

Best practices for virtualizing IBM System Storage DS8000 using IBM FlashSystem V9000

Overview

Challenge

How to virtualize an existing IBM DS8000 storage behind IBM FlashSystem V9000?


Solution

The paper talks about a checklist and best practices that need to be taken care of before performing virtualization and to minimize the impact on performance.

This paper focuses on smarter ways of virtualizing existing IBM® System Storage®DS8000® storage behind IBM FlashSystem®V9000. This can help customers to optimize their storage based on performance and get virtualization benefits of IBM FlashSystem V9000. It covers the prerequisites and the best practices involved for virtualizing IBM DS8000 storage behind IBM FlashSystem V9000.

Target audience

This technical paper is developed to assist IBM Business Partners, field sales representatives, technical specialists, and IBM clients to understand the advantages of virtualizing DS8000 behind FlashSystem V9000 and the best practices involved, including the *state of the art* features such as IBM Real-time Compression™, IBM Easy Tier®, and IBM System Storage SAN Volume Controller (SVC) cache partitioning.



Features

DS8000:

- Enterprise storage
- True four-way and multisite data replication

FlashSystem V9000

- Flash drives
 - Point-in-time flash copies
 - Local and remote mirroring
 - Real-time Compression
 - Easy Tier support (automatically migrates hot extents to the fastest storage and also migrates the fastest storage to hot extents)
-

Problem statement

Consider a customer environment containing IBM FlashSystem V9000 and IBM DS8000 working in isolation.

Probable limitations of the design:

- Flash storage might go out of space.
- Available storage of DS8000 can be used but it is not accessible to hosts connected to FlashSystem. Data on FlashSystem drives is not accessed frequently.

Alternatively, DS8000 has a couple of ranks available but is working in isolation from FlashSystem V9000.

Lab setup

This section describes the setup that was used for collecting data for this paper:

- IBM FlashSystem V9000 includes two IBM SVC nodes connected with FlashSystem (with total space of approximately 22 TB) through 16 Gb FC switches.
- IBM DS8000 consists of approximately 48 TB free storage.
- Both storage arrays are serving I/O from different hosts and are independent of each other.




Figure 1: Separate host I/O operations are served from FlashSystem V9000 and DS8000

Checklist before virtualizing IBM DS8000 behind IBM FlashSystem V9000

Before virtualizing DS8000, the following factors need to be considered:

Cache utilization

If cache utilization is too high on FlashSystem (say 85% to 90%), further division of pools is not recommended as cache will be divided among all pools. So, with high cache utilization, keep lesser storage pools.

```
svcinfolssystemstats | grep cache
```

Processor utilization

This is same as cache and workload. If the FlashSystem is utilizing the CPU upto 80% to 90%, then adding more pools will be an over burden for CPU.

```
svcinfolssystemstats | grep cpu
```

These are instantaneous values and the team should observe it for a period of time.

Note: If the CPU and cache utilization is high, virtualizing DS8000 behind FlashSystem V9000 is not a good option.

Heat maps verification

FlashSystem V9000 has Easy Tier functionality that helps creating heat maps. This heat map should be analyzed on FlashSystem using the Storage Tier Advisor Tool (STAT) to identify the data that is cold enough to be moved to a slower tier than flash drives. Refer ibm.com/support/docview.wss?uid=ssg1S4000935 for more information.

Comprestimator

If post virtualization, users want to use compression to compress volumes currently present on the flash drives, then the Comprestimator utility should be used to know how much compression is feasible, based on the data present on the virtual disks on FlashSystem.

Refer to the Comprestimator utility for more details.

ibm.com/webapp/set2/sas/f/comprestimator/home.html

In the lab setup used for this paper, the following results were observed:

Checkpoint	Observed values on FlashSystem V9000
Cache utilization	50%
CPU utilization	40%
Heat map	Three volumes have cold data which can be moved to DS8000
Comprestimator	70% compression feasibility on SQL volumes. Hence can be compressed. 30% compression feasibility on other volumes

Table 1: Checkpoints on FlashSystem V9000 before virtualizing IBM DS8000 behind it

Proposed solution

Objectives of the proposed solution:

- Low-level design for IBM DS8000 array, rank and extent pool (only for newly upgraded capacity being virtualized behind FlashSystem V9000)
 - Planning and storage low-level design for reorganizing of existing data on FlashSystem V9000
 - Storage-based data migration
-

The ability to virtualize storage arrays makes FlashSystem V9000 a good candidate for connecting DS8000 behind it and utilize the available storage space of DS8000.

The solution focuses on getting a good utilization of available storage by connecting IBM DS8000 behind IBM FlashSystem V9000. After the additional storage from DS8000 is available to FlashSystem V9000, all advanced features of IBM Spectrum Virtualize™ can be used, for example, Easy Tier, Real-time Compression, FlashCopy, and so on.




Figure 2: DS8000 virtualized behind FlashSystem V9000

Storage pool division

- PoolA will have uncompressed volumes.
 - PoolB will have compressed volumes.
 - PoolC can be created if there is a requirement of FlashCopy targets for storing the snapshots.
-

After the criteria in the specified checklist is met, the following solution is proposed:

1. Create another pool (PoolB) with few managed disks (MDisks) coming from DS8000.
2. Add a few MDisks coming from DS8000 to the existing pool, PoolA, to have it converted in to a multi-tier pool (PoolA initially had only Flash Misks coming from FlashSystem V9000).
3. Move few Flash MDisks from PoolA to PoolB.

4. Move all compressed volumes from PoolA to PoolB and limit the creation of compressed volumes only in the new pool (PoolB in this case)
5. Enable Easy Tier on both the pools, if not already.

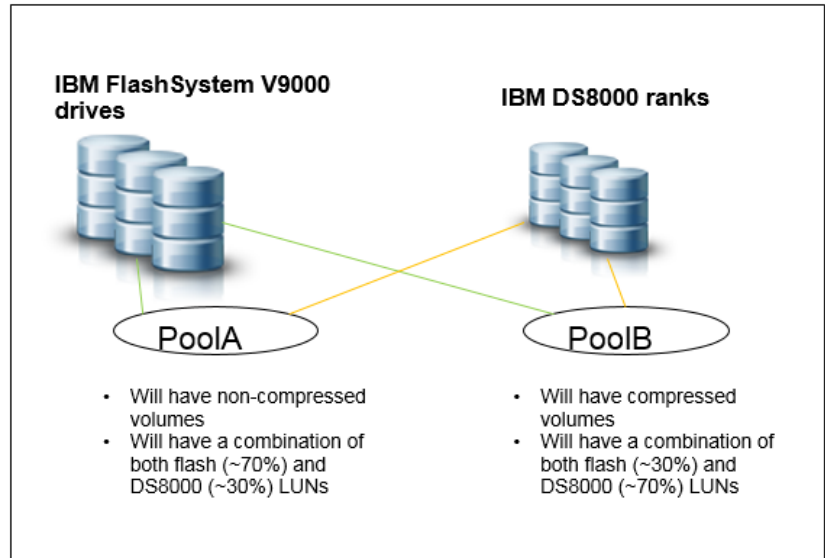


Figure 3: Pools on FlashSystem V9000

Steps to achieve the recommendation:

For virtualizing DS8000 behind FlashSystem V9000, the first step is to connect them together. Connect DS8000 to the same fabric and zone them in such a way that freely available LUNs can be mapped to FlashSystem V9000.

Perform the following steps to carry out I/O operations on DS8000 drives from FlashSystem V9000:

1. After both arrays are zoned together using a Fibre Channel(FC), create a FlashSystem V9000 host on DS8000 and map all volumes created out of available capacity (ranks) that is being virtualized behind FlashSystem V9000
2. Detect all volumes on FlashSystem V9000 using the following command.
`Svctask detectmdisk`

Note: DS8000 volumes will get detected as (MDisks) on FlashSystem V9000.

3. For lab setup, there was only one MDisk group with MDisks coming from FlashSystem V9000 before virtualizing the DS8000 system. Add specific set of MDisks of DS8000 to the existing MDisk group.

```
svctask addmdisk -  
mdisk<mdisk_id_list><mdisk_group_number or  
name>
```

For example:

```
IBM_2145:flashsystemcluster:superuser>svctask addmdisk -mdisk mdisk10:mdisk11:mdisk12 poolA  
IBM_2145:flashsystemcluster:superuser>
```

Figure 4: Adding DS8000 MDisk to an existing pool

4. Remove a few flash MDisk from the existing MDisk group, which will later be used with another MDisk group for compressed volumes.

```
svctask rmdisk -  
mdisk<mdisk_id_list><mdisk_group_number or  
name>
```

For example:

```
IBM_2145:flashsystemcluster:superuser>svctask rmdisk -mdisk mdisk5:mdisk6:mdisk7 poolA  
IBM_2145:flashsystemcluster:superuser>
```

Figure 5: Removing Flash MDisk from the existing pool

5. Create another MDisk group or pool for compressed volumes with MDisk freed up in the previous step and rest of the MDisk from DS8000.

```
svctask mkmdiskgrp -mdisk<mdisk_id_list> -  
ext<extent size>
```

For example:

```
IBM_2145:flashsystemcluster:superuser>mkmdiskgrp -mdisk mdisk5:mdisk6:mdisk7 -ext 1024 -name poolB  
MDisk Group, id [1], successfully created  
IBM_2145:flashsystemcluster:superuser>
```

Figure 6: Create a new pool using freed up MDisk of FlashSystem

Now, new volumes can be created in both the pools.

Validation

The following table shows the results obtained after virtualizing DS8000 behind FlashSystem.

Checkpoints	Before virtualization	After virtualization
Peak cache usage	50%	60%
Peak CPU usage	40%	55%

Table 2: Checkpoints on FlashSystem V9000 after DS8000 is virtualized behind it

Latency and bandwidth for volumes having hot data remained same post virtualization.

Generic best practices

On DS8000:

- Configure a minimum of eight ports per DS8000 system.
- Configure 16 ports per DS8000 system when more than 48 ranks are presented to the SVC cluster.
- Configure adapters across redundant storage area networks (SANs) from different I/O enclosures.
- Configure a maximum of two ports per DS8000 host adapter.
- Create a minimum of two large extent pools, one even and one odd.
- Balance the ranks on a device adapter (DA) pair evenly between odd and even extent pools.
- Configure four to eight ranks per extent pool.

On SVC:

- Use four ports minimum, two from each SVC node of an I/O group (preferably dedicated for DS8000). Refer to the [Port designations section in IBM Redbooks](#) at the following URL to isolate the traffic.
- Configure the SVC storage pool with volumes coming evenly from both servers of DS8000.
- Use STAT before planning a hybrid pool.
- If planning for compression, create separate pools for compressed and non-compressed volumes.

Resources

The following websites provide useful references to supplement the information contained in this paper:

- IBM SVC Knowledge Center
ibm.com/support/knowledgecenter/STPVGU/landing/SVC_welcome.html
- IBM FlashSystem V9000
ibm.com/redbooks/redbooks/pdfs/sg248273.pdf
ibm.com/support/knowledgecenter/en/STKMQV
- IBM DS8000
ibm.com/support/knowledgecenter/ST8NCA/product_welcome/ds8000_kcwelcome.html
- IBM Redbooks on Data Migration to IBM Disk Storage Systems
ibm.com/redbooks/redbooks/pdfs/sg247432.pdf
- IBM FlashSystem V9000 Easy Tier
ibm.com/support/knowledgecenter/en/STKMQV_7.7.1/com.ibm.storage.vflashsystem9000.7.7.1.doc/svc_easy_tier.html
- IBM Comprestimator
ibm.com/webapp/set2/sas/f/comprestimator/home.html

About the authors

Aakanksha Mathur is a software engineer in the SVC team at IBM System Labs, India. You can reach Aakanksha at: aamathur@in.ibm.com.

Akshat Mithal is a development lead for the Spectrum Virtualize (SVC and IBM Storwize®) family of products. You can reach Akshat at: akmithal@in.ibm.com.

Varun Kumar is a storage consultant in IBM Systems Lab Services team in India. He works for various SAN and storage technologies. You can reach Varun at: varun.kumar@in.ibm.com.



©Copyright IBM Corporation 2017
IBM Systems
3039 Cornwallis Road
RTP, NC 27709

Produced in the United States of America

IBM, the IBM logo and ibm.com are trademarks or registered trademarks of the Internal Business Machines Corporation in the United States, other countries, or both. If these and other IBM trademarked items are marked on their first occurrence in the information with a trademark symbol (® or ™), these symbols indicate U.S. registered or common law trademarks owned by IBM at the time this information was published. Such trademarks may also be registered or common law trademarks in other countries. A current list of IBM trademarks is available on the web at “Copyright and trademark information” at ibm.com/legal/copytrade.shtml

Other product, company or service names may be trademarks or service marks of others.

References in the publication to IBM products or services do not imply that IBM intends to make them available in all countries in the IBM operates.

Please recycle

