



Technical Report

NetApp SolidFire Reference Architecture with Commvault Data Platform v11

Erik Kemp, NetApp and Roy Child, Commvault

December 2017 | TR-4636

In partnership with



Abstract

This document outlines the reference architecture and best practices when using NetApp® SolidFire® storage with Commvault Data Platform v11.

TABLE OF CONTENTS

1	Executive Summary	5
1.1	Customer Challenges	6
1.2	Customer Solutions	6
1.3	About NetApp	6
1.4	About Commvault	6
2	Reference Architecture Overview	6
2.1	Solution Overview	6
3	Commvault Data Platform V11	7
3.1	Overview	7
3.2	Commvault Data Platform V11 Requirements	9
4	NetApp SolidFire Overview	9
4.1	NetApp SolidFire SF-Series Hardware	9
4.2	Key NetApp SolidFire Features	10
4.3	Enterprise Reliability and Availability	10
4.4	Element OS Software	11
5	IntelliSnap Technology with NetApp SolidFire	12
5.1	Benefits of Commvault IntelliSnap Technology	12
5.2	IntelliSnap Technology Anatomy	13
5.3	IntelliSnap for SolidFire Requirements	14
5.4	Configure IntelliSnap for SolidFire Using Commvault Array Management	15
5.5	Edit IntelliSnap Properties of a SolidFire Array	15
5.6	Snapshot Configuration Properties for the SolidFire Storage Array	15
5.7	SolidFire Snapshot Limits	16
5.8	Naming Convention for SolidFire Snapshot Copies Created with IntelliSnap	16
6	Commvault Software V11 Design Considerations and Best Practices	16
6.1	Infrastructure Design	16
7	Solution Sizing and Assessment Tools	18
7.1	Backup Storage Capacity Sizing	18
7.2	Assessment Tools	19
7.3	Design Tool: Commvault Value Portal	20
8	NetApp SolidFire Volume Configuration Guidelines	20
8.1	NetApp SolidFire QoS	20

8.2	NetApp SolidFire Storage Configuration Guidelines for Commvault Data Platform Disk Libraries	20
8.3	SolidFire Volumes as VMware Datastores	21
8.4	Directly Connected SolidFire Volume (File System Backups)	21
9	NetApp SolidFire Host Configuration Guidelines	21
9.1	Host Connectivity	21
10	NetApp SolidFire Plug-In for VMware vCenter	22
10.1	NetApp SolidFire Plug-In for VMware vCenter	22
11	Active IQ	22
12	Conclusion	22
	Where to Find Additional Information	23
	Appendix A: QoS Tuning Scripts	23
	libsfs.ps1	23
	pre-backup.ps1	25
	post-backup.ps1	26
	Appendix B: Commvault Data Platform v11 Terminology and Definitions	26
	Version History	30

LIST OF TABLES

Table 1)	Version levels	14
Table 2)	Support protocols	15
Table 3)	Snapshot configuration properties	16
Table 4)	Snapshot naming convention	16
Table 5)	Deduplication terminology	17
Table 6)	Backup storage capacity sizing	18
Table 7)	NetApp SolidFire volume QoS for Commvault data platform disk libraries	20
Table 8)	NetApp SolidFire volume QoS for VMware datastores	21
Table 9)	NetApp SolidFire volume QoS for directly connected SolidFire volumes	21
Table 10)	Commvault Data Platform V11 terminology and definitions	26

LIST OF FIGURES

Figure 1)	NetApp SolidFire and Commvault Data Platform solution overview	7
Figure 2)	Commvault Data Platform	8
Figure 3)	Commvault Data Platform management dashboard	9
Figure 4)	Element OS UI	12
Figure 5)	IntelliSnap with NetApp SolidFire overview	13

Figure 6) IntelliSnap with NetApp SolidFire VMware overview 14

1 Executive Summary

Optimized Backup and Recovery of SolidFire Storage

Commvault IntelliSnap technology enables a modern approach to operational recovery. By integrating SolidFire Snapshot™ copies with application-level protection, Commvault software is able to reduce the cost and time associated with both protection and recovery operations. Capabilities such as automatic storage discovery and one-step recovery minimize the administrative burden for managing recovery based on Snapshot copies.

NetApp SolidFire Storage

NetApp SolidFire all-flash storage was designed for consolidating multiple, mixed workloads onto a single, agile platform. Quality of service (QoS) guarantees predictable performance across all workloads and tenants, while flexible scale-out options meet your requirements for storage capacity. Start with four nodes of any SolidFire model and mix and match nodes as you flexibly scale for capacity, performance, or both, all while taking advantage of the most current, cost-effective flash technology available. Validated enterprise application solutions and deep integration with popular application and cloud management frameworks make your next-generation data center deployment more cohesive, automated, and dynamically scalable.

IntelliSnap Technology

Protecting application databases through array Snapshot copies provides fast access for recovery and many flexible options for data protection. IntelliSnap technology integrates key application and hypervisor awareness together with the array and the Commvault platform to deliver all of the benefits of traditional backup and recovery with all of the Snapshot performance and capabilities on the array.

This application awareness allows true application-consistent “hot” backups rather than risky crash-consistent Snapshot copies for guaranteed recoverability. Leveraging the hardware to roll back data from Snapshot copies gives shorter, more consistent recovery times even for very large datasets. Coordinated log management enables point-in-time database recovery to shorten RPO.

IntelliSnap technology aligns all of the database, log, and hardware operations, using them in concert to provide these recovery point objective (RPO) and recovery time objective (RTO) reductions without burdening you with scripts to manage. Unlike hardware-based Snapshot copies, IntelliSnap technology extends beyond just creating or deleting Snapshot copies. Snapshot contents are cataloged to allow simple recovery of individual files, databases, virtual machines, and other objects.

Snapshot copies are fully integrated into the data lifecycle; backups captured from Snapshot copies allow long-term retention and recovery of the same recovery points, whether from deduplicated disk, cloud, or tape. Catalog and capture are performed away from production systems to avoid production impact. This content-aware process provides rapid recovery options, whether a full system recovery or a single file is required. The scheduling and retention of Snapshot copies are also managed inside of Commvault software to make sure that only the relevant amount of data is retained for recovery, thus minimizing the overhead on the production arrays.

Commvault and SolidFire: The Bottom Line

Thanks to the extremely high and predictable performance of SolidFire, running multiple mixed workloads, data, and application Snapshot copies is more viable than ever before as a tool to protect and recover data quickly. In addition, IntelliSnap technology enables reuse of data for alternate workloads such as reporting, test/dev refresh, and analytics, without compromising primary storage performance.

Commvault IntelliSnap technology provides an application-aware Snapshot management umbrella that automates and orchestrates the creation, retention, and access of SolidFire Snapshot copies, dramatically simplifying protection and recovery procedures as well as data sharing. Commvault software reduces risk and management overhead by eliminating manual, tedious, and error-prone tasks such as creating and maintaining scripts and truly maximizes the value of your SolidFire investment.

The purpose of this document is to serve as reference architecture for enabling a collaborative backup and recovery solution on NetApp SolidFire with Commvault Data Platform v11.

1.1 Customer Challenges

- Inability to scale capacity or performance independently
- Mixed environment of workloads requiring performance levels
- Management challenges of provisioning and managing thousands of VMs
- Manual processes and scripting of management tasks and policies
- Slow application development and test cycles

1.2 Customer Solutions

- Implementing flash storage to improve workload performance
- Consolidating workloads within a single all-flash array
- Modernizing data protection/technology refresh

1.3 About NetApp

NetApp is the data authority for hybrid cloud. We empower customers to simplify and integrate data management across cloud and on-premises environments to accelerate digital transformation. Together with our partners, we provide a full range of hybrid cloud data services to help global organizations unleash the full potential of their data to expand customer touchpoints, foster greater innovation and optimize their operations. For more information, visit www.netapp.com.

1.4 About Commvault

A singular vision—a belief in a better way to address current and future data and information management needs—guides Commvault in the development of Singular Information Management® solutions for high-performance data protection, universal availability, and simplified management of data on complex storage networks. The Commvault exclusive single-platform architecture gives companies unprecedented control over data growth, costs, and risk.

Commvault software modules are designed to work together seamlessly from the ground up, sharing a single code and common function set, to deliver superlative Backup & Recovery, Archive, Replication, Search, and Resource Management capabilities. More companies every day join those who have discovered the unparalleled efficiency, performance, reliability, and control only Commvault can offer. Information about Commvault is available at www.commvault.com. Commvault's corporate headquarters is located in Tinton Falls, New Jersey.

2 Reference Architecture Overview

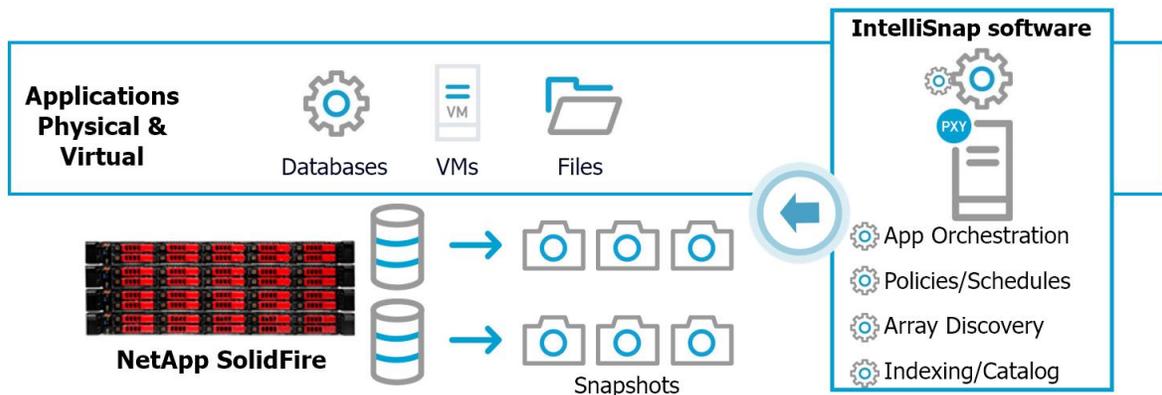
2.1 Solution Overview

For whom is the Commvault and SolidFire solution right?

- Large and medium data centers requiring an easily deployable solution scaling to petabytes of data
- Customers looking to improve their virtual environment performance
- Customers needing to make sure of rapid application recovery

Figure 1 shows an overview of the NetApp SolidFire and Commvault Data Platform solution.

Figure 1) NetApp SolidFire and Commvault Data Platform solution overview.



3 Commvault Data Platform v11

3.1 Overview

The Commvault Data Platform is an enterprise-level, 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, which all 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 properly 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 (UI). End users can access protected data using web browsers or mobile devices.

Key features of the software platform include:

- Complete data protection solution supporting 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.
- Multiple protection methods, including backup and archive, Snapshot management, replication, and content indexing for e-discovery.
- Efficient storage management using deduplication for disk and tape.
- Integrated with the industry's top storage arrays to automate the creation of indexed, application-aware hardware Snapshot copies across multivendor storage environments.
- Complete virtual infrastructure management supporting both VMware and Hyper-V.
- Advanced security capabilities to limit access to critical data, provide granular management capabilities, and provide single sign-on access for Active Directory users, support for OpenID Connect, SAML integration, and open authorization (OAuth).
- Multitenant environment support, which allows a managed service provider (MSP) administrator to mirror real-world organizations by setting up logically separate companies within a single instance of the Commvault software.

- Policy-based data management, transcending limitations of legacy backup products by managing data based on business needs and not physical location.
- Cutting edge end-user experience, empowering them to protect, find, and recover their own data using common tools such as web browsers, Microsoft Outlook, and File Explorer.
- Ability to use third-party screen readers with the web console, admin console, and command-line interface (CLI).

Figure 2) Commvault data platform.

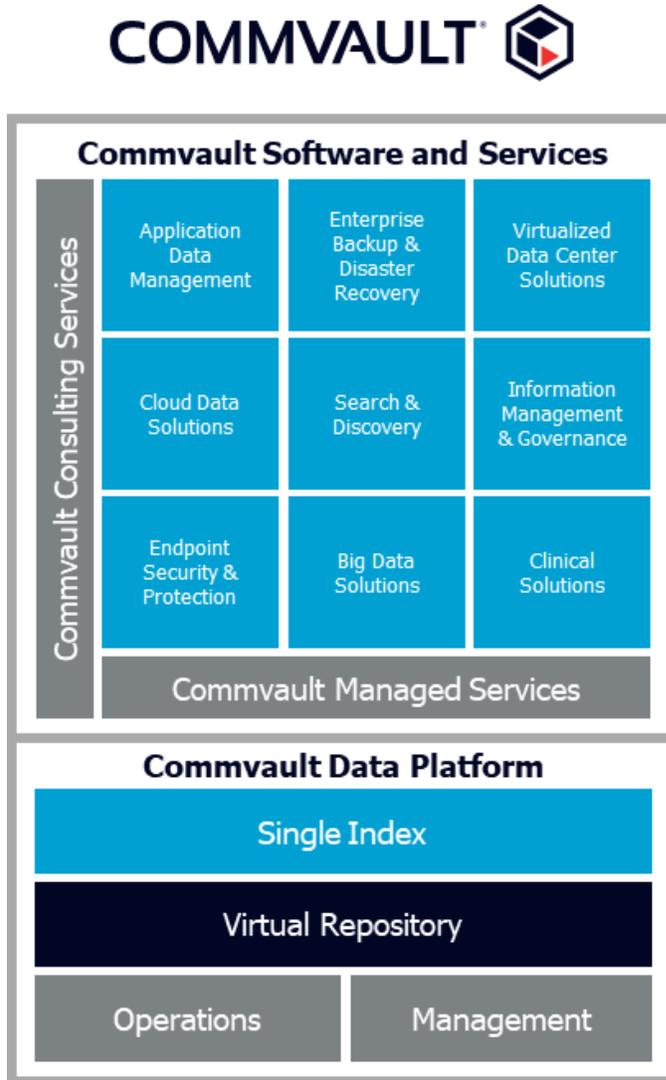
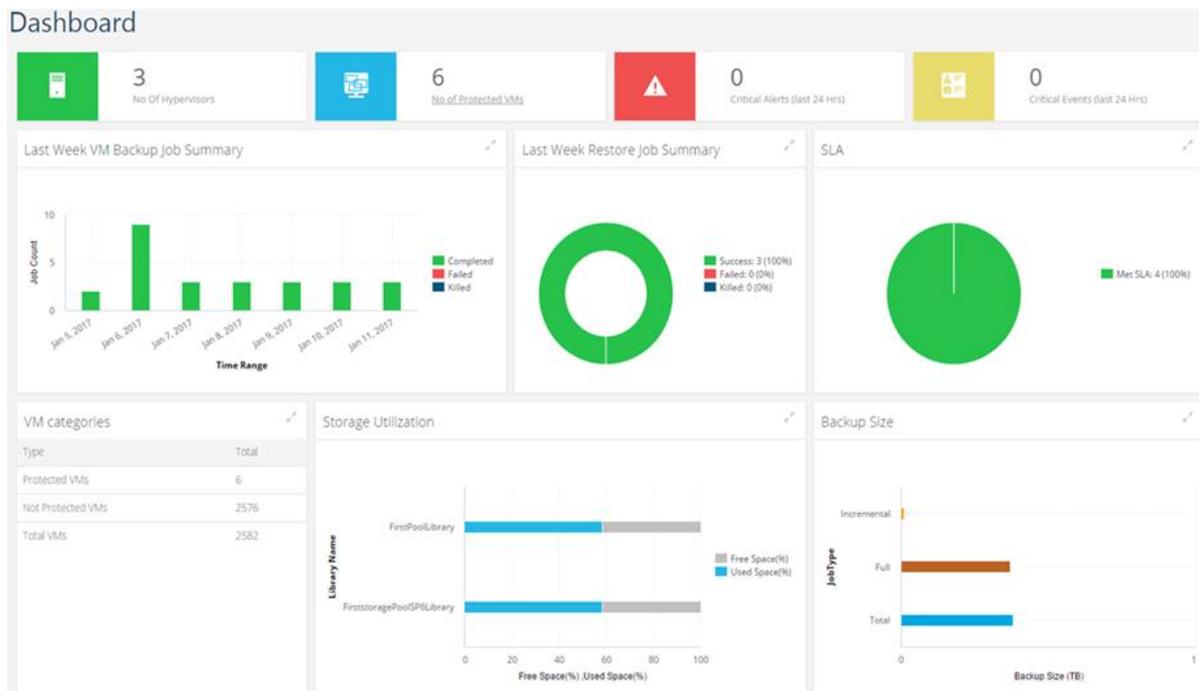


Figure 3 shows the Commvault Data Platform management dashboard.

Figure 3) Commvault Data Platform management dashboard.



3.2 Commvault Data Platform v11 Requirements

For the latest specifications and requirements for the Commvault Data Platform, refer to the following areas of documentation from Commvault:

- Hardware specifications for [CommCell sizing](#)
- [System requirements](#)

4 NetApp SolidFire Overview

4.1 NetApp SolidFire SF-Series Hardware

NetApp SolidFire storage systems are architected from the ground up to be the storage foundation of the next-generation data center's wide-ranging data storage requirements, providing balanced performance. The system is equally adept at handling large sequential I/O for video, analytical, and backup applications, as well as small random input/output (I/O) requirements for small and medium-sized enterprise mixed workloads.

NetApp SolidFire SF-Series storage nodes are the building blocks of a SolidFire platform, allowing you to build a broad range of scale-out, all-flash storage systems that are highly available and easy to control. Mix and match any SolidFire node within a system to take advantage of the most current and cost-effective flash systems available.

Proactive philosophy:

- Scale a single system predictably without performance degradation
- 20TB to petabytes of effective capacity
- 200K to millions of IOPS

The NetApp SolidFire platform compatibility guarantee makes sure of hardware compatibility so customers can seamlessly add new SF-Series nodes to expand their system with the latest flash

technology, rather than being locked into a single generation of flash. This eliminates the hardware replacement, storage migration, and forklift upgrades that are common with many traditional enterprise storage systems, ultimately protecting customers' investment in their infrastructure.

4.2 Key NetApp SolidFire Features

The following are some of the important features of NetApp SolidFire:

Automated management:

- Automatic distribution and load balancing
- Comprehensive API
- Complete management integrations
- Ability to upgrade regularly and with no interruptions

Scale-out:

- Ten of terabytes to multiple petabytes
- Nondisruptive scaling and no downtime
- Ability to mix nodes in same cluster

Guaranteed performance:

- Allocate storage performance independent of capacity
- Manage performance in real time without affecting other volumes
- Guarantee performance to every volume with fine-grained QoS settings

Global efficiency:

- Inline and after compression
- Ability to distribute data automatically (no hot spots) with always-on deduplication
- Global thin provisioning

Data assurance:

- Helix RAID-less data protection
- Real-time replication (synchronous and asynchronous)
- Integrated backup and recovery (cloud)
- Ability to make sure of complete data and performance availability regardless of system condition or application activity
- Ability to safeguard data with 256-bit encryption

4.3 Enterprise Reliability and Availability

- Data assurance
- To provide data redundancy, SolidFire maintains two copies of every data block on two separate nodes: a technology called Helix that is built into our Element® OS software. This allows a cluster to sustain application performance after failures occur. It also takes away the need to have separate storage shelves with shared drive access, making the hardware less complex and less expensive. For instance, SolidFire uses single-attach SSDs instead of more expensive enterprise-grade, dual-attach SSDs.
- The system self-heals quickly, reducing the risk that a second failure will occur before redundancy is restored. Because a cluster responds gracefully to nodes going offline, the feature also facilitates nondisruptive hardware and software upgrades.

- All resources in the system are always in the active pool; there is no need to have spare drives or spare nodes sitting idle in case of failure.
- If a drive fails, the system automatically restores full redundancy by redistributing copies of data using a meshed rebuild process. There are no degraded mode operation and no performance penalty during a rebuild. The process typically completes in 5 minutes or less. Because of the speed with which full redundancy is restored, it provides a level of data protection that exceeds RAID 6 in a typical system.
- If a node fails, because data copies are distributed on separate nodes, all data remains accessible. Connections to the failed node are automatically redirected to other nodes. As with a drive failure, full redundancy is restored quickly and automatically by making sure there are two copies of each block. No matter the failure mode—drive, node, backplane, network failure, software failure—the recovery process is the same. Because the recovery workload is distributed across all nodes in the cluster, redundancy is restored quickly, and no single node (or application workload) takes a performance hit. The more nodes in the cluster, the faster the activity occurs, and the lower the overall impact.

Active support:

- SolidFire Element OS is the storage operating system of the next-generation data center and is included on every SF-Series appliance. The operating system provides all the features you demand from primary storage—total reliability, all-flash performance, end-to-end security, and more—plus all five criteria for delivering agile, predictable storage performance.

Proactive philosophy:

- Secure assist
- 24/7/365 worldwide availability
- Expert tier 3 support engineers only
- Active IQ® assistance

4.4 Element OS Software

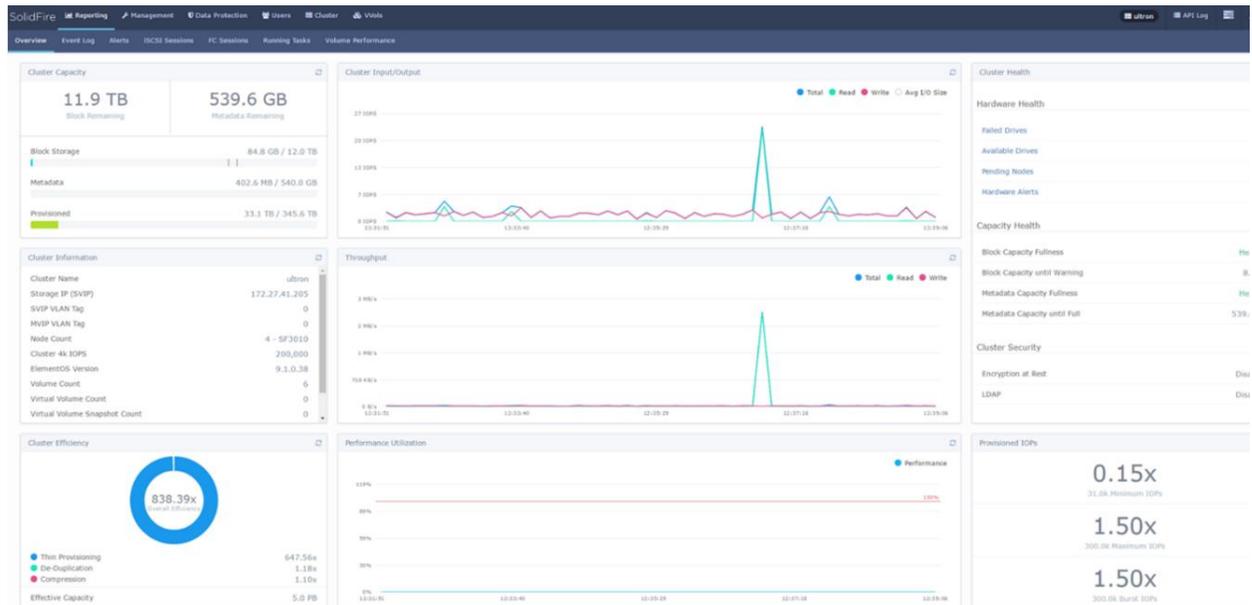
The Element OS continues to deliver the industry’s most comprehensive enterprise feature set for consolidating multiple, mixed workloads onto an agile, secure, and automated infrastructure. By integrating with VMware Virtual Volumes (VVols), increasing Fibre Channel (FC) scalability, and enhancing multitenant network capabilities, the latest Element OS significantly improves flexibility and expands the range of use cases for the SolidFire all-flash scale-out system. Key functionality includes:

- **Integrated VVols.** By integrating SolidFire QoS with VMware VVols, customers can achieve the most granular control over storage performance on a per-virtual machine (VM) basis. This functionality allows them to set min, max, and burst IOPS levels, making sure of exact amounts of capacity and guaranteed performance for even the most performance-sensitive VMs. Both capacity and performance can be changed spontaneously without migrating data or affecting performance.
- **FC scalability.** Connectivity to up to four FC nodes in a SolidFire cluster, allowing users to scale out to 40 storage nodes. 500,000 IOPS per FC node, resulting in 2 million IOPS per a four-FC node cluster.
- **Expansive VLAN features.** Allow for tagging default networks and supporting more flexibility in multitenant networks by enabling customers to use overlapping IP addresses on VLANs through separate routing tables, preserving limited IP address resources.
- **Intuitive UI.** The UI consolidates the storage system data into a single dashboard, saving customers time and resources.

Only SolidFire offers a true shared nothing architecture, which enables customers to upgrade to Element OS nondisruptively. SolidFire has proven the advantages of a shared nothing architecture by offering 100 percent nondisruptive software and hardware upgrades since the SolidFire all-flash scale-out storage platform was released for general availability in November 2012.

Element OS UI

Figure 4) Element OS UI.



5 IntelliSnap Technology with NetApp SolidFire

5.1 Benefits of Commvault IntelliSnap Technology

SolidFire includes hardware-based Snapshot copies, which allow for the creation and deletion of point-in-time views of data, with basic scheduling and retention. IntelliSnap technology significantly extends the value of hardware Snapshot copies:

- IntelliSnap technology adds application awareness and consistency to automate and orchestrate the creation, retention, and access of SolidFire Snapshot copies.
- IntelliSnap software maintains source application context to provide simple, granular recovery. Without IntelliSnap technology, customers have to manually access, locate, and copy data out of Snapshot copies and manage application state during recovery.
- Commvault software preserves source context when backing up from Snapshot copies. With other solutions, the customer must track data across multiple systems and tools during recovery.
- IntelliSnap automation eliminates manual, tedious, and error-prone Snapshot management tasks such as creating and maintaining scripts.
- The Commvault platform provides centralized management to manage protection, retention, search, and reporting for multiple applications, heterogeneous storage arrays, locations, and environments from a single, web-based console.

The benefits of IntelliSnap technology:

First, it's easy. We've seen how it reduces complexity, especially around recovery, and lets you meet your business recovery SLAs. It lets your staff focus on recovery and not on maintaining scripts and doing a lot of manual steps. This can free your IT staff for more value-adding tasks. You really don't want to spend countless hours on script maintenance or keep spending money on expensive professional services engagements to get the work done.

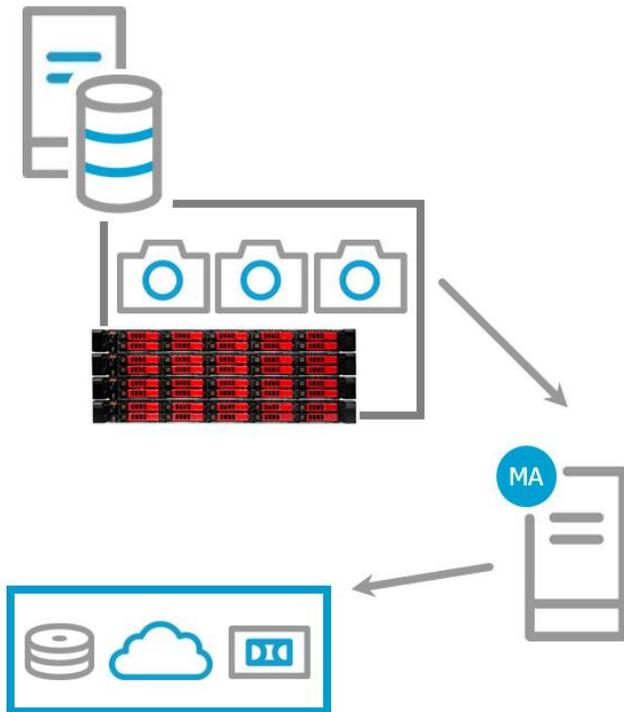
IntelliSnap software is integrated. As we've seen, it understands everything from the application through to the storage and the Snapshot copies. It unifies the two key kinds of recovery: business recovery (when an entire application goes down) and operational recovery (the day-to-day kind of recovery of lost files, deleted e-mails, and so on).

If you're going to use Snapshot copies, we feel that IntelliSnap technology is necessary. It really gives you the things you need. It increases availability of data and visibility into what's going on. It lets you control your data from protecting it to recovering it to how long you want to maintain it. And finally, it reduces the burden on both your infrastructure (by eliminating a lot of the legacy backup needs) and your staff (by dramatically reducing workloads).

5.2 IntelliSnap Technology Anatomy

Figure 5 shows an overview of IntelliSnap with NetApp SolidFire.

Figure 5) IntelliSnap with NetApp SolidFire overview.

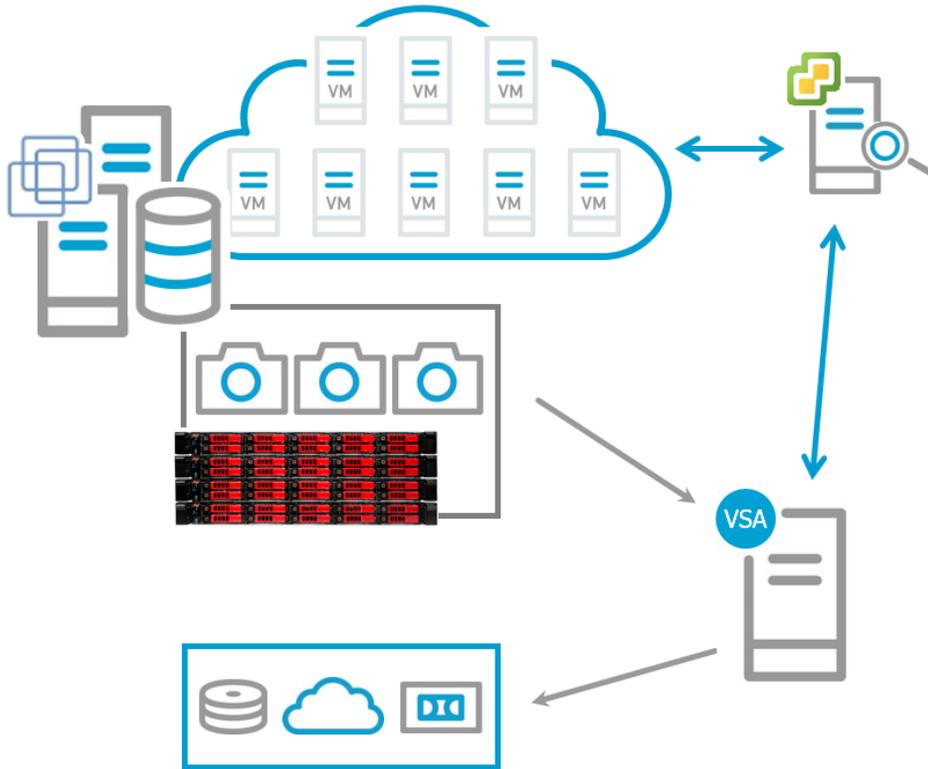


IntelliSnap Operation Workflow:

1. Discover SolidFire array
2. IntelliSnap software communicates with the SolidFire array to take snapshot(s)
3. Index Contents
4. Mount for Backup
5. IntelliSnap software communicates with the storage array to take snapshot(s)

Figure 6 shows an overview of IntelliSnap for NetApp SolidFire in a VMware-specific environment.

Figure 6) IntelliSnap with NetApp SolidFire VMware overview.



IntelliSnap Operation Workflow:

1. vCenter integration provides autodiscovery of new virtual machines for protection.
2. Virtual server agent (VSA) contacts vCenter and creates a consistent VMware Snapshot copy for all VMs being protected.
3. IntelliSnap software communicates with the storage array to take Snapshot copies.
4. Mount Snapshot to ESXi proxy for indexing and LiveBrowse operations; VSA performs indexing operations.
5. Optionally, data can be streamed off array to create a long-term retention copy.

5.3 IntelliSnap for SolidFire Requirements

Table 1 shows the minimum versioning levels of both the SolidFire Element OS and SolidFire REST API necessary for IntelliSnap compatibility.

Table 1) Version levels.

Component	Minimum Version Level
Element OS	SolidFire Element OS 8.0.1.2
API	SolidFire REST API Version 7.0 Note: For initial minimum version level of the Element OS API shown, new API versions add to the SolidFire API, but do not deprecate existing functionality.

Table 2 shows the supported host protocols for IntelliSnap.

Table 2) Support protocols.

SolidFire Array/Node	Protocol
SolidFire array (iSCSI)	iSCSI
SolidFire FC node	FC

5.4 Configure IntelliSnap for SolidFire Using Commvault Array Management

For IntelliSnap backups from Commvault software, add the SolidFire storage array and its configuration information to the CommServe database.

Procedure

1. On the ribbon in the CommCell console, click the Storage tab and then click Array Management. The Array Management dialog box appears.
2. Click Add. The Array Properties dialog box appears.
3. On the General tab, specify the following information:
 - a. From the Snap Vendor list, select SolidFire.
 - b. In the Name box, enter the Management Virtual IP.
 - c. In the Credentials area, click Change. The Enter the User Name and Password dialog box appears.
 - d. Type the credentials of a user who has the cluster admin privileges to access the array and access to reporting and volumes.
 - e. Click OK.
4. To configure additional MediaAgents to perform operations on the storage array, click the Array Controllers tab and then select configuration options. For more information, see [Configuring Array Controllers](#).
5. On the Snap Configuration tab, configure the Snapshot configuration properties for your environment or use the default values. For more information, see [Snap Configuration](#).
6. Click OK.

5.5 Edit IntelliSnap Properties of a SolidFire Array

You can edit the Snap Configuration properties of the storage array to change the behavior of your environment.

Procedure

1. From the CommCell Browser, click the Storage tab and then click the Array Management icon. The Array Management dialog box appears.
2. Select a SolidFire array and then click Edit. The Array Properties dialog box appears.
3. On the Snap Configuration tab, edit the [Snapshot configuration properties for SolidFire](#).
4. Click OK.

5.6 Snapshot Configuration Properties for the SolidFire Storage Array

The Snap Configuration tab contains a set of properties that affects the behavior of the storage array in your environment.

Table 3 shows the configurable Snapshot configuration properties for a SolidFire Snapshot copy.

Table 3) Snapshot configuration properties.

Property	Description
Mount retry interval (in seconds)	The time interval (in seconds) after which the mount operation is no longer retried. The default value is 60.
Mount retry count	The number of retry attempts for a mount operation. The default value is 5.
Create new volume access group if not present	Select this option to create a new volume access group on the array during the map operation, if a volume access group is not already created. The default value is true.
Use iSCSI if FC WWPNS are not connected	Select this option to use iSCSI for the map operation, if FC port worldwide names (WWPNs) are not connected. The default value is false.

5.7 SolidFire Snapshot Limits

There is a maximum of 32 Snapshot copies per volume.

5.8 Naming Convention for SolidFire Snapshot Copies Created with IntelliSnap

IntelliSnap Snapshot copies created on SolidFire have the following naming convention.

Table 4 shows an example of the Snapshot naming convention of a Snapshot copy created on a SolidFire system with the associated attributes from SolidFire and Commvault.

Table 4) Snapshot naming convention.

Snapshot Engine	Configuration	Name Convention	Example
SolidFire Snap	Default	SP-CommCell-JobID-EpochTime	SP-2-5093-1489470

Where:

- SP- is a constant value.
- CommCell is the internal identifier of the CommCell in the CommServe database.
- JobID is the Commvault software JobID.
- EpochTime is a UNIX time stamp.

6 Commvault Software v11 Design Considerations and Best Practices

6.1 Infrastructure Design

Design Questions to Ask Your Customer

- How many sites (that is, the number of sites that exist in a customer's environment)? A site is any location where data exists that needs to be protected or where protected data resides. Don't forget that laptops outside the protected network count as additional sites.

- For each site, where must protection copies reside? You need to know what the protection requirements are for each site, including whether or not a local copy is made and where any remote copies are sent.
- 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.
- 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 5) 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 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 deduplicate together.

Building Block Design Questions to Ask About Sizing

- How much data in each site?
- What is the data composition in each site? (What type of data will be protected?) Breaks down into the following categories:
 - Uncompressed files or file systems.
 - Compressed or uncontrolled data (NDMP data or zipped files fall into this category).
 - Any databases protected at the application level as opposed to dumped by the application are categorized together as well.
- 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 if there are any special retention considerations.)
- 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 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 amount of back-end terabytes exceeds the MediaAgents limits, increase the number of MediaAgents accordingly. Also consider throughput and client count when selecting the number of MediaAgents.
3. Determine whether you will use standalone building block nodes or partitioned deduplication. Size the proper number of MediaAgents 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 best practices at

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

- Hypervisor VMware or Hyper-V?
- How much VM data is being protected at each site? (Get a count of how many guests and data sizing.)
- Are there any physical raw device maps (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 read-only domain controllers (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.

7 Solution Sizing and Assessment Tools

7.1 Backup Storage Capacity Sizing

In a Data Platform v11 environment, backup capacity sizing for a storage system is calculated based on what Commvault refers to as terabyte of front-end protection size, or 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 to plan for storage. The general rule is to plan backup storage capacity at 1.5 to 2 times FET depending on the desired retention time frame.

Table 6) Backup storage capacity sizing.

FET	Backup Storage Capacity Sizing at 1.5x	Backup Storage Capacity Sizing at 2x
15TB	22.5TB	30TB
30TB	45TB	60TB

FET	Backup Storage Capacity Sizing at 1.5x	Backup Storage Capacity Sizing at 2x
60TB	90TB	120TB
120TB	180TB	240TB

7.2 Assessment Tools

Commvault External Data Connector Toolkit

Commvault has made it easy for customers to switch from Veritas NetBackup, EMC Networker, and IBM TSM to Commvault software. 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. Follow these steps to switch:

1. **Import settings from Veritas, EMC, and IBM.** Commvault EDC software connects with Veritas 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.
2. **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 makes sure that customers' configurations and policies are optimized before they start their implementation.
3. **Leverage Commvault's team of product experts.** The Commvault services team is available to help 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 making sure 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 further details, see the [Commvault website](#).

System Discovery and Archive Analyzer Tool

The [System Discovery and Archive Analyzer Tool](#) offers the following functionalities:

- Collects host and file information from a list of servers
- Stages on a virtual machine
- Sizes and reports on the customer environment

Use cases of the tool include:

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

7.3 Design Tool: Commvault Value Portal

The [Commvault Value Portal](#) is an online survey-based configuration tool that generates basic solution design and provides license model comparisons.

8 NetApp SolidFire Volume Configuration Guidelines

8.1 NetApp SolidFire QoS

NetApp SolidFire QoS offers an unprecedented level of control over performance through min IOPS, max IOPS, and burst IOPS. It delivers a specific IOPS and bandwidth.

- Min IOPS is the minimum number of I/O operations per second that an administrator grants to a volume. When an application requests an IOPS value of at least the min IOPS, the SolidFire cluster designates resources to make sure the request is met. The min IOPS level is the guaranteed IOPS and is the focus of most service-level agreement (SLA) provisions. This control is imperative to enable performance-sensitive applications to run in multitenant cloud environments.
- Max IOPS is the maximum number of sustained IOPS that a volume will be allowed to process over an extended period of time. Through this control the noisy neighbor, where one application steals excessive performance from other applications, is eliminated.
- Burst IOPS is the maximum number of IOPS that will be allowed in a short "burst" scenario. If a volume has been running below the max IOPS, burst credits are accumulated. When performance levels become very high and are pushed to maximum levels, short bursts of IOPS will be allowed on the volume. This functionality excels during virtual machine reboots, database flushes, and large file transfers.

8.2 NetApp SolidFire Storage Configuration Guidelines for Commvault Data Platform Disk Libraries

SolidFire volumes make excellent Commvault disk libraries due to their inherent speed and control. For optimal performance, NetApp recommends that you follow these QoS guidelines:

- Run multiple backup jobs to each repository.
- Use multiple volumes per repository for maximum performance.

For NetApp SolidFire volume QoS for backup repositories, see Table 7.

Table 7) NetApp SolidFire volume QoS for Commvault data platform disk libraries.

Parameter	Suggested Tune
Volume QoS	Low minimum IOPS, but sufficient to cover throughput requirements during peak load times. High max IOPS and burst IOPS. Tuning the volume in this manner allows for large volumes that do not consume a lot of the usable IOPS of your SolidFire system but allows for the volume to burst its usage as backups are executed. Example: min:max:burst: 4,000:25,000:25,000
MPIO	SolidFire volumes used as disk libraries should be connected to the backup server using MPIO by following the instructions laid out in the NetApp SolidFire Storage for Microsoft Windows Configuration Guide . This will allow for maximum throughput to the volume.

8.3 SolidFire Volumes as VMware Datastores

SolidFire volumes leveraged as VMware (or other hypervisor) datastores function exactly as any other datastore volumes backed up by Commvault. However, some tuning can be performed on both the SolidFire volume and the systems used for Commvault.

For NetApp SolidFire volume QoS for VMware datastores, see Table 8.

Table 8) NetApp SolidFire volume QoS for VMware datastores.

Parameter	Suggested Tune
Volume QoS	<ul style="list-style-type: none">• Full backups: Increase the max IOPS and burst IOPS to 100,000 for the duration of the backup. Max IOPS and burst IOPS should be returned to their normal levels after the backup completes.• Incremental backups: Increase the max IOPS and burst IOPS by a multiple of four (for example, 1,000:5,000:5,000 would become 1,000:20,000:20,000). Max IOPS and burst IOPS should be returned to their normal levels after the backup completes. <p>Note: These changes can be performed by scripts at the beginning and end of the backup job automatically.</p>

8.4 Directly Connected SolidFire Volume (File System Backups)

This section includes client-side tunes for Windows and Linux systems with directly connected (by iSCSI) SolidFire volumes.

For NetApp SolidFire volume QoS for VMware datastores, see Table 9.

Table 9) NetApp SolidFire volume QoS for directly connected SolidFire volumes.

Parameter	Suggested Tune
Volume QoS	<ul style="list-style-type: none">• Full backups: Increase the max IOPS and burst IOPS to 100,000 for the duration of the backup. Max IOPS and burst IOPS should be returned to their normal levels after the backup completes.• Incremental backups: Increase the max IOPS and burst IOPS by a multiple of four (for example, 1,000:5,000:5,000 would become 1,000:20,000:20,000). Max IOPS and burst IOPS should be returned to their normal levels after the backup completes. <p>Note: These changes can be performed by scripts at the beginning and end of the backup job automatically.</p>

9 NetApp SolidFire Host Configuration Guidelines

9.1 Host Connectivity

iSCSI

All hosts connect to a NetApp SolidFire cluster through the storage virtual cluster IP address. After the connection is established, a redirector service passes the connection to the node that contains the primary metadata copy for the volume for which the host is requesting access. During the node removal process, great care is taken to prevent a host from losing connection to its volume. If there is an active connection from the host to a node that is being removed, the secondary copy is promoted to primary, and a new secondary is created. A seamless iSCSI redirect then occurs, and the application remains connected to its volume. A similar process occurs when adding a node to a cluster. The newly added

node will take on its portion of volumes to help evenly distribute load and capacity on the system. As volumes are rebalanced to provide optimal placement, a seamless iSCSI redirect occurs if the active connection is moved to one of the new nodes added to the system.

FC Nodes

NetApp SolidFire FC nodes provide connectivity to an FC switch, which you can connect to FC clients. FC nodes act as a protocol converter between the FC and iSCSI protocols; this enables you to add FC connectivity to any new or existing NetApp SolidFire cluster.

FC nodes have the following characteristics:

- FC switches manage the state of the fabric, providing optimized interconnections.
- The traffic between two ports flows through the switches only; it is not transmitted to any other port.
- Failure of a port is isolated and does not affect operation of other ports.
- Multiple pairs of ports can communicate simultaneously in a fabric.

10 NetApp SolidFire Plug-In for VMware vCenter

10.1 NetApp SolidFire Plug-In for VMware vCenter

The NetApp SolidFire Plug-in for VMware vCenter is a VMware vCenter Server plug-in that provides integrated management of SolidFire storage arrays from within a VMware vSphere Web Client. The vSphere Web Client is a single management interface that you can use to manage the VMware infrastructure and all your day-to-day storage needs.

Manage SolidFire within vCenter:

- Discover and manage multiple SolidFire clusters
- vSphere datastore create, extend, clone, share, and delete
- SolidFire account create, edit, and delete
- SolidFire volume create, edit, clone, delete, and access group add and remove
- SolidFire access group create, edit, and delete

For the latest download and details, go to

<https://solidfire.brickftp.com/folders/show?path=Downloads%2FvCenter+Plugin>.

11 Active IQ

Active IQ from NetApp SolidFire is a web-based tool that provides continually updated historical views of cluster-wide statistics. Volume, node, drive, and entire cluster data is stored for five years and is updated at regular intervals depending on the data being monitored. Notifications can be set up to alert you about specified events, thresholds, or metrics on a cluster so that they can be addressed quickly when they arise. The Active IQ tool makes monitoring capacity and performance, as well as being informed about cluster health, easy and accessible from anywhere.

For more information about Active IQ, visit <https://activeiq.solidfire.com>.

12 Conclusion

NetApp SolidFire arrays offer extremely high and predictable performance when running multiple, mixed workloads. When combined with the Commvault Data Platform, data and application Snapshot copies are more viable than ever before as a tool to protect and recover data quickly. In addition, IntelliSnap

technology enables reuse of data for alternate workloads such as reporting, test/dev refresh, and analytics, without compromising primary storage performance. Commvault IntelliSnap technology provides an application-aware Snapshot management umbrella that automates and orchestrates the creation, retention, and access of SolidFire Snapshot copies, dramatically simplifying protection and recovery procedures as well as data sharing. The Commvault Data Platform truly maximizes the value of your SolidFire investment.

Where to Find Additional Information

To learn more about the information described in this document, refer to the following documents and/or websites:

- For NetApp SolidFire documentation, see:
 - <https://mysupport.netapp.com/documentation/productlibrary/index.html?productID=62480>
 - <http://mysupport.netapp.com/solidfire/resources>
 - <https://support.solidfire.com>
- For Commvault Data Platform v11 software documentation, see:
 - <http://documentation.commvault.com/commvault/v11/article>

Appendix A: QoS Tuning Scripts

Included are three sample Windows PowerShell scripts that can be used as a basis for creating QoS tuning scripts:

- **libsfs.ps1**. This is a library script that includes helper functions that are called by the other two scripts. It should be saved in the same place as the other two scripts.
- **pre-backup.ps1**. This script increases the maximum and burst QoS settings on a volume for the duration of the backup by the multiplier given. For example, a multiplier of 4 would increase a volume's QoS settings from 1000:4000:4000 to 1000:16000:16000 (min:max:burst). The original QoS settings are saved in temporary file storage (`$env:temp\<MVIP>-<VolumeID>.qos`) in order to restore the original settings after backup. If the multiplier is excluded, it is set to 20, which will set the max and burst IOPS values to at or near the system max of 100000.
- **post-backup.ps1**. This script reads the original QoS settings from the temporary file location and resets the QoS back to the prebackup levels.

libsfs.ps1

```
Function sflog()
{
Param(
[parameter(Position=1)]
[System.String]$Message,
$Color,
[switch] $Deb,
[switch] $Info,
[switch] $Warn,
[switch] $Error
)
if (-not $Message)
{
Write-Host
return
}
if ($Deb) { $fg = "Gray" }
elseif ($Info) { $fg = "White" }
elseif ($Warn) { $fg = "Yellow" }
elseif ($Error) { $fg = "Red" }
```

```

else { $fg = "White" }
if ($Color)
{
$fg = $Color
}
$datestamp = Get-Date -Format "yyyy-MM-dd HH:mm:ss"
$message = $message.Replace("`r`n", "`n")
foreach ($line in $message.Split("`n"))
{
Write-Host -ForegroundColor $fg "$datestamp $line"
}
} Function Ssh-Command()
{
Param(
$Hostname,
$RemoteCommand,
$username = "root",
$password = "solidfire"
) [
reflection.assembly]::LoadFrom((Resolve-Path "
Renci.SshNet.dll") |
Out-Null
{
Param(
$VmHost,
$Rescan = $true
)
sflog -Info "Rescanning iSCSI on $VmHost"
$iscsi_hba = Get-VMHostHba -Type iscsi -VMHost $VmHost
Get-VMHostStorage -VMHost $VmHost -RescanAllHba -RescanVmfs | Out-Null
# Get a list of iSCSI volumes
sflog -Info "Looking for new iSCSI volumes (this may take a few minutes)..."
$luns = Get-ScsiLun -Hba $iscsi_hba
$iqn2volume = @{}
$volume2lun = @{}
foreach ($iscsi_lun in $luns)
{
$path_info = Get-ScsiLunPath -ScsiLun $iscsi_lun
$iqn = $path_info.SanID
$pieces = $iqn.split(".")
$volume_name = ($pieces[4] + "." + $pieces[5])
$iqn2volume[$iqn] = $volume_name
if ($volume2lun.ContainsKey($volume_name))
{
sflog -Error "Duplicate volume name detected. This script cannot handle that case."
exit
} $
volume2lun[$volume_name] =
$
iscsi_lun.CanonicalName
}
# Find the datastore for each LUN
$host_obj = Get-VMHost -Name $VmHost | Get-View
$all_luns = $host_obj.Config.StorageDevice.ScsiLun | ?{$_vendor -match "SolidFir"}
$all_datastores = $host_obj.Config.FileSystemVolume.MountInfo
$lun2ds = @{}
foreach ($ds in $all_datastores | ?{$_Volume.Extent.Count -gt 0})
{
$name = $ds.Volume.Extent[0].DiskName
foreach ($lun in $all_luns)
{ if ($lun.CanonicalName -
eq $
name)
{
$lun2ds[$name] = $ds.Volume.Name
}
}
}
# Create datastores for LUNs
$new_ds = $false
foreach ($iqn in $iqn2volume.keys | Sort-Object)

```

```

{
$volume = $iqn2volume[$iqn]
$lun = $volume2lun[$volume]
if ($lun2ds.ContainsKey($lun))
{
$ds = $lun2ds[$lun]
#sflog -Debug " $iqn -> $volume -> $lun -> $ds"
} else
{
sflog -Info " Creating new datastore '$volume' with iSCSI volume $iqn..."
New-Datastore -VMHost $VmHost -Name $volume -Path $lun -Vmfs | Out-Null
$new_ds = $true
}
}
return $new_ds
}

```

pre-backup.ps1

```

Param( $Mvip = "A.B.C.D",
$SfUsername = "SFQoSadmin",
$SfPassword = "SFPassword",
$VolumeID = "VolumeID",
$IOPSMultiplier = 20
)
#### Usage ####
##
pre-backup.ps1 -Mvip <SF MVIP> -SfUsername <SF User Name> -SfPassword <SF Password> -VolumeID
<Volume ID> [-IOPSMultiplier <Multiplier>]
# Example: pre-backup.ps1 -Mvip 10.10.10.10 -SfUsername qosadmin -SfPassword qospassword -
VolumeID
3 -IOPSMultiplier 4
##
Use in ISE debugger or standard PS shell (not in PowerCLI shell)
#### Helper Functions ####
$workingdir = Split-Path $MyInvocation.MyCommand.Path -Parent
. "$workingdir\libs\ps1"
#### Main Script ####
$ErrorActionPreference = "Stop"
$WarningPreference = "SilentlyContinue"
$ProgressPreference = "SilentlyContinue"
$api_params = New-Object psobject
$api_params | Add-Member -Name "startVolumeID" -Value $VolumeID -MemberType NoteProperty | Out-
Null
$api_params | Add-Member -Name "limit" -Value 1 -MemberType NoteProperty | Out-Null
$vol_details = Call-ApiMethod -Mvip $Mvip -Username $SfUsername -Password $SfPassword -MethodName
ListActiveVolumes -MethodParams $api_params
$vol_details | Export-Clixml $env:temp\$Mvip-$VolumeID.qos
$api_params = New-Object psobject
$api_params | Add-Member -Name "volumeID" -Value $VolumeID -MemberType NoteProperty | Out-Null
$qos_params = New-Object psobject
$qos_params | Add-Member -Name "minIOPS" -Value $vol_details.volumes[0].qos.minIOPS -MemberType
NoteProperty | Out-Null
IF (($vol_details.volumes[0].qos.maxIOPS * $IOPSMultiplier) -gt 100000) {
$qos_params | Add-Member -Name "maxIOPS" -Value 100000 -MemberType NoteProperty | Out-Null } ELSE
{
$qos_params | Add-Member -Name "maxIOPS" -Value ($vol_details.volumes[0].qos.maxIOPS *
$IOPSMultiplier) -MemberType NoteProperty | Out-Null }
IF (($vol_details.volumes[0].qos.burstIOPS * $IOPSMultiplier) -gt 100000) {
$qos_params | Add-Member -Name "burstIOPS" -Value 100000 -MemberType NoteProperty | Out-Null }
ELSE
{ $
qos_params | Add-Member -
Name "
burstIOPS" -
Value ($vol_details.volumes[0].qos.burstIOPS *
$IOPSMultiplier) -MemberType NoteProperty | Out-Null }
$api_params | Add-Member -Name "qos" -Value $qos_params -MemberType NoteProperty | Out-Null
Call-ApiMethod -Mvip $Mvip -Username $SfUsername -Password $SfPassword -MethodName ModifyVolume -
MethodParams $api_params | Out-Null

```

post-backup.ps1

```

Param( $Mvip = "A.B.C.D",
$SfUsername = "SFQoSadmin",
$SfPassword = "SFPassword",
$VolumeID = "VolumeID"
)
#### Usage ####
##
post-backup.ps1 -Mvip <SF MVIP> -SfUsername <SF User Name> -SfPassword <SF Password> -VolumeID
<Volume ID>
# Example: pre-backup.ps1 -Mvip 10.10.10.10 -SfUsername qosadmin -SfPassword qospassword -
VolumeID
3#
# Use in ISE debugger or standard PS shell (not in PowerCLI shell)
#### Helper Functions ####
$workingdir = Split-Path $MyInvocation.MyCommand.Path -Parent
. "$workingdir\libs\ps1"
#### Main Script ####
$ErrorActionPreference = "Stop"
$WarningPreference = "SilentlyContinue"
$ProgressPreference = "SilentlyContinue"
$vol_details = Import-Clixml $env:temp\$Mvip-$VolumeID.qos
$api_params = New-Object psobject
$api_params | Add-Member -Name "volumeID" -Value $VolumeID -MemberType NoteProperty | Out-Null
$qos_params = New-Object psobject
$qos_params | Add-Member -Name "minIOPS" -Value $vol_details.volumes[0].qos.minIOPS -MemberType
NoteProperty | Out-Null
$qos_params | Add-Member -Name "maxIOPS" -Value $vol_details.volumes[0].qos.maxIOPS -MemberType
NoteProperty | Out-Null
$qos_params | Add-Member -Name "burstIOPS" -Value $vol_details.volumes[0].qos.burstIOPS -
MemberType
NoteProperty | Out-Null
$api_params | Add-Member -Name "qos" -Value $qos_params -MemberType NoteProperty | Out-Null
Call-ApiMethod -Mvip $Mvip -Username $SfUsername -Password $SfPassword -MethodName ModifyVolume -
MethodParams $api_params | Out-Null

```

Appendix B: Commvault Data Platform v11 Terminology and Definitions

Table 10 provides a list of Commvault Data Platform v11 terminology and definitions.

Table 10) Commvault Data Platform v11 terminology and definitions.

Commvault Terminology	Definition
Application data management (ADM)	An ADM bundle that provides integrated backup, recovery, and archive in a single solution, eliminating the need to install and manage separate data management applications. The ADM bundle integrates Commvault Simpana 10 software modules for backup and recovery and archive with predefined application reference architectures and operational guidelines to simplify deployment and protection.
Backup series	All of the archive files on given backup media 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 computers	A computer in a CommCell management group that has agent software installed on it.

Commvault Terminology	Definition
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 of the objects in the CommCell management group in a tree structure.
CommCell console	The GUI 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 in the control panel.
Common technology engine	Consists of the CommServe StorageManager and the MediaAgent software modules, which 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 of the information related to the CommCell management group. See System Requirements: CommServe StorageManager for specifications for the CommServe database engine.
CommServe StorageManager	The software module that communicates with all clients and MediaAgents and coordinates operations (data protection, data recovery, 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	Data replication enables you to create 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 of 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

Commvault Terminology	Definition
	hardware, or media. A disaster recovery strategy typically involves the creation and maintenance of a secure disaster recovery site and 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 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 using this storage policy.
Disk library	Storage area network (SAN) 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 of 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 of 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 instances include 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	A feature that enables you to create a point-in-time Snapshot copy 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.

Commvault Terminology	Definition
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	Primary storage refers to data in active use from computer hard disks and/or volumes. See also secondary storage.
Replication policy	A centralized template through which configuration of replication sets/replication pairs within a CommCell management group can be accomplished. A replication policy consists of a common configuration for replication set and replication pairs that can be applied to target replication set/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	Secondary storage refers to backup or archival data moved to storage media, such as tape media, disk volumes, and so on. See also primary storage.
Service-level agreement (SLA)	Measures the data protection coverage aspects in short-term and various long-term intervals in order to determine if the data protection coverage for CommCell client, application, or subclient content is within an acceptable level or not.
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 particular 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/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 that allows you to control all aspects of data management from a single console in virtualized environments for both Microsoft Hyper-V and VMware. This includes data protection, archive, replication, and reporting.

Version History

Version	Date	Document Version History
Version 1.0	October 2017	Initial release.
Version 1.1	December 2017	Added the "Where to Find Additional Information" section to the back of the document and made other minor edits.

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 © 2017 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—without 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.

RESTRICTED RIGHTS LEGEND: Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (c)(1)(ii) of the Rights in Technical Data and Computer Software clause at DFARS 252.277-7103 (October 1988) and FAR 52-227-19 (June 1987).

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-4636-1217