



Technical Report

# In-Place Expansion of 32-Bit Aggregates to 64-Bit Overview and Best Practices

Shree Reddy, NetApp  
September 2011 | TR-3978

## **IN-PLACE EXPANSION OF 32-BIT AGGREGATES TO 64-BIT**

NetApp® Data ONTAP® 8.1 operating in 7-Mode and Cluster-Mode, support in-place and non-disruptive expansion of existing 32-bit aggregates to the 64-bit format. 64-bit aggregates offer a much larger size threshold while providing all the advantages and capabilities of aggregates, including flexibility and storage efficiency. This technical report describes the 32-bit to 64-bit expansion process and how to manage 64-bit aggregates in detail.

## TABLE OF CONTENTS

<b>1</b>	<b>BACKGROUND</b>	<b>4</b>
<b>2</b>	<b>OVERVIEW</b>	<b>4</b>
2.1	DEFAULT AGGREGATE FORMAT IN DATA ONTAP 8.1	4
2.2	MAXIMUM AGGREGATE AND VOLUME SIZES IN DATA ONTAP 8.1	6
<b>3</b>	<b>DETAILS ON THE 64-BIT EXPANSION PROCESS</b>	<b>6</b>
3.1	BACKGROUND	6
3.2	OVERVIEW OF THE EXPANSION PROCESS	6
3.3	HOW TO TRIGGER THE 64-BIT EXPANSION PROCESS	7
3.4	REQUIREMENTS	8
3.5	CAVEATS	8
3.6	SPACE CONSUMPTION	8
3.7	EXAMPLE OF THE EXPANSION PROCESS (CLI)	9
3.8	COMMAND-SET SUMMARY	11
3.9	EXAMPLE OF THE EXPANSION PROCESS (SYSTEM MANAGER 2.0)	12
3.10	FREQUENTLY ASKED QUESTIONS	13
<b>4</b>	<b>INTERACTION WITH DATA ONTAP FEATURES</b>	<b>14</b>
4.1	COMPRESSION	14
4.2	FLEXCLONE	14
4.3	DATA REPLICATION	14
4.4	SUMMARY OF FEATURE COMPATIBILITY WITH THE 64-BIT EXPANSION PROCESS	19
<b>5</b>	<b>PERFORMANCE</b>	<b>21</b>
5.1	READ/WRITE THROUGHPUT AND LATENCY IMPACT DURING EXPANSION	21
5.2	SPC-1 BENCHMARK	21
5.3	FILE-SHARING WORKLOADS	22
5.4	COMPLETION TIME	24
<b>6</b>	<b>KEY BEST PRACTICES</b>	<b>24</b>
<b>7</b>	<b>CONCLUSION</b>	<b>24</b>
<b>8</b>	<b>REFERENCES</b>	<b>25</b>
8.1	NETAPP TECHNICAL REPORTS	25
8.2	DATA ONTAP PRODUCT DOCUMENTATION	25

## LIST OF TABLES

Table 1) Maximum aggregate and FlexVol volume sizes.....	6
Table 2) Block pointer size in 32-bit and 64-bit aggregates. ....	6
Table 3) Command-set summary .....	11
Table 4) Overview of various data replication features in Data ONTAP 8.1 7-Mode and Cluster-Mode. ...	14
Table 5) SnapMirror supportability matrix for 7-Mode. ....	15
Table 6) SnapMirror supportability matrix for Cluster-Mode.....	16
Table 7) Interoperability with Data ONTAP features .....	19
Table 8) Sample completion times for certain configurations.....	24

## LIST OF FIGURES

Figure 1) Existing aggregates on the storage system show up as 32-bit aggregates after upgrading to Data ONTAP 8.1. ....	5
Figure 2) Default format for new aggregates will be 64-bit.....	5
Figure 3) 32-bit to 64-bit expansion .....	7
Figure 4) 64-bit to 32-bit SnapMirror.....	17
Figure 5) 32-bit to 64-bit SnapMirror.....	17
Figure 6) Data replication utilities and their interaction with 32-bit and 64-bit aggregates.....	20
Figure 7) Throughput sustained during 64-bit expansion. ....	22
Figure 8) Latency measurement during 64-bit expansion. ....	22
Figure 9) Throughput sustained during 64-bit expansion. ....	23
Figure 10) Latency measurement during 64-bit expansion. ....	23

## 1 BACKGROUND

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<sup>®</sup> volumes. These FlexVol volumes share the physical resources of the underlying container.

Aggregates can be in 32-bit or 64-bit block format. A FlexVol volume inherits the parent aggregate's format. 32-bit aggregates are limited in size to 16TB. The 16TB limit is an artifact of the maximum addressable space with a 32-bit block pointer in a 32-bit aggregate.

64-bit aggregates contain larger block pointers (64-bit) and therefore can grow beyond 16TB. Hence, they have larger size limits than 32-bit aggregates.

In Data ONTAP 8.0, users could create 32-bit or 64-bit aggregates. Both aggregate types can coexist in a system. However, Data ONTAP 8.0 was still subject to the following limitations:

- An administrator wishing to expand any of the existing 32-bit aggregates to 64-bit is required to copy out the data.
- Existing 32-bit aggregates are unable to grow beyond 16TB seamlessly.
- Physical SnapMirror<sup>®</sup> replication engines are restricted to volumes residing in the same aggregate type.

All of the above limitations have been fixed in Data ONTAP 8.1 which makes the aggregate format and the expansion process seamless to the system administrator. The administrator can now expand existing 32-bit aggregates in place without requiring any data migration.

## 2 OVERVIEW

Data ONTAP 8.1 provides the following new functionality:

- By default, all newly created aggregates (including the root aggregate) in Data ONTAP 8.1 will now be in 64-bit format.
- A 32-bit aggregate and its FlexVol volumes are automatically expanded to 64-bit without requiring a data copy when new disks are added to push the total aggregate size beyond 16TB. There is no disruption to client data access during this process.
- Physical SnapMirror relationships work across volumes residing in any aggregate type.
- Maximum size limits for aggregates have been increased.

One of the key new features is the functionality to expand existing 32-bit aggregates in place, that is, without requiring any data migration or data copy. The expansion process is non-disruptive and occurs while the volumes are online and serving data to the clients. The expansion process is independent of the Data ONTAP upgrade procedure and it works seamlessly with all existing Data ONTAP features. This feature is available in Data ONTAP operating in both 7-Mode and Cluster-Mode.

**Note:** Unless explicitly stated all content is applicable to both 7-Mode and Cluster-Mode.

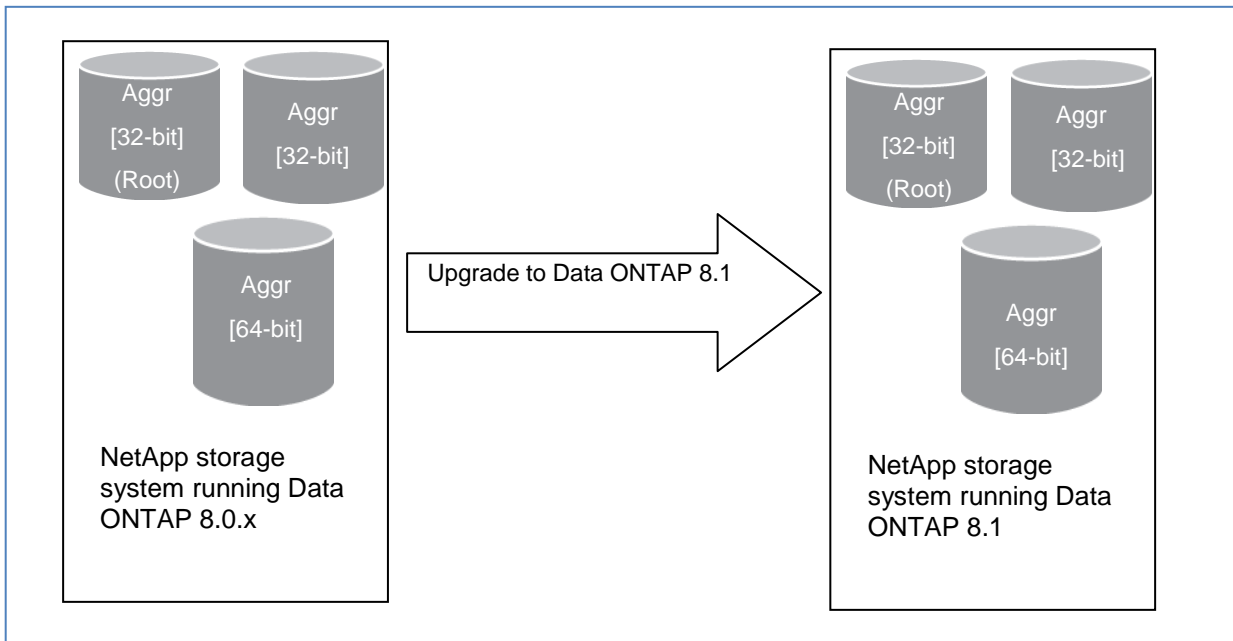
### 2.1 DEFAULT AGGREGATE FORMAT IN DATA ONTAP 8.1

All new aggregates are by default 64-bit aggregates irrespective of the size of the aggregate. This also applies to the root aggregate. An upgrade to Data ONTAP 8.1 does not change the format of any existing aggregates. The administrator will have to add disks to trigger the 64-bit expansion process.

#### SCENARIO 1: UPGRADE TO DATA ONTAP 8.1

On an upgrade to Data ONTAP 8.1, there are no changes. All aggregates will continue to retain the same format. All existing 32-bit aggregates will continue to be 32-bit.


Figure 1) Existing aggregates on the storage system show up as 32-bit aggregates after upgrading to Data ONTAP 8.1.



## SCENARIO 2: NEW INSTALLS

On a factory shipped system, the root aggregate will be 64-bit. Any newly created aggregate will be in the 64-bit format.

Figure 2) Default format for new aggregates will be 64-bit.



## 2.2 MAXIMUM AGGREGATE AND VOLUME SIZES IN DATA ONTAP 8.1

Maximum size limits for aggregates have been increased in Data ONTAP 8.1 and vary based on the platform. Volume sizes stay the same as in Data ONTAP 8.0.

Table 1) Maximum aggregate and FlexVol volume sizes.

FAS/V-Series Platform	Maximum Aggregate Size 8.1 (TB)	Maximum Aggregate Size 8.0 (TB)	Maximum FlexVol Volume Size 8.0 and 8.1 (TB)
6280, 6240, 6080, 6070	162	100	100
6210	162	70	70
6040, 6030, 3270, 3170	105	70	70
3240, 3160	90	50	50
3210, 3140	75	50	50
3070, 3040	50	50	50
2040	50	30	30

**Note:** Maximum file/LUN size is unchanged at 16TB.

Maximum number of Snapshot™ copies in a volume is unchanged at 255.

Maximum number of qtrees is unchanged at 4,995.

## 3 DETAILS ON THE 64-BIT EXPANSION PROCESS

### 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-bit. With the increased block pointer size in 64-bit aggregates, the number of addresses or block pointers that an indirect block can hold has decreased. Consequently, this results in an inode requiring more indirect blocks in a 64-bit volume or aggregate.

Table 2) Block pointer size in 32-bit and 64-bit aggregates.

	32-bit aggregates	64-bit aggregates
Block pointer size	32-bit	64-bit
Number of block pointers in an indirect block in a FlexVol volume	510	255


### 3.2 OVERVIEW OF THE EXPANSION PROCESS

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 rewrites only indirect blocks and does not touch the data blocks.

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

Figure 3) 32-bit to 64-bit expansion



The expansion process occurs in two ways:

1. Inline: As clients are accessing 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.

#### Features of the expansion process

In place  
Online  
Non-disruptive

No data copy or migration required  
Data access is maintained during the expansion  
Existing aggregate configurations including storage efficiency features and data protection relationships are not affected.

## INSTALLATION/LICENSE

No special license is required to enable this feature.

### 3.3 HOW TO TRIGGER THE 64-BIT EXPANSION PROCESS

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 will continue to be 32-bit.

#### How to trigger the expansion process

Add enough storage to a 32-bit aggregate to increase the size beyond 16TB.

For best performance, 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.

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 will fail.

### 64-BIT EXPANSION WORKFLOW

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

1. Establish 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 will be in 64-bit format.
4. Expansion process finishes on all volumes within the aggregate.

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

### 3.4 REQUIREMENTS

- Disks must be added to an existing 32-bit aggregate such that the aggregate's storage capacity exceeds 16TB. This is the only action that is available to customers to trigger the expansion process.
- Enough free space should be available to enable completion of the expansion since 64-bit aggregates consume more metadata. See section 3.6 for more details.
- 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 section 3.6 for more details.

### 3.5 CAVEATS

- If the 32-bit aggregate you wish to expand contains any FlexCache<sup>®</sup> volumes, the cache volumes will have to be destroyed. The cache volumes can be recreated once 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 may experience an initial slowdown in data access due to cache repopulation.
- DataMotion<sup>™</sup> for Volumes for 7-Mode volumes will not work across 32-bit and 64-bit aggregates.
- DataMotion<sup>™</sup> for Vfilers (7-Mode only) will not work across 32-bit and 64-bit aggregates.

**Note:** DataMotion for Volumes in Cluster-Mode does not have the same restriction. Cluster-Mode volumes can be moved between 32-bit and 64-bit aggregates.

#### Limitations

FlexCache volumes need to be destroyed before the expansion process can be triggered.

DataMotion for Volumes in 7-Mode will not work between 32-bit and 64-bit aggregates.

DataMotion for Vfilers (7-Mode only) will not work across 32-bit and 64-bit aggregates.

### 3.6 SPACE CONSUMPTION

64-bit aggregates consume additional metadata with the increase in block pointer size. Hence, the expansion process will need to allocate extra blocks. If the volume is approaching its full capacity, it is possible to run out of space during the expansion process.



If your volume(s) 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 on an upgrade to Data ONTAP 8.1 to gather information on all the space-reserved files in existing volumes so that we can continue to enable space reservations. This 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. As such the expansion process can only be triggered after the completion of the prequalification phase. If the expansion process is triggered in the midst of the pre-qualification phase, it will error out with a more detailed message.

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

### I INCREASED THE VOLUME SIZE BUT THE EXPANSION PROCESS IS STILL OUT OF SPACE.

Even if space was added to a volume before the expansion process, it is still possible to run out of space since client traffic may consume additional space. This scenario is less likely since 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?

If the volume is out of space during the expansion process, the process aborts. Clients may start seeing ENOSPC errors. The appropriate volume autogrow/autodelete configurations will be triggered.

Meanwhile, the administrator can create more space by either:

1. Growing the volume or
2. Adding more disks

The expansion process is restarted automatically when space is available.

### WHAT HAPPENS TO MY QUOTAS?

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

## 3.7 EXAMPLE OF THE EXPANSION PROCESS (CLI)

### STEP 1: ESTABLISH REQUIREMENT TO GROW AN AGGREGATE BEYOND 16TB

In the example below, 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 AN AGGREGATE (SIZE <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 AN AGGREGATE (SIZE >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 will need 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 THE 64-BIT OPTION TO TRIGGER THE EXPANSION

The administrator will now be 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 list of volumes that need to be grown and the amount of space that needs to be added to each of those volumes.

### STEP 5: AGGREGATE IS NOW IN 64-BIT FORMAT

As soon as the expansion process is triggered, the aggregate will be considered to be a 64-bit aggregate. Even if the process is not finished yet, 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
```

### STEP 6: MONITORING THE STATUS OF THE EXPANSION PROCESS

The expansion process can now be tracked if needed. The `status` command will also provide 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
  vol1  upgrading running    10         1719      fbn 0,
inode 32664 of 32781, public
  vol2  upgrading running    10         1736      fbn 0,
inode 32661 of 32781, public
```

## 3.8 COMMAND-SET SUMMARY

64-bit expansion is supported in Cluster-Mode too and the behavior is the same as in 7-Mode. The following table lists the entire command set required for 64-bit expansion.

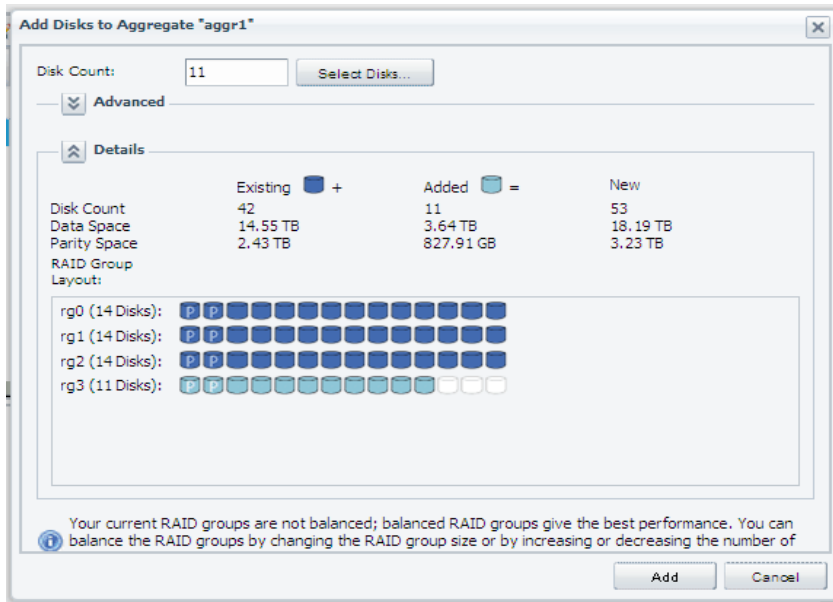
Table 3) Command-set summary

Feature	7-Mode	Cluster-Mode
Trigger the 64-bit expansion	<code>aggr add aggr1 -64bit-upgrade normal</code>	<code>storage aggregate add-disks &lt;aggr-name -64bit-upgrade normal</code>
Monitor the status	<code>aggr 64bit-upgrade status</code>	<code>storage aggregate 64bit-upgrade status &lt;aggr-name&gt;</code>
Check space consumption	<code>aggr add aggr1 -64bit-upgrade check</code>	<code>storage aggregate add-disks &lt;aggr-name -64bit-upgrade check</code>

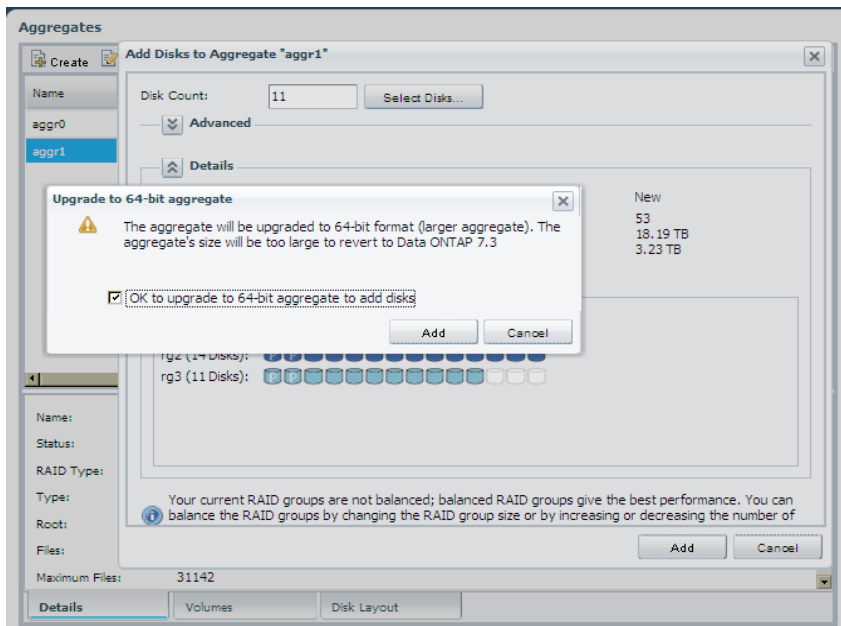
### 3.9 EXAMPLE OF THE EXPANSION PROCESS (SYSTEM MANAGER 2.0)

When total space after disk selection crosses the 16TB limit for a 32-bit aggregate, System Manager will warn the administrator to expand to the 64-bit format. On confirmation, System Manager will launch a wizard to help the administrator expand the aggregate.

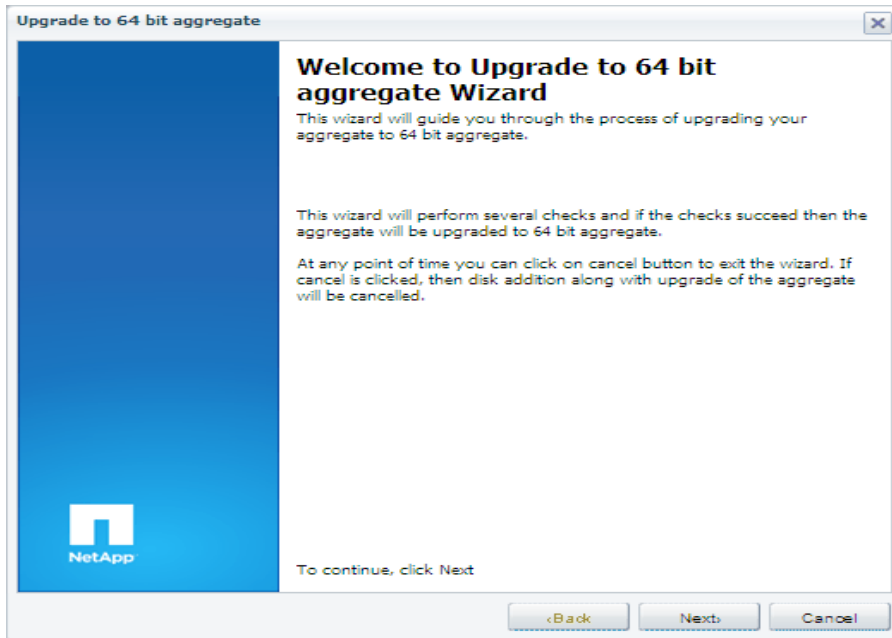
#### STEP 1: ADD DISKS TO GROW AN AGGREGATE (SIZE >16TB)



#### STEP 2: CONFIRM THE 64-BIT EXPANSION



### STEP 3: 64-BIT EXPANSION IS TRIGGERED



### 3.10 FREQUENTLY ASKED QUESTIONS

#### WILL AN UPDATE TO DATA ONTAP 8.1 TRIGGER THE EXPANSION PROCESS?

No. The expansion process can only be triggered with the addition of disks if the size of the aggregate exceeds 16TB.

#### CAN I SHRINK AN AGGREGATE?

No. The size of an aggregate cannot be decreased.

#### CAN I EXPAND MY 32-BIT ROOT AGGREGATE TO 64-BIT?

If there is a strong requirement to expand your root aggregate beyond 16TB, you can add disks and trigger the 64-bit expansion on the root aggregate.

#### CAN I CONVERT MY 64-BIT AGGREGATE TO 32-BIT?

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

#### CAN I EXPAND MY 32-BIT AGGREGATE TO 64-BIT WITHOUT THE ADDITION OF DISKS?

No. The administrator will have to add disks to trigger the 64-bit expansion.

#### ARE BLOCKS IN SNAPSHOT COPIES ALSO CONVERTED TO THE 64-BIT FORMAT?

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.

#### DOES THE EXPANSION PROCESS RESTART FROM THE BEGINNING IF INTERRUPTED?

No. The expansion process maintains checkpoints. If the process is interrupted, it resumes from the latest checkpoint.

## 4 INTERACTION WITH DATA ONTAP FEATURES

### 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

Enable compression after the completion of the expansion process.

### 4.2 FLEXCLONE

If the aggregate that is in the midst of being expanded contains any FlexClone<sup>®</sup> 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.

### 4.3 DATA REPLICATION

Table 4 gives an overview of the various replication engines in 7-mode and Cluster-mode.

Table 4) Overview of various data replication features in Data ONTAP 8.1 7-Mode and Cluster-Mode.

Feature	7-Mode		Cluster-Mode	
	Works across 32-bit and 64-bit aggregates?	Nature of replication	Works across 32-bit and 64-bit aggregates?	Nature of replication
Asynchronous Volume SnapMirror	Yes	Physical	Yes	Physical
Synchronous Volume SnapMirror	Yes	Physical	N/A	Not supported
Aggr Copy	Yes	Physical	N/A	Not supported
Vol Copy	Yes	Physical	Yes	Physical
DataMotion for Volumes	No	Physical	Yes	Physical
Load-Sharing Mirror	N/A	Not supported	Yes	Physical

## 7-MODE

In Data ONTAP 8.0 operating in 7-Mode, all physical-based replication engines like 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 physical-based replication relationships between volumes of different formats.

Please refer to the following supportability matrix for all supported SnapMirror relationships across different formats and Data ONTAP kernels for 7-Mode.

Table 5) SnapMirror supportability matrix for 7-Mode.

		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)
Source Volume	7.3.x (32-bit)	<b>Yes</b>	<b>Yes</b>	No	<b>Yes</b>	<b>Yes</b>
	8.0.x (32-bit)	No	<b>Yes</b>	No	<b>Yes</b>	<b>Yes</b>
	8.0.x (64-bit)	No	No	<b>Yes</b>	<b>Yes</b>	<b>Yes</b>
	8.1 (32-bit)	No	No	No	<b>Yes</b>	<b>Yes</b>
	8.1 (64-bit)	No	No	No	<b>Yes</b>	<b>Yes</b>

## CLUSTER-MODE

Existing 10.0.x and 8.0.x mirror relationships must be deleted and new baseline transfers must occur as part of the upgrade process to Data ONTAP 8.1 operating in Cluster-Mode.

Unlike, Data ONTAP 8.0 operating in 7-Mode, the async replication engine in Data ONTAP 8.0 operating in Cluster-Mode can replicate between a 32-bit and a 64-bit volume.

Starting in Data ONTAP 8.1, you can create physical-based replication relationships between volumes of different formats irrespective of the mode.

Please refer to the following supportability matrix for all the supported SnapMirror relationships across different formats and Data ONTAP kernels for Cluster-Mode.

Table 6) SnapMirror supportability matrix for Cluster-Mode.

		Destination Aggregate				
		10.0.x (32-bit)	8.0.x (32-bit)	8.0.x (64-bit)	8.1.x (32-bit)	8.1.x (64-bit)
Source Volume	10.0.x (32-bit)	<b>Yes</b>	No	No	No	No
	8.0.x (32-bit)	No	<b>Yes</b>	<b>Yes</b>	No	No
	8.0.x (64-bit)	No	<b>Yes</b>	<b>Yes</b>	No	No
	8.1 (32-bit)	No	No	No	<b>Yes</b>	<b>Yes</b>
	8.1 (64-bit)	No	No	No	<b>Yes</b>	<b>Yes</b>

#### CAN AN AGGREGATE BE EXPANDED TO 64-BIT IF IT CONTAINS A SNAPMIRROR SOURCE/DESTINATION?

Yes. An aggregate that contains one or more SnapMirror source/destination volumes when grown beyond 16TB is expanded from 32-bit to 64-bit. All SnapMirror-related activities, 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.

#### DETAILS OF OPERATION

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 since indirect blocks account for a very small percentage of any data-set.


It is now 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 now 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 will result in having a 64-bit destination volume residing in a 32-bit aggregate.

Figure 4) 64-bit to 32-bit SnapMirror



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

Figure 5) 32-bit to 64-bit SnapMirror



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

### **SnapMirror flip resync**

Consider the above 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 will be 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 will fail if the source volume is now bigger than the destination volume.

### **Best Practice for SnapMirror to avoid read performance penalty on the destination volume**

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.

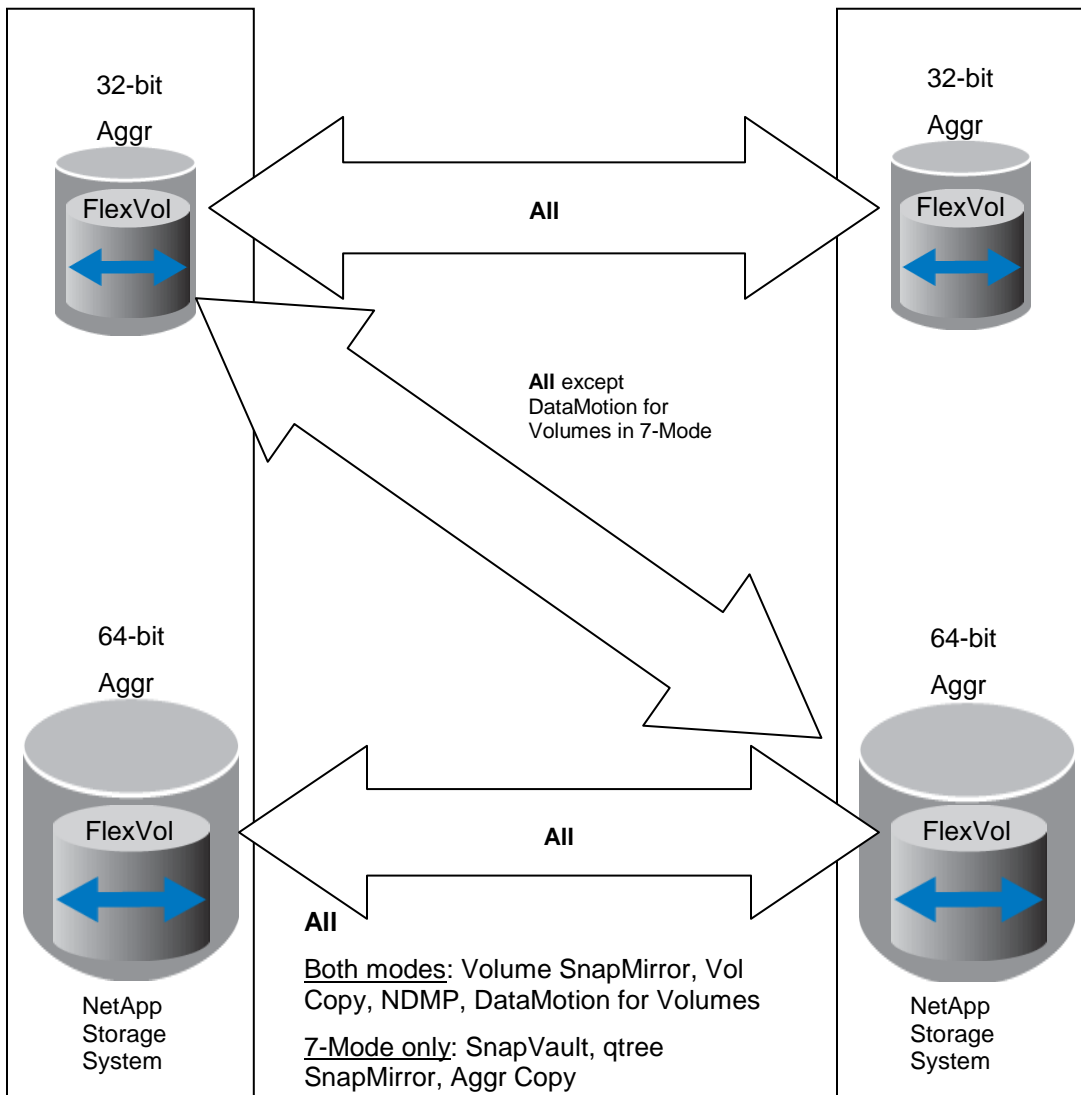
## 4.4 SUMMARY OF FEATURE COMPATIBILITY WITH THE 64-BIT EXPANSION PROCESS

Table 7) Interoperability with Data ONTAP features

Feature	Compatibility
Qtree SnapMirror/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®/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.
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.
DataMotion for Volumes in Cluster-Mode	It can be triggered during the expansion process and also works between 32-bit and 64-bit volumes.
FlexCache	FlexCache volumes need to be destroyed before the expansion process can be triggered. They can be recreated once 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.

Feature	Compatibility
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.

Figure 6) Data replication utilities and their interaction with 32-bit and 64-bit aggregates.



## 5 PERFORMANCE

### 5.1 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. Since the client operations may also trigger the 64-bit expansion, this may increase the latency and reduce throughput. Once 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.

#### 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 impacted 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.

#### DOES THE TIME TAKEN FOR THE EXPANSION PROCESS DEPEND ON THE SIZE OF THE AGGREGATE/VOLUMES?

No. The expansion process updates the format of the indirect metadata in the aggregate and the volume. So, it depends on the amount of indirect metadata.

#### WHAT PERCENTAGE OF THE DATASET COMPRISES INDIRECT BLOCKS?

The amount of indirect metadata can vary based on the number of files and the amount of used data. Typically, it accounts for 2% or less of the dataset.

### 5.2 OLTP BENCHMARK

SPC-1 is a standard benchmark that generates a workload with characteristics of typical business applications such as database and e-mail with random I/O, queries, and updates.

The baseline test was performed with 32-bit aggregates while simulating the SPC-1 benchmark. To measure the impact of the 64-bit expansion process on system latency, the same test was repeated while sustaining the same throughput as the baseline and triggering the expansion process simultaneously.

Figure 7) Throughput sustained during 64-bit expansion.

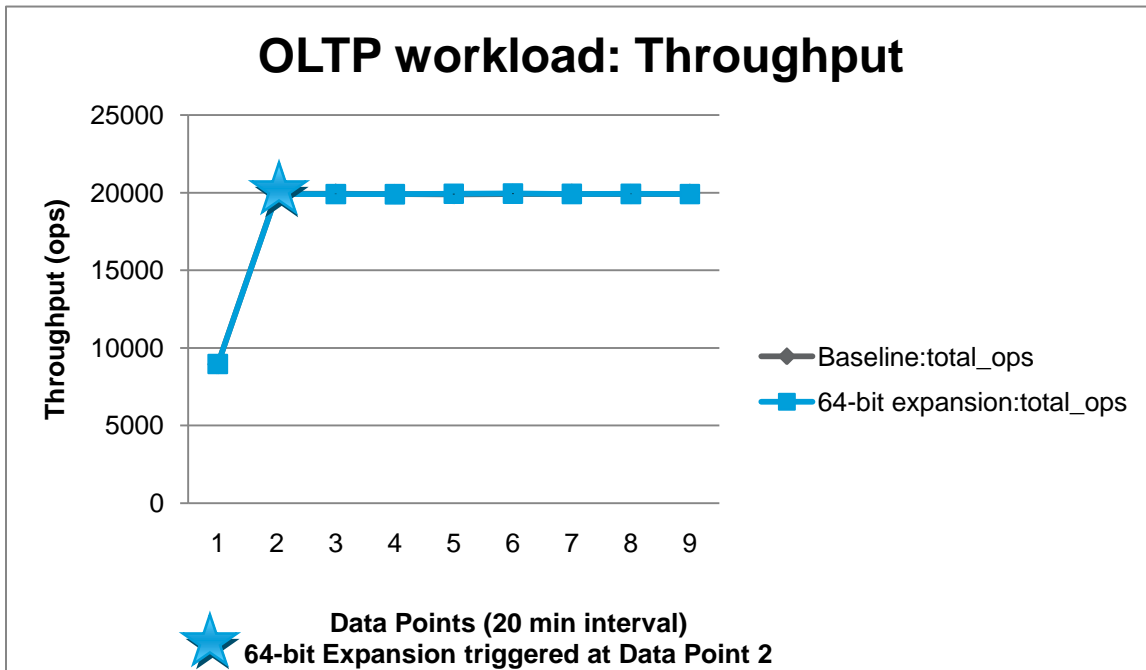
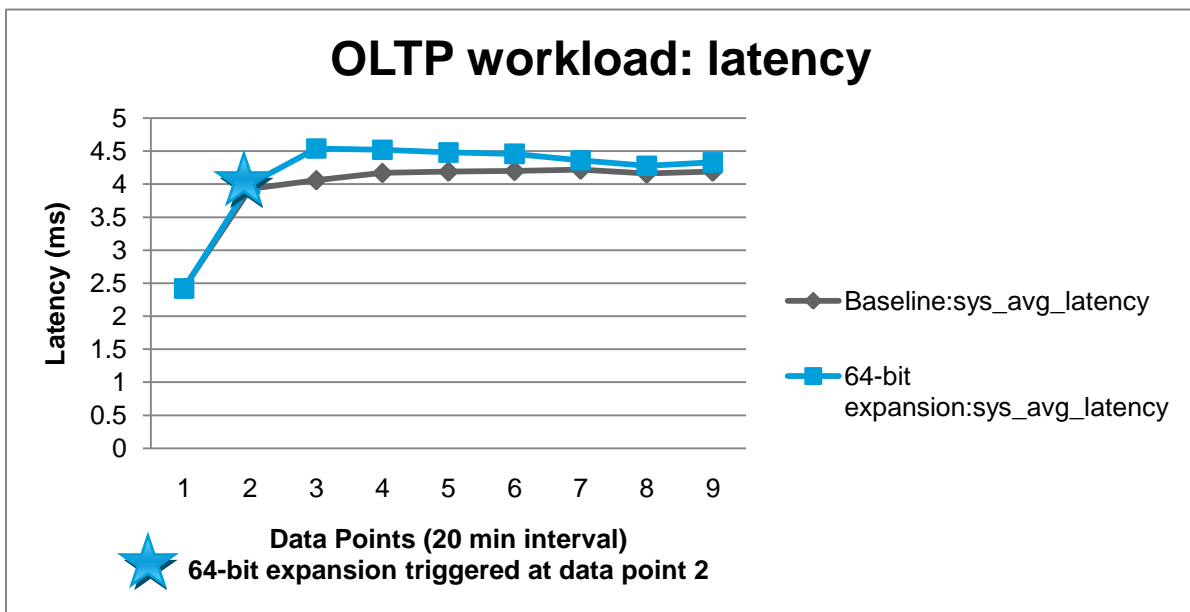


Figure 8) Latency measurement during 64-bit expansion.



### 5.3 FILE-SHARING WORKLOADS

The SFS2008 NFS benchmark allows customers to gauge home directory performance with an industry-standard tool. The following section provides performance data while simulating the SFS2008 NFS benchmark with the 64-bit expansion process. The test measures the latency impact of the 64-bit

expansion process on client traffic while trying to maintain the same throughput as the baseline run with 32-bit aggregates.

The results contain 9 load points and the results of the server at each of the requested load points. Each load point presents a requested number of operations per second.

Figure 9) Throughput sustained during 64-bit expansion.

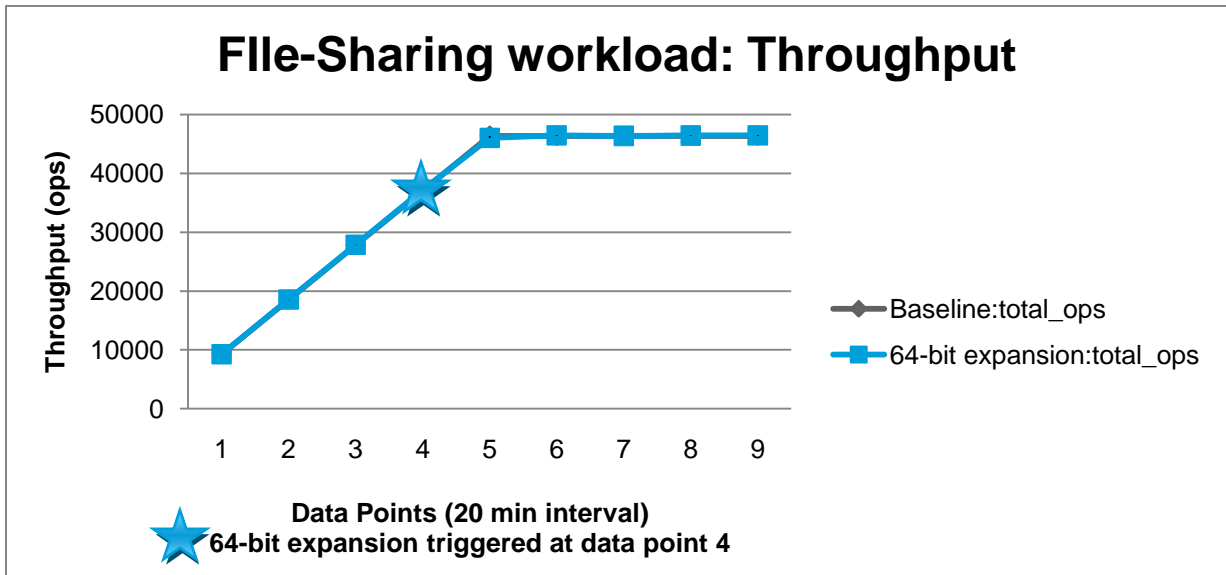
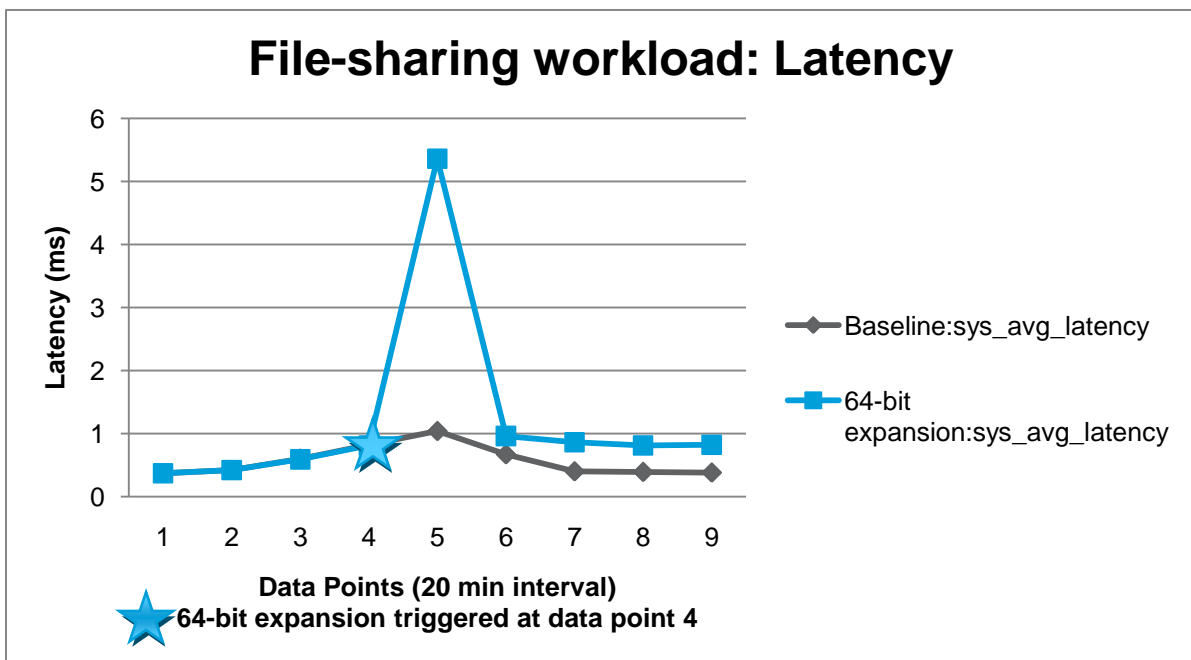


Figure 10) Latency measurement during 64-bit expansion.



## 5.4 COMPLETION TIME

Table 8 can be used as a guideline for how long the expansion process might take to complete on an idle system. Several factors determine the timing, such as system configuration, I/O workload, size of dataset in terms of inodes and blocks, nature of dataset, and so on.

Table 8) Sample completion times for certain configurations.

	Aggregate Size	Volume Size	Number of Volumes	Usage in Each Volume	Completion Time
Config 1	16TB	10G	499	50	~30 hrs
Config 2	16TB	13TB	1	100	~28 hrs
Config 3	13TB	12TB	1	30	~10 hrs
Config 4	13TB	1TB	10	90	~1 hr
Config 5	13TB	100G	100	70	~1 hr
Config 6	13TB	25G	499	66	~1 hr
Config 7	13TB	400G	40	30	~1 hr

## 6 KEY BEST PRACTICES

When triggering 64-bit expansion, keep these best practices in mind.

### Best Practices

Destroy all FlexCache volumes in the aggregate you wish to expand.

Determine that the aggregate will not be a source or destination of DataMotion for Volumes for 7-Mode.

For best performance, 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.

If the customer wishes to enable compression, NetApp recommends waiting until the expansion process completes.

Do not change any autosize settings during the expansion process. Autosize changes will not be effective until the expansion process completes.

Determine that the volumes in the 32-bit aggregate are not close to full.

Determine that the prequalification scanner finished running before triggering the 64-bit expansion.

## 7 CONCLUSION

Data ONTAP 8.1 operating in 7-Mode and Cluster-Mode gives the administrator the flexibility to seamlessly expand existing 32-bit aggregates to 64-bit aggregates without any disruption to client access. These enhancements make the format of the aggregate (32bit or 64bit) largely invisible to the customer.



## 8 REFERENCES

### 8.1 NETAPP TECHNICAL REPORTS

- TR-3505: NetApp Deduplication for FAS Deployment and Implementation Guide  
[www.netapp.com/us/library/technical-reports/tr-3505.html](http://www.netapp.com/us/library/technical-reports/tr-3505.html)
- TR-3786: A Thorough Introduction to 64-Bit Aggregates  
[www.netapp.com/us/library/technical-reports/tr-3786.html](http://www.netapp.com/us/library/technical-reports/tr-3786.html)
- TR-3437: Storage Best Practices and Resiliency Guide  
[www.netapp.com/us/library/technical-reports/tr-3437.html](http://www.netapp.com/us/library/technical-reports/tr-3437.html)

### 8.2 DATA ONTAP PRODUCT DOCUMENTATION

- Data ONTAP System Administration Guide  
[http://now.netapp.com/NOW/knowledge/docs/ontap/ontap\\_index.shtml](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®



© Copyright 2012 NetApp, Inc. All rights reserved. No portions of this document may be reproduced without prior written consent of NetApp, Inc. Specifications are subject to change without notice. NetApp, the NetApp logo, Go further, faster, xxx, and xxx are trademarks or registered trademarks of NetApp, Inc. in the United States and/or other countries. All other brands or products are trademarks or registered trademarks of their respective holders and should be treated as such. TR-3978-0911