



Technical Report

A Thorough Introduction to 64-Bit Aggregates

Shree Reddy, NetApp
September 2011 | TR-3786

CREATING AND MANAGING LARGER-SIZED AGGREGATES

The NetApp® Data ONTAP® 8.0 operating system operating in 7-Mode and NetApp Data ONTAP 8.0.1 operating in Cluster-Mode support a new aggregate type whose maximum size is much bigger than the size that was supported in earlier releases of Data ONTAP. This aggregate type, called 64-bit aggregate, offers a much larger size threshold while providing all the advantages and capabilities of aggregates, including flexibility and storage efficiency. FlexVol® volumes that are created in 64-bit aggregates also have a bigger maximum size threshold. This technical report describes 64-bit aggregates in detail.

TABLE OF CONTENTS

1	INTRODUCTION	4
2	THE NEED FOR 64-BIT AGGREGATES	6
2.1	DELIVERING PERFORMANCE WITH LARGE-SIZED SATA DRIVES	6
2.2	MAINTAINING HIGH STORAGE EFFICIENCY	6
2.3	NEED FOR LARGER FLEXVOL VOLUME SIZES	6
3	MAXIMUM AGGREGATE AND VOLUME SIZES	6
4	MAXIMUM NUMBER OF DISKS IN A MAXIMUM-SIZE AGGREGATE	7
5	CONFIGURATION AND OPERATION	7
5.1	RELEASE	7
5.2	LICENSE	7
5.3	PLATFORMS	7
5.4	CREATING A 64-BIT AGGREGATE	8
5.5	CREATING AND MANAGING FLEXVOL VOLUMES IN A 64-BIT AGGREGATE	8
5.6	IDENTIFYING A 64-BIT AGGREGATE	8
6	POSSIBLE AGGREGATE CONFIGURATIONS	9
7	ADVANTAGES OF 64-BIT AGGREGATES	9
7.1	BETTER PERFORMANCE FROM MORE DISKS IN THE AGGREGATE	9
7.2	MAINTAIN HIGHER STORAGE EFFICIENCY	9
7.3	EASE OF MANAGEMENT	10
7.4	USE ALL EXISTING DATA ONTAP FEATURES	10
7.5	BETTER SCALABILITY FOR FUTURE GROWTH	10
7.6	POWER AND SPACE SAVINGS	10
7.7	CONSOLIDATE DATA	10
8	INTERNAL OPERATION	10
9	DIFFERENCES BETWEEN 64-BIT AND 32-BIT AGGREGATES	12
9.1	ON-DISK ADDRESSING FORMAT	12
9.2	MEMORY USAGE	12
9.3	PERFORMANCE PROFILE	12
10	INTERACTION WITH OTHER DATA ONTAP 8.0.X FEATURES	12
10.1	SNAPMIRROR	12
10.2	NDMP AND DUMP	13
10.3	AGGR COPY AND VOL COPY (7-MODE ONLY)	14
10.4	DEDUPLICATION	16
10.5	MULTISTORE	16
10.6	PERFORMANCE ACCELERATION MODULE	17

10.7	FLEXCACHE.....	17
10.8	SYNCMIRROR	17
10.9	CONTROLLER FAILOVER.....	17
10.10	FLEXSHARE.....	18
11	PERFORMANCE	18
12	DATA MIGRATION.....	18
13	USE CASES.....	19
14	RECOMMENDATIONS.....	19
15	CONCLUSION	20
16	ACKNOWLEDGEMENTS.....	20
17	REFERENCES.....	20
17.1	NETAPP TECHNICAL REPORTS AND WHITE PAPERS	20
17.2	DATA ONTAP PRODUCT DOCUMENTATION.....	21

LIST OF TABLES

Table 1)	Maximum aggregate and FlexVol volume sizes.	6
Table 2)	Maximum number of data drives in 64-bit aggregates.....	7
Table 3)	Volume SnapMirror and SnapMirror sync configurations in Data ONTAP 8.0.x operating in 7-Mode.....	13
Table 4)	Volume SnapMirror configurations in Data ONTAP 8.0.x operating in Cluster-Mode.	13
Table 5)	Qtree SnapMirror and SnapVault configurations.	13

LIST OF FIGURES

Figure 1)	Existing aggregates on the storage system show up as 32-bit aggregates after upgrading to Data ONTAP 8.0 operating in 7-Mode/8.0.1 operating in Cluster-Mode.....	4
Figure 2)	Possible configurations for a NetApp storage system running Data ONTAP 8.0.x with 32-bit and 64-bit aggregates with the default 32-bit root aggregate.....	5
Figure 3)	32-bit addressing format thresholds constrain the aggregate and FlexVol volume size to 16TB.	11
Figure 4)	64-bit addressing makes it feasible to increase the aggregate size.	11
Figure 5)	Data replication utilities and their interaction with 32-bit and 64-bit aggregates (7-Mode only).	15

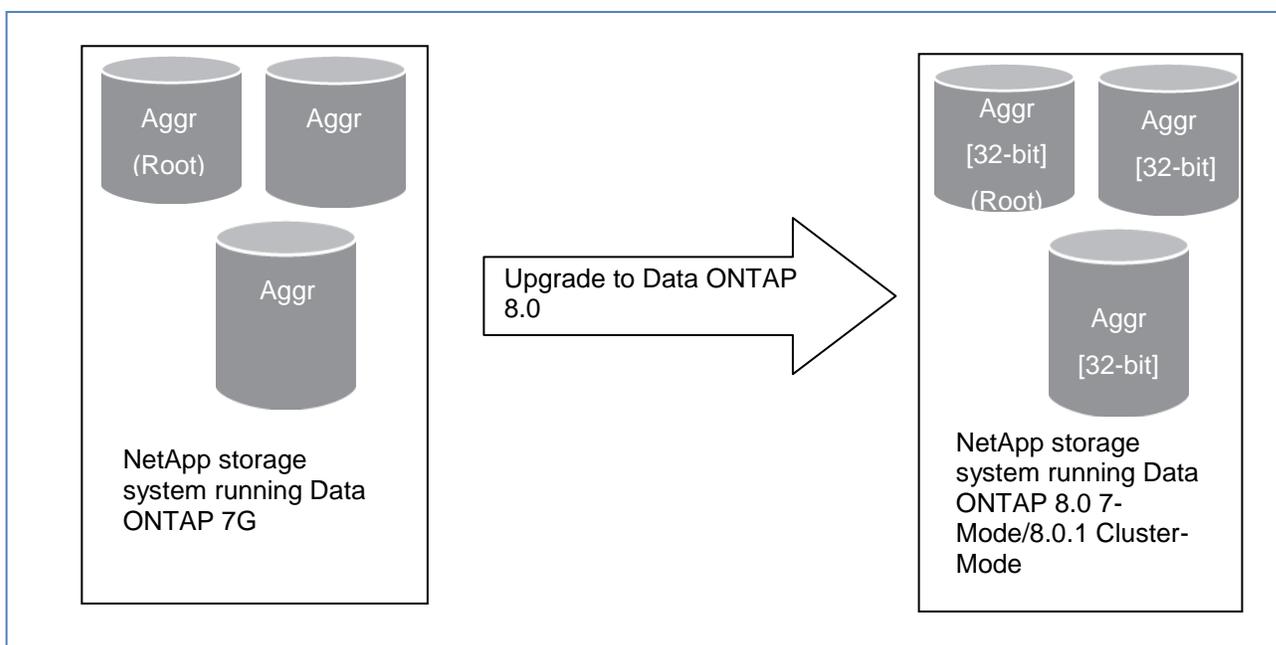
1 INTRODUCTION

Aggregates and FlexVol volumes are 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.

Data ONTAP 8.0 operating in 7-Mode and Data ONTAP 8.0.1 operating in Cluster-Mode support a new aggregate type whose size threshold is much bigger than the 16-terabyte (TB) aggregate size threshold that was supported in previous releases of Data ONTAP. This aggregate type, called 64-bit aggregate, offers a much larger size threshold while providing all the advantages and capabilities of aggregates, including flexibility and storage efficiency. 64-bit aggregates also support larger FlexVol volumes contained inside them. The exact size thresholds for 64-bit aggregates and for FlexVol volumes contained inside them are given in section 3, “Maximum Aggregate and Volume Sizes.” They range from 40TB to 100TB, depending on the model of the storage system.

The aggregates that were supported until now in Data ONTAP 7G, which have a maximum size threshold of 16TB, will be called 32-bit aggregates from Data ONTAP 8.0 on, and they will continue to be supported with the same 16TB threshold. When a storage system is upgraded to a release of Data ONTAP 8.0 onward, all the existing aggregates show up as 32-bit aggregates.

Figure 1) Existing aggregates on the storage system show up as 32-bit aggregates after upgrading to Data ONTAP 8.0 operating in 7-Mode/8.0.1 operating in Cluster-Mode.



The default aggregate type that is created in Data ONTAP 8.0.x is a 32-bit aggregate. Existing and newer 32-bit aggregates created in Data ONTAP 8.0.x onward are still bound by the 16TB size threshold, so you cannot add more disks to a 32-bit aggregate once the threshold is reached. Therefore a 64-bit aggregate gives you the ability and flexibility you need to plan for future data growth.

Once created, 64-bit aggregates behave and can be used just like existing 32-bit aggregates. All the processes—such as creating and managing FlexVol volumes inside the aggregate and accessing the volumes—are the same, with the same commands. All operations that can be performed on FlexVol volumes in 32-bit aggregates are also supported and work with FlexVol volumes in 64-bit aggregates. A 64-bit aggregate of any size can be created as long as it is less than the maximum size threshold stated in this document.

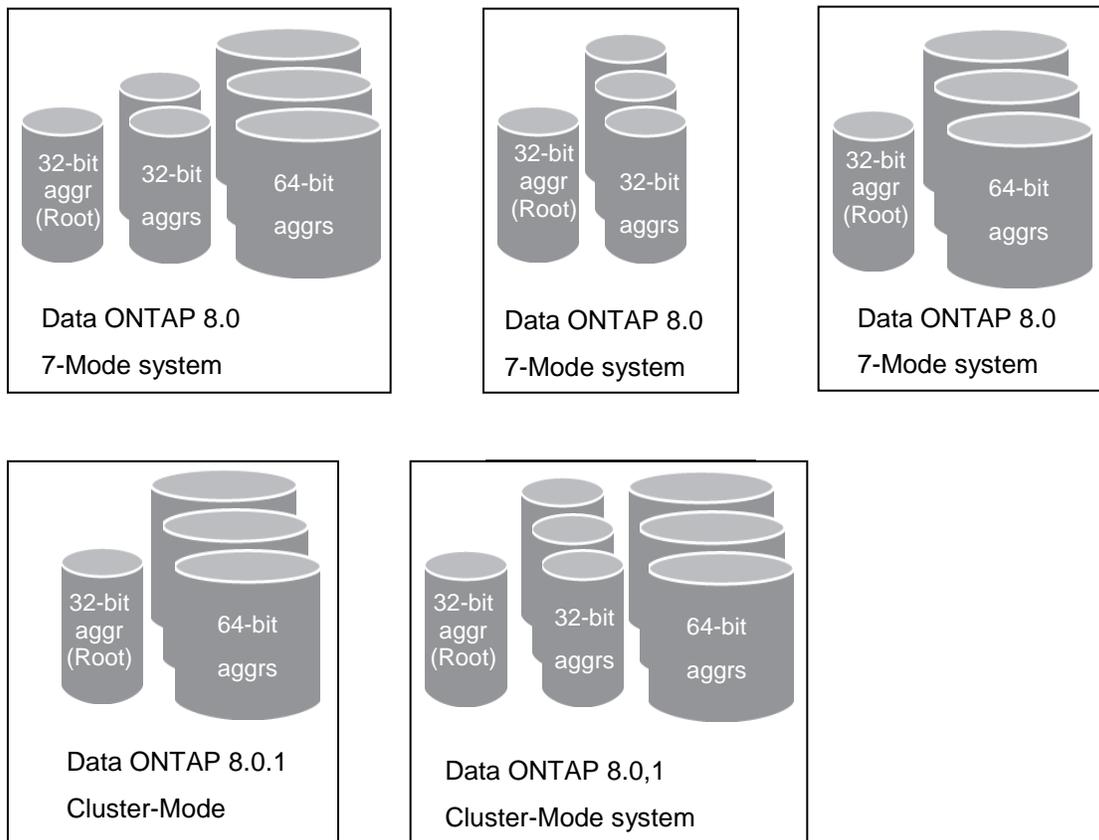
The newer 64-bit aggregates can coexist with new and existing 32-bit aggregates on the storage system. Therefore, you can create a 64-bit aggregate on a storage system that has existing 32-bit aggregates. This coexistence has no impact on the storage system in any way. Once the 64-bit aggregate is created, it is just another aggregate on the storage system, with a higher size threshold, and can be used as a regular aggregate. You can also create 32-bit aggregates, if necessary.

Data ONTAP 8.0.x does not support expanding a 32-bit aggregate to a 64-bit aggregate, or vice versa. Therefore they must exist as two different aggregate types. Data ONTAP 8.1 has the ability to do an in-place expansion of a 32-bit aggregate to 64-bit.

In Data ONTAP 8.0.x, the aggregate containing the root volume is always a 32-bit aggregate. This is true for systems upgraded to and for systems newly installed with Data ONTAP 8.0.x. In Data ONTAP 8.0.1 operating in 7-Mode and in Cluster-Mode, the root can reside in a 64-bit aggregate though the default is always a 32-bit aggregate.

Figure 2 shows the possible aggregate configurations of a storage system running Data ONTAP 8.0 operating in 7-Mode.

Figure 2) Possible configurations for a NetApp storage system running Data ONTAP 8.0.x with 32-bit and 64-bit aggregates with the default 32-bit root aggregate



2 THE NEED FOR 64-BIT AGGREGATES

2.1 DELIVERING PERFORMANCE WITH LARGE-SIZED 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 they must have bigger size thresholds.

2.2 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.

2.3 NEED FOR LARGER FLEXVOL VOLUME SIZES

Certain applications may need the FlexVol volumes in which they store their data to be bigger than the existing 16TB threshold. This means that the underlying aggregate must also be larger.

3 MAXIMUM AGGREGATE AND VOLUME SIZES

The maximum size of a 64-bit aggregate depends on the storage system model. Table 1 lists the maximum aggregate and FlexVol volume sizes for different storage systems in Data ONTAP 8.0.x.

Table 1) Maximum aggregate and FlexVol volume sizes.

FAS/V-Series Platform	Maximum Aggregate Size (TB)	Maximum FlexVol Volume Size (TB)
6080, 6070	100	100
6040, 6030, 3170	70	70
3160, 3070, 3140, 3040	50	50
2040	30	30

The aggregate sizes listed in Table 1 are calculated by using the sum of the space obtained from the used right-sized capacity of all data disks in the aggregate, just as in Data ONTAP 7.3. For details about aggregate space calculation, see the “Data ONTAP Storage Management Guide” on the NetApp Support site: <http://now.netapp.com>.

4 MAXIMUM NUMBER OF DISKS IN A MAXIMUM-SIZE AGGREGATE

Table 2 lists, by size threshold and platform, the maximum number of data drives that can fit in a 64-bit aggregate.

Table 2) Maximum number of data drives in 64-bit aggregates.

FAS/V-Series Platform	Maximum Aggregate Size (TB)	Maximum Number of Data Drives			
		2TB SATA	1TB SATA	750GB SATA	500GB SATA
6080, 6070	100	61	123	164	247
6040, 6030, 3170	70	43	86	115	173
3160, 3070, 3140, 3040	50	30	61	82	123
2040	30	18	37	49	74

FAS/V-Series Platform	Maximum Aggregate Size (TB)	Maximum Number of Data Drives		
		600GB FC or SAS	450GB FC or SAS	300GB FC or SAS
6080, 6070	100	187	250	385
6040, 6030, 3170	70	131	175	269
3160, 3070, 3140, 3040	50	93	125	192
2040	30	56	75	115

5 CONFIGURATION AND OPERATION

5.1 RELEASE

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

5.2 LICENSE

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

5.3 PLATFORMS

All platforms that can run the supported releases can create and manage 64-bit aggregates. For a list of platforms, refer to the NetApp Support site. This is also true for all the V-Series platforms.

5.4 CREATING A 64-BIT AGGREGATE

The command to create a 64-bit aggregate is the same `aggr create` command that is present in Data ONTAP 7G for aggregate creation. In 7-Mode, the `aggr create` command has a new flag, `-B`, to specify the type of aggregate to create. It takes a value of either 32 or 64 and creates the respective type of aggregate. Therefore the `aggr create` command with `-B 64` creates a 64-bit aggregate. The `aggr create` command without the `-B` flag by default creates a 32-bit aggregate in Data ONTAP 8.0.x.

Here are some commands for creating different aggregate types in 7-Mode:

- Creating a 64-bit aggregate:

```
aggr create aggr_64 -B 64 5
```
- Creating a 32-bit aggregate using `-B` flag:

```
aggr create aggr_32 -B 32 5
```
- Creating an aggregate without using `-B` flag:

```
aggr create default_aggr 5
```

This will create a 32-bit aggregate in Data ONTAP 8.0 operating in 7-Mode.

Here are the equivalent commands in Cluster-Mode to create different aggregate types:

- Creating a 64-bit aggregate:

```
storage aggregate create aggr_64 -diskcount 5 -block-type 64-bit
```
- Creating a 32-bit aggregate using the `-block-type` option:

```
storage aggregate create aggr_32 -diskcount 5 -block-type 32-bit
```
- Creating an aggregate without using the `-block-type` option:

```
storage aggregate create default_aggr -diskcount 5
```

This will create a 32-bit aggregate in Data ONTAP 8.0.1 operating in Cluster-Mode.

5.5 CREATING AND MANAGING FLEXVOL VOLUMES IN A 64-BIT AGGREGATE

The process of creating and managing FlexVol volumes in a 64-bit aggregate is exactly the same as it was in previous releases of Data ONTAP 7G. You can use the `vol create` command to create a FlexVol volume inside a 64-bit aggregate. You can use the `vol` family of commands to manage FlexVol volumes.

5.6 IDENTIFYING A 64-BIT AGGREGATE

Starting with Data ONTAP 8.0, the `aggr status` command displays the aggregate type as 32-bit or 64-bit. Here is an example:

```
netapp1> aggr status
```

Aggr	State	Status	Options
aggr0	online	raid_dp, aggr	root
		32-bit	
testaggr_64	online	raid_dp, aggr	
		64-bit	

The equivalent Cluster-Mode command is `storage aggregate show`. It displays the aggregate type as 32-bit or 64-bit.

6 POSSIBLE AGGREGATE CONFIGURATIONS

Once you install or upgrade a storage system to a release that supports 64-bit aggregates, you have the ability to create 64-bit aggregates along with the ability to create 32-bit aggregates; you can choose which type of aggregate you need. This section gives an overview of the different configurations of a storage system running a release that supports 64-bit aggregates. It also provides the possible outcomes of various operations on the aggregates.

- Create a 64-bit aggregate of size less than 16TB.
 - The aggregate is successfully created because 64-bit aggregates can be of any size less than the maximum size for that platform. The system now has a combination of 32-bit and 64-bit aggregates.
- Add disks to the 64-bit aggregate.
 - Disk addition succeeds as long as the aggregate size is less than the maximum size allowed for that platform. Once the size limit is reached, disk addition fails.
- Add disks to an existing 32-bit aggregate.
 - Disk addition succeeds as long as the aggregate size is less than 16TB.
 - Once the 16TB size limit is reached, disk addition fails because 32-bit aggregates are still bound by the 16TB size limit.
- Create a new 32-bit aggregate.
 - Creating a new aggregate succeeds because 32-bit aggregates are still supported.
- Designate a FlexVol volume located in a 64-bit aggregate as the root volume.
 - Designation fails in Data ONTAP 8.0 operating in 7-Mode because the aggregate containing the root volume should always be a 32-bit aggregate. In Data ONTAP 8.0.1 onward, this limitation does not exist and the root can now reside in a 64-bit aggregate. Hence, this operation will succeed in Data ONTAP 8.0.1 onward in both 7-Mode and Cluster-Mode.

7 ADVANTAGES OF 64-BIT AGGREGATES

7.1 BETTER PERFORMANCE FROM MORE DISKS IN THE AGGREGATE

The larger size of 64-bit aggregates makes it possible to add a lot more disks to an aggregate than is feasible with 32-bit aggregates. Therefore, for scenarios in which the disks are the bottleneck in improving performance, 64-bit aggregates with a higher spindle count can give a performance boost. All the FlexVol volumes that are created inside a 64-bit aggregate span across all the data drives in the aggregate, thus providing more disk spindles for the I/O activity on the FlexVol volumes. However, when considering the performance of a storage system, factors beyond just the disk counts often determine the performance being delivered by the storage system. See section 11, “Performance,” for a discussion of the performance characterization of 64-bit aggregates. Also, if you need a performance check on your storage system, contact your NetApp system engineer or NetApp Global Services.

7.2 MAINTAIN HIGHER STORAGE EFFICIENCY

64-bit aggregates give you the ability to add disks to an aggregate in fully populated RAID groups. This gives you maximum storage efficiency while also offering all the data protection benefits of RAID-DP[®] technology.

7.3 EASE OF MANAGEMENT

64-bit aggregates allow more disk shelves to be put into an aggregate because of their larger size. This gives storage administrators the ability to manage a system with fewer aggregates, thus reducing the overhead of managing multiple aggregates.

7.4 USE ALL EXISTING DATA ONTAP FEATURES

64-bit aggregates have all the same abilities, functionality, and features of aggregates that were supported in Data ONTAP 7G and later. Once created, a 64-bit aggregate is just another aggregate that exists in the storage system and can be used as a regular aggregate. All the advantages and features of Data ONTAP, including space efficiency, thin provisioning, deduplication, FlexVol volumes, and many other features, can be used on 64-bit aggregates with the same ease and simplicity as on 32-bit aggregates. See section 10, "Interaction with Other Data ONTAP 8.0.x Features," for details on the interaction of different Data ONTAP features with 64-bit aggregates. You use the same commands for working with 32-bit and 64-bit aggregates. So, you can use your existing scripts and management tools to manage and use 64-bit aggregates and FlexVol volumes contained in 64-bit aggregates. This makes the process of making the transition to 64-bit aggregates seamless and easy.

7.5 BETTER SCALABILITY FOR FUTURE GROWTH

64-bit aggregates offer you bigger aggregate and volume sizes, which means more flexibility in building a storage solution that has greater scope for scalability and management ease in the future as your storage needs grow.

7.6 POWER AND SPACE SAVINGS

64-bit aggregates give you the ability to use large-sized SATA drives in your aggregates while providing the number of disk spindles necessary to maintain performance. You can use fewer large-sized drives, thus reducing the amount of power used by the drives. Fewer drives require less space to house them and also can reduce cooling costs and power consumed by the data center.

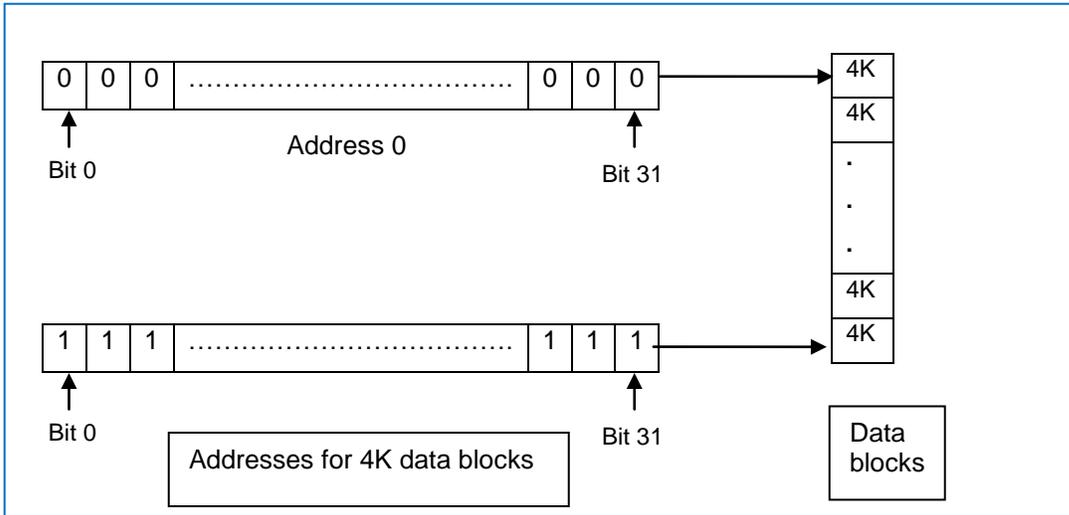
7.7 CONSOLIDATE DATA

With 64-bit aggregates, you can create large-sized FlexVol volumes inside the aggregates, which means that you can consolidate your data into a single FlexVol volume if necessary. This consolidation helps you to simplify your data and volume management.

8 INTERNAL OPERATION

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 3 shows how the 32-bit address space constrained the maximum size to 16TB.

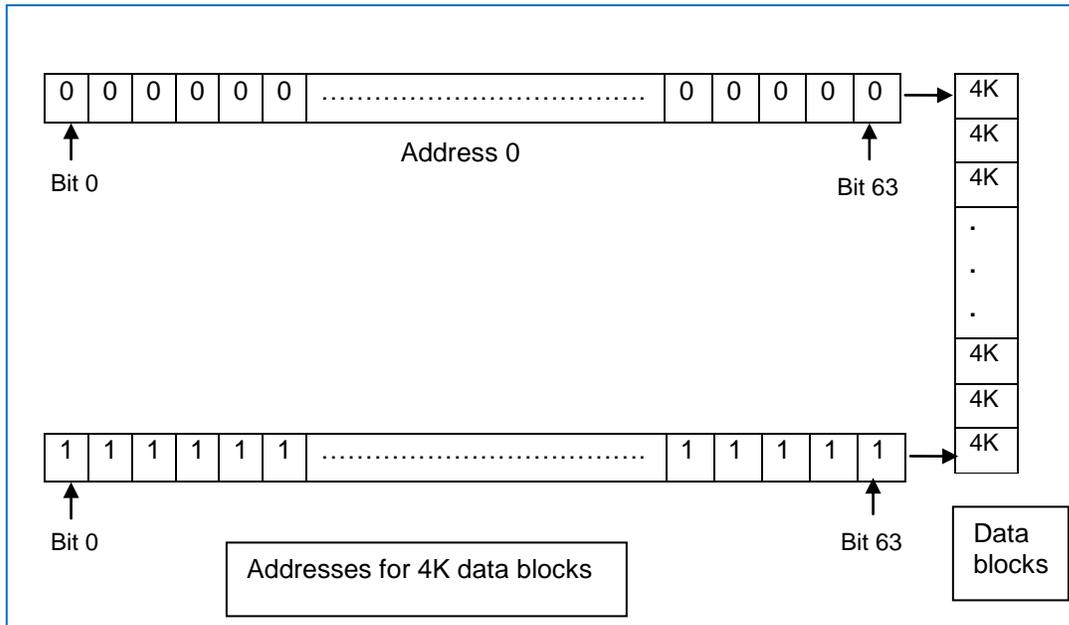
Figure 3) 32-bit addressing format thresholds constrain the aggregate and FlexVol volume size to 16TB.



Because the address pointer is 32 bits in size, the maximum address that can be stored is $(2^{32} - 1)$. Since each block that is referenced with this address is 4KB in size, the maximum size of the aggregate that can be referenced is $(2^{32} - 1) * 4KB = 16TB$.

Data ONTAP 8.0.x has increased the size of the address pointers in 64-bit aggregates to 64 bits, which makes possible a much bigger address space. This bigger address space makes it

Figure 4) 64-bit addressing makes it feasible to increase the aggregate size.



Using a 64-bit addressing scheme, you can have an aggregate that is more than 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 8.0 are platform dependent. These thresholds are listed in section 3, “Maximum Aggregate and Volume Sizes.”

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.

9 DIFFERENCES BETWEEN 64-BIT AND 32-BIT AGGREGATES

9.1 ON-DISK ADDRESSING FORMAT

The way that addresses are stored on disk differs between 32-bit and 64-bit aggregates. The addresses of data blocks for 64-bit aggregates are stored in 64 bits. The addresses for 32-bit aggregates are stored in 32 bits. The bigger addressing format for 64-bit aggregates makes it possible to increase the aggregate size. The addressing format for 64-bit aggregates is discussed in detail in section 8, “Internal Operation.”

9.2 MEMORY USAGE

64-bit aggregates have a bigger address format than 32-bit aggregates. Therefore, depending on the data set and the access pattern of the workload, the memory used for metadata when using 64-bit aggregates might be more than when using 32-bit aggregates. The exact amount of extra memory used in accessing data from 64-bit aggregates depends on the workload.

9.3 PERFORMANCE PROFILE

Given the higher memory usage, 64-bit aggregates may exhibit a different performance profile than 32-bit aggregates, depending on the workload type and the hardware platform. For details about the performance profile of 64-bit aggregates, see section 11, “Performance.”

10 INTERACTION WITH OTHER DATA ONTAP 8.0.X FEATURES

10.1 SNAPMIRROR

This section is an overview of the behavior of volume SnapMirror® and qtrees SnapMirror features with 64-bit aggregates.

VOLUME SNAPMIRROR IN 7-MODE

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 either both be FlexVol volumes in 32-bit aggregates or both be 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 3 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.

Table 3) Volume SnapMirror and SnapMirror sync configurations in Data ONTAP 8.0.x operating in 7-Mode.

Destination 	Volume in 32-bit aggregate	Volume in 64-bit aggregate
Source 		
Volume in 32-bit aggregate	Yes	No
Volume in 64-bit aggregate	No	Yes

REPLICATION IN CLUSTER-MODE

The async replication engine in Cluster-Mode is responsible for all volume replication (data protection and load-sharing mirrors) and volume operations (vol move, vol copy). The Cluster-Mode replication engine is logical and works at a file level. Hence, this works across 32-bit and 64-bit aggregates unlike the engine in 7-Mode.

Table 4) Volume SnapMirror configurations in Data ONTAP 8.0.x operating in Cluster-Mode.

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

QTREE SNAPMIRROR AND SNAPVAULT (7-MODE ONLY)

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 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.

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

10.2 NDMP AND DUMP

Network Data Management Protocol (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.

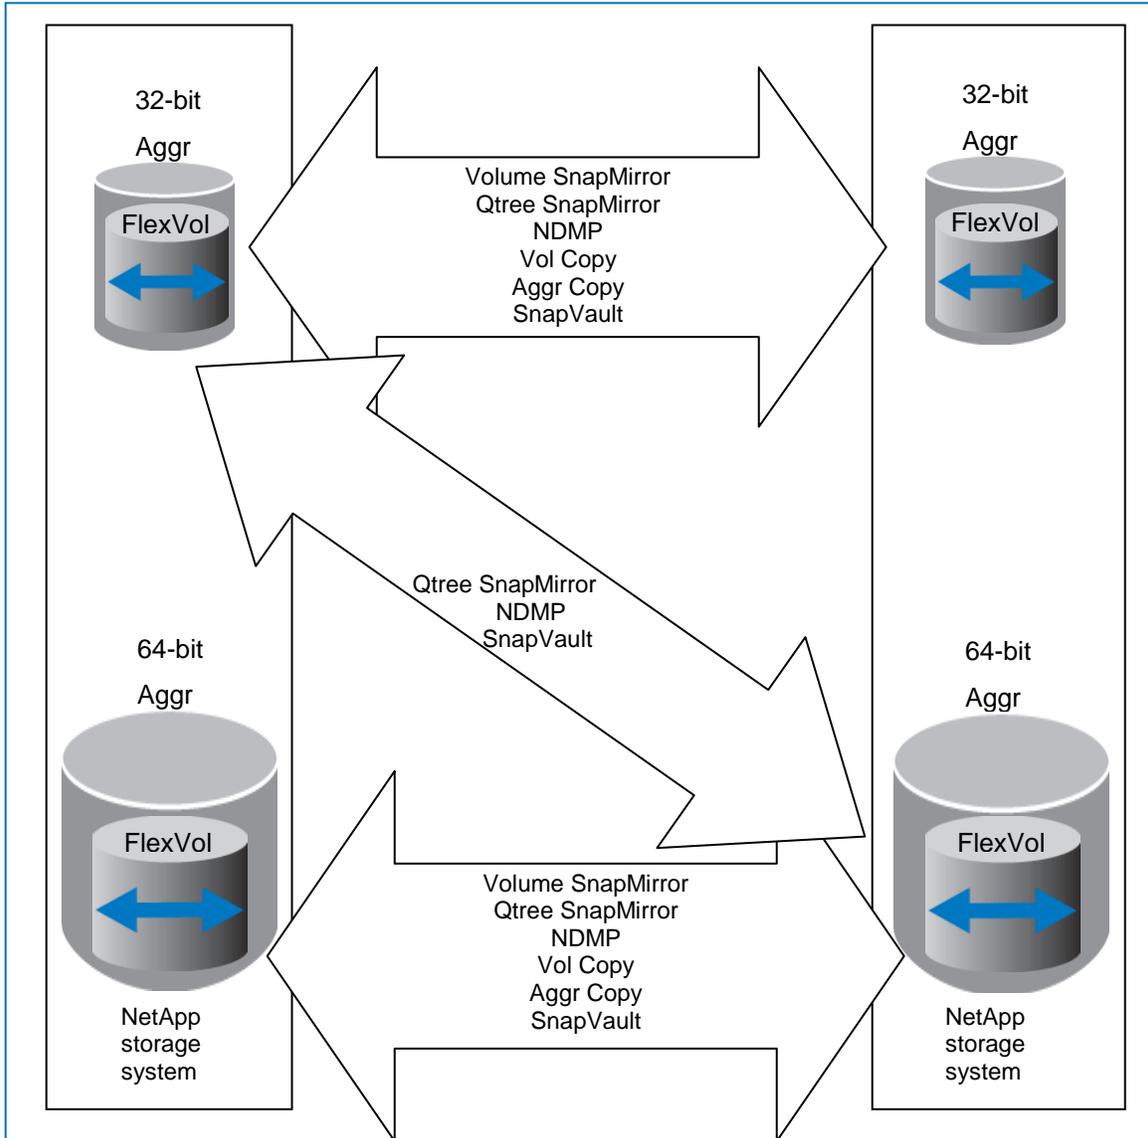
10.3 AGGR COPY AND VOL COPY (7-MODE ONLY)

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 lifts this restriction, allowing you to use vol copy or aggr copy between any-format aggregates.

Figure 5 is an overview of the possibilities for replicating or copying data across 32-bit and 64-bit aggregates using the utilities just discussed.

Figure 5) Data replication utilities and their interaction with 32-bit and 64-bit aggregates (7-Mode only).



10.4 DEDUPLICATION IN DATA ONTAP 8.0.X

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 less than or equal to the sizes shown in the table below.

Table 6) Maximum FlexVol volume size for deduplication-enabled volumes.

FAS/V-Series Platform	Deduplication Maximum FlexVol Volume Size (TB)
2040	3
3040, 3140	4
3070, 3160, 3170, 6030, 6040, 6070, 6080	16

These maximum FlexVol volume sizes for deduplication-enabled FlexVol volumes are less than the maximum supported regular FlexVol volume sizes. If you try to enable deduplication on a volume whose size is bigger than that listed in Table 6, it fails.

When configuring vol autosize on FlexVol volumes that have deduplication enabled, the maximum size threshold for the autosize operation must be less than the maximum allowed volume size for that platform, as listed in Table 6. If the size is more than the maximum allowed size, autosize is not enabled and an error is generated.

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

10.5 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, there are certain restrictions that 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, in Data ONTAP 8.0.x 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.

10.6 PERFORMANCE ACCELERATION MODULE

A Performance Acceleration Module (PAM) can optimize the performance of your random read intensive workloads such as file services and messaging. PAM 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. PAM caches data based on data access regardless of the aggregate type.

The data cached in PAM while the system is in operation depends on the workload, and it can be a combination of data from volumes contained in different aggregates. There is no way to let PAM cache data only from a particular aggregate type. As noted in other sections of this document, 64-bit aggregates have a bigger address space and also take more memory for their metadata than 32-bit aggregates. This might reduce the total amount of effective data that can be cached in PAM when used with 64-bit aggregates present in the system.

10.7 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 Technical Report 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.

10.8 SYNCMIRROR

SyncMirror[®] software enables your data to be available and up to date at all times. By maintaining two copies of data online, SyncMirror protects your data against all types of hardware outages, including triple disk failure.

SyncMirror can be used with 64-bit aggregates. Therefore you can use SyncMirror to create two online copies of the data that is stored in FlexVol volumes contained in 64-bit aggregates.

10.9 CONTROLLER FAILOVER

Controller failover is not affected by the presence of 64-bit aggregates. In case of system failure in an HA pair, controller failover occurs within the HA pair.

Since 64-bit aggregates for Cluster-Mode were introduced in Data ONTAP 8.0.1 and nodes in HA pairs are allowed to run different releases, it is possible to fail over from 8.0 to 8.0.1 and vice versa. If 64-bit aggregates on an 8.0.1 Cluster-Mode node are taken over by an 8.0 node, these 64-bit aggregates will then be running in an unsupported configuration. As such for Cluster-Mode, NetApp recommends

upgrading both nodes in an HA pair in tandem to 8.0.1 before creating any 64-bit aggregates. In addition, 64-bit aggregate creation is disallowed for Cluster-Mode if the partner node is running Data ONTAP 8.0.

10.10 FLEXSHARE

FlexShare[®] software is a powerful quality-of-service tool for Data ONTAP storage systems. It enables you to assign individual priorities to multiple application workloads consolidated on a single system. For details about FlexShare, see NetApp Technical Report TR-3459, “FlexShare Design and Implementation Guide.”

You can use FlexShare to control the priority of volumes that are in 64-bit aggregates. You need to consider FlexVol volumes in 64-bit aggregates and FlexVol volumes in 32-bit aggregates together as the total volume set on the storage system when assigning priorities to the volumes using FlexShare. Once you have the set of all volumes, assign priorities based on your preferences and requirements for the volumes.

11 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.

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 if you can add enough drives such that the workload or storage system is no longer disk spindle constrained in delivering performance.

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 PAM 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, or to figure out whether a 64-bit aggregate is the right choice for your workload, contact your NetApp system engineer.

12 DATA MIGRATION

Data ONTAP 8.0.x does not support converting a 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 need to 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.
- **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 using NDMP. 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.

Data ONTAP 8.1 supports in-place non-disruptive expansion of 32-bit aggregates to 64-bit. Please refer to TR-3978 for more details.

13 USE CASES

64-bit aggregates are in general a good fit for all of your applications and workloads. They offer greater efficiency and also flexibility for planning your future data growth. In general, NetApp recommends that any aggregates created in Data ONTAP 8 should be 64-bit aggregates. For an analysis of your specific environment and for the exact configuration needed to satisfy your workload and requirements, contact your NetApp system engineer.

If you plan to roll out 64-bit aggregates in a phased manner in your data center, here are some use cases for which NetApp recommends that you start with 64-bit aggregates so that you can use the advantages of the larger size thresholds and other efficiency factors discussed in this report.

- **Using large-sized SATA drives.** When using large-capacity SATA drives, NetApp recommends using 64-bit aggregates, because they give you the ability to add enough disks to deliver the required performance and also let you maintain high storage efficiency.
- **Need for large FlexVol volumes to hold huge amounts data.** For applications that store and access huge amounts of data, the larger size thresholds of FlexVol volumes in 64-bit aggregates give you the ability to store all the data in a single FlexVol volume, making it easier to manage the data.
- **Environments that have a high rate of data growth.** For environments in which the rate of data growth is huge, the larger size of 64-bit aggregates offers the ability to manage the data growth while maintaining a small number of aggregates. Also, as discussed earlier in this document, the infrastructure for 64-bit aggregates in Data ONTAP enables the increase of the aggregate size thresholds to even larger sizes in the future without the need for any on-disk or addressing format changes.

14 RECOMMENDATIONS

Here are some NetApp recommendations for configuring and using 64-bit aggregates:

- The recommended RAID group sizes for 64-bit aggregates are the same as the recommendations for 32-bit aggregates. For details, see NetApp Technical Report TR-3437, "Storage Best Practices and Resiliency Guide."
- Be sure to size your FlexVol volumes appropriately to use all the space in the aggregate, since you now have a bigger aggregate size. If you create too many small FlexVol volumes and reach the maximum number of volumes threshold, you will not be able to create any more FlexVol volumes on the storage system, even if the aggregate has free space.

- The recommended size of the root volume has increased in Data ONTAP 8.0. For details, see the “Data ONTAP Storage Management Guide” on the NetApp Support site.
- The maximum possible size for individual files and LUNs has not changed in Data ONTAP 8.0. For the exact size thresholds, refer to the Data ONTAP documentation on the NetApp Support site. Make sure you consider these sizes when you design your storage solution.
- 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 PAM may help improve the performance.
- When migrating data from a FlexVol volume in a 32-bit aggregate to a FlexVol volume in a 64-bit aggregate using qtree SnapMirror or NDMP, only the user-level data is migrated to the volume in the 64-bit aggregate. No metadata related to the volume in the 32-bit aggregate is copied. Therefore you must manually set up any operations, such as SnapMirror relations, LUN mappings, and deduplication schedules, on the new FlexVol volume.
- When using qtree SnapMirror to migrate data from FlexVol volumes in 32-bit aggregates to FlexVol volumes in 64-bit aggregates, you need one qtree SnapMirror relationship per qtree. In this case, be sure to properly schedule your qtree SnapMirror start operations so that you don't have too many transfers running at once.
- If you plan to enable deduplication on a FlexVol volume in the future, make sure that the volume size is within the deduplication volume size thresholds stated in this document.

15 CONCLUSION

Data ONTAP 8.0 operating in 7-Mode and Data ONTAP 8.0.1 operating in Cluster-Mode, with larger 64-bit aggregates and FlexVol volumes, give you the flexibility and efficiency you need to plan for your data growth. 64-bit aggregates seamlessly integrate into your existing processes, and they also help save power, space, and cooling costs for your data center.

16 ACKNOWLEDGEMENTS

The author would like to thank all the people who provided the material for this report and the reviewers for their valuable feedback.

17 REFERENCES

17.1 NETAPP TECHNICAL REPORTS AND WHITE PAPERS

- TR-3437: Storage Best Practices and Resiliency Guide
<http://www.netapp.com/us/library/technical-reports/tr-3437.html>
- TR-3459: FlexShare Design and Implementation Guide
<http://www.netapp.com/us/library/technical-reports/tr-3459.html>

- TR-3505: NetApp Deduplication for FAS Deployment and Implementation Guide
<http://www.netapp.com/us/library/technical-reports/tr-3505.html>
- TR-3669: FlexCache Caching Architecture
<http://www.netapp.com/us/library/technical-reports/tr-3669.html>

17.2 DATA ONTAP PRODUCT DOCUMENTATION

- Data ONTAP System Administration Guide
http://now.netapp.com/NOW/knowledge/docs/ontap/ontap_index.shtml
- MultiStore Management Guide
http://now.netapp.com/NOW/knowledge/docs/ontap/ontap_index.shtml

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®