combines enterprise Flash drive technology with traditional disk drive technology in a single array managed with a single set of software tools, to deliver expanded storage tiering options for ultra-performance and power efficiency. In addition, by utilizing EMC SourceOne, storage management is further optimized and automated to create even more efficiencies in the environment.

References

The following documents can be found on EMC.com and Powerlink®.

White papers

- *An Introduction to EMC CLARiiON Storage Device Technology — Applied Technology*
- *EMC CLARiiON Storage Solutions: Microsoft Exchange 2007 - Best Practices Planning*
- *EMC CLARiiON Virtual Provisioning – Applied Technology*
- *Implementing EMC CLARiiON CX4 with Enterprise Flash Drives for Microsoft SQL Server 2008 Databases — Applied Technology*
- *Microsoft Exchange 2007 and EMC SourceOne Email Management: Unifying the Platforms for Compliance, Storage Management, and Client Productivity — Applied Technology*
- *EMC Symmetrix DMX-4 Enterprise Flash Drives with Microsoft Exchange — Applied Technology*
- *EMC Symmetrix DMX-4 Ultra-Performance Tier 0 Using Flash Drives — Applied Technology*

Appendix: Jetstress results

Microsoft Exchange Server Jetstress: Results for TEST1

28 HDD (24 drives in four RAID 1/0 3+3 for DB and one RAID 1/0 2+2 for all logs)

Performance Test Result Report

Test Summary

Overall Test Result **Pass**

Machine Name WIN-BFZWO4Q9WRX

Test Description

Test Start Time 9/24/2008 1:58:17 AM

Test End Time 9/24/2008 12:10:43 PM

Jetstress Version 08.02.0060.000

Ese Version 08.01.0240.005

Operating System Windows Server (R) 2008 Enterprise Service Pack 1
(6.0.6001.65536)

Performance Log C:\Program Files\Exchange
Jetstress\Performance_2008_9_24_10_4_47.blg
C:\Program Files\Exchange
Jetstress\DBChecksum_2008_9_24_12_10_43.blg

Database Sizing and Throughput

Achieved I/O per Second 4861.93

Capacity Percentage 100%

Throughput Percentage 100%

Initial database size 7376925360128

Final database size 7392811286528

Database files (count) 16

Jetstress System Parameters

Thread count 16 (per-storage group)

Log buffers 9000

Minimum database cache 512.0 MB

Maximum database cache 4096.0 MB

Insert operations 40%

Delete operations 30%

Replace operations 5%

Read operations 25%

Lazy commits 55%

Disk Subsystem Performance

LogicalDisk	Avg. Disk sec/Read	Avg. Disk sec/Write	Disk Reads/sec	Disk Writes/sec	Avg. Disk Bytes/Write
Database (A:)	0.015	0.002	174.110	133.056	(n/a)
Database (B:)	0.014	0.002	172.756	130.344	(n/a)
Database (E:)	0.015	0.002	173.121	131.141	(n/a)
Database (F:)	0.014	0.002	171.194	129.510	(n/a)
Database (G:)	0.015	0.002	174.793	133.140	(n/a)
Database (H:)	0.014	0.002	174.143	133.116	(n/a)
Database (I:)	0.015	0.002	173.093	131.718	(n/a)
Database (J:)	0.015	0.002	173.098	132.561	(n/a)
Database (K:)	0.015	0.002	174.020	131.841	(n/a)
Database (L:)	0.015	0.002	172.267	130.357	(n/a)
Database (M:)	0.014	0.002	172.439	130.423	(n/a)
Database (N:)	0.015	0.002	171.099	127.695	(n/a)
Database (O:)	0.014	0.002	172.438	129.424	(n/a)
Database (P:)	0.015	0.002	172.241	129.928	(n/a)
Database (Q:)	0.015	0.002	173.630	132.428	(n/a)
Database (R:)	0.015	0.002	171.927	128.881	(n/a)
Log (S:)	0.000	0.001	0.000	86.843	4541.388
Log (T:)	0.000	0.001	0.000	86.139	4496.932
Log (U:)	0.000	0.001	0.000	87.057	4491.164
Log (V:)	0.000	0.001	0.000	86.530	4527.673
Log (W:)	0.000	0.001	0.000	88.347	4499.812
Log (X:)	0.000	0.001	0.000	88.149	4525.723
Log (Y:)	0.000	0.001	0.000	87.657	4485.828
Log (Z:)	0.000	0.001	0.000	87.427	4530.073
Log (C:\Log logical drive1)	0.000	0.001	0.000	86.974	4530.429
Log (C:\Log logical drive2)	0.000	0.001	0.000	86.779	4534.687
Log (C:\Log logical drive3)	0.000	0.001	0.000	86.806	4524.633
Log (C:\Log logical drive4)	0.000	0.001	0.000	85.538	4467.911
Log (C:\Log logical drive5)	0.000	0.001	0.000	86.385	4515.169
Log (C:\Log logical drive6)	0.000	0.001	0.000	86.970	4510.505
Log (C:\Log logical drive7)	0.000	0.001	0.000	87.232	4541.267
Log (C:\Log logical drive8)	0.000	0.001	0.000	85.891	4494.489

Microsoft Exchange Server Jetstress: Results for TEST2

5 EFDs (RAID 5 4+1) for the Exchange DBs and 4 fibre drives (RAID 1/0 2+2) for the Exchange logs (Write cache disabled on the EFD LUNs)

Performance Test Result Report

Test Summary

Overall Test Result **Pass**

Machine Name WIN-BFZWO4Q9WRX

Test Description

Test Start Time 3/10/2009 12:07:48 AM

Test End Time 3/10/2009 8:46:45 AM

Jetstress 08.02.0060.000

Version

Ese Version 08.01.0240.005

Operating System Windows Server (R) 2008 Enterprise Service Pack 1
(6.0.6001.65536)

Performance Log C:\Program Files\Exchange
Jetstress\Performance_2009_3_10_6_44_10.blg
C:\Program Files\Exchange
Jetstress\DBChecksum_2009_3_10_8_46_45.blg

Database Sizing and Throughput

Achieved I/O per Second 5287.113

Capacity Percentage 100%

Throughput Percentage 100%

Initial database size 1260199968768

Final database size 1276260745216

Database files (count) 6

Jetstress System Parameters

Thread count 35 (per-storage group)

Log buffers 9000

Minimum database cache 192.0 MB

Maximum database cache 1536.0 MB

Insert operations 40%

Delete operations 30%

Replace operations 5%

Read operations 25%

Lazy commits 55%

Disk Subsystem Performance

LogicalDisk	Avg. Disk sec/Read	Avg. Disk sec/Write	Disk Reads/sec	Disk Writes/sec	Avg. Disk Bytes/Write
Database (E:)	0.010	0.020	479.900	394.012	(n/a)
Database (J:)	0.012	0.024	477.479	406.445	(n/a)
Database (G:)	0.012	0.024	481.558	407.815	(n/a)
Database (H:)	0.012	0.024	479.707	405.454	(n/a)
Database (I:)	0.012	0.024	476.596	397.875	(n/a)
Database (P:)	0.012	0.024	478.800	401.472	(n/a)
Log (K:)	0.000	0.000	0.000	212.318	5961.083
Log (L:)	0.000	0.000	0.000	208.291	6138.170
Log (M:)	0.000	0.000	0.000	208.845	6186.125
Log (N:)	0.000	0.000	0.000	209.624	6144.256
Log (O:)	0.000	0.000	0.000	208.099	6109.364
Log (F:)	0.000	0.000	0.000	207.745	6161.414

Host System Performance

Counter	Average	Minimum	Maximum
% Processor Time	5.941	4.093	8.199
Available MBytes	28418.856	28027.000	28520.000
Free System Page Table Entries	33561460.632	33560545.000	33562069.000
Transition Pages RePurposed/sec	0.000	0.000	0.000
Pool Nonpaged Bytes	102885838.728	101294080.000	103534592.000
Pool Paged Bytes	119439154.343	119209984.000	119869440.000
Database Page Fault Stalls/sec	0.000	0.000	0.000

Microsoft Exchange Server Jetstress: Results for TEST3

5 EFDs (RAID 5 4+1) for the Exchange DBs and 4 fibre drives (RAID 1/0 2+2) for the Exchange logs (Write cache enabled on the LUNs)

Performance Test Result Report

Test Summary

Overall Test Result **Pass**

Machine Name WIN-BFZWO4Q9WRX

Test Description

Test Start Time 3/10/2009 10:50:32 AM

Test End Time 3/10/2009 3:28:49 PM

Jetstress 08.02.0060.000

Version

Ese Version 08.01.0240.005

Operating System Windows Server (R) 2008 Enterprise Service Pack 1 (6.0.6001.65536)

Performance Log C:\Program Files\Exchange Jetstress\Performance_2009_3_10_13_28_3.blg
C:\Program Files\Exchange Jetstress\DBChecksum_2009_3_10_15_28_49.blg

Database Sizing and Throughput

Achieved I/O per Second 8801.851

Capacity Percentage 100%

Throughput Percentage 100%

Initial database size 1260187385856

Final database size 1283502211072

Database files (count) 6

Jetstress System Parameters

Thread count 35 (per-storage group)

Log buffers 9000

Minimum database cache 192.0 MB

Maximum database cache 1536.0 MB

Insert operations 40%

Delete operations 30%

Replace operations 5%

Read operations 25%

Lazy commits 55%

Disk Subsystem Performance

LogicalDisk	Avg. Disk sec/Read	Avg. Disk sec/Write	Disk Reads/sec	Disk Writes/sec	Avg. Disk Bytes/Write
Database (E:)	0.007	0.002	704.490	766.814	(n/a)
Database (J:)	0.007	0.003	703.914	763.730	(n/a)
Database (G:)	0.007	0.003	704.086	763.912	(n/a)
Database (H:)	0.007	0.003	703.828	763.683	(n/a)
Database (I:)	0.007	0.003	700.280	758.085	(n/a)
Database (P:)	0.007	0.003	704.465	764.564	(n/a)
Log (K:)	0.000	0.000	0.000	313.732	6109.394
Log (L:)	0.000	0.000	0.000	313.631	6129.114
Log (M:)	0.000	0.000	0.000	313.642	6086.165
Log (N:)	0.000	0.000	0.000	312.873	6111.874
Log (O:)	0.000	0.000	0.000	310.798	6114.705
Log (F:)	0.000	0.000	0.000	311.130	6159.635

Host System Performance

Counter	Average	Minimum	Maximum
% Processor Time	8.283	4.692	15.102
Available MBytes	28390.528	27967.000	28478.000
Free System Page Table Entries	33561533.720	33561059.000	33562650.000
Transition Pages RePurposed/sec	0.000	0.000	0.000
Pool Nonpaged Bytes	103369529.186	102920192.000	103727104.000
Pool Paged Bytes	120149874.906	119988224.000	120590336.000
Database Page Fault Stalls/sec	0.000	0.000	0.000