



Technical Report

NetApp E-Series storage with Commvault Complete Data Protection

Reference architecture and storage best practices

Mitch Blackburn and Alonso DeVega, NetApp; Eric Nelson and Bryan Clarke, Commvault
December 2022 | TR-4950

In partnership with



Abstract

This document outlines the reference architecture and best practices when using NetApp® E-Series storage in a Commvault® Complete Data Protection environment.

TABLE OF CONTENTS

Executive summary	4
Introduction.....	4
About NetApp	4
About Commvault.....	5
Reference architecture overview	6
Solution configuration options.....	6
Backup storage capacity sizing	17
Backup read performance goals.....	17
NetApp E-Series arrays	17
E2800 and E5700 hybrid arrays	17
EF300 and EF600 all-flash arrays	18
SANtricity OS	19
SANtricity security features.....	19
SANtricity reliability features.....	19
SANtricity data protection features	20
Dynamic Disk Pools (DDP).....	21
Commvault Complete Data Protection	21
Overview	21
Commvault Complete Data Protection components	22
Commvault Complete Data Protection design considerations and best practices	24
Infrastructure design.....	24
NetApp E-Series volume configuration guidelines	26
NetApp E-Series storage configuration guidelines for Commvault Complete Data Protection disk libraries	26
Architectural tools and sizing	27
Assessment tools	27
Software Configurator design tool	28
Summary.....	28
Appendix.....	28
Commvault Complete Data Protection terminology and definitions	28
Where to find additional information	31

Version history.....	31
-----------------------------	-----------

LIST OF TABLES

Table 1) Different configurations of reference architecture.....	6
Table 2) Small configuration: single node.....	7
Table 3) Medium configuration: single node.....	9
Table 4) Large configuration: single node.....	10
Table 5) Extra-large configuration: single node.....	12
Table 6) Extra-large configuration: multinode.....	14
Table 7) Backup storage capacity sizing.....	17
Table 8) Backup read performance goals.....	17
Table 9) E2800 and E5700 controller shelf and drive shelf models.....	18
Table 10) SANtricity features for long-term reliability.....	19
Table 11) Commvault Complete Data Protection components and descriptions.....	23
Table 12) Deduplication terminology.....	25
Table 13) Commvault terminology descriptions.....	28

LIST OF FIGURES

Figure 1) Small configuration: single node	7
Figure 2) Medium configuration: single node.....	8
Figure 3) Large configuration: single node.....	10
Figure 4) Extra-large configuration: single node.....	12
Figure 5) Extra-large configuration: multinode.....	14
Figure 6) Cross-site replication.....	16
Figure 7) Long-term retention.....	16
Figure 8) Commvault Command Center dashboard.....	22
Figure 9) Commvault data management platform.....	24

Executive summary

Streamlined backup and recovery solutions are not on customers' wish lists these days; they are a business imperative. Meeting shrinking backup windows while protecting data is a challenge that all customers struggle with as their need to store, manage, and manipulate that data grows exponentially, whether for competitive gain or for compliance and regulatory mandates. Fast and consistent recovery is paramount.

NetApp and Commvault offer a simple and integrated data storage and data protection solution for custom data fabrics. This extends from core to edge on premises, hybrid or multicloud environments. It gives the ability to deliver business outcomes with risk reduction and security, cost efficiency, operational excellence, and enterprise agility. Unlike point product and appliances, it delivers best-in-class cost efficiency, operational excellence, and protection of business-critical applications, regardless of where the data lives.

This document outlines a reference architecture for enabling a collaborative backup and recovery solution on NetApp E-Series with Commvault Complete™ Data Protection software.

Introduction

Commvault and NetApp have jointly developed this reference architecture to provide guidance for Commvault deployments with NetApp E-Series storage that accelerate time to application for this solution.

E-Series and Commvault Complete Data Protection offer a best-in-class backup and recovery solution through a tightly coupled and thoroughly tested reference architecture from industry leaders NetApp and Commvault. This solution is optimized for disk-to-disk backup and recovery relying on the policy-based management and deduplication features of Commvault while providing high-capacity storage performance on flexible NetApp E-Series storage arrays. This solution gives customers superior performance and functionality at a competitive price.

Features include:

- Eliminate tape completely with a second copy off site
- Scalability
 - Front-end scalability with additional Commvault media servers (deduplication)
 - Back-end scalability with online capacity expansion of E-Series arrays (raw capacity)
- E-Series and Commvault data management platform reduce complexity through off-site replication
- Archiving and stubbing minimize future storage capacity needs
- Joint presales and world-class support from NetApp and Commvault
- Efficiency
 - Direct-to-disk backups are “DASH copied” or replicated in deduplicated fashion to a remote site to provide disaster recovery capabilities for all backed-up data
 - Targeted at opportunities that scale from several terabytes to multiple petabytes of primary data under management

About NetApp

NetApp creates innovative products, including storage systems and software that help customers around the world store, manage, protect, and retain one of their most precious corporate assets: their data. We are recognized throughout the industry for continually pushing the limits of today's technology so that our customers don't have to choose between saving money and acquiring the capabilities that they need to be successful.

We find ways to enable our customers to do things that they couldn't do before and at a speed that they never thought possible. We partner with industry leaders to create efficient and cost-effective solutions that are optimized for customers' IT needs and to deliver to and support these customers worldwide. Leading organizations worldwide count on NetApp for software, systems, and services to manage and store their data. Customers value our teamwork, expertise, and passion for helping them succeed now and into the future ([NetApp.com](https://www.netapp.com)).

About Commvault

Commvault is a global leader in data management. Our Intelligent Data Services help your organization do amazing things with your data by transforming how you protect, store, and use it. We provide a simple and unified data management platform that spans all your data – regardless of where it lives (on-premises, hybrid, or multicloud) from legacy to modern workloads. Commvault solutions are available through any combination of software subscriptions, integrated appliances, partner-managed, or Software-as-a-Service via our Metallic portfolio. Over 25 years, more than 100,000 customers have relied on Commvault to keep their data secure, assessable, and ready to drive business growth. Commvault's corporate headquarters is in Tinton Falls, New Jersey. Information about Commvault is available at [Commvault.com](https://www.commvault.com).

Reference architecture overview

This section details reference architectures that range from those of small environments that protect a few terabytes of data to those in enterprise-size environments with petabytes of data under management.

This section provides recommended structures and integrations of products and services to form a solution. It embodies accepted industry best practices, typically suggesting the optimal delivery method for specific technologies. Reference architectures help project managers, software developers, enterprise architects, and IT managers collaborate and communicate effectively about an implementation project. A reference architecture anticipates – and answers – the most common questions that arise, helping teams avoid errors and delays that might occur without the use of a tested set of best practices and solution approaches. The seven different configurations of the reference architecture are listed in Table 1.

Table 1) Different configurations of reference architecture.

Configuration	Description
Express	Single node, small configuration
Work group	Single node, medium configuration
Data center	Single node, large configuration
Data center XL	Single node, extra-large configuration
Data center multi-node	Multi-node, global deduplication across nodes, large and extra-large configuration
Enterprise extended mode	Multiple copies and long-term retention
With Intellisnap	For LAN-free backup copies

Solution configuration options

Express (S) small configuration: single node

- Targeted for small businesses.
- Remote office or branch office with local backups.
- Initial configuration sufficient for typical 30- to 90-day retention.
- Expansion option: For long-term retention, add more expansion shelves and drives.
- Back-end size (BET) limited to 30TB usable capacity for this configuration. Back-end size is the actual capacity used after deduplication and compression. Front-end (FET) capacity is the size of application data that needs to be protected. For more information, refer to the Commvault documentation at <http://documentation.commvault.com>.

Figure 1 gives a graphical representation of a NetApp E-Series and Commvault single node setup and Table 2 gives in-depth configuration information.

Figure 1) Small configuration: single node.

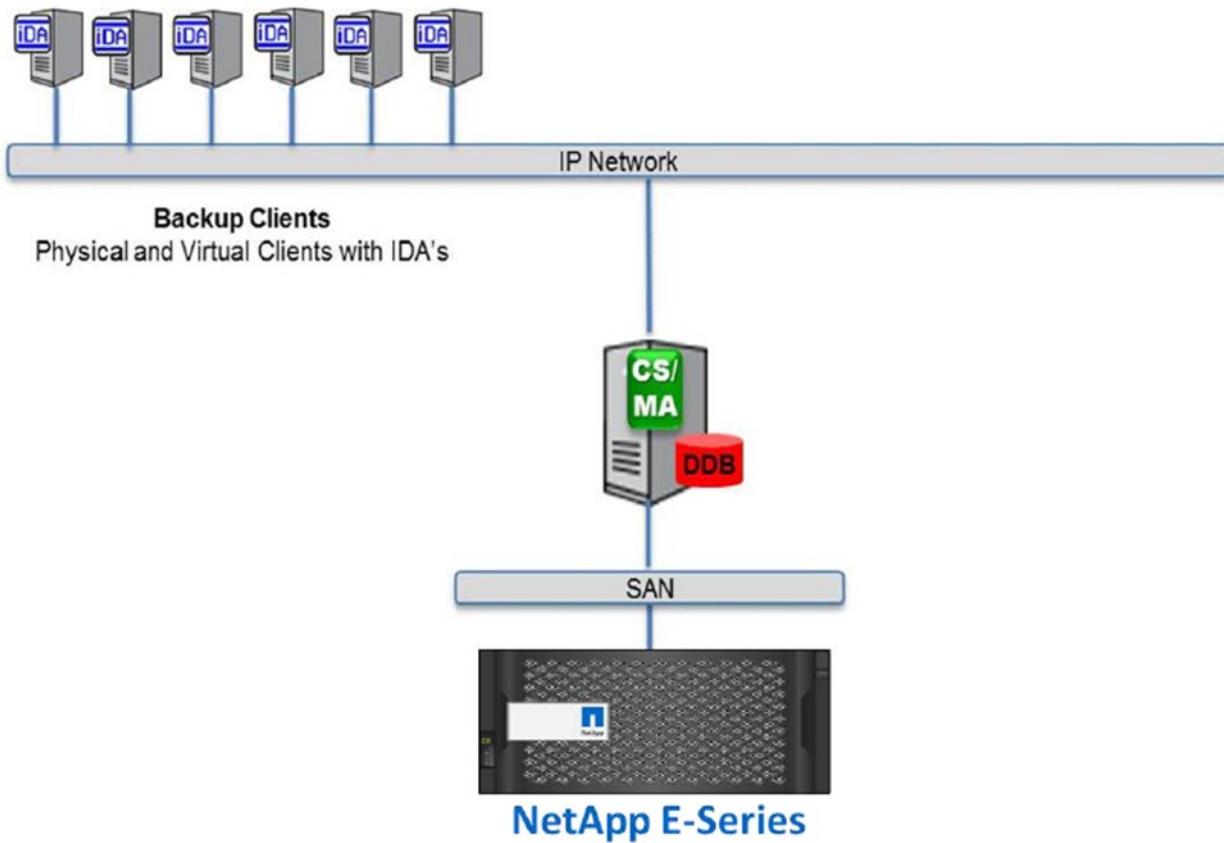


Table 2) Small configuration: single node.

Option	Description
Scope	Around 20TB (FET) of production data based on data type, up to 30TB of back-end size.
Software	OS: 64-bit Windows for CommServe®/MediaAgent (CS/MA) combo OR 64-bit Windows/Linux for MediaAgent only
Commvault	Platform Release 2022E CommServe and MediaAgent OR MediaAgent only

Server configuration (1 Server Required)

CPU/RAM	8 Cores 32GB RAM 2 x 650W power supply
Internal storage	2 x 300GB 15K RPM drives in RAID 1 for OS and Commvault installation. 2 X 300GB 15K RPM drives in RAID 1 for index cache. For IOPS requirements, refer to the Commvault documentation at http://documentation.commvault.com . 2 X 400GB value SSDs in RAID 1 for deduplication database. For IOPS requirements, refer to the Commvault documentation at http://documentation.commvault.com .

Option	Description
NIC	Dual port, 10 GigE card for data ingest and egress Quad, 1 GigE ports for management interface
Host bus adapter cards	Dual port SAS adapter Dual port HBA for tape out (optional)

Back-End Storage Configuration

Storage	NetApp E2800, E5700, EF300, or EF600. For information about these systems and drive shelves, see NetApp E-Series arrays section.
----------------	--

Work group (M) medium configuration: single node

- Targeted for medium businesses and smaller data centers.
- Initial configuration sufficient for typical 30- to 90-day retention.
- Expansion options:
 - For long-term retention, add more expansion shelves and drives.
 - For additional capacity, buy new nodes in the same cell.
- Back-end size (BET) limited to 60TB usable capacity for this configuration. Front-end (FET) capacity can vary from 40 to 50TB, based on data type.

Figure 2 gives a graphical representation of a NetApp E-Series and Commvault single node setup and Table 3 gives in-depth configuration information.

Figure 2) Medium configuration: single node.

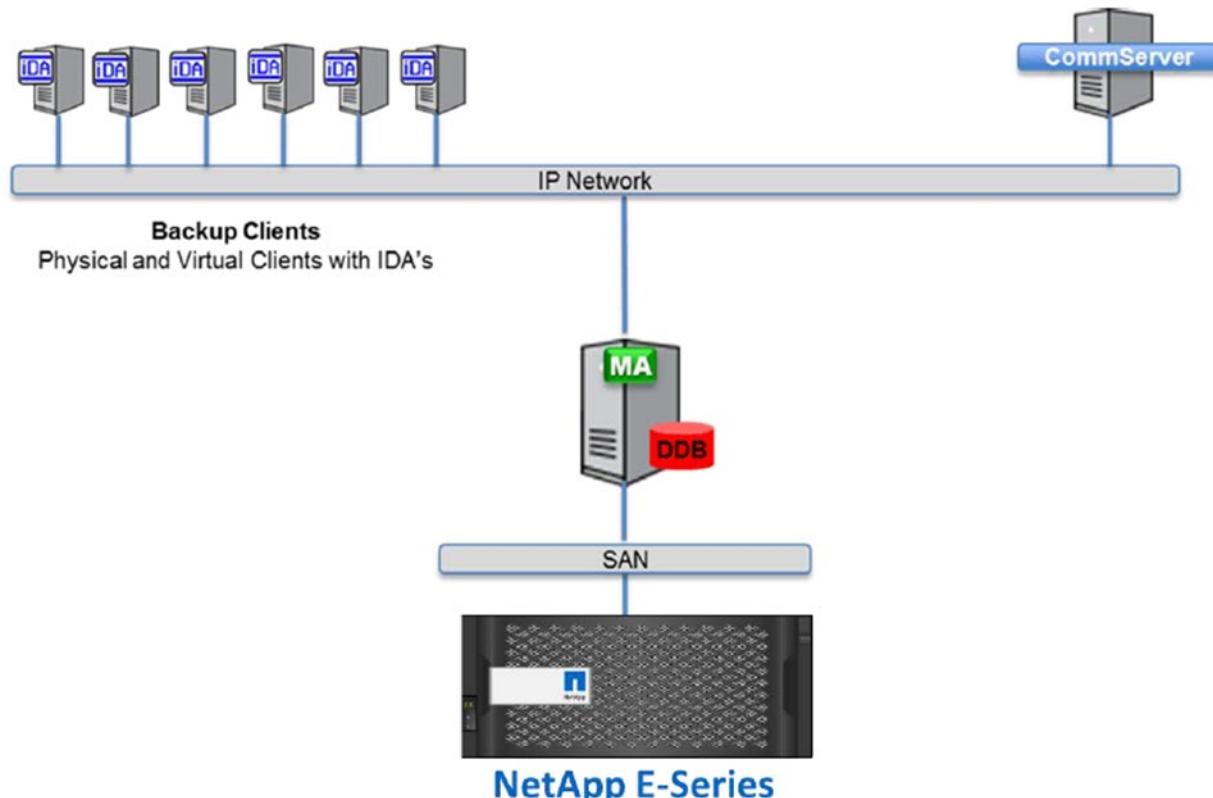


Table 3) Medium configuration: single node.

Option	Description
Scope	40 to 50TB of production data based on data type, up to 60TB of back-end size.
Software	OS: 64-bit Windows for CommServe /MediaAgent (CS/MA) combo OR 64-bit Windows/Linux for MediaAgent only
Commvault	Platform Release 2022E CommServe and MediaAgent OR MediaAgent only

Server Configuration (CS/MA Combo OR Dedicated MA)

Software	OS: 64-bit Windows for CS/MA Combo OR 64-bit Windows/Linux for MediaAgent only
Form factor	2U rack mount with minimum 12 single-form factor (SFF) (2.5") drive bays
CPU/RAM	12 Cores 64GB RAM 2 x 650W power supply
Internal storage	<ul style="list-style-type: none"> • 4 x 400GB+ value SSDs in RAID 5 for OS, Commvault installation, and index cache. For IOPS requirements, refer to the Commvault documentation at http://documentation.commvault.com. • 4 X 400GB+ value SSDs in RAID 5 + 1 hot spare for deduplication database. For IOPS requirements, refer to the Commvault documentation at http://documentation.commvault.com.
NIC	Dual port, 10 GigE card for data ingest and egress Quad 1 GigE ports for management interface
Host bus adapter cards	Dual port SAS adapter Dual port HBA for tape out (optional)

Server Configuration (Dedicated CS – Physical OR Virtual)

Software	OS: 64-bit Windows
CPU/RAM	8 Cores 32GB RAM 2 x 650W power supply
Internal storage	2 x 400+ GB value SSDs in RAID 1 for OS and Commvault installation

Back-End Storage Configuration

Storage	NetApp E2800, E5700, EF300, or EF600. For information about these systems and drive shelves, see NetApp E-Series arrays section.
----------------	--

Data center (L) large configuration: single node

- Targeted for enterprise data centers.
- Initial configuration sufficient for typical 30- to 90-day retention.
- Expansion options:
 - For long-term retention, add more expansion shelves and drives.
 - For additional capacity, buy new nodes in the same cell.

- Back-end size (BET) limited to 150TB usable capacity for this configuration. Front-end (FET) capacity can be up to 100TB, based on data type.

Figure 3 gives a graphical representation of a NetApp E-Series and Commvault single node setup and Table 4 gives in-depth configuration information.

Figure 3) Large configuration: single node.

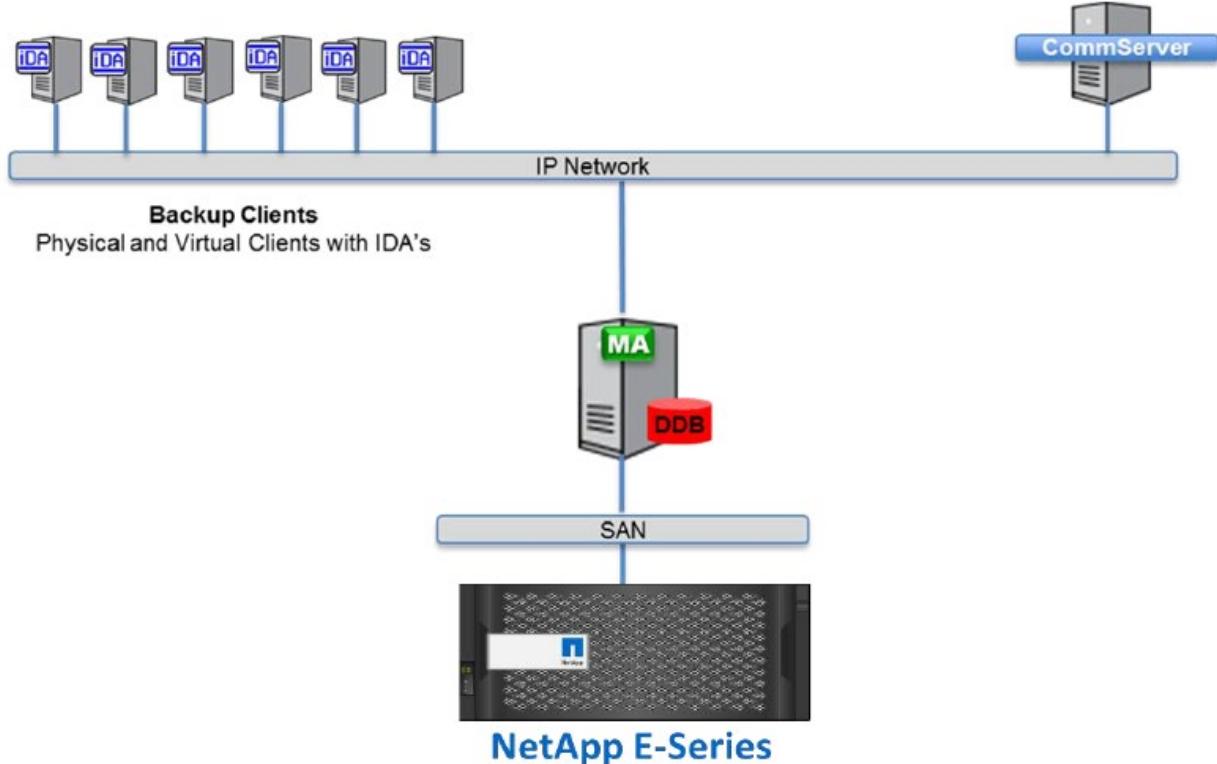


Table 4) Large configuration: single node.

Option	Description
Scope	Around 100TB of production data, based on data type, up to 150TB of back-end size.
Software	OS: 64-bit Windows for CommServe 64-bit Windows/Linux for MediaAgent only
Commvault	Platform Release 2022 CommServe and MediaAgent

Server Configuration (Dedicated MA)

Software	OS: 64-bit Windows/Linux for MediaAgent
Form factor	2U rack mount with minimum 12 SFF (2.5") drive bays
CPU/RAM	12 Cores 64GB RAM 2 x 650W power supply

Option	Description
Internal storage	<ul style="list-style-type: none"> • 400GB usable disk, minimum 4 spindles 15K RPM or higher OR SSD class disk for OS and Commvault installation • 1.2TB usable capacity volume (RAID 1 or RAID 5) SSD class disk/PCIe I/O cards 2GB controller cache memory for index cache. For IOPS requirements, refer to the Commvault documentation at http://documentation.commvault.com. • 1.2TB usable capacity volume (RAID 1 or RAID 5) SSD class disk/PCIe I/O cards 2GB controller cache memory for deduplication database. For IOPS requirements, refer to the Commvault documentation at http://documentation.commvault.com.
NIC	Dual port, 10 GigE card for data ingest and egress Quad 1 GigE ports for management interface
Host bus adapter cards	Dual port SAS adapter Dual port HBA for tape out (optional)

Server Configuration (Dedicated CS – Physical)

Software	OS: 64-bit Windows
Form factor	2U rack mount with minimum 12 SFF (2.5") drive bays
CPU/RAM	8 Cores 32GB RAM 2 x 650W power supply
Internal storage	2 x 400+ GB value SSDs in RAID 1 for OS and Commvault installation

Back-End Storage Configuration

Storage	NetApp E2800, E5700, EF300, or EF600. For information about these systems and drive shelves, see NetApp E-Series arrays section.
----------------	--

Data center (XL) extra-large configuration: single node

- Targeted for enterprise data centers.
- Initial configuration sufficient for typical 30- to 90-day retention.
- Expansion options:
 - For long-term retention, add more expansion shelves and drives.
 - For additional capacity, buy new nodes in the same cell.
- Back-end size (BET) limited to 200TB usable capacity for this configuration. Front-end (FET) capacity can be up to 130TB, based on data type.

Figure 4 gives a graphical representation of a NetApp E-Series and Commvault single node setup and Table 5 gives in-depth configuration information.

Figure 4) Extra-large configuration: single node.

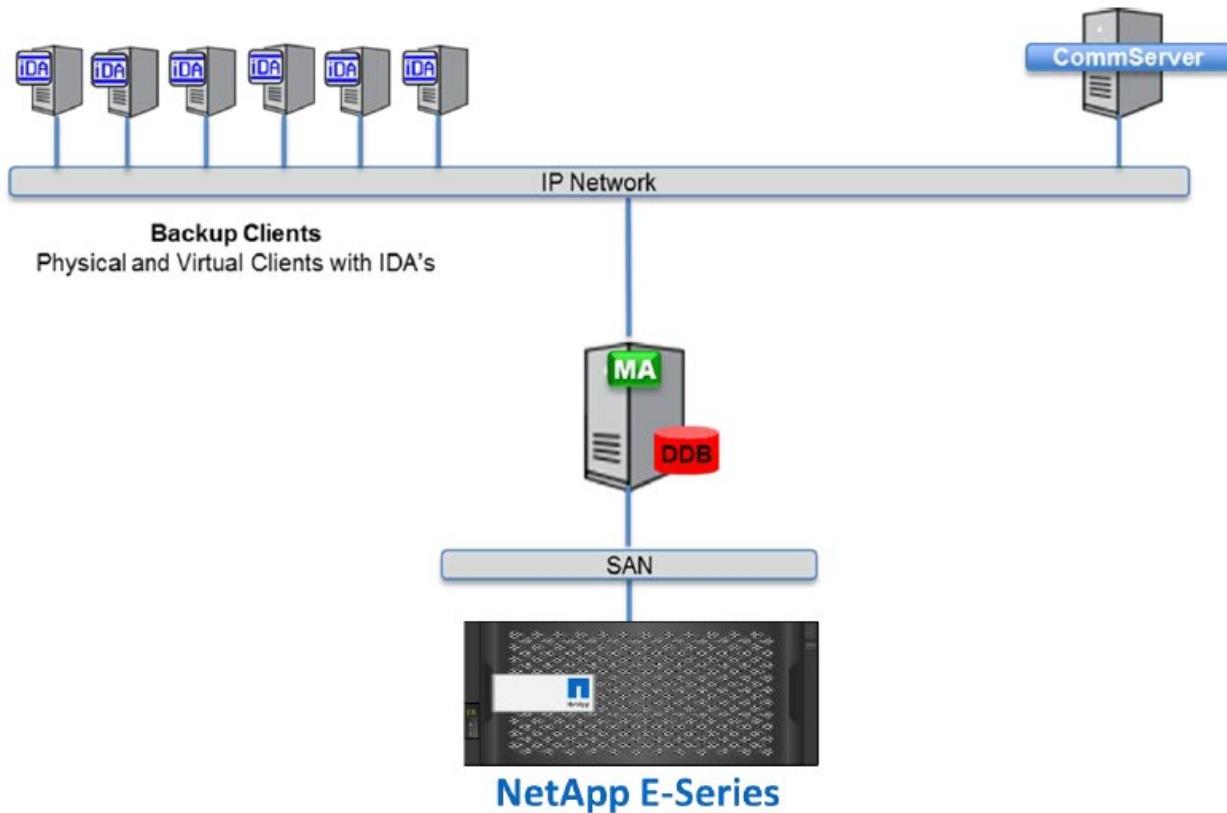


Table 5) Extra-large configuration: single node.

Option	Description
Scope	Around 120TB to 130TB of production data, based on data type, up to 200TB of back-end size.
Software	OS: 64-bit Windows for CommServe 64-bit Windows/Linux for MediaAgent only
Commvault	Platform Release 2022 CommServe and MediaAgent

Server Configuration (Dedicated MA)

Software	OS: 64-bit Windows/Linux for MediaAgent
Form factor	2U rack mount with minimum 12 SFF (2.5") drive bays
CPU/RAM	16 Cores 128GB RAM 2 x 650W power supply
Internal storage	<ul style="list-style-type: none"> 2 x 400GB SSD class disk for OS disk in RAID 1 2TB usable capacity volume (RAID 1 or RAID 5) with SSD class disks/PCIe I/O cards 2GB controller cache memory for index cache database. For IOPS requirements, refer to the Commvault documentation at http://documentation.commvault.com.

Option	Description
	<ul style="list-style-type: none"> • 2TB usable capacity volume (RAID 1 or RAID 5) with SSD class disks/PCIe I/O Cards 2GB controller cache memory for deduplication database. For IOPS requirements, refer to the Commvault documentation at http://documentation.commvault.com.
NIC	Dual port, 10 GigE card for data ingest and egress Quad 1 GigE ports for management interface
Host bus adapter cards	Dual port SAS adapter Dual port HBA for tape out (optional)

Server Configuration (Dedicated CS – Physical)

Software	OS: 64-bit Windows
Form factor	2U rack mount with minimum 12 SFF (2.5") drive bays
CPU/RAM	8 Cores 32GB RAM 2 x 650W power supply
Internal storage	2 x 400+ GB value SSDs in RAID 1 for OS and Commvault installation

Back-End Storage Configuration

Storage	NetApp E2800, E5700, EF300, or EF600. For information about these systems and drive shelves, see NetApp E-Series arrays section.
----------------	--

Data center large or extra-large configuration: multinode – global deduplication across nodes

- Targeted for enterprise data centers.
 - Includes HA and resiliency at the node level.
 - Provides global deduplication across nodes in the cluster.
 - Combine up to four nodes for petabyte scale capacity.
- Initial configuration sufficient for typical 30- to 90-day retention.
- Expansion options:
 - For long-term retention, add more expansion shelves and drives.
 - For additional capacity, buy new nodes in the same cell.

The Commvault V11 data management platform supports clustering together four MediaAgent nodes, each hosting a single deduplication database partition, in a partitioned configuration for global deduplication.

Figure 5 shows a multinode configuration and Table 6 gives in-depth configuration information.

Figure 5) Extra-large configuration: multinode.

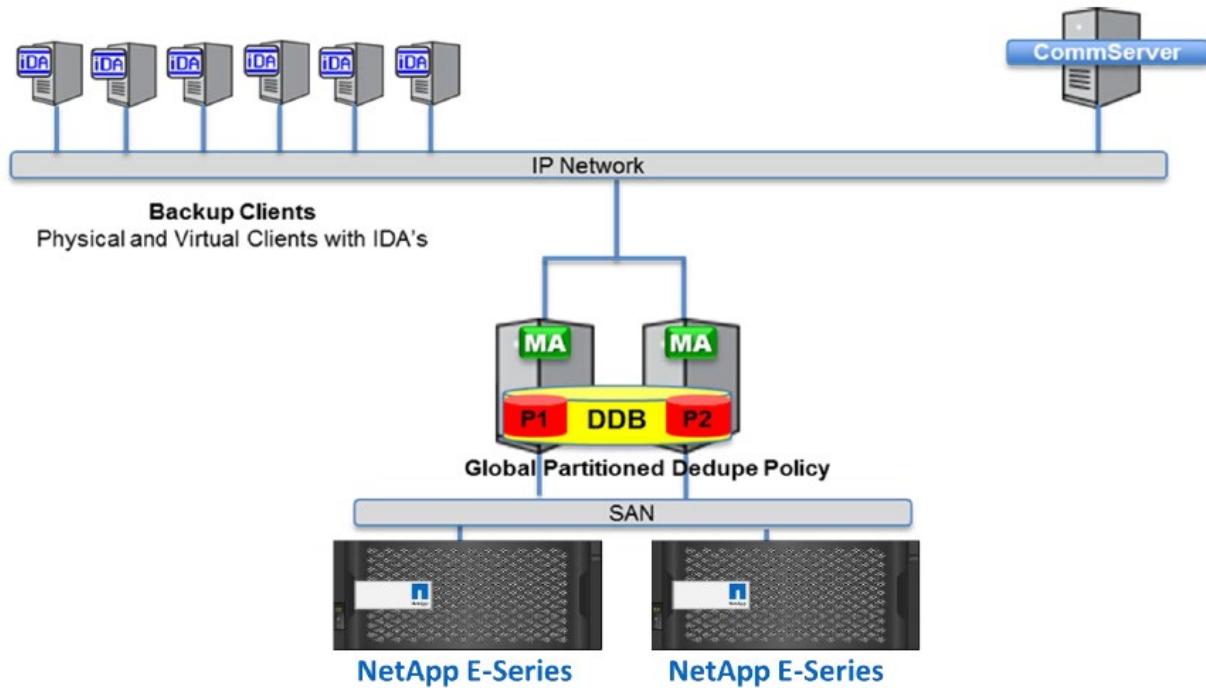


Table 6) Extra-large configuration: multinode.

Option	Description
Scope	Around 200TB of production data, based on data type, up to 400TB of backend size.
Software	OS: 64-bit Windows for CommServe 64-bit Windows/Linux for MediaAgent only
Commvault	Platform Release 2022 CommServe and MediaAgent

Server Configuration (Dedicated MediaAgents)

Software	OS: 64-bit Windows/Linux for MediaAgent
Form factor	2U rack mount with minimum 12 SFF (2.5") drive bays
CPU/RAM	16 Cores 128GB RAM 2 x 650W power supply
Internal storage	<ul style="list-style-type: none"> • 2 x 400GB SSD class disk for OS disk in RAID 1 • 2TB usable capacity volume (RAID 1 or RAID 5) with SSD class disks/PCIe I/O cards 2GB controller cache memory for index cache database. For IOPS requirements, refer to the Commvault documentation at http://documentation.commvault.com. • 2TB usable capacity volume (RAID 1 or RAID 5) with SSD class disks/PCIe I/O Cards 2GB controller cache memory for deduplication database. For IOPS requirements, refer to the Commvault documentation at http://documentation.commvault.com.

Option	Description
NIC	Dual port, 10 GigE card for data ingest and egress Quad 1 GigE ports for management interface
Host bus adapter cards	Dual port SAS adapter Dual port HBA for tape out (optional)

Server Configuration (CommServe)

Software	OS: 64-bit Windows
Form factor	2U rack mount with minimum 12 SFF (2.5") drive bays
CPU/RAM	8 Cores 32GB RAM 2 x 650W power supply
Internal storage	2 x 400+ GB value SSDs in RAID 1 for OS and Commvault installation

Back-End Storage Configuration

Storage	NetApp E2800, E5700, EF300, or EF600. For information about these systems and drive shelves, see NetApp E-Series arrays.
----------------	--

Enterprise extended mode configuration: multiple copies and long-term retention

In extended mode, each MediaAgent hosts two Deduplication Database Backups (DDB), each pointing to a separate copy. Use cases of this configuration are:

- Multiple copy support. One example is cross-site replication, in which each site does local backups and then replicates them over to the other site. The first copy is used for primary backups and the second copy for hosting replicated data coming in from the other site.
- Long-term retention. Primary backups are retrained for the short term, usually 30 to 90 days, but monthly backups need to be hosted for longer periods, usually for several years. The first copy is used for primary backups, and the DASH copy selective backups from primary copy to a second copy which is used for long-term retention.

Note that a MediaAgent of any size (small, medium, large, or extra-large) can be used in this mode, provided that each DDB is hosted on a dedicated volume.

Figure 6 shows a cross-site replication configuration and Figure 7 shows a long-term retention configuration.

Figure 6) Cross-site replication.

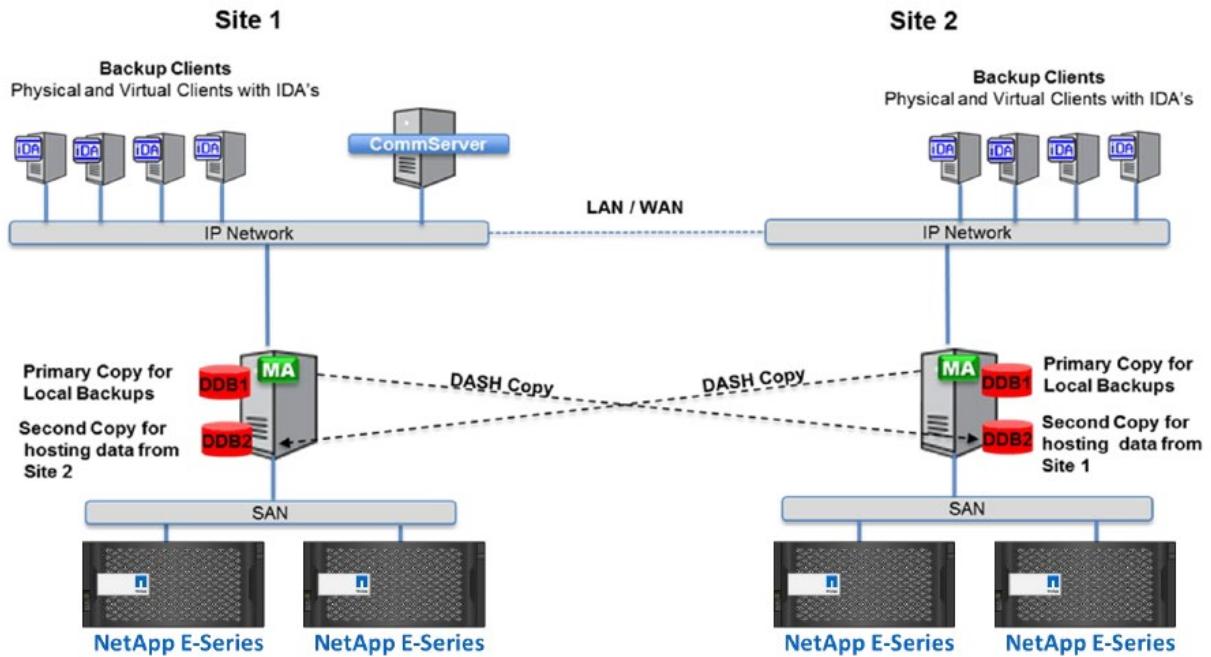
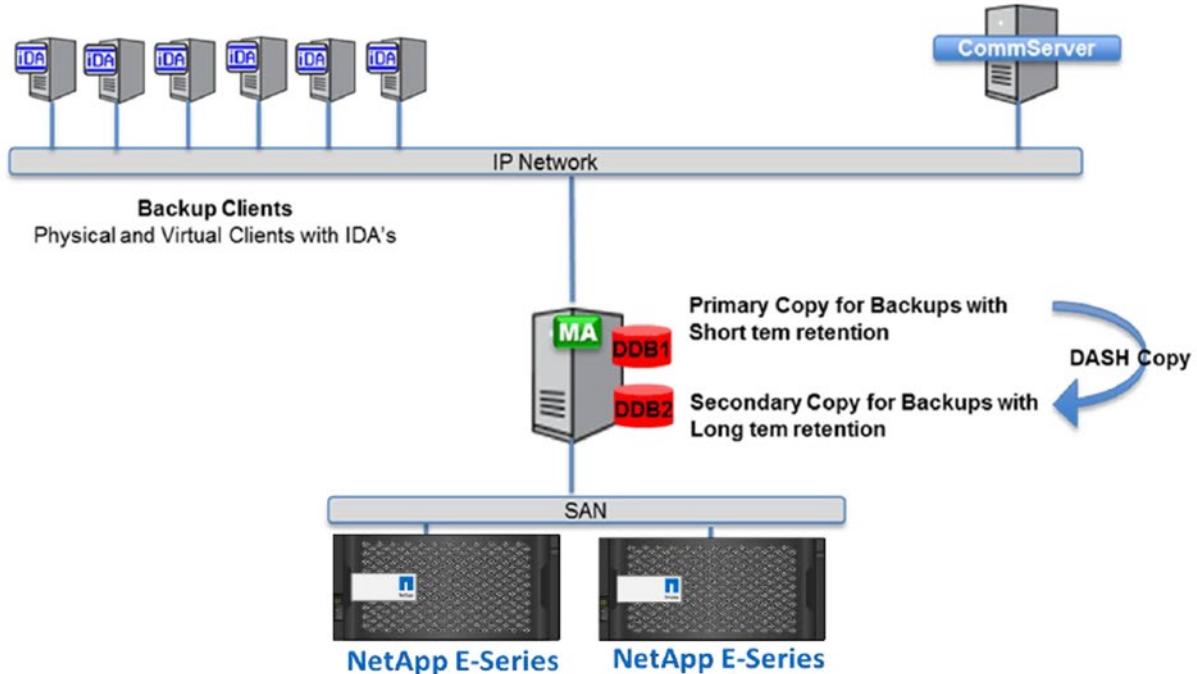


Figure 7) Long-term retention.



For hardware configurations with two deduplication databases, add more storage for the second volume based on the capacity of the node (small, medium, large, or extra-large).

For more information and additional configurations, refer to the Commvault documentation at <http://documentation.commvault.com>.

Configuration with IntelliSnap

In this configuration, the MediaAgent also acts as a proxy to locally mount snapshots created on E-Series primary storage for creating LAN-free backup copies. The MediaAgent requires either iSCSI or FC access to the array to access snapshots. iSCSI access can be provisioned via the 10 GigE network card. FC access requires a dual port 8 Gbps FC card to be added to the server configuration.

Backup storage capacity sizing

In any environment, backup capacity sizing for a storage system is calculated based on what Commvault refers to as “terabyte of front-end protection size” (FET), which is the capacity size of the data being protected. Table 7 provides some examples of front-end data and the resulting back-end capacity sizing required for storage. The general rule is to plan backup storage capacity at 1.5 to 2 times FET, depending on the desired retention timeframe.

Table 7) Backup storage capacity sizing.

Terabyte of Front-End Protection Size (FET)	Backup Storage Capacity Sizing at 1.5X	Backup Storage Capacity Sizing at 2X
15TB	22.5TB	30TB
30TB	45TB	60TB
60TB	90TB	120TB
120TB	180TB	240TB

Backup read performance goals

Table 8 shows backup read performance goals for the different types of configurations.

Table 8) Backup read performance goals.

Configuration	Backup Read Performance Goal
Small configuration: Single Node	500GB/hr restore
Medium configuration: Single node	750GB/hr restore
Large configuration: Single node	1TB/hr restore
Large configuration: Multinode	2TB/hr restore

NetApp E-Series arrays

Our E-Series arrays are the go-to system for so many because of the simplicity and reliability they deliver. From mid-sized businesses driving data-intensive applications like analytics, video surveillance, and disk-based backup, to small enterprises and remote offices needing mixed-workload performance for their dedicated apps.

E2800 and E5700 hybrid arrays

The NetApp E-Series E2800 and E5700 are industry-leading storage systems that deliver high input/output operations per second (IOPS) and bandwidth with consistently low latency to support the demanding performance and capacity needs of backup and recovery applications such as Commvault.

The E2800 and E5700 provide the following benefits:

- Support for wide-ranging workloads and performance requirements
- Fully redundant I/O paths, advanced protection features, and proactive support monitoring and services for high levels of availability, integrity, and security

- Increased IOPS performance by up to 20% compared to the previous high-performance generation of E-Series products
- A level of performance, density, and economics that leads the industry
- Interface protocol flexibility to support FC host and iSCSI host workloads simultaneously

As shown in Table 9, E2800 is available in three shelf options and E5700 is available in two shelf options, which support both hard-disk drives (HDDs) and solid-state drives (SSDs) to meet a wide range of performance and backup application requirements.

Table 9) E2800 and E5700 controller shelf and drive shelf models.

Controller Shelf Model	Drive Shelf Model	Number of Drives	Type of Drives
E2860/E5760	DE460C	60	2.5" and 3.5" SAS drives (HDDs and SSDs)
E2824/E5724	DE224C	24	2.5" SAS drives (HDDs and SSDs)
E2812	DE212C	12	2.5" and 3.5" SAS drives (HDDs and SSDs)

Each controller shelf includes two controllers, with each controller providing two Ethernet management ports for out-of-band management. The system also supports in-band management access and has two 12Gbps (x4 lanes) wide-port SAS drive expansion ports for redundant drive expansion paths. The controllers may also include two built-in host ports, which can be configured as either 16Gb FC or 10Gb iSCSI. The following host interface cards (HICs) can be installed in each controller:

Note: Both controllers in an array must be identically configured.

- 4-port 12Gb SAS Wide Port (SAS-3 connector)
- 4-port 32Gb Fibre Channel
- 4-port 25Gb iSCSI
- 2-port 100Gb InfiniBand (IB) (E5700 controllers only)

Note: E2800 can also be configured having only one controller.

For in-depth information on the E2800, see [TR-4725: Introduction to NetApp E-Series E2800 arrays](#).

For in-depth information on the E5700, see [TR-4724: Introduction to NetApp E-Series E5700 arrays](#).

EF300 and EF600 all-flash arrays

In one powerful all-flash array package, the EF300 and EF600 arrays deliver optimal performance for both random workloads and large sequential workloads. Both arrays can deliver consistent response times for random read IOPS with as few as 24 NVMe SSDs. The same configuration can deliver steady throughput for large sequential reads and cache-mirrored large sequential writes. When your workload meets the criteria, writes can be accelerated by using the built-in full stripe write acceleration feature.

The EF300 and EF600 provide the following benefits:

- Support for wide-ranging workloads and performance requirements
- Fully redundant I/O paths, advanced protection features, and proactive support monitoring and services for high levels of availability, integrity, and security
- A level of performance, density, and economics that leads the industry
- Use the web-based SANtricity System Manager UI to manage individual arrays and SANtricity Unified Manager enables you to organize and manage multiple new-generation E-Series and EF-Series arrays from a central management application

- Built-in web services API integration or the management client-based web services package makes the EF-Series product line easier than ever to integrate with your standard API-driven environment
- SAS expansion shelves for additional SAS SSDs or large capacity NL-SAS HDDs.

Each controller shelf includes two controllers, with each controller providing an Ethernet management port for out-of-band management. The following host interface cards (HICs) can be installed in each controller:

Note: Both controllers in an array must be identically configured.

- 4-port 32Gb Fibre Channel, SCSI over FC or NVMe/FC protocols supported
- 4-port 25Gb iSCSI
- 2-port 100Gb InfiniBand (IB), NVMe/IB, NVMe/RoCE, SRP/IB, or iSER/IB protocols supported
- 2-port 200Gb InfiniBand (IB), NVMe/IB, NVMe/RoCE, and iSER/IB protocols supported (EF600 only)

For in-depth information on the EF300, see [TR-4877: Introduction to NetApp EF300 array](#).

For in-depth information on the EF600, see [TR-4800: Introduction to NetApp EF600 array](#).

SANtricity OS

The NetApp E-Series controllers and SANtricity OS use the on-box, browser-based management interface, SANtricity System Manager.

E-Series storage systems are shipped preloaded with SANtricity OS, which includes SANtricity System Manager. To discover multiple storage systems running SANtricity OS from a central view, download SANtricity Unified Manager (which includes the Web Services Proxy) from the NetApp Support site. Then load it on a management server that has IP access to the storage systems.

For further information about the SANtricity Unified Manager and the SANtricity System Manager, see the [E-Series and SANtricity documentation resources page](#).

SANtricity security features

The new-generation E-Series arrays running the latest SANtricity OS are Common Criteria certified (NDcPP v2 certification). SANtricity security features include LDAP, RBAC, and SSL certificates. For complete details and workflow examples, see the following resources:

- [TR-4712: NetApp SANtricity Management Security Feature Details and Configuration Guide](#)
- [TR-4855: Security Hardening Guide for NetApp SANtricity](#)
- [TR-4813: Managing Certificates for NetApp E-Series Storage Systems](#)

SANtricity drive security technology provides comprehensive security for data at rest without sacrificing system performance or ease of use and supports both internal and external key management.

For more information about disk encryption, see [TR-4474: SANtricity drive security](#).

SANtricity reliability features

Table 10 provides a list of SANtricity reliability features and a brief explanation of each.

Table 10) SANtricity features for long-term reliability.

Reliability features with SANtricity

Proactive drive monitor and data evacuator. Nonresponsive drives are automatically power-cycled to see if the fault condition can be cleared. If the condition cannot be cleared, the drive is flagged as failed. For predictive failure events, the evacuator feature starts to remove data from the affected drive to move the

Reliability features with SANtricity

data before the drive fails. If the drive fails, rebuild resumes where the evacuator was disrupted, reducing the rebuild time.

Automatic drive fault detection, failover, and rebuild. You can perform these tasks by using global hot spare drives for standard RAID and spare pool capacity for DDP.

SSD wear-life tracking and reporting. This metric is found in the hardware tab's Drive Settings dialog box. It indicates the wear life of SSDs and replaces two SSD wear-life metrics (average erase count and spare blocks remaining) that were in previous versions of SANtricity. The metric is Percent Endurance Used; to access it, select a drive from the hardware view and then select Settings.

Online drive firmware upgrade. This feature upgrades one drive at a time and tracks writes to the affected drives during the upgrade window; it should be used only during low write I/O periods.

Parallel drive firmware upgrades are supported offline to upgrade multiple drives more quickly during a maintenance window.

Automatic load balancing. This feature provides automated I/O workload balancing and confirms that incoming I/O traffic from hosts is dynamically managed and balanced across both controllers. The workload of each controller is continually monitored and analyzed in the background. When I/O on one controller significantly exceeds the I/O on the other controller for a prolonged, predictable period, SANtricity can change volume ownership from the busy controller to the less busy controller. The feature does not react to short-term changes in I/O patterns. However, when a change of ownership is needed, SANtricity interacts with the affected host multipath driver to initiate an implicit path failover. Most current server operating systems and associated multipath drivers support implicit failover. For more information, search for "What is automatic load balancing?" in the System Manager online help.

Embedded SNMP agent. For the EF600 controller, SNMP is supported natively. The embedded SNMP agent complies with the SNMP V2C standard and RFC 1213 (MIB-II). For more information, search for "manage SNMP alerts" in the System Manager online help.

Automatic alerts. This feature sends email alerts to notify data center support staff about events on the storage array.

Event Monitor and system log. The SANtricity Storage Manager Event Monitor automatically records events that occur on the storage array. Syslog enables a second level of activity tracking that allows you to connect events with associated changes recorded in the system log.

AutoSupport. E-Series products have supported AutoSupport for several releases.

Ability to enable or disable AutoSupport maintenance window. AutoSupport includes an option for enabling or suppressing automatic ticket creation on error events. Under normal operation mode, the storage array uses AutoSupport to open a support case if there is an issue. To enable or disable the AutoSupport maintenance window, select Support > Access Management > AutoSupport.

SANtricity data protection features

E-Series has a reputation for reliability and availability. Many of the data protection features in these systems can be beneficial in a backup environment.

Background media scan and data assurance (T10 PI)

Media scan is a background process that is performed by the controllers to provide error detection on the drive media. The main purpose of the feature is to detect and repair media errors on disk drives that are infrequently read by user applications and where data loss might occur if other drives in the volume group fail. A secondary purpose is to detect redundancy errors such as data/parity mismatches. A background media scan can find media errors before they disrupt normal drive reads and writes.

The data assurance feature provides controller-to-drive data integrity protection through the SCSI direct-access block device protection information model. This model protects user data by appending protection information to each block of user data. The protection model is sometimes referred to as data integrity field protection or T10 PI. This model makes sure that an I/O has completed without any bad blocks written to or read from disk. It protects against displacement errors, data corruption resulting from

hardware or software errors, bit flips, and silent drive errors, such as when the drive delivers the wrong data on a read request or writes to the wrong location.

You need both data assurance and media scan. They work complementarily to protect your data.

Unreadable sector management

This feature provides a controller-based mechanism for handling unreadable sectors detected both during normal I/O operation of the controller and during long-lived operations such as reconstructions. The feature is transparent to the user and requires no special configuration.

Dynamic Disk Pools (DDP)

With DDP technology, NetApp SANtricity OS and management software allows you to create pools in addition to traditional volume groups (generally referred to as RAID groups). A pool can range in size from a minimum of 11 drives to as large as all the drives in a storage system, which is up to 480 NL-SAS drives in the NetApp E5700 system. Pools can consist of either HDDs or SSDs. In addition, pools and volume groups can coexist in the same system.

For more information about DDP, see [TR-4652: SANtricity OS Dynamic Disk Pools](#).

Commvault Complete Data Protection

Overview

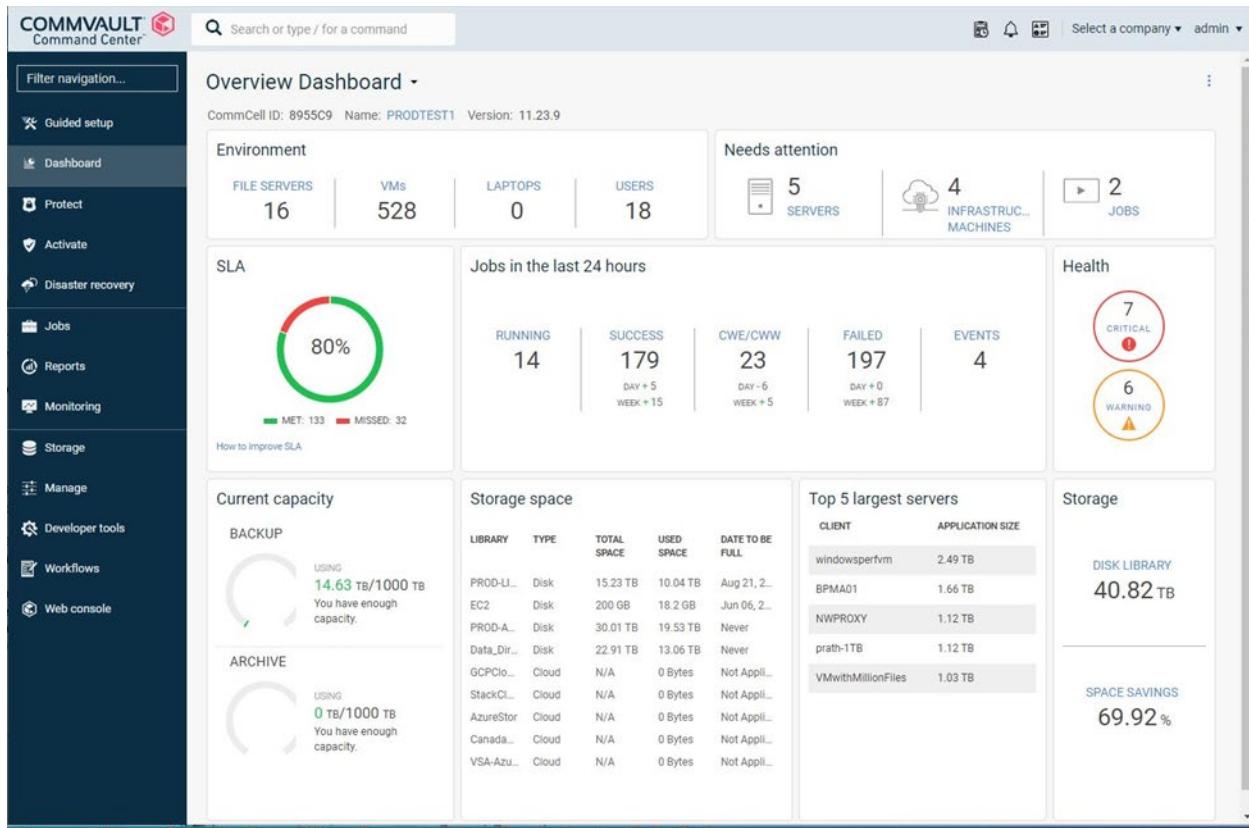
Commvault Complete Data Protection is a highly scalable, integrated data and information management solution, built from the ground up on a single platform and unified code base. All functions share the same back-end technologies to deliver the unparalleled advantages and benefits of a truly holistic approach to protecting, managing, and accessing data. The software contains modules to protect and archive, analyze, replicate, and search your data. The modules share a common set of back-end services and advanced capabilities, seamlessly interacting with one another. This addresses all aspects of data management in the enterprise, while providing infinite scalability and unprecedented control of data and information.

Production data is protected by installing agent software on the physical or virtual hosts, which use operating system or application native APIs to protect data in a consistent state. Production data is processed by the agent software on client computers and backed up through a data manager, the MediaAgent, to disk, tape, or cloud storage. All data management activity in the environment is tracked by a centralized server, the CommServe, and can be managed by administrators through a central user interface. End users can access protected data by using web browsers or mobile devices.

Key features of the software platform:

- Complete data protection solution supports all major operating systems, applications, and databases on virtual and physical servers, NAS shares, cloud-based infrastructures, and mobile devices.
- Simplified management through a single, intuitive, and easily customizable console, the Commvault Command Center™. It allows you to view, manage, and access all data and functions across the enterprise, see Figure 8.

Figure 8) Commvault Command Center dashboard.



- Multiple protection methods include backup and archive, snapshot management, replication, and content indexing for e-discovery.
- Efficient storage management uses deduplication for disk, tape, and cloud. For a list of cloud targets supported, go to the [Commvault Cloud Storage Support](#) page.
- Integration with the industry's top storage arrays automates the creation of indexed, application-aware hardware snapshot copies across multivendor storage environments.
- Complete virtual infrastructure management supports both VMware and Hyper-V.
- Advanced security capabilities limit access to critical data, provide granular management capabilities, and offer single sign-on access for Active Directory users.
- Policy-based data management transcends the limitations of legacy backup products by managing data based on business needs, not physical location.
- Advanced end-user experience empowers users to protect, find, and recover their own data by using common tools such as web browsers, Microsoft Outlook, and File Explorer.
- End users can use third-party screen readers with the Web Console, Admin Console, and command line interface.

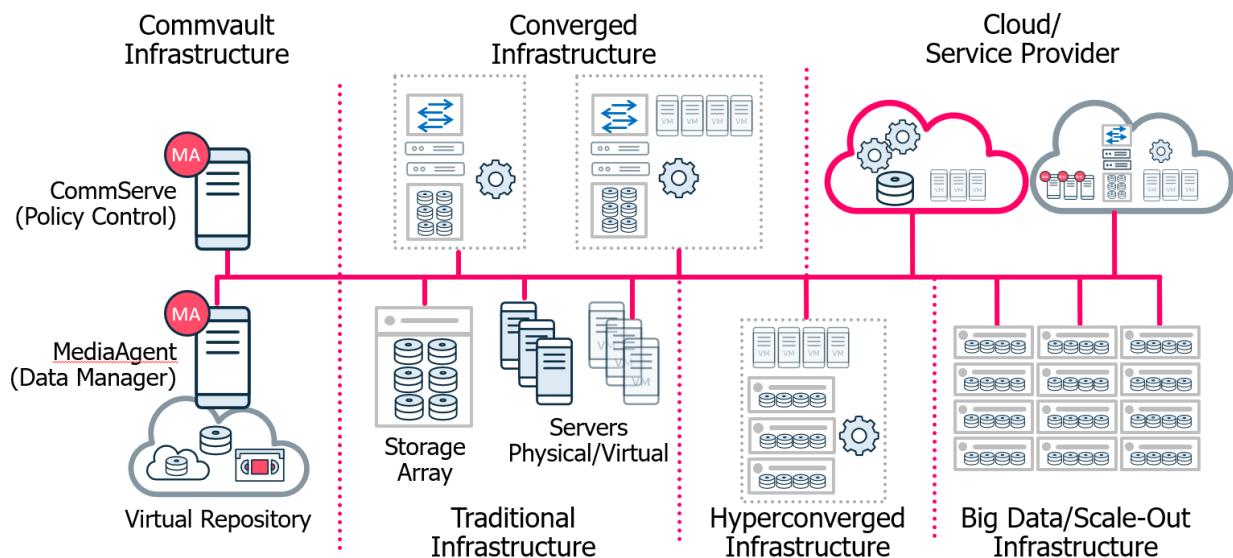
Commvault Complete Data Protection components

Table 11 lists the components of Commvault Complete Data Protection along with descriptions for each of the architecture components. Figure 9 provides a graphical representation of a Commvault architecture.

Table 11) Commvault Complete Data Protection components and descriptions.

Component	Description
CommCell®	A single instance of the Commvault backup environment.
CommServe	“Command and control center” that handles all requests for activity between MediaAgent and iDataAgent components. Administration: Includes configuration, security, licensing, policies, and scheduling. Disaster recovery: Database includes configuration and media association. Reporting: General health and welfare of environment and detailed reporting on the health of the environment, including the health of the data itself. Event Orchestration: Houses the centralized event and job managers that handle the various operations within a CommCell. Includes an instance of Microsoft SQL 2008 Enterprise for the internal database. Requires installation on Windows 2008 or later.
MediaAgent	Data Mover, home of Deduplication Indexer. MediaAgents have broad operating system support, including Windows, Linux, and UNIX options.
CommCell Console	The Commvault management graphical user interface application.
Intelligent Data Agent (iDA)	Agents that control data consistency during backup operations.
CommServe Database	Handles scheduling, storage policies, activity metadata, media management, reporting, security role-based privileges, and encryption key management.
ContentStore	A hardware-agnostic virtual repository. The back-end repository for all Commvault-managed information. ContentStore consolidates protection and archiving of data to eliminate inefficient data silos that waste resources and infrastructure. An intelligent index provides global awareness for your data so you can quickly find what you need, when you need it. Commvault software has a single, deduplicated index across the ContentStore, making it the only enterprise-class solution like it in the market. Having a reusable, common repository facilitates better control of applications, processes, and data workflow across an enterprise. Ultimately, it delivers information at your fingertips for the better productivity, improved collaboration, and smarter decision making that can transform your business.
Clients	Hosts running iDataAgents for which data is protected.
Backup Set	A layer of management within iDataAgents for grouping subclients.
Subclient	A layer of management within a backup set. A client can have multiple subclients, each of which can be associated with different source data.
Disk Library	A storage resource with an associated mount path that is used in the Commvault backup environment to store index information backups.
Storage Policy	A logical object through which a subclient is protected. The storage policy defines how data is backed up and replicated, as well as retention requirements.

Figure 9) Commvault data management platform.



Commvault Complete Data Protection design considerations and best practices

Infrastructure design

Design questions to ask your customer

1. How many sites? That is, the number of sites that exist in the customer's environment. (A site is any location where data exists that needs to be protected or where protected data will reside.) Don't forget that laptops outside the protected network count as additional sites.
2. For each site, where must protection copies reside? You need to know the protection requirements for each site, including whether or not a local copy is made and where any remote copies will be sent.
3. How many clients will be managed? To account for scale, in larger environments it can be helpful to have this information broken down by sites.
4. How much data will be managed in each site? This information helps you address sizing and scale issues.

Infrastructure design: step-by-step design process

1. Determine the proper number and specifications of CommServes, based on client limits.
2. Determine the best location for production CommServe and disaster recovery CommServe.
3. Determine which sites require MediaAgents.
4. Make sure that each site has sufficient MediaAgent quantity and specifications based on client data size or throughput requirements.

Deduplication solution design

Table 12 lists deduplication terminology along with a description of each of the terms.

Table 12) Deduplication terminology.

Deduplication Component	Description
Source-side deduplication	Only unique blocks that are not present in the deduplication database are transferred from the client to the MediaAgent.
Target-side (MediaAgent) deduplication	All data is sent from the client to the MediaAgent. Only unique blocks are transferred to disk.
DASH full	Deduplication-optimized full backup that allows all backups from the client to be incremental-only after initial full backup.
DASH copy	Deduplication-optimized copy of backup jobs from one storage copy to another.
Deduplication database	Database located on a high-speed disk volume that stores information about all deduplicated blocks.
Global deduplication	Policy that allows all storage copies on one or multiple MediaAgents to be deduplicated together.

Building block design questions to ask about sizing

1. How much data in each site?
2. What is the data composition in each site? (What types of data will be protected?)

Breaks down into the following categories:

- a. Uncompressed files or file systems.
- b. Compressed or uncontrolled data (NDMP data or zipped files fall into this category).
- c. Any databases protected at the application level as opposed to dumped by the application are categorized together as well

3. What is the data retention for deduplicated disk copies? (You need to know how long the data will be retained on disk — typically in weeks — and whether there are any special retention considerations.)
4. Will deduplicated secondary copies be made? Where?

Deduplication step-by-step design process

1. Determine how much data will be protected at each site, what type of data is to be protected, how long it will be retained, and whether you will make secondary copies.
2. Use the deduplication calculator to determine how much disk capacity is required. If the number of back-end TB exceeds the MediaAgents limits, increase the number of MediaAgents accordingly. Also consider throughput and client count when determining the number of MediaAgents.
3. Determine whether you will use standalone building block nodes or partitioned deduplication. Size the proper number of Media-Agents based on front-end data protected.
4. SAN disk provides shared disk libraries that are required in order to use partitioned deduplication.

Deduplication to disk best practices

Before using deduplication, review the deduplication best practices:

https://documentation.commvault.com/commvault/v11/article?p=features/deduplication/best_practice.htm

Virtual Server Agent design questions

1. Hypervisor VMware or Hyper-V?
2. How much VM data is being protected at each site? (Get a count of how many guests and data sizing.)

3. Are there any physical raw device mapping (RDM)s or pass-through devices?

Advanced data protection best practices

Microsoft Active Directory protection best practices

4. Install a File System Agent on each domain controller, with the possible exception of read only domain controllers (RODCs). Perform full system protection.
5. Install an Active Directory Agent on a single domain controller in each domain of the forest that requires granular recovery.
6. If you do not want to perform an additional Active Directory backup, use the offline mining feature.

Microsoft Exchange best practices

1. Determine which mailbox servers should be used for backup or recovery and install the File System Agent and Exchange Database Agent.
2. Configure a DAG pseudo-client and set the backup selection rules.
3. If granular recovery is required, install the Exchange Mailbox Agent. If performance or licensing is a concern, use the offline mining feature.

Microsoft SQL protection best practices

1. For single-server deployments, perform all operations locally.
2. For SQL failover clusters, install the agents on all nodes and then configure Windows cluster clients for each cluster group.
3. For Always-On Availability Groups, perform all protection from the primary replica.

Microsoft SharePoint protection best practices

1. Use the File System Agent or Virtual Server Agent to provide full server protection of all SharePoint servers.
2. Use the SharePoint Farm Backup feature to protect all critical SharePoint elements except the content database.
3. Use the SQL Agent to protect the content database.
4. If granular recovery is required, perform document-level and/or site collection backups. If performance is a concern, use Commvault IntelliSnap for NetApp and offline mining.

NetApp E-Series volume configuration guidelines

NetApp E-Series storage configuration guidelines for Commvault Complete Data Protection disk libraries

NetApp recommends not using thin volumes while setting up NetApp E-Series with Commvault.

Small and medium configuration – single node

- Use RAID 6 (8+2) volume groups.
- Create LUNs (not thin) from the volume group, regardless of the spindle size installed in the system.
- Select segment size in 64KB increments. Default segment size (file system typical) is 128KB.

Large and extra-large configurations – single and multi-node

- Create either multiple RAID 6 (8+2) volume groups out of the drives in the system or use Dynamic Disk Pools to maximize ease of use and rebuild times.

- Create LUNs (not thin) from the volume group, regardless of the spindle size installed in the system.
- For RAID 6 Volume Groups, select segment size in 64KB increments. Default segment size (file system typical) is 128KB.
- For a Dynamic Disk Pools, the default segment size (file system typical) is 128KB.

Architectural tools and sizing

Assessment tools

Commvault External Data Connector (EDC) toolkit

Commvault has made it easy for customers to switch from Symantec NetBackup, EMC Networker, and IBM TSM to Commvault software. Just ask any of the more than 16,000 Commvault users who left their costly legacy backup software behind for the singular modern approach of Commvault software to save them time and money while adding value to their business. With the Commvault External Data Connector (EDC) toolkit, customers can switch with confidence and ease. The toolkit is composed of three important components: EDC software, the EDC cloud, and migration services.

Import settings from Symantec, EMC, IBM

Commvault EDC software connects with Symantec NetBackup, EMC Networker, and IBM TSM master servers to collect legacy client attributes, policies, and job history. That metadata can then be imported to the Commvault software platform for reporting, modeling, and automated conversion into the new Commvault software installation.

Safely configure and test in the cloud

Commvault's EDC cloud is a web-based platform that allows customers to trial-run Commvault software in a sandbox before they begin implementation in their environment. This capability helps demonstrate the benefits of using Commvault software and ensures that customers' configurations and policies are optimized before they start their implementation.

Leverage Commvault's team of product experts

The Commvault services team helps customers with each step of their migration to Commvault software. Migration operations services include proactive risk identification and planning, which reduce the time to upgrade while also ensuring that best practices are employed.

External Data Connector

The External Data Connector offers the following functionalities:

- Collect information from non-Commvault backup products
- Stage on a virtual machine
- Size and report the customer environment

The major use case is when an incumbent backup solution is replaced with Commvault. For more information, go to <https://cloud.commvault.com/>.

System Discovery and Archive Analyzer Tool

The System Discovery and Archive Analyzer Tool offers the following functionalities:

- Collect host and file information from a list of servers
- Stage on a virtual machine

- Size and report on the customer environment

Use cases of the tool include:

- Alternative to EDC tool
- Perform detailed file-level archive assessments

The tool can be accessed at <https://cloud.commvault.com/>.

Software Configurator design tool

The Software Configurator is an online survey-based configuration tool that generates basic solution design and provides license model comparisons.

It can be accessed at <https://partners.commvault.com>.

The Software Configurator also:

- Estimates the amount of backup disk required when Commvault deduplication is used.
- Estimates the amount of time it takes to replicate data over the WAN in a DASH copy or Continuous Data Replicator solutions.

Summary

NetApp E/EF-Series and Commvault Complete Data Protection combine to provide the necessary capacity, throughput, IOPS, and response times to meet performance requirements for demanding backup windows while ensuring data reliability in backup environments.

E/EF-Series products offer a modular architecture to meet the most demanding performance, scale, and rack density requirements, through their numerous drive choices, while providing flexibility on drive enclosures. Commvault Complete Data Protection is more than just an upgrade to the industry-leading software solution for protecting, managing, and accessing corporate information; it is an exponential leap forward. It is a data management solution that can scale to meet all the demands of an enterprise of any size. When paired together, the E/EF-Series storage solution and Commvault Complete Data Protection provide an innovative solution to address all data management needs.

Appendix

Commvault Complete Data Protection terminology and definitions

Table 13 provides a list of Commvault terminology and their descriptions.

Table 13) Commvault terminology descriptions.

Commvault Terminology	Description
Backup Series	All the archive files on a given backup medium originating from the same subclient.
Backup Set	A group of subclients that includes all the data backed up by the iDataAgent.
Client Compression	A feature that compresses data on a client computer before sending the data to backup media.
Client Computer	A computer in a CommCell® management group that has agent software installed on it.
Client Computer Group	A logical grouping of client computers in which selected options can apply to all member clients.

Commvault Terminology	Description
CommCell Administration	A user group capability that permits members of the user group to administer a CommCell management group.
CommCell Browser	The window in the CommCell Console that displays all the objects in the CommCell management group in a tree structure.
CommCell Console	The graphical user interface used to access and manage the system.
CommCell Management Group	The basic organizational unit of a data management system. A CommCell management group contains one CommServe StorageManager and at least one client.
CommCell Survey	An automatic reporting service that collects information about the CommCell management group, such as overall wellness of the managed components, license usage, and job statistics, and uploads it to the secure cloud site for customer service monitoring. This service is enabled through the Diagnostics and Usage dialog box, found on the Control Panel.
Common Technology Engine	Consists of the CommServe StorageManager and the MediaAgent software modules that provide the necessary tools to manage and administer the Client Agents and manage the storage media associated with the CommCell management group.
CommServer Database Engine	A SQL server database that is used by the CommServe StorageManager that contains all the information related to the CommCell management group.
CommServer StorageManager	The software module that communicates with all clients and MediaAgents, and coordinates operations (data protection, data recovery, and administration operations, job management, event management, and so on) within a CommCell management group. There is only one CommServe StorageManager per CommCell management group.
Compliance Archiving	An operation that moves data from a Journaling Mailbox on the client computer to secondary storage media for the purpose of complying with legal or business regulations.
Content Indexing	A feature used to search archived data of supported file and message types by their content.
Data Replication	The creation of secondary copies of production data by using a combination of host-based replication and snapshot technologies. These real-time data replication copies can be accessed immediately for fast recovery, used to create multiple recovery points, or used to perform traditional backups without having an impact on server performance.
Data Stream	A data channel through which client data flows to backup media.
Differential Backup	A backup of all the data on a subclient that has changed since the subclient's last full backup.
Disaster Recovery	The planning for and/or the implementation of a strategy to respond to such failures as a total infrastructure loss, or the failure of computers (CommServe StorageManager, MediaAgent, client, or application), networks, storage hardware, or media. A disaster recovery strategy typically involves the creation and maintenance of a secure disaster recovery site, as well as the day-to-day tasks of running regular disaster recovery backups.
Disaster Recovery Backup	Backs up metadata and Windows registry data in two phases. In the first phase, the data to a local or network path is backed up. In the second phase, the data is backed up to media by using a disaster recovery backup storage policy. This data can then be restored by using the CommServe Recovery Tool.
Disaster Recovery Backup Storage Policy	A storage policy that is used to store metadata to media. This metadata contains information about the CommCell database and the backed-up data. In case of a

Commvault Terminology	Description
	system failure, disaster recovery backup data can be retrieved by using this storage policy.
Disk Library	SAN (storage area network) disk is configured as a library to back up data to disk.
Drive Pool	Logical entities used to facilitate the sharing of a library's drives between multiple MediaAgents. See also master drive pool.
Full Backup	A backup of all the data on a subclient. A full backup provides the baseline for subsequent incremental and differential backups. (Known as a level 0 backup in Oracle.)
Incremental Backup	A backup of all the data from a subclient that has changed since the subclient's last full, incremental, or differential backup.
Index Cache	A storage location maintained by a MediaAgent that contains the index data generated by the system when backups are conducted.
Instance	The level in the CommCell Browser tree that represents the database that needs to be backed up. All subclients for the database are defined under an instance.
Instance (File Archiver for Windows)	A File Archiver for Windows instance exists as a level in the CommCell Browser under the client and agent levels and represents the type of file system that needs to be backed up. The four types of instance are Local File System, Celerra, Network File Share, and NetApp FPolicy®. An instance is user defined, rather than created by default after installing the agent.
Intelligent DataAgent (iDataAgent)	A software module that backs up and restores data of a particular application type on a host computer system.
IntelliSnap Backup	Commvault IntelliSnap® for NetApp is a feature that enables the creation of a point-in-time snapshot of the application data to be used for various data protection operations.
Job Controller	The window in the CommCell Console that can be used to monitor and manage the active jobs in the CommCell management group.
Master Drive Pool	A logical entity that is used to facilitate the sharing of a library's drives between multiple MediaAgents. See also drive pool.
MediaAgent	The software module that transmits data between clients and backup media. The MediaAgent manages the data that is stored on the media.
Network Agent	A feature that can be used to increase the data transfer throughput from a client during data protection operations.
Network Bandwidth Throttling	A feature that can be used to control the amount of data transferred in a network during a data protection operation.
Primary Storage	Data in active use from computer hard disks and/or volumes. See also secondary storage.
Replication Policy	A centralized template to configure replication sets or replication pairs within a CommCell management group. A replication policy consists of a common configuration for replication set and replication pairs that can be applied to target replication set or replication pairs within the CommCell management group.
Replication Set	A group of replication pairs.
Schedule Policy	A feature used to associate a schedule or groups of schedules and attach it to any number of clients, backup sets, subclients, or storage policies within the CommCell management group.
Secondary Storage	Backup or archival data moved to storage media, such as tape media, disk volumes, and so on. See also primary storage.

Commvault Terminology	Description
SLA	Measures the data protection coverage aspects in short-term and various long-term intervals to determine whether or not the data protection coverage for CommCell client, application, or subclient content is within an acceptable level.
Snapshot Copy	A snapshot copy of the storage policy is an additional copy of the protected data that is used in IntelliSnap operations. The snapshot backup copy stores the metadata information related to the IntelliSnap feature.
Storage Policy	A logical entity through which data from a subclient is backed up. A storage policy consists of one or more copies that associate data with physical media.
Stubs	Files that point to backed-up and archived data; functionally similar to a Windows shortcut, Macintosh alias, or Unix symbolic link.
Subclient	A logical entity that uniquely defines a unit of data on a client computer.
Subclient Policy	A logical entity through which configuration of multiple file system subclients within a CommCell management group can be accomplished from a centralized template. A subclient policy consists of one or more subclient templates that contain a common configuration that is applied to target subclients within a CommCell management group.
Synthetic Full Backup	An operation that combines the most recent full backup of the selected data with all subsequent incremental or differential backups and stores the result in a single archive file.
Virtualization	Virtual Server iDataAgent, Microsoft Hyper-V backup software, and VMware backup software is a single product for controlling all aspects of data management from a single console in virtualized environments for both Microsoft Hyper-V and VMware. This includes data protection, archiving, replication, and reporting.

Where to find additional information

To learn more about the information that is described in this document, review the following documents and/or websites:

- For Commvault software documentation:
<http://documentation.commvault.com>
- NetApp Product Documentation
[E-Series documentation](#)

Version history

Version	Date	Document version history
Version 1.0	December 2022	Initial release.

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.

Copyright information

Copyright © 2022 NetApp, Inc. All Rights Reserved. Printed in the U.S. No part of this document covered by copyright may be reproduced in any form or by any means—graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system—with prior written permission of the copyright owner.

Software derived from copyrighted NetApp material is subject to the following license and disclaimer:

THIS SOFTWARE IS PROVIDED BY NETAPP "AS IS" AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE, WHICH ARE HEREBY DISCLAIMED. IN NO EVENT SHALL NETAPP BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

NetApp reserves the right to change any products described herein at any time, and without notice. NetApp assumes no responsibility or liability arising from the use of products described herein, except as expressly agreed to in writing by NetApp. The use or purchase of this product does not convey a license under any patent rights, trademark rights, or any other intellectual property rights of NetApp.

The product described in this manual may be protected by one or more U.S. patents, foreign patents, or pending applications.

LIMITED RIGHTS LEGEND: Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (b)(3) of the Rights in Technical Data—Noncommercial Items at DFARS 252.227-7013 (FEB 2014) and FAR 52.227-19 (DEC 2007).

Data contained herein pertains to a commercial product and/or commercial service (as defined in FAR 2.101) and is proprietary to NetApp, Inc. All NetApp technical data and computer software provided under this Agreement is commercial in nature and developed solely at private expense. The U.S. Government has a non-exclusive, non-transferrable, non-sublicensable, worldwide, limited irrevocable license to use the Data only in connection with and in support of the U.S. Government contract under which the Data was delivered. Except as provided herein, the Data may not be used, disclosed, reproduced, modified, performed, or displayed without the prior written approval of NetApp, Inc. United States Government license rights for the Department of Defense are limited to those rights identified in DFARS clause 252.227-7015(b) (FEB 2014).

Trademark information

NETAPP, the NETAPP logo, and the marks listed at <http://www.netapp.com/TM> are trademarks of NetApp, Inc. Other company and product names may be trademarks of their respective owners.

TR-4950-1222