



Technical FAQ

Storage Subsystem

Mohammad Jawwad Memon, NetApp
Feb, 2014

FAQ OVERVIEW

This document addresses the most frequently asked questions about the NetApp® storage subsystem, which includes drives, storage shelves, shelf modules, shelf connectivity, HBAs, RAID, volumes, aggregates, and so on.

TABLE OF CONTENTS

1	Aggregates, raid groups, and volumes	3
1.1	General	3
1.2	32-bit aggregates	9
1.3	64-bit aggregates	10
2	Drives	12
2.1	General	12
2.2	Full Disk Encryption	16
2.3	Spares	17
2.4	SSD	19
2.5	SATA	23
2.6	Fibre Channel	24
2.7	SAS	25
3	Shelves and Shelf Modules	26
3.1	General	26
3.2	Shelf Intermixing	28
3.3	DS2246	30
3.4	DS4243	34
3.5	DS4246	36
3.6	DS4486	38
3.7	Alternate Control Path (acp)	39
3.8	DS14	40
4	Raid	42
5	Software-based Disk Ownership	44
6	Multipath High Availability (MPHA)	45
7	Syncmirror	47
8	Maintenance Center	49
9	Disk Sanitization	49

1 AGGREGATES, RAID GROUPS, AND VOLUMES

1.1 GENERAL

CAN DIFFERENT CAPACITY DISK DRIVES OF THE SAME DISK TYPE BE ADDED TO THE SAME AGGREGATE?

Yes. NetApp recommends adding disk drives into RAID groups of like-capacity disk drives in order to avoid disk downsizing.

CAN DIFFERENT CAPACITY DISK DRIVES OF THE SAME DISK TYPE BE COMBINED IN THE SAME RAID GROUP?

Yes. Disk drives of larger capacity than the smallest capacity disk drive in the existing RAID group will be downsized when added to an existing RAID group. If you are forming a new RAID group out of mixed capacity disk drives, then only disk drives selected to be the parity drives will not be downsized.

It is possible to add smaller capacity disk drives into an existing RAID group of larger capacity drives. The larger capacity drives in this case will not be downsized as they are already being used to store user data. This is a highly suboptimal configuration, and NetApp does not recommend adding smaller capacity drives into RAID groups with larger capacity drives.

SHOULD ALL AGGREGATES ALWAYS USE THE DEFAULT RAID GROUP SIZE?

The previous approach to RAID group and aggregate sizing was to use the default RAID group size. This no longer applies, because the breadth of storage configurations being addressed by NetApp products is more comprehensive than when the original sizing approach was determined. Sizing was also not such a big problem with only 32-bit aggregates, which are limited to 16TB. You can fit only so many drives and RAID groups into 16TB. The introduction of 64-bit aggregates delivers the capability for aggregates to contain a great number of drives and many more RAID groups than were possible before. This compounds the opportunity for future expansion as new versions of Data ONTAP® support larger and larger aggregates.

Aggregates do a very good job of masking the traditional performance implications that are associated with RAID group size. The primary point of this policy is not bound to performance concerns but rather to establishing a consistent approach to aggregate and RAID group sizing that:

- ☐ Facilitates ease of aggregate and RAID group expansion
- ☐ Establishes consistency across the RAID groups in the aggregate
- ☐ Reduces parity tax to help maximize “usable” storage
- ☐ Reduces CPU overhead associated with implementing additional RAID groups that might not be necessary
- ☐ Considers both the time it takes to complete corrective actions and how that relates to actual reliability data available for the drives

These recommendations apply to aggregate and RAID group sizing for RAID-DP®. RAID-DP is the recommended RAID type to use for all NetApp storage configurations. In Data ONTAP 8.0.1, the maximum SATA RAID group size for RAID-DP has increased from 16 to 20.

For HDD (SATA, FC, and SAS) the recommended sizing approach is to establish a RAID group size that is within the range of 12 (10+2) to 20 (18+2); that achieves an even RAID group layout (all RAID groups contain the same number of drives). If multiple RAID group sizes achieve an even RAID group layout, NetApp recommends using the higher RAID group size value within the range. If drive deficiencies are unavoidable, as is sometimes the case, NetApp recommends that the aggregate should not be deficient by more than a number of drives equal to one less than the number of RAID groups. Otherwise you would just pick the next lowest RAID group size. Drive deficiencies should be distributed across RAID groups so that no single RAID group is deficient more than a single drive.

Given the added reliability of SAS and Fibre Channel (FC) drives, it might sometimes be justified to use a RAID group size that is as large as 24 (22+2) if this aligns better with physical drive count and storage shelf layout.

SSD is slightly different. The default RAID group size for SSD is 23 (21+2), and the maximum size is 28. For SSD aggregates and RAID groups, NetApp recommends using the largest RAID group size in the range of 20 (18+2) to 28 (26+2) that affords the most even RAID group layout, as with the HDD sizing approach.

What is maximum raid group size for SATA, SAS, FC and SSD?

Data ONTAP 8.0.1 Default and Maximum RAID Group Size by Drive Type			
Drive Type	RAID Type	Default RAID Group Size	Maximum RAID Group Size
SSD	RAID-DP (default)	23 (21+2)	28 (26+2)
	RAID 4	8 (7+1)	14 (13+1)
SAS/FC	RAID-DP (default)	16 (14+2)	28 (26+2)
	RAID 4	8 (7+1)	14 (13+1)
SATA/NL-SAS	RAID-DP (default)	14 (12+2)	20 (18+2)
	RAID 4	7 (6+1)	7 (6+1)

SHOULD 2TB DRIVES BE USED IN 32-BIT OR 64-BIT AGGREGATES?

2TB drives should be used with 64-bit aggregates in order to take advantage of larger aggregate capacities and, as a result, more drives per aggregate. Two circumstances in which it makes sense to use 2TB drives in 32-bit aggregates are:

- ☐ In FAS2040 configurations using only 12 internal drives (no external storage).
- ☐ FAS2020 is not supported with Data ONTAP 8.0 and therefore has to use 32-bit aggregates regardless of the number of drives attached to the system (internal or external storage).

WHAT ARE THE DIFFERENT PLATFORMS SUPPORTED IN CLUSTERED DATA ONTAP

Clustered Data ONTAP Platform Support			
Platform	Data ONTAP 8.1.x	Data ONTAP 8.2	MAX CLUSTER NODES ²
FAS/V 6290, 6250, 6220	SAN & NAS	SAN & NAS	24 ¹
FAS/V 3250, 3220	SAN & NAS	SAN & NAS	8 ¹
FAS2240 (NFS, CIFS, iSCSI)	SAN & NAS	SAN & NAS	4
FAS2220 (NFS, CIFS, iSCSI)	NAS only	SAN & NAS	4
FAS/V 6240, 6280 & FAS6080	SAN & NAS	SAN & NAS	24 ¹
FAS/V 6210, 3270 & FAS6040	SAN & NAS	SAN & NAS	8 ¹
FAS/V 3240, 3210 & FAS3170, 3160	SAN & NAS	SAN & NAS	4
FAS3140	SAN & NAS	SAN & NAS	2
FAS6070	NAS only	Not Supported	24
FAS6030	NAS only	Not Supported	8
FAS3070	NAS only	Not Supported	4
FAS3040, 2040	NAS only	Not Supported	2

NOTES:

1. CLUSTERS USED FOR SAN ARE LIMITED TO 8 NODES WITH 8.2 AND 6 NODES WITH 8.1.X
2. MAXIMUM NUMBER OF NODES WITHIN A CLUSTER IS DETERMINED BY PLATFORM WHICH SUPPORTS THE FEWEST NUMBER OF NODES

WHY ARE DATA ONTAP 8.0.X ROOT VOLUME SIZES LARGER THAN 7.3.X ROOT VOLUMES?

Data ONTAP 8.0.x supports much larger amounts of memory than 7.3.x. Additional space is needed to account for the increase in memory size—for example, for core dumps.

WHAT IS THE RECOMMENDED SIZE OF A ROOT FLEXVOL VOLUME?

The following table gives the recommended root FlexVol[®] volume sizes for each platform.

Data ONTAP Recommended Root FlexVol Volume Size by Platform		
Platform	Root Volume Size 7.3.x	Root Volume Size 8.0.x
FAS2040	13GiB	133GiB
FAS2240-2	Unsupported	159GiB
FAS2240-4	Unsupported	159GiB
FAS/V3040	16GiB	141GiB
FAS/V3140	16GiB	141GiB
FAS/V3070	23GiB	213GiB
FAS/V3160	16GiB	141GiB
FAS/V3170	40GiB	250GiB
FAS/V3210	17GiB	151GiB
FAS/V3240	22GiB	205GiB
FAS/V3270	44GiB	250GiB
FAS/V6030, 6040	37GiB	250GiB
FAS/V6070, 6080	69GiB	250GiB
FAS/V6200 series	Unsupported	250GiB

When consulting the *Data ONTAP System Administration Guide* (SAG) for the release of Data ONTAP that you run, note that specified minimum root volume sizes might not match the preceding table. The values in the preceding table have been confirmed through actual system testing and are correct. The Data ONTAP documentation (SAGs) being incorrect is a known issue that is being resolved in a future release of the documentation and through release note updates.

DO THE RECOMMENDED ROOT VOLUME SIZES INCLUDE SPACE RESERVED FOR SNAPSHOT COPIES?

No. Recommended root volume size is the amount of capacity that should be available for data to be written.

CAN THE ROOT AGGREGATE BE A 64-BIT AGGREGATE?

In Data ONTAP 8.0.1 and later, you can use a 64-bit root aggregate. In Data ONTAP 8.1 and later, 64-bit aggregates are the default type for all aggregates, including the root aggregate.

SHOULD I USE A DEDICATED ROOT AGGREGATE?

A dedicated root aggregate is an aggregate that contains only a root volume; that is, it does not contain any user data volumes.

Using dedicated root aggregates for Data ONTAP operating in 7-Mode is not required, but it is recommended. The primary reason for using a dedicated root aggregate is that hosting user data volumes on the same aggregate as the root volume can increase the risk of the aggregate becoming inconsistent. Although it is rare for an aggregate to become inconsistent, if it does happen, it requires taking the aggregate offline to restore data consistency. Restoring the root aggregate requires taking the entire system offline, making all data unavailable.

Data ONTAP 7.3.3 and later releases, including Data ONTAP 8.0.x, added the ability to recover an inconsistent root aggregate without having to create an “emergency” root aggregate from hot spares. This reduces the time to recover a root aggregate, but the process still results in a disruption to system availability. If you decide to host user data volumes on the root aggregate, NetApp highly recommends running Data ONTAP 7.3.3 or later.

Data ONTAP operating in Cluster-Mode **requires** the use of a dedicated root aggregate. In addition the points covered above around Data ONTAP operating in 7-Mode, there are additional impacts associated with hosting user data volumes on the root aggregate (for Cluster-Mode). The first aggregate to be returned home during a `sendhome` operation is the root aggregate. The amount of data stored in `mroot` (on the root aggregate) directly affects the time it takes for the failover process to complete. This does not affect separate user data aggregates, because those aggregates are being hosted by the other controller in the HA pair while the root aggregate is being returned. User data volumes within the root aggregate are delayed (and unavailable) in addition, because they cannot be hosted by the other controller in the HA pair. This results in takeover and giveback operations that will exceed the maximum 120 second failover timeframe for NetApp storage systems.

WHAT ARE THE MAXIMUM FLEXVOL SIZES SUPPORTED ON EACH PLATFORM?

The following table lists maximum FlexVol volume size by platform and Data ONTAP release.

Maximum FlexVol Volume Size by Platform and Data ONTAP Version				
Platform	Data ONTAP 7.3.x	Data Ontap 8.0.x	Data Ontap 8.1.x	Data ONTAP 8.2
FAS2040	16TiB	30TiB	30TiB	Not supported
FAS2240-2	Not supported	Not supported	60TiB	60TiB
FAS2240-4	Not supported	Not supported	60TiB	60TiB
FAS/V3040	16TiB	50TiB	50TiB	Not supported
FAS/V3070	16TiB	50TiB	50TiB	Not supported
FAS/V3140	16TiB	50TiB	50TiB	50TiB
FAS/V3160	16TiB	50TiB	50TiB	50TiB
FAS/V3170	16TiB	70TiB	70TiB	70TiB
FAS/V3210	16TiB	50TiB	50TiB	50TiB
FAS/V3220	16TiB	Not supported	50TiB	60TiB
FAS/V3240	16TiB	50TiB	50TiB	50TiB
FAS/V3250	Not supported	Not supported	70TiB	70TiB
FAS/V3270	16TiB	70TiB	70TiB	70TiB
FAS/V6030	16TiB	70TiB	70TiB	Not supported
FAS/V6040	16TiB	70TiB	70TiB	70TiB
FAS/V6070	16TiB	100TiB	100TiB	Not supported
FASV6080	16TiB	100TiB	100TiB	100TiB
FAS/V6210	Not supported	70TiB	70TiB	70TiB
FAS/V6220	Not supported	70TiB	70TiB	70TiB
FAS/V6240	Not supported	100TiB	100TiB	100TiB
FAS/V6250	Not supported	100TiB	100TiB	100TiB
FAS/V6280	Not supported	100TiB	100TiB	100TiB
FAS/V6290	Not supported	100TiB	100TiB	100TiB

HOW SHOULD I ADD CAPACITY (DRIVES) TO AN EXISTING AGGREGATE?

To determine how capacity should be added into an existing aggregate, you should consider the following:

- ☐ **How much capacity do you want to add or can you add based on aggregate size limits?**
Once you have determined how much capacity is being added, you can determine the number of drives that will be added to the aggregate. Once you understand the number of drives being added, you can determine the answer to the following consideration.
- ☐ **Is the current RAID group layout fully optimized (even), or does it lack a number of drives?** There are two ways to tackle additions once you know how many drives are to be added to the storage configuration. Either you can recalculate the optimal RAID group size for the aggregate, or you can add a number of drives by RAID group size. It is easier to add more RAID groups, provided that this does not leave any RAID group deficient of more than one drive per the guidelines for RAID group size. Also, this approach presents the opportunity to assign drives to existing RAID groups that lack some drives. If you change the RAID group size, it should be to a higher value than is currently set, because you will not be able to shrink any existing RAID groups. If you increase the RAID group size, you will need to add drives to existing RAID groups in order to keep them from lacking the number of drives by the delta in RAID group size.

When adding drives to a system, it is very likely that you will not be able to distribute RAID groups across different storage shelves. If storage shelf failure is a concern, NetApp recommends that the storage configuration use NetApp SyncMirror[®] synchronous mirroring software as the method for protecting against shelf failure. When adding drives to a system, remember that you might need to adjust the number of hot spares available in the storage configuration.

HOW CAN I DETERMINE THE USABLE CAPACITY OF THE LARGEST FLEXIBLE VOLUME?

Data ONTAP 8.1 maximum aggregate sizes have increased, but maximum FlexVol volume sizes have remained the same (for example, a FAS6280 can have a maximum aggregate size of 162TiB, but the maximum FlexVol volume size is 100TiB). There are other factors that result in a reduction in the resulting capacity of a FlexVol volume. Additional information on how these numbers are calculated can be found in [TR-3838: Storage Subsystem Configuration Guide](#).

The following table gives the estimated usable FlexVol volume size for the currently supported maximum FlexVol volume capacities by shipping drive capacities for Data ONTAP 8.0.x and 8.1.x. Note that this calculation only takes into account mandatory space reservations and does not consider user-configurable space reservations or storage efficiency features.

Data ONTAP 8.0.x and 8.1.x Estimated Volume Size by Drive Capacity						
Drive Capacity	Maximum FlexVol Volume Size					
	16TiB	30TiB	50TiB	60TiB	70TiB	100TiB
300GB SAS or FC	14.17TiB	26.71TiB	44.60TiB	53.66TiB	62.48TiB	89.43TiB
450GB SAS or FC	14.28TiB	26.77TiB	44.62TiB	53.54TiB	62.47TiB	89.24TiB
600GB SAS or FC	13.87TiB	26.78TiB	44.48TiB	53.56TiB	62.65TiB	89.43TiB
900GB SAS	13.87TiB	26.78TiB	44.48TiB	53.56TiB	62.65TiB	89.43TiB
500GB SATA	14.09TiB	26.74TiB	44.44TiB	53.48TiB	62.51TiB	89.25TiB
1TB SATA	13.75TiB	26.78TiB	44.15TiB	53.56TiB	62.25TiB	89.03TiB
2TB SATA	13.03TiB	26.06TiB	43.44TiB	53.57TiB	62.26TiB	88.32TiB
3TB SATA	13.01TiB	26.02TiB	43.36TiB	52.03TiB	60.70TiB	88.89TiB

Note: The preceding table addresses systems where the maximum FlexVol capacity is equal to the maximum aggregate capacity. With the introduction of Data ONTAP 8.1, many systems' maximum FlexVol capacity is less than the maximum aggregate capacity. On these systems it is possible to yield larger usable volume capacity as some of the aggregate-level capacity reservations are pushed into the remaining aggregate capacity that is beyond the maximum FlexVol capacity. On these systems you are able to create a larger flexible volume (very near full FlexVol size of usable capacity) than is listed in the preceding table.

WHAT IS THE DEFAULT RAID TYPE IN DATA ONTAP?

The following table outlines the default RAID type information.

Default RAID Type		
Storage Container Type	Data ONTAP 7.2.x	Data ONTAP 7.3.x/8.0.x/8.1.x/8.2.x
Traditional volumes (FC drives)	RAID 4	RAID-DP
Traditional volumes (SATA drives)	RAID-DP	RAID-DP
Aggregates	RAID-DP	RAID-DP

SHOULD I ALWAYS DETERMINE THAT NO MORE THAN TWO DRIVES PER RAID GROUP ARE PHYSICALLY LOCATED ON A PARTICULAR STORAGE SHELF?

Although complete storage shelf failures are very rare, they are a possibility. To protect against shelf failures, storage administrators have in some cases spread RAID groups across physical storage shelves so that no more than two drives per RAID group are located on any given storage shelf. As storage shelves become denser and larger, and 64-bit aggregates are implemented, it becomes unreasonable to expect to be able to accomplish this. Most important is the fact that even with this drive layout approach, if a full shelf does fail (without SyncMirror), the system will become extremely degraded to the point that this protection scheme will not work. The recommended method for protecting against storage shelf failures is to use SyncMirror in your storage configuration.

HOW DO I DETERMINE THE TOTAL SIZE OF AN AGGREGATE?

In Data ONTAP 7.2.x and earlier, the total size of aggregates is based on the sum of the Data ONTAP physical capacity of all data and parity drives. In Data ONTAP 7.3.x and later, the total size of aggregates is

based on the sum of the Data ONTAP usable (right-sized) capacity of all data drives only (no parity drives are included).

WHAT IS THE REASON FOR CHANGING THE WAY THE TOTAL SIZE OF AN AGGREGATE IS COMPUTED?

The method used for computing the aggregate size in releases before Data ONTAP 7.3 could result in an aggregate exceeding the maximum allowable size even though the amount of usable space in that aggregate is much smaller than the maximum allowable aggregate size. In Data ONTAP 7.3 and later releases, only the Data ONTAP usable (right-sized) capacity (as reported by the `sysconfig -r` command) of data drives is used to calculate the total aggregate size. If you were prevented from adding drives to an aggregate because of aggregate size constraints in an earlier version of Data ONTAP, you might be able to add more drives to that aggregate after upgrading to Data ONTAP 7.3.

WHY CAN'T I CREATE A VOLUME WHOSE SIZE IS EQUAL TO ALL THE AVAILABLE SPACE IN THE AGGREGATE?

To improve the resiliency of volume space guarantees and WAFL[®] space accounting, Data ONTAP 7.3 consumes an additional 0.5% of every flexible volume's nominal size from the free space on the host aggregate. Therefore, a volume needs free space equal to 100.5% of its specified size from the host aggregate. This reservation applies to new volumes that are created on Data ONTAP 7.3 and is also added to existing volumes upon upgrading to 7.3. One exception to this reservation is with FlexCache[®] volumes, in which only an additional 0.2% of the volume's nominal size is reserved. The Upgrade Advisor tool takes this change into account when recommending upgrades from a release of Data ONTAP earlier than 7.3.

WHAT IS THE RECOMMENDED AGGREGATE AND RAID GROUP CONFIGURATION FOR A FAS2020 OR FAS2040 SYSTEM USING ONLY 12 INTERNAL DRIVES?

In order for the configuration to maximize the number of data drives while maintaining some resiliency, NetApp recommends the following configuration for a FAS2020 or FAS2040 single-controller configuration:

- ☐ Number of aggregates = 1
- ☐ RAID group size = 11
- ☐ Number of data drives = 9
- ☐ Number of parity drives = 2
- ☐ Number of hot spares = 1

In an HA pair configuration there are two approaches that can be taken. The first is an “active-active” approach that allows both controllers in the HA pair to host user data. The second “active-passive” approach has user data stored on one controller, while the second controller is purely a root aggregate/volume providing failover capability for the first controller.

For the “active-active” approach NetApp recommends the following configuration:

- ☐ Controller 1 (RAID-DP):
 - ☐ Number of aggregates = 1
 - ☐ RAID group size = 5
 - ☐ Number of data drives = 3
 - ☐ Number of parity drives = 2
 - ☐ Number of hot spares = 1
- ☐ Controller 2 (RAID-DP):
 - ☐ Number of aggregates = 1
 - ☐ RAID group size = 5
 - ☐ Number of data drives = 3
 - ☐ Number of parity drives = 2
 - ☐ Number of hot spares = 1

For the “active-passive” approach NetApp recommends the following configuration:

- ☐ Controller 1 (RAID-DP):
 - ☐ Number of aggregates = 1
 - ☐ RAID group size = 8

- Number of data drives = 6
 - Number of parity drives = 2
 - Number of hot spares = 1
- Controller 2 (RAID-DP):
- Number of aggregates = 1
 - RAID group size = 3
 - Number of data drives = 1
 - Number of parity drives = 2
 - Number of hot spares = 0*

*As a result of assigning no spare drives to controller 2, warning messages will be generated that can be safely ignored. (`raid.min_spare_count` can be set to 0 to avoid warning messages.) Note that it is still possible to use a RAID 4 aggregate on the “passive” controller. If RAID 4 is used, it is recommended that a single hot spare be used in addition to the drives for the root aggregate (resulting in the same number of drives needed to do RAID-DP with no spares, hence the RAID-DP recommendation). Note that if any RAID 4 aggregates exist on a system, you will not be allowed to set `raid.min_spare_count` to 0 (it must be set to 1 or higher).

It's important to understand aggregate size limits on the FAS2020, as shown in the following table.

FAS2020 Aggregate Size Limit	
Data ONTAP Version	Maximum Aggregate Size
7.3.1 and later	16TB
7.3	6TB
7.2.x and earlier	7TB

Note that 2TB drive capacities are supported starting with Data ONTAP 7.3.2, which is the recommended release for running these configurations.

1.2 32-BIT AGGREGATES

ARE 32-BIT AGGREGATES SUPPORTED IN DATA ONTAP 8.0.X AND LATER RELEASES?

Yes. 32-bit aggregates are currently supported on all Data ONTAP releases. In Data ONTAP 8.0.x and later you have the option of using 32-bit or 64-bit aggregates. In general NetApp always recommends using 64-bit aggregates (although in Data ONTAP 8.1 and later you can do in-place conversion of 32-bit aggregates into 64-bit aggregates simply by growing the aggregate larger than 16TiB).

HOW MANY DRIVES CAN BE USED IN A MAXIMUM SIZE 32-BIT AGGREGATE?

In Data ONTAP 7.2.x and earlier, parity drives and physical drive size are included in the 16TB limit for 32-bit aggregates. In Data ONTAP 7.3.x and 8.0.x, only data drives and usable drive capacity are included in the 16TB limit for 32-bit aggregates. The following recommendations are based on attempting to provide the most optimal RAID group layout, as explained in the answer to the question “Should all aggregates always use the default RAID group size?” earlier.

The column descriptions for the following tables are as follows:

- “Data Drives” is the number of data drives that fit within the maximum aggregate capacity (based on usable drive capacity).
- “RG Size” is the recommended RAID group size to use for the configuration.
- “Number of RGs” is the resulting number of RAID groups the aggregate will contain.
- “Drive Def.” is the number of drives deficient the configuration is from a fully even number of RAID groups.
- “Data + Parity” is the total number of drives used for the aggregate configuration.

The top entries in the following table for the SSD drive show the recommendations for aggregate configurations that are using the maximum number of data drives. In many cases it is better to reduce data drives by a small number in order to achieve a better RAID group layout, as indicated by the bottom numbers, in parentheses.

32-Bit Aggregate Recommendations for All Platforms Data ONTAP 7.3.x, 8.0.x, and 8.1.x						
Capacity	Type	Data Drives	RG Size	Number of RGs	Drive Def.	Data + Parity
100GB	SSD	86 (84)	24 (23)	4 (4)	2 (0)	94 (92)
200GB		86 (84)	24 (23)	4 (4)	2 (0)	94 (92)
300GB	FC	61 (60)	18 (17)	4 (4)	3 (0)	69 (68)
450GB		40 (40)	12 (12)	4 (4)	0 (0)	48 (48)
600GB		29 (28)	17 (16)	2 (2)	1 (0)	33 (32)
300GB	SAS	61 (60)	18 (17)	4 (4)	3 (0)	69 (68)
450GB		40 (40)	12 (12)	4 (4)	0 (0)	48 (48)
600GB		29 (28)	17 (16)	2 (2)	1 (0)	33 (32)
900GB		19 (19)	12 (12)	2 (2)	1 (1)	23 (23)
500GB	SATA	39 (39)	15 (15)	3 (3)	0 (0)	45 (45)
1TB		19 (18)	12 (20)	2 (1)	1 (0)	23 (20)
2TB		9 (9)	11 (11)	1 (1)	0 (0)	11 (11)
3TB		6 (6)	8 (8)	1 (1)	0 (0)	8 (8)

In Data ONTAP 7.3.x the maximum SATA RAID group size is 16. For configuration recommendations stating SATA RAID group sizes larger than 16 it is necessary to file a PVR (which will be approved) so product management can provide the RAID group size override capability with the proper guidance (which allows you to align with best practice recommendations).

The preceding table applies to all platforms with one exception, in regard to the FAS2020 platform. In Data ONTAP 7.3.1 and later, the maximum aggregate capacity for the FAS2020 is 16TB, just as with other platforms. In Data ONTAP 7.3, the actual aggregate size limit for the FAS2020 is 6TB, counting only data drives. Finally, in Data ONTAP 7.2 and earlier, the limit for the FAS2020 is 7TB, counting both data and parity drives. NetApp highly recommends that all FAS2020 systems use Data ONTAP 7.3.1 or later in order to avoid confusion when managing your storage configuration.

Notes for the preceding table:

- 100GB SSD capacity first supported with Data ONTAP 8.0.1
- 200GB SSD capacity first supported with Data ONTAP 8.1.2
- 450GB FC/SAS capacity first supported with Data ONTAP 7.2.5.1
- 600GB FC/SAS capacity first supported with Data ONTAP 7.3.2 and 8.0 RC3/GA
- 900GB SAS capacity first supported in Data ONTAP 7.3.5, 8.0.2, 8.1.0
- 1TB SATA capacity first supported with Data ONTAP 7.2.3
- 2TB SATA capacity first supported with Data ONTAP 7.3.2 and 8.0 RC3/GA
- 3TB SATA capacity first supported with Data ONTAP 8.0.2

1.3 64-BIT AGGREGATES

WHAT ARE THE MAXIMUM 64-BIT AGGREGATE SIZES FOR EACH PLATFORM?

The following table shows the maximum 64-bit aggregate capacities by platform and Data ONTAP release. Note that in Data ONTAP 8.1 the maximum FlexVol volume size is lower than maximum aggregate size on

most platforms (consult the maximum FlexVol table in section 1.1, "General," that shows the maximum FlexVol sizes for each platform).

Maximum 64-Bit Aggregate Size by Platform and Data ONTAP Version			
Platform	Data ONTAP 8.0.x	Data ONTAP 8.1.x	Data ONTAP 8.2.x
FAS2040	30TiB	50TiB	Not supported
FAS/V2240-2/4	Not supported	60TiB	120TiB
FAS/V2220-2/4	Not supported	60TiB	120TiB
FAS/V3040	50TiB	50TiB	Not supported
FAS/V3070	50TiB	50TiB	Not supported
FAS/V3140	50TiB	75TiB	150TiB
FAS/V3160	50TiB	90TiB	180TiB
FAS/V3170	70TiB	105TiB	240TiB
FAS/V3210	50TiB	75TiB	150TiB
FAS/V3220	Not supported	90TiB	180TiB
FAS/V3240	50TiB	90TiB	180TiB
FAS/V3250	Not supported	105TiB	240TiB
FAS/V3270	70TiB	105TiB	240TiB
FAS/V6030	70TiB	105TiB	Not supported
FAS/V6040	70TiB	105TiB	240TiB
FAS/V6070	100TiB	162TiB	Not supported
FAS/V6080	100TiB	162TiB	324TiB
FAS/V6210	70TiB	162TiB	324TiB
FAS/V6220	70TiB	162TiB	324TiB
FAS/V6250	100TiB	162TiB	400TiB
FAS/V6240	100TiB	162TiB	400TiB
FAS/V6280	100TiB	162TiB	400TiB
FAS/V6290	100TiB	162TiB	400TiB

CAN 32-BIT AGGREGATES BE CONVERTED TO 64-BIT AGGREGATES IN DATA ONTAP 8.0?

In Data ONTAP 8.1 and later you can conduct in-place conversion of 32-bit aggregates into 64-bit aggregates. In Data ONTAP 8.0.x, you cannot conduct in-place conversion of 32-bit aggregates into 64-bit aggregates. As a result, NetApp recommends that you upgrade to Data ONTAP 8.1 in order to gain the ability to convert your 32-bit aggregates into 64-bit aggregates.

ARE 64-BIT AGGREGATES SUPPORTED IN DATA ONTAP 7.X.X?

No. 64-bit aggregates were first introduced in Data ONTAP 8.0 and are only supported in Data ONTAP 8.0.x and later releases. Note that both 32-bit and 64-bit aggregates are supported in Data ONTAP 8.0.x and later, whereas only 32-bit aggregates (16TiB maximum aggregate size for all platforms) are supported in Data 7.3.x and earlier.

HOW MANY DRIVES CAN BE USED IN A MAXIMUM SIZE 64-BIT AGGREGATE?

Please find the complete answer to this question here: [64-Bit Aggregate Recommendations](#) (Direct link can be found below)

64-bit aggregates are supported with Data ONTAP 8.0 and later. Each platform has different maximum aggregate capacities for 64-bit aggregates. This section shows the RAID group recommendations for the various maximum aggregate capacities supported across all platforms. The following recommendations are based on attempting to provide the optimal RAID group layout, as explained in the answer to the question "Should all aggregates always use the default RAID group size?" earlier.

The column descriptions for the following tables are as follows:

- ☐ "Data Drives" is the number of data drives that fit within the maximum aggregate capacity (based on usable drive capacity).
- ☐ "RG Size" is the recommended RAID group size to use for the configuration.
- ☐ "Number of RGs" is the resulting number of RAID groups the aggregate will contain.

- “Drive Def.” is the number of drives deficient the configuration is from achieving event RAID group layout.
- “Data + Parity” is the total number of drives used for the aggregate configuration.

The top entries in the following tables show the recommendations for aggregate configurations that are using the maximum number of data drives. In many cases it is better to reduce data drives by a small number in order to achieve a better RAID group layout, as indicated by the bottom numbers, in parentheses.

Please find the 64-Bit Aggregate Recommendations for various Maximum Aggregate Capacities here: [64-Bit Aggregate Recommendations](#)

Please note that we have moved all the tables to this new file “[64-Bit Aggregate Recommendations](#)”

Direct Link to 64-Bit Aggregate Recommendations document:
<https://fieldportal.netapp.com/DirectLink.aspx?documentID=115981&contentID=207731>

2 DRIVES

2.1 GENERAL

1) WHICH DRIVES SUPPORT ADVANCED ZONE CHECKSUMS (AZCS)?

The 3TB (X478A) and 4TB (X480A) MSATA drives in the DS4486 supports AZCS.

2) HOW LONG DOES IT TAKE FOR DISK FIRMWARE UPDATES TO COMPLETE?

When estimating how long it will take to complete a firmware update for a single disk drive use 5 minutes as the general guideline. This estimate is padded and it should complete in less time.

When a nondisruptive disk firmware update is executed the system will update a single drive at a time. For example, if you have 200 disk drives that will be updated via a nondisruptive disk firmware upgrade you are looking at 5 minutes x 200 disk drives for a total time of roughly 16.5 hours. There is little to no performance impact to conducting nondisruptive disk firmware updates. This is the recommended approach for disk firmware updates. Nondisruptive disk firmware updates require the use of RAID-DP.

Disruptive disk firmware updates can be completed faster than nondisruptive but at the cost of taking the system's aggregates and volumes offline (data availability outage). Data ONTAP will do as many disk firmware updates in parallel as possible when conducting a disruptive disk firmware update. How much that can be accomplished in parallel depends on the disk shelf technology being used:

- For DS14 series disk shelves a single disk drive per shelf can be updated at a single time. As there are 14 disk drives in each DS14 disk shelf this means Data ONTAP will conduct 14 passes of disk firmware updates. Each pass will update a single drive in each shelf attached to the system in parallel. This results in an estimated time of 5 minutes x 14 passes for a total time of roughly 1 hour and 10 minutes.
- For SAS based disk shelves or embedded controllers (DS4486, DS4243, DS2246, FAS2200 series, etc.) all drives can be updated in parallel. This means that all drives needing to be updated should be completed within a single 5 minute pass.

3) WHAT IS THE MAXIMUM PER DRIVE IOPS FOR “X” DRIVE?

Per drive IOPS should *never* be used for sizing storage configurations. Sizing should always consider all components of a system as a whole. For this reason, all NetApp storage configuration sizing should be done by using the [Sizer](#) tools.

Per drive IOPS can be helpful when observed from an industry standard perspective to help quickly estimate whether stated performance by a competitor is in the realm of possibility. A standard formula can be used in conjunction with drive specifications to determine what each drive's potential per drive IOPS performance could be. The formula is as follows:

$$1 / ((\text{average seek time} / 1,000) + (\text{average rotational latency} / 1,000)) = \text{estimated IOPS}$$

Average seek time is the average of 100% random read seek time and 100% random write seek time. Both seek time and rotational latency are stated in milliseconds.

The following table is based on publicly available enterprise drive specifications from [Seagate](#) (links to source specifications are embedded in the table), combined with the preceding formula. Specific operating system features, workload, latency, stroke distance per operation, and many more factors ultimately decide what you see in real storage configurations. Note that like drives from different vendors can also vary in performance. The following table uses Seagate to keep the drive manufacturer consistent across the various drive types; however, using the formula you can work out per drive IOPS for any drive manufacturer, provided that you can determine the values needed for the formula.

Estimated Enterprise Drive IOPS					
Type	Speed	Form Factor	Average Seek Time	Average Rotational Latency	Estimated Drive IOPS
SATA	7.2k rpm	3.5"	9.0	4.16	75
NL-SAS	7.2k rpm	3.5"	9.0	4.16	75
FC	10k rpm	3.5"	4.1	2.98	141
FC	15k rpm	3.5"	3.70	2.0	176
SAS	10k rpm	2.5"	4.10	3.0	140
SAS	15k rpm	3.5"	3.70	2.0	176

IOPS latency for the preceding table is a function of the average seek time and average rotational latency, which is the average time it takes the drive to complete an operation without any system-level overhead (for example, RAID).

SSD is not listed in the preceding table as the formula outlined earlier does not apply to solid-state technology (for example, no rotational latency). For reference, testing results have shown NetApp SSD drives are capable of ~8,000 IOPS at 1.5ms or better response times for 4k 100% random read-based workloads. SSD write performance can vary significantly based on many factors. As a general rule for a mixed workload type (significant write portion) it is reasonable to consider SSDs of driving upward of 5,000 IOPS at sub-5.0ms or better response times or 2,000–3,000 IOPS at sub-3.0ms or better response times.

SHOULD SATA DRIVES BE MIXED IN CONFIGURATIONS USING FC, SAS, OR SSD DRIVES?

Scenarios exist in which it is not advisable to create a storage configuration that is actively using both performance drives (FC, SAS, and SSD) and capacity drives (SATA). Any configuration in which maximum controller performance/throughput is a concern should not be mixed.

Two factors must be considered when a SATA drive is added to an aggregate in Data ONTAP:

- Writes directed to a system will end up in NVRAM and from there make their way to drives independent of drive type. When NVRAM flushes data to drives, this is known as a consistency point (CP). CPs targeted at SATA drives take longer than CPs targeted to performance drives such as FC, SAS, or SSD because SATA drives are slower (7.2k rpm). As a result, CPs targeted to performance drives can be “stuck” (waiting) behind CPs being sent to slower SATA drives. Most of the time this is not an issue because writes are normally acknowledged from memory by Data ONTAP, but there are circumstances in which writes are acknowledged from drives—for example, back-to-back CPs.
- When a SATA drive is added to an aggregate, Data ONTAP makes an artificial watermark adjustment to NVRAM (not configurable) to allow CPs to occur more frequently. This adjustment decreases the amount of data in each CP, which in turn allows CPs occurring to performance drives more opportunities to flush to drives when mixed with SATA CPs. The drawback of this adjustment is that maximum controller throughput is reduced slightly as a result of the watermark adjustment.

SSD-based configurations can be sensitive to this, because very few SSDs (less than a single shelf) are capable of driving peak controller performance. This is a write-based factor and does not affect read-based workloads and performance.

WHAT IS THE MTBF OF SAS, FC, SATA, AND SSD DRIVES?

The following table shows the industry standard mean time before failure (MTBF) values for the drive types used by NetApp.

Drive MTBF	
Drive Type	MTBF
FC	1,600,000 hours
SAS 10K	1,600,000 hours
SAS 15K	1,600,000 hours
SATA	1,200,000 hours
SSD (SLC)	2,000,000 hours

WHAT ARE THE DATA ONTAP USABLE (RIGHT-SIZED) CAPACITIES FOR DRIVES?

The following table shows the marketed and right-sized (usable) drive capacities reported by `sysconfig`.

Size	Drive type	Part number	Usable size (MiB)	Physical size(MiB)
300GB	HDD (2.5" 10k rpm, SAS)	X421A-R5	280,104	272,000
450GB		X421A-R5	418,000	420,156
600GB		X422A-R5	560,000	560,208
900GB		X423A-R5	857,000	858,483
1.2TB		X425A-R6	1,142,352	1,144,641
600GB	NSE HDD (2.5" 10k rpm, SAS)	X416A-R5	560,000	560,208
900GB		X417A-R6	857,000	858,483
100GB	SSD (2.5")	X441A-R5	95,146	95,396
100GB		X442A-R5	84,796	84,574
200GB		X446B-R6	190,532	190,782
400GB		X438A-R6	381,304	381,554
800GB		X447A-R6	762,847	763,097
1.6TB		X439A-R6	1,525,935	1,526,185
1TB	SATA (3.5" 7.2k rpm)	X302A-R5	847,555	847,884
2TB		X306A-R5	1,695,466	1,695,702
3TB		X308A-R5	2,538,546	2,543,634
3TB	NSE NL-SAS (3.5" 7.2k rpm)	X309A-R6	2,855,865	2,861,588
4TB		X315A-R6	3,807,816	3,815,447
4TB	NL-SAS (3.5" 7.2k rpm)	X477A-R6	3,807,816	3,815,447
3TB	MSATA	X478A-R5	2,811,241	2,891,588
4TB	MSATA	X480A-R6	3,748,319	3,815,447

IS NEAR-LINE SAS (NL-SAS) THE SAME AS NETAPP'S BRIDGED SAS (BSAS)?

NL-SAS (also referred to as fat SAS) has been made available in the storage industry in order to address the need to have SATA drives that are capable of plugging into SAS shelves directly (with the broad adoption of the SAS shelf architecture). NL-SAS is essentially a SATA drive that has a SAS interface built into the drive itself. ATA-to-SCSI protocol translation is handled internally in the drive. The NetApp bridged-SAS (BSAS) solution uses an external dongle that resides between the drive and shelf enclosure. The bridge provides protocol translation from ATA to SCSI.

There are some performance advantages to using the SAS shelf architecture that translate in some workloads that use both NL-SAS and BSAS drives (increased queue depths and efficiency in driving I/O into the disk subsystem). These same benefits cannot be realized on FC-AL-attached SATA drives.

Key high-level points about NL-SAS and NetApp BSAS are:

- NL-SAS drives are being compared to FC-AL-attached SATA drives. This is not the same as comparing NL-SAS drives to NetApp BSAS drives (which both benefit from the SAS shelf architecture and drive-level SAS feature set).
- From a performance, reliability, and resiliency perspective NL-SAS and BSAS drives are equivalent.
- When NetApp introduced BSAS drives over two years ago, a performance improvement of up to 20+% was observed on some workloads as a result of the benefits of SAS-attached SATA drives (vs. FC-AL-attached SATA drives).
- Manufacturer disk drive specifications available for SATA and NL-SAS disk drives list the same drive-level performance guidance for both drive types.
- Disk rotational speed and areal density continue to be the primary factors that determine disk drive performance capabilities.
- Increased SAS bandwidth to disk and increasing queue depth are software factors that improve how efficiently I/O is driven to disk drives, which can result in improved performance for some workloads.

WHAT ARE THE MAXIMUM SPINDLE COUNT AND MAXIMUM CAPACITY OF A FASXXXX?

Always consult the appropriate system configuration guide to determine the maximum spindle count and system capacity for your storage configuration. The following table provides guidance on these limits for Data ONTAP 7.3.x, 8.0.x, and 8.1.x.

Maximum Spindles and System Capacity by Platform					
Platform	Max Spindles 7.3.x	Max Capacity 7.3.x	Max Spindles 8.0.x/8.1.x	Max Capacity 8.0/8.0.1	Max Capacity 8.0.2/8.1.x
FAS2040	136 drives	136TiB	136 drives	272TiB	408TiB
FAS2240-2	Unsupported		*144 drives	Unsupported	*374TiB
FAS2240-4	Unsupported		*144 drives	Unsupported	*432TiB
FAS/V3040	336 drives	336TiB	336 drives	672TiB	672TiB
FAS/V3140	420 drives	420TiB	420 drives	840TiB	1,260TiB
FAS/V3070	504 drives	504TiB	504 drives	1,008TiB	1,008TiB
FAS/V3160	672 drives	672TiB	672 drives	1,344TiB	2,016TiB
FAS/V3170	840 drives	840TiB	840 drives	1,680TiB	2,520TiB
FAS/V3210	240 drives	240TiB	240 drives	480TiB	720TiB
FAS/V3240	600 drives	600TiB	600 drives	1,200TiB	1,800TiB
FAS/V3270	960 drives	960TiB	960 drives	1,920TiB	2,880TiB
FAS/V6030	840 drives	840TiB	840 drives	1,680TiB	2,520TiB
FAS/V6040	840 drives	840TiB	840 drives	1,680TiB	2,520TiB
FAS/V6070	1,008 drives	1,008TiB	1,008 drives	2,016TiB	3,024TiB
FAS/V6080	1,176 drives	1,176TiB	1,176 drives	2,352TiB	3,528TiB
FAS/V6210	Unsupported		1,200 drives	2,400TiB	3,600TiB
FAS/V6240	Unsupported		1,440 drives	2,880TiB	4,320TiB
FAS/V6280	Unsupported		1,440 drives	2,880TiB	4,320TiB

Notes for the preceding table:

- *FAS2240-2/4 systems have a minimum Data ONTAP release of 8.1.0.
- Maximum spindle counts are halved when using 2TB drives and Data ONTAP 7.3.x.
- FAS6210 maximum spindle count for FC drives is 1,008.
- FAS6240 and FAS6280 maximum spindle count for FC drives is 1,176.
- FAS6200 series systems first supported in Data ONTAP 8.0.1.
- Maximum system capacity increased in Data ONTAP 8.0.2 to support max spindles for 3TB drives.
- 3TB drives are not supported on FAS3000 series systems or as internal FAS2000 series drives.
- 3TB drives are supported as external storage for the FAS2040.

Under normal circumstances, the maximum spindle count is encountered far before the maximum system capacity in “maxcap” systems. Only when a larger drive capacity point is supported by a version of Data ONTAP that has not had its system capacity sized with that capacity point in mind does the maximum system capacity limitation become a real possible limitation.

In many circumstances the term “maximum system capacity” is misused. Some documentation uses a maximum system capacity number that refers to a specific drive technology—for example, a statement that the maximum system capacity of a FAS31XX is Y when using 600GB FC drives. This kind of statement actually refers to full system capacity—that is, the resulting capacity point of a system that is populated with drives of a specific capacity point. This is very different from the maximum system capacity that is the capacity limit enforced by Data ONTAP for a specific platform that is independent of any particular drive capacity point.

WHAT ARE THE ESTIMATED TIMES FOR DRIVE-RELATED ACTIVITIES SUCH AS ZEROING, RECONSTRUCTION, AND RAPID RAID RECOVERY?

Several factors—for example, the drive speed in reading/writing data versus the ability to move that data across the storage stack—can affect the time it takes for any drive-related activity to complete. Also, each drive manufacturer’s drives are different when performing these types of drive activities, so results vary. For more information, see [TR-3838: Storage Subsystem Configuration Guide](#).

The drive information in the following table is based on tests conducted on idle systems, which levels all the factors that affect these drives’ operations. This data is the minimum amount of time these operations should take to complete. Busy systems increase these times. The estimate for systems running at 85% utilization is to multiply the disk reconstruction times by a factor of three (again, this being an estimate).

Drive Operations Minimum Time Estimates					
Capacity	Type	Speed	Zeroing (hrs)	Rapid RAID Recovery (hrs)	Reconstruction (hrs)
100GB	SSD	-	0.1	0.2	0.3
200GB	SSD	-	0.1	0.2	0.3
400GB	SSD	-			
800GB	SSD	-			
1.6TB	SSD	-			
300GB	FC	15k rpm	1.5	1.0	2.0
450GB			2.2	1.5	2.3
600GB			2.5	1.8	2.8
300GB	SAS	15k rpm	1.5	1.0	2.0
450GB			2.2	1.5	2.3
600GB			2.5	1.8	2.8
450GB		10k rpm	2.3	2.5	3.8
600GB			2.6	3.8	4.4
900GB			2.9	5.1	5.0
500GB	SATA	7.2k rpm	2.5	2.8	5.2
1TB			4.3	5.5	9.3
2TB			5.6	12.8	18.5
3TB			7.0	18.3	27.8

DOES NETAPP SUPPORT DRIVE SPINDOWN?

No. One or more of NetApp’s competitors do offer this capability, but because of the risk associated with

spinning down drives, those competitors support drive spindown only on very specific and limited configurations.

Disk drives are not designed with the intention of storing data when the drive is unpowered or is powered down for long periods of time. For similar reasons, cold spares are not considered a replacement for keeping hot spares in your storage configurations.

Drives that are not powered up and spinning offer no guarantee that they will successfully spin back up, and being in the power down or off state could result in a greater likelihood of multiple drive failures occurring at the same time when drives are spun back up to serve data. Multiple drive failures might not occur at the same time physically, but in this case would be discovered at the same time, resulting in significant data availability and resiliency implications for your storage configuration.

There are further implications to data integrity and resiliency. Data ONTAP is continuously using background processes to keep data integral. RAID scrubs and background media scans are just two examples of data integrity background activities that require drives to be spun up and accessible.

Data ONTAP has been designed and implemented to provide high data availability, data integrity, and performance. Drive spindown would be at odds with these objectives.

2.2 FULL DISK ENCRYPTION

WHERE CAN I FIND MORE INFORMATION ON FULL DISK ENCRYPTION?

The NetApp implementation of full disk encryption is called NetApp Storage Encryption (NSE). You can find detailed information in the NSE Technical [FAQ](#) document located on the Field Portal.

2.3 SPARES

CAN A FAILED DISK DRIVE BE REPLACED BY A DISK DRIVE OF SMALLER CAPACITY?

No. A failed disk drive can only be replaced by a disk drive of same or larger capacity.

HOW CAN I CHANGE THE MINIMUM NUMBER OF HOT SPARES REQUIRED FOR MY SYSTEM?

Data ONTAP by default generates warnings about low spare conditions if there is not at least one spare drive that is an appropriate replacement for any file system drive in the system. The option `raid.min_spare_count` can be used to change the minimum number of spares that are checked for warnings for low spare conditions. If there are at least the number of spare drives specified by the option `raid.min_spare_count` that are appropriate replacements for any file system drive, then there will be no warnings for low spares.

This option can be set from 0 to 4; the default setting is 1. Setting this option to 0 means that there will be no warnings for low spares even if no spares are available. This option can be set to 0 only on systems with 16 or fewer attached drives that are running with RAID-DP aggregates. A setting of 0 is not allowed on systems that contain any RAID 4 aggregates.

HOW MANY HOT SPARES SHOULD I KEEP IN MY STORAGE CONFIGURATION?

Recommendations for spares vary by configuration and situation. In the past, NetApp has based spares recommendations strictly on the number of drives attached to a system. This is certainly an important factor, but it's not the only consideration. NetApp storage systems are deployed in a wide range of configurations. This warrants defining more than a single approach to determining the appropriate number of spares to maintain in your storage configuration.

Depending on the requirements of your storage configuration, you can choose to tune your spares policy toward one of the following.

- **Minimum spares:** In configurations where drive capacity utilization is a key concern, you might want to use only the minimum number of spares. This option allows you to survive the most basic failures. If multiple failures occur, it might be necessary to manually intervene to make sure of continued data integrity.
- **Balanced spares:** This configuration is the middle ground between minimum and maximum. It assumes that you will not encounter the worst-case scenario and provides sufficient spares to handle most failure scenarios.
- **Maximum spares:** This option makes sure that enough spares are on hand to handle a failure situation that demands the maximum number of spares that could be consumed by the system at a single time. The term *maximum* doesn't mean that the system might not operate with more than this recommended

number of spares. You can always add additional hot spares within the spindle limits as you deem appropriate.

For RAID-DP configurations, consult the following table for the recommended number of spares.

Recommended Number of Spares		
Minimum	Balanced	Maximum
Two per Controller ¹	Four per Controller	Six per Controller

¹ When using DS4486 shelves, the recommended minimum is four spares.

Special Considerations	
Entry platforms	Entry-level platforms using only internal drives can be reduced to using a minimum of one hot spare.
RAID groups	Systems containing only a single RAID group do not warrant maintaining more than two hot spares for the system.
Maintenance Center	Maintenance Center requires a minimum of two spares to be present in the system.
>48-hour lead time	For remotely located systems, there is an increased chance that they might encounter multiple failures and completed reconstructions before manual intervention can occur. Spares recommendations should be doubled for these systems.
>1,200 drives	For systems using more than 1,200 drives, an additional two hot spares should be added to the recommendations for all three approaches.
<300 drives	For systems using less than 300 drives, you can reduce spares recommendations for a balanced or maximum approach by two.

Selecting any one of the three approaches (minimum, balanced, or maximum) is considered to be the best practice recommendation within the scope of your system requirements. The majority of storage architects will probably choose the balanced approach, although customers who are extremely sensitive to data integrity might warrant taking a maximum spares approach. Given that entry platforms use small numbers of drives, a minimum spares approach would be reasonable for those configurations.

Additional notes about hot spares:

- Spares recommendations are for each drive type installed in the system.
- Larger capacity drives can serve as spares for smaller capacity drives (they will be downsized).
- Slower drives replacing faster drives of the same type affect RAID group and aggregate performance. For example, if a 10k rpm SAS drive (DS2246) replaces a 15k rpm SAS drive (DS4243), this results in a suboptimal configuration.
- Although FC and SAS drives are equivalent from a performance perspective, the resiliency features of the storage shelves in which they are offered are very different. By default, Data ONTAP uses FC and SAS drives interchangeably. This can be prevented by setting the RAID option `raid.disktype.enable` to on.

NetApp does not discourage administrators from keeping cold spares on hand. NetApp recommends removing a failed drive from a system as soon as possible, and keeping cold spares on hand can speed the replacement process for those failed drives. However, cold spares are not a replacement for keeping hot spares installed in a system.

Cold spares can replace a failed part (speeding the return/replace process), but hot spares serve a different purpose: to respond in real time to drive failures by providing a target drive for RAID reconstruction or rapid RAID recovery actions. It's hard to imagine an administrator running into a lab to plug in a cold spare when a drive fails. Cold spares are also at greater risk of being "dead on replacement," as drives are subjected to the increased possibility of physical damage when not installed in a system. For example, handling damage from electrostatic discharge can occur when retrieving a drive to install in a system.

Given the different purpose of cold spares versus hot spares, you should never consider cold spares as a substitute for maintaining hot spares in your storage configuration.

The RAID option `raid.min_spare_count` can be used to specify the minimum number of spares that should be available in the system. This is effective for Maintenance Center users, because when set to the

value 2 it notifies the administrator if the system falls out of Maintenance Center compliance. NetApp recommends setting this value to the resulting number of spares that you should maintain for your system (based on this spares policy) so that the system notifies you when you have fallen below the recommended number of spares.

2.4 SSD

1) WHAT ARE THE USABLE AND PHYSICAL CAPACITIES OF AN SSD?

Size	Shelf	Part number	Usable size	Physical size
100GB	DS4243	X441A-R5	95,146MiB	95,396MiB
	DS4246	X441A-R5	95,146MiB	95,396MiB
200GB	DS2246	X446B-R6	190,532MiB	190,782MiB
	DS4246	X448A-R6	190,532MiB	190,782MiB
400GB	DS2246	X438A-R6	381,304MiB	381,554MiB
800GB	DS2246	X447A-R6	762,847MiB	763,097MiB
1.6TB	DS2246	X439A-R6	1,525,935MiB	1,526,185MiB

2) WHAT IS THE MAXIMUM NUMBER OF SSDS SUPPORTED WITH SHELVES, ONTAP AND PLATFORM?

Size	Shelf	Min ONTAP	PLATFORMS	Max SSDs per system	Max SSDs per stack
100GB	DS4243	8.0.1 +	1,2,3,4,5,6	240 (10 shelves)	96 (4 shelves)
	DS4246	8.0.1 +	1,2,3,4,5,6	240 (10 shelves)	96 (4 shelves)
200GB	DS2246	8.1.1 +	1,3,4,5,6	240 (10 shelves)	96 (4 shelves)
	DS4246	8.1.1 +	1,3,4,5,6	240 (10 shelves)	96 (4 shelves)
400GB	DS2246	8.2.1 +	1,4,6	96 (4 shelves)	96 (4 shelves)
800GB	DS2246	8.1.1 +	1,3,4,6	240 (10 shelves)	96 (4 shelves)
1.6TB	DS2246	8.2.1 +	1,4,6	96 (4 shelves)	96 (4 shelves)

- 1) FAS22xx
- 2) FAS/V 3160/3170
- 3) FAS/V 31xx (FAS3140 Excluded and 3210 only supports Pure SSD)
- 4) FAS/V 32xx
- 5) FAS/V 60xx
- 6) FAS/V 62xx
- 7)

Note:

1) For ONTAP support related to the platform please go to hardware universe.

2) + means mentioned release or higher

3) Mixed Shelf using 800GB not supported on any platform; cache limit is 6Tb per node on 6240, 6280 only

* Please note that Max number of SSDs here means, maximum number of SSDs you can have per system or an HA pair.

3) HOW CAN I DETERMINE THE LEVEL OF WRITE WEAR ON MY SSD DRIVES?

The `storage show disk -a` command can be used to monitor SSD write wear.

4) CAN FLEXSHARE BE USED WITH FLASH POOL SSD'S OR SSD-ONLY AGGREGATES?

No. It is not possible to use FlexShare with SSD's in a Flash Pool or SSD only aggregates.

5) IS THE NETAPP SSD DRIVE SINGLE-LAYER CELL?

Size	Shelf	Part number	Flash Technology
100GB	DS4243	X441A-R5	*SLC
	DS4246	X441A-R5	
200GB	DS2246	X446B-R6	*eMLC
	DS4246	X448A-R6	
800GB	DS2246	X447A-R6	

*SLC: Single-Layer Cell

*eMLC: Enterprise Multi-layer cell

6) IS THERE A TECHNICAL REPORT THAT COVERS THE DETAILS OF HOW SSDS WORK?

Yes, TR-3938: NetApp and Solid State Drive Technology (available from Field Portal) covers the details of how SSDs work and provides detailed explanations around the various SSD-specific features and terminology used today.

7) WHAT IS THE MTBF OF THE NETAPP SSD DRIVE?

All SSD drives offered by NetApp have a 2,000,000-hour MTBF as stated by the drive manufacturer.

8) HOW LONG WILL MY SSD DRIVE OPERATE BEFORE WRITE WEAR BECOMES AN ISSUE?

SSD drives are covered by a five-year warranty (from ship date). SSD drives are fully expected to last for at least five years before any reliability issues are encountered. The MTBF for SSD is 2 million hours (equivalent to FC/SAS drives). NetApp recommends replacing all hardware that is outside of its warranty period.

9) DOES SSD SUPPORT DATA ONTAP 8.0.X OR 8.1.X CLUSTER-MODE?

Yes. Note that SSD support was first introduced in Data ONTAP 8.0.1 for both 7-Mode and Cluster-Mode systems.

10) DOES SSD SUPPORT 32-BIT AGGREGATES?

Yes. SSD is supported in both 32-bit and 64-bit aggregates.

11) DOES THE 10% WAFL RESERVE APPLY TO SSD AGGREGATES?

Yes. The normal 10% WAFL reserve applies to all aggregates created by Data ONTAP.

12) DOES FLASH CACHE ALSO CACHE BLOCKS FOR SSD-BASED VOLUMES?

No. SSD-based volumes are isolated from Flash Cache.

13) SHOULD I USE FLASH CACHE OR SSD FOR MY CONFIGURATION?

The following table addresses whether Flash Cache or SSD is a better fit. Note that mixing Flash Cache and SSD on the same system is supported.

Deciding Between Flash Cache and SSD

SSDs in DS4246 (Persistent Storage) Good Fit When ...	Flash Cache (Intelligent Read Caching) Good Fit When ...
Random read-intensive workload	Random read-intensive workload
Every read must be fast	Improved average response time is adequate
Active dataset is known and fits into SSD capacity	Active dataset is unpredictable or unknown
Active dataset is known, is dynamic, and ongoing administration is OK	Administration-free approach is desired
Upside of write acceleration is desired	Minimizing system price is important
Performance must be consistent across failover events	Want to accelerate an existing HDD storage configuration

14) IS IT TRUE THAT SSD WRITE PERFORMANCE IS THE SAME AS HDD WRITE PERFORMANCE?

No. SSD write performance is better than that of HDD. However, when SSD drives start to fill up and as a result start relying on features such as garbage collection and wear leveling heavily, the write performance of the drive will likely degrade. The return on investment for SSD over HDD for performance is far more significant for read-intensive workloads than for write-intensive workloads.

15) CAN I MIX HDD AND SSD IN THE SAME AGGREGATE OR RAID GROUP?

Yes. In Data ONTAP release 8.1.1 mixing of SSD and HDD is permitted in a Flash Pool. Please refer to the “Flash Pool Design and Implementation Guide [TR-4070](#)” or the “Flash Pool Technical [FAQ](#)” for complete details.

16) CAN I MIX SSD SHELVES WITH HDD SHELVES IN THE SAME STACK?

Yes. In Data ONTAP 8.1.1 SSD and HDD shelves can be mixed in the same stack. Please refer to the “Flash Pool Design and Implementation Guide [TR-4070](#)” or the “Flash Pool Technical [FAQ](#)” for complete details.

17) ARE SSD SHELVES OF FEWER THAN 24 DRIVES AVAILABLE?

Yes. The configured part number for an SSD shelf is for a 24-drive shelf. If you want fewer than 24 SSD drives in a shelf, you can order a quick ship shelf with 6 or 12 SSD drives or you can order the shelf and drives through Add Parts. The following parts can be used to order SSD shelves with fewer than 24 drives:

- **DS4243-SL01-12A-QS-R5:** DS4243 with 12x100GB SSD drives
- **DS4243-SL006-6A-QS-R5:** DS4243 with 6x100GB SSD drives
- **X556A-R6:** DS4243 with 2 PSUs and 24 drive blanks
- **X557A-R6:** DS4243 with 4 PSUs and 24 drive blanks
- **X558A-R6:** DS4243 with no PSUs and no drive blanks
- **X518A-R6:** DS4243 power supply unit (PSU)
- **X5051-R6:** DS4243 blank (drive bay filler)
- **X5712A-R6:** IOM3 shelf module for DS4243
- **X441A-R5:** 100GB SSD drive for DS4243

Note: A minimum of 4 SSD drives is required per shelf in order to properly load the PSUs.

Drive bays on the DS4243 shelf that are not populated by SSD drives are required to use the DS4243 blank drive filler to maintain proper air flow within the shelf. For this reason NetApp recommends using the following part numbers when ordering fewer than 24 SSD drives in a single DS4243 shelf:

- ☐ One X556A-R6 DS4243 with 2 PSUs and 24 drive blanks
- ☐ Two X5712A-R6 IOM3 modules
- ☐ Four or more X441A-R5 SSD drives

The combination of these parts means that you have blanks for all drive bays and the required 2 PSUs for the shelf.

18) WHAT ARE THE PERFORMANCE BOTTLENECKS FOR SSD?

Although drive specifications for SSD can look impressive, the drives are limited by the systems to which they are attached. SSD offers two key benefits:

- ☐ Achieve peak controller performance with fewer drives than would be necessary with HDD

- ☐ Low-latency I/O operations

System CPU power and I/O bandwidth to drives are usually the two limiting factors for maximizing SSD performance.

19) WHICH PLATFORMS AND DATA ONTAP VERSIONS SUPPORT SSD?

The following table shows the minimum versions of Data ONTAP required for each platform that supports SSD.

SSD Minimum Data ONTAP Versions by Platform		
Platform	Data ONTAP 7G	Data ONTAP 8
FAS2240-2/4	Unsupported	Data ONTAP 8.1RC1 or later
FAS/V3160		Data ONTAP 8.0.1 or later 7-Mode or Cluster-Mode
FAS/V3170		
FAS/V3240		
FAS/V3270		
FAS/V6000 series		
FAS/V6200 series		

The following platforms are specifically not supported by SSD:

- ☐ FAS/V2000 series
- ☐ FAS/V3000 series
- ☐ FAS/V3140

20) IS THERE A MINIMUM NUMBER OF SSDS REQUIRED PER SHELF?

Size	Minimum number SSDs
100GB	6
200GB	4

21) WHAT ARE THE DEFAULT AND MAXIMUM RAID GROUP SIZES FOR SSD?

SSD supports both RAID-DP and RAID 4. RAID-DP is highly recommended for all storage configurations. Although a higher parity tax is assumed with RAID-DP, the maximum RAID group size for RAID-DP is double that of RAID 4, which equates to the same parity tax. You get the added resiliency benefits of RAID-DP and the ability to conduct nondisruptive drive firmware updates when using RAID-DP. The following table shows default and maximum RAID group sizes.


SSD RAID Groups		
RAID Type	Default RAID Group Size	Maximum RAID Group Size
RAID-DP (default)	23 (21+2)	28 (26+2)
RAID 4	8 (7+1)	14 (13+1)

22) WHAT IS NOT SUPPORTED WITH SSD THAT IS SUPPORTED WITH HDD?

The following features and configurations are not supported with SSD:

- ☐ Maintenance Center

23) Performance comparison of 100GB, 200GB & 800GB SSDs:



2.5 SATA

WHAT ARE THE REQUIREMENTS FOR USING 3TB SATA DRIVES?

The DS4243/46 supports using 3TB 7.2k rpm 3.5" SATA (X308A-R5) drives starting with Data ONTAP 8.0.2 or later release. The 3TB drive uses the BCS checksum scheme, which results in a Data ONTAP physical capacity of 2,543,634MiB. Data ONTAP usable ("right-sized") capacity is 2,538,546MiB.

The following platforms are supported for use with 3TB drives in a DS4243/46:

- ☐ FAS6200 series
- ☐ FAS6000 series
- ☐ FAS3200 series
- ☐ FAS3100 series
- ☐ FAS2000 series (as external storage in the DS4243 only)

FAS2000 series (internal) and FAS3000 series systems do not support the use of 3TB drives.

HOW DO I DETERMINE THE MAXIMUM NUMBER OF SPINDLES WHEN MIXING 2TB DRIVES WITH OTHER CAPACITY POINTS IN DATA ONTAP 7.3.X?

NetApp has created a simple process for determining spindle count limits when mixing 2TB drives with other capacity points in storage configurations running Data ONTAP 7.3.x. In storage configurations populated completely with 2TB drives, the maximum spindle count for a platform is halved. To determine the number of spindles supported when mixing 2TB drives with other capacity points, do the following:

1. Multiply the number of 2TB drives in the storage configuration by 2.
2. Subtract the result from the maximum number of supported spindles for the platform.

The resulting number is the number of spindles that can be used for non-2TB drive capacities in your storage configuration.

For example, a FAS3170 supports up to 840 spindles (drives). If the storage configuration contains 200 x 2TB drives, the result of step 1 is 400. Subtract 400 from 840, and you get 440 remaining spindles. Those remaining spindles can be used for other capacity points that are less than 2TB.

DOES THE BSAS (SATA IN DS4243) SUFFER A PERFORMANCE HIT FOR GOING FROM ATA TO SCSI?

No. Our bridging technology that is used to bridge between SATA (ATA) and SAS eliminates the performance impact associated with going directly from SATA to SAS without a bridge.

HOW DO THE PERFORMANCE AND RELIABILITY OF SATA DRIVES COMPARE TO THOSE OF FC AND SAS DRIVES?

SAS and FC drives are the same mechanical drives. Programmatically, both SAS and FC drives use the same SCSI command set. The only difference between an SAS and an FC drive is the electronic interface of the drive. Since the electronic interface does not affect performance or reliability, SAS drives show the same characteristics in a side-by-side comparison to FC drives, as outlined in the following table.

Results with SATA drives are included in the chart to highlight the equivalent performance and reliability benefits that both FC and SAS provide over SATA.

Drive Comparison			
	FC	SAS	SATA
Rotational speed	15,000 rpm	10,000 / 15,000 rpm	7,200 rpm
Avg. rotational latency	2.0 ms	2.0 ms	4.2 ms
Seek time avg. read/write	3.5–4.0 ms	3.5–4.0 ms	8.5–9.5 ms
Transfer rate (max)	198MB/sec sustained	198MB/sec sustained	134MB/sec sustained
Drive ports	2	2	1*
MTBF	1.6 million hours	1.6 million hours	1.2 million hours

* NetApp BSAS drives have two ports that are part of the bridge dongle. However, SATA drives are ultimately still single port internally; this tends to be more of a marketing point. For example, competitors might claim that NL-SAS has two ports, but this is still just a marketing point because the drive internally is still a single-port SATA drive. In the case of NL-SAS, the “bridge” with two ports just happens to be inside the drive, as opposed to external like the NetApp BSAS solution.

2.6 FIBRECHANNEL

CAN 10K RPM AND 15K RPM DRIVES BE MIXED IN THE SAME AGGREGATE?

Yes. Set the option `raid.rpm.fcal.enable` to off to allow mixing the two drive types. Note that although this option includes the text “fcal,” it also applies to the mixing of 10k rpm and 15k rpm SAS drives.

WHICH DATA ONTAP VERSIONS SUPPORT 4GB/SEC FC 15K RPM DRIVES?

Data ONTAP 7.0.6, 7.1.2, and 7.2.1 are the minimum releases required for this drive to be recognized as a 15k rpm FC drive. This requirement is independent of loop speed capabilities. Earlier releases of Data ONTAP recognize the 15k rpm 4Gb/sec drives as 10k rpm FC drives. Even though the drive spins at 15k rpm, drive selection policies in older Data ONTAP releases treat the drive as if it were a 10k rpm drive during the following activities:

- ☐ Aggregate creation
- ☐ Drive addition
- ☐ RAID reconstructions
- ☐ Rapid RAID recovery
- ☐ SyncMirror operations
- ☐ Hot spare selection

2.7 SAS

CAN I USE THIRD-PARTY SAS CABLES?

NetApp does not support the use of third-party SAS cables. NetApp extensively qualifies SAS cables for use in NetApp storage configurations. This qualification process maximizes NetApp's confidence in the cables' ability to deliver consistent and sufficient signal strength at all supported lengths. Additionally, NetApp requires that the cable manufacturer put processes in place to test all cables produced to confirm they meet NetApp quality standards. These cables are a critical component of the storage subsystem and NetApp's enterprise-class storage systems.

Data ONTAP also reads information located on the cables' hardware, for example, electronically readable serial numbers. This information is critical in reporting connectivity information and if issues arise. The use of third-party cables disables the ability of Data ONTAP to read this information.

CAN I MIX FC DRIVES AND SAS DRIVES TOGETHER IN MY CONFIGURATION?

By default, Data ONTAP allows FC and SAS drives to be used interchangeably in your storage configuration. This allows you to mix FC and SAS drives in the same aggregate and RAID group. FC and SAS drives are also treated as equivalent for hot spares.

The option `raid.disktype.enable` is set to off by default, which allows the mixing of FC and SAS drives. NetApp recommends setting this option to on to prevent mixing drives together. Mixing FC and SAS drives is not recommended for the following reasons:

- ☐ Although FC and SAS drives are similar at the drive level, the shelf enclosures used for each are very different in both resiliency and performance.
- ☐ As the industry and NetApp transition from FC to SAS (drives and shelves), it will be more challenging to retire FC drives when they are mixed among SAS drives.

HOW IS THE SAS PROTOCOL DIFFERENT FROM THE FC PROTOCOL?

The following table offers a brief overview of the differences between the SAS and FC protocols. For more information about the SAS protocol, see the DS4243 and DS2246 training materials available on the NetApp Field Portal.

SAS and FC Protocol Comparison	
FC-AL	SAS
4Gb/sec per loop (DS14mk4 ESH4)	12Gb/sec aggregate with 4x-wide ports (DS4243) 24Gb/sec aggregate with 4x-wide ports (DS4246 and DS2246)
Drives are not isolated from each other	Drives are isolated from each other
Discovery/management and I/O can conflict	SAS protocol discovery/management and workload I/O can coexist
Limited addressability	Addressability is limited only by performance considerations

IS A SPECIAL COMMAND SET USED TO MANAGE THE SAS SUBSYSTEM?

Yes. A new command called `sasadmin` has been added to facilitate SAS subsystem management.

3 SHELVES AND SHELF MODULES

3.1 GENERAL

ARE SAS (QSFP) CABLES LONGER THAN 5M AVAILABLE?

Yes. There is a 10m active copper SAS cable that is now available. This cable is supported for use DS4243 and DS2246 disk shelves. Note that when using the 10m cable there are specific minimum Data ONTAP versions required and some hardware exceptions:

- ☐ Data ONTAP 7.3.5.1P1, 8.0.1P5, 8.0.2, and 8.1 or later
- ☐ Not supported for use in the FAS2040 onboard SAS port or the X2062 miniSAS dual-port SAS HBA that is used with the FAS2050

HOW CAN I DETERMINE HOW MANY SAS AND ACP CABLES ARE REQUIRED FOR MY NETAPP STORAGE SYSTEM?

Two formulas have been created to make sure you have enough cables to complete the SAS and ACP connections for your NetApp storage system. The formulas only result in the number of cables needed; it is up to the installer to make sure the proper lengths of cables are used.

For SAS cables the formula must be applied individually to each stack of shelves that will be attached to the system:

- ☐ $(\text{number of shelves in stack} * 2) + 2$

For example, if a system were to use 1 stack of 3 shelves and a second stack of 10 shelves, you would need a total of 30 SAS cables to complete the data path connections:

- ☐ $(3 * 2) + 2 = 8$
- ☐ $(10 * 2) + 2 = 22$
- ☐ $8 + 22 = 30$

For ACP cables the formula is applied to the entire system once:

- ☐ $(\text{total number of shelves} * 2) + 1$

For example, if using the same example of a 3-shelf stack and a 10-shelf stack, you would need a total of 27 ACP Ethernet cables to complete the ACP connections:

- ☐ $(13 * 2) + 1 = 27$

WHEN MY NETAPP STORAGE SYSTEM ARRIVES FROM NETAPP, IS THERE A LABEL ON THE SHELF THAT TELLS ME TO WHICH LOOP/STACK THE SHELF BELONGS?

Each shelf that is shipped with configured systems should contain a label that is placed on the top of the disk shelf itself and on the shipping box. The label lists the master ID, system serial number, loop or stack number, and shelf chassis serial number.

An example of the label is as follows:



ARE THERE ANY DETAILED EXAMPLES OF PROPER SAS AND ACP CABLING FOR NETAPP STORAGE SYSTEMS?

Yes. SAS and ACP cabling examples exist for the FAS3200 series and FAS6200 series storage systems:

- [FAS3200 Series SAS and ACP Cabling Examples](#)
- [FAS6200 Series SAS and ACP Cabling Examples](#)

WHICH DISK SHELVES ARE SUPPORTED IN DATA ONTAP 8.1?

Data ONTAP 8.1 supports the following disk shelves:

- DS14mk2 AT using FCX modules
- DS14mk2 FC using ESH4 modules
- DS14mk4 FC using ESH4 modules
- DS4243 using IOM3 modules
- DS2246 using IOM6 modules
- DS4246 using IOM6 modules
- DS4486 using IOM6 modules

ESH2 shelf modules are not supported in Data ONTAP 8.1.

CAN DIFFERENT CAPACITY DRIVES BE MIXED IN THE SAME SHELF?

Drives of the same type (for example, SAS, SATA, FC, or SSD) but of different capacities can be mixed in the same storage shelf. For example, 900GB, 600GB and 450GB 2.5" 10k rpm SAS drives can be mixed in the same DS2246 storage shelf.

WHICH SHELVES AND CONFIGURATIONS SUPPORT NONDISRUPTIVE SHELF REPLACEMENT?

Nondisruptive shelf replacement (NDSR) is supported for DS14mk2 FC, DS14mk2 AT, DS14mk4, DS4243, and DS2246 storage shelves that are in SyncMirror configurations. SyncMirror is a requirement for NDSR.

Note that NDSR applies to replacement of an existing shelf and not the removal or relocation of a shelf. Nondisruptive relocation or removal of storage shelves is a process referred to as single-shelf removal (SSR) and is not currently supported by NetApp.

WHICH SHELF MODULES ARE SUPPORTED ON THE VARIOUS NETAPP STORAGE SHELVES?

The following table shows which shelf modules are supported on each NetApp storage shelf available today.

Shelf Module Support by Storage Shelf	
Storage Shelf	Shelf Module
DS14Mk1	LRC, ESH, and ESH2
DS14Mk2	LRC, ESH, ESH2, and ESH4
DS14Mk2AT	AT-FC, AT-FC2, and AT-FCX
DS14Mk4	ESH4
DS4243	IOM3
DS2246	IOM6
DS4246	IOM6
DS4486	IOM6

IS CONNECTING A SINGLE SHELF SUPPORTED ON ALL AVAILABLE CONTROLLERS?

This is supported but not recommended. NetApp highly recommends not splitting drive ownership between controllers in an HA pair configuration, because this can compromise aggregate- and volume-level resiliency. In some entry-level configurations, you might have no choice but to split drive ownership. In this case, NetApp highly recommends not locating either controller's root aggregates or volumes on the same drive shelf, because this would compromise the ability of the HA pair to conduct successful takeover and giveback operations.

CAN I USE THIRD-PARTY SAS CABLES FOR MY DS4243, DS4246 OR DS2246 DATA PATH CONNECTIONS?

No. NetApp extensively qualifies SAS cables for use with our SAS shelf family. SAS cables are a high-performance and critical component of the SAS architecture. Therefore, only official NetApp SAS cables are supported for use in SAS data path connections.

WHERE CAN I FIND POWER, BTU, AND OTHER ENVIRONMENTAL INFORMATION FOR STORAGE SHELVES AND DRIVES?

The site requirements guide that is available on the NetApp Field Portal has detailed information about each storage shelf and drive configuration available today.

ARE OUR SAS SHELVES SAS 1.0 OR 2.0?

The DS4243 is SAS 1.1 compliant, and DS2246, DS4246, DS4486 are SAS 2.0 compliant.

IS THERE A WAY TO PREVENT MULTIPLE DRIVE RECONSTRUCTIONS IN THE EVENT OF LOST CONNECTIVITY TO A SHELF OR A SHELF FAILURE?

A hidden option can be used to initiate a panic in the event of multiple simultaneous failed drives indicative of lost connectivity to a shelf. The option `raid.panic.missing.disks`, available starting in Data ONTAP 7.1, initiates a panic if the number of simultaneous drive failures equals or exceeds the set threshold. By default, this option is disabled and set to 0. If you enable this capability, NetApp recommends setting this option to 2.

Changing this hidden option in Data ONTAP 7.2.x also changes the related firmware environment variable without needing a reboot to take effect. In Data ONTAP 7.1.x, the variable `halt-missing-disks` must be set from the firmware prompt, which requires a reboot. Data ONTAP releases earlier than 7.1 do not support this capability.

3.2 SHELF INTERMIXING

CAN DS2246, DS4246 AND DS4243 BE MIXED IN THE SAME STACK?

DS224, DS4246 & DS4243 can be mixed into the same stack. However, when going from IOM3 to IOM6, you can't have more than one transition.

PLEASE CONSIDER THESE POINTS WHEN YOU ARE MIXING DIFFERENT SHELF TYPES:

- (1) There should be no more than a single transition between shelves using different IOM module type i.e IOM3 and IOM6. **Note: this is a rule and not a recommendation**
- (2) Any shelf that contains SSDs should be located as the "first/top" or "last/bottom" shelf in the stack so that it is a minimal number of SAS expander hops away from the controller.
- (3) FAS2240-4, FAS2240-2, and FAS2220 storage systems using these shelves as external storage can mix but cannot exceed more than a single transition between shelf types within the storage stack.
- (4) Stack depth limits (maximum number of shelves in a stack) are different for stacks containing SSDs.
- (5) It is recommended to keep storage tiers in a different stack. Therefore SSD/SAS/SATA each should have its own stack. This approach allows for scalability and better performance.

IN WHAT WAYS CAN THE DS4243 AND DS4246 BE MIXED IN THE SAME STACK?

DS4243	DS4246 SATA	DS4246 SATA + SSD
DS4243 SAS	✓	✓
DS4243 SATA	✓	✓
DS4243 SSD	✓	✓

IN WHAT WAYS CAN THE DS4243 AND DS2246 BE MIXED IN THE SAME STACK?

DS4243	DS2246 SAS	DS2246 SSD
DS4243 SAS	✓	✓
DS4243 SATA	✓	✓
DS4243 SSD	✓	✓

IN WHAT WAYS CAN THE DS2246 AND DS4246 BE MIXED IN THE SAME STACK?

DS4246	DS2246 SAS	DS2246 SSD
DS4246 SATA	✓	✓
DS4246 SATA+SSD	✓	✓

MIXING DIFFERENT DRIVE TYPES WITHIN THE SAME SHELF

DS2246	DS2246 SAS	DS2246 SSD
DS2246 SAS	✓	✓

DS4246	DS4246 SSD
DS4246 SATA	✓

WHEN USING IOM3 AND IOM6 IN THE SAME STACK, WILL THE IOM6 THROTTLE DOWN TO IOM3 SPEED?

Yes. Link speed is determined per port based on the shelf module connected. There are two important points to remember in mixing different shelf types:

- 1) When mixing shelf modules in the same stack in supported configurations, determine that the speed is set to match the slowest component speed for the shelf, shelf module, SFP, drive, or HBA port.
- 2) IOM3 and IOM6 can be mixed into the same stack. However, when going from IOM3 to IOM6, you can't have more than one transition.

CAN 15K RPM SAS DRIVES AND 10K RPM SAS DRIVES BE MIXED IN THE SAME AGGREGATE?

Yes. However, NetApp does not recommend mixing 15k rpm and 10k rpm in the same aggregate; because

this will limit the performance of the 15k rpm drives.

HOW MANY DS2246 STORAGE SHELVES ARE SUPPORTED IN EACH STACK?

Each DS2246 storage shelf stack can contain up to 10 shelves (except FAS2040, which is limited to 4). The 10-shelf limit applies only to the number of storage shelves in a stack and is not a system limit, because systems can contain multiple stacks. Note that for very large block sequential workloads it is recommended not to exceed 4 shelves per stack to make sure maximum bandwidth is available for the drives in the stack.

CAN DS2246 (IOM6) AND DS4243 (IOM3) BE CONNECTED TO THE SAME SAS HBA?

Yes. Link speed is determined per port based on the shelf module connected.

NOTE: When mixing shelf modules in the same stack in supported configurations, determine that the speed is set to match the slowest component speed for the shelf, shelf module, SFP, drive, or HBA port.

CAN THE DS4243 BE UPGRADED TO A DS4246?

No. The DS4243 is not upgradeable to the DS4246

HOW MANY DS424x STORAGE SHELVES ARE SUPPORTED IN EACH STACK?

Each DS4243 storage shelf stack can contain up to 10 shelves (except FAS2040, which is limited to 4). The 10-shelf limit applies only to the number of storage shelves in a stack and is not a system limit, because systems can contain multiple stacks. Note that for very large block sequential workloads it is recommended not to exceed 4 shelves per stack to make sure maximum bandwidth is available for the drives in the stack.

3.3 DS2246

1) WHAT ARE THE DRIVES SUPPORTED WITH DS2246?

Size	Drive type	Part number	Usable size (MiB)	Physical size(MiB)
450GB	HDD (2.5" 10k rpm)	X421A-R5	418,000	420,156
600GB		X422A-R5	560,000	560,208
900GB		X423A-R5	857,000	858,483
1.2TB		X425A-R6	1,142,352	1,144,641
600GB	NSE HDD (2.5" 10k rpm)	X416A-R5	560,000	560,208
900GB		X417A-R6	857,000	858,483
200GB	SSD (2.5")	X446B-R6	190,532	190,782
400GB		X438A-R6	381,304	381,554
800GB		X447A-R6	762,847	763,097
1.6TB		X439A-R6	1,525,935	1,526,185MiB

2) What are the supported DATA ONTAP and platforms?

Size	Drive type	Min ONTAP Support	Platform Support
450GB	HDD (2.5" 10k rpm)	7.3.4 +	1,2,3,4,5,6,7
600GB		7.3.4 +	1,2,3,4,5,6,7
900GB		7.3.5 +	1,2,4,5,7
1.2TB		8.1.2 +	2,4,5,6,7
600GB	NSE HDD (2.5" 10k rpm)	8.1.1 +	2,5,7
900GB		8.1.2 +	2,5,7
200GB	SSD	8.1.1 +	2,3,4,5,6,7
400GB		8.2.1 +	2,5,7

800GB	8.1.1 +	2,3,4,5,6,7
1.6TB	8.2.1 +	2,5,7

1. FAS/V 2040, 2050
2. FAS2220, FAS2240-2, FAS2240-4
3. FAS/V 3040, 3070
4. FAS/V 31xx
5. FAS/V 32xx
6. FAS/V 60xx (6000 series is not supported with 900GB)
7. FAS/V 62xx

Note:

- 1) For ONTAP support related to the platform please go to hardware universe.
- 2) + means "or greater"
- 3) 800GB and 1.2TB are not supported on FAS3140 4) 1.2TB is not supported with 3210

3) WHAT IS THE MTBF OF THE DS2246?

There are four main components that make up the DS2246. Those components are the drives, backplane, shelf module, and power supplies. Excluding drive-specific MTBF (which is available in section 2 of this document), the remaining component MTBF values are as follows:

- ☐ DS2246 backplane = 3,000,000 hours MTBF
- ☐ IOM6 = 750,000 hours MTBF
- ☐ DS2246 power supply = 1,500,000 hours MTBF

4) DOES THE DS2246 WORK WITH THE EXISTING SAS HBAS THAT ARE AVAILABLE WITH DS4243 AND DS4246?

Yes. The X2065 quad-port SAS HBA is 6Gb/sec ready and supports operating at both 3Gb/sec SAS and 6Gb/sec SAS link speeds. The onboard SAS port for the FAS2040 and the X2062 dual-port SAS HBA for the FAS2050 are supported with the DS2246 (IOM6) and DS4246(IOM6), but operate at a link speed of 3Gb/sec.

5) WHAT ARE THE COMMONALITIES BETWEEN DS2246, DS4246 AND DS4243?

There are several commonalities between the DS2246, DS4246 and the DS4243, because both are SAS-based shelves. The following components are compatible with the DS2246, DS4246 and DS4243:


- ☐ Quad small form-factor pluggable (QSFP) copper SAS cables
- ☐ X2065 quad-port SAS host bus adapter (HBA)
- ☐ X2062 dual-port SAS HBA (FAS2050 only)
- ☐ Alternate control path (ACP) Ethernet cables

6) HOW DOES DS2246 AND DS4246 PERFORMANCE COMPARE TO THAT OF THE DS4243?


There are two primary performance differences between the DS2246, DS4246 and the DS4243:

- ☐ DS2246 and DS4246 support 10k rpm SAS drives, whereas DS4243 supports 15k rpm SAS drives.
- ☐ DS2246 and DS4246 use IOM6 shelf modules, whereas DS4243 uses IOM3 shelf modules.

Although the aggregate bandwidth is more with IOM6 than with IOM3, the primary performance difference is in the rotational speed of the drives. The following graph compares OLTP workload testing results on a FAS3070 using DS2246 with 600GB 10k rpm SAS drives and DS4243 with 600GB 15k rpm SAS drives.



The Following graph shows an example of sequential read performance between the DS2246 and DS4243.



Sequential read performance is similar between the two shelves at low to mid throughput. Although DS2246 uses 10k rpm drives, the primary benefit for the DS2246 is storage density when compared to the equivalent DS4243 configuration that uses 3.5" 15k rpm SAS drives.

The following table compares a full 42U cabinet of DS2246 shelves using 600GB SAS 2.5" 10k rpm drives to a full cabinet of DS4243 shelves using 600GB SAS 3.5" 15k rpm drives.

DS2246 and DS4243		
Comparison	DS2246 Fully Populated 42U Cabinet 600GB Drives	DS4243 Fully Populated 42U Cabinet 600GB Drives
Shelves per cabinet	21 shelves	10 shelves
Drives per cabinet	504 drives	240 drives
GB per rack unit*	7,200GB	3,600GB
Power consumption per cabinet**	34.44 amps 5,586 watts 19,068 BTU/hr	28.7 amps 5,470 watts 18,650 BTU/hr

□ Based on marketed capacity of the 600GB drive.

** Typical case power consumption for line voltage 200-240V.

7) WHICH PLATFORMS AND DATA ONTAP VERSIONS SUPPORT DS2246?

The following table shows the minimum version of Data ONTAP required for each platform.

DS2246 Minimum Data ONTAP Version by Platform			
Platform	Port Type/HBA	Data ONTAP 7G	Data ONTAP 8
FAS2040	Onboard	Data ONTAP 7.3.4 or later*	Data ONTAP 8.0.1 or later*
FAS2050	X2062 dual-port SAS HBA		Unsupported
FAS/V3040 and FAS/V3070	X2065 quad-port SAS HBA	Data ONTAP 7.3.3P2 or later	Data ONTAP 8.0P1 or later
FAS/V3100 series			
FAS/V6000 series			
FAS/V3200 series	X2065 quad-port SAS HBA and onboard	Data ONTAP 7.3.5 or later	Data ONTAP 8.0.1 or later
FAS/V6200 series		Unsupported	

- Stack bandwidth operates at 3Gb/sec (aggregate 12Gb/sec) SAS, not 6Gb/sec (aggregate 24Gb/sec) SAS, because the onboard SAS port for the FAS2040 and the X2062 dual-port SAS HBA for the FAS2050 only support 3G/sec SAS connectivity.

8) WHAT ARE THE PHYSICAL AND USABLE CAPACITIES OF THE SAS DRIVES OFFERED WITH DS2246?

The following table shows the usable and physical drive capacities as reported by `sysconfig -r` for the 2.5" SFF 10k rpm SAS drives offered with the DS2246.

DS2246 Drive Capacities		
Drive	Usable "Right-Sized"	Physical
450GB, 2.5" SFF, 10k rpm, SAS	418,000MiB	420,156MiB
600GB, 2.5" SFF, 10k rpm, SAS	560,000MiB	560,208MiB* 572,325MiB*
900GB, 2.5" SFF, 10k rpm, SAS	857,000MiB	858,483MiB
1.2TB, 2.5" SFF, 10k rpm, SAS	1,142,352	1,144,641
600GB, NSE 2.5" SFF, 10k rpm, SAS	560,000MiB	560,208MiB

* The two physical capacities listed are a result of NetApp's standard policy of multisourcing drive suppliers.

9) DOES THE DS2246 SUPPORT THE IOM3 SHELF MODULE?

No. The DS2246 does not support the IOM3 shelf module. There are currently no plans to support the IOM3 shelf module in the DS2246.

10) ARE 15K RPM 2.5" DRIVES SUPPORTED WITH DS2246?

No. The 2.5" small form factor (SFF) 15k SAS drives available today have a quarter to half of the capacity available with 2.5" SFF 10k rpm SAS drives and are double the cost per GB. For example, two 300GB 2.5" 15k rpm SAS drives cost twice as much as a single 600GB 2.5" 10k rpm SAS drive. Given these considerations, implementing 2.5" 15k rpm SAS drives with DS2246 would completely eliminate the primary density benefit (GB/U) the shelf provides.

The DS2246 provides excellent density per rack unit for performance drives, allowing you to serve workloads with more drives in less space. Due to the density and lower power consumption per drive, the DS2246 offers higher throughput per cabinet, amps, watts, and BTU/hr than the DS4243. Customers looking for additional performance per drive should consider accelerating their HDD configuration with Flash Cache, the deployment of SSDs, or the use of DS4243 15k rpm SAS drives.

Key points in regard to 2.5" 15k rpm SAS drives versus 2.5" 10k rpm SAS drives are:

- The industry is moving to 2.5" 10k rpm SAS disk drives and eMLC SSD. This is not being decided by NetApp. NetApp is aligning with the industry changes, not dictating them.
- Sequential disk drive performance as compared between 3.5" 15k rpm SAS disk drives and 2.5" 10k rpm SAS disk drives is very similar.
- Small block random read workloads have the largest performance delta when comparing 3.5" 15k rpm SAS and 2.5" 10k rpm SAS disk drives (on the order of 30-35% delta). SSD is best for addressing random read workloads in general (far better than even 15k rpm disk drives), so consider using SSD in these environments.
- Performance deltas can be addressed through the use of Flash Cache or additional spindles (since the DS2246 uses half the space of a DS4243).
- NetApp will continue to offer 3.5" 15k rpm SAS disk drives in the DS4243 for as long as they are commonly available from enterprise disk manufacturers.

11) ARE 2.5" SFF SATA DRIVES SUPPORTED WITH DS2246?

No. The 2.5" small form factor (SFF) SATA drives available today have a quarter to half of the capacity available with 3.5" SATA drives, at a higher cost per GB. Therefore, a SATA option presents the same or worse GB/U at a higher cost than a DS4243 solution that offers the same total capacity.

12) IS THERE A MINIMUM DRIVE REQUIREMENT FOR THE DS2246?

No. NetApp recommends using a single drive installed in any drive bay in order to provide some load to the PSUs in the shelf.

13) WHERE CAN I FIND POWER CONSUMPTION INFORMATION FOR THE DS2246?

The Site Requirements Guide, available from the NetApp Support (formerly NOW[®]) site, contains power consumption information for all NetApp storage shelves.

14) HOW MANY POWER SUPPLIES ARE REQUIRED FOR DS2246?

The DS2246 ships with two power supply units (PSUs). Although the shelf will operate with only a single PSU (in the event of a PSU failure), NetApp recommends using two PSUs for shelf power resiliency.

15) ARE THE POWER SUPPLY UNITS FOR DS2246 AND DS4243 INTERCHANGEABLE?

No, the DS4243 and DS2246 use different power supply units.

16) DOES THE DS2246 SUPPORT SATA DRIVES?

No. The DS2246 supports the following drives:

- ☐ X421, 450GB, 2.5", 10k rpm SAS drive
- ☐ X422, 600GB, 2.5", 10k rpm SAS drive
- ☐ X423, 900GB, 2.5", 10k rpm SAS drive

The DS4243 is the storage shelf that will continue to support 3.5" form factor SATA drives. Although 2.5" SATA drives are available today, they don't make sense from a cost and capacity perspective, much in the same way that 2.5" 15k rpm SAS drives don't make sense to implement at this time.

3.4 DS4243

WHAT IS THE MTBF OF THE DS4243?

There are four main components that make up the DS4243. Those components are the drives, backplane, shelf module, and power supplies. Excluding drive-specific MTBF (which is available in section 2 of this document), the remaining components MTBF values are as follows:

- ☐ DS4243 backplane = 3,000,000 hours MTBF
- ☐ IOM3 = 750,000 hours MTBF
- ☐ DS4243 power supply = 1,500,000 hours MTBF

CAN THE DS4243 BE UPGRADED TO A DS4246?

No. The DS4243 is not upgradeable to the DS4246

CAN SAS AND SATA DRIVES BE MIXED IN THE SAME DS4243 DISK SHELF?

No. Although the DS4243 shelf enclosure supports both SAS and SATA drives, it must contain only a single drive type at any given point in time. The DS4243 disk shelf must have either all SATA or all SAS drives. You can change between using either drive. For example, it is possible to take a shelf of DS4243 using 24 SATA drives, remove the SATA drives, and install 24 SAS drives without the need to do anything other than physically remove and install the drives.

CAN THE DS4243 BE USED IN A STRETCH METROCLUSTER CONFIGURATION?

Yes. The DS4243 is supported in a stretch MetroCluster[™] configuration within the supported SAS cable distance (5m).

CAN FAS2000 SERIES INTERNAL DRIVES BE REMOVED AND INSTALLED IN A DS4243?

No. FAS2000 series internal drives are a separate FRU, and the drive carrier is not compatible with the DS4243 disk shelf.

IS THERE A MINIMUM DRIVE REQUIREMENT FOR THE DS4243?

When you are using SSD drives, the DS4243 must use a minimum of four drives. There is no requirement for which drive bays must be used as long as at least four drives are installed somewhere in the shelf.

There is no minimum for rotating drives (SATA or SAS). This is at odds with the minimum requirements for SSD drives. Technically, NetApp does support an empty DS4243 shelf. However, NetApp recommends using a single drive installed in any drive bay in order to provide some load to the PSUs in the shelf.

HOW DOES DS4243 STORAGE SHELF PERFORMANCE COMPARE TO THAT OF DS14 SERIES STORAGE SHELVES?

FC and SAS performance is equivalent to that of DS4243. Up to a 20% improvement (workload dependent) has been seen when comparing DS4243 SATA to DS14Mk2AT SATA configurations. Performance is not any worse than that of DS14-based solutions and might be (or rather should be) better than that of DS14. The SATA improvement is due to features such as native command queuing. DS4243 stack bandwidth is better, although it is up to the drives and system resources such as the CPU to take advantage of that increased bandwidth under various workloads.

CAN THE POWER CONSUMPTION OF A DS4243 STORAGE SHELF BE HALVED WITH A HALF-POPULATED SHELF?

DS4243 storage shelves can be ordered in full (24-drive) and half-full (12-drive) configurations. Shelf power consumption information is the same for a half-loaded shelf as it is for a fully loaded shelf. There is nothing to prevent someone from plugging in additional drives and exceeding power parameters. Also, shelf power consumption with a half-loaded shelf is not a 50%-type equation because there are several shelf infrastructure components that are on and that require power no matter how many drives are in the shelf. Full shelf power requirements should always be used for both full and half-loaded shelves.

Reducing power requirements for any NetApp storage shelf based on drive population is unsupported.

HOW MANY DS4243 STORAGE SHELVES ARE SUPPORTED IN EACH STACK?

Each DS4243 storage shelf stack can contain up to 10 shelves (except FAS2040, which is limited to 4). The 10-shelf limit applies only to the number of storage shelves in a stack and is not a system limit, because systems can contain multiple stacks. Note that for very large block sequential workloads it is recommended not to exceed 4 shelves per stack to make sure maximum bandwidth is available for the drives in the stack.

HOW MANY CROSSOVER POINTS BETWEEN SATA AND SAS CAN I IMPLEMENT IN MY DS4243 STACK?

The term "crossover point" refers to the number of times a DS4243 storage shelf makes the transition from SATA to SAS (or vice versa) within a single stack. Multiple crossover points per stack are supported. NetApp recommends minimizing the number of crossover points when initially deploying a storage configuration. As storage configurations grow, it is likely that additional crossover points will be introduced into the storage configuration. This is acceptable; however, by reducing the number of crossover points in your stacks, you reduce the variability of the storage configuration. As is true in all computer systems today, increased variability can add difficulty when troubleshooting issues and also raises the possibility of encountering unknown issues.

WHAT ARE THE DIFFERENCES IN PSU REQUIREMENTS FOR THE DS4243 WHEN POPULATED WITH SATA OR SAS DRIVES?

DS4243 power supply unit (PSU) requirements for a shelf populated with SATA drives are different than those for a shelf populated with SAS drives. DS4243 shelves shipped with SAS drives have four PSUs included per shelf. When shipped with SATA drives, two PSUs are included, along with two blocking plates. When operating with only two PSUs, the blocking plates are required in order to provide proper airflow in the shelf.

Note that 2 PSUs only are required for DS4243 shelves using SATA drives. This is the normal configuration. Staggered drive spin-up is not a "bad" condition but simply a mechanism that is built into the shelf. There is no need to order 2 additional PSUs for DS4243 shelves using SATA drives.

The following table describes the effects of running with different numbers of PSUs.

DS4243 Power Supply Unit Behavior		
Type	PSUs	Expected Behavior
SAS	4	Normal operation. All drives spin up at the same time.

	3	Staggered drive spin-up. Six drives per column (four columns) spin up at a time, with 12-second intervals between each column.
	2	Staggered drive spin-up as with 3 PSUs.
	1	The shelf still functions, but with several issues. Any drives that are currently powered up continue to function when reduced to 1 PSU. From a cold boot, some drives (undetermined) will not spin up due to a lack of power once a certain number of drives have spun up. This is an unsupported operating condition.
SATA	4	Normal operation. All drives spin up at the same time.
	3	Staggered drive spin-up. Six drives per column (four columns) will spin up at a time, with 12-second intervals between each column.
	2	Staggered drive spin-up as with 3 PSUs.
	1	Staggered drive spin-up, with one difference. The first three columns, six drives each, spin up with the normal 12-second interval. The fourth column spins up five of six drives, with the sixth drive (bay 23) spinning up 12 seconds later.

DOES THE DS4243 SUPPORT THE IOM6 SHELF MODULE?

No. The DS4243 does not support using the IOM6 shelf module.

IS A DS4243 THAT USES 3GB/SEC SLOWER THAN A DS14MK4 THAT USES 4GB/SEC?

No. The 3Gb/sec SAS ports used with DS4243 storage shelves are SAS wide ports, each of which contains four lanes. Each lane operates at 3Gb/sec for a total port bandwidth of 12Gb/sec.

WHAT IS THE BEST PRACTICE FOR ASSIGNING SHELF IDS TO DS4243 STACKS?

DS4243 shelf IDs can range from 00 to 99. NetApp recommends assigning intervals of 10 to each stack present in the system. For example, the first stack would have IDs 10 through 19 reserved, stack two would have IDs 20 through 29 reserved, and so on. Those IDs are reserved even if the stacks are not fully populated. For example, if the first stack has four DS4243 shelves, the shelf IDs would be 10, 11, 12, and 13. The second stack would start at shelf ID 20, even though IDs 14 through 19 are not currently used. This helps storage administrators easily identify which shelves are assigned to each stack attached to the system. Unused IDs can be used for future stack expansion.

IS MPHA REQUIRED WHEN USING DS4243 STORAGE SHELVES?

All DS4243-supported platforms, except for the FAS2050 and FAS2040, that are an HA pair require MPHA. Due to onboard SAS port limitations, the FAS2040 is supported in a single-path configuration. Single-controller configurations require dual-path connectivity when using DS4243.

WHICH RACKMOUNT KITS SHOULD BE USED FOR DS4243 STORAGE SHELVES IN NETAPP CABINETS?

Many are familiar with the X5515A-R6 universal rackmount kit used for mounting DS14 into two- or four-post racks. Use X877B-R6 for mounting DS4243 storage shelves into cabinets. Part numbers X877B-R6 and X5515A-R6 are not equivalent parts.

The X5515A-R6 kit has been replaced by two part numbers for DS4243. The previous universal approach was convenient but wasteful, since by definition only about half the content was being used in any deployment. The two new part numbers for DS4243 are:

- ☐ X5525A-R6: Rackmount Kit,2-Post,DS4243,R6 (which includes midmount)
- ☐ X5526A-R6: Rackmount Kit,4-Post,Universal,R6

The previous X5515A-R6 kit is not compatible or suitable for use with DS4243. When choosing the appropriate new part, be aware of whether the DS4243 is being installed in a two-post or four-post rack.

The X877B-R6 kit is the follow-on for the older X877-R6. The X877B-R6 is backward compatible with the previous part, but the X877-R6 is not forward compatible with the DS4243 due to length/weight and mounting point issues.

3.5 DS424

1) WHAT ARE THE DRIVES SUPPORTED WITH DS4246?

DS4246 DRIVE INFORMATION				
Size	Drive type	Part number	Usable size	Physical size
100GB	SSD	X441A-R5	95,146MiB	95,396MiB
200GB		X448A-R6	190,532MiB	190,782MiB
1TB	SATA (3.5" 7.2k rpm)	X302A-R5	847,555MiB	847,884MiB
2TB		X306A-R5	1,695,466MiB	1,695,702MiB
3TB		X308A-R5	2,538,546MiB	2,543,634MiB
3TB	NSE NL-SAS (3.5" 7.2k rpm)	X309A-R6	2,855,865Mib	2861588MiB
4TB	NL-SAS (3.5" 7.2k rpm)	X477A-R6	3,807,816MiB	3,815,447MiB

2) What are the supported DATA ONTAP and platforms?

DS4246 ONTAP & PLATFORM SUPPORT			
Size	Drive type	Min ONTAP Support	Platform Support
100GB	SSD	8.0.1RC1 +	1,2,3,4,5,6,7
200GB		8.1.1 +	2,3,4,5,6,7
1TB	SATA (3.5" 7.2k rpm)	7.3.4 +	1,2,3,4,5,6,7
2TB		7.3.4 +	1,2,3,4,5,6,7
3TB		7.3.4 +	1,2,3,4,5,6,7
3TB	NSE NL-SAS (3.5" 7.2k rpm)	8.1.1 +	2,5,7
4TB	NL-SAS (3.5" 7.2k rpm)	8.1.2 +	2,3,4,5,6,7

1. FAS/V 2040, 2050 (Supported in external shelves with SSDs)
2. FAS2220, FAS2240-2, FAS2240-4
3. FAS/V 3040, 3070
4. FAS/V 31xx
5. FAS/V 32xx
6. FAS/V 60xx
7. FAS/V 62xx

Note:

- 1) For ONTAP support related to the platform please go to hardware universe.
- 2) + means mentioned release or higher

3) WHAT IS THE MTBF OF THE DS4246?

There are four main components that make up the DS4246. Those components are the drives, backplane, shelf module, and power supplies. Excluding drive-specific MTBF (which is available in section 2 of this document), the remaining components MTBF values are as follows:

- DS4246 backplane = 3,000,000 hours MTBF

- ☐ IOM6 = 750,000 hours MTBF
- ☐ DS4246 power supply = 1,500,000 hours MTBF

4) CAN THE DS4246 BE USED IN A STRETCH METROCLUSTER CONFIGURATION?

Yes. The DS4246 is supported in a stretch MetroCluster™ configuration within the supported SAS cable distance (5m).

5) HOW MANY DS4246 STORAGE SHELVES ARE SUPPORTED IN EACH STACK?

Each DS4246 storage shelf stack can contain up to 10 shelves (except FAS2040, which is limited to 4). The 10-shelf limit applies only to the number of storage shelves in a stack and is not a system limit, because systems can contain multiple stacks. Note that for very large block sequential workloads it is recommended not to exceed 4 shelves per stack to make sure maximum bandwidth is available for the drives in the stack.

6) WHAT IS THE BEST PRACTICE FOR ASSIGNING SHELF IDS TO DS4246 STACKS?

DS4246 shelf IDs can range from 00 to 99. NetApp recommends assigning intervals of 10 to each stack present in the system. For example, the first stack would have IDs 10 through 19 reserved, stack two would have IDs 20 through 29 reserved, and so on. Those IDs are reserved even if the stacks are not fully populated. For example, if the first stack has four DS4246 shelves, the shelf IDs would be 10, 11, 12, and 13. The second stack would start at shelf ID 20, even though IDs 14 through 19 are not currently used. This helps storage administrators easily identify which shelves are assigned to each stack attached to the system. Unused IDs can be used for future stack expansion.

7) IS MPHA REQUIRED WHEN USING DS4246 STORAGE SHELVES?

All DS4246-supported platforms, except for the FAS2050 and FAS2040, that are an HA pair require MPHA. Due to onboard SAS port limitations, the FAS2040 is supported in a single-path configuration. Single-controller configurations require dual-path connectivity when using DS4246.

3.6 DS4486

CAN THE DS4486 BE MIXED WITH OTHER SHELVES IN A STACK?

No. The DS4486 must be used in a homogenous stack.

CAN THE TWO DRIVES IN A SINGLE CARRIER BE IN A SINGLE RAID GROUP?

No. The Data ONTAP will not allow you to configure the drives on the same carrier in the same RAID group.

IF ONE OF DRIVES ON THE CARRIER FAIL, WHEN CAN I REMOVE THE CARRIER TO REPLACE THE DRIVES?

Upon a drive failure, Data ONTAP will begin a Well Disk Copy to move the data from the functional drive so a complete evacuation of the carrier. When one of the drives fails on the carrier, you must wait for status LED to become solid amber before you can remove the carrier.

HOW LONG DOES IT TAKE FOR A WELL DISK COPY TO COMPLETE?

The time will vary from 8 to 12 hours based on the storage controller load. Well Disk Copy runs at a lower priority than RAID reconstructs..

HOW MANY DS4486 STORAGE SHELVES ARE SUPPORTED IN EACH STACK?

Each DS4486 storage shelf stack can contain up to 5 shelves. The 5-shelf limit applies only to the number of storage shelves in a stack and is not a system limit, because systems can contain multiple stacks. Note that for very large block sequential workloads it is recommended not to exceed 4 shelves per stack to make sure maximum bandwidth is available for the drives in the stack.

WHAT IS THE BEST PRACTICE FOR ASSIGNING SHELF IDS TO DS4486 STACKS?

DS486 shelf IDs can range from 00 to 99. NetApp recommends assigning intervals of 10 to each stack present in the system. For example, the first stack would have IDs 10 through 19 reserved; stack two would have IDs 20 through 29 reserved, and so on. Those IDs are reserved even if the stacks are not fully populated. For example, if the first stack has four DS486 shelves, the shelf IDs would be 10, 11, 12, and 13.

The second stack would start at shelf ID 20, even though IDs 14 through 19 are not currently used. This helps storage administrators easily identify which shelves are assigned to each stack attached to the system. Unused IDs can be used for future stack expansion.

CAN THE DS4486 BE USED IN AGGREGATES OR RAID GROUPS WITH OTHER SHELVES TYPES?

No. The DS4486 must be used in a homogenous aggregate or RAID group.

IS MPHA REQUIRED WHEN USING DS4486 STORAGE SHELVES?

All DS4486-supported platforms that are an HA pair require MPHA. Single-controller configurations require dual-path connectivity when using DS4486.

IS THE DS4486 SUPPORTED IN A METROCLUSTER SOLUTION?

No. The DS4486 does not support MetroCluster.

HOW DOES DS4486 PERFORMANCE COMPARE TO THAT OF THE DS4246?

DS4486 is positioned for backup/archival, with that in mind, drives in DS4486 use Advanced Zone Checksum scheme (AZCS) vs. DS4246 uses Block checksum scheme (BCS). From the performance perspective, if the workload is sequential then we won't see much difference in performance. However, for random workloads, there can be a difference of 30%. This comes from AZCS in DS4486. AZCS uses 63/64th checksum scheme, which results in higher capacity per drive than BCS, which uses 8/9th method. Thus, AZCS gives us back the 10% capacity but there is trade off on performance. This aligns with our message of utilizing DS4486 for higher density, back up and archival.

3.7 ALTERNATE CONTROL PATH (ACP)

SHOULD THE ACP NETWORK ADDRESS BE CHANGED IF IT CONFLICTS WITH NETWORKS BEING USED ON CLIENT NETWORK INTERFACES?

Yes. If the default ACP network address is the same network address as a network being used on a client network interface than the ACP network should be changed to avoid conflicts. The ACP network address can be changed by running the `setup` command in Data ONTAP.

DOES THE ATTO FIBREBRIDGE PRODUCT THAT IS USED FOR METROCLUSTER SYSTEMS SUPPORT ACP?

No. ACP is not supported for use with the ATTO FibreBridge product that is used to connect SAS shelves (DS4243 and DS2246) in MetroCluster systems.

AM I REQUIRED TO USE THE OFFICIAL NETAPP ETHERNET CABLES FOR ACP?

No. Although NetApp recommends using the shielded Ethernet cables provided by NetApp for ACP connections, it is not required.

WHICH VERSIONS OF DATA ONTAP CONTAIN THE FIXES FOR THE ACP ISSUES SEEN IN DATA ONTAP 7.3.2?

The fixes for the ACP bugs seen in Data ONTAP 7.3.2 are in the following releases:

- ☐ Data ONTAP 7.3.3P2 and later
- ☐ Data ONTAP 8.0P1 and later

NetApp recommends using one of these releases if you are using DS4243 or DS2246 storage shelves.

WHAT DO THE VALUES FOR THE ACP OPTIONS ACP.DOMAIN AND ACP.NETMASK MEAN?

NetApp does not recommend changing your storage configuration's ACP configuration through the ACP options. The recommended approach to changing the ACP configuration is to run the `setup` command. There are three options for ACP:

- ☐ **acp.domain.** This is the decimal representation of the ACP domain (network). For example, if set to 108736, this converts to hexadecimal 0x1a8c0. The value "a8" equates to 192, and "c0" equates to 168 (in binary), which indicates that the ACP domain is set to 192.168.0.0.

- **acp.netmask.** This is the decimal representation of the ACP subnet mask. For example, if set to 16777215, this converts to hexadecimal 0x1ffffff. The value “ff” equates to 255 in binary and is repeated three times in the result, which indicates that the ACP subnet mask is set to 255.255.255.0.
- **acp.port.** This is the Ethernet port assigned to the ACP network for the controller.

IS THERE A SEPARATE FIRMWARE FILE FOR ACP?

SAS storage shelves using ACP should download both IOM firmware and ACP firmware updates. The ACP firmware update file should be placed in the `/etc/acpp_fw` directory before the `storage download acp` command is run. ACP firmware updates are nondisruptive.

3.8 DS14

IS THE DS14MK2 FC SHELF SUPPORTED IN DATA ONTAP 8.1?

The DS14mk2 FC shelf is supported in Data ONTAP 8.1 with the exception that it must use ESH4 shelf modules. The ESH2 shelf module is not supported in Data ONTAP 8.1 and later.

WHAT IS THE HIGH-EFFICIENCY POWER SUPPLY UNIT, AND WHERE CAN IT BE USED?

The HE PSU is intended for use with DS14Mk2 and DS14Mk4 storage shelves. All DS4243 storage shelves are shipped with HE PSUs. These PSUs are at least 10% more efficient than the non-HE PSUs. HE PSUs started shipping in storage shelves in mid-November 2009. The Site Requirements Guide available on the NetApp Support (formerly NOW) site contains additional power consumption information for the HE PSU.

HOW CAN I DETERMINE WHETHER MY STORAGE SHELF HAS HIGH-EFFICIENCY POWER SUPPLY UNITS?

To determine whether your shelf contains any HE PSUs, look at the environmental information. The following is an example of environmental information for an HE PSU:

Power Supply information by element:

- [1] Serial number: PMA929320000408 Part number: <N/A>
Type: HE Power Reset Capable (86)
Firmware version: <N/A> Swaps: 0
- [2] Serial number: PMA929320000348 Part number: <N/A>
Type: HE Power Reset Capable (86)

HE PSUs display that they are HE power-reset capable.

CAN I MIX DS14MK2 AT SHELVES WITH DS14MK2 FC OR DS14MK4 FC SHELVES IN THE SAME LOOP?

No. DS14mk2 AT shelves must be in a loop that does not contain FC shelves.

WHAT ARE THE DIFFERENCES BETWEEN LRC, ESH, ESH2, AND ESH4 SHELF MODULES?

The loop resiliency circuit (LRC) was introduced with the DS14 shelf. FC-AL loop speed over LRC to back-end storage was 1Gb/sec. The embedded switching hub (ESH) was the first generation of NetApp shelf controllers that had embedded intelligence to deal with malfunctioning drives and the ability to bypass them from the loop. FC-AL loop speed to back-end storage with ESH remained at 1Gb/sec.

The next-generation ESH, the ESH2, introduced faster loop speed and resiliency capabilities. In conjunction with DS14Mk2 shelves, FC-AL loop speed to back-end storage using ESH2 modules increased to 2Gb/sec.

The ESH4 was combined with the DS14Mk4 shelf to enable 4Gb/sec FC operation and resiliency features such as power control functionality for the drives using the active backplane on the DS14Mk4 shelf. Provided that 4Gb/sec drives and DS14Mk4 shelves are used, FC-AL loop speed to back-end storage using ESH4 is also 4Gb/sec.

Starting with Data ONTAP 7.3.x, there is no support for LRC, ESH, and DS14Mk1.

WHAT IS REQUIRED TO ACHIEVE 4GB/SEC BACK-END STORAGE LOOP SPEED WITH ESH4?

All components on the loop must run at 4Gb/sec before the effective loop speed reaches 4Gb/sec. This includes all shelves on a loop, all HBAs connected to shelves, and all drives. Data ONTAP 7.2.1 is the minimum release that supports 4Gb/sec back-end loop speed.

The following six components are required for 4Gb/sec operation:

- ☐ Data ONTAP 7.2.1 or later
- ☐ 4Gb/sec-capable HBA (X2052A, X2054A, X2055A, or embedded 4Gb/sec-capable QLogic 2462 found on FAS3040 and FAS3070)
- ☐ 4Gb/sec-capable shelf (DS14Mk4/ESH4/4Gb SFP)
- ☐ 4Gb/sec drives (X278 or X279)
- ☐ 4Gb/sec SFPs in HBA and shelf (X6539-R6-C)
- ☐ Shelf loop speed switch set to 4Gb/sec

WHICH DRIVES CAN BE USED IN THE DS14MK4 SHELF?

FC drives in the X23 series are usable only in the DS14Mk1 storage shelves. FC drives in the X27 series are universal drives supported in DS14Mk1, DS14Mk2, and DS14Mk4 storage shelves. In general, all drives supported for use in the DS14Mk2 shelf are also supported in the DS14Mk4 shelf.

WHAT ARE SOME BUSINESS BENEFITS OF 4GB/SEC BACK-END STORAGE PERFORMANCE?

Some of the benefits of 4Gb/sec back-end storage performance are as follows:

- ☐ It provides twice the maximum loop bandwidth of DS14Mk2 and ESH2.
- ☐ Higher bandwidth is available for heavy sequential workloads.
- ☐ It takes fewer HBAs/slots to achieve bandwidth needs.
- ☐ Achieving 4Gb/sec FC bandwidth helps enable maximum storage performance.

IN WHICH SHELVES CAN THE ESH4 MODULES BE USED?

The ESH4 module can be used in the DS14Mk4 shelf and the DS14Mk2 shelf.

The ESH4 module is not supported in the DS14Mk1 shelf. ESH4 technology enhancements are incompatible with the DS14Mk1 ops panel. When an RMA is needed for an ESH module in a DS14Mk1 shelf, an ESH2 module will be sent as a replacement.

WHAT IS THE DIFFERENCE BETWEEN AT-FC, AT-FC2, AND AT-FCX MODULES?

All three are ATA controllers. They are used in conjunction with DS14Mk2-AT shelves. The differences are as follows:

- ☐ AT-FC is used with R150 and R200 platforms. It is offered in single-path configurations only.
- ☐ AT-FC2 is used with R150 and R200 platforms. It is offered in single-path configurations only.
- ☐ AT-FCX is used for SATA on FAS and NearStore[®] platforms. It is offered in single-path, dual-path, and multipath HA configurations.

WHICH MODULES CAN BE USED WITH THE DS14MK4 SHELF?

The DS14Mk4 shelf supports only the ESH4 module.

WHAT DOES THE DRIVE POWER-CYCLING FEATURE IN DS14MK4 SHELVES HELP WITH, AND HOW IS IT ENABLED?

Individual drive power cycling, introduced with DS14Mk4 shelves and supported starting with Data ONTAP 7.2.1, is a new tool in the NetApp storage error recovery process. In Data ONTAP 7.2.1 and 7.2.2, this feature is off by default but can be enabled by changing the option `disk.powercycle.enable`. In Data ONTAP 7.3 and later, this feature is on by default.

WHICH MODULES CAN BE USED IN THE SAME LOOP?

NetApp highly recommends that all storage loops contain homogeneous shelf modules, because mixing shelf modules might result in unexpected issues.

WHAT IS THE MAXIMUM NUMBER OF DS14MK4 SHELVES SUPPORTED ON A LOOP FOR THE FAS2040?

The FAS2040 supports a maximum of six (6) shelves per loop.

HOW SHOULD LOOP SPEEDS BE CONFIGURED IN A MIXED SHELF/DRIVE ENVIRONMENT?

DS14Mk4 loop speed must be set to the speed supported by all other components in the loop. Loop speed is set with the loop speed switch on the rear panel of the DS14Mk4 or DS14Mk2 shelf.

Note: There is no loop speed switch on the DS14Mk1 shelf, because it supports only 1Gb/sec loop speed.

If the loop contains drives that are only 2Gb/sec capable, all storage shelves are required to be set to run at 2Gb/sec.

CAN 2GB/SEC SFPs BE USED IN 4GB/SEC LOOPS?

No. Although the 4Gb/sec SFP is tested and supported in 1Gb/sec, 2Gb/sec, and 4Gb/sec loops, the 2Gb/sec SFP is not supported in 4Gb/sec loops. With 4Gb/sec FCP connections, use only the 4Gb/sec SFP (MFG part number 332-00011).

Note: All onboard adapter FCP ports and shelf module ports ship with 4Gb/sec SFPs, even if their speeds are limited to 2Gb/sec.

4 RAID

WILL THE LOCATION OF A DISK DRIVE WITHIN A SHELF IMPACT THE RAID RECONSTRUCTION OR RAPID RAID RECOVERY TIME?

In a FC-AL shelf architecture the disk location within a shelf could have a very small impact on RAID reconstruction or Rapid RAID Recovery times. In the SAS shelf architecture used in NetApp storage shelves today the location of a disk within a shelf will certainly have no impact on RAID reconstruction or Rapid RAID Recovery times. The primary factors in any RAID reconstruction or Rapid RAID Recovery time to completion is a combination of drive rotational speed, drive capacity and available stack/loop bandwidth.

IF I MOVE A DISK DRIVE TO A SYSTEM RUNNING AN OLDER VERSION OF DATA ONTAP, WILL THE RAID LABELS NEED TO BE CHANGED?

Any disk drive that is a data or parity drive (for example, part of an aggregate) has a RAID label that is appropriate for that Data ONTAP release. Moving these types of drives to an earlier Data ONTAP release will most likely result in RAID label conflicts. Spares or unassigned disk drives contain a universal RAID label and as such can be freely moved between systems without encountering RAID label conflicts.

NetApp recommends only moving spares or unassigned drives between systems that are running dissimilar Data ONTAP versions.

WHAT IS THE MAXIMUM NUMBER OF RAID GROUPS I CAN HAVE IN A SINGLE AGGREGATE?

The maximum number of RAID groups a single aggregate can contain is 150 RAID groups. SyncMirror aggregates can contain double the number of RAID groups (300).

WHAT RAID RECONSTRUCTION OPTIONS ARE AVAILABLE THAT CAN MANAGE RAID RECONSTRUCTION TIME?

Using default RAID options is generally the right thing to do. However, there are circumstances that might warrant your changing those options. In storage configurations with a large number of spindles, these options can be adjusted to meet the needs of tier 1 type storage deployments. Adjusting these options can result in greater application performance during reconstructions, but with increased reconstruction time; or they can minimize reconstruction times at the expense of application performance. The following three options can be used to manage RAID reconstruction time:

- **raid.reconstruct.threads** (hidden). In Data ONTAP 7.3.2 and later, there is a hidden option that allows administrators to increase or decrease the number of RAID reconstruction threads that can be spawned on a system. NetApp recommends that you use the default setting for this option (default is 2). This option allows storage administrators to manage how efficiently corrective I/O competes for system resources (within the `raid.reconstruct.perf_impact` setting) with foreground I/O. Hidden options should only be set or changed under the direction of NetApp personnel.

- **raid.reconstruct.perf_impact.** The `raid.reconstruct.perf_impact` option manages the percentage of system resources (CPU, I/O bandwidth, and drive utilization) that corrective I/O can compete with foreground I/O over when the system is under load. This option has little to no impact on idle systems, because corrective I/O consumes idle system resources as needed. The default setting is medium. The option can be set to low (consume only idle and background resources) or high (corrective I/O and foreground I/O compete with each other for up to 90% of system resources).
- **raid.reconstruct.max_groups** (hidden). Adjusting the `raid.reconstruct.max_groups` option does not help an individual reconstruction to finish any sooner, but it can allow multiple reconstructions to complete sooner. The default setting for this is 2, with a maximum value of 8. For example, by default up to two RAID groups can reconstruct at a time. This means that up to four drives, two doubly degraded RAID groups can be in the process of reconstructing. If you have more than four failed drives in more than two RAID groups, you have to wait for the two initial RAID groups to complete reconstruction before any other RAID groups can reconstruct. Increasing this value allows additional RAID groups to reconstruct at the same time without the need to wait. However, reconstructing more than two RAID groups at a time has a significant impact on system performance. NetApp recommends using the default setting for this option. Hidden options should only be set or changed under the direction of NetApp personnel.

IF A DRIVE IS NOT ZEROED BEFORE RECONSTRUCTION, WILL IT BE ZEROED AS PART OF THE RECONSTRUCTION?

No. The console might report a status that says zeroing is in process during a reconstruction, but this is displayed only while labels are being removed from the spares. Reconstruction is a process that does not require the spares to be zeroed. If a spare is not zeroed, Data ONTAP does not zero the drive before starting the reconstruction, because it is not necessary.

WHAT ARE THE DEFAULT AND MAXIMUM RAID GROUP SIZES FOR SAS, FC, SATA, AND SSD DRIVES?

The following table outlines default and maximum RAID group sizes by drive type. NetApp recommends using RAID-DP for all storage configurations.

Data ONTAP 8.0.1 Default and Maximum RAID Group Size by Drive Type			
Drive Type	RAID Type	Default RAID Group Size	Maximum RAID Group Size
SSD	RAID-DP (default)	23 (21+2)	28 (26+2)
	RAID 4	8 (7+1)	14 (13+1)
SAS/FC	RAID-DP (default)	16 (14+2)	28 (26+2)
	RAID 4	8 (7+1)	14 (13+1)
SATA	RAID-DP (default)	14 (12+2)	20 (18+2)
	RAID 4	7 (6+1)	7 (6+1)

CAN RAID RECONSTRUCTION BE DISABLED?

No. RAID reconstruction is not a feature that can be enabled or disabled. If your customer has concerns around the impact of RAID reconstructions on foreground I/O, there are a number of RAID options that can be used to manage RAID reconstruction activity.

WHICH OPTIONS ARE AVAILABLE TO MANAGE RAID SCRUBS?

There are several options that can be used to configure and manage RAID scrubs:

- **raid.scrub.enable:** By default, RAID scrubs are enabled. RAID scrubs can be disabled by setting the option `raid.scrub.enable` to off. NetApp highly recommends keeping RAID scrubs enabled.
- **raid.scrub.duration:** RAID scrubs will run for six hours each week as indicated by the default value of 360 (minutes). NetApp recommends using the default value for this option. The option value is a numeric value that represents the number of minutes that will be used for the scrub duration. You can set the value to -1, which will result in the scrub operation running to completion.
- **raid.scrub.schedule:** This allows you to schedule RAID scrubs to occur more frequently than once per week. For systems using large-capacity SATA drives, NetApp recommends running RAID scrubs more than once per week. This option can also be used to schedule RAID scrubs to occur during off-peak hours to minimize system performance effect. This option specifies the weekly schedule (day, time, and duration) for scrubs started by the `raid.scrub.enable` option. By default scrubs initiate at 1 a.m.

on each Sunday. If an empty string is specified, it will delete the previously entered scrub schedule and revert to the default scrub schedule (1 a.m. on Sundays). The syntax for a single schedule entry is `duration[h|m]@weekday@start_time`. Multiple entries can be made using commas to separate each entry. If you do not specify duration in minutes or hours, then the `raid.scrub.duration` value will be used instead. Weekday is the day(s) when the scrub operation should occur (sun, mon, tue, wed, thu, fri, and sat are valid values). Start time is specified in a 24-hour format (0–23 are valid entries).

Examples of `raid.scrub.schedule`:

(single entry) `options raid.scrub.schedule 240m@tue@2`

(multiple entry) `options raid.scrub.schedule 8h@wed@13,10h@thu@23`

- **`raid.scrub.perf_impact`:** RAID scrubs are intended to have low effect on system performance, reflected in the default setting of low for this option. If you want RAID scrubs to be more aggressive, you can set this option to medium or high. NetApp recommends using the default value for this option, but in some storage configurations, for example, in archival environments, it might make sense to change this option. The option allows a user to manage how system resources will be consumed by the scrubbing process while the system is under load. Data ONTAP will always consume idle resources that are needed when conducting background operations (such as scrubs). A `low` setting means only use idle resources. `Medium` means allow foreground I/O and scrubbing to compete with up to 40% of the system's resources. A setting of high means allow foreground I/O and scrubbing to compete with up to 90% of the system's resources (note that this works just like the `raid.reconstruct.perf_impact` option, as described in [TR-3838: Storage Subsystem Configuration Guide](#)).
- **`raid.scrubbers` (hidden):** This options value represents the number of RAID groups that will be scrubbed at the same time (minimum value is 1, maximum is 4, and default is 4). NetApp recommends that you use the default setting for this option. Hidden options should only be set or changed under the direction of NetApp personnel.
- **`raid.scrub.threads` (hidden):** This allows administrators to increase or decrease the number of RAID scrub threads that can be spawned on a system. NetApp recommends that you use the default setting for this option (default is 4). This option allows storage administrators to manage how efficiently corrective I/O competes for system resources (within the `raid.scrub.perf_impact` setting) with foreground I/O. This option can be set to a minimum of 1 or a maximum of 8. Hidden options should only be set or changed under the direction of NetApp personnel.

5 SOFTWARE-BASED DISK OWNERSHIP

HOW DO I ENABLE SOFTWARE DISK OWNERSHIP?

A system can be converted to software disk ownership by using the `disk upgrade_ownership` command while in maintenance mode. This assigns ownership based on the controller configuration rules.

DOES AUTOMATIC ASSIGNMENT OF DRIVES WORK WITH SPLIT OWNERSHIP STACKS?

No. If within the same storage stack you have assigned drives to both controllers in an HA pair, then automatic assignment of drives using software-based disk ownership does not work. This is because there is no way to determine which controller should have the drives assigned.

HOW CAN I TELL IF SOFTWARE DISK OWNERSHIP IS ON?

It's probably easier to use the opposite approach to find out if software disk ownership is not enabled. You can tell that software disk ownership is not on by running the `disk show` command. If the resulting output contains a message stating that the option Show was not recognized, then software drive ownership is not enabled. Additionally, the command `storage show` tells you at the end of its output that SANOWN is not enabled.

WHICH PLATFORMS AND DRIVES SUPPORT SOFTWARE DISK OWNERSHIP?

Software disk ownership is on by default for the following platforms, including fabric MetroCluster (if supported):

- FAS270
- FAS2000 series

- ☐ FAS3040
- ☐ FAS3070
- ☐ FAS3100 and 3200 series
- ☐ FAS6000 and 6200 series
- ☐ Any V-Series system

It can be enabled for the following platforms:

- ☐ FAS900 series
- ☐ FAS3020
- ☐ FAS3050

Software disk ownership is supported for all DS14 FC shelves and on DS14Mk2-AT shelves with AT-FCX (CPLD version 24) modules, as well as DS4243 and DS2246 shelves.

DOES MPHA REQUIRE SOFTWARE DISK OWNERSHIP?

Yes. All HA pair controllers configured for MPHA must have software disk ownership turned on.

HOW DOES DISK_FW_UPDATE WORK ON A SOFTWARE DISK OWNERSHIP SYSTEM?

The `disk_fw_update` command updates the drive firmware only on those drives that are owned by the controller on which the command is entered. On an HA pair configuration the user must enter the command on both heads. Currently, unowned drives are not updated.

HOW DO I PERFORM A CONTROLLER HEAD SWAP WHEN USING SOFTWARE DISK OWNERSHIP?

The recommended method for doing a controller head swap assumes that the other controller head has done a takeover when the broken controller head failed. The user should run the `disk_reassign -d old-filersystemid` command in `priv set advanced` mode on the node that is in takeover.

If the other controller head is not in takeover mode, then boot the newly swapped head into maintenance mode and use the `disk_reassign -o old-filename` command. This command determines what the new serial number is and assigns it to the drives. When the node boots, it updates the partner node with the serial number information, provided it is in an HA pair configuration.

6 MULTIPATH HIGH AVAILABILITY (MPHA)

WHAT IS MPHA?

MPHA cabling means there are multiple and redundant connections from your controller to your storage.

IS MPHA CABLING REQUIRED FOR ALL NETAPP STORAGE SYSTEMS?

All NetApp HA pair systems now require the use of MPHA cabling, with just one exception. HA pair FAS2040 configurations only have a single SAS port and as such only support single-path connectivity to SAS storage shelves.

All NetApp single-controller systems are required to use dual-path connectivity, with the exception of the FAS2040 using SAS storage shelves (due to the single SAS data port on the system).

IS MPHA SUPPORTED WHEN USING ONLY A SINGLE HBA OR ONBOARD PORTS ONLY?

NetApp does support this, but continues to recommend using redundant components, because using only onboard ports or a single HBA does introduce a single point of failure into the storage configuration.

MPHA is supported for configurations using only a single HBA. NetApp recommends splitting connections between the HBA's ASICs. On a quad-port HBA, this generally means that ports A and B are using one ASIC, and ports C and D are using another. Therefore, ports should be paired across the ASICs on the HBA (for example, A and C ports and B and D ports). For SAS connections, see the Universal SAS Cabling Guide, available on the NetApp Support site.

Using only onboard ports is also supported. For example, the FAS3200 series includes two onboard SAS ports for storage connectivity. For systems using only a single stack of storage, NetApp does support using

only the two onboard ports for MPHA connectivity. NetApp still recommends using an HBA in addition to the onboard ports for additional hardware resiliency.

Although supported, you might see messages from Data ONTAP indicating that your MPHA storage connections are not redundant from a hardware perspective. If you choose to use this approach, you can ignore these messages.

HOW CAN I DETERMINE WHETHER MY MPHA CONFIGURATION IS PROPERLY CABLED IN ANY SUPPORTED CONFIGURATION?

A best practice cabling method has been developed that enables ASIC, HBA, and port isolation for all configurations. This method also enables proper top/bottom cabling for your stacks. NetApp recommends using this cabling method for all NetApp storage configurations. The [Universal SAS Cabling Guide](#) is the official source for all SAS MPHA cabling; it is available on the NetApp Support site. This same approach can also be used for FC cabling.

NetApp's [Wiregauge](#) tool is a software application that can be used to assess the current state of cabling on any deployed NetApp storage system.

WHICH SHELF MODULES ARE SUPPORTED WITH MPHA?

ESH2, AT-FCX (with CPLD version 24), ESH4, IOM3, and IOM6 shelf modules are supported with MPHA. LRC and ESH module shelves are not supported because they do not have autotermination.

WHICH VERSIONS OF DATA ONTAP ARE REQUIRED TO USE MPHA?

Data ONTAP 7.1.1 is required for FAS900 series, FAS3020, and FAS3050 platforms. Data ONTAP 7.2.1 or later is required for all other platforms (with the context of platform minimum supported release).

CAN I UPGRADE EXISTING NETAPP STORAGE DEPLOYMENTS IN THE FIELD TO MULTIPATH HA?

Provided that all the preceding conditions are met in the storage solution and enough ports are available on the controllers, then upgrading to multipath HA is straightforward. Multipath HA kits that include extra HBAs, cables, and SFPs can be ordered if additional parts are required.

Software disk ownership must be enabled on both controllers. If it is not currently enabled, boot both controllers into maintenance mode. Issue the `disk upgrade_ownership` command before adding the multipathing cables. The command takes a couple of seconds to run and assigns drives automatically.

If the controllers already use software disk ownership, then the customer might just need additional HBA cards for additional FCP port connectivity. Additional HBAs can be added nondisruptively by using a series of takeover/giveback actions, as outlined in the [HA Pair Configuration Guide](#).

ARE METROCLUSTER ENVIRONMENTS SUPPORTED WITH MULTIPATH HA?

Both stretch and fabric MetroCluster configurations are supported for use with multipath HA.

IS SOFTWARE-BASED DISK OWNERSHIP REQUIRED WITH MULTIPATH HA?

Yes.

CAN MPHA INCREASE PERFORMANCE?

Because MPHA doubles the number of connections from each controller to its primary storage, the available bandwidth increases, and performance improvement is possible. However, if the storage configuration was not loop/stack bandwidth bound, then MPHA probably will not help improve performance.

ARE THERE ISSUES WITH USING MULTIPATH HA ON SOME STACKS AND NOT ON OTHERS?

It is possible to multipath only some stacks on the storage system and not others. An EMS message is generated saying that some stacks are not multipathed. NetApp recommends using multipath for all stacks.

WHAT ARE THE SOME OF THE OTHER BENEFITS OF USING THE MULTIPATH HA CONFIGURATION?

Some benefits of using multipath HA are:

- ☐ Multipath HA provides added protection against storage faults, such as:
 - HBA or port failure

- Controller-to-shelf cable failure
- Shelf module failure
- Dual intershell cable failure
- Secondary path failure
- When running Data ONTAP 7.1.2 or 7.2.1 and later, multipath HA helps avoid storage stack instability from the failure of a single stack without the need for a takeover/giveback or system downtime.
- Data ONTAP mirrors bypasses on a single-HA configuration, and as a result failures are propagated to the alternate path. In this scenario, a multidrive panic due to the multidrive bypass leads to controller downtime. Running multipath HA helps avoid storage instability due to multidrive bypasses because bypasses are not mirrored on multipath HA systems.
- In the event of a takeover/giveback, stateful protocols such as CIFS and iSCSI experience a session disruption. From an application point of view, takeover/giveback time includes the time to restore the stateful sessions. Because multipath HA reduces the chance of a takeover/giveback, it reduces the number of session disruptions for stateful protocols running on the system.

DO I NEED MULTIPATH HA IF I HAVE CONTROLLER FAILOVER SOFTWARE INSTALLED?

The multipath HA storage configuration enhances the resiliency of HA pair controller configurations. Although controller failover software provides high availability by providing fault tolerance in the event of controller failure, storage-triggered events often result in unneeded failovers or prevent successful takeovers. Multipath HA storage enhances storage resiliency by reducing unnecessary takeover by a partner node due to a storage fault. This improves overall system availability and promotes higher performance consistency.

7 SYNCMIRROR

WHAT ARE SOME KEY SYNCMIRROR CONCEPTS?

Some key SyncMirror concepts are as follows:

- Only aggregates or traditional volumes can be mirrored with SyncMirror.
- SyncMirror cannot be implemented on flexible volumes, LUNs, qtrees, directories, or files.
- Mirroring is layered on top of RAID 4 or RAID-DP.
- Both plexes are updated synchronously on writes.
- Reads can be served from either plex, just like with RAID 1.
- There is no loss of data availability from a failure that brings down a plex. Data is still accessible from the remaining plex.

IS SYNCMIRROR SUPPORTED WITH DATA ONTAP 8.0.X OR 8.1.X CLUSTER-MODE?

No.

IS SYNCMIRROR SUPPORTED WITH THE DS4243 OR DS2246?

Yes. Both the DS4243 and the DS2246 support SyncMirror.

CAN I SET UP SYNCMIRROR BETWEEN MY FAS2000 SERIES INTERNAL DRIVES AND EXTERNAL DRIVES?

Yes. SyncMirror can be configured using FAS2000 internal drives and external drives.

IS SYNCMIRROR THE SAME AS RAID 1?

It is similar but not exactly the same. SyncMirror mirrors—or rather, the data on one plex that consists of one or more RAID 4, RAID-DP, or RAID 0 (V-Series only) RAID groups is duplicated onto another plex. RAID 1 involves duplicating the data on one drive onto a different drive. If one drive fails, the other continues to function as a single drive until the failed drive is replaced.

WHEN IS IT SAFE TO REDUCE THE AGGREGATE RESYNCSNAPTME OPTION?

This applies to automatic aggregate Snapshot™ reserve growth during mirror breaks. By default, the `aggr options resyncsnaptme` option is set to 60 (minutes). This option can be safely decreased to 15 minutes to reduce the duration of level-1 resynchronization events.

IS SYNCMIRROR A LICENSED PRODUCT?

Yes. The license is available at no extra cost, although activation does require a reboot.

IS SYNCMIRROR THE SAME AS SYNCHRONOUS SNAPMIRROR?

No. SyncMirror occurs at the RAID level, while SnapMirror® happens at the WAFL level. In general, SyncMirror provides higher levels of storage resiliency compared to synchronous SnapMirror.

WHAT DOES SYNCMIRROR PROTECT AGAINST THAT HA PAIRS DO NOT?

SyncMirror and RAID-DP layer in additional protection for HA pairs as follows:

HA pair availability protection:

- ☐ Storage processor failure
- ☐ Single-path failure
- ☐ Single-shelf module failure
- ☐ Single-HBA/port failure

RAID-DP availability protection:

- ☐ Single-drive failure
- ☐ Double-drive failure
- ☐ Single-drive failure plus uncorrectable bit read error that the drive ECC cannot correct

SyncMirror availability protection:

- ☐ Shelf failure (backplane, ops panel)
- ☐ Triple+ drive failure
- ☐ Dual-shelf failures (controllers, PSU/fans, cables)
- ☐ FC loop or SAS stack path failure

WHAT ARE THE PERFORMANCE EFFECTS OF SYNCMIRROR?

SyncMirror can improve reads significantly (very good for random read workloads) if the workload is drive bound by not having enough spindles. Writes can take slightly longer because data has to be written to both plexes in SyncMirror. Provided that the controller is not drive or CPU bound, the customer generally will not experience any write performance impact from SyncMirror.

ARE THERE SPECIAL CONSIDERATIONS TO KEEP IN MIND ABOUT AGGREGATE-LEVEL SNAPSHOT SPACE WHEN USING SYNCMIRROR?

SyncMirror uses the aggregate's reserve Snapshot space to synchronize mirrored plexes. When using SyncMirror, to make ample space available, other aggregate Snapshot schedules should be deleted, including those that occur by default. It is also important to monitor available aggregate Snapshot reserve space when using SyncMirror and, if necessary, to increase the aggregate Snapshot reserve space above the 5% default setting.

WHAT IS THE MAXIMUM NUMBER OF MIRRORED AGGREGATES OR TRADITIONAL VOLUMES WHEN USING SYNCMIRROR?

The maximum number of mirrored aggregates and traditional volumes that NetApp recommends is 64 per storage controller.

DO BOTH PLEXES HAVE TO BE IDENTICAL TO USE SYNCMIRROR?

Yes. The RAID group configuration and number of drives per RAID group must be identical.

WHAT IS THE MAXIMUM SIZE FOR EACH PLEX IN SYNCMIRROR?

The maximum size for each plex depends on the type of aggregate being used. For 32-bit aggregates, the maximum size is 16TB. 64-bit maximum plex size depends on the controller model being used and is equal to the maximum supported 64-bit aggregate size for the hosting controller.

8 MAINTENANCE CENTER

WHAT IS MAINTENANCE CENTER?

Maintenance Center is the newest addition to the storage resiliency software suite, which in the past introduced such storage resiliency features as RAID-DP and rapid RAID recovery.

HOW DOES MAINTENANCE CENTER WORK?

Maintenance Center performs proactive health monitoring of drives. When certain event thresholds are reached, it attempts preventive maintenance on the suspect drive. Maintenance Center takes advantage of onboard drive diagnostics to determine whether a drive error detected by Data ONTAP was a transient event or a real underlying drive issue. Drives that pass their onboard diagnostics are placed back in the spares pool, while true failures are flagged as such, and an RMA is triggered.

IS THERE AN OPTION TO ENFORCE MY SYSTEM TO BE IN MAINTENANCE CENTER COMPLIANCE?

Although there is not a specific Maintenance Center option, NetApp recommends that you set the RAID option `raid.min_spare_count` to a minimum value of 2 in order to be informed if there are insufficient spares available in the system for Maintenance Center. Maintenance Center also requires rapid RAID recovery to be enabled on the system (`raid.disk.copy.auto.enable` should be set to on).

IS MAINTENANCE CENTER CONFIGURABLE?

Yes. Maintenance Center can be configured, including being disabled, with the `options` command. NetApp recommends leaving Maintenance Center enabled with the default options settings.

HOW MANY SPARE DRIVES ARE REQUIRED TO TAKE ADVANTAGE OF MAINTENANCE CENTER?

Maintenance Center requires that a minimum of two equivalent spares be available for each drive attached to the system. Spares don't need to be an exact match; in several cases Data ONTAP considers some drives equivalent.

By default, SAS and FC drives are considered equivalent, although NetApp recommends that SAS and FC drives not be mixed in the same aggregate or RAID group. This can be disabled by setting the option `raid.disktype.enable` to off. Larger capacity drives of the same type (SATA, FC, SAS, and SSD) can replace smaller capacity drives.

WHICH VERSION OF DATA ONTAP IS REQUIRED TO USE MAINTENANCE CENTER?

Maintenance Center is enabled by default with Data ONTAP 7.1.0 and later.

9 DISK SANITIZATION

WHAT IS DISK SANITIZATION?

Disk sanitization is the process of obliterating stored data by means of overwriting patterns on NetApp system drives in a manner that prevents recovery of that data by any known recovery methods. The process consists of a drive format operation followed by the specified overwrite patterns repeated for the specified number of cycles. For SATA drives, the drive format operation is skipped because the `SCSI format` command is not supported on these drives.

DO I NEED A LICENSE TO RUN DISK SANITIZATION? HOW MUCH DOES IT COST?

A license is required to run disk sanitization; however, the license is free.

IS DISK SANITIZATION SUPPORTED IN DATA ONTAP OPERATING IN CLUSTER-MODE?

Yes. Data ONTAP operating in Cluster-Mode does include the disk sanitization feature.

WHY WOULD I WANT TO SANITIZE A DISK?

There are several reasons to sanitize a drive. For example, compliance or security protocols might require that drives that have been used to store sensitive data undergo a process that guarantees that the sensitive

data is not recoverable from the drive. Perhaps drives are being repurposed for a different system and need to be sanitized before being deployed. Whatever the reason, once a drive is sanitized, the data is not recoverable.

HOW DO I KNOW IF DISK SANITIZATION IS SUPPORTED ON A PARTICULAR DRIVE?

To determine whether disk sanitization is supported on a specific drive, run the `storage show disk` command. If the vendor of the drive in question is listed as NetApp, then drive sanitization is supported.

Disk sanitization support was introduced in Data ONTAP 8.1.1 for the NetApp 100GB SSD drive, X441A.

WHAT SECTORS OR DATA DOES DISK SANITIZATION OVERWRITE?

Disk sanitization overwrites the entire contents of the drive.

HOW DOES SANITIZATION DIFFER FROM STANDARD ZEROING?

With sanitization, any pattern can be written, including all zeroes. Random patterns can also be written. Zeroing a drive clears only the file data areas; it does not clear the entire drive.

HOW ARE BAD BLOCKS AFFECTED BY DISK SANITIZATION?

Bad blocks are removed from the grown defect list during formatting. If the bad block can be written to, then it is also sanitized.

WHAT ARE THE DIFFERENCES, IF ANY, BETWEEN FC/SAS AND SATA DISK SANITIZATION?

For SATA drives, the drive format operation is skipped because the `SCSI format` command is not supported.

IS IT ADVISABLE TO BYPASS THE FORMAT OPERATION ON FC/SAS DRIVES?

Yes. There is no real advantage to the format operation if there are no grown defects.

CAN YOU SANITIZE INDIVIDUAL FLEXVOL VOLUMES, LUNS, OR FILES?

No. You can only sanitize individual drives.

NetApp provides no representations or warranties regarding the accuracy, reliability, or serviceability of any information or recommendations provided in this publication, or with respect to any results that may be obtained by the use of the information or observance of any recommendations provided herein. The information in this document is distributed AS IS, and the use of this information or the implementation of any recommendations or techniques herein is a customer's responsibility and depends on the customer's ability to evaluate and integrate them into the customer's operational environment. This document and the information contained herein may be used solely in connection with the NetApp products discussed in this document.

Go further, faster®



© 2012 NetApp, Inc. All rights reserved. No portions of this document may be reproduced without prior written consent of NetApp, Inc. Specifications are subject to change without notice. NetApp, the NetApp logo, Go further, faster, Data ONTAP, FlexCache, FlexVol, MetroCluster, NearStore, NOW, RAID-DP, SnapLock, SnapMirror, Snapshot, SyncMirror, and WAFL are trademarks or registered trademarks of NetApp, Inc. in the United States and/or other countries. All other brands or products are trademarks or registered trademarks of their respective holders and should be treated as such.