



Technical Report

## Introduction to NetApp Infinite Volume

Sandra Moulton, Reena Gupta, NetApp  
April 2013 | TR-4037

### Summary

This document provides an overview of NetApp<sup>®</sup> Infinite Volume, a new innovation in scalable NetApp storage systems introduced in the clustered Data ONTAP<sup>®</sup> 8.1.1 architecture. Scalability allows customers to nondisruptively grow their storage as needed in order to address enormous business growth. This document also discusses the advantages of Infinite Volume and its features and provides a high-level overview of its architecture.

## TABLE OF CONTENTS

<b>1</b>	<b>Introduction</b> .....	<b>3</b>
1.1	Business Challenge .....	3
1.2	Solution .....	3
<b>2</b>	<b>Advantages of Infinite Volume</b> .....	<b>4</b>
<b>3</b>	<b>Overview of Infinite Volume</b> .....	<b>4</b>
<b>4</b>	<b>Target Workloads and Use Cases</b> .....	<b>5</b>
<b>5</b>	<b>Specifications and Supported Configuration</b> .....	<b>6</b>
5.1	Supported Configuration for Infinite Volume .....	6
<b>6</b>	<b>Features of Infinite Volume</b> .....	<b>6</b>
6.1	Manageability.....	6
6.2	Data Protection .....	7
6.3	Native Storage Efficiency .....	8
6.4	Performance Acceleration with Flash Cache .....	8
6.5	Performance Acceleration with Flash Pool.....	8
6.6	Resiliency.....	9
6.7	Unified Security Style.....	9
<b>7</b>	<b>Conclusion</b> .....	<b>9</b>
	<b>Reference</b> .....	<b>9</b>
	<b>Version History</b> .....	<b>9</b>

## LIST OF TABLES

Table 1)	Specifications for Infinite Volume. ....	6
----------	--	---

## LIST OF FIGURES

Figure 1)	Infinite Volume architecture.....	5
-----------	-----------------------------------	---

# 1 Introduction

Current market trends for data storage indicate that file-based storage is the fastest growing market segment in the storage industry and is one of the most influential drivers of storage industry development and market expansion. This growth in file-based storage brings with it scale requirements that make it increasingly challenging for administrators to manage their data. *Big data* is a term applied to datasets whose size is beyond the ability of commonly used software tools to capture, manage, and process data within a tolerable elapsed time.

Big data sizes are a constantly moving target, ranging from a few dozen terabytes to many petabytes of data in a single dataset. *Big content* applies to the category of big data workloads that address the challenges associated with data management for high capacity and file counts. At the extremes, big content includes petabyte-scale file systems, billions of objects, and geographically distributed data. Traditional network-attached storage (NAS) systems are evolving into large enterprise content repositories to meet the increasing demand to manage the unprecedented growth of unstructured data.

## 1.1 Business Challenge

An explosion of unstructured content growth in the market for active deep archives and content repositories creates new challenges in the storage and retrieval of large volumes of data. Typical storage footprints for applications are growing at a 50% to 90% compound annual growth rate and have unique retention periods based on certain types of metadata, such as business unit, personal data, and so on.

Staying ahead of the data tide is difficult when considering traditional volume-based storage constructs for file management. Customers require scalable file and object-data repositories that eliminate the need to worry about volumes and volume replication pairs as they become too numerous to manage costs effectively.

Cost, scalability, and complexity are the prominent concerns for customers today. Customers require petabytes (PBs) of storage, billions of files, high availability, storage efficiency, and reduced operational overhead. They also need a single mountpoint with no boundaries in order to seamlessly grow storage space without disrupting their active operations.

Further, their entire enterprise requires storage for many different types of applications, each with different storage needs, different security requirements, and different access methods. Customers struggle with needing to have different hardware with different management for each different type of solution.

## 1.2 Solution

Big content storage solutions can be categorized into three main categories based on the storage, management, and retrieval of the data into file services, enterprise content repositories, and distributed content repositories. NetApp addresses the business challenges of big content by providing the appropriate solution for all of these different environments.

- File services represent the portion of the unstructured data market in which NetApp has traditionally been a leader, including project shares and home directory use cases.
- The enterprise content repository market, by contrast, is less driven by direct end users and more by applications that require large container sizes with an increasing number of files.
- Distributed content repositories take advantage of object protocols to provide a global namespace that spans numerous data centers.

Infinite Volume addresses the enterprise content repository market and is optimized for scale and ease of management. Infinite Volume is a cost-effective large container that can grow to PBs of storage and billions of files. It is built on NetApp's reliable fabric-attached storage (FAS) and V-Series systems, and it inherits the advanced capabilities of clustered Data ONTAP.

By providing a single large container for unstructured data, e-mail, video, and graphics, Infinite Volume eliminates the need to build data management capabilities into applications with big content requirements. For these environments, Infinite Volume takes advantage of native storage efficiency features, such as deduplication and compression, to keep storage costs low.

Further, since Infinite Volume is built into clustered Data ONTAP, the customer is able to host both Infinite Volume(s) and FlexVol<sup>®</sup> volumes together in a unified scale-out storage solution. This provides the customer with the ability to host a variety of different applications in a multi-tenancy environment, with nondisruptive operations and the ability to use both SAN and NAS in the same storage infrastructure leveraging the same hardware.

## 2 Advantages of Infinite Volume

Infinite Volume offers many business advantages for enterprise content repositories. For example, an Infinite Volume for an enterprise content repository solution can be used to:

- Reduce the cost of scalability
  - Lower the effective cost per GB
  - Efficiently ingest, store, and deliver large amounts of data
- Reduce complexity and management overhead
  - Simplify and automate storage management operations
  - Provide seamless operation and data and service availability

Infinite Volume leverages dense storage shelves from NetApp with the effective use of large-capacity storage disks. The solution is built on top of the proven foundation of Data ONTAP with storage efficiency features like deduplication and compression.

Infinite Volume gives customers a single, large, scalable container to help them manage huge amounts of growth in unstructured data that might be difficult to manage by using several containers. Data is automatically load balanced across the Infinite Volume at ingest. This manageability allows storage administrators to easily monitor the health state and capacity requirements of their storage systems.

Infinite Volumes are configured within a Data ONTAP cluster and do not require dedicated hardware. Infinite Volumes can share the same hardware with FlexVol volumes.

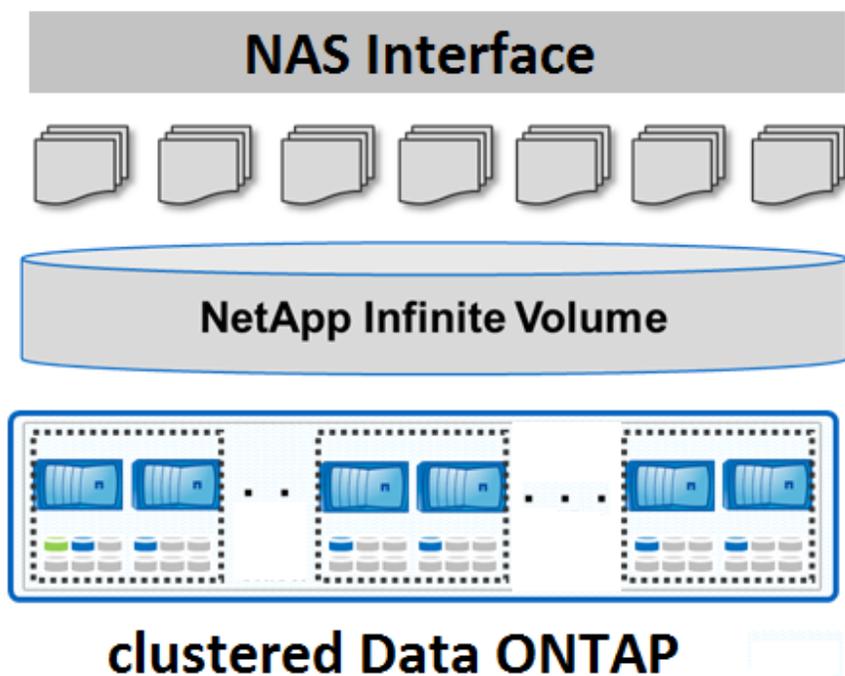
## 3 Overview of Infinite Volume

NetApp Infinite Volume is a software abstraction hosted over clustered Data ONTAP. It provides a single mountpoint that can scale to 20PB and 2 billion files, and it integrates with NetApp's proven technologies and products, such as deduplication, compression, and NetApp SnapMirror<sup>®</sup> replication technology.

Infinite Volume writes an individual file in its entirety to a single node but distributes the files across several controllers within a cluster.

Figure 1 shows how an Infinite Volume appears as a single large container with billions of files stored in numerous data constituents.

Figure 1) Infinite Volume architecture.



In the first version of Infinite Volume, data access was provided over the NFSv3 protocol. Starting in clustered Data ONTAP 8.2, Infinite Volume added support for NFSv4.1, pNFS, and CIFS. Like a FlexVol volume, Infinite Volume data is protected by using NetApp Snapshot™, RAID-DP®, and SnapMirror technologies, and NFS or CIFS mounted tape backups.

## 4 Target Workloads and Use Cases

Infinite Volume was developed to provide a scalable, cost-effective solution for big content workloads. Specifically, Infinite Volume addresses the requirements for large unstructured repositories of primary data, which are also known as enterprise content repositories.

Enterprise content repositories can be subdivided into workloads with similar access patterns, data protection requirements, protocol requirements, and performance requirements.

Infinite Volume is focused on use cases that can be characterized by input/output (I/O) patterns in which data is written once and seldom changed. However, this data is used for normal business operations, and therefore content must be kept online for fast retrieval, rather than being moved to secondary storage.

One example of this type of workload is a video file archive. Libraries of large video files are kept in a repository from which they are periodically retrieved and sent to broadcast sites. These repositories typically grow as large as 5PB.

Another example is enterprise content management storage. This can be used to store large amounts of unstructured content such as documents, graphics, and scanned images. These environments commonly can contain a million or more files.

## 5 Specifications and Supported Configuration

The capabilities for the general availability release of Infinite Volume are optimized for enterprise content repositories. Table 1 describes the specifications for Infinite Volume.

Table 1) Specifications for Infinite Volume.

Attribute	NetApp Infinite Volume
File system type	Enterprise content repository
Value propositions	Large-scale single file system with availability and storage efficiency
Best suited for	Workloads that are write once, update rarely, with an average file size >100KB
Scale	<ul style="list-style-type: none"><li>• 20PB volumes</li><li>• 2 billion files</li></ul>
Data protection	<ul style="list-style-type: none"><li>• NetApp RAID-DP</li><li>• Snapshot technology, asynchronous SnapMirror replication technology</li><li>• NFS or CIFS mounted tape backup</li></ul>
Important ecosystem need	Scalable grow-as-you-go NAS repository for enterprise content

### 5.1 Supported Configuration for Infinite Volume

Infinite Volume in clustered Data ONTAP 8.2 supports the following:

- Each Infinite Volume can span a maximum of 10 nodes in a cluster; the total number of nodes in the cluster can exceed 10
- NFSv3, NFSv4.1, pNFS, CIFS (SMB 1.0) protocols
- Maximum 20PB capacity
- Maximum 2 billion files
- FAS and V-Series 3000 and 6000 series systems that support Data ONTAP 8.2; for a current list of platforms that support Infinite Volumes, see the Hardware Universe (formerly the Systems Configuration Guide) at [support.netapp.com/knowledge/docs/hardware/NetApp/syscfg/index.shtml](http://support.netapp.com/knowledge/docs/hardware/NetApp/syscfg/index.shtml)

## 6 Features of Infinite Volume

Infinite Volume features are grouped into the following categories:

- Manageability
- Data protection
- Native storage efficiency
- Performance acceleration with Flash Cache™ and Flash Pool™ intelligent caching
- Resiliency
- Unified security style

### 6.1 Manageability

An Infinite Volume is an easily managed clustered storage solution. Infinite Volume manageability features include capacity reporting, create-time automatic capacity balancing across nodes, nondisruptive scalability, nondisruptive upgrades, simplified setup, single namespace, and single mountpoint.

## Capacity Reporting

An Infinite Volume can serve up to 2 billion files to a maximum capacity of 20PB. A single file can grow as large as 16TB. The capacity of the volume includes used disk space and available disk space. Used disk space is the amount of space used in the volume by data files. Available disk space is the amount of free space left in the volume. Administrators can designate a threshold in the volume to track the amount of remaining available disk space. When the threshold is reached, a message is sent to the storage administrator indicating the need to add more storage.

## Automatic Capacity Balancing Across Nodes

When new files are written to an Infinite Volume, they are distributed evenly across storage resources on a round-robin basis with a preference for storage resources that have the most available space. When new capacity is added to the volume, data is balanced across nodes. This is an automatic background function that helps storage administrators avoid problems created by uneven capacity balancing.

## Scalability

An Infinite Volume can scale up to 20PB. Storage capacity can be added nondisruptively to the volume. This allows storage administrators to start small and make additional investments in storage on an add-as-you-grow basis, which helps them to manage changing requirements as the business develops.

## Simplified Setup

The setup of an Infinite Volume is extremely simple. It uses virtually the same setup as is used to create a virtual storage server (VServer) with FlexVol and a FlexVol volume. The setup includes:

- Creating a Vserver for Infinite Volume
- Creating an Infinite Volume
- Setting up NFS and/or CIFS

## Single Namespace and Mountpoint

An Infinite Volume offers a single namespace. The entire volume can be mounted as a single mountpoint or share and does not require any additional changes on the application side, even when more storage is added. This reduces the management complexity associated with managing numerous containers and numerous mountpoints/shares from the client side.

## 6.2 Data Protection

Infinite Volume supports several methods of data protection, including Snapshot technology, SnapMirror replication technology, and NFS or CIFS mounted tape backup.

### Snapshot Technology

NetApp Snapshot at the Infinite Volume level offers the ability to manage the creation of Snapshot copies across the entire volume. Compared to a scenario in which there are several smaller limited-capacity flexible volumes with several Snapshot copies, Infinite Volume has the advantage of having a single Snapshot copy that runs across a single large-capacity container.

### SnapMirror Replication

Infinite Volume supports SnapMirror storage replication to a disaster recovery site. Although an Infinite Volume is a composite of several storage resources, there is no need to manage individual replication relationships. A single SnapMirror relationship (at the Infinite Volume level) is maintained for the entire repository.

## Tape Backup

Tape backup for Infinite Volume can be done 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.

### 6.3 Native Storage Efficiency

Infinite Volume supports storage efficiency features such as deduplication, compression, and thin provisioning.

#### Deduplication

NetApp deduplication removes duplicate data blocks, storing only unique blocks. This results in storage space savings and cost savings.

Deduplication runs on a schedule that can be customized.

#### Compression

NetApp data compression works by replacing repeating patterns within a subset of a file. It is complementary with deduplication, and compression and deduplication can be used together to provide optimal savings.

NetApp data compression has the ability to run either as an inline process as data is written to disk or as a scheduled process.

#### Thin Provisioning

NetApp thin provisioning is the most efficient way to provision storage, because the storage is not all preallocated up front. In other words, when a volume is created by using thin provisioning, no space on the storage system is used. The space remains unused until data is written to the volume; therefore at any given time only enough space to store the data is used.

Infinite Volume can be thin provisioned and planned for a larger size. The size can be increased as needed; when the physical space runs low, more storage can be added to the Infinite Volume nondisruptively. Using this technique, storage administrators can plan for a large-sized container, but they don't need to preallocate all the storage space and can avoid the initial large investment in disk drives.

### 6.4 Performance Acceleration with Flash Cache

NetApp Flash Cache is a second-generation controller module with software that offers improved performance in NetApp storage systems without adding more high-performance disk drives. Flash Cache offers 256GB, 512GB, and 1TB memory configurations, and it is scalable. Infinite Volume works seamlessly with Flash Cache to improve system performance.

### 6.5 Performance Acceleration with Flash Pool

A Flash Pool is the newest addition to the NetApp Virtual Storage Tier. It is a technology that allows flash technology in the form of solid-state disks (SSDs) and traditional hard disk drives (HDDs) to be combined to form a single Data ONTAP aggregate. When SSD and HDD technologies are combined in a Data ONTAP aggregate, the NetApp storage system takes advantage of the latency and throughput benefits of SSD while maintaining the mass storage capacity of HDD. Infinite Volume works seamlessly with Flash Pool and can improve system performance.

## 6.6 Resiliency

Infinite Volume is designed for storage resiliency and takes advantage of NetApp technologies such as storage failover, high-availability pairs, and IP failover. It inherits the advantages of clustered Data ONTAP including SnapMirror, RAID-DP, and hardware features that enhance its focus on resiliency.

## 6.7 Unified Security Style

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<sup>®</sup> and UNIX<sup>®</sup> users and groups, allowing more simplified security management.

Another part of unified security is support for unified ACLs, which consolidates file permissions for both UNIX and Windows users and groups in a single access control list. Unified ACLs facilitate access checks using both Windows and UNIX credentials. They provide multiprotocol access to files without relying on user mapping between UNIX and Windows users.

## 7 Conclusion

As capacity requirements for file-based content continue to grow exponentially, storage administrators seek enterprise content repositories that offer the ability to scale data volumes that can store billions of files. NetApp FAS and V-Series storage systems with support for NetApp Infinite Volume offer an excellent solution for storing content that requires a large continuous repository. Infinite Volume offers up to 20PB of capacity with industry-leading storage efficiency, data protection, and file systems, which makes it a leading solution for enterprise content repository workloads.

## Reference

The following reference was used in this technical report:

“Clustered Data ONTAP: An Introduction”  
<http://media.netapp.com/documents/tr-3982.pdf>

## Version History

Version	Date	Document Version History
Version 1.0	June 2012	Initial release supporting clustered Data ONTAP 8.1.1
Version 2.0	April 2013	Updated to include clustered Data ONTAP 8.2 updates.

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, FlexVol, RAID-DP, SnapMirror, and Snapshot are trademarks or registered trademarks of NetApp, Inc. in the United States and/or other countries. NetBackup and Symantec are trademarks of Symantec Corporation. Windows is a registered trademark of Microsoft Corporation. UNIX is a registered trademark of The Open Group. All other brands or products are trademarks or registered trademarks of their respective holders and should be treated as such. TR-4037-0413