



Technical Report

64-Bit Aggregates

Overview and Best Practices

GV Govindasamy, NetApp
April 2015 | TR-3978

Abstract

Earlier releases of NetApp® Data ONTAP® used data block pointers in 32-bit format which limited the addressable space of an aggregate to 16TB. Starting with Data ONTAP 8.0 software operating in 7-Mode and clustered Data ONTAP in 8.0.1, larger aggregates support was introduced by using 64-bit block pointers. Larger aggregates simplify administration by allowing many large volumes in a single container and can improve IO performance with a large number of larger disk drives. This technical report gives an overview and best practices for using larger aggregates and the recommended procedures to expand 32-bit aggregates to 64-bit aggregates in place, nondisruptively.

Data Classification

Public

Version History

Version	Date	Document Version History
Version 2.1	April 2015	GV Govindasamy: <ul style="list-style-type: none">• The support for expanding aggregates without adding disks in 8.1.4P4 updated.• Minor customary edits in section 4 for snapmirror related content.
Version 2.0	March 2015	GV Govindasamy: <ul style="list-style-type: none">• TR-3786: A Thorough Introduction to 64-Bit Aggregate discontinued and relevant content merged with this document.• Content updated for clustered Data ONTAP 8.3.
Version 1.2	February 2014	Roy K. Scaife: <ul style="list-style-type: none">• Added further updates for 32-bit to 64-bit aggregate expansion in Data ONTAP 8.2.1
Version 1.1	December 2013	Roy K. Scaife: <ul style="list-style-type: none">• Added updates for 32-bit to 64-bit aggregate expansion in Data ONTAP 8.2.1
Version 1.0	September 2011	Shree Reddy: Initial version

TABLE OF CONTENTS

Version History	2
1 Overview	5
2 Requirements	7
2.1 Release.....	7
2.2 License.....	7
2.3 Platforms.....	7
2.4 Requirements for 64-bit expansion process.....	7
2.5 Caveats.....	8
3 Architecture	9
3.1 Background.....	9
3.2 Data ONTAP Upgrades and Aggregate Formats.....	10
3.3 Overview of Expansion Process.....	11
4 Interoperability	15
4.1 Compression.....	15
4.2 FlexClone.....	15
4.3 NDMP and Dump.....	15
4.4 Deduplication.....	15
4.5 FlexCache.....	15
4.6 Flash Pool.....	16
4.7 Flash Cache.....	16
4.8 Data Replication.....	17
4.9 Transitioning from 7-Mode to Clustered Data ONTAP 8.3.....	22
4.10 Summary of Feature Compatibility with the 64-Bit Expansion Process.....	23
4.11 Interoperability Matrix Tool (IMT).....	24
5 Performance	25
5.1 Better Performance with More Disks.....	25
5.2 Memory Usage.....	25
5.3 Consideration During 64-Bit Expansion Process.....	25
6 Best Practices	27
6.1 Best Practices Table.....	27
6.2 Disclaimer.....	29

7 Frequently Asked Questions (FAQ)	29
Additional Resources	32
Contact Us	32
Appendix	33
Expansion Process Examples Using Command Line Interface (CLI).....	33

BEST PRACTICES

Best Practice 1: Use 64-Bit Aggregates.....	27
Best Practice 2: RAID Type and RAID Group Sizes	27
Best Practice 3: Maximum Aggregate Size	27
Best Practice 4: Recommended Root Volume Size.....	27
Best Practice 5: Maximum File and LUN Size.....	27
Best Practice 6: Consider Flash Cache for Highly Random Workloads	27
Best Practice 7: Enable Compression Only After Expansion Process Is Complete	28
Best Practice 8: Ensure Adequate Space for Volumes Before Expansion Process Is Started	28
Best Practice 9: Considerations to Start the Expansion Process.....	28
Best Practice 10: Expand Source Aggregate Before Expanding Destination Aggregate	28

LIST OF FIGURES

Figure 1) 32-bit addressing format thresholds constrain aggregate and FlexVol volume size to 16TB	9
Figure 2) 64-bit addressing makes it feasible to increase aggregate size	10
Figure 3) Upgrade with 32-bit aggregates.....	11
Figure 4) 32-bit to 64-bit expansion.....	13
Figure 5): 64-bit to 32-bit SnapMirror relationship.	19
Figure 6): 32-bit to 64-bit SnapMirror relationship.	19

TABLES

Table 1) Expansion process features.	12
Table 2) Data replication overview.	17
Table 3) SnapMirror supportability matrix for clustered Data ONTAP.....	18
Table 4) SnapMirror supportability matrix for 7-Mode.....	20
Table 5) Qtree SnapMirror and SnapVault configurations in 7-Mode.....	21
Table 6) Interoperability with Data ONTAP features.	23

1 Overview

Aggregates and FlexVol® are the ground-breaking technology that was introduced in Data ONTAP 7G to give storage administrators greater flexibility in managing ever-changing storage size requirements while maximizing storage efficiency. Aggregates provide storage to the volumes they contain. Each aggregate has a set of assigned disks or array LUNs. An aggregate can hold one or more FlexVol volumes. These FlexVol volumes share the physical resources of the underlying container.

There has been growing demand for larger volumes and better performance with larger disk drives. The first implementation of aggregates was limited to 16TB size because it was using 32-bit block pointers. The 16TB limit is an artifact of the maximum addressable space with a 32-bit block pointer in a 32-bit aggregate. Data ONTAP 8.0 operating in 7-Mode and clustered Data ONTAP 8.0.1 introduced larger aggregates, implemented using 64-bit block pointers. With 64-bit block pointers the aggregate size threshold was not limited by the data format and offered flexibility for future growth. Larger size aggregates have several benefits including the ability to increase the number of larger disks for better performance and simplified management with fewer aggregates. Starting with clustered Data ONTAP 8.3, support for 32-bit aggregates is deprecated and only 64-bit aggregates are supported.

This technical report gives an overview of the benefits of using 64-bit aggregates and the recommended ways to expand 32-bit aggregates to 64-bit aggregates seamlessly.

The Journey of NetApp Aggregates

7.0	The introduction of aggregates. An aggregate is a container that can hold one or more FlexVol® volumes. The initial format is 32-bit and limited to maximum size of 16TB.
8.0 7-Mode 8.0.1 clustered	The introduction of support for 64-bit (larger) aggregates which can support larger FlexVol volumes. Both 32-bit and 64-bit aggregates can coexist in the system.
8.1	The default aggregate type is changed from 32-bit to 64-bit for new aggregates.
8.2	In-place expansion of 32-bit aggregates to 64-bit introduced.
8.2.1	In-place expansion procedure for 32-bit aggregates to 64-bit improved without requiring additional disks.
8.3	Deprecation of support for 32-bit aggregates. Only 64-bit aggregates are supported.

Benefits of Larger Aggregates

Delivering Performance with Large, SATA drives

The size of disk drives is constantly increasing. As the drives get bigger, the maximum number of disks that can fit in a 16TB aggregate decreases. The low disk count and therefore the low spindle count can

become a performance bottleneck, especially in workloads that are spindle bound for meeting their performance requirements. To alleviate this problem, aggregates that can have more spindles are needed, which means the aggregates size threshold must increase.

Maintaining High Storage Efficiency

The 16TB size thresholds of existing 32-bit aggregates can hinder the ability to add new disk drives to the aggregate in fully populated RAID groups, especially when using large-sized SATA drives. In such cases, in which the aggregate is not composed of fully populated RAID groups, the storage system does not provide the best possible storage efficiency. 64-bit aggregates with larger size thresholds will let you configure your storage systems to deliver the maximum storage efficiency.

Ease of Management

More disks can be part of a single aggregate that allows system administrators to manage fewer larger aggregates. Administrators also have the flexibility of mixing different workloads and hosting more volumes in a larger aggregate to take advantage of the additional throughput achieved by the increased number of spindles per aggregate.

Better Scalability for Future Growth

As storage needs grow in the future, 64-bit addressing allows having even larger aggregates without changing any data format changes or conversion to the existing data. This will make it easier to grow the aggregates and FlexVol sizes in the future.

Power and Space Savings

Because larger aggregate threshold sizes are constrained by a 32-bit addressing limitation, where performance is not a criterion, a small number of larger SATA drives can be considered to create a larger aggregate. Having fewer disks saves space as well as power.

Data Consolidation

Because FlexVol volume sizes can be much larger, scattered data living in various volumes can be consolidated under a single larger FlexVol volume. This will simplify the volume management.

2 Requirements

2.1 Release

64-bit aggregates are supported in Data ONTAP 8.0 operating in 7-Mode and clustered Data ONTAP 8.0.1 onward.

Starting in clustered Data ONTAP 8.3, only 64-bit aggregates are supported and 32-bit aggregates have been deprecated.

2.2 License

No extra licenses are needed to enable and use the 64-bit aggregate feature in the supported releases.

2.3 Platforms

All platforms that can run the supported releases can create and manage 64-bit aggregates. For a list of platforms, refer to the Hardware Universe on the NetApp Support site at <https://hwu.netapp.com>.

2.4 Requirements for 64-bit expansion process

Data ONTAP 8.1 and 8.2

- Disks must be added to an existing 32-bit aggregate such that the aggregate's storage capacity exceeds 16TB. This is the only method that is available for customers to trigger the expansion process.
- By upgrading to Data ONTAP 8.2.1 or later, 32-bit aggregates can be expanded to 64-bit without adding additional storage. This may be a preferred way for many users.
- Because 64-bit aggregates consume more space for metadata, enough free space should be available to enable completion of the expansion process. Expansion on any 32-bit aggregate created in prior releases can be triggered only after the completion of the prequalification scanner. The prequalification scanner is triggered on an upgrade to Data ONTAP 8.1. See the "Architecture" section for more details.

Data ONTAP 8.2.1+ and 8.1.4P4+

- Starting in Data ONTAP 8.2.1+ or 8.1.4P4+, in-place expansion without adding disks is supported. The following commands can be used to trigger the expansion process. Please refer the Appendix and Storage Management Guide for exact steps to be followed.

```
cDOT::> aggr 64bit-upgrade start -aggregate <aggregate name>
```

```
7-Mode:> aggr 64bit-upgrade start <aggregate name>
```

- Because 64-bit aggregates consume more space for metadata, enough free space should be available to enable completion of the expansion process. See the "Space Consumption in Architecture" section for more details.
- When planning to upgrade to Data ONTAP 8.3, verify the aggregate does not have any 32-bit Snapshot@ copies. Starting in Data ONTAP 8.2.1, you can use the following commands to identify any 32-bit Snapshot copies in the volume or aggregate. However, earlier releases such as 8.1.4P4 do not display "fs-block-format".

```
cDOT::> snap show -fields fs-block-format <aggr_name or vol_name>
```

```
7-Mode:> snap list -fs-block-format <aggr_name or vol_name>
```

Clustered Data ONTAP 8.3

- Clustered Data ONTAP 8.3 supports only 64-bit aggregates. The support for 32-bit aggregates is deprecated. Upgrading to clustered Data ONTAP 8.3 requires all existing aggregates to be 64-bit and no 32-bit aggregates, volumes or Snapshot copies are allowed to exist.

2.5 Caveats

- If the 32-bit aggregate you want to expand contains any FlexCache[®] volumes, the cache volumes have to be destroyed. The cache volumes can be recreated after the expansion process has started. This causes FlexCache clients to lose their existing connections. These clients must reconnect to the newly created FlexCache volumes and they may experience an initial slowdown in data access due to cache repopulation.
- Prior to upgrading to Data ONTAP 8.3 the aggregate must be free of all 32-bit Snapshot copies. 32-bit Snapshot copies must either be deleted manually by the administrator or wait to be cycled off and replaced by newer 64-bit Snapshot copies. Starting in Data ONTAP 8.2.1, the CLI “`snap list`” in 7-Mode and “`snap show`” in clustered Data ONTAP displays “`fs-block-format`” field to help identify 32-bit snapshot copies.
- DataMotion[™] for volumes for 7-Mode volumes does not work across 32-bit and 64-bit aggregates.
- DataMotion for vFiler[®] units (7-Mode only) does not work across 32-bit and 64-bit aggregates.
- Starting with clustered Data ONTAP 8.2, SnapVault[®] is supported only on 64-bit aggregates which requires 64-bit aggregates on both source and destinations to establish the vault relationship.

Note: DataMotion for volumes in clustered Data ONTAP does not have the same restriction. Clustered Data ONTAP volumes can be moved between 32-bit and 64-bit aggregates.

3 Architecture

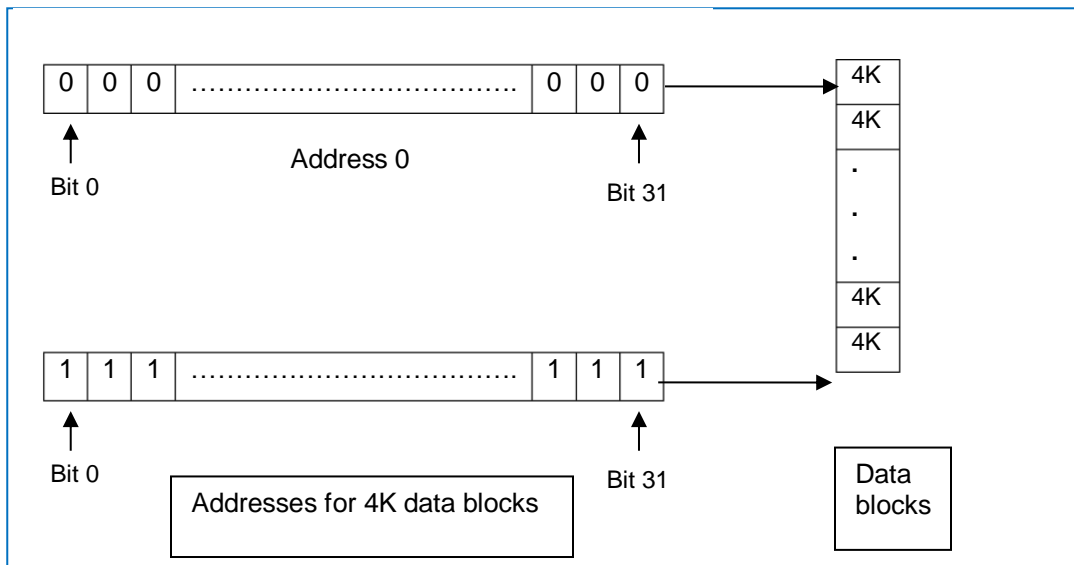
3.1 Background

A file is made up of individual data blocks, and large files have additional layers of indirection between the inode and the actual data blocks. These additional layers are referred to as indirect blocks.

The block addresses for 64-bit aggregates are stored in 64 bits. The block addresses for 32-bit aggregates are stored in 32 bits. With the increased block pointer size in 64-bit aggregates, the number of addresses or block pointers that an indirect block can hold decreases from 510 to 255. This results in an inode requiring more indirect blocks in a 64-bit volume or aggregate.

Data ONTAP versions earlier than Data ONTAP 8.0 used pointers that were 32 bits in length to address disk blocks. Data ONTAP addresses disk space in blocks of 4K. Figure 1 shows how the 32-bit address space constrained the maximum size to 16TB. Because the address pointer is 32 bits in size, the maximum address that can be stored is $(2^{32} - 1)$. Because each block that is referenced with this address is 4KB in size, the maximum size of the aggregate that can be referenced is $(2^{32} - 1) * 4KB = 16TB$.

Figure 1) 32-bit addressing format thresholds constrain aggregate and FlexVol volume size to 16TB

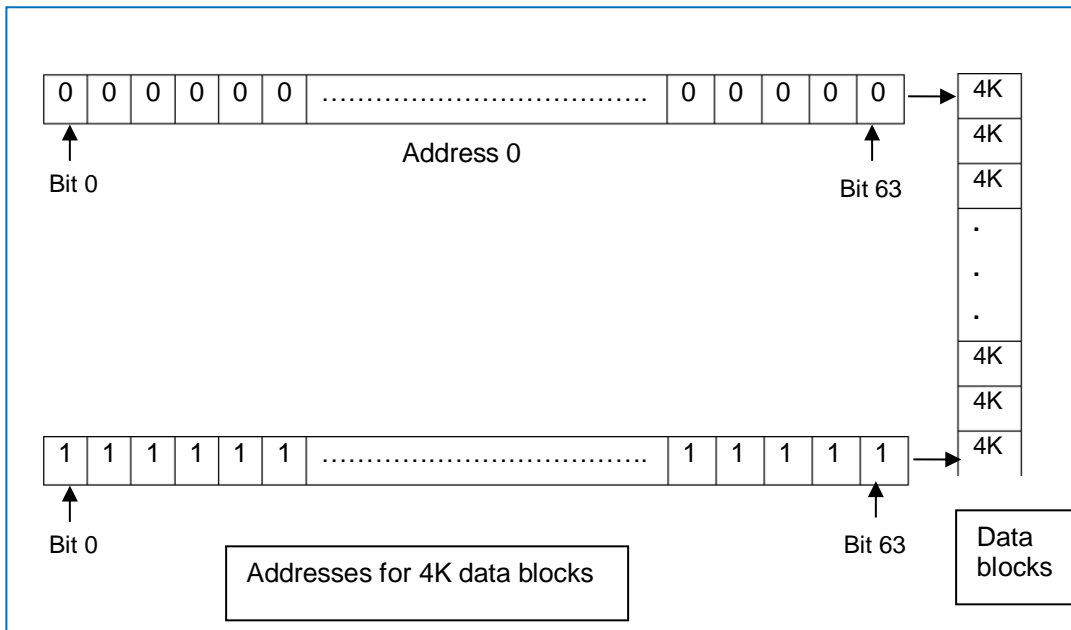


Using a 64-bit addressing scheme, aggregate size can be a million terabytes in size. However, taking into consideration other factors such as system capacity, memory, and hardware configuration, the maximum aggregate size thresholds for Data ONTAP are platform dependent. These thresholds are listed in NetApp Hardware Universe at <http://hwu.netapp.com/>.

FlexVol volumes that are contained in an aggregate inherit the addressing scheme of the containing aggregate. So, FlexVol volumes in 64-bit aggregates have a 64-bit addressing scheme.

The advantage of the new 64-bit addressing scheme is that the large address space allows increasing the aggregate size to an even larger size in the future without any changes to the underlying on-disk data format or memory addressing format.

Figure 2) 64-bit addressing makes it feasible to increase aggregate size.

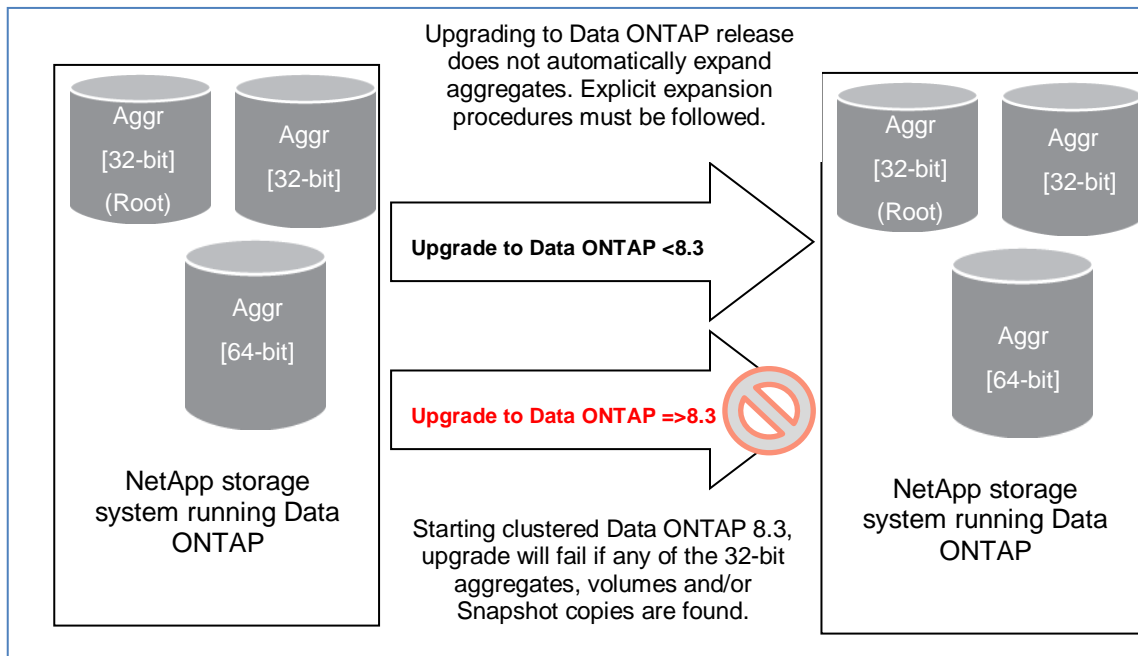


3.2 Data ONTAP Upgrades and Aggregate Formats

Upgrading Data ONTAP releases will not have any change to existing aggregates. All aggregates continue to retain the same format. All existing 32-bit aggregates continue to be 32-bit after upgrading to the supported Data ONTAP releases. To expand 32-bit aggregates to 64-bit aggregates, an administrator must perform explicit expansion procedures depending on the Data ONTAP release. See section “Expansion Process” for details.

Starting with Data ONTAP 8.1, the default aggregate format changed from 32-bit to 64-bit for all new aggregates. Starting with clustered Data ONTAP 8.3, 32-bit data formats will not be supported. Upgrading to Data ONTAP 8.3 will fail if any 32-bit aggregates, volumes or Snapshot copies exist. The administrator must expand all 32-bit aggregates to 64-bit aggregates and remove all 32-bit Snapshot copies or wait for them to cycle off prior to upgrading to clustered Data ONTAP 8.3.

Figure 3) Upgrade with 32-bit aggregates.



3.3 Overview of Expansion Process

64-bit aggregates consume additional metadata with the increase in block pointer size. Therefore, the expansion process needs to allocate extra blocks. The expansion process expands all existing 32-bit indirect blocks within the aggregate and its containing FlexVol volumes to the 64-bit format. It walks through all the indirect blocks in every inode tree, allocating additional blocks when needed to hold the larger 64-bit pointers.

The expansion process flags the inode tree as it completes the expansion and can be resumed naturally if the expansion process gets interrupted or stopped before the completion. When a space issue is encountered, the expansion process will be stopped. When the space issue is fixed by allocating more space, the expansion process will resume automatically.

After the expansion process is initiated, all new writes will always write blocks in the 64-bit format.

Nondisruptive Process

The expansion process occurs in two ways:

1. **Inline.** As clients access their data, existing 32-bit indirect blocks are expanded to 64-bit.
2. **Background.** A background process expands all 32-bit indirect blocks to 64-bit.

The process is seamless, online, and in place, and there is no disruption to client access.

Table 1) Expansion process features.

Features of Expansion Process	
In-place	No data copy or migration is required.
Online	Data access is maintained during the expansion.
Nondisruptive	Existing aggregate configurations including storage efficiency features and data protection relationships are not affected.

Completion Time

The expansion process rewrites only indirect blocks and does not modify the data blocks. This means the completion time for expansion process will vary depending on the data layout in the aggregates and volumes. As an example, an aggregate with a lot of metadata (indirect) blocks can take longer to complete the expansion process than an aggregate with fewer metadata blocks. This expansion process runs in the background and does not affect front-end I/O performance directly. However, when a client is accessing data that is still in 32-bit format and has not been converted to 64-bit format, the expansion process starts in line which can cause higher latency for this particular I/O.

Space Consumption

During the expansion process, metadata blocks are rewritten in 64-bit format. Data ONTAP uses 4k blocks. Each 4k block can contain 510 32-bit block pointers or 255 64-bit block pointers. When metadata blocks are converted from 32-bit to 64-bit block pointers, the space requirements for metadata blocks approximately double. Typical datasets contain 2% of metadata blocks in an aggregate.

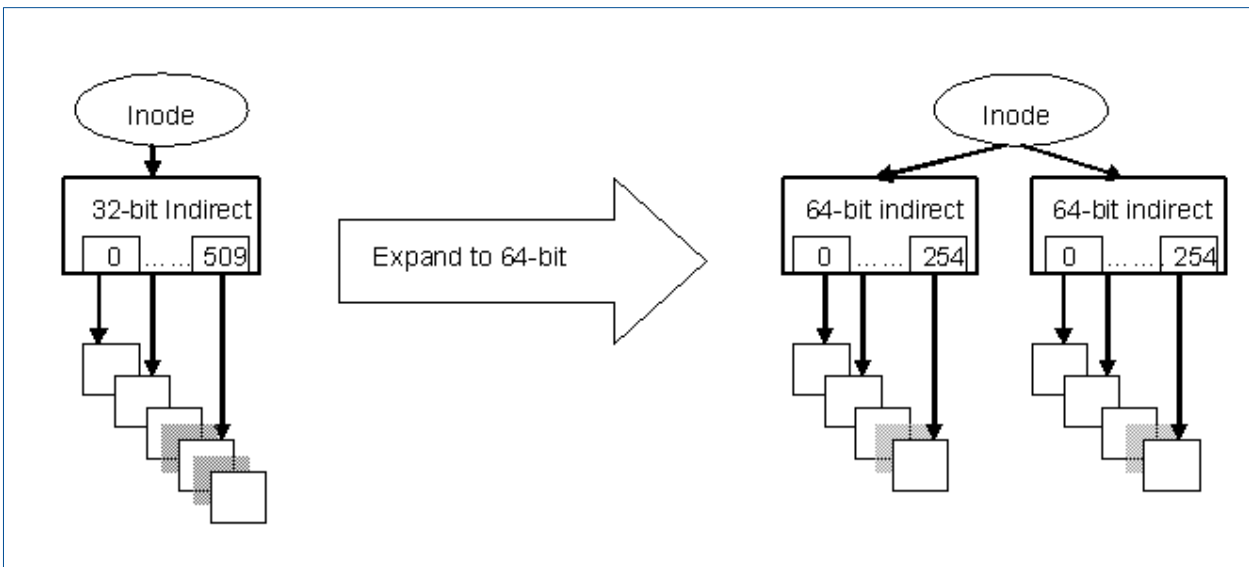
If your volumes are close to full, you can increase the size of the volumes beforehand. When triggering the expansion process the administrator will be notified about the list of volumes that will require additional space for the expansion process.

Prequalification Scanner

Because the expansion process requires additional space, a prequalification process is triggered first during the expansion process. The prequalification process is also triggered during a Data ONTAP upgrade to Data ONTAP 8.1 and 8.2 to gather information on all the space-reserved files in existing volumes. This ensures we can continue to enable space reservations. The scanner provides a quick estimate of additional blocks needed to expand a volume. The expansion process uses this estimate to determine if there is sufficient space in the volumes to successfully trigger the expansion and maintain space guarantees. Therefore, the expansion process can be triggered only after the completion of the prequalification phase. If the expansion process is triggered in the midst of the prequalification phase, it errors out with a more detailed message.

The scanner is not triggered for any new volumes created in Data ONTAP 8.1 or 8.2, because the required space accounting is maintained from the beginning.

Figure 4) 32-bit to 64-bit expansion.



Clustered Data ONTAP 8.3

Starting with clustered Data ONTAP 8.3, the support for 32-bit addressing is discontinued. This means that aggregates, volumes, Snapshot copies and SnapMirror® relations with 32-bit aggregates will not be supported. The upgrade procedure will prevent the system being upgraded to clustered Data ONTAP 8.3 when 32-bit aggregates, volumes or Snapshot copies are found in the system. Expanding all 32-bit aggregates to 64-bit aggregates before upgrading to 8.3 is recommended.

Data ONTAP 8.2.1+ and 8.1.4P4+

Starting in Data ONTAP 8.2.1+ or 8.1.4P4+, the expansion process can be triggered from the command line interface (CLI) without adding disks to provide more than 16TB of aggregate capacity.

Starting in Data ONTAP 8.2.1 or later, the CLI “`snap show`” in clustered Data ONTAP and “`snap list`” in 7-Mode can be used to identify the filesystem formats of the snapshots. This is useful to handle the existing 32-bit snapshots, as clustered Data ONTAP 8.3 does not support 32-bit formats.

When the conversion process begins, Data ONTAP checks that there is enough free space for the extra metadata blocks required for the 64-bit expansion. If all the volumes have sufficient free space, the aggregate expands to 64-bit. If any of the volumes are too full to accommodate the 64-bit expansion, the command fails.

64-Bit Expansion Workflow

The administrator should execute the following sequence of steps to expand 32-bit aggregates.

1. Establish the requirement to expand from a 32-bit aggregate to a 64-bit aggregate.
2. Complete the command line commands to begin the expansion process. All new writes are in 64-bit format.
3. The expansion process finishes on all volumes within the aggregate.

Clients can continue to access their data during the entire process.

For the steps and the commands, refer to Data ONTAP 8.2 Physical Storage Management Guide or to an example in the appendix in this document.

Data ONTAP 8.1 and 8.2

The expansion process is triggered only when you add disks to a 32-bit aggregate such that the resulting aggregate capacity is greater than 16TB. When disks are added, if the resulting capacity is less than 16TB, the aggregate continues to be 32-bit.

For best performance, add a complete RAID group to prevent the new disks from becoming a performance bottleneck because all new client writes are directed to the newly added disks.

When disks are added, Data ONTAP checks that there is enough free space for the extra metadata blocks required for the 64-bit expansion. If all the volumes have sufficient free space, the disks are added and the aggregate is expanded to 64-bit. If any of the volumes are too full to accommodate the 64-bit expansion, the command fails.

64-Bit Expansion Workflow

The administrator should execute the following sequence of steps to expand 32-bit aggregates

1. Establish the requirement to grow an aggregate beyond 16TB.
2. Add enough disks to grow the aggregate beyond 16TB.
3. Background expansion starts on all volumes (excluding read-only SnapMirror destinations) within the aggregate. All new writes are in 64-bit format.
4. The expansion process finishes on all volumes within the aggregate.

Clients can continue to access their data during the entire process.

For the steps and the commands, refer to Data ONTAP 8.1 Physical Storage Management Guide.

Data ONTAP 8.0.x

Data ONTAP 8.0.x does not support expansion of an existing 32-bit aggregate to a 64-bit aggregate. If your FlexVol volume is in a 32-bit aggregate and you need a volume larger than 16TB to hold the data in that volume, you have to either upgrade to latest Data ONTAP releases or move the data to a FlexVol volume located in a 64-bit aggregate. You can use the following tools to migrate data from a FlexVol volume in a 32-bit aggregate to a FlexVol volume in a 64-bit aggregate:

Qtree SnapMirror. Qtree SnapMirror can be used to copy data from a qtree or a FlexVol volume in a 32-bit aggregate to a qtree located in a FlexVol volume in a 64-bit aggregate. If the data on the source volume is laid out in qtrees, you need to have one qtree SnapMirror relationship per qtree from the source to the destination to move all the data from the source to the destination.

Network Data Management Protocol (NDMP). You can use `ndmpcopy` command to copy volumes, qtrees, directories, or individual files from the source volume located in a 32-bit aggregate to the destination volume located in the 64-bit aggregate. If you have a single large directory or file that you want to copy from a volume in a 32-bit aggregate to a volume in a 64-bit aggregate, you can use this tool.

Both of these tools can be used to copy data that is present only in the active file system of the source volume. However, any Snapshot copies that are present in the source volume cannot be copied to the destination using either of these tools. To migrate data along with all the Snapshot copies to the destination, contact your NetApp Professional Services engineer.

4 Interoperability

4.1 Compression

Compression is not supported on 32-bit volumes. Compression can be enabled as soon as you trigger the 32-bit to 64-bit expansion process on the aggregate.

Best Practice

Best Practice 7: Enable Compression Only After Expansion Process Is Complete

4.2 FlexClone

If the aggregate that is in the midst of being expanded contains any FlexClone[®] volumes, the 64-bit expansion process will expand the indirect blocks in the FlexClone volume too. This will cause parts of the FlexClone volume to be split from its parent because it will overwrite all indirect blocks shared with the parent volume.

If FlexVol volumes contain a lot of FlexClone files and/or LUNs, there will be more metadata blocks in them. This can cause the expansion process to take slightly longer to complete and may require more space for expansion.

4.3 NDMP and Dump

NDMP and dump operation work at the logical file level and are not aware of the underlying physical storage structure or memory addressing scheme. Therefore, they work across different aggregate types.

You can use NDMP to copy data from a FlexVol volume in a 32-bit aggregate to a FlexVol volume in a 64-bit aggregate or a 32-bit aggregate.

You can create a dump of a volume in a 32-bit aggregate to create a backup copy and then restore from that backup to a volume in a 64-bit aggregate, and vice versa.

4.4 Deduplication

Deduplication is supported on FlexVol volumes that reside in 64-bit aggregates. However, to enable deduplication on a FlexVol volume, the maximum FlexVol volume size must be within the supported range. Refer to NetApp Hardware Universe at <http://hww.netapp.com>.

For more information about deduplication, see NetApp TR-3505 NetApp Deduplication for FAS and V-Series Deployment and Implementation Guide.

4.5 FlexCache

FlexCache software creates a flexible, horizontally scalable caching layer in your storage infrastructure that increases performance by automatically adapting to changing usage patterns to eliminate bottlenecks. FlexCache volumes store the most frequently used data from origin volumes

on other NetApp storage systems. When servers or users request data, FlexCache serves it locally whenever possible, shortening storage response time and reducing the workload on the origin system. Full cache consistency and file locking enable the integrity of your data at all times. For details about FlexCache, see NetApp TR-3669: FlexCache Caching Architecture.

FlexCache relationships between FlexVol volumes are independent of the underlying aggregate types. Therefore you can use a FlexVol volume in a 64-bit aggregate as the cache for an origin FlexVol volume that is located in a 32-bit aggregate and vice versa.

Although FlexCache relationships are independent of the underlying aggregate types, the in-place expansion of 32-bit aggregates to 64-bit aggregates is not supported with the presence of FlexCache volumes. FlexCache volumes must be removed to start the expansion process. After the process is started, a new FlexCache volume can be created and the clients can be reconnected to the FlexCache volumes.

4.6 Flash Pool

A FlashPool™ aggregate configures solid-state drives (SSDs) and hard disk drives (HDDs) into a single aggregate, with the SSDs providing fast-response-time cache for volumes that are provisioned on the Flash Pool aggregate. Flash Pool aggregates are in general helpful for transactional workloads and databases, because SSDs in Flash Pool aggregates gives faster response time by reducing latency for random reads and overwrites. Flash Pool aggregates are supported only in 64-bit aggregate types and unsupported with 32-bit aggregate types. For additional details, please refer to TR-4070: Flash Pool Design and Implementation Guide.

4.7 Flash Cache

NetApp Flash Cache integrates caching software with PCI-Express (PCIe) cards to improve storage system performance at the node level. Flash Cache offers millisecond access to cached read data, which is 8 to 20 times faster than typical read latencies from HDDs. The additional cache can optimize the performance of your random read intensive workloads such as file services and messaging. It works with 64-bit aggregates as well as 32-bit aggregates and caches data that is coming from volumes located in both types of aggregates. Flash Cache caches data based on data access regardless of the aggregate type.

The data cached in Flash Cache depends on the workload, and it can be a combination of data from volumes contained in different aggregates (except Flash Pool). As noted in other sections of this document, 64-bit aggregates have a bigger address space and require more memory for their metadata than 32-bit aggregates. Flash Cache may be considered to achieve better performance for highly intensive random read workloads with relatively larger working data sets. For additional details, please refer to TR-3832: Flash Cache Best Practices Guide.

Best Practice

Best Practice 6: Consider Flash Cache for Highly Random Workloads

4.8 Data Replication

Table 2 gives an overview of the various DR features supported in Data ONTAP operating in 7-Mode and clustered Data ONTAP with the aggregate format compatibility.

Table 2) Data replication overview.

Feature	7-Mode	Clustered Data ONTAP
	Works across 32-Bit and 64-Bit Aggregates?	Works across 32-Bit and 64-Bit Aggregates?
Asynchronous volume SnapMirror	Yes	Yes
Synchronous volume SnapMirror	Yes	N/A
Aggr copy	Yes	N/A
Vol copy	Yes	N/A
DataMotion for volumes	No	Yes
Load-sharing mirrors	N/A	Yes

Clustered Data ONTAP

Starting in clustered Data ONTAP 8.1, you can create block-based asynchronous replication relationships between volumes of different formats.

Starting in clustered Data ONTAP 8.3, only 64-bit aggregate format is supported and the support for 32-bit format is deprecated.

Refer to the supportability matrix in Table 3 for all the supported SnapMirror relationships across different formats and Data ONTAP kernels for clustered Data ONTAP. The asynchronous replication engine in clustered Data ONTAP is responsible for all volume replication (data protection and load-sharing mirrors) and volume operations (vol move, vol copy). Hence, this works across 32-bit and 64-bit aggregates unlike the engine in 7-Mode.

Table 3) SnapMirror supportability matrix for clustered Data ONTAP.

Source Volume	Destination Aggregate				
	8.1.x (32-Bit)	8.1.x (64-Bit)	8.2.x (32-Bit)	8.2.x (64-Bit)	8.3 (64-Bit)
8.1.x (32-bit)	Yes	Yes	Yes	Yes	No
8.1.x (64-bit)	Yes	Yes	Yes	Yes	Yes
8.2.x (32-bit)	No	No	Yes	Yes	No
8.2.x (64-bit)	No	No	Yes	Yes	Yes
8.3 (64-bit)	No	No	No	No	Yes

Expansion Process in SnapMirror Volumes (8.2.x)

On clustered Data ONTAP versions prior to 8.3, an aggregate that contains one or more SnapMirror source/destination volumes can be expanded from 32-bit to 64-bit format. All activities related to SnapMirror, such as SnapMirror break, resync, and so on, remain functioning on the SnapMirror source/destination volumes during and after the expansion process. The SnapMirror source will be expanded to 64-bit first. The SnapMirror destination will always be skipped. SnapMirror destination volumes follow the format of their source volumes.

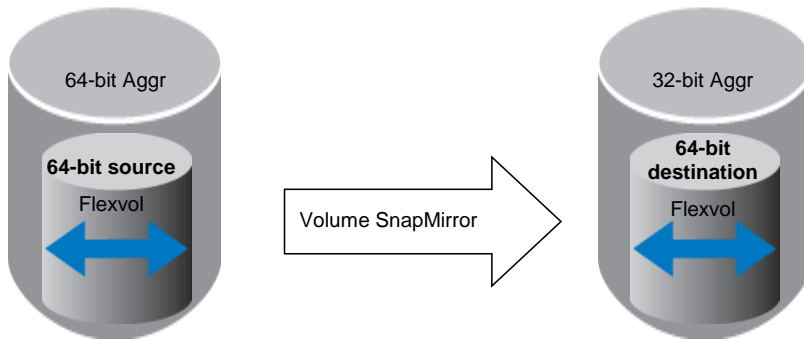
When the expansion process is triggered on a SnapMirror source volume, it will result in increased SnapMirror traffic while the expansion process rewrites all the indirect blocks in the new 64-bit format. This increase should be small because indirect blocks account for a very small percentage of any dataset.

Starting in Data ONTAP 8.1, it is possible to support SnapMirror relationships across different volume formats. The SnapMirror destination volume inherits the format of the source volume. In addition, aggregates can be expanded from 32-bit format to 64-bit format without interrupting existing SnapMirror relationships.

There are two possible configurations:


1. It is possible to create a SnapMirror relationship from a 64-bit volume to a volume in a 32-bit aggregate as long as the volume size is less than 16TB. This results in having a 64-bit destination volume residing in a 32-bit aggregate.

Figure 5): 64-bit to 32-bit SnapMirror relationship.



2. It is also possible to create a SnapMirror relationship from a 32-bit volume to a volume in a 64-bit aggregate. This results in having a 32-bit destination volume residing in a 64-bit aggregate.

Figure 6): 32-bit to 64-bit SnapMirror relationship.



Note: A SnapMirror destination is skipped when the aggregate is expanded to 64-bit. The 32-bit SnapMirror destination is expanded to 64-bit only on a SnapMirror break or if the source gets expanded to 64-bit.

SnapMirror Flip Resync

Consider the preceding scenario in which a 32-bit source volume is mirrored with SnapMirror to a volume in a 64-bit aggregate. 64-bit expansion is triggered on the 32-bit destination volume when the SnapMirror relationship is broken. If the administrator attempts to do a flip resync from the destination volume back to the original 32-bit source volume, the original 32-bit source volume is converted to a 64-bit volume in a 32-bit aggregate.

Existing SnapMirror Relationship May Be Interrupted Due to Volume Size Change

If a 64-bit expansion on an aggregate containing a 32-bit SnapMirror source volume results in having the size of the source volumes increased, SnapMirror transfers fail if the source volume is then bigger than the destination volume.

SnapVault

Starting with clustered Data ONTAP 8.2, SnapVault is supported only on 64-bit aggregates. This means both source and destination volumes must be in 64-bit format and on 64-bit aggregates to establish relationships.

7-Mode

In Data ONTAP 8.0 operating in 7-Mode, the features such as volume SnapMirror, aggr copy, and vol copy, are restricted to aggregates and volumes of the same type. Starting in Data ONTAP 8.1, you can create replication relationships between volumes of different formats. Data ONTAP 8.3 do not support operation in 7-Mode.

Refer to the supportability matrix in Table 4 for all supported SnapMirror relationships across different formats and Data ONTAP kernels for 7-Mode.

Table 4) SnapMirror supportability matrix for 7-Mode.

Source Volume	Destination Aggregate						
	7.3.x (32-Bit)	8.0.x (32-Bit)	8.0.x (64-Bit)	8.1.x (32-Bit)	8.1.x (64-Bit)	8.2.x (32-Bit)	8.2.x (64-Bit)
7.3.x (32-bit)	Yes	Yes	No	Yes	Yes	Yes (7.3.7 onwards)	Yes (7.3.7 onwards)
8.0.x (32-bit)	No	Yes	No	Yes	Yes	Yes	Yes
8.0.x (64-bit)	No	No	Yes	Yes	Yes	Yes	Yes
8.1.x (32-bit)	No	No	No	Yes	Yes	Yes	Yes
8.1.x (64-bit)	No	No	No	Yes	Yes	Yes	Yes
8.2.x (32-bit)	No	No	No	No	No	Yes	Yes
8.2.x (64-bit)	No	No	No	No	No	Yes	Yes

Volume SnapMirror (VSM)

Volume SnapMirror is an efficient data replication product that transfers only the modified data disk blocks to the destination after the initial base transfer. The information in this section is applicable to volume SnapMirror in asynchronous mode, semi-synchronous mode, and synchronous mode. Because volume SnapMirror works at the block level, the source and destination must both be FlexVol volumes in the same aggregate type. For Data ONTAP 8.0.x operating in 7-Mode, the source and destination of a volume SnapMirror relationship must be either FlexVol volumes in 32-bit aggregates or both FlexVol volumes in 64-bit aggregates. Therefore you cannot create a volume SnapMirror relationship from a FlexVol volume in a 32-bit aggregate to a FlexVol volume in a 64-bit aggregate or vice versa. Table 4 shows the volume SnapMirror combinations that are possible and not possible in Data ONTAP 8.0.x operating in 7-Mode. This restriction is lifted in Data ONTAP 8.1, allowing you to mirror with SnapMirror between any format aggregates.

Qtree SnapMirror (QSM) and SnapVault

Qtree SnapMirror enables replication of the source volume or qtree to a destination qtree. SnapVault software is a product for protecting data against loss and preserving old versions of data. SnapVault replicates data in primary system paths to qtrees on a SnapVault secondary storage system. SnapVault uses the qtree replication engine to perform replication.

Qtree SnapMirror and SnapVault in 7-Mode work at the logical file level and are not aware of the underlying physical storage structure or memory addressing scheme. Therefore, they can be used to transfer data between FlexVol volumes located in any kinds of aggregates. You can establish a qtree SnapMirror relationship by using a qtree or a FlexVol volume in a 32-bit aggregate as the source and a qtree inside a FlexVol volume in a 64-bit aggregate as the destination, and vice versa.

Table 5) Qtree SnapMirror and SnapVault configurations in 7-Mode.

Source	Destination	
	Volume in 32-Bit Aggregate	Volume in 64-Bit Aggregate
Volume in 32-bit aggregate	Yes	Yes
Volume in 64-bit aggregate	Yes	Yes

SyncMirror

SyncMirror® is an optional feature in Data ONTAP operating in 7-Mode. It is used to mirror data to two separate identical aggregates that are connected to the same storage system. It provides synchronous mirroring of data at the RAID level. Therefore data gets replicated in the aggregate level. The SyncMirror feature can be used with both 32-bit and 64-bit aggregates.

MultiStore

MultiStore® software is an optionally licensed feature of Data ONTAP 8.0 operating in 7-Mode that enables you to partition the storage and network resources of a single storage system so that it appears as multiple storage systems on the network. Each storage system created as a result of the partitioning is called a vFiler® unit. For details, see the “MultiStore Management Guide” on the NetApp Support site.

MultiStore works with 64-bit aggregates. However, certain restrictions must be followed when setting up a vFiler unit for disaster recovery or data migration relationships.

- For a vFiler unit that uses qtrees as its primary storage unit and for all other storage units associated with it, vFiler disaster recovery and data migration operations can be performed to any other qtree in

any FlexVol volume, regardless of the underlying aggregate type. Therefore, for a vFiler unit that uses a qtree located in a FlexVol volume in a 32-bit aggregate as its primary storage unit and some other qtrees located in FlexVol volumes in 32-bit aggregates as its storage units, you can perform disaster recovery and migration operations to qtrees located in FlexVol volumes in either 32-bit aggregates or 64-bit aggregates.

- If a vFiler unit has a qtree as its primary storage unit but has at least one other FlexVol volume as its storage unit, you can perform disaster recovery and data migration operations only to another FlexVol volume that is located in the same aggregate type as the aggregate type in which the storage unit FlexVol volume is located. Therefore, if the vFiler unit has a qtree as its primary storage unit and a FlexVol volume in a 32-bit aggregate as one of its storage units, you can perform disaster recovery and data migration operations only to another FlexVol volume that is located in a 32-bit aggregate.
- If a vFiler unit has a FlexVol volume as its primary storage unit, you can perform disaster recovery and data migration operations only to another FlexVol volume in the same aggregate type as the aggregate type in which the primary storage unit FlexVol volume is located. Therefore, if the vFiler unit has a FlexVol volume in a 32-bit aggregate as its primary storage unit, you can perform disaster recovery and data migration operations only to another FlexVol volume that is located in a 32-bit aggregate.

To summarize, if a vFiler unit has a FlexVol volume as its primary storage unit or as one of its storage units, you can perform disaster recovery and data migration operations only to a FlexVol volume that is located in the same aggregate type as its storage unit.

Aggr Copy and Vol Copy

Aggr copy and vol copy operations work at the physical level. They work across similar types of aggregates and FlexVol volumes in Data ONTAP 8.0.x operating in 7-Mode. Therefore, when using an aggr copy operation in Data ONTAP 8.0.x 7-Mode to copy data from one aggregate to another, both the source and destination aggregates must be of the same type. Similarly, for vol copy in Data ONTAP 8.0.x operating in 7-Mode, both the source and destination must be FlexVol volumes contained in the same aggregate type.

Data ONTAP 8.1 and newer releases lifts this restriction, allowing you to use vol copy or aggr copy between any format aggregates.

4.9 Transitioning from 7-Mode to Clustered Data ONTAP 8.3

Starting in clustered Data ONTAP 8.3, 32-bit aggregates are deprecated and not supported. Copy-based migration methods can be used to migrate data regardless of the source and destination aggregate types. However, replication-based migration methods cannot migrate a 32-bit aggregate from a source 7-Mode storage system to a 64-bit aggregate in clustered Data ONTAP 8.3. If the destination cluster is running clustered Data ONTAP 8.3 it will be necessary to convert the source aggregates to 64-bit and age out (or delete) 32-bit Snapshot copies prior to executing the data migration. For details, refer to TR-4052: Successfully Transitioning to Clustered Data ONTAP. If you are unsure if you have 32-bit Snapshot copies present in your 64-bit aggregate, contact NetApp Support for assistance.

4.10 Summary of Feature Compatibility with the 64-Bit Expansion Process

Table 6 provides additional feature interoperability details for 64-bit aggregates.

Table 6) Interoperability with Data ONTAP features.

Feature	Compatibility
Qtree SnapMirror and SnapVault (supported in 7-Mode only)	No restriction. It can be triggered during the expansion process and also works between 32-bit and 64-bit volumes.
NDMP	No restriction. It can be triggered during the expansion process. The data can be restored to any format volume.
Synchronous SnapMirror (supported in 7-Mode only)	No restriction. It can be triggered during the expansion process and also works between 32-bit and 64-bit volumes.
Asynchronous SnapMirror	No restriction. It can be triggered during the expansion process and also works between 32-bit and 64-bit volumes.
SnapLock® (supported in 7-Mode only)	No restriction. Expansion can be triggered on an SLE/SLC aggregate.
RAID SyncMirror and MetroCluster™	No restriction. When disks are added to a RAID SyncMirror configuration, the disks are split between the two plexes.
Vol copy	No restriction. It can be triggered during the expansion process and also works between 32-bit and 64-bit volumes.
Aggr copy (supported in 7-Mode only)	No restriction. It can be triggered during the expansion process and also works between 32-bit and 64-bit aggregates.
Compression	No restriction. Compression can be enabled only on 64-bit volumes. Compression can be triggered during the expansion process. Best Practice 7: Enable Compression Only After Expansion Process Is Complete
Deduplication	No restriction. The expansion process does not touch data blocks so deduplication is not affected. It can also be triggered during the expansion process.
DataMotion for volumes in 7-Mode	It can be triggered during the expansion process. It does not work across aggregate types. Starting in Data ONTAP 8.1.1, this has been enhanced to support a <code>vol move</code> from a 32-bit source aggregate to a 64-bit destination aggregate.
DataMotion for Volumes in clustered Data ONTAP	It can be triggered during the expansion process and also works between 32-bit and 64-bit volumes.

Feature	Compatibility
FlexCache	FlexCache volumes need to be destroyed before the expansion process can be triggered. They can be recreated after the expansion process has started. There is no restriction on the format of the origin and cache volumes.
Reallocate	No restrictions.
FlexClone	If an aggregate containing a FlexClone volume is expanded, the FlexClone volume will also be expanded to 64-bit, causing some minimal loss of savings for the metadata blocks of the clone.
SnapRestore®	No restrictions. Restoring to a 32-bit Snapshot copy automatically triggers format change.
Volume transition	No restriction. Volumes can be transitioned while the expansion process is in progress.
Autosize	Any autosize changes are not honored until the completion of the expansion process.
Failover/failback	No restrictions. The expansion process has no impact on failover/failback.
DataMotion for vFiler (7-Mode)	This will not work across aggregate types.

4.11 Interoperability Matrix Tool (IMT)

Refer to the [Interoperability Matrix Tool \(IMT\)](#) on the NetApp Support site to validate that the exact product and feature versions described in this document are supported for your specific environment. The NetApp IMT defines the product components and versions that can be used to construct configurations that are supported by NetApp. Specific results depend on each customer's installation in accordance with published specifications.

5 Performance

This section is an overview of the performance characteristics of 64-bit aggregates. It also offers recommendations on the generic workloads that have better performance and those that will be affected by 64-bit aggregates.

5.1 Better Performance with More Disks

64-bit aggregates enable you to add more data disks to an aggregate than do 32-bit aggregates. When you create a FlexVol volume inside a 64-bit aggregate, the FlexVol volume is striped across all the disks in the aggregate. Therefore the FlexVol volume has the capability to use more spindles. This means that for workloads and storage system configurations that were constrained by the number of available disk drive spindles, 64-bit aggregates can provide better performance provided you can add enough drives such that the workload or storage system is no longer disk spindle constrained in delivering performance.

5.2 Memory Usage

64-bit aggregates have a larger address space and need more memory for their metadata, compared to 32-bit aggregates. This extra memory usage reduces the amount of memory available for user data. Therefore, for workloads that are memory intensive, you may experience a slight performance impact when running the workload from a FlexVol volume contained in a 64-bit aggregate compared to running the workload from a volume contained in a 32-bit aggregate.

Workloads that are highly random in nature typically access more metadata over a given period of time compared to sequential workloads. Random read workloads with a very large active dataset size may experience a performance impact when run on a FlexVol volume in a 64-bit aggregate, compared to when run on a FlexVol volume in a 32-bit aggregate. This is because the dataset size combined with the increased metadata size can increase memory pressure on the storage system and result in an increased amount of on-disk I/O. In such scenarios, if you want to run the random workload from a volume contained in a 64-bit aggregate, using FlexCache improves the performance delivered by the storage system and helps alleviate any performance impact seen with 64-bit aggregates.

Note that just having a 64-bit aggregate on the storage system does not result in any sort of performance degradation. The effects on performance, if any, are seen when data in any FlexVol volume in the 64-bit aggregate starts to be accessed.

The overall performance of a storage system depends on a variety of factors, such as the configuration of the system, the workload, and activities on the storage system itself. For an analysis of your specific storage system performance, contact your NetApp system engineer.

5.3 Consideration During 64-Bit Expansion Process

Read/Write Throughput and Latency Impact During Expansion

For client operations on inodes that are not yet fully expanded, additional block lookups are sometimes required to perform the 64-bit expansions inline. Because the client operations may also trigger the 64-bit expansion, this may increase the latency and reduce throughput. After an inode is fully expanded to 64-bit, the extra block lookups are not required. The performance has the highest hit when the expansion is first triggered, then improves gradually as the expansion process progresses. Our SPC-1 lab results showed 10% increase in latency when expansion process is triggered and improved gradually as expansion process progressed over time.

Performance Impact on SnapMirror Destinations

If an aggregate containing a 32-bit SnapMirror destination volume is expanded to 64-bit before the 32-bit source aggregate, read throughput may be affected on the destination.

This problem only exists until the source volumes are expanded to 64-bit format and the new blocks are transferred to the destination. Until then, there can be a performance hit for read workloads due to the extra lookups required to resolve physical addresses. The performance degradation for sequential read workloads on such destinations is no more than 15% for various datasets in the worst case.

Completion Time of Expansion Process

The expansion process updates the format of the indirect metadata in the aggregate and the volume. Therefore, the completion time for the expansion process depends on the amount of indirect metadata. The amount of indirect metadata will vary based on the number of files and the amount of used data. Typically, it accounts for 2% or less of the dataset. Depending on the number of indirect blocks in the aggregate, completion time vary.

6 Best Practices

6.1 Best Practices Table

Best Practice 1: Use 64-Bit Aggregates

The 64-bit aggregates support larger threshold sizes which allows an increased number of larger sized disks in a single aggregate. This improves performance with simplified management. Starting with clustered Data ONTAP 8.3, the support for 32-bit aggregates is deprecated.

Best Practice 2: RAID Type and RAID Group Sizes

The recommended RAID type (RAID-DP®) and RAID group sizes for 64-bit aggregates are the same as recommendations for 32-bit aggregates. For details, refer to NetApp TR-3437: Storage Subsystem Resiliency Guide.

Best Practice 3: Maximum Aggregate Size

The maximum aggregate size for 32-bit aggregates is 16TB. 64-bit aggregates can be much larger in size than 32-bit aggregates and vary with respect to the storage controller and the maximum number of volumes supported. For exact values, refer to NetApp Hardware Universe, <http://hwu.netapp.com/>.

Best Practice 4: Recommended Root Volume Size

For the recommended size of the root volume, see NetApp Hardware Universe, <http://hwu.netapp.com/>.

Best Practice 5: Maximum File and LUN Size

For the exact size thresholds, refer to NetApp Hardware Universe, <http://hwu.netapp.com/>. Make sure you consider these sizes when you design your storage solution.

Best Practice 6: Consider Flash Cache for Highly Random Workloads

When running a highly random workload with a huge active dataset on a FlexVol volume located in a 64-bit aggregate, be sure to monitor the performance of the storage system during periods of high activity for a few days. 64-bit aggregates might cause more memory usage than 32-bit aggregates. If you observe any performance impact, adding Flash Cache may help improve the performance. For details, refer to TR-3832: Flash Cache Best Practices Guide.

Best Practice 7: Enable Compression Only After Expansion Process Is Complete

32-bit aggregates do not support compression. Compression can be enabled on the volumes as soon as the expansion process is started on the hosting aggregate. Enabling compression only after the expansion process is completed is recommended. The progress of the expansion process can be tracked using the `aggr 64bit-upgrade status` command.

Best Practice 8: Ensure Adequate Space for Volumes Before Expansion Process Is Started

When the expansion process converts 32-bit block pointers to 64-bit block pointers, some additional space is required for expansion process to be successful. When the expansion process is started, a prequalification check is initiated and will report the space requirements. The prequalification check will fail if volumes do not have sufficient space for expansion. Volumes may continue to be assessed by clients and continue their I/O while the expansion process is taking place on an aggregate. It is important to adequately size the volumes considering all factors upfront for the expansion process to complete. When the expansion process fails due to space limitation, additional space can be allocated to the volume and the expansion process can be resumed as space becomes available.

Best Practice 9: Considerations to Start the Expansion Process

- Destroy all FlexCache volumes in the aggregate you want to expand.
- Verify that the aggregate will not be a source or destination of DataMotion for volumes for 7-Mode.
- Verify the system has available space prior to triggering the expansion process.
- For best performance, when adding disks to trigger expansion prior to Data ONTAP 8.2.1, add a complete RAID group to prevent the new disks from becoming a performance bottleneck as all new client writes are directed to the newly added disks.
- Ensure the volumes in the 32-bit aggregate are not close to full.
- Verify the prequalification scanner has finished running before triggering the 64-bit expansion.
- Do not change any autosize settings during the expansion process. Autosize changes are not effective until the expansion process completes.

Best Practice 10: Expand Source Aggregate Before Expanding Destination Aggregate

Expand the source aggregate to 64-bit before expanding the destination aggregate.

If the destination aggregate needs to be expanded first, follow up by expanding your source aggregate as soon as possible.

After expansion of the source aggregate, the destination volume gets automatically converted to the format of the source volume during the next SnapMirror sync. If the destination aggregate gets expanded first, the destination volume is skipped from the expansion and may stay in 32-bit format (of the source) while the underlying aggregate is in 64-bit format. The reads in destination volume requires translating 32-bit pointers to new 64-bit pointers which could cause additional delay in response. It is a best practice to expand the source first and then expand the destination aggregate.

6.2 Disclaimer

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.

7 Frequently Asked Questions (FAQ)

What are 64-bit aggregates? What are larger aggregates?

Answer: They both mean the same. These aggregates utilize 64-bit indirect block pointers, allowing support for aggregate capacity greater than 16TB. Furthermore, 64-bit aggregates provide substantial improvements over 32-bit aggregates, such as compression, Flash Pool™ intelligent caching, storage-efficient vaulting, and more.

Can 64-bit aggregates and 32-bit aggregates coexist in Data ONTAP 8.1 and 8.2?

Answer: Yes. In Data ONTAP 8.1 and 8.2, 32-bit aggregates and 64-bit aggregates can coexist. Starting with clustered Data ONTAP 8.3, 32-bit aggregates will not be supported and can not exist.

Can 64-bit aggregates and 32-bit aggregates coexist in clustered Data ONTAP 8.3?

Answer: No. Starting in clustered Data ONTAP 8.3, 32-bit aggregates will not be supported and can not exist.

What is the default format of the aggregate in Data ONTAP 8.1, 8.2 and 8.3?

Answer: Any new aggregate in Data ONTAP 8.1 or newer is always created in the 64-bit format.

What is the format of a volume within an aggregate?

Answer: FlexVol volumes inherit the block format from the aggregate. A 32-bit aggregate contains 32-bit volumes. A 64-bit aggregate contains 64-bit volumes. This rule applies to all volumes except SnapMirror® destinations. SnapMirror destinations inherit the block format of the SnapMirror source volume irrespective of the block format of the hosting aggregate.

Can I shrink an aggregate?

Answer: No. The size of an aggregate cannot be decreased.

How does the 64-bit expansion process work? Will it affect client access?

Answer: The expansion process scans the aggregate and all the volumes within the aggregate and expands all indirect blocks in the 32-bit format to 64-bit format. The process is seamless for clients. The 64-bit expansion is in place and does not require any data movement or data migration. It expands the 32-bit data to 64-bit while the volumes are online. Clients can continue to access their data without any disruption.

Can 64-bit expansion be triggered in Data ONTAP operating in 7-Mode and clustered Data ONTAP?

Answer: Yes. This feature is supported in both modes. Starting clustered Data ONTAP 8.3, this feature is disabled since this release does not support 32-bit aggregates.

Can I expand my 32-bit aggregate to 64-bit without the addition of disks?

Answer: Starting in Data ONTAP 8.2.1, additional disks are not required and the expansion process can be triggered with advanced-level CLI commands. Prior to Data ONTAP 8.2.1, additional disks are required to trigger the 64-bit expansion.

Will 32-bit aggregates be supported in future Data ONTAP releases?

Answer: No. Starting in Data ONTAP 8.3, only 64-bit aggregates will be supported. 64-bit aggregates provide substantial improvements over 32-bit aggregates, including compression, Flash Pool™, storage-efficient vaulting, and new Data ONTAP features. These enhancements, and the Data ONTAP 64-bit architecture, have made 64-bit aggregates the preferred format for today's storage environments.

Can I upgrade Data ONTAP 8.x with 32-bit aggregates to clustered Data ONTAP 8.3?

Answer: No. 32-bit aggregates must be converted to 64-bit aggregates prior to upgrading to Data ONTAP 8.3. When upgrading to Data ONTAP 8.3 only upgrades from Data ONTAP 8.2.1, or newer, will be supported.

Will an update to Data ONTAP 8.1 or 8.2 trigger the expansion process?

Answer: No. The expansion process does not get triggered automatically and it must be explicitly started by the administrator. In Data ONTAP 8.1 and 8.2, the expansion process can be triggered only with the addition of disks if the size of the aggregate exceeds 16TB. Starting in Data ONTAP 8.2.1, the expansion process can be triggered by administrators using CLI commands in advanced-privilege mode.

Can I expand my 32-bit root aggregate to 64-bit?

Answer: If there is a strong requirement to expand your root aggregate beyond 16TB in Data ONTAP 8.1 or 8.2, you can add disks and trigger the 64-bit expansion on the root aggregate. However, NetApp recommends that a separate aggregate be created to host the data volumes. Starting with Data ONTAP 8.2.1, you can expand the root aggregate to 64-bit without adding disks.

Can I convert my 64-bit aggregate to 32-bit?

Answer: No. This is not supported irrespective of the size of the aggregate.

Can 64-bit expansion be triggered on an individual volume?

Answer: No. The only supported mechanism is to perform the expansion at the aggregate level. It then proceeds to expand all the volumes within the aggregate.

Are blocks in Snapshot copies also converted to the 64-bit format?

Answer: No. Snapshot copies are read-only and are not updated by the expansion process. The expansion process updates indirect blocks in the active file system.

How can I find out if I have 32-bit Snapshot copies?

Answer: Starting in Data ONTAP 8.2.1, you can use `snap list -fs-block-format <aggr_name or vol_name>` to identify any 32-bit Snapshot copies in the volume or aggregate. Data ONTAP 8.1 and 8.2 do not have a command to identify 32-bit Snapshot copies.

How do I convert 32-bit Snapshot copies to 64-bit?

Answer: 32-bit Snapshot copies cannot be converted to 64-bit and must be expunged from the system.

- 32-bit Snapshot copies must either be deleted manually by the administrator or wait to be cycled off and replaced by a newer 64-bit Snapshot copies.
- For 32-bit SnapMirror destination volumes, the 32-bit Snapshot copies need to be cycled off on the SnapMirror source volumes, followed by a number of SnapMirror updates. This can take weeks to months, depending on the Snapshot copy schedule.
- If a 32-bit Snapshot copy is a base Snapshot copy for a clone volume and locked down, the user must split the clone volume from the parent first, then remove the 32-bit Snapshot copy.

Does the expansion process restart from the beginning if interrupted?

Answer: No. The expansion process maintains checkpoints. If the process is interrupted due to CFO or SFO events and/or space issues, it resumes from the latest checkpoint.

Why would the expansion process run out of space even though volume size was increased earlier?

Answer: Even if space was added to a volume before the expansion process, it is still possible to run out of space because client traffic may consume additional space. This scenario is less likely in Data ONTAP 8.1 because the expansion is triggered by the addition of disks, which results in more usable capacity.

What happens if I run out of space during the expansion process?

Answer: If the volume runs out of space during the expansion process, the process aborts. Clients may see ENOSPC errors. The appropriate volume autogrow/autodelete configurations will be triggered.

Meanwhile, the administrator can create more space by doing on the following:

1. Growing the volume
2. Adding more disks

The expansion process is restarted automatically when space is available.

Best Practice

Best Practice 8: Ensure Adequate Space for Volumes Before Expansion Process Is Started

What happens to quotas during the expansion process?

Answer: If the quotas are close to the limit, it is possible to exceed quota limits during the expansion process due to additional blocks being required.

Can an aggregate be expanded to 64-bit if it contains a SnapMirror source/destination?

Answer: Yes. An aggregate that contains one or more SnapMirror source/destination volumes is expanded from 32-bit to 64-bit. All activities related to SnapMirror, such as SnapMirror break, resync, and

so on, remain functioning on the SnapMirror source/destination volumes during and after the expansion process. The SnapMirror source will be expanded to 64-bit. The SnapMirror destination will always be skipped. SnapMirror destination volumes follow the format of their source volumes.

When are SnapMirror 32-bit destinations expanded to 64-bit?

Answer: A 32-bit read-only SnapMirror destination in a 64-bit aggregate is expanded to 64-bit only when the SnapMirror relationship is broken or when the source gets expanded to 64-bit.

What happens if I use SnapRestore to restore the volume to a Snapshot copy in 32-bit format?

Answer: If the parent aggregate is 64-bit, then restoring a volume with SnapRestore® technology to a 32-bit Snapshot copy will automatically trigger the 64-bit format change.

Additional Resources

- NetApp Hardware Universe
<https://hwu.netapp.com>
- TR-3437: Storage Best Practices and Resiliency Guide
<http://www.netapp.com/us/library/technical-reports/tr-3437.html>
- TR-4052: Successfully Transitioning to Clustered Data ONTAP
<http://www.netapp.com/us/system/pdf-reader.aspx?m=tr-4052.pdf>
- TR-4070: Flash Pool Design and Implementation Guide
<http://www.netapp.com/us/system/pdf-reader.aspx?m=tr-4070.pdf>
- TR-3832: Flash Cache Best Practices Guide
<http://www.netapp.com/us/system/pdf-reader.aspx?m=tr-3832.pdf>
- Data Protection Online Backup and Recovery Guide for 7-Mode
https://library.netapp.com/ecm/ecm_download_file/ECMP1635994
- Data ONTAP 8 Product Library
<https://mysupport.netapp.com/documentation/productlibrary/index.html?productID=30092>

Contact Us

Let us know how we can improve this technical report.

Contact us at docfeedback@netapp.com.

Include TECHNICAL REPORT 3978 in the subject line.

Appendix

Expansion Process Examples Using Command Line Interface (CLI)

Data ONTAP 8.1 and 8.2 Operating in 7-Mode

Step 1: Establish Requirement to Grow Aggregate Beyond 16TB.

In the following example, aggr1 is a 32-bit aggregate and is less than 16TB.

```
elis22b> aggr status aggr1
      Aggr State      Status      Options
      aggr1 online    raid_dp, aggr
                        32-bit

      Volumes: vol1, vol2

elis22b> df -A aggr1
Aggregate      kbytes      used      avail      capacity
aggr1          9358087680  2159414732  7198672948      23%
```

Step 2: Add Disks to Grow Aggregate with Capacity Size Less than 16TB.

When trying to add disks, if the resulting size of the aggregate is less than 16TB, the aggregate will continue to be in 32-bit format.

```
elis22b> aggr add aggr1 14
Note: preparing to add 12 data disks and 2 parity disks.
Continue? ([y]es, [n]o, or [p]review RAID layout) y

elis22b> aggr status aggr1
      Aggr State      Status      Options
      aggr1 online    raid_dp, aggr
                        32-bit

elis22b> df -A aggr1
Aggregate      kbytes      used      avail      capacity
aggr1          14037131520  2159414748  11877716772      15%
```

Step 3: Add Disks to Grow Aggregate with Capacity Size More than 16TB.

When trying to add disks, if the resulting size of the aggregate is greater than 16TB, the command will fail and the administrator will be warned that the aggregate needs to be expanded to 64-bit with the new size.

```
elis22b> aggr add aggr1 12
Note: preparing to add 10 data disks and 2 parity disks.
Continue? ([y]es, [n]o, or [p]review RAID layout)

File system size 18.56 TB exceeds maximum 15.99 TB
Addition of disks failed. To proceed with this operation, and upgrade the aggregate from 32-bit
format to 64-bit format, follow these steps:
1. Run the "aggr add" command with the "-64bit-upgrade check" option to determine whether there
is enough free space for the 64-bit upgrade.
2. Resolve any space issues identified by the "-64bit-upgrade check" option.
```

3. Run the "aggr add" command with the "-64bit-upgrade normal" option to add the disks and trigger 64-bit upgrade on the aggregate aggr1.

aggr add: Cannot add specified disks to the aggregate because aggregates in 32-bit block format cannot be larger than 16TB.

Step 4: Add Disks with 64-Bit Option to Trigger Expansion Process.

The administrator is now required to specify the `-64bit-upgrade` option when adding disks to trigger the 64-bit expansion.

```
elis22b> aggr add aggr1 -64bit-upgrade normal 12
Note: preparing to add 10 data disks and 2 parity disks.
Continue? ([y]es, [n]o, or [p]review RAID layout) y
File system size 18.56 TB exceeds maximum 15.99 TB
Checking for additional space required to upgrade all writable 32-bit
volumes in aggregate aggr1 (Ctrl-C to interrupt).....
Addition of 12 disks to the aggregate has completed.
[elis22b:waf1.scan.64bit.upgrade.start:notice]: The 64-bit upgrade scanner has started running on
volume voll.
[elis22b:waf1.scan.64bit.upgrade.start:notice]: The 64-bit upgrade scanner has started running on
aggregate aggr1.
[elis22b:waf1.scan.64bit.upgrade.completed:notice]: The 64-bit upgrade scanner has completed
running on volume voll.
[elis22b:waf1.scan.64bit.upgrade.completed:notice]: The 64-bit upgrade scanner has completed
running on aggregate aggr1..
```

If any of the volumes are too full, the command will error out. The administrator can run the space check command `aggr add aggr1 -64bit-upgrade check` to get more information on the volumes that need to be grown and the amount of space that needs to be added to each of those volumes.

Step 5: Monitor Status of Expansion Process.

The expansion process can be tracked, if needed. The `status` command also provides a time estimate in addition to its current progress.

```
elis22b*> aggr 64bit-upgrade status aggr1 -all
Volume  Format      Scanner Status  % Completed  Time to Completion  Progress
aggr1   upgrading  running      6            2976          fbn 0, inode 79 of 32781,
private
  voll  upgrading  running     10            1719          fbn 0, inode 32664 of 32781,
public
  vol2  upgrading  running     10            1736          fbn 0, inode 32661 of 32781,
public
```

Step 6: Aggregate Now in 64-Bit Format.

As soon as the expansion process is triggered, the aggregate is considered to be a 64-bit aggregate. Even if the process has not finished, the aggregate is a 64-bit aggregate with a mix of 32-bit and 64-bit indirect blocks.

```
elis22b*> aggr status aggr1
      Aggr State      Status      Options
aggr1 online         raid_dp, aggr
                    parity uninit'd!
                    64-bit
```

Clustered Data ONTAP 8.2.1

Step 1: Establish Requirement to Expand from a 32-Bit Aggregate to 64-Bit Aggregate.

In the following example, aggr1 is a 32-bit aggregate.

```
cluster::> aggr show aggr1 -fields aggregate, state, raidtype, block, volcount, size
aggregate block-type raidtype size state volcount
-----
aggr1 32-bit raid_dp 168.1MB online 3

cluster::> vol show -aggregate aggr1 -fields block
(volume show)
vserver volume block-type
-----
vs1 aggr1_root 32-bit
vs1 vol1 32-bit
vs1 vol2 32-bit
3 entries were displayed.

cluster::> df -A aggr1
Aggregate kbytes used avail capacity
aggr1 172112 145272 26840 84%
```

Step 2: Complete Command Line Command to Begin Expansion Process.

This step initiates the 32-bit to 64-bit aggregate expansion process. Starting in Data ONTAP 8.2.1, administrators can convert 32-bit aggregates to 64-bit aggregates without the addition of new disks. Administrators must have advanced privileges to perform the commands for in-place expansion. All new writes are in 64-bit format.

```
cluster::> set -privilege advanced

Warning: These advanced commands are potentially dangerous; use them only when
        directed to do so by NetApp personnel.
Do you want to continue? {y|n}: y

cluster::*> aggr 64bit-upgrade start -aggregate aggr1

Checking for additional space required to upgrade all writable 32-bit volumes
in aggregate aggr1 (Ctrl-C to interrupt).....
Started 64-bit upgrade on aggregate aggr1 and all contained writable volumes.
```

Step 3: Monitor Status of Expansion Process

The expansion process can be tracked, if needed. The `status` command also provides a time estimate in addition to its current progress.

```
cluster::*> aggr 64bit-upgrade status aggr1
Volume Format Scanner Status % Completed Time to Completion Progress
aggr1 upgrading running 80 29 minutes fbn 0, inode 79 of 32781, private
vol1 upgrading running 74 12 minutes fbn 0, inode 28664 of 32781, public
vol2 upgrading running 92 9 minutes fbn 0, inode 306611 of 32781, public
```

Step 4: Aggregate Now in 64-Bit Format

As soon as the expansion process is triggered, the aggregate is considered to be a 64-bit aggregate. Even if the process has not finished, the aggregate is a 64-bit aggregate with a mix of 32-bit and 64-bit indirect blocks.

```
cluster::*> aggr show aggr1 -fields aggregate, state, raidtype, block, volcount, size
aggregate block-type raidtype size state volcount
-----
aggr1 64-bit raid_dp 168.1MB online 3
elis22b*> vol show -aggregate aggr1 -fields block
(volume show)
vserver volume block-type
-----
vs1 aggr1_root 64-bit
vs1 vol1 64-bit
vs1 vol2 64-bit
3 entries were displayed.
```

Running Check Before Starting Expansion Process

This example shows a volume with a lot of metadata. The message shows that additional space is required for the volume for the expansion process to be successful.

```
cluster::*> aggr 64bit-upgrade start -aggregate aggr32_2 -mode check

Checking for additional space required to upgrade all writable 32-bit volumes
in aggregate aggr32_2 (Ctrl-C to interrupt).....

Upgrading a volume to 64-bit consumes additional free space in the volume. The
following table shows the space usage after each volume is upgraded to 64-bit:
```

Vserver	Volume	Total	Used	Available	Capacity
vs1	gvvol	18GB	25.65GB	0B	100%

To initiate the upgrade of aggregate aggr32_2 to 64-bit, run this command with the "normal", "grow-reserved", or "grow-all" option.

Before initiating 64-bit upgrade with the "normal" option, increase volume sizes by the following amounts to enable successful upgrade of the aggregate:

Vserver	Volume	Additional Space Required for Upgrade
vs1	gvvol	7.65GB

If you use the "grow-reserved" or "grow-all" option, volume sizes will be increased automatically as needed during the upgrade.

```
cluster::*> aggr 64bit-upgrade status -aggregate aggr32_2
```

```
64-bit upgrade is not currently in progress on aggregate aggr32_2 or its
contained volumes.
```

```
cluster::*> aggr 64bit-upgrade start -aggregate aggr32_2
```

```
Error: command failed: 64-bit upgrade failed. Reason: One or more volumes do
not have enough free space to support the 64-bit upgrade. Run this
command with the "check" option to check for additional space required.
```

Check for Status and Allocate More Space if Interrupted Due to Space Issue

The expansion process resumes as space becomes available.

```
cluster::*> aggr 64bit-upgrade status -aggregate aggr32_2
```

Aggregate	Vserver	Volume	Format	Scanner	Status	Done	Time to Complete	Progress
aggr32_2	vs1	gvvol	upgrading	stopped	99%	00:00:00.000		fbn 635970, inode 16612 of 32781, public

```
cluster::*> df -h gvvol
```

Filesystem	total	used	avail	capacity	Mounted on	Vserver
/vol/gvvol/	25GB	25GB	0B	100%	---	vs1
/vol/gvvol/.snapshot	1382MB	6107MB	0B	442%	---	vs1

2 entries were displayed.

```
cluster::*> snap delete -volume gvvol -snapshot hourly.2014-11-26_0205
```

```
Warning: Deleting a Snapshot copy permanently removes any data that is stored
only in that Snapshot copy. Are you sure you want to delete Snapshot
copy "hourly.2014-11-26_0205" for volume "gvvol" in Vserver "vs1" ?
{y|n}: y
```

```
cluster::*> snap delete -volume gvvol -snapshot hourly.2014-11-26_0305
```

```
Warning: Deleting a Snapshot copy permanently removes any data that is stored
only in that Snapshot copy. Are you sure you want to delete Snapshot
copy "hourly.2014-11-26_0305" for volume "gvvol" in Vserver "vs1" ?
{y|n}: y
```

```
cluster::*> df -h gvvol
```

Filesystem	total	used	avail	capacity	Mounted on	Vserver
/vol/gvvol/	25GB	21GB	4724MB	82%	---	vs1
/vol/gvvol/.snapshot	1382MB	0B	1382MB	0%	---	vs1

2 entries were displayed.

```

cluster::*> aggr 64bit-upgrade status -aggregate aggr32_2
                Scanner      Time to
Aggregate Vserver  Volume      Format      Status Done Complete Progress
-----
aggr32_2 vs1      gvvol      upgrading running 99% 00:00:05.000
                                                fbn 713490,
                                                inode 16625
                                                of 32781,
                                                public

```

```

cluster::*> event log show -node gvgcluster-1-01 -time <2m
Time           Node           Severity      Event
-----
11/26/2014 03:27:44 gvgcluster-1-01 NOTICE      waf1.scan.64bit.upgrade.nospace: The 64-bit
upgrade scanner running on volume gvvol@vserver:73956df3-7407-11e4-878d-123478563412 was aborted
because there is insufficient free space.
11/26/2014 03:16:46 gvgcluster-1-01 NOTICE      waf1.scan.64bit.upgrade.nospace: The 64-bit
upgrade scanner running on volume gvvol@vserver:73956df3-7407-11e4-878d-123478563412 was aborted
because there is insufficient free space.

```

Refer to the [Interoperability Matrix Tool \(IMT\)](#) on the NetApp Support site to validate that the exact product and feature versions described in this document are supported for your specific environment. The NetApp IMT defines the product components and versions that can be used to construct configurations that are supported by NetApp. Specific results depend on each customer's installation in accordance with published specifications.

Copyright Information

Copyright © 1994–2015 NetApp, Inc. All rights reserved. Printed in the U.S. No part of this document covered by copyright may be reproduced in any form or by any means—graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system—without prior written permission of the copyright owner.

Software derived from copyrighted NetApp material is subject to the following license and disclaimer:

THIS SOFTWARE IS PROVIDED BY NETAPP "AS IS" AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE, WHICH ARE HEREBY DISCLAIMED. IN NO EVENT SHALL NETAPP BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

NetApp reserves the right to change any products described herein at any time, and without notice. NetApp assumes no responsibility or liability arising from the use of products described herein, except as expressly agreed to in writing by NetApp. The use or purchase of this product does not convey a license under any patent rights, trademark rights, or any other intellectual property rights of NetApp.

The product described in this manual may be protected by one or more U.S. patents, foreign patents, or pending applications.

RESTRICTED RIGHTS LEGEND: Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (c)(1)(ii) of the Rights in Technical Data and Computer Software clause at DFARS 252.277-7103 (October 1988) and FAR 52-227-19 (June 1987).

Trademark Information

NetApp, the NetApp logo, Go Further, Faster, ASUP, AutoSupport, Campaign Express, Cloud ONTAP, Customer Fitness, Data ONTAP, DataMotion, Fitness, Flash Accel, Flash Cache, Flash Pool, FlashRay, FlexArray, FlexCache, FlexClone, FlexPod, FlexScale, FlexShare, FlexVol, FPolicy, GetSuccessful, LockVault, Manage ONTAP, Mars, MetroCluster, MultiStore, NetApp Insight, OnCommand, ONTAP, ONTAPI, RAID DP, SANtricity, SecureShare, Simplicity, Simulate ONTAP, Snap Creator, SnapCopy, SnapDrive, SnapIntegrator, SnapLock, SnapManager, SnapMirror, SnapMover, SnapProtect, SnapRestore, Snapshot, SnapValidator, SnapVault, StorageGRID, Tech OnTap, Unbound Cloud, and WAFL are trademarks or registered trademarks of NetApp, Inc., in the United States and/or other countries. A current list of NetApp trademarks is available on the Web at <http://www.netapp.com/us/legal/netapptmlist.aspx>.

Cisco and the Cisco logo are trademarks of Cisco in the U.S. and other countries. All other brands or products are trademarks or registered trademarks of their respective holders and should be treated as such.
