VERITAS

Access Appliance with NetBackup

Long-Term Retention Solution.

White Paper | August 2022

Contents

Introduction . . . 5

Executive Summary . . . 5

Scope . . . 5

Target Audience . . . 5

Solution Value . . . 5

Solution Key Features . . . 6

Seamless Solution Integration with NetBackup . . . 6

Immutability and Ransomware Protection . . . 6

Data Encryption . . . 7

Storage Efficiency . . . 7

Support for Multiple NetBackup Domains . . . 7

Appliance CallHome and AutoSupport . . . 8

Monitoring and Detection . . . 8

Solution Architecture . . . 8

NetBackup . . . 9

Veritas Data Deduplication . . . 9

MSDP Cloud Tiering (MSDP-C) . . . 10

Veritas Data Deduplication on Access Appliance . . . 10

Comparison Highlights of MSDP-C with Access and Veritas Data Deduplication . . . 11

Traditional Duplication (Without Deduplication) . . . 11

Access Appliance . . . 11

Solution Data Flow . . . 13

Deduplication Data Flow with MSDP-C . . . 13

Veritas Data Deduplication Flow . . . 14

Traditional Duplication Data Flow . . . 16

Contents

Disaster Recovery 17

AIR and Optimized Duplication (MSDP-C) to Access S3 Bucket 17

AIR for Veritas Data Deduplication (VDD) 18

AIR With Access S3 Bucket (Without Deduplication) 20

Cloud Support 21

Best Practices and Recommendations. 21

Data Layout on Access Appliance 21

Deduplication 21

Compression 22

Encryption 22

Network Connectivity 22

Multiple NetBackup Domains 22

Load Balancing for MSDP-C 22

Monitoring 23

Sizing Guidance 23

Samples of Capacity Sizing for Access Appliance 23

Sizing for MSDP-C 24

Sizing for Veritas Data Deduplication 24

Sizing for Traditional Duplication of Data 25

Conclusion 25

References. 26

Date:	Updates
March 2018	Initial version
March 2018	Minor corrections on support
April 2018	Minor updates
October 2018	Updates based on 7.4.2 release
March 2019	Updates to level of support Veritas Data Deduplication
July 2019	Updates to memory requirements for Veritas Data Deduplication
February 2021	Updates based on 7.4.2.301 release and NetBackup 8.3
May 2021	Updates based on 7.4.3 release and NetBackup 9
June 2022	Updates based on 8.0 release and NetBackup 10

Introduction

Executive Summary

Veritas Technologies is a leader in developing data resiliency solutions that focus on protection and management of companies' digital assets critical for their success and business continuity. NetBackup, one of Veritas' flagship products, is designed to protect datacenters as well as hybrid- and multi-cloud environments. Adding to Veritas' portfolio and legacy of creating stable solutions that customers have trusted and relied on is Access Appliance, a turn-key storage appliance created to address the long-term retention needs of organizations. An Access Appliance acts as an on-premises, very dense, highly available and WORM-capable storage target for data that has been backed up using NetBackup. Access software has been optimized and designed to work seamlessly as part of a fully featured, ransomware-resilient NetBackup data protection solution, and this integration provides a compelling offering for a long-term retention (LTR) use case.

Scope

The purpose of this document is to provide technical details to assist in understanding Access Appliance with NetBackup as a solution for long-term retention of backup data. It describes the components of this solution, its value, sizing guidance, and some best practices. It is advised to refer to Veritas product documentation for installation, configuration, and administration of each of the products discussed in this whitepaper.

Target Audience

This document is targeted for customers, partners, and Veritas field personnel interested in learning more about the Veritas Access Appliance with the NetBackup solution for long-term retention. It provides a technical overview of this solution, guidance in sizing, and highlights some best practices.

Solution Value

Companies usually have a strategy to protect data in case of a failure, disaster, or crisis and NetBackup is an industry leader in this area. However, as data increases at an accelerating pace, companies are striving to determine the best strategy in the management, preservation, and retention of their valued data for long-term. There are several challenges that come to mind when talking about a long-term solution which include cost, complexity, control, and visibility. Traditionally, the solution for long-term retention has been tape because of its low cost. However, the complexity in tape management in addition to the time to restore has been an issue. Recently companies have looked to the public cloud for a possible solution, however, issues in total cost of ownership and control become a concern.

To address all these challenges, Veritas has designed the Access Appliance as a purpose-built, on-premises storage appliance for long-term retention use cases. Together with NetBackup, the Access Appliance provides a resilient and cost-effective solution for the preservation of data backups that companies want to retain and have readily available for further use.

As this document unfolds the architecture and features of this solution, it will showcase the following key values:

- Cost Minimization – Access Appliance provides a low-cost, disk-based solution that is easy to manage. With NetBackup deduplication feature, the amount of storage space is reduced by saving only one copy of the data blocks and having the duplicates point to that one copy, thus providing a more storage efficient solution and reducing overall costs. When using the Veritas Data Deduplication (VDD) feature introduced in version 7.4.2, further reduction can be observed in a multi-domain NetBackup environment as Access supports global deduplication of data across the domains or across multiple media servers.
- Simplified Management – Access Appliance with NetBackup has deep integration features that simplify the configuration and administration, such as invoking policy-based storage management and intelligent data movement between tiers.
- Increased Visibility and Control – more and more companies would like to leverage the data that has been archived and retained for IT or business analysis and investigations so having the data on-premises under the company's control and visibility allows for quick restores to conduct these studies.

Solution Key Features

There are certain key features that companies look for in a long-term retention solution product: flexibility, storage efficiency, and ease of management. The Access Appliance with NetBackup provides these features to assist customers in preserving their most valued data.

Seamless Solution Integration with NetBackup

Access Appliance closely integrates with a larger NetBackup domain in three areas:

- As a target open storage server providing a Veritas Data Deduplication storage pool that will receive and store backup data from multiple NetBackup domains, participate in storage lifecycle policies, and implement WORM retention periods
- As an S3/object storage provider for MSDP-C

The Access Appliance has been integrated as a cloud provider in NetBackup 8.1 (with updates) and later, and as a target open storage server for the Veritas Data Deduplication feature. For instance, during configuration of a cloud storage server on the NetBackup primary server's administration graphical user interface (GUI), Access is listed as one of the cloud storage providers that can be selected.

A container-based NetBackup client add-on package is also available to be installed on the Access Appliance. A good use case for having the NetBackup client integrated into the Access Appliance is the protection of the deduplication catalog when using a Veritas Data Deduplication pool as a target. It can be backed up to another location for added protection.

In addition, the Access Appliance GUI provides a configuration wizard and policies to make it easier to provision an S3 bucket for NetBackup or Veritas Data Deduplication pool. Provisioning can also easily be done using "Quick Actions" to provision another S3 bucket or configure a Veritas Data Deduplication storage pool.

Immutability and Ransomware Protection

NetBackup Appliances offer a lockdown mode of operation that provides a secure, immutable infrastructure in support of archival data that requires WORM (write-once-read-many) storage. Data written to WORM storage by definition cannot be altered or deleted before its scheduled retention time has elapsed. Access (and other appliances that run in lockdown mode) use a compliance clock to maintain time consistency and prevent tampering.

Both S3 buckets on Access used by NetBackup for MSDP-C object storage and Access running a Veritas Data Deduplication storage pool are capable of WORM storage. Access must have WORM enabled for the filesystem containing an S3 bucket to enforce object locking.

In the case of the Veritas Data Deduplication storage pools both in Access and on NetBackup appliances, the minimum and maximum retention periods are first configured when entering lockdown mode. A NetBackup data protection policy that sends data to Access via storage lifecycle policy therefore must ensure that the retention period configured for the second copy on Access fits into this retention window.

NOTE: You can enter lockdown mode with a Veritas Data Deduplication service already configured in Access, but the existing storage pool will not gain WORM / immutability support. A new storage pool must be created and a retention policy set by NetBackup.

For more information about running Access Appliance in lockdown mode, see the section Support for immutability in Access Appliance in the Veritas Access Appliance Administrator's Guide.

Data Encryption

NetBackup has security features that protect all NetBackup components and operations at different security implementation levels such as datacenter, world, and enterprise. For enhanced security, NetBackup offers encryption of data. When NetBackup encrypts backup data, it remains encrypted when transferred to long-term retention on Access Appliance. NetBackup sends data over dedicated and secure network ports to Access.

Additional security that is employed for this solution is the requirement to use Access user keys and credentials when configuring Access as a cloud storage destination. When using MSDP-C, data sent to Access can also be secured by enabling the SSL encryption feature on Access. When SSL is enabled, certificates are generated by Access and placed in NetBackup primary and media servers and data is sent via HTTPS. Refer to the NetBackup Security and Encryption Guide for further details on how NetBackup performs encryption. When using NetBackup deduplication technology, there is encryption for deduplicated data which is separate and different from NetBackup policy-based encryption. For more information on the implementation, refer to the NetBackup Deduplication Guide.

Storage Efficiency

Long-term storage solutions benefit greatly from any increase in storage efficiency, as an increase in utilization will maximize storage space while assisting in reducing overall cost. Backup images stored in the Access Appliance can be deduplicated using NetBackup Media Server Deduplication Pool (MSDP) technology. Data is sent to Access either via MSDP cloud tiering (MSDP-C) or directly into the Veritas Data Deduplication storage server, and in either case, data deduplication is maintained.

NetBackup also supports compression prior to sending data to the Access Appliance. So, Access stores the compressed format of the data. Compression improves storage utilization by reducing the number of bits required to represent data. The type of data defines the degree a file can be compressed. Data types that compress well include text files or unstripped binaries. Data that is already compressed and stripped binaries are not good candidates for compression.

Support for Multiple NetBackup Domains

A single Access Appliance deployment can receive data from multiple NetBackup domains or environments, even ones in different geographical locations or network contexts. For example, multiple instances of NetBackup protecting multiple clients from different sites, locations, or data centers can use a single Access Appliance for long-term retention as shown in Figure 1. The Access Appliance must be reachable from either site and data is sent to Access from the NetBackup instances using the S3 protocol on HTTP or HTTPS transport when using MSDP-C or with a proprietary protocol when using Veritas Data Deduplication. When using Veritas Data Deduplication in a multi-domain NetBackup environment, more storage efficiency may be observed since only one copy is saved if a duplicate block is encountered from both domains.

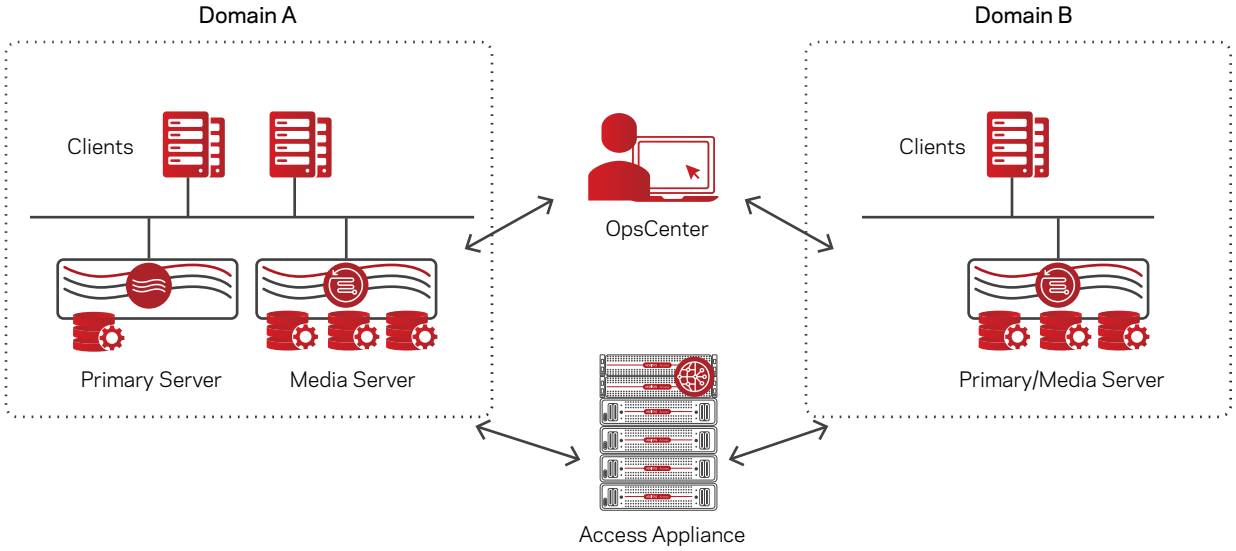

Figure 1. Utilizing Access Appliance for long-term retention

Appliance CallHome and AutoSupport

Veritas Appliances such as the Access Appliance and NetBackup Appliance can call home if their health monitoring services observe hardware or software issues. The advantages of using Veritas appliances for the entire solution are the ability to automate support case management and leverage guided workflows for faster resolutions of issues and mitigation of risks. Veritas AutoSupport service provides continuous proactive monitoring and alerts for appliance health. This feature alerts customers and/or service engineers to quickly handle the issue and reduce further risks. Enabling this feature can be done simply by registering the appliances at the Veritas NetInsights portal, and enabling call-home functionality to send appliance telemetry.

Monitoring and Detection

Available on the Access Appliance is Symantec Data Center Security (SDCS), an intrusion detection system. SDCS is a real-time monitoring and auditing software. It performs host intrusion detection, file integrity monitoring, configuration monitoring, user access tracking and monitoring, and produces logs and event reports. SDCS adds security hardening and monitoring for the Access Appliance to reduce security risks and

attacks. For more information on the Access Appliance intrusion detection system, refer to the section Access Appliance Initial Configuration Guide.

Solution Architecture

The high-level architecture of a long-term retention solution based on NetBackup and Access consists of:

- Data sources to back up (such as databases, applications, virtual machines, files, and email)
- NetBackup components (either appliances, or build your own server (BYOS) deployments)
- Access Appliance

As pictured in Figure 2, data is backed up to NetBackup for short-to- mid-term retention and then moved to Access Appliance for long-term retention.

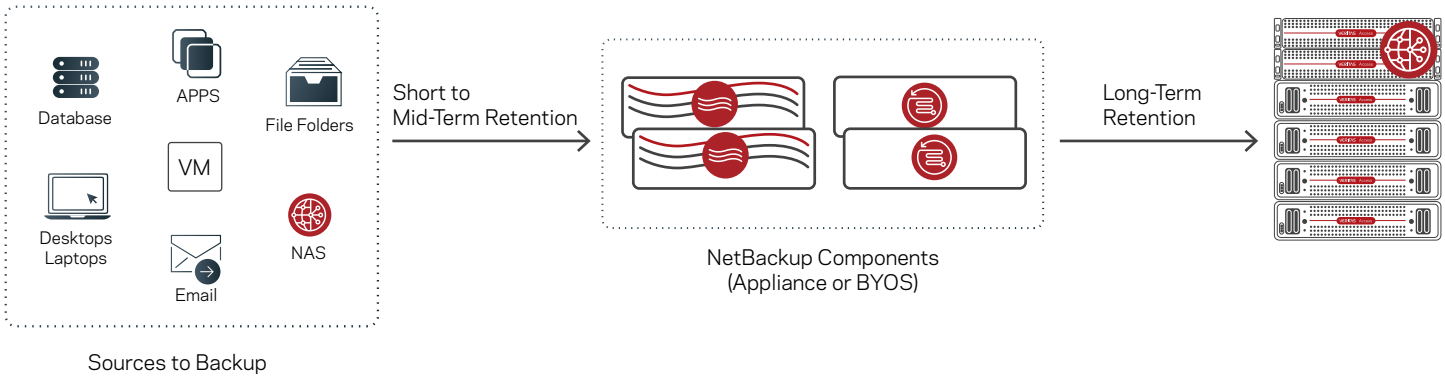

Figure 2. Utilizing Access Appliance for long-term retention

Two ways to store backup images on Access from NetBackup include:

- Deduplication (Optimized Duplication) – deduplication using NetBackup MSDP deduplication technology
- Without Deduplication (Traditional Duplication) – backup images are duplicated to the Access Appliance from NetBackup

The next section will show in detail how Access integrates with NetBackup, with a view of describing the participation of all involved solution components. It will also cover information on different ways to deploy NetBackup, along with deduplication options.

NetBackup

Veritas NetBackup provides protection for a wide variety of data and platforms such as operating systems, virtual systems, databases and applications, files, and all kinds of content. It has many add-on features to speed up backups, snapshot management, backup automation, and provide insights on where the active and inactive backups are located. It has the capability to backup data to tape, SAN, NAS, public or private cloud. Schedules, retention periods, and the ability to tier to different types of storage are defined in policies or storage lifecycle polices (SLP).

A typical NetBackup environment consists of three types of components:

- **Primary Server** – manages and controls the backup and recovery activities and hosts the catalog that contains policies and schedules, metadata about the backup jobs, and media, device, and image metadata information.
- **Media Server(s)** – writes client data as backup images to varying types of storage – such as local disks, storage area network (SAN) volumes, tapes, or network attached storage (NAS) shares – and later restores the data to the client as instructed by the primary server.
- **Clients** – NetBackup client components are installed on hosts that have the data to be backed up and responsible for sending and receiving data to and from media server for backup and recovery.

The primary and media server components can coexist in a single system or be distributed to several, depending on the quantity of clients, variety of workload, and data protection schedule. For instance, a small environment could consist of a primary and media service sharing a single server, while a larger environment with more data to protect would be spread across a primary and multiple media services on separate hardware. NetBackup environments may also span physical sites and participate in cross-domain replication relationships. Figure 3 illustrates a simple configuration of a set of clients with a dedicated primary server and one media server sharing a common network.

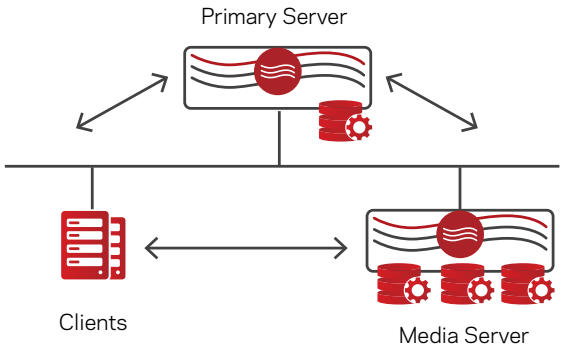

Figure 3. A NetBackup simple configuration

NetBackup is very flexible in its deployment. It can be deployed using NetBackup appliances, commodity servers (described as BYOS, or Bring Your Own Server), or both of these when necessary. It can also be deployed as a containerized instance on a NetBackup FLEX Appliance, or as a scale-out cluster with a shared storage pool using NetBackup Flex Scale.

Veritas Data Deduplication

Backup images are generally ideal for deduplication since the probability of encountering duplicated blocks of data can be high depending on the workload being protected. The space savings of deduplication is expressed using a deduplication ratio, a measurement of the data's original size compared to its on-disk size after deduplication. The higher the deduplication ratio, the more space is saved.

Workload and data type, data change rate, required retention period, and backup policy are all factors in the effectiveness of deduplication for backup data. For instance, encrypted data is inherently unique and will not benefit from deduplication savings.

Data that has a high change rate will not take advantage of the savings long enough to justify the overhead imposed by deduplication. In the context of backup images, daily full backups will have higher deduplication ratios when compared with incremental or differential backups.

NetBackup permits inline deduplication of backup images on either the client, using the NetBackup client plugin, or on a media server. Client-side deduplication uses available resources on clients and reduces overall network traffic, as only deduplicated data is sent over the network. In either case, backup images are retained in a deduplicated storage pool.

Veritas Data Deduplication is composed of the following architectural components (with their general areas of responsibility):

- Deduplication Plugin: turn a data stream into segments, use a hash algorithm to calculate segment fingerprints, compare with known data fingerprints, send needed data to deduplication engine
- Deduplication Engine (spoold): manage and store the fingerprint database and metadata, store unique segments or update references to stored data, conduct integrity checks
- Deduplication Manager (spad): maintain the storage pool configuration, control and dispatch internal processes, security, and events handling

Note that these components are commonly distributed and sometimes duplicated in a NetBackup environment for sizing, resiliency, or security reasons.

NetBackup MSDP utilizes SHA-2 (SHA256) for the hash algorithm. The chunk segment size unit used to compute fingerprints is by default a fixed length of 128 KB or configurable to variable-length size based on chunk boundary. The storage pool also compresses deduplicated data when possible to further increase storage efficiency. There is also an option to encrypt deduplicated data. For more information on the architecture of NetBackup MSDP deduplication technology, refer to the NetBackup Deduplication Guide.

There are two methods to send deduplicated data to Access without rehydration:

- MSDP Cloud Tiering (MSDP-C) – deduplication is done by media server and deduplicated data is sent to Access using MSDP cloud tiering. S3 protocol is used in this scenario.
- Veritas Data Deduplication (VDD) – deduplication done by NetBackup media server or client prior to sending to Veritas Data Deduplication pool. MSDP proprietary protocol is used to send data to Access.

MSDP Cloud Tiering (MSDP-C)

MSDP has support to send deduplicated data without rehydration directly to Access, referred to as MSDP Cloud (MSDP-C) tiering. Configuring a NetBackup appliance or BYOS as a MSDP-C storage server allows data to be sent via S3 protocol to one local storage target and one or more cloud storage targets. However, there is a combined capacity support of 1.2 PB for both block and object.

Veritas Data Deduplication on Access Appliance

To support NetBackup MSDP in Access, the NetBackup MSDP technology components ported to Access Appliance include the deduplication engine (spoold) and deduplication manager (spad). This feature was introduced in Access Appliance version 7.4.2. The primary functions of these components are to manage and store unique data, fingerprints, metadata, and the associated logs and journals. Both nodes share the storage used as a deduplication pool, but only one runs the management processes at a time, in an active/passive configuration. If the active node fails or is unreachable, Access cluster management will automatically start the necessary components on the passive node and resume the processing and storing of deduplicated data.

NOTE: A NetBackup media server or client with the deduplication plug-in is still required to do the actual deduplication of data which involves segmentation of data, calculating fingerprints and comparison with existing fingerprints. During configuration of the Access Appliance as a storage server in NetBackup, a media server is required and needed during restores.

An advantage of VDD is in a multi-domain NetBackup configuration or when utilizing multiple media servers in one domain, data is "globally" deduplicated in Access such that only one copy of the data is stored between the domains.

Comparison Highlights of MSDP-C with Access and Veritas Data Deduplication

Some essential highlights of the differences between the two options to support NetBackup data storage with Access are enumerated in Table 1.

Features:	MSDP-C	Veritas Data Deduplication
Transfer Protocol	S3 (HTTP/HTTPS)	Proprietary (TCP/IP)
Ports	8143	10082 and 10102
Destination Type	Bucket	Deduplicated storage pool
Access storage layout	CFS (clustered file system) hosted buckets, 1 bucket per filesystem recommended	1 catalog (catfs) filesystem deduplicated data store in concatenated CFS filesystems
Cloud Support	Via NetBackup MSDP-C	Via NetBackup MSDP-C
Data Stored in Access	Metadata, unique data segments, encryption key (if enabled)	Metadata, fingerprints database, unique data segments, and encryption key (if enabled), journals and logs
Maximum pool size	1.2 PB	1.4 PB
Global Deduplication in Multi-Domain	No	Yes
WORM capability	Yes	Yes

Traditional Duplication (Without Deduplication)

In some cases, deduplication of backup images is not ideal. Backups that have a strict time limit for restores, have a high rate of change, or encrypted are not good candidates for deduplication. Another case is when incremental backups are done instead of full. For these types of data or backup, images are best sent to Access without deduplication. Data is sent to Access using the S3 protocol and stored in an S3 bucket.

Access Appliance

NetBackup can send backup images to various storage types (disk, tape, cloud, etc.) for long-term retention. For those seeking an on-premises disk-based solution for faster recovery times, control and/or simplicity when compared to tape or cloud, Veritas has developed the Access Appliance for ease of acquisition, management, and support. Access Appliance is a turn-key storage solution designed for high capacity and cost optimization, making it well suited for long-term retention. Access Appliance is comprised of two clustered nodes and one primary storage shelf and up to three additional expansion storage shelves. The appliance can scale up to 2,800 TB of usable space.

Highlights of Access Appliance specifications are shown in Table 2.

Specification	Access Appliance 3340	Access Appliance 3350
Number of nodes	2	2
CPU	Intel Xeon 4108	Intel Xeon Silver 4210
Memory (per node)	384, 768, 1152, or 1536 GB	
Rack units	Node: 2U Storage shelf: 5U Fully expanded system: 24U	
Expansion shelves	1 – 4 per appliance	
Shelf storage	82 drives per shelf (all 4TB or all 10TB drives)	
10/25Gb network ports	2	2 or 4
Storage total	280 TB – 2800 TB usable per appliance	

The two nodes are clustered in active/active configuration such that each node can handle I/O requests. Storage shelves are connected to each node and configured with dynamic multipathing so I/O can be sent to either node for performance and availability.

When using Access as an S3 target, deduplicated data written to Access from NetBackup is placed in an S3 bucket. A bucket maps to a single filesystem of the default CFS (clustered filesystem) type. For availability reasons these filesystems are available to either node of an Access Appliance. When using the S3 protocol to receive backup images from NetBackup, the S3 object URL presented to clients is of the form s3.<clustername>:8143. Clients such as NetBackup can use this URL as an S3 endpoint for reading and writing to the Access bucket. Clients simply map this S3 object URL to one of the Access virtual IPs. A dedicated S3 communication port, 8143, is required for both HTTP and HTTPS and thus firewalls must keep this port open.

Veritas Data Deduplication (VDD) support is also a feature of Access. When deploying VDD, the fingerprints, metadata, and deduplicated data are placed in purpose-created clustered filesystems, so it is not required that the entire storage pool be dedicated for the data deduplication filesystems created. As a proprietary protocol is implemented to communicate between NetBackup MSDP and VDD engine and manager, ports 10082 for Access deduplication engine and 10102 for the Access deduplication manager must be open (that is, not filtered by a firewall) for this communication. The active node is determined when specifying the virtual IP during configuration and provisioning of VDD. Access will start deduplication services on the node that owns the physical interface associated with the virtual IP (VIP) specified during configuration.

NOTE: Qualified maximum data deduplication pool size is 1.4 PB. Also, maximum supported concurrent jobs or streams is 80.

For management, the appliance can be managed by the command-line shell referred to as the CLISH and a web-based graphical user interface (GUI) where one can create and provision Access as an S3 target and/or data deduplication pool. A configuration wizard is also available on the GUI for quick provisioning of an Access S3 bucket and VDD for NetBackup. The wizard walks user through creation of a storage pool of disks, activation of the appropriate policies and provisioning of an S3 bucket or data deduplication pool for NetBackup.

NOTE: For examples of how to deploy and configure the Access Appliance with NetBackup with MSDP-C refer to the Quick start section of Veritas NetBackup Deduplication Guide. For NetBackup with VDD, refer to the Veritas Access Solutions Guide for NetBackup.

Solution Data Flow

This section explores how all these components integrate and how data flows through each component. The flow of data depends on whether deduplication using MSDP-C, VDD or no deduplication is employed. Also, the definitions within the NetBackup protection policies and/or Storage Lifecycle policies (SLP) also affect the flow.

In all data flow scenarios, an Open Storage Technology (OST) plugin sends data to Access Appliance. OST is an API developed by Veritas in order to allow third-party vendors to develop a software plugin module that tightly integrates their products with NetBackup software. This plugin is installed on the media servers to communicate with the vendor's storage device. OST plugins were developed to send data to S3 compatible cloud providers and VDD pool. Both plugins are by default installed on NetBackup media and primary servers to send data to Access.

Deduplication Data Flow with MSDP-C

NetBackup polices and/or SLP define the path or flow of data. Regular policies can be defined if the data needs to be sent to a single target and SLP can be set up to backup, duplicate, and/or replicate the data in different storage types or destinations. For instance, an SLP can send backup first to a media server with faster disks and then duplicate to another media server with slower disks or to an MSDP and/or to secondary storage for long-term retention. Data is sent to an Access S3 bucket utilizing the S3 protocol with the S3 OST cloud plugin installed on a media server. Examples of NetBackup MSDP deduplication technology with MSDP-C and Access data paths are explained below and illustrated in Figure 4:

A. Data from clients are initially backed up and deduplicated to an MSDP local storage disk pool residing on a media server for short or mid-term retention (Copy 1). This deduplicated data is sent to Access Appliance via MSDP-C for long-term retention (Copy 2). For restores, a copy of the data can be restored from either the first copy on MSDP or the second copy residing on Access S3 bucket. However, by default, restores are retrieved from Copy 1. The media server is required to rehydrate the deduplicated data from the MSDP copy.

B. Data is backed up on an advanced disk on the media server for short or mid-term retention and then data is deduplicated and unique data is sent to the Access Appliance S3 bucket via MSDP-C. In case of restores, the data by default will be restored from the advanced disk (Copy 1) unless specified to restore from the Access Appliance copy 2. The media server will rehydrate the deduplicated data before sending to the client.

C. The maximum combined MSDP size for local and object is 1.2 PB per MSDP-C storage server. To utilize a 1 PB Access S3 bucket and adhere to this requirement, another media server can be deployed to act mainly as a data mover. In example C, data from clients are initially backed up and deduplicated to an MSDP local storage disk pool residing on a media server for short or mid-term retention (Copy 1). Another media server is deployed to be used as a data mover to send deduplicated data to Access Appliance via MSDP-C for long-term retention (Copy 2). As in the other examples, to restore data, a copy of the data can be restored from either the first copy on MSDP local storage or the second copy residing on Access S3 bucket.

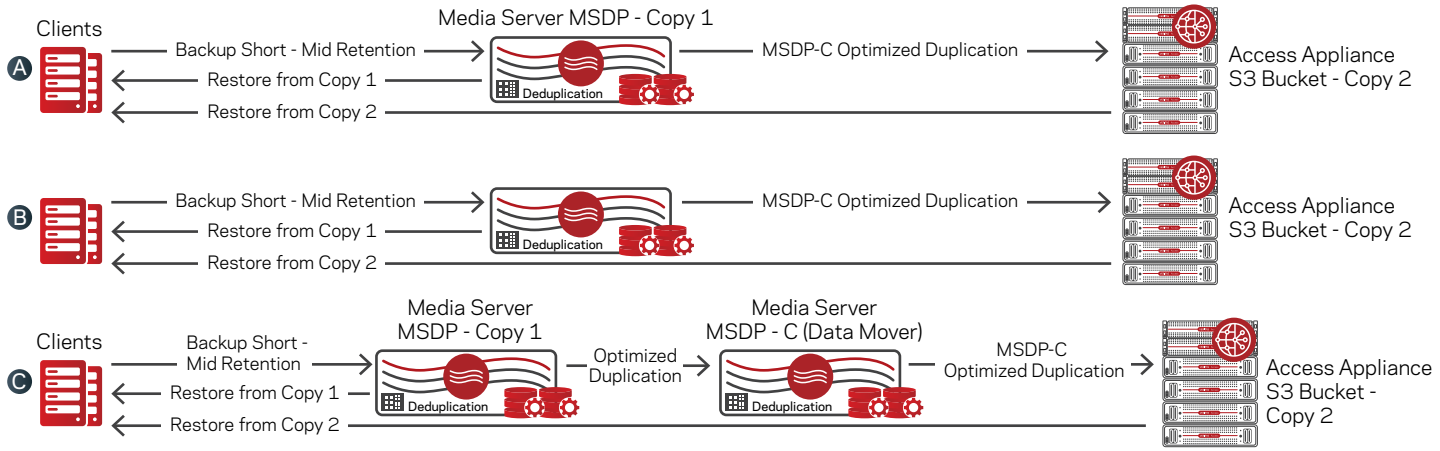

Figure 4 – Examples of MSDP Deduplication with MSDP-C to Access Appliance S3 Bucket Data Flows

In addition to the MSDP metadata, MSDP separates a backup image into container files and adds a header for each container. Thus, there are two files for each MSDP container consisting of a data container file holding unique data with fingerprint and header file for each container. If encryption is enabled, there is additional information relating to the keys and header for the keys. A sample view using the S3Browser application shows the data in an Access S3 bucket after being sent from MSDP-C as shown in Figure 5.

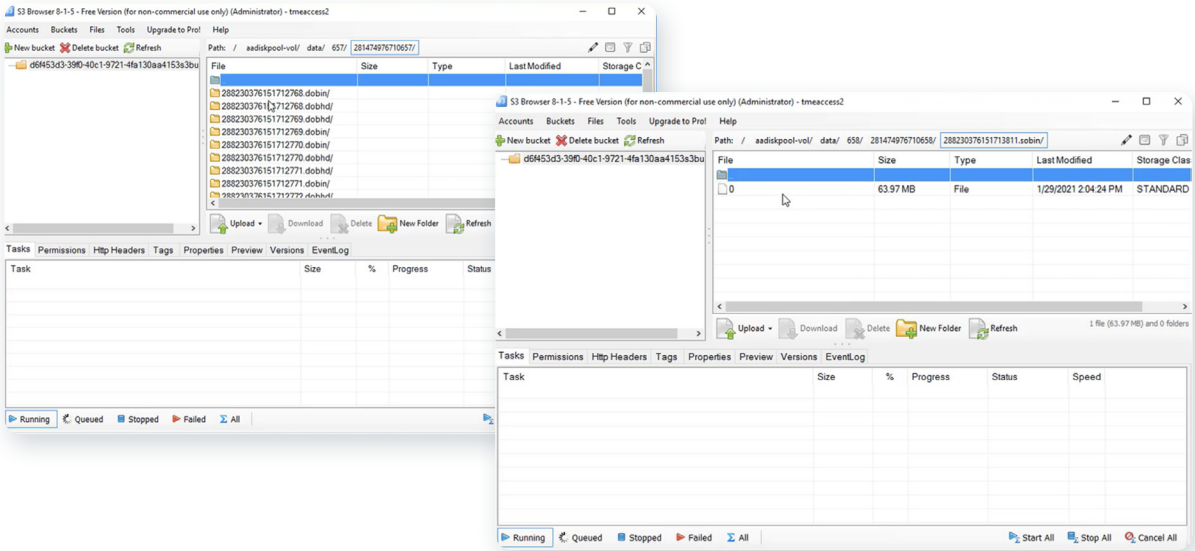

Figure 5. Sample View of the Data in an S3 Bucket Sent from MSDP-C.

Veritas Data Deduplication Flow

As discussed, data protection policies and storage lifecycle policies define how backup images are stored and tiered to other storage platforms and the path data takes when transferred from primary to long-term archival storage. When utilizing VDD, a media server is required to do the deduplication of the data. Access main role in the path is to store the unique data, fingerprints database, metadata in addition to supporting files such as journals and logs. As shown in the example in Figure 6, data from clients are initially backed up to an MSDP on a media server where it is deduplicated and stored for short to mid-term retention and is the primary copy. An SLP is defined to duplicate the unique data blocks to the Access Appliance for the second copy. Restores can be done from Copy 1 (default) residing on an MSDP in media server or from copy 2 residing in Access. When restoring from copy 2, a media server is required to retrieve the data from Access, rehydrate and send to client.

As previously discussed, data is sent to Access using NetBackup MSDP proprietary protocol based on the OST framework.

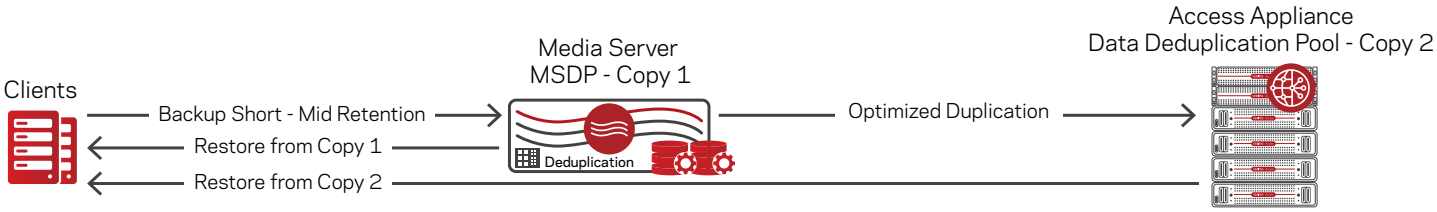

Figure 6. Example of NetBackup Data Flows to VDD Pool

One advantage of VDD is the ability to do global deduplication in a multi-domain environment or multiple media servers in one domain. A single Access Appliance can be the target storage for multiple media servers from same or different domains. Figure 7 illustrates the flow of how global deduplication is achieved with NetBackup and VDD in a multi-domain environment. A description of flow is as follows in either domain:

1. The local media server is responsible for segmentation and fingerprinting of the backup data into data blocks. It will check the media server fingerprint cache for duplicates and stores only unique blocks in local media server media server deduplication pool.

2. A duplication is conducted to create a second copy on the Access Appliance. The local media server fetches and compares the fingerprints from previous backup of image and then on the fingerprint cache on the Access Appliance. If data block does not exist, then the unique data block is sent to the Access Appliance.

In this case it's necessary for the Access Appliance to be reachable from both domains. If an address on another network context is needed, the dedupe addip command may be used to set up additional listening IP addresses for the Veritas Data Deduplication service.

NOTE: Conditions in which deduplication and global deduplication is not achievable include a) the data's fingerprint is not in local media server or Access fingerprint cache (size is based on physical memory) during comparison and b) fragmentation where the blocks are in different containers that are far apart such that it affects performance of restores. Frequency of backups and/or duplication jobs can also affect the differences in the deduplication rate observed on media server and the Access Appliance.

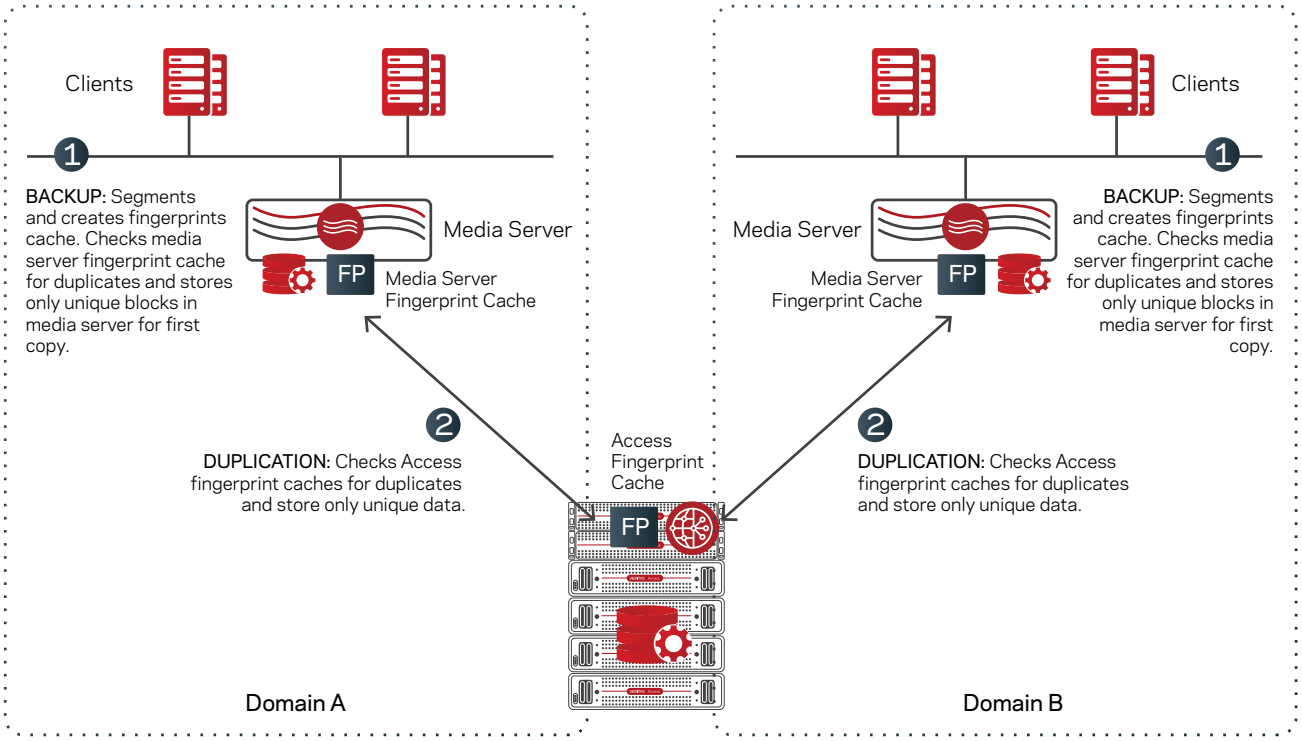

Figure 7 - Global Deduplication in Multi-Domain Environment

Figure 8 provides a view of how the data is stored on Access. In this example, there are multiple 100 TB filesystems for the deduplicated data (i.e. /vx/D3_*) and another filesystem for the MSDP catalog (i.e./vx/CAT_*). Listing the contents of the MSDP catalog directory one can see the fingerprint database, journals, logs and other associated directories. The data directory holds the actual the backup image container (i.e. *64.bin) and the header information (i.e. *64.bhd) associated with the backup image container. The image containers hold multiple segmented data objects and by default, the maximum size of the containers is 64 MB. If encryption is enabled on NetBackup, then there are also files for the encryption key and header of key.

```
/dev/mapper/isw_biffdgjhg_Baseboard   56G  506M   53G   1% /baseboard
tmpfs                                4.0K     0  4.0K   0% /dev/vx
tmpfs                                 38G     0   38G   0% /run/user/0
/dev/vx/dsk/sfsdg/D3_4132244         100T   82T   19T  82% /vx/D3_4132244
/dev/vx/dsk/sfsdg/D3_2575593         100T   93T  7.9T  93% /vx/D3_2575593
/dev/vx/dsk/sfsdg/D3_9976090         100T   89T   12T  89% /vx/D3_9976090
/dev/vx/dsk/sfsdg/D3_6157082         100T   79T   22T  79% /vx/D3_6157082
/dev/vx/dsk/sfsdg/D3_9319176         100T   87T   14T  87% /vx/D3_9319176
/dev/vx/dsk/sfsdg/D3_2570877         100T   85T   16T  85% /vx/D3_2570877
/dev/vx/dsk/sfsdg/D3_8938215         100T   86T   15T  86% /vx/D3_8938215
/dev/vx/dsk/sfsdg/D3_5478620         100T   93T  7.9T  93% /vx/D3_5478620
/dev/vx/dsk/sfsdg/D3_7630366         100T   93T  7.9T  93% /vx/D3_7630366
/dev/vx/dsk/sfsdg/D3_1350171         100T   93T  7.9T  93% /vx/D3_1350171
/dev/vx/dsk/sfsdg/_nlm_              1.0G   37M  927M   4% /shared
/dev/vx/dsk/sfsdg/CAT_7035928        5.0T  210G  4.8T   5% /vx/CAT_7035928
/dev/vx/dsk/sfsdg/D3_4411258         100T   90T   11T  90% /vx/D3_4411258
/dev/vx/dsk/sfsdg/D3_7034605         100T   89T   12T  89% /vx/D3_7034605
-bash-4.2# ls /vx/CAT_7035928/dedupe/databases
catalog  catalogshadow  datacheck  refdb  spa  task
-bash-4.2# ls -lh /vx/D3_4132244/12283/12577996*
-rw-r----- 1 root root 55K Jan 25 09:30 /vx/D3_4132244/12283/12577996.bhd
-rw-r----- 1 root root 64M Jan 25 09:29 /vx/D3_4132244/12283/12577996.bin
```

Figure 8. Example View of the Data Stored on VDD Pool

Traditional Duplication Data Flow

For data that are not deduplicated, a copy of the backup images is sent to Access appliance. Depending on policies and/or SLP defined, examples of traditional duplication to Access appliance are as follows and pictured in Figure 9:

A. Data are first stored in an advanced disk on a media server for short to mid-term retention and later copied to Access. Restores can be done from either the advanced disk (Copy 1, the default) or from Access S3 bucket (Copy 2).

B. Data are sent from clients to media server to create a backup images and directly copied to the Access Appliance S3 bucket. Restores from Access (Copy 1) goes through media server to assemble and then sent to client.

Data is sent to Access Appliance using the S3 OST Cloud plugin to send data to Access S3 bucket.

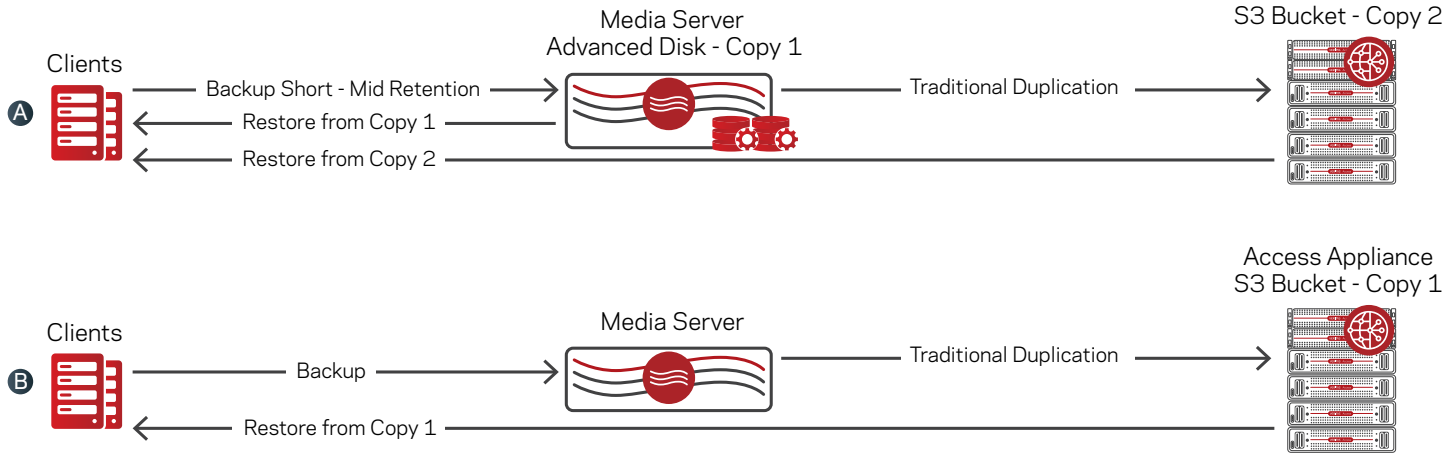

Figure 9. Traditional Duplication Data Flow from NetBackup to Access Appliance

A sample view of data on an Access S3 bucket after traditional duplication is shown in Figure 10. Backup images and associated header information are stored in a directory structure. Inside each directory are image properties, block map file, and the actual image. The header directory contains the header information, properties of header information, and block map file for the header. The S3 OST Cloud plugin will send data to Access S3 bucket into configurable fixed object sizes.

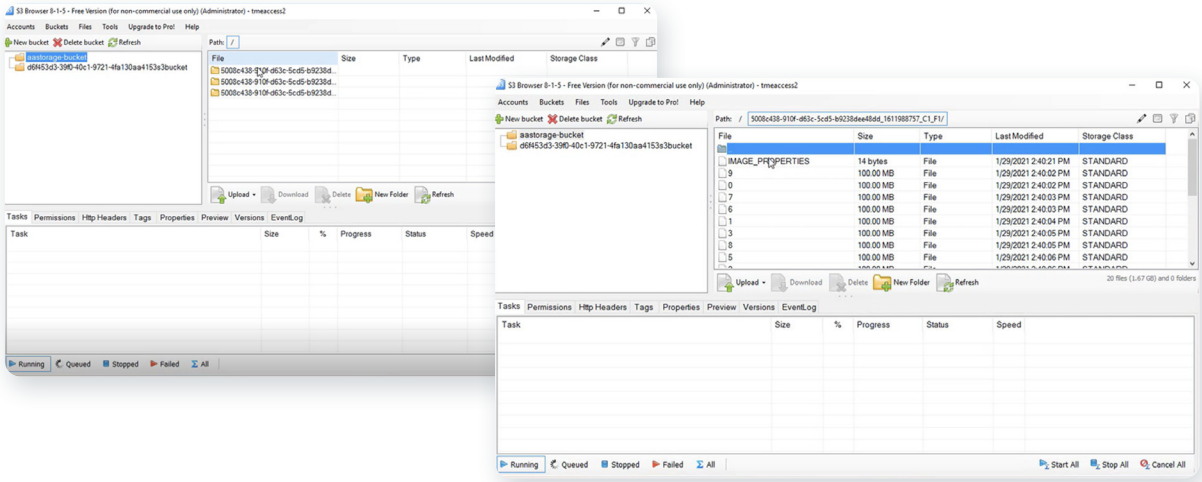

Figure 10. Example View of Data on Access Appliance After Traditional Duplication.

Disaster Recovery

Having a disaster protection plan is imperative for business continuity. The NetBackup catalogs and data are key components to protect. NetBackup Auto Image Replication (AIR) is a mechanism in which NetBackup replicates required data from NetBackup domain to another NetBackup domain for disaster recovery purposes, assuming the domains are configured in separate geographical locations. At the replicated site, the data is first replicated to an MSDP on media server and then optimized duplicated to the Access Appliance. Each domain should contain an Access Appliance with a bucket or VDD pool defined. The size of the target bucket and VDD pool should be the same size or greater than the source.

NOTE: Replication using Auto Image Replication is performed asynchronously by the storage services.

AIR is only one approach to protect against storage failure and site loss in a NetBackup environment with Access Appliance and allows for rapid recovery. Protection of the NetBackup catalog is critical, and it is recommended to do regular backups of the catalogs to protect against corruption or accidental deletion. More information on other ways on how to protect NetBackup environments from disaster can be found in the Veritas NetBackup in Highly Available Environments Administrator's Guide.

The following sections illustrate sample flows of NetBackup AIR configuration where Access Appliance is used for long term retention in the deployment. NOTE: A single Access Appliance cannot store or act as both the source and destination of AIR. It is expected that the destination of the replication is a different Access Appliance.

AIR and Optimized Duplication (MSDP-C) to Access S3 Bucket

An a MSDP on a media server needs to be configured on both source and target for the replication. Storage lifecycle polices (SLP) is where replication can be scheduled to occur after backup and/or duplication. A sample flow of NetBackup AIR where MSDP-C and Access S3 bucket is utilized involves (also illustrated in Figure 11):

1. A trust relationship is established between the NetBackup servers in the domains where credentials and certificates are required for authentication.

2. A sample SLP on the source is defined as follows:

   a) Backup to an MSDP local disk storage pool (Copy 1)

   b) Duplicate to Access S3 bucket via MSDP-C (Copy 2)

   c) Replicate to MSDP on target (Domain B). There is no rehydration of optimized data.

3. An SLP on the target is defined as follows:

   a) Import to an MSDP local disk storage pool (Copy 1)

   b) Duplicate to Access S3 bucket via MSDP-C (Copy 2)

4. The primary server at the target domain automatically creates the entries in the NetBackup catalog as the data is being imported.

NOTE: NetBackup AIR is only occurring between the MSDP on the media servers on source and target. The media server is responsible for the control and orchestration of the entire lifecycle policies in each domain. Each site maintains its own metadata and fingerprint database and thus the fingerprints are compared at the target site and only unique data is sent to target Access Appliance.

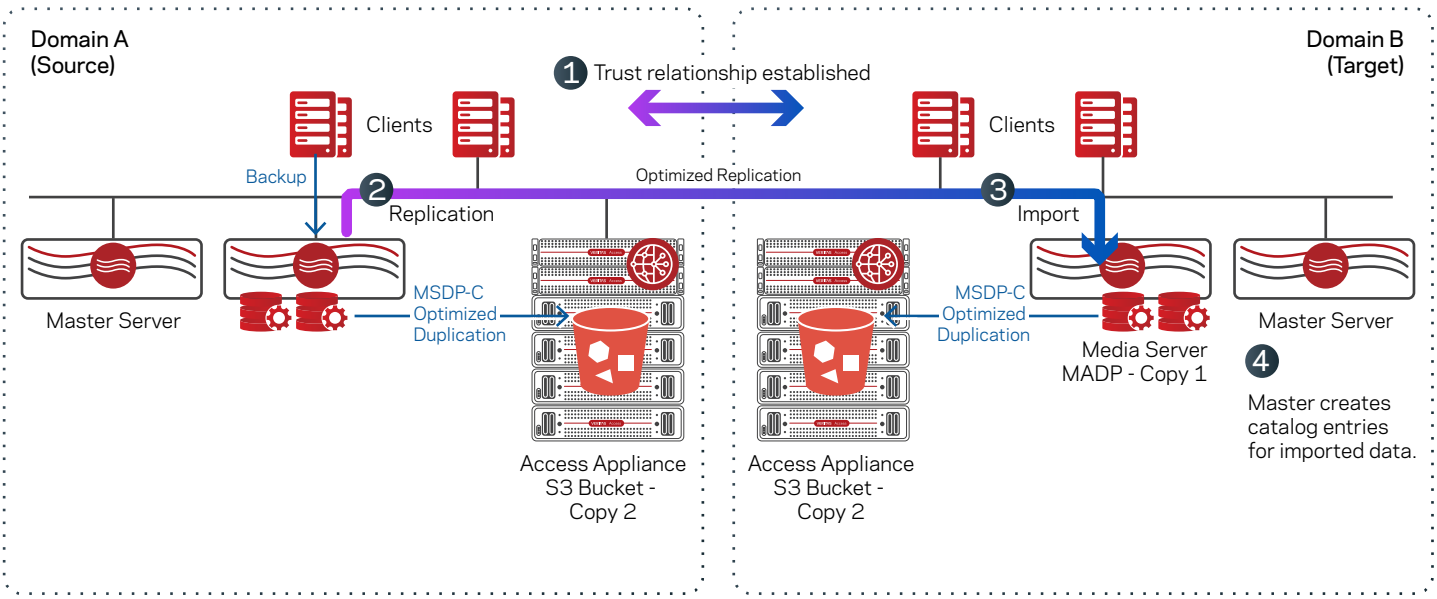

Figure 11. Sample Flow of NetBackup AIR with MDSP-C and Access S3 Bucket in the Environment

AIR for Veritas Data Deduplication (VDD)

When using VDD with AIR, the Access Appliance acts as a duplication target for MSDP on target domain. After the replication and import processes are done by the media servers MSDP, the data is duplicated to Access Appliance data deduplication pool. A sample flow of NetBackup AIR with VDD illustrated in Figure 12 is as follows:

1. A trust relationship is established between the NetBackup servers in the domains where credentials and certificates are required for authentication.

2. A sample SLP on the source is defined as follows:

   a) Backup to an MSDP local disk storage pool (Copy 1 on source)

   b) Duplicate from MSDP to Access Appliance deduplication pool (Copy 2).

   c) Replicate to target MSDP local disk storage pool (Copy 1) to MSDP local disk storage pool (Copy 1) in Domain B.

The media server will initiate the backup and replication processes on source. There is no rehydration of optimized data.

1. A sample SLP on the target is defined as follows:

   a) Import to an MSDP local disk storage pool (Copy 1)

   b) Duplicate from MSDP local disk storage pool to VDD pool (Copy 2 on target)

   The media server will initiate the import process on target.

2. The media server receives optimized data sent, compares the fingerprints in its cache, and stores only unique data on disk and then it is duplicated to Access Appliance where fingerprints are also checked, and only unique data is stored. The fingerprint and metadata (catalogs) on the target side is updated as data is being written.

3. The primary server at the target domain automatically creates the entries in the NetBackup catalog as the data is being imported.

NOTE: Each site maintains its own metadata and NetBackup catalogs. Metadata and catalogs are updated as data is imported.

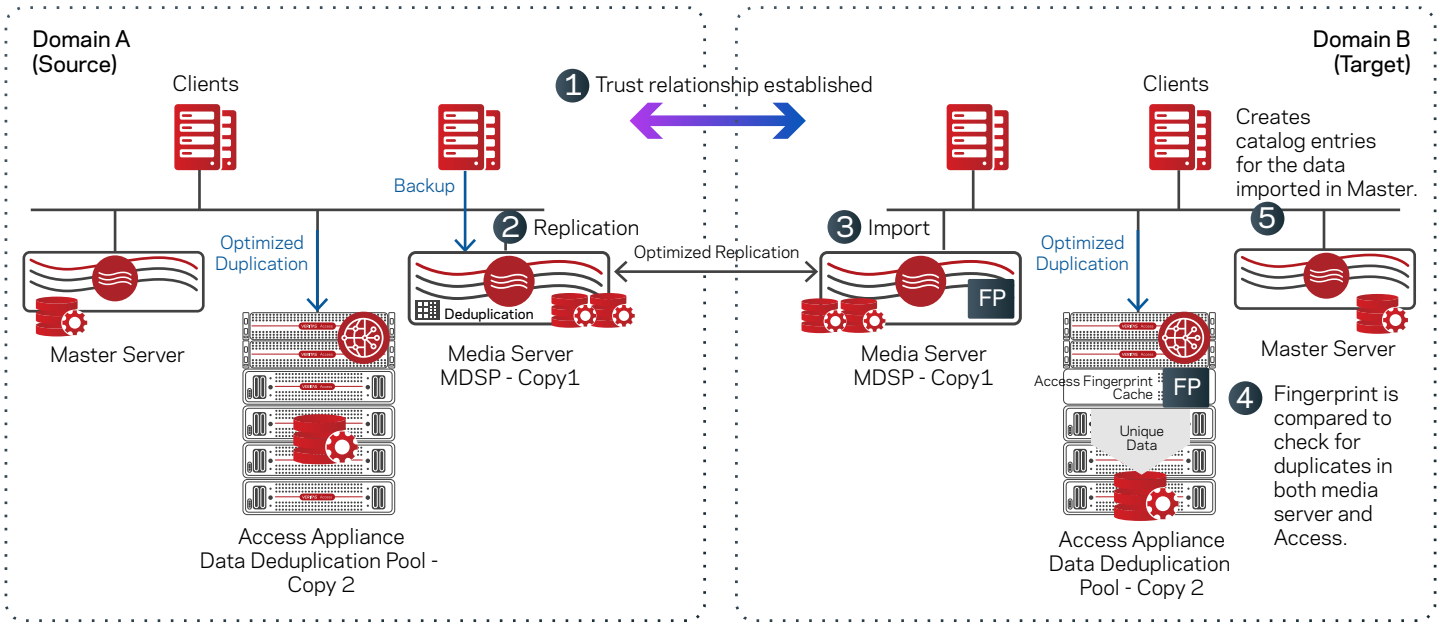

Figure 12. NetBackup AIR with VDD

Another example of AIR with VDD is shown in Figure 13. In this example, there is only an Access Appliance in the target domain. In this example, the sample flow is as follows:

1. A trust relationship is established between the NetBackup servers in the domains where credentials and certificates are required for authentication.

- A sample SLP on the source is defined as follows:

  a. Backup to an MSDP local disk storage pool (Copy 1 on source)

  b. Replicate from source MSDP local disk storage pool (Copy 1) in Domain A to target MSDP Pool (Copy 1) in Domain B.

  The media server will initiate the backup and replication processes on source. There is no rehydration of optimized data.

2. A sample SLP on the target is defined as follows:

   a. Import to an MSDP local disk storage pool (Copy 1) in Domain B

b. Duplicate from MSDP local disk storage pool to VDD pool (Copy 2 on target)

The media server will initiate the import process on target.

3. The media server receives optimized data sent, compares the fingerprints in its cache, and stores only unique data on disk and then it is duplicated to Access Appliance where fingerprints are also checked, and only unique data is stored. The fingerprint and metadata (catalogs) on the target side is updated as data is being written.

4. The primary server at the target domain automatically creates the entries in the NetBackup catalog as the data is being imported.

NOTE: As in previous example, each site maintains its own metadata and NetBackup catalogs. Metadata and catalogs are updated as data is imported.

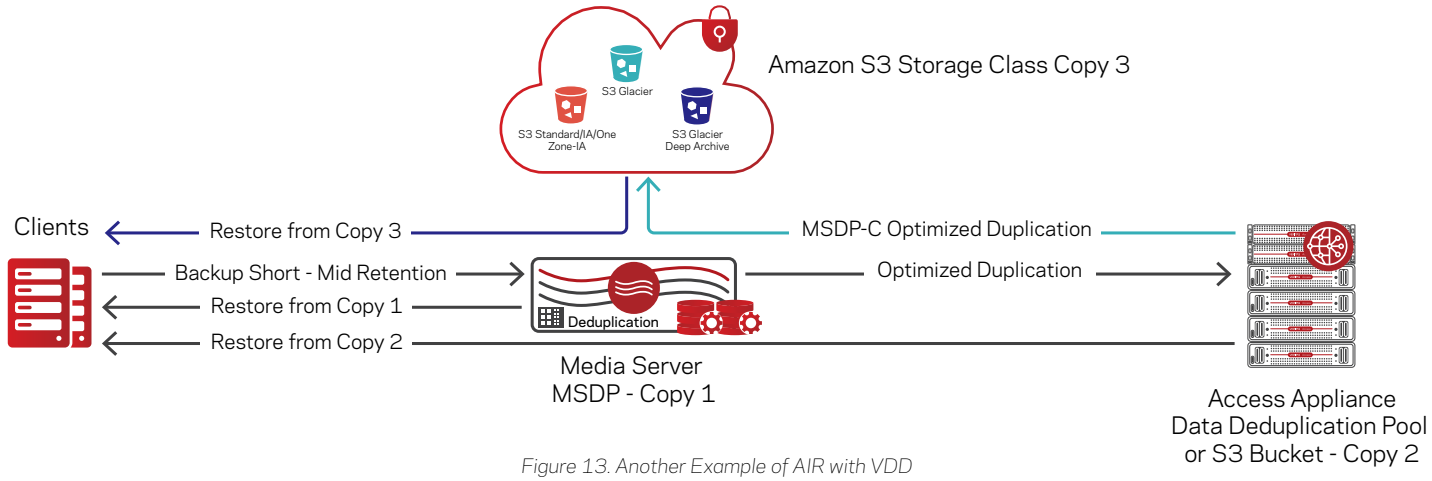

Figure 13. Another Example of AIR with VDD

AIR With Access S3 Bucket (Without Deduplication)

Using NetBackup AIR is not supported for the S3 connector/plugin and for advanced disk as a replication target. One can protect the data on Access Appliance by having several copies of the data and NetBackup catalog on-premises or one can duplicate the data and NetBackup catalogs to a remote site.

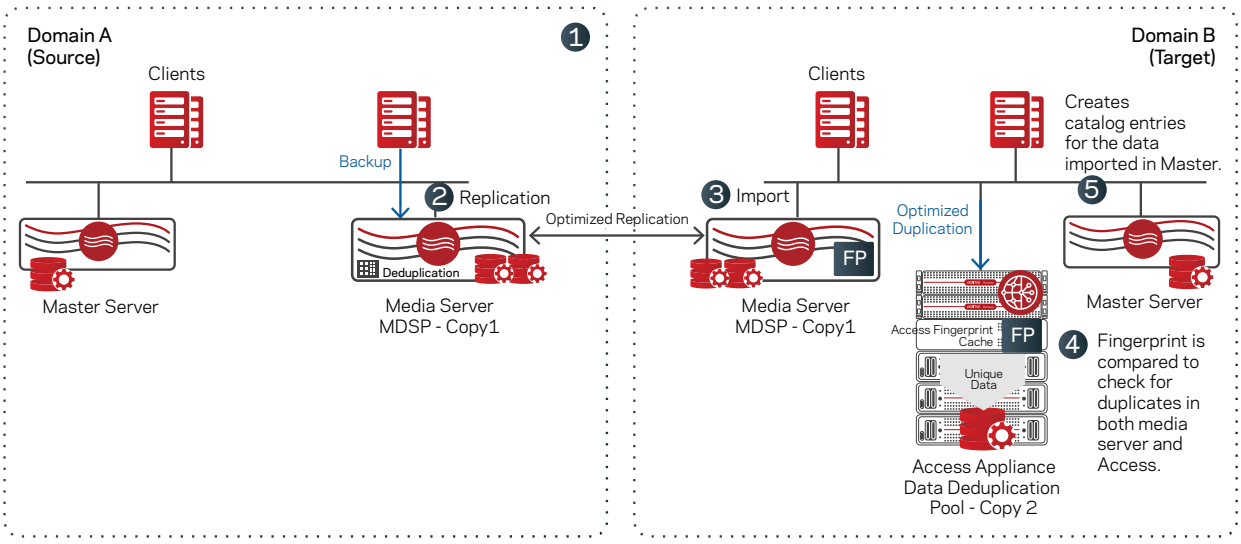

Figure 14.

NOTE: When duplicating non-deduplicated data to a remote site within the same domain, keep in mind that transferring large amounts of data can consume the network bandwidth and may take a long time.

Cloud Support

When sending MSDP data from Access Appliance to the cloud, one can duplicate data to the cloud utilizing NetBackup SLP policies and MSDP-C. The SLP would specify to duplicate the data from Access Appliance using MSDP-C to duplicate the data to the cloud. Figure 17 provides a view of how data is sent to the cloud from VDD or S3 bucket. In this example, an optimized duplication is sent to the Access Appliance and then it does an optimized duplication to the cloud via MSDP-C. The role of the media server during the optimized duplication is to control and orchestrate the transfer between Access Appliance and MSDP-C. The actual I/O is between the Access Appliance and MSDP-C. A restore can be done either from Access (Copy 2) or from public cloud (Copy 3). By default, copy 1 is used for restores unless the restore job specifies otherwise.

Best Practices and Recommendations

Following best practices is important in creating an optimum deployment. This section covers some best practices relating to the Access Appliance as a long-term retention solution for NetBackup.

Data Layout on Access Appliance

Access Appliance provides resiliency at the level of storage devices by using RAID (redundant array of independent disks) controllers in its first shelf of disks. The disks are configured to be multiply redundant using a RAID 6 layout, so that a volume consisting of 16 disks can lose 2 and continue operating without data loss.

The redundant hardware RAID controller in the primary storage shelf configures and presents the shelves' physical disks into disk groups (volumes) protected by a RAID 6 storage layout. With a RAID 6 configuration, data with dual parity is striped across the configured volumes (5 volumes per storage shelf, with each volume containing 16 disks). Each data volume can remain operational despite two concurrent disk failures.

One or more filesystems will be created based on the size of the deduplication pool defined; a pool that is configured to be over 100 TB in size will result in multiple filesystems of 100 TB that logically belong to the same pool. For ease of consistency checking and concurrency, the Veritas Data Deduplication software on Access creates and manages these filesystems. Data is distributed equally among the file system, and the fingerprint database and metadata will be held in a separate filesystem.

Access can dynamically expand the storage available to its filesystems as shelves are added in sequence (up to a total of 4 per Access Appliance deployment).

Deduplication

If more than 250 TB of MSDP is required, use the NetBackup appliances as opposed to the BYOS version. NetBackup BYOS has a limitation of 250 TB for the size of MSDP, whereas the size of MSDP on an appliance can be up to 442 TB or 1 PB depending on the model. The maximum MSDP capacity depends on the NetBackup appliance used as a media server.

For MSDP-C, multiple buckets are supported per storage server. By default, the storage server defines cache properties in the contentrouter.cfg file with a default total of 1 TB. Hence, as a rule of thumb, when configuring the storage server, configure 1 TB of disk storage per cloud target. Also, it is best practice to assign a bucket to each MSDP-C storage server.

As previously stated, Veritas Data Deduplication runs in an active/passive configuration on one of an Access deployment's two nodes. If an Access Appliance deployment will also support other workloads, it is recommended to distribute the workload by using the passive node for other workloads by situating endpoint virtual IP addresses on that node.

In general, for resiliency purposes, Access Appliance should be sized appropriately to handle multiple workloads and single-node failure scenarios.

Another consideration when utilizing VDD is not to send daily incremental backups but instead just duplicating weekly full and/or monthly backups to VDD. Expiration of numerous daily incremental backups from the Access Appliance can overload the garbage collection process and cause performance degradation.

NOTE: The size of the VDD pool can be increased but cannot be decreased.

Compression

For better storage utilization, using NetBackup compression might be an option when deduplication is not ideal, and the data type being backed up is compressible. Although compression can reduce the size of a backup, it can consume server resources. As a best practice, the media server should be sized appropriately for compression. For detailed information on NetBackup compression attributes and considerations, refer to the NetBackup Administration Guide, Volume I, and compression for cloud storage targets and deduplication, refer to NetBackup Cloud Administrators Guide and NetBackup Deduplication Guide respectively. NOTE: For deduplicated data, compression is enabled by default and different from the NetBackup compression policy attribute.

Encryption

NetBackup encryption can be enabled at the NetBackup attribute policy level or at the deduplicated storage pool level via the pd.conf file. By default, encryption is off in both levels. If enabling encryption at the NetBackup attributes policy, data will be encrypted prior to the deduplication which would reduce the deduplication rate. However, enabling encryption at the MSDP level will not affect overall deduplication rates since the data is encrypted after deduplication. NOTE: NetBackup KMS is not supported with the Access Appliance.

Network Connectivity

An Access Appliance 3340 has 2 high-speed (10/25 GbE) ports per node, and an Access Appliance 3350 has either 2 or 4 high-speed ports. As a best practice present the fully qualified domain name mapping to the virtual IP so it will automatically transition to the other node if one node fails or the physical links on one node fails or is unreachable. Using the fully qualified domain name as opposed to the virtual IP is beneficial if the virtual IP changes for scenarios like migration. For instance, map one of the virtual IPs to the S3 object URL s3.<clustername> when using the S3 protocol. When using Veritas Data Deduplication, you should specify the fully qualified domain name (FQDN) mapping to Access Appliance virtual IP during NetBackup configuration.

Data port bonding is also an option on Access Appliance. Joining or bonding multiple network interfaces on the Access appliances into a single interface improves the bandwidth and network throughput through the combined single interface. As a best practice, the switch that the uplinks of the Access Appliance are connected to must be configured appropriately for the link aggregation being used.

Multiple NetBackup Domains

As previously mentioned, the Access Appliance can store data from multiple NetBackup domains. When using MSDP-C with Access, a single bucket can be used to store data for both domains. However, it is advisable when using MSDP-C with Access, to have each NetBackup domain with their own bucket to avoid name collisions in the different domains and for better categorization and identification. However, this best practice is not to be applied when using VDD in which a single data deduplication pool is configurable and global deduplication is observed between the domains.

NOTE: As previously mentioned, when using AIR with VDD, the Access Appliance should not be both the source and destination. A separate Access Appliance is needed in each domain.

Load Balancing for MSDP-C

An Access Appliance deployment consists of two nodes in an active/active configuration. For services such as NFS, CIFS, and S3, the use of multiple IP addresses can serve to separate and/or balance workloads across an Access deployment's nodes. Load balancing can be achieved using any of the following techniques:

- External load balancing – using an external load balancer such as HAProxy or F5, allows for more algorithms to distribute load across nodes such as least connections or weights. It also frees the Access nodes from the proxy handling and balances the network traffic between the nodes.
- Manual load balancing – virtual IP addresses of the nodes can be manually assigned to applications in a distributed manner. The disadvantage of this approach is that even distribution might be difficult to gauge since applications are not all equal in sense of workload.
- DNS load balancing – the S3 object URL name for an Access S3 bucket, s3.<clustername> is created in DNS and includes all the virtual IP addresses of the nodes. DNS round-robins through the virtual IP addresses. The disadvantage of using DNS is in case of connectivity issues, the virtual IP is still in rotation until it is manually removed.

NOTE: Load balancing does not apply to Veritas deduplication services running on Access.

Monitoring

It is important to monitor or be aware of the alerts especially storage utilization warnings and hardware critical alerts. The AutoSupport features assists in this manner, but as a best practice, it is advisable to be pro-active instead of re-active. For instance, once the capacity reaches 60%, it might be a good time to revisit the storage utilization or plan for growth.

Sizing Guidance

In planning for data protection, two considerations come to mind: recovery point objectives (RPO) and recovery time objectives (RTO). From a backup and recovery standpoint, the RPO and RTO determine which policies are implemented, and therefore the resources required by a NetBackup deployment in terms of the necessary systems, appliances, and storage. Other considerations include the number of users and applications, amount of data that is backed up, the frequency, and how long to keep the data. When planning for a long-term retention solution for backup images there are two factors:

- Capacity – how many backup images can be stored
- Performance – how much workload (backup streams and bandwidth) the storage platform can handle

Veritas can assist in the sizing of the appliance based on your requirements using these factors. Some parameters that might enter in the equation when estimating long-term storage requirements include:

1. Volume of source data.
2. Daily data change ratio
3. Annual storage growth
4. Retention period for daily and incremental backups
5. Retention for weekly, monthly, and yearly full backups
6. Estimated deduplication ratio for initial backup and daily incremental
7. Estimated deduplication ratio for weekly, monthly, and yearly full backups
8. Performance and/or service level requirements

Samples of Capacity Sizing for Access Appliance

Deduplication affects the actual used capacity of a long-term retention archive in multiple ways. In addition to the storage of image data, overhead is incurred by deduplication functionality such as the creation of an encapsulating directory structure, metadata storage, and a fingerprint or hash table. This overhead should be accounted for when calculating available archive capacity.

NOTE: The sample calculations in this section deal with overhead relating to capacity, leaving out other parameters previously discussed such as annual growth rate, daily change rate, performance, etc. It is best to work with the Veritas Account teams who have the tools and expertise to determine the optimum configuration based on your requirements.

Sizing for MSDP-C

NetBackup MSDP deduplication technology places the backup images into 64 MB containers. Using the default block size of the Access Appliance for the cluster file system is 8 KB. The overhead of Access, NetBackup header, and keys (if encryption is enabled) for unique data includes:

- Without encryption enabled: 0.18%
- With encryption enabled: 0.26%

Example without encryption for 800 TB of data with an 8:1 deduplication ratio is as follows:

- Logical data stored on disk: 800 TB
- Unique data stored on disk: 100 TB
- Access and NetBackup overhead (no encryption):
  - 0.18% × 100 TB = 0.18 TB
- Total volume storage requirements = unique data + overhead (Access and NetBackup)
  - 100 TB + 0.18 TB = 100.18 TB

Example with encryption for the same size:

- Access and header overhead (with encryption):
  - 0.26% × 100 TB = 0.26 TB
- Total Volume Storage Requirements = Unique Data + Access and NetBackup Header and Encryption Overhead
  - 100 TB + 0.6 TB = 100.26 TB

NOTE: The above calculation does not include the MSDP meta-data on Access Appliance.

Sizing for Veritas Data Deduplication

For VDD, the fingerprint database, the metadata, unique blocks, journals and logs are stored on Access. If defining an 800 TB data deduplication pool, eight 100 TB file systems is created. Each file system created incurs a 0.1% file system overhead, thus, with eight file systems, the maximum would be 0.8%. In addition, the database, the metadata, logs and journals consume up to 4% of the storage. Example for an 800 TB of data to backup with 8:1 deduplication ratio is as follows:

- Fingerprint database, metadata, journal, and logs overhead: 4%
- File system overhead: 0.8 % for 800 TB filesystem
- Logical data stored on disk: 800TB
- Unique data stored on disk: 100TB
- (Fingerprint database, metadata, journal, logs overhead) = 4% × (logical data stored on disk)
  - 4% × 800 TB=32 TB
- File system overhead for one file system = 0.8% × Unique Data
  - 0.8% × 100TB = 0.8 TB
- Total volume storage requirements = (unique data) + (fingerprint database, metadata, journal, logs overhead) + (file system overhead)
  - 100 TB + 32 TB + 0.8 TB = 132.8 TB

NOTE: The above calculation does not include the 5 TB catfs filesystem created for the Veritas Data Deduplication catalog on Access when initially configuring storage and starting the service.

Sizing for Traditional Duplication of Data

For traditional duplication, as previously explained, the S3 OST cloud plugin shards the backup image into 16 MB fixed object and sends it to Access. There is one header information per 50 GB of backup image and keys if encryption is enabled. The overhead values are miniscule and thus, the size of backup image for traditional duplication is about the same size of backup image being stored. Take for example a 150 GB, Access and NetBackup header and keys (if encryption is enabled) overhead are:

- Without encryption enabled: 0.000082%
- With encryption enabled: 0.00014%

Example without encryption for 150 GB backup image size:

- Data stored on disk 150 GB
- Access and with Overhead = 0.000082% × data size
  - 0.000082% × 150 GB = 0.000012 GB
- Total volume storage requirements = Data + Access and NetBackup Header Overhead
  - 150 GB + 0.000012 GB = 150 GB

Example with encryption for the same size:

- Access, Header and Encryption Overhead = 0.00014%
  - 0.00014% × 150 GB = 0.00021 GB
- Total volume storage requirements = Data + Access and NetBackup Header and Encryption Overhead
  - 150 GB+0.00021 GB = 150 GB

Conclusion

The addition of an Access Appliance to a Veritas appliance-based solution provides a competitive, ransomware-resilient, disaster recovery-capable method of long-term retention of data. Integrated with NetBackup, it becomes an even more compelling option in data protection, disaster planning and recovery. Implementing Access Appliance with NetBackup as a long-term retention solution simplifies management and support, minimizes costs, and improves control and visibility.

.

References

- Veritas Access Appliance Administrator's Guide
  https://sort.veritas.com/DocPortal/pdf/146126550-151725528-1
- Veritas Access 3340 Appliance Product Description
  https://sort.veritas.com/DocPortal/pdf/125460431-134247411-1
- Veritas Access 3350 Appliance Product Description
  https://sort.veritas.com/DocPortal/pdf/125460431-154188696-1
- Veritas Access Appliance Initial Configuration Guide
  https://sort.veritas.com/DocPortal/pdf/129305376-151823867-1
- Veritas Access Appliance Solutions Guide for NetBackup
  https://sort.veritas.com/DocPortal/pdf/146127092-151725537-1
- Veritas NetBackup in Highly Available Environments Administrator's Guide
  https://sort.veritas.com/DocPortal/pdf/39129704-152913080-1
- Veritas NetBackup Deduplication Guide
  https://sort.veritas.com/DocPortal/pdf/25074086-151874762-1
- Veritas NetBackup Security and Encryption Guide
  https://sort.veritas.com/DocPortal/pdf/21733320-149123528-1
- Veritas NetBackup in Highly Available Environments Administrator's Guide
  https://sort.veritas.com/DocPortal/pdf/39129704-152913080-1
- Veritas NetBackup Administrator's Guide, Volume I
  https://sort.veritas.com/DocPortal/pdf/18716246-151251642-1
- Veritas NetBackup Cloud Administrator's Guide
  https://sort.veritas.com/DocPortal/pdf/58500769-152291633-1

About Veritas

Veritas Technologies is a global leader in data protection and availability. Over 80,000 customers—including 95 percent of the Fortune 100—rely on us to abstract IT complexity and simplify data management. The Veritas Enterprise Data Services Platform automates the protection and orchestrates the recovery of data everywhere it lives, ensures 24/7 availability of business-critical applications, and provides enterprises with the insights they need to comply with evolving data regulations. With a reputation for reliability at scale and a deployment model to fit any need, Veritas Enterprise Data Services Platform supports more than 800 different data sources, over 100 different operating systems, more than 1,400 storage targets, and more than 60 different cloud platforms. Learn more at www.veritas.com. Follow us on Twitter at @veritastechllc.

VERITAS™

2625 Augustine Drive
Santa Clara, CA 95054
+1 (866) 837 4827
veritas.com

For global contact information visit:
veritas.com/company/contact

Copyright © 2022 Veritas Technologies LLC. All rights reserved. Veritas, the Veritas Logo, and NetBackup are trademarks or registered trademarks of Veritas Technologies LLC or its affiliates in the U.S. and other countries. Other names may be trademarks of their respective owners.

V1283 08/22