



Implementing Pure Storage[®] with IBM[®] SAN Volume Controller

Simon Dodsley, Global Solutions Architect

Version: 2.2
16 April 2018

Contents

Notices	3
Summary.....	4
Audience	4
Prerequisites	4
Preparation for Pure Storage Integration	4
Confirm Support for the Environment	4
Concurrent Host connections to the Pure Storage and IBM SVC.....	4
Quorum Disk Support.....	5
Record the SVC WWPNs.....	5
Visibility of Pure Storage under SVC.....	6
Integration with IBM Easy Tier	8
SVC Storage Pool Extent Size.....	8
Volume Count and Sizing.....	8
SVC Cache Considerations	9
IBM Real-Time Compression (RtC), Thin Provisioning and Encryption	9
SVC and SCSI UNMAP	10
SVC VDisk Formatting	10
Using SVC for Data Migration to Pure Storage	11
About the Author	12

Notices

© 2016 Pure Storage, Inc. All rights reserved. Pure Storage and the P Logo are trademarks or registered trademarks of Pure Storage, Inc. in the U.S. and other countries. IBM is a trademark of International Business Machines Corporation, registered in many jurisdictions worldwide. All other trademarks are registered marks of their respective owners.

The Pure Storage products and programs described in this documentation are distributed under a license agreement restricting the use, copying, distribution, and decompilation/reverse engineering of the products. 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. Pure Storage may make improvements and/or changes in the Pure Storage products and/or the programs described in this documentation at any time without notice.

THIS 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

Summary

This document will cover the integration of a Pure Storage FlashArray into an existing IBM SAN Volume Controller (SVC) environment. It will provide recommended configuration options and implementation settings.

Audience

The audience for this document is Pure Storage System Engineers, Certified Partners and Customers who will be deploying Pure Storage behind an IBM SVC.

Prerequisites

This document assumes that the IBM SVC has already been deployed in the environment and has been configured using IBM's Best Practices for the environment.

It is also assumed that the audience has a good knowledge and understanding of fibre channel storage environment, the Pure Storage Flash Array and the IBM SVC management interfaces and command line structure.

Preparation for Pure Storage Integration

Integration of any backend storage array into an IBM SVC environment is relatively straightforward but it is essential that initial checks are performed to ensure that the integration will succeed.

Confirm Support for the Environment

The IBM SVC supports all models of the Pure Storage FlashArray, provided the array has the minimum Purity code level as detailed in the table below.

Pure Storage Model	Minimum Purity Code Level
FA-400 Series	3.2.1
FlashArray//m Series	4.5.1

At the time of writing the minimum SVC code level required to support Pure Storage Flash Arrays is 7.2.0.2 and note that IBM supports concurrent code upgrades for the Pure Storage Flash Array with the SVC.

For the latest supported models and code levels please refer to the IBM support website for the SVC.

Concurrent Host connections to the Pure Storage and IBM SVC

The following restrictions apply when sharing a host connection to a Pure Storage and an IBM SVC:

- The same host cannot be connected to both an SVC and a Pure Storage FlashArray if the host is using IBM's SDD multipath drivers as this cannot coexist with the native multipath drivers required by the Pure Storage array. This restriction is removed if the host is using native multipath drivers for the SVC
- Volumes on the Pure Storage cannot be shared between a host and an SVC
- The SVC switch zone must include at least one physical port from each of the Pure Storage controllers in order to have no single point of failure
- The Pure Storage volumes must be presented to all the SVC initiator ports in the storage fabric zone between SVC and Pure Storage

Quorum Disk Support

The IBM SVC can choose managed disks (MDisks) that are presented from the Pure Storage as quorum disks.

Record the SVC WWPNS

You need the SVC worldwide port names (WWPNs) for each of the IBM SVC controller nodes. This will enable the configuration of the SVC host object on the Pure Storage to allow volumes to be presented from the Pure Storage to the SVC.

Obtain the WWPNs for each SVC controller node from the SVC GUI or by using the following example command from the SVC command line interface:

```
IBM_2145:ITSO_SVC2:superuser> svcinfo lsnode 1
id 1
name 1
UPS_serial_number 1000739004
WWNN 50050768010027E2
status online
IO_group_id 0
IO_group_name io_grp0
partner_node_id 2
partner_node_name 2
config_node no
UPS_unique_id 10000000000027E2
port_id 50050768014027E2
port_status active
port_speed 8Gb
port_id 50050768013027E2
port_status active
port_speed 8Gb
port_id 50050768011027E2
port_status active
port_speed 8Gb
port_id 50050768012027E2
port_status active
port_speed 8Gb
```

Use the recorded SVC WWPNs for all controllers to create a host object on the Pure Storage array. This can be done using the Pure Storage GUI or from the command line using the following example command:

```
# purehost create --wwnlist <wwpn1>,<wwpn2>,<wwpn3>,...<wwpnx> SVC
```

Visibility of Pure Storage under SVC

After you have zoned a Pure Storage array to an SVC you will see the Pure Storage array appear in the SVC as a number of controllers. The number of controllers visible will depend on the number of FC ports in your Pure Storage array. Each FC port and associated WWPN will appear as a separate controller under the SVC.

The SVC creates a separate controller based on unique WWNNs therefore, as the fibre channel implementation used by Pure Storage creates a unique WWNN (which are the same as the WWPN) for each FC port in the array a Pure Storage array with 4 FC ports will show as 4 distinct controllers, and an 8 port array will show as 8 distinct controllers.

This screenshot shows an SVC with an 8 port Pure Storage array from the SVCs perspective

Name	Status	Capacity	Mode	Site	Storage Pool	Storage System	LUN
IBM FlashSystem-9840	Online	Serial Number: [REDACTED]		Site: Unassigned	WWNN: [REDACTED]		
IBM FlashSystem-9840	Online	Serial Number: [REDACTED]		Site: Unassigned	WWNN: [REDACTED]		
IBM FlashSystem-9840	Online	Serial Number: [REDACTED]		Site: Unassigned	WWNN: [REDACTED]		
IBM FlashSystem-9840	Online	Serial Number: [REDACTED]		Site: Unassigned	WWNN: [REDACTED]		
controller0	Online	PURE FlashArray		Site: Unassigned	WWNN: 524A9374A971E100		
controller1	Online	PURE FlashArray		Site: Unassigned	WWNN: 524A9374A971E102		
controller2	Online	PURE FlashArray		Site: Unassigned	WWNN: 524A9374A971E110		
controller3	Online	PURE FlashArray		Site: Unassigned	WWNN: 524A9374A971E112		
controller4	Online	PURE FlashArray		Site: Unassigned	WWNN: 524A9374A971E101		
controller5	Online	PURE FlashArray		Site: Unassigned	WWNN: 524A9374A971E103		
controller6	Online	PURE FlashArray		Site: Unassigned	WWNN: 524A9374A971E111		
controller7	Online	PURE FlashArray		Site: Unassigned	WWNN: 524A9374A971E113		

The SVC command line that can be used to see the controllers is as follows:

```
# svcinfo lscontroller
id controller_name ctrl_s/n vendor_id product_id_low product_id_high
0 controller0 PURE FlashArr ay
1 controller1 PURE FlashArr ay
2 controller2 PURE FlashArr ay
3 controller3 PURE FlashArr ay
4 controller4 PURE FlashArr ay
5 controller5 PURE FlashArr ay
6 controller6 PURE FlashArr ay
7 controller7 PURE FlashArr ay
```

More detail for each controller, such as the WWPN fibre channel port on the Pure Storage array, can be obtained using the following example SVC command:

```
# svcinfo lscontroller 1
id 1
```

```

controller_name controller1
WWNN 524A93744B570E13
mdisk_link_count 16
max_mdisk_link_count 16
degraded no
vendor_id PURE
product_id_low FlashArr
product_id_high ay
product_revision 100
ctrl_s/n
allow_quorum yes
fabric_type fc
site_id
site_name
WWPN 524A93744B570E13
path_count 32
max_path_count 32

```

When used in conjunction with the Pure Storage CLI command `pureport list`, which gives the WWPN of each fibre channel port on the array you can, if you wish, make these SVC controller names more user-friendly by renaming them using the `chcontroller` command. For example:

```
# svctask chcontroller -name Pure_CT0_FC0 controller0
```

When volumes are presented to an SVC from a Pure Storage although the volume is mapped from all the Pure Storage ports to the SVC the volumes will only show under one of the Pure Storage controllers. This is perfectly normal. In the event of a hardware failure in the fabric or the Pure Storage array, the internal multi-pathing of the SVC will automatically move the volume to a different controller.

The following screenshot shows 2 different 4 port Pure Storage arrays connected to an SVC with one array presenting a 30GB volume and the other a 20GB volume. Each volume only appears visible to one of the controllers per backend array.

Nombre	Estado	Capacidad	Modo	Sitio	Agrupación d...	Sistema de alm...	LUN
controller0	✓ En línea	NETAPP LUN		Sitio: Sin asignar	WWNN: [REDACTED]		
controller1	✓ En línea	PURE FlashArray		Sitio: Sin asignar	WWNN: 524A937D40A64D03		
mdisk19	✓ En línea	30,00 GiB	Sin gestionar			controller1	000000
controller2	✓ En línea	PURE FlashArray		Sitio: Sin asignar	WWNN: 524A937D40A64D13		
controller3	✓ En línea	PURE FlashArray		Sitio: Sin asignar	WWNN: 524A937D40A64D01		
controller4	✓ En línea	PURE FlashArray		Sitio: Sin asignar	WWNN: 524A937D40A64D11		
controller5	✓ En línea	PURE FlashArray		Sitio: Sin asignar	WWNN: 524A937D595A0001		
mdisk20	✓ En línea	20,00 GiB	Sin gestionar			controller5	000000
controller6	✓ En línea	PURE FlashArray		Sitio: Sin asignar	WWNN: 524A937D595A0011		
controller7	✓ En línea	PURE FlashArray		Sitio: Sin asignar	WWNN: 524A937D595A0003		
controller8	✓ En línea	PURE FlashArray		Sitio: Sin asignar	WWNN: 524A937D595A0013		
CSA1840P01	✓ En línea	IBM FlashSystem-9840	Número de serie: [REDACTED]	Sitio: Sin asignar	WWNN: [REDACTED]		

Integration with IBM Easy Tier

IBM Easy Tier is a function within the SVC that performs automated data placement throughout SVC disk pools to improve the efficiency and performance of the SVC storage system. It attempts to relocate data across up to three tiers automatically and without I/O disruption. Prior to SVC 7.3, there were only two tiers available to Easy Tier, but here we will deal with the three-tier model available from SVC code level 7.3.

Easy Tier attempts to store the ‘hottest’ data on SSDs at an SVC extent level. Pure Storage volumes presented to an SVC create MDisks that are fully supported within the Easy Tier system, but there are two parameters which must be set to ensure that Easy Tier correctly recognises the Pure Storage MDisks as the correct tier:

- Ensure all Pure Storage MDisks have the `easttierload` parameter set to `high`
- Set Pure Storage MDisks to be of type `ssd`

Both of these settings can be achieved in a combined SVC command:

```
> svctask chmdisk -tier ssd -easytierload very_high <mdisk_id>
```

SVC Storage Pool Extent Size

The extent size of a storage pool has no effect on the performance of a storage pool; it is more of a management factor. If there are existing storage pools it is recommended to create a storage pool for the Pure Storage MDisks with the same extent size as the extent size of existing storage pools. This will facilitate easier migration of VDIs between storage pools should that be required.

If there is never going to be a requirement to migrate VDIs between volumes then it is recommended to create the Pure Storage storage pool with the default extent size, which at the time of writing is 1GB.

As per IBM's Best Practices only create one storage pool per backend storage system and place all volumes from this system into the pool. The only exception to this is if Easy Tier is being implemented, in which case you must put the Pure Storage MDisks being used as Easy Tier devices into a storage pool with MDisks of other characteristics (preferably SAS).

Volume Count and Sizing

To ensure that all the resources of the IBM SVC are fully utilised, at least eight volumes should be created per Pure Storage FlashArray. This way all CPU cores, nodes and FC ports are fully utilised.

A further recommendation, should you require more than 8 MDisks from the Pure Storage array, is to make the number of volumes presented to be multiples of eight.

Prior to SVC code level 7.1.0.2 there was no support for volumes presented to the SVC greater than 2TB, however as the minimum supported SVC code level for Pure Storage is greater than this you can create Pure Storage volumes of greater than 2TB and successfully present these to the SVC. In general a 2TB volume size is recommended if you are creating storage pools for permanent usage, however, this may not be appropriate where there your applications may be performing excessive delete cycles (discussed later in this

document in the SCSI UNMAP section). If volumes are to be used for migration purposes then create volumes of the correct size to map to the source VDisks on the SVC.

SVC Cache Considerations

Each SVC node contains a read/write cache that can be enabled or disabled on a per-volume basis. The cache is used to help hide read and write latency experienced by traditional spinning disk drives.

With the introduction of NAND-based flash devices I/O operations now complete with sub-millisecond latencies, rather than the 10 or more milliseconds to/from spinning disk media. This implies that the cache could no longer provide a substantial benefit to the I/O operations from flash devices.

Whilst the cache is DRAM based and therefore fast in nature, it has been observed that disabling the cache on flash backed VDisks can provide better performance than having the cache turned on.

For this reason, we recommend disabling the cache on all Pure Storage backed VDisks. This can be achieved using the following SVC command line command:

```
> svctask chvdisk -cache none <vdisk_name>
```

An added benefit of disabling the cache for these VDisks is that the cache can be better utilised on other backend storage arrays with spinning disk media.

There are conditions where it is advisable to turn the cache on for Pure Storage backed VDisks as the SVC requires the cache to perform internal operations correctly. This would be the case where you are using some of the advanced operations provided by the SVC such as Global and Metro Mirroring.

IBM Real-Time Compression (RtC), Thin Provisioning and Encryption

It is not recommended to implement Real-time Compression, Thin Provisioning or Encryption (available from SVC code level 7.4) on any VDisks created from a Pure Storage storage pool.

The Pure Storage array is already thinly-provisioned and has always on, inline and post-processing compression. Implementing RtC or thin-provisioning on the SVC VDisks will negate the Pure Storage equivalents and will reduce the effectiveness of the array to be able to provide its class beating TCO models. Encrypted data cannot be effectively reduced by the Pure Storage array and the data reducing benefits of the array will be lost. Any data on the Pure Storage array will be data encrypted at rest by default.

Enabling these features will reduce the performance of the SVC in general due to the additional memory and CPU cycles required to implement these features. Thin provisioning at the SVC layer also requires more I/Os due to the requirement of reading and writing to the compression directory. In some circumstances, you may require double the I/Os with thin provisioning enabled at the SVC layer. If RtC is enabled at all in the SVC then a large percentage of SVC controller cache is dedicated to this feature and takes away from the useful read/write cache better utilised by VDisks backed by storage pools created from legacy spinning disk arrays.

SVC and SCSI UNMAP

At the time of writing the SVC does not support the SCSI UNMAP primitive and therefore any backend array that is thinly provisioned will require regular administrative maintenance to ensure that deleted blocks from the host side are correctly recovered on the backend array.

There is an open IBM RFE for the support of SCSI UNMAP which we would recommend customers contact IBM about to get a more current status on this request. The open IBM RFE reference is 26193.

It is recommended that hosts connected to the SVC with VDisks from any thinly provisioned backend array, including Pure Storage, run regular space reclamation tasks. Examples of these tasks for different operating systems or applications would be:

- Windows: sdelete
- UNIX: dd
- VMware: vmkfstools
- Oracle: ASRU

An alternative method for space recovery that will force the Pure Storage array to run space reclamation and garbage collection would be as follows:

- Use 1:1 mappings of VDisk to MDisk
- Create a new VDisk/MDisk from the Pure Storage storage pool
- Use `vdiskmigrate` to move data from old VDisk to new VDisk
- Delete the old Vdisk/MDisk pair on SVC and then on the Pure Storage array

WARNING: If this method is considered then please ensure you fully understand the procedure and that the data migration process has completely correctly before deleting the old VDisk/MDisk pair as this will potentially make your data inaccessible to your host and application.

SVC VDisk Formatting

From the v7.5 release of the SVC code, there has been added a new, default, feature which will, slowly and in the background, format a fully allocated disk by filling it with zeros. Be aware this feature can literally take weeks to complete if your VDisk is very large.

Any newly created VDisk uses free extents in the pool it is created within. These extents may well have been used by previous VDisks that have now been deleted. To ensure that these extents do not contain data which may be read by the host presented with the new vdisk the SVC will perform this formatting process.

The process will only format extents that have not been written by the host that currently VDisk from the host.

From a Pure Storage perspective, this is advantageous as these zeros will be detected by our array and consolidated using our data reduction technology. This in a small way will negate the lack of SCSI UNMAP provided by the SVC, but remember that this space recovery mechanism is extremely slow. A test by IBM showed a 100GB volume took 13 hours to complete this formatting procedure.

It is possible to increase the performance of this background process on a per disk basis by increasing the “Mirror Sync Rate” of the Vdisk using the command (where *rate* is a percentage of the peak synchronisation rate):

```
# svctask chvdisk -syncrate <rate>
```

This parameter is also useful during migration processes as detailed in the next section.

Using SVC for Data Migration to Pure Storage

The IBM SVC can be used to facilitate data migrations from third-party storage arrays, or from SVC VDIsks to Pure Storage but there are a number of things which must be noted if either of these migration methodologies is going to be utilised.

Where the SVC is to be used as a migration appliance the target Pure Storage VDIsks should be created from *image-mode* MDIsks and there should be a one-to-one mapping of VDisk to MDisk.

If you are performing a migration from an SVC to Pure Storage, then you must ensure that you do not migrate either thin and/or compressed SVC volumes directly to the Pure Storage array. The method for performing this action is to create a thick and uncompressed VDisk from a Pure Storage MDisk and use SVC volume mirroring features. This will allow the VDisk to be migrated successfully.

Any data migration performed using the SVC will not only perform slower than expected due to the limitation of the source array volumes but will also, due to the implementation method of the SVC cache, perform large block writes for the duration of the migration. These large block writes are not optimal for Pure Storage arrays, although they are handled better by the //m series over the FA400 series, and therefore this will affect the write performance of the Pure Storage array.

If using volume mirroring to perform a migration task be aware that by default, the SVC software sets the default background copy rate, referred to as the `sync_rate` to a value of 50. This value can be set dynamically on a scale between 0 to 100.

The `sync_rate` limits the copy speed as detailed in the following table. These values are independent of the grain size of a Vdisk, with the exception of v5.1.

To ensure that your volume copies complete in the minimum time it is recommended to increase the `sync_rate` to 100 per Vdisk using the following command:

```
# svctask chvdisk -syncrate 100
```

The `syncrate` parameter limits the data copy speeds based on the following table, but it should be noted that is write IO are still being performed to the during the mirroring process, these can be impacted by a larger `syncrate` value.

syncrate value	Data copy rate maximum /sec
1-10	128 KB
11-20	256 KB
21-30	512KB
31-40	1 MB
41-50	2 MB
51-60	4 MB
61-70	8 MB
71-80	16 MB
81-90	32 MB
91-100	64 MB

As soon as volume mirroring is in sync or volume migration tasks have completed the target VDisk will perform as would be expected from an All Flash Array.

About the Author



As Global Solutions Architect, Simon specialises in Data Migration methodologies and has experience with a wide range of virtualization technologies. These virtualization technologies can be used to assist in data migration to Pure Storage or as a permanent virtualization layer supporting Pure Storage as a high-performance backend.

With over 25 years of storage experience across all aspects of the discipline, from administration to architectural design, Simon has worked with all major storage vendors' technologies and organisations, large and small, across Europe and the USA as both customer and service provider. He also specialises in OpenStack methodologies assisting customers with integrating Pure Storage into their private and public clouds.

Blog: <http://www.purestorage.com/blog/author/simon>



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