

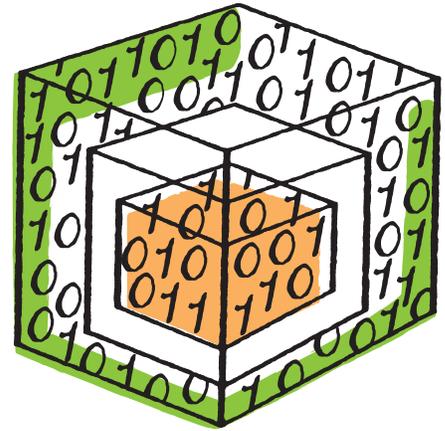


NetApp®

Datasheet

# NetApp Infinite Volume

Big volume for big data



## KEY BENEFITS

### Manage Large Datasets

Seamlessly scale to billions of files and petabytes of capacity in a single volume, reducing operational complexity.

### Simplify Data Access and Retrieval

Rapidly access large datasets from a single file system that is efficient and easy to manage.

### Increase Storage Efficiency

Store massive amounts of data using less physical capacity.

### Provide Continuous Availability

Avoid operational downtime due to system upgrades and routine management tasks with nondisruptive operations.

### Protect Your Data Easily

Retain and rapidly recover your data with efficient local copies and remote mirrors, all from a single storage platform.

## The Growth of Big Content

The data universe is expanding rapidly. Organizations of all sizes face the challenge of collecting and storing massive amounts of data to address the requirements of large unstructured repositories of primary data. Digital technologies are moving to denser media, photos have gone digital, video and medical imaging systems are using higher resolution, and advanced analytics require significantly more storage. Retaining information is critical for ongoing business operations.

This explosive increase in big data is putting enormous pressure on existing infrastructures—especially the storage platform. Traditional approaches can't scale to the level needed to retain all of the data, provide quick access to large datasets, and store the relevant information efficiently for extended periods of time. Scalable, efficient, and highly available storage solutions are needed to derive the maximum business benefits from these dramatically larger datasets.

## Addressing the Challenges of Enterprise Content Repositories

Enterprise content repositories must provide the ability to store multi-petabytes of data for long periods of time and to retrieve it quickly when needed, without losing enterprise-class

functionality such as storage efficiency and data protection. Content repository challenges include:

- **Effectively addressing long-term retention policies.** As the amount of stored data grows, data retention becomes more complex—not only in terms of volume, but also in implementing corporate policies for how long to store data, when to delete it, and where to store it. Industry regulations require enterprises to store data for decades or forever, with the need to make sure that data is not compromised as the storage infrastructure evolves over time.
- **Quickly finding and retrieving content from long-term repositories.** Once retention policies are set, the next challenge is to be able to quickly find and retrieve relevant content from the long-term repositories. Data that is stored for long periods of time has historically been stored on tape libraries. However, retrieving this data for analysis or for new projects can waste valuable time—if the data can be found at all. These processes become even more difficult as the storage infrastructure evolves and grows more complex over time.
- **Managing dedicated hardware with unique software tools and features.** Managing a content repository with purpose-built hardware introduces

yet more management complexity. Additionally, single-purpose solutions affect operational efficiency, because resources are commonly sharable with other applications.

### NetApp Infinite Volume

The new era of big data is breaking traditional architectures—there’s simply too much content to store and manage efficiently. NetApp is addressing these challenges by integrating scalable data storage with enterprise functionality. The NetApp® Infinite Volume feature enables enterprises to store tens of petabytes and billions of files in a single, scalable, easy-to-manage repository.

With NetApp clustered Data ONTAP®, organizations no longer need to compromise on full enterprise functionality, including storage efficiency, integrated data protection, multi-tenancy, application integration, and seamless scalability. With Infinite Volume, clustered Data ONTAP delivers both scale and enterprise functionality, enabling customers to start small and grow incrementally in a shared Data ONTAP cluster.

The NetApp Infinite Volume feature effectively addresses the key operational requirements of enterprise content repositories:

- **Reduce operational complexity across enterprise storage environment.** NetApp Infinite Volume can share the same storage cluster with other business applications, such as Oracle®, SAP®, VMware®, and Microsoft®. Increases in storage utilization reduce operational cost and complexity.
- **Easily handle massive data growth.** Clustered Data ONTAP with Infinite Volume offers the ability to seamlessly scale to billions of files and tens petabytes of capacity in a single volume.

| FEATURE               | CAPABILITY   |
|-----------------------|--|
| Max Capacity          | 20PB raw<br>2 billion files<br>Max file size = 16TB  |
| Supported Systems     | <ul style="list-style-type: none"> <li>• FAS/V 3000 and 6000 series (requires support for Data ONTAP 8.2)</li> <li>• Support for mixing different models of storage controllers</li> </ul> |
| Protocol Support      | NFSv3, NFSv4.1, pNFS, SMB 1.0  |
| Multi-tenancy Support | Support for Infinite Volume and FlexVol® volumes in same storage cluster   |
| Node Maximum          | Infinite Volume can span up to 10 nodes in a cluster, regardless of cluster size   |
| Data Protection       | RAID-DP®, Snapshot, asynchronous volume SnapMirror   |
| Storage Efficiency    | Deduplication, compression, thin provisioning  |
| Tape Backup           | NFS, SMB, or Differential Support (SnapDiff)   |

- **Increase storage efficiency.** Storage efficiency features, such as deduplication, compression, and thin provisioning, reduce the amount of physical disk storage required, offering significant economic benefits for organizations with large datasets.
- **Provide continuous availability.** You can seamlessly scale, load balance, and upgrade your storage environment while retaining access to your data.
- **Quickly find and retrieve content from long-term repositories.** With clustered Data ONTAP, enterprises can retrieve content rapidly from efficient, easy-to-manage, active archives as opposed to storing on tape.
- **Provide integrated data protection.** Clustered Data ONTAP provides integrated data protection, including Snapshot™ copies, SnapMirror® replication, and integration with data protection software partners, so that data is protected, even across multiple data centers.

### Summary

The world has entered a new era of scale, where the amount of data is rapidly breaking existing storage architectures. The NetApp Infinite Volume feature offers boundless, secure, scalable data storage for big data. With NetApp, enterprises can efficiently and securely drive successful business outcomes—all at the speed of today’s business.

### About NetApp

NetApp creates innovative storage and data management solutions that deliver outstanding cost efficiency and accelerate business breakthroughs. Discover our passion for helping companies around the world go further, faster at [www.netapp.com](http://www.netapp.com).

Go further, faster®



[www.netapp.com](http://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, FlexVol, RAID-DP, SnapMirror, and Snapshot are trademarks or registered trademarks of NetApp, Inc. in the United States and/or other countries. Microsoft is a registered trademark of Microsoft Corporation. Oracle is a registered trademark of Oracle Corporation. SAP is a registered trademark of SAP AG. VMware is a registered trademark of VMware, Inc. All other brands or products are trademarks or registered trademarks of their respective holders and should be treated as such. DS-3458-0413

Follow us on:     