

QNAP



QNAP NAS Solution

QTS SSD Extra Over-Provisioning

White Paper

Table of Contents

01	Summary	3
02	SSD Storage and Challenges	3
03	Understanding SSD Over-Provisioning	4
	3.1. SSD Operation Theory	4
	3.2. SSD Over-Provisioning Space	5
	3.3. Over-Provisioning and Write Amplification	6
04	QTS SSD Extra Over-Provisioning	7
	4.1. Supported Hardware and RAID Configurations	7
	4.2. Using SSD Extra Over-Provisioning	7
	4.3. SSD Extra Over-Provisioning Benefits	8
	4.3.1. Test 1: Extra Over-Provisioning With a Single Drive	8
	4.3.2. Test 2: Extra Over-Provisioning with RAID	9
	4.4. SSD Extra Over-Provisioning Limitations	9
05	SSD Profiling Tool	10
	5.1. Requirements	10
	5.2. Using SSD Profiling Tool	10
	5.3. Reference Figures	11
06	Case Study: Using SSD Extra Over-Provisioning to Improve Performance on a NAS Used for Video Editing	11
	6.1. Background	11
	6.2. Problem	11
	6.3. Solution	13
	6.4. Conclusion	13

1. Summary

Solid state drives (SSDs) store data persistently in solid state flash memory. Compared to mechanical hard disk drives (HDDs), SSDs provide increased I/O performance, reliability and energy efficiency. The popularity of SSDs for both home and enterprise users is increasing as their price per gigabyte has fallen by up to 70% between 2013 and 2018.

There are two challenges for IT professionals purchasing SSDs. The first is how to configure them to maximize performance. The second is how to prevent performance degradation over time. SSD performance degrades because they must perform garbage collection when full to free up blocks. Garbage collection requires a large number of write commands, meaning that one write to a full SSD can lead to several actual writes to the SSD's flash memory. This phenomenon is known as write amplification. Write amplification reduces an SSD's write performance and lifespan.

Over-provisioning is a solution to write amplification. Over-provisioning is reserved space that the SSD uses to write incoming data and for performing garbage collection. SSDs reserve 7.37% of their storage space for over-provisioning by default. Some SSD manufacturers also reserve additional space. However, reserving more space can help improve random write performance and lifespan. Studies show that an additional 28% of reserved space decreases write amplification by more than 50% and doubles the SSD's endurance.

QTS SSD Extra Over-Provisioning is a complete solution designed by QNAP. This new feature enables NAS administrators to configure and monitor additional SSD reserved space using the QTS utilities SSD Profiling Tool and Storage & Snapshots. SSD Profiling Tool enables you to determine the optimum amount of over-provisioning for your SSD configuration. Storage & Snapshots allows you to reserve 1-60% of space for over-provisioning when creating an SSD RAID group.

Our lab tests show that SSD Extra Over-Provisioning can dramatically improve consistent random write performance. Setting extra over-provisioning to 40% on a single Samsung 850 SSD increased the average consistent random write performance by 100%, from 25,000IOPS to 50,000IOPS. For a RAID 1 group containing two Micron 512GB SSDs, setting extra over-provisioning to 20% increased consistent write performance from 15,000IOPS to 36,000IOPS. A TV station case study also demonstrates how SSD Extra Over-Provisioning provided benefits in a real production environment.

QTS now has several dedicated features and utilities, including SSD Extra Over-Provisioning, SSD Profiling Tool, SSD Cache Acceleration, and Qtier™ Auto Tiering. This makes buying a QNAP NAS device the first choice for maximizing your SSD storage.

2. SSD Storage and Challenges

A solid state drive (SSD) stores data persistently in solid state flash memory. Compared to mechanical hard disk drives (HDDs), SSDs provide better I/O performance, reliability and energy efficiency. While SSDs are still more expensive than HDDs, prices are falling every year. A 2018 report published by Sandisk (<https://www.sandisk.com/business/datacenter/resources/white-papers/the-ssd-enabled-pc-total-cost-of-ownership>) shows that the price per GB of a 256GB SSD fell 70% between 2013 and 2018, making it the same price per GB as a 500GB HDD (Figure 1).




Figure 1: SanDisk Statistics for Global SSD/HDD Price Trends

The biggest challenges for IT professionals purchasing solid state drives is how to configure the SSDs to ensure maximum storage efficiency and return on investment over the lifetime of the drives.

- SSD performance can vary between similar drives using the same interface. These variations are due to several factors such as type of NAND flash (MLC, SLC, TLC), device architecture, design optimization, and the environment in which they are used.
- SSD performance can degrade over time.
- SSDs have a limited lifespan.

3. Understanding SSD Over-Provisioning

3.1. SSD Operation Theory

Solid state drives use an array of semiconductor memory and transistors to store information without relying on moving parts and spinning disks.

In order to store more data, flash memory cells are arranged in blocks, which are subdivided into pages. Each block contains 128 to 256 pages, and each page stores 2KB to 16KB of data (Figure 2).

An SSD reads and writes data at the page level. Before a page containing invalid data can be overwritten, its cells voltage threshold must be adjusted. This operation requires a high amount of voltage, which means the whole block will be erased during the operation.

Whenever an SSD is full with blocks contain both valid and invalid pages in a random pattern, the SSD must perform garbage collection to free up blocks. During garbage collection the SSD copies all valid pages from a block to reserved storage space, and then erases the whole block so that it can store new data. This lead to an undesirable effect called write amplification, which mean that one write to an SSD might cause garbage collection, requiring many actual writes to the SSD's flash memory.

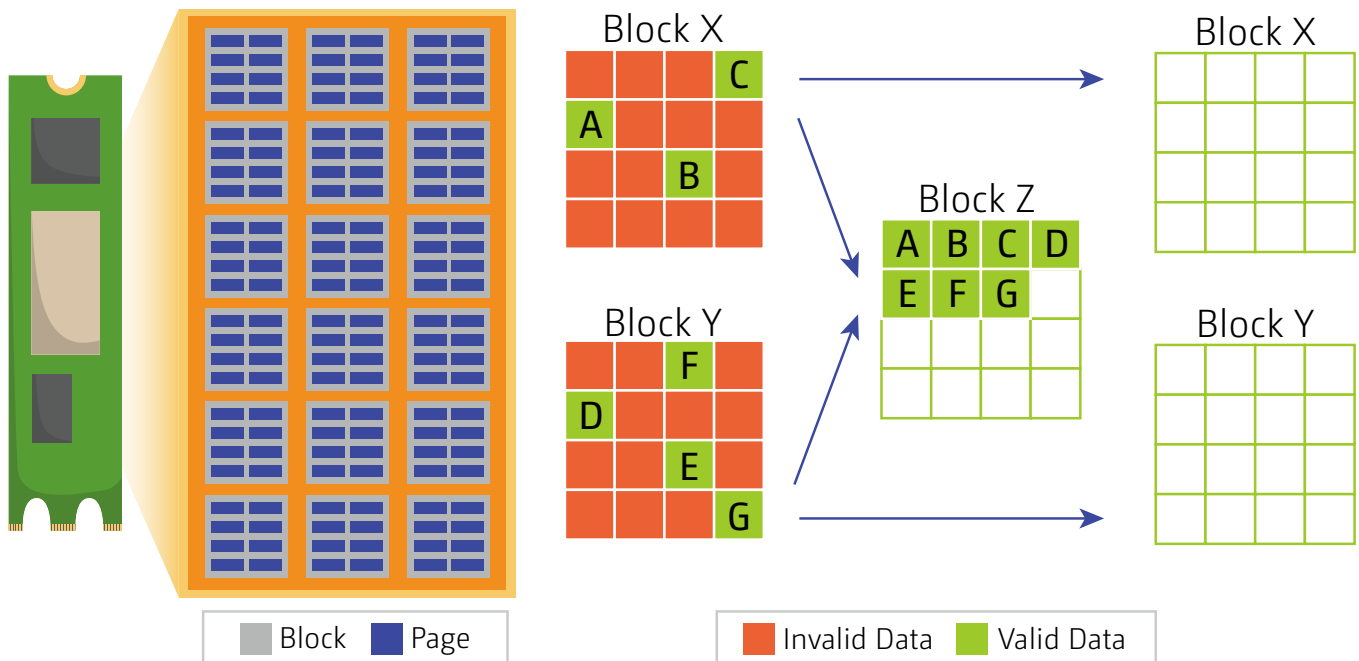


Figure 2: SSD Blocks and Pages

Write amplification affects both the write performance and the lifespan of the SSD. Each erase causes a small amount of physical damage to an SSD block, which means each block can only be erased and filled a limited number of times. This number is known as P/E (Program Erase) cycles. Write amplification increases the number of P/E cycles, which reduces drive endurance and can cause a specific block to wear out more quickly than other blocks.

SSD manufacturers have continually improved the architecture and design of their drives to maintain consistent SSD performance. For example, most SSDs have a command called TRIM which enables them to perform garbage collection more efficiently to free up blocks before the drive is completely full. However, a certain amount of reserved space is still needed for garbage collection and to control wear leveling. This reserved space is called over-provisioning space.

Understand more:

https://en.wikipedia.org/wiki/Solid-state_drive

https://en.wikipedia.org/wiki/Write_amplification

<https://www.extremetech.com/extreme/210492-extremetech-explains-how-do-ssds-work>

3.2. SSD Over-Provisioning Space

Like HDDs, SSDs use two different systems for representing storage capacity:

- Capacity seen by computers: 1 GiB = 2^{30} = 1,073,741,824 bytes
- Capacity displayed to users: 1 GB = 10^9 = 1,000,000,000 bytes

The difference between these two values means that every 1GB of SSD storage contains 73,741,824 bytes, or about 7.37% of the total capacity, which is free to be used for over-provisioning.

Certain SSD manufacturers reserve additional over-provisioning space to ensure that the drives can handle periods of high burst writes when the SSD is under heavy load.

3.3. Over-Provisioning and Write Amplification

Figure 3 show a simplified SSD with 4 blocks.

1. The system wants to write 6 new blocks. There are 6 invalid pages, but the SSD must perform garbage collection before they are usable. The system must wait.
2. The SSD performs garbage collection. It copies 4 valid pages to a free block (4 writes), erases the two freed up blocks, and writes the new data to them (6 writes).

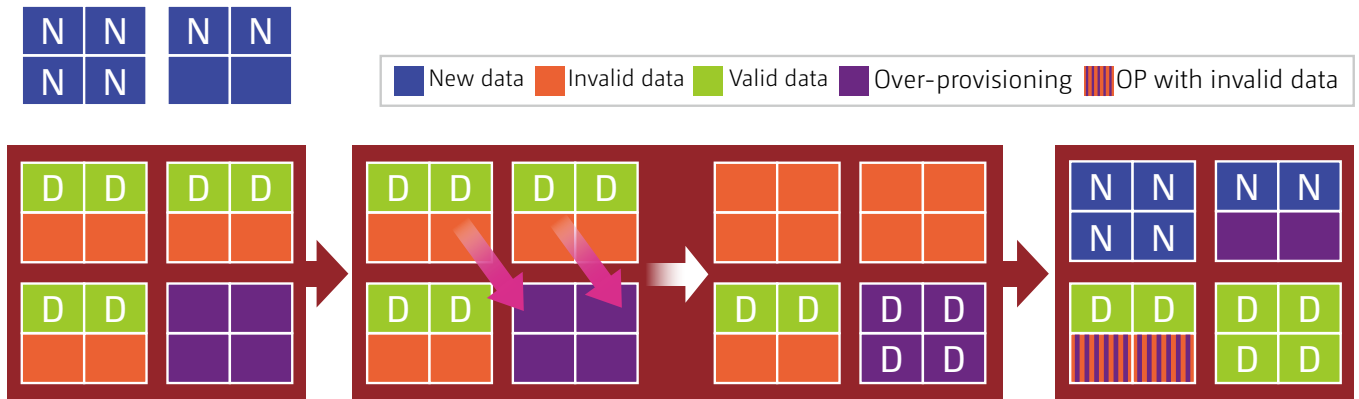


Figure 3: A Simple SSD with 25% Over-Provisioning Space

In this example, 6 new pages of data results in 10 total writes, giving a write amplification ratio of 1.66.

Figure 4 shows the same SSD with additional over-provisioning space.

1. The system wants to write 6 new blocks. There are 6 invalid pages, but the SSD must perform garbage collection before they are usable. However, this time there are 2 blocks reserved for over-provisioning, containing 8 pages.
2. Instead of performing garbage collection to free up blocks, the new data is written immediately into over-Provisioning space, giving a write amplification ratio of 1.
3. The SSD does not need to perform garbage collection until new data is written to the SSD or the SSD performs a TRIM.

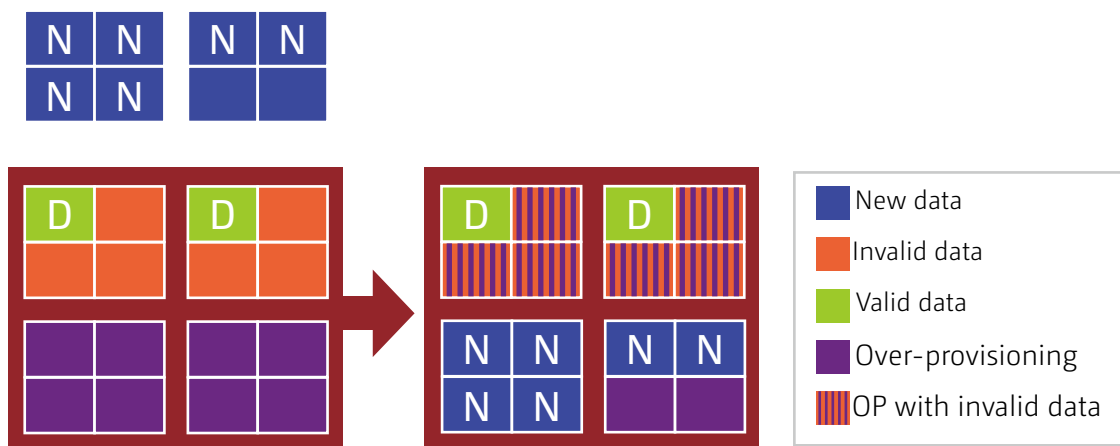


Figure 4: A Simple SSD with 25% Over-Provisioning Space

In a study by flash memory manufacturer Atpinc (<http://www.atpinc.com/technology/ssd-over-provisioning-benefits-configure>), adding an additional 28% of reserved space:

- Decreased the write amplification ratio from 4.7 to 1.9.
- Doubled the SSD's endurance (measured in Terabytes Written).

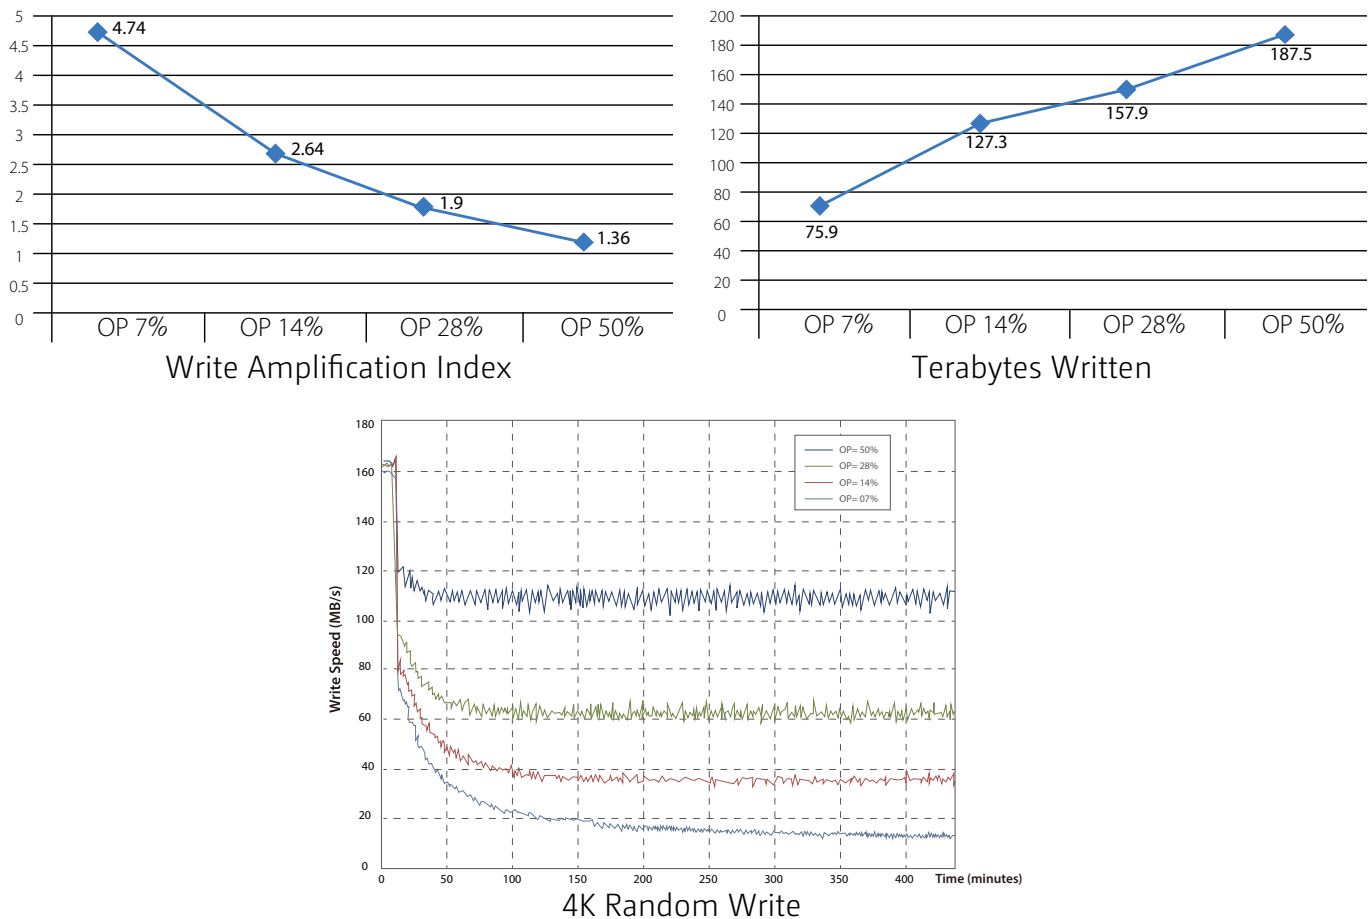


Figure 5: An Atpinc Study Showing How Increasing Over-Provisioning Decreases Write Amplification and Increases Endurance

4. QTS SSD Extra Over-Provisioning

SSD Extra Over-Provisioning is a complete software-defined over-provisioning solution for QTS. You can now reserve additional over-provisioning space when creating SSD RAID groups in Storage & Snapshots, and then easily manage this space. The new QTS utility SSD Profiling Tool helps to determine the optimal amount of SSD over-provisioning for your specific SSD models and configuration.

4.1. Supported Hardware and RAID Configurations

- Available on all QNAP NAS models running QTS 4.3.5 or later.
- Support for all SSD form-factors and interfaces including SATA, SAS, M.2, and PCIe NVMe.
- Available when creating SSD cache, all-SSD storage pools, Qtier™ auto-tiering storage pool, all-SSD static volumes.
- Supported for the following RAID types:
 - SSD cache: Single, JBOD, RAID 0, RAID 1, RAID 5, RAID 6, RAID 10, RAID 50, RAID 60.
 - Storage pools, Qtier pools, static volumes: Single, RAID 1, RAID 5, RAID 6, RAID 10, RAID 50, RAID 60.

4.2. Using SSD Extra Over-Provisioning

You can set an extra over-provisioning value of 1% to 60% when creating an SSD RAID group in Storage & Snapshots. QTS then uses this space to create a partition at the RAID level when creating the RAID group. This feature cannot be enabled on existing storage pools or static volumes, and after creation the amount of SSD Extra Over-provisioning can be decreased but not increased.

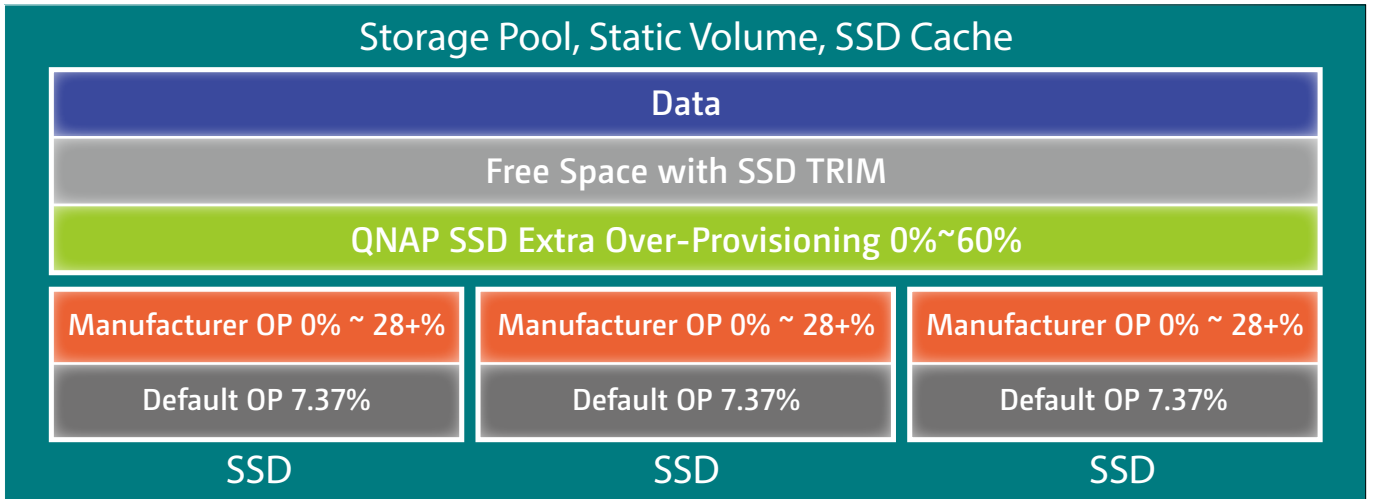


Figure 6: RAID Level Partition with SSD Extra Over-Provisioning

4.3. SSD Extra Over-Provisioning Benefits

SSD Extra Over-provisioning can enhance the performance of both enterprise-level SSDs and consumer-level SSDs, specifically TLC (Triple-level cell) drives which provide higher SSD capacities but with limited endurance and consistent performance.

4.3.1. Test 1: Extra Over-Provisioning With a Single Drive

QNAP labs selected one Samsung 850 Pro 512GB SSDs to represent consumer-level hardware. The drive was filled with data and then tested using FIO configured with the following settings:

runtime=1800, ioengine=libaio, direct=1, rw=randwrite, bs=4k, numjobs=1, iodepth=32

The FIO test was run with extra over-provisioning set to 0%, 20% and 40%.

Results

For full results, see Figure 7. With extra over-provisioning set to 20%, the Samsung 850 Pro gave an average consistent random write speed of 25,000 IOPS. With extra over-provisioning set to 40%, average consistent random write speed increased to 50,000 IOPS. These figures show that a consumer-level Samsung drive can reach a similar level of performance as an enterprise Intel drive: https://ark.intel.com/products/120525/Intel-SSD-DC-S4500-Series-960GB-2_5in-SATA-6Gbs-3D1-TLC




Figure 7: Samsung 850 SSDs Tested with Extra Over-Provisioning of 0%, 20%, 40%

4.3.2. Test 2: Extra Over-Provisioning with RAID

This test illustrates the benefits of extra over-provisioning with RAID. A QNAP NAS model TVS 882 NAS was configured with 4 HDDs and 2 Micron 1100 512GB SSDs, configured as a RAID 1 read-write cache. System was tested using FIO with the following settings:

```
runtime=180, ramp_time=30, size=32GB, ioengine=libaio, direct=0, rw=randwrite, bs=4k numjobs=64, iodepth=1
```

Over-provisioning varied during the test from 0% to 60%.

Results

Figure 8 shows the final average IOPS for each amount of over-provisioning. With extra over-provisioning set to 20%, SSD Cache constant random write performance increases more than 100% from 15000 IOPS to 36000 IOPS. With extra over-provisioning set to 50%, average consistent write performance reaches 50,000IOPS.

Consider the case of a NAS used in a company with few employees. As the company grows, the SSDs in the NAS will be under more load. With no extra over-provisioning, SSD cache performance can degrade by more than 75% as cache usage increases. SSD Extra Over-Provisioning helps reduce this degradation and maximize the SSD's return on investment (ROI).




Figure 8: Two Micron 110 SSDs in a RAID 1 Read-Write Cache Tested With Over-Provisioning of 0-60%

Item	Consumer SSD with Default 7% Over-Provisioning	Consumer SSD with QNAP Extra Over-Provisioning	Enterprise SSD with Default Manufacturer Over-Provisioning	Enterprise SSD with QNAP Extra Over-Provisioning
Consistent Random Write Performance	Low	Medium ↑	Medium	High ↑
Cost index (TBW)	High	Medium ↓	Medium	Low ↓

4.4. SSD Extra Over-Provisioning Limitations

There are some hardware conditions and limitations to consider when using QTS Extra Over-Provisioning.

Access patterns

Over-provisioning space improves SSD consistent performance and endurance for random access only. SSD Extra Over-Provisioning will probably provide limited to no benefits in the following situations:

- The NAS is mainly used for sequential access (for example, a home NAS with a single user).
- NAS storage performance is limited by other bottlenecks, such as network throughput.
- The NAS is not used for higher performance applications.

SSD Hardware and Controller

QTS Extra Over-Provisioning has been tested and shown to provide benefits with several major brands of SSDs, including Intel, Kingston, Micron, Samsung, and Sandisk(WD). However, some SSD controllers might not be able to recognize and utilize the extra over-provisioning space. Also certain new types of SSDs, such as the Intel® Optane™ series, do not require over-provisioning space to maintain consistent performance.

SSD Capacity

QTS Extra Over-Provisioning increases random write performance only. If the SSDs will be used for mainly read operations, for example a read cache, then you can limit the amount of the extra over-provisioning. Refer back to the SSD read cache test in Figure 8. The test showed no improvement in read performance when using different amounts of extra over-provisioning amounts. This means that you should not enable extra over-provisioning in this case.

You can use SSD Profiling Tool to ensure that your SSDs support and will benefit from QTS Extra Over-Provisioning.

5. SSD Profiling Tool

SSD Profiling Tool is a new utility available in QTS 4.3.5 and later. It enables you to identify whether your SSDs support QTS Extra Over-Provisioning, and also helps to determine the optimum amount of over-provisioning for your SSD configuration. SSD Profiling Tool can test a single SSD or multiple SSDs in a RAID group.

5.1. Requirements

- NAS must be running QTS 4.3.5 or later
- At least one free SSD

5.2. Using SSD Profiling Tool

SSD Profiling Tool first measures the initial performance of the SSDs. Then it fills the disks with random data and tests consistent random access with different amounts of over-provisioning. Performance is measured using FIO configured with "Data Blocks: 4KB" and "Queue Depth: 32".

Each incremental test with a different amount of over-provisioning is called a phase. SSD Profiling Tool will skip a phase if it detects that the average random write speed is too low.




Figure 9 shows an example SSD Profiling Tool test report. The report shows that performance increased and time required to finish decreased as the amount of over-provisioning was increased. This shows that the test SSD was able to utilize the extra over-provisioning space. If your application requires a specific minimum IOPS, the report also enables you to determine the optimal amount of over-provisioning required to reach it.

Tip: To best utilize SSD Profiling Tool, use a performance monitoring application such as QTS Resource Monitor or Q'center to identify the random IOPS requirements of an application before running a test.

Note:

- The test results provided by SSD Profiling Tool are specific to the NAS being tested. SSD Profiling Tool should not be used as a general drive benchmarking tool.
- SSD Profiling Tool tests might run for several days and may reduce the SSD's endurance. You can reduce test time by using a lower test data size and reducing the over-provisioning range.

5.3. Reference Figures

You should always use SSD Profiling Tool on your NAS to determine the most appropriate amount of extra over-provisioning for your specific hardware configuration. However, if it is not possible to run a test in your environment, then you can use the reference configuration settings below.

Use Case	SSD Configuration	SSD RAID Type	Extra Over-Provisioning
File Server	Qtier™	RAID 10	10%
Web Server, Application Server, Virtualization Server	Qtier™	RAID 10	20%
Video Editing	Read-Write Cache	RAID 5	20%
Backup/Log Server	Write-Only Cache	RAID 10	30%
Business-Critical Database	All-SSD Storage Pool or Static Volume	RAID 10	30%

6. Case Study: Using SSD Extra Over-Provisioning to Improve Performance on a NAS Used for Video Editing

6.1 Background

Parkmedia Audiovisual, a television station in Spain, is using a QNAP TS-EC2480U R2 NAS as their production storage server. The NAS is used by video recorders and editors for online editing and storing recorded media. To satisfy their high performance requirements, the NAS contains a read-write cache consisting of 4 Samsung 850 SSDs and is connected to a 10GbE switch.

6.2. Problem

Editors reported that while new files are being uploaded to the NAS, their Adobe Premiere Pro CC editing jobs showed delays when previewing videos.

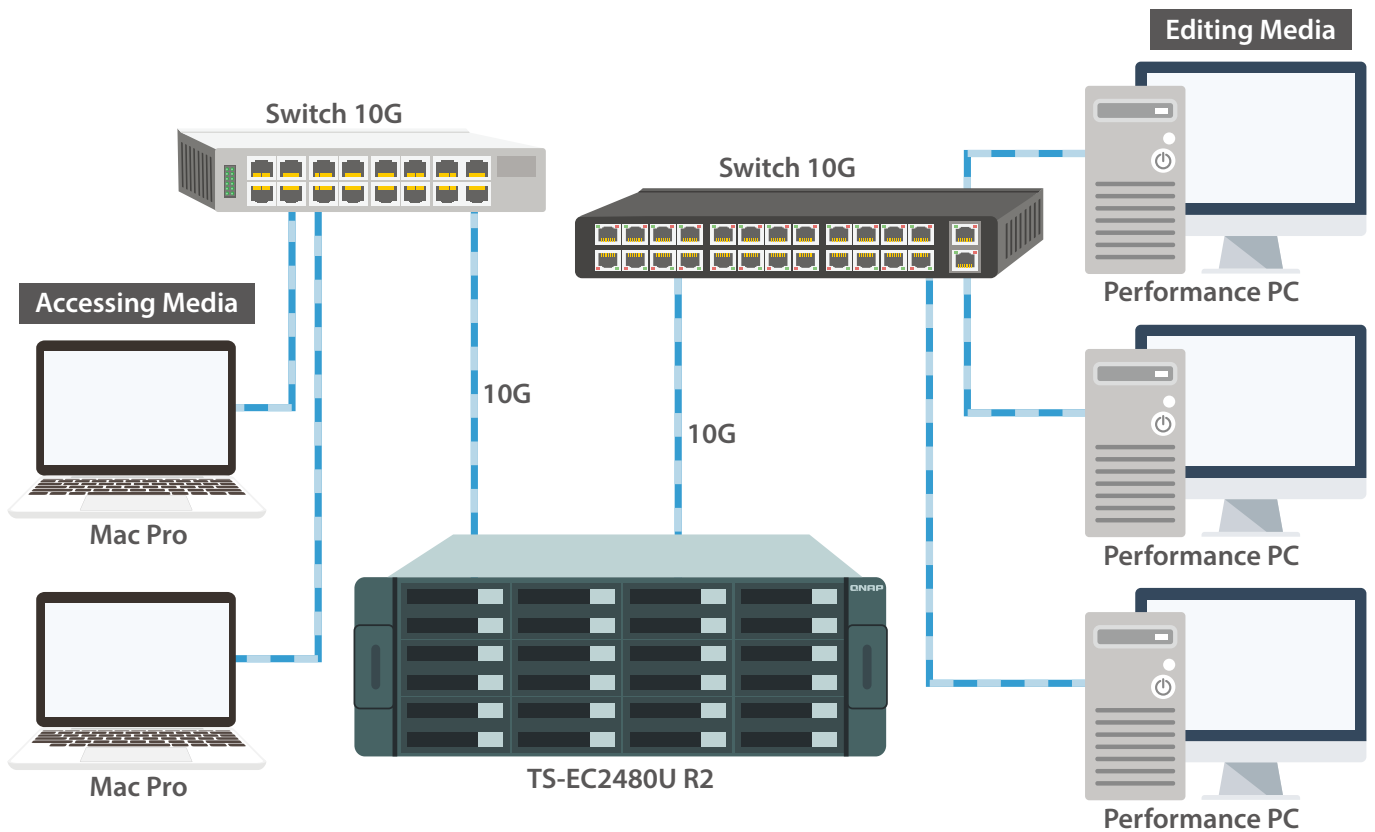


Figure 10 :QNAP Lab's Reproduction of the Hardware and Network Environment at Parkmedia Audiovisual

The Parkmedia Audiovisual IT manager contacted QNAP and reported the problems their editors were having. QNAP reproduced the hardware and network configuration of Parkmedia Audiovisual in a lab environment, and then tested the hardware in similar usage scenarios.

The tests found that when running an editing job on the NAS with a throughput of 300MBps, a separate upload operation caused the editing job's read performance to drop. This problem occurred even when the two jobs were run on different network switches.

Further investigation showed that the problem was caused by existing data and newly written data competing to utilize the SSD cache. Large file operations are generally performed sequentially. However, if multiple clients perform sequential write operations simultaneously, the data saved to the SSD cache becomes scattered and is stored in a random pattern. When the cache is full, the SSDs must continually perform garbage collection to re-allocate blocks causing read and write delays.

Figure 11 shows the sent and received bandwidth of two network interfaces on the test NAS during the simulated editing and uploading test. At around 331 seconds, the SSD cache became full and bandwidth dropped.




Figure 11: Bandwidth of two network interfaces during upload and online editing tests. No extra over-provisioning.

6.3. Solution

QNAP ran the tests again with differing amounts of software-defined over-provisioning (Figure 12). With over-provisioning set to 40%, sent (read) throughput for the editing job remained stable at 300 MBps. Received (write) throughput increased from 450MBps to 750MBps, a 66% improvement.




Figure 12: Bandwidth of two network interfaces during upload and online editing tests with 40% extra over-provisioning enabled on the SSD cache.

After the successful lab tests, SSD Extra Over-Provisioning was enabled on the QNAP TS-EC2480U R2 production NAS at Parkmedia Audiovisual. Editors reported that they no longer encountered delayed during video editing work. Their feedback can be viewed at the following link: <https://www.youtube.com/watch?v=bG4KtvGTO-w&t=103s>

6.4. Conclusion

This case demonstrates that SSD Extra Over-Provisioning can improve SSD performance in a real production environment. It also shows that evaluating, deploying, and maintaining an SSD storage solution can be a challenging task.

QNAP's complete solution of SSD Profiling Tool and SSD Extra Over-Provisioning helps IT managers to evaluate the performance of SSDs before production deployment, avoiding unforeseen performance bottlenecks with no additional cost. As SSDs become cheaper and more ubiquitous, we believe buying a QNAP NAS device is the first choice for maximizing your SSD storage.

QNAP SYSTEMS, INC.

TEL : +886-2-2641-2000 FAX: +886-2-2641-0555 Email: qnapsales@qnap.com

Address : 3F, No.22, Zhongxing Rd., Xizhi Dist., New Taipei City, 221, Taiwan

QNAP may make changes to specification and product descriptions at any time, without notice.

Copyright © 2018 QNAP Systems, Inc. All rights reserved.

QNAP® and other names of QNAP Products are proprietary marks or registered trademarks of QNAP Systems, Inc.

Other products and company names mentioned herein are trademarks of their respective holders.

Netherlands (Warehouse Services)

Email: nlsales@qnap.com

TEL: +31(0)107600830

China

Email: cnsales@qnap.com

TEL: +86-400-028-0079

Thailand

Email: thsales@qnap.com

TEL: +66-2-5415988

Japan

Email: jpsales@qnap.com

FAX: 03-6435-9686

US

Email: usasales@qnap.com

TEL: +1-909-595-2782

India

Email: indiasales@qnap.com

Germany

Email: desales@qnap.com

France

Email: frsales@qnap.com