

IBM System Storage SAN Volume Controller
IBM Storwize V7000



Information Center Errata

Version 6.3.0
April 27, 2012

Contents

Introduction	4
Who should use this guide	4
Last Update	4
Change History	4
iSCSI Limits	5
iSCSI Limits with Multiple I/O Groups	5
Definition of terms	5
Limits that take effect when using iSCSI	6
Single I/O Group Configurations	6
iSCSI host connectivity only	6
Mixed iSCSI and Fibre Channel host connectivity	6
Multiple I/O Group Config	7
Symptoms of exceeding the limits	7
Configuring the HP 3PAR F-Class and T-Class Storage Systems	8
Minimum Supported STORWIZE V7000 Version	8
Configuring the HP 3PAR Storage System	8
Supported models of HP 3PAR Storage Systems	8
Support firmware levels of HP 3PAR storage arrays	8
Concurrent maintenance on HP 3PAR storage arrays	8
HP 3PAR user interfaces	8
HP 3PAR Management Console	9
HP 3PAR Command Line Interface (CLI)	9
Logical units and target ports on HP 3PAR storage arrays	9
LUNs	9
LUN IDs	9
LUN creation and deletion	10
LUN Presentation	10
Special LUNs	10
LU access model	11
LU grouping	11
LU preferred access port	11
Detecting Ownership	11
Switch zoning limitations for HP 3PAR storage arrays	11
Fabric zoning	11
Target port sharing	11
Controller splitting	12
Configuration settings for HP 3PAR storage array	12
Logical unit options and settings for HP 3PAR storage array	12
Creation of CPG	12
Set up of Ports	13
Setup of Host	14
LUN creation	15
Host options and settings for HP 3PAR storage array	16

Quorum disks on HP 3PAR storage arrays	16
Clearing SCSI reservations and registrations.....	17
Copy functions for HP 3PAR storage array.....	17
Thin Provisioning for HP 3PAR storage array	17
Recommended Settings for Linux Hosts.....	18
Multipath settings for specific Linux distributions and Releases	19
Udev Rules SCSI Command Timeout Changes	21
Editing the udev rules file	22

Introduction

This guide provides errata information that pertains to release 6.3.0 of the *IBM System Storage SAN Volume Controller Information Center and the IBM Storwize V7000 Information Center*.

Who should use this guide

This errata should be used by anyone using iSCSI as a method to connect hosts, Connecting Linux hosts using Fibre Channel or when connecting HP 3PAR Storage to IBM System Storage SAN Volume Controller or IBM Storwize V7000 .

Last Update

This document was last updated: April 27, 2012.

Change History

The following revisions have been made to this document:

Revision Date	Sections Modified
Nov 18, 2011	New publication
Apr 27 2012	Linux Host Settings

Table 1: Change History

iSCSI Limits

iSCSI Limits with Multiple I/O Groups

The information is in addition to, and a simplification of, the information provided in the Session Limits pages at the following links:

http://publib.boulder.ibm.com/infocenter/Storwize V7000/ic/index.jsp?topic=/com.ibm.storage.Storwize V7000.console.doc/Storwize V7000_iscsisessionlimits.html

http://publib.boulder.ibm.com/infocenter/storwize/ic/topic/com.ibm.storwize.v7000.doc/Storwize V7000_iscsisessionlimits.html

Definition of terms

For the purposes of this document the following definitions are used:

IQN: an iSCSI qualified name – each iSCSI target or initiator has an IQN. The IQN should be unique within the network. Recommended values are of the form `iqn.<date>.<reverse domain name>:<hostname>.<unique id>` e.g. `iqn.03-1996.com.ibm.hursley:host1.1`

initiator: an IQN that is used by a host to connect to an iSCSI target

target: an IQN on an STORWIZE V7000 or V7000 node that is the target for an iSCSI login

target portal: an IP address that can be used to access a target IQN. This can be either an IPv4 or an IPv6 address.

Limits that take effect when using iSCSI

Single I/O Group Configurations

iSCSI host connectivity only

1 target IQN per node

2 iSCSI target portals (1xIPv4 and 1xIPv6) per network interface on a node

4 sessions per initiator for each target IQN

256 defined iSCSI host object IQNs

512 host iSCSI sessions per I/O group **

256 host iSCSI sessions per node (this is to allow the hosts to reconnect in the event of a failover)

** e.g. if a single initiator logs in 3 times to a single target count this as 3. If a single initiator logs in to 2 targets via 3 target portals each count this as 6.

Only the 256 defined iSCSI IQN limit is enforced by the GUI or CLI commands.

Mixed iSCSI and Fibre Channel host connectivity

512 total sessions per I/O group where:

1 defined FC host object port (WWPN) = 1 session

1 defined iSCSI host object IQN = 1 session

1 additional iSCSI session to a target = 1 session

If the total number of defined FC ports & iSCSI sessions in an I/O group exceeds 512, some of the hosts may not be able to reconnect to the STORWIZE V7000/V7000 targets in the event of a node IP failover. See above section for help on calculating the number of iSCSI sessions.

Multiple I/O Group Config

If a host object is defined in more than one I/O group then each of its host object port definitions is counted against the session limits for every I/O group it is a member of. This is true for both FC and iSCSI host objects. By default a host object created using the graphical user interface is created in all available I/O groups.

Symptoms of exceeding the limits.

The following list is not comprehensive. It is given to illustrate some of the common symptoms seen if the limits defined above are exceeded.. These symptoms could also indicate other types of problem with the iSCSI network.

- The host reports a time out during the iSCSI login process
- The host reports a time out when reconnecting to the target after a STORWIZE V7000/V7000 node IP failover has occurred.

In both of the above cases no errors will be logged by the STORWIZE V7000/V7000 system.

Configuring the HP 3PAR F-Class and T-Class Storage Systems

Minimum Supported STORWIZE V7000 Version

6.2.0.4

Configuring the HP 3PAR Storage System

This portion of the document covers the necessary configuration for using an HP 3PAR Storage System with an IBM Storwize V7000 cluster.

Supported models of HP 3PAR Storage Systems

The HP 3PAR F-Class (Models 200 and 400) the HP 3PAR T-Class (Models 400 and 800) are supported for use with the IBM STORWIZE V7000. These systems will be referred to as HP 3PAR storage arrays. For the latest supported models please visit <http://www-01.ibm.com/support/docview.wss?uid=ssg1S1003907>

Support firmware levels of HP 3PAR storage arrays

Firmware revision HP InForm Operating System 2.3.1 (MU4 or later maintenance level) is the supported level of firmware for use with IBM STORWIZE V7000. For support on later versions, consult <http://www-01.ibm.com/support/docview.wss?uid=ssg1S1003907>

Concurrent maintenance on HP 3PAR storage arrays

Concurrent Firmware upgrades (“online upgrades”) are supported as per HP procedures.

HP 3PAR user interfaces

Users may configure an HP 3PAR storage array with the 3PAR Management Console or HP 3PAR Command Line Interface (CLI).

HP 3PAR Management Console

The management console accesses the array via the IP address of the HP 3PAR storage array. All configuration and monitoring steps are intuitively available through this interface.

HP 3PAR Command Line Interface (CLI)

The CLI may be installed locally on a Windows or Linux host. The CLI is also available through SSH.

Logical units and target ports on HP 3PAR storage arrays

For clarification, partitions in the HP 3PAR storage array are exported as Virtual Volumes with a Virtual Logical Unit Number (VLUN) either manually or automatically assigned to the partition.

LUNs

HP 3PAR storage arrays have highly developed thin provisioning capabilities. The HP 3PAR storage array has a maximum Virtual Volume size of 16TB. A partition Virtual Volume is referenced by the ID of the VLUN.

HP 3PAR storage arrays can export up to 4096 LUNs to the STORWIZE V7000 Controller (STORWIZE V7000's maximum limit). The largest Logical Unit size supported by STORWIZE V7000 under PTF 6.2.0.4 is 2TB, STORWIZE V7000 will not display or exceeded this capacity.

LUN IDs

HP 3PAR storage arrays will identify exported Logical Units through SCSI Identification Descriptor type 3.

The 64-bit IEEE Registered Identifier (NAA=5) for the Logical Unit is in the form; 5-OUI-VSID .

The 3PAR IEEE Company ID of 0020ACh, the rest is a vendor specific ID.

Example 50002AC000020C3A.

LUN creation and deletion

Virtual Volumes (VVs) and their corresponding Logical Units (VLUNs) are created, modified, or deleted through the provisioning option in the Management Console or through the CLI commands. VVs are formatted to all zeros upon creation.

To create a VLUN, highlight the Provisioning Menu and select the Create Virtual Volume option. To modify, resize, or destroy a VLUN, select the appropriate Virtual Volume from the window, right click when the specific VLUN is highlighted.

*** Note: Delete the mdisk on the STORWIZE V7000 Cluster before deleting the LUN on the HP 3PAR storage array.

LUN Presentation

VLUNs are exported through the HP 3PAR storage array's available FC ports by the export options on Virtual Volumes. The Ports are designated at setup and configured separately as either Host or Target (Storage connection). Ports being identified by a node : slot : port representation.

There are no constraints on which ports or hosts a logical unit may be addressable. To apply Export to a logical unit, highlight the specific Virtual Volume associated with the Logical Unit in the GUI and right click and select Export.

Special LUNs

There are no special considerations to a Logical Unit numbering. LUN 0 may be exported where necessary.

Target Ports

A HP 3PAR storage array may contain dual and/or quad ported FC cards. Each WWPN is identified with the pattern 2N:SP:00:20:AC:MM:MM:MM where N is the node, S is the slot and P is the port number on the controller and N is the controller's address. The MMMMMM represents the systems serial number.

Port 2 in slot 1 of controller 0 would have the WWPN of 20:12:00:02:AC:00:0C:3A. The last 4 digits of serial number 1303130 in hex (3130=0x0C3A). This system has a WWNN for all ports of 2F:F7:00:02:AC:00:0C:3A.

LU access model

All controllers are Active/Active. In all conditions, it is recommended to multipath across FC controller cards to avoid an outage from controller failure. All HP 3PAR controllers are equal in priority so there is no benefit to using an exclusive set for a specific LU.

LU grouping

LU grouping does not apply to HP 3PAR storage arrays.

LU preferred access port

There are no preferred access ports on the HP 3PAR storage arrays as all ports are Active/Active across all controllers.

Detecting Ownership

Detecting Ownership does not apply to HP 3PAR storage arrays.

Switch zoning limitations for HP 3PAR storage arrays

There are no zoning limitations for HP 3PAR storage arrays.

Fabric zoning

When zoning an HP 3PAR storage array to the STORWIZE V7000 backend ports, be sure there are multiple zones or multiple HP 3PAR storage array and STORWIZE V7000 ports per zone to enable multipathing.

Target port sharing

The HP 3PAR storage array may support LUN masking to enable multiple servers to access separate LUNs through a common controller port. There are no issues with mixing workloads or server types in this setup.

Host splitting

There are no issues with host splitting on an HP 3PAR storage array.

Controller splitting

HP 3PAR storage array LUNs that are mapped to the Storwize V7000 cluster cannot be mapped to other hosts. LUNs that are not presented to STORWIZE V7000 may be mapped to other hosts.

Configuration settings for HP 3PAR storage array

The management console enables the intuitive setup of the HP 3PAR storage array LUNs and export to the Storwize V7000 cluster.

Logical unit options and settings for HP 3PAR storage array

From the HP 3PAR storage array Management Console the following dialog of options are involved in setting up of Logical Units.

Creation of CPG

The set up of Common Provisioning Groups (CPGs). If Tiering is to be utilised, it should be noted it is not good practice to mix different performance LUNs in the same STORWIZE V7000 mdiskgrp.

Action->Provisioning->Create CPG (Common Actions)

General

General

System	3PAR (1303130)	
Domain	<none>	
Name	R1_CPG	
Template	<none>	
Allocation Warning	85	% <input checked="" type="checkbox"/> Enabled
Growth Increment	<default>	GiB <input type="checkbox"/> Enabled
Growth Warning		GiB <input type="checkbox"/> Enabled
Growth Limit		GiB <input type="checkbox"/> Enabled
Device Type	FC	
Device RPM ?	15 K	
RAID Type	RAID 1 (default)	
Set Size	2 (default)	

Set up of Ports

Shown is on a completed 8 node STORWIZE V7000 cluster.
Each designated Host ports should be set to Mode; point.

Connection Mode: Host

Connection Type: Point

System->Configure FC Port (Common Actions)

General	
System	3PAR (1303130)
Port	0:1:1
Connected Device Type	Host
Connected Device	svc, svc, svc, svc, svc, svc, svc, svc, svc, svc, svc, svc, svc, svc, svc, svc
Mode Change	Prohibited
State	● Ready

Settings	
Connection Mode	Host
Connection Type	Point
Unique Node WWN	Disabled
Configured Rate	Auto (default)
VLUN Change Notification (VCN)	Disabled
Interrupt Coalesce	Enabled

Setup of Host

Host Persona should be: 6 – Generic Legacy.
 All STORWIZE V7000 ports need to be included.

Actions->Hosts->Create Host (Common Actions)

General	
General	
System	3PAR (1303130)
Domain	<none>
Name	svc_
Set Name	<none>
Persona	6 - Generic-legacy

LUN creation

Size limitations: 256 MiB minimum

2TB maximum (STORWIZE V7000 limit)

Provisioning: Fully Provision from CPG

Thinly Provisioned

CPG: Choose provisioning group for new LUN, usually R1,R5,R6 or drive specific.

Allocation Warning: Level at which warning is given, optional [%]

Allocation Limit: Level at which TP allocation is stopped, optional [%]

Grouping: For creating multiple sequential LUNs in a set [integer values, 1-999]

Actions->Provisioning->Create Virtual Volumes (Common Actions)

General

General

System: 3PAR (1303130) ▼

Domain: <none> ▼

Name: Name_of_LUN

Template: <none> ▼

Comments:

User Space

Size: 200.000 GiB ▼

Provisioning: Thinly Provisioned ▼

Thinly Provisioned Virtual Volume. The user space and copy space are separate and allocated on demand from one or more CPGs.

CPG: CPG_R1 ▼
(RAID 1, FC, 1,238.250 GiB Used, No Limit)

Allocation Warning: 85 % Enabled

Allocation Limit: % Enabled

Grouping

Count (1-999): 1

Attempt to share LDs

Set Name: <none> ▼

Exporting LUNs to STORWIZE V7000

Host selection: choose host definition created for STORWIZE V7000

Actions->Provisioning->Virtual Volumes->Unexported (Select VV and right click)

Settings

Settings

Host Host Set

reagen (--)
svc (1:1:1, 1:1:2, 0:1:2, 0:1:1)

Multiple host selection is allowed when only one volume or volume set is selected and the selected volume or volume set doesn't have a one host policy.

Port: <none>

LUN: 1965 Auto

Allowed values are: n, m-n, or n+.

Override lower priority templates

After export, do not issue VLUN change notification (no LIP or RSCN)

Host options and settings for HP 3PAR storage array

The host options required to present the HP 3PAR storage array to Storwize V7000 clusters is, “6 legacy controller”.

Quorum disks on HP 3PAR storage arrays

The Storwize V7000 cluster selects disks that are presented by the HP 3PAR storage array as quorum disks. To maintain availability with the cluster, ideally each quorum disk should reside on a separate disk subsystem.

Clearing SCSI reservations and registrations

You must not use the HP 3PAR storage array to clear SCSI reservations and registrations on volumes that are managed by Storwize V7000. The option is not available on the GUI.

Note; the following CLI command should only be used under qualified supervision, “setvv –clrsv”.

Copy functions for HP 3PAR storage array

The HP 3PARs copy/replicate/snapshot features are not supported under STORWIZE V7000.

Thin Provisioning for HP 3PAR storage array

The HP 3PAR storage array provides extensive thin provisioning features. The use of these thin provisioned LUNs is supported by STORWIZE V7000.

The user should take notice of any warning limits from the Array system, to maintain the integrity of the STORWIZE V7000 mdisks and mdiskgrps. An mdisk will go offline and take its mdiskgroup offline if the ultimate limits are exceeded. Restoration will involve provisioning the 3PAR Array LUN, then including the mdisk and restoring any slandered paths.

Recommended Settings for Linux Hosts

The following details the recommended multipath (DMMP) settings and udev rules for the attachment of Linux hosts to SAN Volume Controller and Storwize V7000. The settings are recommended to ensure path recovery in failover scenarios and are valid for x-series, all Intel/AMD based servers and Power platforms.

A host reboot is required after completing the following two steps

- Editing the multipath settings in `etc/multipath.conf`
- Editing the udev rules for SCSI command timeout

For each Linux distribution and releases within a distribution please reference the default settings under `[/usr/share/doc/device-mapper-multipath.*]` for Red Hat and `[/usr/share/doc/packages/multipath-tools]` for Novell SuSE. Ensure that the entries added to `multipath.conf` match the format and syntax for the required Linux distribution. Only use the `multipath.conf` from your related distribution and release. Do not copy the `multipath.conf` file from one distribution or release to another.

Note for some OS levels the "polling_interval" needs to be located under defaults instead of under device settings.

If "polling_interval" is present in the device section, comment out "polling_interval" using a # key

Examples

Under Device Section

```
#          polling_interval 30,
```

Under Defaults Section

```
defaults {
    user_friendly_names yes
    polling_interval 30
}
```

Multipath settings for specific Linux distributions and Releases

Edit `/etc/multipath.conf` with the following parameters and confirm the changes using `“multipathd -k”show config`.

RHEL61

```
device {
    vendor "IBM"
    product "2145"
    path_grouping_policy group_by_prio
    getuid_callout "/lib/udev/scsi_id --whitelisted --device=/dev/%n"
    features "1 queue_if_no_path"
    prio alua
    path_checker tur
    failback immediate
    no_path_retry "5"
    rr_min_io 1
#    polling_interval 30
    dev_loss_tmo 120
}
```

RHEL56

```
device {
    vendor          "IBM"
    product         "2145"
    path_grouping_policy group_by_prio
    prio_callout    "/sbin/mpath_prio_alua /dev/%n"
    path_checker    tur
    failback        immediate
    no_path_retry   5
    rr_min_io       1
#    polling_interval 30
    dev_loss_tmo    120
}
```

RHEL57

```
device {
    vendor      "IBM"
    product     "2145"
    path_grouping_policy group_by_prio
    prio_callout "/sbin/mpath_prio_alua /dev/%n"
    path_checker tur
    failback    immediate
    no_path_retry 5
    rr_min_io   1
    dev_loss_tmo 120
}
```

SLES10SP4

```
device {
    vendor "IBM"
    product "2145"
    path_grouping_policy "group_by_prio"
    features "1 queue_if_no_path"
    path_checker "tur"
    prio "alua"
    failback "immediate"
    no_path_retry "5"
    rr_min_io "1"
#   polling_interval 30
    dev_loss_tmo 120
}
```

SLES11SP1

```
device {
    vendor      "IBM"
    product     "2145"
    path_grouping_policy group_by_prio
    prio        alua
    features     "0"
    no_path_retry 5
    path_checker tur
    rr_min_io   1
    failback    immediate
#   polling_interval 30
    dev_loss_tmo 120
}
```

```
}
```

SLES11SP2

```
device {  
    vendor "IBM"  
    product "2145"  
    path_grouping_policy "group_by_prio"  
    prio "alua"  
    path_checker "tur"  
    failback "immediate"  
    no_path_retry "5"  
    rr_min_io 1  
    dev_loss_tmo 120  
}
```

Udev Rules SCSI Command Timeout Changes

Set the udev rules for SCSI command timeout

Set SCSI command timeout to 120s

OS Level	Default	Required Setting
RHEL61	30	120
RHEL62	30	120
RHEL56	60	120
RHEL57	60	120
SLES10SP4	60	120
SLES11SP1	60	120
SLES11SP2	30	120

Creating a udev rules file

Create the following udev rule that increases the SCSI command timeout for SVC and V7000 block devices

udev rules file

```
cat /etc/udev/rules.d/99-ibm-2145.rules
```

```
# Set SCSI command timeout to 120s (default == 30 or 60) for IBM 2145 devices
SUBSYSTEM=="block", ACTION=="add", ENV{ID_VENDOR}=="IBM",
ENV{ID_MODEL}=="2145", RUN+="/bin/sh -c 'echo 120 >
/sys/block/%k/device/timeout'"
```

Reconfirm the settings following the system reboot.