



Technical FAQ

NetApp Flash Pool

Skip Shapiro, NetApp

Version 4.1

June 2013

Overview

This technical FAQ addresses the common questions about NetApp® Flash Pool™ with an emphasis on information that is not already covered in TR-4070: Flash Pool Design and Implementation Guide, which is a primary reference for this FAQ. The questions addressed cover a range of topics, including a basic understanding about Flash Pool, feature interoperability, performance, and more.

TABLE OF CONTENTS

1	GENERAL	5
	What is a Flash Pool aggregate?	5
	How does the read and write caching work in Flash Pool?	5
	What is the granularity of data movement?	5
	What type of application workloads are a good or bad fit for Flash Pool?	6
	Are the SSDs used in a Flash Pool aggregate considered storage capacity, or are they a cache?	6
	Can read and write caching be tuned?	6
	Can data movement policies be customized to specify how data is or is not cached?	6
	Can read or write caching be turned off for a specific Flash Pool aggregate?	6
	Can data be pinned in the SSD cache and never moved to the HDD portion of the aggregate?	6
	Can Flash Pool be disabled or removed from an aggregate?	7
	When and how often does the eviction scanner run?	7
	Is there a performance impact to the system when the eviction scanner is running?	7
	Does Data ONTAP use the TRIM command when managing the SSDs?	7
	How does wear leveling work with SSDs in a Flash Pool aggregate?	7
2	CONFIGURATION AND ADMINISTRATION	7
	What are the general requirements for using Flash Pool?	7
	How do I create a new Flash Pool aggregate?	8
	Can an HDD-only aggregate be converted into a Flash Pool aggregate?	8
	Can an all-SSD aggregate be converted into a Flash Pool aggregate by adding HDDs?	8
	Can a root aggregate be a Flash Pool aggregate?	9
	Can a Flash Pool aggregate be converted into an HDD-only Flash Pool aggregate?	9
	Which commands are used to manage Flash Pool?	9
	Can more than one HDD disk type be configured in a Flash Pool aggregate?	9
	Which SSD models are supported for use in a Flash Pool aggregate?	9
	How is Flash Pool cache size determined?	10
	How much Flash Pool cache capacity is available to cache data?	10
	Which command shows the used and available SSD cache space in a Flash Pool aggregate?	10
	Does the capacity of SSDs in a Flash Pool aggregate count toward the maximum aggregate storage capacity?	10
	Do the SSDs in Flash Pool aggregates count toward the maximum drive count for a system?	10
	Does the capacity of SSDs in Flash Pool aggregates count toward the maximum storage capacity of a system?	10
	Is a portion of configured Flash Pool cache capacity allocated for aggregate-level WAFL reserve?	11
	Is Flash Pool SSD cache capacity decremented by aggregate Snapshot reserve?	11
	What is the recommended minimum number of SSDs when configuring a Flash Pool aggregate?	11
	What is the recommended minimum number of SSDs to add when expanding Flash Pool cache?	11
	Can Flash Cache and Flash Pool be used on the same system?	12
	What is the maximum Flash Pool cache size supported on a single controller or HA pair system?	12
	How many Flash Pool aggregates are supported on a single system?	12
	What happens if the maximum supported cache capacity is exceeded on a system?	13
	What RAID type is recommended for the SSD RAID group in a Flash Pool aggregate for a system that is running Data ONTAP 8.1.1 or a later 8.1.x release?	13

What RAID type is recommended for the SSD RAID group in a Flash Pool aggregate for a system that is running Data ONTAP 8.2 or a later release?.....	14
What is the best practice for SSD RAID group size in a Flash Pool aggregate?.....	14
Are the SSDs and HDDs in a Flash Pool aggregate configured in the same RAID group?.....	14
Can a Flash Pool aggregate be configured with SyncMirror RAID protection?	15
What is the impact when a Flash Pool SSD RAID group is in degraded mode?	15
What is the impact when a Flash Pool SSD RAID group fails?	15
How many shelves containing SSDs can be configured in a single SAS stack?.....	15
Can different shelf types and drive types be configured in the same storage stack?	15
How many shelves containing SSDs can be configured in a single SAS stack?.....	16
Can a Flash Pool aggregate contain both SSDs and HDDs configured in a mixed drive shelf?.....	16
Does the FAS2240-2, FAS2240-4, or FAS2220 support mixed SSD + HDD or SSD-only internal storage options?.....	17
Are spare SSDs required for Flash Pool configurations?	17
Is the DS14mk2 AT shelf supported with Flash Pool?.....	17
Is Flash Pool supported with V-Series systems?.....	17
Why do storage systems support more SSDs than can be used for Flash Pool cache?	17
How can I determine the level of write wear on the SSDs in a Flash Pool aggregate?	17
Is NetApp Storage Encryption (NSE) or full disk encryption (FED) supported with Flash Pool?	18
3 FEATURE INTEROPERABILITY.....	18
Does Flash Cache reduce the amount of Flash Pool caching capacity allowed on a system?.....	18
Can FlexShare be used to prioritize specific volumes for caching in Flash Pool?	18
Does Flash Pool work with NetApp storage efficiency features?	18
How do Snapshot copies work with Flash Pool?.....	18
Does Flash Pool work with SnapMirror?	18
Does Flash Pool work with MetroCluster?	18
Can mixed SSD-HDD storage shelves be used in a MetroCluster configuration?	19
In a MetroCluster or SyncMirror configuration, do the SSDs in both plexes of a Flash Pool aggregate count toward the maximum Flash Pool cache size for the system?.....	19
Does SnapLock work with Flash Pool?	19
4 PERFORMANCE AND SIZING.....	19
How can Flash Pool performance be monitored?.....	19
How does Flash Pool affect system performance?	19
Can System Performance Modeler (SPM) be used to size system configurations with Flash Pool?	20
Can Predictive Cache Statistics (PCS) be used to estimate how much Flash Pool cache might be beneficial to add to a system?	20
Are quality of service (QoS) objectives used to manage Flash Pool cache data insertion and retention?.....	20
How soon will performance improve after converting an aggregate to a Flash Pool aggregate?	20
Do all flexible volumes in a Flash Pool aggregate benefit from Flash Pool cache?.....	21
Does system performance degrade as the SSDs in Flash Pool cache age?	21
Which counters in Data ONTAP will help determine if implementing Flash Pool is a good option?.....	21
Does adding Flash Pool accelerate random write performance?.....	21
5 FLASH CACHE AND FLASH POOL CACHE SIZE TABLES	22

6	REFERENCES	25
7	REVISION HISTORY.....	25

LIST OF TABLES

Table 1)	SSD part numbers, usable capacities and shelf support for support.	10
Table 2)	Recommended minimum number of SSD data drives per Flash Pool aggregate.	11
Table 3)	Recommended minimum number of SSDs to add when expanding Flash Pool cache.	11
Table 4)	Maximum number of SSD data drives in Flash Pool aggregate by maximum cache size value.	13
Table 5)	Default and maximum RAID group sizes for disk types.	14
Table 6)	Mixed stack storage configurations.	16
Table 7)	Data ONTAP 8.1.1 Flash Cache, Flash Pool, and combined maximum cache sizes.	22
Table 8)	Data ONTAP 8.1.2 Flash Cache, Flash Pool, and combined maximum cache sizes.	23
Table 9)	Data ONTAP 8.2 Flash Cache, Flash Pool, combined maximum cache sizes.	24
Table 10)	Revision information.	25

1 General

What is a Flash Pool aggregate?

Answer: A NetApp Flash Pool aggregate configures solid state drives (SSDs) and hard disk drives (HDDs)—either performance disk drives (often referred to by their interface, SAS or FC) or capacity disk drives (often called SATA)—into a single aggregate. The SSDs and HDDs are configured into separate RAID groups, and the SSD RAID group acts as a cache for data stored in volumes provisioned on the HDDs within the same Flash Pool aggregate. Deploying a hybrid SSD-HDD Flash Pool aggregate can improve application and storage system performance relative to HDD-only systems at a cost that is comparable or lower.

How does the read and write caching work in Flash Pool?

Answer: Caching random read data significantly improves read I/O response times for the volumes provisioned in a Flash Pool aggregate that have read caching enabled. The default Flash Pool cache insertion policy for reads copies random read operations of any block size into the cache immediately before that data is evicted from read buffer memory in the storage controller. There is no option to cache data from sequential read operations.

Data is inserted into Flash Pool cache using a Data ONTAP® consistency point (CP) write operation, which assures consistency between data in memory and data on storage. Flash Pool cache is SSD storage acting as a cache. The data in Flash Pool read cache is a copy of what is already stored on HDD, and thus the cache and stored copies are identical and consistent.

Flash Pool also can cache a certain type of random data that is written to volumes provisioned on a Flash Pool aggregate. If the default write cache policy is enabled, random write operations of 16KB size or smaller that are updates to previously written data (these are called overwrites) will be written into Flash Pool write cache based on the age of the previously written data. Caching small block overwrites is beneficial for workloads that consist of frequently updated data (for example, trading or online processing applications), because it avoids writing short-lived data to HDDs, which is inefficient. There is an additional benefit when the overwrite data is also read shortly after it was written; the data is read from SSD cache with the resultant low latency.

Unlike with cached reads, blocks that are cached as random overwrites are not copies of data that is stored on disk; the data in write cache is the only copy of that data. This is a key reason why Flash Pool SSD cache is configured with RAID protection.

Flash Pool keeps track of all cached data—whether it originates from read or write operations—to determine how frequently each 4KB block is accessed (in other words, how “hot” the data is). Data that is less frequently used is demoted in priority, and the least frequently used data is evicted from cache when room is needed for new hot data. Data in read cache is a copy of data on disk, so it is simply ejected. Data in write cache is written to disk before it is evicted from cache, because the cached data is the only copy that exists.

More information about Flash Pool cache policies and cache insertion and eviction processes can be found in TR-4070.

What is the granularity of data movement?

Answer: Data ONTAP stores data in WAFL® blocks that are 4kB in size. Flash Pool caches WAFL blocks (4kB).

What type of application workloads are a good or bad fit for Flash Pool?

Answer: Flash Pool is most effective at caching random reads, which offloads the HDD from these expensive operations. Likewise, Flash Pool provides benefit in offloading random writes (specifically, random overwrites) from disk as well. Any application workload that presents random reads or writes as a large part of its mix might benefit from Flash Pool. A key benefit to Flash Pool SSD cache is that in the event of a failover the data is persistent and therefore is immediately available to the node that has taken over the Flash Pool aggregate (no cache rewarming).

Are the SSDs used in a Flash Pool aggregate considered storage capacity, or are they a cache?

Answer: SSDs in a Flash Pool aggregate are persistent storage drives that are being used to cache data for the HDDs that are present in the same Flash Pool aggregate. The capacity of the SSDs does not count as part of the storage capacity of the aggregate; only the sum of the data HDD capacities counts as the aggregate's storage capacity.

Can read and write caching be tuned?

Answer: Yes. The `priority` command can be used to set caching policy values for a Flash Pool aggregate. This command requires advanced mode privileges for Data ONTAP operating in 7-Mode or can be run from the node-scoped CLI for clustered Data ONTAP. The syntax of the `priority` command is as follows:

```
priority hybrid-cache set <volume name> <read-cache>=<value> <write-cache>=<value>
```

Note: It is not required to set values for both the “read-cache” and “write-cache” parameters when executing this command.

Valid values for the “read-cache” parameter are as follows:

- “none”: disable read caching.
- “random-read”: (default) cache randomly read data as determined by Data ONTAP.
- “random-read-write”: cache random written blocks if they have not already been cached by the write cache of the Flash Pool aggregate.
- “meta”: all metadata blocks of the volumes in the aggregate plus directory blocks will be cached.

Valid values for the “write-cache” parameter are as follows:

- “none”: disable write caching.
- “random-write”: (default) cache random overwrites as determined by Data ONTAP.

Can data movement policies be customized to specify how data is or is not cached?

Answer: There are volume-level policies that allow you to set the kind of workload data that can be inserted into the SSD for reads and writes.

Can read or write caching be turned off for a specific Flash Pool aggregate?

Answer: Yes. The `priority` command can be used to set the “read-cache” or “write-cache” policy to the value “none,” which disables caching of either reads, writes, or both.

Can data be pinned in the SSD cache and never moved to the HDD portion of the aggregate?

Answer: No. There is no pinning of data in a Flash Pool SSD cache. Data is maintained in SSD based on use. For reads this means that recently or frequently accessed data has a higher chance of being cached

or remaining in cache. For writes this means that recently written (and overwritten) data has a higher chance of being cached or remaining in cache.

Can Flash Pool be disabled or removed from an aggregate?

Answer: The command “`priority hybrid-cache disable <aggr>`” (diag privilege command) can be used to disable Flash Pool on a specific aggregate. This should only be used for testing purposes. SSD RAID groups (like any RAID group in an aggregate) cannot be surgically removed from an existing aggregate. In order to remove Flash Pool from an existing aggregate, you would need to migrate the data, then destroy the aggregate, which would return the SSDs to the spares pool for the system.

When and how often does the eviction scanner run?

Answer: The eviction scanner “wakes up” every 60 seconds to determine if it needs to start running. If the eviction scanner is already running, it will continue to run; otherwise, it will start running if one of the following two conditions is met:

- The SSD cache capacity is 75% or more used.
- System heuristics determine that the SSD cache will exceed 75% used within the next 60 minutes.

If the scanner is not already running and neither of the preceding two conditions is met, the process remains inactive (goes back to sleep).

Is there a performance impact to the system when the eviction scanner is running?

Answer: The eviction scanner process is an extremely low-impact process. Extensive testing has shown that the eviction scanner process has no measurable impact to system performance. What might be visible to system performance is very active cache destaging or cache overruns that are associated to the working set size exceeding the available cache capacity (cache overrun).

Does Data ONTAP use the TRIM command when managing the SSDs?

Answer: No. The function of freeing up stale file system blocks is performed by a native background process within Data ONTAP. This process eliminates the need to use the TRIM feature.

How does wear leveling work with SSDs in a Flash Pool aggregate?

Answer: A characteristic of all flash-based SSDs is a specified finite number of cycles that data can be written to any cell within the media. Wear leveling is a standard feature of SSDs that assures data is written evenly across all the cells such that all cells wear at roughly the same rate. Wear leveling is done with SSDs used for Flash Pool cache just as it is for SSDs used as storage. For additional information on wear leveling and NetApp SSD, see [TR-3938: NetApp and Solid State Drive Technology](#).

2 Configuration and Administration

What are the general requirements for using Flash Pool?

Answer: There are system-level and aggregate-level requirements for using Flash Pool.

The system requirements for Flash Pool are as follows:

- Data ONTAP 8.1.1 or later operating in 7-Mode or clustered Data ONTAP 8.1.1 or later.
- Both controllers in the HA pair must be running the same version of Data ONTAP.

- Supported platforms:
 - FAS22x0: 2220, 2240-2, and 2240-4
 - FAS and V3xx0: 3220, 3240, 3250, and 3270; 3160 and 3170
 - FAS and V6xx0: 6210, 6220, 6240, 6250, 6280, and 6290; 6030, 6040, 6070, and 6080
- Note:** V-Series systems must use only NetApp SSD and HDD in a Flash Pool aggregate.
- Supported storage shelves: DS4243, DS4246, DS2246, DS4486, and DS14mk4 FC
- Note:** A Flash Pool aggregate with DS4486 shelves is not recommended as a best practice. DS4486 shelves are intended for secondary storage use cases that typically don't need nor benefit from storage caching.
- Note:** Flash Pool with DS14mk2 FC shelves is not supported with MetroCluster™ systems.
- Supported HDD disk types: performance (10,000 or 15,000 rpm) or capacity (7200 rpm)
- Note:** Flash Pool is not supported on systems that use NetApp Storage Encryption (NSE). Encrypting SSDs is currently not available with NetApp storage systems.
- Supported SSDs: 100GB, 200GB, and 800GB models (see Table 1 for part numbers and capacities).

The aggregate requirements for Flash Pool are as follows:

- Must be a 64-bit aggregate (cannot be a traditional volume, 32-bit aggregate, or striped aggregate).
 - Note:** In general, with Data ONTAP 8.1 and later releases, a 32-bit aggregate can be converted into a 64-bit aggregate, and the 64-bit aggregate can then be converted into a Flash Pool aggregate. However, there are situations in which a 64-bit aggregate that was converted from a 32-bit aggregate cannot become a Flash Pool aggregate. Refer to the *Storage Management Guide* for version of Data ONTAP in use to confirm whether an aggregate can be converted into a Flash Pool aggregate.
- Must be an HDD aggregate (cannot be an SSD aggregate or an aggregate using encrypting HDDs).
- The aggregate cannot be in a FAILED, LIMBO, offline, or foreign state.
- Supported RAID configurations: RAID 4, RAID-DP® technology, and SyncMirror® software (including local SyncMirror and MetroCluster).

How do I create a new Flash Pool aggregate?

Answer: A Flash Pool aggregate must be created with 64-bit addressing; that is, it must be a 64-bit aggregate. A Flash Pool aggregate is created first as an HDD-only aggregate, using System Manager or the "aggr create" command. Then, System Manager or the option "hybrid_enabled" is used to designate the aggregate as a Flash Pool aggregate. Finally, an SSD RAID group is added to the aggregate for the Flash Pool cache. The process of creating a Flash Pool aggregate is more completely described in TR-4070: Flash Pool Design and Implementation Guide.

Can an HDD-only aggregate be converted into a Flash Pool aggregate?

Answer: Yes. In fact, a Flash Pool aggregate is created by first creating an HDD-only aggregate and then converting the aggregate into a Flash Pool aggregate and adding an SSD RAID group. An HDD-only aggregate that is already provisioned with active volumes can be converted into a Flash Pool aggregate while those volumes are being accessed for read and write I/O.

Can an all-SSD aggregate be converted into a Flash Pool aggregate by adding HDDs?

Answer: No. In an all-SSD aggregate, the SSDs are used to store data. These SSDs cannot be used as Flash Pool cache without first migrating all the data from the SSDs to another aggregate, destroying the

all-SSD aggregate and releasing the drives to the spares pool. Only then can the SSDs be used in a Flash Pool aggregate.

Can a root aggregate be a Flash Pool aggregate?

Answer: Yes, provisioning node root volumes in a Flash Pool aggregate is supported with Data ONTAP 8.1.2 and later releases. However, this is not supported with Data ONTAP 8.1.1.

Can a Flash Pool aggregate be converted into an HDD-only Flash Pool aggregate?

Answer: No. The SSDs are provisioned as a RAID group, and a RAID group cannot be removed or destroyed within an active aggregate. If the objective is to reuse the SSDs that are currently part of a Flash Pool aggregate elsewhere in the same or another system, all data must first be migrated from the Flash Pool aggregate to a different aggregate. Then the source Flash Pool aggregate and all RAID groups are destroyed, and the SSDs and HDDs are returned to the spare drives pool.

Which commands are used to manage Flash Pool?

Answer: Systems running Data ONTAP operating in 7-Mode use the “`aggr create`” command to create a 64-bit aggregate (this is the normal method). The “`aggr option`” command is used to set the option “`hybrid_enable,`” which converts the aggregate into a Flash Pool aggregate. Adding SSDs to the aggregate into a new SSD-only RAID group is accomplished with the “`aggr add`” command. The status of the Flash Pool aggregate is shown in the output of the “`aggr status`” command. The “`priority hybrid-cache`” command allows you to modify and display Flash Pool read and write cache policies. The “`stats`” command is used to show Flash Pool performance and diagnostics counters.

For systems running clustered Data ONTAP use the “`storage aggregate`” command to create an aggregate (this is the normal method). The “`storage aggregate -modify`” command is used to set the “`-hybrid_enabled true`” option. SSD disks are added to an SSD-only RAID group in the Flash Pool aggregate using the “`storage aggregate add-disks`” command. Use the “`priority hybrid-cache`” command (node scope) to modify and display Flash Pool read and write cache policies.

Can more than one HDD disk type be configured in a Flash Pool aggregate?

Answer: No. A Flash Pool aggregate consists of SSDs and only one type of HDD: 7200 rpm, 10,000 rpm, or 15,000 rpm drives. Disk drives that have different rotational speeds cannot be used in a Flash Pool aggregate. The SSDs and HDDs are configured in separate RAID groups.

Note: SAS and FC disk drives that have the same capacity and rotational speed look as if they are the same HDD type; however, it is not recommended to the different interface drives in the same aggregate, including a Flash Pool aggregate. The reason for not using SAS and FC drives in the same aggregate is the SAS shelf and stack architecture is significantly more resilient than FC shelf and loop architecture.

Which SSD models are supported for use in a Flash Pool aggregate?

Answer: Table 1 lists the part numbers and corresponding right-sized usable capacities of the SSDs that are supported for use in a Flash Pool aggregate. The storage shelves with which each model is supported are also listed.

Table 1) SSD part numbers, usable capacities, and shelf support for support.

SSD Part Number and Capacity	Right-Sized Usable Capacity	Supported Shelves
X441A-R5 100GB	95,146MiB	DS4243 and DS4246
X446A-R6 200GB	190,532MiB	DS2246
X446B-R6 200GB	190,532MiB	DS2246
X448A-R6 200GB	190,532MiB	DS4246
X447A-R6 800GB	762,847MiB	DS2246

Note: The 800GB SSD is supported for Flash Pool use only with FAS/V62x0 controllers.

How is Flash Pool cache size determined?

Answer: The cache size of a Flash Pool aggregate is sum of the right-sized usable capacity of all the data SSDs in the aggregate. For example, the cache size of a Flash Pool aggregate that has an SSD RAID group of six 200GB data drives is 6 x 190,532MiB, which equals 1,243,192MiB (or 1,116GiB, or 1TiB). The same calculation method is used to determine the total Flash Pool cache size for a controller or HA pair. Table 1 lists the right-sized usable capacities of the SSDs that are supported for Flash Pool use.

Note: Cache and storage capacities in Data ONTAP are stated as binary (base2) values. [TR-3838: Storage Subsystem Configuration Guide](#) has more information about calculating capacities.

How much Flash Pool cache capacity is available to cache data?

Answer: Approximately 75% of the total size of a Flash Pool cache is available to cache data. For example, if the size of a Flash Pool cache is 1TiB, then roughly 0.75TiB (or 768GiB) is available to cache data. The remaining 25% of the cache capacity is used to store Flash Pool metadata and to assure consistent system performance as data is inserted into and evicted from a full cache.

Which command shows the used and available SSD cache space in a Flash Pool aggregate?

Answer: The “df -AH -hybrid” command shows the available and used SSD space in a Flash Pool aggregate.

Does the capacity of SSDs in a Flash Pool aggregate count toward the maximum aggregate storage capacity?

Answer: No. The SSD capacity is treated as cache capacity, and thus it does not count toward the storage capacity of the Flash Pool aggregate. Only the data HDDs in the Flash Pool aggregate count toward the maximum aggregate storage capacity.

Do the SSDs in Flash Pool aggregates count toward the maximum drive count for a system?

Answer: Yes. All drives, including SSDs that are used for cache capacity in Flash Pool aggregates, count toward the maximum supported drive count for a system.

Does the capacity of SSDs in Flash Pool aggregates count toward the maximum storage capacity of a system?

Answer: Yes. Because the SSDs in Flash Pool aggregates count toward the maximum drive count for a system, their presence in a system reduces the maximum amount of HDD storage capacity that can be configured on the system.

Is a portion of configured Flash Pool cache capacity allocated for aggregate-level WAFL reserve?

Answer: No. The SSDs in a Flash Pool aggregate are considered to be cache, and not storage; the cache capacity is excluded from space that is allocated for aggregate-level WAFL reserve.

Is Flash Pool SSD cache capacity decremented by aggregate Snapshot reserve?

Answer: No. The SSDs in a Flash Pool aggregate are considered to be cache, and not storage; the cache capacity is excluded from space that is reserved for aggregate Snapshot™ copies.

Note: Except for Flash Pool aggregates that are configured with local SyncMirror or MetroCluster, aggregate Snapshot copies should be disabled.

What is the recommended minimum number of SSDs when configuring a Flash Pool aggregate?

Answer: The recommended minimum number of data SSDs in a Flash Pool aggregate is a function of the controller model in the system. Although SSDs are capable of many more IOPS per drive than an HDD, a small number of them are required to assure that I/O to and from Flash Pool cache can be properly handled. Table 2 lists the recommended minimum number of data SSDs per Flash Pool aggregate to assure good performance.

Table 2) Recommended minimum number of SSD data drives per Flash Pool aggregate.

System Family	100GB SSD	200GB SSD	800GB SSD
FAS2200 series	1	1	Not supported
FAS/V3xxx series	3	2	Not supported
FAS/V6xxx series	9	5	5

Note: The 800GB SSD is supported for Flash Pool use only with FAS/V62x0 controllers.

What is the recommended minimum number of SSDs to add when expanding Flash Pool cache?

Answer: The recommended minimum number of SSDs to add when expanding a Flash Pool cache is a function of controller model with the goal of assuring good cache performance. Table 3 lists the recommended minimum number of data SSDs to add when expanding a Flash Pool cache.

Table 3) Recommended minimum number of SSDs to add when expanding Flash Pool cache.

System Family	100GB SSD	200GB SSD	800GB SSD
FAS2200 series	1	1	Not supported
FAS/V3xxx series	3	2	Not supported
FAS/V6xxx series	6	3	3

Note: The 800GB SSD is supported for Flash Pool use only with FAS/V62x0 controllers.

Note: Although it is common to use reallocate after expanding an HDD-only aggregate, it is not necessary to run reallocate on the SSD RAID group in a Flash Pool aggregate when adding SSDs. Running reallocate on the HDD RAID groups in a Flash Pool aggregate after adding HDDs might be advisable; follow the guidelines in the Storage Subsystem Tech FAQ for when to run reallocate on HDDs.

Can Flash Cache and Flash Pool be used on the same system?

Answer: Yes. Flash Cache™ modules and Flash Pool aggregates can be used on the same controller or one on each node of an HA pair system. There are specific configuration requirements when using both caching products on the same system that are covered in later in this FAQ. Blocks from volumes that are provisioned in a Flash Pool aggregate are only cacheable in Flash Pool cache; they are automatically excluded from using Flash Cache when Flash Cache modules are configured on the same system.

Note: While use of both Flash Cache and Flash Pool on the same system is supported, the best practice recommendation is to use only one caching product per controller or system, whenever possible. Flash Cache and Flash Pool each use memory and CPU, and thus combining both cache products uses more of these resources than using only one product. Use of Flash Cache and Flash Pool should be avoided on systems that have a maximum per node cache size that is less than 2.0TiB.

What is the maximum Flash Pool cache size supported on a single controller or HA pair system?

Answer: The maximum supported Flash Pool cache capacity varies by platform controller model and Data ONTAP release. Maximum cache sizes for Flash Cache, Flash Pool, and the combined use of Flash Cache and Flash Pool by Data ONTAP release can be found in the three tables in section 5 of this FAQ.

The total Flash Pool cache capacity configured on a controller or HA pair system is the sum of the right-sized usable capacities of all data SSDs in all Flash Pool aggregates assigned to the controller or present on the HA pair system. When multiple Flash Pool aggregates are configured on the same controller or HA pair system, each one can have a different cache size and different HDD storage capacity.

When both Flash Cache and Flash Pool are configured on a single controller or HA pair system, the combined total configured cache capacity must not exceed the combined per controller, or the combined per HA pair, maximum cache values as shown in the tables in section 5.

Note: The amount of Flash Pool cache configured on each controller in an HA pair system can be different; however, the cache size on a controller must not exceed the per controller maximums. Per controller Flash Pool cache maximum values apply to systems running Data ONTAP 8.1.1 and later 8.1.x releases. For systems running Data ONTAP 8.2 and later releases, there is only an HA pair cache maximum value; that means up to the maximum supported cache can be configured on one controller within the HA pair.

How many Flash Pool aggregates are supported on a single system?

Answer: There are no specified maximum numbers of Flash Pool aggregates that can be configured on a system, although maximum can be calculated based on the recommended minimum number of data SSDs per Flash Pool aggregate (as shown in Table 2) and the maximum number of data SSDs derived from the maximum Flash Pool cache capacity per system (as shown in Table 4).

However, to assure best utilization and storage efficiency from the SSDs in Flash Pool aggregates, the recommendation is to configure a small number of Flash Pool aggregates, each of which has relatively large HDD storage capacity. This results in more of the configured SSDs being used as data drives and larger Flash Pool caches for volumes to use.

Table 4) Maximum number of SSD data drives in Flash Pool aggregate by maximum cache size value.

Flash Pool Max Cache Size	Number of 100GB Data SSDs	Number of 200GB Data SSDs	Number of 800GB Data SSDs
300GiB	3	1	Not supported
400GiB	4	2	Not supported
0.5TiB	5	2	Not supported
0.6TiB	6	3	Not supported
1.0TiB	11	5	Not supported
1.2TiB	13	6	Not supported
1.6TiB	17	8	Not supported
2.0TiB	22	11	Not supported
4.0TiB	44	22	5
6.0TiB	66	33	8
12.0TiB	132	66	16
18.0TiB	198	99	24
24.0TiB	240 ¹	132	32

Note: (1) The maximum number of SSDs supported on midrange and high-end systems is 240.

What happens if the maximum supported cache capacity is exceeded on a system?

Answer: Data ONTAP does enforce cache capacity maximums. However, when both Flash Cache and Flash Pool are used on the same system, there are circumstances when the cache capacity maximum can be exceeded. If the maximum cache capacity for a controller is exceeded for, EMS messages similar to the following will periodically appear on the system.

```
[localhost:raid.hybrid.SSDTotExceed:error]: The sum of sizes of SSD disks of
hybrid aggregates <actual capacity> exceeds the <HA pair limit> maximum.

[localhost:raid.hybrid.SSDHAExceed:error]: The sum of all sizes of SSD disks
of hybrid aggregates <actual capacity>, exceeds <node limit> (half of the
maximum in an HA pair). If the sum of all sizes of SSD disks of hybrid
aggregates on this node and its partner exceeds the <HA pair limit> maximum,
takeover might fail.
```

What RAID type is recommended for the SSD RAID group in a Flash Pool aggregate for a system that is running Data ONTAP 8.1.1 or a later 8.1.x release?

Answer: For systems that run Data ONTAP 8.1.1 and later 8.1.x releases, the SSD and HDD RAID groups in a Flash Pool aggregate must have the same RAID policy, which means RAID protection type and RAID group size. Since RAID-DP is the recommended RAID type for HDDs, the SSD RAID groups in Flash Pool aggregates will also be RAID-DP protected. The SSD RAID group will be the same size, in terms of the number of data and parity drives, as was applied to the HDD RAID groups.

Note: The option "raid.min_spare_count" does not apply to SSD RAID groups in a Flash Pool aggregate. This enables a FAS2200 series that has a Flash Pool aggregate configured with RAID-DP protection to operate with no SSD spare, without generating a low spares error message. However, the absence of a hot SSD spare means Flash Pool performance will be affected if an SSD fails until the failed SSD can be replaced and the RAID group rebuild completes.

Note: The RAID group size for an HDD-only aggregate that is being converted into a Flash Pool aggregate can be increased before configuring the SSD RAID group. However, the maximum size of the RAID group is governed by the maximum size supported for the type of HDD in the aggregate (for example, 20 drives for a SATA RAID-DP group). In addition, if an HDD RAID group is added after the larger SSD RAID group was configured, the updated larger group size will apply to the new HDD RAID group. Following the best practice of configuring all HDD RAID groups to the same size (plus or minus one drive) in the aggregate requires manual attention when adding new HDD RAID groups.

Table 5) Default and maximum RAID group sizes for disk types.

Drive Type	RAID Type	Default RAID Group Size	Maximum RAID Group Size
SAS, FC, or SSD	RAID-DP (default)	16 (14 data + 2 parity)	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)

What RAID type is recommended for the SSD RAID group in a Flash Pool aggregate for a system that is running Data ONTAP 8.2 or a later release?

Answer: For systems that run Data ONTAP 8.2 or a later release, the SSD and HDD RAID policies (RAID protection type and group size) can be different. This means a Flash Pool aggregate can be configured with the HDDs in RAID-DP groups—per best practice—while RAID 4 is used for the SSD RAID group. Or, if RAID-DP is used for both the SSD and HDD RAID groups, the group sizes for each drive type can be different; in particular, the SSD RAID group can be larger.

RAID 4 protection for the SSDs in a Flash Pool aggregate is recommended when only one SSD RAID group per Flash Pool aggregate meets the cache capacity needs and at least one hot spare SSD per controller is configured to assure immediate RAID rebuild should an SSD fail. For example, on a FAS2200 system that has one Flash Pool aggregate, a two data drive RAID 4 SSD RAID group plus one hot spare is a better configuration than a two data drive RAID-DP SSD RAID group and no hot spare.

Nevertheless, for best resiliency and performance and lower emphasis on cost implications, the best practice recommendation for SSD configuration in a Flash Pool aggregate is RAID-DP protection with at least one hot SSD spare.

What is the best practice for SSD RAID group size in a Flash Pool aggregate?

Answer: The smallest RAID group size is based on the recommended minimum number of data SSDs per aggregate (refer to Table 2). When RAID 4 protection is used, the RAID group will have one drive more than the recommended minimum number of data SSDs. When RAID-DP is used, the group size will be two drives larger than the recommended minimum number of data SSDs.

The best practice recommendation is to expand the SSD RAID group up to maximum number of drives supported before adding a second SSD RAID group.

Note: When using SSDs with RAID 4 protection, at least one hot SSD spare per controller must always be configured to assure adequate data protection should an SSD fail.

Are the SSDs and HDDs in a Flash Pool aggregate configured in the same RAID group?

Answer: No. SSDs and HDDs are different drive types and must always be configured in separate RAID groups.

Can a Flash Pool aggregate be configured with SyncMirror RAID protection?

Answer: Yes. The RAID groups in both plexes of a Flash Pool aggregate that is protected with SyncMirror must be configured identically. The SSD RAID groups in each plex must have the same RAID type, drive count, and drive capacity; the same is true for the HDD RAID groups in each plex. The SSD and HDD RAID policies can be different from each other if the system is running Data ONTAP 8.2 or a later release. If the system is running Data ONTAP 8.1.1 or a later 8.1.x release, the SSD and HDD RAID groups must have the same RAID policies.

What is the impact when a Flash Pool SSD RAID group is in degraded mode?

Answer: A degraded SSD RAID group affects the performance of the Flash Pool cache; for example, every cache read will require parity reconstruction. The rebuild of an SSD RAID group is very fast (typically less than 30 min) compared to HDD RAID rebuilds, which can take many hours. However, if no SSD hot spare is available on the controller, Flash Pool cache performance will be adversely affected until the failed SSD is replaced and RAID rebuild is completed. Consequently, the best practice recommendation is to provision at least one SSD hot spare per controller.

What is the impact when a Flash Pool SSD RAID group fails?

Answer: If the SSD RAID group in a Flash Pool aggregate fails, the entire Flash Pool aggregate fails and is taken offline. A RAID group failure consists of two (or more) drives in a RAID 4 protected group, or three or more drives in a RAID-DP protected group, simultaneously either having uncorrectable errors or being in a failed state. However, if the SSD RAID group fails (exceeding the RAID type's ability to protect the data in the RAID group), the entire aggregate will be in a failed state and unable to serve data. Using RAID-DP is the best protection against the possibility of this failure event.

How many shelves containing SSDs can be configured in a single SAS stack?

Answer: A maximum of three shelves containing SSDs can be configured in a single storage stack with FAS22x0 systems. A maximum of four shelves containing SSDs can be configured in a single SAS stack for all systems that support Flash Pool. These stack shelf limits apply regardless of how many SSDs are configured in each shelf.

Can different shelf types and drive types be configured in the same storage stack?

Answer: Yes, the requirements and recommendations that apply to storage stacks that are configured with different shelf and drive types are covered in Table 6. These rules and guidelines apply equally when configuring mixed storage stacks for storage or Flash Pool use.

Note: The DS4486 is not listed because it does not support any mixed stack or drive configurations.

Table 6) Mixed stack storage configurations.

Shelves	DS4246 SATA	DS4243 SSD	DS4246 SATA + SSD	DS2246 SAS	DS4243 SATA	DS4243 SAS	DS2246 SSD
DS4246 SATA	Yes	Yes ^{1, 2, 4}	Yes ^{1, 2, 4}	Yes ¹	Yes ¹	Yes ¹	Yes ^{1, 4}
DS4243 SSD	Yes ^{1, 2, 4}	Yes ⁴	Yes ^{1, 2, 4}	Yes ^{1, 2, 4}	Yes ^{1, 2, 4}	Yes ^{1, 2, 4}	Yes ^{1, 2, 4}
DS4246 SATA + SSD	Yes ^{1, 2, 4}	Yes ^{1, 2, 4}	Yes ⁴	Yes ^{1, 2, 4}	Yes ^{1, 2, 4}	Yes ^{1, 2, 4}	Yes ^{1, 2, 4}
DS2246 SAS	Yes ¹	Yes ^{1, 2, 4}	Yes ^{1, 2, 4}	Yes	FAS22x0 ³	FAS22x0 ³	Yes ^{1, 4}
DS4243 SATA	Yes ¹	Yes ^{1, 2, 4}	Yes ^{1, 2, 4}	FAS22x0 ³	Yes	Yes ¹	FAS22x0 ^{3, 4}
DS4243 SAS	Yes ¹	Yes ^{1, 2, 4}	Yes ^{1, 2, 4}	FAS22x0 ³	Yes ¹	Yes	FAS22x0 ^{3, 4}
DS2246 SSD	Yes ^{1, 4}	Yes ^{1, 2, 4}	Yes ^{1, 2, 4}	Yes ^{1, 4}	FAS22x0 ^{3, 4}	FAS22x0 ^{3, 4}	Yes ⁴

Note: (1) The best practice recommendation is no more than one transition within a stack between shelves that are configured with different drive and shelf module types.

Note: (2) The best practice recommendation is that any shelf configured with SSDs should be located as the “first/top” or “last/bottom” shelf in the stack. This enables the most direct data path from SSDs to the controller.

Note: (3) FAS2240-4, FAS2240-2, and FAS2220 systems using these shelves as external storage can mix shelf types but must not exceed more than one transition between shelf types within the storage stack. This is a requirement, not a recommendation.

Note: (4) Maximum stack depth (the maximum number of shelves supported in one stack) is shallower (lower) for stacks that contain shelves configured with SSDs. The SSD stack limit for the FAS2220, FAS2240-2, and FAS2240-4 is 3 external shelves; for all other platforms the limit is 4 shelves. This is a requirement, not a recommendation.

How many shelves containing SSDs can be configured in a single SAS stack?

Answer: A maximum of three shelves containing SSDs can be configured in a single storage stack with FAS22x0 systems. A maximum of four shelves containing SSDs can be configured in a single SAS stack for all systems that support Flash Pool. These stack shelf limits apply regardless of how many SSDs are configured in each shelf.

Can a Flash Pool aggregate contain both SSDs and HDDs configured in a mixed drive shelf?

Answer: Yes. In fact, mixed SSD and HDD shelves are supported and offered explicitly for use with Flash Pool; they are not intended for use where the SSDs and HDDs will be used in standard (other than Flash Pool) aggregates. The DS4246 disk shelf supports mixed configurations of SSDs and 7200 rpm HDDs. The DS2246 shelf supports mixed configurations of SSDs and 10,000 rpm HDDs. The DS4243 and DS4486 shelves do not support mixed SSD and HDD configurations.

Are FAS2200 system configurations that have a mix of internal SSDs and HDDs supported?

Answer: Yes. As of May 2013, the following mixed internal drive configurations are available as orderable FAS2200 systems:

- FAS2220 HA with 4x 200GB SSD and 8x 900GB 10k rpm HDD
- FAS2220 HA with 4x 200GB SSD and 8x 2TB 7.2k rpm HDD
- FAS2240-2 HA with 4x 200GB SSD and 20x 900GB 10k rpm HDD
- FAS2240-4 HA with 4x 200GB SSD and 20x 2TB 10k rpm HDD
- FAS2240-4 HA with 6x 100GB SSD and 18x 1TB 7.2k rpm HDD
- FAS2240-4 HA with 6x 100GB SSD and 18x 3TB 7.2k rpm HDD

Are spare SSDs required for Flash Pool configurations?

Answer: Typically, yes. For most system configurations, at least one SSD hot spare per system—preferably one hot spare per node (controller)—is recommended. For FAS2200 series systems, which support a small SSD cache for Flash Pool, it is acceptable to configure the SSD RAID group with RAID-DP and not have a SSD hot spare. However, if there is no hot spare when an SSD fails, it is likely that Flash Pool cache performance will be worse until the failed SSD is replaced and the RAID group has been rebuilt with the replacement drive. Systems that use Flash Pool aggregates with RAID 4 protection for the SSD RAID groups should always maintain a least one SSD hot spare.

Is the DS14mk2 AT shelf supported with Flash Pool?

Answer: No. The DS14mk2 AT shelf is not supported for use with Flash Pool.

Is Flash Pool supported with V-Series systems?

Answer: Yes. Flash Pool is supported on V-Series systems; however, only NetApp SSDs and HDDs can be configured in a Flash Pool aggregate. Neither SSDs nor HDDs from third-party storage arrays can be used in a Flash Pool aggregate.

Why do storage systems support more SSDs than can be used for Flash Pool cache?

Answer: SSDs can be used for storage in a homogeneous SSD aggregate or as cache within a Flash Pool aggregate that uses HDDs for storage. The number of SSDs supported for Flash Pool cache use is generally lower than the maximum allowed for storage. Flash Pool cache uses system memory resources that regular storage-only aggregates do not, and this is why the SSD maximums for Flash Pool cache use are lower. On a system that uses the maximum allowed SSDs for Flash Pool cache, additional SSDs can be used in SSD-only aggregates up to the maximum number of SSDs supported with the controller platform.

How can I determine the level of write wear on the SSDs in a Flash Pool aggregate?

Answer: The “`storage show disk -a`” command reports write wear information for all SSDs attached to the system.

Is NetApp Storage Encryption (NSE) or full disk encryption (FED) supported with Flash Pool?

Answer: No. There are currently no encrypting SSDs qualified. Mixing nonencrypting and encrypting drives on the same storage system is currently unsupported.

3 Feature Interoperability

Does Flash Cache reduce the amount of Flash Pool caching capacity allowed on a system?

Answer: Yes. When Flash Cache and Flash Pool are combined on the same controller, the combined capacity of both solutions must fall within the cache capacity limits for the platform (covered by another topic in this FAQ in detail).

Can FlexShare be used to prioritize specific volumes for caching in Flash Pool?

Answer: NetApp does not recommend using FlexShare® with Flash Pool because the Flash Pool cache policy options enable per volume cache policies. The “`priority hybrid-cache`” command is used to change Flash Pool read and write caching policies. FlexShare should not be used with any system that has a Flash Pool aggregate. The FlexShare priority setting has no effect on Flash Pool caching activities; however, the use of FlexShare volume priorities can limit the effectiveness of Flash Pool cache.

Does Flash Pool work with NetApp storage efficiency features?

Answer: All storage efficiency features are supported for use with Flash Pool. Deduplicated and cloned blocks that match Flash Pool cache insertion policy will be inserted into Flash Pool cache. Flash Pool is supported with data compression; however, only uncompressed blocks are eligible for cache insertion. compressed blocks are part of a compression group, and compression groups are not cached.

How do Snapshot copies work with Flash Pool?

Answer: All of the standard Snapshot functionality applies and works with Flash Pool. However, data blocks that belong only to a Snapshot copy are not inserted into Flash Pool cache. Therefore, access to data that exists only in a Snapshot copy, and not also in the active file system, is not accelerated.

Does Flash Pool work with SnapMirror?

Answer: Yes. Data cached from a source volume that is provisioned in a Flash Pool aggregate is not transferred as part of a SnapMirror® replication, even if the destination volume is also provisioned in a Flash Pool aggregate. The destination volume can be provisioned in an HDD-only or Flash Pool aggregate; there is no requirement that the destination volume be in a Flash Pool aggregate. Asynchronous and synchronous SnapMirror replication are both supported with Flash Pool.

Note: I/O that results from a SnapMirror replication operation does not affect the data content of a Flash Pool cache.

Does Flash Pool work with MetroCluster?

Answer: Yes. Flash Pool is supported with MetroCluster. The following list describes important considerations and requirements that are specific to deploying Flash Pool with MetroCluster:

- Mixed stack (and shelf) support is the same for both MetroCluster and standard HA systems.
- Storage stack depth when using SSDs (containing SSD-only or mixed SSD-HDD shelves) is four disk shelves (which is exactly the same as with standard HA systems).

- DS14mk4 FC disk shelves and FC disk drives are not supported with Flash Pool in MetroCluster configurations.

Can mixed SSD-HDD storage shelves be used in a MetroCluster configuration?

Answer: Yes. With the exception of FC drive support, MetroCluster currently supports all combinations of storage shelves (stack and shelf configurations) that are supported with standard HA systems.

In a MetroCluster or SyncMirror configuration, do the SSDs in both plexes of a Flash Pool aggregate count toward the maximum Flash Pool cache size for the system?

Answer: No. The second plex of a mirrored aggregate does not count toward the Flash Pool cache limit for a platform.

Does SnapLock work with Flash Pool?

Answer: No. A Flash Pool aggregate cannot use the SnapLock[®] software feature. Likewise, an aggregate with SnapLock data may not be converted into a Flash Pool aggregate.

4 Performance and Sizing

How can Flash Pool performance be monitored?

Answer: The “stats” command shows Flash Pool performance and diagnostics counters for systems that are running Data ONTAP operating in 7-Mode. The “stats” command is used from the node CLI for systems that are running clustered Data ONTAP. A stats preset “stats show -p hybrid_aggr” shows interactive statistics from the command line.

How does Flash Pool affect system performance?

Answer: Flash Pool improves the average response time (latency) of operations for all volumes provisioned in a Flash Pool aggregate (when the default read cache setting is used). Randomly read blocks are inserted into Flash Pool cache after first access and before eviction from main memory read buffers; that allows subsequent read requests to be served from cache instead from HDD. Random writes are inserted into Flash Pool cache if they are 16KB or smaller in size, and they arrive soon after a previous write of the same logical blocks (where soon enough is a function of how actively all the HDD blocks in the Flash Pool aggregate are being accessed). Caching of random overwrites avoids consuming HDD write operations for data that is likely to change again soon thereafter. The primary goal when using Flash Pool is to achieve equal or better individual workload and aggregate storage system performance with a lower cost storage system configuration. Realizing this goal typically means replacing a storage system that is configured with many performance (high-rpm) HDDs with one that is configured with a small quantity of SSDs and fewer HDDs (and often slower and lower cost per GB HDDs as well).

Does Flash Pool increase the maximum throughput capability of a storage controller?

Answer: No, using Flash Pool on a system will not increase the maximum throughput capability of the controller. The maximum throughput capability of a controller, for a given mix of workloads, is determined by the controller’s memory, CPU cores, and the version of Data ONTAP. Implementing Flash Pool on a system where throughput performance is constrained by the number and type of configured HDDs will often increase throughput performance; however, the maximum performance is still bounded by controller memory and CPU resources.

Note: A system on which performance is constrained by memory or CPU utilization will not benefit from implementing Flash Pool. Upgrading to a more capable controller or moving some

workloads to another system is the recommended course of action to remediate performance constraints due to high memory or CPU heavy utilization.

Can System Performance Modeler (SPM) be used to size system configurations with Flash Pool?

Answer: Yes. An option to size with Flash Pool is provided in SPM. The output of the sizing will compare the configuration with Flash Pool to a configuration that has no storage cache (that is neither Flash Pool nor Flash Cache).

Note: In SPM, the field Flash Pool capacity requests the total amount data SSD cache capacity per controller. SPM treats the value as fixed input; it does not increase the cache size if more cache per controller (up to the supported maximum) would be beneficial, which could result in a sizing output with more controllers than required for performance or capacity reasons. In addition, SPM does assure the cache capacity that is input does not exceed the maximum allowed for the selected controller or system; this must be done manually by checking the table using the link in SPM or the tables in section 5 of this FAQ.

Can Predictive Cache Statistics (PCS) be used to estimate how much Flash Pool cache might be beneficial to add to a system?

Answer: PCS projects the performance benefit of adding Flash Cache modules to a controller that currently has no extended cache based on the workloads the controller is servicing. PCS only considers the benefit of caching data for read requests, and it was designed specifically for Flash Cache prediction. Therefore, the effectiveness of PCS in estimating the benefit of converting an HDD-only aggregate to a Flash Pool aggregate is limited. PCS can be used to get a rough estimate of the read cache benefit of Flash Pool for a system that isn't configured with Flash Pool or Flash Cache if all the aggregates will be converted to Flash Pool aggregates or if the aggregates that are under consideration for conversion are known to contain volumes that have only sequential read/write workloads. Follow these guidelines when using PCS to estimate Flash Pool benefits:

- PCS can only be used on a system that has neither Flash Cache nor Flash Pool configured.
- Use a PCS cache size that is 75% of the total SSD data drive capacity that is being considered; 25% of Flash Pool cache capacity is needed to store metadata and assure consistent performance.
- If the workloads being analyzed by PCS are likely to have a high rate of random overwrites (for example, transaction processing), adjust the resultant PCS hit rate down by 20% because PCS only accounts for read activity. This adjustment adds conservatism in sizing the HDDs needed to support the write workload.

Are quality of service (QoS) objectives used to manage Flash Pool cache data insertion and retention?

Answer: No. Insertion policies are determined by volume- or aggregate-level cache policies, and data retention (or eviction) is controlled by a modified version of a least recently used (LRU) algorithm.

How soon will performance improve after converting an aggregate to a Flash Pool aggregate?

Answer: Flash Pool cache needs to be populated with data (often called warming the cache) before the benefits of rereading the cached data can be realized. The amount of time until there is a noticeable improvement in read response time or reduction in HDD utilization depends on I/O activity and cache size. It might take as little as 30 minutes or as long as a few hours to see a dramatic improvement in performance.

Do all flexible volumes in a Flash Pool aggregate benefit from Flash Pool cache?

Answer: Using the default cache policies, all volumes provisioned in a Flash Pool aggregate are enabled to use Flash Pool cache. Cache policies can be modified at the volume level if there is a need to prioritize cache capacity use for data from one or more volumes in a Flash Pool aggregate.

Does system performance degrade as the SSDs in Flash Pool cache age?

Answer: No. Read performance is not affected by the amount of data that has been written to SSDs, and write performance will achieve and maintain a steady state with a full cache.

Which counters in Data ONTAP will help determine if implementing Flash Pool is a good option?

Answer: It is important to understand two key aspects of current system performance before deciding whether creating Flash Pool aggregates will improve performance. First you should know how random the workload is; to see this, use the following two counters:

- `stats show readahead:readahead:seq_read_req`
- `stats show readahead:readahead:rand_read_req`

Next, you should know how much I/O is being served from each resource on the system. That is done by collecting the following statistics:

- `stats show wafl:wafl:read_io_type.cache` (system memory)
- `stats show wafl:wafl:read_io_type.ext_cache` (Flash Cache module)
- `stats show wafl:wafl:read_io_type.disk` (HDD aggregates)
- `stats show wafl:wafl:read_io_type.bamboo_ssd` (SSD aggregates)
- `stats show wafl:wafl:read_io_type.hya_hdd` (Flash Pool aggregate HDDs)
- `stats show wafl:wafl:read_io_type.hya_cache` (Flash Pool aggregate SSD cache)

If the majority of the workload is handled by system memory, then adding cache—either Flash Pool SSDs or Flash Cache modules—will provide little to no improvement in overall system performance. In contrast, if a high percentage of total I/O is coming from disk (HDDs), then adding cache might improve, depending on how much of the read workload is random (as reflected by the readahead counters).

If there is a high degree of random write activity, it is important to know the write transfer (block) sizes. Only random writes, or more precisely overwrites, that are 16KB or smaller are candidates to be cached by Flash Pool. The protocol-specific “write size” histogram can be used to display the average block or transfer size. As an example, to view the histogram for NFS writes, use the “`show stats nfs:stats:nfs:write_size_histo`” counter.

In addition to using the counters outlined earlier, a Perfstat can be collected from the system during a time period that includes peak workloads. The Perfstat can then be uploaded into System Performance Modeler (SPM). SPM will process the Perfstat information and will break out the workload characteristics.

Does adding Flash Pool accelerate random write performance?

Answer: No. Data ONTAP systems handle random writes very well by caching write data in memory and protecting it with a copy in NVRAM. As a result, write operations are acknowledged back to hosts and clients at memory speeds. Flash Pool write cache does not accelerate this process.

Situations such as back-to-back consistency points (CPs) result from the inability of HDDs to keep up with the amount of write data being flushed from memory. For random write workloads that are disk bound (or close to it), Flash Pool cache can help by offloading overwrites from HDDs to SSD cache. (Flash Pool only caches random overwrites that are 16KB and smaller.) This frees up the HDDs to handle other operations, such as sequential reads or write appends.

5 Flash Cache and Flash Pool Cache Size Tables

Table 7) Data ONTAP 8.1.1 Flash Cache, Flash Pool, and combined maximum cache sizes.

Data ONTAP 8.1.1: Clustered Data ONTAP and Data ONTAP Operating in 7-Mode						
Platform Model	Flash Cache per Node	Flash Cache per HA Pair	Flash Pool per Node	Flash Pool per HA Pair	Combined per Node	Combined per HA Pair
FAS/V6280	8.0TiB	16.0TiB	6.0TiB	12.0TiB	8.0TiB	16.0TiB
FAS/V6240	3.0TiB	6.0TiB	6.0TiB	12.0TiB	6.0TiB	12.0TiB
FAS/V6210	1.5TiB	3.0TiB	2.0TiB	4.0TiB	2.0TiB	4.0TiB
FAS/V6080	2.0TiB	4.0TiB	2.0TiB	4.0TiB	2.0TiB	4.0TiB
FAS/V6070	2.0TiB	4.0TiB	2.0TiB	4.0TiB	2.0TiB	4.0TiB
FAS/V6040	1.0TiB	2.0TiB	1.0TiB	2.0TiB	1.0TiB	2.0TiB
FAS/V6030	Not supported	Not supported	1.0TiB	2.0TiB	Not supported	Not supported
FAS/V3270	1.0TiB	2.0TiB	1.0TiB	2.0TiB	1.0TiB	2.0TiB
FAS/V3240	0.5TiB	1.0TiB	0.5TiB	1.0TiB	0.5TiB ¹	1.0TiB ¹
FAS/V3210	256GiB ³	512GiB ³	Not supported	Not supported	Not supported	Not supported
FAS/V3170	1.0TiB	2.0TiB	1.0TiB	2.0TiB	1.0TiB	2.0TiB
FAS/V3160	0.5TiB	1.0TiB	0.5TiB	1.0TiB	0.5TiB	1.0TiB
FAS/V3140	256GiB ⁴	512GiB ⁴	Not supported	Not supported	Not supported	Not supported
FAS/V3070	0.5TB	1.0TB	Not supported	Not supported	Not supported	Not supported
FAS2240	Not supported	Not supported	300GiB ²		Not supported	Not supported
FAS2220	Not supported	Not supported	300GiB ²		Not supported	Not supported

Note: (1) When configuring Flash Cache and Flash Pool on a FAS/V3240, only the 256GB Flash Cache card can be used with up to two 100GB SSDs or one 200GB data SSD per controller. Both controllers in the HA pair should have the same (this is, symmetric) configuration.

Note: (2) For FAS2240 and FAS2200, the maximum Flash Pool cache capacity per HA pair and per node (including single-controller systems) is the same; that is, the maximum cache capacity can be configured on one node of an HA pair with no cache on the other node.

Note: (3) Only FAS/V3210 systems running Data ONTAP 8.1.1 operating in 7-Mode and that had Flash Cache cards installed before January 2012 are supported.

Note: (4) Only FAS/V3140 systems running Data ONTAP 8.1.1 operating in 7-Mode and that had Flash Cache cards installed before January 2012 are supported. The 256GB Flash Cache card is supported with a maximum of 128GiB of data cached per node.

Table 8) Data ONTAP 8.1.2 Flash Cache, Flash Pool, and combined maximum cache sizes.

Data ONTAP 8.1.2: Clustered Data ONTAP and Data ONTAP Operating in 7-Mode						
Platform Model	Flash Cache per Node	Flash Cache per HA Pair	Flash Pool per Node	Flash Pool per HA Pair	Combined per Node	Combined per HA Pair
FAS/V6290	8.0TiB	16.0TiB	6.0TiB	12.0TiB	8.0TiB	16.0TiB
FAS/V6280	8.0TiB	16.0TiB	6.0TiB	12.0TiB	8.0TiB	16.0TiB
FAS/V6250	3.0TiB	6.0TiB	6.0TiB	12.0TiB	6.0TiB	12.0TiB
FAS/V6240	3.0TiB	6.0TiB	6.0TiB	12.0TiB	6.0TiB	12.0TiB
FAS/V6220	1.5TiB	3.0TiB	2.0TiB	4.0TiB	2.0TiB	4.0TiB
FAS/V6210	1.5TiB	3.0TiB	2.0TiB	4.0TiB	2.0TiB	4.0TiB
FAS/V6080	2.0TiB	4.0TiB	2.0TiB	4.0TiB	2.0TiB	4.0TiB
FAS/V6070	2.0TiB	4.0TiB	2.0TiB	4.0TiB	2.0TiB	4.0TiB
FAS/V6040	1.0TiB	2.0TiB	1.0TiB	2.0TiB	1.0TiB	2.0TiB
FAS/V6030	Not supported	Not supported	1.0TiB	2.0TiB	Not supported	Not supported
FAS/V3270	1.0TiB	2.0TiB	1.0TiB	2.0TiB	1.0TiB	2.0TiB
FAS/V3250	1.0TiB	2.0TiB	1.0TiB	2.0TiB	1.0TiB	2.0TiB
FAS/V3240	0.5TiB	1.0TiB	0.6TiB	1.2TiB	0.6TiB ¹	1.2TiB ¹
FAS/V3220	0.5TiB	1.0TiB	0.6TiB	1.2TiB	0.6TiB	1.2TiB
FAS/V3210	256GiB ³	512GiB ³	Not supported	Not supported	Not supported	Not supported
FAS/V3170	1.0TiB	2.0TiB	1.0TiB	2.0TiB	1.0TiB	2.0TiB
FAS/V3160	0.5TiB	1.0TiB	0.6TiB	1.2TiB	0.6TiB	1.2TiB
FAS/V3140	256GiB ⁴	512GiB ⁴	Not supported	Not supported	Not supported	Not supported
FAS/V3070	0.5TiB	1.0TiB	Not supported	Not supported	Not supported	Not supported
FAS2240	Not supported	Not supported	400GiB ²		Not supported	Not supported
FAS2220	Not supported	Not supported	400GiB ²		Not supported	Not supported

Note: (1) When configuring Flash Cache and Flash Pool on a FAS/V3240, only the 256GB Flash Cache card can be used with up to two 100GB SSDs or one 200GB data SSD per node. Both controllers in the HA pair should have the same (that is, symmetric) configuration.

Note: (2) For FAS2240 and FAS2200, the maximum Flash Pool cache capacity per HA pair and per node (including single-controller systems) is the same; that is, the maximum cache capacity can be configured on one node of an HA pair with no cache on the other node.

Note: (3) Only FAS/V3210 systems running Data ONTAP 8.1.2 operating in 7-Mode and that had Flash Cache cards installed before January 2012 are supported.

Note: (4) Only FAS/V3140 systems running Data ONTAP 8.1.1 operating in 7-Mode and that had Flash Cache cards installed before January 2012 are supported. The 256GB Flash Cache card is supported with a maximum of 128GiB of data cached per node.

Table 9) Data ONTAP 8.2 Flash Cache, Flash Pool, combined maximum cache sizes.

Data ONTAP 8.2: Clustered Data ONTAP and Data ONTAP Operating in 7-Mode				
Platform Model	Flash Cache per Node	Flash Cache per HA Pair	Flash Pool per Node and per HA Pair ⁴	Combined Flash Cache and Flash Pool ⁵
FAS/V6290	8.0TiB	16.0TiB	24.0TiB	24.0TiB
FAS/V6280	8.0TiB	16.0TiB	18.0TiB	18.0TiB
FAS/V6250	4.0TiB	8.0TiB	18.0TiB	18.0TiB
FAS/V6240	3.0TiB	6.0TiB	12.0TiB	12.0TiB
FAS/V6220	3.0TiB	6.0TiB	12.0TiB	12.0TiB
FAS/V6210	1.5TiB	3.0TiB	4.0TiB	4.0TiB
FAS/V6080	2.0TiB	4.0TiB	4.0TiB	4.0TiB
FAS/V6070	Not supported ²	Not supported ²	Not supported ²	Not supported ²
FAS/V6040	1.0TiB	2.0TiB	2.0TiB	2.0TiB
FAS/V6030	Not supported ²	Not supported ²	Not supported ²	Not supported ²
FAS/V3270	1.0TiB	2.0TiB	2.0TiB	2.0TiB
FAS/V3250	1.0TiB	2.0TiB	4.0TiB	4.0TiB
FAS/V3240	0.5TiB	1.0TiB	1.2TiB	1.2TiB ¹
FAS/V3220	0.5TiB	1.0TiB	1.6TiB	1.6TiB
FAS/V3210	Not supported ³	Not supported ³	Not supported	Not supported
FAS/V3170	1.0TiB	2.0TiB	2.0TiB	2.0TiB
FAS/V3160	0.5TiB	1.0TiB	1.2TiB	1.2TiB
FAS/V3140	Not supported ³	Not supported ³	Not supported	Not supported
FAS/V3070	Not supported ²	Not supported ²	Not supported ²	Not supported ²
FAS2240	Not supported	Not supported	400GiB	Not supported
FAS2220	Not supported	Not supported	400GiB	Not supported

Note: (1) When configuring Flash Cache and Flash Pool on a FAS/V3240, only the 256GB Flash Cache card can be used with up to two 100GB SSDs or one 200GB data SSD per node. Both nodes in the HA pair should have the same (that is, symmetric) configuration.

Note: (2) FAS/V6070, FAS/V6030, and FAS/V3070 are not supported with Data ONTAP 8.2.

Note: (3) FAS/V3140 and FAS/V3140 are not supported with Data ONTAP 8.2 when Flash Cache cards are installed. Both platforms are supported running Data ONTAP 8.2 without Flash Cache.

Note: (4) For platforms, the maximum Flash Pool cache capacity per HA pair and per node (including single-controller systems) is the same; that is, the maximum cache capacity can be configured on one node of an HA pair with no cache on the other node.

Note: (5) Cache maximums when both Flash Cache and Flash Pool are configured on an HA pair, or on a single node, apply at the node level. For example, for an HA pair with Nodes A and B, the maximum combined cache allowed for Node A = Flash Cache on Node A + total Flash Pool

cache on Node A + total Flash Pool cache on Node B. Node B's Flash Pool cache must be factored in to account for controller takeover scenarios when Node A manages all cache.

6 References

This FAQ used the following references:

- [TR-4070: NetApp Flash Pool Design and Implementation Guide](#)
- [NetApp Flash Pool Sales and SE Technical Presentation](#)
- [Product Bulletin: Pure SSD and Flash Pool with DS424X Disk Shelf](#)
- [DS4246 SE Technical Presentation](#)
- [Storage Subsystem Technical FAQ](#)

7 Revision History

Table 10) Revision information.

Date	Version	Comments
June 2013	4.1	Skip Shapiro: FAS2200 mixed internal drive configuration and format/branding updates.
May 2013	4.0	Skip Shapiro: Updated based Data ONTAP 8.2 enhancements and additional SSD support. Replaced content that appears in TR-4070 with references.
December 2012	3.2	Jay White: Corrected SyncMirror information concerning the mirrored plex counting toward the cache limit.
December 2012	3.1	Jay White: Updated FAS2200 series SSD limits.
December 2012	3.0	Jay White: Major update. Added 200GB SSD drive, DS2246 SSD offerings, updated cache limits for 8.1.1 and 8.1.2, POC recommendations, MetroCluster updates, performance graphs for OLTP and SFS, updated shelf mixing information, fixed some typos, eviction scanner information, pure SSD support information, and additional command information (cache used, disable cache, and more). Corrected the X441A-R5 usable capacity.
August 2012	2.1	Jay White: Fixed references link to the Storage Subsystem Technical FAQ.
August 2012	2.0	Jay White: Restructured the content and added several new items based on a total product review (hardware and software).
July 2012	1.1	Paul Updike: Added an item to address combined Flash Pool and Flash Cache limits.
May 2012	1.0	Paul Update: Initial document created and published.

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®



www.netapp.com

© 2013 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, Flash Cache, Flash Pool, FlexShare, MetroCluster, 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.