

Fujitsu Storage
ETERNUS AB series All-Flash Arrays,
ETERNUS HB series Hybrid Arrays

ETERNUS AB/HB series and Commvault Data Protection
Platform Reference Architecture and Storage Best Practices



In partnership with Commvault

Table of Contents

1. Executive Summary	7
Introduction	7
2. Reference Architecture Overview	8
Solution Configuration Options	8
Express (S) Small Configuration: Single Node	8
Work Group (M) Medium Configuration: Single Node	9
Data Center (L) Large Configuration: Single Node	10
Data Center (XL) Extra-Large Configuration: Single Node	12
Data Center Large or Extra-Large Configuration: Multinode – Global Deduplication Across Nodes	14
Enterprise Extended Mode Configuration: Multiple Copies and Long-Term Retention	15
Configuration with IntelliSnap	16
Backup Storage Capacity Sizing	17
Backup Read Performance Goals	17
3. Enterprise Reliability and Availability	18
4. Commvault Data Protection Platform	19
Overview	19
5. Data Protection Platform Design Considerations and Best Practices	21
Infrastructure Design	21
Design Questions to Ask Your Customer	21
Infrastructure Design: Step-by-Step Design Process	21
Deduplication Solution Design	21
Virtual Server Agent (VSA) Design Questions	22
Advanced Data Protection Best Practices	23
6. ETERNUS AB/HB series Volume Configuration Guidelines	24
ETERNUS AB/HB series Configuration Guidelines	24
Small and Medium Configuration - Single Node	24
Large and Extra-Large Configurations - Single and Multinode	24

7. Architectural Tools and Sizing.....	25
Assessment Tools	25
Commvault External Data Connector (EDC) Toolkit	25
System Discovery and Archive Analyzer Tool	25
Design Tool: Software Configurator	26
8. Summary	27
A. Appendix.....	28
Commvault Data Protection Platform Terminology and Definitions	28
ETERNUS AB/HB series Software Documentation	31
Commvault Data Protection Platform Software Documentation	31

List of Figures

Figure 1	Small configuration: Single node	8
Figure 2	Medium configuration: Single node	9
Figure 3	Large configuration: Single node	11
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 management dashboard	19

List of Tables

Table 1	Small configuration: Single node	8
Table 2	Medium configuration: Single node	10
Table 3	Large configuration: Single node.....	11
Table 4	Extra-large configuration: Single node.....	13
Table 5	Extra-Large configuration: 2-node or larger.....	14
Table 6	Backup storage capacity sizing	17
Table 7	Backup read performance goals.....	17
Table 8	Deduplication terminology.....	21
Table 9	Commvault Data Protection Platform terminology and definitions.....	28

Preface

This document outlines the reference architecture and best practices when using the ETERNUS AB/ HB series storage system in a Commvault Data Protection Platform environment.

Copyright 2023 Fujitsu Limited

First Edition
October 2023

Trademarks

Third-party trademark information related to this product is available at:

<https://www.fujitsu.com/global/products/computing/storage/eternus/trademarks.html>

Trademark symbols such as ™ and ® are omitted in this document.

About This Manual

Intended Audience

This manual is intended for system administrators who configure and manage operations of the ETERNUS AB/HB, or field engineers who perform maintenance. Refer to this manual as required.

Related Information and Documents

The latest information for the ETERNUS AB/HB is available at:

<https://www.fujitsu.com/global/support/products/computing/storage/manuals-list.html>

Document Conventions

■ Notice Symbols

The following notice symbols are used in this manual:

Caution

Indicates information that you need to observe when using the ETERNUS AB/ HB. Make sure to read the information.

Note

Indicates information and suggestions that supplement the descriptions included in this manual.

1. 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.

This document outlines a reference architecture for enabling a collaborative backup and recovery solution on the ETERNUS AB/HB with Commvault Data Protection Platform data protection software.

Introduction

This reference architecture provides guidance for Commvault Data Protection Platform deployments with the ETERNUS AB/HB series that will accelerate time to application for this solution.

The ETERNUS AB/HB series and Data Protection Platform offer a best-in-class backup and recovery solution. This solution is optimized for drive-to-drive backup and recovery relying on the policy-based management and deduplication features of Data Protection Platform while providing high-capacity storage performance on flexible the ETERNUS AB/HB series storage systems. This solution gives customers superior performance and functionality at a competitive price.

Features:

- 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 the ETERNUS AB/HB series storage system (raw capacity)
- The ETERNUS AB/HB and Commvault data management platform reduce complexity through off-site replication
- Archiving and stubbing minimize future storage capacity needs
- Efficiency
 - Direct-to-drive 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 primary data that scales from several terabytes to multiple petabytes under management

2. Reference Architecture Overview

This chapter describes reference architectures ranging from small environments that protect a few terabytes of data to enterprise-scale environments with petabytes of data under management.

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 Small configuration: Single node

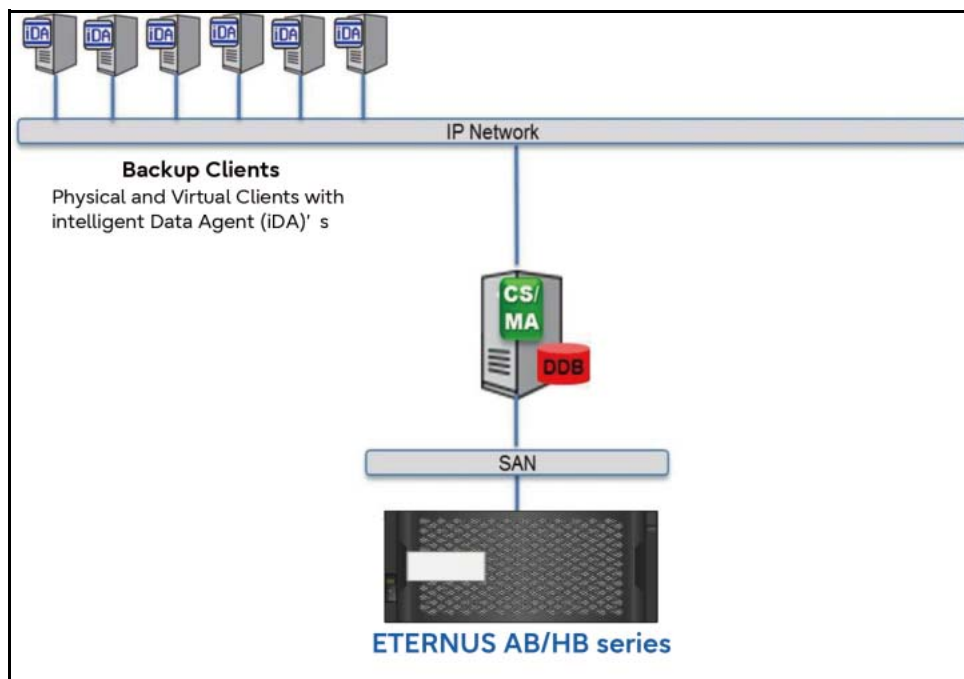


Table 1 Small configuration: Single node

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	V11 CommServe and MediaAgent or MediaAgent only

2. Reference Architecture Overview
Solution Configuration Options

Server Configuration (1 Server Required)	
CPU/RAM	8 Cores 32GB RAM 2 x 650W power supply
Internal storage	<ul style="list-style-type: none"> • 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.
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	ETERNUS AB/HB series

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 Medium configuration: Single node

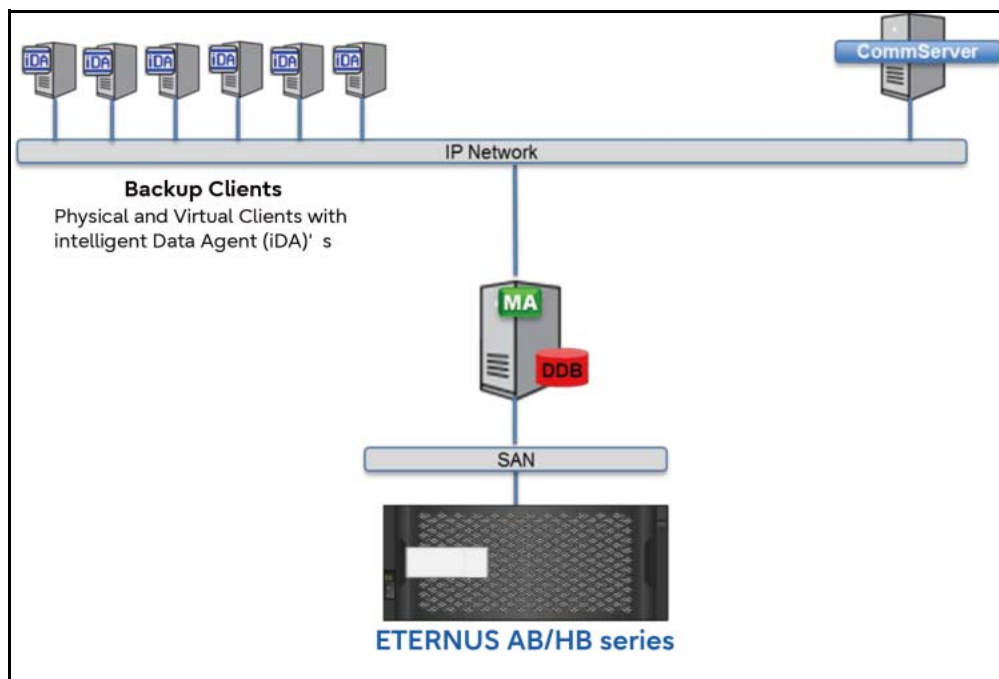


Table 2 Medium configuration: Single node

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	V11 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 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"> 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 data-base. 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	ETERNUS AB/HB series

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 Large configuration: Single node

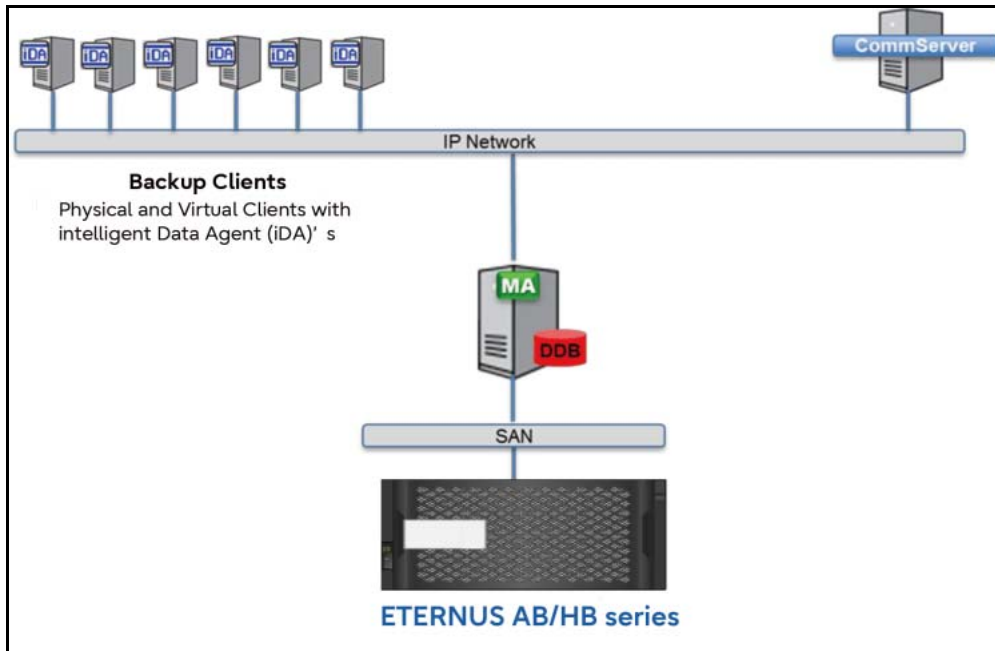


Table 3 Large configuration: Single node

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 (CS) 64-bit Windows/Linux for MediaAgent only
Commvault	V11 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
Internal storage	<ul style="list-style-type: none"> 400GB usable drive, minimum 4 spindles 15K RPM or higher or SSD class drive for OS and Commvault installation 1.2TB usable capacity volume (RAID 1 or RAID 5) SSD class drive/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 drive/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)

2. Reference Architecture Overview
Solution Configuration Options

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	ETERNUS AB/HB series

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 Extra-large configuration: Single node

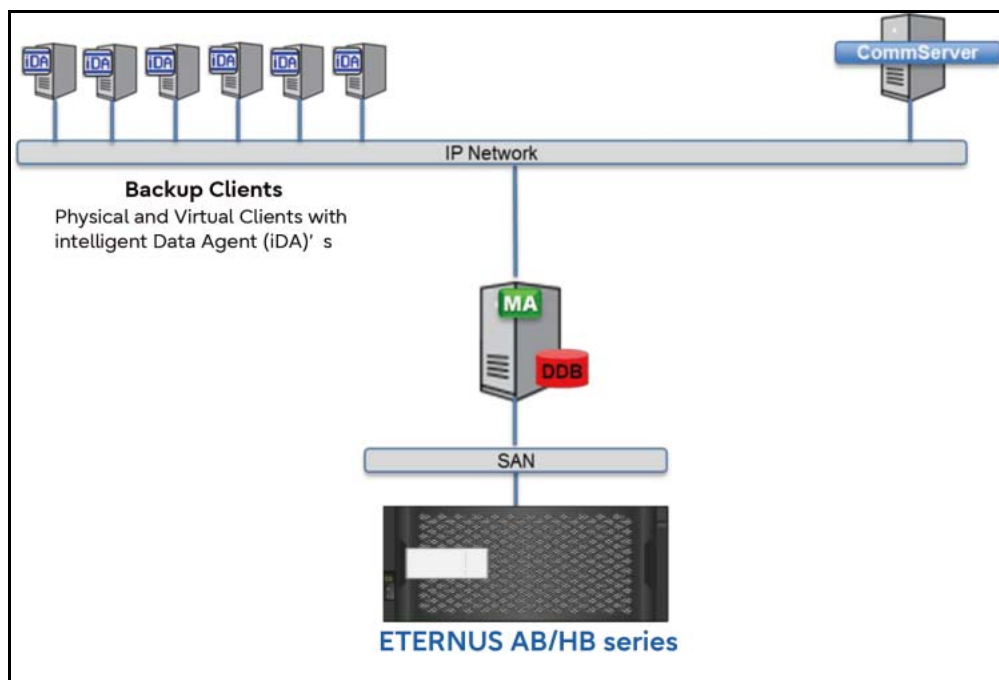


Table 4 Extra-large configuration: Single node

Option	Description
Scope	Around 120 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	V11 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 drive for OS drive in RAID 1 • 2TB usable capacity volume (RAID 1 or RAID 5) with SSD class drives/ 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 drives/ 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	ETERNUS AB/HB series

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 option:
 - 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 two-node configuration.

Figure 5 Extra-large configuration: Multinode

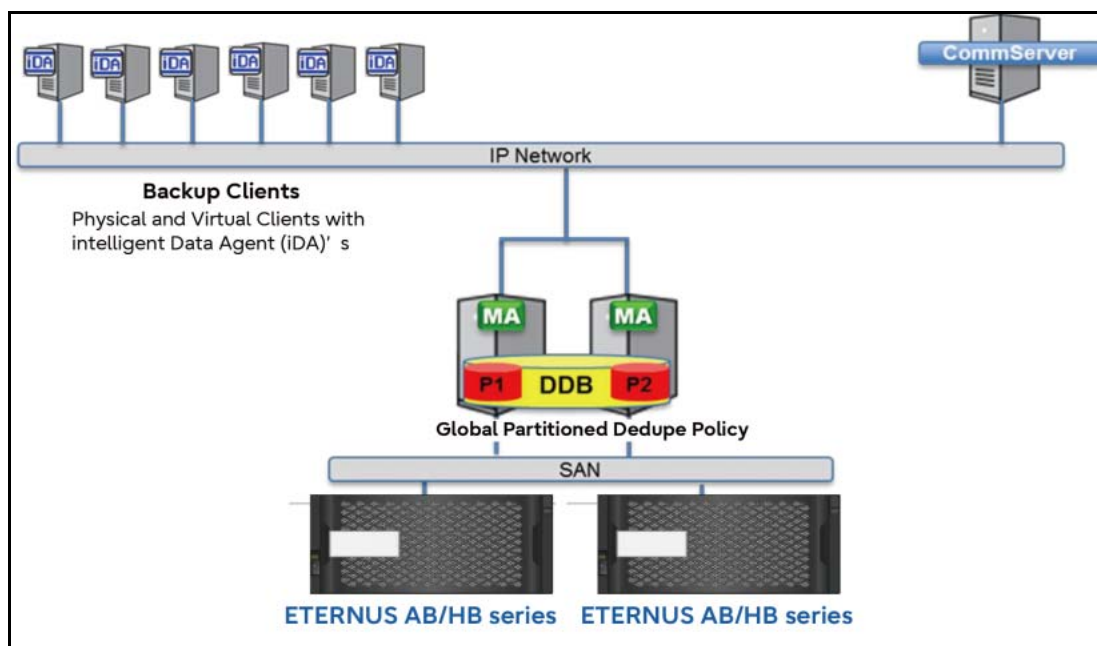


Table 5 Extra-Large configuration: 2-node or larger

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	V11 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 drive for OS drive in RAID 1 • 2TB usable capacity volume (RAID 1 or RAID 5) with SSD class drives/ 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 drives/ 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 (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	ETERNUS AB/HB series

Enterprise Extended Mode Configuration: Multiple Copies and Long-Term Retention

In extended mode, each MediaAgent hosts two Deduplication Database Backup (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 retained 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 second copy are 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 Cross-site replication

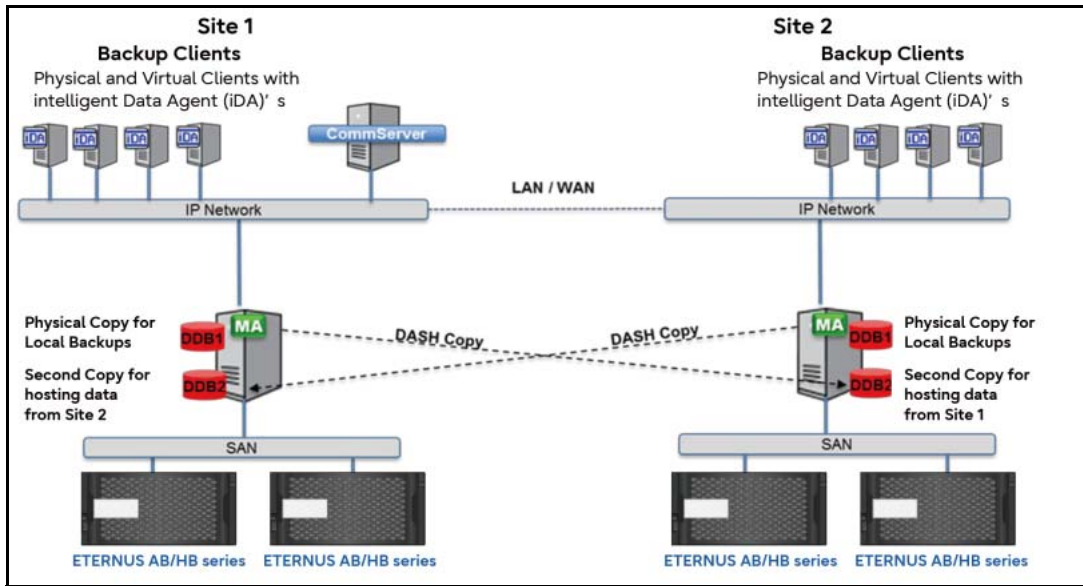
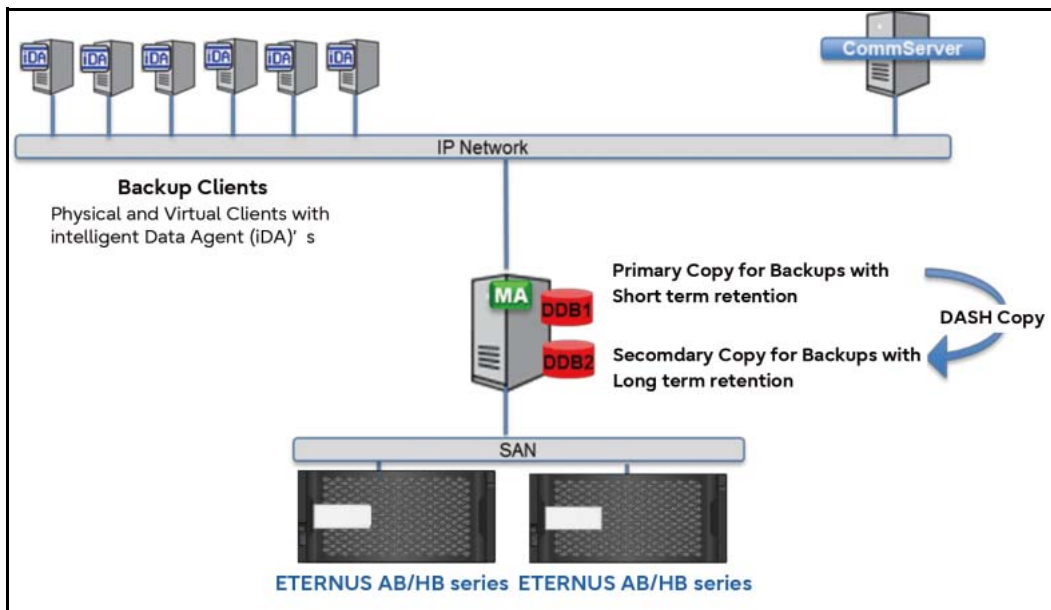


Figure 7 Long-term retention



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 the ETERNUS AB/HB series storage system for creating LAN-free backup copies. The MediaAgent requires either iSCSI or FC access to the storage system to access snapshots. iSCSI access can be provisioned via the 10GigE 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 6](#) 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 6 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 7 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

3. Enterprise Reliability and Availability

Field-proven technology protects your valuable data.

- The ETERNUS AB/HB series is the right choice for peace of mind because it:
 - Leverages knowledge from 1 million systems
 - Is backed by a worldwide support organization
- The ETERNUS AB/HB series is architected for the highest reliability and availability and includes:
 - A fully redundant I/O path with automated failover
 - Online configuration, expansion, and maintenance
 - Advanced monitoring and diagnostic features that enable fast problem resolution
 - Proactive tracking of SSD wear life and sending of alert messages
- The ETERNUS AB/HB series provides enterprise data protection, including:
 - Robust disaster recovery (sync and async)
 - Local protection with high-efficiency Snapshot copies

4. Commvault Data Protection Platform

Overview

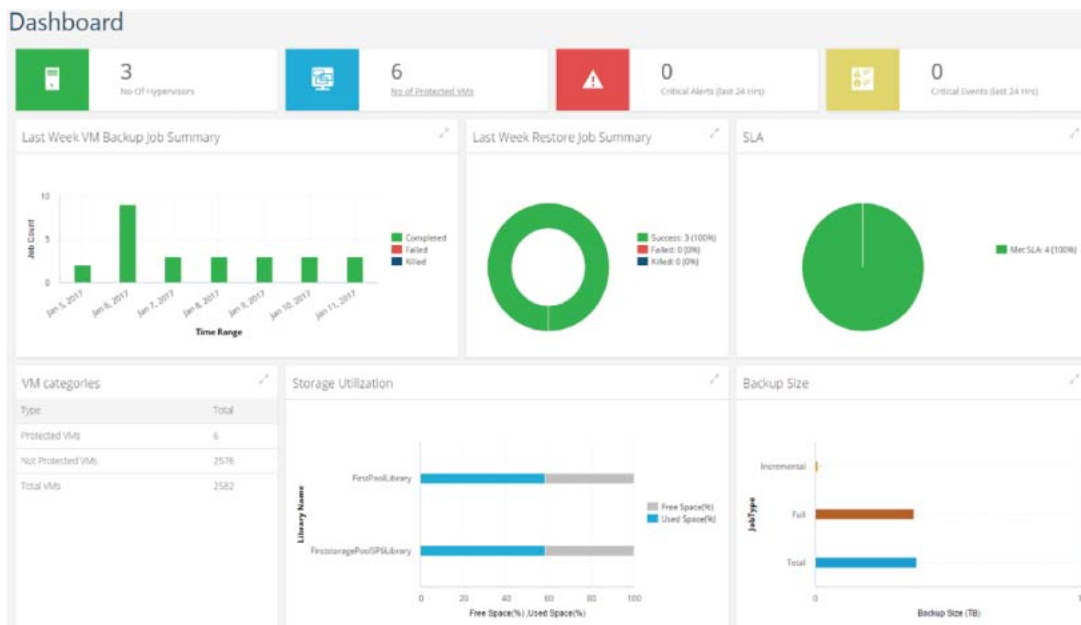
The Commvault Data Protection Platform is a highly scalable, integrated data and information management solution, built from the ground up on a single platform and unified code base. All functions share the same back-end technologies to deliver the unparalleled advantages and benefits of a truly holistic 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 drive, 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 console; view, manage, and access all functions and all data and information across the enterprise.

Figure 8 Commvault management dashboard



- Multiple protection methods include backup and archive, snapshot management, replication, and content indexing for e-discovery.
- Efficient storage management uses deduplication for drive, tape, and cloud. For a list of cloud targets supported, go to the [Commvault Supported Cloud Storage Products 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.
- Cutting-edge 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.

Refer to the Commvault website for details.

<https://www.commvault.com/>

5. Data Protection Platform 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 are counted 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 CommServe and DR 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 8 Deduplication terminology

Deduplication Component	Description
Source-side deduplication	Only unique blocks 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 drive.
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 drive 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 drive copies?
You need to know how long the data will be retained on drive — 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 drive 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 drive provides shared drive libraries that are required in order to use partitioned deduplication.

■ Deduplication to Drive: Best Practices

Before using deduplication, review the deduplication best practices:

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

For better deduplication performance, use the settings described in the rest of this section.

Virtual Server Agent (VSA) 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 RDMs or pass-through devices?

Advanced Data Protection Best Practices

■ Microsoft Active Directory Protection Best Practices

- (1) Install a File System Agent on each domain controller, with the possible exception of RODCs. Perform full system protection.
- (2) Install an Active Directory Agent on a single domain controller in each domain of the forest that requires granular recovery.
- (3) 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 pseudoclient 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 Fujitsu and offline mining.

6. ETERNUS AB/HB series Volume Configuration Guidelines

ETERNUS AB/HB series Configuration Guidelines

Caution

It is not recommended to use thin volumes while setting up the ETERNUS AB/HB series with Commvault.

Small and Medium Configuration - Single Node

- Use RAID 6 (8+2) volume groups.
- Create LUNs (not thin) from the volume group of 9TB if possible, 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 Multinode

- 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 of 9TB if possible, 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 Dynamic Disk Pools, the default segment size (file system typical) is 128KB.

7. Architectural Tools and Sizing

Assessment Tools

Commvault External Data Connector (EDC) Toolkit

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 backup software manufactured by other vendors
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/>.

Design Tool: Software Configurator

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 drive required when Commvault deduplication will be used.
- Estimates the amount of time it will take to replicate data over the WAN in a DASH copy or Continuous Data Replicator solutions.

8. Summary

The ETERNUS AB/HB series and Commvault Data Protection Platform 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.

The ETERNUS AB/HB 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. Data Protection Platform 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 ETERNUS AB/HB series storage solution and Data Protection Platform provide an innovative solution to address all data management needs.

A. Appendix

Commvault Data Protection Platform Terminology and Definitions

Table 9 Commvault Data Protection Platform terminology and definitions

Commvault Terminology	Definition
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 agentsoftware installed on it.
Client Computer Group	A logical grouping of client computers in which selected options can apply to all member clients.
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.
CommServe Database Engine	A SQL server database that is used by the CommServe StorageManager that contains all the information related to the CommCell management group.
CommServe 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.

A. Appendix
Commvault Data Protection Platform Terminology and Definitions

Commvault Terminology	Definition
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 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 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 Fujitsu 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.

A. Appendix
Commvault Data Protection Platform Terminology and Definitions

Commvault Terminology	Definition
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, drive volumes, and so on. See also primary storage.
SLA (Service-Level Agreement)	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.

Commvault Terminology	Definition
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.

ETERNUS AB/HB series Software Documentation

For the ETERNUS AB/HB series SANtricity software documentation, refer to the [Fujitsu manual site](#).

Commvault Data Protection Platform Software Documentation

For Commvault Data Protection Platform software documentation:
<http://documentation.commvault.com>

Fujitsu Storage
ETERNUS AB series All-Flash Arrays,
ETERNUS HB series Hybrid Arrays
ETERNUS AB/HB series and Commvault Data Protection Platform
Reference Architecture and Storage Best Practices

C140-0051-01ENZ3

Date of issuance: October 2023
Issuance responsibility: Fujitsu Limited

- The content of this manual is subject to change without notice.
- This manual was prepared with the utmost attention to detail. However, Fujitsu shall assume no responsibility for any operational problems as the result of errors, omissions, or the use of information in this manual.
- Fujitsu assumes no liability for damages to third party copyrights or other rights arising from the use of any information in this manual.
- The content of this manual may not be reproduced or distributed in part or in its entirety without prior permission from Fujitsu.


FUJITSU