

IBM System Storage SAN Volume Controller



Software Installation and Configuration Guide – Errata

Version 5.1.0
February 11, 2010

SC23-6628-05 Errata

Contents

Who should use this guide.....	3
Last Update.....	3
Change History.....	3
Chapter 3. SAN fabric and LAN overview.....	4
iSCSI configuration rules.....	4
SSD configuration rules for nodes, I/O groups and clusters.....	4
SSD configuration rules for VDisks.....	4
Chapter 5. Using the SAN Volume Controller Console.....	5
Configuring the iSNS server address.....	5
Configuring cluster iSCSI authentication.....	5
Adding a mirrored copy to a VDisk.....	5
Chapter 6. Using the CLI.....	6
Creating VDisks using the CLI.....	6
Adding a copy to a VDisk using the CLI.....	7
Collecting SSD dump files using the CLI.....	7
Chapter 7. Backing up and restoring the cluster configuration.....	8
Restoring the cluster configuration using the CLI.....	8
Chapter 8. Upgrading the SAN Volume Controller software.....	11
Upgrading clusters with internal SSD drives.....	11
Upgrading solid-state drive (SSD) software.....	12
Chapter 11. Configuring and servicing storage systems.....	12
Configuring EMC Symmetrix and Symmetrix DMX systems.....	12
Global settings for IBM System Storage DS5000 or IBM DS4000 systems.....	12

Who should use this guide

This errata should be used by anyone using the *IBM System Storage SAN Volume Controller Software Installation and Configuration Guide*. You should review the errata contained within this guide and note the details with respect to the copy of the *Software Installation and Configuration Guide* supplied with your SAN Volume Controller.

Last Update

This document was last updated: Feb 11, 2010

Change History

The following revisions have been made to this document:

Revision Date	Sections Modified
November 3, 2009	New publication
February 11, 2010	SSD configuration rules for nodes, I/O groups and clusters Global settings for IBM System Storage DS5000 or IBM DS4000 systems

Table 1: Change History

Chapter 3. SAN fabric and LAN overview

The following corrections should be noted.

iSCSI configuration rules

Page 91. The following paragraph is incorrect and should be removed.

Each I/O group can map VDisks to the same total maximum number of host objects (256), which could include fibre-channel attachments, iSCSI attachments, or both.

The following information replaces the above paragraph.

See the following Web site for the latest maximum configuration support:
www.ibm.com/storage/support/2145

SSD configuration rules for nodes, I/O groups and clusters

Page 95. The following bullet points are incorrect and should be removed.

- Do not combine nodes that contain SSDs and nodes that do not contain SSDs in a single I/O group. However, while upgrading an earlier SAN Volume Controller node to a SAN Volume Controller 2145-CF8 node, you can temporarily combine the two node types in a single I/O group.
- Nodes in the same I/O group must share the same number of SSDs.

SSD configuration rules for VDisks

Page 95/96. The following additional information is provided.

The synchronization rate must be set such that the VDisk copies will resynchronize quickly after loss of synchronization. Synchronization will be lost if one of the nodes goes offline either during a concurrent code upgrade or because of maintenance. During code upgrade the synchronization must be restored within 30 minutes or the upgrade will stall. Unlike VDisk copies from external storage subsystems, during the period that the SSD VDisk copies are not synchronized access to the VDisk depends on the single node containing the SSD storage associated with the synchronized VDisk copy. The default synchronization rate is typically too low for SSD VDisk mirrors; instead it should be set to 80 or above.

Note: To increase the amount of time between the two nodes containing VDisk copies going offline during the normal upgrade process, consider using the User-paced Software upgrade procedure.

Chapter 5. Using the SAN Volume Controller Console

The following corrections should be noted.

Configuring the iSNS server address

Page 149. The following information is misleading and should be removed.

After you configure the iSNS server address for the cluster, you can configure cluster iSCSI authentication.

Note: To help in problem determination, this step can be delayed until after the first one or two hosts have been configured and their connectivity has been tested without authentication configured.

Configuring cluster iSCSI authentication

Page 149. The following information replaces this section in the original document.

You can use the SAN Volume Controller Console to configure the Challenge-Handshake Authentication Protocol (CHAP) to authenticate the SAN Volume Controller cluster to the iSCSI-attached hosts. After the CHAP secret is set for the cluster, any attached hosts which have target authentication enabled must use this CHAP secret to authenticate.

This task assumes that you have already launched the SAN Volume Controller Console. To configure authentication between the SAN Volume Controller cluster to the iSCSI-attached hosts, follow these steps:

1. In the portfolio, click **Configure iSCSI Authentication**. The Configure Cluster iSCSI Authentication panel is displayed.
2. On the Configure Cluster iSCSI Authentication panel, enter the shared passphrase that is used to authenticate the SAN Volume Controller to the host in the CHAP Authentication Secret field.
3. Click **OK**.

After you configure the CHAP secret for the SAN Volume Controller cluster, ensure that the cluster CHAP secret is added to each iSCSI-attached host which has target authentication enabled.

Adding a mirrored copy to a VDisk

Page 164. The following additional information is provided.

The rate at which the VDisk copies will resynchronise after loss of synchronization can be specified. The following table defines the rates:

Synchronization rate	Data copied/sec
1-10	128KB
11-20	256KB

21-30	512KB
31-40	1MB
41-50	2MB
51-60	4MB
61-70	8MB
71-80	16MB
81-90	32MB
91-100	64MB

The default setting is 50. The synchronization rate must be set such that the VDisk copies will resynchronize quickly after loss of synchronization.

In the case of a mirrored VDisk that uses disk extents on a solid-state drive (SSD) that is located on a SAN Volume Controller node, synchronization will be lost if one of the nodes goes offline either during a concurrent code upgrade or because of maintenance. During code upgrade the synchronization must be restored within 30 minutes or the upgrade will stall. Unlike VDisk copies from external storage subsystems, during the period that the SSD VDisk copies are not synchronized access to the VDisk depends on the single node containing the SSD storage associated with the synchronized VDisk copy. The default synchronization rate is typically too low for SSD VDisk mirrors; instead it should be set to 80 or above.

Note: To increase the amount of time between the two nodes containing VDisk copies going offline during the normal upgrade process, consider using the User-paced Software upgrade procedure.

Chapter 6. Using the CLI

The following corrections should be noted.

Creating VDIs using the CLI

Page 242. The following additional information is provided for point 5.

The rate at which the VDisk copies will resynchronise after loss of synchronization can be specified using the **syncrate** parameter. The following table defines the rates:

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

The default setting is 50. The synchronization rate must be set such that the VDisk copies will resynchronize quickly after loss of synchronization.

In the case of a mirrored VDisk that uses disk extents on a solid-state drive (SSD) that is located on a SAN Volume Controller node, synchronization will be lost if one of the nodes goes offline either during a concurrent code upgrade or because of maintenance. During code upgrade the synchronization must be restored within 30 minutes or the upgrade will stall. Unlike VDISK copies from external storage subsystems, during the period that the SSD VDisk copies are not synchronized access to the VDisk depends on the single node containing the SSD storage associated with the synchronized VDisk copy. The default synchronization rate is typically too low for SSD VDisk mirrors; instead it should be set to 80 or above.

The following is an example of the CLI command that you can issue to create a VDisk with two copies using the I/O group and MDisk group name and specifying the synchronization rate:

```
svctask mkvdisk -iogrp io_grp1 -mdiskgrp grpa:grpb -size 500 -vtype striped -copies 2 -syncrate 90
```

where *io_grp1* is the name of the I/O group that you want the VDisk to use, *grpa* is the name of the MDisk group for the primary copy of the VDisk and *grpb* is the name of the MDisk group for the second copy of the VDisk, and 2 is the number of VDisk copies and the synchronization rate is 90 which is equivalent to 32MB per second.

Adding a copy to a VDisk using the CLI

Page 243. The following additional information is provided

The rate at which the VDisk copies will resynchronise after loss of synchronization can be specified using the **syncrate** parameter. See the Creating VDIsks topic for a description of this parameter.

Collecting SSD dump files using the CLI

Page 238. The following information replaces this section in the original document.

You can use the command-line interface (CLI) to collect dump files from solid-state drives (SSDs). To collect internal log files from solid-state drive (SSD) MDisks, run the `triggerdiskdump` command. Subsequently, you can list, delete or copy the dump files.

The `triggerdiskdump` command generates a dump file and saves it in the `/dumps/mdisk` directory on the node that contains the SSD.

1. Issue the **svctask triggerdiskdump** CLI command.

The following example shows the CLI format for generating a dump file for the specified SSD MDisk:

```
svctask triggerdiskdump mdisk_id | mdisk_name
```

2. Issue the **svcinfolsmdiskdumps** command to list files in the /dumps/mdisk directory on the specified node.

The following example shows the CLI format for listing the dump files for the specified node:

```
svcinfolsmdiskdumps node_id | node_name
```

3. Issue the **svctask cleardumps** command to delete all files from the /dumps directory and all subdirectories on the specified node. To delete files from a subdirectory of /dumps only, specify the -prefix parameter.

The following example shows the CLI format for deleting all dump files from the specified node:

```
svctask cleardumps node_id | node_name
```

The following example shows the CLI format for deleting only the dump files in the specified /elogs/ directory:

```
svctask cleardumps -prefix "/dumps/elogs/*"
```

4. Issue the **svctask cpdumps** command to copy dump files to the configuration node. If the /dumps directory on the configuration node becomes full before the copy completes, no message is returned. To avoid this scenario, clear the /dumps directory after migrating data from the configuration node.

The following example shows the CLI format for copying all dump files from the specified node to the configuration node:

```
svctask cpdumps -prefix /dumps node_id | node_name
```

Chapter 7. Backing up and restoring the cluster configuration

The following corrections should be noted.

Restoring the cluster configuration using the CLI

Page 299. The following information replaces this section in the original document. Steps 6 and 7 have been modified. Step 13 has been added.

You can restore your cluster configuration data using the command-line interface (CLI).

Important: There are two phases during the restore process: prepare and execute. You must not make any changes to the fabric or cluster between these two phases.

Complete the following steps to restore your cluster configuration data:

1. Select delete cluster from the front panel on each node in the cluster that does *not* display Cluster : on the front panel. If the front panel of the node displays Cluster :, the node is already a candidate node.
2. Create a new cluster from the front panel of any node in the cluster. If possible, use the node that was originally the configuration node for the cluster.
3. Generate an SSH key pair for all of the hosts to use to access the CLI.

4. Start the SAN Volume Controller Console.
5. On the Viewing Clusters panel, select the cluster that you are recovering from the list, select **Remove the Cluster** from the task list and click **Go**. The Remove Cluster panel displays. Click **Yes** to confirm the removal of the cluster. The Viewing Cluster panel displays.
6. On the Viewing Cluster panel, select **Add a Cluster** from the task list and click **Go**. The Adding a Cluster panel displays. From this panel, you need to initialise the new cluster by completing these steps:
 - a. Enter the IP address for the cluster that you are recovering. Select **Create (Initialize) Cluster**. Click **OK**.
 - b. The Sign on to Cluster panel appears. On this panel, enter superuser and the initial password when the cluster had been created in step 2.
 - c. Configure the new cluster with required settings as described in Chapter 5.
7. To work with the command-line interface to finish restoring the cluster, you also need to assign an SSH key to the user that has Security Administrator role on the cluster by completing these steps:
 - a. On the Viewing Cluster panel, select the new cluster and select Launch SAN Volume Controller Console from the task list and click **Go**.
 - b. Click **Manage Authentication** → **Users** in the portfolio. The Modifying User superuser panel is displayed.
 - c. To optionally modify the password, enter a new password in the **New Password** field. In the **Re-Enter New Password** field, re-type the new password.
 - d. To assign the SSH key that you generated in Step 3 to the user, enter the name of SSH key file in the SSH Key Public File field or click **Browse** to select the file.
 - e. Click **OK**.
8. Using the command-line interface, issue the following command to log onto the cluster:


```
ssh -l admin your_cluster_name -p 22
```

 Where *your_cluster_name* is the name of the cluster for which you want to restore the cluster configuration.

Note: Because the RSA host key has changed, a warning message displays when connecting to the cluster using SSH.
9. Issue the following CLI command to ensure that only the configuration node is online:


```
svcinfolsnode
```

 The following is an example of the output that is displayed:


```
id name status IO_group_id IO_group_name config_node
1 node1 online 0 io_grp0 yes
```
10. Verify that the most recent version of your **/tmp/svc.config.backup.xml** configuration file has been copied to your SSPC. The most recent file is located on your configuration node in the **/tmp** or **/dumps** directory. In addition, a **/dumps/svc.config.cron.xml_node_name** configuration file is created daily on the configuration node. In certain cases, you might prefer to copy an earlier configuration file. If necessary, back up your configuration file as described in “Backing up the cluster configuration using the CLI” on page 297.

11. Issue the following CLI command to remove all of the existing backup and restore cluster configuration files that are located on your configuration node in the /**tmp** directory:

```
svcconfig clear -all
```

12. Copy the svc.config.backup.xml file from the IBM System Storage Productivity Center or master console to the /tmp directory of the cluster using the PuTTY pscp program. Perform the following steps to use the PuTTY pscp program to copy the file:

a. Open a command prompt from the IBM System Storage Productivity Center or master console.

b. Set the path in the command line to use pscp with the following format:

```
set PATH=C:\path\to\putty\directory;%PATH%
```

c. Issue the following command to specify the location of your private SSH key for authentication:

```
pscp <private key location> source [source] [user@]  
host:target
```

13. If the cluster contains any SAN Volume Controller 2145-CF8 nodes with internal Solid State Disks, then these nodes must be added to the cluster now. In order to do this, determine the panelname, name and I/O groups of any such nodes from the configuration backup file. To add the nodes to the cluster, issue this command:

```
source svctask addnode -panelname <panelname> -iogrp  
<iogrp name/id> -name <node name>
```

14. Issuing the following CLI command to compare the current cluster configuration with the backup configuration data file:

```
svcconfig restore -prepare
```

This CLI command creates a log file in the /tmp directory of the configuration node. The name of the log file is svc.config.restore.prepare.log.

Note: It can take up to a minute for each 256-MDisk batch to be discovered. If you receive error message CMMVC6119E for an MDisk after you enter this command, all the managed disks (MDisks) might not have been discovered yet. Allow a suitable time to elapse and try the svcconfig restore -prepare command again.

15. Issue the following command to copy the log file to another server that is accessible to the cluster:

```
pscp -i <private key location> [user@]host:source target
```

16. Open the log file from the server where the copy is now stored.

17. Check the log file for errors.

- If there are errors, correct the condition which caused the errors and reissue the command. You must correct all errors before you can proceed to step 17.
- If you need assistance, contact the IBM Support Center.

18. Issue the following CLI command to restore the cluster configuration:

```
svcconfig restore -execute
```

Note: Issuing this CLI command on a single node cluster adds the other nodes and hosts to the cluster.

This CLI command creates a log file in the /tmp directory of the configuration node. The name of the log file is svc.config.restore.execute.log.

19. Issue the following command to copy the log file to another server that is accessible to the cluster:

```
pscp -i <private key location> [user@]host:source target
```

20. Open the log file from the server where the copy is now stored.

21. Check the log file to ensure that no errors or warnings have occurred.

Note: You might receive a warning that states that a licensed feature is not enabled. This means that after the recovery process, the current license settings do not match the previous license settings. The recovery process continues normally and you can enter the correct license settings in the SAN Volume Controller Console at a later time.

The following output is displayed after a successful cluster configuration restore operation:

```
IBM_2145:your_cluster_name:admin>
```

22. After the cluster configuration is restored, verify that the quorum disks are restored to the MDisks that you want by using `svcinfo lsquorum` command. To restore the quorum disks to the correct MDisks, issue the appropriate `svctask setquorum` CLI commands.

You can remove any unwanted configuration backup and restore files from the `/tmp` directory on your configuration by issuing the `svconfig clear -all` CLI command.

Note: The recovery process does not recreate the superuser password and SSH keys. Ensure those are recreated before managing the recovered cluster.

Chapter 8. Upgrading the SAN Volume Controller software

The following corrections should be noted.

Upgrading clusters with internal SSD drives

Page 303/304. The following additional information is provided.

The SAN Volume Controller upgrade process reboots each node in the cluster in turn. Before the upgrade commences and before each node is upgraded, the upgrade process checks for dependent VDIs. Dependent VDIs can occur on nodes with internal SSDs drives if a VDisk has been recently created or a copy of the VDisk has been offline recently so the node has the only up-to-date copy of the data.

The upgrade process takes each node offline temporarily to perform the upgrade. While the node containing an internal SSD is offline, any data written to VDIs with a mirrored copy on the offline node will only be written to the other online copy. Once the upgraded node rejoins the cluster, data will be resynchronized from the copy that remained online. The upgrade process will delay approximately 30 minutes before starting the upgrade on the partner node, the synchronization must complete within this time or the upgrade will stall and require manual intervention. Any mirrored VDisk that uses disk extents on a solid-state drive (SSD) that is located

on a SAN Volume Controller node for one or both of its VDisk copies should have its synchronization rate set to 80 or above to ensure that the resynchronization completes in time.

Note: To increase the amount of time between the two nodes containing VDisk copies going offline during the normal upgrade process, consider using the User-paced Software upgrade procedure.

The following table defines the synchronization rates:

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

Upgrading solid-state drive (SSD) software

Page 309. The following additional information is provided.

Note: This procedure upgrades firmware on a solid-state drive (SSD) that is internal to a supported SAN Volume Controller node. If the upgrade could cause any VDIs to go offline, the **force** button must be used. For example, a firmware update to a managed MDisk requires the **force** option.

Chapter 11. Configuring and servicing storage systems

The following corrections should be noted.

Configuring EMC Symmetrix and Symmetrix DMX systems

Page 378. The following additional information is provided.

On some versions of Symmetrix and Symmetrix DMX, the setting of SPC-2 can be configured. This is set either on a per port basis or on a per initiator basis. LUs mapped to SAN Volume Controller must be configured with SPC-2 disabled.

Note: changing the value of the SPC-2 setting on a live system can cause errors. If you have a live system running with SPC-2 enabled on LUs mapped to SAN Volume

Controller, contact IBM support for guidance on how to proceed. Do not disable SPC-2 on a live system before taking guidance from IBM support.

Global settings for IBM System Storage DS5000 or IBM DS4000 systems

Page 398. The following information replaces this section in the original document.

Global settings apply across IBM System Storage DS5000 or IBM DS4000 systems. Table 40 lists the global settings that can be used with SAN Volume Controller clusters.

Table 40. IBM System Storage DS5000 and DS4000 system global options and recommended settings

Option	IBM DS5000 or IBM DS4000 system recommended setting
Start flushing	80%
Stop flushing	80%
Cache block size	4Kb (for systems running 06.x or earlier) 8Kb or 16Kb (for systems running 07.x or later)

Attention: Refer to DS5000/DS4000 documentation for details on how to modify the settings

Depending on the IBM DS5000 or IBM DS4000 model, use a host type of IBM TS SAN VCE or SAN Volume Contr to establish the correct global settings for the SAN Volume Controller cluster. Either set this as the system default host type or, if partitioning is enabled, associate each SAN Volume Controller port with this host type.