

Using VMware ESX Server with IBM System Storage SAN Volume Controller

ESX Server 3.0.2

This technical note discusses using ESX Server hosts with an IBM System Storage SAN Volume Controller (SVC). An SVC storage system uses advanced storage virtualization technology. The SVC system uses two or more nodes in clustered mode to support failover. See the *SAN Compatibility Guide* for precise information on the supported devices. If you want to use an SVC with ESX Server hosts, you must perform different setup tasks than if you want to set up a non-virtualized SAN storage array.

The technical note starts with an introduction to IBM SVC technology, and then steps you through setting up your storage array and your ESX Server hosts. The last section discusses how you can resolve the issue that arises because the LUN identifiers change when the LUNs are virtualized.

Understanding SVC Storage Virtualization

An ESX Server host can be connected to virtualized LUNs, which are storage arrays that export LUNs from back-end internal RAID devices, as discussed in detail in the *SAN Configuration Guide*.

Virtualized LUNs hide the actual LUN allocation details from the ESX Server host. Storage virtualization provides added flexibility because it brings disparate storage arrays under single management control and allows data migration across tiered storage arrays.

IBM SVC is an appliance-based device that can connect external storage arrays as back-end devices and virtualize the storage from these back-end devices. An IBM SVC device virtualizes storage and offers high availability with node clustering. During SVC setup, administrators set up two SVC nodes as a cluster to avoid a single point of failure. For more information, see the IBM System Storage SAN Volume Controller documentation, "[Configuration Example](#)" on page 2, and "[System Setup](#)" on page 3.

During SVC setup, administrators set up pairs of nodes to form a cluster and use zoning to create two different zones on each switch. One zone allows certain SVD ports to talk to certain storage port, the other zone allows the ESX Server host to see certain SVD ports. This setup allows the ESX Server host to continue to have access to the data on the backend array in case of path failure because the SVD can use a path that is not directly visible to the ESX Server host.

Configuration Example

Figure 1 shows an example of a typical configuration.

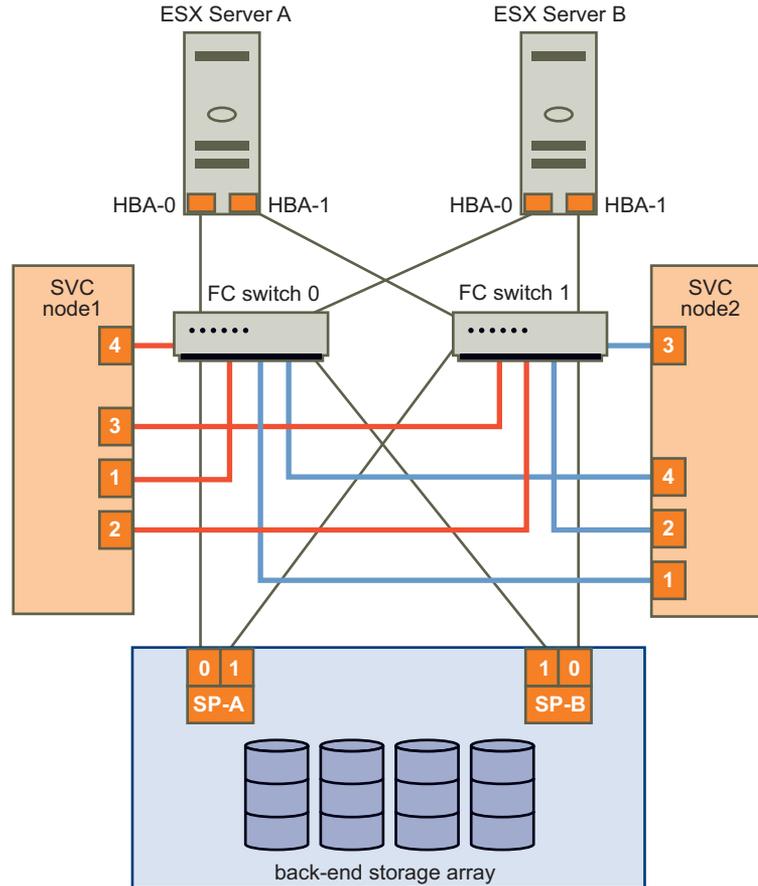


Figure 1. Using a Clustered Storage Virtualization Device

The diagram shows the following elements from top to bottom:

- Two ESX Server hosts (ESX Server A and ESX Server B) are connected through two FC switches to an SVC node pair.
- The SVC consists of two clustered nodes. Each node has four ports. On each node, two of the ports connect to one switch, the other two connect to the other switch.
- Each switch also connects to the storage array on two different ports to allow path failover.

The rest of this document gives an overview of the storage array setup and explains how to set up your ESX Server host to use the virtualized LUNs.

NOTE Consult the IBM documentation (<http://www.ibm.com/storage/support/2145/>) for detailed information on the storage array setup.

Terminology

The following terminology is used in this document (see Figure 1 for an illustration).

Storage Virtualization Device (SVD) – A storage controller, appliance, or network-based device that can export LUNs from external storage arrays. The external storage arrays are connected to the SVD as back-end devices.

IBM SAN Volume Controller (SVC) – A storage virtualization device that is deployed as one or more pairs of nodes. Each pair of nodes is called an IO Group. Write data are cached and mirrored across the node pair.

Virtualized SAN LUN — A LUN that is exported by an SVD. Such a LUN is allocated on the back-end storage device but accessible through the SVD. During the virtualization process, the SVD presents the back-end storage LUNs to the hosts using a new LUN ID (UUID or serial number).

System Setup

System setup consists of these tasks:

- [“Task 1: Setting Up SVC and Zoning”](#) on page 3
- [“Task 2: Allocating New Storage”](#) on page 3
- [“Task 3: Making Storage Available to the ESX Server Host”](#) on page 3

Task 1: Setting Up SVC and Zoning

The first step is to set up the SVC cluster. The SVC Console is a separate machine, on which the Windows operating system and the SVC management tool are available.

NOTE It is essential that you set up the SVC and the four zones correctly. See the IBM SVC documentation for details.

This task consists of the following steps.

Step	Description	Performed by...	Documentation
1	Set up the SVC cluster.	Using IBM SVC Console.	IBM SVC Documentation.
2	Set up fibre channel connections between <ul style="list-style-type: none"> ■ ESX Server hosts and Switch ■ Switch and SVC nodes ■ Switch and backend storage array. 	Manually providing FC connections.	No documentation.
3	Set up zoning between: <ul style="list-style-type: none"> ■ ESX server initiators and SVC node ports ■ SVC node ports and back-end array target ports. <p>NOTE: ESX Server host initiators should not be in the same zone as backend target array ports.</p>	Logging into the switch.	Switch management software documentation.

Task 2: Allocating New Storage

This section discusses allocating new virtualized storage from the SVC. For information on migrating back-end storage with VMFS volume or RDM to virtualized storage, see [“Making VMFS Volumes or RDMs Visible After LUN ID Changes”](#) on page 4. Allocating new storage consists of the following steps:

Step	Description	Performed from...	Documentation
1	Create LUNs in the back-end storage array and map each LUN to SVC ports.	Backend storage array's management software.	Backend storage array documentation.
2	Instruct SVC cluster to discover all LUNs it has been assigned by the backend-storage array.	IBM SVC console	IBM SVC documentation.
3	Create virtual disks and map virtual disks to the ESX Server hosts.	IBM SVC console	IBM SVC documentation.

Task 3: Making Storage Available to the ESX Server Host

After you've completed the tasks above, you can make the storage available to the virtual machines running on your ESX Server host. The process is the same as for any other FC storage, and is discussed in the *ESX Server 3 Configuration Guide*, and the *SAN Configuration Guide*.

Making VMFS Volumes or RDMs Visible After LUN ID Changes

VMFS-3 metadata identify storage volumes using several properties, including the LUN number and the LUN ID (UUID or Serial Number). Virtualized LUNs have different UUIDs than the corresponding LUNs on the back-end array. Because of the resulting metadata mismatch, VMware Logical Volume Manager (LVM) identifies the SVD volumes as snapshots. This section discusses how to make an ESX Server host recognize a virtualized LUN containing either RDM storage or a VMFS volume.

Accessing VMFS Volumes on Virtualized LUNs Without Volume Resignaturing

This section discusses how you can make VMFS volumes accessible without volume resignaturing. In this approach, you turn off the `LVM.DisallowSnapshotLun` option. This setting allows your ESX Server host to mount VMFS volumes even though the LUN ID has changed. If you expect to use array-based LUN snapshots in the same configuration, you should use the volume resignaturing approach instead. See [“Accessing VMFS Volumes on Virtualized LUNs with Volume Resignaturing.”](#)



CAUTION When `LVM.DisallowSnapshotLun` is set to 0, no snapshot LUNs should be presented to the ESX Server host or data corruption can result. For details, see the section on volume resignaturing in the *SAN Configuration Guide*.

To Access VMFS Volumes on Virtualized LUNs Without Volume Resignaturing

- 1 Power off all virtual machines.
- 2 Present the virtualized LUN to the ESX Server hosts and perform two rescan operations on each host.
The virtualized LUN is detected as a snapshot.
- 3 From a VI Client connected to a VirtualCenter Server, select an ESX Server host in the Inventory panel and click the Configuration tab.
- 4 Click **Advanced Settings**.
- 5 In the `LVM.DisallowSnapshotLun` field, replace the default (1) with 0 and click **OK**.
- 6 Click **Rescan**, and then click **Rescan** again.
- 7 Make sure the VMFS volume is mounted successfully.
- 8 Repeat the steps for each ESX Server host that shares the same virtualized LUN.

Accessing VMFS Volumes on Virtualized LUNs with Volume Resignaturing

The second approach for presenting VMFS volumes on virtualized LUNs involves volume resignaturing.

To Access VMFS Volumes on Virtualized LUNs With Volume Resignaturing

- 1 Power off all virtual machines.
- 2 Present the virtualized LUN to only one ESX Server host and perform two rescan operations.
The virtualized LUN is detected as a snapshot.
- 3 In a VI Client connected to a VirtualCenter Server, select the ESX Server host and click the Configuration tab.
- 4 Click **Advanced Settings**.
- 5 In the `LVM.EnableResignature` field, replace the default (0) with 1 and click **OK**.
- 6 Click **Rescan**, and then click **Rescan** again.
- 7 Make sure the VMFS volume is mounted successfully.
- 8 Change `LVM.EnableResignature` back to 0 and click **OK**.
- 9 Reregister all existing virtual machines from the virtualized LUN on the ESX Server host.

- 10 Present the virtualized LUN to the next ESX Server hosts and repeat step 8 until you have reregistered all virtual machines.

Accessing RDM Storage on Virtualized LUNs

If you use RDM (raw device mapping) storage on the back-end storage array, then virtualize the corresponding LUN, the result is a stale RDM pointer, which you have to remove.

To Access RDM Storage on Virtualized LUNs

- 1 Power off all virtual machines.
- 2 Present the virtualized RDM LUN to the ESX Server host and perform two rescans.
- 3 Locate the virtual machine that has the RDM mapped to the back-end LUN (which is now virtualize) and power off the virtual machine.
- 4 Remove the stale RDM pointer.
- 5 Recreate the RDM to point to the virtualized LUN.
- 6 Repeat steps 1-5 for each RDM that was previously mapped to back-end storage.

Caveats

When working with SVDs, consider the following caveats:

- VI Client Might Incorrectly Display RDM LUNs in Physical Compatibility Mode as Available Disk when Storage Virtualization Devices Are Used. See <http://kb.vmware.com/kb/1002564>.