



Technical FAQ

NetApp Infinite Volume in Clustered Data ONTAP

Sandra Moulton, NetApp
November 2013

Version 2.2

Abstract

This document answers some common technical questions about the Infinite Volume feature in the clustered Data ONTAP[®] 8.2 operating system. There is a separate section at the end of this FAQ on Infinite Volume in Data ONTAP 8.1.x.

This document does not cover troubleshooting Infinite Volume.

TABLE OF CONTENTS

1	Infinite Volume Overview	9
	What is an Infinite Volume?	9
	What are the components of an Infinite Volume?	9
	What are the major features of Infinite Volume in clustered Data ONTAP?	9
	What are the capacity limits of Infinite Volume?	10
	How does Infinite Volume work?	10
	How many Infinite Volumes can I have in a cluster?	10
	Which NetApp platforms does Infinite Volume support?	11
	Which versions of Data ONTAP support Infinite Volume?	11
	Can I configure Infinite Volume on 7-Mode systems?	11
	Are heterogeneous nodes supported in an Infinite Volume (that is, can FAS3270 and FAS6280 be mixed)?	11
	Are 32-bit and 64-bit aggregates supported with Infinite Volume?	11
	Why is the number of files supported in Infinite Volume limited to 2 billion?	11
	Does Infinite Volume allow deeper directories than a FlexVol volume?	11
	Does Infinite Volume have the same high-file-count limitations that FlexVol has?	11
	How many redirector files are created in the namespace constituent?	11
	Does the admin have to manually increase the number of inodes in the namespace constituent?	11
	How is metadata handled in Infinite Volume?	12
	How do I size an Infinite Volume in Data ONTAP 8.2 for capacity?	12
	How do I size an Infinite Volume in Data ONTAP 8.2 for performance?	15
1.1	Storage Classes and Data Policies	16
	What is a storage class?	16
	What is a data policy?	16
	Will modifications of files already written to the Infinite Volume result in it being moved to a different storage class?	16
	Why would I create a storage class and data policy?	16
	What version of clustered Data ONTAP supports storage classes and data policies?	16
	What tool do I use to create and manage a storage class?	16
	What tool do I use to create and manage a data policy?	16
	Does an Infinite Volume have storage classes by default?	17
	Can I add a storage class to an Infinite Volume that was created without storage classes?	17
	Does OnCommand WFA have storage class workflows available by default?	17
	How can I identify which data constituents are part of each storage class?	17
	Can storage classes ever be deleted or reduced in size?	17
1.2	Disk Configuration Recommendations for Data ONTAP 8.2	17
	Can you use different disk drive types for the namespace and the data constituents within an Infinite Volume?	17
	Before I create an Infinite Volume, is there anything I should do with my disks?	18
	What is the recommended configuration for an Infinite Volume in Data ONTAP 8.2?	18
	What is the recommended configuration for the namespace constituent?	18

Can I specify which aggregate or node to use for the namespace constituent?	18
Does the namespace constituent always need to be 10TB?	19
Can I limit which aggregates or nodes the Infinite Volume can use for the data constituents?	19
1.3 Flexible Volume and Infinite Volume Comparison.....	19
What are some major differences between an Infinite Volume and a FlexVol volume in clustered Data ONTAP?	19
Can an Infinite Volume and a FlexVol volume coexist in a cluster?	20
What is the difference between the clustered Data ONTAP namespace and the data stored in the Infinite Volume namespace constituent?.....	20
2 Infinite Volume Setup	20
At what size will an Infinite Volume be created by default?	20
How do I select which nodes in my cluster to use for Infinite Volume?.....	20
Does the Infinite Volume feature require a license?	20
Is Infinite Volume generally available?	20
What is the recommended way to configure an Infinite Volume?	20
Can I use System Manager to manage Infinite Volumes in Data ONTAP 8.2?	21
How do I create an Infinite Volume in Data ONTAP 8.2?	21
How many SVMs for Infinite Volume can I create?.....	21
How do I create and manage a storage class?.....	21
How do I create and manage a data policy?	22
What additional actions can I perform in admin mode?	22
For what additional actions do I require CSS support?.....	23
What login levels are used for CLI management of Infinite Volume?	23
Are qtrees allowed in Infinite Volume? Would an error message appear if the volume qtree create command is executed?.....	23
2.1 NFS Configuration of Infinite Volume	23
How does the NFS client access an Infinite Volume resource?	23
How do I set up and configure a network interface to access the Infinite Volume?	23
2.2 SMB1.0 Configuration of Infinite Volume (Available Starting in Data ONTAP 8.2).....	23
How does the CIFS client access an Infinite Volume resource?.....	23
How do I set up and configure a network interface to access the Infinite Volume?	24
2.3 Namespace Constituents	24
When is a namespace constituent created?	24
What size are namespace constituents?	24
How many namespace constituents exist in an Infinite Volume?	24
2.4 Namespace Mirror Constituents in Data ONTAP 8.2	24
When are namespace mirror constituents created?	24
What size are namespace mirror constituents?.....	24
How many namespace mirror constituents will exist in an Infinite Volume?	24
Are namespace mirror constituents replicated as part of the SnapMirror replication for Infinite Volume?	25
Are namespace mirror constituents created on the SnapMirror destination for Infinite Volume?.....	25

2.5	Data Constituents in Data ONTAP 8.2	25
	When are data constituents created?	25
	What determines how many data constituents are created and their size?	25
	Are there any requirements for data constituents if I have a SnapMirror relationship for Infinite Volume?	25
	Can I access the data constituents? If yes, how?	26
2.6	SnapDiff, Available Starting in Data ONTAP 8.2	26
	When is SnapDiff configured?	26
	What requirements are there for enabling SnapDiff?	26
	Where can I restore to with SnapDiff?	26
2.7	Capacity Balancing	26
	Which nodes are included in Infinite Volume load balancing?	26
	Do the aggregates of the Infinite Volume have to be of equal size?	26
	What algorithm does Infinite Volume use for capacity balancing when placing files?	26
	Does Infinite Volume support automatic capacity balancing?	27
3	Manageability	27
	What management user interfaces are available for Infinite Volume?	27
	What operations are managed at the Infinite Volume level, and what is managed at the storage class level?	27
	Does Infinite Volume support Manage ONTAP?	27
	Can I delete an Infinite Volume?	27
3.1	Infinite Volume Space Management	28
	What happens when the data constituent gets filled up?	28
	What happens when a storage class gets filled up?	28
	Can the size of an Infinite Volume be increased?	28
	Can the size of an Infinite Volume be decreased?	29
	Can a storage class within an Infinite Volume be deleted?	29
	What happens if I append to an existing file but the data constituent containing the file becomes full? ..	29
	What type of notifications or alerts can be set for an Infinite Volume?	29
	Can two different aggregates, one with SAS disks and another with SATA disks, be added to an Infinite Volume?	29
	Can the file system size be increased when adding new controllers to the cluster?	29
4	Infinite Volume Security	29
	What security style is supported for Infinite Volume?	29
	What is the difference between mixed and unified security style?	29
	What type of locking is implemented on Infinite Volume?	31
5	Migration	31
	Can data be migrated from a FlexVol volume in clustered Data ONTAP or 7-Mode or 7G to an Infinite Volume?	31
	Does Infinite Volume support Volume Move?	31
	How are data files moved from one data constituent to another in an Infinite Volume?	31

How are data files moved between storage classes with an Infinite Volume?	32
6 Reliability and Availability	32
Is Infinite Volume capable of detecting a corrupted file?	32
How does a constituent failure affect Infinite Volume?	32
Does the SVM for Infinite Volume support any high-availability options for the SVM root?	32
7 Data Protection	33
7.1 Snapshot.....	33
How does Snapshot work for Infinite Volume?	33
What is the recommended Snapshot schedule for Infinite Volume?.....	33
How does the Infinite Volume–level Snapshot maintain consistency across all the data constituents? ..	34
What new Snapshot copy states are introduced in Infinite Volume?	34
What is the maximum number of Snapshot copies possible for an Infinite Volume?.....	34
7.2 SnapMirror	34
How can I create a volume SnapMirror volume of an Infinite Volume?	34
Can I set different volume SnapMirror update schedules for different storage classes?	34
Can I do a resize of an Infinite Volume that is configured with volume SnapMirror?	34
Can I use SnapMirror to mirror an Infinite Volume within a cluster (intracluster)?	34
Does Infinite Volume support both synchronous and asynchronous SnapMirror?	34
What are the SnapMirror requirements for Infinite Volume?.....	35
How is a failed or missing constituent recovered on the primary site from the secondary site?	35
Can the cluster that is used for Infinite Volume disaster recovery (DR) use a system model that is different from the primary (that is, can FAS3270 be used on either source or destination and FAS6280 on the other)?	35
Can the namespace mirror constituent be a load-sharing mirror?	35
Can the Infinite Volume be configured in a fan-out or bidirectional SnapMirror configuration?	35
How does the SnapMirror function in a FlexVol volume differ from that in an Infinite Volume?	36
Does Infinite Volume support QoS parameters and transfer priorities?	36
7.3 Tape Backup.....	37
What type of tape backup solution is applicable for Infinite Volume?	37
What type of tape backup solution is applicable for Infinite Volume?	37
Is NDMP backup or SMTape backup supported for Infinite Volume?	37
8 Storage Efficiency	37
Does Infinite Volume support compression and deduplication?	37
Can I monitor progress of the deduplication or compression for the Infinite Volume?	37
Can I modify volume efficiency policy for when deduplication and compression run on the Infinite Volume?	37
Can I set different deduplication and compression schedules or policies for individual data constituents in the Infinite Volume?	37
Does Infinite Volume support thin provisioning?.....	38
Can I use the Space Savings Estimation Tool (SSET) to estimate compression and deduplication savings?	38

9	Performance Characteristics.....	38
	Does the read/write performance increase if more controllers are added?	38
	Does the read/write performance increase if Flash Cache or Flash Pool is used?	38
	Can Flash Cache or Flash Pool improve performance of namespace constituent activity?	38
	Can SAS disk drives improve performance of constituent activity?	38
	What is the default setting for access time update?	38
	Does Infinite Volume support QoS policies?	39
	Are there performance benefits when using pNFS?	39
	What are the performance characteristics of Infinite Volume in clustered Data ONTAP?	39
	If SAS disk drives are used for the namespace constituent, is there any difference in the recommended configuration for the namespace constituent?	39
10	Feature Interoperability.....	40
	What features does Infinite Volume support starting with Data ONTAP 8.2?	40
	What features does Infinite Volume not support in Data ONTAP 8.2?	41
11	Specifics of Infinite Volume in Data ONTAP 8.1.x.....	41
	Is Infinite Volume generally available in Data ONTAP 8.1.x?	41
	Does the Infinite Volume feature require a license in Data ONTAP 8.1.x?	41
	What is the reason for putting a PVR in the sales process for Infinite Volume in Data ONTAP 8.1.x?	41
	Why is a PVR key required in Data ONTAP 8.1.x?	42
	How can I get a PVR key for Infinite Volume in Data ONTAP 8.1.x?	42
	How is a PVR key used for Infinite Volume in Data ONTAP 8.1.x?	42
	If a storage controller is replaced, how can I get a new PVR key for Infinite Volume in Data ONTAP 8.1.x?	42
	How many Infinite Volumes can I have in a cluster in Data ONTAP 8.1.x?	42
	Are heterogeneous nodes supported in an Infinite Volume (that is, can FAS3270 and FAS6280 be mixed)?	42
	Which NetApp platforms does Infinite Volume support in Data ONTAP 8.1.x?	42
	Which protocols does Infinite Volume support in Data ONTAP 8.1.x?	42
	What security type does Infinite Volume support in Data ONTAP 8.1.x?	42
	Can an Infinite Volume and a FlexVol volume coexist in a cluster in Data ONTAP 8.1.x?	42
	Can I use System Manager to manage Infinite Volume in Data ONTAP 8.1.x?	42
	How do I size an Infinite Volume in Data ONTAP 8.1.x?	43
11.1	Disk Configuration Recommendations for Data ONTAP 8.1.x	43
	Before I create an Infinite Volume is there anything I should do with my disks in Data ONTAP 8.1.x? ...	43
	Can you use different disk drive types for the namespace and the data constituents within an Infinite Volume in Data ONTAP 8.1.x?	43
	Can two different aggregates, one with SAS disks and another with SATA disks, be added to an Infinite Volume in Data ONTAP 8.1.x?	43
	What is the recommended configuration for an Infinite Volume in Data ONTAP 8.1.x?	43
	Why shouldn't I create an aggregate larger than 100TB for FAS6280 or 70TB for FAS3270 for an Infinite Volume in Data ONTAP 8.1.x?	45
	What size aggregate should I create if the recommended size listed in Table 7 is not possible for an Infinite Volume in Data ONTAP 8.1.x?	45

What is the maximum size aggregate recommended for an Infinite Volume in Data ONTAP 8.1.x?	45
Should I create several smaller aggregates or fewer larger aggregates for an Infinite Volume in Data ONTAP 8.1.x?	46
Can I specify which aggregate or node to use for the namespace constituent in Data ONTAP 8.1.x?	46
What is the recommended configuration for the namespace constituent in Data ONTAP 8.1.x?	46
11.2 Infinite Volume Setup in Data ONTAP 8.1.x.....	46
What management user interfaces are available for Infinite Volume in Data ONTAP 8.1.x?	46
How do I select which nodes in my cluster to use for Infinite Volume in Data ONTAP 8.1.x?	46
How do I create an Infinite Volume in Data ONTAP 8.1.x?	47
Is there any difference in configuring Infinite Volume with FAS3270 and FAS6280 with Data ONTAP 8.1.x?.....	47
What are the SnapMirror requirements for Infinite Volume in Data ONTAP 8.1.x?	47
Can the Infinite Volume be configured in a fan-out or bidirectional SnapMirror configuration in Data ONTAP 8.1.x?	47
Can the cluster that is used for Infinite Volume disaster recovery (DR) use a different model of systems from the primary in Data ONTAP 8.1.x?	48
What type of tape backup solution is applicable for Infinite Volume in Data ONTAP 8.1.x?	48
11.3 Capacity Balancing in Data ONTAP 8.1.x.....	48
Do the aggregates of the Infinite Volume have to be of equal size in Data ONTAP 8.1.x?	48
In Data ONTAP 8.1.x, can I increase the size of an Infinite Volume by increasing the size of the aggregates beyond 100TB for FAS6280 and 70TB for FAS3270?.....	48
11.4 Space Management in Data ONTAP 8.1.x.....	49
Can the size of an Infinite Volume be increased in Data ONTAP 8.1.x?	49
12 Technical Support	49
How do I get Infinite Volume support?	49
13 Terminology	49
References.....	50
Version History	50

LIST OF TABLES

Table 1) Comparing mixed mode to unified security.	30
Table 2) Comparison of FlexVol volume and Infinite Volume Snapshot.	33
Table 3) SnapMirror functionality comparison of FlexVol volume and Infinite Volume.	36
Table 4) Features supported with Infinite Volume.	40
Table 5) Features not supported with Infinite Volume in Data ONTAP 8.2.	41

LIST OF FIGURES

Figure 1) Infinite Volume overview.	10
--	----

FAQ Details

This FAQ answers basic questions about NetApp® Infinite Volume. This document is focused on Infinite Volume in clustered Data ONTAP 8.2. There is a later section for [Infinite Volume in Data ONTAP 8.1.x](#) specifics.

Note: All references to NetApp FlexVol® volumes in this document refer to their clustered Data ONTAP implementation.

1 Infinite Volume Overview

What is an Infinite Volume?

Answer: An Infinite Volume is a feature in clustered Data ONTAP that provides capacity up to 20PB and scalability across numerous aggregates and nodes. It uses a single junction path and integrates with NetApp technologies such as deduplication, data compression, asynchronous volume SnapMirror® technology, and OnCommand® System Manager.

NetApp Infinite Volume is a single large volume in a clustered Data ONTAP Storage Virtual Machine (SVM); data can be accessed by using the NFSv3, NFSv4.1, pNFS, and SMB 1.0 protocols. This volume can be managed through the Data ONTAP CLI, similar to a FlexVol volume or OnCommand Workflow Automation starting with version 2.1 and OnCommand Unified Manager starting with version 6.0. Its data is protected by the use of NetApp Snapshot™ technology and asynchronous volume SnapMirror and is backed up with NFS- or SMB-mounted tape backup.

What are the components of an Infinite Volume?

Answer: The basic components of an Infinite Volume are listed structurally from the bottom layer:

- **Aggregate.** A collection of disk drives configured with RAID.
- **Namespace constituent.** A type of container that contains the directory hierarchy and file names. The file names are associated with redirector (pointer) files, each of which refers to a data file in a data constituent. There is only one namespace constituent per Infinite Volume.
- **Namespace mirror constituent.** An intracluster volume SnapMirror copy of the namespace constituent. It provides data protection for the namespace constituent, and starting in clustered Data ONTAP 8.2, supports SnapDiff for incremental tape backup support.
- **Data constituent.** A type of container that stores data files. Numerous data constituents are allowed in an Infinite Volume; however, a minimum of two data constituents is required.
- **SVM with Infinite Volume.** A storage virtual machine with a special attribute that hosts an Infinite Volume.
- **Storage class.** A grouping of storage within an Infinite Volume that can be defined by disk type and volume settings. This is optional.

What are the major features of Infinite Volume in clustered Data ONTAP?

Answer: Infinite Volume provides the following capabilities:

- A single namespace
- Aggregated throughput optimized for large I/O
- A single large pool of storage scalable up to 20PB
- Cluster-wide capacity balancing for all the data constituents at the time of ingest, based on the capacity threshold values of the data constituents
- Data protection at the Infinite Volume level using Snapshot and SnapMirror
- Inherited system-level high availability (HA)

What are the capacity limits of Infinite Volume?

Answer: Infinite Volume is optimized to handle the following capacity limits:


- Up to 20PB raw capacity
- Up to 2 billion files
- File size up to 16TB
- Spanning up to a maximum of 10 nodes or 5 HA pairs in a cluster; however, the cluster can contain more nodes

How does Infinite Volume work?

Answer: An Infinite Volume is a compound volume that is composed of numerous containers like FlexVol volumes known as *constituents*. These constituents are grouped as a namespace constituent and data constituents. The core architecture is based on a software layer of redirection that distributes file data among all data constituents running in a clustered Data ONTAP solution. The NAS protocols provide a single junction path for the Infinite Volume. The namespace constituent stores the redirector file names that contain details as to which data constituents own which data files.

Figure 1 shows the NetApp Infinite Volume as a single large container that can store billions of files in numerous data constituents. NS represents the namespace constituent, NSm represents the namespace mirror constituent, and DC1 and DC2 represent data constituents.

Figure 1) Infinite Volume overview.



How many Infinite Volumes can I have in a cluster?

Answer: Starting in Data ONTAP 8.2, the number of Infinite Volumes you can have in a cluster is limited to the number of SVMs the cluster can support. However, as a best practice NetApp recommends that you limit the number of Infinite Volumes to no more than the number of nodes you have in the cluster.

Which NetApp platforms does Infinite Volume support?

Answer: Starting in clustered Data ONTAP 8.2, Infinite Volume is supported on all FAS3000 and FAS6000 series controllers that support this release of Data ONTAP. Equivalent V-Series controllers are also supported.

Which versions of Data ONTAP support Infinite Volume?

Answer: Infinite Volume is available starting with clustered Data ONTAP 8.1.1 under PVR and is GA with clustered Data ONTAP 8.2. Infinite Volume cannot be configured on a 7-Mode system.

Can I configure Infinite Volume on 7-Mode systems?

Answer: No, Infinite Volume is supported only with clustered Data ONTAP 8.1.1 and later. Infinite Volume cannot be configured on a 7-Mode system.

Are heterogeneous nodes supported in an Infinite Volume (that is, can FAS3270 and FAS6280 be mixed)?

Answer: Yes, starting in clustered Data ONTAP 8.2, an Infinite Volume will support a mix of different controllers.

Are 32-bit and 64-bit aggregates supported with Infinite Volume?

Answer: No, only 64-bit aggregates are supported with Infinite Volume.

Why is the number of files supported in Infinite Volume limited to 2 billion?

Answer: This is the standard FlexVol volume limit that applies to the namespace constituent where the complete Infinite Volume namespace hierarchy is stored. Therefore, the maximum number of possible data files supported is 2 billion.

Does Infinite Volume allow deeper directories than a FlexVol volume?

Answer: No, the entire directory structure is stored in the namespace constituent. Because this is a single container it inherits the same limitation that a single FlexVol volume has. For recommendations on how to best architect your high file count and directory infrastructure, refer to TR-3537.

Does Infinite Volume have the same high-file-count limitations that FlexVol has?

Answer: Yes. For recommendations on how to best architect your high-file-count and directory infrastructure, refer to TR-3537.

How many redirector files are created in the namespace constituent?

Answer: Every file stored in an Infinite Volume gets a redirector file along with the data file. Therefore, two files are created for each file: one redirector file in the namespace constituent and one file containing data in the data constituent.

Does the admin have to manually increase the number of inodes in the namespace constituent?

Answer: No, whenever we create or expand an Infinite Volume, we automatically set the “files” property on it to the largest value it can be, based on the current size. The admin doesn’t have to worry about it.

How is metadata handled in Infinite Volume?

Answer: Metadata is part of the file and is stored with the file in any one of the data constituents. However, the namespace constituent also contains some parts of the metadata, including the name and file handle of the data file. Table 1 describes the metadata split for the namespace redirector between the namespace constituent and the data constituent.

Table 1) Metadata split between namespace and data constituents.

Namespace Constituent	Data Constituent
Name of the file	Content of the file
Name of the stream(s) attached to the file	Content of the stream(s) attached to the file
Name and content of the stream(s) attached to a directory	
Filehandle	
Files_Total	
	Size of the file
	Locks
Space_Free	Space_Free
Hardlinks	Hardlinks
Mode bits	Mode bits
ACL	ACL


How do I size an Infinite Volume in Data ONTAP 8.2 for capacity?

Answer: Starting in Synergy version 5.4, it should be used to size an Infinite Volume for capacity. Synergy will size things the same as it does for FlexVol volumes with the following exceptions:


- Includes an additional 20TB for the namespace, or 10TB/node if you plan to use SnapDiff
- Isolates 9% of the aggregate sizes to not be available for constituents
- Does not exceed 10 nodes per Infinite Volume
- Provides equal free aggregate space across nodes

To configure an Infinite Volume in Synergy, use Storage Design Studio.


- 1) After adding your nodes and disks and configuring your aggregates, right-click the cluster and select Add Vserver.




- 2) Right-click the SVM (Vserver) you just created and select Add Infinite Volume. This configures this SVM to be available for an Infinite Volume and creates an SVM root volume.




- 3) Assign the appropriate aggregates that you want the Infinite Volume to use.



- 4) Create an Infinite Volume by either right-clicking the SVM or selecting +Infinite Vol on the bottom-left corner.



- 5) This shows the largest possible Infinite Volume for the aggregates that were included in the configuration. It also shows the Infinite Volume name, junction path, associated SVM, constituent breakout, total Infinite Volume size, and data usable capacity. You can reduce the Infinite Volume size, enable/disable thin provisioning, change the snapshot reserve, add/remove SnapDiff, set a maximum data constituent or namespace constituent size, and assign a specific aggregate for the namespace constituent.




You can continue to create as many SVMs with FlexVol volumes and SVMs with Infinite Volume as required for your cluster.

How do I size an Infinite Volume in Data ONTAP 8.2 for performance?

Answer: Infinite Volume is optimized for large (average file size 100KB or larger) unstructured static data. If your workload meets this criteria, you can use the system performance modeler tool (SPM) to size Infinite Volume for performance. To accomplish this make the following changes in your sizing:

- ensure capacity reserve is at least 10%



Advanced Configuration

CPU Headroom Percent: 30.0

Map to Full Shelves: ☐

Capacity Reserve: 10 %

- select “Custom Application” as the workload type




Step 4

Choose Workload Type

Workload Type: Custom Application

- increase your performance requirements by ~10% and your capacity requirements by the size of the namespace constituent and namespace mirror constituent(s), default is 20TB



Step 5

Workload Description:


Protocol: NFS v3

Throughput*: 550

Throughput Type: MB/s

Required Capacity (TB)*: 220

- select “All nodes in the cluster” for how the workload traffic arrives into the cluster.




Step 5

Layout Hints

Can this workload be placed on a shared aggregate? : ☒

Can the workload be split on different aggregates across the cluster?: ☒

How will the workload traffic arrive into the cluster?: ☐ On a Single node ☒ All nodes in the cluster



Note: If your workload doesn't meet this criteria it may not be valid for Infinite Volume and this sizing methodology should not be used.

1.1 Storage Classes and Data Policies

What is a storage class?

Answer: A storage class is an optional way of grouping different disk types and/or different volume settings together within an Infinite Volume. As an example, you may create different storage classes for SATA and SAS drives. You may also create different classes of storage for deduplication and compression settings.

What is a data policy?

Answer: A data policy contains a set of rules that direct placement of new file writes during ingest to specific storage classes. The client writes a new file to the Infinite Volume as a whole, and Data ONTAP uses the rules in a data policy to direct to which storage class the file will be written. Only new writes will use data policies; modifications to files already written in the Infinite Volume will not redirect within which storage class the file is contained unless the file is saved as a new file with a new name.

A data policy can contain rules based on file type, file owner, and directory path.

Will modifications of files already written to the Infinite Volume result in it being moved to a different storage class?

Answer: No, the data policy is used only on the first write of a file. If someone else modifies the file and saves it with the same name, it will remain in its current location. If someone modifies the file and saves it with a new name, it could be written to a new storage class.

Why would I create a storage class and data policy?

Answer: An example of when you may create different storage classes could be for different deduplication and compression settings. This can be used in conjunction with a data policy to direct compressible data into a storage class with compression enabled and noncompressible data into a class that doesn't have compression enabled. This can reduce the amount of load on the system from attempting to compress data that isn't compressible. The same can be done for deduplicatable and nondeduplicatable data.

You could also create different storage classes for different disk types or systems with and without Flash Cache™ or Flash Pool™. This could be used by a data policy to group your more latency-sensitive or frequently accessed files on faster disks or systems with additional cache.

What version of clustered Data ONTAP supports storage classes and data policies?

Answer: Starting in clustered Data ONTAP 8.2, Infinite Volume supports storage classes and data policies.

What tool do I use to create and manage a storage class?

Answer: Storage classes must be created and managed using OnCommand Workflow Automation 2.1 or higher. They can be monitored using OnCommand Unified Manager version 6.0 or higher.

What tool do I use to create and manage a data policy?

Answer: To create a data policy to make use of more than one storage class, you must use OnCommand Unified Manager version 6.0 or higher. OnCommand Unified Manager will allow you to create a data policy with rules specific to file type, owner, and directory location and assign each rule to a specific storage class within an Infinite Volume.

Does an Infinite Volume have storage classes by default?

Answer: No, by default when you create an Infinite Volume, it does not contain storage classes. If you create the Infinite Volume using OnCommand Workflow Automation (WFA) 2.1 or higher, there is an option to create it with or without storage classes.

Can I add a storage class to an Infinite Volume that was created without storage classes?

Answer: Yes, you can add a storage class to an Infinite Volume that was created without storage classes. This will create new data constituents in a new storage class; the preexisting storage class type will be legacy. Any modifications such as a resize operation to the preexisting storage class would require a custom workflow.

Does OnCommand WFA have storage class workflows available by default?

Answer: Yes, starting with OnCommand WFA 2.0, it contains two storage class workflows, capacity storage class and performance storage class.

Default Settings	Capacity Storage Class	Performance Storage Class
Drive type	SATA	SAS
Deduplication enabled	Yes	Yes
Compression enabled	Yes	No
Inline compression enabled	Yes	No
Volume efficiency policy	Midnight–6 a.m. daily	Weekends 6 a.m.–midnight
Volume efficiency priority	Background	Background
Space guarantee	Volume	Volume

Additional storage class workflows can be created to meet your customers' specific needs.

How can I identify which data constituents are part of each storage class?

Answer: Data constituents will be labeled in a format similar to <Infinite Volume name>_1024_data001. The designation "1024" will be used for the first storage class, "1025" for the second, and so forth.

Can storage classes ever be deleted or reduced in size?

Answer: No, a storage class can never be deleted without deleting the entire Infinite Volume. Like the Infinite Volume itself, it also can never be reduced in size.

1.2 Disk Configuration Recommendations for Data ONTAP 8.2

Can you use different disk drive types for the namespace and the data constituents within an Infinite Volume?

Answer: Yes, starting in clustered Data ONTAP 8.2, Infinite Volume supports mixed disks. As an example, you could use SAS disks for the namespace constituent and SATA disks for the data constituents. Optionally, you can also use different disk types when creating storage classes. By default the capacity storage class uses SATA drives, and the performance storage class uses SAS drives.

Before I create an Infinite Volume, is there anything I should do with my disks?

Answer: Yes. NetApp recommends that you create aggregates prior to creating an Infinite Volume. By default, Infinite Volume will use all available nodes in the cluster. Starting in Data ONTAP 8.2, you are able to limit the number of nodes in your Infinite Volume to something less than the number of nodes in your cluster. To accomplish this, after you create the SVM for Infinite Volume, you should assign aggregates to the SVM aggregate list by using the `vserver modify -aggr-list` command. This is also an option in OnCommand Workflow Automation when using the “Create and configure a Vserver with Infinite Volume” workflow.

What is the recommended configuration for an Infinite Volume in Data ONTAP 8.2?

Answer: In clustered Data ONTAP 8.2, Infinite Volume does not require dedicated aggregates. However, it is better to create larger data constituents rather than smaller ones. To reduce the likelihood of creating small data constituents, it is better to create aggregates at close to the maximum supported data constituent size +10% or multiples of it in your Infinite Volume. For example, if the maximum data constituent size is 100 TB, aggregates should be near to 110TB or 220TB, so that 100TB data constituents can fit on the aggregates.

The maximum possible size of a data constituent corresponds to the lowest of all platforms with aggregates in the Infinite Volume’s maximum FlexVol volume size. For example, a platform might have a maximum FlexVol volume size of 100TB, which means the maximum data constituent size on that platform is 100TB.

What is the recommended configuration for the namespace constituent?

Answer: For optimal usable capacity and performance, NetApp recommends the following for the namespace constituent:

- Verify usable space of at least 11TB.
- SAS drives might help reduce latency, especially for smaller files or directory lookups.
- Flash Cache or Flash Pool intelligent caching might help for namespace activities such as directory lookups.
- Create the namespace constituent on the node with the most memory. You can accomplish this by assigning an aggregate from this node to the namespace constituent by using the `-namespace-aggregate` option when creating your Infinite Volume.
- Clusters containing 2–6 nodes:
 - o Minimum of 8 disks configured with RAID-DP® technology (verify that usable space is at least 11TB)
- Clusters containing 8–10 nodes:
 - o Minimum of 12 disks configured with RAID-DP (verify that usable space is at least 11TB)
 - o Limit the amount of other activity on the node containing the namespace constituent, for example, no client-accessible LIFs, no data constituents, assign QoS policies on applications running on other SVMs. The limited amount of activity on this node is to dedicate as much CPU and memory as possible to the namespace constituent activities.
- If you have numerous Infinite Volumes in your cluster, NetApp recommends no more than one namespace constituent per node. This is the namespace constituent, not the namespace mirror constituent(s). You can assign a particular aggregate on a node to the namespace constituent using the `-namespace-aggregate` option when creating your Infinite Volume.

Can I specify which aggregate or node to use for the namespace constituent?

Answer: Yes, starting in Data ONTAP 8.2, you can accomplish this by assigning an aggregate from a specific node to the namespace constituent by using the `-namespace-aggregate` option when creating your Infinite Volume.

Does the namespace constituent always need to be 10TB?

Answer: No, if you are configuring a smaller Infinite Volume with less than 2 billion files, you can manually limit the size of the namespace constituent using the advanced option `-max-namespace-constituent-size`. To determine the minimum size of namespace constituent that you will need, use $(\text{max number of files desired}) / 2B * 10TB$. Example, for 1 billion files $((1 \text{ billion} / 2 \text{ billion}) * 10TB)$ you would need a total of 5TB. Remember that you need the same size for each of the namespace mirrors.

Can I limit which aggregates or nodes the Infinite Volume can use for the data constituents?

Answer: Yes, by default the Infinite Volume will use one or more aggregates from all of the nodes in the cluster. To limit from which aggregates or nodes the Infinite Volume will select, you can use the `vserver aggr-list` option when creating or resizing your Infinite Volume. This is also an option in OnCommand Workflow Automation when using the “Create and configure a Vserver with Infinite Volume” workflow.

1.3 Flexible Volume and Infinite Volume Comparison

What are some major differences between an Infinite Volume and a FlexVol volume in clustered Data ONTAP?

Answer: Infinite Volume observes all major attributes and characteristics of FlexVol volumes and is in essence a big container that includes all underlying constituents as containers like FlexVol volumes.

Table 2) Comparison of FlexVol volume and Infinite Volume features for clustered Data ONTAP.

Attributes	FlexVol Volumes	Infinite Volume
Max size with 10 FAS6280s	Up to 100TB	Up to 20PB raw
Protocol support	NFSv3/v4, SMB1/2, iSCSI, FCP	NFSv3/v4.1/pNFS, SMB1
Junction path	Each FlexVol volume is attached to the SVM at a junction path defined by the administrator, which appears as a directory to NAS clients	The SVM for Infinite Volume contains a single junction path for the Infinite Volume; individual constituents within the Infinite Volume cannot be mounted The junction path for an Infinite Volume must be rooted under <code>/.</code> and can only be one level, ex. <code>/junction</code> , not <code>/junction/path</code>
df space reporting	Reports space available for an individual FlexVol volume under one junction	Reports the total space available for the entire Infinite Volume from the root of the namespace
Parent volume capability	Yes, with several junctions	No, because no child junctions are allowed
Capability to behave as an SVM root volume	Yes	No
Data placement	Data is placed in respective volume	Autoplacement at ingest
Security type	Mixed, NTFS or UNIX	Unified
Load share mirror, FlexCache®, SyncMirror®, FlexClone®	Yes	No
Antivirus, NDMP, quotas, Single File SnapRestore® support	Yes	No
Volume move functionality	Yes	• No, at the Infinite Volume level • Yes, for the data constituent but only with the help of Customer Success Services

There are certain differences in some features and functionality compared to those in FlexVol volumes. For example:

- **Security type.** Starting in Data ONTAP 8.2, Infinite Volume only supports unified security style.
- **Partial availability.** In an Infinite Volume, in case of volume degradation, a majority of data files are accessible, while some data files are not accessible from an offline data constituent. In contrast, in a FlexVol volume, no files are accessible in case of degradation.
- **Performance.** In comparison to a FlexVol volume, the creation, deletion, and renaming operations on files are more expensive on the Infinite Volume because of the namespace constituent, which mandates two operations: one to create the namespace redirector file and one to create the data file. Therefore, small file access to unique files within an Infinite Volume is an expensive operation, because it is optimized for large (>100KB) file access.

Can an Infinite Volume and a FlexVol volume coexist in a cluster?

Answer: Yes, starting in Data ONTAP 8.2 you can have a cluster that supports multi-tenancy. This can include several SVMs for Infinite Volume as well as several SVMs for FlexVol volumes.

What is the difference between the clustered Data ONTAP namespace and the data stored in the Infinite Volume namespace constituent?

Answer: The clustered Data ONTAP namespace is built by using junctions (a hierarchy of junctions) under the root volume, and the FlexVol volumes can join other FlexVol volumes by using these junctions. In the case of Infinite Volume, a single junction namespace, with /NS as the root of the namespace, maps redirector path names to the data files stored in the data constituents by using a redirector file.

2 Infinite Volume Setup

At what size will an Infinite Volume be created by default?

Answer: If you don't specify a size for the Infinite Volume it will be created at the minimum possible size which is approximately 1.33TB times the number of nodes with aggregates assigned to the Infinite Volume.

How do I select which nodes in my cluster to use for Infinite Volume?

Answer: Starting in Data ONTAP 8.2, you are able to limit the number of nodes in your Infinite Volume to something less than the number of nodes in your cluster. To accomplish this, after you create the SVM for Infinite Volume, you should assign aggregates to the vservers aggregate list by using the `vserver modify -aggr-list` command. This is also an option in OnCommand Workflow Automation when using the "Create and configure a Vserver with Infinite Volume" workflow.

Does the Infinite Volume feature require a license?

Answer: No, Infinite Volume does not require a separate license to be installed, but in Data ONTAP 8.1.x a special hash key is required to configure. This key is available through a feature product variance request (PVR) for Data ONTAP 8.1.x.

Is Infinite Volume generally available?

Answer: Yes, starting in Data ONTAP 8.2, Infinite Volume is GA. It will remain under PVR in Data ONTAP 8.1.x.

What is the recommended way to configure an Infinite Volume?

Answer: In Data ONTAP 8.2 you can use the CLI or OnCommand WFA version 2.1 or higher to configure an Infinite Volume.

Can I use System Manager to manage Infinite Volumes in Data ONTAP 8.2?

Answer: No, not at this time. A new version of System Manager that will support Infinite Volume is under development. If your customer requires this, send an e-mail to smoulton@netapp.com to request an early build.

How do I create an Infinite Volume in Data ONTAP 8.2?

Answer: The Infinite Volume is a compound volume that includes containers like FlexVol volumes that serve as a namespace constituent, namespace mirror constituent(s), and several data constituents.

Starting in Data ONTAP 8.2, after the cluster is created, the following steps are used to create an Infinite Volume:

- 1) Create an SVM for Infinite Volume by including the `--is-repository true` option in the `vserver create` command.
- 2) Create LIFs for the SVM for Infinite Volume to use.
- 3) If they are not already created, you should create aggregates.
- 4) Recommended: Assign aggregates to the Infinite Volume by using the `vserver modify` command with the `--aggr-list` option.
- 5) Create the Infinite Volume.
- 6) Add NFS and/or SMB1.0.

These steps can be performed using either OnCommand WFA version 2.1 or higher or the CLI.

With OnCommand WFA 2.1 you can create the Infinite Volume using the default workflow labeled "Create and configure a Vserver with Infinite Volume."

With the command line you can create the Infinite Volume using commands similar to those used when creating FlexVol volumes in clustered Data ONTAP. See TR-4178: Infinite Volume Deployment and Implementation Guide for more details.

Note: If you want to create storage classes, you must use OnCommand WFA version 2.1 or higher.

How many SVMs for Infinite Volume can I create?

Answer: Starting in Data ONTAP 8.2, the number of Infinite Volumes you can have in a cluster is limited to the number of SVMs the cluster can support. However, as a best practice NetApp recommends that you limit the number of Infinite Volumes to no more than the number of nodes you have in the cluster.

How do I create and manage a storage class?

Answer: Storage classes must be created and managed using OnCommand Workflow Automation 2.1 or higher. There are default workflows to add or expand a storage class within an Infinite Volume, or you can create a custom workflow.

Storage classes can be monitored using OnCommand Unified Manager version 6.0 or higher.

How do I create and manage a data policy?

Answer: The default data policy places all incoming data in the Infinite Volume as a whole if there are no storage classes created. If there are one or more storage classes, it places all the data into the first storage class that was created. To create custom data policies to make use of more than one storage class, you must use OnCommand Unified Manager version 6.0 or higher.

To create a storage class in OnCommand Unified Manager:

1. Select storage > Vservers and select the appropriate Vserver that contains multiple storage classes.
2. Select the Data Policy tab.
3. Create a data policy rule:
 - a. Click Create.
 - b. Choose an appropriate rule template from the drop-down list (file type, owner, or directory path).
 - c. Add the necessary conditions. For example, if you selected file type, place the specific type you for which want to filter. Select Add button and type the file extension under “List of file names ending with” section.
 - d. Select the storage class to which you want content that matches the criteria you select to be written.
 - e. Select Create.
4. Repeat the data policy rule creation step as many times as necessary to create all the necessary rules required for the data policy.
5. Select the default rule in the data policy tab and select Edit.
 - a. Select the storage class to which you want files to be written that do not meet the criteria of other rules within this data policy.
6. When you are finished adding rules, scroll to the bottom of the rule template and select Activate.

What additional actions can I perform in admin mode?

Answer: The following list describes some of the tasks you can perform on the Infinite Volumes in admin mode.

Admin mode:

- Add aggregates to increase size (nondisruptive action).
- Add disks to the aggregate to increase size (nondisruptive action).
- Increase the volume size (nondisruptive action).
- Create a Snapshot copy.
- Create and manage an Infinite Volume SnapMirror relationship.
- Mount/unmount the Infinite Volume.
- Enable/disable SnapDiff (available starting in Data ONTAP 8.2).
- Manage volume efficiency.
- Assign/change volume guarantee (thin provisioning).
- View space usage and savings.
- Use aggregate relocate.

For what additional actions do I require CSS support?

Answer: The following list describes some of the tasks that require Customer Success Services to perform them.

Customer Success Services required:

- Fail over and recover the namespace constituent from the namespace mirror constituent.
- Check the file handle for a specific file.
- Delete a constituent and its containing aggregate.
- Move a data constituent to a different aggregate.
- Constituent level management.
- Recover from SVM root volume failure.
- REALLOC (can only be run at the constituent level, not at the Infinite Volume level).
- Consistency scan.

What login levels are used for CLI management of Infinite Volume?

Answer: CLI management supports both the cluster admin and SVM admin levels.

Are qtrees allowed in Infinite Volume? Would an error message appear if the volume qtree create command is executed?

Answer: You cannot create a qtree or quota on an Infinite Volume or its constituents. If you try, you receive an error message.

2.1 NFS Configuration of Infinite Volume

How does the NFS client access an Infinite Volume resource?

Answer: When deploying Infinite Volume, a single junction point (for example, /NS) is created that points to the namespace constituent. The NFS clients can mount the Infinite Volume at or below this junction path.

How do I set up and configure a network interface to access the Infinite Volume?

Answer: Infinite Volume network management follows the network management best practices of clustered Data ONTAP. Setting up the network interface to access the Infinite Volume is a three-step process:

- 1) Create data logical interfaces (LIFs) at the SVM level to have a minimum of one LIF per physical node in the cluster.
- 2) Enable NFS access at the SVM level.
- 3) Create a network routing group to route data LIFs to outbound traffic.

2.2 SMB1.0 Configuration of Infinite Volume (Available Starting in Data ONTAP 8.2)

How does the CIFS client access an Infinite Volume resource?

Answer: When deploying Infinite Volume, a single junction point (for example, /NS) is created that points to the namespace constituent. The CIFS clients can map shares that point to this junction path or below to access the Infinite Volume over the SMB1.0 protocol.

How do I set up and configure a network interface to access the Infinite Volume?

Answer: Infinite Volume network management follows the network management best practices of clustered Data ONTAP. Setting up the network interface to access the Infinite Volume is a four-step process:

- 1) Create data logical interfaces (LIFs) at the SVM level.
- 2) Enable CIFS access at the SVM level.
- 3) Create a network routing group to route data LIFs to outbound traffic.
- 4) Create a CIFS share that points to the junction path for CIFS data access to the Infinite Volume.

These steps can be performed in the CLI or with OnCommand Workflow Automation.

2.3 Namespace Constituents

When is a namespace constituent created?

Answer: The namespace constituent is automatically created as part of an Infinite Volume.

What size are namespace constituents?

Answer: In most situations the namespace constituent will be 10TB. Exceptions include when your Infinite Volume is less than 80TB or in Data ONTAP 8.2 and you set a `-max-namespace-constituent-size` (advanced mode option). Setting a `-max-namespace-constituent-size` will limit the number of files and directories that can be written to the Infinite Volume. To determine the minimum size of namespace constituent that you will need use $(\text{max number of files desired}) / 2B * 10TB$. Example, for 1 billion files $((1 \text{ billion} / 2 \text{ billion}) * 10TB)$ you would need a total of 5TB. Remember you need the same size for each of the namespace mirrors.

How many namespace constituents exist in an Infinite Volume?

Answer: There is only one namespace constituent per Infinite Volume.

2.4 Namespace Mirror Constituents in Data ONTAP 8.2

When are namespace mirror constituents created?

Answer: The namespace mirror constituent for data protection is created when the Infinite Volume is created. Additional namespace mirror constituents will be created either at the time SnapDiff is enabled or when aggregates from a new node are added to an Infinite Volume as part of a resize operation and SnapDiff is enabled.

What size are namespace mirror constituents?

Answer: Namespace mirror constituents will always be equal in size to the namespace constituent.

How many namespace mirror constituents will exist in an Infinite Volume?

Answer: Each Infinite Volume will contain at least one namespace mirror constituent for a data protection copy of the namespace constituent. If you have SnapDiff enabled it will create one namespace mirror constituent for each node that doesn't contain either the namespace constituent or the data protection namespace mirror constituent. As an example, if you have a six-node configuration, you would have one namespace constituent, one namespace mirror constituent for data protection, and four namespace mirror constituents for SnapDiff.

Are namespace mirror constituents replicated as part of the SnapMirror replication for Infinite Volume?

Answer: No, only the namespace constituent and data constituents are replicated during an Infinite Volume SnapMirror transfer.

Are namespace mirror constituents created on the SnapMirror destination for Infinite Volume?

Answer: No, there are no namespace mirror constituents created on the SnapMirror destination unless SnapDiff is enabled on the destination. If the SnapMirror destination Infinite Volume becomes source during a failover and SnapDiff is not enabled NetApp recommends that you manually create one.

2.5 Data Constituents in Data ONTAP 8.2

When are data constituents created?

Answer: Data constituents will automatically be created during Infinite Volume creation and, if required, during a resize operation.

What determines how many data constituents are created and their size?

Answer: Data constituents will be created as follows:

- 1) Namespace and namespace mirror constituents are created.
- 2) The remaining Infinite Volume size (total Infinite Volume size requested – sum (size of namespace constituent + size of namespace mirror(s))) is divided equally across the number of nodes that have aggregates assigned in the aggregate list.
- 3) Look at the system model with aggregates assigned to the Infinite Volume and determine the smallest supported data constituent size (equal to the maximum 64-bit FlexVol size). Use this as the maximum for all nodes in the Infinite Volume unless a max-data-constituent-size is manually set to a lower value. For example, if you have a six-node cluster, four FAS6280s that have a max constituent size of 100TB and two FAS3270s that have a maximum constituent size of 70TB, all constituents will have a maximum constituent size of 70TB, even on the FAS6280 systems.
- 4) For each node, a data constituent will be created on the aggregate with the most free space. The size will attempt to be the largest possible to meet the needs of the data space required from Step 2 up to the maximum allowed from Step 3. We repeat this step until the node has created the data constituent space as required.

Note: If this was a resize operation we would attempt to grow any constituents that weren't already at the maximum configuration before creating new ones.

Are there any requirements for data constituents if I have a SnapMirror relationship for Infinite Volume?

Answer: Yes, the SnapMirror destination requires at least as much available space to support the source Infinite Volume. The destination systems must also be able to create constituents of an equal size to the source. This is especially important if you are trying to replicate from a source that supports larger constituent sizes than a destination; for example, the source contains a FAS6280 that has a max constituent size of 100TB and the destination has a FAS3270 that has a maximum constituent size of 70TB. In this circumstance, when you create the source Infinite Volume, you would have to specify a max data-constituent size of 70TB or the SnapMirror operation would never succeed. The size of a constituent can never be shrunk, so you need to determine which system model the destination will contain before creating the source.

Can I access the data constituents? If yes, how?

Answer: Yes, the storage administrator can view constituents. However, management at the data constituent level is only available in diag mode with the assistance of Customer Success Services.

Note: The Infinite Volume is set up to prevent NFS and SMB clients from directly accessing data constituents.

2.6 SnapDiff, Available Starting in Data ONTAP 8.2

When is SnapDiff configured?

Answer: SnapDiff can be enabled at any time, including at Infinite Volume creation.

What requirements are there for enabling SnapDiff?

Answer: There needs to be enough space to accommodate one namespace mirror constituent per node that doesn't already have a namespace mirror constituent or namespace constituent. In other words, assuming the namespace constituent is 10TB you need to have $10TB * (\text{number of nodes in the Infinite Volume} - 2)$.

Where can I restore to with SnapDiff?

Answer: Although you can back up from the SnapMirror destination, restoration from SnapDiff is only allowed to the source Infinite Volume.

2.7 Capacity Balancing

Which nodes are included in Infinite Volume load balancing?

Answer: Any node that has at least one aggregate in the aggregate list of the SVM with Infinite Volume is included in node balancing. If you also use the advanced `-data-aggr-list` parameter when creating the Infinite Volume, only nodes that have an aggregate in the data aggregate list are used for node balancing.

Do the aggregates of the Infinite Volume have to be of equal size?

Answer: No, the aggregates of the Infinite Volume don't have to be equal in size; capacity balancing across nodes is automatic. For ideal capacity balancing, and for getting balanced throughput across the aggregates, they should have similar available capacity and layout. For ideal balancing across aggregates NetApp recommends providing aggregates that are large enough to support creation of the largest supported data constituent size in the Infinite Volume.

What algorithm does Infinite Volume use for capacity balancing when placing files?

Answer: Infinite Volume provides capacity balancing across the nodes as data is ingested. When new files are written to an Infinite Volume, they are distributed evenly across data constituents in a round-robin fashion, with preference for data constituents that have the most available space.

In general, incoming files are created in the following manner:

- Data constituents with more than 80% used space generally do not receive incoming files.
- Data constituents with less than 80% used space receive incoming files in a round-robin fashion.
- Data constituents that are empty or nearly empty receive most of the incoming files.

Over time, this method of creating files tends to balance space usage across the data constituents. This is an automatic background function that helps storage administrators to avoid problems created by uneven capacity balancing. If a file cannot be created on a data constituent, Data ONTAP tries to create the file on other data constituents before returning an error to the client.

Does Infinite Volume support automatic capacity balancing?

Answer: Yes, Infinite Volume supports automatic capacity balancing by the preferential placement of data in underutilized data constituents at the time of data ingestion. The capacity utilization of the data constituents can become unbalanced over time due to reasons such as:

- The existing files in the data constituents might grow or shrink at different rates, resulting in some data constituents being underutilized in comparison to others.
- The deletion of files or Snapshot copies can result in freeing more spaces in some constituents than in others.
- Files may be compressed and/or deduplicated to different degrees in different data constituents.
- When new storage is added to the cluster and then incorporated into the NetApp Infinite Volume, the constituents created in the new storage are empty, causing underutilization of data constituents.

The underutilization problem is solved by preferentially distributing new data files to the underutilized data constituents. This is an automatic process, and the users do not notice any disruption. There is no postcreate capacity balancing; therefore, the imbalance caused by size change or deletion of existing files is not addressed by this preferential distribution.

3 Manageability

What management user interfaces are available for Infinite Volume?

Answer: Infinite Volume can be managed in the Data ONTAP CLI, with OnCommand Workflow Automation version 2.1 and higher and OnCommand Unified Manager version 6.0 and higher. An Infinite Volume is managed as a whole or as a storage class; however, constituent-level information can be viewed. Constituent-level management is done only with assistance from Customer Success Services.

What operations are managed at the Infinite Volume level, and what is managed at the storage class level?

Answer: Without storage classes, all Infinite Volume operations are managed at the Infinite Volume level. Storage class operations such as Snapshot copies, SnapMirror, and tape backup are done at the Infinite Volume level. Storage class operations such as resize and storage efficiency are performed at the storage class level.

Monitoring can be done at either the Infinite Volume or storage class level.

Does Infinite Volume support Manage ONTAP?

Answer: Infinite Volume introduced a set of new Manage ONTAP® (ZAPI) units specifically for Infinite Volume management that, starting in Data ONTAP 8.2, are officially supported.

Can I delete an Infinite Volume?

Answer: Yes, an Infinite Volume can be removed after the volume is taken offline.

Caution

When the Infinite Volume is deleted, the associated data is lost.

3.1 Infinite Volume Space Management

What happens when the data constituent gets filled up?

Answer: There might be numerous data constituents in an Infinite Volume, but not all of them will be full at the same time. However, if all data constituents are filled up, you must add more disks to an existing aggregate or add new aggregates. As a best practice, NetApp recommends adding more capacity to the Infinite Volume if the constituents are reaching the reserved capacity threshold, or about 70% of the capacity.

Infinite Volume will write files to data constituents that have the most available space. In rare occurrences that all data constituents are >80% full, a client might get the “not enough space” error while copying the file, even if the Infinite Volume as a whole has enough space to store a large file. To remedy this you can do one of the following:

- If all the existing data constituents are not already at maximum size and the containing aggregate has sufficient space, resize the Infinite Volume to grow the existing data constituent sufficiently.
- If the existing data constituents that are not already at maximum size have aggregates that are full but not at maximum size, add space to these aggregates and resize the Infinite Volume.

When you are updating an existing file, if there is insufficient space in the data constituent that stores the file and the data constituent is already at maximum size, you need to contact NetApp Support to help place this file to avoid this error.

As a best practice, NetApp recommends creating the aggregates (big enough to contain the largest possible data constituent) for an Infinite Volume. For details on the maximum and largest recommended configuration, see the section on [Disk Configuration Recommendations](#).

What happens when a storage class gets filled up?

Answer: If all data constituents in a storage class are full, new writes that the data policy sends to this storage class will fail; writes to other storage classes continue to work if they have available space. Writes will not be redirected to another storage class if the class to which the rule sends them is full.

To avoid this situation you should monitor the capacity usage of the storage class using OnCommand Unified Manager (minimally version 6.0) and set appropriate thresholds. If a storage class is reaching 70% full, you should increase the size of the storage class and if required add more disks to an existing aggregate assigned to the storage class or add new aggregates of a type that meets the criteria of the storage class.

Can the size of an Infinite Volume be increased?

Answer: Yes, an Infinite Volume can be increased in size without disrupting client access. Infinite Volume capacity can be increased by using one of these three methods.

- 1) If the constituents are not at maximum size and there is room in the aggregates simply resize the Infinite Volume.
- 2) Add disks to an existing aggregate assigned to the Infinite Volume and resize the Infinite Volume.
- 3) Add additional aggregates either on an existing node or a new node to the aggr-list for the Infinite Volume with the new set of disks and resize the Infinite Volume.

While the capacity is being increased, users can continue to add new files with no downtime or interruption.

Note: If the Infinite Volume contains storage classes, then the resize operation can only be performed at the storage class level using OnCommand WFA. You must also make sure if adding storage that it meets the criteria of the storage class you want to increase.

Can the size of an Infinite Volume be decreased?

Answer: No, neither the Infinite Volume as a whole nor a storage class can be shrunk.

Can a storage class within an Infinite Volume be deleted?

Answer: No, a storage class can't be deleted without deleting the entire Infinite Volume.

What happens if I append to an existing file but the data constituent containing the file becomes full?

Answer: The storage administrator receives a "not enough space" error while trying to save the file. The workaround for this would be to increase the size of the data constituent containing the file. If the data constituent is already at maximum size, then the user will have to save the update to a new file. This new file will get placed in a constituent that has enough space. The user would then have to manually delete the original file.

What type of notifications or alerts can be set for an Infinite Volume?

Answer: OnCommand Unified Manager 6.0 can perform capacity monitoring of the Infinite Volume, storage classes and individual data constituents. The storage administrator can set several thresholds for alerts, including warning and critical alarms.

Can two different aggregates, one with SAS disks and another with SATA disks, be added to an Infinite Volume?

Answer: Yes, starting in Data ONTAP 8.2 you can mix disk types within an Infinite Volume. However, NetApp recommends having aggregates of the same drive type assigned for data constituents. You may choose to create different storage classes for different types of disks or use one type for the namespace constituent and another type for the data constituents.

Can the file system size be increased when adding new controllers to the cluster?

Answer: Yes, the file system size can be increased by adding more storage controllers (nodes) to the Infinite Volume. However, an Infinite Volume should not span more than 10 nodes. You would do this by adding aggregates from the new node to the aggr list for the Infinite Volume and doing a resize. The standard clustered Data ONTAP operations allow seamless coordination between the Infinite Volume level of operation and the clustered Data ONTAP level of operation. There is no disruption to end-user interactions.

4 Infinite Volume Security

What security style is supported for Infinite Volume?

Answer: Starting in Data ONTAP 8.2, Infinite Volumes always use unified security style. Unified security style allows all users to view and set file permissions regardless of whether they are CIFS or NFS clients. Further, it allows the file permissions to include both Windows® and UNIX® users and groups, allowing more simplified security management.

Note: Although the SVM root volume might have been created with a different security style, this does not affect the Infinite Volume.

What is the difference between mixed and unified security style?

Answer: The main differences are ability for nonnative protocols to view permissions, how nonnative ACL changes occur, ability to merge mode bits, and ability for file permissions to contain both NFS and CIFS users and groups.

In mixed mode a client can only view native permissions. Further applying a new ACL causes the existing ACL to be dropped if it is different than the initial type of ACL (NFS or CIFS). In other words, applying an NFS ACL causes existing CIFS ACLs to be dropped; applying a CIFS ACL causes existing NFS ACL to be dropped. While mode bits set from a UNIX or NFS client could be merged with an existing NFSv4 ACL, it is not possible to merge UNIX mode bits into an NTFS ACL. It allows only user mapping (no group mapping) to the NFSv4 domain.

In unified security style any client can view and/or modify any ACL regardless of which protocol was used to previously set permissions (mode bits or ACL). Thus Windows users can view/modify an NFSv4 ACL and UNIX users can view/modify an NTFS ACL. It adds the option to merge UNIX mode bits into an NTFS ACL. It allows for both NFS and CIFS users and groups. It allows both user and group mapping to the NFSv4 domain.

Table 1) Comparing mixed mode to unified security.

	Mixed Mode	Unified Security
Can NFS clients view an existing CIFS ACL?	No	Yes
Can CIFS clients view an existing NFSv4 ACL?	No	Yes
What is the on-disk security type?	Whichever wrote it last	If an NFSv4 client performs the last Set ACL operation, the on-disk ACL becomes an NFSv4 ACL that provides unified access by allowing a mix of ACEs for UNIX and Windows users/groups. If a CIFS client performs the last Set ACL operation, the on-disk ACL becomes an NTFS ACL that provides unified access by allowing a mix of ACEs for UNIX and Windows users/groups. A chown/chmod/chgrp operation performed from an NFS client will result in the mode bits being merged with whatever ACL is present on disk (the on-disk ACL could be an NFSv4 ACL or an NTFS ACL).
What happens when an NFS client modifies an existing CIFS ACL?	NFS clients can only blindly overwrite an existing CIFS ACL causing the existing ACL to be dropped.	The changes will be merged into the existing ACL. UNIX users can view/modify a CIFS ACL without dropping the previous ACL.
What happens when a CIFS client modifies an existing NFS ACL?	CIFS clients can only blindly overwrite an existing NFS ACL causing the existing ACL to be dropped.	The changes will be merged into the existing ACL. Windows users can view/modify an NFSv4 ACL

Can UNIX mode bits be merged into a CIFS/NTFS ACL?	No, while mode bits from a UNIX or NFS client could be merged with an existing NFS ACL, it is not possible to merge UNIX mode bits into an NTFS ACL.	Yes, it allows both the option to merge UNIX mode bits into an NTFS ACL as well as merging the mode bits from a UNIX or NFS client with an existing NFSv4 ACL.
What type of users can be represented in an NFS ACL?	Only UNIX users	Both UNIX and Windows users and groups. Every Windows user/group must be mapped to a UNIX user/group.
What type of users can be represented in a CIFS ACL?	Only Windows users	Both UNIX and Windows users and groups. UNIX principals would appear in CIFS ACLs prefixed by unix-user or unix-group.
What type of mapping is allowed to the NFSv4 domain?	User mapping only	Both user and group mapping

What type of locking is implemented on Infinite Volume?

Answer: Locking is supported via NFSv3 for read/write functions; NFSv4 and CIFS share locks are supported. NFSv4.1 delegations are not supported with Infinite Volume. Starting in Data ONTAP 8.2, traditional oplocks are also supported and can be enabled on a CIFS share. Access on the active file system of the Infinite Volume SnapMirror destination is supported with NFSv3 protocol only. Snapshot file system on DP volumes can be accessed with NFS and CIFS protocols.

5 Migration

Can data be migrated from a FlexVol volume in clustered Data ONTAP or 7-Mode or 7G to an Infinite Volume?

Answer: No, data migration from a FlexVol volume in clustered Data ONTAP or 7-Mode or 7G to an Infinite Volume is not supported. The data can be manually copied from a FlexVol volume to an Infinite Volume by mounting both of the volumes on an NFS or a CIFS (in Data ONTAP 8.2) client and then performing the copy.

Does Infinite Volume support Volume Move?

Answer: Infinite Volume can't be moved as a whole; however, a constituent within the Infinite Volume can be moved but this is a diag mode command, and therefore the administrator must work with Customer Success Services to perform this action. This operation involves a few additional checks that are done for a vol move operation of FlexVol volumes. The extra checks include making sure the destination aggregate isn't residing on a 2000 series system and that if SnapDiff is enable, the rules about each node that contains a data constituent has a namespace mirror constituent are adhered to.

Note: Moving the Infinite Volume as a single container is not supported.

How are data files moved from one data constituent to another in an Infinite Volume?

Answer: Infinite Volumes are viewed as a whole to the user, and as such this moving should not be necessary. However, if required the administrator must work with Customer Success Services to perform this task. CSS will walk the administrator through stopping all access from the client to the data file, finding the data file in the constituent, and moving the file to another constituent.

How are data files moved between storage classes with an Infinite Volume?

Answer: Data policies are only used at the time of ingest. There is no way to move data between storage classes automatically. If a data policy is changed, only new files written after the new rule is activated will be placed in a different storage class; existing data on disk will not be moved. For example, if you have a data policy that places *.doc files in a performance storage class and you have written file.doc and file2.doc, these would be written to the performance storage class. If later you change the data policy to have *.doc files written to the capacity storage class, then if a client writes file3.doc, it would be written in the capacity storage class. File.doc and file2.doc would remain in the performance storage class even if they are modified. The only workaround would be to save file2.doc as a new file and manually delete the old file.

6 Reliability and Availability

Is Infinite Volume capable of detecting a corrupted file?

Answer: There are three protection layers in Infinite Volume to intercept and correct file corruption:

- Infinite Volume layer
- WAFL[®] layer
- RAID layer

Infinite Volume redirectors contain a magic number and a checksum that help to detect corrupted redirector or back pointer files. The Infinite Volume–level CRiron utility can be run with Customer Success Services to detect and fix file handle mismatch errors. Infinite Volume relies on WAFL and RAID to detect or correct byte corruption. CRiron must be executed after every WAFL iron, thereby enabling a complete check to rule out file system inconsistencies.

How does a constituent failure affect Infinite Volume?

Answer: Constituent failure is highly unlikely, but it might occur as a result of the failure of both nodes of an HA controller pair, software defects, aggregate failures, and so on. The impact of such failures on the Infinite Volume can be described as follows:

- **Namespace constituent failure.** The Infinite Volume's namespace (directory structure and file names) is no longer visible. However, clients with open file handles can still use those file handles without interruption.
- **Data constituent failure.** If one or more data constituents is unavailable, the Infinite Volume is shown as being in a mixed state. In this situation, you can still read the data that exists on the data constituents that are available, but the data that exists on the failed data constituents is not available. It is not easy to determine the set of files affected by a data constituent going off line.

If the constituent failure is permanent, the administrator must recover the Infinite Volume from a mirror site.

Does the SVM for Infinite Volume support any high-availability options for the SVM root?

Answer: A SVM for Infinite Volume root is protected via RAID-DP and HA failover. Infinite Volume does not support a volume SnapMirror or load sharing mirror for the SVM root. If the SVM root gets corrupted, the administrator will have to call CSS Support and have personnel walk him or her through recreating the SVM root.

Note: Losing the root volume of an SVM that contains an Infinite Volume is not the same as losing the root volume of an SVM that contains FlexVol volumes. The root volume of an SVM that contains an Infinite Volume only contains a single junction path and therefore can simply be recreated. No database is required to be restored.

7 Data Protection

7.1 Snapshot

How does Snapshot work for Infinite Volume?

Answer: Snapshot copies are managed at the Infinite Volume level, not at the individual data constituent level or the storage class level. Similarly to the way he uses a FlexVol volume, an administrator can create, delete, and restore data files by using Snapshot on the Infinite Volume along with the similar (hourly, nightly, weekly) Snapshot schedule management.

The Snapshot copy of Infinite Volume differs from that of a FlexVol volume in terms of performance, because all the constituents have to be fenced, Snapshot copies created, and unfenced together. The latency that this incurs for file operations increases linearly with the number of data constituents.

- A Snapshot copy can incur latency of up to 30 seconds if the number of data constituents is ≤ 50 .
- It can incur latency up to 2 minutes if the number of data constituents is 200.

The `.snapshot` directory for a FlexVol volume shows the files in the Snapshot copies only for that specific FlexVol volume, not for the entire namespace that is using numerous junctions in that hierarchy. A Snapshot copy of an Infinite Volume shows the files for the entire namespace at the Infinite Volume level. Table 2 compares the Snapshot functionality for the FlexVol volumes with that for the Infinite Volume.

Table 2) Comparison of FlexVol volume and Infinite Volume Snapshot.

Snapshot Functionality	FlexVol Volume	Infinite Volume
create	Yes	Yes
delete	Yes	Yes
restore	Yes	Yes
Single File SnapRestore	Yes	No
Compute space reclaimable	Yes	No
rename	Yes	No
show	Yes	Yes
Latencies less than 2 minutes for commands	Yes	No
Snapshot policy (includes any of the defined policies in <i>volume snapshot policy</i>)	Yes	Yes
Subvolume Snapshot copies (partial behavior)	N/A	No
Auto Snapshot delete	Yes	No
Manage ONTAP (ZAPI) support	Yes	Yes

What is the recommended Snapshot schedule for Infinite Volume?

Answer: NetApp recommends setting the Snapshot schedule to an hourly interval at a minimum. This is because a Snapshot copy of an Infinite Volume is created at the Infinite Volume level, rather than at an individual data constituent level. An Infinite Volume Snapshot copy spans numerous constituents across numerous nodes. During the Snapshot copy creation, the client I/O is fenced for a few seconds, which can affect client latency.

How does the Infinite Volume–level Snapshot maintain consistency across all the data constituents?

Answer: Infinite Volume–level Snapshot copies are crash-consistent Snapshot copies. The Snapshot process coordinates with each individual data constituent on each node to start the Snapshot creation process by fencing all I/Os while performing the Snapshot creation.

What new Snapshot copy states are introduced in Infinite Volume?

Answer: Infinite Volume introduces three states of Snapshot copies:

- **Valid.** This state indicates that the Snapshot copies of all constituents are intact.
- **Partial.** This state indicates that the Snapshot copies of one or more data constituents are missing.
- **Invalid.** This state indicates that the Snapshot copy of the namespace constituent is missing.

What is the maximum number of Snapshot copies possible for an Infinite Volume?

Answer: The maximum number of Snapshot copies for an Infinite Volume is slightly less than the 255 supported on a FlexVol volume. You need to subtract 2 for each namespace mirror constituent and another 2 if you have a SnapMirror relationship between Infinite Volumes.

7.2 SnapMirror

How can I create a volume SnapMirror volume of an Infinite Volume?

Answer: If your Infinite Volume doesn't have storage classes, you can use either the CLI or OnCommand Workflow Automation. If your Infinite Volume does have storage classes, you must use OnCommand Workflow Automation; the CLI will not work.

Can I set different volume SnapMirror update schedules for different storage classes?

Answer: No, volume SnapMirror operates at the Infinite Volume level, not the storage class level.

Can I do a resize of an Infinite Volume that is configured with volume SnapMirror?

Answer: Yes. If your Infinite Volume doesn't have storage classes you can use either the CLI or OnCommand Workflow Automation to resize the Infinite Volume. You should resize the destination Infinite Volume before resizing the source.

If your Infinite Volume does have storage classes, you must use OnCommand Workflow Automation; the CLI will not work. Since the resize operation with this workflow doesn't increase the size of the data constituents on the destination until the next SnapMirror update, you must make sure there is space on the destination aggregates that meet the criteria of the storage class being increased before you run this workflow.

Can I use SnapMirror to mirror an Infinite Volume within a cluster (intracluster)?

Answer: No, you must use SnapMirror to mirror an Infinite Volume to another cluster. The namespace mirror constituent is the only thing that can be replicated within the Infinite Volume in the same cluster.

Does Infinite Volume support both synchronous and asynchronous SnapMirror?

Answer: No, Infinite Volume supports only asynchronous SnapMirror in Data ONTAP.

What are the SnapMirror requirements for Infinite Volume?

Answer: SnapMirror is also based on Snapshot; it follows some major rules of creation and management (update, resync, break, reverse resync, and so on).

- The SnapMirror destination must be able to create an equal-size namespace space constituent as well as an equal number and size of data constituents as the source's Infinite Volume.
- The SnapMirror destination Infinite Volume space should be equal to or greater than the source.
- SnapMirror requires an SVM for Infinite Volume at both the source and the destination.
- In Data ONTAP 8.2, the SnapMirror destination can have a different model system than the source so long as the destination can create the same maximum data constituent size on the source. You can limit the size on the source using the `-max-data-constituent-size` option with the `volume create` command.
- The default workflows in OnCommand WFA require that the SnapMirror destination must be able to create equal-size constituents that meet the source storage class criteria. For example, if the source has two 50TB data constituents on SATA drives associated with the capacity class and four 60TB data constituents on SAS drives associated with the performance class, then the destination must be able to create the same two 50TB data constituents on SATA drives and four 60TB data constituents on SAS drives. You can create a custom workflow that does not have this requirement if you like.

Note: SnapMirror cascading is not supported for Infinite Volume.

How is a failed or missing constituent recovered on the primary site from the secondary site?

Answer: If the primary Infinite Volume has any failed or missing constituents (including the namespace constituent), the `snapmirror resync` operation can still be used to restore the primary Infinite Volume. The resync operation creates new data protection (DP) constituents and performs initialization for the missing constituents. However, for the rest of the constituents, `snapmirror resync` performs the resync operation at the constituent level, thereby avoiding full initialization of the constituent. The resync operation determines the deleted constituents based on the current Infinite Volume configuration.

Note: The administrator is prompted with a warning message before continuing the resync operation.

Can the cluster that is used for Infinite Volume disaster recovery (DR) use a system model that is different from the primary (that is, can FAS3270 be used on either source or destination and FAS6280 on the other)?

Answer: Yes, starting in clustered Data ONTAP 8.2 the system model used in the DR site can be different than the system model used on the source. The requirement for this is that the destination systems be able to create the same size constituents that are on the source. The maximum supported constituent size is equal to the maximum supported FlexVol volume size of all the system models used in the Infinite Volume. If necessary you can limit the size of the constituents with the `-maximum-constituent-size` option.

Note: Constituent sizes can be limited only upon creation; they can never be shrunk after they are created. If your source Infinite Volume contains constituents larger than what the destination supports, the replication will fail.

Can the namespace mirror constituent be a load-sharing mirror?

Answer: Namespace mirror constituent is a type of data protection mirror; it is not used for load balancing.

Can the Infinite Volume be configured in a fan-out or bidirectional SnapMirror configuration?

Answer: Yes, starting in Data ONTAP 8.2, Infinite Volume supports both a fan-out and a bidirectional SnapMirror configuration.

How does the SnapMirror function in a FlexVol volume differ from that in an Infinite Volume?

Answer: Table 3 compares the differences between the SnapMirror function in a FlexVol volume and an Infinite Volume.

Table 3) SnapMirror functionality comparison of FlexVol volume and Infinite Volume.

SnapMirror Functionality	FlexVol Volume		Infinite Volume	
	Data ONTAP 8.1	Data ONTAP 8.2	Data ONTAP 8.1	Data ONTAP 8.2
abort (destination)	Yes		Yes	
abort (source)	Yes		No	Yes
quiesce/resume (destination only)	Yes		No	Yes
delete/release (source only)	Yes		No	Yes
break/create (destination only)	Yes		Yes	
restore (destination only)	Yes		N/A	
show	Yes		Yes	
update/initialize/resync/modify	Yes		Yes	
check	Yes		No	
SMB1.0 read-only access for destination volume	Yes		No	Only Snapshot access
NFSv3 read-only access for destination volume	Yes		Yes	
NFSv4.1 read-only access for destination volume	Yes		No	Only Snapshot access
locking	Yes		No (only on Snapshot)	
RPO/RTO < 1hr	Yes		No	
fan out (one to several)	Yes		No	Yes
bidirectional mirroring	Yes		No	Yes
cascading	Yes		No	
load share mirrors	Yes		No	
SnapVault	No	Yes	No	
System Manager integration	Yes		No	

Does Infinite Volume support QoS parameters and transfer priorities?

Answer: No, there is no support for QoS with Infinite Volume SnapMirror.

7.3 Tape Backup

What type of tape backup solution is applicable for Infinite Volume?

Answer: Starting in Data ONTAP 8.2, tape backup for Infinite Volume can be performed by using external backup applications such as CommVault Simpana and Symantec™ NetBackup™ over CIFS or NFS mounts. You can also use NetApp SnapDiff to perform differential tape backups. SnapDiff is an application programming interface that identifies new, changed, and deleted files in an Infinite Volume for data management applications to incrementally back up the data to tape.

What type of tape backup solution is applicable for Infinite Volume?

Answer: Tape backup can only be performed at the Infinite Volume level, not the storage class or constituent level.

Is NDMP backup or SMTape backup supported for Infinite Volume?

Answer: Infinite Volume does not natively support NDMP for data backup. NetApp supports off-box third-party backup solutions that mount the Infinite Volume over SMB1.0 or NFS for backup and restore.

Both the primary and secondary (SnapMirror) clusters can be backed up. The restore operation can be performed only on the primary or secondary volume that has been promoted to a read/write volume.

8 Storage Efficiency

Does Infinite Volume support compression and deduplication?

Answer: Yes, Infinite Volume supports both of these technologies. However, storage efficiency is enabled only on the data constituents and not on the namespace constituent because the redirector files in the namespace constituent are unique individually. Therefore, employing storage efficiency technologies does not yield much savings.

Note: Deduplication is done in the Infinite Volume data constituents and not across the Infinite Volume. For example, if identical blocks are stored twice in the Infinite Volume on two different data constituents, those duplicate blocks between the data constituents are not deduplicated. If there are duplicated blocks within a data constituent, those are deduplicated.

Can I monitor progress of the deduplication or compression for the Infinite Volume?

Answer: Yes. However, this is done at the data constituent level, not the Infinite Volume level. To view the progress or status of postprocess compression or deduplication operations of individual data constituents, run the `volume efficiency show -is-constituent true` command (notice that you are not using the `-volume` option).

Can I modify volume efficiency policy for when deduplication and compression run on the Infinite Volume?

Answer: Yes, the volume efficiency policy is assigned at the Infinite Volume level if you are not using storage classes. Otherwise, it is assigned at the storage class level. You can modify the volume efficiency policy options with the `volume efficiency modify` command.

Can I set different deduplication and compression schedules or policies for individual data constituents in the Infinite Volume?

Answer: In admin mode, only a single policy or schedule can be modified for the entire Infinite Volume. Although it is possible to change the volume efficiency policy on individual data constituents while in diag mode NetApp does not recommend it.

Does Infinite Volume support thin provisioning?

Answer: Yes. Infinite Volume supports thin provisioning only on the individual data constituents in the Infinite Volume and not on the namespace constituent. It's possible to create an Infinite Volume that exceeds the size of actual aggregate space available. For example, you can create an Infinite Volume of 100TB while the actual space available from the underlying aggregates is only 80TB. If you have an Infinite Volume with storage classes, then space-guarantee will be configured at the storage class level instead of the Infinite Volume level and will be assigned when the storage class is created; it can't be changed after it is created except in diag mode.

Can I use the Space Savings Estimation Tool (SSET) to estimate compression and deduplication savings?

Answer: SSET will provide accurate compression-only savings since compression is done within a file. Deduplication savings might not be valid since deduplication is done within a data constituent on a 4KB basis and files are distributed across data constituents.

9 Performance Characteristics

Does the read/write performance increase if more controllers are added?

Answer: It does for the new data being ingested. As more space is made available, newly ingested data files can be spread across more aggregates and nodes. Therefore, the read/write throughput for any new files can scale when the number of controllers is increased; however, for Infinite Volumes this scaling is not completely linear.

Does the read/write performance increase if Flash Cache or Flash Pool is used?

Answer: When determining if Infinite Volume will benefit from Flash Cache or Flash Pool, use the same tools as those used with standard FlexVol volumes to determine what the benefit will be with Infinite Volume constituents.

Can Flash Cache or Flash Pool improve performance of namespace constituent activity?

Answer: Yes, if your namespace activity is mainly repeated directory lookups to many different files, then adding cache can benefit namespace constituent performance. It will not have any benefit for new file writes, rename operations, or directory lookups not already in cache.

Can SAS disk drives improve performance of constituent activity?

Answer: Yes, like with FlexVol volume access, reading data off SAS disks can be faster than off SATA disks. When determining if Infinite Volume data constituents will benefit from SAS disks, use the same recommendations as those used with standard FlexVol volumes to determine what the benefit will be for Infinite Volume data constituents.

If your namespace activity is mainly done in memory such as directory lookups that do not exceed the working set size, then using SAS disks for your namespace constituent will have no performance improvement.

What is the default setting for access time update?

Answer: By default this is set to true. When determining if you should set this to false, use the same recommendations as you would for FlexVol volumes. If you determine you want to disable access time updates, you can set it at the Infinite Volume level using advanced mode. For example, volume modify -vserver vs1v1 -atime-update false -volume InfiniteVol1.

Does Infinite Volume support QoS policies?

Answer: No, there is no support for QoS with the Infinite Volume or Infinite Volume SnapMirror.

Are there performance benefits when using pNFS?

Answer: Yes. When accessing files indirectly within a cluster there is a variable effect on throughput that is related to the physics of putting data over the cluster network. pNFS can help to reduce the amount of indirect access. It does this by returning metadata details, including which node contains the file, back to the client. The client caches this information and, as a result, any subsequent requests the client makes for the same file will be made directly to the node containing the file. This can reduce the amount of data that needs to be read indirectly.

What are the performance characteristics of Infinite Volume in clustered Data ONTAP?

Answer: Because Infinite Volume is based on clustered Data ONTAP, it can apply a scale-out strategy by adding more nodes or storage capacity. However, for Infinite Volumes this scaling is not completely linear. The same scale-out strategy cannot be applied to the namespace volume. An Infinite Volume contains a single namespace constituent, which has performance implications for new create and lookup requests. However, the caching of the file handles and the redirectors on the local nodes help mitigate some of the performance impact, because the subsequent requests for any existing files do not go through the namespace constituent.

The overall strategy is to target >100KB I/O workloads to amortize the processing and accessing time of the namespace constituent to where they are negligible factors in the overall performance of the Infinite Volume.

Note: Workloads that do not target large I/O can aggravate these performance characteristics.

Testing has shown Infinite Volume's performance are as follows:

- **I/O performance.** Large I/O workloads are the target profile for Infinite Volume. In this profile, for example, by using files with 300MB size, which are sufficiently large files, the read and write performance can approach approximately 90% of a FlexVol volume's performance.

Note: This is an approximation of performance and depends on the number of concurrent activities and other factors that might prevent the client from achieving these rates.

- **Node performance.** The node that contains the namespace volume is generally more processing intensive because every new read, write, and all directory activity must be accessed through a single namespace constituent.
- **Aggregated performance.** Infinite Volume performance scales with the addition of cluster nodes, but at a less-than-linear rate. For Infinite Volume workloads, the preferred number of nodes is influenced more by capacity considerations than by performance.

If SAS disk drives are used for the namespace constituent, is there any difference in the recommended configuration for the namespace constituent?

Answer: No. To reduce the possibility of the namespace constituent becoming a bottleneck, NetApp recommends creating the namespace constituent by using the recommended configuration regardless of the type of disk being used. See the [disk configuration section](#) for details on the recommended configuration.

10 Feature Interoperability

What features does Infinite Volume support starting with Data ONTAP 8.2?

Answer: The following features are supported with Infinite Volume; the list is not all inclusive.

Table 4) Features supported with Infinite Volume.

Feature	Infinite Volume Support in Data ONTAP 8.1.x	Infinite Volume Support in Data ONTAP 8.2
Protocol	NFSv3	NFSv3, NFSv4.1, pNFS, and SMB1.0
Security style	UNIX only	Unified security style (ACL management from Windows/UNIX platforms)
Tenancy	One SVM per cluster dedicated to a single Infinite Volume	Multi-tenancy—SVMs for Infinite Volume and SVMs for FlexVol volumes can share the same cluster
Platforms	• FAS/V6280 and FAS/V3270 • All models in the cluster must be identical models • 10 nodes maximum	• All systems that support clustered Data ONTAP 8.2 with the exception of the 2000 series • Can mix system models within an Infinite Volume • Each Infinite Volume can span a maximum of 10 nodes; cluster can contain more than 10 nodes
Disk type	SATA or SAS	SATA, SAS, SSD
Disk layout	No sharing of aggregates, each aggregate is dedicated to a single constituent	Aggregates do not have to be dedicated—they can share with numerous constituents from numerous Infinite Volumes as well as FlexVol volumes
Load balancing	Files are distributed and balanced across aggregates at ingest time	Files are distributed equally across nodes at ingest times
Vol move	Disruptive, requires CSS support	Nondisruptive, requires CSS support
SMB1.0 restricted change notify	NA, no support for SMB	Yes*
Cache support	Flash Cache	Flash Cache and Flash Pool
SnapMirror support	Limited source-side commands	Full SnapMirror support (command completion, vserverization, transfer priorities)
Tape backup	NFSv3 mounted external backup applications	NFS or SMB mounted external backup applications and differential tape backup using NetApp SnapDiff

*SMB1.0 restricted change notify is off by default for Infinite Volume. If it is enabled, the primary effect is that Windows Explorer will automatically update when a file is created, deleted, or renamed, but the user has to manually refresh to see the size change if he or she modified a file. The exception is ACL changes, for which Infinite Volume does generate a change notification.

What features does Infinite Volume not support in Data ONTAP 8.2?

Answer: Some of the features not supported with Infinite Volume are included in Table 6.

Table 5) Features not supported with Infinite Volume in Data ONTAP 8.2.

- SMB2.x, SMB 3.0, and NFSv4.0
- SMB/NFSv4.1 referrals
- SMB1.0 and NFSv4.1 on active file system of read-only volumes
- NFS 4.1 delegations
- LUNs
- QoS
- NDMP
- Qtrees
- Quota
- Home directories
- Fpolicy
- SFMod
- Watches or change notifies
- FlexCache
- Autosnapshot delete
- Autogrow
- Single file snapshot restore
- Cascading
- FAS2000 series
- SMB tools like sectrace, fsecurity
- SMTape
- SyncMirror
- MetroCluster™ technology
- SnapVault® technology

11 Specifics of Infinite Volume in Data ONTAP 8.1.x

Is Infinite Volume generally available in Data ONTAP 8.1.x?

Answer: Infinite Volume is available starting with clustered Data ONTAP 8.1.1 under a product variance request (PVR) and is generally available (GA) with clustered Data ONTAP 8.2. Infinite Volume cannot be configured on a 7-Mode system.

Does the Infinite Volume feature require a license in Data ONTAP 8.1.x?

Answer: No, Infinite Volume does not require a separate license to be installed, but in Data ONTAP 8.1.x a special hash key is required to configure it. This key is available through a PVR for Data ONTAP 8.1.x.

What is the reason for putting a PVR in the sales process for Infinite Volume in Data ONTAP 8.1.x?

Answer: Infinite Volume in Data ONTAP 8.1.x is a complete feature, but a PVR is included in the sales process to closely monitor the sales opportunities for Infinite Volume and to provide approval for appropriate use cases.

Why is a PVR key required in Data ONTAP 8.1.x?

Answer: A PVR hash key is required to enable creation of an SVM for Infinite Volume. Without the PVR hash key, the SVM for Infinite Volume cannot be created.

How can I get a PVR key for Infinite Volume in Data ONTAP 8.1.x?

Answer: The sales team will file a PVR after the CMET team qualifies the sales opportunity for Infinite Volume. The product manager will provide the PVR hash key after approval from Customer Success Services.

How is a PVR key used for Infinite Volume in Data ONTAP 8.1.x?

Answer: A PVR hash key is associated with the serial number of each storage controller node. The storage administrator needs to set the boot arguments at the `LOADER` prompt for each node in a cluster by using the PVR hash key.

If a storage controller is replaced, how can I get a new PVR key for Infinite Volume in Data ONTAP 8.1.x?

Answer: The sales team provides the product manager with a new serial number for the controller. The product manager in turn provides the customer with a new PVR hash key.

How many Infinite Volumes can I have in a cluster in Data ONTAP 8.1.x?

Answer: In Data ONTAP 8.1.x, you can only have one Infinite Volume in a cluster, and the cluster cannot contain any other SVMs.

Are heterogeneous nodes supported in an Infinite Volume (that is, can FAS3270 and FAS6280 be mixed)?

Answer: No, in clustered Data ONTAP 8.1.x, heterogeneous nodes are not supported in an Infinite Volume. In other words, all nodes need to be the same model.

Which NetApp platforms does Infinite Volume support in Data ONTAP 8.1.x?

Answer: In clustered Data ONTAP 8.1.x, Infinite Volume is only supported on FAS6280 and FAS3270 controllers with DS4243 and DS4486 storage shelves.

Which protocols does Infinite Volume support in Data ONTAP 8.1.x?

Answer: In clustered Data ONTAP 8.1.x, Infinite Volume is only supported with NFSv3.

What security type does Infinite Volume support in Data ONTAP 8.1.x?

Answer: In Data ONTAP 8.1.x, Infinite Volume only supports UNIX-style security and UNIX-mode bit permissions.

Can an Infinite Volume and a FlexVol volume coexist in a cluster in Data ONTAP 8.1.x?

Answer: No, in Data ONTAP 8.1.x this is not available. It only supports one Infinite Volume for the entire cluster.

Can I use System Manager to manage Infinite Volume in Data ONTAP 8.1.x?

Answer: In Data ONTAP 8.1.x, you are able to configure Infinite Volume using System Manager 2.1. However, this functionality was removed in Data ONTAP 8.2. For instructions on configuring Infinite Volume in Data ONTAP 8.1.x using System Manager, see TR-4089i.

How do I size an Infinite Volume in Data ONTAP 8.1.x?

Answer: As part of the feature PVR process for Infinite Volume we provide assistance with sizing recommendations. At a high level some of the considerations include:

- 500MB/s per node for FAS6280
- Space requirements:
 - o Usable capacity requirements
 - o + 10TB for dedicated aggregate for namespace constituent
 - o + 10TB for dedicated aggregate for namespace mirror constituent
 - o + 8% aggregate reservation space on each aggregate
 - o + 10% WAFL reserve on each constituent
 - o + 5% snapshot reserve on each constituent
 - o Disk right-sizing on each disk
 - o Optimal aggregate configurations as detailed in the disk configuration section that follows

11.1 Disk Configuration Recommendations for Data ONTAP 8.1.x

Before I create an Infinite Volume is there anything I should do with my disks in Data ONTAP 8.1.x?

Answer: Yes. NetApp recommends creating aggregates prior to creating an Infinite Volume. In Data ONTAP 8.1.x, Infinite Volume required all nodes in the cluster to be dedicated to the Infinite Volume.

Can you use different disk drive types for the namespace and the data constituents within an Infinite Volume in Data ONTAP 8.1.x?

Answer: No, in clustered Data ONTAP 8.1.x Infinite Volume does not support mixing of disk drives. That includes using different disks such as SAS disks for the namespace and SATA disks for the data constituents; in other words, all disks must either be SATA disks or SAS disks.

Can two different aggregates, one with SAS disks and another with SATA disks, be added to an Infinite Volume in Data ONTAP 8.1.x?

Answer: No, in Data ONTAP 8.1.x, Infinite Volume does not support the mixing of disk drive types.

What is the recommended configuration for an Infinite Volume in Data ONTAP 8.1.x?

Answer: In clustered Data ONTAP 8.1.x, Infinite Volume requires one constituent per aggregate. (This is not the case for Data ONTAP 8.2.) Table 7 provides recommendations to enable the best usable capacity, performance, and resiliency.

Table 7) Recommended Infinite Volume configurations for Data ONTAP 8.1.x.

Recommendation	FAS6280 with DS424x	FAS6280 with DS4486	FAS3270 with DS424x	FAS3270 with DS4486
RAID configuration for aggregates that contain namespace or data constituents	51 3TB disks configured in 3 RAID-DP groups of 17 disks (15 data + 2 parity)	40 3TB disks configured in 2 RAID-DP groups of 20 disks (18 data + 2 parity)	20 3TB disks configured in 1 RAID-DP group (18 data + 2 parity)	
Maximum number of aggregates per HA pair	28	35	47	
Recommended disk configuration per HA pair	1,428 disks for namespace and data constituents 6 disks for node root aggregates 6 spares 0 left over	1,400 disks for namespace and data constituents 6 disks for node root aggregates 12 spares 22 left over	940 disks for namespace and data constituents 6 disks for node root aggregates 12 spares 2 left over	940 disks for namespace and data constituents 6 disks for node root aggregates 6 spares 8 left over
Aggregate Snapshot reserve	0%			
Volume Snapshot reserve	5%			
Estimated constituent usable capacity given the configuration listed above	~93TB	~82TB	~38TB	~42TB
Maximum number of aggregates for Infinite Volume	1 namespace 1 namespace mirror 138 data constituents	1 namespace 1 namespace mirror 173 data constituents	1 namespace 1 namespace mirror 233 data constituents	
Maximum usable storage for Infinite Volume	~13PB	~14PB	~9PB	~10PB
Maximum raw storage for Infinite Volume	~20PB	~20PB	~13PB	

Why shouldn't I create an aggregate larger than 100TB for FAS6280 or 70TB for FAS3270 for an Infinite Volume in Data ONTAP 8.1.x?

Answer: In Data ONTAP 8.1.x, only one constituent per aggregate is supported, and each individual data constituent cannot exceed the maximum capacity of a FlexVol volume. Given these limitations, creating a larger aggregate will result in wasted space. In clustered Data ONTAP 8.2, Infinite Volume does not require dedicated aggregates so you can create larger aggregates without wasting space.

What size aggregate should I create if the recommended size listed in Table 7 is not possible for an Infinite Volume in Data ONTAP 8.1.x?

Answer: In Data ONTAP 8.1.x, to achieve the maximum Infinite Volume size with the best performance and resiliency in Data ONTAP 8.1.x, NetApp recommends that the aggregate used for the namespace and data constituents be configured as follows.

- **FAS6280 with DS424x:** 51 3TB disks configured in 3 RAID-DP groups of 17 disks (15 data + 2 parity)
- **FAS6280 with DS4486:** 40 3TB disks configured in 2 RAID-DP groups of 20 disks (18 data + 2 parity)
- **FAS3270:** 20 3TB disks configured in 1 RAID-DP group (18 data + 2 parity)

Note: By default, the System Manager Vserver with Infinite Volume Wizard creates aggregates with a fixed RAID group size of 16. To create a different RAID group size—such as the recommended maximum configuration, which includes 20-drive or 17-drive RAID groups—aggregates should first be created by using the standard cluster administration aggregate create tool. After that, the Create Vserver with Infinite Volume Wizard uses all the available aggregates to create the Infinite Volume spread across all of these aggregates.

If this configuration is not possible, NetApp recommends creating an aggregate that is 10% larger than the desired constituent size, up to the maximum FlexVol size supported on the platform. For the FAS6280, the maximum aggregate size for a constituent should be 100TB; for the FAS3270, it should be 70TB. The reason is that in Data ONTAP 8.1.x, only one constituent per aggregate is supported, and each individual data constituent cannot exceed the maximum capacity of a FlexVol volume. Given these limitations, creating a larger aggregate results in wasted space. For example, if the desired constituent size is 50TB for optimal space usage, create an aggregate no larger than 55TB.

To see the size of your aggregates, look at the following:

- System Manager, Configure Aggregates portion of the wizard, Available Space value
- System Manager Aggregates tab, Total Space value for the aggregate
- Command line, the value shown for Size from the output of the `aggr show` command

In clustered Data ONTAP 8.2, Infinite Volume does not require dedicated aggregates so there are no recommendations for disk configurations different than those for FlexVol volumes.

What is the maximum size aggregate recommended for an Infinite Volume in Data ONTAP 8.1.x?

Answer: For Data ONTAP 8.1.x, for a FAS/V6280, NetApp recommends not creating an aggregate larger than 100TB. For Data ONTAP 8.1.x, for a FAS/V3270, NetApp recommends not creating an aggregate larger than 70TB. Any space larger than this is wasted because it is not possible to create a constituent larger than the maximum supported FlexVol size, and only one constituent per aggregate is supported in Data ONTAP 8.1.x.

In clustered Data ONTAP 8.2, Infinite Volume does not require dedicated aggregates, so you can create larger aggregates without wasting space.

Should I create several smaller aggregates or fewer larger aggregates for an Infinite Volume in Data ONTAP 8.1.x?

Answer: It is better to create aggregates at least as large as the maximum supported data constituent size +10% in your Infinite Volume so that you can create a few larger data constituents. In Data ONTAP 8.1.x, only one constituent per aggregate is supported. For a FAS/V6280, NetApp recommends creating aggregates no larger than 100TB each. For a FAS/V3270, NetApp recommends creating aggregates no larger than 70TB each. In Data ONTAP 8.2, there is no recommendation to limit aggregates to this size.

Can I specify which aggregate or node to use for the namespace constituent in Data ONTAP 8.1.x?

Answer: No, in Data ONTAP 8.1.x the namespace constituent will be automatically placed on the largest aggregate. If you want to use a specific aggregate you have to verify that it is the largest aggregate during Infinite Volume creation. Then, after the Infinite Volume is created, add disks to the aggregates used for the data constituents. Another option is to work with CSS on diag mode options to specify the aggregate for the namespace constituent.

What is the recommended configuration for the namespace constituent in Data ONTAP 8.1.x?

Answer: For optimal usable capacity and performance, NetApp recommends the following for the namespace constituent.

- Verify usable space of at least 11TB
- Flash Cache might help for namespace activities such as directory lookups
- Clusters containing 2–6 nodes:
 - o Minimum of 8 disks configured with RAID-DP (verify that usable space is at least 11TB)
- Clusters containing 8–10 nodes:
 - o Minimum of 12 disks configured with RAID-DP (verify that usable space is at least 11TB)
 - o Node dedicated to namespace constituent only (no client accessible LIFs, no data constituents); the dedicated node is to dedicate as much CPU as possible to the namespace constituent activities

Note: In Data ONTAP 8.1.x, by default, the System Manager Vserver with Infinite Volume Wizard creates aggregates with a fixed RAID group size of 16. To create a different RAID group size, aggregates should first be created by using the standard cluster administration aggregate create tool. After that, the Create Vserver with Infinite Volume Wizard uses all the available aggregates to create the Infinite Volume spread across all of these aggregates.

11.2 Infinite Volume Setup in Data ONTAP 8.1.x

What management user interfaces are available for Infinite Volume in Data ONTAP 8.1.x?

Answer: Infinite Volume in Data ONTAP 8.1.x can be managed by System Manager 2.1 or the Data ONTAP CLI. The admin CLI follows the Infinite Volume–specific command hierarchy, which is modeled closely after the standard Data ONTAP hierarchy. However, the constituent level of information can be viewed but not managed in admin mode. Constituent-level management is done only with assistance from Customer Success Services.

How do I select which nodes in my cluster to use for Infinite Volume in Data ONTAP 8.1.x?

Answer: In Data ONTAP 8.1.x, Infinite Volume requires all nodes in the cluster to be dedicated to the Infinite Volume.

How do I create an Infinite Volume in Data ONTAP 8.1.x?

Answer: The Infinite Volume is a compound volume that includes FlexVol volume-like containers that serve as a namespace constituent, namespace mirror, and several data constituents. In Data ONTAP 8.1.x, after the cluster is created, in addition to the aggregate used for the SVM root a minimum of four additional aggregates is required to create an Infinite Volume. One aggregate (the largest) is used for the namespace constituent, one (the smallest) for the SVM root, and at least two aggregates are needed to host data constituents. Once the aggregates are created, the SVM for Infinite Volume must be enabled with a PVR hash key, as described previously. Using System Manager 2.1, create the SVM for Infinite Volume and then create an Infinite Volume; the system automatically sets up the namespace and data constituents. NetApp recommends creating an additional aggregate after the Infinite Volume is created and using that to create a namespace mirror.

Is there any difference in configuring Infinite Volume with FAS3270 and FAS6280 with Data ONTAP 8.1.x?

Answer: In Data ONTAP 8.1.x, when creating the Infinite Volume on a FAS3270, you must manually verify that the maximum Infinite Volume that you create does not exceed $((62.9\text{TB} * (\text{<\# aggregates in the Infinite Volume> - 1})) + 10\text{TB})$. You can modify this value in the Create Vserver with Infinite Volume Wizard on the Create Infinite Volume page by changing the value of Total Size. Not manually modifying this value can lead to an error when trying to create the Infinite Volume. See TR-4089i for details.

Also, in Data ONTAP 8.1.x, when resizing an Infinite Volume on FAS3270 systems, the maximum size is calculated incorrectly in the tool. Do not set this value larger than:

- $((62.9\text{TB} * (\text{<\# aggregates in the Infinite Volume> - 2})) + 10\text{TB})$ if you have a namespace mirror
- $((62.9\text{TB} * (\text{<\# aggregates in the Infinite Volume> - 1})) + 10\text{TB})$ if you don't have a namespace mirror for a FAS3270 configuration

In Data ONTAP 8.1.x, attempting to increase the Infinite Volume size on the FAS3270 beyond the maximum causes the resize to fail. The maximum size is calculated correctly for the FAS6280.

Note: This is fixed in Data ONTAP 8.1.2.

What are the SnapMirror requirements for Infinite Volume in Data ONTAP 8.1.x?

Answer: SnapMirror is based on Snapshot; it follows some major rules of creation and management (update, resync, break, reverse resync, and so on).

- In Data ONTAP 8.1.x, the SnapMirror destination must have an equal number and size of aggregates assigned to the SVM at both ends.
- The SnapMirror destination Infinite Volume space should be equal to or greater than the source.
- SnapMirror requires an SVM with Infinite Volume at both the source and the destination.
- In Data ONTAP 8.1.x, the SnapMirror destination must be the same system model as the source. In other words, if the source Infinite Volume resides on a FAS/V6280, then the SnapMirror destination must also be a FAS/V6280.

Note: SnapMirror cascading is not supported for the Infinite Volume.

Can the Infinite Volume be configured in a fan-out or bidirectional SnapMirror configuration in Data ONTAP 8.1.x?

Answer: No, in Data ONTAP 8.1.x, Infinite Volume does not support fan-out or bidirectional SnapMirror configurations.

Can the cluster that is used for Infinite Volume disaster recovery (DR) use a different model of systems from the primary in Data ONTAP 8.1.x?

Answer: No, an Infinite Volume running in clustered Data ONTAP 8.1.x does not support different model platforms on the DR source and DR destinations. That means that if using FAS3270 on the source the destination must also contain FAS3270 systems; it can't contain FAS6280 systems.

What type of tape backup solution is applicable for Infinite Volume in Data ONTAP 8.1.x?

Answer: In Data ONTAP 8.1.x, tape backup for Infinite Volume can be done by using external backup applications such as CommVault Simpana over the NFSv3 mounts.

11.3 Capacity Balancing in Data ONTAP 8.1.x

Do the aggregates of the Infinite Volume have to be of equal size in Data ONTAP 8.1.x?

Answer: No, the aggregates of the Infinite Volume don't have to be equal in size; however, for capacity balancing to work well for larger files (relative to the aggregate sizes) and for getting balanced throughput across all the aggregates, they should have similar size and layout.

Here are some recommendations for configuring aggregates to achieve the maximum performance, resiliency, and Infinite Volume size in Data ONTAP 8.1.x.

- For FAS/V3270, use 1 RAID-DP group of 20 disks.
- For FAS/V6280 with DS424x, use 3 RAID-DP groups of 17 disks each.
- For FAS/V6280 with DS4486, use 2 RAID-DP groups of 20 disks each.
- By default, the largest aggregate is selected for the namespace constituent.
- By default, the smallest aggregate is selected for the SVM root aggregate.

In Data ONTAP 8.1.x, by default, the System Manager Vserver with Infinite Volume Wizard creates aggregates with a fixed RAID group size of 16. To create a different RAID group size—such as the recommended maximum configuration, which includes 20-drive or 17-drive RAID groups—aggregates should first be created by using the standard cluster administration aggregate create tool. After that, the Create Vserver with Infinite Volume Wizard uses all the available aggregates to create the Infinite Volume spread across all of these aggregates.

In Data ONTAP 8.1.x, can I increase the size of an Infinite Volume by increasing the size of the aggregates beyond 100TB for FAS6280 and 70TB for FAS3270?

Answer: No, in Data ONTAP 8.1.x only one constituent per aggregate is supported, and each individual data constituent cannot exceed the maximum capacity of a FlexVol volume. Given these limitations, creating a larger aggregate would result in wasted space. Starting in Data ONTAP 8.2, Infinite Volume no longer limits aggregate usage to a single constituent.

11.4 Space Management in Data ONTAP 8.1.x

Can the size of an Infinite Volume be increased in Data ONTAP 8.1.x?

Answer: Yes, an Infinite Volume can be increased in size without disrupting client access. Infinite Volume capacity can be increased with these steps:

- Increase the capacity of the Infinite Volume by using one of these three methods:
 - 1) If the constituents are not at maximum size and there is room in the aggregates simply resize the Infinite Volume.
 - 2) Add disks to an existing aggregate assigned to the Infinite Volume and resize the Infinite Volume.
 - 3) Assign additional aggregates either on an existing node or a new node with the new set of disks and resize the Infinite Volume.

While the capacity is being increased, users can continue to add new files with no downtime or interruption.

Note: In Data ONTAP 8.1.x, when resizing an Infinite Volume on FAS3270 systems, the maximum size is calculated incorrectly in the tool. Do not set this value to larger than:

- With a namespace mirror: $((62.9\text{TB} * (<\# \text{ aggregates in the InfiniteVol} > - 2)) + 10\text{TB})$
- No namespace mirror: $((62.9\text{TB} * (<\# \text{ aggregates in the InfiniteVol} > - 1)) + 10\text{TB})$

Attempting to increase the Infinite Volume size on the FAS3270 beyond the maximum causes the resize to fail. The maximum size is calculated correctly for the FAS6280. This is fixed in Data ONTAP 8.2.

12 Technical Support

How do I get Infinite Volume support?

Answer: Customer Success Services handles all post sales customer support issues. For presales technical queries, send an e-mail to dl-cmet.

13 Terminology

Table 8) Terminology.

Term	Definition
NetApp Infinite Volume	A large compound volume that includes a number of constituents
Constituent	A data or namespace container that is a member of a NetApp Infinite Volume
Data constituent	A constituent that contains the data for one or more regular files; in other words, data files reside in this constituent
Namespace constituent	A constituent that contains the directory hierarchy and file names; the file names are associated with redirector (pointer) files, each of which refers to a data file in a data constituent
Namespace redirector	The module that implements the redirection between the user-visible path names and their corresponding data files in data constituents
Redirector file	A regular WAFL leaf file in the namespace constituent that contains the mapping of a file name to a data file
Redirector file handle	The file handle of the redirector file

References

The following references were used in this FAQ.

- TR-4037: Introduction to NetApp Infinite Volume
<http://www.netapp.com/us/media/tr-4037.pdf>
- Data ONTAP 8.2 Logical Storage Management Guide
https://library.netapp.com/ecm/ecm_get_file/ECMP1155398
- Data ONTAP 8.2 Cluster-Mode Data Protection Guide
https://library.netapp.com/ecm/ecm_get_file/ECMP1155389
- Data ONTAP 8.1.1 Cluster-Mode Documentation
<http://support.netapp.com/documentation/docweb/index.html?productID=61476>

Version History

Version	Date	Document Version History
Version 1.0	September 2012	Initial release
Version 2.0	April 2013	Updated to include clustered Data ONTAP 8.2
Version 2.1	August 2013	Updated to include unified security, storage classes and data policies
Version 2.2	November 2013	Updated to include performance sizing

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.

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

[Go further, faster®](#)



www.netapp.com

© 2013 NetApp, Inc. All rights reserved. No portions of this document may be reproduced without prior written consent of NetApp, Inc. Specifications are subject to change without notice. NetApp, the NetApp logo, Go further, faster, Data ONTAP, Flash Cache, Flash Pool, FlexCache, FlexClone, FlexVol, Manage ONTAP, MetroCluster, OnCommand, RAID-DP, SnapMirror, SnapRestore, Snapshot, SnapVault, SyncMirror, and WAFL are trademarks or registered trademarks of NetApp, Inc. in the United States and/or other countries. UNIX is a registered trademark of The Open Group. Windows is a registered trademark of Microsoft Corporation. Symantec and NetBackup are trademarks of Symantec Corporation. All other brands or products are trademarks or registered trademarks of their respective holders and should be treated as such.