

TOSHIBA



Enterprise Hard Drives

Designed for your business

Replacement of Enterprise
Hard Drives in legacy systems –
what are the key issues and
how much flexibility is available?

Replacement of Enterprise Hard Drives in
legacy systems – key issues and how much
flexibility is available?

Introduction

Today, enterprise hard drives are highly reliable. For example, Toshiba's MG06 series has increased its statistical reliability – based on Mean Time to Failure (MTTF) – from 2 million hours to 2.5 million hours. Statistically, in real systems and during the five-year warranty period, this means that in an installation of 1000 hard drives only two to three disks will fail each year instead of four to five as previously.

Therefore, much longer life cycles are possible compared to desktop hard drives. Despite the higher reliability, IT operators need to anticipate the occasional failure of individual disks in storage systems and take appropriate measures for their replacement.

This causes no problems as long as the originally used drive model is available on the market. A faulty drive is then usually replaced by a new drive or replaced by the manufacturer with an RMA-equivalent model. If the production of a model is discontinued and is no longer available, the customer relies on the manufacturer or the retail market to acquire the obsolete product/model. But eventually, this source will dry out too. Since lifecycles will become longer due to the higher reliability of the drives, today's HDD users may reach this point more often.

If a drive is replaced by the next generation, models of the old series not being covered by a 1:1 replacement scheme within the new generation will be produced as long as there is a significant demand. This is especially true for smaller storage capacities within a series. These models are then available with the extended lifecycle of the old series or as next generation "optimized low-end" models.

Ultimately, this support also comes to an end at some time and the older models are no longer available. Now the customer is faced with the question: can I replace my old model with a new but different model? To help answer this question, failure-replacement scenarios were examined in different environments.

In areas where hard drives or SSDs are "software-defined" managed as storage capacities, different models can usually be used in parallel. Here it is important to ensure that the interface (usually a host bus adapter HBA) supports the new (replacement) model.

With RAID-based storage systems, replacement is more critical because RAID controllers aggregate a number of similar physical drives into one logical drive, assuming identical physical components. This was the focus of Toshiba's research.

The most common RAID controllers being used are Broadcom's LSI "MegaRAID SAS 9380" and the new Microware Adaptec "SmartRAID 3154", which are considered functionally compatible with Adaptec's "Series7" and "Series8" models.

Firstly, the following question arises: can a failed drive from a RAID set of HDDs with 512 bytes of "native" block size (512n) be replaced by a successor with 512 bytes of "emulated" block size (512e)? New generations typically use the so-called "advanced format" with a disk-internal block length of 4k bytes. If a 4k byte block length is expected at the interface, it is called "4kn". However, since this format has not yet widely established itself in the enterprise area, the interface is usually emulated for enterprise hard drives by a 512 byte block length (512e).

In each test a RAID5 configuration was used consisting of four identical hard drives connected to a Broadcom LSI or a Microsemi Adaptec RAID controller. One hard drive was then replaced and checked as to whether this drive had been accepted by the controller, whether a rebuild had been successfully completed and system performance is still maintained.

In test one, a 3TB 512n SATA HDD (MG03ACA300) was replaced by an equally sized 512e model (MG04ACA300E). Both RAID controllers were able to accept the new hard drive, perform a rebuild and no performance losses were detected.

Test 1	Interface	Size	Block Size	Type	Model	MegaRAID SAS 9380	SmartRAID 3154
Original	SATA	3TB	512n	HDD	MG03ACA300	OK	OK
Replacement	SATA	3TB	512e	HDD	MG04ACA300E		

With the change from generation MG03 (512 bytes native block size) to MG04 (512 bytes emulated block size) Toshiba launched its common models with 1, 2 and 4TB storage capacity in a special 512n variant in addition to the mainstream format (512e). Therefore, there should be no need to replace a 512n with a 512e drive. However, it is good to know that this is possible without restriction, such as in cases where there are problems with the availability of 512n drives.

But if the drive with its original capacity is no longer available the question arises, can a compatible drive with a higher capacity be used? In test two a 1TB 512n SATA hard drive was replaced by a drive with double the capacity and the equivalent, but emulated, block size.

Test 2	Interface	Size	Block Size	Type	Model	MegaRAID SAS 9380	SmartRAID 3154
Original	SATA	1TB	512n	HDD	MG03ACA100	OK	OK
Replacement	SATA	2TB	512e	HDD	MG04ACA200E		

This replacement also caused no problems. However, only a capacity of 1TB is used from the new 2TB disk.

For systems and drives with SAS interfaces, many current installations are still based on the SAS 6Gbit/s interface speed. New drives (and all new controllers) consistently support 12Gbit/s - with purported backwards compatibility to 6Gbit/s SAS. The open question is, in a RAID set of 6Gbit/s SAS HDDs, could a failed drive be replaced by a 12Gbit/s SAS drive?

Test 3	Interface	Size	Block Size	Type	Model	MegaRAID SAS 9380	SmartRAID 3154
Original	SAS 6Gbit/s	600GB	512n	HDD	AL13SEB600	OK	OK
Replacement	SAS 12Gbit/s	1.8TB	512e	HDD	AL14SEB18EQ		

Test three, based on a 10krpm 2.5” Enterprise HDD shows that with SAS drives not only the interface is compatible but also a differing block size definition (512n vs 512e) and a larger disk can also be deployed as a replacement. In test four, the 7200rpm 3.5” SAS model (Enterprise Capacity “Nearline”) class from Toshiba was also tested and returned positive results.

Test 4	Interface	Size	Block Size	Type	Model	MegaRAID SAS 9380	SmartRAID 3154
Original	SAS 12Gbit/s	4TB	512n	HDD	MG04SCA40EN	OK	OK
Replacement	SAS 12Gbit/s	6TB	512e	HDD	MG04SCA60EE		

One question remains: can a failed drive be replaced by a “randomly available” HDD with a 4kn block format in a RAID using drives with a block length of 512bytes? Unfortunately, the fifth test shows that 512byte and 4kByte block formats cannot be combined in a RAID configuration.

Test 5	Interface	Size	Block Size	Type	Model	MegaRAID SAS 9380	SmartRAID 3154
Original	SAS 12Gbit/s	2TB	512e	HDD	MG04SCA20EE	X	X
Replacement	SAS 12Gbit/s	2TB	4kn	HDD	MG04SCA20EA		

With “software-defined” storage the operator should consult the software manufacturer and check if 4kn is supported.

Test five raised another question: how far can formats, interfaces, type, block size and capacity be mixed in general? To answer this, another combination was examined in test six: Can a failed SATA hard drive be replaced by an SAS hard drive?

These tests answer the question of where the boundaries of limitation really are. It should be noted that in real life settings and strict product environments, these kind of configurations are not recommended.

Test 6	Interface	Size	Block Size	Type	Model	MegaRAID SAS 9380	SmartRAID 3154
Original	SATA	1TB	512n	HDD	MG03ACA100	OK	X
Replacement	SAS 12Gbit/s	1.8TB	512e	HDD	AL14SEB18EQ		

The results show that the MegaRAID controller uses all available components independent of the interface. The operator could set up a RAID with a mix of SAS and SATA, and if a failed SATA drive is replaced by an SAS model, the controller can restore the RAID set. The SmartRAID is much more selective: neither is it possible to build a RAID set from a mix of SAS and SATA drives, nor can individual drives be exchanged for models with a differing interface. Which strategy makes more sense remains to be seen - the MegaRAID allows maximum flexibility while the SmartRAID prohibits configurations (e.g. SAS/SATA) that are not generally found in practice or are not recommended.

Another configuration is also conceivable, which was tested in test seven. Can a failed drive from a RAID set be replaced by an SSD drive?

Test 7	Interface	Size	Block Size	Type	Model	MegaRAID SAS 9380	SmartRAID 3154
Original	SATA	1TB	512n	HDD	MG03ACA100	OK	X
Replacement	SATA 6Gbit/s	1.6TB	-	eSSD	THNSN81Q60		

Similar results were found with an SAS hard drive in a SATA array. In MegaRAID, the SSD was accepted and a rebuild successfully completed while the SmartRAID was blocked by the unsuitable configuration. The result of test eight, based on a RAID built of SAS drives in which a failed drive is replaced by an SAS SSD, is surprising.

Test 8	Interface	Size	Block Size	Type	Model	MegaRAID SAS 9380	SmartRAID 3154
Original	SAS 6Gbit/s	600GB	512n	HDD	AL13SEB600	OK	X
Replacement	SAS 12Gbit/s	800GB	-	eSSD	PX04SMB080		

With MegaRAID it would even be possible to install an SAS SSD in a RAID set of SATA hard drives or vice versa. But this experimental setup remains more of academic interest.

In practice it is interesting to see whether failed hard drives in storage systems can be replaced by models from another manufacturer. For this purpose, test nine determined whether a RAID set of 2TB SATA hard drives from a well-known third-party manufacturer of enterprise HDDs accepted Toshiba's corresponding hard drives as a replacement for failed disks. Here both RAID controllers could handle the replacement with ease.

Test 9	Interface	Size	Block Size	Type	Model	MegaRAID SAS 9380	SmartRAID 3154
Original	SATA	2TB	512n	HDD	Manufacturer 'X'	OK	OK
Replacement	SATA	2TB	512e	HDD	MG04ACA200E		

After almost all possible combinations had been tried, there is still one question left to be answered. Do the controllers notice if a replacement HDD or SSD with a smaller capacity is installed? Test 10 shows that both controllers detect the installation of the lower capacity drive and refuse to restore the RAID set.

Test 10	Interface	Size	Block Size	Type	Model	MegaRAID SAS 9380	SmartRAID 3154
Original	SATA 6Gbit/s	2TB	512e	HDD	MG04ACA200E	X	
Replacement	SATA 6Gbit/s	1TB	512n	HDD	MG04ACA100N		

Conclusion

There is some flexibility when replacing failed disk drives in a RAID set. You can replace 512n drives with 512e models as well as larger capacity or third-party models. Additionally, outdated SAS 6Gbit/s hard drives can be replaced with new models featuring a 12Gbit/s interface as the 12Gbit/s model is backwards-compatible with the 6Gbit/s standard. However, it should be noted that 512n/512e models are not compatible with 4kn models. In addition, certain RAID controllers even allow replacement of SATA drives with SAS models and replacement of hard drives with SSDs. However, just because these configurations are supported by the RAID controller this does not mean they should be implemented in practice.

Contact us for more information:

www.toshiba-storage.com/contact/



TOSHIBA

Toshiba Electronics Europe GmbH

Hansaallee 181
40549 Düsseldorf
Germany

info@toshiba-storage.com
toshiba-storage.com

Copyright © 2019 Toshiba Electronics Europe GmbH. All rights reserved. Product specifications, configurations, prices and component / options availability are all subject to change without notice. Product design, specifications and colours are subject to change without notice and may vary from those shown. Errors and omissions excepted.

05/2018