



**Commvault® IntelliSnap®
Technology FlashArray Integration
for Microsoft SQL Server® 2014**

October, 2015

Contents

Executive Summary	4
Goals and Objectives	4
Audience	4
Pure Storage Overview	6
FlashArray//m Specifications	8
Purity Operating Environment	8
Pure1	9
Experience Evergreen Storage	9
Commvault® IntelliSnap™ Technology Overview	10
Commvault Data Protection Overview	11
Commvault IntelliSnap Technology and Microsoft SQL Server	11
System Requirements	12
Pre-Requisites	12
FlashArray Configuration for Commvault	13
FlashArray Credentials Setup	13
Commvault IntelliSnap Software Array Management Setup	17
Commvault IntelliSnap Storage Resources Library Setup	20
Commvault IntelliSnap Storage Policy Setup	22
Commvault IntelliSnap Technology and SQL Server Technology Functions	28
Use Case 1: Create IntelliSnap technology array-based snapshot of SQL Server	29
Use Case 2: Mount an IntelliSnap Array-based Snapshot	41
Use Case 3: Create a Transaction Log and Recover to a Point-in-Time	49
Use Case 4: Recover to a Point-in-Time using Hardware Revert	56

Summary.....	62
References.....	62
Appendix A – Create API Token.....	63
Appendix B – Retrieve API Token	63
Appendix C – PureStorage_Commvault_TestDb	64
About the Author.....	65

Executive Summary

Enterprises today increasingly turn to array-based snapshots to augment or replace legacy data protection solutions that have been overwhelmed by exponential data growth. Management and automation are an integral part of being able effectively leverage this technology. Efficient and integrated use of array-based snapshots are key requirements to protecting virtualized applications.

Pure Storage FlashArray snapshots are an intrinsic part of the way the Purity Operating Environment reduces complexity and maintains efficiency. FlashArray snapshots delivers superior space efficiency, high scalability, and simplicity of management. FlashArray snapshots are always thin provisioned with no dedicated space allocated upon creation. As new or changed data is written to the source volume, new capacity is assigned to the newly-written or overwritten blocks while unchanged blocks are shared between snapshots and volumes. Additionally, the Purity Operating Environment data structures allow snapshots to preserve the granular data reduction efficiencies of volumes through global deduplication and compression, thus volume snapshots require minimal physical capacity on flash drives. Since FlashArray snapshots are entirely metadata constructs, they can be created from a volume in a matter of milliseconds regardless to the size or written capacity of the source volume.

Commvault adds orchestration and functional value on top of the Pure Storage FlashRecover Snapshot technology with the IntelliSnap™ technology snapshot management feature set. IntelliSnap technology streamlines and simplifies snapshot management by centralizing snapshot management across one or many storage arrays; automating object, application and database recovery; and linking snapshots to backup processes. The tight coupling of managed snapshots along side of data protection and recovery operations enables Commvault software to provide a complete view into data across applications, devices, operating systems and locations, cutting administrative overhead and improving access, availability and IT efficiency.

Goals and Objectives

This paper provides an overview of the IntelliSnap technology's technical capabilities in a Microsoft SQL Server environment. Initial configuration of the FlashArray object in Commvault and a walkthrough of various recovery option use cases and their workflows.

In-depth instructions on the configuration, setup and use of the Commvault IntelliSnap technology is beyond the scope of this paper. For additional details, please refer to [Commvault documentation](#). Backup and recovery performance testing is out of scope this this document.

Audience

This paper is written for database administrators, storage or backup administrators interested in Commvault IntelliSnap technology integration with the Pure Storage FlashArray snapshot technology. Familiarity with Microsoft SQL Server, Pure Storage and Commvault is recommended.

© 2015 Pure Storage, Inc. All rights reserved. Pure Storage, the "P" Logo, and Pure1 are trademarks or registered trademarks of Pure Storage, Inc. in the U.S. and other countries. IntelliSnap are registered trademarks of Commvault in the U.S. and other countries. Windows and SQL Server are registered trademarks of Microsoft Corporation in the U.S. and other countries. The Pure Storage product described in this documentation is distributed under a license agreement and may be used only in accordance with the terms of the agreement. The license agreement restricts its use, copying, distribution, decompilation, and reverse engineering. No part of this documentation may be reproduced in any form by any means without prior written authorization from Pure Storage, Inc. and its licensors, if any.

THE DOCUMENTATION IS PROVIDED "AS IS" AND ALL EXPRESS OR IMPLIED CONDITIONS, REPRESENTATIONS AND WARRANTIES, INCLUDING ANY IMPLIED WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, OR NON-INFRINGEMENT ARE DISCLAIMED, EXCEPT TO THE EXTENT THAT SUCH DISCLAIMERS ARE HELD TO BE LEGALLY INVALID. PURE STORAGE SHALL NOT BE LIABLE FOR INCIDENTAL OR CONSEQUENTIAL DAMAGES IN CONNECTION WITH THE FURNISHING, PERFORMANCE, OR USE OF THIS DOCUMENTATION. THE INFORMATION CONTAINED IN THIS DOCUMENTATION IS SUBJECT TO CHANGE WITHOUT NOTICE.

Pure Storage, Inc. 650 Castro Street, Mountain View, CA 94041
<http://www.purestorage.com>

Pure Storage Overview

Who knew that moving to all-flash storage could help reduce the cost of IT? FlashArray//m makes server and workload investments more productive, while also lowering storage spend. With FlashArray//m, organizations can dramatically reduce the complexity of storage to make IT more agile and efficient, accelerating your journey to the cloud.



FlashArray//m's performance can also make your business smarter by unleashing the power of real-time analytics, driving customer loyalty, and creating new, innovative customer experiences that simply weren't possible with disk. All by Transforming Your Storage with FlashArray//m.

FlashArray//m enables you to transform your data center, cloud, or entire business with an affordable all-flash array capable of consolidating and accelerating all your key applications.

Mini Size—Reduce power, space and complexity by 90%

- 3U base chassis with 15-120+ TBs usable
- ~1kW of power
- 6 cables

Mighty Performance—Transform your datacenter, cloud, or entire business

- Up to 300,000 32K IOPS
- Up to 9 GB/s bandwidth
- <1ms average latency

Modular Scale—Scale FlashArray//m inside and outside of the chassis for generations

- Expandable to ~½ PB usable via expansion shelves
- Upgrade controllers and drives to expand performance and/or capacity

Meaningful Simplicity—Appliance-like deployment with worry-free operations

- Plug-and-go deployment that takes minutes, not days
- Non-disruptive upgrades and hot-swap everything
- Less parts = more reliability

The FlashArray//m expands upon the FlashArray's modular, stateless architecture, designed to enable expandability and upgradability for generations. The FlashArray//m leverages a chassis-based design with customizable modules, enabling both capacity and performance to be independently improved over time with advances in compute and flash, to meet your business' needs today and tomorrow.

The Pure Storage FlashArray is ideal for:

Accelerating Databases and Applications—Speed transactions by 10x with consistent low latency, enable online data analytics across wide datasets, and mix production, analytics, dev/test, and backup workloads without fear.

Virtualizing and Consolidating Workloads—Easily accommodate the most IO-hungry Tier 1 workloads, increase consolidation rates (thereby reducing servers), simplify VI administration, and accelerate common administrative tasks.

Delivering the Ultimate Virtual Desktop Experience—Support demanding users with better performance than physical desktops, scale without disruption from pilot to >1000's of users, and experience all-flash performance for under \$100/desktop.

Protecting and Recovering Vital Data Assets—Provide an always-on protection for business-critical data, maintain performance even under failure conditions, and recover instantly with FlashRecover.

Pure Storage FlashArray sets the benchmark for all-flash enterprise storage arrays. It delivers:

Consistent Performance—FlashArray delivers consistent <1ms average latency. Performance is optimized for the real-world applications workloads that are dominated by I/O sizes of 32K or larger vs. 4K/8K hero performance benchmarks. Full performance is maintained even under failures/updates.

Less Cost than Disk—Inline de-duplication and compression deliver 5 – 10x space savings across a broad set of I/O workloads including Databases, Virtual Machines and Virtual Desktop Infrastructure.

Mission-Critical Resiliency—FlashArray delivers >99.999% proven availability, as measured across the Pure Storage installed base and does so with non-disruptive everything without performance impact.

Disaster Recovery Built-In—FlashArray offers native, fully-integrated, data reduction-optimized backup and disaster recovery at no additional cost. Setup disaster recovery with policy-based automation within minutes. And, recover instantly from local, space-efficient snapshots or remote replicas.

Simplicity Built-In—FlashArray offers game-changing management simplicity that makes storage installation, configuration, provisioning and migration a snap. No more managing performance, RAID, tiers or caching. Achieve optimal application performance without any tuning at any layer. Manage the FlashArray the way you like it: Web-based GUI, CLI, VMware vCenter, REST API, Windows PowerShell, Python or OpenStack.



FlashArray//m Specifications

	//m20	//m50	//m70
Capacity	- Up to 120+ TBs effective capacity* 5 – 40TBs raw capacity (base chassis)	- Up to 250+ TBs effective capacity* 30 – 88TBs raw capacity (w/shelves)	- Up to 400+ TBs effective capacity* 44 – 136TBs raw capacity (w/shelves)
Performance	- Up to 150,000 32K IOPS** <1ms average latency - Up to 5 GB/s bandwidth	- Up to 220,000 32K IOPS** <1ms average latency - Up to 7 GB/s bandwidth	- Up to 300,000 32K IOPS** <1ms average latency - Up to 9 GB/s bandwidth
Connectivity	8 Gb/s Fibre Channel 10 Gb/s Ethernet iSCSI - Management and Replication ports	16 Gb/s Fibre Channel 10 Gb/s Ethernet iSCSI - Management and Replication ports	16 Gb/s Fibre Channel 10 Gb/s Ethernet iSCSI - Management and Replication ports
Physical	3U 742 Watts (nominal draw) 110 lbs. (49.9 kg) fully loaded 5.12" x 18.94" x 29.72" FlashArray//m chassis	3U – 7U 1007 - 1447 Watts (nominal draw) 110 lbs. (49.9 kg) fully loaded + 44 lbs. per expansion shelf 5.12" x 18.94" x 29.72" FlashArray//m chassis	5U – 11U 1439 – 2099 Watts (nominal draw) 110 lbs. (49.9 kg) fully loaded + 44 lbs. per expansion shelf 5.12" x 18.94" x 29.72" FlashArray//m chassis

* Effective capacity assumes HA, RAID, and metadata overhead, GB-to-GiB conversion, and includes the benefit of data reduction with always-on inline deduplication, compression, and pattern removal. Average data reduction is calculated at 5-to-1, below the global average of the FlashArray user base.

** Why does Pure Storage quote 32K, not 4K IOPS? The industry commonly markets 4K IOPS, but real-world environments are dominated by IO sizes of 32K or larger. FlashArray//m adapts automatically to 512B-32KB IO for superior performance, scalability, and data reduction.

Table 1. Pure Storage FlashArray//m Series.

Purity Operating Environment

Purity implements advanced data reduction, storage management and flash management features, and all features of Purity are included in the base cost of the FlashArray//m.

Storage Software Built for Flash—The FlashCare technology virtualizes the entire pool of flash within the FlashArray, and allows Purity to both extend the life and ensure the maximum performance of consumer-grade MLC flash.

Granular and Adaptive—Purity Core is based upon a 512-byte variable block size metadata layer. This fine-grain metadata enables all of Purity's data and flash management services to operate at the highest efficiency.

Best Data Reduction Available—FlashReduce implements five forms of inline and post-process data reduction to offer the most complete data reduction in the industry. Data reduction operates at a 512-byte aligned variable block size, to enable effective reduction across a wide range of mixed workloads without tuning.

Highly Available and Resilient—FlashProtect implements high availability, dual-parity RAID-3D, non-disruptive upgrades, and encryption, all of which are designed to deliver full performance to the FlashArray during any failure or maintenance event.

Backup and Disaster Recovery Built-In—FlashRecover combines space-saving snapshots, replication, and protection policies into an end-to-end data protection and recovery solution that protects data against loss locally and globally. All FlashProtect services are fully-integrated in the FlashArray and leverage the native data reduction capabilities.



Pure1 Manage—By combining local web-based management with cloud-based monitoring, Pure1 Manage allows you to manage your FlashArray wherever you are – with just a web browser.

Pure1 Connect—A rich set of APIs, plugin-is, application connectors, and automation toolkits enable you to connect FlashArray//m to all your data center and cloud monitoring, management, and orchestration tools.

Pure1 Support—FlashArray//m is constantly cloud- connected, enabling Pure Storage to deliver the most proactive support experience possible. Highly trained staff combined with big data analytics help resolve problems before they start.

Pure1 Collaborate—Extend your development and support experience online, leveraging the Pure1 Collaborate community to get peer-based support, and to share tips, tricks, and scripts.


Experience Evergreen Storage



Tired of the 3-5 year array replacement merry-go-round? The move to FlashArray//m can be your last data migration. Purchase and deploy storage once and once only – then expand capacity and performance incrementally in conjunction with your business needs and without downtime. Pure Storage's vision for Evergreen Storage is delivered by a combination of the FlashArray's stateless, modular architecture and the ForeverFlash business model, enabling you to extend the lifecycle of storage from 3-5 years to a decade or more.

Commvault® IntelliSnap™ Technology Overview

Commvault IntelliSnap technology integrates with native storage array snapshot engines to provide consistent point-in-time recovery copies for large data sets and enterprise applications. IntelliSnap technology quiesces applications or file systems, triggers the storage array-based snapshot, and returns the system to a fully operational state within seconds. By incorporating and linking snapshots with backup and archive operations, software makes more online and offline copies available for recovery while reducing data protection's impact on production systems. IntelliSnap technology harnesses the power of array-based snapshots to accelerate backup and recovery.



IntelliSnap technology integrates with array-specific APIs in order to execute snapshot management functions. These functions include configure, create, retire, mount, mine, dismount, monitor, retain, revert and restore — and are managed and executed in the same way regardless of hardware platform. Thus, IntelliSnap technology can consolidate and standardize snapshot management and snapshot-based recovery across nearly all leading storage platforms.

IntelliSnap technology enables a modernized approach to data protection by merging storage system hardware snapshots directly into the data protection process. IntelliSnap technology integrates tightly with both host applications and with the system software specific to each hardware array. As the central orchestration point between the two, the IntelliSnap software drives snapshot creation, indexes the contents and can then push application-consistent and deduplicated backup, archive or DR copies to secondary storage, tape or cloud. IntelliSnap technology normalizes snapshot operations so they look the same and operate the same way regardless of application or storage platform. For longer-term retention copies, Commvault software offloads deduplication, backup and encryption to a separate (proxy) host to minimize impact to production systems. By automatically integrating application intelligence with hardware snapshots, Commvault software is able to reach through the application and file systems into the storage array, discover volume/disk configurations for the snapshot operations, and coordinate these operations with proper application awareness and log management, minimizing administrative configuration and eliminating any scripting requirements.

The Commvault snapshot menu enables granular retention options, such as hourly snapshots retained for a day or daily snapshots retained for a week. The snapshot menu is also enabled with an option to retain a set number of snapshots, which can help eliminate the days/cycles conversation with storage and application administrators. These indexing and retention changes help align storage snapshot retention with standard data protection operations, and are available for all applications and storage platforms. Smart use of IntelliSnap technology to manage native snapshots should ensure the creation of more recovery points for fast recovery without adding complexity.

Commvault Data Protection 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 platform 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. The Commvault software platform addresses all aspects of data management in the enterprise, while providing infinite scalability and unprecedented control of data and information.




Figure 1. Commvault software.

Commvault IntelliSnap Technology and Microsoft SQL Server

Commvault IntelliSnap Technology integrates with native storage array snapshot engines to provide consistent point-in-time recovery copies for large data sets and enterprise applications. IntelliSnap technology quiesces applications or systems, triggers the storage array-based snapshot, and returns the system to a fully operational state within minutes. By incorporating and linking snapshots with backup and archive operations, software makes more online and offline copies available for recovery while reducing data protection's impact on production systems. IntelliSnap technology harnesses the power of array-based snapshots to accelerate backup and recovery.

IntelliSnap software integration with SQL Server allows you to protect large databases (even those in the extreme TB size range) within a few minutes. The SQL software agent provides consistent backups by

quiescing the database for a few seconds while taking the snap. Multiple point-in-time snapshots enable you to recover the database to any specific point in time.

Rapid recovery is available using an application-aware revert operation. By keeping the recovery within the storage array there is no need to transfer blocks over the network and through a backup server. This provides far faster restores. In addition, many revert operations are delta-block based, meaning only the changed blocks need to be restored to return a volume to a previous state.

IntelliSnap technology takes advantage of SQL Server capabilities and provides advanced features, such as taking snapshots simultaneously on multiple databases; opening or mounting database snapshots on other clients without actual restore; and performing restore operations from snapshots.

System Requirements

The following are required in order to use Commvault IntelliSnap technology with the Pure Storage FlashArray:

- FlashArray 400 series or FlashArray//m
- Purity v4.1.1 (or higher)
- REST API: v1.4 (found in Purity v4.x and higher)
- Fibre Channel or iSCSI Protocol
- Commvault software version 10 service pack 12 or later with Commvault IntelliSnap software license

Pre-Requisites

The following requirements need to be completed before walking through the configuration steps and use cases presented in this document.

1. **Pure Storage Volume** – A single volume needs to be connected and ready to the Commvault IntelliSnap virtual host. Figure 2 illustrates a volume named **E:\Commvault Library** which is connected to the Commvault virtual machine instance.




Figure 2. Pure Storage volume used for the Disk Library.

2. **Microsoft SQL Server** – A version of SQL Server needs to be installed and running with a sample database to be used in the various use cases discussed in this document. This paper uses Microsoft SQL Server 2014 SP1.
3. **Commvault IntelliSnap Technology** – The software related to Commvault IntelliSnap Technology should be setup and configured before proceeding. This includes the iDataAgent for SQL Server has been deployed to the Microsoft SQL Server instance. Please refer to the [References](#) section for more details.

FlashArray Configuration for Commvault

The Pure Storage FlashArray includes everything required to perform IntelliSnap software operations—there is no special licensing, configuration or management appliance needed. A minimal amount of configuration is required inside of Commvault to add and authorize a FlashArray. The following section describes how to register and configure a FlashArray object into the IntelliSnap software.

Please note that this is a one-time per-array configuration that will enable the use of the Pure Storage FlashArray within the Commvault environment.

FlashArray Credentials Setup

To manage a Pure Storage FlashArray from Commvault the credentials must have “storage admin” privileges or higher. The default Pure Storage FlashArray account, “pureuser”, has these privileges and can be used for testing. It is highly recommended to use a specific account for production management.

For this configuration Pure Storage Directory Service integration is being used with an Active Directory account named “Commvault Administrator”. Figure 3 illustrates the Active Directory configuration for the different security groups for integration with the Pure Storage FlashArray.




Figure 3. Active Directory Users.

Once you have determined the credentials to be used for configuration log into the Pure Storage FlashArray Web Management interface.

After logging into the FlashArray Web Management interface an API Token can be created for the Commvault administrator as illustrated in Figure 4.

1. Click **System** tab
2. Click **Me** under Users
3. Click the  (gear icon)
4. Click **Create API Token**




Figure 4. Creating API Token for the logged in user to the FlashArray.

Now an API Token has been created for the Commvault administrator account we need to retrieve the token to be used in the Commvault Array Management setup. Retrieving the API Token requires the same steps involved in creation of the token.

1. Click **System** tab
2. Click **Me** under Users
3. Click the  (gear icon)
4. Click **Show API Token**
5. Highlight the token and copy (CTRL+C)




Figure 5. Viewing the API Token.




Figure 6. Copy the API Token.



Creating and retrieving API Tokens is also possible with Windows PowerShell and the Pure Storage PowerShell SDK. See Appendix A for more details.

Commvault IntelliSnap Software Array Management Setup

Once the FlashArray has been configured with the credentials the next step is to setup the Array Management with Commvault software.

1. Start and login to Commvault console. This does not have to be the same credentials as previously setup with the Pure Storage FlashArray, but those same credentials can be used if desired, and if Commvault external authentication is configured. Once Commvault is running you will be presented the view shown in Figure 7.




Figure 7. Commvault software.

2. Select the **Storage** tab to begin the Array Management setup. Follow the numbered steps to begin adding the Pure Storage FlashArray.




Figure 8. Array Management from Commvault software.

3. After clicking **Add** in Step 2 select the **Snap Vendor** dropdown and pick **PURE Storage**.




Figure 9. Selecting PURE Storage for Snap Vendor.

4. In the **Name** field enter either the IP Address or Fully Qualified Domain Name (FQDN) of the Pure Storage FlashArray.




Figure 10. Enter IP Address or FQDN for the Pure Storage FlashArray.

5. Next click **Change** shown in Step 4 above to enter the credentials from the FlashArray Credentials Setup section. The **User Name** is equivalent to the **logon name** from Active Directory. For example this configuration the Commvault Administrator account's logon is **commvault-admin**.

The API Token from the FlashArray Credentials Setup should be used for the **Password** and **Confirm Password**. The password in this dialog box is not the logon accounts password.



The API Token can also be retrieved using Windows PowerShell and the Pure Storage PowerShell SDK. See Appendix A, Appendix B – Retrieve API Token for more details.

Enter User Name and Password

User Name: commvault-admin

Password: [masked]

Confirm Password: [masked]

OK Cancel Help

Figure 11. Enter storage admin and API Token.



The Pure Storage FlashArray Users API Tokens does not expire, but they can be deleted and re-created by a Pure Storage FlashArray storage admin. If deleted and re-created, make sure to update the Array Management information for the specific FlashArray.

6. Optionally, enter in a **Description** for the FlashArray to provide more detail on the specific system. Now click on the **Snap Configuration** tab. This tab provides the user with advanced settings for the behavior of IntelliSnap technology and FlashArray interaction. Pure Storage and Commvault recommend **leaving the default settings**.

Array Properties

General Snap Configuration Security

1

Configuration

Mount Retry Interval seconds 30

Mount Retry Count 5

Connect to a Host Group ☐

Use Host if Host Group is not available ☐

Enable Diagnostic Logging ☐

Figure 12. Snap Configuration default values.

There are no changes required to be made in the **Security** tab.



Leave all advanced settings at the default values—these settings should only be changed on an as-needed basis ideally at the functional level only. Please note that these settings can be overridden for specific operations.

Commvault IntelliSnap Storage Resources Library Setup

A Storage Resource needs to be created for storage the output of operations that are performed using the IntelliSnap and Pure Storage integration. There are three library types that can be created. For the Pure Storage integration, a Disk Library will be created in this section.

Steps

1. From the Commvault Administrative Interface select the **Storage Resources** → **Libraries** → **Add** → **Disk Library...**



Figure 13. Creating a Disk Library.

2. Enter a meaningful **Name** to describe the Disk Library then select the [...] to select the **Disk Device**.




Figure 14. Add new Disk Library.

3. The **commvault Media Agent** will query the host for any volumes that are available. Figure 15 shows the volume that will be used for the Disk Library (E:\). For this setup create a **New Folder** called **Library**, then click **OK**.




Figure 15. Browse for Folder and create a New Folder.

Figure 16 shows the newly created Disk Library named Pure Storage Library.




Figure 16. New Pure Storage Disk Library.

Commvault IntelliSnap Storage Policy Setup

IntelliSnap functionality requires a Storage Policy in order to control retention and other operations. For the use cases presented in this document a single Storage Policy will be created to execute the various IntelliSnap operations. Multiple client computers can utilize a single Storage Policy depending on the configuration options.

Steps

1. From the Commvault Administrative interface expand the **Policies** and right-click the **Storage Policies** node in the navigation treeview and click **New Storage Policy** to launch the wizard.




Figure 17. New Storage Policy Wizard.

2. Select **Data Protection and Archiving** as this Storage Policy will be used for protecting a Microsoft SQL Server 2014 database. Click **Next**.




Figure 18. Storage Policy Type.

3. Entering a meaningful name for the Storage Policy, **Pure Storage IntelliSnap** and click **Next**.




Figure 19. Storage Policy Name.

4. Select the newly created Disk Library, **Pure Storage Library**, from the dropdown list that was completed in the previous section. This is the library that will be used for backup copies of the snapshots. Click **Next**.




Figure 20. Default Library for Primary Copy.

5. Select **MediaAgent** from the dropdown list, **commvault**, then click **Next**.




Figure 21. Select the MediaAgent.

6. Leave all of the **default values** for the streams and retention management criteria. Click **Next**.




Figure 22. Streams and Retention Criteria, leave default values.

7. Deselect **Yes** as this setup is focused on using the IntelliSnap integration with the Pure Storage FlashArray which provides the data reduction features through the Purity Operating Environment. The deduplication setting only applies if backup copies of the snapshots are being created.




Figure 23. Do not enable deduplication.

8. The final step in the **Storage Policy Wizard** is to review the settings and then click **Finish** to create the policy.




Figure 24. Storage Policy Wizard settings review.

The newly created Pure Storage IntelliSnap storage policy can be seen in the Commvault Administrative Interface.




Figure 25. Pure Storage IntelliSnap Storage Policy shown in the Commvault Administrative Interface.

Right-click the newly created **Pure Storage IntelliSnap** policy and choose **All Tasks > Create New Snapshot Copy**.





Figure 26. Create New Snapshot Copy for the Pure Storage IntelliSnap policy.

Enter a name for the new **Snap Copy Properties** and then select the **Library** and **MediaAgent** from the respective dropdown lists. The **Library** to select is the one created earlier, **Pure Storage Library**, this is the library that will be used for index operations for the snapshots. The **MediaAgent** to select is the one that was used in the Storage Policy Configuration Wizard, **commvault**.



FIG

The final step is to click on the **Retention** tab and see that the **Enable Data Aging** is checked and the only modification that is necessary is to update the **Basic Retention Rule for All Backups** from **Infinite** to **Retain for** a set numbers of days, for example 15 days.




Figure 27. Snap Copy Properties.

All of the prerequisites have been completed and now we can begin walking through several different use cases that illustrate common activities that database or line of business administrators perform.

Commvault IntelliSnap Technology and SQL Server Technology Functions

Commvault software in combination with the IntelliSnap technology and Pure Storage FlashArray integration offer an extensive feature set with rich controls and behaviors. The following section will outline main features of IntelliSnap technology integration with the Pure Storage FlashArray. For detailed information concerning setup and use of the Commvault IntelliSnap software please refer to Commvault documentation in the References section at the end of this document.

The following functions are the main operations directly involved with IntelliSnap technology for Microsoft SQL Server 2014:

1. Create an IntelliSnap technology array-based snapshot of a SQL Server database
2. Mount an IntelliSnap technology array-based snapshot to a SQL Server host
3. Create a transaction log backup and recover to a specific point-in-time

The subsequent use cases are not meant to be an exhaustive list of features but instead is meant to demonstrate common backup and recovery workflows for Commvault IntelliSnap technology and Microsoft SQL Server environments.

The Implemented system configuration used for all of the use cases is based on the following:

- Commvault IntelliSnap software has been deployed as a single CommServe running as a virtual machine based on VMware vSphere 6.0.
- Microsoft SQL Server 2014 has been deployed as a standalone default instance to a physical host named HYPERV-NODE1. This host is a member of a three node Windows Server Failover Cluster.
- The Pure Storage FlashArray is a FA-420 with two 5.5TB shelves running Purity 4.5.8. Connectivity fabric is Fibre Channel (8GB).
- Connectivity for all of the components runs through two Cisco MDS 9148 switches.

In the next several sections several use cases will be outlined with procedures to deploy, configure, manage, backup, mount and recover to point-in-time instances for Microsoft SQL Server. In order to provide an example for backup and recovery a simple database was created, PureStorage_Commvault_TestDb. The details for creating this database can be found in Appendix C. This is a very simple database that provides timestamps as records to show backup and recovery examples.

Use Case 1: Create IntelliSnap technology array-based snapshot of SQL Server

In this use case we will setup the necessary components and create a Pure Storage FlashRecover snapshot for a SQL Server 2014 database.

Procedure

For this use case there is a physical host named **HYPERV-NODE1** that has a three volumes connected: SQL2014-Data-01, SQL2014-Sys-01 and SQL2014-Tempdb-01. Each of these volumes is mapped as a Mount Point in Microsoft Windows Server 2012 R2.




Figure 28. Connected Volumes to HYPERV-NODE1.

The database that will be backed up is named **PureStorage_Commvault_TestDb**. This is a very simple database with one table, **PointInTime**. This is a simple example just to show how backups, restores and point-in-time recovery can be used with Pure Storage and Commvault IntelliSnap Technology.




Figure 29. Database properties for tpce.



As mentioned in the Pre-requisites section all of the components should be installed on the instance of Microsoft SQL Server to be able to setup the Commvault IntelliSnap technology. These include VSS Provider, MediaAgent and MSSQL Server iDataAgent from Commvault IntelliSnap software version 10 R2 or later setup.

Expand the **Client Computers** node in the CommCell Browser. The **hyperv-node1** should be visible, if it is not perform a **Refresh** using the Commvault Administrative interface **View** tab **Refresh** button. Right-Click the **Client Computer, hyperv-node1** to to access the properties.




Figure 30. Client Computer Properties for hyperv-node1.

Click on the **Advanced** button at the bottom of the dialog. With the **Advanced Client Properties** dialog open check the **Enable IntelliSnap** option.




Figure 31. Advanced Client Properties for hyperv-node1.

Expand the **Client Computers** node in the CommCell Browser. The **hyperv-node1** should be visible, if it is not perform a **Refresh** using the Commvault Administrative interface **View** tab **Refresh** button. Expand the **hyperv-node1** and **SQL Server** node.




Figure 32. HYPERV-NODE1 within the Client Computer node.

The next step is to add a New Subclient Database. Right-click on **HYPERV-NODE1** under the SQL Server node and select **New Subclient** and **Database**.




Figure 33. Creating a New Subclient.

This next step involves setting up the **New SQL Subclient**. There are multiple tabs and options that need to be configured for the subclient. The tabs that will be focused on are **General, Content, Storage Device, SQL Settings, and IntelliSnap Operations**.

The first is the **General** tab to set the **Subclient** name and **Description**.




Figure 34. New SQL Subclient Properties.

Click the **Content** tab to configure the **Database List**. The first step is to add the PureStorage_Commvault_TestDb database to the list. Click **Configure** and then in the **Database Configuration** dialog click **Discover**, all of the databases that are the client will be auto-discovered. Next is to select the **Subclient Name** for the auto-discovered databases. As Figure 36 illustrates the master, model, msdb and PureStorage_Commvault_TestDb databases were discovered.

To select the **Subclient Name** for each of the databases click the drop-down in the Subclient Name column and select the **PureStorage_Commvault_TestDb** as the Subclient Name. For master, model and msdb set those to **Do Not Backup** then click **OK**.




Figure 35. New SQL Subclient Properties.




Figure 36. Results after clicking Discover.

Clicking OK navigates back to the main New SQL Subclient Properties dialog. Click on the **Storage Device** tab and then click on the **Data Storage Policy** tab. Click on the **Storage Policy** drop-down list and select the **Pure Storage IntelliSnap** policy that was created previously. Perform the same **Storage Policy** selection on the **Log Storage Policy** tab.




Figure 37. Settings for Data Storage Policy and Log Storage Policy.

Click the **Data Transfer Option** and select **Use Storage Policy Settings**.




Figure 38. Data Transfer Option.

Now select the **IntelliSnap Operations** tab and check to enable the **IntelliSnap (Applicable only to Full and Differential Jobs)**, then using the **Available Snap Engines** dropdown select the **PURE Storage Snap**.




Figure 39. IntelliSnap Operations.

Click the **SQL Settings** tab to update the final settings for the New SQL Subclient. The default settings should be updated to follow the [Pure Storage SQL Server Best Practices](#) to the following values:

- Block Size = 32768 Bytes
- Maximum Transfer Size = 65536 Bytes




Figure 40. SQL Settings, Block Size and Maximum Transfer Size follow Pure Storage Best Practices.

The following tasks have been completed up to this point:

1. Added the Pure Storage FlashArray to Storage Management
2. Created a Library
3. Created a Storage Policy
4. Created and configured a Snapshot Policy Copy
5. Created and configured a Client
6. Created and configured a Subclient

Now with all of the various components configured backups can be taken of the Microsoft SQL Server instance. In these final procedures for Use Case 1 a Commvault IntelliSnap technology snapshot will be created for the PureStorage_Commvault_TestDb database which resides on the Pure Storage FlashArray.

In the CommCell Browser expand the **Client Computers > hyperv-node1 > SQL Server > HYPERV-NODE1** to display the **Subclient** list. **PureStorage_Commvault_TestDb** will be in the list, right-click on the item and select **Backup**.




Figure 41. Backup the subclient TPCE-Database.

After clicking the **Backup Options** dialog with display. By default, the **Full** and **Immediate** options are selected, click **OK** to begin the backup.




Figure 42. Backup Options for PureStorage_Commvault_TestDb.

After starting the **Job** click the **Job Controller** tab to view the running job. Figure 43 shows the running Job ID is **4449**.




Figure 43. Job Controller view of Job ID 4449.

Figure 44 shows the corresponding FlashRecover Snapshot, **SQL2014-Data-01.SP-2-4449-1446812042**, for volume **SQL2014-Data-01**.




Figure 44. Pure Storage FlashArray view of the snapshot (Job ID 4449).

Use Case 2: Mount an IntelliSnap Array-based Snapshot

In this use case the focus is mounting a Pure Storage FlashRecover Snapshot that was created with Commvault IntelliSnap Technology. The objectives are as follows:

1. View existing snapshots of the PureStorage_Commvault_TestDb from HYPERV-NODE1
2. Mount a select snapshot of the PureStorage_Commvault_TestDb back to the same host, HYPERV-NODE1

Procedures

One of the first tasks that can be performed is to check the available snapshots. All of this can be done directly from the Commvault Administrative Interface without the need to use the Pure Storage Management interface.

Right-click on the **HYPERV-NODE1** select **All Tasks > List Snaps** to display the lists of snapshots that have been created for the Client Computer.




Figure 45. Listing all the snaps for HYPERV-NODE1.

Figure 46 and Figure 47 shows the correlation between what is seen in the Commvault and Pure Storage interfaces.

Source Client	Source Path	Mount Host	Mount Path	Application	Job ID	Mount Status	Mount Stat...	Creation Ti...	Storage Po...
hyperv-node1	C:\PURE STORAGE\SQL2014-DATA\			SQL Server	4451	Successfully created snap	11/5/15 9:0...	11/5/15 9:0...	Pure Storag...
hyperv-node1	C:\PURE STORAGE\SQL2014-DATA\			SQL Server	4450	Successfully created snap	11/5/15 9:0...	11/5/15 9:0...	Pure Storag...
hyperv-node1	C:\PURE STORAGE\SQL2014-DATA\			SQL Server	4449	Successfully created snap	11/5/15 8:1...	11/5/15 8:1...	Pure Storag...
hyperv-node1	C:\PURE STORAGE\SQL2014-DATA\			SQL Server	4448	Successfully created snap	11/5/15 7:0...	11/5/15 7:0...	Pure Storag...
hyperv-node1	C:\PURE STORAGE\SQL2014-DATA\			SQL Server	4442	Successfully unmounted	11/5/15 6:4...	11/5/15 6:3...	Pure Storag...
hyperv-node1	C:\PURE STORAGE\SQL2014-DATA\			SQL Server	4441	Successfully created snap	11/5/15 6:1...	11/5/15 6:1...	Pure Storag...

Figure 46. Snapshots created with IntelliSnap.

NAME	PGROUP SNAPSHOT	SNAPSHOTS CREATED
SQL2014-Data-01.SP-4451	446815164	0 GB 2015-11-06 10:55:08
SQL2014-Data-01.SP-4450	446815073	0 GB 2015-11-06 10:53:37
SQL2014-Data-01.SP-4449	446812042	0 GB 2015-11-06 10:03:06
SQL2014-Data-01.SP-4448	446807736	0 GB 2015-11-06 08:51:21
SQL2014-Data-01.SP-4442	446806017	0 GB 2015-11-06 08:22:42
SQL2014-Data-01.SP-4441	446805141	0 GB 2015-11-06 08:08:05

Figure 47. Corresponding IntelliSnap snapshots shown in Pure Storage Management interface.

The two columns highlighted with boxes in Figure 46 and Figure 47 show the snapshot identification for each interface. To mount a specific snapshot, follow the procedures in Figure 45 and Figure 46 and pick one. For this use case example **Job ID 4449** is selected, right-click and select **Mount**.




Figure 48. Mount selected snapshot.

After clicking on **Mount** the **Mount Path** dialog is displayed. Using the dropdown box select the **Destination Client**, for this example the selected snapshot will be mounted back to the same client, **hyperv-node1**. Next select the **Destination Path** using the **Browse** button. The destination path will reflect the available paths of the destination client selected.




Figure 49. Mount Path, hyperv-node1.

For this example, the Destination Path for the snapshot mount will be placed under the existing **Pure Storage** mount path of other volumes being used by Microsoft SQL Server. After determining the mount path location click **New Folder...** to create a unique folder location for the mount. In this example **COMMVault-MOUNT** is used. The folder name can be set to any meaningful name.




Figure 50. Browse for Mount Path.




Figure 51. Create new folder, COMMVault-MOUNT.




Figure 52. New folder created.

Once the Destination Path has been selected or a new folder created click **OK**. This will display the **Mount Path** dialog with the Destination Path entered, click **OK** to begin the mounting process onto the Destination Client, hyperv-node1.




Figure 53. Completed Mount Path setup.

Figure 54 shows the Destination Path and mounted snapshot for the PureStorage_Commvault_TestDb (MDF and LDF) files. At this point the database files can be attached to Microsoft SQL Server and used. The only detail to be aware of for this example is that mounting a database from a volume that has been created from a snapshot will result in a database name collision since the database already exists. This is a simple problem to fix but appending a "-1" to the name of the database file. This problem would not occur if the database was attached to a different host.

Figure 55 shows Microsoft SQL Server Management Studio after selecting to **Attach Database** has been initiated. Navigate to the **Destination Path** and selected the database then click **OK**.




Figure 54. Destination Path and mounted snapshot on hyperv-node1.




Figure 55. Select the PureStorage_Commvault_TestDb to attach.

As mentioned earlier Figure 56 shows the **Attach As** textbox with the “-1” appended, **PureStorage_Commvault_TestDb-1**.




Figure 56. Attach as PureStorage_Commvault_TestDb-1.

Figure 57 shows both the original and snapshot copy of the database with results returned from both databases.




Figure 57. Original and snapshot copy of database.

Moving back over to the Commvault Administration interface the the dialog will show that the snapshot has been mounted successfully with the **Mount Host** and **Mount Path** details.




Figure 58. Successfully mounted snapshot.

To unmount the snapshot from the client right-click the item and choose **Unmount** or **Force Unmount**. Force unmount should only be used in extreme cases.




Figure 59. Unmount snapshot.

Use Case 3: Create a Transaction Log and Recover to a Point-in-Time

So far we have taken snapshots and then mounted entire databases. There are many times administrators need to restore to a previous point-in-time to recovery more quickly and one way of accomplishing that is by applying transaction logs. In this use case a database will be recovered to a specific point-in-time using a transaction log backup taken by Commvault IntelliSnap Technology.

The objectives are as follows:

1. Create several transaction log backups of the PureStorage_Commvault_TestDb while actively adding records. To setup this use case a number of transaction log backups were taken of the PureStorage_Commvault_TestDb while running a simple insert statement for the PointInTime table.
2. Recover a selected transaction log to the PureStorage_Commvault_TestDb to a specific point-in-time.

Procedures

The first step is to create a transaction log backup for the PureStorage_Commvault_TestDb. Navigate to the **HYPERV-NODE1** Client Computer and right-click the **PureStorage_Commvault_TestDb** Subclient Name and click **Backup**.




Figure 60. Begin a new backup.

After selecting **Backup** and prior to actually creating a backup, the **Backup Options** dialog will be displayed. Previously the backup type was Full which created a snapshot of the entire volume. For this use case a **Transaction Log** backup will be selected. There is an option to **Do Not Truncate Log** and this can be used as it applies to your business backup strategy. For this example, the log will not be truncated.

In order to create a backup, select the **Job Initiation** as **Immediate** and click **OK**.




Figure 61. Backup Options for Transaction Log.

After starting the backup navigate to the **Job Controller** tab to view the running job.




Figure 62. Transaction log running.

Once the transaction log backup has completed right-click the **HYPERV-NODE1 > View > Backup History**. Figure 64 shows all of the different Transaction Log backup that are available.




Figure 63. View Backup History.

Job ID	Status	Operation Type	Subclient	Storage Policy	Job Type	Failed Folders	Failed Files	Skipped Files	Start Time	End Time	Duration
4458	✓	C... Backup	PureStor...	Pure Storage IntelliSnap	Transaction Log	0	0	0	11/6/2015 2:16:26 AM	11/6/20... 00:00:29	
4457	✓	C... Backup	PureStor...	Pure Storage IntelliSnap	Transaction Log	0	0	0	11/5/2015 10:15:42 PM	11/5/20... 00:00:28	
4456	✓	C... Backup	PureStor...	Pure Storage IntelliSnap	Transaction Log	0	0	0	11/5/2015 10:00:11 PM	11/5/20... 00:00:27	
4455	✓	C... Backup	PureStor...	Pure Storage IntelliSnap	Transaction Log	0	0	0	11/5/2015 9:46:00 PM	11/5/20... 00:00:29	
4453	✓	C... Backup	PureStor...	Pure Storage IntelliSnap	Transaction Log	0	0	0	11/5/2015 9:31:04 PM	11/5/20... 00:00:30	
4452	✓	C... Backup	PureStor...	Pure Storage IntelliSnap	Transaction Log	0	0	0	11/5/2015 9:16:03 PM	11/5/20... 00:00:30	
4446	✓	C... Backup	PureStor...	Pure Storage IntelliSnap	Transaction Log	0	1	0	11/5/2015 6:40:47 PM	11/5/20... 00:20:39	
4444	✓	C... Backup	PureStor...	Pure Storage IntelliSnap	Transaction Log	0	0	0	11/5/2015 6:37:25 PM	11/5/20... 00:00:28	
4443	✓	C... Backup	PureStor...	Pure Storage IntelliSnap	Trans. Log No Truncation.	0	0	0	11/5/2015 6:35:06 PM	11/5/20... 00:00:27	

Figure 64. Backup history for Transaction Logs.

Before recovering the transaction log from Job ID 4446 the current PureStorage_Commvault_TestDb should be queried to see the maximum point-in-time now. In order to accomplish this navigate to SQL Server Management Studio and run the following query:

```
SELECT MAX(PIT)
FROM [PureStorage_Commvault_TestDb].[dbo].[PointInTime]
```

Figure 65 shows the point-in-time that has been inserted last.




Figure 65. Current Point-in-Time.

Navigate back to the Commvault Administrative interface and select a transaction log backup to recover. For this example, **Job ID 4455** has been selected. Right-click Job ID 4455 and click **Browse and Restore**.




Figure 66. Browse and Restore a Transaction Log backup.

The **Browse and Restore Options** dialog will be displayed. The **Absolute Time** option has been selected automatically since the specific Job ID 4455 was selected from the Backup History. Click **View Content** to continue.




Figure 67. Browse and Restore Options for Transaction Log.

Select **HYPERV-NODE1** checkbox, then **PureStorage_Commvault_TestDb** checkbox and **Recover All Selected...**




Figure 68. Recover All Selected...

The **SQL Restore Options** dialog is displayed next. To restore to a point-in-time select **Point-in-Time**, note that the **Restore Time** is set to the specific time based on the original select from the Backup History. Next select **Recovery** and **Unconditionally overwrite existing database or files**, click **OK** to being the operation.




Figure 69. SQL Restore Options.

Figure 70 shows the transaction log recovery in progress.




Figure 70. Job Controller view of Transaction Log recovery in progress.

Once the transaction log recovery has completed a new inspection of the PureStorage_Commvault_TestDb database to check the new maximum point-in-time record. Figure 71 shows the newly recovered database with a PIT in the past.

- Original Point-in-Time Record = **2015-11-06 02:20:31.993**
- Transaction Log Recovered Point-in-Time Record = **2015-11-05 21:46:23.757**




Figure 71. Updated Point-in-Time from Transaction Log recovery.

Use Case 4: Recover to a Point-in-Time using Hardware Revert

Use case 3 illustrated the ability to recovery to a specific Point-in-Time (PiT) by performing a transaction log backup and restore. Use case 4 will show how to recover to a specific PiT using an advanced restore option, hardware reverts. Using this feature with a Pure Storage FlashRecover Snapshot provides the ability to restore a volume in place by performing an overwrite operation of the volume.

The same procedural steps will be used as illustrated in use case 3 with one additional step to use the hardware revert option via the advanced settings.

The objectives are as follows:

1. Recover a selected backup of the PureStorage_Commvault_TestDb to a specific point-in-time.

Procedures

As part of use case 3 several transaction log backups were created. Using Figure 64 as a reference for the backup history the PureStorage_Commvault_TestDb will be hardware reverted back to the point-in-time of November 5, 2015 at 9:17pm. Figure 72 shows the maximum point-in-time record from the database (2015-11-06 02:16:49.570).




Figure 72. PureStorage_Commvault_TestDb before hardware revert.

As shown in Figure 73, from the CommCell browser expand the **Client Computers** node elements until the **SQL Server** and **HYPERV-NODE1** can be seen. Right-click the **HYPERV-NODE1** select **All Tasks > Browse and Restore**.




Figure 73. CommCell Browser view for All Tasks on HYPERV-NODE1.

The next step is to select a specific **Time Range** as shown in Figure 74. For this example, an **Absolute Time** with a specific **Start Time** and **End Time** are being set. As mentioned earlier the objective is to restore to the point-in-time of November 5 at 9:17 PM. Once the start and end time's have been configured click **View Content**.




Figure 74. Browse and Restore Options, setting a specific Start and End Time.

Select **HYPERV-NODE1** and **PureStorage_Commvault_TestDb** then **Recover All Selected...**




Figure 75. Selecting the Client Computer and Database Name to begin recovery process.

Once the Recover All Selected button has been clicked that will present the **SQL Restore Options** dialog. For this example, a **Point-in-time** recovery will be performed, notice the **Restore Time** has automatically been set. Select **Recovery** and **Unconditionally overwrite existing database or files**. The final step to set this action as a hardware revert is to click the **Advanced** button.




Figure 76. SQL Restore Options.

Clicking the **Advanced** button presents the dialog in order to enable the use of hardware revert. Click the checkbox to enable **Use hardware revert capability if available**, Figure 77. Important to note that using this option with Pure Storage FlashArray is available for any backup because of the integration with FlashRecover Snapshots.




Figure 77. Use hardware revert option.

Once the hardware revert option has been enabled a **Warning** message will be displayed explaining that the operation will overwrite the entire disk or volume. Click **Yes** to continue, then **OK** for the **Advanced Restore Options** and finally **OK** for the **SQL Restore Options** dialog to being the restore operation.




Figure 78. Warning message regarding selecting hardware revert option.

After starting the backup navigate to the **Job Controller** tab to view the running job.

Job ID	Operation	Client Computer	Agent Type	Subclient	Phase	Storage Policy	MediaAgent	Status	Progress	Errors	Delay Reason
4488	Restore	hyperv-node1	SQL Server		Database Restore			Running	95%		

Figure 79. Job controller for hardware revert operation.

Once the hardware revert has completed a new inspection of the PureStorage_Commvault_TestDb database to check the new maximum point-in-time record. Figure 80 shows the newly recovered database with a PIT in the past.

- Original Point-in-Time Record = **2015-11-06 02:20:31.993**
- Transaction Log Recovered Point-in-Time Record = **2015-11-05 21:46:23.757**

The screenshot shows the SQL Server Management Studio interface. The Object Explorer on the left displays the database structure for 'HyperV-Node1'. The main query window shows the following SQL query:

```

/***** Script for SelectTopNRows command from SSMS *****/
SELECT MAX(PIT)
FROM [PureStorage_Commvault_TestDb].[dbo].[PointInTime]

```

The Results pane shows the output of the query:

	(No column name)
1	2015-11-05 21:16:59.490

An orange box highlights the result, and an arrow points to the timestamp **2015-11-05 21:16:59.490** displayed below the table.

Figure 80. Restored point-in-time.

Summary

The combination of the Commvault IntelliSnap Technology and Pure Storage FlashRecover Snapshot technology makes for a powerful and flexible backup mechanism for physical or virtualized Microsoft SQL Server environments. The deployment process of development/test environments and virtual machine/file backup and recovery is simplified and fully-integrated in Commvault providing a complete enterprise-class backup solution.

References

[Commvault New Features – Service Pack 12](#)

http://documentation.commvault.com/commvault/v10/article?p=new_features/service_pack_12.htm

IntelliSnap – Pure Storage – Overview

http://documentation.commvault.com/commvault/v10/article?p=features/snap_backup/pure/overview.htm

IntelliSnap – SQL Server iDataAgent – Deployment

http://documentation.commvault.com/commvault/v10/article?p=features/snap_backup/sql/deployment.htm

Modern Snapshot Management: Enhancing Application Protection and Recovery

<http://www.commvault.com/resource-library/54c910ae4180aa814d0001fc/enhancing-application-protection-and-recovery-with-a-modern-approach-to-snapshot-management-whitepaper.pdf>

Pure1 Community

<https://community.purestorage.com/>

Programming Interfaces Community Page

<http://community.purestorage.com/t5/Interoperability-Best-Practice/Programming-Interfaces-Community-Page/ta-p/5788>

Appendix A – Create API Token

Creating the API Token for Commvault Administrator (commvault-admin) using Windows PowerShell and Pure Storage PowerShell SDK:

Script

```
$Creds = Get-Credential
$FlashArray = New-PfaArray -EndPoint 'csg-fa420-2' -Credentials $Creds -IgnoreCertificateError
New-PfaApiToken -Array $FlashArray -User 'commvault-admin'
```

Output

```
created          name          api_token
-----
2015-10-19T19:53:01Z commvault-admin c42dcd1e-f227-68fa-9c42-36cc3231c128
```

Appendix B – Retrieve API Token

Retrieving the API Token for Commvault Administrator (commvault-admin) using Windows PowerShell and Pure Storage PowerShell SDK:

Script

```
$Creds = Get-Credential
$FlashArray = New-PfaArray -EndPoint 'csg-fa420-2' -Credentials $Creds -IgnoreCertificateError
Get-PfaApiToken -Array $FlashArray -User 'commvault-admin'
```

Output

```
created          name          api_token
-----
2015-10-19T19:53:01Z commvault-admin c42dcd1e-f227-68fa-9c42-36cc3231c128
```

Appendix C – PureStorage_Commvault_TestDb

```
USE PureStorage_Commvault_TestDb
GO

IF NOT EXISTS (SELECT name FROM sysobjects WHERE name = 'PointInTime'
AND TYPE = 'U')
CREATE TABLE PointInTime (PIT datetime)
WHILE ( 1 = 1 )
BEGIN
    INSERT PointInTime SELECT GETDATE()
    WAITFOR DELAY '00:00:001'
END
```

About the Author



As a Solutions Architect, Barkz is creating the foundation knowledgebase for implementing Microsoft server technologies on Pure Storage. Those core items include best practices, reference architectures, management tasks, automation scripts and examples for application extensibility. With more than 20 years of experience with Microsoft solutions, Barkz has been part of all aspects from architecture, user design, development, test, release and administration. Barkz has experience in Windows PowerShell, Windows Server, Microsoft SQL Server (Admin & Development), Microsoft SharePoint Server (Admin & Development), Microsoft Hyper-V, Visual Studio, C# and REST API.

Barkz blog: <http://www.purestorage.com/blog/author/barkz/>

Twitter: [@purepowershell](https://twitter.com/purepowershell)

Demonstration Videos
[YouTube](#)

Pure1 Community
[Programming Interfaces Community Page](#)

Pure Storage PowerShell Toolkit
<https://github.com/barkz/PureStoragePowerShellToolkit>



Pure Storage, Inc.
Twitter: [@purestorage](https://twitter.com/purestorage)
www.purestorage.com

650 Castro Street, Suite #260
Mountain View, CA 94041

T: 650-290-6088
F: 650-625-9667

Sales: sales@purestorage.com
Support: support@purestorage.com
Media: pr@purestorage.com
General: info@purestorage.com