

## Solid State Drive (SSD) Cache: Enhancing Storage System Performance

### Application Notes

***Version: 1.2***

**Abstract:**

This application note introduces Storageflex HA3969's Solid State Drive (SSD) Cache technology's key benefits and detail operations to help readers better understand SSD Cache. SSD Cache leverages the high speed and low latency of solid state drives in delivering faster read performance. The requirement details for SSD Cache are also provided in this paper in order to assist readers to implement this feature as well as integrated SSD wear leveling technology that minimizes the impact of the limited write life cycle of SSD more efficiently and effectively.



# Contents

<b>Contents.....</b>	<b>2</b>
<b>Introduction .....</b>	<b>3</b>
What is SSD Cache? .....	3
Why HA3969 SSD Cache? .....	3
<b>Storageflex SSD Cache .....</b>	<b>5</b>
How HA3969 SSD Cache works.....	5
SSD Cache Requirements.....	6
<b>How to use SSD Cache for HA3969 via SANWatch.....</b>	<b>10</b>
<b>How to use SSD Cache for HA3969 via EonOne .....</b>	<b>12</b>
<b>About SSD Wear Leveling.....</b>	<b>13</b>
<b>SSD Remaining Life Notification .....</b>	<b>14</b>
<b>Conclusion.....</b>	<b>15</b>



# Introduction

With rapid development of businesses, the performance of workload-intensive applications is often constrained by existing hard disk drives (HDDs). While the storage capacities of HDDs have increased dramatically, their performance has not. However, it is possible to accelerate the performance of read-intensive workloads, such as On-Line Transaction Processing (OLTP), Web Server, File Server or Database. Deploying Storageflex's SSD cache technology that leverages the high throughput and low latency of solid state drives in delivering faster read performance for frequently accessed high demand and vital data can dramatically boost system performance and reduce response time.

## What is SSD Cache?

Cache is a component that transparently stores data so that future requests for the data can be served more efficiently. It is fundamental to storage, especially for read-intensive applications. Because the capacity of controller memory cache is limited, SSD Cache allows fast SSD to be used to extend the cache pool capacity of storage systems and store frequently accessed data. With increased capacity of SSD Cache, more hot data can be stored in the SSD Cache so future requests for the data can be served more efficiently and the read performance is dramatically enhanced.

## Why Storageflex HA3969 SSD Cache?

In most cases, if an application's workload has high a percentage of reads relative to writes and a small amount of data is repetitively read, SSD Cache may have the following benefits:

### 1. Read performance enhancement

SSD Cache leverages an intelligent algorithm to accelerate IO performance of applications with read-intensive workloads such as On-Line Transaction Processing (OLTP) and Database, having large numbers of small random read I/Os. In these scenarios, SSD Cache can greatly improve overall read performance. For example, a system with SSD cache can deliver a 2.5 times better IOPS performance compared to the same system without SSD Cache in the case of OLTP workload and latency is also reduced. The level of performance improvement is based on the application's actual workload and user's access behavior.

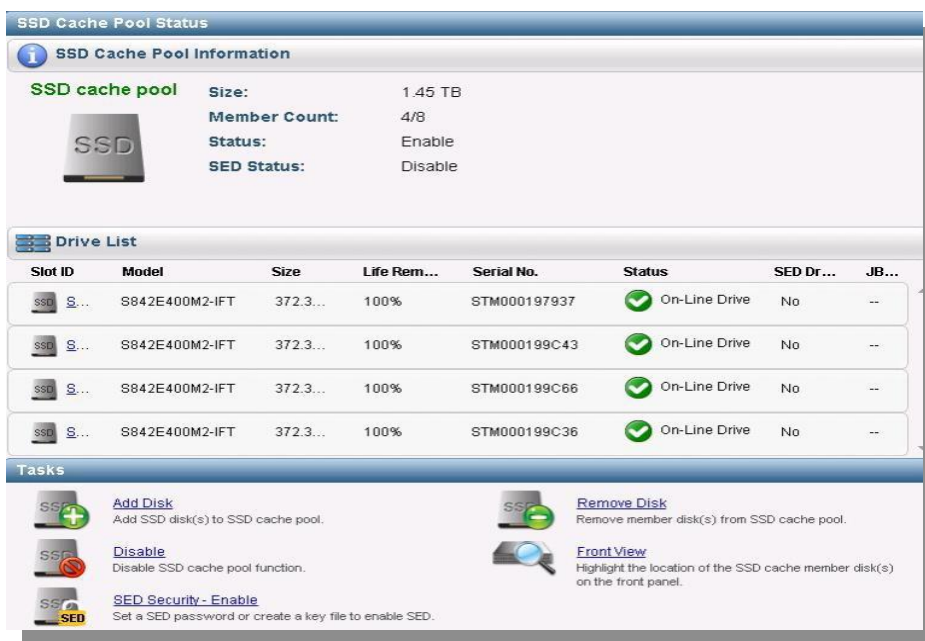
### 2. Intelligent firmware and management algorithm

The Intelligent firmware automatically analyzes data access behavior and recognizes sequential or random read/write I/O behavior. Data from sequential reads or writes are not written to the SSD Cache pool and only random read data is stored in the SSD Cache pool in order to make the most effective use of SSDs. More precisely, the firmware will automatically move a copy of the most frequently-accessed

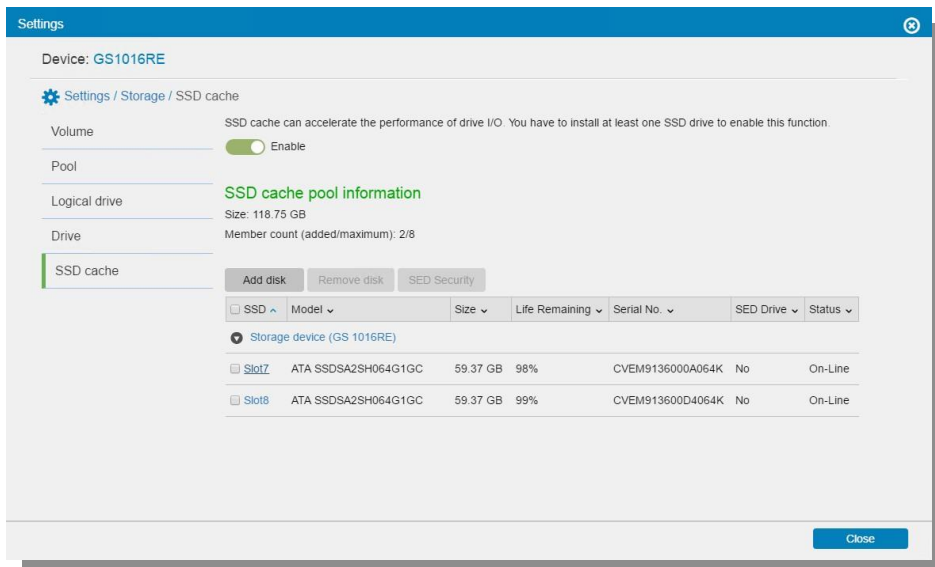
read data from the controller memory cache to the SSD Cache pool at an appropriate time. The hot data will be read from the SSD Cache pool if the system receives the same read request in the future. A proprietary write algorithm is used to copy data to SSDs and this optimizes the cycle duty. This solution not only improves read performance but also extends the life cycle of hard drives by reducing hard drive reads and writes.

### 3. Simple and intuitive GUI design

The SSD Cache function is fully integrated into the intuitive SANWatch and embedded RAIDWatch GUI for HA3969 or HA3969U, for simple setup, management and maintenance. Users can also monitor SSD Cache pool status and check the remaining life cycle for each SSD.



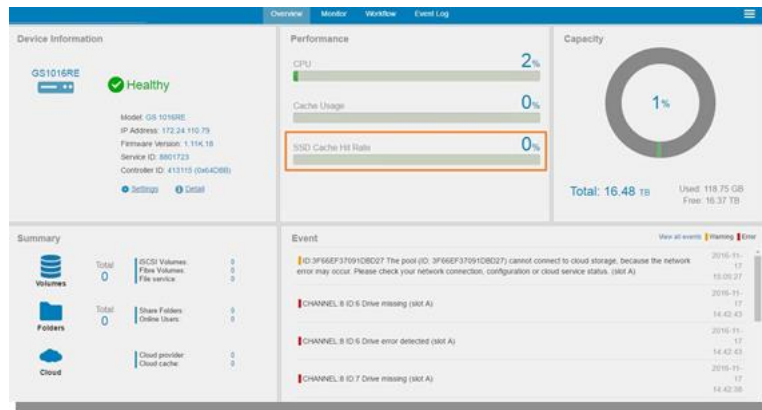
SANWatch GUI design for SSD Cache



EonOne GUI design for SSD Cache



EonOne Performance Quick Monitor shows the SSD Cache Hit Rate for a connected HA3969 device. The Cache hit rate is usually expressed as a percentage. If the requested data is found in the SSD cache, i.e. a read cache hit, the data is sent directly to the host without any disk operations. This provides a fast response time to the host. A higher read hit rate means improved read performance.



SSD Cache Hit Rate

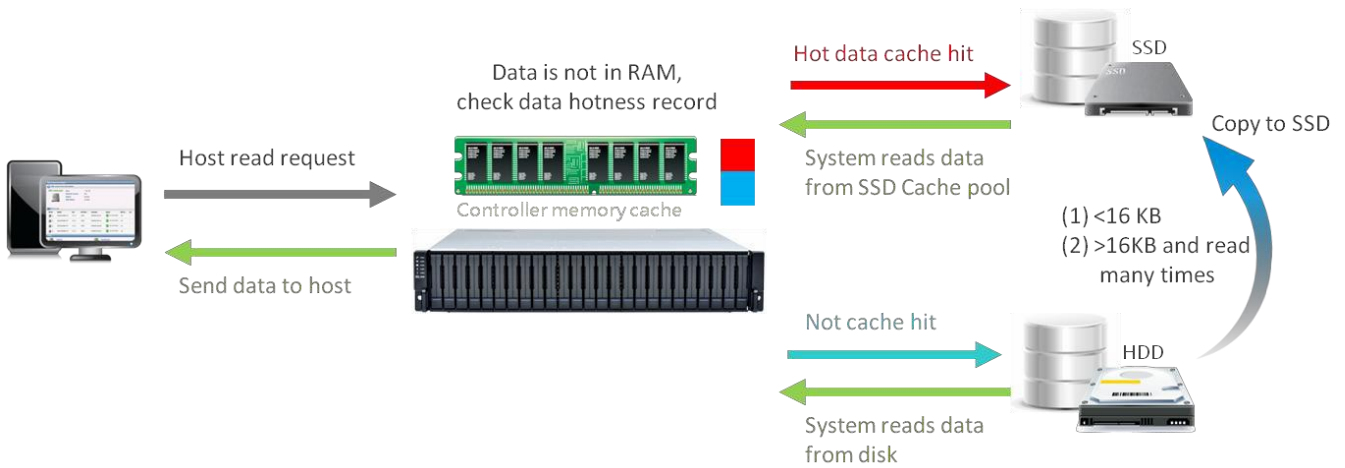
# Storageflex SSD Cache

## How Storageflex SSD Cache works

Once SSD Cache has been enabled and business applications have been running for a while, the intelligent firmware records statistical data and immediately update the hotness record into the controller memory cache. Based on the hotness record, the firmware will automatically copy the corresponding hot data from the disk to the SSD Cache pool at an appropriate time using the sequential write I/O method to avoid intensively writing data onto SSDs and thereby improving SSDs' life cycle. Before the SSD Cache Pool is filled up with hot data generated by the host applications, a pre-copy method to SSD via zone blocks prediction speeds up read performance. For data of a size that is less than or equal to 16KB, the data will be copied to the SSD Cache pool directly even if it has been read only once. When the read data size is larger than 16KB and the firmware has recognized it as hot data (been read a couple of times), it will then be classified as frequently accessed hot data and cached in the SSD Cache pool. This hot data will have two copies, one in the SSD Cache and one in the hard drives.

Basically, on receiving a read request, the system will check if relevant data is contained in the controller memory cache. If the requested data is stored in the memory cache, the system will directly return this data to the host. On the other hand, if the requested data is not cached in the controller memory cache, the

system will check the SSD Cache pool. If the requested data has been copied into the SSD Cache pool based on the real-time hotness record, the storage will directly read the data from the SSD Cache pool and return the data to the host. Otherwise, the requested data will be accessed from a slower device. Therefore, as more cache hits occur, more frequent requests can be served by SSD Cache and overall performance and average response time can improve.



How SSD Cache works

## SSD Cache Requirements

### 1. Firmware and SANWatch Requirements

HA3969 family	Version
Firmware	512F12 or later
SANWatch	3.0.h.14 or later

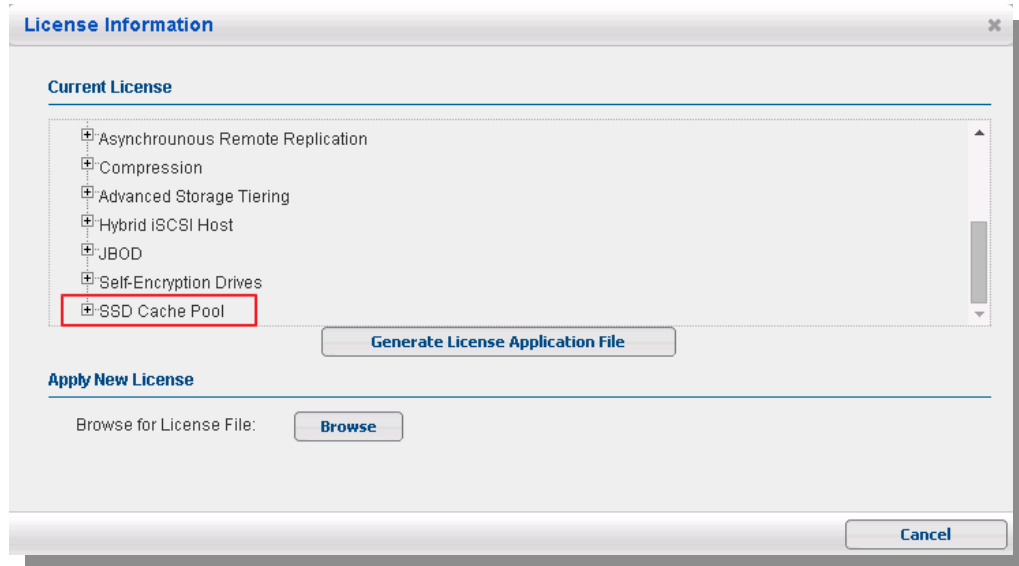
HA3969U family	Version
Firmware	All available
EonOne / Central EonOne	All available

Firmware and GUI supported versions

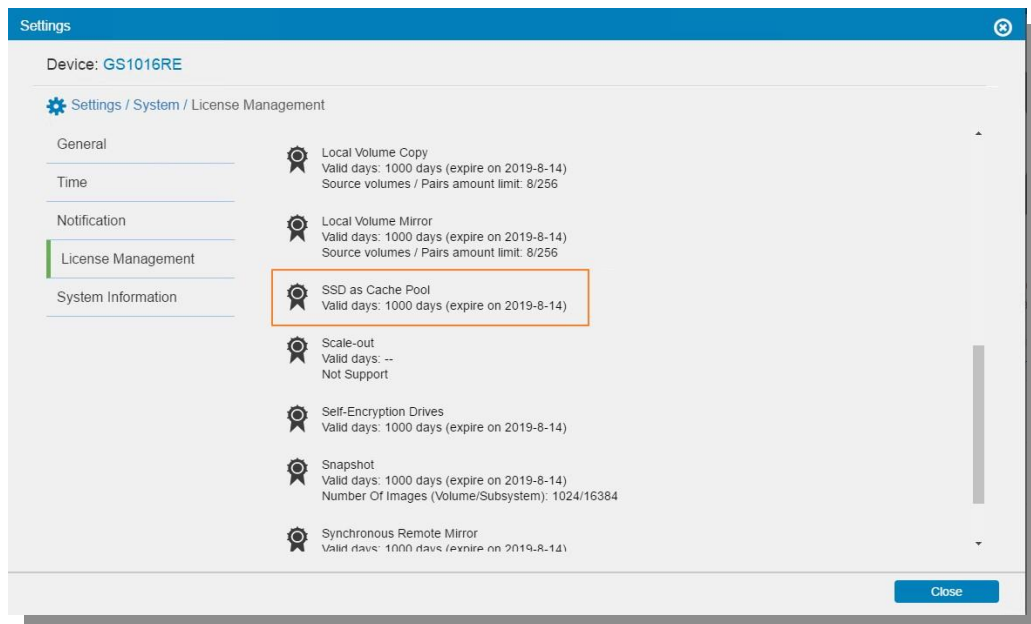


## 2. SSD Cache License

SSD Cache is a licensed feature. Storageflex also provides a 30-day trial license.



HA3969 license information



HA3969 license information



### 3. Controller Memory Cache Capacity V.S. Maximum SSD Cache Pool Capacity

Once SSD Cache is enabled, the controller memory cache will use a certain percentage of space to store the hotness record, and the maximum SSD Cache Pool size is determined based on the capacity for the hotness record in the controller memory cache.

DRAM	Max SSD Cache pool size
2GB	150GB
4GB	400GB
8GB	800GB
16GB	1,600GB
32GB	3,200GB
64GB	6,400GB
128GB	6,400GB
256GB	6,400GB

*Reference Table for HA3969 DIMM capacity supported for Maximum SSD Cache pool*

DRAM	Max SSD Cache pool size
8GB	400GB
16GB	600GB
32GB	1,000GB
64GB	1,600GB
128GB	3,200GB
256GB	3,200GB

*Reference Table for HA3969 DIMM capacity supported for Maximum SSD Cache pool*

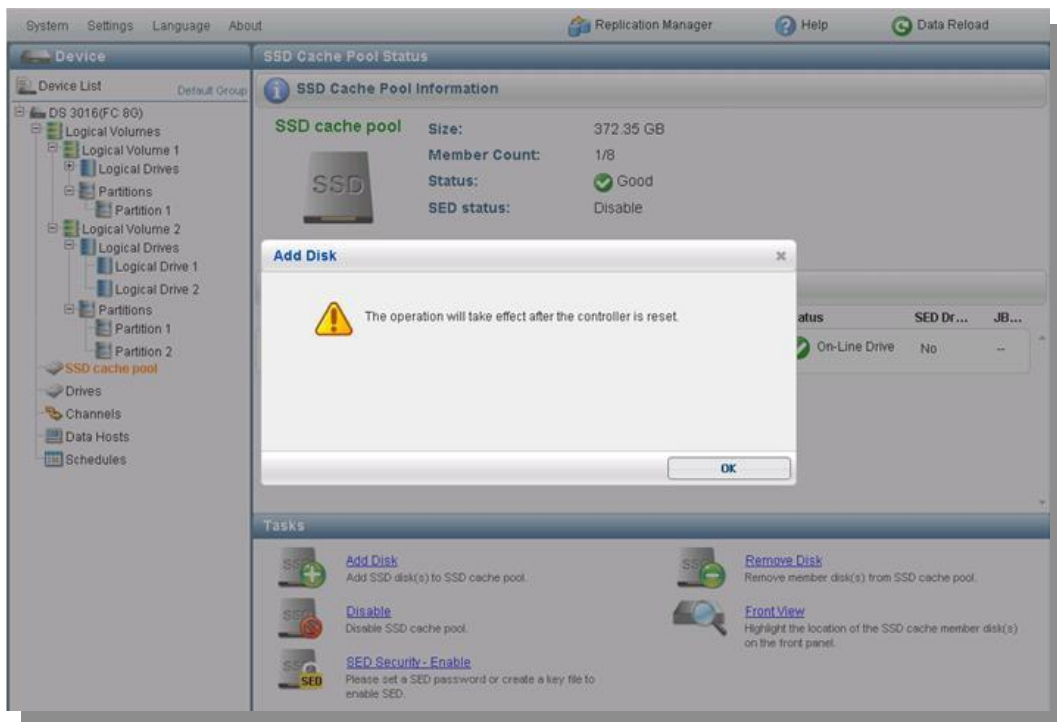
Compared to SSD, the controller memory cache is more economical. In addition, write cache can also be stored in the controller memory cache. Therefore, we recommend users upgrade their controller memory cache to max capacity and purchase suitable SSDs for the cache pool (according to needs and budget) to gain maximum overall performance benefits.

### 4. Resetting Controller(s) for Action

The final step to enable the SSD Cache function involves resetting the controller(s). By default, the controller memory cache does not allocate any space to store the hotness record. Therefore, the controllers need to be reset and initiated to allocate appropriate space in the controller memory cache to store the hotness record. After the initial controller reset and SSD Cache function activation, its management is easy. There's no need to reset or restart the system when adding or removing SSDs to



and from the SSD Cache pool. This can be done through the intuitive GUI via SANWatch / Embedded RAIDWatch.



*HA3969 Controller(s) needs to be reset for the first time*

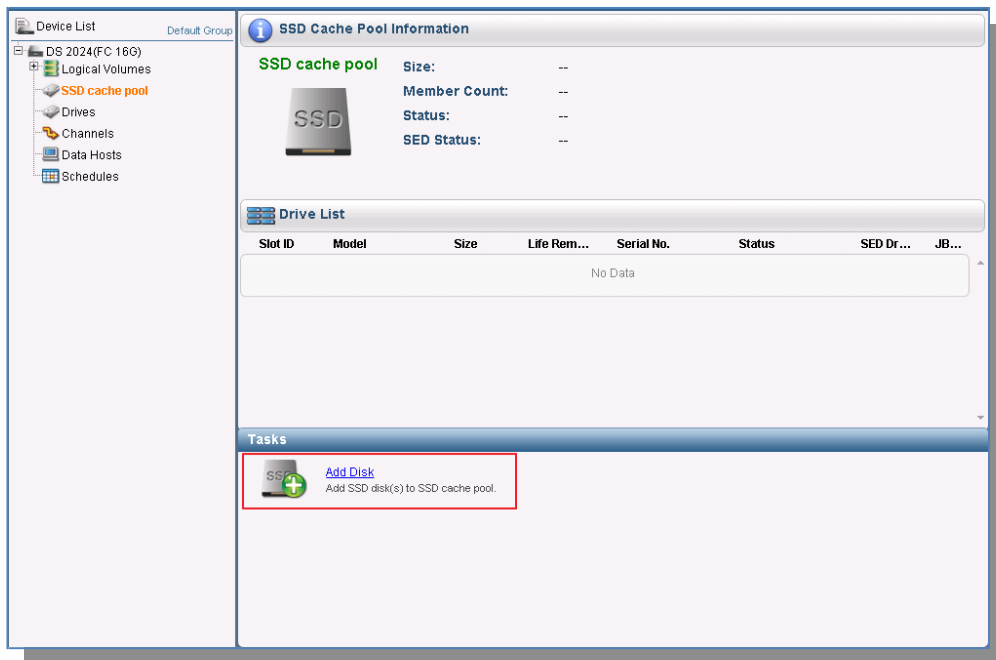
## 5. SSD requirement

Currently, one controller can support up to 4 SSDs. We strongly recommend using an even number of SSDs in redundant models. For HA3969 redundant models, at least 2 SSDs are required. If you wish to implement the SSD Cache function, please check whether the specific SSD model being used is on the Storageflex Qualified Vendor List (QVL). Only SSD products on our QVL can be applied to enhance performance as described in this document.

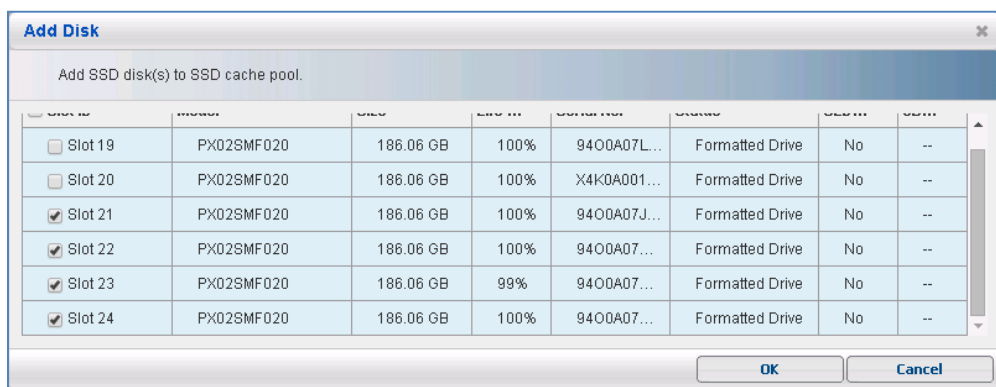


# How to use SSD Cache for HA3969 via SANWatch

1. Once you have purchased and applied the SSD Cache license, click SSD cache pool and click "AddDisk."

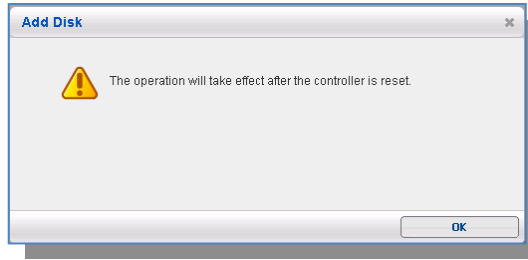


2. Select and add the SSDs to the SSD Cache Pool, and click OK when ready.





3. A window prompt will appear asking to reset the controllers for cache to take effect.



4. After the reset, you should see the SSD cache pool status showing "Enable." Users can add or remove SSD(s) from the pool or disable SSD Cache function without controller reset. In addition, the user can check each SSD's remaining life in percentages on the SSD Cache Pool Status page.

**SSD Cache Pool Status**

**SSD Cache Pool Information**

**SSD cache pool** Size: 744.18 GB  
Member Count: 4/8  
**Status:** Enable  
SED Status: Disable

**Drive List**

Slot ID	Model	Size	Life Rem...	Serial No.	Status	SED Dr...	JB...
SSD S...	PX02SMF020	186.0...	100%	9400A07JTY96	On-Line Drive	No	--
SSD S...	PX02SMF020	186.0...	100%	9400A07HTY96	On-Line Drive	No	--
SSD S...	PX02SMF020	186.0...	99%	9400A07DTY96	On-Line Drive	No	--
SSD S...	PX02SMF020	186.0...	100%	9400A07BTY96	On-Line Drive	No	--

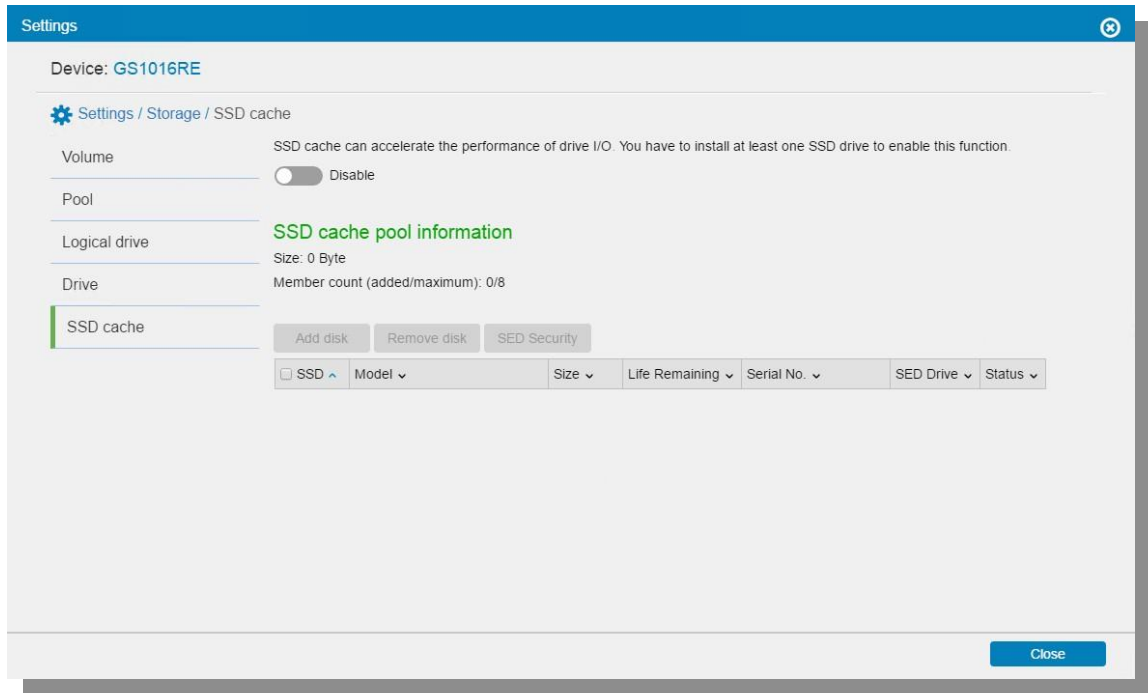
**Tasks**

- Add Disk**: Add SSD disk(s) to SSD cache pool.
- Remove Disk**: Remove member disk(s) from SSD cache pool.
- Disable**: Disable SSD cache pool function.
- Front View**: Highlight the location of the SSD cache member disk(s) on the front panel.
- SED Security - Enable**: Set a SED password or create a key file to enable SED.




# How to use SSD Cache for HA3969 via EonOne

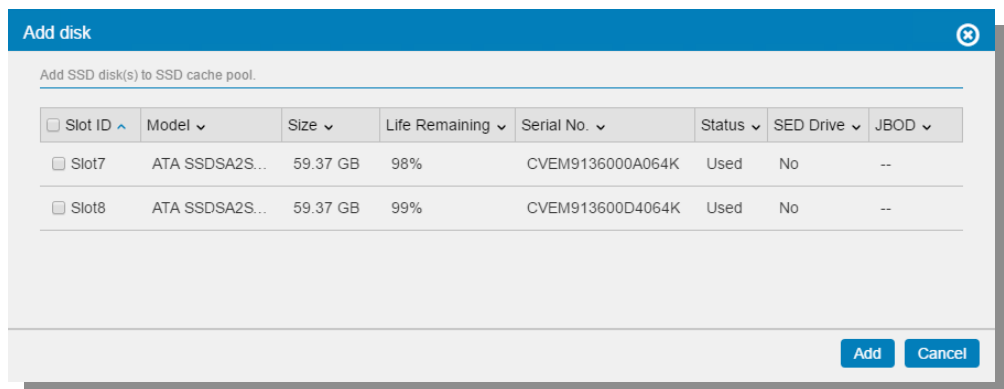
1. Once you have purchased and applied the SSD Cache license, go to Settings → Storage → SSD Cache.



2. Click the switch bar to enable/disable the SSD cache function



3. You will be asked to add new SSD disks into the SSD cache pool if no SSDs have been added into the SSD cache pool before.



You can see the list of installed SSDs and SSD Cache Pool Information. To add more SSDs, click the **Add disk** button.

SSD cache can accelerate the performance of drive I/O. You have to install at least one SSD drive to enable this function.

Enable

**SSD cache pool information**

Size: 118.75 GB

Member count (added/maximum): 2/8

<input type="checkbox"/> SSD ^	Model v	Size v	Life Remaining v	Serial No. v	SED Drive v	Status v
▼ Storage device (GS 1016RE)						
<input type="checkbox"/> Slot7	ATA SSSDA2SH064G1GC	59.37 GB	98%	CVEM9136000A064K	No	On-Line
<input type="checkbox"/> Slot8	ATA SSSDA2SH064G1GC	59.37 GB	99%	CVEM913600D4064K	No	On-Line

Click the link on the slot to see the information of the SSD on the slot.

**Drive Details**

View the parameters of the selected drive.

Slot: 7

Model: ATA SSSDA2SH064G1GC

Serial number: CVEM9136000A064K

Status: On-Line

Size: 59.37 GB

Speed: 0.3 GB

Revision Number: 8862

Logical Drive: CACHE

## About SSD Wear Leveling

Unlike HDD, SSD technology's overall write performance is limited by its endurance. In other words, it only supports a fixed number of writes. In order to assist users to easily monitor the remaining life of SSDs and plan ahead, Storageflex storage solutions provide an intuitive GUI design with easy-to-understand native



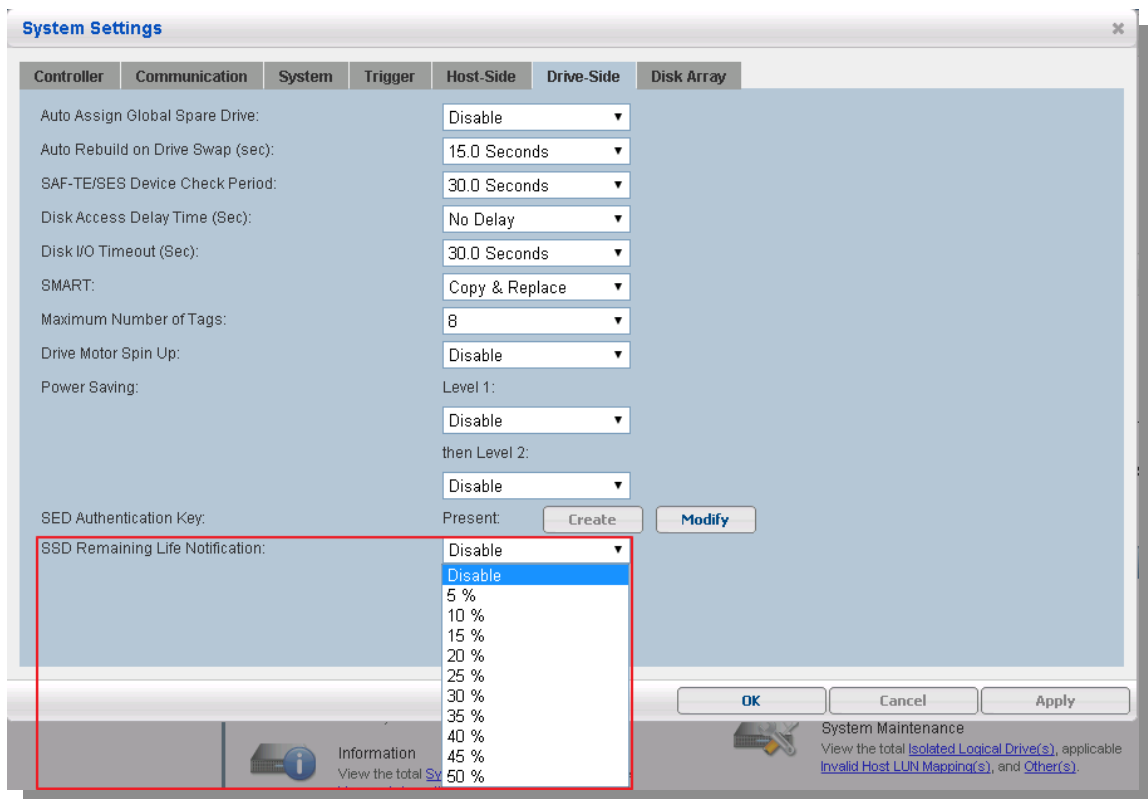
endurance characteristics from SSD manufacturers to notify users and minimize the impact of the limited write life cycle of SSDs.

Wear leveling applied in NAND flash is a technique developed by SSD vendors to extend SSD life cycle by distributing the use among the blocks. NAND flash memory is susceptible to wear due to repeated write and erase cycles that are common in everyday business applications. In order to prevent continuous writes and erases in the same block, wear leveling works continuously to re-allocate SSD's logical blocks to different pages and evenly distribute wear to each block to prolong the overall endurance and read/write life cycles.

# SSD Remaining Life Notification

The user can define a threshold for the remaining life of SSD. When this threshold is reached, the system will notify the user and actions such as data replication, data migration or SSD replacement can be taken. The threshold is configurable to Disable/5%/10%/20%/25%/30%/35%/40%/45%/50% and 10% is the default. The user will get an event notification which will be also logged when any SSD's remaining life reaches the user-defined threshold.

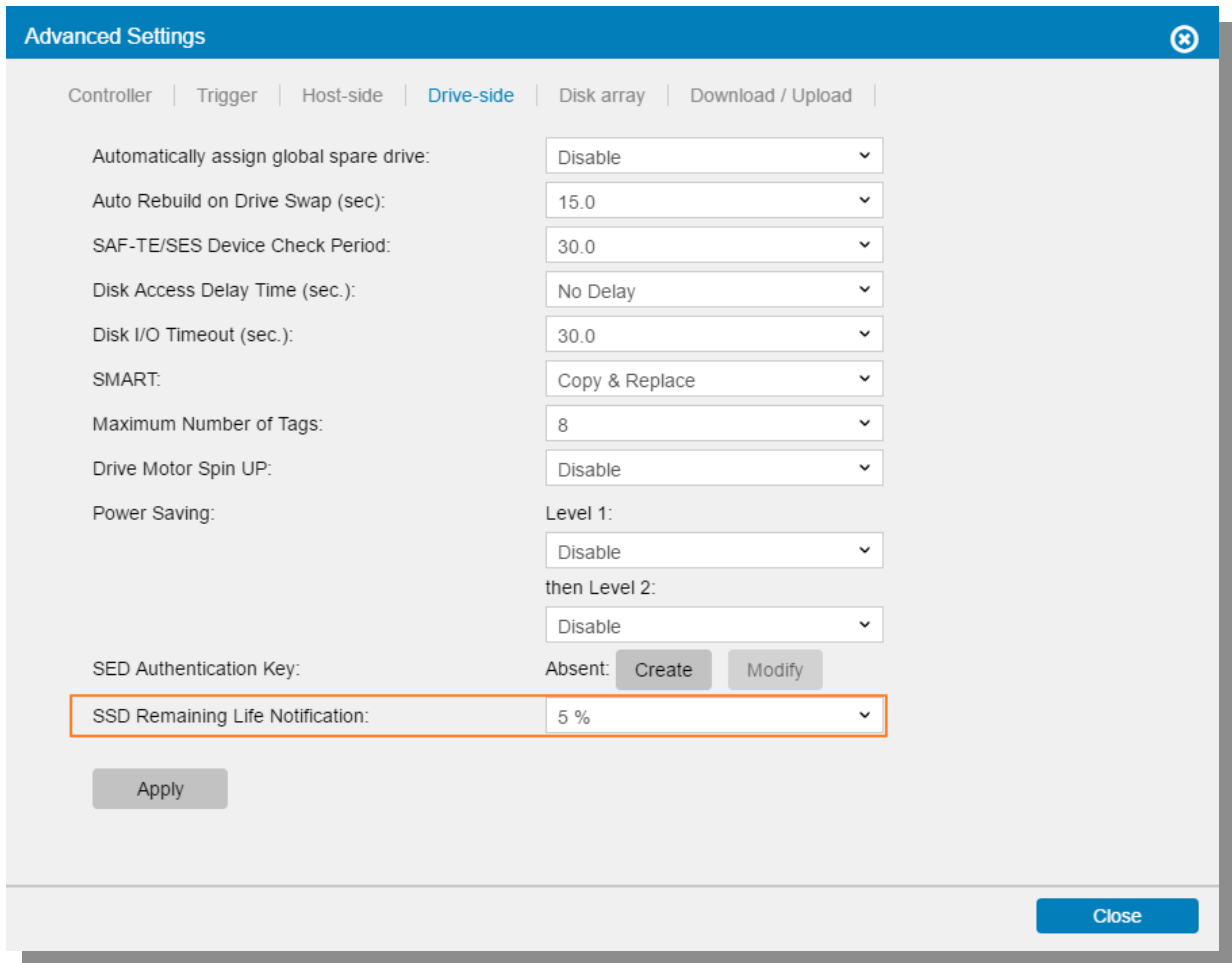
SANWatch for HA3969: Go to System Settings -> Drive Side -> SSD Remaining life Notification



*User defined notification threshold of SSD remaining life*



EonOne for HA3969: Go to Settings → System → General → Advanced Options → Manage → Drive-side



*User defined notification threshold of SSD remaining life*

## Conclusion

Storageflex SSD Cache is an intelligent solution that dramatically improves read performance for read-intensive applications, greatly reduces latency and supports large capacity cache pool. It is simple to set up, manage and maintain on Storageflex’s intuitive GUIs. It is highly recommended for use with system application workload that involves high-volume read operations and intense and repetitive small-volume read actions.

With monitoring features that warn users when an SSD has reached its lifespan limit, users can effectively employ the power of SSD while being well prepared for its limitation and confident that their valuable data is always well protected with actions such as data replication, data migration or SSD replacement.