


SP17 Newsletter

September 15, 2019



Contents

Complete Backup And Recovery	3
Optimized Scan and Multi-Node Backup for Nutanix Files Shares	3
Configure the Job Start Time at the Server Level	3
Back Up and Restore Hive Data Warehouse Application Data by Using Workflows	3
Protect Splunk Clusters in the Command Center	4
Use Salesforce Sandbox Data Masking to Protect Sensitive Data	4
Stage-Free Recovery of Cassandra Databases	5
Back Up and Restore Applications in OpenShift	5
Configuring Cascading Network Gateway Connections Using Predefined Network Topologies	6
Protect the Index Server Cloud	6
Cross-Platform Migration of Oracle Databases Using a Workflow	6
ContentStore Mail Server (IMAP) For Exchange Mailbox Agent	7
Filter Backups of CIFS Shares and NFS Exports Using Wildcards	7
Use 1-Touch for Linux to Perform Bare Metal Recovery of Computers That Contain Software RAID.....	7
Onboard and Offboard Tenants and Users with Backup Subscriptions Using REST APIs	8
Configure Log File Settings for Modules in the CommCell Console	8
Configuring a Syslog Server to Forward the System Logs	8
View Operating Instances (OIs) and Virtual Operating Instances (VOIs) in the License Summary Report	9
Journey To The Cloud	9
Back Up and Restore an Azure SQL Database Managed Instance	9
For Google Cloud Platform, You Can Configure a Linux Instance as a Virtual Server Agent Proxy	9
For Microsoft Azure, You Can Configure a Linux Instance as a Virtual Server Agent Proxy	10
Replicate Oracle Cloud Infrastructure Instances	10
Convert Virtual Machines from Azure to VMware	11
IAM Role-Based Authentication for Snapshot Copies of Encrypted Amazon EC2 and Amazon RDS.....	11
Back Up and Restore Object Storage for Amazon S3, Azure Blob, Azure File, and Google Cloud Using... the Command Center	12
Modern Infrastructures	12
Use Horizontal Scaling of Deduplication Databases to Improve Deduplication Performance	12
Manage Snapshots During Dell EMC VPLEX Migration	13
Install Windows OS and Microsoft SQL Server Updates When You Install Commvault Service Pack or.....	13
Use the Credentials Manager feature to Manage Credentials for Disk Storage Libraries and Cloud.....	13
Configure a Remote Snap MediaAgent for the Dell EMC TimeFinder SnapVX Engine	14

Create Snapshots in SRDF/Metro Configuration Using the Dell EMC TimeFinder SnapVX Engine 14

Complete Backup And Recovery

Optimized Scan and Multi-Node Backup for Nutanix Files Shares

You can back up Nutanix Files shares, either CIFS or NFS, using distributed backups that run on multiple data access nodes. Depending on your needs, you can recover Nutanix Files data either in-place or out-of-place.

We have integrated the Nutanix Files Change File Tracking (CFT) API to optimize scan times, which can help reduce incremental backup windows.

Applicable Agents

- Windows File System
- UNIX File Systems

Licenses

Server File System

More Information

- [Nutanix Files](#)

Configure the Job Start Time at the Server Level

You can configure a job start time at the server level. The Commvault software honors the backup window configuration, regardless of the job start time, and runs the next job based on the RPO (recovery point objective) setting at the plan level and the most recent successful backup job for the associated subclient.

More Information

- [Configuring the Job Start Time for a Server](#)

Back Up and Restore Hive Data Warehouse Application Data by Using Workflows

You can back up and restore Hive data warehouse application data by using workflows.

Applicable Agents

Hadoop

Setup Requirements

- Hadoop must be configured on data access nodes.
- The Beeline utility must be installed on the Hive node.

More Information

- [Hive Overview](#)

Protect Splunk Clusters in the Command Center

You can use the Command Center to back up and recover Splunk index files.

Key Features

You can use the Command Center as a self-service portal to perform the following tasks:

- Quickly configure servers, storage, and subclients to back up
- Back up Splunk index files immediately or on a schedule
- Restore Splunk index files to the most recent time, or to a point in time
- Monitor jobs, events, and alerts

More Information

- [Splunk](#)

Use Salesforce Sandbox Data Masking to Protect Sensitive Data

Use Salesforce data masking to change sensitive information when restoring production data to a sandbox. This is useful when populating or refreshing a sandbox for development or test purposes. Data masking allows realistic production data to be used without exposing sensitive information.

The Commvault software applies data-masking policies when you run a restore operation. Restore operations with data masking must be cross-instance restores.

More Information

- [Data Masking for Salesforce](#)

Stage-Free Recovery of Cassandra Databases

You can recover Cassandra databases without restoring the data to a staging location. The stage-free recovery process uses the Commvault 3DFS solution to export the backup data as a network share, and then recovers the data from the share to the Cassandra cluster.

Key Features

The stage-free recovery process includes the following key features:

- Because the data does not need to be staged, the recovery process is faster.
- The recovery process can automatically invoke the Cassandra bulk loader (also called sstableloader) to load the data into the database.

Applicable Agents

Cassandra Agent

Setup Requirements

The stage-free recovery process uses a Linux MediaAgent. If you used a Windows MediaAgent for the backup, then use a Linux MediaAgent as a proxy for the recovery.

More Information

- [Cassandra Restore Operations](#)

Back Up and Restore Applications in OpenShift

You can use the Virtual Server Agent to back up and restore OpenShift applications that run as containerized workloads in the Red Hat OpenShift Container Platform. Organizations that use OpenShift as a platform for rapid development and deployment can use Commvault to protect application-specific data that works with objects that are provided by the OpenShift platform.

This solution protects application metadata and related objects, such as secrets, configuration maps, and persistent volumes. Restoring an application combines that application-specific information with platform-provided objects such as containers.

From the Command Center, you can configure an OpenShift hypervisor and perform backups and restores.

Applicable Agents

Virtual Server Agent

More Information

- [Adding a Red Hat OpenShift Hypervisor](#)

Configuring Cascading Network Gateway Connections Using Predefined Network Topologies

You can use predefined network topologies to configure cascading network gateway connections. In this topology, two groups of clients are each configured to communicate one-way to their own network gateway. Then each network gateway communicates one-way to each other, negotiating connections between the two groups of clients. This is a configuration often used in multi-tenant environments for service providers.

More Information

- [Configuring Cascading Network Gateway Connections Using Predefined Network Topologies](#)
- [Configuring Cascading Network Gateway Connections Using Predefined Network Topologies \(Command Center\)](#)

Protect the Index Server Cloud

You can back up and recover the Index Server Cloud. The data collection is backed up for each node in the cloud.

More Information

- [Backing Up the Index Server Cloud](#)
- [Restoring the Index Directory for the Index Server Cloud](#)

Cross-Platform Migration of Oracle Databases Using a Workflow

You can use the **Unix to Linux Oracle Migration Workflow** workflow to migrate Oracle databases from an AIX or Solaris operating system to a Linux operating system, with minimal impact on production applications.

Setup Requirements

- You can use the workflow for Oracle 11 and more recent versions
- The source and destination hosts must have the same directory structure

More Information

- [Performing Oracle Cross-Platform Migration Using Workflow](#)

ContentStore Mail Server (IMAP) For Exchange Mailbox Agent

You can configure the ContentStore Mail Server (IMAP) for the Exchange Mailbox Agent, and can give end users access to their backed up or archived mailbox messages on their mobile devices that support an IMAP client.

More Information

- [ContentStore Mail Server \(IMAP\) For Exchange Mailbox Agent](#)

Filter Backups of CIFS Shares and NFS Exports Using Wildcards

You can use wildcards to back up CIFS shares and NFS exports with the IntelliSnap feature.

Key Features

- You can browse for CIFS shares when you add content to a subclient.
- You can browse for NFS exports when you add content to a subclient.
- You can back up the default subclient.
- For CIFS shares, you can add `||fileserver|` as the path to the default subclient so that all CIFS shares from a file server are automatically discovered.
- For NFS exports, you can add `fileserver:|` as the path to the default subclient so that all NFS exports from a file server are automatically discovered.

More Information

- [Network Share](#)

Use 1-Touch for Linux to Perform Bare Metal Recovery of Computers That Contain Software RAID Devices

On supported Red Hat Enterprise Linux, SuSE Linux, and Ubuntu operating system versions, you can use 1-Touch for Linux to perform a bare metal recovery of computers that contain software RAID devices.

More Information

- [1-Touch for Linux: System Requirements](#)

Onboard and Offboard Tenants and Users with Backup Subscriptions Using REST APIs

You can now onboard and offboard users to manage the data on clients using REST APIs.

The onboard operation assigns a role to the specific user and adds the user as an owner to the client. The API also assigns a server plan that matches the specified Recovery Point Objectives (RPO) and storage pool settings to the subclient. This operation prepares the subclients and VM clients for data management based on the service level subscription.

The offboard operation removes the user as an owner of the client and disables further data management operations on the subclient or VM client. You can also offboard a user completely and stop data management operations on all the onboarded subclients or VM clients.

Applicable Agents

- File System
- Virtualization

More Information

- [Onboard Subclient](#)
- [Onboard Virtual Machine Client](#)
- [Offboard Subclient](#)
- [Offboard Virtual Machine Client](#)
- [Offboard Tenant User](#)

Configure Log File Settings for Modules in the CommCell Console

You can configure log file settings for individual Commvault software modules at the CommServe level, client computer group level, agent level, or MediaAgent level.

More Information

- [Configuring Log File Settings in the CommCell Console](#)

Configuring a Syslog Server to Forward the System Logs

You can configure a Syslog server to forward the system logs for alerts, audit trails, and events to the server. You can also use a REST API to create a Syslog server and to view the Syslog server information.

More Information

- [Configuring a Syslog Server](#)
- [Syslog Server REST APIs](#)

View Operating Instances (OIs) and Virtual Operating Instances (VOIs) in the License Summary Report

You can view usage information about both the OIs license and the VOIs (VMs) license in the License Summary Report. The **Commvault Complete OI Licenses** table displays the number of licenses that your organization purchased and used. You can view the Operating Instances view or the Virtual Operating Instances (VMs) view to see information about each client computer or VM.

In environments with both VOIs (VMs) and OIs, the software applies surplus OI licenses to VOI (VM) license overages before any license overage conditions occur.

More Information

- [Operating Instances in the License Summary Report](#)
- [Virtual Operating Instances in the License Summary Report](#)

Journey To The Cloud

Back Up and Restore an Azure SQL Database Managed Instance

You can migrate a SQL database from an on-premises SQL server to an Azure SQL database managed instance. The migration is accomplished by backing up the on-premises SQL database, and then restoring it to the Azure SQL database managed instance. After the database is migrated, you can run a full backup operation to back up the Azure SQL database managed instance in the Azure cloud. You can restore an Azure SQL database managed instance from an Azure cloud to either the same Azure cloud instance or to another Azure cloud instance.

More Information

- [Back Up and Restore an Azure SQL Database Managed Instance](#)

For Google Cloud Platform, You Can Configure a Linux Instance as a Virtual Server Agent Proxy

You can configure a Linux instance to act as a Virtual Server Agent (VSA) proxy for Google Cloud Platform. You can use the Linux proxy for backups and restores of Google Cloud Platform instances.

Key Features

With a Linux instance as a VSA proxy, you can perform the following operations:

- Streaming backups
- Full virtual machine restores
- Guest file and folder restores

More Information

- In the Command Center: Google Cloud Platform [System Requirements](#)
- In the CommCell Console: Google Cloud Platform [System Requirements](#)

For Microsoft Azure, You Can Configure a Linux Instance as a Virtual Server Agent Proxy

You can configure a Linux instance to act as a Virtual Server Agent (VSA) proxy for Microsoft Azure. You can use the Linux proxy for backups and restores of Microsoft Azure virtual machines.

Key Features

You can use the Linux proxy for the following operations:

- Streaming backups
- IntelliSnap backups
- Full virtual machine restores
- Guest files and folders restores (for streaming backups and backup copy jobs only)
- Attach disk restores

More Information

- In the Command Center: Microsoft Azure [System Requirements](#)
- In the CommCell Console: Microsoft Azure [System Requirements](#)

Replicate Oracle Cloud Infrastructure Instances

Use Live Sync to replicate instances running in Oracle Cloud Infrastructure. You can replicate instances from a primary site to a destination site, which can be in the same region or a different region. Both the primary site and the destination site must be Oracle Cloud Infrastructure regions.

You can perform replication using the Command Center.

Key Features

- Replication from streaming backups
- Planned or unplanned failover
- Undo failover

Applicable Agents

Virtual Server Agent

Setup Requirements

Deploy at least one VSA proxy in the destination region and in each availability domain where an instance is replicated.

More Information

- [Replication for Virtual Machines](#)

Convert Virtual Machines from Azure to VMware

When restoring an Azure Resource Manager virtual machine from backup, you can choose to restore the VM as a VMware VM. You can perform conversion from streaming backups or IntelliSnap backup copies.

This feature is available from the Command Center and from the Commcell Console.

Applicable Agents

Virtual Server Agent

More Information

- [Converting from Azure to VMware](#) (Command Center)
- [Converting from Azure to VMware](#) (CommCell Console)

IAM Role-Based Authentication for Snapshot Copies of Encrypted Amazon EC2 and Amazon RDS Instances

You can use Amazon EC2 instances that have an IAM role configured to run snapshot replication of encrypted Amazon RDS instances.

Previously, you could run snapshot replication of encrypted Amazon RDS instances only by using the AWS secret key and access key.

More Information

- [Creating an Amazon RDS Snapshot Copy or an Aurora Database Cluster Snapshot Copy in a Different Region](#)

Back Up and Restore Object Storage for Amazon S3, Azure Blob, Azure File, and Google Cloud Using the Command Center

Using Object Storage in the Command Center console, you can back up Amazon S3, Azure Blob, Azure File, and Google Cloud data.

You can restore from an object storage to its original location, to any other object storage, or to disk.

More Information

- [Object Storage](#)
- [Amazon S3](#)
- [Azure Blob](#)
- [Azure File](#)
- [Google Cloud](#)

Modern Infrastructures

Use Horizontal Scaling of Deduplication Databases to Improve Deduplication Performance

Use horizontal scaling of deduplication databases to create separate deduplication databases for three different data types (virtual machines, databases, and file system agents) when you create a storage policy copy.

Horizontal scaling improves deduplication efficiency because similar data types deduplicate more efficiently than dissimilar data types.

Also, the Commvault software automatically creates a new deduplication database when either the free space on a deduplication database partition disk reaches the threshold or the query and insertion (QI) time exceeds the threshold.

More Information

- [Media Management Configuration: Deduplication](#)

Manage Snapshots During Dell EMC VPLEX Migration

When migrating storage within the same cluster of Dell EMC VPLEX storage array, you can select a backend array to create snapshots.

Applicable Agents

[All agents that are supported by the Dell EMC VPLEX storage array](#)

More Information

- [Snapshots During Dell EMC VPLEX Migration](#)

Install Windows OS and Microsoft SQL Server Updates When You Install Commvault Service Pack or Hotfix Updates

You can install Windows OS and Microsoft SQL Server updates when you install Commvault service pack or hotfix updates on a backup appliance, CommServe server, Web Server, or MediaAgent client.

More Information

- [Install Service Pack and Hotfixes Options \(General\)](#)

Use the Credentials Manager feature to Manage Credentials for Disk Storage Libraries and Cloud Storage Libraries

With Credentials Manager, you can store credentials, such as Windows and cloud credentials. Then, instead of manually entering the credentials every time you log on to the Command Center or set up disk libraries or cloud storage libraries, you can simply select the saved credentials.

The Credentials Manager offers the following benefits:

- You can store the credentials in one place and use them wherever they are needed.
- You can share the credentials with other users without revealing the credentials.
- The passwords are encrypted using a CommServe-specific encryption key, a passphrase, or a key provided via an external key management server.

More Information

- [Store Account Information with Credentials Manager](#)
- [Creating a Credential to Access a Network Drive \(Command Center\)](#)
- [Creating a Credential to Access a Cloud Library \(Command Center\)](#)

Configure a Remote Snap MediaAgent for the Dell EMC TimeFinder SnapVX Engine

A remote snap MediaAgent is a remote MediaAgent computer that runs snapshot operations on the storage array instead of running snapshot operations on the MediaAgent computer in the local host.

You can configure a remote snap MediaAgent for Dell EMC TimeFinder SnapVX engine.

More Information

- [Remote Snap MediaAgent](#)

Create Snapshots in SRDF/Metro Configuration Using the Dell EMC TimeFinder SnapVX Engine

With the SRDF/Metro configuration, the TimeFinder SnapVX engine backs up both the local copy and the synchronous remote copy of data. The engine creates snapshots on both sides of the SRDF pair of storage arrays, so if one array is not available, you can use snapshots from the other side of the pair.

Key Features

The primary snapshot and the secondary snapshot are independent, but have the same retention rule as the primary snapshot copy.

Licenses

Install the SRDF/Metro licenses on both storage arrays of the SRDF/Metro pair.

More Information

- [MetroCluster Solutions](#)
- [SRDF/Metro for the Dell EMC TimeFinder SnapVX Engine](#)

©2019 Commvault Systems, Inc. All rights reserved. Commvault, Commvault and logo, the "CV" logo, Commvault Systems, Solving Forward, SIM, Singular Information Management, Simpana, Simpana OnePass, Commvault Galaxy, Unified Data Management, QiNetix, Quick Recovery, QR, CommNet, GridStor, Vault Tracker, Innervault, QuickSnap, QSnap, Recovery Director, CommServe, CommCell, IntelliSnap, ROMS, Commvault Edge, and Commvalue, are trademarks or registered trademarks of Commvault Systems, Inc. All other third party brands, products, service names, trademarks, or registered service marks are the property of and used to identify the products or services of their respective owners. All specifications are subject to change without notice.

The development release and timing of future product releases remains at Commvault's sole discretion. Commvault is providing the following information in accordance with Commvault's standard product communication policies. Any resulting features, functionality, and enhancements or timing of release of such features, functionality, and enhancements are at the sole discretion of Commvault and may be modified without notice. All product roadmap or other similar information does not represent a commitment to deliver any material, code, or functionality, and should not be relied upon in making a purchasing decision.

Visit the [Commvault Documentation](#) website for complete documentation of Commvault products.



[COMMVault.COM](#) | 888.746.3849 | [GET-INFO@COMMVault.COM](#)