Synology

White Paper

Using Synology SSD Technology to Enhance System Performance

Synology Inc.

Synology_WP_ 20121112

Table of Contents

Chapter 1

Enterprise Challenges and SSD Cache as Solution

Enterprise Challenges

High information system productivity demands low latency, and requirements for I/O latency are most stringent for 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 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-read workloads 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-read workloads, a larger than strictly 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 read-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

Synology SSD Cache technology provides a solution to enterprise challenges which boosts read-speed without adding to the overall number of disks by leveraging the superior random access performance of Solid State Drives (SSD).

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-only buffer while maintaining total cost of ownership at a minimum.

SSD Cache on Synology NAS

Synology's proprietary read-only SSD Cache technology, which supports up to two SSD disks, is available on all XS+ series products. By attaching each SSD to a single storage volume or iSCSI LUN (block-level), Synology SSD Cache creates a read-only buffer which greatly enhances system performance.

Chapter 2

How Synology SSD Cache Works

Data reading operations and Hot Data

Typically, on receiving a read request, severs first check if relevant data is located in the system memory cache-buffer known as RAM, which stores the most recently access information. If the requested data is absent, read processes on disks will be triggered. However, 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.

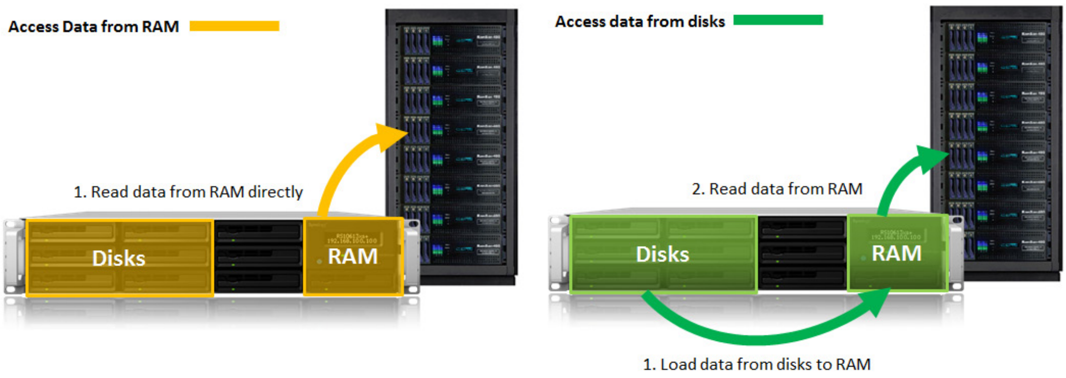

Figure 2 Accessing Data from RAM and disks

In most of applications, there are observable patterns in data retrieval and workload due to the particular I/O characteristics of application behavior. For instance, in an OLTP database workload, some tables in the database are more frequently read than others. The mostly 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 most business critical workloads, the most recently accessed data is also the most relevant and therefore in need of timely retrieval.

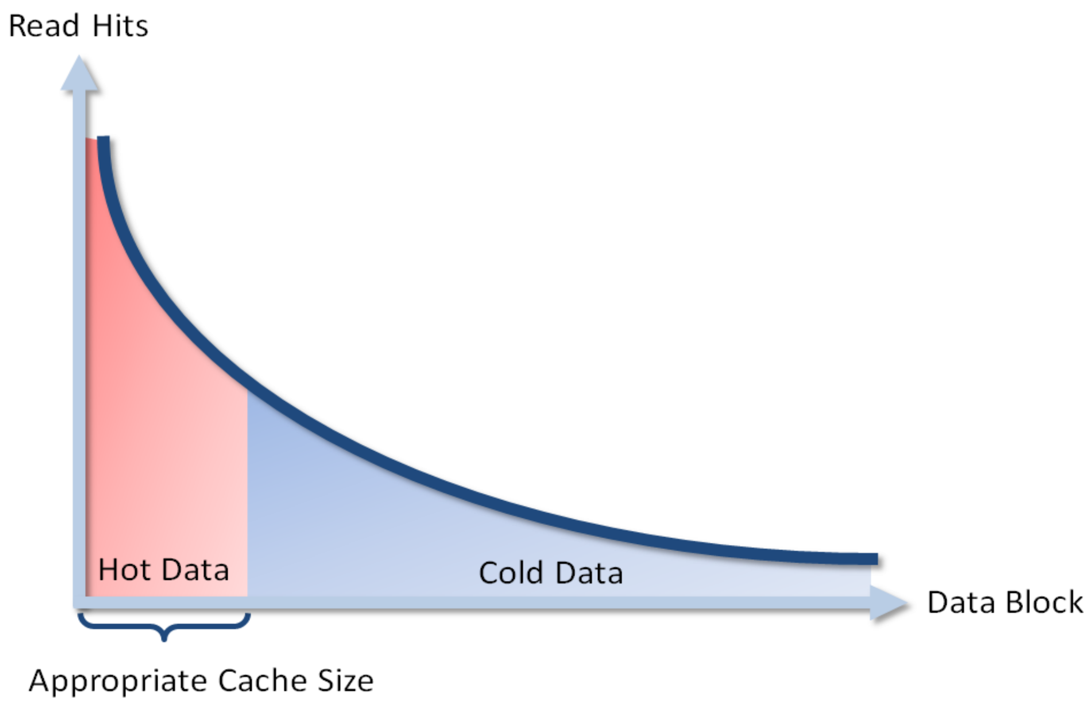

Figure 3 Hot data, cold data and appropriate cache size

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

Synology SSD Cache is a read-only cache which optimizes data retrieval. To enable SSD Cache on a Synology NAS, two SSDs must be installed in the SSD Cache slots of the server[1]. For best reading performance, the SSDs is configured in RAID 0. Once installed, each SSD can be attached to any one volume or iSCSI LUN (block level) on the system. If a volume is SSD Cache enabled, the iSCSI LUN (file level) on the volume will also benefit.

[1] The SSDs should be of the same size, brand, and model, and must be 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.

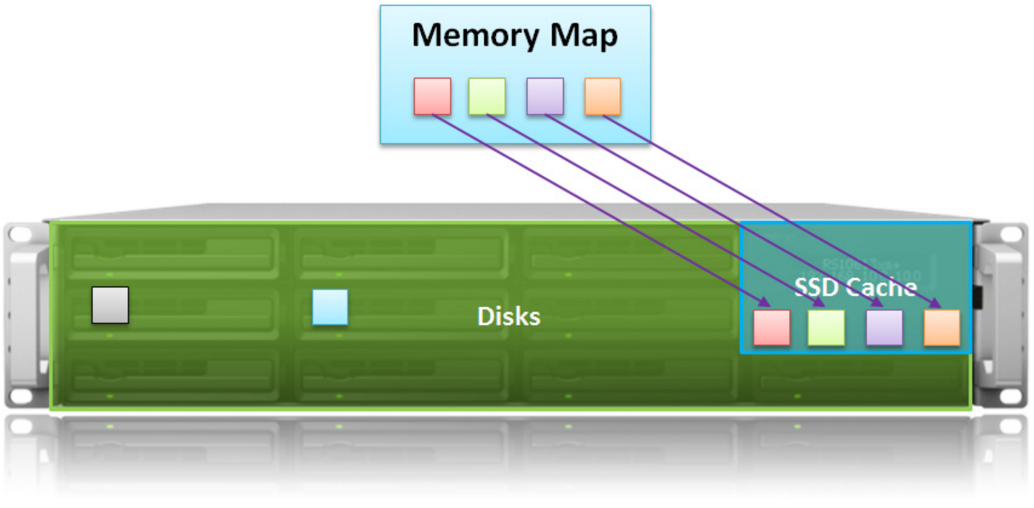

Figure 4 Memory mapping of cached blocks

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

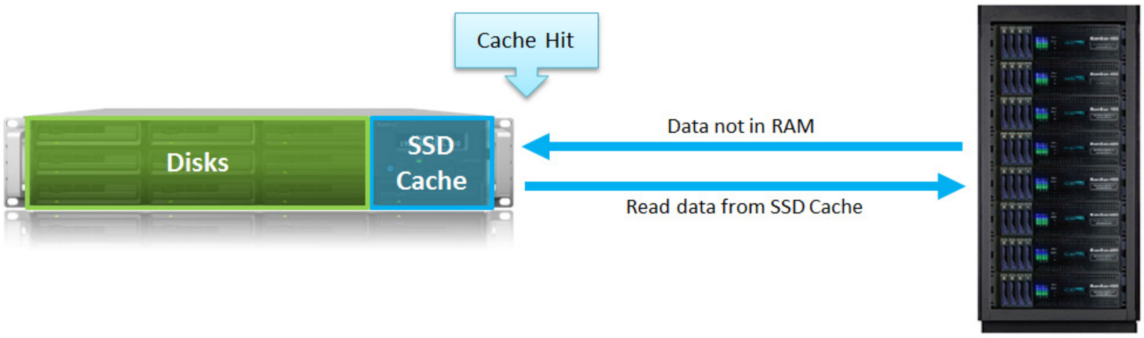

Figure 5 Data access from cache as “cache hit”

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.

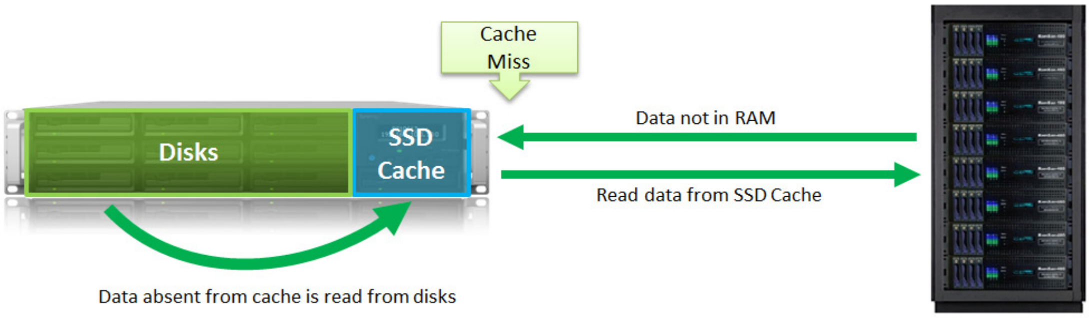

Figure 6 Data from disks as "cache miss"

The SSD Cache manipulates all data in 512-byte blocks. For instance, to read a 4 KB piece of data from cache, 8 x 512-byte data blocks will be fetched. In addition, any data requests will be monitored at block-level. For instance, when reading a 400 KB piece of data from a 4 GB file, only the relevant 400 KB 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.

Removing and Rotating Data from SSD Cache

Once the SSD Cache is full, any cache misses will trigger a data-rotation mechanism which removes some of the data from the cache to make room for more relevant data. SSD Cache uses a "First In, First Out" (FIFO) replacement algorithm for data-rotation.

The FIFO 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

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 reading 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.

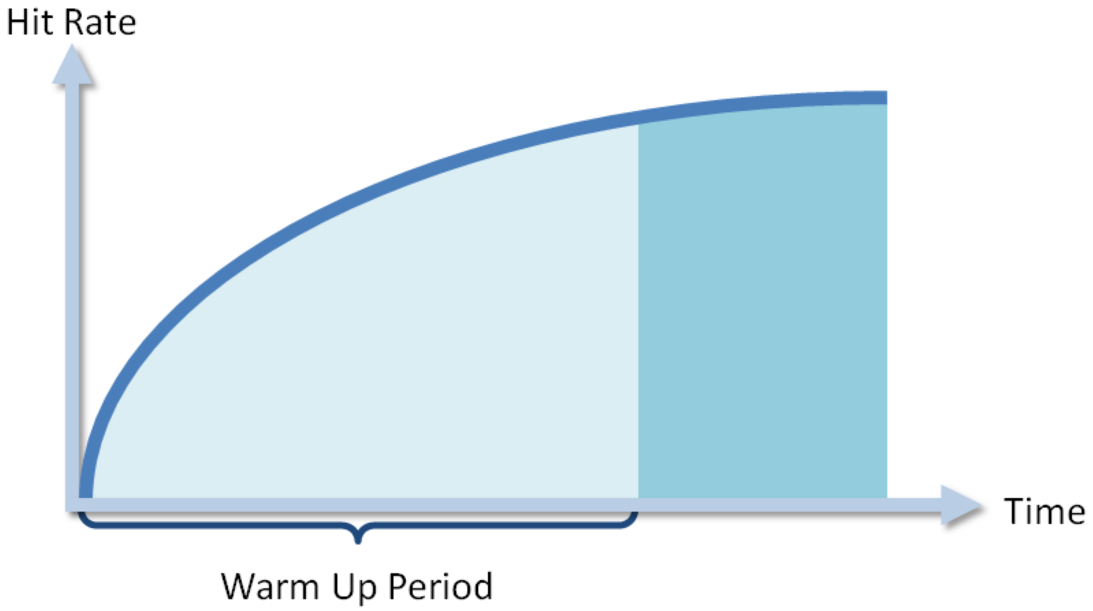

Figure 7 Warm-up period and cache hit rate

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 the "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 designing a 500 GB SSD Cache configuration would be limited.

Read Acceleration and Maximum System Performance

SSD Cache is a read-only mechanism which is designed to accelerate reading performance alone. Despite this, some performance enhancement on intensive writing workloads may be observed as disks are freed for writing operations while SSD Cache handles most reading requests. However, under normal circumstances, write performance will not be directly improved by SSD caching.

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 modules.

Accelerated Protocols and Non-Accelerated Workloads

As data is cached in blocks, all block-level protocols will benefit from SSD Cache, e.g. 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 workload pattern based on email servers.

In certain workload patterns, due to low re-read rate, benefits of SSD caching will be very limited. Workloads which consist mostly of data writing operations (e.g. data archiving or backup solutions) draw only residual benefits from SSD Caching, as writing cannot be accelerated by this method.

Large sequential reads (e.g. single channel HD video streaming) or entirely random data reading patterns (e.g. music streaming services with playlist randomizer in a large music library), both of which lack re-reading patterns, will also not benefit significantly from SSD Caching.

Determining the SSD Cache and RAM Module Size 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 built-in tools, where the number of unique files[1] 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.

Additionally, a mapping table will be created in the RAM module to track SSD Cache data. Therefore RAM size must be proportional to size of SSD Cache. The following table provides a guideline of appropriate RAM to SSD Cache ratio.

Table 1 Recommended RAM to SSD Cache size ratio

Cache Size	RAM Size
Less than 512 GB	8 GB
512 GB 1024 GB	16 GB

Synology DiskStation Manager's built-in Performance Monitor can be used to oversee the performance of an SSD cached volume. Enhancements can also be observed by comparisons to historical statistics collected prior to enabling SSD Cache.

[1] As differentiated by path and file name.

Chapter 3

Test Results

Test Case

In the performance test, Microsoft Exchange 2010 was used as a workload example, as email service is a typical application which requires intensive I/O and low latency, with many re-reads in a small portion of a whole data-set. For example, important email and most recent exchanges are re-read more regularly. For these reasons, enabling SSD Cache is expected to improve performance.

Testing Configuration

SSD Cache performance enhancement was assessed using Jetstress as the workload generator and testing tool. Jetstress simulates disk I/O load on a test server running Microsoft Exchange to verify storage system performance and stability.

Storage Server Configurations:

Model: RS10613xs+
Hard drive: WD4000FYYZ x 10
SSD cache: Intel 520 Series SSDSC2CW24 240GB 6Gbit/s
RAID Type: Raid 5
RAM: 8GB

Workload-generating Server Configurations:

CPU: Intel Core i5-3550 3.30GHz
RAM: 8GB DDRIII
OS: Windows Server 2008 R2

Jet Stress settings:

Mailbox size: 2GB
IOPS: 0.12
Database: 2
Users: 500
Thread: 5
Test time: 2 Hours (warm-up 12 hours)

Results and Analysis

The total achieved database IOPS, including both read and write operations, are shown in the following figures. Without SSD Cache, the total achieved database IOPS was 155; with SSD Cache, the figure was raised to 241, at a hit rate of 17%. Therefore the total achieved database IOPS was 55.48% higher once SSD Cache was enabled.

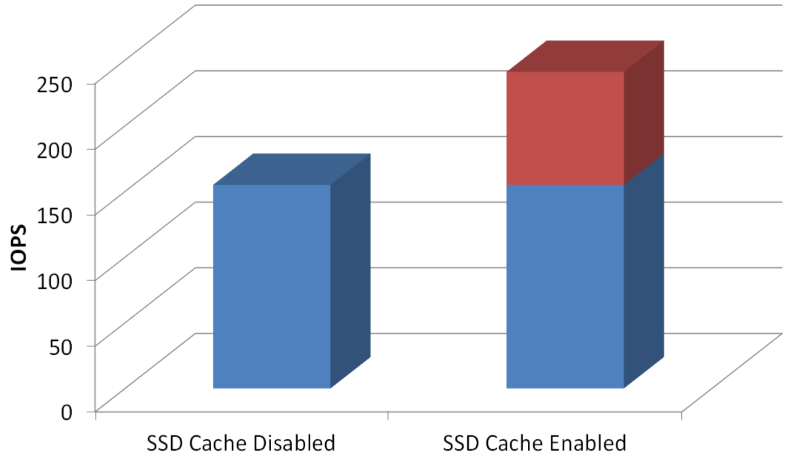

Figure 8 Total achieved database IOPS (higher is better)

Chapter 4

Conclusion

Test results demonstrated that performance can be significantly improved by leveraging SSD Cache technology on Microsoft Exchange 2010, even without adding more HDDs to boost IOPS capability. SSD Cache can improve IOPS and provides more system capacity to boost service performance. As a 17% hit rate was capable of generating a 55.48% improvement in IOPS, higher performance gains can be expected in an environment with greater hit rate.

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