
Using Synology SSD Technology to Enhance System Performance

Based on DSM 5.2

Table of Contents

Chapter 1: Enterprise Challenges and SSD Cache as Solution

Enterprise Challenges.....	3
SSD Cache as Solution	3
SSD Cache on Synology NAS	3

Chapter 2: How Synology SSD Cache Works

Data Operations and Hot Data	4
SSD Cache on Synology NAS	5
Reading and Buffering Data with SSD Cache	6
Writing Data to SSD cache.....	7
Removing and Flushing Data from SSD Cache	7
Hit Rate, Warm-Up Period, and Cache Size.....	8
Acceleration and Maximum System Performance.....	9
Accelerated Protocols and Non-Accelerated Workloads	9
Determining the SSD Cache and RAM Module Size and Monitoring SSD Cache Performance	9

Chapter 3: Test Results

Test Case.....	11
Testing Configuration	11
Results and Analysis.....	12

Chapter 4: Conclusion

Enterprise Challenges and SSD Cache as Solution

Enterprise Challenges

High information system productivity demands low latency, and the requirements for I/O latency are most stringent when running mission-critical business applications. For any IT deployment, the greatest challenge is achieving a balance between low latency, high efficiency, and optimized system-utilization rate.

Degree of I/O latency in a storage system is determined by two factors: I/O workload pattern and storage media capabilities. The majority of business applications (e.g. OLTP databases or email services) involve random IOPS, which access data stored non-contiguously on system disks. As required bits of data are not located within physical proximity to one another, the system produces numerous seeking processes, thereby increasing I/O latency.

Traditionally, to overcome high I/O latency as a result of random IOPS workloads, a larger than necessary number of disks may be deployed to increase the number of heads and reduce the chance of two consecutive reads on the same disk, thereby boosting access performance. However, there are several drawbacks to over-deployment, including lower efficiency and overall system-utilization. More specifically, increasing the number of disks necessarily involves increasing number of enclosures, amount of space required, power consumed for operating and cooling, and ultimately leads to higher maintenance costs. Moreover, system-utilization rate may diminish as unnecessary capacity is added to reach the requisite amount of heads.

SSD Cache as Solution

By leveraging the superior random access performance of Solid State Drives (SSD), Synology SSD Cache technology provides a solution to enterprise challenges which boosts read and write speeds without adding to the overall number of disks.

In terms of read operations, statistically only a small portion of data in any given storage system is accessed frequently. System performance can therefore be improved by storing frequently accessed data on the SSD Cache to create a read buffer while maintaining total cost of ownership at a minimum.

Small, random write operations are also common in enterprise workloads. With excellent random access, SSD read-write cache can accelerate the performance of volumes and iSCSI LUNs, reduce the latency of the random write operations, and also greatly reduce any impact on the performance of other data transfers.

SSD Cache on Synology NAS

Synology's read-write and read-only SSD Cache technology is available on all XS+, XS, and selected Plus series products¹. By attaching each SSD cache to a single storage volume or iSCSI LUN (block-level), Synology SSD Cache can create a read and write buffer which greatly enhances the system performance of your Synology NAS.

¹ Please visit Synology website (www.synology.com) for applied models supporting SSD cache.

How Synology SSD Cache Works

Data Operations and Hot Data

Typically, on receiving a read request, servers first check if relevant data is located in the system memory cache known as RAM, which stores the most recently access information. If the requested data is absent, read processes on disks will be triggered. As the RAM size is severely limited compared to the working data-set, most retrieval requests necessarily result in reading from disks and therefore increased latency.




Figure 2 Accessing Data from RAM and disks

In most applications, there are observable patterns in data retrieval and workload due to the particular I/O characteristics of the application's behavior. For instance, in an OLTP database workload, some tables in the database are more frequently read than others. The most frequently accessed data is termed "hot data."

Of this "hot data" subset, the most recent data has an even higher probability of being frequently accessed. In the majority of critical, business workloads, the most recently accessed data is also the most relevant and therefore in need of timely retrieval.




Figure 3 Hot data, cold data, and appropriate cache size

As the hot data is merely a portion of the whole data set with the most intensive I/O requests, a small number of SSDs can be used to cache all hot data, thereby leveraging its superior I/O capabilities to significantly improve system performance.

SSD Cache on Synology NAS

Read cache can be created with only one SSD. If larger read cache or read-write is required, two SSDs of the same model and make must be installed in the server¹. For read-only cache, the SSDs are configured using RAID 0. For read-write cache, SSDs are configured using RAID 1 to ensure data integrity in case one SSD fails.

Once installed, each SSD Cache can be attached to any one volume or iSCSI LUN (block level) in the system. If a volume has SSD Cache enabled, the iSCSI LUN (file level) on the volume will also benefit from the increased performance.

¹ To ensure the best performance, the SSDs should be of the same size, brand, and model, and listed on Synology's official Hardware Compatibility List

Reading and Buffering Data with SSD Cache

Once SSD Cache has been enabled, data requested from disks is continually copied and stored onto the SSD disks. A memory map will be created in the system RAM to record which data blocks are on the SSD Cache. Therefore SSD Cache size and system RAM size are proportionally correlated.




Figure 4 Memory mapping of cached blocks

If further requests for the same data are generated, read operations will be conducted on SSD in what is termed a “cache hit.” As data is retrieved from cache, reading performance is enhanced.




Figure 5 Data access from cache as “cache hit”

Conversely, when a reading request is sent for data and the data is not located on the SSD Cache, the situation is termed “cache miss.” A disk reading operation will be triggered after cache miss, and a copy of the data requested will be made on the SSD Cache to accelerate the reading speed of any future requests.





Figure 6 Data from disks as “cache miss”

The SSD Cache manipulates all data in block-level. For instance, when reading a 400 KB piece of data from a 4 GB file, only the relevant 400 KB of data will be accessed. Furthermore, if these 400 KB of data are absent, the system will read from disks and copy it into SSD Cache.

Writing Data to SSD cache

Synology SSD Cache employs a write-back cache approach in read-write cache to accelerate write operations. When the Synology NAS receives a write request, it stores the data on SSD cache blocks and then acknowledges the write operation to the application server.



Removing and Flushing Data from SSD Cache

Once SSD Cache is full, any cache misses will trigger a flushing mechanism which removes some of the data from the cache to make room for more relevant data. SSD Cache uses a “Least Recently Used” (LRU) replacement algorithm for data-rotation.

The LRU algorithm considers not only the sequence of data accessed, but also the time of its last read. The least recently accessed data is removed first for new insertions. Data blocks which are accessed more frequently (i.e. hot data) will be kept for longer periods and remain in the cache loop as long as there is continued benefit to be gained from its caching.

Hit Rate, Warm-Up Period, and Cache Size

For read operations, the SSD Cache “hit rate” represents the ratio of cache hits, whereby data requests can be fulfilled by the SSD Cache. A higher hit rate indicates fewer read operations on disks and therefore lower I/O latency. Conversely, a lower hit rate indicates that most data being read is fetched from disks, where response time is comparable to reading from volumes without the benefit of SSD Cache. As the SSD Cache hit rate increases, response times will reflect the I/O latency of SSD disks, which is much lower than that of HDD disks.




Figure 7 Warm-up period and cache hit rate

As the memory map in the SSD Cache is empty at the start, almost every data reading operations will cause a cache miss. Meanwhile, copies of data will be added to the SSD Cache continually. This period is called “warm-up” and is mainly composed of copy operations. The warm-up period can also occur if the working data set changes drastically, so that currently cached data is no longer requested. A high cache hit rate indicates that the SSD Cache is being fully utilized. This indicator will grow through the warm-up period.

System hit rate is dependent on two factors: the size of SSD Cache and size of hot data. Higher hit rate requires that more hot data is stored on the SSD Cache. For instance, in a 2 TB file server with 100 GB of frequently accessed data, the recommended cache size would be slightly above 100 GB. However, as the size of hot data is 100 GB, the benefit of setting up a 500 GB SSD Cache configuration would be limited.

You can view read hit rate information in the management user interface of Storage Manager.




Figure 8 Cache read hit rate information displayed in management user interface

Acceleration and Maximum System Performance

SSD Cache is designed to improve the performance of disk subsystems. It does not accelerate the maximum system performance where the bottleneck is equivalent to CPU capacity, outbound connection bandwidth, and other system limitations. In other words, if a system can perform 100,000 IOPS at best, it will not be able to increase IOPS through SSD Cache.

Accelerated Protocols and Non-Accelerated Workloads

As data is cached in blocks, all file and block-level storage services will benefit from SSD Cache, including CIFS, iSCSI and so on. Virtualization applications, cloud computing, OLTP databases, mail servers and file services which have predominantly re-read workload patterns will also benefit from SSD Caching. In the following section, we will compare the testing results from a workload pattern based on email servers.

In certain workload patterns, the benefits of SSD caching will be very limited. For example, large sequential read or write operations (e.g. single channel HD video streaming or surveillance video recording) or entirely random data reading patterns (e.g. music streaming services with a playlist randomizer in a large music library), both of which lack re-reading patterns, will not benefit significantly from SSD Caching.

Determining the SSD Cache and Monitoring SSD Cache Performance

The aggregate size of hot data and appropriate cache size can be determined with the help of Synology DiskStation Manager's SSD cache Advisor, where the number of unique files¹ accessed in the last seven days and their aggregate size will be displayed. For example, if the accessed file size is 200 GB daily, the corresponding SSD Cache size should be 200 - 300 GB to ensure that all hot data can be accommodated.

¹ As differentiated by path and file name.




Figure 9 DSM's built-in SSD cache advisor

Additionally, a mapping table will be created in the RAM module to track the SSD cache data. Therefore, RAM size must be proportional to the size of SSD cache. Every 1 GB of SSD cache requires approximately 400 KB of system memory. For Intel-based Synology NAS, only 1/4 of the pre-installed system memory can be used for SSD cache. For Annapurna-based Synology NAS, the total size of SSD cache is limited to 1 TB.

You can also measure the IOPS in Windows and VMware environment using built-in Performance Monitor tools. For more information, please refer the blog posts:

- How to measure IOPS for Windows
<http://blog.synology.com/blog/?p=2086>
- How to measure IOPS for VMware
<http://blog.synology.com/blog/?p=2225>

Test Results

Test Case

In this performance test, IOMeter was used to generate workload to simulate applications or virtual machines which require intensive I/O and low latency, with many re-read operations in a small portion of a larger data-set. As a result, enabling SSD cache is expected to improve performance.

Testing Configurations

Tests were conducted with the following configurations.

Storage Server Configurations #1:

Model: RS3614xs+
Hard drive: WD4000FYYZ x 10
SSD cache: Intel 520 Series SSDSC2CW24 240GB
RAID type: RAID 5

Storage Server Configurations #2:

Model: RS2414+
Hard drive: WD4000FYYZ x 10
SSD cache: Intel 520 Series SSDSC2CW24 240GB
RAID type: RAID 5

IOMeter Settings:

The number of outstanding I/Os: 32 per target
Worker: 1 (per share)
Running time: 3 minutes (3 times)
Ramp up time: 30 seconds
Data size: 20 GB
SSD cache size: 10 GB and 20 GB
Workload: Random 4KB IOPS

Results and Analysis

The total achieved IOPS, including both read and write operations, is shown in the following figures. The performance of data size equaling to cache size is tested for reference only.

Test Configuration #1	No Cache	10 GB SSD Cache (Data Size = 2 x Cache Size)	20 GB SSD Cache (Data Size = Cache Size)
100% Write	562.023	1373.807	11077.01
50% Read 50% Write	586.03	1336.607	9785.493
100% Read	1015.73	1434.947	14425.953




Figure 8.1 Total achieved IOPS for RS3614xs+ (higher is better)

Test Configuration #2	No Cache	10 GB SSD Cache (Data Size = 2 x Cache Size)	20 GB SSD Cache (Data Size = Cache Size)
100% Write	446.547	1209.883	4321.457
50% Read 50% Write	587.557	1452.51	4811.27
100% Read	1164.66	1634.89	6468.373




Figure 8.2 Total achieved IOPS for RS2414+ (higher is better)

Conclusion

The test results demonstrated that performance can be significantly improved by leveraging SSD Cache technology, without the need to add more HDDs to boost IOPS capability. SSD cache can improve IOPS and provide more system capacity to boost service performance up to 2.7x faster.

Please note that testing results were derived through restricted conditions and specific configurations of our testing lab and that the results might vary in different environments.