



XCubeSAN Series White Paper

SSD Cache 2.0



QSAN Technology, Inc.
www.QSAN.com



Copyright

© Copyright 2018 QSAN Technology, Inc. All rights reserved. No part of this document may be reproduced or transmitted without written permission from QSAN Technology, Inc.

February 2018

This edition applies to QSAN XCubeSAN Series. QSAN believes the information in this publication is accurate as of its publication date. The information is subject to change without notice.

Trademarks

QSAN, the QSAN logo, XCubeSAN, and QSAN.com are trademarks or registered trademarks of QSAN Technology, Inc.

Microsoft, Windows, Windows Server, and Hyper-V are trademarks or registered trademarks of Microsoft Corporation in the United States and/or other countries.

Linux is a trademark of Linus Torvalds in the United States and/or other countries.

UNIX is a registered trademark of The Open Group in the United States and other countries.

Mac and OS X are trademarks of Apple Inc., registered in the U.S. and other countries.

Java and all Java-based trademarks and logos are trademarks or registered trademarks of Oracle and/or its affiliates.

VMware, ESXi, and vSphere are registered trademarks or trademarks of VMware, Inc. in the United States and/or other countries.

Citrix and Xen are registered trademarks or trademarks of Citrix Systems, Inc. in the United States and/or other countries.

Other trademarks and trade names used in this document to refer to either the entities claiming the marks and names or their products are the property of their respective owners.

Notices

This XCubeSAN series white papers applicable to the following XCubeSAN models:

XCubeSAN Storage System 4U & 6U Rack Mount Models

Model Name	Controller Type	Form Factor, Bay Count, and Rack Unit
XS5224D	Dual Controller	LFF 24disk 4U Chassis
XS3224D	Dual Controller	LFF 24disk 4U Chassis
XS3224S	Single Controller	LFF 24disk 4U Chassis
XS1224D	Dual Controller	LFF 24disk 4U Chassis
XS1224S	Single Controller	LFF 24disk 4U Chassis

XCubeSAN Storage System 3U 19 Rack Mount Models

Model Name	Controller Type	Form Factor, Bay Count, and Rack Unit
XS5216D	Dual Controller	LFF 16disk 3U Chassis
XS3216D	Dual Controller	LFF 16disk 3U Chassis
XS3216S	Single Controller	LFF 16disk 3U Chassis
XS1216D	Dual Controller	LFF 16disk 3U Chassis
XS1216S	Single Controller	LFF 16disk 3U Chassis

XCubeSAN Storage System 2U 19 Rack Mount Models

Model Name	Controller Type	Form Factor, Bay Count, and Rack Unit
XS5212D	Dual Controller	LFF 12disk 2U Chassis
XS5212S	Single Controller	LFF 12disk 2U Chassis
XS3212D	Dual Controller	LFF 12disk 2U Chassis
XS3212S	Single Controller	LFF 12disk 2U Chassis
XS1212D	Dual Controller	LFF 12disk 2U Chassis
XS1212S	Single Controller	LFF 12disk 2U Chassis
XS5226D	Dual Controller	SFF 26disk 2U Chassis
XS5226S	Single Controller	SFF 26disk 2U Chassis
XS3226D	Dual Controller	SFF 26disk 2U Chassis
XS3226S	Single Controller	SFF 26disk 2U Chassis
XS1226D	Dual Controller	SFF 26disk 2U Chassis

XS1226S	Single Controller	SFF 26disk 2U Chassis
---------	-------------------	-----------------------

Information contained in this document has been reviewed for accuracy. It could include typographical errors or technical inaccuracies. Changes are made to the document periodically. These changes will be incorporated in new editions of the publication. QSAN may make improvements or changes in the product's features, functionality, and product specifications, which are subject to change without prior notice or obligation. All statements, information, and recommendations in this document do not constitute a warranty of any kind, express or implied.

Any performance data contained herein was determined in a controlled environment. Therefore, the results obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurements may have been estimated through extrapolation. Actual results may vary. Users of this document should refer to the applicable data for their specific environment.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products.

All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

Table of Contents

Notices	i
SSD Cache 2.0	1
Executive Summary.....	1
Audience.....	1
Overview.....	2
Theory of Operation.....	5
SSD Cache Pool Architecture.....	5
RAID Level of SSD Cache Pool.....	7
Read/Write Cache Cases.....	9
Populating the Cache.....	13
SSD Cache Tuning.....	14
Configure SSD Cache.....	16
Enable SSD Cache License.....	16
Create an SSD Cache Pool.....	17
List SSD Cache Pools.....	20
Operations on SSD Cache Pools.....	23
Monitor SSD Cache Statistics.....	25
Test Results.....	28
Test Case 1: SSD Read Cache with 1 / 2 / 4 / 8 SSDs.....	28
Test Case 2: SSD Write Cache with 2 / 4 / 8 SSDs.....	31
Test Case 3: Simulate Database Application.....	33
Test Case 4 Best Practice of SSD Read Cache.....	35
Test Case 5: Best Practice of SSD Write Cache.....	38
Conclusion.....	41
Apply To.....	41
Reference.....	41
Appendix.....	42
Related Documents.....	42
Technical Support.....	42

SSD Cache 2.0

Executive Summary

SSD cache is a large capacity secondary cache that uses enterprise SSD positioned between the RAID controller's primary DRAM memory cache and hard disks (HDD). SSD read/write cache boosts random IOPS of read and write I/O for the system by copying frequently accessed random data to SSD drives, which are faster than HDDs. Therefore increasing the overall random IOPS and lowering the costs by using only few SSDs. this technology QCache 2.0 can improve random read performance by up to 92 times and random write by up to 171 times.

SSD cache is a feature available on XCubeSAN series and requires license to activate. This document discusses the SSD cache technology and describes its features, functions, management and best practice



INFORMATION:

SSD cache 2.0 with read and write cache is available in SANOS firmware 1.1.0.

Audience

SSD cache function needs system memory to store metadata. The usable capacity of SSD cache is in proportion to the size of the controller system memory. The following table is the relationship between the system memory per controller and the maximum SSD cache capacity.

System Memory per Controller	Maximum SSD Cache Capacity per System
4GB	X (No Support)
8GB ~ 15GB	2TB
16GB ~ 31GB	4TB
32GB ~ 63GB	8TB
64GB ~ 127GB	16TB
128GB	32TB



CAUTION:

SSD cache function is not supported when the system memory is under 8GB per controller

The SSD cache is optional; the SSD cache license can be extended to SSD cache 2.0 with read and write cache. The data structure of the SSD cache 2.0 is different from the previous SSD cache. After upgrade to firmware 1.1.0 with SSD cache 2.0, you have to configure SSD cache again if you have already enabled SSD cache in firmware 1.0.x.



CAUTION:

Execute COLD REBOOT after upgrade from firmware 1.0.x to 1.1.0.
Notice that it CAN NOT be downgraded to firmware 1.0.x after upgrade to 1.1.0.

This document is applicable for QSAN customers and partners who are familiar with QSAN products and considering using auto tiering function. Any settings which are configured with basic operations will not be detailed in this document. If there is any question, please refer to the user manuals of products, or contact QSAN support for further assistance.

Overview

Traditionally, data has been stored on traditional rotating memory, HDDs (Hard Disk Drives) and SSDs (Solid-State Drives) are mainly used for mission-critical applications that demand high-speed storage systems, however tend to be costly. In recent years, the capacity of HDDs has increased, but their random I/O (Input/Output) has not kept pace. For

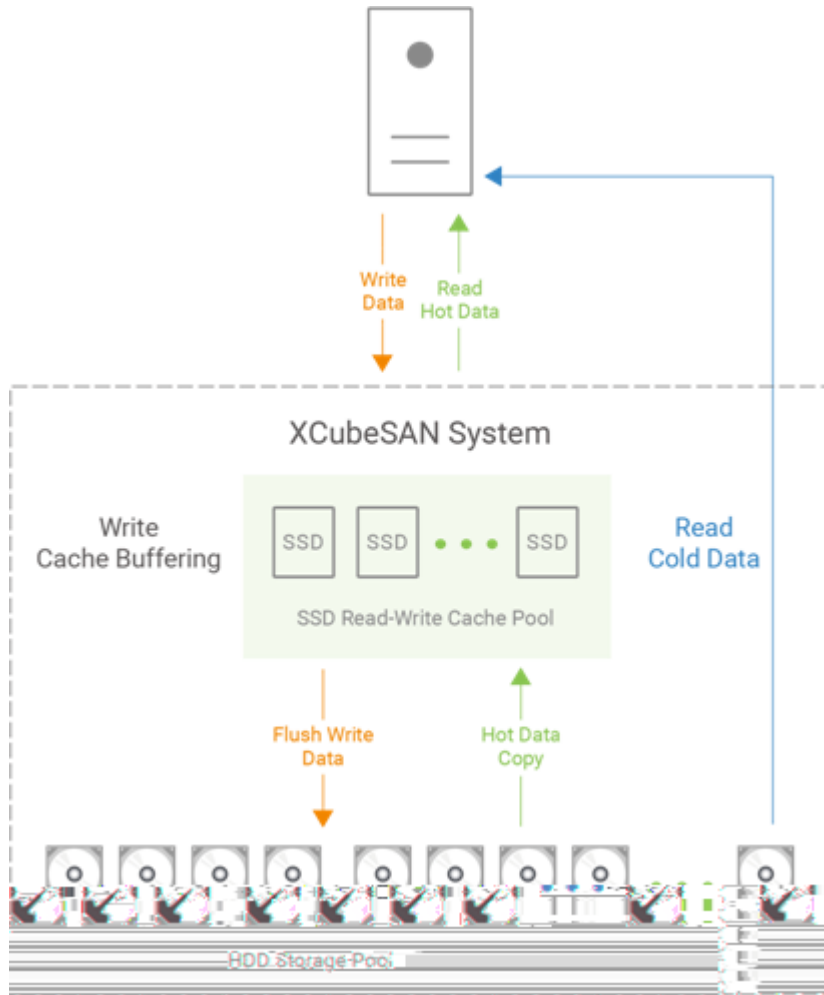
some applications such as webcommerce, clouds, and virtualization that require both high capacity and performance. HDDs, though capacious, simply are not fast enough. SSDs have increased capacity and have declined in cost, making them more attractive for caching in SAN storage networks.

Smart Response Technology (also known as SSD cache technology) leverages the strengths of both HDDs and SSDs, to cost-effectively meet the capacity and performance requirements of enterprise applications. Data is primarily stored on HDDs while SSDs serve as an extended HDD memory cache for many I/O operations. One of the major benefits of using SSD cache is the improved application performance, especially for workloads with frequent I/O activity. The read data of an application that is frequently accessed is copied to the SSD cache; the write data is stored to the SSD cache temporarily and then flush to HDDs in bulk. So the application receives an immediate performance boost. QSAN SSD cache enables applications to deliver consistent performance by absorbing bursts of read/write loads at SSD speeds.

Another important benefit is improved TCO (Total Cost of Ownership) of the system. SSD cache copies the hot or frequently used data to SSDs in chunks. By offloading many if not most of the remaining IOPS to SSD cache, the user can fill the remainder of their storage needs with low cost, high capacity HDDs. This ratio of a small amount of SSD paired with a lot of HDD offers the best performance at the lowest cost with optimal power efficiency.

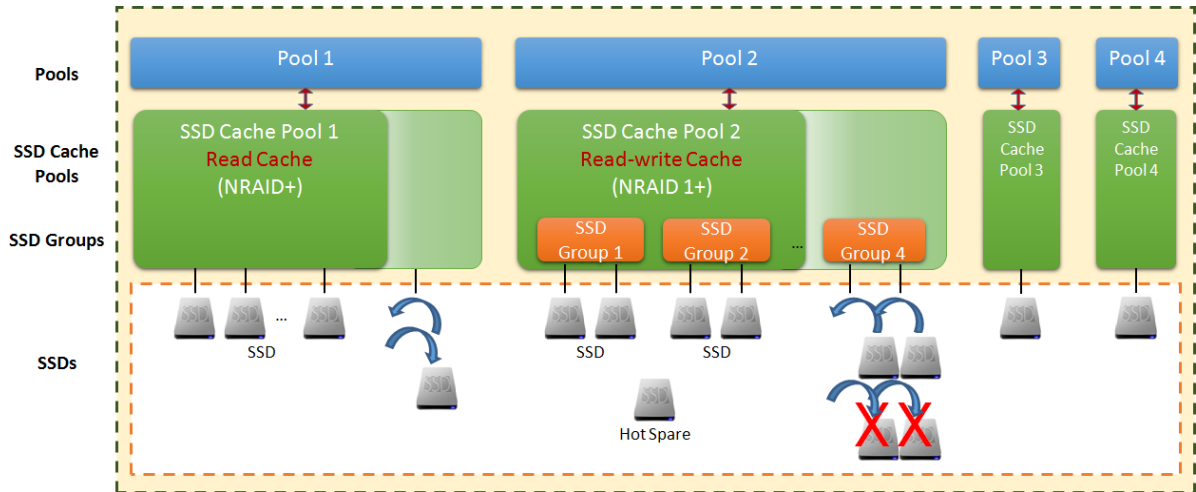
Generally, SSD read cache is particularly effective when:

- ☒ Reads are far more common than writes in the production environment, common in live database or web service applications
- ☒ The inferior speeds of HDD reads cause performance bottlenecks
- ☒ The amount of repeatedly accessed data is small than the capacity of the SSD cache.



SSD read-write cache is particularly effective when:

- ☒ Reads and writes mix in the production environment, common in file service applications
- ☒ The inferior speeds of HDD reads and writes cause performance bottlenecks
- ☒ Same as SSD read cache case; the size of repeatedly accessed data is smaller than the capacity of the SSD cache.
- ☒ Willing to take a little risk to increase performance because its write cache buffering at SSD cache pool. Of course, these write data in SSD cache can be used at the next read.



Take an example of SSD read cache; one or more SSDs with NRAID+ (which is described in the next section) are grouped into the SSD Cache Pool 1 and assigned to the Pool 1 for use. The maximum SSD quantity in a SSD cache pool is 8. When the SSD read cache is performing the capacity of the SSD cache can be increased by adding an SSD or decreased by removing an SSD.

Another example of SSD read/write cache; two SSDs with NRAID 1+ (which is also described in the next section) are grouped into an SSD group. One or more SSD groups combine to the SSD Cache Pool 2 and assigned to the Pool 2 for use. When the SSD read/write cache is performing the capacity of the SSD cache can be increased only by adding an SSD group with two SSDs at a time. The SSD read/write cache pool can set SSDs as dedicated hot spares. The maximum dedicated spare SSD quantity in a SSD cache pool is 4. The following table is the summary of the SSD cache parameters.

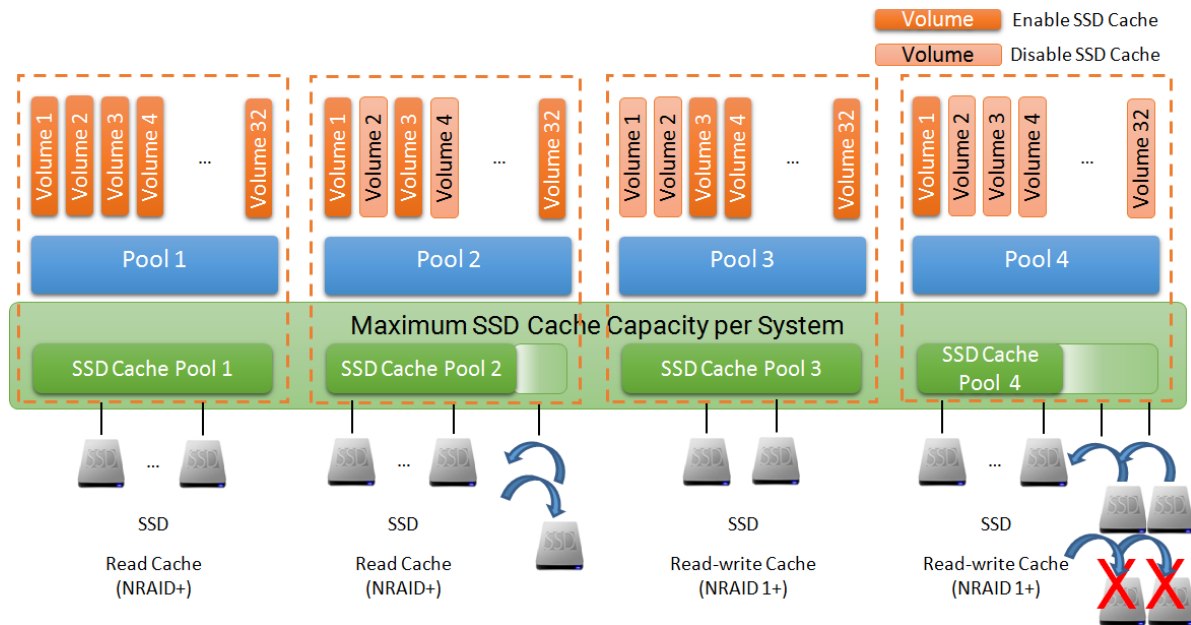
Item	Value
Maximum SSD cache pool quantity per system (either dual controller or single controller)	4
Maximum SSD quantity in a SSD cache pool	8
Maximum capacity of a SSD cache pool	32TB
Maximum quantity of volume shared in a SSD cache pool	32
Maximum dedicated spare SSD quantity in a SSD cache pool	4



TIP:

Note that the capacity allocated to SSD cachepool is not counted in the regular data storage

The volumes in pool can be selected to enable SSD cache function. The following is relationship between SSD cache pool and storage pool.

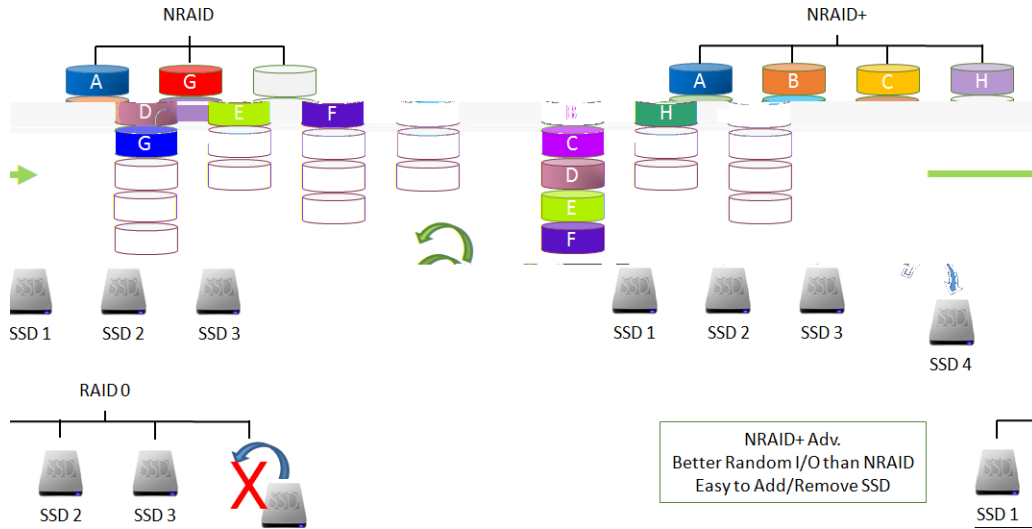


The volumes enabled SSD cache with hot data will consume the capacity of the SSD cache pool. Users have to consider which volumes are enabled according to the SSD cache resources. When the SSD cache performing you can enable or disable the volume. The maximum quantity of volume shared in SSD cache pools 32.

RAID Level of SSD Cache Pool

SSD read cache with NRAID+

Generally, SSD read cache uses NRAID (NRAID) or RAID 0 without data protection to create the SSD cache storage space. Our SSD read cache technology uses NRAID+ which is parallel NRAID without striping.




NRAID is a method of combining the free space on multiple disk drives to create a spanned capacity. A NRAID is generally a spanned volume only, and it contains disk drives of various types and various sizes. RAID 0 consists of stripes without mirroring or parity. Adding a disk drive into RAID 0 needs to stripe the data. For SSD cache, it will have terrible performance when performing SSD cache and migrating RAID 0 at the same time. Let alone remove a disk drive.

Compare to the NRAID or RAID 0, NRAID+ distributes cache data over SSDs. This NRAID+ technology combines with the advantages of NRAID and has better random I/O than NRAID. It also has the advantage of easy to add/remove SSDs from the SSD cache pool to increase/decrease the capacity.

Although the SSD cache is a technique of duplicating data. If the SSD cache pool corrupts, the original data will be safe but stop SSD read cache.

SSD read-write cache with NRAID 1+

SSD read-write cache needs data protection, so it usually uses RAID 1 or RAID 10 level to create the SSD cache storage space. Our SSD read-write cache technology uses NRAID 1+ which is parallel NRAID with mirror.



RAID 10 creates a striped set from a series of mirrored drives. The same, adding pair disk drives into RAID 10 needs to stripe the data. So it is hard to add SSDs into SSD cache pool.

Compare to the RAID 10, NRAID 1+ distributes cache data over mirrored SSDs. This NRAID 1+ technology has the data protection and better random I/O than RAID 10. It also has the advantage of easy to add SSDs into the SSD cache pool to increase the capacity. In this case, we disallow the SSDs to be removed from SSD cache pool because it has the risk of losing write cache data.

If one SSD failure occurs in the SSD cache pool, the status of the SSD cache pool will be changed to degrade. It will stop write cache and start flushing write cache data into HDDs. At this time, read cache is still working but no more new read cache will be added into SSD cache pool. The SSD read-write cache pool can set SSDs as dedicated hot spares to prevent SSD cache pool failure. After inserting an SSD as a hot spare, the SSD read cache pool will be rebuilt until it is completed. And then it reverts to SSD read cache service.

Read/Write Cache Cases

The following describes the cases of read/write cache. In SSD read cache, the processes of read/write data will be disassembled to the following cases.

- ☒ Read Data with Cache Miss

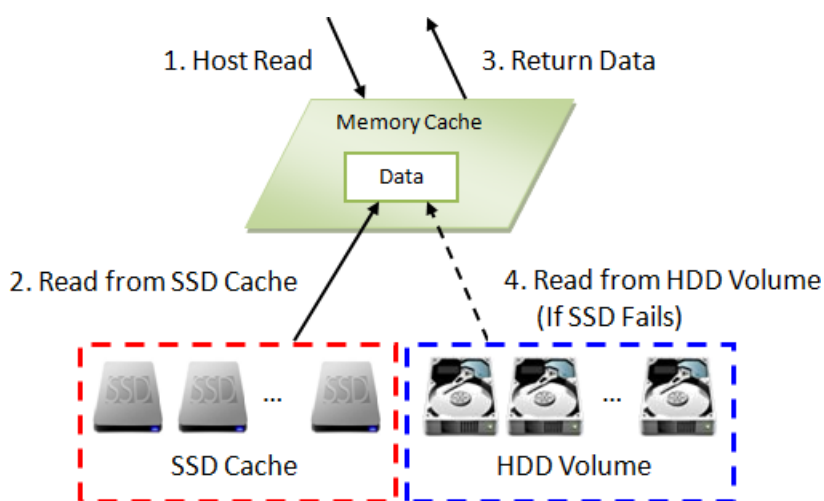


INFORMATION:

The actions that read data from the HDD and then write to the SSD are called populating the cache. For more information, please refer to [Populating the Cache](#) in the next section.

Read Data with Cache Hit

The following figure shows how the controller handles a host read request when data is in the SSD cache.




These steps are:

1. A host requests a read data. The system finds that the data is in SSD cache, so it is called cache hit.
2. Data is read from the SSD cache.
3. The requested data is returned to the host.
4. If there is an SSD cache error, data is read from the HDD volume.

Write Data in SSD Read Cache

The following figure shows how the controller handles a host write request in SSD read cache. The write data can also be frequently accessed data and populated into SSD cache for next read.




These steps are:

1. A host requests to write data.
2. Data is written to the HDD volume.
3. The status is returned to the host.
4. The SSD cache is populated if the write threshold is reached.

Write Data in SSD Read-write Cache

The following figure shows how the controller handles a host write request in SSD read-write cache. The write data stays in the SSD cache for a while, and will be flushed to the HDD volume at the appropriate time.



These steps are:

1. A host requests to write data.
2. Data is written to SSD cache
3. The status is returned to the host.
4. Data will be flushed to the HDD volume at the appropriate time

Flush Write Data to HDD Volume

In SSD readwrite cache, the write data will be flushed to the HDD volume in the following situations.

- ☒ Flush data to the HDD volumes before the SSD cache pool is deleted.
- ☒ Flush data to the HDD volumes before the related volume is disabled.
- ☒ Flush data when the system is idle.



CAUTION:

It will take time to flush data from SSD to HDD volumes. Please do not remove SSDs before completing flushing; otherwise, it may take risk of losing data.

In addition, to avoid taking a long time to complete flushing, please make sure the front I/O affords not heavy when deleting SSD cache pool.



TIP:

It will flush data to the HDD volumes when system goes shutdown or reboots.

Populating the Cache

The actions that read data from the HDD and then write to the SSD are called populating the cache. Typically, this is a background operation that immediately follows a host read write operation. As the goal of the cache is to store frequently accessed data, not every I/O operation should trigger a cache population, but only ones that pass a certain threshold, implemented as a counter. There are both a population-read threshold, and a population-write threshold.

Populate-on-read Threshold

When the same data block to be read over the threshold, it is called hot data and populated to the SSD cache. The threshold must be greater than or equal to 1 in SSD read cache or SSD read/write cache. The value is forbidden to set 0 because of non-artificial read cache. The maximum of the threshold is 4. If larger than 4, the frequently accessed data hard into SSD cache, so there is no obvious effect.

Populate-on-write Threshold

When the same data block to be written over the threshold, it is called hot data and populated to the SSD cache. The threshold must be greater than or equal to 1. If it is set to 0, no action is performed for a write cache. The value must be greater than or equal to 1 in SSD read/write cache. The value is forbidden to set 0 because of no action in write cache. The same as above, the maximum of the threshold is 4. If larger than 4, there is no effect.

Operation Process

Each cache block on a HDD volume has a read and write count associated. When a host requests to read data located in that cache block, the read count is increased. If the data is not found in the cache and the read count is greater than or equal to the populate threshold, then a cache populate operation is performed concurrently with the host read operation. If a cache hit occurs, the data is immediately returned from the SSD cache and a populate operation is not performed. If the read count is smaller than the threshold, a populate operation is not performed.

Write cases are the same scenarios as read.

SSD Cache Tuning

The SSD cache can be tuned to maximize its efficiency based on application usage. Cache block size, populate-on-read threshold and populate-on-write threshold are the main parameters.

Cache Block Size

A large cache block suits applications where frequently accessed data is close to each other, known as a high locality of reference. A large cache block will also fill up the SSD cache quickly - this is known as the warm-up time. After the cache is warmed up, the performance would be quite good for applications with high locality of reference. Such as

the file system or web service usage, the frequently accessed data are based on some concentrated files which are usually in large block size. However large cache blocks will also generate large I/O overhead, increasing response time, especially for cache misses.

A smaller cache block size suits applications with data that is less localized, meaning the data is accessed more randomly such as database usage. The SSD cache will fill up slower, but with more cache blocks, there is greater chance of a cache hit, especially for data with less locality of reference. With a smaller cache block size, cache usage is usually less than with a larger cache block size, but overhead is less, so the penalty for misses is less severe.

Population Threshold

The population threshold is the quantity of accesses at which point that cache block is copied to the SSD Cache. A higher number ensures that the cache only stores frequently accessed data so there will not be much cache turnover however it also means the cache will take longer to warm up and fully effective. A lower number means the cache is warmed up quickly, but may cause excessive cache populations. A populate on read threshold of 2 is sufficient for many applications. Populate on write is useful when data that is written to is often read soon after. This is often the case in file systems. Other applications, such as database software, do not have this tendency so populate on write may sometimes even be disabled.

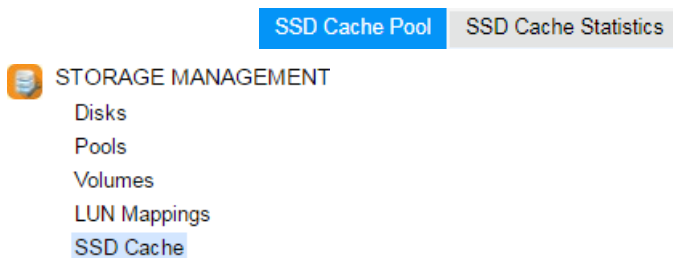
I/O Type	Block Size (Sectors)	Populate-on-Read Threshold	Populate-on-Write Threshold
Database	1MB (2,048)	2	0
File System	2MB (4,096)	2	2
Web Service	4MB (8,192)	2	0
Customization	1MB/2MB/4MB	"&VcY")	"%VcY")

I/O Type	Block Size (Sectors)	Populate-on-Read Threshold	Populate-on-Write Threshold
Database	1MB (2,048)	2	1
File System	2MB (4,096)	2	1
Web Service	4MB (8,192)	2	1
Customization	1MB/2MB/4MB	"&VcY" ")	"1 VcY" ")

As you can see there are tradeoffs for increasing or decreasing each parameter. Understanding the data locality of the application is essential and it can be useful to do some field testing to see what works best.

Configure SSD Cache

The **SSD Cache** provides **SSD Cache Pools** and **SSD Cache Statistics** function tabs to configure and monitor SSD cache. This section will describe the operations of configuring SSD cache.



Enable SSD Cache License

The SSD cache function is optional. Before using it, you have to enable SSD cache license. Select the **Update** function tab in the **Maintenance** function submenu, download **Request License** file and send to your local sales to obtain a License Key. After getting the license key, click the **Choose File** button to select it, and then click the **Apply** button to enable it. When the license is enabled, please reboot the system. Each license key is unique and dedicated to a specific system. If you have already enabled it, this option will be invisible.

SSD Cache License

Download Request License file and send to your local sales to get a License Key.

Select the license file to enable SSD Cache: Choose File No file chosen

Apply
Request License

Create an SSD Cache Pool

Here is an example of creating an SSD cache pool with 2 SSDs configured in SSD read cache

1. Select the **SSD Cache** functions submenu click the **Create SSD Cache Pool** button. It will scan available SSDs first.

Create SSD Cache Pool

General

Enable SSD Cache

Summary

SSD Cache Pool

Please enter a pool name and select the cache type.

SSD Cache Pool Name : ⓘ

Cache Type : Read Cache (NRAID+) Read-write Cache (NRAID 1+)

I/O Type :

Cache Block Size :

Populate-on-read Threshold :

Populate-on-write Threshold :

Select SSDs

Enclosure ID:

<input type="checkbox"/>	Enclosure ID	Slot	Health	Capacity	Disk Type	Manufacturer	Model
<input checked="" type="checkbox"/>	0	1	Good	372.36 GB	SAS SSD 12.0Gb/s	SEAGATE	ST400FM0053
<input checked="" type="checkbox"/>	0	2	Good	372.36 GB	SAS SSD 12.0Gb/s	SEAGATE	ST400FM0053
<input type="checkbox"/>	0	3	Good	372.36 GB	SAS SSD 12.0Gb/s	SEAGATE	ST400FM0053

Next
Cancel

2. Enter a **SSD Cache Pool Name** for the pool. Maximum length of the pool name is 16 characters. Valid characters are [A-Z | a-z | 0-9 | _ | < >].

3. Select the **Cache Type** as Read Cache (NRAID+).
4. Select an **I/O Type** from the dropdown list according to your application
5. Please select SSDs for SSD cache pool. Maximum quantity of disks is 8. Select an **Enclosure** from the dropdown list to select SSDs from expansion enclosures.
6. Click the **Next** button to continue

General
Enable SSD Cache

Summary
Select a Pool

Please select a pool to enable SSD cache. The pool can not be change after selection.

Pool Name:

Select Volumes

Within the pool, check volumes to enable SSD cache, uncheck volumes to disable.

	Volume Name	Status	Health	Capacity	
<input checked="" type="checkbox"/>	Vol-1	Online	Optimal	100.00 GB	
<input checked="" type="checkbox"/>	Vol-2				

Online
Optimal
100.00 GB

Vol-2

7. Select a **Pool Name** from the dropdown list which lists available pools to be assigned to.
8. Within the pool, check **Volumes** to enable SSD cache, uncheck **Volumes** to disable.
9. Click the **Next** button to continue

	SSD Cache Pool Name	Status	Health	Total	Cache Type	RAID	Pool Name	I/O Type
▼	RCPool-1	Online	Good	744.72 GB	Read Cache	NRAID+	Pool-1	Database
▼	RWCPool-2	Online	Good	372.36 GB	Read Write Cache	NRAID 1+	Pool-2	File System

SSDs

No.	Enclosure ID	Slot	Status	Health	Capacity	Disk Type	Manufacturer	Model
1	0	1	Online	Good	372.36 GB	SAS SSD 12.0Gb/s	SEAGATE	ST400FM0053
2	0	2	Online	Good	372.36 GB	SAS SSD 12.0Gb/s	SEAGATE	ST400FM0053

Enabled Volumes

Volume Name	Status	Health	Capacity	Total Cache Used	Read Cache Used	Write Cache Used
Vol-2	Online	Optimal	100.00 GB	0 MB	0 MB	0 MB
Vol-1	Online	Optimal	100.00 GB	0 MB	0 MB	0 MB

This table shows the column descriptions.

Column Name	Description
SSD Cache Pool Name	The SSD cache pool name.
Status	The status of the pool ☒ Online : The SSD cache pools online. ☒ Offline : The SSD cache pools offline. ☒ Rebuilding : The SSD cache pools being rebuilt. ☒ Flushing : The SSD cache pools being flushed.
Health	The health of the pool ☒ Good : The SSD cache pools good. ☒ Failed : The SSD cache pools failed. ☒ Degraded : The SSD cache pools not healthy and not complete. The reason could be missing or failed SSDs.
Total	Total capacity of the SSD cache pool.
Cache Type	The SSD cache type: ☒ Disabled. ☒ Enabled.
RAID	The RAID level of the SSD cache pool:

Add SSDs

SSD Cache Pool

SSD Cache Pool Name : RWCPool-2
 Cache Type : Read-write Cache (NRAID 1+)
 Capacity : 372.36 GB

Select SSDs

Please select two SSDs at a time to add into the SSD cache pool.

Enclosure ID: 0 (Head Unit: XS5216) ▼

<input checked="" type="checkbox"/>	Enclosure ID	Slot	Health	Capacity	Disk Type	Manufacturer	Model
<input checked="" type="checkbox"/>	0	7	Good	744.96 GB	SAS SSD 12.0Gb/s	MICRON	S630DC-800
<input checked="" type="checkbox"/>	0	8	Good	744.96 GB	SAS SSD 12.0Gb/s	MICRON	S630DC-800

Delete an SSD Cache Pool

Click **z -> Delete** to delete the SSD cache pool. If deleting an SSD read-write cache pool, it will take times to flush data from SSDs to HDDs. The SSD cache pool will be deleted after flushed, otherwise the data may be corrupt.

More Information of the SSD Cache Pool

Click **z -> More Information** to show the detail information of the SSD cache pool

Monitor SSD Cache Statistics


This section will describe the monitoring of SSD cache statistics.

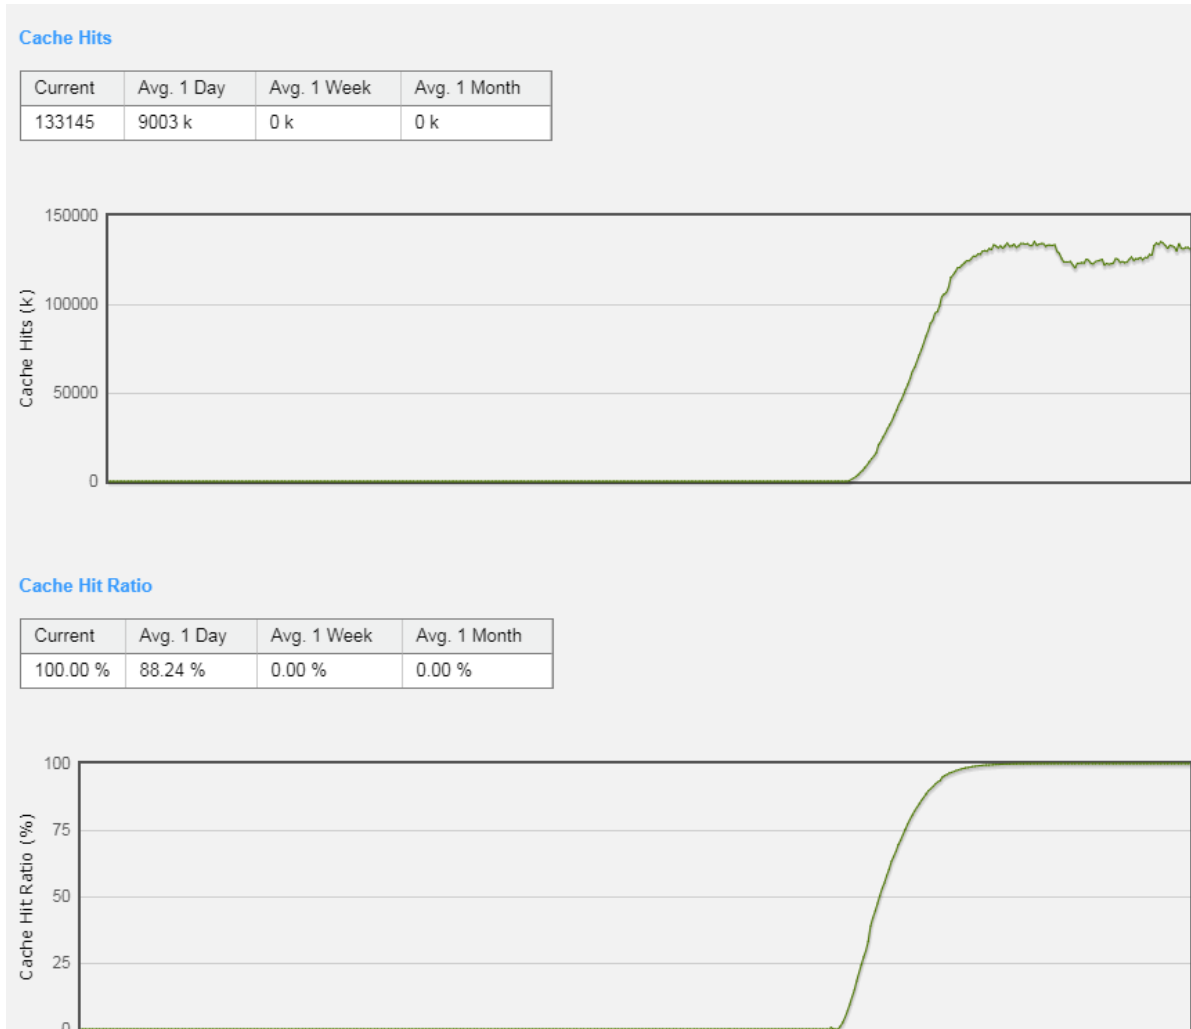
STORAGE MANAGEMENT

- Disks
- Pools
- Volumes
- LUN Mappings
- SSD Cache**

Select the **SSD Cache Statistics** function tab in the **SSD Cache** function sub-menu to monitor the SSD cache statistics. Here is an example of monitoring the volume of the SSD cache pool.

1. Select an **SSD Cache Pool Name** and select an **Enabled Volumes** which you want to monitor.





- The SSD cache statistics include **Cache Used**, **Cache Hits**, and **Cache Hit Ratio**. Each item has a current value, average 1 day value, average 1 week value, and average 1 month value. The diagram displays the current value with 30 seconds refresh and total 4 hours.

Column Name	Description
Cache Used	How much capacity of the SSD cache pool is used by volume? Unit is MB, GB, or TB.
Cache Hits	How many cache hits occurred by the volume? Unit is the number of times
Cache Hit Ratio	How many cache hit ratio occurred by the volume? Unit is percentage.

Test Results

Test Case 1: SSD Read Cache with 1 / 2 / 4 / 8 SSDs

This test verifies the dramatic performance gains offered by SSD cache. We test the SSD read cache with 1/2/4/8 SSDs. According to the design structure of RAID level in SSD read cache pool, in theory, the more SSDs, the better SSD read cache. We also set the populate-on-read threshold to 1 which means the data hits once and is populated to the SSD.

Test Equipments and Configurations

☒ Server

- Model: HP Z840 (CPU: 2 x Xeon E52620v3 2.4Hz / RAM: 16GB)
FC HBA: QLogic QLE2699R
OS: Windows Server 2012 R2
- Model: Dell E25S (CPU: 2 x Xeon E52620v3 2.4Hz / RAM: 32GB)
FC HBA: QLogic QLE2699R
OS: Windows Server 2012 R2

☒ Storage

- Model: XCubeSAN XS5224
Memory: 16GB (2 x 8GB in bank 1 & 3) per controller
Firmware 1.1.2
HDD: 16 x Seagate Constellation ES5T500NM0001 500GB SAS 6Gb/s
SSD: 1 / 2 / 4 / 8 x HGST Ultrastar SSD800MH.B, HUSMH180SS200, 100GB, SAS 12Gb/s
- HDD Pool: 1 x RAID 5 Pool with 16 x SAS HDDs in Controller 1

- HDD Volume: 2 x 45GB in Pool
- FC Session: 2 per Volume

☒ SSD Cache

- **SSD Cache Pool: Read Cache (NRAID+) with 1 / 2 / 4 / 8 x SAS SSDs**
- I/O Type: Customization
Cache Block Size 4MB
Populate-on-read Threshold: 1
Populate-on-write Threshold: 0

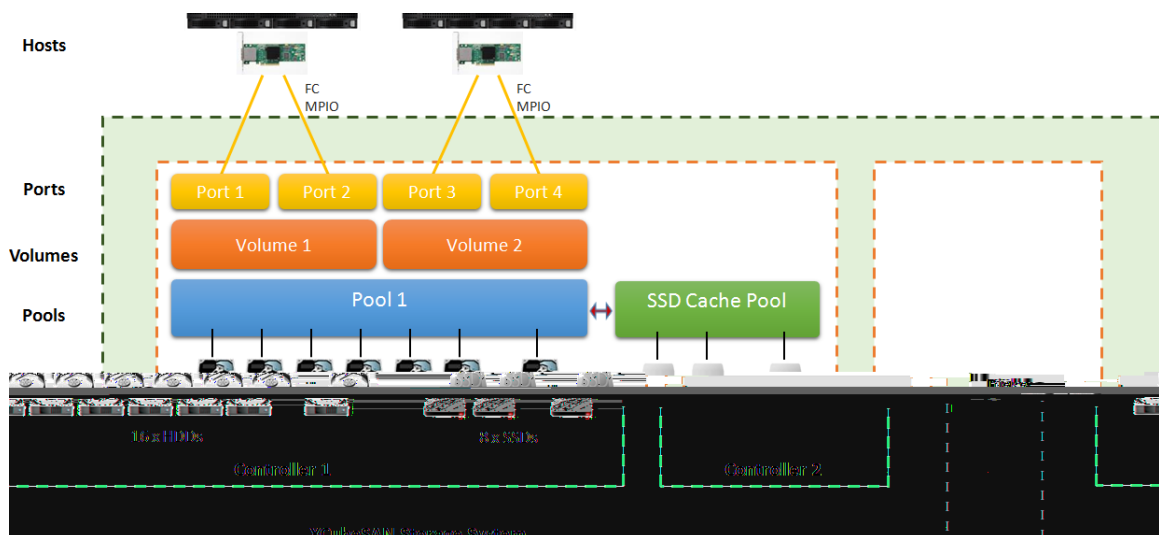
☒ I/O Pattern

- Tool: IOMeter V1.1.0
- Workers: 1
- Outstanding (Queue Depth): 128
- **Access Specifications: 4KB, 100% Read, 100% Random**

☒ Test Scenario

- HP server runs 1 x Volume (45GB) of RAID 5 Pool in Controller 1 via 2x FC ports with MPIO
- Dell server runs 1 x Volume (45GB) of RAID 5 Pool in Controller 2 via 2x FC ports with MPIO

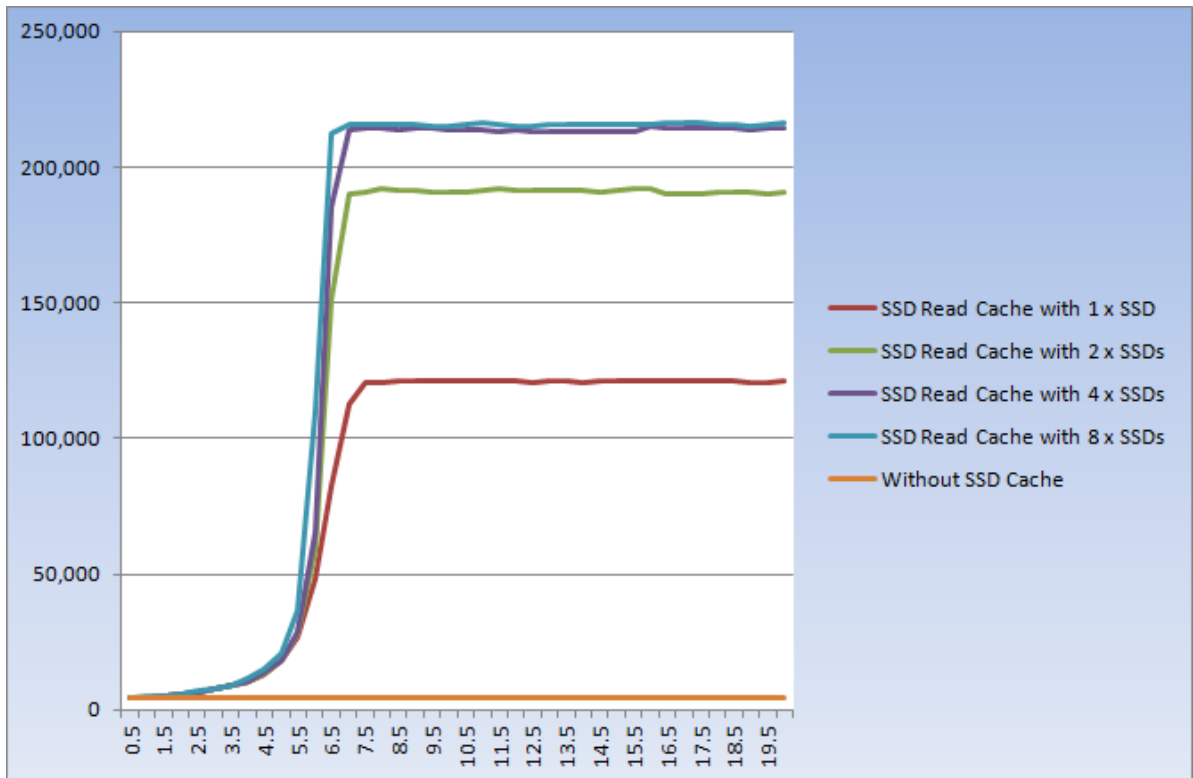
Test Diagram



Test Result

Volume Capacity	Disable SSD Cache (IOPS)	Enable SSD Cache			After Populating (IOPS)	Improved
		No. of SSDs	SSD Capacity	Warm up Time		
90GB	4,512	1	100GB	7.5 min.	121,326	2,589%
90GB	4,512	2	200GB	7 min.	192,374	4,164%
90GB	4,512	4	400GB	7 min.	214,986	4,665%
90GB	4,512	8	800GB	7 min.	216,434	4,697%

Without SSD cache, the average IOPS is 4,512. Enable SSD cache with 8 SSDs, IOPS increases to 216,434. It improves $(216,434 - 4,512) / 4,512 = 46.968 = 4,697\%$. The warm up time is about 7 minutes.



According to the chart for the relationship between IOPS and warm up time, the orange line is the IOPS without SSD cache and the others are enabled SSD cache. At the period of 7 minutes, the data is hit and read from the SSD. If it is not hit, the data is read from HDD and populated to the SSD. After 7 minutes, the IOPS climbs to the top and stays stable. This means that the data are all in the SSD.

Summary

- ☒ SSD Cache 2.0 can improve random read performance by 4.7 times.
- ☒ The more SSDs used, the better SSD read cache is.



TIP:

You can also watch the **SSD Cache -> SSD Cache Statistics** page in Web UI to monitor the SSD cache changes, **Performance -> Disks** by enabling the HDD and SSD slots to monitor the HDD and SSD/O changes.

Test Case 2: SSD Write Cache with 2 / 4 / 8 SSDs

In this test, we tested the SSD read/write cache with 2 / 4 / 8 SSDs. As before, we set the populate-on-write threshold to 1, which means the data hits once and is populated to SSD.

Test Equipments and Configurations

☒ Server

- Model: HP Z840 (CPU: 2 x Xeon E5-2620v3 2.4Hz / RAM: 32GB)
FC HBA: QLogic QLE2690R
OS: Windows Server 2012 R2
- Model: Dell E25S (CPU: 2 x Xeon E5-2620v3 2.4Hz / RAM: 32GB)
FC HBA: QLogic QLE2690R
OS: Windows Server 2012 R2

☒ Storage

- Model: XCubeSAN XS5224D
Memory: 16GB (2 x 8GB in bank 1 & 3) per controller
Firmware 1.1.2
HDD: 16 x Seagate Constellation ES ST500NM0001 500GB SAS 6Gb/s

SSD: 2 / 4 / 8 HGST Ultrastar SSD800MH.B, HUSMH8010BSS200, 100GB, SAS 12Gb/s

- HDD Pool: 1 x RAID 5 Pool with 16 x SAS HDDs in Controller 1
- HDD Volume: 2 x 3GB in Pool
- FC Session: 2 per Volume

☒ SSD Cache

- **SSD Cache Pool: Read-write Cache (NRAID 1+) with 2 / 4 / 8 x SAS SSDs**
- I/O Type: Customization
Cache Block Size 4MB
Populate on-read Threshold: 1
Populate-on-write Threshold: 1

☒ I/O Pattern

- Tool: IOMeter V1.1.0
- Workers: 1
- Outstanding (Queue Depth): 128
- **Access Specifications: 4KB, 100% Write, 100% Random**


☒ Test Scenario

- HP server runs 1 x Volume (3GB) of RAID 5 Pool in Controller 1 via 2x FC ports with MPIO
- Dell server runs 1 x Volume (3GB) of RAID 5 Pool in Controller 1 via 2x FC ports with MPIO

Test Result

Volume Capacity	Disable SSD Cache (IOPS)	Enable SSD Cache				Improved
		No. of SSDs	SSD Capacity	Warm up Time	After Populating (IOPS)	
60GB	1,660	2	100GB	13 min.	41,358	2,391%
60GB	1,660	4	200GB	9 min.	84,405	4,985%
60GB	1,660	8	400GB	6.5 min.	143,898	8,569%

Without SSD cache, the average IOPS is 1,660. Enable SSD cache with 8 SSDs, IOPS increases to 143,898. It improves $(143,898 - 1,660) / 1,660 = 85.685 = 8,56\%$. The warm up time is about 6.5 minutes.



Summary

- ☒ SSD Cache 2.0 can improve random write performance by 8.6 times.
- ☒ The more SSDs used, the better SSD write cache is.

Test Case 3: Simulate Database Application

This test simulates the database application. We use the Database I/O type in the configuration of SSD cache pool, and use the database access pattern (8KB, 67% read, 100% random) to test.

Test Equipments and Configurations

- ☒ Server
 - Model: HP Z840
 - CPU: 2 x Xeon E5-2620v3 2.4Hz / RAM: 32GB
 - FC HBA: QLogic QLE2690R
 - OS: Windows Server 2012 R2

- Model:Dell E25S (CPU2 xXeon E52620v3 2.4Hz / RAM:32GB)
FC HBA: QLogic QLE269SR
OS: Windows Server 2012 R2

☒ Storage

- Model: XCubeSAN XS5224D
Memory: 16GB (2 x 8GB in bank 1 & 3) per controller
Firmware 1.1.2
HDD: 16 xSeagateConstellation ES5T500NM0001500GBSAS 6Gb/s
SSD: 2 xHGST Ultrastar SSD800MH.B, HUSMH8010BSS200, 100GB, SAS 12Gb/s
- HDD Pool:1 xRAID 5Pool with 16 x SAS HDDs in Controller 1
- **HDD Volume: 2 x 30GB / 2 x 45GB / 2 x 60GB / 2 x 90GB in Pool**
- FC Session: 2 per Volume

☒ SSD Cache

- **SSD Cache Pool: Read-write Cache (NRAID 1+) with 2 x SAS SSDs**
- **I/O Type: Database**
Cache Block Size:1MB
Populateon-read Threshold:2
Populateon-writeThreshold:1

☒ I/O Pattern


- Tool: IOmeter V1.1.0
- Workers: 1
- Outstanding (Queue Depth): 128
- **Access Specifications: 8KB, 67% Read, 100% Random** (Database Access Pattern)

☒ TestScenario

- HP server run 1 x Volume (30GB / 45GB / 60GB / 90GB) RAID 5 Pool in Controller 1 via 2x FC ports with MPIO
- Dell server run 1 x Volume (30GB / 45GB / 60GB / 90GB) RAID 5 Pool in Controller 1 via 2x FC ports with MPIO

Test Result

Volume Capacity	Disable SSD Cache (IOPS)	Enable SSD Cache			After Populating (IOPS)	Improved
		No. of SSDs	SSD Capacity	Warm up Time		
60GB	2,799	2	100GB	19 min.	64,809	2,215%
90GB	2,799	2	100GB	23 min.	63,660	2,174%
120GB	2,799	2	100GB	N/A min.	3,999	43%



The result is very good when the amount of the hot data is less than the capacity of SSD cache pool.

Summary

☒ User has to accurately estimate the amount of data used to achieve the best result

Test Case 4: Best Practice of SSD Read Cache

The cases above are tested on single controller. The following test provides the best practice of SSD read cache on dual controller. We assume that the performance can be twice that of the test on single controller

Test Equipments and Configurations

CE Server

- Model: HP Z840 (CPU: 2 xXeon E52620v3 2.4Hz / RAM:32GB)
FC HBA: QLogic QLE2699R
OS: Windows Server 2012 R2
- Model:Dell E25S (CPU: 2 xXeon E52620v3 2.4Hz / RAM:32GB)
FC HBA: QLogic QLE2699R
OS: Windows Server 2012 R2

CE Storage

- Model: XCubeSANXS5224D
Memory: 16GB (2 x 8GB in bank 1 & 3) per controller
Firmware 1.1.2
HDD: 16 xSeagateConstellation ES5T500NM0001500GBSAS 6Gb/s
SSD: 8 xHGST Ultrastar SSD800MH.B, HUSMH8010BSS200, 100GB, SAS 12Gb/s
- HDD Pool:1 xRAID 5Pool 1 with 8 x SAS HDDs in Controller 1
1 xRAID 5Pool 2 with 8 x SAS HDDs in Controller 2
- HDD Volume:2 x 45GB in Pool
2 x 45GB in Pool
- FC Session: 2 per Volume

CE SSD Cache

- **SSD Cache Pool: Read Cache (NRAID+) with 8 x SAS SSDs**
- I/O Type: Customization
Cache Block Size4MB
Populate-on-read Threshold: 1
Populateon-writeThreshold:0

CE I/O Pattern

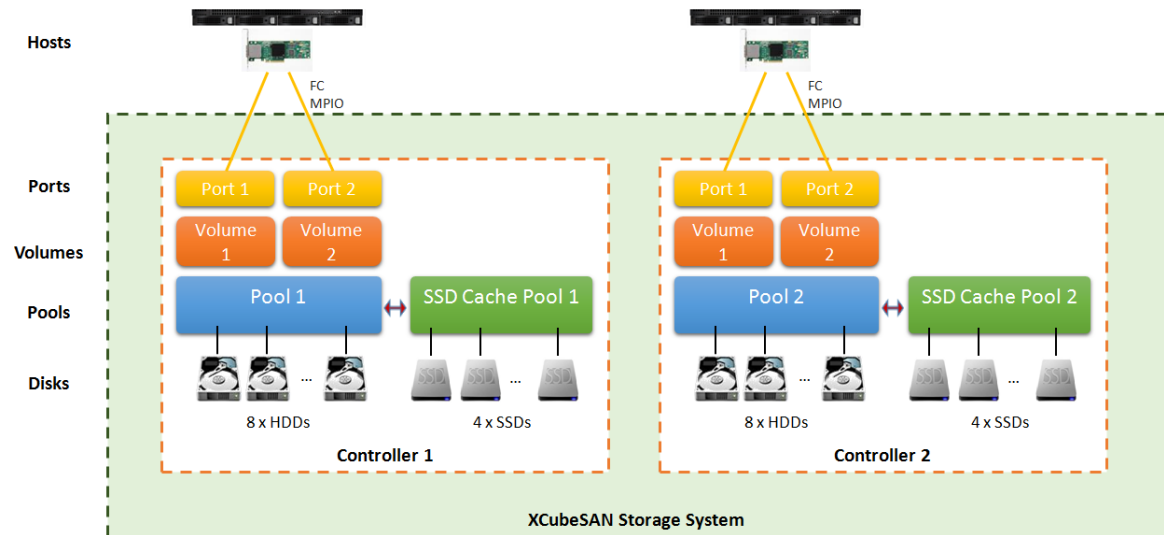
- Tool: IOmeter V1.1.0
- Workers: 1
- Outstanding (Queue Depth): 128
- **Access Specifications: 4KB, 100% Read, 100% Random**

CE TestScenario

- HP server runs 1 x Volume (45GB) of RAID 5 Pool1 in Controller 1 via 2x FC ports withMPIO

- Dell server runs 1 x Volume (45GB) RAID 5 Pool2 in Controller 2 via 2x FC ports with MPIO

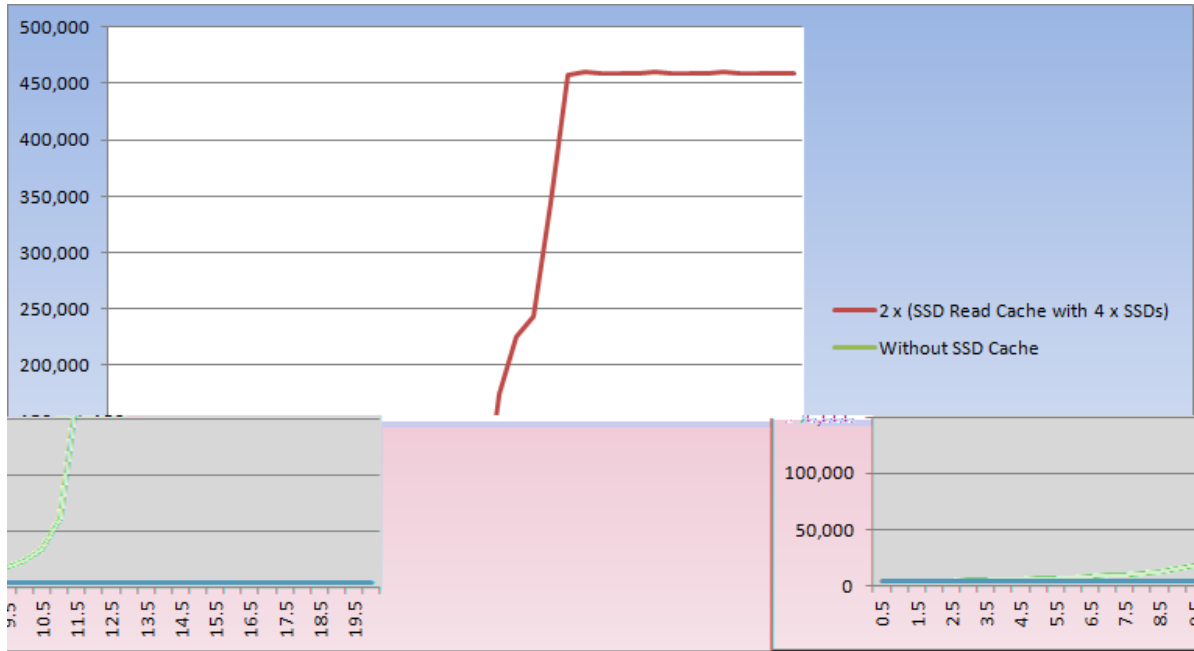
Test Diagram



Test Result

Volume Capacity	Disable SSD Cache (IOPS)	Enable SSD Cache			Improved
		No. of SSDs	SSD Capacity	Warm up Time	
2 x 90GB = 180GB	4,986	2 x 4 = 8	2 x 400GB = 800GB	13.5 min.	461,037
					9,147%

Without SSD cache, the average IOPS is 4,986. After enabling SSD cache with 8 SSDs, IOPS increases to 461,037. Improves $(461,037 / 4,986) = 91.466 = 9,147\%$ warm up time is about 13.5 minutes.



Summary

- SSD Cache 2.0 can improve random read performance by up to 46 times.
- The test is the highest SSD cache performance of a system

Test Case 5: Best Practice of SSD Write Cache

This test provides the best practice of SSD cache on dual controllers before we assume that the performance can be twice that of the test on single controller.

Test Equipments and Configurations

- Server
 - Model: HP Z840 (CPU: 2 x Xeon E52620v3 2.4Hz / RAM: 32GB)
FC HBA: QLogic QLE2693R
OS: Windows Server 2012 R2
 - Model: Dell E25S (CPU: 2 x Xeon E52620v3 2.4Hz / RAM: 32GB)
FC HBA: QLogic QLE2693R
OS: Windows Server 2012 R2

Conclusion

The hybrid storage concept of storage acceleration uses the idea of hot data to accelerate I/O performance of an entire storage system. When hardware and IT administration costs are taken into consideration, it turns out that SSDs are available in modern SAN systems is generally the best way for most businesses to gain the benefits of the faster performance from flash based storage without sacrificing the reliability of their data.

Apply To

☒ XCubeSAN XS5200 / XS3200 / XS1200 1.1.2 and later

Reference

Obsolete QSAN White Paper

☒ QWP201403SSD_Caching.v1.5.pdf

Appendix

Related Documents

There are related documents which can be downloaded from the website.

- ☒ [All XCubeSAN Documents](#)
- ☒ [XCubeSAN QIG \(Quick Installation Guide\)](#)
- ☒ [XCubeSAN Hardware Owner Manual](#)
- ☒ [XCubeSAN Configuration Worksheet](#)
- ☒ [XCubeSAN SANOS 4.0 User Manual](#)
- ☒ [Compatibility Matrix](#)
- ☒ [White Papers](#)
- ☒ [Application Notes](#)

Technical Support

Do you have any question or need help troubleshooting a problem? Please contact QSAN Support we will reply to you as soon as possible.

- ☒ Via the Web <https://qsan.com/support>
- ☒ Via Telephone: +886-77202118 extension 136
(Service hours: 09:30-18:00, Monday - Friday, UTC+8)
- ☒ Via Skype Chat, Skype ID: qsan.support
(Service hours: 09:30-18:00, Monday - Friday, UTC+8, Summer time: 09:30-17:00)
- ☒ Via Email: support@qsan.com