



Storage and Backup Engineering
Citi Architecture & Technology Engineering (CATE)

HDS VSP/HUS-VM Solutions Guide

CSBE-10071-SG

Solutions Guide

Version 2.5

1 Version Control

Version	Date	Author	Reviewed	Comments
0.1	08/10/2011	Gary Johnson	Ludwig Derodel Fibish Madathil Chris Clayton	Peer Review
1.0	08/31/2011	Gary Johnson	Global Storage Community	RFT
1.1	09/30/2011	Gary Johnson	Global Storage Community	RFP
1.2	03/12/2012	Gary Johnson		Updated Package List
1.3	07/10/2012	Siddhraj Choksi Gary Johnson	Fibish Madathil Gary Johnson	Updated Packages SCSI Reserves NDM Migrations. Microcode 70-03-33 and HCS 7.2.1
2.0	02/04/2013	Gary Johnson	Gary Johnson Jeff Hernandez	Added HUS-VM Array and 14+2
2.1	03/11/2013	Gary Johnson	Jeff Hernandez	Host Mode 55 HUS- VM
2.2	04/29/2013	Gary Johnson Jeff Hernandez	Jeff Hernandez Gary Johnson	SOM 803/729 and DRU
2.3	9/27/2013	Jeff Hernandez	Ludwig Derodel Gary Johnson	Standard Configurations, LRR, Reliability, LUN sizes
2.4	11/25/2013	Gary Johnson	Jeff Hernandez	TMEA and HTnM
2.5	5/20/2014	Jeff Hernandez Gary Johnson	Ludwig Derodel Gary Johnson	FMD Configuration Options, TMEA, HTnM
2.5.1	04/15/2015	Gary Johnson	Jeff Hernandez	Updated FID List

For details on changes to these production versions see the document [Version Notes](#).

2 Regional Reviewers

The following regional reviewers provided feedback for the current version of this document:

Region	Version Last Reviewed
Asia Pacific	
EMEA	
Japan	
Latin America / Mexico	
North America	

3 Legal Notice

Information contained herein is for Internal Use and may be used only for business purposes authorized by Citi.

CITIGROUP, CITIBANK, CITICORP, CITI and the Umbrella Device are trademarks and service marks of Citicorp and are used and registered throughout the world. Citibank, Citicorp, Citi and their subsidiaries also claim rights in certain other trademarks and service marks contained in these pages.

4 Contact

Please address any comments or questions regarding the information contained in this document via email to the “CATE Storage and Backup Engineering” distribution list or alternatively to gdse@citi.com.

5 Latest Versions

It is recommended that the latest version of the document be used from the GDSE website: <https://catecollaboration.citigroup.net/domains/platstor/gdse/products/techstdshmpgs/tagmastore.aspx>

Newer versions may include expanded scope, or change in policy on issues discovered after deployment.

6 Table of Contents

1	VERSION CONTROL	2
2	REGIONAL REVIEWERS	2
3	LEGAL NOTICE.....	3
4	CONTACT	3
5	LATEST VERSIONS	3
6	TABLE OF CONTENTS.....	4
7	INTRODUCTION.....	11
8	RELEASE STATUS	11
9	PREREQUISITES FOR A HDS VSP/HUS-VM ENVIRONMENT	12
9.1	REQUIRED LICENSING FOR ARRAY (BASED ON RAW ARRAY STORAGE)	12
9.2	REQUIRED LICENSING FOR REPLICATED ARRAY (BASED ON RAW ARRAY STORAGE)	12
9.3	REQUIRED LICENSING PER SITE (BASED ON TOTAL RAW STORAGE)	12
10	ARCHITECTURE.....	14
10.1	OVERVIEW	14
10.2	VSP HARDWARE	14
10.2.1	Logical Components of VSP	14
10.2.2	Back-End SAS Links.....	15
10.2.3	BED to DKU Organization.....	17
10.2.4	Dual Chassis (Dual DKC).....	17
10.2.5	Slot Names	18
10.3	HUS-VM HARDWARE.....	20
10.3.1	HUS-VM Hardware Overview	20
10.3.2	Logical Components of HUS-VM Different from VSP.....	20
10.3.3	Cache and Control Memory.....	24
10.3.4	Front-End Design	24
10.3.5	High-Performance SAS Back-End Design.....	24
10.3.6	HUS-VM and VSP Component differences	26
10.4	DISK RAID MECHANISM	27
10.4.1	RAID 5 (7+1).....	27
10.4.2	RAID 6 (6+2).....	28
10.4.3	RAID 6 (14+2).....	28
10.4.4	RAID 10 (2+2).....	28
10.5	LIMIT COMPARISON OF USPV, VSP AND HUS-VM	28
11	INSTALLATION.....	29
11.1	REQUIRED INSTALLATION STEPS.....	29
11.2	SOFTWARE PACKAGES	30
11.3	SOFTWARE PACKAGE DETAILS / NOTES	30
11.3.1	HDvM CLI version 7.2.....	30
11.3.2	HDvM CLI version 7.2.1 for Windows 2K8.....	30
11.3.3	jre16016.....	30
11.3.4	CITI_WIN_2K8_HDSHCS_74_01.....	31
11.3.5	HTnM_AFR_740.....	33
12	VSP BEST PRACTICES.....	35
12.1	STORAGE HOST SETUP (HOST MODES)	35
12.2	SPECIAL HOST MODE OPTION 55 FOR HUS-VM AND POWERPATH	36
12.2.1	Raidcom Command for Host Mode Option 55.....	36

12.2.2	Storage Navigator.....	36
12.2.3	DMP Hosts.....	37
12.2.4	Removal of Setting	37
12.3	RAID CONFIGURATIONS	37
12.3.1	VSP RAID (Parity) Group Configurations.....	38
12.3.2	HUS-VM RAID (Parity) Group Configurations.....	38
12.4	DRIVE TYPES.....	38
12.4.1	SAS vs. FC.....	39
12.4.2	SATA vs. SAS.....	39
12.5	DRIVE REBUILD TIMES.....	39
12.6	ARRAY CONFIGURATION TYPES.....	40
12.6.1	VSP Single Tier Non Replicated Arrays.....	40
12.6.2	HUS-VM Single Tier Non Replicated Arrays for Production.....	42
12.6.3	HUS-VM Single Tier Arrays for Development and Archive	43
12.6.4	Replication Array Types for TC/S and HUR.....	44
12.6.5	HUS-VM All FMD Non Replicated Arrays for Production	46
12.6.6	Dynamic Tiering Arrays (Arrays using HDT).....	47
12.7	SPARE DRIVES.....	48
12.8	LUN EXPANSION	48
12.8.1	Expansion with a New LUN.....	48
12.8.2	Expansion with Host Volume Management	48
12.8.3	Expansion of Thin Volumes.....	49
12.9	HOST PORT BALANCING AND REDUNDANCY	49
12.9.1	WWN Numbering for Host Ports	49
12.9.2	VSP Port Options and Slot Names.....	50
12.9.1	VSP Host Ports and Replication Ports.....	50
12.9.2	HUS-VM Port Options and Slot Names.....	52
12.9.3	HUS-VM Host Ports and Replication Ports.....	52
12.10	TOPOLOGY	52
12.11	MAXIMUM HOSTS PER PORT	52
12.12	MAXIMUM PORTS PER HOST	53
12.13	MANAGEMENT SOFTWARE (HCS vs. SN).....	53
12.14	SAN MONITORING	53
12.15	ARRAY MONITORING: ALERTS AND REMOTE CONNECTIONS	53
12.15.1	SNMP.....	54
12.15.2	Vendor Notification (Dial-Out / HiTrack)	54
12.15.3	HDS Remote Access Application (RACC).....	54
12.16	NETWORK TIME PROTOCOL (SNTP OR NTP) – TIME SYNCHRONIZATION	54
12.17	MULTIPLE ARRAY MANAGEMENT.....	55
12.18	POLLING INTERVALS FOR DEVICE MANAGER.....	55
12.19	ENTERPRISE ALERTING (SNMP).....	56
12.20	RUNNING HCS ON VMWARE.....	58
12.21	HCS MANAGEMENT SERVER SIZING.....	58
12.22	RUNNING MULTIPLE VERSIONS OF SERVER SOFTWARE SIMULTANEOUSLY	58
12.23	ELEMENT MANAGER (RUNNING SN FROM HCS)	58
12.24	BACKUP, RESTORE, AND MIGRATION OF DEVICE MANAGER AND RELATED DATABASES	59
12.24.1	Backup HDvM Databases.....	59
12.24.2	Restore HDvM Databases.....	59
12.24.3	Restrictions to Backup and Restores.....	59
12.24.4	Migrations (Export and Import of HDvM databases).....	60
12.25	MAXIMUM LUN SIZE	61
12.26	INCREASED HEAP SIZE (HDvM MEMORY EXPANSION)	61
12.27	SYSTEM OPTIONS (SOM)	62
12.28	CITI THIN PROVISIONING (HDP) BEST PRACTICES.....	63
12.28.1	Virtual to Physical Capacity (Thin Provisioning Rules).....	63
12.28.2	Pool Composition.....	63

12.28.3	Pool Naming	64
12.28.4	Pool Creation.....	64
12.28.5	Pool Growth.....	64
12.28.6	Migrations.....	65
12.28.7	DP-VOL Sizing	65
12.28.8	Thresholds.....	65
12.28.9	HDP Alerts.....	65
12.28.10	Pool Full Condition	68
12.28.11	Application Migrations.....	69
12.28.12	Application and Host Operating System Considerations.....	69
12.29	ZERO DATA REMOVAL TOOL	71
12.29.1	Caveats and Restrictions.....	71
12.29.2	Using the Tool.....	71
12.30	TOOL PANEL	73
12.31	SCSI RESERVATIONS ON VSP (VIEW / MODIFY).....	74
12.32	NDM (SEAMLESS HOST MIGRATION) ON VSP – USPv SPOOFING	75
12.33	SHARED MEMORY REQUIREMENTS	76
12.34	SAN “PORT FLAPPING” ALERTING	76
13	SECURITY PRACTICES.....	77
13.1	INTERFACES TO THE USP/USPV/VSP.....	77
13.2	DEFAULT USERIDs	77
13.3	SPECIAL USER ACCOUNTS (FUNCTIONAL IDs).....	78
13.4	AUDIT LOGGING VIA SEMS.....	78
13.4.1	Configuring Device Manager Auditing.....	79
13.5	USER ENTITLEMENTS REPORTING	79
13.5.1	Storage Navigator User Entitlements for non AD accounts	79
13.5.2	Device Manager User Entitlements.....	80
13.6	DECOMMISSIONING OF STORAGE	80
13.7	SETTING SECURITY POLICY IN DEVICE MANAGER	80
13.7.1	Password Policy.....	80
13.7.2	Failed Login Attempts.....	81
13.7.3	Security Banner.....	81
13.8	CCI SECURITY BIT	82
13.9	HDvM CLI SECURITY	82
13.10	HCS ARRAY NAME	82
13.11	ACTIVE DIRECTORY AUTHENTICATION WITH HCS	82
13.11.1	Acquire a functional id used for binding to the AD environment.....	82
13.11.2	Obtain the SSL certificate used by AD.....	83
13.11.3	Add certificate to the keystore on the HCS server	83
13.11.4	Verify the keystore contents on the HCS server	84
13.11.5	Configure the exauth.config file.....	84
13.11.6	Defining AD Groups on HCS.....	85
13.12	ACTIVE DIRECTORY AUTHENTICATION VIA RADIUS WITH HCS.....	87
13.12.1	IAS Installation and Setup.....	87
13.12.2	HCS Configuration of Radius	96
13.12.3	HCS User Setup	98
13.13	CONFIGURING IE FOR USE WITH HCS AND SN	99
13.14	IMPORTING SERVER CERTIFICATES FOR HCS / ELEMENT MANAGER	100
13.14.1	Exporting the Certificates from the Arrays.....	100
13.14.2	Importing the Certificates from the Arrays to the HCS Server	103
13.15	CONFIGURING HCS FOR HTTPS	107
13.15.1	Create SSL keys for secure CLI access.....	108
13.15.2	Generate private key	109
13.15.3	Create a Certificate Signing Request.....	110
13.15.4	Creating a Self-Signed Certificate	110

13.15.5	Copy new key files to prod SSL directory	110
13.15.6	Modify conf file to use HTTPS and find the key files	111
13.15.7	Change SSL URL to HTTPS.....	111
13.15.8	Setup GUI/CLI on other Windows Hosts	111
14	REFERENCES.....	112
15	APPENDIX A: ARRAY PHYSICAL REQUIREMENTS	113
15.1	VSP AND HUS-VM POWER	113
15.1.1	Single Phase Power	113
15.1.2	Three Phase Power	113
15.1.3	PDU Plugs	113
15.2	VSP DIMENSIONS AND WEIGHT	113
15.2.1	Dimensions.....	114
15.2.2	Rack Naming.....	114
15.2.3	Service Clearances.....	115
15.3	HUS-VM DIMENSIONS AND WEIGHT.....	116
15.3.1	Dimensions.....	116
15.3.2	Rack Naming.....	116
15.3.3	Service Clearances.....	117
15.4	HVAC (COOLING)	118
15.5	NETWORK CONNECTIVITY TO SVP(s).....	118
15.6	PHONE LINE CONNECTIVITY	118
15.7	FIBER CONNECTIVITY	118
16	APPENDIX B: CONFIGURATION OF ARRAY VIA HCS (HDVM).....	119
16.1	ACCESSING HCS / HDvM	119
16.2	HCS SERVICES	119
16.3	LICENSE KEYS FOR HCS	120
16.4	STORAGE SYSTEM / ARRAY CONFIGURATION	121
16.4.1	Add Storage System.....	122
16.4.2	Refresh Storage Systems	123
16.4.3	Edit Storage Systems.....	123
16.4.4	Remove Storage System	124
16.5	HOST CONFIGURATION.....	124
16.5.1	Add Hosts.....	124
16.5.2	Allocate Volumes	125
16.5.3	Refresh Hosts.....	128
16.5.4	Edit Hosts.....	128
16.5.5	Edit LUN Paths.....	129
16.5.6	Import Hosts from Storage Navigator (Scan Hosts).....	131
16.5.7	Merge Hosts.....	133
16.5.8	Removing Hosts (Deleting Hosts).....	136
17	APPENDIX C: CONFIGURATION OF ARRAY VIA STORAGE NAVIGATOR.....	137
17.1	ACCESSING STORAGE NAVIGATOR.....	137
17.2	INFORMATION SCREEN	137
17.3	UPDATES AND LOCK MANAGEMENT	138
17.4	MODIFY MODE (SOME OLDER SCREENS)	138
17.5	SECURITY IN STORAGE NAVIGATOR.....	139
17.6	LOGIN BANNER	141
17.7	SYSLOG BASED AUDIT LOG	142
17.8	SYSLOG AUDIT RECORD FORMAT	143
17.8.1	Syslog Server Setup.....	146
17.9	PORT STATUS	147
17.10	SYSTEM EVENTS	148

17.11	PORT DEFINITION	150
17.12	RAID (PARITY) GROUP DEFINITION	151
17.13	LUSE VOLUMES	151
17.14	DISK SHREDDING	151
17.15	REUSE OF LUNs (DELETING LUNs FROM HOST FOR FUTURE RE-ALLOCATION).....	152
17.15.1	Reuse of LUNs: Procedures for Storage Navigator:.....	152
17.15.2	Reuse of LUNs: Procedures for RAIDCOM:.....	155
17.16	DISPLAYING WWNs LOGIN (LIST LOGINS FUNCTION)	155
18	APPENDIX D: CLI FOR HCS (HDVM CLI).....	157
18.1	HiCOMMAND CLI SETUP (SOLARIS)	157
18.2	HiCOMMAND CLI SETUP (WINDOWS)	157
18.3	SECURE OPTION	157
18.4	SECURE PASSWORD PROMPTING (USING JRE 1.6 OR LATER).....	157
18.5	PROPERTIES FILE	158
18.5.1	Using properties file to store read-only username/password	159
18.6	CLI STRUCTURE.....	159
18.7	HOST MODES	161
18.8	ARRAY DISCOVERY.....	161
18.9	ADDING LUNS.....	162
18.10	CLI EXAMPLE USE CASE	166
19	APPENDIX E: COMMAND CONTROL INTERFACE / SHADOWIMAGE	172
19.1	CCI SECURITY	172
19.2	HORCM FILES.....	172
19.2.1	HORCM_MON	172
19.2.2	HORCM_CMD	172
19.2.3	HORCM_DEV	173
19.2.4	HORCM_INST.....	173
19.3	CCI CONFIGURATION FILES	173
19.4	GETTING THE LUN INFORMATION.....	175
19.5	STARTING THE INSTANCES	175
19.6	STARTING REPLICATION.....	176
19.7	RESYNCHRONIZING VOLUMES	177
19.8	DELETING PAIR RELATIONSHIPS	177
20	APPENDIX F: FIREWALL PORTS	178
21	APPENDIX G: CITI MCAFEE VIRUS SCAN AND EPO AGENT.....	179
21.1	INSTALLATION OF ANTI-VIRUS SOFTWARE	179
21.1.1	Prior To Beginning Installation.....	179
21.1.2	Installation Steps.....	179
21.2	GLOBAL EPO POLICY EXCLUSIONS	180
21.3	UNINSTALLING ANTI-VIRUS SOFTWARE	181
21.4	REPORTING AND TREE STRUCTURE.....	181
22	APPENDIX H: SETUP OF HITRACK SERVICE FOR HTTPS.....	182
22.1	SETUP OF HI-TRACK SOFTWARE	182
23	APPENDIX I: HDS TIERED STORAGE MANAGER (MIGRATIONS).....	186
23.1	WHAT IS HTSM?	186
23.2	MIGRATION STEPS.....	187
23.2.1	Example Task.....	190
23.2.2	Example Email Notification.....	191
23.2.3	Steps in the actual migration process (for DP vols targets only).....	192
23.3	HTSM CONFIGURATION	192

23.3.1	Licensing.....	192
23.3.2	User Management setup	193
23.3.3	Configuration of server.properties files.....	193
23.3.4	Audit Log Configuration - event log	194
23.4	CITIGROUP HTSM BEST PRACTICE AND OPERATIONAL NOTES	194
23.4.1	I/O loading during migrations	194
23.4.2	Delayed execution of a task	194
23.4.3	Disk shredding	195
23.5	USING HTSM WITH REMOTELY-REPLICATED STORAGE.....	195
23.5.1	General	195
23.5.2	TrueCopy Sync (TC) specific comments	196
23.5.3	Hitachi Universal Replicator (HUR) specific comments	197
24	APPENDIX J: SNMP MIB DEFINITIONS	198
25	APPENDIX K: PERFORMANCE MONITOR DATA EXTRACT TOOL	199
25.1	WHAT IS PERFORMANCE MONITOR DATA EXTRACT TOOL?	199
25.2	PREPARING TO USE THE EXTRACT TOOL.....	199
25.2.1	Preparing the Array.....	199
25.2.2	Choosing a Server.....	199
25.2.3	Temporary User.....	199
25.2.4	Software	199
25.3	USING THE EXTRACT TOOL	199
25.3.1	Creating the Command File.....	199
25.3.2	Optional Replication Commands.....	200
25.3.3	Optional Data Ranges (Longrange and Shorrange Settings).....	200
25.3.4	Creating a Batch File.....	200
25.3.5	Running the Export Tool.....	201
25.3.6	Timing the Export Tool.....	201
25.3.7	Error Codes from the Export Tool.....	201
25.3.8	Optional Data Ranges (Longrange and Shorrange Settings).....	202
25.4	TRANSMITTING THE DATA TO HDS.....	202
26	APPENDIX L: SERVER PRIORITY MANAGER (QOS).....	203
26.1	CONSIDERATIONS AND RISKS	203
26.2	SETUP AND USE OF SPM	203
27	APPENDIX M: TUNING MANAGER (HTNM)	205
27.1	INSTALLATION OF SOFTWARE	205
27.2	TUNING MANAGER SETUP / CONFIGURATION	205
27.2.1	Creating Array Agents (Agents for RAID).....	205
27.2.2	Verify Agents.....	216
27.2.3	Polling Times.....	217
27.2.4	Manual Polling (Starting or Interrupting Polling)	218
27.2.5	Performance Reporter Configuration (60-day data storage)	219
27.3	PERFORMANCE REPORTING METHODOLOGY.....	224
27.3.1	Basic Terminology	224
27.3.2	Review of a Typical Application Transaction	224
27.3.3	Response Time Factors.....	224
27.3.4	General Performance Monitoring.....	225
27.3.5	Performance Analysis Methodology	225
27.4	CITI PERFORMANCE CRITERIA FOR HDS ARRAYS.....	226
27.5	CITI STANDARD REPORTS AND ALARMS	227
27.5.1	Pre-Defined Reporting Setup.....	227
27.5.2	Pre-Defined Reporting.....	227
27.5.3	Alarms Setup.....	231

27.5.4	Email Alerting Setup.....	231
27.5.5	RAID Alarms.....	234
27.6	THIN POOL (HDP) TREND REPORTING	236
27.7	VIRTUAL VOLUME TREND REPORTING (DETAIL OF POOL TREND REPORT)	239
27.8	GENERAL ARRAY REPORTING	241
28	APPENDIX N: TMEA REPORTING.....	246
28.1	SETUP AND CONFIGURATION.....	246
28.1.1	Use of Proxy.....	246
28.1.2	Root Access Setup (Powerbroker).....	246
28.1.3	Storage for TMEA.....	246
28.1.4	Installation of RMLIB	247
28.1.5	Installing TMEA Collector Package	248
28.1.6	License Files	248
28.1.7	Command Devices	249
28.1.8	Configuring TMEA Collector.....	249
28.1.9	Starting and Stopping Instances	251
28.1.10	Current Status.....	252
28.1.11	Manual SFTP Transfer	253
28.1.12	Removing an Instance.....	253
28.1.13	Removing TMEA Collector	254
28.1.14	TMEA Space Cleanup Maintenance (Cleanup Script).....	254
28.1.15	Process Impact to RHEL Server	255
28.1.16	Debug Mode.....	255
28.1.17	Automatic Process Management (Keeping TMEA Running).....	256
28.2	TMEA REPORTING WEBSITE	257
28.2.1	URL for TMEA.....	257
28.2.2	Accounts for Accessing TMEA.....	257
28.2.3	Distribution Lists for Reports	257
28.3	TMEA SUPPORT	257
29	APPENDIX O: SVP PERFORMANCE COLLECTION.....	259
29.1	CONFIGURE EXPORTS ACCOUNT	259
29.2	INITIAL COLLECTION AND POST MICROCODE UPDATE STEPS	263
29.3	CREATE DAILY SCHEDULED TASKS	265
29.4	CONFIGURATION OF BROWSER.....	272
30	APPENDIX P: SETUP OF AD FOR STORAGE NAVIGATOR.....	274
31	APPENDIX Z: VERSION NOTES	282

7 Introduction

In order to provide Citi with more choices for Enterprise Class storage, Global Distributed Storage Engineering is certifying the latest Enterprise storage arrays from Hitachi Data Systems. The HDS Virtual Storage Platform (VSP) and Hitachi Unified Storage (HUS-VM) arrays are the latest generation of enterprise-class storage that are certified for use globally within Citi Technology Infrastructure (CTI).

This guide will cover the new installations and additions of components necessary for connecting and supporting the VSP and HUS-VM SAN Storage in a heterogeneous server environment.

8 Release Status

This product is engineering ready. Security approvals (ISRP and any related processes) are being explored. This product is very similar to the previous Hitachi USPV Enterprise Storage Array and subsequently similar approvals are expected.

9 Prerequisites for a HDS VSP/HUS-VM Environment

The following items are required in order to implement a HDS VSP or HUS-VM according to CSBE standards:

- A Storage Management Workstation (Windows Server 2k8). This physical (or VMware hosted) server will run all of the GUI interfaces for HCS (Hitachi Command Suite) that includes Device Manager (HDvM), Tiered Storage Manager (HTSM), and Replication Manager (HRpM). This windows server should also host the Tuning Manager (HTnM) software that will eventually be fully incorporated into HCS. These applications should run exclusively of other major applications on the host.
- If command line functionality is desired, then an additional management system (Solaris or RH Linux only) will be required to run Device Manager CLI, and optionally the ShadowImage/Replication Command Control Interface (CCI).
- A Citi Extranet Platform (CEP) or firewall that encompasses and selectively isolates the HDS SVP devices (IP connections to array).
- Dial-In capability for hardware support / diagnostics via HDS RACC.
- SEMS audit logging infrastructure environment.
- NetApp OnCommand Insight (formerly SANScreen) SAN auditing environment is required (via HCS).
- ECC for chargeback and reporting (via HCS).
- Appropriate physical floor space, cooling, power, and redundant storage SAN fabric connectivity.

9.1 Required Licensing for Array (based on raw array storage)

- Hitachi Basic System (BOS)
 - o Device Manager (HDvM)
 - o Provisioning Manager
 - o Hitachi Dynamic Provisioning (HDP)
 - o Hitachi Open Volume Management / Virtual LUN Software
 - o Hitachi LUN Manager
 - o Hitachi Performance Monitor
 - o Hitachi Server Priority Manager
 - o Hitachi Storage Navigator 2
 - o Hitachi SNMP Agent
 - o Hitachi Virtual LVI
 - o Hitachi Volume Shredder
 - o Hitachi Volume Security
 - o Hitachi Volume Security Port Option
 - o Others not certified or recommended for use at Cit
- Hitachi Basic System and Virtualization (BOSV)
 - o Hitachi Universal Volume Manager (UVM)
- ShadowImage (Internal array replication)
- Replication Manager (External Replication Control)

9.2 Required Licensing for Replicated Array (based on raw array storage)

- ShadowImage (Internal array replication)
- Hitachi Universal Replicator (HUR)

9.3 Required Licensing per site (based on total raw storage)

- Tuning Manager (HTnM)
- Tiered Storage Manager (HTSM)

- Replication Manager (External Replication Control) if Replication Used

10 Architecture

10.1 Overview

The VSP (Virtual Storage Platform) and HUS-VM (Hitachi Unified Storage) Arrays are the latest generation of Hitachi Enterprise Storage. There have been several key upgrades from their USPV arrays that are discussed below and many of the core features are maintained. For example, all major elements are installed in pairs across the power domain. The built-in service processor (SVP), a blade PC running Windows, still hosts the Storage Navigator web interface and is the maintenance platform for on-site CEs.


The basic architecture of the VSP array has been upgraded in comparison to the USPV. However, much of the basic structure and division of work is similar to the USPV. The exception being the introduction of the new Virtual StorageDirectors and their interaction with the front end and back end directors. Overall the speed of the system has been increased, but many of the initial microcode limitations are similar to the USPV. This is expected to change with future releases of the next few quarters. The number of drive slots has increased along with the drive density per floor tile.

10.2 VSP Hardware

10.2.1 Logical Components of VSP

The major components of the VSP are modularized into functional roles:

- Host/external storage adapters (FED)
- Back-end adapters (BED)
- Cache (DCA)
- Data processors (VSD)
- Grid switches (GSW)



10.2.1.1 Grid Switch (GSW)

The HiStar-E PCI Express Switched Grid is the communication backbone of the VSP. This is the interconnection between the FED, BED, DCA, and VSD boards. There are 24 high speed ports per board, where each port supports a full-duplex rate of 2048MB/s (1024MB/s send and 1024MB/s receive concurrently). The Grid Switch (GSW) is also the interconnection point

between the new dual chassis (large) configuration. This allows the dual-chassis array to function as a single integrated array, not as a pair of units operating as a cluster. There can be 2 or 4 GSWs per chassis. This component is also called ESW in some documentation.

10.2.1.2 Virtual Storage Director (VSD)

Virtual Storage Directors manage all I/O by sets of assigned LDEVs in conjunction with the updated FEDs and BEDs. There are either 2 or 4 VSDs installed per chassis. These boards execute all of the internal I/O jobs (dispatched from the FEDs and BEDs) and manage the control memory tables.

10.2.1.3 Front End Director (FED)

The processors on the FEDs that were 800MHz NEC RISC “Y-Wing” CPUs have been replaced with 2.33GHz Intel Xeon Core Duo quad core CPUs (Tachyon) on the VSP. The FEDs now have Data Directors that work with the VSDs. The FEDS auto negotiates between 2Gbps, 4Gbps, and 8Gbps.

10.2.1.4 Back End Director (BED)

The BEDs are the disk controllers for the array. The processors on the BEDs are 2.33GHz Intel Xeon Core Duo quad core CPUs on the VSP. The BED provides 8x backend 6Gbps SAS links via four SAS 2W ports (16 SAS links per BED pair, up to 32 SAS links per chassis). The BEDs now have Data Directors that work with the VSDs. The BED component is also called a DKA in some documentation.

10.2.1.5 Data Cache Board (DCA)


Each DCA cache board has eight DDR3-800 DIMM slots, organized as four banks of RAM (thus 32 independent banks of RAM in a dual chassis array). Each DCA board can support 8 to 32GB of RAM using the 4GB DDR3 DIMMs. When the 8GB DDR3 DIMMs are released, this can grow up to 64GB. Each DCA has four 2GB/s full duplex GSW ports that provide for 8GB/s of aggregate read-write bandwidth per board to the FEDs and BEDs. There are two on-board SSDs (and batteries) on the DCA that provide a de-staging area for cache during a power outage.

10.2.1.6 Disk Storage


The disk storage basic building block has changed significantly on this platform. There are no FC disks in the array. The primary disk drive type is a SAS drive that is available in 146GB, 300GB, 600GB, and 900GB (2.5” enclosure). SATA drives are also available in a 3TB form (3.5” enclosure). SSDs are also available in 200GB (2.5”), 400GB (2.5”), and 1.6TB FMD.

10.2.2 Back-End SAS Links

In the performance model configuration that Citi uses, there are up to 4 BED pairs. Each BED pair is connected to all of the DKUs in the same chassis. In the dual chassis model this is repeated from the other chassis (see below).




Each disk is dual connected to 2 SAS channels. Each channel is comprised of 4x SAS Links.



These connections would be identical on 3.5" (LFF) drives such as SATA and SSDs, though the number of slots would be reduced from 16 to 10.

10.2.3 BED to DKU Organization


A chassis containing a DKC can be expanded from the 2 DKUs in the same DKC Rack to have up to 2 additional DKU Racks (each with 3 DKUs). Under the performance model configuration, there will be 4x BEDs. The SAS links from the 4x BED pairs are extended to cover the additional drives as follows:



The first rack (Rack 00) will contain DKU-00 and DKU-01. The second rack (Rack 01) will contain DKU-02, DKU-03, and DKU-04. The third rack (Rack 02) will contain DKU-05, DKU-06, and DKU-07. The SAS cables depicted above are actually SAS 2W (2-Wide) cable, each with two independent 6Gbps SAS links. Thus


10.2.4 Dual Chassis (Dual DKC)

When there is a need for more than 8x DKUs, then the DKC can be expanded into another chassis. The second chassis will function much like the original DKC but is fully integrated with inter-component communication between the two sides. The grid switches (GSW) are cross attached in a one to one relationship with a partner GSW in the other chassis.




10.2.5 Slot Names

In order to properly identify failed or expanded components, it may be useful to understand the names associated with the individual component slots on the VSP. The following are the Logic Box slot names for Chassis #1 (DKC 00):



Once a second Chassis is added to the configuration (DKC 10), the following are the slot names for the additional components:




10.3 HUS-VM Hardware

10.3.1 HUS-VM Hardware Overview

The HUS-VM is essentially the same architecture as the VSP with some consolidation of components. The HUS-VM is based on the HiStar architecture which provides a shared resource system where all the primary components, including front-end ports, back-end ports, cache modules and processors, are connected through an internal switched network. The internal processors share a global cache along with the front-end and back-end I/O modules, which creates a common pool of storage resources and disk capacity. This disk capacity can include external storage, which is connected through host initiator fiber channel ports.

The controller chassis has dual controllers, and each controller has a cache blade and a processor blade. All data passes through the cache blade ASIC to cache. The processor blade executes all I/O jobs and manages cache.




10.3.2 Logical Components of HUS-VM Different from VSP

HUS-VM hardware is composed of the following major components that are different from the VSP:

- Controller Chassis
- Drive trays – Support for large form factor (LFF) and small form factor (SFF)
- Rack – HDS or standard 19-inch rack using supplied rail kits


HUS-VM

**10.3.2.1 Controller Chassis**


The controller chassis contains dual HiStar-E based controller modules, 2 main modules, 2 microprocessor modules (MPB), an integrated service processor, and dual power supplies with built-in cooling fans. Each cache blade has slots for the system cache, contains an ASIC that includes most of the discrete logic previously found on the VSP FED and BED boards, and a battery. The processor blade contains an Intel Xeon 8-core processor and local RAM.

Main components can be added, removed and replaced in the block module without shutting down a device while the storage system is in operation. The microcode can also be upgraded without shutting down the storage system. An integrated service processor is mounted in the controller chassis that monitors the running condition of the storage system.

Controller chassis




Controller chassis rear detail




10.3.2.2 Drive Trays (Drive Boxes)

There are four different drive trays (boxes) that are supported with the HUS-VM. All drive type use a 6Gb/sec SAS interface. These drive boxes are:

DBS -- A 2U box containing up to 24 vertically mounted 2.5" small form factor (SFF) drives. This box contains two redundant power supplies and two ENC adapter boards for controller connectivity. The connection to the controller may be direct or via daisy-chain through other ENC adapters in other boxes.




DBL – A 2U box containing up to 12 horizontally mounted 3.5" large form factor (LFF) drives. This box contains two redundant power supplies and two ENC adapter boards for controller connectivity. The connection to the controller may be direct or via daisy-chain through other ENC adapters in other boxes.



DBF – A 2U high box containing up to 12 horizontally mounted 5-1/4" Hitachi flash module drives (FMD). This box contains two redundant power supplies and two ENC adapter boards for controller connectivity. The connection to the controller may be direct or via daisy-chain through other ENC adapters in other boxes.



DBX – A 4U "high-density" box containing up to 48 vertically mounted 3.5" large form factor (LFF) drives. The box is functionally separated into two separate drive units (A and B), each containing a maximum of 24 drives. Each drive unit contains two redundant power supplies and two ENC adapter boards for controller connectivity. The connection to the controller may be direct or via daisy-chain through other ENC adapters in other boxes.



10.3.3 Cache and Control Memory

Hitachi Unified Storage VM uses partitioned cache rather than a shared global cache. A portion of this cache space (4GB to 20GB per blade, mirrored) is reserved for the control memory region. There is also a transfer buffer region where the host ports and drive ports transfer data blocks. There is also space for 2 or 4 user data cache directories (500MB each).

The control memory and cache directory regions are backed up to each cache blade's SSD on a planned shutdown. On power loss, all of cache per blade is backed up to its local SSD on power loss.

The remaining cache space is managed as a global pool. This space is evenly divided into four partitions for the 4 microprocessor units. A microprocessor unit may borrow some space from another partition under certain circumstances. A microprocessor unit cannot access its own data partition – only the cache blades can do so with DMA transfers, but the microprocessors manage their partitions. Each processor blade microprocessor unit uses DMA to access control memory and the cache directories on the cache blades, but the cache blades do not.

10.3.4 Front-End Design

Each host I/O module contains one 4-port Tachyon QE8 8Gb/sec Fibre Channel controller. Each module is connected to one PCIe x8 link on the ASIC on the cache blade. HUS-VM supports Fibre Channel for open systems and SAN connectivity. HUS VM supports up to 32 Fibre Channel ports from up to 8 host I/O modules per controller chassis. Additional ports, up to 48, can be added by replacing 1 to 4 drive modules per controller chassis with host I/O modules. The Fibre Channel ports support industry standard 2Gb/sec, 4Gb/sec and 8Gb/sec transfer rates.


10.3.5 High-Performance SAS Back-End Design

The back end of Hitachi Unified Storage VM is a SAS engine that is a cross between the HUS 100 family and Hitachi VSP. The SAS engine elements include the drive I/O modules, up to 48 disk trays with their dual 4-link SAS switches (ENCs), and the SAS drives. The cache blade controls all communications between the drive modules and the processor blades and all DMA transfers to cache. Each drive module is connected to one PCIe x8 link on the ASIC in the main blade. Each drive I/O module contains one SAS controller device with a pair of 4-wide ports (4 x SAS 6Gb/sec lanes).


To optimize performance, it is recommended to create parity groups so that as many back-end (BE) paths can be used as possible. This can also be called vertical striping the parity groups. Horizontal striping of parity groups is not recommended for higher performance environments. Because the HUS-VM uses a 4-link SAS switch to communicate with the drive trays, it is recommended to have a minimum of 4 drive trays in order to maximize performance on the BE paths.

In a dense chassis, different DB numbers are assigned to the left and right sides, and different BE paths are defined for the BE numbers. To optimize performance in a HUS-VM with dense trays, create parity groups by using the left and right drives evenly, again so that the load of BE paths is balanced.

Standard Tray Configuration (DBS, DBL, DBF)



Dense Tray Configuration (DBX)



10.3.6 HUS-VM and VSP Component differences

The HUS-VM and VSP share the same logical components, but are physically implemented with some slight differences. This results in a potential larger performance impact to component failure on the HUS-VM, but the array remains fully redundant and will continue to service IO during the failure period.


Component	Impact of Failure on VSP	Impact of Failure on HUS-VM	Replacement Urgency
Front-End Adapter	Loss of 8 FC Ports. Path failure for attached hosts.	Loss of 4 FC Ports. Path failure for attached hosts.	Medium
MP	Back-up MP assumes control of failed MP's LDEV's. High CPU utilization can occur if combined LDEV workload exceeds 100% of remaining CPU's limits	Back-up MP assumes control of failed MP's LDEV's. High CPU utilization can occur if combined LDEV workload exceeds 100% of remaining CPU's limits.	High
Grid Switch	Remaining Grid switches handle all internal traffic requests until failed GSW is replaced.	If the Control Module main ASIC fails then all associated LDEV management is failed over to the other control module. Path loss for all multi-pathed hosts.	High
Memory Module	The VSP uses mirrored write cache so the system would operate in degraded cache mode	Each of the two HUS-VM control modules has two banks of cache. When a failure occurs in DIMM in	High

	until failed module is replaced.	CMG1, only CMG1 is blocked. Note that CMG0 is also used for shared memory information. If a DIMM in CMG0 has a failure, MAIN Blade is blocked (CHBs/DKBs in the cluster are also blocked).	
DKC Adapter	Loss of up to 8 8Gb SAS Links, A fully configured system will have up to 56 remaining links to balance the workload across.	Loss of up to 8 8Gb SAS Links, A fully configured system will have up to 24 remaining links to balance the workload across.	Medium
DKU	All DKU's have multiple power supplies, fans, and SAS interconnects. Each DKU contains 8 LFF or SFF HDU's. Failed DKU's are extremely rare, and raid groups are generally striped across equal number of HDU's allowing for greater redundancy.	HUS-VM uses 24 slot and 48 slot drive trays. Each drive tray has multiple connections to each of the two control modules. All trays are fault tolerant with redundant power supplies, fans, and interconnects. Failures in HDS drive trays are extremely rare. Horizontal striping is allowed for lower tier configurations.	Medium
SVP - Service Processor	All Citigroup VSP systems are configured with dual SVP for failover.	HUS-VM systems are only available with a single SVP. Impact includes loss of SN connectivity, Call home reporting, and out-of-band Raidcom connectivity. SVP should be replaced as soon as possible.	Medium / High

10.4 Disk RAID Mechanism

10.4.1 RAID 5 (7+1)

RAID 5 (7+1) is only used for SSDs at Citi on the VSP or HUS-VM. In traditional RAID 5 schemas, the stripe width (row size in diagram) by multiplying the RAID chunk (stripe) size by 7. For Open-V volumes (open systems), the VSP uses a native chunk size of 256KB. Traditionally this would be a row size of 1,792KB. However, the VSP manages RAID 5 somewhat differently than other vendors. Two adjacent 256KB RAID chunks are taken from the same disk before switching to the next disk in a row, so the effective chunk size is 512KB. Thus, there are twice as many chunks per row than expected. The VSP stripe width is 3,584KB (7 x 512KB).



10.4.2 RAID 6 (6+2)

RAID 6 (6+2) is an allowed storage RAID protection used on the VSP at Citi. The VSP manages RAID 6 similar to RAID 5 where there are two adjacent RAID chunks taken from the same disk before switching to the next disk in the row. The VSP RAID 6 6+2 stripe width is 3,072KB (6 x 512KB).



10.4.3 RAID 6 (14+2)

RAID 6 (14+2) is the default RAID configuration for Bronze tier on both the VSP and HUS-VM. RAID 6 (14+2) is basically the same as RAID 6 (6+2) but uses twice as many data disks in the RAID set and again only 2 parity disks. The advantage from using RAID 6 (14+2) is a wider stripe across more disks giving better performance and more capacity from the same number of disks as two RAID 5 (7+1) parity groups..

10.4.4 RAID 10 (2+2)

RAID 10 (2+2) is the primary storage RAID protection for Silver and Platinum tiers used on the VSP and HUS VM respectively at Citi. RAID 10 (2+2) is mirrored disks, primary and secondary drives, compose a RAID pair and the identical data is written to the primary and secondary drives. RAID 10 (2+2) provides reliability and higher performance.

10.5 Limit Comparison of USP, VSP and HUS-VM

There are many new aspects to the VSP and HUS-VM arrays. The key features that were added since the USP are:

- Improved Storage Navigator Interface (Task Oriented)
- Pool shrink (take storage away from pools)
- Multiple tiers within same pool (HDT - Hitachi Dynamic Tiering)

- Page Level Tiering (HDT)

The following is a comparison of the limits from the VSP and HUS-VM arrays. These limits are expected to be changed over the next few quarters as newer microcode releases are made.

Limits	VSP Single Chassis	VSP Dual Chassis	HUS-VM	USPV
Data Cache (GB)	512	1024	256	512
Raw Cache Bandwidth	64 GB/s	128 GB/s	--	68 GB/s
Control Memory (GB)	8-48	8-48	4-20	24
Host Port Speed (or Ext Storage)	1,2,4,8 Gb/s	1,2,4,8 Gb/s	1,2,4,8 Gb/s	1,2,4 Gb/s
SSD Drives	128	256	128	128
2.5" SAS Disks	1024	2048	1152	-
3.5" SAS Disks	640	1280	576	-
3.5" FC Disks	-	-	-	1152
3.5" SATA Disks	640	1280	-	1152
Flash module drives	192	192	96	N/A
Logical Volumes	65,280	65,280	16,384	65,280
Max Volume Size (Internal)	2.99 TB	2.99 TB	2.99 TB	2.99 TB
Max Volume Size (External)	4 TB	4 TB	4 TB	4 TB
IO Request Limit per Port	2048	2048	2048	2048
Nominal Queue Depth per LUN	32	32	32	32
HDP Pools	128	128	128	128
Max Pool Capacity	5 PB	5 PB	2.1 PB	1.1 PB
Max Total Pool Capacity	5 PB	5 PB	2.1 PB	1.1 PB
LDEVs per Pool	1024	1024	1024	1024
Max Pool Volume Size	3.99 TB	3.99 TB	2.99 TB	2.99 TB
DP Volumes per Pool	Approx 62000	Approx 62000	16,384	8192
DP Volume Size (w/o replication)	46 MB – 60 TB	46 MB – 60 TB	46 MB – 60TB	46 MB – 4 TB
DP Volume Size (w/ replication)	46 MB – 4 TB	46 MB – 4 TB	46 MB – 4 TB	46 MB – 4 TB

11 Installation

11.1 Required Installation Steps

Major Steps:

- Physical Installation of Array into Enterprise Class Data Facility
- Network Installation of SVPs
- Installation of Management server (Windows)
- (Optional) Installation of Management server for CLI (RH or Solaris)
- Redundant management server(s)
- Firewalled Network Segment (CEP) Isolation of SVPs
- Installation of HCS and setup of Device Manage software
- Configuration of Array via Device Manager and Storage Navigator (on Mgmt Server)
- Configuration of HTTPS-based dial-home for alerting to HDS
- RACC software configuration (HDS Remote Access)
- Configure SYSLOG to send records to SEMS for auditing
- Create userid for OnCommand Insight (formerly SANScreen) for auditing
- Create userid for ECC (Prosphere)
- Configuration of host ports and SAN

- Allocation of storage to storage hosts
- Installation of HTnM and setup of remaining HCS software (HTnM, HRpM, ...)
- Setup of Replication (external or internal if applies)

11.2 Software Packages

The following packages are the current packages for the VSP array:

Software	Platform	Package Name
Hitachi Command Suite (HCS) version 7.4.1 (HDvM,HTSM,HRpM,HTnM)	Windows 2008	CITI_WIN_2K8_HDSHCS_74_01
Hitachi Device Manager CLI version 7.2.0	RHEL	HDvMCLI_7.2_RHEL5.x.tar.gz
Hitachi Device Manager CLI version 7.2.1	Windows 2008	HDvMCLI_7.2.1_WIN2K8
Hitachi Device Manager Agent version 7.4.1 (for CCI/Raidcom/HRpM)	Solaris 10	SOL_HdvMAgt_741
Hitachi Device Manager Agent version 7.4.1 (for CCI/Raidcom/HRpM)	RHEL	RHEL_HdvMAgt_741
Hitachi Tuning Manager Agent for RAID version 7.4.0	Solaris 10	HTnM_AFR_740
Hitachi Tuning Manager Agent for RAID version 7.4.1	RHEL 5.x	HTnM_AFR_740
Hitachi RAID Manager CCI /Raidcom version 01-29-03/06	Solaris10	SOL_RMHORC_01-29-03
Hitachi RAID Manager CCI /Raidcom version 01-29-03/06	RHEL 5.x	RHEL_RMHORC_01-29-03.tar
Hitachi TMEA version 9-5-3b12	RHEL 5.9	HDS_TMEA_RHEL_9-5-3b12
Sun Java JRE 1.6.0_16	Solaris 10	jre16016 (directory jre160161.6.0_16)

11.3 Software Package Details / Notes

When installing on a windows machine, you should always disable the On-Access Virus Scan to prevent conflicts with the setup software.

11.3.1 HDvM CLI version 7.2

This package (HDvMCLI_7.2_RHEL5.x.tar.gz) can be downloaded from GDS Unix website. This package is not really a Unix package format and it a tar file of the required software and must be untar'd to a local directory (such as /opt/HiCommandCLI). If you are going to upgrade your HDvM CLI package, then you will want to preserve a copy the HiCommand CLI.properties file manually. If you are unable to change the path (as described below for JRE) and the default JRE is a previous version, then the JRE version must be specified to hide the password (when prompted).

Two lines can be added to the HicommandCLI.sh script specifying this JRE version if changing the \$PATH variable is not desirable:

```
HDVM_CLI_JRE_PATH=/opt/jre/1.6.0_16/bin
export HDVM_CLI_JRE_PATH
```

For more information, please see [Appendix D: Command Line Interface for Device Manager \(HDvM CLI\)](#).

11.3.2 HDvM CLI version 7.2.1 for Windows 2K8

This package (HDvMCLI_7.2.1_WIN2K8) can be downloaded from GDS website. This package is not really a Windows package format – it is a self-extracting compressed file for the required software and must be executed in a local directory. If you are going to upgrade your HDvM CLI package, then you will want to preserve a copy the HiCommand CLI.properties file manually. If you are unable to change the path (as described below for JRE) and the default JRE is a previous version, then the JRE version must be specified to hide the password (when prompted).

11.3.3 jre16016

Using JRE 1.6 or later will enable the password hiding for the HDvMCLI. This JRE package will require the manual add of the /opt/jre/1.6.0_16/bin directory to the \$PATH variable in your profile. Remove any previous JRE versions from \$PATH. You may also need to delete a symbolic link that points to the wrong version (or nonexistent directory) in /bin or /usr/bin. Running “whence

java” should find the first instance of java following your path. Running “java –version” should tell you the current version. Do not assume that the package install has resolved these issues.

If you are unable to change the path (as described above) and the default JRE is a previous version, then the JRE version must be specified to hide the password (when prompted). Two lines can be added to the HicommandCLI.sh script specifying this JRE version if changing the \$PATH variable is not desirable:

```
HDVM_CLI_JRE_PATH=/opt/jre/1.6.0_16/bin
export HDVM_CLI_JRE_PATH
```

11.3.4 CITI_WIN_2K8_HDSHCS_74_01

It is critical that you disable on-access Virus Scanning, disable DEP for non-windows applications (reboot required), and enable the Application Experience service to successfully install the software. These changes can be reversed once installation is successful. These packages only place the original software into an install directory.

Notes:

1. Flash 10.1 is required to run HCS.
2. Application Experience Service is required to install HCS (HiRDB Setupper has stopped working error will occur without this service running).
3. During the new installation or upgrade of HCS the IPV6 agent will get installed and enabled on management server by default. Administrator will have to disable IPV6 manually after new installation or upgrade of HCS.
4. On the windows 2008 Citi build, there was an SNMP listening trap port conflict experienced during installation. If SNMP traps are used to/from HDvM, then this will need to be resolved.

HCS Server Hardware and Operating System requirement

The following are the table of minimum hardware and operating system requirements for the management server

Item	Requirement
Processor	Minimum: Dual-Core Processor Recommended: At least Quad-Core Processor
Physical memory	Minimum: 3 GB Recommended: At least 6GB*2
Disk space	Minimum: 5 GB Recommended: At least 20 GB

OS	Edition	Service pack	Hitachi Command Suite		Ipv6
			HDvM HTsM HRpM	HDC	
Windows Server 2003 R2 (x64)	Datacenter Edition Enterprise Edition Standard Edition	SP2	Y	Y	Y
Windows Server 2008 (x64)	Datacenter Edition Datacenter without Hyper-V Edition Enterprise Edition Enterprise without Hyper-V Edition Standard Edition Standard without Hyper-V Edition	No SP SP1 SP2 SP2 SP2 SP2	Y Y Y Y Y Y	-- -- Y Y Y Y	Y Y Y Y Y Y
Windows Server 2008 R2 (x64)	Datacenter Edition Enterprise Edition Standard Edition	No SP SP1 SP1	Y Y Y	Y Y Y	Y Y Y

Prerequisite JDK

For Window, In Hitachi Command Suite products, the following Oracle Sun JDK versions are available besides the bundled JDK 5 Update 03 (1.5.0_03) or later except in Windows XP or Windows Server 2003 R2 environment: – Oracle Sun JDK 5 (1.5.0), An Update 17 or later version is required for SSL communication between the external authentication server (LDAP directory server) and the management server.

If the OS of the management server is Windows, the file system of the management server must be NTFS in order to use Oracle Sun JDK. Use Oracle Sun JDK for 32 bit because Hitachi Command Suite products are 32 bit applications. After new installation, the bundled JDK in the product will be set to use. Exchange the JDK for Hitachi Command Suite products to use Oracle Sun JDK after the installation.

Server installation prerequisites

- Stop Tuning Manager services:
 - o If a version earlier than 6.3 is installed on the same management server, stop the Tuning Manager Agent for SAN Switch service.
 - o If Tuning Manager is remotely connected, stop the Tuning Manager server service.
- Run Tiered Storage Manager tasks when upgrading from 7.1.0 or earlier
 - o If there are incomplete tasks (standby, running, or being cancelled), goto the Tasks & Alerts tab and execute or cancel the tasks. After upgrading, register the canceled tasks as new.
- Determine values that are set during installation
 - o The installation destination path
 - o The database storage destination path
 - o Management server information
 - o Memory heap size
 - o The database backup destination

Prerequisite checker

The prerequisite checker verifies whether the installation destination satisfies the requirements.

- In Windows, use the package **CITI_WIN_2K8_HDSHCS_74_01** to run the **prereqchk.exe**. The file is in **C:\HCS_7.4.0_Install\HCS** directory of the package.

Installing HCS in Windows

Install the HCS in Windows with the package available for HCS 7.2.1.

- Log on to Windows as an administrator.

- Go to **C:\HCS_7.4.1_Install\HCS\setup.exe**, double click on **setup.exe** to run the setup for HCS 7.4.1
- Click Install in the window that appears.
- When you are prompted, enter the required information.
- In the Installation Completed window, select the When installation completes; open the Device Manager GUI check box.
- Click Finish to get the GUI login window
- Proceed with the install of HTnM (Tuning Manager) via the setup.exe in that subdirectory.

Verifying the installation

You can verify that the installation was successful by accessing the Web client from a browser. To access the HCS Web client following are the steps.

- Start a Web browser.
- In the address bar, specify the URL for HCS in the following format:
<http://Management-server-IP-address-or-hostname:23015/DeviceManager/>

RMHORC_01-29-03 This package is a software repository and does not run the installation.

- Download the RAID Manager CCI software package to a temporary directory (e.g. "/tmp/CCI") on the server
- Change the current directory to the 'Specification Directory' (Ex: /opt). Ensure that there is at least 100MB of free space on this filesystem.)
cd /'Specification Directory'
- Install the CCI software to 'Specification Directory' using "**cpio**" command
cpio -idmu < /tmp/CCI/SOL_RMHORC_01-29-03
- Makes symbolic link for /HORCM.
ln -s /'Specification Directory'/HORCM /HORCM
- Verify installation of the proper version using the "**raidqry**" command
raidqry -h
- Executes the RAID Manager installation command
/HORCM/horcminstall.sh

Note: If you are upgrading CCI from an existing older version, you will need to preserve any existing HORCM configuration files before uninstalling the old package and installing the new package.

11.3.5 HTnM_AFR_740

This package is not a package but a directory that contains an installation script. You will want to run the installation script (install.sh) from the Solaris or the Linux subdirectories. That script will add the package to your system or upgrade the existing agent package. For new installs on Solaris, a RAID Manager library should also be added. However, updating the RAID Manager

library from previous versions is not necessary. See [Setup of Agent for RAID for Solaris 10](#) for detailed instructions on using this package.

12 VSP Best Practices

It is assumed that the audience is familiar with the configuration of the array (via Storage Navigator, Device Manager, and/or CLI). For additional detailed information on these subjects, see the following appendixes:

- [Appendix B: Configuration of Array via HCS / HDvM](#)
- [Appendix C: Configuration of Array via Storage Navigator](#)
- [Appendix D: Command Line Interface for Device Manager](#)
- [Appendix L: Tiered Storage Manager](#)

12.1 Storage Host Setup (Host Modes)

For the latest information regarding which platforms have been certified for use with the HDS array and the appropriate path management software, refer to the GDSE Website: [CSBE Supported MultiPath Configuration Matrix](#).

Each supported host OS has its own stack information relating to the configuration and use with SAN-attached storage.

For Host modes in setting up Host Groups on the VSP or HUS-VM, use the following chart:

Operating System Type	Host Mode	Option	Array
AIX	OF		VSP (and HUS-VM)
AIX with VCS	OF	2, 22	VSP (and HUS-VM)
AIX with Oracle AMS	OF		VSP (and HUS-VM)
AIX with Oracle RAC & AMS	OF		VSP (and HUS-VM)
AIX with HACMP	OF	15	VSP (and HUS-VM)
AIX with VIO vSCSI	OF		VSP (and HUS-VM)
ESX (VMWare)	0		VSP (and HUS-VM)
HP-UX	3	12	VSP (and HUS-VM)
HP-UX with VCS	3	2,12,22	VSP (and HUS-VM)
HP-UX with Oracle AMS	3	12	VSP (and HUS-VM)
Linux (RH)	0		VSP (and HUS-VM)
Linux (RH) with RHCS (Red Hat Linux Cluster)	0		VSP (and HUS-VM)
Linux (RH) with VCS	0	2, 22	VSP (and HUS-VM)
Linux (RH) with VDBE/AC for Oracle RAC	0	2	VSP (and HUS-VM)
Solaris with VM	9	2	VSP (and HUS-VM)
Solaris with VCS	9	2, 22	VSP (and HUS-VM)
Solaris with VDBE/AC for Oracle RAC	9	2	VSP (and HUS-VM)
Solaris with Oracle AMS	9	2	VSP (and HUS-VM)
Solaris with ZFS	9	2	VSP (and HUS-VM)
Solaris with Sun Cluster	9		VSP (and HUS-VM)
Windows	2C		VSP (and HUS-VM)
Windows with VCS	2C	2, 22	VSP (and HUS-VM)
Windows with VDBE/AC for Oracle RAC	2C	2	VSP (and HUS-VM)
Windows Cluster	2C		VSP (and HUS-VM)
Windows Cluster with Veritas Volume Mgr	2C		VSP (and HUS-VM)

12.2 Special Host Mode Option 55 for HUS-VM and PowerPath

Until an update can be made to all versions of PowerPath, there is a critical host mode option that needs to be set on all host groups for hosts attached to HUS-VM arrays. Host Mode Option 55 (fifty five) must be set in order for Powerpath to manage HUS-VM LUNs. This Host Mode Option forces the HUS-VM array to respond as it were a VSP or USPV array. This option only applies to HUS-VM arrays. ***It is recommended that this be set for all hosts attached to the HUS-VM (using Powerpath or not) as it may be difficult for a SAN Administrator to know if a host is using Powerpath for path management and load balancing.***

It is critical to set this mode option consistently between host groups for the same server (or cluster of servers), otherwise the storage will appear to be from two different arrays on different paths.

To make Host Mode Option 55 visible under Storage Navigator or CCI, System Mode 847 (SOM847) must be enabled by HDS Support.

Note: In this VSP/USPV mode, the array will only respond with the last five digits of the six-digit serial number. Therefore, any HUS-VM arrays deployed need to be unique based on their last 5-digits at Citi.


12.2.1 Raidcom Command for Host Mode Option 55

The following syntax is the CLI command to change the host group option:

```
raidcom modify host_grp -port <port#> [<host group name>]
      -host_mode < host mode> [-host_mode_opt <host mode option> ... ]
```

12.2.2 Storage Navigator

To modify the host mode setting in SN (or Element Manager), you would go into each host group and select host mode option 55. With SOM847 enabled, you will see more options than normal. Select Edit Host Groups. Scroll down the list and enable 55. Select Finish and Apply. The changes take place as soon as the task has completed.



12.2.3 DMP Hosts

For hosts using DMP instead of Powerpath, this host mode option should not be set as it is not needed. DMP versions 5.1 and later fully recognize the HUS-VM Array however, setting it will not impede the functionality of DMP.

12.2.4 Removal of Setting

Once an updated Powerpath has been released and installed on the host, Host Mode Option 55 can be removed. However, this requires a rescan/reconfigure of storage (as it will appear to be a different array serial number). No data loss should occur, but the host will have a disruption of connectivity during the rescan. Once all hosts have been updated, System Mode 847 (SOM847) can be removed.

12.3 RAID Configurations

The following supported RAID Groups (also called Parity Groups) have the best performance characteristics on the HDS VSP and HUS VM:

Disk	RAID Configuration	Relative Capacity of Array	Relative Performance Level
200GB 6Gb SSD (2.5")	R5 (7+1)	Lowest	Highest
400GB 6Gb SSD (2.5")	R5 (7+1)	...	...
1.6TB 6Gb FMD	R5 (3+1)	...	...
1.6TB 6Gb FMD	R10 (2+2)	...	...
3.2TB 6Gb FMD	R10 (2+2)	...	...
300GB 6Gb SAS 15k HDD (2.5")	R10 (2+2)	...	...
300GB 6Gb SAS 10k HDD (2.5")	R10 (2+2)	...	...
600GB 6Gb SAS 10k HDD (2.5")	R10 (2+2)	...	...
300GB 6Gb SAS 15k HDD (2.5")	R6 (14+2)	...	...
600GB 6Gb SAS 10k HDD (2.5")	R6 (14+2)	...	...
900GB 6Gb SAS 10k HDD (2.5")	R6 (14+2)	...	...
3TB 6Gb SAS 7.2k (3.5")	R6 (14+2)	...	...

3TB 6Gb NL-SAS 7.2k (3.5")*	R6 (14+2)	Highest	Lowest
-----------------------------	-----------	---------	--------

*3TB NL-SAS is supported on HUS VM only

Note: RAID 10 is required for all enterprise storage systems at Citi supporting the Silver and Platinum tiers. The only exceptions to this are legacy systems that utilize RAID 6 configurations and Bronze tier arrays. SSDs which do not have long rebuild times, should be configured as RAID 5 (7+1) or (3+1). All other RAID Configurations are allowed only by exceptional approval from CATE.

Having multiple tiers of storage (heterogeneous configuration) within the same array is allowed and is appropriate for data centers and where applications need several tiers of storage available. RAID groups must be comprised of identical drives (size/speed). RAID groups should be in an 8-disk or 16-disk configuration. RAID groups should be combined into large homogeneous (or heterogeneous when certified) pools for thin (dynamic) provisioning.

12.3.1 VSP RAID (Parity) Group Configurations

The default RAID 6 14+2 configuration is supported on legacy VSP's and Bronze tier configurations only. The use of RAID10 2+2 is now the standard configuration supported for pool storage on the Silver tier. SSD and FMDs use a R5 (7+1) or R5 (3+1) configuration.

Different RAID configurations are allowed only by exceptional approval from CSBE.

12.3.2 HUS-VM RAID (Parity) Group Configurations

Storage arrays designated for development or archive storage may use RAID 6 14+2 with horizontal stripping (raid groups with members from same disk tray).

Storage arrays designated for production storage should use RAID 10 2+2 with vertical stripping (raid group members on different drive trays) for maximum performance. Any arrays with partial production and development should use RAID 10 2+2.

Different RAID configurations are allowed only by exceptional approval from CSBE.

12.4 Drive Types

The following drive types are supported on the VSP & HUS-VM:

Disk Type	Back-End Speed
200GB SSD (2.5")	6Gb
400GB SSD (2.5")	6Gb
1.6TB FMD	6Gb
3.2TB FMD	6Gb
300GB SAS 15k HDD (2.5") J300	6Gb
300GB SAS 10k HDD (2.5") J300	6Gb
600GB SAS 10k HDD (2.5") J600	6Gb
900GB SAS 10k HDD (2.5") J900	6Gb
3TB SAS 7.2k HDD (3.5") H3R0	6Gb
3TB NL-SAS 7.2k HDD (3.5") H3R0*	6Gb

*3TB NL-SAS is supported on HUS VM only

Note: The 146GB 15k drive, 1TB, 2TB NL-SAS, and all SATA drives have been removed from this list.

12.4.1 SAS vs. FC

The SAS drive has superior seek times and significantly better back-end (6Gb) when compared to the 15K FC standard. This allows the SAS 10k drive to perform within 10% of the FC 15k drive at a much lower cost and power impact. The 15k variant of the SAS drive is actually superior to the FC equivalent but it has a much higher cost/power than the 10K SAS so the future of the drive is likely not much longer. The 146GB 15K SAS drive is no longer on the CSBE supported drive list.

12.4.2 SATA vs. SAS

It is expected that over time, the 7.2k SAS drives will have sufficient capacity to edge out the SATA drives for high capacity. The SAS drives have superior connectivity (dual for SAS, single for SATA), better performance, and do not have the native formatting issue (512 on SATA without error checking, so error check is additional overhead). These high capacity SAS drives are sometimes referred to as near-line SAS (or NL-SAS). SATA disk offer very high capacity at an "economy" level of performance. These are best suited for archival duty and are not suitable for high levels of random workloads. SATA drives have been removed from the CSBE supported drives list.

12.5 Drive Rebuild Times

The following chart shows approximate rebuild times for each of the drive types in the supported RAID configurations. These times are based on an idle array (no other IO) and could be longer on a more active array.

VSP/HUS-VM Drive Type	Rebuild Time for RAID10 (2+2)	Rebuild Time for RAID5 (3+1)*	Rebuild Time for RAID5 (7+1)*	Rebuild Time for RAID6 (6+2)*	Rebuild Time for RAID6 (14+2)
200GB 6Gb SSD (2.5")	0hr 50min	0hr 50min	0hr 50min	0hr 50min	1hr 10min
400GB 6Gb SSD (2.5")	1hr 50min	1hr 50min	1hr 50min	1hr 50min	2hr 10min
300GB 6Gb SAS 10k HDD (2.5") J300	1hr 30min	1hr 30min	1hr 30min	1hr 30min	2hr 10min
600GB 6Gb SAS 10k HDD (2.5") J600	2hr 50min	2hr 50min	2hr 50min	2hr 50min	4hr 35min
900GB 6Gb 10k HDD (2.5") J900	3hr 20min	3hr 20min	3hr 20min	3hr 20min	5hr 10min
1.6TB 6Gb FMD	5hr 40min	5hr 40min	5hr 40min	5hr 40min	11hr 20min
2TB 6Gb NL-SAS 7.2k HDD (2.5") H2R0	11hr 00min	11hr 00min	11hr 00min	11hr 00min	11hr 00min
3TB 6Gb NL-SAS 7.2k HDD (2.5") H3R0	17hr 00min	17hr 00min	17hr 00min	17hr 00min	17hr 00min

*RAID 5 (3+1), RAID 5 (7+1) are not allowed configurations for SAS and SATA at Citi. The only supported RAID configurations for SAS and SATA are RAID 10 (2+2), RAID 6 (6+2) or RAID 6 (14+2). This is due to the long rebuild times and subsequent risk that this places on RAID Groups and storage pools. RAID 5 (7+1) or RAID 5(3+1) are the only configurations allowed for the SSD series drives due to their increased cost and faster rebuild times. Different RAID configurations are allowed only by exceptional approval from CSBE.

12.6 Array Configuration Types

12.6.1 VSP Single Tier Non Replicated Arrays

There are two configurations of the VSP used at Citi that do not utilize HDT (Dynamic Tiering). These models use the 2.5" slots exclusively (SAS only) and are considered a 'Silver' tier offering.

Summary of Configurations

Config #	Usable Capacity (TB)	Configuration Objective	Number of working disks/spares	RAID	Latency Response Rating (IOPS)	Reliability Rating (yrs)	Cache Per Chassis/Total	Number of Bays (include system bay)	IOPS Per uTB
VSP - Single Chassis – Single Tier Non Replicated									
1	100	Medium	760/8	R10 2+2	62,100	216,294	160GB/256GB	3	621
VSP - Dual Chassis – Single Tier - Non Replicated									
2	200	Large	1528/8	R10 2+2	124,800	107,581	160GB/256GB	5	624

Terminology Reference

- **Latency Response Rating:** Number of IOPS sustained for a specific workload where all array components are within a "Greenzone" providing acceptable performance and response time. The OLTP2HW (Online transaction processing release 2-heavy write) workload was used for these configurations and 8 milliseconds was the target response time.
- **Reliability Rating:** Equates to the thin pool failure probability which is a measure in years of the likelihood that a thin pool unavailable or data loss condition will occur due to hard drive failures. Takes into consideration MTBF, MTBR, number of disks in pool, and RAID protection.

12.6.1.1 1: 100TB 3 Bay 768 x 300GB 10K SAS RAID 10

This VSP configuration includes 6 DKU's with a single DKC chassis. The physical footprint for this option is 3 racks and can grow to the larger 5 rack configuration. This 100 TB configuration is targeted for medium capacity requirements.

Specifications

Component	Medium (00, 01, 02)
Number working drives/spares	760/8
DKU x 2.5" (SSD/SAS)	6
Drive Type	300GB 10K SAS
Protection Type	RAID10 2+2
Usable Capacity in TB	100
DKC / DKU Racks	3
DKC Chassis	1
Drive Slots 2.5"	768
Cache in GB	160
BED Pairs	2

FED Ports	64
VSD Pairs	2

Physical Front Layout with 768 drives

Instructions: Edit configuration, then click "OK" to see Weight & Power in the spreadsheet

Frame configuration

Frame 12	Frame 11	Frame 10	Frame 00	Frame 01	Frame 02
☐ DKU-17	☐ DKU-14	☐ DKU-11	☒ DKU-01	☒ DKU-04	☐ DKU-07
☒ 2.5 SFF	☒ 2.5 SFF	☒ 2.5 SFF	☒ 2.5 SFF	☒ 2.5 SFF	☒ 2.5 SFF
☐ 3.5 LFF	☐ 3.5 LFF	☐ 3.5 LFF	☐ 3.5 LFF	☐ 3.5 LFF	☐ 3.5 LFF
☐ 5.5 NFF	☐ 5.5 NFF	☐ 5.5 NFF	☐ 5.5 NFF	☐ 5.5 NFF	☐ 5.5 NFF
☐ DKU-16	☐ DKU-13	☐ DKU-10	☒ DKU-00	☒ DKU-03	☐ DKU-06
☒ 2.5 SFF	☒ 2.5 SFF	☒ 2.5 SFF	☒ 2.5 SFF	☒ 2.5 SFF	☒ 2.5 SFF
☐ 3.5 LFF	☐ 3.5 LFF	☐ 3.5 LFF	☐ 3.5 LFF	☐ 3.5 LFF	☐ 3.5 LFF
☐ 5.5 NFF	☐ 5.5 NFF	☐ 5.5 NFF	☐ 5.5 NFF	☐ 5.5 NFF	☐ 5.5 NFF
☐ DKU-15	☐ DKU-12	☐ DKC 1	☐ DKC 0	☒ DKU-02	☒ DKU-05
☒ 2.5 SFF	☒ 2.5 SFF	☐ See below	☐ See below	☒ 2.5 SFF	☒ 2.5 SFF
☐ 3.5 LFF	☐ 3.5 LFF			☐ 3.5 LFF	☐ 3.5 LFF
☐ 5.5 NFF	☐ 5.5 NFF			☐ 5.5 NFF	☐ 5.5 NFF

6 SFF DKUs, max 768 SFF drives, 0 LFF DKUs, max 0 LFF drives, 0 NFF DKUs, max 0 FMDs, max total 768 drives.

Power Cord Type
☒ USA 1 phase ☐ EMEA/APAC 1 phase ☐ China 1 phase
☐ USA 3 phase ☐ EMEA/APAC 3 phase ☐ China 3 phase

Disk Drives

Drive type	Number of 4-drive groups	# of spares
SAS SFF 10K 300 GB	190 x 4 = 760 drives	8 drives
Total raw capacity 230,400 GB Raw capacity 228,000 GB Raw 2,400 GB		
-	-	-
-	-	-
-	-	-
-	-	-

Drives: 768 SFF, 0 LFF, 0 SSD, 0 FMD, 0 15K, 768 10K, 0 7200, total 768 drives.
Raw GB: SSD/FMD=0 GB, 15K=0 GB, 10K=230,400 GB, 7200=0 GB

(Disk drive power & heat are for typical activity)

12.6.1.2 2: 200TB 5 Bay 1536 x 300GB 10K SAS RAID 10

This VSP configuration includes 12 DKU's with dual DKC chassis. The physical footprint for this option is 5 racks and should not grow any larger. This 200TB configuration is targeted for large capacity requirements.

Specifications

Component	Large (11,10,00,01,02)
Number working drives/spares	1528/8
DKU x 2.5" (SSD/SAS)	12
Drive Type	300GB 10K SAS
Protection Type	RAID10 2+2
Usable Capacity in TB	200
DKC / DKU Racks	5
DKC Chassis	2
Drive Slots 2.5"	1536
Cache in GB	320
BED Pairs	4
FED Ports	128
VSD Pairs	4

Physical Front Layout with 1536 drives

Instructions: Edit configuration, then click "OK" to see Weight & Power in the spreadsheet

Frame configuration

Frame 12	Frame 11	Frame 10	Frame 00	Frame 01	Frame 02
☐ DKU-17 2.5 SFF 3.5 LFF 5.5 NFF	☒ DKU-14 2.5 SFF 3.5 LFF 5.5 NFF	☒ DKU-11 2.5 SFF 3.5 LFF 5.5 NFF	☒ DKU-01 2.5 SFF 3.5 LFF 5.5 NFF	☒ DKU-04 2.5 SFF 3.5 LFF 5.5 NFF	☐ DKU-07 2.5 SFF 3.5 LFF 5.5 NFF
☐ DKU-16 2.5 SFF 3.5 LFF 5.5 NFF	☒ DKU-13 2.5 SFF 3.5 LFF 5.5 NFF	☒ DKU-10 2.5 SFF 3.5 LFF 5.5 NFF	☒ DKU-00 2.5 SFF 3.5 LFF 5.5 NFF	☒ DKU-03 2.5 SFF 3.5 LFF 5.5 NFF	☒ DKU-06 2.5 SFF 3.5 LFF 5.5 NFF
☐ DKU-15 2.5 SFF 3.5 LFF 5.5 NFF	☒ DKU-12 2.5 SFF 3.5 LFF 5.5 NFF	☒ DKC 1 See below	☐ DKC-0 See below	☒ DKU-02 2.5 SFF 3.5 LFF 5.5 NFF	☒ DKU-05 2.5 SFF 3.5 LFF 5.5 NFF

12 SFF DKUs, max 1536 SFF drives,
0 LFF DKUs, max 0 LFF drives,
0 NFF DKUs, max 0 FMDs,
max total 1536 drives.

Power Cord Type
☒ USA 1 phase
☐ USA 3 phase
☐ EMEA/APAC 1 phase
☐ EMEA/APAC 3 phase
☐ China 1 phase
☐ China 3 phase

Disk Drives

Drive type	Number of 4-drive groups	# of spares
SAS SFF 10K 300 GB	382 x 4 = 1528 drives	8 drives
Total raw capacity 460,800 GB Raw capacity 458,400 GB Raw 2,400 GB		
SAS SFF 15K 146 GB	-	-
SAS SFF 10K 600 GB	-	-
-	-	-
-	-	-

Drives: 1536 SFF, 0 LFF, 0 SSD, 0 FMD, 0 15K, 1536 10K, 0 7200, total 1536 drives.
 Raw GB: SSD/FMD=0 GB, 15K=0 GB, 10K=460,800 GB, 7200=0 GB
 (Disk drive power & heat are for typical activity)

These models maximize the slot and floor space available to the array. Adherence to these models is recommended.

12.6.2 HUS-VM Single Tier Non Replicated Arrays for Production

There are two production configuration options of the HUS-VM used at Citi that do not use HDT (Dynamic Tiering). These models use the 2.5" slots exclusively (SAS only). These RAID 10 options are considered a 'Silver' tier offering.

Summary of Configurations

Config #	Usable Capacity (TB)	Configuration Objective	Number of working disks/spares	RAID	Latency Response Rating (IOPS)	Reliability Rating (yrs)	Cache Per Chassis/Total	Number of Bays (include system bay)	IOPS Per uTB
HUS-VM - 600GB 10K SAS 2.5" - Single Tier Non Replicated									
1	50	Small	192/2	R10 2+2	24,500	489,237	128GB/256GB	1	488
2	100	Medium	384/2	R10 2+2	49,000	244,618	128GB/256GB	1	488

12.6.2.1 1: 50TB Single Rack 194 x 600GB 10K SAS RAID 10

This HUS VM configuration includes 9 SFF drive boxes with a single controller. The physical footprint for this option is 1 rack and can grow to the medium configuration while maintaining the same floor space. This 50 TB configuration is targeted for in country deployments and small capacity requirements.

Specifications

Component	Small (00)
DKC/DKU Racks	1
HUS-VM Controllers	1
8Gb FC Ports	16

Cache in GB	128
BE I/O Module	4
24 Slot SFF Drive Box	9
Drive Slots	216
Number working drives/spares	192/2
Drive Type	600GB 10K SAS Drives
Protection Type	RAID10 2+2
Usable Capacity in TB	50

12.6.2.2 2: 100TB Single Rack 386 x 600GB 10K SAS RAID 10

The 100TB configuration includes 17 SFF drive boxes with a single controller. The physical footprint for this option is a single rack which provides floor and power advantages. This configuration option is targeted for in country deployments and medium capacity requirements.

Specifications

Component	Medium (00)
DKC/DKU Racks	1
HUS-VM Controllers	1
8Gb FC Ports	32
Cache in GB	128
BE I/O Module	4
24 Slot SFF Drive Box	17
Drive Slots	408
Number working drives/spares	384/2
Drive Type	600GB 10K SAS Drives
Protection Type	RAID10 2+2
Usable Capacity in TB	100

12.6.3 HUS-VM Single Tier Arrays for Development and Archive

There is one development model of the HUS-VM used at Citi . This option uses the 2.5" slots exclusively (SAS only).

Summary of Configurations

Config #	Usable Capacity (TB)	Configuration Objective	Number of working disks/spares	RAID	Latency Response Rating (IOPS)	Reliability Rating (yrs)	Cache Per Chassis/Total	Number of Bays (include system bay)	IOPS Per uTB
HUS-VM - 900GB 10K SAS 2.5" - Single Tier Non Replicated									
1	300	Large Dev	432/4	R6 14+2	33,000	271,798,217	256GB/256GB	2	109

12.6.3.1 1: 300TB 2 Rack 436 x 900GB 10K SAS RAID 6

This 300TB configuration includes 19 SFF drive boxes with a single controller. The physical footprint for this option is two racks. The 300TB configuration is targeted for development environments and considered a 'Bronze' tier.

Specifications

Component	Large (00,01)
DKC/DKU Racks	2
HUS-VM Controllers	1
8Gb FC Ports	32
Cache in GB	256
BE I/O Module	4
24 Slot SFF Drive Box	19
Drive Slots	456
Number working drives/spares	432/4
Drive Type	900GB 10K SAS Drives
Protection Type	RAID6 14+2
Usable Capacity in TB	300

12.6.4 Replication Array Types for TC/S and HUR

12.6.4.1 True Copy/Synchronous Configuration Options

TC/S production sites are limited to 50TB or 100TB capacity configurations. The target DR array must be the same model type (VSP or HUS VM), have the same disk profile, and number of remote connections.

Summary of Configurations

Config #	Usable Capacity (TB)	Configuration Objective	Number of working disks/spares	RAID	Latency Response Rating (IOPS)	Reliability Rating (yrs)	Cache Per Chassis /Total	#Bays (include system bay)	IOPS Per uTB
VSP - Single Chassis - 300GB 10K SAS 2.5" - Replicated TC/S									
1	100	Medium - PVOL	760/8	R10 2+2	62,100	216,294	160GB/256GB	3	621
		Medium - SVOL	760/8	R10 2+2	62,100	216,294	160GB/256GB	3	621
HUS-VM - 600GB 10K SAS 2.5" - Single Tier 'Silver' Replicated TC/S									
2	50	Small - PVOL	192/2	R10 2+2	24,500	489,237	128GB/256GB	1	488
		Small - SVOL	192/2	R10	24,500	489,237	128GB/	1	488

				2+2			256GB		
3	100	Medium - PVOL	384/2	R10 2+2	49,000	244,618	128GB/ 256GB	1	488
		Medium - SVOL	384/2	R10 2+2	49,000	244,618	128GB/ 256GB	1	488

12.6.4.2 HUR Configuration Options

HUR production sites are limited to 50TB or 100TB capacity configurations. The DR site requires twice the capacity of the production site for Shadow Image copies. The target DR array must be the same model type (VSP or HUS VM), have the same disk profile, and number of remote connections.

Summary of Configurations

Config #	Usable Capacity (TB)	Configuration Objective	Number of working disks/spares	RAID	Latency Response Rating (IOPS)	Reliability Rating (yrs)	Cache Per Chassis/Total	Number of Bays (include system bay)	IOPS Per uTB
VSP - Single Chassis - 300GB 10K SAS 2.5" - Replicated HUR									
1	100	Medium - PVOL	760/8	R10 2+2	62,100	216,294	160GB/512GB	3	621
2	200	Medium - SVOL 100TB Medium - SI 100TB	1528/8	R10 2+2	124,800	107,581	160GB/512GB	5	624
HUS-VM - 600GB 10K SAS 2.5" - Replicated HUR									
3	50	Small - PVOL	192/2	R10 2+2	24,500	489,237	128GB/256GB	1	488
4	100	Medium - SVOL 50TB Medium - SI 50TB	384/2	R10 2+2	49,000	244,618	128GB/256GB	1	488

* Cache and journal group requirements must be verified based on UR and SI review

12.6.5 HUS-VM All FMD Non Replicated Arrays for Production

There are three all flash production configuration options for the HUS-VM. These models use the DBF drive box with 5.25" slots exclusively (FMD only). These RAID 10 options are considered a 'Platinum' tier offering for high performance application requirements.

Summary of Configurations

Config #	Usable Capacity (TB)	Configuration Objective	Number of working disks/spares	RAID	Latency Response Rating (IOPS)	Reliability Rating (yrs)	Cache Per Chassis/Total	Number of Bays (include system bay)	IOPS Per uTB
HUS VM – 1.6TB FMD – Single Tier 'Platinum' Non Replicated									
1	13.1	Small	16/2	R10 2+2	99,000	5,096,217	128GB/256GB	1	7557
2	26.2	Medium	32/2	R10 2+2	134,400	2,548,108	128GB/256GB	1	5130
HUS VM – 3.2TB FMD – Single Tier 'Platinum' Non Replicated									
3	39.3	Large	24/2	R10 2+2	134,400	1,683,706	128GB/256GB	1	3420

12.6.5.1 1: 13.1TB 1 Rack 16 x 1.6TB FMD RAID 10(2+2)

This HUS VM configuration includes 4 DBFs with a single controller. The physical footprint for this option is a single rack or 13U. This 13.1TB configuration is targeted for small high performance capacity requirements.

Specifications

Component	Small (00)
DKC/DKU Racks	1
HUS-VM Controllers	1
8Gb FC Ports	16
Cache in GB	128
BE I/O Module	4
12 Slot DBF Drive Tray	4
Drive Slots	48
Number working drives/spares	16/2
Drive Type	1.6TB FMD
Protection Type	RAID10 2+2
Usable Capacity in TB	13.1

12.6.5.2 2: 26.2TB 1 Rack 32 x 1.6TB FMD RAID 10(2+2)

This HUS VM configuration includes 4 DBFs with a single controller. The physical footprint for this option is a single rack or 13U. This 26.2TB configuration is targeted for medium high performance capacity requirements.

Specifications

Component	Medium (00)
DKC/DKU Racks	1
HUS-VM Controllers	1
8Gb FC Ports	16
Cache in GB	128
BE I/O Module	4
12 Slot DBF Drive Tray	4
Drive Slots	48
Number working drives/spares	32/2
Drive Type	1.6TB FMD
Protection Type	RAID10 2+2
Usable Capacity in TB	26.2

12.6.5.3 3: 39.3TB 1 Rack 24 x 3.2TB FMD RAID 10(2+2)

This HUS VM configuration includes 4 DBFs with a single controller. The physical footprint for this option is a single rack or 13U. This 39.3TB configuration is targeted for large high performance capacity requirements.

Specifications

Component	Large (00)
DKC/DKU Racks	1
HUS-VM Controllers	1
8Gb FC Ports	16
Cache in GB	128
BE I/O Module	4
12 Slot DBF Drive Tray	4
Drive Slots	48
Number working drives/spares	24/2
Drive Type	3.2TB FMD
Protection Type	RAID10 2+2
Usable Capacity in TB	39.3

12.6.6 Dynamic Tiering Arrays (Arrays using HDT)

Note: The use of HDT has not yet been certified at Citi.

These models maximize the slot and floor space available to the array. Adherence to these models is recommended, but purchases can be made on configurations between the models. There may also be significant variation on the percentage of NL-SAS slots required based on the HDT characteristics and the level of performance tier required.

12.7 Spare Drives

Spares on the VSP and HUS VM are permanent and global. They are no longer based on the size of the back-end. The following charts are a summary of the number of recommended spare drives (2 spare per drive type is the minimum number):

Drive Type	Minimum	>100 Drives	>450 Drives	>600 Drives	>1000 Drives
SSD	2	2	3	4	6
FMD	2	4			
146GB 15k SAS	2	2	3	4	6
300GB, 600GB, 900GB 10k SAS	2	3	5	6	8
All NL-SAS	2	3	5	6	8

Having more spares than required is not a problem. However each set of 8 drives takes away a potential parity group. The spare drives reside in a DKC in the place of an 8-disk parity group.

Optionally, a larger drive can be used as a spare for a smaller drive (assuming the same speed and connectivity type). Therefore, on larger arrays, it may be beneficial to have spares of only a larger drive type allowing for more flexibility. For example, a large array with equal numbers of 300GB and 600GB drives may benefit from having 6x 600GB drives rather than 3 of each drive as spares.

For all FMD configurations the best practice is 1 spare for every 24 FMDs. With the current all FMD configuration options the maximum number of drives are 96 on the HUS VM. For any configuration options less than 24 drives, 2 spares should be available.

Note: if a spare drive is currently being used to rebuild a lost disk, and the lost disk has not totally failed, the rebuild process must complete before the drive can be replaced.

12.8 LUN Expansion

In order to grow the LUNs for a given host there are three methods that should be used with HDS volumes. The exact method to use depends on the business requirements of accommodating the growth.

The three principal methods of LUN expansion are:

12.8.1 Expansion with a New LUN

Allocate a new LUN that is of proper size to accommodate the original LUN plus the required growth of the data. Once the allocation is complete and the host can access the new LUN, then the host will copy the contents of the LUN to the new LUN. Once the data has been migrated to the new LUN, then the old LUN can be destroyed or allocated to another host. This migration typically requires application downtime and may not be appropriate for large or many volumes.

12.8.2 Expansion with Host Volume Management

Allocate a new LUN that is of proper size to accommodate the required growth of data. Once the allocation is complete and the host can access the new LUN, then the host volume manager should be configured to grow the original volume to span (stripped or concatenated) the original LUN and the new LUN. Data loss and application downtime may occur on some operating systems. A backup of the data is recommended.

Depending on the host volume manager this may be a simple task for the host administrator to perform.

Note: Windows based systems will need to use dynamic disk and have a host group type of 2C.

12.8.3 Expansion of Thin Volumes

Thinly provisioned volumes can be expanded in size as they are only a logical reporting of the amount of storage to the host owning the LUN. Assuming the (recent OS) host can support this, expansion of a LUN is an acceptable means of growth. Not all hosts can support this, so caution must be exercised.

12.9 Host Port Balancing and Redundancy

On the VSP (like with USPV), any port can be configured to allow access to any device on the system. Any given host should be connected to at least two host ports (via redundant and independent fabrics). All devices on both storage groups should be identical. The port attributes (and host storage group attributes) should also be the same to ensure proper load balancing from the host.

All host ports should be shared with different OS types. There should not be dedicated host ports for OS types unless there is a need to troubleshoot or isolate a given application (rare).

The array has absolute redundancy across the cluster barrier (Cluster 1 <CL-1> and Cluster 2 <CL-2>). Each cluster has its own power, replicated cache, and internal communication and comprises the overall array redundancy.

Like with previous HDS arrays, The A and B Fabrics should be connected to ports that are in the same position, but on different clusters. For example, Fabric A should be connected to port CL1A and Fabric B should be connected to port CL2A.

Fabric A: Odd Numbered Port (CL# N)
Fabric B: Even Numbered Port (CL# N+1)

There should not be any ports skipped for performance reasons on the VSP array. A wide range of ports should be used only to minimize disruption due to future field replacements.

Spreading the LUNs over each host port is ideal to properly load-balance the ports. There should be less than 256 to 275 LUNs per port.

12.9.1 WWN Numbering for Host Ports

The following is the VSP (and USPV, USP) WWN pattern for external fiber ports:

WWN Component	Value
Prefix (HDS Enterprise Array Prefix + Array Type)	50060e80+(05-USPV or 06-VSP or 13-HUS-VM)
Serial Number (Convert Last Four Digits of SN Decimal to Hex)	x x x x
Port Card Number:	CL# (zero-based, so subtract 1)
Port Card Bank:	A=0

	B=1 C=2 D=3 E=4 F=5 G=6 H=7 J=8 K=9 L=A M=B
--	---

For example, a VSP with a serial number of 53009 (CF11 Hex) would have a WWN of 50060e8006**cf1130** for port CL-4A and a WWN of 50060e8006**cf1123** for port CL-3D.

A HUS-VM with a serial number 210060 (Last 4 digits 10060 = 274C Hex) would have a WWN of 50060e8013**274c00** for port CL-1A and a WWN of 50060e8013**274c71** for port CL-8B.

12.9.2 VSP Port Options and Slot Names

When using the 16-port fiber “option” of 8Gbps ports (16UFC FCP), there will be two cards of 8 ports each in the initial slots 1EU and 2QU. Subsequent pairs of card are added to the system named by options 1-5 on DKU0.

DKC	Option Name	Cluster 1		Cluster 2	
		CHA:	Port Name	CHA:	Port Name
0	Basic	1EU:	1A, 3A, 5A, 7A, 1B, 3B, 5B, 7B	2QU:	2A, 4A, 6A, 8A, 2B, 4B, 6B, 8B
0	Option 1	1EL:	1C, 3C, 5C, 7C, 1D, 3D, 5D, 7D	2QL:	2C, 4C, 6C, 8C, 2D, 4D, 6D, 8D
0	Option 2	1FU:	1E, 3E, 5E, 7F, 1F, 3F, 5F, 7F	2RU:	2E, 4E, 6E, 8F, 2F, 4F, 6F, 8F
0	Option 3	1FL:	1G, 3G, 5G, 7G, 1H, 3H, 5H, 7H	2RL:	2G, 4G, 6G, 8G, 2H, 4H, 6H, 8H
0	Option 4	1AL:	1L, 3L, 5L, 7L, 1M, 3M, 5M, 7M	2ML:	2L, 4L, 6L, 8L, 2M, 4M, 6M, 8M
0	Option 5	1AU:	1J, 3J, 5J, 7J, 1K, 3K, 5K, 7K	2MU:	2J, 4J, 6J, 8J, 2K, 4K, 6K, 8K
1	Basic	1GU:	9A,BA, DA, FA, 9B, BB, DB, FB	2TU:	AA,CA, EA, GA, AB, CB, EB, GB
1	Option 1	1GL:	9C,BC, DC, FC, 9D, BD, DD, FD	2TL:	AC,CC, EC, GC, AD, CD, ED, GD
1	Option 2	1HU:	9E,BE, DE, FE, 9F, BF, DF, FF	2UU:	AE,CE, EE, GE, AF, CF, EF, GF
1	Option 3	1HL:	9G,BG, DG, FG, 9H, BH, DH, FH	2UL:	AG,CG, EG, GG, AH, CH, EH, GH
1	Option 4	1LL:	9L,BL, DL, FL, 9M, BM, DM, FM	2XL:	AL,CL, EL, GL, AM, CM, EM, GM
1	Option 5	1LU:	9J,BJ, DJ, FJ, 9K, BK, DK, FK	2XU:	AJ,CJ, EJ, GJ, AK, CK, EK, GK


The card names are named according to physical slot on the array starting with 1E on cluster 1 and 2Q on cluster 2. There is an upper slot used first then a lower slot used second for each. Thusly, under the first chassis (or DKC 0), 1EU and then 1EL for Cluster 1. For Cluster 2 it would be 2QU and then 2QL. The next slots used are 1FU/L then 1AL/U (reversed) for Cluster 1 and 2RU/L then 2ML/U (reversed) for Cluster 2. There will also be same basic and options on DKC 1 (second chassis).

12.9.1 VSP Host Ports and Replication Ports

When planning for non-replicated arrays, all SAN connected and available ports should be used for hosts.

When planning for external replication, there will be 2x to 8x Initiator ports and 2x to 8x RCU Target ports. These ports will be last port for on each option 7X and 8X for each port series. For example, ports 7A/8A and 7B/8B would be used for replication, leaving 1A/2A,3A/4A,5A/6A,1B/2B,3B/4B,5B/6B for host use. On Dual DKC machines, all ports on the secondary DKC are used for hosts.

Single DKU with Option 1




→ RCU Target Ports (Min 2, Max 8 per VSP)


→ Initiator Ports (Min 2, Max 8 per VSP)

→ Host Ports

Single DKC with Option 3



Dual DKC with Option 3




12.9.2 HUS-VM Port Options and Slot Names

The HUS-VM array uses 4-port cards instead of the 8-port cards that are used in the VSP. In addition to having fewer slots for host port cards, this lowers the total number of host ports on the array significantly.

Option Name	Cluster 1		Cluster 2	
	CHB:	Port Name	CHB:	Port Name
Basic	1A:	1A, 3A, 5A, 7A	2A:	2A, 4A, 6A, 8A
Option 1	1B:	1B, 3B, 5B, 7B	2B:	2B, 4B, 6B, 8B
Option 2	1C:	1C, 3C, 5C, 7C	2C:	2C, 4C, 6C, 8C
Option 3	1D:	1D, 3D, 5D, 7D	2D:	2D, 4D, 6D, 8D

The card names (“CHB” rather than “CHA” on VSP) are named according to physical slot on the array starting with 1A on cluster 1 and 2A on cluster 2.



12.9.3 HUS-VM Host Ports and Replication Ports

When planning for non-replicated arrays, all SAN connected and available ports should be used for hosts. When utilizing HUR or TC/S a minimum of 2 Initiator and 2 RCU Target ports must be configured for remote replication. Refer to HDS VSP/HUS VM Remote Replication Solutions Guide for specific details on configuring remote connections.

12.10 Topology

The VSP should be connected to the core directors from each redundant fabric. There should be at least 8x 2Gb/4Gb/8Gb (more if only 2Gb) and up to 64x (for VSP, 16x for HUS-VM) fiber connections from each of the fabrics. Additional ports may be connected if these ports become overloaded or to spread out the storage over more ports. Ports on the fabric are generally more restricted than additional ports on the array.

12.11 Maximum Hosts Per Port

HDS has a concept called **T**ags that it uses to describe **Q**ueue depth x **S**ervers x **L**UNs on a port. They limit the recommended Tags for a given port to prevent excessive IO on any single port.

$$T = Q * S * L$$

The HDS VSP Tag recommendation is 2048 per port. As a comparison, the HDS USPV Tag recommendation is 4096 per MP (shared between certain ports)

$$S = T / (Q * L)$$

If our average Queue Depth is 32, and the average number of LUNs is 2 (per host), then the maximum number of hosts per port would be

$$S = 2048 / (32 * 2) = 32$$

Note: This is a typical assumption for most data centers.

If our average Queue Depth is 32, and the average number of LUNs is 4 (per host), then the maximum number of hosts per port would be

$$S = 2048 / (32 * 4) = 16$$

12.12 Maximum Ports Per Host

In order to maximize pre-fetch algorithms (esp. for sequential reads), the default number of array host ports a given host is zoned to is 4. Going beyond this is by exception only.

12.13 Management Software (HCS vs. SN)

Storage Navigator is designed to manage one array (VSP or USPV Series) directly. If there are other types of HDS arrays to be managed or if multiple arrays need a common management interface, then Hitachi Command Suite (HCS) will be needed. In addition, if command line functionality is desired, then the HDvM environment along with the CLI and CCI/Raidcom components (if appropriate) will need to be installed on multiple systems. At Citi we limit the server to be installed on Windows. The CCI and CLI are best suited for UNIX (RH Linux or Solaris).

Note: To keep the HCS and SN configurations synchronized, the best way to launch SN is from HCS (using the Element Manager). In addition, this prevents the need for local accounts on the array.

HCS is required for most third-party toolsets to monitor, report, or manage the VSP. This includes NetApp OnCommand Insight (formerly SANScreen) and EMC Control Center (ProSphere).

The VSP array requires the use of HCS version 7.0 or later. THE HUS-VM array requires use of HCS 7.4 or later. The USPV array does not have a version restriction.

12.14 SAN Monitoring

NetApp OnCommand Insight (formerly SANScreen) should be the primary SAN monitoring tool with end-to-end path change monitoring. OCI requires the use of HCS.

12.15 Array Monitoring: Alerts and Remote Connections

The array will be monitored for status and component failure via a dial-out process directly to HDS. If the vendor needs to remotely connect to the array, this will be accomplished via a VPN-based connection process.

12.15.1 SNMP

The USPV and VSP series arrays support SNMP event broadcasting. SNMP alerts can be configured to send traps to a HCS server and SNMP Trap Servers (ESS). For more information see [Enterprise Alerting \(SNMP\)](#).


12.15.2 Vendor Notification (Dial-Out / HiTrack)

The Hitachi HiTrack remote maintenance application will periodically send configuration and status information to the HDS/Sun support center(s) including SIM alerts. The HiTrack software uses HTTPS via a webproxy to send alerts and configuration information to HDS. For more information see [Appendix K: Setup of HiTrack Service for HTTPS](#).

12.15.3 HDS Remote Access Application (RACC)

As of 2014, RACC will reside on chosen SVPs. Generally it will be configured to run on at least 2x arrays per major data center by HDS staff for redundancy. RACC uses the application proxy for firewall access. Any given RACC can access any HDS array at that site (within the same CEP), but cannot access HDS arrays at other sites.


Citi staff will be responsible for establishing firewall rules for the RACC SVP to access HDS Support sites via the HTTP Application Proxy. On these SVPs, TUF will also be a destination due to RACC using IE for firewall configurations. Non-RACC SVPs should also use the application proxy for consistency and to avoid confusion. This provides for flexibility on moving the RACC software during array retirements.



12.16 Network Time Protocol (SNTP or NTP) – Time Synchronization

In order to keep the time reported in the logs on the array accurate and consistent, it is advised that the array be configured for network-based time synchronization. This is an

HDS CE function at this time, though Citi SNTP servers should be used. With the current microcode, only standard time (not daylight savings time) is supported. It may be beneficial to set all arrays to the same time zone as where the logs are sent or to use a universal time zone such as GMT. There is a Time Settings screen on the SVP application (Modify Mode → Install → Time Settings → Time Synchronization). The following is the screen that the CE uses to configure SNTP (or NTP) on the array:



The IP address and port of the SNTP server are required. The Time Zone and synchronization check time are selectable.

Note: Because the SVPs are behind a firewall, there must be a port (default is 123) opened between the primary SVP and the SNTP server. See [Firewall Ports](#) for more information.

12.17 Multiple Array Management

All of the arrays should be fully managed via the HDvM instance for each major data facility. If Storage Navigator is needed then the element manager function can be selected for a given array.

As a backup method, desktop icons/shortcuts/bookmarks (saved for each primary SVP) could be used to connect to each array. This would allow the user to run multiple copies of Storage Navigator at the same time. Desktop icons/shortcuts could be copied from one management server to another on a manual or automated basis.

12.18 Polling Intervals for Device Manager

To prevent too frequent polling of the array from the Device Manager server(s), it is recommended that the polling interval be reduced from once every 5 minutes to once every 60 minutes. This can be accomplished by changing the following parameter file:

```
C:\Program Files\HiCommand\DeviceManager\HiCommandServer\config\dispatcher.properties
# polling interval (minutes) for Component status and ConfigurationVersion daemons
# (a value of zero disables these polling agents)
server.dispatcher.daemon.pollingPeriod=60
```

To make this change effective, a restart of the Device Manager service is required.

12.19 Enterprise Alerting (SNMP)

The array should be configured such that SNMP traps are being sent to an Enterprise Management System. The Simple Network Management Protocol (SNMP) environment can be configured to send traps to more than one trap receiver. For example, an instance of HDvM and the ESS system can be configured as trap servers such that all traps are sent to both environments.

To configure SNMP traps, go to environmental settings → SNMP. This screen has the legacy interface that features the modify mode button.

Once in modify mode, add the new destination ip address to the Manager and *Community & Trap* window. Once the address has been entered, select the Set button to add it to the manager list. The new ip address should be in blue italics. If a new community string is to be defined (such as public community – this is common) then right-click the community listing and select add community. Once public (or any other community) is listed, right-click the community name to add an ip-address to that community. The community can be renamed (right click: modify). Under the system group, the name of the array should be modified to include the serial number or other proper array identification.

Note: the SNMP trap destination should be listed in both the *Manager* and *Community & Trap* lists to enable two-way messaging and queries and may be needed for future releases of Tivoli. The Community & Trap list is needed for outbound (trap only) messaging.


Click Apply to make the changes effective.

The community name of Public is discouraged by the ESS (SNMP) team at Citi as it is unsecure. However, there are no known implementations of Tivoli that accept other community names at this time.

Use the Trap Test button to send a test trap like the following:

Date	Time	Priority	Hostname	Message
08-30-2011	18:30:01	Local7.Debug	10.40.42.60	community=Public, enterprise=1.3.6.1.4.1.116.3.11.4.1.1, uptime=103638720, agent_ip=10.40.53.13, generic_num=6, specificTrap_num=4, specificTrap_name=eventTrapPartsID, version=Ver1, 1.3.6.1.4.1.116.5.11.4.2.1=53112, 1.3.6.1.4.1.116.5.11.4.2.2=RAID700, 1.3.6.1.4.1.116.5.11.4.2.3=7FFFFFFF, 1.3.6.1.4.1.116.5.11.4.2.4=1.3.6.1.4.1.116.5.11.4.1.1.6.1.9, 1.3.6.1.4.1.116.5.11.4.2.5=2011/08/30, 1.3.6.1.4.1.116.5.11.4.2.6=17:30:02, 1.3.6.1.4.1.116.5.11.4.2.7="This is a test code."

You will get a confirmation of the sent trap when using the test button:



The SNMP trap ports (161, 162 UDP) must also be configured if the SVP is behind a CEP/firewall. For more information on firewall settings see [Appendix F: Firewall Ports](#).

For a detailed listing of all of the defined traps in the USP/USPV Mib definition see [Appendix M: SNMP Mib Definitions](#).

The information "1.3.6.1.4.1.116.5.11.4.2.3" provides a reference to SIM events in the array.

The information “1.3.6.1.4.1.116.5.11.4.2.7” (if provided) describes the reason for the event very briefly.

The information “1.3.6.1.4.1.116.5.11.4.2.2” provides for the **array** type being USP (RAID500) or USPV (RAID600) or VSP (RAID700).

The information “1.3.6.1.4.1.116.5.11.4.2.1” provides the serial number of the array.

The information “1.3.6.1.4.1.116.5.11.4.2.5” provides the date of the event.

The information “1.3.6.1.4.1.116.5.11.4.2.6” provides the time of the event.

The information “1.3.6.1.4.1.116.5.11.4.2.7” (if provided) is a description of the event.

12.20 Running HCS on VMWare

VMWare is a supported platform to run the HDS server and client software. VMWare with Windows 2003/2008 hosts was tested with the following software:

HDvM (Device Manager) 7.x

HTnM (Tuning Manager) 7.x

HTSM (Tiered Storage Manager) 7.x

Client software tested:

Storage Navigator

HDvM (with HTSM ,HTnM) GUI Clients

HDvM CLI

VMware can use the command devices (as SAN attached LUNs). This is required for Tuning Manager Agent for RAID. However, many VMWare Clusters are not SAN capable and it may be simpler to run SAN attached systems on different servers.

12.21 HCS Management Server Sizing

When combining the servers (HDvM, HTnM, HTSM , HRpM) to run on a single instance of Windows 2003, the following sizing guideline should be used:

Data Center with 15x LARGE VSP Frames (with all 300GB or larger drives):

Four Quad-Core Intel® Xeon™ X7350 Processor (2.93 GHz)

8 GB of Memory

40GB Disk space

12.22 Running Multiple Versions of Server Software Simultaneously

It may be necessary to run two versions of server software at the same time esp. with the deployment of new versions. For example, if new features are required, but ECC only works with the previous version, then running two instances may be the only viable solution for a period of time. HDS and CSBE support running two versions of HDvM only if once instance is used for making changes. The other version must be used only for read-only purposes.

12.23 Element Manager (Running SN from HCS)

Running Element Manager from HCS launches an instance of the Storage Navigator (SN) client. This is the preferred way to run Storage Navigator because it synchronizes the HCS database with the changes made under SN. This also allows the array to be configured with no local accounts (other than functional ids).

To be able to run Element Manager (formerly Physical View) successfully, do the following:

- Open port 2001 from the SVP to the HDvM server (different direction than most ports). For more information on all of the ports see [Appendix F: Firewall Ports](#).
- Add the SVPs to the ACL list for HDvM. For more information see [Limiting Client IPs to Known Management Systems \(ACL for HDvM\)](#).

12.24 Backup, Restore, and Migration of Device Manager and Related Databases

12.24.1 Backup HDvM Databases

In order to backup the HDvM and Common Services Databases, the following command should be issued:

```
C:\Program Files\HiCommand\Base\bin> hcmdsbackups /dir c:\hdvm_bkup01 /auto
```

```
KAPM05309-I Backup of CommonComponent started.
```

```
KAPM05328-I The hcmdsbasebackup processing has succeeded.
```

```
KAPM05313-I Backup of CommonComponent successfully ended.
```

```
KAPM05309-I Backup of Database started.
```

```
KAPM06902-I Processing to stop the service has started.
```

```
KAPM06906-I Processing to stop the DBMS has started.
```

```
KAPM06907-I The DBMS was stopped successfully.
```

```
KAPM06903-I The service was stopped successfully.
```

```
KAPM06904-I Processing to start the DBMS has started.
```

```
KAPM06905-I The DBMS was started successfully.
```

```
KAPM05530-I Backup of the database has started.
```

```
KAPM05531-I Backup of the database has ended.
```

```
KAPM06900-I Processing to start the service has started.
```

```
KAPM06904-I Processing to start the DBMS has started.
```

```
KAPM06905-I The DBMS was started successfully.
```

```
KAPM06901-I The service was started successfully.
```

```
KAPM05313-I Backup of Database successfully ended.
```

```
KAPM05320-I Backup successfully ended : CommonComponent Database
```

Notes: /auto will automatically stop and start required services for backup
 /dir <directory> is an empty directory to place the backup dump file
 If the export fails due to service stopping issues, try running the command again.

12.24.2 Restore HDvM Databases

In order to restore the entire HDvM database and Common Services, the following command should be issued:

```
C:\Program Files\HiCommand\Base\bin> hcmdsdb /restore c:\hdvm_bkup01\database\backup.hdp /type ALL  

/auto
```

Note: /auto will automatically stop and start required services for restore

12.24.3 Restrictions to Backup and Restores

The following are restrictions to backup and restores:

- The backup will halt the database and will restart once completed.
- Backups can only be restored on the same installation paths and version of software.
- Backups can only be restored on the same servername and ip address.

- No files or subdirectories can be in the /dir directory or the backup will fail.

12.24.4 Migrations (Export and Import of HDvM databases)

Migrations are more flexible than backup/restores. Migrations support:

- Destination can be a **newer** (or the same) version software than source
- Destination can have different directory paths
- Destination can be a different operating system (like Solaris)

To export the database run the following command:

```
C:\Program Files\HiCommand\Base\bin>hcmsdbtrans /export /workpath C:\temp\trans_work /file
c:\hdvm_trans\db_arc /auto
```

```
KAPM05900-I The hcmsdbtrans command has started.
KAPM06902-I Processing to stop the service has started.
KAPM06906-I Processing to stop the DBMS has started.
KAPM06907-I The DBMS was stopped successfully.
KAPM06903-I The service was stopped successfully.
KAPM06904-I Processing to start the DBMS has started.
KAPM06905-I The DBMS was started successfully.
KAPM05930-I The Application data will now be exported.
KAPM06525-I The data for DeviceManager will now be moved.
KAPM06531-I Processing to export data has started.
KAPM06532-I The table definitions of the database are being exported.
KAPM06533-I The data is being exported.
KAPM06537-I The data is being exported. (progress = 01/85)
KAPM06537-I The data is being exported. (progress = 02/85)
...
KAPM06537-I The data is being exported. (progress = 160/161)
KAPM06537-I The data is being exported. (progress = 161/161)
KAPM05931-I The Application data has been exported.
KAPM05930-I The Authentication data will now be exported.
KAPM05821-I Processing to export data has started.
KAPM05822-I Processing to export data has ended.
KAPM05931-I The Authentication data has been exported.
KAPM05933-I The archive file will now be created.
KAPM05934-I The archive file has been created.
KAPM06900-I Processing to start the service has started.
KAPM06904-I Processing to start the DBMS has started.
KAPM06905-I The DBMS was started successfully.
KAPM06901-I The service was started successfully.
KAPM05901-I The hcmsdbtrans command ended normally.
```

Each application database will be exported. The exports will take 10-15 minutes or longer depending on the amount of data. Due to the length of output, you can also redirect the output to a logfile (with no screen output) by adding "> output.log" to the end of the command string.

To import the database (on another system) do the following:

1. Run the command:

```
C:\Program Files\HiCommand\Base\bin>hcmsdbtrans /import /workpath C:\temp\trans
_exp /file C:\hdvm_exp\db_arc /type DeviceManager /auto
```

```
KAPM05900-I The hcmsdbtrans command has started.
KAPM05935-I The archive file will now be extracted.
KAPM05936-I The archive file has been extracted.
KAPM06902-I Processing to stop the service has started.
KAPM06906-I Processing to stop the DBMS has started.
KAPM06907-I The DBMS was stopped successfully.
KAPM06903-I The service was stopped successfully.
KAPM06904-I Processing to start the DBMS has started.
KAPM06905-I The DBMS was started successfully.
KAPM05937-I The Application data will now be imported.
```

```

KAPM06525-I The data for DeviceManager will now be moved.
KAPM06526-I Processing to import data has started.
KAPM06527-I The table definitions of the database are being imported.
KAPM06528-I The data is being imported.
KAPM06536-I The data is being imported. (progress = 01/85)
KAPM06536-I The data is being imported. (progress = 02/85)
...
KAPM06536-I The data is being imported. (progress = 84/85)
KAPM06536-I The data is being imported. (progress = 85/85)
KAPM06540-I The view is being imported.
KAPM06530-I Processing to import data has ended.
KAPM05938-I The Application data has been imported.
KAPM05937-I The Authentication data will now be imported.
KAPM05819-I Processing to import data has started.
KAPM05820-I Processing to import data has ended.
KAPM05938-I The Authentication data has been imported.
KAPM06902-I Processing to stop the service has started.
KAPM06906-I Processing to stop the DBMS has started.
KAPM06907-I The DBMS was stopped successfully.
KAPM06903-I The service was stopped successfully.
KAPM05901-I The hcmdsdbtrans command ended normally.

```

Notes: /ALL will import all of the databases
 /DeviceManager will only import Device Manager data
 /workpath <directory> must be empty
 /file -- on the archive file location you must use have a filename at the end of the path
 If the export fails due to service stopping issues, try running the command again.
 Due to the length of output, you can also port the output to a logfile (no screen output) by adding "> output.log" to the end of the command string

2. You will then need to synchronize the data. Specify the following parameter in the server.properties file:

```
server.base.initialsynchro = true
```

Since, other than user information, the hcmdsdbtrans command does not migrate the Common Component repository, you need to synchronize the repository information with the database information of the imported Device Manager. This will take several minutes. Any attempt to login to HDvM during the synchronization will result in error.

3. Start the database with common services. Go to Start → All Programs → Hitachi Command Suite → Device Manager → Start Server with Common Services
4. You can log into HDvM after 5-10 minutes to allow for synchronization. You may have to refresh some data.
5. After several minutes, change the value in server.properties back to its original value:

```
Server.base.initialsynchro = false
```

12.25 Maximum LUN size

The LUN size limit for thin provisioning is now 4TB (may be increased in future). The maximum traditional thick LUN size (assignable a host) under Citi standards remains 277.52 GB

12.26 Increased Heap Size (HDvM Memory Expansion)

In order for HDvM to handle more objects, the following memory changes are necessary:

1. Run the stop_all command on the HDvM server to shut everything down
2. Edit server.ini in "c:\Program Files\HiCommand\DeviceManager\HiCommandServer\
Change server.ini option: JVM_XOPT_HEAP_MAX=-Xmx256m to
JVM_XOPT_HEAP_MAX=-Xmx1024m
3. Edit HiCommandCLI.bat in
"c:\Program Files\HiCommand\DeviceManager\HiCommandCLI\
Change HiCommandCLI.bat option: HDVM_CLI_MEM_SIZE=256m to
HDVM_CLI_MEM_SIZE=1024m
4. Run command "hcmdschgheap /set Medium" in "c:\Program Files\HiCommand\Base\bin"
5. Run the start_all command on the HDvM server to restart all the HDvM and HSCS components

12.27 System Options (SOM)

There are several system option modes (occasionally referred to as SOM) that are needed for the array to function as expected. These options must be configured by HDS staff. HDS can provide a report of these options for every array they receive update data from (dial-home).

The following system options should be enabled on all VSP and HUS-VM systems:

System Option	Value	Description	Notes
729*	ON	This setting limits writes and reads to a full pool if the setting is ON.	Prevents data corruption when pool is full.
803*	ON	This setting limits writes and reads when a pool volume error is encountered if the setting is ON. (Data Retention Utility is enabled.)	Prevents data corruption when pool volume error occurs.
734	ON	Repeat SIMs every 8 hours on pool threshold alerts.	
741	ON	Repeated pool threshold alerts are reported to HDS.	
749	OFF	The pool rebalance functionality is activated.	
87	ON	Mode 80 = ON: Normal Restore (Reverse Copy) is performed for Shadow Image. Mode 80 = OFF (default): Quick Restore is performed for ShadowImage.	
917	ON	Pools are balanced by Parity Group rather than by the LDEVs. This allows for variability in the LDEV size / number from each Parity Group.	

The following system options should be enabled on only replicated (source and destination) VSP and HUS-VM (when certified for replication) systems:

System Option	Value	Description	Notes
448	OFF	When the SVP detects a blocked path and when the path does not recover within a specified period of time, an error is assumed and the remote replication is suspended.	
449	ON	Configures the SVP to ignore all HUR path failure. In the event of link failure, HUR will continue to operate and write all update I/Os to the journal volumes.	

760	OFF	Host I/O can decrease when the UR pair is resynchronized. Set mode 760 to ON to suppress an increase in the microprocessor utilization rate when resynchronizing the pair. Mode 760 ON: Decrease in host I/O response during the UR resync operation is reduced. Mode 760 OFF: Operations run normally.	
-----	-----	---	--

Notes:

*System Option 729 and 803 should be set to prevent updates to already allocated pages when a full pool condition is met or pool volume errors are encountered. Data corruption could occur if update is allowed even though new pages cannot be allocated (default behavior). See [Pool Full Condition](#) for Data Retention Utility recovery procedures.

12.28 Citi Thin Provisioning (HDP) Best Practices

Thin Provisioning (also called HDP) is the default configuration for HDS frames at Citi and is required for new deployments. Thin storage should be used on all applications for significant gains in storage utilization. Thin provisioning, when properly balanced across the array, provides for wide-stripping of data and can improve performance over traditional storage.

12.28.1 Virtual to Physical Capacity (Thin Provisioning Rules)

There are two basic rules that need to be considered before deployment of Thin Provisioning. These rules may change over time as we will continue to shape the thin storage environments.

CSBE Thin Provisioning Rule #1: The maximum supported amount of virtual storage that CSBE supports is no more than 100% of the physical capacity of the array. Therefore, if all applications were at 100% utilization, the array will need to be fully expanded, but could handle the physical representation. This value may change in the future as Citi's experience in thin provisioning grows over time. Higher values are allowed only by exceptional approval from CSBE.

CSBE Thin Provisioning Rule #2: The maximum supported ratio of current allocated virtual storage to current physical pool storage is 2x. This value is very likely to change in the near future as Citi's experience in thin provisioning grows over time. Higher values are allowed only by exceptional approval from CSBE.

Example scenario:

If there is 100TB in a pool, then no more than 200TB should be allocated to hosts assuming that there is potentially 100TB in additional open drive slots based on the current pool components (drive type/config).

12.28.2 Pool Composition

- Each pool should consist of 1x LDEVs from each parity group, encompassing all of the storage from the parity group (dedicated Parity Groups). Older pools may have been created from 4 LDEVs per parity group.

LDEV size = Total usable storage from Parity Group

- High capacity drives may require an additional LDEV for capacity that exceeds the max LDEV size. The parity group should be evenly split between two LDEVs in this case.

- There should be 4-8 Parity Groups (each from different BED pairs) to maximize performance of the pool storage. If there are fewer BEDs on the array, then each pool should have PGs from all of those BEDs.
- Pools must consist of the same drive type/size, drive speed, and raid configuration.
- There should be 1-2 pools per drive type/size to minimize monitoring, reporting, maintenance. You may want to consider different pools for production and development. However, the more pools that are present, the more pools need to be monitored and the fewer resources a given pool will have in terms of performance.
- Pools should be built from PGs that are from each BED(DKA)
- The following are supported maximums:

Note: The max LUN size for Citi approved configurations is 4TB

Limits	VSP Single Chassis	VSP Dual Chassis	HUS-VM	USPV
HDP Pools	128	128	128	128
Max Pool Capacity	5 PB	5 PB	2.1 PB	1.1 PB
Max Total Pool Capacity	5 PB	5 PB	2.1 PB	1.1 PB
LDEVs per Pool	1024	1024	1024	1024
Max Pool Volume Size	3.99 TB	3.99 TB	2.99 TB	2.99 TB
DP Volumes per Pool	Approx 62000	Approx 62000	16,384	8192
DP Volume Size (w/o replication)	46 MB – 60 TB	46 MB – 60 TB	46 MB – 60TB	46 MB - 4 TB
DP Volume Size (w/ replication)	46 MB – 4 TB	46 MB – 4 TB	46 MB – 4 TB	46 MB – 4 TB

12.28.3 Pool Naming

We have taken away the LDEV numbering and POOL numbering on the VSP. We no longer name the LDEVs or POOL IDs based on the drive type of storage. You should name the pools according to their composition and role (i.e. R6_6+2_600GB_PROD)

12.28.4 Pool Creation

- Create Pool with 4-8 PGs (based on approximately 1 LDEVs per PG).
- Optimize after all growth is complete – do not optimize at any other time (this will resort free-list)
- Create Virtual Volumes (formerly called DP-Vols / V-Vols).

12.28.5 Pool Growth

- Upgrade should be multiples of original size of pool (n): n, n*2, n*3, but can occur with any number of raid groups as automatic rebalancing will occur over the new and old raid groups. However, the pool should remain balanced over the back-end wherever possible, so the number of raid groups should reflect the number of BEDs in the system.
- Upgrade should be the same drive type / speed / RAID configuration as the existing pool
- Dedicated large HDP arrays should have all pools spread over all BED/DKAs

12.28.6 Migrations

- Host-based file transfer (LUN to LUN) should be used in order to take advantage of the thin storage savings.
- Shadow Image and UVM copies (or other block level copies) will not take advantage of HDP
- ShadowImage and TrueCopy/HUR copies of HDP-Volumes are supported

12.28.7 DP-VOL Sizing

The standard volume sizes below can be used on LDEVs from DP-Vols (V-Vols). This should make any migrations easier as it is the same size as other volumes on the array (or other arrays).

Standard LUN Sizes				
Cyl	Blocks (512 byte)	Bytes	GB	Unit
1092	2096640	1073479680	1.000	
36828	70709760	36203397120	33.717	1
110484	212129280	1.0861E+11	101.151	3
294624	565678080	2.89627E+11	269.736	8
589248	1131356160	5.79254E+11	539.473	
1473120	2828390400	1.44814E+12	1348.682	
4202688	8069160960	4.13141E+12	3847.676	Reserved for future

If business needs demand custom volumes sizes, then custom DP-Vols are very easy and fast to create. The maximum size for a DP-VOL (V-VOL) under current Citi standards is 4TB. For more information on this restriction, see [Maximum LUN size](#).

12.28.8 Thresholds

User definable (Pool Threshold 1)

- Large Pools: 70% (> 10TB)
- Small Pools: 50% (=< 10TB)

Fixed (Pool Threshold 2) – 80%

12.28.9 HDP Alerts

12.28.9.1 HDP Alerts via SNMP Traps

The standard SNMP interface from the SVP will send out traps for all system events (SIMs). Thresholds are SIM events. The following is a sample SNMP Trap generated by exceeding the user definable threshold (Threshold 1) for a pool. The Pool ID is imbedded in the SIM code as a hexadecimal value (i.e. 620001 – error 620 on pool ID 001 Hex).

Date	Time	Priority	Hostname	Message
04-22-2008	17:01:42	Local7.Debug	10.40.53.25	community=public, enterprise=1.3.6.1.4.1.116.3.11.4.1.1, uptime=173625422, agent_ip=10.40.53.25, generic_num=6, specificTrap_num=3, specificTrap_name=eventTrapREFCODE, version=Ver1, 1.3.6.1.4.1.116.5.11.4.2.1=10102, 1.3.6.1.4.1.116.5.11.4.2.2=RAID600, 1.3.6.1.4.1.116.5.11.4.2.3=620001,

				1.3.6.1.4.1.116.5.11.4.2.4=1.3.6.1.4.1.116.5.11.4.1.1.6.1.9, 1.3.6.1.4.1.116.5.11.4.2.5=2008/04/22, 1.3.6.1.4.1.116.5.11.4.2.6=17:04:17, 1.3.6.1.4.1.116.5.11.4.2.7="The DP POOL threshold 1 was exceeded"
--	--	--	--	---

In a typical Citi Tivoli or similar alert environment, you may not see the full message above. You are more likely to see a message like the following:

"On 1240104394 a WARNING message occurred:

CCC/OT/GCTI/INF/SAN/SNMP P SWDC ouh028n8667-spa UNKOS 420
raideventUsermoderate RAID600 620032 The impact of this event on the
subsystem is moderate."

This is only a summary of the trap information in the chart above, but it does contain the key points.

CCC/OT/GCTI/INF/SAN/SNMP P SWDC	Site= SWDC
ouh028n8667-spa	Frame= ouh028n8667 (primary SVP)
620032	Error 620 on Pool ID 032 Hex (50 Dec)

The impact of this event on the subsystem is moderate

Error 620 is "threshold 1 (user-threshold) has been exceeded. For interpretation of the error number see [HDP Alerts via SIM on array](#) below.

For more information on setting up SNMP see [Enterprise Alerting \(SNMP\)](#).

12.28.9.2 HDP Alerts via Email (via HDvM)

Using email for alerts is less reliable than SNMP traps. However, the email alerts are generated by HDvM when it receives a trap message from the USPv array. Therefore, the email alerting can be used as a secondary (or backup) means of alerting. The following is a sample email message from HDvM:

From: GDSE@citigroup.com
Subject: [DVM] Alert Notification

The following alert occurred.

MessageID: KAID10008-W
Alert Type: Trap
Source: USP_V@10.40.53.25
Status: Moderate
Component: DKC Environment
Description: Pool error DP Pool utilization threshold excess 1 Pool number 001(1).
Recommended Action: Contact your service representative.
Additional Info: 620001
Occurrence Time: 2008/05/01 15:56:11


This message was sent automatically by the Device Manager server.

The source field listed in the sample email is the array name as configured in Device Manager. You should change this description to incorporate the serial number or local array name.

It is important to understand that setting up SNMP traps is important for the email functionality to work. Device Manager needs to be one of the SNMP trap designations in order for the email notification to function properly and to deliver all of the needed information. For more information on setting up SNMP see [Enterprise Alerting \(SNMP\)](#).

12.28.9.3 HDP Alerts via SIM on array

All SIMs are reported to HDS via the HTTPS dial-home HiTrack interface. SIMs can be viewed in Storage Navigator under status. The following is an example SIM for exceeding the user-defined threshold (Threshold 1) for a pool:



The following chart is a summary of the SIM codes generated by the array relative to HDP:

SIM (Ref) Code	Error Condition	Error Level
620xxx	Pool exceeded threshold 1 (user-defined)	Warning
621xxx	Pool exceeded threshold 2 (fixed 80%)	Warning
622xxx	Pool is full	Critical
623xxx	Error occurred in pool	Critical
624xxx	No space in the shared memory	Critical
625xxx	Pool usage level continues to exceed pool thresholds (repeated every 8 hours)	Critical
630xxx	V-Vol threshold exceeded	Warning
640xxx	The V-Vol management area cannot be saved to a pool	Critical


*Note: The xxx refers to the pool ID that has the error. For example, 620032 refers to error 620 (pool threshold 1) on pool ID=50 (32 in Hex). The ID in the code is a hexadecimal value of the pool ID.

12.28.10 Pool Full Condition


In order to prevent any potential data corruption, the array needs to prevent any writes to disk once a pool has met its full state. This will mean that any requests for new storage or updates to existing storage will not be allowed. Reads can continue without issue.

There may be portions of pages that have been allocated, but are not entirely filled. Allowing for updates to these pages while preventing new pages to be allocated could result in a corruption of host data as only some updates would be allowed.

To implement this required behavior, system option 729 must be set. For more information see [System Options and System Disks](#). With SOM 729 turned on the DPVOL will become inaccessible if a full pool I/O error is encountered. The Data Retention Utility (DRU) attribute will be enabled on the DPVOL that experienced the I/O error and subsequent reads or writes are refused. The recovery procedure in this case is to first expand the pool to resolve the full pool condition then launch Storage Navigator. Select Actions > Other Function > Data Retention... and locate the LDEV with the Protect Attribute enabled.



Change to modify mode > select the protected LDEV and right mouse click to change Attribute > Read/Write > Apply



This recovery procedure will also be required when SOM 803 is enabled and pool volume errors are encountered. Pool volumes are not affected by DRU, only DPVOLs are impacted. For more information see [System Options and System Disks](#).

12.28.11 Application Migrations

Applications can be migrated to thin storage by using the block level migration tools such as UVM and HTSM. Once the migration is complete, the Zero Page Reclaim utility will need to be run on those LUNs to free up any blocks not in use by the application. For more information see [Zero Data Removal Tool](#).

12.28.12 Application and Host Operating System Considerations

Most applications can be classified as either Thin-Friendly or Not Thin-Friendly based on how they use storage. The sweet spot of thin provisioning are systems where over allocation of storage is typical.

Operating systems and file systems that tend to spread data over all available disk space generally are Not Thin-Friendly. For example, if there is metadata written about the file system throughout the LUN, then the array will consider all of those blocks written to, shifting the used pages to appear much higher than what the operating system reports. Many file systems do not reuse deleted file space.

The following chart shows the operating system and file system combinations and how a basic volume fares for thin compatibility in terms of writing to a high percentage of blocks when formatting. This does not consider the reuse of deleted files.

Host OS	File System	Thin Device Friendly
AIX	JFS	Low (write every 64MB – 65%)
	JFS2	High
	VxFS	High

HP-UX	HFS JFS (VxFS)	Low (write by 10MB) High
Linux	VxFS Ext2 Ext3 OCFS2 XFS	High Medium (write every 128MB – 30%) Medium (write every 128MB – 30%) High Write by 2GB (2%)
Solaris	UFS VxFS ZFS	Medium (write by 52MB) High (Not determined)
Windows 2000/2003	NTFS	High
Windows 2008	NTFS	Low (full format – 100%)

Windows 2000/2003 (NTFS) – use of quick format is faster but not necessary. There are no thin allocation differences between the traditional and quick formats. The traditional format simply adds an additional scan in search of media errors. NTFS uses a Least Recently Used free list management. If application has static files there is no issue (i.e. SQL Server Data and Logs). If application creates/deletes files then minimize the file system to meet exact needs (this create/delete area will quickly become 100% utilized from the array perspective). Disk Defragmenter is not recommended as it has a lot of block movement that will quickly allocate all of the thin storage.

Windows 2008 (NTFS) – use of quick format is mandatory. A normal format is not thin friendly.

Databases should use an automated growth of tablespaces (datafiles) with little buffer space between the tablespace and the tables that are within the tablespaces. In addition, if there are settings to prevent the database from write formatting blocks on all of the storage, these should be employed.

Oracle – Autoextend is recommended. Manual file management is acceptable assuming that there will be an increase number or size of tablespaces over time and the space is well-monitored. Under most conditions, using the same thin pool for logs, data, temp, and archives/FRA is recommended because of the very wide stripping nature of the pool devices (increased number of disk spindles). Separate pools will only decrease the stripe width and lower the number of potential spindles. When a separate storage tier is required (such as for FRA and/or archive logs) a different storage pool be used (different tiers will have different pools).

SQL Server – With SQL Server 2005 and later, use instant file initialization. Note: Transaction logs will be fully allocated even with this setting (may be appropriate for non-thin storage). The temporary tablespace “TEMPDB” will quickly become fully allocated due to nature of frequent writes and deletes (arrays cannot track re-use of allocated blocks). Auto-grow may not be appropriate for all environments due to unbalanced file growth (not all files in a file group are expanded at the same time). Manual file

management is acceptable assuming that there will be an increase number or size of tablespaces over time (grow all files in a file group) and the space is well-monitored.

When migrating data, check for wasted space in source files – this can be reclaimed before migration. Do not mix SQL server files with other files, especially general file space or file shares. Ideally have separate file systems for data and log files (best practice anyway, for performance, for flexibility – Hot Split Backup). Make sure logs are backed up and truncated. Do not copy files without thinking about the consequences. For volatile environments (i.e. test, development) it is suggested to limit the file system to less than the volume size.

The choice of initial SIZE and FILEGROWTH increments can have some impact. Larger values will result in better contiguity and smaller values will result in less storage waste. Contiguity is less important with enterprise storage and HDP.

Instant file initialization can be used with SQL Server 2005 on Windows 2003. The user account must have the OS privilege SE_MANAGE_VOLUME_NAME. If this setting is used, you should not need to tune SIZE and FILEGROWTH.

Sybase – The temporary tablespace “tempdb” will quickly become fully allocated due to nature of frequent writes and deletes (arrays cannot track re-use of allocated blocks). Other aspects of Sybase have not been explored.

12.29 Zero Data Removal Tool


This tool will remove zero-page data from thin storage volumes. This is intended for use against volumes that have been recently migrated to HDP and that were not 100% utilized when FAT. Typically these volumes will be 100% utilized upon a block level copy to the HDP volume.

12.29.1 Caveats and Restrictions

You cannot use this tool on volumes that are being used by volume shredding.

12.29.2 Using the Tool


While in Storage Navigator (Element Manager), go into Pools, select the appropriate pool, select the virtual volumes desired, then under more actions, select Reclaim Zero Pages. This is also an option when unallocating volumes.




12.30 Tool Panel

In order to be able to pull trace files (diagnostics and troubleshooting), there is a tool panel available on the SVP for the VSP. To access the tool panel, the following address should be used:

<https://SVP-IP-Address/cgi-bin/utility/toolpanel.cgi>



You can download trace files for diagnostics (at the request of HDS Support). You can also manage the certificates for the SVP and manage HTTP blocking. From this screen, you can also access the control panel. This is where the default time-outs and some polling period values are configured. You can also export the local account information on the download tab.



12.31 SCSI Reservations on VSP (View / Modify)

To view or change the current SCSI reservations on a host LUN, you will need to select (while on Storage Navigator) one of the host groups that contain the LUN, then select *View Host-Reserved LUNs* under *More Options*. You do not need to select specific LUNs.


Note: SCSI Reservations can be present and different types on some or all paths for a given LUN. Do not assume that a reservation is the same on all paths.

The screenshot displays the Hitachi Storage Navigator web interface. The left sidebar shows the 'Explorer' pane with 'Storage Systems' and 'Ports/Host Groups'. The main area shows the 'c14-p560-viocl4 (25)' storage system. Below the summary, the 'LUNs' tab is selected, showing a table of LUNs with reservations.

Port ID	LUN ID	LDEV ID	LDEV Name	Emulation Type	Capacity	Provisioning Type	Attribute	Number of Paths
CL1-A	0000	00:00:07		OPEN-V CVS	34.68 GB	DP	-	2
CL1-A	0001	00:00:08		OPEN-V CVS	34.68 GB	DP	-	2

At the bottom right, a context menu is open, showing options: 'View Host-Reserved LUNs', 'Delete LUN Paths', 'Edit UUIDs', and 'Export'.

You will see a list of LUNs with reservations. If there are no reservations the list will be blank (empty).



To release (force clear) the reservations, select the LUNs you wish to manipulate, then select the *Release Host-Reserved LUNs* button. This will generate a task that you will need to hit Apply to initiate. You may have to repeat this for any other paths (or host groups).

Note: SCSI Reservations can be present and different types on some or all paths for a given LUN. Do not assume that a reservation is the same on all paths.

12.32 NDM (Seamless Host Migration) on VSP – USPV Spoofing

HDS has a new system feature and process that allow a VSP to simulate the WWNs of USPV arrays (spoofing ports) for migration purposes. This allows a host to be migrated from a USPV to a VSP without making target changes that typically require host/application down time. There is a considerable amount of work and complexity with the current microcode levels, so this should be used only on hosts that have difficult or non-existent green zones. This requires access to raidcom and storage navigator on the USPV and VSP arrays.

The high level view of the process is to externalize the LUNs from the USPV to the VSP, present the LUNs down new unused ports from the VSP (but pretending to be USPV ports). Once the new paths are present and working like the original paths, the original paths are dropped. The LUN is then copied to internal VSP disk and there is no perceivable difference from the host.

Unfortunately, once a host has been migrated to the VSP using this method, it may be difficult to see in Storage Navigator (and raidcom) that a given set of ports are simulating another array. This will become more complex once multiple USPV arrays can be migrated to a VSP. The

best practice until this configuration is more obvious is to rename the 1st (default) host group to something similar to “VSP_NDM_SIMULATE_USPV-10102-CL9A”. This host group is only for documentation and should not be used to hold LUNs or WWNs. This should be performed on all ports that being used for NDM (USPV Spoofing). It is also advised to name the ports on the SAN as they will appear to be from the USP array, though they are really from the VSP.

Limitations:

- Storage Navigator does not indicate the port status as simulating a USP array, so alternative documentation is needed.
- Raidcom is the primary means of migration control.
- SCSI Reserves are not propagated from the USP to the VSP (with current microcode).
- Dropping the old paths should be accomplished via host groups on the USP array (not from zoning change).
- Professional services must perform the migration, but eventually this will become customer driven.
- The new paths will appear to be on USP host ports that are not connected to the SAN, The best ports to use are beyond what is physically installed on the USP (i.e. CL-9A, CL-AA or higher).
- There must be at least one pair of (SAN attached) interconnects between the USP and VSP. This can be on any unused ports for both arrays. This is for the externalization of storage.

12.33 Shared Memory Requirements

The following table shows the shared memeory requirements for each software product:

Array Feature	Required Shared Memory Capacity
Shadow Image	8 GB
TrueCopy	16 GB
Universal Replicator	16 GB
FlashCopy V2	16 GB
Dynamic Tiering	24 GB

12.34 SAN “Port Flapping” Alerting

Under special circumstances, SAN ports may exhibit multiple failures (and recoveries) over a short period of time. This condition is known as port flapping and indicates failing or failed components that generally reside outside of the array. However, it is possible to alert on these events. It is important to use conservative values to prevent frequent and unnecessary alerting.

To configure these alerts, remote session access to the SVP will be needed (RDP).

The screenshot shows the 'Threshold' configuration window. The 'Port' section, specifically the 'Fibre' subsection, is highlighted with a red rectangle. It contains the following settings:

- ☐ Loss Of Signal Count: Threshold 0 Cnt./sec., Over 0 sec., Term 0 sec.
- ☐ Bad Received Character Count: Threshold 0 Cnt./sec., Over 0 sec., Term 0 sec.
- ☐ Loss of Synchronization Count: Threshold 0 Cnt./sec., Over 0 sec., Term 0 sec.
- ☐ Link Failure Count: Threshold 0 Cnt./sec., Over 0 sec., Term 0 sec.
- ☐ Received EOFa Count: Threshold 0 Cnt./sec., Over 0 sec., Term 0 sec.

Other sections include 'Cache' (Cache Use Rate, Cache Write Pending Rate, Cache MCU Side File Rate), 'MP' (MP Processing Rate), 'LDEV' (Read Hit Rate), 'FCoE' (Link Failure Count, Virtual Link Failure Count, Symbol Error Count, FCS Error Count), and 'FICON' (HTP/FNP Ex Multiple, HTP/FNP Read Data Transfer Rate, HTP/FNP Write Data Transfer Rate, HTP/FNP Processing Rate).

The recommended values are:

Bad Received Character Count	1000	900 secs
Loss of Synchrnization Count	1000	900 secs

These recommended values are based on observed issues at Citi. Once either threshold has been met, dial-home events will be generated. HDS will relay these events to Citi if they occur. No SIMs or SNMP traps will be generated.

13 Security Practices

13.1 Interfaces to the USPI/USPV/VSP

The following is a descriptive list of the various interfaces into the VSP. All of these interfaces were considered for security controls.

Interface	Abbreviation	Type
Storage Navigator	SN	HTTP
HCS Device Manager GUI	HDvM (or DM)	HTTP/HTTPS
HCS Device Manager CLI	CLI	TCP
SVP Application & OS	SVP-APP	TCP
Remote Desktop	RD	TCP
Command Control Interface	CCI	TCP
Audit records sent via Syslog	SYSLOG	UDP
Configuration and Alerts sent to HDS	HiTrack	HTTPS
HCS Tiered Storage Manager GUI	HTSM	HTTP/HTTPS
HCS Tiered Storage Manager CLI	HTSM CLI	RMI

13.2 Default UserIDs

The default user id and password for Storage Navigator is "root"/"root". This password should be changed immediately and other users should be added to Storage User Account List. "Root" should not be used after an initial setup of userids. For more information see [Security in Storage Navigator](#).

The default userid and password for Device Manager is “system”/”manager”. This password should be changed immediately and other users should be added to the user database. “System” should not be used after an initial setup of users. For more information see [User Management](#).

For information regarding the setup of other user accounts, see below.

The default userid for the SVP is “maintenance” with a password of “raid-mainte” (2nd password of “etniam-diar”). The passwords should be changed by HDS staff.

13.3 Special User Accounts (Functional IDs)

The following is a list of accounts that are to be used for integrating the VSP (and HUS-VM) array and its components into the infrastructure required for secure and efficient operation. There are several rules around these functional ids that must be followed:

- These userids are restricted and to be used only by the system processes described.
- They are not to be used for a person to directly login into the system they are being created in.
- There must be an operational owner assigned to each of these userids that is fully responsible for the maintenance, use and entitlement review of these accounts.
- The passwords on these accounts should be changed at least every 12 months.

For more information regarding functional id requirements see [Citi Information Security Standards \(CISS\)](#)

Account	Application Function (Location)	Role	Interactive / Non-Interactive	Privileged / Non-Privileged	Password Change Frequency
HDS_HCS_SN_Link (local)	HCS to Array (SN)	Read-Only	Interactive	Privileged	12 Months
HDS_SRM_HCS_Link (AD)	ViPR SRM to HCS (HCS)	Read-Only	Interactive	Privileged	12 Months
HDS_ViPRC_HCS_Link (AD)	ViPR Controller to HCS (HCS)	Admin	Interactive	Privileged	12 Months
HDS_HPNA_HCS_Link (AD)	HPNA to HCS (HCS)	Read-Only	Interactive	Privileged	12 Months
Exports (local)	Export performance data from array for HDS support (SN)	Performance Data Only (Read-Only)	Non-Interactive	Privileged	12 Months
	ECC to HCS (HCS)	Retired	N/A	N/A	N/A
	SANScreen to HCS (HCS)	Retired	N/A	N/A	N/A

The legacy applications ECC and SANScreen may still have functional accounts on HCS. These accounts should be removed when the applications have been retired.

13.4 Audit Logging via SEMS

In order integrate into the EMS logging system, log files must be provided to a SEMS server via a syslog stream. As new events are added to the log file they are processed into the SEMS master logs. In addition, HDvM logs are written to the local Windows application log. These events are pushed to the SEMS repository via local agents.

In order to automatically send audit records to SEMS repository, the SVP should be configured to send the audit trail to a syslog server. This syslog server is then used to send records to a SEMS repository for reporting and record keeping. For details on setting up syslog on the SVP, refer to [Syslog Based Audit Log](#).

It is also recommended that time synchronization be used to make the information being sent to SEMS more useful. For more information see [Network Time Protocol \(SNTP\) – Time Synchronization](#). The time reported in SEMS will be based on the time set on the array SVP(s).

13.4.1 Configuring Device Manager Auditing

Authentication (e.g. login, logout and login failures) and Access Control (e.g. failure due to lack of permissions) are output to the Audit log, which is the event log (application log) on Windows. To generate Device Manager audit log information, one must check the environment settings file for the HCS Common Component is as follows:

```
C:\Program Files\HiCommand\Base\conf\sec\auditlog.conf
```

```
Log.Event.CategoryStartStop,Authentication,AccessControl, ConfigurationAccess
Log.Level 6
```

If you need to edit this file, it will be necessary to restart Hitachi Common Component to enable the new settings.


HDvM information will be logged as part of the Hitachi integrated logging (including events from other components such as HTSM), and hence the source in the event log for HTSM messages will be “HBase Storage Mgmt Log”.

13.5 User Entitlements Reporting

The user entitlements reporting for the VSP is handled by exports/queries of the Active Directory database. This is because the majority of the users accessing Device Manager (and subsequently Storage Navigator) will do so via their AD account.

13.5.1 Storage Navigator User Entitlements for non AD accounts

The best reporting for Storage Navigator users and their entitlements is to manually review each user with the Account Management screen. For more information and how to access this screen, see [Security in Storage Navigator](#).



13.5.2 Device Manager User Entitlements

Most users will access HDvM (and subsequently SN) via AD accounts. The best reporting of AD accounts is to perform a direct query on the AD groups associated with SAN management.

The best reporting for local accounts (including functional ids) and their entitlements is to run the user reporting that is built into Device Manager. For more information and how to access this functionality, see [Device Manager Reports](#). For a sample of this user reporting see [Report: Users and Permissions \(HTML\)](#).

13.6 Decommissioning of Storage

In order to return a storage array back to the vendor, the disks must be scrubbed of any data. To perform this for an entire array the following procedures should be followed:

Until an entire-array Disk Shredding function is developed, each volume on the array must be shredded.

Remove all LUNs from all of the host groups. For details on how to perform this function, refer to [Removing LUNs from Host Groups \(Deleting Paths\)](#). There should not be any paths to the existing volumes once this has been completed.

Create additional volumes to account for any free space on the parity groups. For details on how to perform this function, refer to [Volume \(LDEV\) Definition](#).

Use Volume Disk Shredding on every volume on the array. This may take a considerable amount of time. For details on how to perform this function, refer to [Disk Shredding](#).

13.7 Setting Security Policy in Device Manager

In order to comply with the security requirements for Citi, there are several settings in Device Manager which need to be made. The security screens in Device Manager define the Password policy (length, content, etc.), failed login attempt lockouts (Account Lock), and the Warning Banner. These screens can be accessed from the Security page (Administration → Security)

Keep in mind that this is only for local accounts as use of active directory based accounts are the primary method of authentication for Device Manager.

13.7.1 Password Policy

The following password policy (Administration → Security → Password) changes should be made:

- The password length should be set to 6.
- The password content should be set to case alphanumeric (at least one numeral, at least one uppercase, at least one lowercase).
- The password and user ID cannot be the same.


13.7.2 Failed Login Attempts

The following failed login lockout changes (Administration → Security → Account Lock) should be made:

- The user will be locked out after 6 failed login attempts

13.7.3 Security Banner

In addition there should be a banner (Administration → Security → Warning Banner) with information about the limited use of the software or any other restrictions. The following is an example banner:



13.8 CCI Security Bit

To turn on security, the CM Security Bit should be turned on (right click command device to select security on). With this bit set to 1, only LDEVs within the same host group as the command device can be managed by the server with its WWNs associated with the Host Group. This is a very important setting if application servers will be controlling ShadowImage.

13.9 HDvM CLI Security

Do not use a read-only username/password in the properties file. Provide the username on the command. Wait for the secure password prompt (requires JRE 1.6 or later).

If using a properties file for read only username/password is strongly desired, then use a different file when running modify commands. For example: copy the directory /opt/HiCommandCLI (or the directory used to install) to a different directory and rename the executable (HiCommandCLI.sh) to be able select the different paths.

In addition, the secure (-s) option should be used to encrypt the session.

For more information see [Appendix D: Command Line Interface for Device Manager](#).

13.10 HCS Array Name

HCS uses a default name of "VSP@ipaddress" or "USPV@ipaddress" when an array is added to the database. This name should be changed as this name will show up in reports as well as external systems like NetApp OnCommand Insight (formerly SANScreen) and EMC Control Center (ECC). Using IP addresses in this way may be considered a security violation. It is suggested that a more descriptive name such as a local array name or serial number be used instead.

13.11 Active Directory Authentication with HCS

Active Directory should be used to authenticate users for HCS. This method directly configures HCS to communicate with AD and does not need a Radius interface. Users are not defined locally with this method, only AD groups. This is the recommended configuration.

13.11.1 Acquire a functional id used for binding to the AD environment

This functional id should have a non-expiring password. This functional id is sometimes called the bind-id and is used strictly to query AD for member/group queries. Any existing storage functional/bind AD account can be used if one is already created for other storage equipment (Ex:


```
3 -file "C:\Program Files (x86)\HiCommand\Base\conf\sec\i10-hds-vm03.crt" -keystore "C:\Program Files (x86)\HiCommand\Base\conf\sec\ldapcacerts" -storepass hdskey
```

```
Owner: O=WAR 159, C=US
Issuer: O=WAR 159, C=US
Serial number: 44b80703
Valid from: Fri Jul 14 16:35:10 EDT 2006 until: Tue Jul 14 17:05:10 EDT 2026
Certificate fingerprints:
    MD5: 98:FB:DC:C0:CE:29:26:95:DB:98:46:6C:DD:AA:45:5B
    SHA1: EA:8F:F2:2B:F6:7A:DC:FA:66:1C:99:6D:23:28:34:3A:69:39:68:08
Trust this certificate? [no]: yes
Certificate was added to keystore
```

13.11.4 Verify the keystore contents on the HCS server

Verify the contents of the keystore with the following command:

```
C:\Program Files (x86)\HiCommand\Base\jdk\bin>keytool -list -v -keystore "C:\Program Files (x86)\HiCommand\Base\conf\sec\ldapcacerts" -storepass hdskey
```

```
Keystore type: jks
Keystore provider: SUN
```

Your keystore contains 1 entry

```
Alias name: i10-hds-vm03
Creation date: Dec 7, 2012
Entry type: trustedCertEntry
```

```
Owner: O=WAR 159, C=US
Issuer: O=WAR 159, C=US
Serial number: 44b80703
Valid from: Fri Jul 14 16:35:10 EDT 2006 until: Tue Jul 14 17:05:10 EDT 2026
Certificate fingerprints:
    MD5: 98:FB:DC:C0:CE:29:26:95:DB:98:46:6C:DD:AA:45:5B
    SHA1: EA:8F:F2:2B:F6:7A:DC:FA:66:1C:99:6D:23:28:34:3A:69:39:68:08
```

13.11.5 Configure the exauth.config file

In the exauth config file, there are definitions for each domain controller for the AD environment as well as specifications on the protocols used to communicate with each server. The following is an example exauth.config file:

```
auth.server.type=ldap
auth.server.name=namdev003.namdev.nsrootdev.net,namdev004.namdev.nsrootdev.net,namdev001.namdev.nsrootdev.net
auth.group.mapping=true
auth.ldap.namdev003.namdev.nsrootdev.net.protocol=ldaps
auth.ldap.namdev003.namdev.nsrootdev.net.host=namdev003.namdev.nsrootdev.net
auth.ldap.namdev003.namdev.nsrootdev.net.port=636
auth.ldap.namdev003.namdev.nsrootdev.net.timeout=15
auth.ldap.namdev003.namdev.nsrootdev.net.attr=sAMAccountName
auth.ldap.namdev003.namdev.nsrootdev.net.basedn=DC=namdev,DC=nsrootdev,DC=net
auth.ldap.namdev003.namdev.nsrootdev.net.retry.interval=1
```

```

auth.ldap.namdev003.namdev.nsrootdev.net.retry.times=20
auth.ldap.namdev003.namdev.nsrootdev.net.domain.name=namdev.nsrootdev.net
auth.ldap.namdev003.namdev.nsrootdev.net.dns_lookup=false
auth.ldap.namdev004.namdev.nsrootdev.net.protocol=ldaps
auth.ldap.namdev004.namdev.nsrootdev.net.host=namdev004.namdev.nsrootdev.net
auth.ldap.namdev004.namdev.nsrootdev.net.port=636
auth.ldap.namdev004.namdev.nsrootdev.net.timeout=15
auth.ldap.namdev004.namdev.nsrootdev.net.attr=sAMAccountName
auth.ldap.namdev004.namdev.nsrootdev.net.basedn=DC=namdev,DC=nsrootdev,DC=net
auth.ldap.namdev004.namdev.nsrootdev.net.retry.interval=1
auth.ldap.namdev004.namdev.nsrootdev.net.retry.times=20
auth.ldap.namdev004.namdev.nsrootdev.net.domain.name=namdev.nsrootdev.net
auth.ldap.namdev004.namdev.nsrootdev.net.dns_lookup=false
auth.ldap.namdev001.namdev.nsrootdev.net.protocol=ldaps
auth.ldap.namdev001.namdev.nsrootdev.net.host=namdev001.namdev.nsrootdev.net
auth.ldap.namdev001.namdev.nsrootdev.net.port=636
auth.ldap.namdev001.namdev.nsrootdev.net.timeout=15
auth.ldap.namdev001.namdev.nsrootdev.net.attr=sAMAccountName
auth.ldap.namdev001.namdev.nsrootdev.net.basedn=DC=namdev,DC=nsrootdev,DC=net
auth.ldap.namdev001.namdev.nsrootdev.net.retry.interval=1
auth.ldap.namdev001.namdev.nsrootdev.net.retry.times=20
auth.ldap.namdev001.namdev.nsrootdev.net.domain.name=namdev.nsrootdev.net
auth.ldap.namdev001.namdev.nsrootdev.net.dns_lookup=false

```


13.11.6 Defining AD Groups on HCS

Log into HCS. Go to Administration → Users and Permissions.
You should see the domain(s) defined in `exauth.config` above.

You will need to add an entry for each permission group you are defining from AD.
The following is a simple example of three different groups that correlate to three different roles.



Administrative Role



Hitachi Command Suite 7 - Windows Internet Explorer

Logged in as: Gary Johnson

Users and Permissions

per_gdse_hds_admins
Users and Permissions > Groups > namdev.nsrootdev.net > per_gdse_hds_admins


Group Detail

Distinguished Name: CN=per_gdse_hds_admins,OU=Dept_Grps,OU=StorageEng,DC=namdev,DC=nsrootdev,DC=net

Granted Permission

Application	Admin	Modify	Execute	View
User Management	✓			
HTnM	✓			✓
HRpM	✓	✓		✓
HTSM	✓	✓	✓	✓

ISA Role



Hitachi Command Suite 7 - Windows Internet Explorer

Logged in as: Gary Johnson

Users and Permissions

per_gdse_hds_isa
Users and Permissions > Groups > namdev.nsrootdev.net > per_gdse_hds_isa


Group Detail

Distinguished Name: CN=per_gdse_hds_isa,OU=Dept_Grps,OU=StorageEng,DC=namdev,DC=nsrootdev,DC=net

Granted Permission

Application	Admin	Modify	Execute	View
User Management	✓			
HTnM				
HRpM				
HTSM				

ReadOnly Role



Hitachi Command Suite 7 - Windows Internet Explorer

Logged in as: Gary Johnson

Users and Permissions

per_gdse_hds_readonly
Users and Permissions > Groups > namdev.nsrootdev.net > per_gdse_hds_readonly

Group Detail

Distinguished Name: CN=per_gdse_hds_readonly,OU=Dept_Grps,OU=StorageEng,DC=namdev,DC=nsrootdev,DC=net

Granted Permission


Application	Admin	Modify	Execute	View
User Management				
HTnM				✓
HRpM				✓
HTSM				✓

13.12 Active Directory Authentication via Radius with HCS


Active Directory should be used to authenticate users for HCS. This method uses a Radius Interface to communicate with Active Directory. Users are locally defined, but are configured to point to a Radius interface (that points to configured Active Directory Groups) to obtain current password authentication. This authentication method is recommended over the use of only local accounts because it achieves some improvement in the overall security of the system and reduces the workload associated with entitlement reviews. This method should be used only if direct Active Directory connectivity is not possible. In order to configure HDvM with Active Directory via Radius, the following steps must be completed.

13.12.1 IAS Installation and Setup


1. Install Internet Authentication Service (IAS) if not previously installed. Go to add-remove programs; select windows components; check networking services then click the details button; uncheck all boxes except IAS; click ok and next to install




2. Start IAS (start→programs→administrative tools→internet authentication service). Right click the "radius client" folder and select new radius client. Add the HCS server name and IP address and secret password (for example "hitachi").




3. Click on the "Remote Access Logging" folder then right click on the "Local File" entry in the right pane and choose properties. Check all the logging boxes and click OK.




4. Click on the "Remote Access Policies" folder; delete the existing policies in the right pane and add new policies; right click on the right pane and select new remote access policy. This will launch the remote access policy wizard. Click Next.



5. Select "Set up a custom policy" and provide a policy name like "HCS User Group 1" and click Next.




6. Click Add and select “Windows-Groups” at the bottom of the list and click add.




7. Click Add




8. Type in the name of the active directory group that contains your user list such as “nam\san_support” and click ok.




9. The group will be displayed if it was found in AD. Click OK.




10. Click next




11. Select "grant remote access permission" and click next.




12. Click “edit profile”




13. Click on the “authentication” tab and select all the boxes except the last one.




14. Click the “advanced” tab. Click on “service type”; click on edit; select login from the pull down menu. Remove the “framed-protocol” and select remove. Click on Add. Select “Ignore-User-Dialin-Properties” from the list and then click Add. Select true and click OK. Click close on the add attribute box. Your advanced setting should look like the following.




15. Click OK and then confirm “no” to the following.



16. Click Next




17. Click "Finish"



18. Repeat the above steps for each new group you need to define.

19. Your finished remote policies should look similar to the following.



IAS setup is now complete.

13.12.2 HCS Configuration of Radius

1. Edit the file C:\Program Files\HiCommand\Base\conf\exauth.properties

Note: Save a copy of the file; edit the word "ServerName" to the name of the server you installed IAS on. Also change "localhost" (if any present) to the same name as servername; change 127.0.0.1 to IP of the IAS server (even if it is the same server). Uncomment all lines (remove #) below (except the nas-ipv6-address line).

```
# Select the authentication server type.
# Enter "internal", "ldap", or "radius".
auth.server.type=radius

# Specify the authentication server identification name.
auth.server.name=namshds01

# Environment Settings for the RADIUS Server:
# The authentication protocol for the RADIUS server.
auth.radius.namshds01.protocol=PAP

# The connected server name.
auth.radius.namshds01.host=namshds01

# The number of the port used for the server connection. (default: 1812)
auth.radius.namshds01.port=1812

# The connection timeout period (in seconds).
auth.radius.namshds01.timeout=1

# The connection retry times.
auth.radius.namshds01.retry.times=3
```

```
# The client IPv4 address.
auth.radius.namshds01.attr.NAS-IP-Address=153.40.205.71

# The client IPv6 address.
#auth.radius.namshds01.attr.NAS-IPv6-Address=

# The client identifier.
auth.radius.namshds01.attr.NAS-Identifier=namshds01
```

Define the IAS server and secret password (entered on the IAS client panel) to HCS. Open an command prompt and go to the directory C:\Program Files\HiCommand\Base\bin. Issue the following command to see what IAS servers have been registered.

```
hcmdsradiussecret /list
```

```
D:\Program Files\HiCommand\Base\bin>hcmdsradiussecret /list
KAPM05296-W A server has not been registered.
```

2. Register your server (use secret password and IAS server name)
hcmdsradiussecret /set hitachi /name namshds01

```
D:\Program Files\HiCommand\Base\bin>hcmdsradiussecret /set hitachi
/name namshds01
KAPM05280-I Registration of a secret has succeeded.
```

3. Validate your server is defined
hcmdsradiussecret /list

```
D:\Program Files\HiCommand\Base\bin>hcmdsradiussecret /list
[ServerName]
namshds01
```


4. Next validate your setup and IAS connection; run the hcmdscheckauth utility providing your NAM userid and password. If successful all 3 phases will work. Note: Phase 1 & 2 are HDvM settings; Phase 3 is IAS setup/settings.

```
hcmdscheckauth /id kh64535 /pass password
```

```
D:\Program Files\HiCommand\Base\bin>hcmdscheckauth /id kh64535
/pass password
KAPM15003-I The configuration check of Phase1 will now start.
type : radius
server : namshds01
KAPM15004-I The result of the configuration check of Phase1 was
normal.
KAPM15003-I The configuration check of Phase2 will now start.
KAPM15006-I The configuration of the server namshds01 will now be
checked.
KAPM15035-I Authentication information has been acquired.
KAPM15007-I The result of the configuration check of the server
namshds01 was normal.
KAPM15004-I The result of the configuration check of Phase2 was
normal.
KAPM15003-I The configuration check of Phase3 will now start.
KAPM15010-I The connection to the server namshds01 will now be
checked. (host = namshds01, port = 1812, protocol = PAP)
KAPM15011-I The server namshds01 can be connected to normally.
KAPM15004-I The result of the configuration check of Phase3 was
normal.
```

KAPM15039-I A connection was successfully established. (server name = namshds01)

You can also verify IAS but checking the windows event log...




```

Event Type: Information
Event Source: IAS
Event Category: None
Event ID: 1
Date: 7/10/2009
Time: 5:08:38 PM
User: N/A
Computer: NAMSHDS01
Description:
User kh64535 was granted access.
Fully-Qualified-User-Name = NAM\kh64535
NAS-IP-Address = 153.40.205.71
NAS-Identifier = namshds01
Client-Friendly-Name = namshds01
Client-IP-Address = 153.40.205.71
Calling-Station-Identifier = <not present>
NAS-Port-Type = <not present>
NAS-Port = <not present>
Proxy-Policy-Name = Use Windows authentication for all users
Authentication-Provider = Windows
Authentication-Server = <undetermined>
Policy-Name = Device Manager User Group 1
Authentication-Type = PAP
EAP-Type = <undetermined>


```

13.12.3 HCS User Setup

1. Next change the accounts defined in HCS (Administration → Users and Permissions → Users) to use "Radius" instead of "internal". Select the check box next to the userid(s) you want to change. Click the "Change Auth" button at the bottom of the screen.



2. Select RADIUS and click okay..



3. Validate RADIUS by logging on using the domain userid and password. HDS setup is now complete. You will not specify the domain when logging in. Use the local account name (typically just the SOE ID). The IAS configuration should point it to the correct AD account.

13.13 Configuring IE for use with HCS and SN

There are several security configuration changes required in order for HCS and SN to work properly on a Citi desktop or server build.


1. Adobe Flash 10.1 or later should be installed.
2. Oracle JAVA JRE 1.6.0 or later should be installed.
3. For IE, go to Tools → Internet Options → Advanced Tab
 Uncheck *Do not save encrypted pages to disk*
 Uncheck *Warn about certificate address mismatch*
4. Add the SVP IPs (or subnet) to the trusted sites list
5. Bypass the proxy for local addresses

13.14 Importing Server Certificates for HCS / Element Manager

Importing Array Certificates to HCS Server (Element Manager Certificate Error)

To eliminate certificate / Apache errors when running Element Manager from HCS, the server certificates need to be imported to the HCS server.


An example error would be a certificate error from Apache (HTTP Status 404):




13.14.1 Exporting the Certificates from the Arrays

You will need to obtain a certificate from each array type, but not each individual array. For example, if you had multiple USPVP and VSP arrays, you would need to obtain a certificate from 1 USPVP array and 1 VSP array.

1. Launch the element manager for the array from HCS. Proceed to the screen after the initial certificate warning. Click on the Certificate Error box and Select View Certificates link at the bottom of the pop-up.




2. Go to the Details Tab




3. Select "Copy to File" button



4. Select Next



5. Select Next. Use a simple filename here for all of the same array type (like VSP.cer or HUS-VM.cer). These certificates can be used as imports on other HCS systems.



6. Select Next.




7. Select Finish




13.14.2 Importing the Certificates from the Arrays to the HCS Server

You will need to import a certificate for each array type.


1. Run MMC from the command line (Start --> Run --> MMC)
2. From the File Option on the Menu bar, select Add/Remove Snap-In




3. From the list of available Snap-Ins, select the Certificates and move it to the selected Snap-Ins List.




4. Select Computer Account




5. Select Local computer and then Finish.




6. Expand Certificates and then Trusted Root Certification Authorities.




7. Right Click Certificates and then select All Tasks → Import




8. Select Next




9. Specify the name of the certificate you exported for each array type.




10. Select Next and then Place all certificates in the following store (Trusted Root Certification Authorities)



11. Select Finish



12. The certificate you added should now be in the Trusted Root Certificates Authorities trust store.



This should be repeated for USPV and any other array types. The VSP certificates show up as issued to/by "SVP". The HUS-VM certificates show up as issued to/by "Hitachi, Ltd."

13.15 Configuring HCS for HTTPS

In order to secure web-based access for HCS, HTTPS needs to be configured.

13.15.1 Create SSL keys for secure CLI access

Run HiKeytool in x:\Program Files\HiCommand\DeviceManager\HiCommandServer
Enter appropriate responses otherwise take the default

```

=====
=====
HiKeytool v6.0.0-05
=====
=====
1) SSL configuration for Device Manager Server
2) SSL configuration for SMI-S
3) Exit

>1

1) Make KeyPair/Self-Signed Certificate
2) Set Device Manager Server Security Level
3) Generate CSR
4) Import Digitally Signed Certificate
5) Display contents of Device Manager Server KeyStore
6) Display verbose contents of Device Manager Server KeyStore
7) Delete an entry from the Device Manager Server KeyStore
8) Change Device Manager Server KeyPair/Self-Signed Certificate Keypass
9) Change Device Manager Server KeyStore Password
10) Import Certificate to Device Manager Server TrustStore
11) Display contents of Device Manager Server TrustStore
12) Display verbose contents of Device Manager Server TrustStore
13) Delete an entry from the Device Manager Server TrustStore
14) Change Device Manager Server TrustStore Password
15) Exit

>1

Enter Server Name [default=namghds01]:
Enter Organizational Unit [default=Device Manager Administration]:
Enter Organization Name [default=namghds01]:Citi
Enter your City or Locality:Georgetown
Enter your State or Province:Texas
Enter your two-character country-code [default=US]:
Enter Key Alias [default=namghds01]:
Passwords must only contain characters (A-Z,a-z), digits (0-9) and
whitespaces.
Do not enter special characters for your password!
This may render your keystore damaged or unusable!
Enter Key Password (6 characters minimum) [default=passphrase]:
Enter Key Algorithm [default=RSA]:
Enter Key Size (512 min, 2048 max) [default=2048]:
Enter Signature Algorithm [default=MD5withRSA]:
Enter number of days valid [default=365]:9999
Passwords must only contain characters (A-Z,a-z), digits (0-9) and
whitespaces.
Do not enter special characters for your password!
This may render your keystore damaged or unusable!
Enter KeyStore Password (6 characters minimum) [default=passphrase]:

Creating new X509Name for
    namghds01...
Creating the Device Manager Server KeyPair for namghds01 at:
    D:\Program Files\HiCommand\DeviceManager\HiCommandServer\keystore
    <this can take up to a minute>

```

All done.

(A)nother command or E(x)it?**a**

- 1) Make KeyPair/Self-Signed Certificate
- 2) Set Device Manager Server Security Level
- 3) Generate CSR
- 4) Import Digitally Signed Certificate
- 5) Display contents of Device Manager Server KeyStore
- 6) Display verbose contents of Device Manager Server KeyStore
- 7) Delete an entry from the Device Manager Server KeyStore
- 8) Change Device Manager Server KeyPair/Self-Signed Certificate Keypass
- 9) Change Device Manager Server KeyStore Password
- 10) Import Certificate to Device Manager Server TrustStore
- 11) Display contents of Device Manager Server TrustStore
- 12) Display verbose contents of Device Manager Server TrustStore
- 13) Delete an entry from the Device Manager Server TrustStore
- 14) Change Device Manager Server TrustStore Password
- 15) Exit

>**2**

Current Device Manager Server Security Level = User Logon (Basic Authentication)

Options:

- 1) User Logon (Basic Authentication)
- 2) TLS/SSL (Secure Sockets)

Enter selection: [default=1]:**2**

Device Manager Server Security level set to: TLS/SSL Secure Socket
You must restart the Device Manager Server for this change to take effect.

(A)nother command or E(x)it?**x**

Exiting...

13.15.2 Generate private key

Run sslc; key will be stored in the demoCA directory.

```
D:\Program Files\HiCommand\Base\httpsd\sslc\bin>sslc genrsa -out
demoCA\httpsdkey.pem 2048
Loading 'entropy' into random state - unable to load 'random state'
warning, not much extra random data, consider using the -rand option
Generating 2 prime RSA private key, 2048 bit long modulus
.....+++++
.....+++++

e is 65537 (0x10001)
```

13.15.3 Create a Certificate Signing Request

Use the SSLC utility to create a self signed certificate request. This step will change in the future if we need to use CA-signed certificates.

```
D:\Program Files\HiCommand\Base\httpsd\sslc\bin>sslc req -config
demoCA\sslc.cnf -new -key demoCA\httpsdkey.pem -out demoCA\httpsd.csr
```

```
Using configuration from demoCA\sslc.cnf
You will be prompted to enter information to incorporate
into the certificate request.
```

```
This information is called a Distinguished Name or a DN.
```

```
There are many fields however some can remain blank.
```

```
Some fields have default values.
```

```
Enter '.', to leave the field blank.
```

```
-----
```

```
Country Name (2 letter code) []:US
```

```
State or Province Name (full name) []:Texas
```

```
Locality Name (eg, city) []:Georgetown
```

```
Organization Name (eg, company) []:Citi
```

```
Organizational Unit Name (eg, section) []:San Planning
```

```
Common Name (eg, YOUR name) []:Kery Hines
```

```
Email Address []:kery.hines@citi.com
```

```
Please enter the following 'extra' attributes
to be sent with your certificate request
```

```
A challenge password []:citigroup
```

```
An optional company name []:Citi
```

13.15.4 Creating a Self-Signed Certificate

Use the SSLC utility to create a self signed certificate. This step will change in the future if we need to use CA-signed certificates.

```
D:\Program Files\HiCommand\Base\httpsd\sslc\bin>sslc x509 -in
demoCA\httpsd.csr -out demoCA\selfcert.pem -req -signkey
demoCA\httpsdkey.pem -days 9999
```

```
Signature OK
```

```
subject=/C=US/ST=Texas/L=Georgetown/O=Citi/OU=San Planning/CN=Kery
```

```
Hines/EMAIL=kery.hines@citi.com
```

```
Obtaining Private key
```

13.15.5 Copy new key files to prod SSL directory

Copy httpsdkey.pem and selfcert.pem from

x:\Program Files\HiCommand\Base\httpsd\sslc\bin\demoCA to

x:\Program Files\HiCommand\Base\httpsd\conf\ssl\server

13.15.6 Modify conf file to use HTTPS and find the key files

Modify the file "httpsd.conf" in x:\Program Files\HiCommand\Base\httpsd\conf\ Uncomment all of the commented lines (#) at the end of the file regarding the SSL configuration. Note that the Listen IPV6 and SSLCertificateFile lines remain commented out.

```
SSLSessionCacheSize 0
Listen 23016
#Listen [::]:23016
<VirtualHost namghds01:23016>
    ServerName namghds01
    SSLEnable
    SSLProtocol SSLv3 TLSv1
    SSLRequireSSL
    SSLCertificateFile "D:/Program
Files/HiCommand/Base/httpsd/conf/ssl/server/selfcert.pem"
    SSLCertificateKeyFile "D:/Program
Files/HiCommand/Base/httpsd/conf/ssl/server/httpsdkey.pem"
# SSLCertificateFile "D:/Program
Files/HiCommand/Base/httpsd/conf/ssl/cacert/anycert.pem"
    SSLSessionCacheTimeout 3600
</VirtualHost>
```

13.15.7 Change SSL URL to HTTPS

Use the hcmdschurl to change the URL.

x:\Program Files\HiCommand\Base\bin\hcmdschgurl.bat /change http:// <cluster manager IP address>:port number http:// <host name for the cluster manager IP address>:port number
To find out the current setting, use: hcmdschgurl.bat /print

```
D:\Program Files\HiCommand\Base\bin>hcmdschgurl /print
http://namghds01:23015 Hitachi Tuning Manager Hitachi Tiered Storage
Manager Hitachi Replication Manager Hitachi Provisioning Manager Hitachi
Device Manager
```

```
D:\Program Files\HiCommand\Base\bin>hcmdschgurl /change
http://namghds01:23015 https://namghds01:23016
KAPM06111-I The URL was changed from "http://namghds01:23015" to
"https://namghds01:23016".
```

```
D:\Program Files\HiCommand\Base\bin>hcmdschgurl /print
https://namghds01:23016 Hitachi Tiered Storage Manager Hitachi Tuning
Manager Hitachi Replication Manager Hitachi Provisioning Manager Hitachi
Device Manager
```

Have users log into <https://server:23016> instead of the http address. Reboot the server to make these changes effective (or at least restart services).

13.15.8 Setup GUI/CLI on other Windows Hosts

Download the HiCommandCerts file from the Device Manager server from the following (http) URL:

<http://<devicemanagerserver>:2001/service/HiCommandCerts>

Save in directory x:\Program Files\HiCommand\DeviceManager\HiCommandCLI

14 References

Document	Location
CSBE HDS Storage Product Page	http://catecollaboration.citigroup.net/domains/platstor/gdse/products/techstdshmpgs/tagmastore.aspx
CSBE Supported SAN Stacks	https://catecollaboration.citigroup.net/domains/platstor/gdse/products/techstdshmpgs/SANStacks.aspx

15 Appendix A: Array Physical Requirements

15.1 VSP and HUS-VM Power

The VSP and HUS-VM arrays can run on either single-phase or three-phase power. Most Citi Data centers are prepared for and prefer to use three-phase power. Either configuration is acceptable and is determined by local facility capability.

15.1.1 Single Phase Power

Location	PDU Plug	Max Voltage Rating	Max Current Rating	No. of CB per PDU	Breaker Rating	Power Supply Receptacle
USA	NEMA L6 30P twistlock 2 pole, 3 wire A + B + Gr	250V	32A	2	16A, 20A Trip	2 pole, 3 wire 220V AC (A + B + Gr)
Outside USA	IEC 309 blue 2 pole, 3 wire A + B + Gr	250V	32A	2	16A, 20A Trip	2 pole, 3 wire 220V AC (A + B + Gr)

15.1.2 Three Phase Power

Location	PDU Plug	Max Voltage Rating	Max Current Rating	No. of CB per PDU	Breaker Rating	Power Supply Receptacle
USA	NEMA L15 30P twistlock 3 pole, 4 wire A + B + C + Gr	250V	32A per phase	6 UL489	15A 2 pole	2 pole, 3 wire 220V AC (A-B or B-C + Gr)
APAC	IEC 309 blue 3 pole, 4 wire A + B + C + Gr	250V	32A per phase	6 UL489	15A, 2 pole	2 pole, 3 wire 220V AC (A-B, A-C, or B-C + Gr)
EMEA	IEC 309 red 4 pole, 5 wire A + B + C + N + Gr	415V	16A per phase	3	16A 2 pole	1 pole, 3 wire 200V AC (A-N or B-N or C-N + Gr)

15.1.3 PDU Plugs


Description	Controller Chassis	Drive Chassis
Required Number of Plugs in PDU per chassis	4	4
Required Number of PDUs per rack	4	4

15.2 VSP Dimensions and Weight

The VSP uses standard rack width racks and has the same rack width/depth on each rack no matter if it contains a DKC and DKUs or DKUs. These racks have special weighting and depth to accommodate the array components. Standard server racks would not be compatible.


15.2.1 Dimensions

Dimension	Single Module (0) 1-Rack	Single Module (0) 3-Rack	Dual Module (0,1) 6-Rack
Width	24.0 in (610 mm)	71.3 in (1810 mm)	142.1 in (3610 mm)
Height	73.1 in (2105 mm)	73.1 in (2105 mm)	73.1 in (2105 mm)
Depth	43.3 in (1100 mm)	43.3 in (1100 mm)	43.3 in (1100 mm)
Max Weight	862 lbs (389 kg)	4592 lbs (2071 kg)	9784 lbs (4413 kg)




15.2.2 Rack Naming

There can be from 1 to 6 racks in an array, depending on the number of drives. Each rack is the standard 42" (610mm) VSP rack, no matter its role. There are up to two DKC racks (named DKC 00 and DKC 01) and up to 4 DKU racks (named DKU 12, DKU 11, DKU 01, DKU 02). The growth order starts at DKC 00, DKU 01, DKU 02 (grows right). If additional drives are needed, then another DKC (DKC 10) is added next to DKC 00 and up to 2x DKUs are added to the left of DKC 10 (DKU 11 and DKU 12).



15.2.3 Service Clearances

In order to properly access the array for service, the following additional clearances are necessary:




The additional 576mm of clearance front and rear applies to each rack module (DKC or DKU). All modules are the same width (610mm).

15.3 HUS-VM Dimensions and Weight

The HUS-VM uses standard rack width racks and has the same rack width/depth on each rack no matter if it contains the DKC or DKUs. These racks have special weighting and depth to accommodate the array components. Standard server racks would not be compatible.


15.3.1 Dimensions

Dimension	Single-Rack	2-Rack	3-Rack
Width	23.5 in (600 mm)	47 in (1200 mm)	70.5 in (1800 mm)
Height	79 in (2100 mm)	79 in (2100 mm)	79 in (2100 mm)
Depth	43.3 in (1100 mm)	43.3 in (1100 mm)	43.3 in (1100 mm)
Max Weight	2166 lbs (984 kg) w/ 384 LFF drives	4165 lbs (1887 kg) w/768 LFF drives	6165 lbs (2793 kg) w/1152 LFF drives




15.3.2 Rack Naming

There can be from 1 to 3 racks in an array, depending on the number of drives. Each rack has the same dimensions, no matter its role. There is a single DKC rack (named R00) and up to 2 DKU racks (named R01 and R02). The growth order starts at R00, R01, R02 (grows right).



15.3.3 Service Clearances

In order to properly access the array for service, the following additional clearances are necessary:



The additional 576mm of clearance front and rear applies to each rack module (DKC or DKU). All modules are the same width (600mm).

15.4 HVAC (Cooling)

The recommended operational room temperature is 70–75°F (21–24°C). The recommended operational relative humidity is 50% to 55%.

The VSP and HUS-VM arrays are air-cooled. Air must enter the subsystem through the airflow intakes at the front of the frames and must be exhausted out of the back, so it is very important that the air intakes and outlets remain clear. Hitachi Data Systems recommends that under-floor air-cooling has a positive pressure and meets the specifications listed in the following table. For subsystems located at elevations from 3000 to 7000 feet (900 to 2100 meters) above sea level, decrease the maximum air temperature by two degrees for each 1000 feet (300 meters) above 3000 feet (900 meters).

The following table assumes the majority use of 3.5" LFF disks in the DKUs:

Array Configuration	DKUs	Heat Output (kW)	Heat Output (BTUs)	Air Flow (m3/min)
VSP Small and HUS-VM	2	3.7	12,485	25
VSP Medium	8	11.3	38,607	79
VSP Large	13	19.2	65,393	131
VSP Extra Large	16	23	78,523	158

15.5 Network Connectivity to SVP(s)

On the VSP, there are two SVPs (1x 10/100/1000 network drops each) within CEP / firewall. On the HUS-VM, there is only a single SVP (1x 10/100/1000 network drop). This network is used for both management and alerting.

15.6 Phone Line Connectivity

There should be no phone lines attached to the array. All communication is via TCP/IP.

15.7 Fiber Connectivity

LC Connector to (Patch Panel Connection) 50 micron multimode (8x to 128x connections)

16 Appendix B: Configuration of Array via HCS (HDvM)

Device Manager (HDvM) is the primary application and database element of Hitachi Command Suite. They are somewhat synonymous terms, though technically HCS also encompasses the HTSM and HTnM applications as well.

16.1 Accessing HCS / HDvM

To obtain access to Hitachi Command Suite / Device Manager (HDvM), the following URLs can be used:

<http://HDvMserver:23015/DeviceManager/> (if running client locally or remotely)
or
<http://127.0.0.1:23015/DeviceManager/> (if running client on same system)

After SSL has been configured:

<https://HDvMserver:23016/DeviceManager/> (if running client locally or remotely)
or
<https://127.0.0.1:23016/DeviceManager/> (if running client on same system)

Alternatively, there is a Login icon (on the HCS server) in the start menu under Programs → Hitachi Command Suite → Login. This will access the second http link above.

16.2 HCS Services

In order to check the status, stop, or start the HCS services it is recommended that you use the Start bar icons rather than directly managing the services from the windows services screens.


Status of HCS Services

To check the status of the HCS services, go to Start → Programs → Hitachi Command Suite → Manage Services → Status




Stop HCS Services

To stop the HCS services, go to Start → Programs → Hitachi Command Suite → Manage Services → Status




Start HCS Services

To start the HCS services, go to Start → Programs → Hitachi Command Suite → Manage Services → Status




16.3 License Keys for HCS

Initially logging into HCS may require licensing if not upgrading a previous instance. You will see alerts on the login page if you are not licensed or under-licensed.



License keys are best distributed via files. The HDS Support Manager should provide temporary and permanent keys for the software. Select the license button and then the file button to be able to browse for the license key files (.plk extension) and then select Save. Once you have added the files (typically one for Device Manager, one for Replication Manager, one for Tiered Storage

Manager, and one for Tuning Manager), the messages should go away. However, you may still be under licensed as in the example below:



16.4 Storage System / Array Configuration

From the Administration area of HCS, you will be able to manage the logical aspects of the arrays. You can Add, Refresh, Edit, or Remove Storage Arrays.

Hitachi Command Suite 7 Administration

Managed Resources

Storage Systems | Hosts | File Servers | Data Collection Tasks

Filter: On | Off | Select All | Column Settings | Rows/page: 30 | Page: 1 / 1

Storage System	Model	IP Address	Serial No.	Refresh Status	Last Refreshed
USP_V@10.40...	USP_V	10.40.53.25	10102	Completed	2011-06-23 08:0...
USP_VM@10.4...	USP_VM	10.40.53.5	25519	Completed	2011-06-23 08:1...
VSP@10.40.53...	VSP	10.40.53.13	53112	Completed	2011-06-23 08:1...

Selected: 1 of 3 | Add Storage System | Refresh Storage Systems | Edit Storage Systems | Remove Storage Systems

Waiting: 0 | In Progress: 0 | Completed: 0 | Failed: 0

Done | Trusted sites

16.4.1 Add Storage System

To add an array to the database, go to Administration -> Storage Systems tab, select Add Storage System. You will have to select the array type (VSP, USP, AMS, ...), provide the IP Address of the primary SVP, and a username/password. The username should be a read-only functional ID. Once HCS is linked to the array, the rights of users within HCS will supersede the rights of this user and will allow for changes on the array.

Add Storage System

To add a storage system to the list of the managed resources, specify its type, IP address, and other parameters.

Storage System Type: VSP

IP Address:

User ID:

Password:


☐ Go to Data Collection Task Status

OK Cancel Help


HCS will create a generic name for the array based on the IP provided (for example USPVM@10.40.53.5). You should change the name to reflect its serial number and/or location (for example HDS10102@Warren). This name will be reflected in all exported data and reports. See Edit Storage Systems below.

16.4.2 Refresh Storage Systems

Select the storage system you which to refresh data. Then from the Administration -> Storage Systems Tab, select the Refresh Storage Systems button. This is a refresh of the HCS database from the array.




You will not be able to make array changes (or similar activities) during this refresh. For example, if you attempted to perform a host scan, you would see an error:



16.4.3 Edit Storage Systems

Select the storage system you which to change. Then from the Administration -> Storage Systems Tab, select the Edit Storage Systems button. You will be able to change the Name, IP address, and/or the Userid/Password. The name should reflect a serial number and location rather than the default of the series and IP address.



Edit Storage Systems

To edit the storage systems, modify the parameters you want to change.

Name:

IP Address:

☐ Change User ID and Password

User ID:


Password:

☐ Go to Data Collection Task Status

OK Cancel Help

16.4.4 Remove Storage System

To remove an array from the database, select the array(s) to delete. Then go to Administration -> Storage Systems tab, select Remove Storage System. This will not affect any existing host configurations other than to remove any LUNs (and reported capacity) from the database. This can be easily reversed by re-adding the array.



Remove Storage Systems

Do you want to remove the selected storage systems?

☐ Go to Data Collection Task Status

OK Cancel Help

16.5 Host Configuration

From the Administration area of HCS, you will be able to manage the major logical aspects of the hosts. You will be able to add hosts, refresh hosts (with agent or without agent), edit hosts, import hosts, merge hosts, or remove hosts.

16.5.1 Add Hosts

Any host groups added to Storage Navigator on the array will have a "0000" added to end of the hostname used in HCS. This will not occur until you actually allocate storage to a new host.

To add a host to the database manually, click the Add Hosts button on the Administration → Hosts Tab.

Add Hosts

To manually add a host, specify its name, operating system, and WWNs. To add WWNs, specify a WWN and click [Add]. The WWN list includes WWNs that are known to the system but are not associated with any hosts.

☐ Discover Hosts ☒ Add Hosts Manually

Host Name:

OS Type:

Enter WWN:

Select WWN:

Known WWNs	
Type	WWN
WWN	12.34.56.78.90.AA.AA.AA
WWN	12.34.56.78.90.BB.BB.BB

Add >

< Remove

Selected WWNs	
Type	WWN
WWN	11.22.33.44.55.66.77.88

OK Cancel Help

You can choose the HBAs from the list or enter them manually using the Add/Remove buttons. Select OK when you are done.

16.5.2 Allocate Volumes

To allocate volumes to a host, you will want to select the host under Resources → Hosts, then hit the Allocate volumes button.

If Creating Storage

- Choose a capacity that is appropriate in TB, GB, MB, or Blocks.
- Select a storage system that will have the host zoned.
- Select Dynamic Provisioning for thin LUNs (Dynamic Tiering for page-level tiering LUNs).
- Volume location should be internal.
- Selecting Advanced options will display the criteria (drive/speed/raid) or specific pool if automatically creating them.

If Selecting Existing LUNs

- Select a storage system that will have the host zoned.
- Selecting Advanced options will display the Volume Selection option. Set it to manual.
- You will want to select the add volumes button to be able to select volumes from the array.
- Select the desired volumes from the Volumes list and click the Add button to bring them to the Selected Volumes list.
- Hit OK when you are done.

Add Volumes

Select volumes to allocate.

Volumes

Filter ☐ On ☒ Off Select All Column Settings Rows/page: 30 Page: 1 / 20

Volume	Label	Host	WWN	Port	Host Group	LUN	Capacity	Parity Group	Pool	RAID Level	Volume Attribute	Drive Type	Drive Speed (RPM)
00:00:00		HCMD...	10.00...	CL1-A...	sangs1	000A	33.40 ...	-	DP 51	-	CVS, ...	-	-
00:00:01		L15a1...	10.00...	CL1-A...	L15a1...	0000	34.68 ...	-	DP 0	-	CVS, ...	-	-
00:00:02		L15a1...	10.00...	CL1-A...	L15a1...	0001	34.68 ...	-	DP 0	-	CVS, ...	-	-
00:00:03		-	-	-	-	-	34.68 ...	-	DP 0	-	CVS, ...	-	-
00:00:04		-	-	-	-	-	34.68 ...	-	DP 0	-	CVS, ...	-	-
00:00:05		-	-	-	-	-	34.68 ...	-	DP 0	-	CVS, ...	-	-
00:00:06		-	-	-	-	-	34.68 ...	-	DP 0	-	CVS, ...	-	-
00:00:07		-	-	-	-	-	34.68 ...	-	DP 0	-	CVS, ...	-	-

Selected: 0 of 30

▼ Add ▲ Remove

Selected Volumes

Select All Column Settings Rows/page: 30 Page: 1 / 1

Volume	Label	Host	WWN	Port	Host Group	LUN	Capacity	Parity Group	Pool	RAID Level	Volume Attribute	Drive Type	Drive Speed (RPM)
00:00:...		-	-	-	-	-	34.68 ...	-	DP 0	-	CVS, D...	-	-
00:00:...		-	-	-	-	-	34.68 ...	-	DP 0	-	CVS, D...	-	-

Selected: 0 of 2

OK Cancel Help

Under Path Options, select the edit paths to manually configure (or to confirm) the LUN paths.

The Selection Table is recommended rather than Topological Graph.

Edit Paths

You can edit the paths graphically or by using the selection table. To create paths, choose ports to connect. First, select storage ports on one end, then connect them to the ports on the other end. You can also select paths to remove them.

Selected Hosts: test_host_3
 Selected Storage System: VSP@10.40.53.13
 No. of Paths: 1
 Edit Path By: ☐ Topological Graph ☒ Selection Table

Storage Port: CL2-A ▼

Host Ports

Host	Host Port
test_host_3	11.22.33.44.55.66.7...
test_host_3	11.22.33.44.55.66.7...

Add > < Remove

Selected Paths

Storage Port	Controller	Host	Host Port
CL1-A	CHA-1EU	test_host_3	11.22.33.44.55.66.77.88
CL2-A	CHA-2QU	test_host_3	11.22.33.44.55.66.77.89

OK Cancel Help

Remove any incorrect paths. Select the correct storage port and then WWN of the appropriate HBA. Select Add to combine these choices into the selected path. Repeat this for any other WWN / storage port combinations. Click OK when done.

Host Group and LUN Settings are where you select the host mode and any host mode options. LU Number can also be selected here (LUN number from host perspective).

Allocate Volumes

Host Group and LUN Settings

Host Groups:

Name: test_host_3

Host mode: Standard

Host mode options:

- VMware
- HP
- Open VMS
- Tru64
- Solaris
- Netware
- Windows
- AIX
- VMware Extension
- Windows Extension

☐ REC command support
☐ Set/Report Device Identifier enable
☐ V-Vol expansion
☐ Prioritized device recognition command
☐ Prevent "OHUB PCI retry"
☐ Queue Full Response
☐ HAM Svol Read Option
☐ Support Option for the EXTENDED COPY command

LU Number:

☒ auto: Start from 0 (hex)

☐ manual: Select LU Number

Show Plan Cancel Help

Click the Show Plan button to display a summary of the changes and confirm before executing.

Allocate Volumes

Pool	Storage System	RAID Level	Drive Type	Drive Speed (RPM)	Capacity	Used Capacity	Subscription	Subscription Limit	Used %	Used Threshold 1
DP 0	VSP@10.40.53.13	RAID6(6D+2P)	SAS	10000	3.14 TB	52.70 ...	37	-	1	70

Volume Information

Volume	Pool/Parity Group	Capacity	New/Available
00:00:0A	DP 0	35522.00 MB	Available
00:00:09	DP 0	35522.00 MB	Available

Path Information

Volume	Host WWN	Storage System Port	LUN	Host Group	Host Mode	Other Options
00:00:09	11.22.33.44.55.66.77...	CL1-A	0000	test_host_3	Standard	
00:00:09	11.22.33.44.55.66.77...	CL2-A	0000	test_host_3	Standard	
00:00:0A	11.22.33.44.55.66.77...	CL1-A	0001	test_host_3	Standard	
00:00:0A	11.22.33.44.55.66.77...	CL2-A	0001	test_host_3	Standard	

Task name: * Allocate-1

Description:

Schedule

☐ Now

☒ Later:

Date: 2011-06-27

Time: 18:32

* Required

Back Submit Cancel Help

You can alter the schedule of when this will be executed and Submit the task.

Warning: If the execution time is during a locked session of the array (i.e. SN is in modify mode) any task will fail and you will not be able to resubmit it. It is advised to pick a time that has been verified as clear for large jobs.

16.5.3 Refresh Hosts

This can only be performed on discovered (agent less but with credentials) or hosts with a HDvM agent. Manually created hosts or hosts imported (host scan) from Storage Navigator cannot be refreshed.

16.5.4 Edit Hosts

To edit a host in the HDvM database, go to Administration → Hosts tab, select the host and then click on the Edit Hosts button.

Edit Hosts

You can edit hosts by adding or removing WWNs and by modifying the operating system.

Host Name: * test-do-not-use

OS Type: Windows

Enter WWN:

Select WWN:

Known WWNs	
Type	WWN
No Data	

Selected WWNs	
Type	WWN
WWN	12.34.56.78.90.12.34.56
WWN	12.34.56.78.90.AB.CD.EF

* Required

OK Cancel Help

You can rename a host, change its OS type, or add/subtract HBAs.

Note: changes here are for the HCS database. *To affect important changes such as changing a WWN to a host that has LUNs assigned, you will also need to go to Edit LUN Paths to make changes to existing host groups on the array. Changes on this screen will only affect future allocations.*

16.5.5 Edit LUN Paths

In order to make changes on existing host groups on an array, you will need to the LUN path. Go to Resources -> Hosts then click on the desired host. Select any LUNs that need to be changed and then select the Edit LUN Path button at the bottom of the page.

Edit LUN Paths

To optimize the redundancy and the load balance, edit the paths between the host and storage system.

Selected Hosts: test-do-not-use

Selected Storage System: USP_V@10.40.53.25

No. of Existing Paths per Host: 2

Selected Volumes:

Volume Information				
Volume	Capacity	Pool/Parity Group	RAID Level	Drive Type
00:00:1F	34.69 GB	DP 123	-	-
00:70:26	34.68 GB	DP 123	-	-

Edit Path By: ☒ Topological Graph ☐ Selection Table

Storage Ports		
Port	Controller	N...
CL1-A	CHA-1EU	47
CL1-B	CHA-1EU	0
CL1-E	CHA-1FU	8
CL1-F	CHA-1FU	28
CL2-A	CHA-2QU	47
CL2-B	CHA-2QU	0
CL2-E	CHA-2QU	0

Path Editor		
Storage Ports	Paths	Host Ports
CL1-A (1 path)		test-do-not-use (2 paths)
CL2-A (1 path)		12.34.56.78.90.AB.CD.EF

Show Plan Cancel Help

It is recommended that you use the Selection Table for modifying HBAs rather than the Topological Graph.

Edit LUN Paths

To optimize the redundancy and the load balance, edit the paths between the host and storage system.

Selected Hosts: test-do-not-use

Selected Storage System: USP_V@10.40.53.25

No. of Existing Paths per Host: 2

Selected Volumes:

Volume Information				
Volume	Capacity	Pool/Parity Group	RAID Level	Drive Type
00:00:1F	34.69 GB	DP 123	-	-
00:70:26	34.68 GB	DP 123	-	-

Edit Path By: ☐ Topological Graph ☒ Selection Table

Storage Port: CL1-A

Host Ports	
Host	Host Port
test-do-not-use	12.34.56.78.90.12.3...
test-do-not-use	12.34.56.78.90.AB.C...
test-do-not-use	12.34.56.78.90.AB.C...

Add >

< Remove


Selected Paths				
Storag. ...	Control ...	Host	Host Port	State/...
CL1-A	CHA-...	test-do-no...	12.34.56.78.90.12.3...	In Use
CL2-A	CHA-...	test-do-no...	12.34.56.78.90.AB.C...	In Use

Show Plan Cancel Help


You can choose which WWNs should be used and can remove by managing the select paths with the Add and Remove buttons. The Show Plan will give you a summary of any changes before executing.

16.5.6 Import Hosts from Storage Navigator (Scan Hosts)


To bring all of the host groups from the array element manager (Storage Navigator), go to Administration -> Hosts and then select the Host Scan under More Actions.



You will get a list of current arrays and will need to acknowledge that the host group names should be the same as the host names and then to select which arrays to pull host groups from. You cannot select OK to start until you have checked the acknowledgement box at the top of the page. This can be repeated as often as desired without creating duplicate database entries in HDvM.



When the scanning has completed for all arrays selected without errors, you will see a completion screen:



16.5.7 Merge Hosts

To combine host groups into a single host logically, select a host (that is the merge destination) and then select the Merge Host option under More Actions from the Administration -> Hosts tab. This is a logical change and only affects the HCS database. It is not a way to manage host groups on an array.

Hitachi Command Suite 7

File Actions Tools Help

Address: http://127.0.0.1:23015/DeviceManagerWebService/index.jsp

Logged in as: System Log Out

Dashboard Resources Tasks & Alerts Search Administration

Quick Find

Administration

Managed Resources

Storage Systems Hosts File Servers Data Collection Tasks

Filter On Off Select All Column Settings Rows/page: 500 Page: 1 / 1

Host	IP Address	OS	Data Collector	Refresh Status	Last Refreshed
intl11_Linux_B	-	-	-	-	-
namgdse06	-	-	-	-	-
namgdse11	-	Windows	-	-	-
namgdse12	-	-	-	-	-
namgdse14	-	-	-	-	-
namgdse15_B	-	-	-	-	-
sangs1	-	Solaris	-	-	-
sangs10	-	Solaris	-	-	-
sangs11	-	Solaris	-	-	-
sangs2	-	Solaris	-	-	-
sangs3	-	Solaris	-	-	-
sangs4	-	Solaris	-	-	-
sangs6	-	Solaris	-	-	-
sangs8	-	Solaris	-	-	-
sangs9	-	Solaris	-	-	-
test-do-not-use	-	Windows	-	-	2011-06-23 15:...
test-do-not-use0...	-	Windows	-	-	-
test-do-not-use2	-	Solaris	-	-	-



Selected: 1 of 95

Add Hosts Refresh Hosts Edit Hosts More Actions

Waiting: 0 In Progress: 0 Completed: 0 Failed: 0

Done Trusted sites

Select the host(s) to be merged into the merge destination (previously selected) and then Confirm.




You will get a confirmation screen that the merge host listed above has been deleted. Note: The merge is a logical deletion and will not delete host groups on the array. It will however, tie


together different host groups on arrays for future allocations and for reporting purposes.

16.5.8 Removing Hosts (Deleting Hosts)

To remove a host from the HCS database, select the host from Administration → Hosts tab and Remove Hosts under More Actions.




You will get a confirmation of the action:



17 Appendix C: Configuration of Array via Storage Navigator

17.1 Accessing Storage Navigator


The best way to access Storage Navigator is to launch it from Device Manager. This is the “Element Manager” option available for each storage array that is defined in the HDvM database. Right click the array under Resources → Storage Systems.



If you consistently fail to load the Java Applet, then you may need to adjust your proxy settings such that a proxy server is not used for the private IP range that the SVP is in. This is a common issue with the default Internet Explorer settings at Citi.

17.2 Information Screen

The main Information page (SN → Information Option) shows basic system information including the name of the array (modifiable), contact information (modifiable), serial number, IP address of the SVP, and the current major software versions. There is also a dashboard with the capacity and allocation information for the array in graphic and table forms.



VSP SN Main Screen

On the USP, installed options will be available on the left hand list. On the USPV, installed options will be under the Go menu. Options not installed will be grayed out and not accessible.

17.3 Updates and Lock Management

As with HCS, SN updates are generally task oriented. Therefore locks on the array are held for a short period of time and are automatically managed. Normally an array will not have locks present and will have the following status:




If an array modify lock is present, the resource locked status will appear:



17.4 Modify Mode (some older screens)

To make a change to modifiable fields on some legacy screens, select the write-mode icon  in the upper right hand section of the screen. A confirmation dialogue box will appear.



Once the mode changes to modify mode, then a session time-out bar will appear at the top of the screen, the modify icon will be depressed , and the lock status will change to red .

To return to view-mode, select the modify icon  again and confirm the mode change.


To make any changes effective throughout Storage Navigator while in Modify mode, the *Apply* button must be pressed. Items that will be changed generally are displayed in blue italics and then change back to a normal font once the change has been effective.

17.5 Security in Storage Navigator


The default user id and password direct access for SN on the VSP is “maintenance/raid-mainte”. However this password should be changed immediately and other users should be added to Storage User Account List. Maintenance should not be used after the initial configuration of any local usersids.

The User Groups page (Administration → User Groups) shows the user group information. To modify /add specific users, select a group to drill down into. Each user group represents series of roles on the system (i.e. Storage Administrator, Audit-Log Administrator, etc.). There is a list of users for each group and a list of roles. You cannot modify the roles assigned to a built-in user group and there are only built-in user groups.

Note: You will not be able to access this screen from Settings → User Management as expected.



On the Administration User Group screen, you will see user name, status, and authentication method. All authentications will be local for local accounts. The Account screen allows the user to create, delete, and modify user accounts.



Removing Users


This removes the selected user from the current group. Each user must be a member of at least one group, so you will not be able to remove the user from its last group.

Deleting Users

Select the user you wish to delete from the database then select delete user. There will be a task created for this item (after you select Apply).


Creating Users

Select the Create User button to specify a new user in this group. You will need to provide a username and password. Select Finish and then a task will be created after you select apply.



Editing Users


When you edit a user, you will need to select the field(s) that you are going to modify. Otherwise no changes will be made, when you select finish. There will be a task created after you select Apply.



You cannot edit the username (you will have to delete and re-add the user with a different name).

17.6 Login Banner

To affect the contents of the login banner, go to Settings → Security → Login Banner
This screen uses the legacy modify mode to control a lock on the array.



17.7 Syslog Based Audit Log

The VSP system can be configured to automatically send all audit records from the SVP to one or two syslog servers. Syslog messages should be configured to be sent to one or more SEMS servers for audit log checking and repository. For more information on SEMS, please see [Audit Logging \(SEMS\)](#).

The Syslog screen (Settings → Security → Syslog) shows the enabled/disabled status of two syslog server destinations. Note: This screen has the legacy edit method of switching to modify mode. The *Primary Server Setting* should be configured with an ip address and port (514 is the default UDP port). The *Secondary Server Setting* should be used if there is a second syslog server destination for redundancy. The *Location Identification Name* should briefly describe the array and its location. All audit records transmitted will have this field within the record. Since the serial number is already present in the record, a suggested format is [array_name@location](#). This field can be easily updated if the equipment name changes or is moved. To get a copy of all of the available audit records in the syslog format, there is a *Download* button. This produces a .tgz compressed file.

17.8 Syslog Audit Record Format

The format of the syslog audit log is:

Field Name	Value (example)	Description
Priority	"<14>"	Facility * 8 + Severity
Timestamp	"Mmm dd hh:mm:ss"	date and time in the format of "Mmm dd hh:mm:ss"
Hostname	"SVP" fixed	hostname of the originator of the message
Program Name	"Storage" fixed	name of the program or process that generated the message
Log Format ID	"CELFSS" fixed	"CELFSS" means "Common Event Log Format for Storage System"
Log Format Revision Number	"1.1"	revision number of Audit Log Message format
Sequence Number	"1234"	sequence number of Audit Log Message (assigned by each storage or software)
Message ID	Not used in V06	identifier of Audit Log Message
Date and Time	2006-01-16T01:23:45.6+09:00	date and time in the format of RFC 3339 (ISO 8601)
Detected Entity ID	"Storage" fixed	name of the entity that detected and reported the event
Detected Point	"SVP" fixed	hostname of the originator of the Audit Log Message
Event Category	Login/Logout : "Authentication" Configuration Change : "ConfigurationAccess" Log Full etc. : "AnomalyEvent"	category of the event (Authentication, ConfigurationAccess, AnomalyEvent, etc...)
Event Code	Normal End : "Success" Error : "Failed"	result of the event (Success, Failed, or Occurred)
Subject ID	"uid=root"	"uid=" + "user name"
Hardware ID	"R500:65307"	hardware serial number
Hardware Component ID	Not used	hardware subcomponent where the event occurred

Location ID	"Tokyo-center01"	location or site (customer defined)
FQDN	Not used	Fully Qualified Domain Name
HA ID	Not used	identifier for redundant components in HA (HighAvailability) system
Agent Information	Not used	hostname, IPv4 address, or IPv6 address of the host where the agent program detected the event
Source Host	Not used	hostname, IPv4 address, or IPv6 address of the host that sent a request message and caused the event
Source Port	Not used	source tcp/udp port number of the request message
Destination Host	Not used	hostname, IPv4 address, or IPv6 address of the host that received a request message and caused the event
Destination Port	Not used	destination tcp/udp port number of the request message
Batch Operation ID	"1234"	identifier of batch operation. When one operation outputs the log of two or more lines, this number is used in order to identify that it is the log of the same operation.
Log Type	"BasicLog" or "DetailLog"	type of log ("BasicLog" means basic information, and "DetailLog" means detailed information)
Reserved #1	Not used	reserved field
Reserved #2	Not used	reserved field
Detailed Information		detailed information of each operation, command, etc... (name, operand, and parameter of operation, meaningful message, etc...) The basic information or detail information in audit.log/login.log are output in this field.

The following are some sample Syslog Audit outputs:

Login:

Oct 27 11:49:34 [10.40.53.13.19.81] Oct 27 10:49:34 SVP Storage: CELFSS,1.1,268313,,2006-10-27T10:49:34.7Z,Storage,SVP,Authentication,Success,uid=gj06529,R500:10192,,Local_USP@Warren_Lab10,,,,,,,,,268047,BasicLog,,,99,[BASE],Login ,IP=10.40.43.15*,Normal end,Seq.=386899

Mode Change:

Oct 27 12:19:28 [10.40.53.13.20.8] Oct 27 11:19:28 SVP Storage: CELFSS,1.1,268329,,2006-10-27T11:19:28.1Z,Storage,SVP,Authentication,Success,uid=gj06529,R500:10192,,Local_USP@Warren_Lab10,,,,,,,,,268061,BasicLog,,,99,[BASE],Login ,IP=10.40.43.15*,Normal end,Seq.=386905

Failed Login:

Oct 27 11:45:12 [10.40.53.13.19.58] Oct 27 10:45:12 SVP Storage: CELFSS,1.1,268311,,2006-10-27T10:45:12.7Z,Storage,SVP,Authentication,Failed,uid=gj06529,R500:10192,,Local_USP@Warren_Lab10,,,,,,,,,268045,BasicLog,,,99,[BASE],Login ,IP=10.40.43.15*,Error(290-5014),Seq.=386897
Oct 27 11:45:12 [10.40.53.13.19.58] Oct 27 10:45:12 SVP Storage: CELFSS,1.1,268311,,2006-10-27T10:45:12.7Z,Storage,SVP,Authentication,Failed,uid=gj06529,R500:10192,,Local_USP@Warren_Lab10,,,,,,,,,268045,BasicLog,,,99,[BASE],Login ,IP=10.40.43.15*,Error(290-5014),Seq.=386897
Oct 27 11:45:13 [10.40.53.13.19.60] Oct 27 10:45:13 SVP Storage: CELFSS,1.1,268312,,2006-10-27T10:45:13.7Z,Storage,SVP,Authentication,Failed,uid=gj06529,R500:10192,,Local_USP@Warren_Lab10,,,,,,,,,268046,BasicLog,,,99,[BASE],Login ,IP=10.40.43.15*,Error(290-5014),Seq.=386898
Oct 27 11:45:13 [10.40.53.13.19.60] Oct 27 10:45:13 SVP Storage: CELFSS,1.1,268312,,2006-10-27T10:45:13.7Z,Storage,SVP,Authentication,Failed,uid=gj06529,R500:10192,,Local_USP@Warren_Lab10,,,,,,,,,268046,BasicLog,,,99,[BASE],Login ,IP=10.40.43.15*,Error(290-5014),Seq.=386898

Modify Mode, Change User Entitlements for JimBob, Modify Mode Off:

Oct 27 12:19:28 [10.40.53.13.20.8] Oct 27 11:19:28 SVP Storage: CELFSS,1.1,268329,,2006-10-27T11:19:28.1Z,Storage,SVP,Authentication,Success,uid=gj06529,R500:10192,,Local_USP@Warren_Lab10,,,,,,,,,268061,BasicLog,,,99,[BASE],Login ,IP=10.40.43.15*,Normal end,Seq.=386905
Oct 27 12:23:37 [10.40.53.13.20.57] Oct 27 11:23:37 SVP Storage: CELFSS,1.1,268

[illegible]

```

Login, Modify Mode, Change Host Group (Add LUNs), Modify Mode Off, Logout:
Aug 9 10:52:01 [10.40.53.13.75.193] Aug 9 14:53:34 SVP Storage: CELFSS,1.1,20
4076,,2006-08-09T14:53:34.8Z,Storage,SVP,Authentication,Success,uid=gj06529,R500
:10192,,Local_USP@Warren_Lab10,,,,,,,,,203928,BasicLog,,,99,[BASE],Login ,IP=10.4
0.42.14*,Normal end,Seq.=359189
Aug 9 10:54:08 [10.40.53.13.75.247] Aug 9 14:55:41 SVP Storage: CELFSS,1.1,20
4082,,2006-08-09T14:55:41.1Z,Storage,SVP,ConfigurationAccess,Success,uid=gj06529
,R500:10192,,Local_USP@Warren_Lab10,,,,,,,,,203934,BasicLog,,,99,[BASE],ChangeMod
e,Modify,Normal end,Seq.=479876
Aug 9 10:56:37 [10.40.53.13.76.20] Aug 9 14:58:09 SVP Storage: CELFSS,1.1,204
087,,2006-08-09T14:58:09.8Z,Storage,SVP,ConfigurationAccess,Success,uid=gj06529,
R500:10192,,Local_USP@Warren_Lab10,,,,,,,,,203939,BasicLog,,,99,[LUNM],Add LU Pat
h,,Normal end,Seq.=479879
Aug 9 10:56:37 [10.40.53.13.76.21] Aug 9 14:58:09 SVP Storage: CELFSS,1.1,204
088,,2006-08-09T14:58:09.8Z,Storage,SVP,ConfigurationAccess,Success,uid=gj06529,
R500:10192,,Local_USP@Warren_Lab10,,,,,,,,,203939,DetailLog,,,+{Port,HostGrpID,LU
N,CU:LDEV}=[{1A,0x03,0x0001,0x10:0x04}],Num. of Paths=1
Aug 9 10:57:47 [10.40.53.13.76.39] Aug 9 14:59:19 SVP Storage: CELFSS,1.1,204
093,,2006-08-09T14:59:19.9Z,Storage,SVP,ConfigurationAccess,Success,uid=gj06529,
R500:10192,,Local_USP@Warren_Lab10,,,,,,,,,203944,BasicLog,,,99,[BASE],ChangeMode
,View,Normal end,Seq.=479882
Aug 9 10:58:27 [10.40.53.13.76.49] Aug 9 14:59:59 SVP Storage: CELFSS,1.1,204
094,,2006-08-09T14:59:59.9Z,Storage,SVP,ConfigurationAccess,Success,uid=gj06529,
R500:10192,,Local_USP@Warren_Lab10,,,,,,,,,203945,BasicLog,,,99,[BASE],ChangeMode
,View,Normal end,Seq.=479883
Aug 9 10:58:27 [10.40.53.13.76.50] Aug 9 15:00:00 SVP Storage: CELFSS,1.1,204
095,,2006-08-09T15:00:00.6Z,Storage,SVP,ConfigurationAccess,Success,uid=gj06529,
R500:10192,,Local_USP@Warren_Lab10,,,,,,,,,203946,BasicLog,,,99,[BASE],ChangeMode
,View,Logout,Normal end,Seq.=479884
Aug 9 10:58:27 [10.40.53.13.76.51] Aug 9 15:00:00 SVP Storage: CELFSS,1.1,204
096,,2006-08-09T15:00:00.7Z,Storage,SVP,Authentication,Success,uid=gj06529,R500:
10192,,Local_USP@Warren_Lab10,,,,,,,,,203947,BasicLog,,,99,[BASE],Logout,,Normal
end,Seq.=359196

```

Note: The priority field <14> translates to “user.info” on Unix system syslog configurations. Additional detail on the security changes has been requested.

It is also recommended that time synchronization be used to make the information being sent to eTrust more useful. For more information see [Network Time Protocol \(SNTP\) – Time Synchronization](#). The time reported in eTrust will be based on the time set on the array SVP(s).

17.8.1 Syslog Server Setup


In order for SYSLOG messages to be stored on a local server and routed to a log messaging repository such as SEMS, a server must have its standard Citi configuration modified. This modification will configure the SYSLOG services to receive external messages. The default behavior of a Citi server is to reject external SYSLOG messages.

Note: Windows/AIX/Linux/HP syslog platforms are not discussed here. Only Solaris platforms are discussed.

As syslog is a common service among Solaris platforms, it is a simple process to configure a server to receive syslog messages from an HDS array.

17.8.1.1 Syslog on Solaris 8

There are two configuration files that need to be modified (on Solaris 8) to enable it for remote logging.



17.10 System Events

From the Information Menu, select the Status tab to display the current system status events. These events generally reflect major system events for the system. However, not all events and alerts for the array show up on this display.

The VSP generates service information messages (SIMs) to identify normal operations as well as service requirements and errors or failures. SIMs can be generated by the front-end and back-end directors and by the SVP(s).

Ref.Code	Error Level	Status	Date
DC0C00	Serious	Not Complete	2011/07/25 04:41:58
DC0C00	Serious	Not Complete	2011/07/21 09:01:52
21804D	Moderate	Not Complete	2011/07/21 08:58:05
DC0B18	Serious	Not Complete	2011/07/21 01:46:31
21804D	Moderate	Not Complete	2011/07/21 01:42:44
621065	Moderate	Not Complete	2011/07/15 04:19:58
21804D	Moderate	Not Complete	2011/06/28 18:21:06
DC0B1A	Serious	Not Complete	2011/06/28 17:08:34
21804D	Moderate	Not Complete	2011/06/28 17:04:48
DC6B16	Serious	Complete	2011/06/02 03:01:22
21804D	Moderate	Complete	2011/06/02 02:57:35
21804D	Moderate	Complete	2011/05/24 09:33:06
21804D	Moderate	Complete	2011/05/23 10:28:13
BF5002	Moderate	Complete	2011/05/23 09:12:14
BF5001	Moderate	Complete	2011/05/23 09:12:13
21804D	Moderate	Complete	2011/05/23 09:10:27
AC5004	Moderate	Complete	2011/05/23 09:10:01
AC5003	Moderate	Complete	2011/05/23 09:10:01
AC5002	Moderate	Complete	2011/05/23 09:10:01
AC5001	Moderate	Complete	2011/05/23 09:10:00
21804D	Moderate	Complete	2011/05/20 13:50:46
21804D	Moderate	Complete	2011/05/19 15:20:22
21804D	Moderate	Complete	2011/05/18 11:19:47
DC5B00	Serious	Complete	2011/04/04 00:41:51
BF86A3	Moderate	Complete	2011/03/09 13:26:41
DC0C14	Serious	Complete	2011/01/31 21:46:24
21804D	Moderate	Complete	2011/01/31 00:46:39
DC0C04	Serious	Complete	2011/01/27 03:19:48
21804D	Moderate	Complete	2011/01/27 03:15:56
DC0C04	Serious	Complete	2011/01/27 02:01:55

Total Rows : 60

PP Status : Normal

- Reference Code – Also known as SIM
- Error Level -- Serious, Moderate, or Service (CE initiated)
- Status -- Not Complete or Complete (acknowledged by HDS/Sun CE)
- Date – Date/Time stamp of event

Selecting an event, then right clicking on the event will give an option for details. Selecting the details will display more detailed information about the event.

Ref.Code	DC0C00
Error Level	Serious
Status	Not Complete
Date	2011/07/25 04:41:58
Error Section	PAIR VOLUME STATUS ERROR
Error Detail	PAIR SUSPEND(RIO PATH CLOSE)
Location	-

1/1


Previous Next Close

A status of “Complete” shows that a CE has acknowledged or corrected the error. A status of “Not Complete” shows that a CE has not acknowledged or corrected the error. SIMs are classified according to severity: service, moderate, serious, or acute. The service and moderate SIMs (lowest severity) do not require immediate attention and are addressed during routine

maintenance. Severe (highest severity) require immediate attention. All SIMs are reported to HDS/Sun Support via HiTrack remote maintenance tool.

17.11 Port Definition

1. Select the Port Definition screen (Resources → Array → Port/Hosts Groups → Port tab)



2. Select the port(s) that should be changed and then select the edit ports button.

Edit Ports

1. Edit Ports > 2. Confirm

This wizard lets you edit one or more properties. Check the box in front of the property you want to edit, and then enter the new value.

☒ Port Attribute : Initiator

☐ Port Security : Enable

☐ Port Speed : Auto

☐ Address (Loop ID) : B2 (32)

☐ FC Switch : Enable

☐ Connection Type : P-to-P

Back Next Finish Cancel Help

You can alter the:

- o Port Attribute (Target, RCU Target, Initiator, External)
- o Port Security (Enable, Disable)

- o Port Speed (2Gbps, 4Gbps, 8Gbps, Auto)
- o Address (Loop ID)
- o FC Switch (Enable, Disable)
- o Connection Type (P-to-P, FC-AL)

Any fields to be changed must have the checkbox selected to make the change effective once the finish button is selected.

3. Make sure that all ports to be connected to the SAN (for host connectivity) are P-to-P (Point to Point) Connection, FC Switch Enabled, Auto (for 2Gbps, 4Gbps, and 8Gbps SANs), and Target Attribute. If the Target attribute is set incorrectly (i.e. it was used for another option) then refer to:

Normal FC Host Port:	Target
External:	Universal Volume Manger
Initiator:	TrueCopy and HUR Communication Link
RCU Target:	TrueCopy and HUR Communication Link

17.12 RAID (Parity) Group Definition

Creating RAID groups is a CE activity. RAID groups (also called parity groups) are created and organized throughout the array by “box” designations. Each RAID group is numbered according to its box number followed by the order of the RAID Group creation.

17.13 LUSE Volumes

LUSE volumes are concatenated sets of volumes that are presented as a single volume to the host. These volumes are used to create larger than normal volumes or to pull together smaller pieces of unused disk to make a useable volume.

LUSE Volumes are not compatible with Thin provisioning and are therefore not recommended for use at Citi.

17.14 Disk Shredding

To fully destroy the data contents of logical volumes use the volume shredding function. The shredding function will perform a write operation on any LDEVs with unspecified data with specified number of times.

Volume Shredder is part of resource manager bundle it is a Virtual LVI/LUN function. Its API can be called by other functions such as Tiered Storage Manager CLI. Volume shredder overwrites all of the data on the volume with dummy data. The default is to overwrite three times, which is the minimum that is recommended. If you define the settings, you can overwrite the data up to eight times.

Note: You cannot perform volume shredder on the following types of volumes:

- Pool volumes
- Journal volumes

The time required for a shredding operation is influenced by I/O load, but can be estimated by the following formula:

RAID 1

$0.9 * \text{volume capacity [GB]} * \text{number of times data is overwritten} = \text{number of seconds required}$

RAID 5 (7D+1P)


$0.3 * \text{volume capacity [GB]} * \text{number of times data is overwritten} = \text{number of seconds required}$

17.15 Reuse of LUNs (Deleting LUNs from Host for Future Re-Allocation)

This procedure preserves the LDEV ID such that it can be reused with other hosts without recreating a new LUN with the old ID. You can either format the LUNs or Shred them. Shredding is preferred for security purposes, but takes more time to execute.


17.15.1 Reuse of LUNs: Procedures for Storage Navigator:

1. Login to the VSP Array SN via HCS (Element Manager)
2. Got to the *Ports/Host Groups*, select the Host Group, and then select the LUNS to be removed from the host. Under *More Actions*, select Delete LUN Paths.




- a. As an optional for unallocation (Deleting LUN Paths), you can select to shred the LDEVs as a part of the task. This will perform a shred and ZPR to reclaim all of the used pages. This will preserve the LDEV id. Shredding takes longer than a format, but this option makes it very easy from the GUI. The


blocking/format/ZPR steps below are not needed if you Shred the LUNs. You will need to select the *Delete all defined LUN paths to above LDEVs* checkbox to be able to Shred the LDEVs. Otherwise, you will only be able to select the *Finish* button. If you do not want to Shred the LDEVs, then select *Finish* (and then *Apply*) and proceed to step 3.



- b. Select Next. You will get a warning message about “Shred LDEVs”. You can modify the Shredding Data pattern (00 FF 00 is the default and current Citi standard). Select Finish (then Apply). Once the task has completed, you can reallocate the LDEV.



3. Block the LDEVs
4. Format the LDEVs (There will no longer be data on the LDEV, but the used pages have not been released from the pool).
5. Reclaim Zero Pages (all or nearly all of the pages will be released back to the pool). The LDEV will be unblocked after ZPR has completed such that the LUN can be allocated to another host.
6. You can also unallocate the LUN, and then at a later time, Shred the LDEV (under *More Options* after selecting specific LDEVs in Logical Devices) to skip the block/format/ZPR steps. This may take a little more time to execute, but has fewer steps.



17.15.2 Reuse of LUNS: Procedures for RAIDCOM:

1. Start the horcm instance and export the HORCMINST variable. Login to the array (raidcom –login).
2. Remove the LDEVs from the host groups (can identify host group by name or id; id used here for simplicity)


```
raidcom delete lun –port CL1-A-1 –ldev_id 0x3500
raidcom delete lun –port CL2-A-1 –ldev_id 0x3500
```
3. Block the LDEVs



```
raidcom modify ldev –status blk –ldev_id 0x3500
```
4. Format the LDEVs


```
raidcom modify ldev –initialize –ldev_id 0x3500 –operation fmt
```
5. Reclaim Zero Pages (ZPR)



```
raidcom modify ldev –ldev_id 0x3500 –status discard_zero_page
```
6. The LDEV can be allocated to another host.

17.16 Displaying WWNs Login (List Logins Function)

To determine what WWNs have logged into the array, go to Ports/Host Groups. Select the “Login WWNs” tab and a list of recently logged in WWNs is displayed.



To show if the WWNs status of being logged in or logged out, select the “View Login WWN Status” button at the bottom of the screen.



Alternatively, you can obtain this same information via the RAIDCOM CLI.

```
>raidcom get port -port CL2-A -IH0
PORT      LOGIN_WWN      Serial#      -
CL2-A     10000000c9392531  210060      -
CL2-A     10000000c938f1f1  210060      -
CL2-A     10000000c937faaf  210060      -
CL2-A     10000000c9423972  210060      -
CL2-A     10000000c9872367  210060      -
CL2-A     c0507602cbef002c  210060      -
```

```
CL2-A 10000000c9ee5ea7 210060 -
CL2-A 10000000c95c6233 210060 -
CL2-A c0507602cbef002e 210060 -
CL2-A 10000000c9e53295 210060 -
```

(This example is for a specific port, but you could leave out the port value and it will display for all ports).

For more details on how to setup and use RAIDCOM, refer to:

<https://catecollaboration.citigroup.net/domains/platstor/gdse/products/platforms/Tagmastore/CS/BE-11097-SG.pdf>

18 Appendix D: CLI for HCS (HDvM CLI)

CLI (Command Line Interface) is the component of HCS (Device Manager) installed on an open systems host to enable command line functions. CLI allows you to control the configuration and storage allocation changes to the USP array from a scripting based interface.

18.1 HiCommand CLI Setup (Solaris)

Install the Solaris package for HDS Device Manager CLI

The software will be placed in:

/opt/HiCommandCLI (or /export/opt/HiCommandCLI)

18.2 HiCommand CLI Setup (Windows)

Install the Windows Package for HDS Device Manager CLI.

18.3 Secure Option

To avoid the transmission of a clear text password, use the `-s` (or `-secure`) option when invoking the command line. This must be configured after SSL has been configured for the server.

18.4 Secure Password Prompting (Using JRE 1.6 or later)

To get the CLI to ask for the password in a secure way, JRE 1.6 or later needs to be installed on the CLI client machine.

Note: the HDvM server JRE version does not affect this feature, only the client.

The 1.6 JRE package will require the manual add of the `/opt/jre/1.6.0_16/bin` directory to the `$PATH` variable in your profile. Remove any previous JRE versions from `$PATH`. You may also need to delete a symbolic link that points to the wrong version (or nonexistent directory) in `/bin` or `/usr/bin`. Running “`whence java`” should find the first instance of java following your path. Running “`java -version`” should tell you the current version.

JRE version must be specified if the default JRE is a previous version is below version 1.6 and the path modification is not desirable. Two lines can be added to the `HicommandCLI.sh` script specifying this JRE version if changing the `$PATH` variable is not desirable:

```
HDVM_CLI_JRE_PATH=/opt/jre/1.6.0_16/bin
export HDVM_CLI_JRE_PATH
```

If this modification is implemented, it will have to be repeated with every subsequent upgrade.

Example of password prompting with JRE = 1.6.0_16

```
C:\HiCommandCLI>hicommandcli getstoragearray -u system
Password required; Password:
RESPONSE:
An instance of StorageArray
```

```

objectID=ARRAY.R600.10102
name=USP_V@10.40.53.25
description=USP_V (10102) at 10.40.53.25
serialNumber=10102
arrayFamily=R600

```

Example of password prompting with JRE = 1.6.0_16

```

[/opt/HDvMcli]
sangs6# HiCommandCLI GetStorageArray -u gj06529
Password required; Password:
RESPONSE:
An instance of StorageArray
  objectID=ARRAY.R600.10102
  name=HDS10102_USPV
  description=USP_V (10102) at 10.40.53.25
...

```

18.5 Properties File

Command line parameters can be added to the properties file to prevent them from being entered at every command. The properties file is in the directory where the package was installed. The file name is:

`/export/opt/HiCommandCLI/HiCommandCLI.properties`

Set the HDvM (HiCommand Device Manager) URL location in the properties file as the example output indicates

```

...
# set the server url
#HiCommandCLI.serverurl=http://localhost:1963/service
HiCommandCLI.serverurl=http://emcsrv07:2001/service
#HiCommandCLI.serverurl=https://emcsrv07:2443/service

##### OPTIONS #####

#user=khalsa
#user=system

# password can be provided directly, or from a password file
#password=
#password=@D:\\HiCommand\\passwd

# provide a copy of the raw xml request & response in MessageTrace.log file
#messagetrace=true

##### COMMAND PARAMETERS #####

# set the Array model, for commands that use this parameter
model=USP
serialnum=10192

```

Since the CLI client system is likely a shared system, do not use the “user” entry in the properties file. Instead, specify the username with the `-u username` option. Specify the password with the `-`

p password option. It is suggested that for scripting, the username and password are passed to the script using variables rather than hard coding these values into the scripts.

Set the model number and serial number of the array in this file as well.

18.5.1 Using properties file to store read-only username/password

If you use the properties file to store a read-only username/password, then you can avoid the additional parameters on the command line. However, when you wish to use a modify-capable userid, you will have to remove the password from the properties file or you will get the following error:

```
C:\HiCommandCLI>Hicommandcli getstoragearray -u system
Error performing command: KAIC90084-E Server returned the following HTTP Respons
e: "HTTP/1.0 401 Unauthorized"; no additional information available
```

This can be avoided by specifying the username and password on the command, but that will put the modify username and password on the screen in clear text. If Powerbroker is used, PB will capture this full command in its log.

CSBE recommendation: Do not use a read-only username/password in the properties file. Provide the username on the command. Wait for the secure password prompt (requires JRE 1.6 or later). If using a properties file for read only username/password is strongly desired, then use a different file when running modify commands (i.e. copy \HiCommandCLI or /opt/HiCommandCLI to a different directory and rename the executable to select the different paths).

18.6 CLI Structure

The structure of CLI is as follows:

HiCommandCLI {server-url} {server-command} [options] [parameters]

Available Commands:

- AddStorageArray
- GetStorageArray
- DeleteStorageArray
- AddLunScan
- AddLogicalGroup
- DeleteLogicalGroup
- ModifyLogicalGroup
- GetLogicalGroup
- AddObjectForLogicalGroup
- DeleteObjectForLogicalGroup
- AddLunGroup
- ModifyLunGroup
- AddWWNForLUNGroup
- DeleteWWNForLUNGroup
- AddWWNGroup
- ModifyWWNGroup
- DeleteLunGroup
- AddLun
- DeleteLun
- AddHostStorageDomain
- DeleteHostStorageDomain
- AddWWNForLun
- DeleteWWNForLun

AddWWNForHostStorageDomain
 DeleteWWNForHostStorageDomain
 DeleteWwn
 DeleteWwnGroup
 AddHost
 DeleteHost
 ModifyHost
 GetHost
 AddHostRefresh
 AddHostInfo
 DeleteHostInfo
 ModifyHostInfo
 GetHostInfo
 GetServerInfo
 GetLogFile
 ModifyDebugLevel
 GetDebugLevel
 GetAlerts
 DeleteAlerts
 AddArrayGroup
 DeleteArrayGroup
 AddLogicalUnit
 DeleteLogicalUnit
 ModifyLogicalUnit
 FormatLU
 AddSpareDrive
 DeleteSpareDrive
 ModifyPort
 ModifyPortController
 AddURLLink
 DeleteURLLink
 GetURLLink
 RefreshStorageArrays
 AddLUSE
 DeleteLUSE
 AddReplication
 DeleteReplication
 ModifyReplication
 AddConfigFileForReplication
 GetReplicationControllerPair
 AddArrayReservation
 ModifyArrayReservation
 DeleteArrayReservation
 GetArrayReservation

Options:

-u {userid} or --user {userid}	login user ID for HiCommand Server
-p {password} or --password {password}	login password for HiCommand Server
-s or --secure	presence indicates secure connection (https)
-i {filename} or --input {filename}	take input from the specified file, which must contain the complete xml request
-o {filename} or --output {filename}	send output to the specified file, instead of the console
-b {filename} or --batch {filename}	specify the batch file, and execute the batch
-t or --messagetrace	record raw XML request and response in MessageTrace.log file

18.7 Host Modes

Set operating system specific options is similar to the GUI, but you must know the command text.

WARNING: Changing host group mode options on the USP and USPV subsystems that are already installed and configured is disruptive and requires the server to be rebooted.

For a list of the valid supported host modes, see [Storage Host Setup \(Host Modes\)](#).

Example Configurations

- HP-UX
 - Hostmode="HP"
 - Hostoption=12
- Solaris
 - Hostmode="Solaris"

18.8 Array Discovery

After a new install of HDvM, the array(s) must first be discovered in order to manage them. This step is not necessary if the array has already been configured in Device Manager. To discover the array, the following must be known, the IP address of the array, the userid and password that is required for Storage Navigator and the array type.

```
HCLI.sh AddStorageArray ipaddress=192.168.94.101 family=USP userid=demolab
arraypasswd=demolab
```

RESPONSE:

```
An instance of StorageArray
objectID=ARRAY.USP.10192
name=10192
description=USP (10192) at 192.168.94.101
serialNumber=10192
arrayFamily=USP
arrayType=USP
microcodeVersion=50-03-98/00
agentVersion=04_03_10
productName=RAID500
controllerVersion=50-03-99-00/00
numberOfControllers=2
capacityInGB=4,319
cacheInMB=16,384
sharedMemoryInMB=-1
numberOfSpareDrives=-1
freeCapacityInGB=3,821
allocatedCapacityInGB=497
hihsmCapacityInGB=0
onDemandCapacityInGB=0
totalFreeSpaceInGB=623
largestFreeSpaceInGB=203
capacityInKB=4,529,375,808
freeCapacityInKB=4,007,614,848
allocatedCapacityInKB=521,760,960
hihsmCapacityInKB=0
onDemandCapacityInKB=0
totalFreeSpaceInKB=653,520,896
```

```

largestFreeSpaceInKB=213,679,104
multipathSupport=1
securityStatus=2
sequenceNumber=10192
displayArrayFamily=USP
displayArrayType=USP
numberOfLUs=261
numberOfAllocatedLUs=24
numberOfUnallocatedLUs=237
slprStatus=0

```

The response provides quite a bit of information about the array, such as the number of devices, firmware version, number of unallocated devices, number of allocated devices, subsystem serial number, array type, etc.

18.9 Adding LUNS

Once the array has been added, the CLI can now be used to add LDEVs to storage ports on the array.

1. Find free devices
2. Determine which ports will be used and find the corresponding port IDs
3. Add to existing Host Storage Domain
4. Add LDEVs to ports

1. Find free devices

We need to first determine which devices or “free” or unallocated. Unallocated in this case means not attached to a port.

Note: The example command uses the filter of free to display unallocated devices. The “- - output” option allows you to save the output to a file.

```
HCLI.sh GetStorageArray --output freelist subtarget=LogicalUnit lufilter=Free
```

```

RESPONSE:
An instance of StorageArray
objectID=ARRAY.USP.10192
name=10192
description=USP (10192) at 192.168.94.101
serialNumber=10192
arrayFamily=USP
arrayType=USP
microcodeVersion=50-03-98/00
agentVersion=04_03_10
productName=RAID500
controllerVersion=50-03-99-00/00
numberOfControllers=2
capacityInGB=4,319
cacheInMB=16,384
sharedMemoryInMB=-1
numberOfSpareDrives=-1
freeCapacityInGB=3,821
allocatedCapacityInGB=497
hihsmCapacityInGB=0
onDemandCapacityInGB=0
totalFreeSpaceInGB=623
largestFreeSpaceInGB=203
capacityInKB=4,529,375,808
freeCapacityInKB=4,007,614,848

```

```

allocatedCapacityInKB=521,760,960
hihsmCapacityInKB=0
onDemandCapacityInKB=0
totalFreeSpaceInKB=653,520,896
largestFreeSpaceInKB=213,679,104
multipathSupport=1
securityStatus=2
sequenceNumber=10192
displayArrayFamily=USP
displayArrayType=USP
numberOfLUs=261
numberOfAllocatedLUs=24
numberOfUnallocatedLUs=237
slprStatus=0
List of 237 Lu elements:
An instance of LogicalUnit
  objectID=LU.USP.10192.10
  devNum=10
  emulation=OPEN-V
  devCount=1
  devType=
  capacityInKB=33,554,880
  path=false
  commandDevice=false
  commandDeviceSecurity=false
  chassis=1
  arrayGroup=54
  raidType=RAID5(7D+1P)
  currentPortController=-1
  defaultPortController=-1
  isComposite=0
  trueCopyVolumeType=Simplex
  shadowImageVolumeType=Simplex
  quickShadowVolumeType=Simplex
  universalReplicatorVolumeType=Simplex
  sysVolFlag=0
  externalVolume=0
  differentialManagement=false
An instance of LogicalUnit
  objectID=LU.USP.10192.1028
  devNum=1,028
  emulation=OPEN-V
  devCount=1
  devType=
  capacityInKB=33,554,880
  path=false
  commandDevice=false
  commandDeviceSecurity=false
  chassis=3
  arrayGroup=18
  raidType=RAID1(2D+2D)
  currentPortController=-1
  defaultPortController=-1
  isComposite=0
  trueCopyVolumeType=Simplex
  shadowImageVolumeType=Simplex
  quickShadowVolumeType=Simplex
  universalReplicatorVolumeType=Simplex
  sysVolFlag=0
  externalVolume=0
  differentialManagement=false
...

```

The first block of information is information regarding the array. What you want from the output is the LDEV number (devNum) and possibly the capacity. To keep things simple, we will add this device as shown below

```
devNum=16
capacityInKB=61,088,640
```

In the above example, the first available device address is 19 and it is approximately 61 GB. Note that these device addresses are decimal values. The values are represented in Hexidecimal within the Storage Navigator interface.

2. Determine ports

Use the following command to determine the portIDs of the ports for use. In this example, the device address 19 will be added to port 1B and 2B. In order to specify these commands when adding LUNs, the portIDs must be determined. By issuing the following:

```
HCLI.sh GetStorageArray --output hsd.log subtarget=HostStorageDomain
Simplified excerpt from response:
...
An instance of HostStorageDomain
  objectID=HSDOMAIN.HDS9970V.20414.1.0
  portID=1
  domainID=0
  hostMode=AIX
  displayName=CL1-B-0
  nickname=Prod
...
...
...
An instance of HostStorageDomain
  objectID=HSDOMAIN.HDS9970V.20414.17.1
  portID=17
  domainID=1
  hostMode=Solaris
  displayName=CL2-B-1
```

The responses are usually quite lengthy, however UNIX shell commands can be used like, grep and awk to filter specific information from these files

As the above example shows, the CL1-B has a corresponding port ID of 1 and CL2-B shows a port ID of 17.

3. Add the new LUN to a HSD

Add LUNs by issuing the following command:

```
HCLI.sh AddLun --output addpath.log port=17 domain=1 lun=0 devnum=16
HCLI.sh AddLun --output addpath.log port=1 domain=1 lun=0 devnum=16
```

The above command will add device number 16 to domain 1 with a Lun number of 0. This device will be added to both paths; 1B and 2B and assigned to Lun number 0.

Now that the Lun has been allocated, the administrator may want to create and add this Lun to a logical group for management purposes. HDvM enables you to create a hierarchical structure for collectively grouping storage based on applications, functions, etc.

For example, an administrator may desire to have a logical group called IT Infrastructure named for a departmental group within the company. He may then want to add all servers for this “IT Infrastructure” that house storage to this group and further divide them by platform type. The below text illustrates the scenario.

IT Infrastructure >	Windows>	ServerA's storage
		ServerB's storage
	Solaris>	ServerC's storage
	AIX>	ServerD's storage

With such a scenario, a storage administrator can group servers/storage according to company departments, or functions and track the amount of storage being used for each department or group. For additional details on Logical Groups, reference the *HDvM (HiCommand Device Manager) Web Client User's Guide*.

In order to create the structure illustrated above with the HDvM CLI, Add top level logical group by doing the following:

```
HCLI.sh AddLogicalGroup groupname=ITInfrastructure iconfile=group_0.gif
```

RESPONSE:

```
An instance of LogicalGroup
objectID=GROUP.1
name=ITInfrastructure
logicalPath=ITInfrastructure
icon=group_0.gif
capacity=0
capacityInKB=0
realCapacityInKB=0
percentUsed=0
numberOfLUNs=0
```

```
HCLI.sh AddLogicalGroup groupname=solaris iconfile=group_0.gif
```

RESPONSE:

```
An instance of LogicalGroup
objectID=GROUP.3
name=solaris
parentID=GROUP.1
logicalPath=ITInfrastructure/solaris
icon=group_0.gif
capacity=0
capacityInKB=0
realCapacityInKB=0
percentUsed=0
numberOfLUNs=0
```

Add storage to the “Solaris” group as follows. Solaris group will be converted from a logical group to a storage group:

```
HCLI.sh AddobjectForLogicalGroup objectID=GROUP.3
groupelement=PATH.HDS9970V.20414.17.1.16
```

Other Useful Commands:

Provide information about all paths that have LUNs assigned to them:

```
HCLI.sh GetStorageArray --output path subtarget=Path
```

Provides information about a given port, i.e., topology, hostmode, etc.

```
HCLI.sh GetStorageArray --output port subtarget=Port
```

Provides information pertaining to logical units. Can be filtered to display specific about the LUN, like free or allocated

HCLI.sh GetStorageArray --output Lun subtarget=LogicalUnit lufilter=FREE

18.10 CLI Example Use Case

The following use case is a typical scenario that one would follow in allocating LUNs to a host via the DM CLI:

- Add a Storage Array to HDvM Server. Adding an array via the CLI requires an ip address and username by which to access the array.

```
hc.bat addstoragearray ipaddress=172.16.100.1 family=USP userid=root
```

RESPONSE:

```
An instance of StorageArray
objectID=ARRAY.USP.14053
name=USP@172.16.100.1
description=USP (14053) at 172.16.100.1
serialNumber=14053
arrayFamily=USP
arrayType=USP
microcodeVersion=50-04-01/00
agentVersion=04_05_00
productName=RAID500
controllerVersion=50-04-01-00/00
numberOfControllers=1
capacityInGB=5,096
cacheInMB=32,768
sharedMemoryInMB=-1
numberOfSpareDrives=-1
freeCapacityInGB=2,682
allocatedCapacityInGB=2,390
hihsmCapacityInGB=22
onDemandCapacityInGB=0
totalFreeSpaceInGB=873
largestFreeSpaceInGB=499
capacityInKB=5,343,728,320
freeCapacityInKB=2,812,868,240
allocatedCapacityInKB=2,507,059,040
hihsmCapacityInKB=23,801,040
onDemandCapacityInKB=0
totalFreeSpaceInKB=916,350,304
largestFreeSpaceInKB=523,783,680
multipathSupport=1
securityStatus=2
sequenceNumber=14053
displayArrayFamily=USP
displayArrayType=USP
numberOfLUs=1,547
numberOfAllocatedLUs=818
numberOfUnallocatedLUs=729
slprStatus=0
openTotalCapacity=5,343,728,320
openAllocatedCapacity=2,507,059,040
openFreeCapacity=2,812,868,240
openHiHsmCapacity=23,801,040
openOnDemandCapacity=0
imTotalCapacity=0
imAllocatedCapacity=0
imFreeCapacity=0
imHiHsmCapacity=0
```

```

imOnDemandCapacity=0
mfTotalCapacity=0
mfHiHsmCapacity=0
mfOnDemandCapacity=0
mfAllocatedCapacity=0
mfUnallocatedCapacity=0
numberOfOpenAllocatedLUs=818
numberOfOpenUnallocatedLUs=729
numberOfImAllocatedLUs=0
numberOfImUnallocatedLUs=0
numberOfMfLDEVs=0
numberOfAllocatedMfLDEVs=0
numberOfUnallocatedMfLDEVs=0
productCode=2
List of 1 CommParameters elements:
  An instance of CommParameters
    userID=root
    userPassword=root
    ipAddress=172.16.100.1

```

- Create a Solaris Host

```

hc.bat addhost hostname="citigroup01.citi.com" ipaddress="10.15.7.126"
      wwwlist="21.00.00.E0.8B.0D.B8.D3,21.00.00.E0.8B.12.4A.BB"

```

RESPONSE:

```

An instance of Host
  objectID=HOST.7
  name=citigroup01.citi.com
  ipAddress=10.15.7.126
  capacityInKB=0
  hostType=-1

```

- Find Available LUNs

```

hc.bat getstoragearray model=USP serialnum=14053 subtarget=logicalunit lufilter=free lusubinfo=LDEV -o
      free_lun

```

edited

```

An instance of LogicalUnit
  objectID=LU.USP.14053.4096
  devNum=4,096
  displayName=10:00
  emulation=OPEN-V
  devCount=1
  devType=
  capacityInKB=36,375,360
  path=false
  commandDevice=false
  commandDeviceSecurity=false
  chassis=1
  arrayGroup=0
  raidType=RAID5(7D+1P)
  currentPortController=-1
  defaultPortController=-1
  isComposite=0
  trueCopyVolumeType=Simplex
  shadowImageVolumeType=Simplex
  quickShadowVolumeType=Simplex
  universalReplicatorVolumeType=Simplex

```

- Find HSD in use on Ports. Examine port 5 CL3B to find HSD in use.

```
hc.bat getstoragearray subtarget=Port port=5 model=USP serialnum=14053
portsubinfo=Hoststoragedomain
```

List of 1 Port elements:

An instance of Port

```
objectID=PORT.USP.14053.5
serialNumber=14053
arrayType=USP
portID=5
portType=Fibre
fibreAddress=E0
topology=Fabric(off), FC-AL
displayName=CL3-B
lunSecurityEnabled=true
controllerID=1
worldWidePortName=50.06.0E.80.04.36.E5.21
channelSpeed=0
portRole=Target
slprNumber=0
```

List of 2 HostStorageDomain elements:

An instance of HostStorageDomain

```
objectID=HSDOMAIN.USP.14053.5.0
portID=5
domainID=0
hostMode=Standard
hostModeOption=13
displayName=CL3-B-0
nickname=3B-G00
```

An instance of HostStorageDomain

```
objectID=HSDOMAIN.USP.14053.5.1
portID=5
domainID=1
hostMode=Standard
hostModeOption=
displayName=CL3-B-1
nickname=HCMD0500
```

Examine port 133 CL4-B to find HSD in use.

```
hc.bat getstoragearray subtarget=Port port=133 model=USP serialnum=14053
portsubinfo=Hoststoragedomain
```

List of 1 Port elements:

An instance of Port

```
objectID=PORT.USP.14053.133
serialNumber=14053
arrayType=USP
portID=133
portType=Fibre
fibreAddress=C7
topology=Fabric(on), Point-to-Point
displayName=CL4-B
lunSecurityEnabled=true
controllerID=9
worldWidePortName=50.06.0E.80.04.36.E5.31
channelSpeed=0
portRole=Target
slprNumber=0
```

List of 1 HostStorageDomain elements:

```
An instance of HostStorageDomain
objectID=HSDOMAIN.USB.14053.133.0
portID=133
domainID=0
hostMode=Standard
hostModeOption=13
displayName=CL4-B-0
nickname=4B-G00
```

- Create new HDS for Solaris Server. Create HSD for port 5 CL3-B and port 133 CL4-B

```
hc.bat addhoststoragedomain serialnum=14053 port=5 domain=2 nickname="citigroup01-3B"
hostmode="Solaris"
```

List of 1 HostStorageDomain elements:

```
An instance of HostStorageDomain
objectID=HSDOMAIN.USB.14053.5.2
portID=5
domainID=2
hostMode=Solaris
displayName=CL3-B-2
nickname=citigroup01-3B
```

```
hc.bat addhoststoragedomain serialnum=14053 port=133 domain=1 nickname="citigroup01-4B"
hostmode="Solaris"
```

List of 1 HostStorageDomain elements:

```
An instance of HostStorageDomain
objectID=HSDOMAIN.USB.14053.133.1
portID=133
domainID=1
hostMode=Solaris
displayName=CL4-B-1
nickname=citigroup01-4B
```

- Add WWN and LUN to HSD that was defined.

```
hc.bat addwwnforhoststoragedomain domain=1 port=133 wwn="21.00.00.E0.8B.12.4A.BB"
serialnum=14053
```

List of 1 HostStorageDomain elements:

```
An instance of HostStorageDomain
objectID=HSDOMAIN.USB.14053.133.1
portID=133
domainID=1
hostMode=Solaris
displayName=CL4-B-1
nickname=citigroup01-4B
```

List of 1 WWN elements:

```
An instance of WWN
WWN=21.00.00.E0.8B.12.4A.BB
```

```
hc.bat addwwnforhoststoragedomain domain=2 port=5 wwn="21.00.00.E0.8B.0D.B8.D3" serialnum=14053
```

List of 1 HostStorageDomain elements:

```
An instance of HostStorageDomain
objectID=HSDOMAIN.USB.14053.5.2
portID=5
domainID=2
hostMode=Solaris
```

```

displayName=CL3-B-2
nickname=citigroup01-3B
List of 1 WWN elements:
  An instance of WWN
    WWN=21.00.00.E0.8B.0D.B8.D3

```

Add LU 4096 to HSD on port 5, 133

```
hc.bat addlun serialnum=14053 model=USP port=5 domain=2 lun=000 devnum=4096
```

```

List of 1 Path elements:
  An instance of Path
    objectID=PATH.USP.14053.5.2.4096
    devNum=4,096
    portID=5
    domainID=2
    scsiID=15
    LUN=0
    wwnSecurityValidity=true

```

```
hc.bat addlun serialnum=14053 model=USP port=133 domain=1 lun=000 devnum=4096
```

```

List of 1 Path elements:
  An instance of Path
    objectID=PATH.USP.14053.133.1.4096
    devNum=4,096
    portID=133
    domainID=1
    scsiID=15
    LUN=0
    wwnSecurityValidity=true

```

- Create Logical Groups and Assign Host/Paths to Logical Groups

Create Logical Group Hierachy

```
hc.bat addlogicalgroup groupname=Warren iconfile=group_0.gif
```

```

RESPONSE:
An instance of LogicalGroup
  objectID=GROUP.1
  name=Warren
  logicalPath=Warren
  icon=group_0.gif
  capacity=0
  capacityInKB=0
  realCapacityInKB=0
  percentUsed=0
  numberOfLUNs=0

```

```
hc.bat addlogicalgroup groupname="citigroup01.citi.com" iconfile=group_0.gif parent=GROUP.1
```

```

RESPONSE:
An instance of LogicalGroup
  objectID=GROUP.2
  name=citigroup01.citi.com
  parentID=GROUP.1
  logicalPath=Warren/citigroup01.citi.com
  icon=group_0.gif
  capacity=0
  capacityInKB=0
  realCapacityInKB=0

```

```
percentUsed=0
numberOfLUNs=0
```

Add all host and path to logical group.

```
hc.bat addobjectforlogicalgroup
objectid=GROUP.2groupelements="HOST.7,PATH.USP.14053.5.2.4096,PATH.USP.14053.133.1.4096"
RESPONSE:
An instance of LogicalGroup
  objectID=GROUP.2
  name=citigroup01.citi.com
  parentID=GROUP.1
  logicalPath=Warren/citigroup01.citi.com
  icon=group_0.gif
  capacity=69
  capacityInKB=72,750,720
  realCapacityInKB=36,375,360
  percentUsed=0
  numberOfLUNs=2
List of 3 GroupElement elements:
  An instance of Host
    objectID=HOST.7
    name=citigroup01.citi.com
    ipAddress=10.15.7.126
    capacityInKB=36,375,360
    hostType=-1
  An instance of Path
    objectID=PATH.USP.14053.5.2.4096
    devNum=4,096
    portID=5
    domainID=2
    scsiID=15
    LUN=0
    wwnSecurityValidity=true
  An instance of Path
    objectID=PATH.USP.14053.133.1.4096
    devNum=4,096
    portID=133
    domainID=1
    scsiID=15
    LUN=0
    wwnSecurityValidity=true
```

19 Appendix E: Command Control Interface / ShadowImage

CCI (Command Control Interface) is the component of ShadowImage installed on an Open Systems host to enable command line functions. CCI in conjunction with ShadowImage allows you to control internal array replication operations from the host enabling the use of scripting on the Open Systems host to automate ShadowImage tasks.

Warning: Never send standard write I/O to a command device. It is not intended to be used as a storage device.

With previous versions of this document there was a restriction of the use of CCI to storage administrators. This was due to some security limitations of the product. HDS now supports additional security based on the SM security bit. Subsequently, CCI can be used by application servers to control Shadow Image.

19.1 CCI Security

To turn on security, the CM Security Bit should be turned on (right click command device and select security on). With this bit set to 1, only LDEVs within the same host group as the command device can be managed by the server with its WWNs associated with the Host Group.

19.2 HORCM Files

HORCM stands for Hitachi Open Remote Copy Manager. These are the configuration files that are used to control the copies of LDEVs internally to an array (Shadow Image) or externally (True Copy).

The configuration files, `/etc/horcm0.conf`, and `/etc/horcm1.conf` control the operation of each CCI instance. Two instances are configured to control each half of the pairs. Each instance reads its configuration file at startup. If changes are made, these will become active only if the instances have been restarted after the changes have been made. It is recommended that instances be shutdown while changes are made.

There are four main sections of the configuration file. The components of the configuration file are HORCM_MON, HORCM_CMD, HORCM_DEV, and HORCM_INST.

19.2.1 HORCM_MON

Describes the IP address, port number, and the paired volume error-monitoring interval of the local host. The following parameters are defined in the HORCM_MON section.

- **ip_address:** The IP address used by the local host.
- **Service:** The port name assigned to the CCI service, (registered in the `/etc/services` file). The port number may be used.
- **Poll:** The interval for monitoring the paired volumes.
- **timeout:** The time-out period of communication with the remote server.

19.2.2 HORCM_CMD

Describes the device path of the command device of the USP. The command device must be mapped to the SCSI interface using the SVP or the remote console LUN Manager software. Use the "command device" option in LUN manager of Storage navigator to identify the command device. Multiple command devices can be defined allowing a second command device to take over should the initial device fail.

Never attempt to use the command device as a storage device after it has been designated as a command device.

19.2.3 HORCM_DEV

Describes the addresses in the USP of the physical volumes corresponding to the paired logical volume names. Each group name must have a unique name. Each device must also have a unique name. Care must be taken to ensure the correct Port Numbers are used. There is no checking to verify that a host interface corresponds to the USP port. The following parameters are defined in the HORCM_DEV parameter:

- **dev_group**: Names a group of paired logical volumes. A command is executed for all corresponding volumes according to this group name.
- **dev_name**: Unique name that controls the paired logical volume within a group.
- **Port #**: Describes the port number of the volume on the USP, which corresponds to the dev_name volume. The port described is the SCSI or fibre port sending or receiving remote copy information.
- **target ID**: Describes the SCSI target ID number of the physical volume on the port specified in the port # parameter. When defined on a local host, this is the target ID for the primary volume of a pair, and when defined on a remote host, this is the target ID for the secondary volume of a pair.
- **LU #**: Describes the SCSI logical unit number (LU#) of the physical volume specified in the target ID parameter.
- **MU #**: Describes the remote copy number (0 - 2) of ShadowImage volumes. If this number is omitted it is assumed to be zero.

19.2.4 HORCM_INST

Describes the network address (IP address) of the remote server (active or standby). It is used to refer to or change the status of the paired volume in the remote server (active or standby). When the primary volume is shared by two or more servers, there are two or more remote servers using the secondary volume. Thus, it is necessary to describe the addresses of all the servers. The following parameters are defined by the HORCM_INST parameter:

- **dev_group**: The name of the pair group shared by other CCI instances.
- **ip_address**: The network address of the remote server.
- **service**: The port name assigned to the CCI communication path (registered in the /etc/services file). If a port number is specified instead of a port name, the port number will be used.

19.3 CCI Configuration Files

Configuration files must be edited after installing the CCI software. The default configuration file is /etc/horcm.conf. The default /etc/horcm.conf was copied to /etc/horcm0.conf and /etc/horcm1.conf as there were two instances of horcm to be started on the CCI Server. Each instance corresponds to the primary (source) and secondary (target) volumes in the USP. A description of these files is provided below.

Sample Citi configuration files:

/etc/horcm0.conf

```
HORCM_MON
#ip_address      service      poll(10ms)  timeout(10ms)
146.128.135.141  horcm0      1000        3000
#/****** For HORCM_CMD *****/

HORCM_CMD
#dev_name        dev_name      dev_name
/dev/rds/c3t0d0s2

#
#/****** For HORCM_DEV *****/
```

```

HORCM_DEV
#dev_group    dev_name    port#    TargetID    LU#    MU#
#VG01         oradb1      CL1-A    1           1
#VG02         oradb2      CL1-A    1           1    0
group1        group1_dev1  CL1-A    0           14
group1        group1_dev2  CL1-A    0           15

HORCM_INST
#dev_group    ip_address    service
group1        146.128.135.141 horcm1

```

NOTES:

HORCM_MON- ip_address can be either the hostname or the IP address. If using the hostname, the host name must be in either /etc/hosts or in the DNS server.

HORCM_CMD - This is the command device specified on the USP, there may be more than one command device for redundancy.

HORCM_DEV - This configuration file is for the primary (source) volumes. Dev_names in a configuration file have to be unique.

HORCM_INST - This indicates the other horcm instance that identifies the host where the secondary (target) volumes reside.

/etc/horcm1.conf

```

HORCM_MON
#ip_address    service    poll(10ms)    timeout(10ms)
146.128.135.141    horcm1    1000    3000
#/****** For HORCM_CMD *****/

HORCM_CMD
#dev_name      dev_name      dev_name
/dev/rdisk/c3t0d0s2

#
#/****** For HORCM_DEV *****/

HORCM_DEV
#dev_group    dev_name    port#    TargetID    LU#    MU#
#VG01         oradb1      CL1-A    1           1
#VG02         oradb2      CL1-A    1           1    0
group1        group1_dev1  CL1-A    0           18
group1        group1_dev2  CL1-A    0           19

HORCM_INST
#dev_group    ip_address    service
group1        146.128.135.141 horcm0

```

HORCM_MON- ip_address is the same as in the /etc/horcm0.conf. The service has been changed.

HORCM_CMD - This is the same command device specified in the /etc/horcm0.conf.

HORCM_DEV - This configuration file is for the secondary (target) volumes. Notice that the dev_names in this config file, horcm1.conf, are the same as device names in horcm0.conf. The

device names create the copy relationship between P-Vols and S-Vols. For example, the device name “group1_dev1” tells Shadow Image that LUN 14 in horcm0.conf will be replicated to LUN 18 in horcm1.conf.

HORCM_INST - This indicates the other horcm instance that identifies the host where primary (source) volumes.

After the configuration files are defined and rechecked, the HORCM instances must be started on the CCI server. Any errors in the configuration files will prevent the horcm instances from starting. Check the log files in /HORCM/logx to find details for problem determination. Ensure the environmental variables are set prior to starting the instances. It may be a good practice to place the following variables in the .profile file.

19.4 Getting the LUN information

The LUN number and target ID must be determined for adding the primary LUNs to /etc/horcm0.conf and the secondary LUNs to /etc/horcm1.conf. These numbers must be values seen by the storage subsystem. To obtain the subsystem LUN numbers and target IDs, the “raidscan” command should be run as follows:

```
#raidscan -p <port> <HostStorageDomain> -fx -CLI
#raidscan -p cl1-A 2 -fx -CLI
```

The above command does a “raidscan” on storage port 1A, Host Storage Domain 2. The HSD would be the group that the storage admin defined the P-VOLs (Primary volumes) and S-VOLs (Secondary Volumes) to. Column 5 from this output will give the LUN numbers managed by the subsystem. The “raidscan” command can also be run without the “HSD” option. However the “HSD” option enables you to filter for a specific group of LUNs. Without this option, all LUNs attached to the port will be displayed. Details about “raidscan” are documented in the “*Command Control Interface User’s Guide*”.

Ensure that each LUN added to for a given ShadowImage group has the same **dev_group** identifier in the first field of each line and a unique **dev_name** in the second field of each line in the HORCM_DEV section of the /etc/horcm0.conf and /etc/horcm1.conf file.

CCI uses the dev_name of a device as defined the horcm.conf files to determine the pairing of specific P-VOLs with specific S-VOLs within a ShadowImage volume group. The dev_name field for each device must be unique within /etc/horcm0.conf, and must be the same for a given pair in on the CCI server in both /etc/horcm0 and /etc/horcm1.conf files.

19.5 Starting the Instances

Begin by first setting the environment variables as follows:

```
export HORCMINST=0
```

The default, gives control to instance 0, which manages the primary volumes

```
export HORCC_MRCF=1
```

Specifies ShadowImage, if not set CCI will expect commands that relate to TrueCopy.

Use the following command to start the instances:

```
#horcmstart.sh 0 1
```

This will start two daemons, one for each instance

Note: These instances are not set up to start automatically by default. In order to start them up automatically, scripts will be required in /etc/init.d or run-level directories.

Use the “pairedisplay” command to display the pairs prior to creating the pairs. These commands can be run from either node in the cluster but cannot be run from both.

#pairedisplay -g group1 -fcx

```
Group PairVol(L/R) (Port#,TID, LU-M) ,Seq#,LDEV#.P/S,Status, % ,P-LDEV# M
group1 group1_dev1(L) (CL1-A , 0, 14-0 )10192 10b.SMPL ----,---- -
group1 group1_dev1(R) (CL1-A , 0, 18-0 )10192 133.SMPL ----,---- -
group1 group1_dev2(L) (CL1-A , 0, 15-0 )10192 10d.SMPL ----,---- -
group1 group1_dev2(R) (CL1-A , 0, 19-0 )10192 131.SMPL ----,---- -
```

Notice the local (L) and remote (R) designations. Use this to verify the intended direction of the copy process. The command to initiate the copy process is “paircreate”. This command with the -vl option specifies that the direction in which to copy is from local to remote.

19.6 Starting Replication

When assured of the mapping, issue the following command to create a pair:

#paircreate -g group1 -vl

After issuing the paircreate command, the system prompt will return immediately. Use the pairevtwait command to monitor pair creation or issue the pairedisplay command to view status.

#pairevtwait -g group1 -s pair -t 9999 &

-or-

#pairedisplay -g group1 -fcx (the “fcx” option ‘c’ is for copy rate and ‘x’ displays LDEVs in hex)

```
Group PairVol(L/R) (Port#,TID, LU-M) ,Seq#,LDEV#.P/S,Status,% ,P-LDEV# M
group1 group1_dev1(L)(CL1-A , 0, 14-0 )10192 10b.P-VOL COPY,25 133 -
group1 group1_dev1(R)(CL1-A , 0, 18-0 )10192 133.S-VOL COPY,25 10b -
group1 group1_dev2(L)(CL1-A , 0, 15-0 )10192 10d.P-VOL COPY,25 131 -
group1 group1_dev2(R)(CL1-A , 0, 19-0 )10192 131.S-VOL COPY,25 10d -
```

The above output shows that the P-VOLs (Primary volume) to S-VOLs (Secondary Volume) copy rates are at 25%.

Once the volumes have been completely synchronized, the volumes will show a “pair” status. This status can be viewed by issuing “pairedisplay”. The below example illustrates.

#pairedisplay -g group1 -fcx

```
Group PairVol(L/R)(Port#,TID, LU-M),Seq#,LDEV#.P/S,Status, % ,P-LDEV# M
group1 group1_dev1(L)(CL1-A , 0, 14-0 )10192 10b.P-VOL PAIR, 88 133 -
group1 group1_dev1(R)(CL1-A , 0, 18-0 )10192 133.S-VOL PAIR, 88 10b -
group1 group1_dev2(L)(CL1-A , 0, 15-0 )10192 10d.P-VOL PAIR, 88 131 -
group1 group1_dev2(R)(CL1-A , 0, 19-0 )10192 131.S-VOL PAIR, 88 10d -
```

Preparing Secondary Volumes for Mounting

The following lists steps required to import the ShadowImage copy of the disk group on each node in the cluster.

1. Split the pairs:

#pairsplit -g group1

2. Issue the pairedisplay command to check the status:

#pairedisplay -g group1 -fcx

```

Group PairVol(L/R)(Port#,TID,LU-M),Seq#,LDEV#.P/S,Status,%,P-LDEV# M
group1 group1_dev1(L)(CL1-A , 0, 14-0 )10192 10b.P-VOL PSUS,90 133 W
group1 group1_dev1(R)(CL1-A , 0, 18-0 )10192 133.S-VOL SSUS,90 10b -
group1 group1_dev2(L)(CL1-A , 0, 15-0 )10192 10d.P-VOL PSUS,90 131 W
group1 group1_dev2(R)(CL1-A , 0, 19-0 )10192 131.S-VOL SSUS,90 10d -

```

Note: The pairs are suspended and the secondary (target) volumes are in read/write status.

Use UNIX mount command to mount the volumes on the target server. The target server is the server attached to the S-VOLs (Secondary Volumes). If Veritas Volume Manager is used, the volume group will need to be imported and then the volumes mounted.

Begin backups using backup utility

19.7 Resynchronizing Volumes

Resynchronize ShadowImage volumes after backups have been completed. The resynchronization does not have to take place immediately after the backups. The resynchronization should start prior to the next backup cycle. Only the changes will be copied. The display below shows the pairs in a “copy” state indicating that the changes are being replicated to the S-VOLs (Secondary volume). Once the status shows “PAIR”, the S-VOLs (Secondary Volume) have been replicated completely. A “pairsplit” can now be issued. This will suspend the pairs. Pairs should only be suspended for mount after a “PAIR” status has been reached.

#pairresync -g group1**#pairedisplay -g group1 -fcx**

```

Group PairVol(L/R)(Port#,TID,LU-M),Seq#,LDEV#.P/S,Status,%,P-LDEV# M
group1 group1_dev1(L)(CL1-A , 0,14-0 )10192 10b.P-VOL COPY,88 133 -
group1 group1_dev1(R)(CL1-A , 0,18-0 )10192 133.S-VOL COPY,88 10b -
group1 group1_dev2(L)(CL1-A , 0,15-0 )10192 10d.P-VOL COPY,88 131 -
group1 group1_dev2(R)(CL1-A , 0,19-0 )10192 131.S-VOL COPY,88 10d -

```

19.8 Deleting Pair Relationships

In order to delete pair relationships:

#pairsplit -S -g group1**#pairedisplay -g group1 -fcx**

```

Group PairVol(L/R)(Port#,TID,LU-M),Seq#,LDEV#.P/S,Status,%,P-LDEV# M
group1 group1_dev1(L)(CL1-A , 0,14-0 )10192 10b.SMPL ----,----- -
group1 group1_dev1(R)(CL1-A , 0,18-0 )10192 133.SMPL ----,----- -
group1 group1_dev2(L)(CL1-A , 0,15-0 )10192 10d.SMPL ----,----- -
group1 group1_dev2(R)(CL1-A , 0,19-0 )10192 131.SMPL ----,----- -

```

Notice that the status of the pairs is now “SMPL”. This indicates that a pair relationship no longer exists between the volumes. The data on the volumes do not change, but the replication no longer takes place.

20 Appendix F: Firewall Ports

The recommended configuration for the VSP is to have the SVPs located behind a firewall or CEP. HDvM should reside on an SOE Wintel server on the corporate network. The CLI client for HDvM should reside on a Linux or Solaris server. The HDvM server and the SN client will go through the firewall to communicate with the SVPs. The following ports must be configured:

Description	Type	Port	Protocol
Device Manager to SVP	TCP	1099	
	TCP	51099	
	TCP	21	FTP
	TCP	51100	
SVP to Device Manager	TCP	2001	Physical View
	UDP	161	SNMP
	UDP	162	SNMP
	TCP	51099	
	TCP	51100	
Device Manager CLI	TCP	2001	Does not go through firewall
CCI for Replication		User Config	Does not go through firewall
Storage Navigator to SVP	TCP	80	HTTP
	TCP	443	HTTPS
	TCP	21	FTP
	TCP	1099	
	TCP	51099	
	TCP	51100	
HDS Support to SVP	TCP	3389	MS-RDP
HDS Dial Home	TCP	8080	Webproxy
	TCP UDP	53	DNS
Audit Records to External Syslog Server	UDP	514	Syslog
SVP to SNTP	UDP	123	SNTP
SVP to ePO	TCP	18080	
	TCP	18083	
	TCP	18085	
	TCP	8801	
ePO to SVP	TCP	18081	
	TCP	18082	
	TCP	18443	New Port
ESS/SNMP to SVP	UDP	161	SNMP
SVP to ESS/SNMP	UDP	162	SNMP

21 Appendix G: Citi McAfee Virus Scan and ePO Agent

21.1 Installation of Anti-Virus Software

In order to use Citi approved virus scanning software on the SVP, any previously installed virus scanning software must first be removed. The Citi packaged McAfee Virus Scan software and ePO agent are then installed. The ePO agent will keep VirusScan up to date. The following procedures require Administrator access to the SVP workstation and therefore should be performed by HDS personnel.

Note: These procedures have been tested on both the Windows XP and the Windows Vista variant of the SVP.

21.1.1 Prior To Beginning Installation

1. Firewall ports should be opened if there is a firewall between the SVPs and the ePO server(s). These TCP ports include 18080, 18083, 18085, and 8801. The ePO server(s) will need TCP ports 18081, 18082, and 18443 to communicate to the SVPs. For a complete listing of firewall ports see [Appendix F: Firewall Ports](#).
2. Obtain a SiteList.xml file that describes the local ePO server locations from regional ePO / Virus Scan administrators. In recent versions of ePO, the sitelist has been bundled with the FramePackage.exe. The ePO installation will require this file to point the client to the correct ePO server. The [Citi Host Security Website](#) has valuable information including the latest version information. Regional Administrators can be found at the [ePO Useful Information and Contacts](#) website. Obtain the files for a non-soe installation of Virus Scan 8.0i (for example VSE85_Non-SOE.zip) from the [Citi Host Security \(Administrator Toolkit\)](#). To obtain the 4.X ePO agent (FramePackage.exe) for non-soe systems go to the [Citi Host Security ePO](#) website.

Note: different versions of this software may be named differently and may have slightly different locations and filenames.

3. It is critical that you contact the regional ePO administrator to make sure the exclusions are configured and assigned to the SVPs before installing the software. Without these exclusions in place, the SVP will likely not function correctly and may inhibit changes on the array. This is due to the resource conflicts between McAfee and the HDS software running on the SVP. There are detailed instructions on the configuration of [McAfee VSE Exclusion Policy for HDS Storage Arrays](#).

21.1.2 Installation Steps

4. Remove any previously installed versions of virus scanning software. This includes Symantec Virus Scan and the Symantec Autoupdate applications. To uninstall these applications go to Control Panel → Add/Remove Programs and select to uninstall these applications.
5. ePO / McAfee Client Install. Run the FramePackage.exe file provided. Recent versions will include the sitelist.xml file within the executible, so be careful that it is pointing to the right ePO location for your region.

In older versions, go to the directory that contains Frmlnst.exe (or FramePkg.exe). The variability in the name of the executable is version dependent. Run the following command (from the command line) to install the product:

```
FramePkg.exe /install=agent /siteinfo=<path to>sitelist.xml
```

For example:

```
FramePkg.exe /install=agent /siteinfo=c:\temp\virus\sitelist.xml
```

6. Repeat steps 4-8 for the secondary SVP.

Once the McAfee agent (ePO agent) has been successfully installed, confirm that the regional ePO administrator can configure the SVPs from the ePO server. The ePO administrator will also have the ability to push the appropriate version of VSE to the system. This is the best method of deployment as it ensures the correct policies are in place and the ePO system can communicate with the SVP before you run VSE.


For more information on McAfee ePO and Anti-Virus at Citi, see the [Host Security Engineering Website](#).

If a SVP has to be replaced, it is important to follow the installation steps mentioned above. Any non-communicating ePO client will be purged from the ePO database after four weeks. In addition, alerting may not be configured unless specially requested from the regional ePO administrators.

21.2 Global ePO Policy Exclusions

In order to prevent contention on the SVP with the SVP Applications, the following exclusions have been approved for global deployment. For more details on how to properly configure ePO for HDS, see [McAfee VSE Exclusion Policy for HDS Storage Arrays](#).


On-Access Default Processes.Settings for Workstation.Detection.What not to scan.Exclusions



This will exclude *.DAT, *.DBG, *.DMP, *.LOG, *.TRC files from the appropriate directories.

Note: Do not check the exclude subfolders option as it will conflict with the pattern and the ePO agent will not exclude any files.

On-Access Low Risk Processes Policies.Settings for Workstation.Processes



Adding svp_chk.exe and ctfmon.exe to the list will prevent undesirable the blocking of undesirable ports (and registry modification). The SVP uses ports 6666 and 6667 on its internal network side (isolated from the Citi network). In addition, these applications may try to change the registry entries to make sure they are set to autorun.

Note: Without these exclusions in place, the SVP will likely not function correctly and may inhibit changes on the array. This is due to the resource conflicts between McAfee and the HDS software running on the SVP.

21.3 Uninstalling Anti-Virus Software

The uninstall will follow the reverse order of the installation procedures. It is critical to remove the ePO agent first as it will automatically have Virus Scan reinstalled during the next update cycle.

To uninstall the ePO agent, go to the directory that contains FrmInst.exe. Run the following command (from the command line) to install the product:

```
FrmInst.exe /remove=agent
```

To uninstall Virus Scan, go to the control panel and select add/remove programs. You should be able to remove the application (like any other windows-based application) from this screen.

21.4 Reporting and Tree Structure

For each region, a tree structure should be requested from the regional ePO Administrators that will contain all of the HDS USP SVP units. This will simplify change requests and reporting. SAs in conjunction with the Site Admins in ePO need to keep investigating out of date machines based upon daily DAT/Definition Deployment Summary reports.

22 Appendix H: Setup of HiTrack Service for HTTPS

HiTrack is the outbound alert and configuration information system that is from Citi to HDS.


Firewall ports must be opened in order for the HTTP dial-home functionality to work. Once configured and tested, any analog modem should be removed from the configuration of the array and disconnected from any phone lines.

The address for the proxy-server must be resolved via dns if not using an ipaddress. Use of a server name instead of ipaddress is recommended, but this will require that the SVP have access to a dns server (port 53 TCP/UDP). The transmission of data back to HDS and the lookup of the HDS servers are accomplished via the webproxy server(s) (port 8080 TCP) or the HTTP Application proxy service. For a more detailed listing of the firewall ports see [Appendix F: Firewall Ports](#).


Notes: The following information is provided for information only. HDS CEs are supposed to make these changes. The information being transmitted to HDS is non-confidential and is encrypted. This applies to all HDS arrays.

22.1 Setup of Hi-Track Software

1. Confirm that the SVP has internet access by running IE with the proxy server information described below (webproxy.ssmb.com, port 8080).
2. Open the HiTrack application (by selecting the HiTrack  icon from the system tray) on the SVP desktop.



3. Go to the Center Tab to configure the Contact Attempt Sequence.



4. Primary FTP Server. Select Add to create a new entry.

Sample entries (each site/region may vary)

Method Type: HTTPS

HDS host: ushtinet01.hds.com

Username: ht_http_row

Password: httpw0425

Proxy Server Type: HTTP Tunnel

Host or IP: webproxy.wlb.nsroot.net (or HTTP Application Proxy)

Port: 8080 (or HTTP Application Proxy port)

Note: the address for the proxy-server must be resolved via dns if not using an ipaddress. Use of a server name instead of ipaddress is recommended, but this will require that the SVP have access to a dns server (port 53 TCP/UDP). The transmission of data back to HDS and

the lookup of the HDS servers are accomplished via the webproxy server(s) (port 8080 TCP). For a more detailed listing of the firewall ports see [Appendix F: Firewall Ports](#).

Select the Close button to save the entry.

5. Testing. From the General Tab, select the *Site Manager Log* button to bring up the Site Manager Log. Return to the Center Tab, select the new entry and then select the *Properties* button. Select the *Test* button and monitor the activity in the Site Manager Log. A successful session will look like:

```

Hi-Track Site Manager version B.8 Operations Log
Hi-Track is a registered trademark of Hitachi Data Systems Corporation.
Copyright (c) Hitachi Data Systems Corporation 1995-2007
All rights reserved.
-----
07/27/2007 09:47:45 GMT Standard Time
07/27/2007 09:47:45 GMT is 7/27/2007 9:47:45 AM
07/27/2007 09:47:45 GMT offset is 00:00:00
07/27/2007 09:47:45 Daylight savings time is not in effect.
07/27/2007 09:47:45 Culture is English (United States)
07/27/2007 09:47:45 Testing HTTPS Method: ushtinet01.hds.com
07/27/2007 09:47:45 Initiating transfer.
07/27/2007 09:47:45 Retrieving files from SVP agents.
07/27/2007 09:47:45 Fetching files from back 00
07/27/2007 09:47:46 Fetch successful for C:\HDS\Dat.zip
07/27/2007 09:47:48 Fetch successful for C:\HDS\Config.zip
07/27/2007 09:47:48 Environ file written for 10192
07/27/2007 09:47:48 Setting HTTP proxy to webproxy.ssmb.com:8080.
07/27/2007 09:47:48 Transferring files via HTTPS:
07/27/2007 09:47:48 ..transferring Config.R054665_10192 to
ushtinet01.hds.com/ht_http_row/Config.R054665_10192
07/27/2007 09:47:49 SSL Cert acceptable on host basis.
07/27/2007 09:47:49 SSL Cert acceptable on host basis.
07/27/2007 09:47:49 SSL Cert acceptable on host basis.
07/27/2007 09:47:49 SSL Cert acceptable on host basis.
07/27/2007 09:47:49 SSL Cert acceptable on host basis.
07/27/2007 09:47:49 SSL Cert acceptable on host basis.

```

...This message may be repeated several times.

```

07/27/2007 09:47:49 SSL Cert acceptable on host basis.
07/27/2007 09:47:49 SSL Cert acceptable on host basis.
07/27/2007 09:47:49 SSL Cert acceptable on host basis.
07/27/2007 09:47:49 SSL Cert acceptable on host basis.
07/27/2007 09:47:49 SSL Cert acceptable on host basis.
07/27/2007 09:47:49 SSL Cert acceptable on host basis.
07/27/2007 09:47:49 SSL Cert acceptable on host basis.
07/27/2007 09:47:49 SSL Cert acceptable on host basis.
07/27/2007 09:47:49 SSL Cert acceptable on host basis.
07/27/2007 09:47:49 SSL Cert acceptable on host basis.
07/27/2007 09:47:49 SSL Cert acceptable on host basis.
07/27/2007 09:48:07 ..transferring Dat.R054665_10192 to
ushtinet01.hds.com/ht_http_row/Dat.R054665_10192
07/27/2007 09:48:10 ..transferring Environ.R054665_10192 to
ushtinet01.hds.com/ht_http_row/Environ.R054665_10192
07/27/2007 09:48:11 Transfer completed successfully

```

6. Repeat steps 3,4,5 for the Secondary FTP server

Sample entries (each site/region may vary)

Method Type: HTTPS

HDS host: ushtinet02.hds.com

Username: ht_http_row

Password: httpw0425

Proxy Server Type: HTTP Tunnel

Host or IP: webproxy.wlb.nsroot.net (or HTTP Application Proxy)

Port: 8080 (or HTTP Application Proxy port)

7. Delete modem-based entries. From the Center Tab, select one of the modem-based entries. Then select the *Remove* button. Verify the deletion confirmation dialog. Repeat for any other modem entries. Select *Ok* when done. *This will remove the analog modem as a dial-home method.*

23 Appendix I: HDS Tiered Storage Manager (Migrations)

23.1 What is HTSM?

HDS Tiered Storage Manager (HTSM) allows the transparent migration of data (no loss of access to the data during the migration) from one storage tier to another within the same array (Storage Domain).

Key features certified by Citigroup include:


- **Migration of data between tiers within an array without taking storage offline.** During the migration there is no loss of access to the data and the migration is transparent to the host, with the exception of degraded performance during the actual migration. All migration tasks are handled by the Storage Management team and no involvement is required from host or application owners.
- **Support for both thick and thin (Dynamically Provisioned) storage.**
- **Storage Tiers defined in terms of RAID type and disk size.** HTSM can manage Storage Tiers in terms of LDEV characteristics, rather than lists of devices. Once these definitions are setup, HTSM will automatically manage LDEV membership in each tier. Storage tiers are the targets of migrations, and are not a fixed number of LDEVs; they are a pool of devices.
- **Automatic erasure of release storage after migration.** Storage returned to a Storage Tier after a migration exercise can be automatically erased ready for reallocation to other hosts. For thin devices, erased storage will be released and returned to its DP pool for re-use.
- **Interactive Web Client interface.** Allows control and monitoring of HTSM from a browser. This is the interface that users are expected to use for most operations, as some tasks (e.g. migration target device selection within a storage tier) are interactive.
- **Command Line interface.** Alternative interface that can be used to write scripts for HTSM control. Tasks created with the Web Client can be executed via the CLI and vice versa. *The CLI is no longer supported by CSBE.*

23.2 Migration Steps


HTSM functionality has been integrated into HCS along with HDvM and no longer requires a separate GUI. The HTSM screens can be accessed from the HCS Tools menu or from the Mobility tab.

The major steps of migration for internal storage are:

1. In the Resources or Mobility tab, select a migration source host or logical group (and then the specific source volumes). You can also select migrate data from the General tasks on the bottom left of the screen. You can also right click the DP Pool and then select Migrate Data.



2. Click the Migrate Data button if not already selected.
3. Verify the correct migration source volumes are selected, and then click the Next button.



4. Confirm that all migration source volumes and target volumes (for thick devices) or DP pools (for thin devices) are paired properly and then click the Next button. If the pairs need to be modified, select the volume pair. Then from the Other Candidates for Migration Target table, select the target volume to change to and then select the Map Volume button.

Hitachi Command Suite - Windows Internet Explorer

http://ecctest01ms:23015/DeviceManager/WebService/index.jsp

File Edit View Favorites Tools Help

Hitachi Command Suite

Hitachi Command Suite 7

File > Migrate Data

1. Migration Source > 2. Migration Target > 3. Migration Pair > 4. Show Plan and Option Setting

Edit migration pairs

Warning: Some migration pairs contain volumes that will not be migrated. Confirm the volumes that will not be migrated, and then execute the migration task. (KAIC17027-W)

Migration Pairs

Select All Column Settings Rows/page: 30 Page: 1 / 1

Pair Status	Resource	Label	Location	Parity Group	Drive Type	Drive Speed (RPM)	Busy %	Max Busy %	Capacity	Pool	Resource Group
Move	Source	-	Internal	X210-1	-	-	-	-	33.87...	DP 99	-
	Target	-	-	-	-	-	-	-	33.87...	DP 100	-
Not Migrate	Source	-	Internal	X64-1	-	-	-	-	34.69...	DP 99	-
	Target	-	-	-	-	-	-	-	-	-	-

Selected: 1 of 2 Do Not Migrate

Other Candidates for Migration

Column Settings Rows/page: 30 Page: 1 / 1

Pool	Pool Name	Capacity	Used %	Used Threshold 1	Used Threshold 2	Subscription	Subscription Warning	Subscription Limit
DP 100		15.75 GB	44.00 %	70.00 %	80.00 %	44.00 %	Disabled	Disabled


Back Next Submit Cancel Help

Virtualize Volumes Selected: 0 of 0 Migrate Data Edit Tiering Policy Create Logical Group Edit Logical Group Delete Logical Group

Waiting: 0 In Progress: 0 Completed: 0 Failed: 0

Done Trusted sites 100%


- Specify options (shredding, ZPR, and email notification) then execute the task.



6. In the tasks list, confirm the execution result of the task. After a task is created, you must not perform operations on volumes included in the task until the task finished executing.

23.2.1 Example Task

The following is an example task of a partially successful task. The migration occurred, but the task failed because shredding failed.



23.2.2 Example Email Notification

If the email option is selected, the task will generate an email to the provided email destination. The following is an example email of a task that completed the migration, but failed on a scrubbing procedure:

```
From: gdse@citigroup.com [mailto:gdse@citigroup.com]
Sent: Monday, August 29, 2011 5:38 PM
To: Johnson, Gary G [CCC-OT_IT]
Subject: [TSM] A migration task has finished.(Task : Migrate-Volume-1)
```

The following migration task has finished:

```
Task : Migrate-Volume-1
Task Type : Migrate Volume
Status : Failed
Description : -
User : gj06529
Scheduled Time : 2011-08-29 17:51:18
Completed Time : 2011-08-29 18:38:12
Storage System : HDS10102_USPV
Migration Source (Host) : Unspecified
Migration Source (LG) : Unspecified
Migration Source (Pool) : DP 99
Migration Target (Tier) : -
Migration Target (Pool) : DP 100
Advanced Option : Shredding, Zero page reclaim
Migrated Volumes : 00:00:9E
```

This message was sent automatically by the Tiered Storage Manager.

To make migrations easier, storage tiers and logical groups can be defined prior to making any migrations.

Creation of Storage Tiers. These are groups of free (typically thick) storage to which migrations take place. These will usually be set up once when HTSM is configured for a Storage Domain. Once setup, Storage Tiers are “self-managing” i.e. as LDEVs no longer meet the tier definition requirements (e.g. as migrations and LDEV number swaps take place) they are automatically removed from the storage tier. Also, as LDEVs are allocated to hosts, or are placed into a Migration Group, they are automatically filtered out of the Storage Tier as they are no longer candidates for migration target volumes. One can use DP pool as a target instead of having to use a Storage tier.

Creation of a Logical Group. The Logical Group (LG) defines the source volumes for the migration. LGs can be populated with source by volumes by filtering on a wide range of LDEV attributes, including actual Device Number, Logical Group, I/O consumer (i.e. host), Volume Status (i.e. if used or allocated volumes), Host Storage Domain, etc., with further manual selection of LDEVs from these filtered searches. Very easy to use and flexible when undertaken with the web client.

23.2.3 Steps in the actual migration process (for DP vols targets only)

Migration is the act of relocating data from one storage subsystem volume to another. When Tiered Storage Manager performs this copying operation, **it swaps the LDEV numbers**. It is critical that this is understood. The following figure shows the steps in the migration, including the swapping of LDEV numbers.

The key points to note are:

- LDEV 1:01 begins in Parity Group A (or DP Pool A) but after migration is part of Parity Group B (or DP Pool B). Hence, sequential LDEV numbers may be in different Parity Groups (or DP Pools) when HTSM is used in an array.
- Although the Migration Source volume is 1:01 and the Migration Target Volume is 11:01, disk scrubbing takes place on device 11:01 (originally the target) as erasure occurs **AFTER** the LDEV numbers are swapped. This is important to understand if you ever want to execute erasure as a separate task to migration; you will need to erase the LDEVs that were originally the **target** LDEVs (not source) of the migration. It is recommended that erasure always be done automatically as part of the migration process to avoid mistakes that might be made with manual scrubbing.

23.3 HTSM Configuration

23.3.1 Licensing

The following licenses key need to be installed on the HDS Array:


- Tiered Storage Manager (apply via HCS)
- Volume Migrator2 (apply via Storage Navigator or Element Manager)
- Volume Shredder (apply via Storage Navigator or Element Manager)

As HTSM depends on Device Manager, this of course must also be installed and licensed.

Array keys must be obtained from HDS as they are locked to the HDS array serial number.

23.3.2 User Management setup

HTSM (like HRpM and HTnM) is considered different application for permission purposes. Each user should have its rights based on their role. The following is a sample permission list for a user.



There are four permissions are provided for performing HTSM operations. User management permission can be combined with any other permission. Admin rights are recommended for SAN administrators performing migrations. Note however, that the User Management permission does not allow you to change your own permissions; another user-id with User Management permission must be used.

Each user of HTSM should have their own account with the appropriate permissions for their role.

HTSM Permission	Description
Admin	Dedicated account permission for managing the storage domain. Permits user to manage storage domains and view all HTSM resources
Modify	Permits user to use all functions except creating and deleting storage domains. Needed for daily operation of HTSM.
Execute	Permits user to execute existing tasks and view all HTSM resources.
View	Permits user to view all HTSM resources

If a user does not have as permission for an operation, the relevant button or icon will not appear in the web client. CLI commands will fail, indicating permissions of executing user and required permissions.

23.3.3 Configuration of server.properties files

The following change should be made to the HTSM server properties file:

C:\Program Files\HiCommand\TieredStorageManager\conf\server.properties

server.migration.dataErase.defaultValue=true

Specifies whether to delete the source volume data by default after migration. This can always be overridden by the web client (or CLI) when the task is created.

server.migrationPlan.candidateVolumeCountLimit=false

The number of volumes that appear when creating a migration plan is usually limited. If you want to remove this limit to allow more specific selection of target volumes, then set this entry to false. This is an optional change.

The following optional change can be made if the HTSM CLI needs to be used:

```
server.rmi.secure=2
```

```
server.securty.port=24500
```

These changes allow the HTSM CLI client to connect using SSL communication. This is an optional setting as the CLI is no longer encouraged for use. Note that setting `server.rmi.secure=2` enables SSL communications but does not disable non-SSL communications via port 20352.

23.3.4 Audit Log Configuration - event log

The following change should be made to the HTSM logger properties file to enable logging of audit information (e.g. task creation/execution/deletion, users logging in, etc) to the Windows event log. This output level matches that used for the HTSM `serverMessagesN.log` files.

```
C:\Program Files\HiCommand\TieredStorageManager\conf\logger.properties
logger.sysloglevel=20
```

HTSM information will be logged as part of the HiCommand integrated logging, and hence the source in the event log for HTSM messages will be "HBase Storage Mgmt Log".

23.4 Citigroup HTSM Best Practice and Operational Notes

23.4.1 I/O loading during migrations

There are no specific limits on how much I/O can be sustained to a volume during migration. However, high write rates to the source volumes will elongate the migration can cause the migration to fail. Failures are typically due to rapid invalidation of blocks during the attempt to complete the pair swap (the final migration step). Extended sustained heavy I/O at this stage can cause the original swap attempt and subsequent retries to all fail, leading to a failure of the migration process; this final swap stage does required a period of reduced I/O to succeed .

Hence, migrations should be scheduled for periods of lower I/O. Application green zones would be ideal times.

If a migration is having an unacceptable impact on I/O performance, or taking too long it can be stopped. However, once a migration task is stopped. It cannot be restarted. It must be deleted and a new task created.

23.4.2 Delayed execution of a task

It is possible to create a task without immediately executing it. This has the advantage that the tasks can be created upfront and then executed at a more suitable time (e.g. during a green zone), or by another member of staff. Once a task is created, the target devices to which data is to be migrated are made unavailable and hence it is not possible to set up conflicting migrations.

However, there is a potential danger in creating a task and not executing immediately. Should the underlying storage configuration be modified after the task is created, then the data might be "migrated to an unexpected area". If you are using delayed migrations, you must be sure that the underlying storage is not being modified by another member of your team, or an HDS CE in the meantime.

23.4.3 Disk shredding

HDS uses the terms shredding and erasing interchangeably.

Under normal circumstances, disk shredding should be undertaken as part of the migration task. This is achieved by selecting the appropriate check box when using the web client to create a migration task. This will execute a ZERO-ONCE shredding pattern.

The PSE BISO has indicated that DoD shredding is not needed when erasing an LDEV after a migration operation. The default ZERO-ONCE option is acceptable.

DoD-style shredding (only possible with the CLI via a separate shredding task) would only need to be executed if disks are leaving Citi or ownership of the LUNs is moving to a business area on the other side of a Chinese Wall.

23.5 Using HTSM with remotely-replicated storage

At present, it is usually necessary to suspend replication before migrating a P-VOL and/or its corresponding S-VOL. The exception is when migrating TC P-VOLs (only). Following Citi testing, certain caveats have been identified, and the following Best Practice is recommended.

23.5.1 General

To use Tiered Storage Manager to migrate volume data, the migration source volumes and migration target volumes must satisfy the following conditions:

- If the volumes are configured as a TrueCopy or Universal Replicator pair, the CLPR of the migration source and target volumes must be the same.
- If the migration source volumes are normal (non-DP) volumes and configured as a TrueCopy or Universal Replicator pair, the migration target volumes must be normal (non-DP) volumes, or the migration must be to a DP pool. Migration from non-DP volumes to DP volumes is not supported.
- TrueCopy or Universal Replicator pair DP volumes cannot be migrated to normal (non-DP) volumes. The replication relationship must be removed first.

Migration Source	Migration Target	Supported with TC?	Supported with HUR?
Normal / Non-DP	Normal/Non-DP	For P-VOLs, TC Pairs must be in a PAIR, SUSPEND or SPLIT status For S-VOLs, TC Pairs must be in SUSPEND or SPLIT status	UR Pairs must be in a SUSPEND or SPLIT status
Normal / Non-DP	DP Vol	Not supported	Not supported
Normal / Non-DP	DP Pool	TC Pairs must be in a PAIR, SUSPEND or SPLIT status. ZPR cannot be used after migration (as this requires replication to be suspended). For S-VOLs, TC Pairs must be in SUSPEND or SPLIT status.	UR Pairs must be in a SUSPEND or SPLIT status.

DP Vol	Normal/Non-DP	Not supported	Not supported
DP Vol	DP Vol	<i>GDSE recommends migration to DP Pool instead.</i>	<i>CSBE recommends migration to DP Pool instead.</i>
DP Vol	DP Pool	For P-VOLs, TC Pairs must be in a PAIR, SUSPEND or SPLIT status. For S-VOLs, TC Pairs must be in SUSPEND or SPLIT status	UR Pairs must be in a SUSPEND or SPLIT status

Migration support matrix for remotely replicated storage

- Tiered Storage Manager is sometimes unable to determine whether some volumes can be used as migration sources or targets. Accordingly, a volume might not be suitable as a migration source or target even if **Yes** is displayed for **Can Migrate** and **Free** is displayed for **Volume Status**. For such volumes, the reason that migration is not possible is not displayed in **Migration Restrictions**. If such volumes are selected as migration sources, an error occurs when the migration task is executed. If such volumes are selected as migration targets, an error occurs when the migration task is created. The volumes that satisfy the conditions listed below are in this category.
 - a) The migration source is configured as a TrueCopy Synchronous or TrueCopy Asynchronous pair, and the migration target is a DP volume.
 - b) The migration source volume is configured as a Universal Replicator pair, and the migration target volume is a DP volume.
- From the time you create a task until the task ends, do not change the configuration of the volumes included in the task. If you change the volume configuration, the task might fail or produce unexpected results. If you perform any of the following operations for the migration source volume, migration will be interrupted:
 - a) A TrueCopy operation that changes the volume status to a status other than Suspended
 - b) A Universal Replicator operation that changes the volume status to COPY

23.5.2 TrueCopy Sync (TC) specific comments

- For TrueCopy Sync, migration of P-VOLs (only) can take place for volumes while still in PAIR state.
- If using an earlier version of HTSM or microcode, then for TC volumes, use the process described below for UR should be followed.
- P-VOL migration is supported while the TC Sync source volume is in a PAIR status. However, testing has shown that migration of the S-VOL is NOT supported while the TC Sync target volume (migration source) is in a PAIR status. To migrate a both a TC P-VOL and it associated S-VOL, TC replication must be suspended, and the process documented under UR used instead.
- If the migration source volume is a thick volume and part of a TC Sync pair, then the target can be a HDP Pool.

- If the migration source volume is a thin volume and part of a TC Sync pair, then the target **must be** a HDP Pool. The target volume **cannot be** a thick volume.

23.5.3 Hitachi Universal Replicator (HUR) specific comments

- HUR Pairs must be in a Suspend or Split status to be migrated. Migrations cannot take place for volumes in PAIR or COPY state with HTSM.
- Split off the ShadowImage copy of the S-VOL after replication is suspended but before migration is undertaken.
- Set up the migration in HTSM. Note that it is not possible to set up a migration task before replication is suspended and the DM and HTSM databases have been refreshed. Failure to refresh will result in the following error message if a migration is attempted:

“MG does not contain any migratable volumes”

- The P-VOLs and S-VOLs can be migrated in parallel to new volumes in each of their local arrays.
- Once both the P-VOL and S-VOL migrations have completed:
 1. Resume replication (pairresync or use replication manager)
 2. Refresh DM and then HTSM for P-VOL array
 3. Refresh DM and then HTSM for S-VOL array
 4. Once the replication state is PAIR, resync your ShadowImage copy of the S-VOL

24 Appendix J: SNMP Mib Definitions

The following table lists the covered (by ESS) Enterprise IDs from the USP/USPV Mib definition.

Class Name	Enterprise ID	Specific ID	Event Name	EventText: Summary as it appears on the console	Severity	State - (Notify / Clear)	Clear Trap (Yes / No)	Category	Event Type (Momentary / Durable)	Forward to Ops display?
SNMPTrap	.1.3.6.1.4.1.116.3.11.4.1.1.0.1	1	raidEventUseracute	\$V1\$ The impact of this event on the subsystem is acute.	Minor	Notify	No	Configuration	Durable	Yes
SNMPTrap	.1.3.6.1.4.1.116.3.11.4.1.1.0.2	2	raidEventUserserious	\$V1\$ The impact of this event on the subsystem is serious.	Critical	Notify	No	Configuration	Durable	Yes
SNMPTrap	.1.3.6.1.4.1.116.3.11.4.1.1.0.3	3	raidEventUsermoderate	\$V1\$ The impact of this event on the subsystem is moderate.	Major	Notify	No	Configuration	Durable	Yes
SNMPTrap	.1.3.6.1.4.1.116.3.11.4.1.1.0.4	4	raidEventUserservice	\$V1\$ The impact of this event on the subsystem is low.	Minor	Notify	No	Configuration	Durable	Yes
SNMPTrap	.1.3.6.1.4.1.116.3.11.4.1.1.0.100	.1.0.100	raidEventCommandFinished	\$V1\$ Required command was finished.	Normal	Clear	Yes	Misc	Momentary	Yes
Host	.1.3.6.1.4.1.116.3.11.1.2.0.1	.0.1	System down occurred.	\$V1\$ System Down	Critical	Notify	No	Configuration	Durable	Yes
SYS	.1.3.6.1.4.1.116.3.11.1.2.0.2	.0.2	driveFailure	\$V1\$ Drive blocking occurred.	Minor	Notify	No	Configuration	Durable	Yes
SYS	.1.3.6.1.4.1.116.3.11.1.2.0.3	.0.3	fanFailure	\$V1\$ Fan failure occurred.	Minor	Notify	No	Temperature	Durable	Yes
PowerSupply	.1.3.6.1.4.1.116.3.11.1.2.0.4	.0.4	powerSupplyFailure	\$V1\$ Power supply failure occurred.	Major	Notify	No	PowerSupply	Durable	Yes
PowerSupply	.1.3.6.1.4.1.116.3.11.1.2.0.5	.0.5	batteryFailure	\$V1\$ Battery failure occurred.	Minor	Notify	No	Resource	Durable	Yes
SYS	.1.3.6.1.4.1.116.3.11.1.2.0.6	.0.6	cacheFailure	\$V1\$ Cache memory failure occurred.	Major	Notify	No	Performance	Durable	Yes
PowerSupply	.1.3.6.1.4.1.116.3.11.1.2.0.7	.0.7	upsFailure	\$V1\$ UPS failure occurred.	Major	Notify	No	Resource	Durable	Yes
SYS	.1.3.6.1.4.1.116.3.11.1.2.0.8	.0.8	inboxFailure	\$V1\$ AC line or inbox failure occurred.	Major	Notify	No	PowerSupply	Durable	Yes
SYS	.1.3.6.1.4.1.116.3.11.1.2.0.9	.0.9	backupCircuitFailure	\$V1\$ Cache backup circuit failure occurred.	Major	Notify	No	Performance	Durable	Yes
SYS	.1.3.6.1.4.1.116.3.11.1.2.0.10	.0.10	otherControllerFailure	\$V1\$ Other controller failure occurred.	Major	Notify	No	Performance	Durable	Yes
SNMPTrap	.1.3.6.1.4.1.116.3.11.1.2.0.11	.0.11	warning	\$V1\$ Warning occurred.	Major	Notify	No	Misc	Durable	Yes
SYS	.1.3.6.1.4.1.116.3.11.1.2.0.12	.0.12	spareDriveFailure	\$V1\$ Spare drive failure occurred.	Minor	Notify	No	Availability	Durable	Yes
SYS	.1.3.6.1.4.1.116.3.11.1.2.0.13	.0.13	microprogramReplacementExecuted	\$V1\$ Microprogram Replacement Executed.	Minor	Notify	No	Availability	Durable	Yes
SYS	.1.3.6.1.4.1.116.3.11.1.2.0.14	.0.14	encFailure	\$V1\$ Enclosure controller failure occurred.	Minor	Notify	No	Performance	Durable	Yes
SYS	.1.3.6.1.4.1.116.3.11.1.2.0.15	.0.15	loopFailure	\$V1\$ Loop failure occurred.	Minor	Notify	No	Availability	Durable	Yes
Partition	.1.3.6.1.4.1.116.3.11.1.2.0.16	.0.16	pathFailure	\$V1\$ Path failure occurred.	Major	Notify	No	Availability	Durable	Yes
Partition	.1.3.6.1.4.1.116.3.11.1.2.0.21	.0.21	drivePreMaintenance	\$V1\$ The count has exceeded the drive preventive maintenance threshold.	Major	Notify	No	Performance	Durable	Yes
SYS	.1.3.6.1.4.1.116.3.11.1.2.0.200	.0.200	nasServerFailure	\$V1\$ NAS Server failure occurred.	Critical	Notify	No	Operational	Durable	Yes
SYS	.1.3.6.1.4.1.116.3.11.1.2.0.201	.0.201	nasPathFailure	\$V1\$ NAS Path failure occurred.	Major	Notify	No	Availability	Durable	Yes
SYS	.1.3.6.1.4.1.116.3.11.1.2.0.202	.0.202	nasUpsFailure	\$V1\$ NAS UPS failure occurred.	Major	Notify	No	Availability	Durable	Yes

25 Appendix K: Performance Monitor Data Extract Tool

25.1 What is Performance Monitor Data Extract Tool?

Performance Monitor Data Extract Tool is a utility that can pull performance data from a HDS storage array. This data is typically requested by HDS support staff in order to diagnose a performance related issue with an array.

25.2 Preparing to Use the Extract Tool

25.2.1 Preparing the Array

The array must be configured to save performance data in order for the extract tool to function.

25.2.2 Choosing a Server

The utility needs to be run from a server that has network connectivity to the SVP and can run Storage Navigator. A good server for this is usually the HDvM server. No fiber connectivity is required.

25.2.3 Temporary User

A "Storage Administrator" user (a role defined in Storage Navigator) must be used to extract and save all of the data. The username and password must be placed in the configuration file (command.txt). It is critical that this file be changed after use to prevent a security violation. Use of Windows encryption on the command.txt is recommended.

25.2.4 Software

Each version of microcode requires a different version of the executable. A separate package will be provided for each version of microcode. For information about the specific package names see [Software Packages](#).

The software will be installed into the following subdirectory:

C:\app_srv\instpkgsl <package name>

You can move the software with the subdirectories to any location desired.

25.3 Using the Extract Tool

25.3.1 Creating the Command File

The example command.txt file demonstrates most of the common commands:

```
.....
;Sample Command.txt file
.....
;
svpip 172.31.254.96           ;change IP to your SVP IP
login xx12345 "pass1234"      ;use the Storage Navigator login/pwd that used for exporting data
show                          ;displays the date ranges that will be used (short and long )

group PhyProc                 ;Statistics about channel processors, disk proc, and data recovery
proc
group PhyCSW                  ;Statistics about access paths, write pending rate, and cache
group PhyPG                   ;Statistics about parity groups
group PhyLDEV                 ;Statistics about volumes
group PhyExG                  ;Usage conditions about external volume groups
group PhyExLDEV               ;Usage conditions about external volumes
group Port                    ;Statistics about ports
group PortWWN                 ;Statistics about host bus adapters connected to ports
group PG                      ;Statistics about parity groups, external vol groups, or V-VOL
groups
```

group LDEV	;Statistics vols in parity groups, extern vol groups, or V-VOL
groups	
group LU	;Statistics about LUs
group PPCG	;Statistics about SPM groups
group PPCGWWN	;Statistics about host bus adapters belonging to SPM groups
outpath out	;This is the folder named out where the .zip files will go
option compress	;This basically says to zip output files
apply	

25.3.2 Optional Replication Commands

The following options can be added for replication environments:

group RemoteCopy	;Statistics about remote copy operations by TrueCopy
group RCLU	;Statistics about remote copy operations by TrueCopy
group RCLDEV	;Statistics about remote copy operations by TrueCopy
group UniversalReplicator	;Statistics about Universal Replicator
group URJNL	;Statistics about remote copy operations by Universal Replicator
group URLU	;Statistics about remote copy operations by Universal Replicator
group URLDEV	;Statistics about remote copy operations by Universal Replicator

25.3.3 Optional Data Ranges (Longrange and Shorrange Settings)

The shorrange and longrange subcommands specify the term of monitoring data that you want to export. Use these subcommands when you want to narrow the export target term (range) within the stored data. You can specify both the shorrange and longrange subcommands at the same time. The difference between these subcommands is as follows:

- The **shorrange** subcommand is valid for monitoring data in short range. You can use this subcommand to narrow the export-target term for all the monitoring items you can specify by the **group** subcommand.
- The **longrange** subcommand is valid for monitoring data in long range. You can use this subcommand only when you specify the PhyPG, PhyLDEV, PhyProc, or PhyCSW operand with the **Long** option in the **group** subcommand. (The items that can be saved by these operands are the monitoring data displayed in the **Physical** tab of the **Performance Monitor** window with selecting **longrange**.)

Examples:

Shorrange [[yyyyMMddhhmm][{+|-}hhmm]:[yyyyMMddhhmm][{+|-}hhmm]]

shorrange 200401010930:200401011730

The Export Tool saves monitoring data within the range of Jan. 1 2004, 9:30-17:30.

Longrange [[yyyyMMddhhmm][{+|-}ddhhmm]:[yyyyMMddhhmm][{+|-}ddhhmm]]

longrange 200401010930:200401011730

The Export Tool saves monitoring data within the range of Jan. 1 2004, 9:30-17:30.

If you run the Export Tool without specifying the shorrange or longrange subcommand, the monitoring data in the whole storing period (data in the period displayed by the **show** subcommand) will be exported.

25.3.4 Creating a Batch File

The example RunWin.bat file should be ready to go without significant modification.

```
Java -classpath ".\lib\JSanExport.jar;.\lib\JSanRmiServerSx.jar" -Xmx536870912 -Dmd.command=command.txt -
Dmd.logpath=log sanproject.getmondar.RJMdMain
pause
```

Note the reference to "command.txt". If you used another name for the command file, then you would reference it here. The line that starts with Java is a single line that wraps around the screen in the editor.

25.3.5 Running the Export Tool

To initiate the export tool, run the batch file that was create/modified above (i.e. runWin.bat).

When the Export Tool starts exporting monitoring data, dots (...) are issued to the command prompt. The dots increment as export processing continues. If an internal error occurs, the exclamation mark (!) is issued and the Export Tool attempts to restart exporting data (up to 3 times). If the export processing restarts, dots reappear and increment until export processing finishes.

When the Export Tool finishes successfully, monitoring data are compressed in ZIP-format archive files. If you want to obtain CSV files, you need to decompress ZIP files and extract CSV files out of the ZIP files.

25.3.6 Timing the Export Tool

The following is an estimate of the amount of time required to run the export tool. These times are based on an array with a normal workload with approximately 8,192 volumes and 12,288 LUs during a period of 24 hours.

Operand for the Group Command	Estimated Time
Port	5 minutes
PortWWN	5 minutes
PPCG	5 minutes
LDEV	60 minutes
LU	60 minutes

25.3.7 Error Codes from the Export Tool

The following chart lists the error codes returned by the Export Tool.

Termination Code	Meaning
0	The Export Tool finished successfully
1	An error occurred when the set subcommand (see set subcommand) executed, because an attempt to switch to Modify mode failed. Some other user might have been logged on in Modify mode.
2	An error occurred due to some reason unrelated to system option modes (i.e., View mode and Modify mode)
3	An error occurred due to more than one reason. One of the reasons is that an attempt to switch to Modify mode failed when the set subcommand (see set subcommand) executed. Some other user might have been logged on in Modify mode.
4	The user ID has none of permissions for Performance Monitor, TrueCopy and Universal Replicator.

Using these error codes can help with any scripting that is used to pull data on an automated basis. The following example is given to show how this can be used:

```
Java -classpath ".\lib\JSanExport.jar;.\lib\JSanRmiServerSx.jar" -Xmx536870912 -
Dmd.command=command.txt -Dmd.logpath=log sanproject.getmondar.RJMdMain
if %errorlevel%==1 echo THE SET SUBCOMMAND FAILED
if %errorlevel%==3 echo THE SET SUBCOMMAND FAILED
pause
```

25.3.8 Optional Data Ranges (Longrange and Shortrange Settings)

The shortrange and longrange subcommands specify the term of monitoring data that you want to export. Use these subcommands when you want to narrow the export target term (range) within the stored data. You can specify both the shortrange and longrange subcommands at the same time. The difference between these subcommands is as follows:

- The `shortrange` subcommand is valid for monitoring data in short range. You can use this subcommand to narrow the export-target term for all the monitoring items you can specify by the `group` subcommand.
- The `longrange` subcommand is valid for monitoring data in long range. You can use this subcommand only when you specify the `PhyPG`, `PhyLDEV`, `PhyProc`, or `PhyCSW` operand with the `Long` option in the `group` subcommand. (The items that can be saved by these operands are the monitoring data displayed in the **Physical** tab of the **Performance Monitor** window with selecting **longrange**.)

Shortrange `[[yyyyMMddhhmm]][{+|-}hhmm]:[yyyyMMddhhmm][{+|-}hhmm]`

`shortrange 200401010930:200401011730`

The Export Tool saves monitoring data within the range of Jan. 1 2004, 9:30-17:30.

Longrange `[[yyyyMMddhhmm]][{+|-}ddhhmm]:[yyyyMMddhhmm][{+|-}ddhhmm]`

`longrange 200401010930:200401011730`

The Export Tool saves monitoring data within the range of Jan. 1 2004, 9:30-17:30.

If you run the Export Tool without specifying the shortrange or longrange subcommand, the monitoring data in the whole storing period (data in the period displayed by the `show` subcommand) will be exported.

25.4 Transmitting the Data to HDS

The performance data should be sent to HDS via the TUF Website (<https://tuf.hds.com>). You will need a case number from HDS support in order to be able to upload the data (upload password is "truenorth"). You can upload a single file or multiple files. You will receive an email once the file upload is complete.

26 Appendix L: Server Priority Manager (QoS)

Server Priority Manager is a Quality of Service (QoS) function within Storage Navigator that enables the storage manager to limit the amount of data that a given host can send to the array. This is a valuable tool for limiting the impact of one or more hosts may have on the whole array or on other others. This could be used to manage:

- An unusual application condition that is causing a high degree of random writes (thus filling up write cache)
- A development application that is affecting performance of production applications
- An application that is processing large amounts of data on a periodic basis (for example a data warehouse that is transforming and loading data)
- A server that is sending large amounts of data that is overwhelming the host port (or processors) and is limiting the ability of other applications use of the port
- A very high write pending cache percentage that is affecting the entire array

26.1 Considerations and Risks

There are some risks and important considerations associated with the use of this feature that should be examined before any implementation:

- It is imperative that you have the correct WWNs of the host that needs to be limited. Using the wrong WWN will result in the wrong host being affected.
- If you over-restrict a host, you may cause an I/O failure on that host (timeout).
- By default all hosts (WWNs) and/or ports have the same priority. SPM will identify hosts (or ports) that should have a lower priority.
- Limiting host performance should be done by WWN rather than port if possible. Port level restrictions will affect all hosts (WWNs) attached to the array port.

26.2 Setup and Use of SPM

The following are the steps required to setup and use SPM for the TagmaStore Series Arrays:

1. Log into the array via Storage Navigator (Element Manager)
2. Under Performance Management, go to the Monitoring Options tab. While in Modify mode, change the current status to Enable (if not already set).
3. Go to the Performance Management (1st) tab. On the left side there is another set of smaller tabs. Select the Port-LUN tab. Select the SPM button on the right side of the screen. The Server Priority Manager dialog will appear. There is a Port tab and WWN tab to limit hosts by port (all hosts on a port) or by WWN. WWN is recommended because is specific to a host.
4. Before you can limit a given host (by WWN), you must move it from the “non-monitor” list to the “monitor” list for each port. By default, all hosts are in a non-monitor state. To move a WWN, select it and then drag the WWN to the Monitor list. You will be prompted for an SPM name (this should be the hostname or similar identifying information). Select Apply. The current performance numbers for the WWN should appear. This should be repeated for any other desired WWNs from the host.

5. To limit a given host, right-click the entry and select Prio ->> Non-Prio. Under thresholds, select the checkbox and fill in a maximum value for IOPS. Alternatively, the value can be provided in MB/s (select MB/s instead of IOPS). Hit Apply.
6. The limiting value will be implemented slowly over a period of several minutes. While the host(s) is in a Non-Priority mode,

From the host perspective:

- Response time will gradually increase
- IOPS (or MB/s) throughput will gradually decrease

From the array perspective:

- The write pending % should gradually decrease, depending on the amount of impact this host had on the overall write-cache.

7. If the desired effect on the array or other hosts is not achieved,
 - the limiting value may need to be decreased to a lower value
 - a different host (or more hosts) should be limited

27 Appendix M: Tuning Manager (HTnM)

Tuning Manager is a powerful reporting engine that collects detailed data about the arrays and presents it in a meaningful format. In addition, Tuning Manager can examine the data and generate alerts based on data that falls outside of defined criteria. These alerts can be email, SNMP traps, or script-based actions.

27.1 Installation of Software

Tuning Manager should be installed on the same server as HCS and is typically setup at the same time HCS is configured. Command device access is currently limited to non-vmware based hosts so the HTnM Agent for RAID will need to reside on a separate host. This host is typically a standalone RHEL server (legacy configurations include Solaris 10 and Windows). The RHEL server also runs TMEA and Raidcom as they also require command device access. Each application should have its own dedicated command device for a given array.

27.2 Tuning Manager Setup / Configuration

Tuning Manager is a large reporting and analysis application that requires agents to function. In order to pull data from arrays, an array agent (agent for RAID) must be created to retrieve and store performance information. The array should be configured to save local performance data. For instructions on how to set this feature in Storage Navigator, see section [Configuring Performance Data Collection](#).

Note: Citi does not use the host or switch agents at this time.

27.2.1 Creating Array Agents (Agents for RAID)

RAID agents can run on the Tuning Manager server or on other servers. These agents should run on servers that are local to the array (not over WAN). The Tuning Manager server could connect to the agents over the WAN, but ideally each site would have its own Tuning Manager instance.

RAID agents need fiber based (in-band) access to the array. Therefore any server that is running the agent for RAID would need a command device provisioned to it to gather information about the array. The RAID agent can reside on the Tuning Manager server or it can reside on a different server (like an existing CLI server that already has access to command devices / gatekeepers).

27.2.1.1 Setup of Agent for RAID on Windows

The HTnM server and the server that will run the agent need to be time synchronized (NTP).

Install the RAID agent software on the server that has access to the command device. The command device (like an EMC gatekeeper) needs to be allocated and visible to the host that will run the agent. You will need to access the command device via the SAN. Use `jpctdlistraid` to list the command devices visible to the server.

```
C:\Program Files\HiCommand\TuningManager\jp1pc\tools>jpctdlistraid
KAVF18700-I The detection of the monitorable storage subsystem has begun.
"PRODUCT"          , "SERIAL"   , "LDEV"      , "SLPR", "PORT" , "DEVICE_FILE"
, "VOLUME_GUID"
"USP V"             , "10102"   , "00:3E:00", "0"   , "CL1-A", "\\.\PhysicalDrive1"
, ""
```

```
"USP" , "10192" , "3E:40" , "0" , "CL1-A", "\\.\PhysicalDrive2"
, ""
KAVF18701-I The detection of the monitorable storage subsystem has ended.
```

For windows, there should be a (primary) partition on the command device. Do not assign it a drive letter (do not mount it; do not assign it to a folder). Do not format the partition. Use `jpctdlistraid` to list the command devices visible to the server. If no partition has been created, then the last column (Volume GUID) will not have a value (as above). Once the partition has been created, a valid GUID should appear. GUID will be used to refer to the device (will account for changes to disk order on server). Deleting the partition will remove the GUID.

```
C:\Program Files\HiCommand\TuningManager\jp1pc\tools>jpctdlistraid
KAVF18700-I The detection of the monitorable storage subsystem has begun.
"PRODUCT" , "SERIAL" , "LDEV" , "SLPR", "PORT" , "DEVICE_FILE"
, "VOLUME_GUID"
"USP V" , "10102" , "00:3E:00", "0" , "CL1-A", "\\.\PhysicalDrive1"
, "\\?\Volume{3a5ccf4e-0feb-11dd-b4ef-000bcdffd04d}"
"USP" , "10192" , "3E:40" , "0" , "CL1-A", "\\.\PhysicalDrive2"
, "\\?\Volume{3a5ccf50-0feb-11dd-b4ef-000bcdffd04d}"
KAVF18701-I The detection of the monitorable storage subsystem has ended.
```

Create an instance of the Agent For RAID for every array attached to this host. To create the agent instance, the `jpminssetup` command should be used (if this command is missing then you have not installed the RAID agent software). When creating the RAID agent instance, you specify the instance name. The instance name should be the numerical serial number of the array. You will then have a series of questions that will be asked interactively (Storage Model=2, Command Device= local drive for command device – the GUID is recommended, Unassigned Open Volume Monitoring =N, Mainframe Volume Monitoring=N, Store Version =2.0).

```
C:\Program Files\HiCommand\TuningManager\jp1pc\tools>jpminssetup agtd -inst
99999
Storage Model (1:Thunder/AMS or 2:Lightning/USP/USP V) : 2
Command Device File Name [\\.\PhysicalDrive0] : \\?\Volume{3a5ccf4e-0feb-
11dd-b4ef-000bcdffd04d}
Unassigned Open Volume Monitoring (Y/N) [N] : N
Mainframe Volume Monitoring (Y/N) [N] : N
Store Version [2.0] : 2.0
KAVE05080-I The instance environment is now being created. (servicekey=agtd,
inst=99999)
KAVE05081-I The instance environment has been created. (servicekey=agtd,
inst=99999)
```

To obtain a list of the configured agent instances, use the `jpminslist` command. You must specify the agent type (RAID agent – `agtd`)

```
C:\Program Files\HiCommand\TuningManager\jp1pc\tools>jpminslist agtd
10102
10192
99999
```

To check the status of an agent instance, use the `jpctdchkinst` command should be used. You must specify the agent instance number.

```
Working agent instances:
C:\Program Files\HiCommand\TuningManager\jp1pc\tools>jpctdchkinst -inst 10102
KAVF18800-I The verification of the agent instance settings will now start.
(instance name=10102)
[Instance parameters]
Storage Model : Lightning/USP/USP V
Command Device File Name : \\?\Volume{3a5ccf4e-0feb-11dd-b4ef-000bcdffd04d}
Unassigned Open Volume Monitoring : y
Mainframe Volume Monitoring : n
[Check result]
```

```

KAVF18850-I No error was found during verification of the agent instance
setting.
[Monitored storage subsystem information]
PRODUCT   : USP V
SERIAL    : 10102
FIRMWARE  : 60-02-25/00
SLPR      : 0
KAVF18801-I The verification of the agent instance settings will now end.

Non-working agent instance:
C:\Program Files\HiCommand\TuningManager\jp1pc\tools>jpctdchkinst -inst 99999
KAVF18800-I The verification of the agent instance settings will now start.
(instance name=99999)
[Instance parameters]
Storage Model : Lightning/USP/USP V
Command Device File Name : \\.\PhysicalDrive99
Unassigned Open Volume Monitoring : N
Mainframe Volume Monitoring : N
[Check result]
KAVF18851-E An attempt to access the device set by the agent instance
parameter has failed. (parameter name=Command
Device File Name, parameter value=\\.\PhysicalDrive99
KAVF18801-I The verification of the agent instance settings will now end.

```

Non-working instances are typically not correct due to the wrong or missing command device.

To check the status of the communication processes, run the jpcctrl command.

```

Correctly functioning processes:
C:\Program Files\HiCommand\TuningManager\jp1pc\tools>jpcctrl list "*"
Host Name      ServiceID      Service Name    PID      Port      Status
-----
ecctest01ms    PT1ecctest01ms  Status Server   2328     22350     Active
ecctest01ms    PN1001         Name Server     4396     22285     Active
ecctest01ms    PM1001         Master Manager  5048     1154      Active
ecctest01ms    PS1001         Master Store    4872     1175      Active
ecctest01ms    PE1001         Correlator      4608     1192      Active
ecctest01ms    PC5ecctest01ms  Trap Generator  6384     1225      Active
ecctest01ms    PP1ecctest01ms  View Server     6636     1268      Active
ecctest01ms    PH1ecctest01ms  Action Handler  5156     1200      Active
ecctest01ms    DS110102[ecctest01ms] Agent Store    6540     4715      Busy
ecctest01ms    DS110192[ecctest01ms] Agent
Store          5408     1194      Active
ecctest01ms    DA110102[ecctest01ms] Agent
Collector      6856     4720      Active
ecctest01ms    DA110192[ecctest01ms] Agent
Collector      6308     1221      Active
KAVE06003-I List processing of the service information terminated normally.

Non-working example:
C:\Program Files\HiCommand\TuningManager\jp1pc\tools>jpcctrl list *
Host Name      ServiceID      Service Name    PID      Port      Status
-----
ecctest01ms    PN1001         Name
Server         Inactive*
ecctest01ms    PT1ecctest01ms  Status
Server         Inactive*
ecctest01ms    PC5ecctest01ms  Trap
Generator      Inactive*
ecctest01ms    PP1ecctest01ms  View
Server         Inactive*
ecctest01ms    PM1001         Master
Manager        Inactive*
ecctest01ms    PS1001         Master
Store          Inactive*
ecctest01ms    PE1001         Correlator
*

```

```

ecctest01ms PH1ecctest01ms Action
Handler Inactive*
KAVE06021-W The detailed information cannot be displayed because Status Server
i
s not running. (host=ecctest01ms)
KAVE06003-I List processing of the service information terminated normally.

```

A reboot of the HTnM and/or RAID agent server(s) may be required to synchronize the processes.

To manually start the processes, run the `jpcstart` command.

```

C:\Program Files\HiCommand\TuningManager\jp1pc\tools>jpcctrl list "*"
Host Name      ServiceID      Service Name    PID      Port      Status
-----
ecctest01ms    PT1ecctest01ms Status Server    13244    22350    Active
ecctest01ms    PN1001         Name Server      456      22285    Active
ecctest01ms    PM1001         Master Manager   10236    4686     Active
ecctest01ms    PS1001         Master Store     8336     4772     Active
ecctest01ms    PE1001         Correlator       15576    4811     Active
ecctest01ms    PC5ecctest01ms Trap Generator   16108    4819     Active
ecctest01ms    PP1ecctest01ms View Server      16076    4893     Active
ecctest01ms    PH1ecctest01ms Action Handler   Inactive
ecctest01ms    DS110102[ecctest01ms] Agent Store
Inactive
ecctest01ms    DS110192[ecctest01ms] Agent Store
Inactive
ecctest01ms    DA110102[ecctest01ms] Agent Collector
Inactive
ecctest01ms    DA110192[ecctest01ms] Agent Collector
Inactive
KAVE06003-I List processing of the service information terminated normally.

C:\Program Files\HiCommand\TuningManager\jp1pc\tools>jpcstart all
KAVE06007-I The service will now start. (service=Action Handler)
KAVE06007-I The service will now start. (service=Agent Store for RAID,
inst=10102)
KAVE06007-I The service will now start. (service=Agent for RAID, inst=10102)
KAVE06007-I The service will now start. (service=Agent Store for RAID,
inst=10192)
KAVE06007-I The service will now start. (service=Agent for RAID, inst=10192)

C:\Program Files\HiCommand\TuningManager\jp1pc\tools>jpcctrl list "*"
Host Name      ServiceID      Service Name    PID      Port      Status
-----
ecctest01ms    PT1ecctest01ms Status Server    13244    22350    Active
ecctest01ms    PN1001         Name Server      456      22285    Active
ecctest01ms    PM1001         Master Manager   10236    4686     Active
ecctest01ms    PS1001         Master Store     8336     4772     Active
ecctest01ms    PE1001         Correlator       15576    4811     Active
ecctest01ms    PC5ecctest01ms Trap Generator   16108    4819     Active
ecctest01ms    PP1ecctest01ms View Server      16076    4893     Active
ecctest01ms    PH1ecctest01ms Action Handler   10232    4773     Active
ecctest01ms    DS110102[ecctest01ms] Agent Store    14192    4835
Active
ecctest01ms    DS110192[ecctest01ms] Agent Store    12112    4943
Active
ecctest01ms    DA110102[ecctest01ms] Agent Collector 15252    4848
Active
ecctest01ms    DA110192[ecctest01ms] Agent Collector 1376     4956
Active
KAVE06003-I List processing of the service information terminated normally.

```

To delete an agent instance, use the `jpcinsunsetup` command. You will need to specify the agent type and instance name.

```
C:\Program Files\HiCommand\TuningManager\jp1pc\tools>jpcinsunsetup agtd -inst
99999
KAVE05083-I The instance environment is being deleted. (servicekey=agtd,
inst=99999)
KAVE05084-I The instance environment was deleted. (servicekey=agtd,
inst=99999)
```

27.2.1.2 Setup of Agent for RAID on RHEL (or Solaris 10)

Typically this setup is performed on a RHEL (or Solaris 10) host when the HCS / HTnM server does not have SAN connectivity to a given array.

The HCS / HTnM server and the server that will run the agent need to be time synchronized (NTP).

Install the RAID agent software on the server that has access to the command device.

The RHEL (and Solaris) package for the HTnM Agent For RAID software does not start the install. There are two pieces of software to install in Solaris: Agent for RAID and the RAID Manager Library (RMLIB).

Note: For initial installations, there will be a warning message about not being able to install Tuning Manager. The warning message is about installing the Tuning Manager database and should be ignored.

Tuning Manager (also Raidcom and TMEA Collector) requires the installation of RMLIB to be able to communicate with HDS arrays. On UNIX systems, RMLIB is a separate installation.

To install the RAID Manager Library (RMLIB):

Run RLinst.sh from the package directory (you have to provide the directory where to place the RMLIB directory. This directory must already exist. It is not necessary to update the RAID Manager Library unless specific errors are encountered. Leaving the existing version is recommended for agent upgrades.

```
[/mnt/gj06529/HDS/HTnM_Sol/Agent_for_RAID]
sangs6# mkdir /opt/HTnM

[/mnt/gj06529/HDS/HTnM_Sol/Agent_for_RAID]
sangs6# RLinst.sh /opt/HTnM
***** Confirmation for New Introduction of the
RMLIB.*****
Please specify a directory(recommends except '/') for the installation.
For continue -> please enter a 'directory'.
For cancel -> please enter 'exit'.
/opt/HTnM
cpio -idmu < /mnt/gj06529/HDS/HTnM_Sol/Agent_for_RAID/SOLARIS/RMLIB
1264 blocks
The following model was installed to '/opt/HTnM/RMLIB' under as
'/usr/lib/RMLIB'.
When you have to uninstall of the RMLIB,please executes the following RUninst
command.
When you also have to be installing by floppy or tape,please executes the
following RLinstsh command.
[ Model :RAID-Manager/LIB/Solaris Ver&Rev:01-15-03/04 for RMLIB ]
-----
--
bin                lib                rmlibuninstall.sh
include            rmlibinstall.sh    sample
-----
```

If the RLibinst.sh script is not available, then run the following CPIO command to extract the files to the desired location. It will force a destination directory name of RMLIB, so initiate the CPIO extract from the /opt directory.

```
[root@namgdse17 opt]# cpio -idmu < /mnt/gj06529/HDS/RMLIB
1999 blocks

[root@namgdse17 opt]# cd RMLIB
[root@namgdse17 RMLIB]# ls
bin include lib RLSYSLK rmlibinstall.sh rmlibuninstall.sh sample

[root@namgdse17 RMLIB]# ls -al
total 32
drwxr-xr-x  6 root sys  4096 Nov 18 11:52 .
drwxr-xr-x 41 root root 4096 Nov  6 17:51 ..
drwxr-xr-x  2 root sys  4096 Nov 18 11:52 bin
drwxr-xr-x  2 root sys  4096 Nov 18 11:52 include
drwxr-xr-x  2 root sys  4096 Nov 18 11:52 lib
-rw-r--r--  1 root root    0 Oct  2 14:24 RLSYSLK
-r-xr-xr-x  1 root sys  4006 Apr 10 2012 rmlibinstall.sh
-r-xr-xr-x  1 root sys 3284 Apr 10 2012 rmlibuninstall.sh
drwxrwxrwx  2 root sys  4096 Nov 18 11:52 sample
```

Run /opt/HTnM/RMLIB/rmlibinstall.sh after the setup to make sure the symbolic link is in place.

```
[/opt/HTnM/RMLIB]
sangs6# rmlibinstall.sh
The symbolic link of '/usr/lib/libsvrrm.so' was successful.

If previously run:
[root@namgdse17 RMLIB]# ./rmlibinstall.sh
The symbolic link of '/usr/lib/libsvrrm.so' was successful.
The symbolic link of '/lib64/libsvrrm.so' was successful.
ln: creating symbolic link '/usr/lib/RMLIB/bin/whatrmver' to
'/usr/lib/RMLIB/bin/whatrmverx64': File exists
```

To verify the symbolic links are functional and that RMLIB is working, run the version command (/opt/RMLIB/bin/whatrmver):

```
[root@namgdse17 bin]# cd /opt/RMLIB/bin
[root@namgdse17 bin]# ./whatrmver
Model :RAID-Manager/LIB/Linux
Ver&Rev:01-17-03/01
```

You can uninstall the software by running rmlibuninstall.sh from the /opt/RMLIB directory.

To install or upgrade the Agent for RAID on RHEL or Solaris:

You must run this script by executing the full path. You may also need to give it a temporary directory to use for unpacking. You should execute the borne shell before starting. (The following example is an upgdgrade):

```
#!/bin/sh

# /mnt/gj06529/HDS/HTnM_Sol/Agent_for_RAID/install.sh
Begin installation of Hitachi Tuning Manager - Agent for RAID - 7.0.0(7.0.0-00)
Fri Jun 10 16:31:58 EDT 2011
> A current version [085003]
> A premise version [091003]
> Install version ok [ 060000 -> 070000 ]
> Shutdown all Tuning Manager processes.
```

```

    The service will now stop.
KAVE06008-I The service will now stop. (service=Agent for RAID,
inst=uspv10102)
KAVE06008-I The service will now stop. (service=Agent Store for RAID,
inst=uspv10102)
KAVE06008-I The service will now stop. (service=Agent for RAID,
inst=uspv25519)
KAVE06008-I The service will now stop. (service=Agent Store for RAID,
inst=uspv25519)
KAVE06008-I The service will now stop. (service=Action Handler)
KAVE06008-I The service will now stop. (service=Status Server)
> /opt/jp1pc/tools/jpcstop: Completed.
Installing The Common Files.

Processing package instance <HTMRaid> from
</mnt/gj06529/HDS/HTnM_Sol/Agent_for_RAID>

Hitachi Tuning Manager - Agent for RAID(sparc) 7.0.0.0
Hitachi, Ltd.
## Executing checkinstall script.
>
#####
> # : RAID Agent for Solaris
> # Script starts: before installation script
> # date: Fri Jun 10 16:34:33 EDT 2011
> # Script completed: before installation script
>
#####
Using </opt> as the package base directory.
## Processing package information.
## Processing system information.
1 package pathname is already properly installed.

Installing Hitachi Tuning Manager - Agent for RAID as <HTMRaid>

## Executing preinstall script.
## Installing part 1 of 1.
/opt/jp1pc/HTMAGT4D.tar
[ verifying class <none> ]
## Executing postinstall script.
>
#####
> # : RAID Agent for Solaris
> # Script starts: after installation script
> # date: Fri Jun 10 16:34:36 EDT 2011
> mkdir /opt/jp1pc/jp1pc_AGTDTMP: Completed.
x ., 0 bytes, 0 tape blocks
x ./opt, 0 bytes, 0 tape blocks

...output is abbreviated...

> Change Agent File Permission: Completed.
> # Script completed: after installation script
>
#####

Installation of <HTMRaid> was successful.
Start IPv6 Setup.
KAVE05932-I The function for IPv6 communication will now be set up. (enable)
KAVE05933-I The function for IPv6 communication was successfully set up.
(enable)
End IPv6 Setup.
Fri Jun 10 16:37:58 EDT 2011
Done installing Hitachi Tuning Manager - Agent for RAID.

```

If there is a problem upgrading the agent for raid, then uninstalling the software and then reinstalling it is the best course of action.

To uninstall the Agent for RAID, run the uninstall.sh script

```
sangs10# /mnt/gj06529/HDS/HTnM_Sol/Agent_for_RAID/uninstall.sh
Removing Hitachi Tuning Manager - Agent for RAID - 7.4.0(7.4.0-00)
Wed Feb 13 15:57:50 EST 2013

Do you wish to start the removal?
To continue, enter Y. (Default Y) Y/N > Y
>
#####
> #          :   RAID Agent for Solaris
> # Script starts: removal script
> # date:      Wed Feb 13 15:57:57 EST 2013
> Shutdown all Tuning Manager processes.
>   pfm_inst delhains
> Remove all instance.
uspv10102
vsp53112
KAVE05083-I The instance environment is being deleted. (servicekey=agtd,
inst=uspv10102)
KAVE05084-I The instance environment was deleted. (servicekey=agtd,
inst=uspv10102)
KAVE05083-I The instance environment is being deleted. (servicekey=agtd,
inst=vsp53112)
KAVE05084-I The instance environment was deleted. (servicekey=agtd,
inst=vsp53112)
> # Script completed: removal script
>
#####
>
#####
> #          :   RAID Agent for Solaris
> # Script starts: removal script
> # date:      Wed Feb 13 15:58:16 EST 2013
> Start removing installed PP files.
#   /bin/rm -rf /opt/jp1pc/agtd
> Completed.
#   /bin/rm -f /opt/jp1pc/instagtd.ini
> Completed.
> Completed.
#   /bin/rm -rf /opt/jp1pc/tools/jpctdchkinst
#   /bin/rm -rf /opt/jp1pc/tools/jpctdlistraid
#   /bin/rm -rf /opt/jp1pc/tools/jpctdrefresh
> Finished removing installed files.
> # Script completed: removal script
>
#####

Removal of <HTMRaid> was successful.
Removing The Common Files.
Wed Feb 13 15:58:20 EST 2013
Done removing Hitachi Tuning Manager - Agent for RAID.
```

The command device (like an EMC gatekeeper) needs to be allocated and visible to the host that will run the agent. You will need to access the command device via the SAN.

Use `jpctdlistraid` to list the command devices visible to the server.

```
[root@namgdse04 tools]# ./opt/jp1pc/tools/jpctdlistraid
KAVF18700-I The detection of the monitorable storage system has begun.
"PRODUCT"          , "SERIAL"      , "LDEV"      , "SLPR" , "PORT" , "DEVICE_FILE"
"USP V"            , "10102"      , "00:3E:06" , "0"    , "CL1-A" , "/dev/sda"
"USP V"            , "10102"      , "00:3E:09" , "0"    , "CL1-A" , "/dev/sdb"
"USP V"            , "10102"      , "00:3E:17" , "0"    , "CL1-A" , "/dev/sdc"
"VSP"              , "53112"      , "00:3E:17" , ""     , "CL1-A" , "/dev/sdu"
"VSP"              , "53112"      , "00:3E:16" , ""     , "CL1-A" , "/dev/sdt"
"VSP"              , "53112"      , "00:3E:18" , ""     , "CL1-A" , "/dev/sdv"
"USP V"            , "10102"      , "00:3E:09" , "0"    , "CL2-A" , "/dev/sdai"
"USP V"            , "10102"      , "00:3E:06" , "0"    , "CL2-A" , "/dev/sdah"
```

```
"USP V"      , "10102"    , "00:3E:17", "0"    , "CL2-A", "/dev/sdaj"
"VSP"       , "53112"    , "00:3E:17", ""     , "CL2-A", "/dev/sdbb"
"VSP"       , "53112"    , "00:3E:16", ""     , "CL2-A", "/dev/sdba"
"VSP"       , "53112"    , "00:3E:18", ""     , "CL2-A", "/dev/sdbc"
"HUS VM"    , "10060"    , "00:01:68", ""     , "CL1-A", "/dev/sdae"
"HUS VM"    , "10060"    , "00:01:69", ""     , "CL1-A", "/dev/sdaf"
"HUS VM"    , "10060"    , "00:01:FC", ""     , "CL1-A", "/dev/sdag"
"HUS VM"    , "10060"    , "00:01:68", ""     , "CL2-A", "/dev/sdbe"
"HUS VM"    , "10060"    , "00:01:69", ""     , "CL2-A", "/dev/sdbf"
"HUS VM"    , "10060"    , "00:01:FC", ""     , "CL2-A", "/dev/sdbg"
KAVF18701-I The detection of the monitorable storage system has ended.
```

Note: If this host is also used as a cli server for other vendor arrays or if there are non-command devices allocated, then you may see errors:

```
KAVF18742-W An error occurred during system call or OS command execution.
(name=ioctl(SG_IO), rc=-1, error=25)
KAVF18743-W An attempt to get the device information has failed. (device
name=/dev/sda1)
KAVF18742-W An error occurred during system call or OS command execution.
(name=ioctl(SG_IO), rc=-1, error=25)
KAVF18743-W An attempt to get the device information has failed. (device
name=/dev/sdak3)
KAVF18742-W An error occurred during system call or OS command execution.
(name=ioctl(SG_IO), rc=-1, error=25)
KAVF18743-W An attempt to get the device information has failed. (device
name=/dev/sdak7)
```

To account for the lack of persistence on REHL, you should use the scsi id names for the devices rather than the /dev/sdxxx format. Take te command devices identified above and search the by-id directory:

```
ls -al /dev/disk/by-id | grep "3e09"
[root@namgdse04 by-id]# ls -al | grep 3e09
lrwxrwxrwx 1 root root 10 Dec 3 18:15 scsi-
360060e80052776000000277600003e09 -> ../../sdai
lrwxrwxrwx 1 root root 10 Dec 3 18:15 wwn-
0x60060e800527760000000277600003e09 -> ../../sdai
```

This address format (i.e. /dev/disk/by-id/scsi-360060e80052776000000277600003e09) is compatible with HTnM.

Create an instance of the Agent for RAID for every array attached to this host. To create the agent instance, the `jpcinssetup` command should be used. When creating the RAID agent instance, you specify the instance name. The instance name should be the numerical serial number (or alphanumeric name + serial) of the array. You will then have a series of questions that will be asked interactively (Storage Model=2, Command Device= /dev/rdsd/..., Unassigned Open Volume Monitoring = Y, Mainframe Volume Monitoring=N, Store Version =2.0). Take the defaults other than Unassigned Volume Monitoring. *The Unassigned Volume Monitoring (set to Y) option will include the Pool volumes that are assigned to a pool but not any hosts.* This value (and other agent settings) can be modified without loss of data by using the `jpcinssetup` command (see below).

```
[root@namgdse04 tools]# ./jpcinssetup agtd -inst 10102
Storage Model (1:Thunder/AMS/HUS100 or 2:Lightning/USP/USP V/VSP/HUS VM) : 2
Command Device File Name [/dev/sda] : /dev/disk/by-id/wwn-
0x60060e800527760000000277600003e09
Unassigned Open Volume Monitoring (Y/N) [N] : Y
Mainframe Volume Monitoring (Y/N) [N] : N
Store Version [2.0] :
```

```

KAVE05080-I The instance environment is now being created. (servicekey=RAID,
inst=10102)
KAVE05081-I The instance environment has been created. (servicekey=RAID,
inst=10102)
[root@namgdse04 tools]# ./jpcinssetup agtd -inst 53112
Storage Model (1:Thunder/AMS/HUS100 or 2:Lightning/USP/USP V/VSP/HUS VM) : 2
Command Device File Name [/dev/sda] : /dev/disk/by-id/scsi-
360060e80052776000000277600003e17
Unassigned Open Volume Monitoring (Y/N) [N] : Y
Mainframe Volume Monitoring (Y/N) [N] : N
Store Version [2.0] :
KAVE05080-I The instance environment is now being created. (servicekey=RAID,
inst=53112)
KAVE05081-I The instance environment has been created. (servicekey=RAID,
inst=53112)
[root@namgdse04 tools]# ./jpcinssetup agtd -inst 210060
Storage Model (1:Thunder/AMS/HUS100 or 2:Lightning/USP/USP V/VSP/HUS VM) : 2
Command Device File Name [/dev/sda] : /dev/disk/by-id/scsi-
360060e8013274c0000000274c00000169
Unassigned Open Volume Monitoring (Y/N) [N] : Y
Mainframe Volume Monitoring (Y/N) [N] : N
Store Version [2.0] :
KAVE05080-I The instance environment is now being created. (servicekey=RAID,
inst=210060)
KAVE05081-I The instance environment has been created. (servicekey=RAID,
inst=210060)

```

To obtain a list of the configured agent instances, use the `jpcinlist` command. You must specify the agent type (RAID agent – `agtd`)

```

[/opt/jp1pc/tools]
sangs9# jpcinlist agtd
10102
53112
25519
99999

```

To delete an agent instance, use the `jpcinsunsetup` command. You will need to specify the agent type and instance name.

```

[/opt/jp1pc/tools]
sangs9# jpcinsunsetup agtd -inst 99999
KAVE05083-I The instance environment is being deleted. (servicekey=agtd,
inst=99999)
KAVE05084-I The instance environment was deleted. (servicekey=agtd,
inst=99999)

```

To modify an instance, use the `jpcinssetup` command. You will be prompted to update the existing instance and to terminate the services. You must specify the instance name and type. You will need to start or restart the services after the modification (see below).

```

[/opt/jp1pc/tools]
sangs10# jpcinssetup agtd -inst uspv10102
KAVE05912-Q The instance environment already exists. Do you want to update?
(Y/N)Y
KAVE05089-Q Running services will be forcedly terminated. Do you want to
continue? (Y/N)Y
KAVE05090-I Service is stopped.
Storage Model (1:Thunder/AMS or 2:Lightning/USP/USP V) [Lightning/USP/USP V]
:
Command Device File Name [/dev/rdisk/c3t50060E8005277600d5s2] :
Unassigned Open Volume Monitoring (Y/N) [N] : Y
Mainframe Volume Monitoring (Y/N) [N] : N
KAVE05913-I The instance environment is being updated. (servicekey=agtd,
inst=uspv10102)
KAVE05914-I The instance environment was updated. (servicekey=agtd,
inst=uspv10102)

```

To check the status of an agent instance, the `jpctdchkinst` command should be used. You must specify the agent instance number.

```

Working agent instances:
[/opt/jp1pc/tools]
sangs9# jpctdchkinst -inst 25519
KAVF18800-I The verification of the agent instance settings will now start.
(instance name=25519)
[Instance parameters]
Storage Model : Lightning/USP/USP V
Command Device File Name : /dev/rdisk/c1t50060E800563AF00d1s2
Unassigned Open Volume Monitoring : N
Mainframe Volume Monitoring : N
[Check result]
KAVF18850-I No error was found during verification of the agent instance
setting.
[Monitored storage subsystem information]
PRODUCT   : USP_VM
SERIAL    : 25519
FIRMWARE  : 60-03-06/00
SLPR      : 0
KAVF18801-I The verification of the agent instance settings will now end.

```

```

Non-working agent instance:
[/opt/jp1pc/tools]
sangs1# jpctdchkinst -inst uspv10102
KAVF18800-I The verification of the agent instance settings will now start.
(instance name=uspv10102)
[Instance parameters]
Storage Model : Lightning/USP/USP V/VSP
Command Device File Name : /dev/rdisk/c1t50060E8005
Unassigned Open Volume Monitoring : Y
Mainframe Volume Monitoring : N
[Check result]
KAVF18851-E An attempt to access the device set by the agent instance
parameter has failed. (parameter name=Command Device File Name, parameter
value=/dev/rdisk/c1t50060E8005)
KAVF18801-I The verification of the agent instance settings will now end.

```

Non-working instances are typically not correct due to the wrong or missing command device. If the RMLIB was not installed, a similar result will occur.

You must configure the location of the HTnM server (if not the same host).

```

Show current value
[/opt/jp1pc/tools]
sangs9# jpcnshostname
sangs9

Set new value
[/opt/jp1pc/tools]
sangs9# jpcnshostname -s ecctest01ms
KAVE05089-Q Running services will be forcibly terminated. Do you want to
continue? (Y/N)Y
KAVE05090-I Service is stopped.
KAVE05087-I The connecting host has changed. (sangs9 -> ecctest01ms)

```

To check the status of the communication processes, run the `jpcctrl` command.

```

The services will most likely not be started.

[/opt/jp1pc/tools]
sangs9# jpcctrl list "*"
KAVE05024-E The specified service is not registered.

```

On Solaris systems, you must start the services manually. We will configure them to start automatically in a later step.

```

[/opt/jp1pc/tools]
sangs9# jpcstart all

```

```
KAVE06007-I The service will now start. (service=Status Server)
KAVE06007-I The service will now start. (service=Action Handler)
KAVE06007-I The service will now start. (service=Agent Store for RAID,
inst=25519)
KAVE06007-I The service will now start. (service=Agent for RAID, inst=25519)
```

Check the services again. They should be “Active” or “S Active”.

```
[/opt/jp1pc/tools]
sangs9# jpcctrl list "*"

Host Name      ServiceID      Service Name    PID      Port      Status
-----
sangs9         PT1sangs9      Status Server   25410    22350     Active
sangs9         PH1sangs9      Action Handler  25411    61373     S Active
sangs9         DS125519[sangs9] Agent Store     25492    61441     S Active
sangs9         DA125519[sangs9] Agent Collector 25543    61445     S Active
KAVE06003-I List processing of the service information terminated normally.
```


To configure the services to run automatically at startup, do the following. Use the Service Management Facility built into Solaris 10:

- start the services with /opt/jp1pc/jpc_start.model (you may want to rename script)
- stop the services with /opt/jp1pc/jpc_stop.model (you may want to rename script)

Alternatively you could use these scripts renamed with S (for start) and K (for kill) prefixes in the /etc/rc3.d directory. However, Solaris will send a “start” (or “stop”) as a parameter, so the script will need to be modified. The existing scripts expect a hostname instead of “start” or “stop”.


27.2.2 Verify Agents

To see if the agent communication is working before a polling period has passed, **Go to Performance Reporter → Products → RAID**. You should see all active agents listed with the “DA” at the beginning of the instance name. The hostname of where the agent resides is in brackets.




Alternatively, to view the agents from Tuning Manager after one polling period has passed, go to **Administration → Data Polling**. You should see all agents listed with the last polling time after a successful polling period (default is 1x per 24 hours, recommended settings are 1x every 1-2 hours). The agents may show up under the

related agents column. Refer to the Tuning Manager Installation Guide for additional troubleshooting.



27.2.3 Polling Times

To alter the Polling Times, go to Administration under Explorer, then select Data Polling. There is a Default Polling Schedule as well as individual schedules. To change the Default Polling schedule, select Edit General Polling Settings.




Select all desired hours for Polling Times. Select at least every two hours (even hours) for Aggregation Times. Enable the retry settings (use approximately 5 minutes as the retry interval and 3 as the retry times).

Select OK when you are done with the changes.


You can also go to the previous screen and change the polling times of any group of agents from a given server. Select the agent(s) you wish to modify and then select the Edit Polling Schedule button on the bottom right of the screen. You cannot modify the retry settings or backdate offset here.

27.2.4 Manual Polling (Starting or Interrupting Polling)

To manually start an agent polling session, go to Administration under Explorer, then select Data Polling.



Select the agent you want to manually start and then select the Start Polling button on the bottom of the screen.




You can also interrupt polling that is taking too long or causing other issues. *It is advised to refresh the screen before deciding that it is taking too long as the Data Polling screen can quickly get out of date.* During an active polling, you will have the Interrupt Polling button enables and the Start Polling button grayed out.

27.2.5 Performance Reporter Configuration (60-day data storage)

The default settings for HTnM (Performance Reporter) provides only a small amount of historical data. These settings need to be changed in order to expand the reporting data storage to 60 days and to make sure all critical components are measured equally. The estimated amount of storage that the HTnM Agent for RAID will need is approximately 3GB per array.


27.2.5.1 Expansion of data set (data retention)

Go into HCS → Tools → Tuning Manager → Go → Performance Reporter → Services → Machines → Select the agent host. Then select the **DS1 instance name[agent host]** as shown below on left panel. Then select (on the right panel) **Properties** and then **RetentionEx**.




On each Product Interval Record (PI, LDS, LDS1, ...), change the Period - Minute Drawer (Day) from 1 to 60 and the Period Hour Drawer (Day) from 7 to 60. Press OK when you have changed the last record.


To make this change effective on all other instances (i.e. other arrays on same Tuning Manager), select the Distribute Property on the right panel after you have updated the periods above.



Select all of the other instances (USPV arrays, VSP arrays, etc.). Then Select Next.




Go into the RetentionEx on the right panel. For each of the interval records, hit the select all button. After going through all of the interval records, select Finish. The setting changes will be distributed for all of the other instances (arrays).



Note: This will increase the amount of storage that the agent requires on the agent host, but will not affect the HCS/Tuning Manager server. The estimated amount of storage required is 3GB per array.

27.2.5.2 Expansion of Collection Interval

The collection interval needs to be modified to make all of the components consistent. The default values do not have this consistency. This will further improve reporting accuracy. Go into Performance Reporter → Services → Machines → Select the agent host. Then select the DA1 instance name [agent host] as shown below. Then select Properties and then Interval Records.



The following Collection Interval records need to be changed:

- LDS Logical Device Summary 300 to 60
- LDA Logical Device Aggregation 300 to 60
- PDS Physical Device Summary 300 to 60
- RGS RAID Group Summary 300 to 60

You will want to distribute these settings to the other agent instances (arrays). Select the Distribute Properties from the right panel and distribute similar to previous step.

Note: This will also affect the agent disk usage as you will be collecting records more frequently. The total amount of storage needed is 3GB per array.

27.2.5.3 Updating max records allowed in a report (config.xml parameter change)

This change (made on the HCS server) allows the reports to include more data in the report table and graph.

Update the config.xml under

\Program Files (x86)\HiCommand\TuningManager\PerformanceReporter\conf

There are two locations to change. You can search for "1440" and replace the value with "2147483647".

Location #1

```
</vserver-connection>
```

```
<!-- The maximum limits of the records.
      Specifiable values: 1 to 2,147,483,647
      Default           : 1,440
```

```
-->
```

```
<param name="maxFetchCount" value="2147483647"/>
```

Location #2

```
<export>
```


```
<!-- The maximum limits of the records.
      Specifiable values: 1 to 2,147,483,647
      Default           : The value of maxFetchCount of the <vsa> tag is applied.
```

```
-->
```

```
<param name="maxFetchCount" value="2147483647"/>
```

```
</export>
```

From the HCS server, restart the HiCommand Performance Reporter services (from services.msc) to make this change effective.



27.3 Performance Reporting Methodology

HDS uses the term Subsystem Flow Analysis to describe the process of tracing the performance issue from the host to the back-end disks. At Citi, we do not use host agents for servers as there is resistance to placing additional agents on our systems. Therefore the Workflow Analysis starts at the array.

27.3.1 Basic Terminology

- Read hit: data found in cache
- Read miss: data must be read from disk
- Write: always write/update to cache. Acknowledge (ack) occurs after written to cache.
- Write through mode: when physical maintenance is being performed on cache boards, the ack occurs only after destaged to disk.

27.3.2 Review of a Typical Application Transaction

The following is a summary of the typical application I/O transaction (OLTP example):

```
|----CPU Time (25%)-----||--Queue Time (15%)--||-----I/O Time (60%)-----|
|-----Total Time (Excluding LAN or WAN) = Aprox. 0.03 sec-----|
```

Email, Decision Support System (Warehouse), and Rich Media (video, audio) applications have a higher percentage of I/O throughput.

Online Analytical Processing, Customer Relationship Management, and Enterprise Resource Management applications have a higher percentage of processor utilization.

27.3.3 Response Time Factors

- Port load balancing is important. Performance problems could occur because of:
 - o Too many HSDs on the same physical port.
 - o Too many active LUNs on the same physical port.
- Protocol time on physical port is very low (microseconds)

- A physical port can process thousands of I/O requests before reaching saturation point. This saturation point will be reached faster in cases where the average block size is large (64K and higher)
- Applications can be throttled using Server Priority Manager (see [Appendix P: Server Priority Manager](#)).

27.3.4 General Performance Monitoring

- Look at historical trend first. Identify trouble spots by looking at longer period of trend.
- Focus on trouble area with finer granularity. After identifying trouble areas, zoom into these time periods with finer granularity (minute-level).
- For response-time centric applications such as On-Line Transaction Processing (OLTP) applications:
 - o You want low utilization to ensure CPU cycle availability
 - o You want low Q-Depth to ensure no wait time
- For throughput centric applications such as backups:
 - o You want high utilization to ensure maximum throughput
 - o You want high Q-Depth to let CPU manage queues smarter

27.3.5 Performance Analysis Methodology

Expected Peak I/O Rates – A general rule of thumb is that peak I/O can be estimated at 60% over the average I/O rate.

27.3.5.1 Workflow is from the Outside to the Inside

Start at the server that is reporting the problem

- This will have to come from the SA or application administrator as we do not have host agents for HTnM
- Analyze its I/O characteristics of the problem devices
- Response Time
- IO Rate
- Read/Write Ratio
- Connected Storage

Look at the ports for problems:

- Port and Port CPU Utilization (Data Transfer Rate)
- Other connected servers

Look deeper in the Array for problems:

- Utilization
- Other connected servers
- Cache Read Hit percentage
- Read/Write Ratio
- Random vs. Sequential

Look at the Logical Device

- Response Time
- Cache Read Hit%
- Read/Write Ratio
- Random vs. Sequential

Look at Overall Storage System

- Ports
- Array Groups
- Logical Devices
- Cache

Look at related or contending server workloads

- This will have to come from the SA or application administrator as we do not have host agents for HTnM
- Analyze its I/O characteristics of the problem devices
- Response Time
- IO Rate
- Read/Write Ratio
- Connected Storage

27.4 Citi Performance Criteria for HDS Arrays

The following is a summary of expected performance limits for normal array behavior. It is possible to beyond these limits without experiencing user impact, but that is not likely. The two thresholds listed are for warning and a critical level. These values should be the basis for alarms and determining performance problems. These criteria apply to all Citi VSP and HUS-VM arrays.

Array Component	Metric	Warning Level	Critical Level
Cache	Write Pending %	30%	75%
FED	Transfer Rate	600MB/sec	800MB/sec
VSD	MBP Busy %	40%	70%
	MP Busy %	40%	70%
Pool	Average Read Response Time	8ms	>10ms
Parity Group	PG Max Busy %	50%	70%
LDEV	Average Response Time	20ms	>30ms

Cache Write Pending

A high write pending percentage (>30%) indicates that cache is not being destaged to disk fast enough. Sustained high levels indicate more of an issue than occasional spikes.

Read hit percentage indicates the randomness of data (low percentage).

FED Transfer Rate

A high sustained transfer rate (>800MB/sec) indicates that the host port is being overloaded. Migration to a different port pair or additional port spreading is likely needed.

VSD Busy Percentage

Less than or equal to 40% should be the busy percentage objective to allow for failover to the other processors on the same board. LUN ownership by any MP can be dynamically reassigned.

Pool Response Time

Pool read response time (average) indicates the overall health of the pool. If the average is higher than 8-10ms, then the contributing parity groups should be examined.

Parity Group (RAID Group) Busy %

A high (sustained) parity group busy percentage (>50%) indicates that there may not be enough drives to support the workload.

LDEV


Read response time indicates overall health of the LDEV. A historical perspective is important.

27.5 Citi Standard Reports and Alarms

There are several standard reports and alarms that are part of the Citi HTnM Reporting standard. These are obtained by installing the latest Hitachi Predfined Reports export.

27.5.1 Pre-Defined Reporting Setup

- Import the reports from the latest version of the Citi HTnM reports. There will be a separate file for Reports and Alarms. Import the Reports using Performance Reporter and select the import button after selecting reports. Select the filename of the reports file from the package.



27.5.2 Pre-Defined Reporting

There are several canned reports that have been developed in coordination with HDS in order to pull the most common performance data out of HTnM. The Hitachi Predfined Reports (V7.6.0) are provided in each HCS instance at Citi. If these user reports are not available, please contact CATE Storage Engineering. The following categories each contain relevant reports.

If in doubt on where to begin: start with the Array Health Check reports below as they provide likely bottleneck reports that are part of the individual categories.

- ▼ Hitachi Predfined Reports (V7.6.0)
 - ▼ VSP/HUSVM/USPV Reports
 - ▶ 1. Cache
 - ▶ 2. Front End Port
 - ▶ 3. Processor
 - ▶ 4. Parity Group
 - ▶ 5. LDEV
 - ▶ 6. Subsystem
 - ▶ 7. HDP/HDT
 - ▶ 8. TC/HUR
 - ▶ 9. Array Health Check

27.5.2.1 Cache Reporting

- ▼  **1. Cache**
 -  [Cache Read Hit %](#)
 -  [Cache Usage %](#)
 -  [Max Write Pending Usage %](#)
 -  [Multi Array Cache Write Pending Usage %](#)
 -  [Write Pending Usage %](#)

Available reports in this category include Cache Read Hit %, Cache Write Pending Utilization, and Overall Cache Utilization. There is a Multi Array Cache Write Pending Utilization report that is useful for combining data from multiple arrays in the same report (for array consolidation).

27.5.2.2 Front End Port Reporting

- ▼  **2. Front End Port**
 -  [Port Avg IOPS](#)
 -  [Port Avg Transfer](#)
 -  [Port Configuration \(Tabular\)](#)
 -  [Port Max IOPS](#)
 -  [Port Max Transfer](#)

Available reports in this category include average and maximum IOPS and Throughput (Transfer). There is also a configuration report.

27.5.2.3 Processor Reporting

- ▼  **3. Processor**
 -  [USP\(V\) BE Processor Busy %](#)
 -  [USP\(V\) FE Processor Busy %](#)
 -  [VSD Busy % by MPB](#)
 -  [VSD Busy % by Processor ID](#)

For USP(V) arrays, the Back End (BE) and Front End (FE) Processor Utilization reports depict potential problems with over taxed processors that handle host requests and disk requests. For VSD systems, the VSD Utilization reports (sorted two ways) show potential problems with over taxed processors that handle data flow (host and disk combined).

27.5.2.4 Parity Group Reporting

▼ 4. Parity Group

- PG Busy %
- PG Random Read IOPS
- PG Random Read Transfer
- PG Random Write IOPS
- PG Random Write Transfer
- PG Read Hit %
- PG Sequential Read IOPS
- PG Sequential Read Transfer
- PG Sequential Write IOPS
- PG Sequential Write Transfer
- Parity Group Configuration (Tabular)

The Parity group reporting includes IOPS and throughput reports for random read, random write, sequential read, sequential write as well as overall parity group busy and read hit percentage. There is also a parity group configuration report available.

27.5.2.5 LDEV Reporting

▼ 5. LDEV

▼ LDEVs with CU 00-3F

- CU 00-3F Daily IOPS for LDEV
- CU 00-3F Daily Transfer for LDEV
- CU 00-3F LDEV Configuration (Tabular)
- CU 00-3F LDEV Performance Summary Report
- CU 00-3F LDEV Random IOPS
- CU 00-3F LDEV Random Transfer
- CU 00-3F LDEV Read Hit %
- CU 00-3F LDEV Read IOPS
- CU 00-3F LDEV Read Response Time
- CU 00-3F LDEV Read Transfer
- CU 00-3F LDEV Sequential IOPS
- CU 00-3F LDEV Sequential Transfer
- CU 00-3F LDEV Total Response Time
- CU 00-3F LDEV Write IOPS
- CU 00-3F LDEV Write Response Time
- CU 00-3F LDEV Write Transfer
- CU 00-3F Up to 20 LDEV Bars for Yesterday's IOPS
- CU 00-3F Up to 20 LDEV Bars for Yesterday's Transfer

- ▶ LDEVs with CU 40-7F
- ▶ LDEVs with CU 80-BF
- ▶ LDEVs with CU C0-FF

There are several reports demonstrating combinations of IOPS and Throughput (Transfer) for Read, Write, Daily (W/R), Sequential (W/R), Random (R/W), Bar Charts for the top 20 LDEVs

per MPB, This category is organized by the logical CU range and is repeated for each range beyond 00-3F (40-7F, 80-BF, CO-FF).

27.5.2.6 Subsystem Reporting

▼ 6. Subsystem

-  Multi Array Total Read IOPS (Stacked)
-  Multi Array Total Write IOPS (Stacked)
-  Total Read/Write IOPS (Line)
-  Total Read/Write IOPS (Stacked)
-  Total Read/Write Transfer (Line)
-  Total Read/Write Transfer (Stacked)

The subsystem category includes reports for Read/Write IOPS and Throughput (Transfer) with Line and stacked area variations. The Muti Array reports for Read IOPS and Write IOPS can be used for combining data from multiple arrays (used for array consolidation).

Thin Provisioning Pools (HDP and HDT) Reporting

▼ 7. HDP/HDT

-  Pool Read IOPS
-  Pool Read Response Times
-  Pool Space Usage %
-  Pool Write IOPS
-  Pool Write Response Times
-  Tier IOPS
-  Tier IOPS Usage % from Monitor
-  Tier IOPS by DP-VOL
-  Tier Pages Demoted
-  Tier Pages Promoted
-  Tier Pages Relocation Status
-  Tier Space Usage %
-  Tier Space Usage % by DP-VOL

The HDP (Thin Pool) reporting includes information on a pool level for IOPS, Space Utilization, Read Response, and Write Response. The tier reports are used for dynamic tiering (HDT) reporting and include page promotion/demotion, and tier space utilization.

27.5.2.7 Replication (TC/HUR) Reporting

▼ 8. TC/HUR

-  TC/HUR Initiator Port List
-  TC/HUR RCU Target Port List

The replication reporting category includes reports about the replication port configuration.


27.5.2.8 Array Health Check Reporting

- ▼ 9. Array Health Check
 - Cache Read Hit %
 - Max Write Pending Usage %
 - PG Busy %
 - Pool Read Response Times
 - Port Max Transfer
 - Total Read/Write IOPS (Line)
 - VSD Busy % by MPB

This category is likely the best place to start to get a good overview and likely indicator of performance bottlenecks. All of these reports are preconfigured with a selection of minute data for 7 days.

27.5.3 Alarms Setup

1. Import the Alarms from the latest copy of the Citi HTnM alarms. There will be a separate file for Reports and Alarms. Import the Alarms using Performance Reporter and select the import button after selecting alarms. Select the filename of the alarms file from the package.



2. Once the alarms are present in Performance Reporter, you will need to activate the alarms. Once activated, the icon on the alarms will change to .
3. After the alarms have been activated, you will need to bind them to the agents for each of the arrays. Select an array-agent instance under agents (Products → RAID → DA1xxxxx) and then select Bind Alarm Table. If the Alarms are active, then it should allow you to select the appropriate Alarm Table (RAID → Citi Alarms).

27.5.4 Email Alerting Setup

27.5.4.1 User.properties

There is a properties file that needs to be updated. This file is located in C:\Program Files\HiCommand\TuningManager\conf.

```
#pr.incontextlaunch.mode=
alert.email.smtpport=25
alert.email.from=GDSE@CITIGROUP.COM
```

```

alert.email.subject=Tuning Manager (HTnM) Alert
#table.filter.default.unit.datasize=
#table.filter.default.unit.transferRate=
#portNumberCLI=
#poller.offset=
csv.encode.type=UTF-8

```

Uncomment the alert.email lines and place values at the end of each line. The smtpport should be 25, the email.from should be a valid address, and the email.subject should be meaningful information for the alert email.

27.5.4.2 System Alerts Page

The System Alerts page should also be updated (Explorer → Administration → System Reports). Select the Edit System Alerts button.

Edit System Alerts

Alert Definitions

Select and specify alert definitions.

- ☒ Alerts for agent status
This function sends alerts to an administrator if agents could not collect data.
Threshold: 24 Hour
- ☐ Alerts for polling skip
This function sends alerts to an administrator if polling was skipped.
- ☐ Alerts for time consumed in polling
This function sends alerts to an administrator if polling takes long time.
Threshold: 6 Hour
- ☐ Alerts for error in polling
This function sends alerts to an administrator if polling task in polling was failed.
- ☒ Alerts for Db space
This function sends alerts to an administrator if polling was failed.

Alert Notification

Specify notification settings.

*E-mail To: GARY.G.JOHNSON@CITIGROUP.COM

*Mail Server: mailhub-ss01.nam.nsroot.net

☐ SMTP

UserName: _____

Password: _____ Edit

*Required field

Send Test E-mail

OK Cancel


1. (Optional) Uncheck the alerts for polling skip, alerts for time consumed in polling, and alerts for error in polling.
2. Fill out the E-mail To and Mail server fields. The E-mail to field should be a distribution list for the HTnM administrators (must be an external address distribution list like "gdse@citi.com"). The mail server should be a hostname or ip address.
3. Hit the Send Test Email button to test the ability to send test emails.

27.5.4.3 SNMP Alerting Setup


This functionality has not been tested with the ESS team. Implementing SNMP traps for alarms is not recommended at this time.

27.5.4.4 Agent Email Enable

In order to enable email alerting, start Performance Reporter (HTnM → Go → Performance Reporter). Select Services, Machines. Select the server that is running HTnM. Under that machine will be a list of processes, select the PH1_____ process. Select Properties. Under Properties select Capabilities. Change the script and email values to be YES.




Under Mail, you will need to update several entries.




Fill out the SMTP host (Mail server on previous screens), SMTP Sender (valid email address coming from), Charset to utf-8, and Protocol to IPv4. Press OK when done.

27.5.4.5 Modifying Alarms (Enabling Email Notification)

After email capability has been allowed, each of the alarms will need to be modified to use email for notification and to define the destination email address. Go to each Alarm and select Edit. Press Next until you get to the Actions page. Select the checkboxes for Email under Abnormal and Warning. Select Next.




Fill out the email address that the alarm will be sent to. Select Finish when done. Repeat this for all of the alarms.



27.5.5 RAID Alarms

To find these alarms start Performance Reporter (HTnM → Go → Performance Reporter). Select Alarms, RAID, and then USPv.




Note: The alarms listed above show that they have been activated as the icon has changed to .

27.5.5.1 Pool Utilization Alarms

There are two Pool Utilization Alarms. The primary alarm is Pool Utilization %. It has a warning level of 70% and an abnormal level of 80%. Additionally there is a Pool Utilization Early % alarm. This alarm is triggered with a warning of 50% and an abnormal value of 60%.

27.5.5.2 Triggering Alarms


Once the criterion for an alarm has been met, the icon for that alarm will change. To observe which alarm has been triggered for a given agent instance, click on the agent instance and then select Display Alarm Status. The following is an example of a triggered alarm (Abnormal level):




If configured as above, there should have been an entry in the log file as well as an email alert sent.

27.6 Thin Pool (HDP) Trend Reporting


To access the monthly trend report on pool storage, go to the Performance Reporter area (Go → Performance Reporter). Products → RAID → Select appropriate Service ID of the agent for the array. The service ID should be named according to the serial number of the array and the [location of the agent]. For example, the agent running against array 10102 is named “DA110102[ecctest01ms]”.



Once you select the service ID of the RAID agent, you will enable the links on the right side of the screen. On the right side of the screen, select Display Reports, then System Reports → RAID → Monthly Trend → Pool Usage Trend (7.1).



Adjust the time period and then select OK:





Each Pool ID should be selectable so that further drill down reports can be generated. See the [Virtual Volume Trend Reporting](#) section that follows this section.

This report displays a graph of the data points over the time period and a table that lists all of the values. The data fields that comprise this report are:


Data Field	Description
Agent Host	This is the name of the host running PFM - Agent.
Agent Instance	This is the PFM - Agent instance name.
Agent Type	This is the PFM - Agent product ID (1-byte identifier).
Attribute	Type of the pool
Collection Time	Time when the data was collected from the storage subsystem (GMT)
Date	This is the record creation date (GMT).
Date and Time	This is the combination of the Date (DATE) and Time (TIME) fields.
Emulation Type	Emulation type of the pool
Free Capacity	Free space in the pool (GB)
GMT Offset	This is the difference between Greenwich Mean Time (GMT) and the local time in seconds.
Interval	Interval in seconds during which the information is collected
Pool ID	Pool ID of the pool
Pool Volume Count	Number of pool volumes in the pool
Record Time	Time when the record was stored (GMT)

Record Type	Record type identifier, which is always set to PLC
Status	Status of the pool - NORMAL (the status is normal) - WARNING (the usage rate of the pool capacity has exceeded the value of the Threshold or Warning Threshold field) - BLOCKED (the usage rate of the pool capacity becomes 100%, and the pool is blocked)
Threshold	Pool threshold value 1(%) set for the usage rate of the pool capacity (the value of the Usage % field).
Time	This is the record creation time (GMT).
Total Actual Capacity	Actual capacity of the pool (GB)
Total Managed Capacity	Virtual capacity of all V-VOLs mapped to the pool (GB)
Usage %	Of the total actual capacity of the pool, the percentage of the pool used. Usage rate of the pool capacity.
Used Capacity	Size of the pool used (GB)
Virtual Volume Count	Number for the V-VOLs mapped to the pool
Warning Threshold	Pool threshold value 2(%) set for the usage rate of the pool capacity (the value of the Usage % field).

You can export this table (in CSV format) by selecting the Export link in the upper right hand corner of the report.




If you selected a RAID agent for an array that is not HDP capable (i.e. a USP) or that does not have any HDP pools, then you will get an error message when you try to select the Pool Trend Report:




27.7 Virtual Volume Trend Reporting (Detail of Pool Trend Report)

To obtain a report displaying the virtual volumes that come from a given pool, you can click on the Pool ID of the desired pool in the table below the graph. This will bring up another options screen to allow you to select the time range.





For this report, it is suggested that you change the time range from monthly to a week or less. This is due to maximum number of records (1440) in the report. Otherwise your report will reflect only the first 1440 records, which will not cover the whole month.

Note: the name of the report has changed from "Pool Usage Trend" to "Virtual Volume Usage Trend".




You can export this table (in CSV format) by selecting the Export link in the upper right hand corner of the report. You will be given the option to save the file locally.






27.8 General Array Reporting


Tuning Manager also provides non-HDP reporting. These reports include Write I/O, Read I/O, Write IOPS, Read IOPS, etc. To get to these reports, follow the same method as the HDP reports. To access the monthly trend reports, go to the Performance Reporter area (Go → Performance Reporter). Products → RAID → Select appropriate Service ID of the agent for the array. The service ID should be named according to the serial number of the array and the [location of the agent]. For example, the agent running against array 10102 is named "DA110102[ecctest01ms]".




Once you select the service ID of the RAID agent, you will enable the links on the right side of the screen. On the right side of the screen, select Display Reports, then System Reports → RAID → Monthly Trend.




The following is a sample of the Subsystem Read I/O Rate Trend:




The following is a sample of the Subsystem Write I/O Rate Trend:



The following is a sample of the Subsystem Read I/O Transfer Rate Trend (Read IOPS):



The following is a sample of the Subsystem Read I/O Transfer Rate Trend (Write IOPS):



28 Appendix N: TMEA Reporting

TMEA is a data collection and reporting service that Citi has purchased from HDS to supplement the array reporting. The nature of the data being transmitted to HDS is non-confidential as it is only performance and configuration information about the array.

TMEA accesses the array data similar to how HTnM works. It makes inquiries to the array frequently via a command device. Instead of sending the data back to the HCS server, the data is sent to HDS and added to an extensive reporting database. Reports are then generated automatically and sent back to Citi staff.

28.1 Setup and Configuration

The following steps are needed to configure the TMEA Collector that is installed on a RHEL 5.9 server with command device access to arrays at a given facility.

TMEA uses a command device to access the array (similar to HTnM and Raidcom). Command device access is currently limited to non-vmware based hosts so the TMEA Collector will need to reside on a stand-alone host with fiber access to the arrays at a given site. This host is typically a standalone RHEL server. The RHEL server also runs HTnM and Raidcom as they also require command device access. Each application should have its own dedicated command device for a given array.

28.1.1 Use of Proxy

TMEA also sends data via sftp (TCP port 22) to external HDS servers (datamart.hds.com) for data processing and report generation. This requires that the servers running TMEA Collector have access to the HTTP (or SOCKS) Proxy. This may require an ACL update to allow this server to access the proxy.

To obtain the best response time and to avoid unnecessary WAN traffic, each site may have a preferred HTTP proxy server to use. For APAC / EMEA, the Network Proxy team should provide ACL access to a local HTTP proxy server if available. For NAM / LATAM, there may need to be multiple ping queries to determine the fastest response time for the DNS entries under "webproxy.wlb.nsroot.net".

28.1.2 Root Access Setup (Powerbroker)

The setup of the software package requires the use of root. In addition, extraction of the RMLIB software and symbolic link setup also requires the use of root. Powerbroker should be used to allow for root access from other accounts. The following commands should be added to the powerbroker configuration:

- /opt/TMEA/bin/manage
- /opt/TMEA/bin/volinfo

Note: In the examples below, sudo was used instead of pbrun as they have similar behavior. The pbrun command was inserted to avoid confusion.

28.1.3 Storage for TMEA

TMEA requires 10GB per array as buffer space in case there is a disruption of data flow between Citi and HDS. This storage should be presented to the application as a single file system (for example \HDS_data\TMEA). If this storage is hosted by an HDS array, then

statistics can be reported on the growth and performance on its use. It is important to configure the clean-up script to ensure the number of archive files is kept in check.

28.1.4 Installation of RMLIB

TMEA Collector depends on RMLIB to be able to communicate with the HDS arrays. RMLIB is included with the package software, but is not a part of the .rpm package. RMLIB may have been previously setup with the installation of HTnM or Raidcom. This will require root access to install (CPIO to /opt directory and creating symbolic links). Here are some important information about RMLIB:

- RMLIB requires a separate installation (hopefully before TMEA, HTnM, or Raidcom is configured)
- RMLIB is provided as a CPIO file (similar to zip). It is not a package. It has an installation script once the software has been extracted.
- RMLIB should be extracted to the directory /opt/RMLIB (or similar). The CPIO file expects to put it into a directory named RMLIB, so specify /opt (or the appropriate parent directory) as the destination.
- RMLIB is the library that is called to communicate with the arrays and is critical for TMEA, HTnM Agent, and Raidcom.

Note: If in doubt, reinstall the new version of RMLIB on top any existing version.

Run the following CPIO command to extract the files to the desired location. It will force a destination directory name of RMLIB, so initiate the CPIO extract from the /opt directory.

```
[root@namgdse17 opt]# cpio -idmu <
/net/stealth/export1/home1.localhost/sw/PKG/HDS_TMEA_RHEL_9-5-
3b12/RMLIB_RHEL_01-17-03

1999 blocks

[root@namgdse17 opt]# cd RMLIB

[root@namgdse17 RMLIB]# ls -al
total 32
drwxr-xr-x  6 root sys  4096 Nov 18 11:52 .
drwxr-xr-x 41 root root 4096 Nov  6 17:51 ..
drwxr-xr-x  2 root sys  4096 Nov 18 11:52 bin
drwxr-xr-x  2 root sys  4096 Nov 18 11:52 include
drwxr-xr-x  2 root sys  4096 Nov 18 11:52 lib
-rw-r--r--  1 root root    0 Oct  2 14:24 RLSYSLK
-r-xr-xr-x  1 root sys  4006 Apr 10 2012 rmlibinstall.sh
-r-xr-xr-x  1 root sys  3284 Apr 10 2012 rmlibuninstall.sh
drwxrwxrwx  2 root sys  4096 Nov 18 11:52 sample
```

Run /opt/ /RMLIB/rmlibinstall.sh after the setup to setup the symbolic link.

```
[/opt/ RMLIB]
sangs17# rmlibinstall.sh
The symbolic link of '/usr/lib/libsvrrm.so' was successful.

If previously run:
[root@namgdse17 RMLIB]# ./rmlibinstall.sh
The symbolic link of '/usr/lib/libsvrrm.so' was successful.
The symbolic link of '/lib64/libsvrrm.so' was successful.
ln: creating symbolic link '/usr/lib/RMLIB/bin/whatrmver' to
'/usr/lib/RMLIB/bin/whatrmverx64': File exists
```

To verify the symbolic links are functional and that RMLIB is working, run the version command (/opt/RMLIB/bin/whatrmver):

```
[root@namgdse17 bin]# cd /opt/RMLIB/bin
[root@namgdse17 bin]# ./whatrmver
Model :RAID-Manager/LIB/Linux
```

```
Ver&Rev:01-17-03/01
```

You can uninstall the software by running `rmlibuninstall.sh` from the `/opt/RMLIB` directory.

28.1.5 Installing TMEA Collector Package

Installing the TMEA package will require root access. This activity requires root access.

Run the `rpm -ivh` command to install the package from the source location.

```
[root@namgdse04 TMEA_Linux]# rpm -ivh --nodeps --nomd5 TMEA-9_5_0b12.rpm
Preparing... ##### [100%]
   1:TMEA          ##### [100%]
```

Verify the package is installed.

```
[root@namgdse04 TMEA_package]# rpm -qa | grep TMEA
TMEA-9_5_0b12-0.el5.i386

[root@namgdse04 TMEA_package]# rpm -qi TMEA-9_5_0b12-0.el5.i386
Name           : TMEA                      Relocations: (not relocatable)
Version        : 9_5_0b12                 Vendor: (none)
Release        : 0.el5                     Build Date: Thu 14 Nov 2013
02:47:56 PM EST
Install Date: Thu 14 Nov 2013 04:03:55 PM EST   Build Host:
localhost.localdomain
Group          : HDS                      Source RPM: TMEA-9_5_0b12-
0.el5.src.rpm
Size           : 22633445                  License: TmeaCollector
Signature      : (none)
Summary        : Meta package for TMEA
Description    :
Tmea Collector
```

28.1.6 License Files

Each array will have its own license file that should last for 24-36 months. This license file will need to be obtained from HDS (Datamart / TMEA support). The license files for TMEA are simple text files. However, the TMEA application is sensitive to the Windows to Linux translations. You can either convert the file using a dos to UNIX converter or you can copy the file contents and paste them into a new file via a SSH session running a UNIX editor. The license keys should be placed in a common directory (`/HDS_data/TMEA/license_keys`) for consistency. The following is an example license file. It was named with a simple name for ease of reference. Do not use this example file.

```
[root@namgdse04 license_keys]# pwd
/HDS_data/TMEA/license_keys
[root@namgdse04 license_keys]# ls -al
total 20
drwxr-xr-x 2 root root 4096 Oct 31 12:11 .
drwxr-xr-x 6 root root 4096 Nov 14 17:44 ..
-rw-r--r-- 1 root root 247 Oct 31 12:10 TMEA_10102.lic
-rw-r--r-- 1 root root 248 Oct 31 12:11 TMEA_210060.lic
-rw-r--r-- 1 root root 247 Oct 31 12:09 TMEA_53112.lic
[root@namgdse04 license_keys]# cat TMEA_53112.lic
## license
devices = 1
expires = 12/31/2014
good_for = 429
hostid = 53112
issued = 10/28/2013
key[0] =
V4rJc02aV0KMqC1jIjFbgAoQesTSnLL53YYw4nPL8ACI6gnBZ5sisSi2Xi3sj/KkAb0/ff2bYZ0=
key_version = 2.0
nproducts = 1
```

```
product[0] = HiPERFC
version = 2.x
```

28.1.7 Command Devices

Command devices should be allocated from each array to the RHEL host. These command devices should be unique and not shared by any other application (i.e. HTnM or Raidcom). These are simple command devices and do not have the security options turned on.

Linux native devices (/dev/sd[a-z] format) are not persistent and therefore should not be used. You should search for the scsi- (or wwn-) based name to allow for persistence between reboots of the server.

For example, look for the command device 00:01:68

```
[root@namgdse04 usr]# ls -al /dev/disk/by-id/ | grep 0168
lrwxrwxrwx 1 root root 10 Nov 19 17:32 scsi-
360060e8013274c000000274c00000168 -> ../../sdae
lrwxrwxrwx 1 root root 10 Nov 19 17:32 wwn-
0x60060e8013274c0000000274c00000168 -> ../../sdae
```

This is the format you will refer to the command device in later steps:
/dev/disk/by-id/scsi-360060e8013274c000000274c00000168

Note: you can also use the wwn name, but there is not always a wwn name in addition to the scsi name.

Verify the command device is valid and is for the correct array.

```
[testuser1@namgdse04 bin]# pbrun /opt/TMEA/bin/volinfo /dev/disk/by-id/wwn-
0x60060e8013274c0000000274c00000168
Command Device:/dev/disk/by-id/scsi-360060e8013274c000000274c00000168
vendor id: HITACHI
product id:OPEN-V-CM
storage type:HUSVM
serial number:210060
(/dev/sdbc)
```

28.1.8 Configuring TMEA Collector

Run the “manage” (pbrun /opt/TMEA/bin/manage) to initiate the configuration manager for TMEA. Add an instance with the “cv” command when prompted. This will create a new instance. An instance is needed for every array at the datacenter to be included in the TMEA reporting.

Note: It is important that you reference the command device in the format described previously to maintain consistency between reboots of the operating system.

Array Specific Configuration Parameters

- | | |
|---------------------------|---|
| • Instance Name | <serial number of array> |
| • Storage Subsystem Alias | use global name (i.e. XXXXX) |
| • Array Type | VSP or HUSVM or USPV |
| • Command Device | use /dev/disk/by-id/wwn... format |
| • License File | use /HDS_data/TMEA/licence_keys/ OR similar |
| • Output directory | use /HDS_data/TMEA/<serial num> OR similar |
| • Log directory | use same as output directory |

Common Configuration parameters:

- Collection Interval (in Hours) 24
- Number of collections 0
- Perf: Collection Interval (in Minutes) 1
- Perf: Number of collections 0
- Perf: Roll file after x collections 60
- Start Collection Time 0 (1st array, +5 mins for next array...)
- Sftp Proxy Machine (use an IP of local HTTP Proxy)
- Protocol for sftp proxy xfer HTTP
- Sftp Proxy Port 8080
- Machine ip Sftp transfer 207.126.248.84
- Username for sftp xfer citi
- Password for sftp xfer citi123
- Directory for sftp xfer /<serial number of array>

```
[testuser1@namgdse04 TMEA]$ pbrun /opt/TMEA/bin/manage

Enter 'cv' to create VSP instance and 'ca' to create AMS instance.
? cv
Adding new instance ...

Instance Name (one word): 53112
Storage Subsystem Alias: HDS53112-VSP
Storage Subsystem Type(USP/USPV/VSP/HUSVM): VSP
Specify a command device in /dev/sd[a-z] format: /dev/disk/by-id/wwn-
0x60060e800
6cf7800000cf7800003e16
License File: /HDS_data/TMEA/license_keys/TMEA_53112.lic
Output Directory: /HDS_data/TMEA/53112
/HDS_data/TMEA/53112 does not exist; would you like to create it [y] y
making /HDS_data/TMEA/53112 directory...
Log Directory: [/HDS_data/TMEA/53112]

***** Configuration Setup
Collection Interval (in HOURS): 24
Number of collections (0 for continue until stopped.): 0
No limit being set on number of collections.

***** Performance Setup
Collection Interval (in MINUTES): 1
Number of performance collections (0 for continue until stopped.): 0
No limit being set on number of performance collections.
Roll file after x collections: 60
Enter Start Collection Time for performance (enter between 0-55,in the order
of
5): 5
Enter sftp proxy machine address for sftp xfer or 'none' to skip sftp proxy
? 199.67.138.49

Enter protocol(SOCKS4/SOCKS5/HTTP) for sftp proxy xfer
? HTTP
Enter sftp proxy port for HTTP
? 8080

Enter machine address for sftp xfer or 'none' to skip sftp xfer
? 207.126.248.84
Enter user name for sftp xfer
? citi
Enter password for sftp xfer
?
Transfers will be made to 207.126.248.84 using sftp account for citi
```

```

Enter path to directory relative to the citi sftp account
where the data directory will be created for the transferred files.
Use forward slashes (/) instead of back slashes.
? /53112
Creating /53112 directory for citi on 207.126.248.84...
Added row to col2anal.config
TMEA config for 53112 started
TMEA dkcbusy for 53112 started
TMEA ldevio for 53112 started
TMEA portio for 53112 started
TMEA pairstat for 53112 started
TMEA clprbusy for 53112 started
TMEA poolstat for 53112 started
TMEA for 53112 started
/opt/TMEA/bin/autoxfer started at 11/14/13 16:34:09
Auto Transfer for 53112 started
53112 set up completed.

Enter number to select instance or 'cv' to create VSP instance and 'ca' to
crea
te AMS instance.
1 53112                      running

```

Repeat this configuration for any remaining arrays at the same site. You will be able to hit enter for the sftp and proxy values instead of having to enter these each time. All array instances should be in a running status after they have been configured.

28.1.9 Starting and Stopping Instances

It may be necessary to start or stop the processes for a specific array. Run the manage command (pbrun /opt/TMEA/bin/manage). Select the id of the instance you want to start or stop.

stop = stop the processes for the array

```

[testuser1@namgdse04 opt]$ pbrun /opt/TMEA/bin/manage

Enter number to select instance or 'cv' to create VSP instance and 'ca' to
create AMS instance.
1 53112                      running
2 10102                      running
3 210060                    running
? 1

Instance 53112 -- Collector: running, Autoxfer: enabled
Last ftp transfer: succeeded at Thu Nov 14 18:07:51 EST 2013

Enter operation code:
stat -- Display detailed status of Collector instance
stop -- Stop the Collector
enddebug -- Enable the debug logs
disdebug -- Disable the debug logs
dis -- Disable auto transfer of data
fc -- Force collection of configuration data
n -- Select new instance
!CMD -- Execute CMD in sh
q -- Quit
? stop
TMEA config for 53112 stopped
TMEA dkcbusy for 53112 stopped
TMEA ldevio for 53112 stopped
TMEA portio for 53112 stopped
TMEA pairstat for 53112 stopped
TMEA clprbusy for 53112 stopped
TMEA poolstat for 53112 stopped
TMEA for 53112 stopped
Auto Transfer for 53112 stopped

```

```
Instance 53112 -- Collector: stopped, Autoxfer: disabled
Last ftp transfer: succeeded at Thu Nov 14 18:07:51 EST 2013
```

start = start the processes for the array

```
Enter number to select instance or 'cv' to create VSP instance and 'ca' to
create AMS instance.
  1 53112                stopped
  2 10102                stopped
  3 210060              stopped
? 1

Instance 53112 -- Collector: stopped, Autoxfer: disabled
Last ftp transfer: succeeded at Thu Nov 14 18:07:51 EST 2013

Enter operation code:
start -- Start the Collector
stat  -- Display detailed status of Collector instance
stop  -- Stop the Collector
endebug -- Enable the debug logs
disdebug -- Disable the debug logs
en    -- Enable auto transfer of data
fc    -- Force collection of configuration data
xd    -- Transfer all data to Analyzer
rm    -- Remove this instance
n     -- Select new instance
!CMD  -- Execute CMD in sh
q     -- Quit
? start
TMEA config for 53112 started
TMEA dkcbusy for 53112 started
TMEA ldevio for 53112 started
TMEA portio for 53112 started
TMEA pairstat for 53112 started
TMEA clprbusy for 53112 started
TMEA poolstat for 53112 started
TMEA for 53112 started
/opt/TMEA/bin/autoxfer started at 11/14/13 18:14:02
Auto Transfer for 53112 started

Instance 53112 -- Collector: running, Autoxfer: enabled
Last ftp transfer: succeeded at Thu Nov 14 18:07:51 EST 2013
```

28.1.10 Current Status

You may want to check the status of a given array (more than general running or stopped status). Run the manage command (pbrun /opt/TMEA/bin/manage). Select the id of the instance you want to check the status.

Use the stat command to check the status.

```
Enter number to select instance or 'cv' to create VSP instance and 'ca' to
create AMS instance.
  1 53112                running
? 1

Instance 53112 -- Collector: running, Autoxfer: enabled

Enter operation code:
stat  -- Display detailed status of Collector instance
stop  -- Stop the Collector
endebug -- Enable the debug logs
disdebug -- Disable the debug logs
dis   -- Disable auto transfer of data
fc    -- Force collection of configuration data
n     -- Select new instance
```

```
!CMD -- Execute CMD in sh
q    -- Quit
? stat
TMEA config for 53112 running
TMEA dkcbusy for 53112 running
TMEA ldevio for 53112 running
TMEA portio for 53112 running
TMEA pairstat for 53112 running
TMEA clprbusy for 53112 running
TMEA poolstat for 53112 running
Local output directory for csv files is /HDS_data/TMEA/53112
Owners should be set in /HDS_data/TMEA/53112/data/config/ldevmap.csv
ftp set to the /53112 directory of the citi ftp account on 207.126.248.84

Instance 53112 -- Collector: running, Autoxfer: enabled
```

28.1.11 Manual SFTP Transfer

It may be necessary to manually force the sftp transfer for a specific array. Run the manage command (pbrun /opt/TMEA/bin/manage). Select the id of the instance you want to manually force the transfer.

Run the fc command to reinitiate the transfer of data to HDS.

```
Enter number to select instance or 'cv' to create VSP instance and 'ca' to
create AMS instance.
  1 53112                      running
? 1

Instance 53112 -- Collector: running, Autoxfer: enabled

Enter operation code:
stat -- Display detailed status of Collector instance
stop -- Stop the Collector
endebug -- Enable the debug logs
disdebug -- Disable the debug logs
dis -- Disable auto transfer of data
fc -- Force collection of configuration data
n -- Select new instance
!CMD -- Execute CMD in sh
q -- Quit
? fc
TMEA config for 53112 stopped
TMEA config for 53112 started
waiting 60 seconds for configuration collection to complete

Instance 53112 -- Collector: running, Autoxfer: enabled
```

28.1.12 Removing an Instance

Upon retirement of an array (or removal from TMEA reporting), you will want to delete the instance. Run the manage command (pbrun /opt/TMEA/bin/manage). Select the id of the instance you want to delete and then type the command "rm".

```
Enter number to select instance or 'cv' to create VSP instance and 'ca' to
crea
te AMS instance.
  1 53112                      stopped
  2 10102                      stopped
  3 210060                     stopped
? 3

Instance 210060 -- Collector: stopped, Autoxfer: disabled
Last ftp transfer: succeeded at Thu Nov 14 14:53:52 EST 2013

Enter operation code:
```

```

start -- Start the Collector
stat -- Display detailed status of Collector instance
stop -- Stop the Collector
enddebug -- Enable the debug logs
disdebug -- Disable the debug logs
en -- Enable auto transfer of data
fc -- Force collection of configuration data
xd -- Transfer all data to Analyzer
rm -- Remove this instance
n -- Select new instance
!CMD -- Execute CMD in sh
q -- Quit
? rm
If you are sure you want to remove instance 210060, enter 'YES': YES
Removing instance ...

Instance Name: 210060 removed successfully

```

Note: The /opt/TMEA/properties directory contains all of the configuration files for each instance. Removing these files would be needed if a manual deletion is necessary. The output directory (i.e. /HDS_data/TMEA/<serial num>) would need to be empty before trying to re-add the instance.

28.1.13 Removing TMEA Collector

When retiring a collector server (or for other reasons recommended by HDS), it may be desirable to remove the TMEA Collector software. This activity requires root access. The rpm command is used to manage the package.

Run the rpm -qa command to find out the current package name.

```

[root@namgdse04 ~]# rpm -qa | grep TMEA
TMEA-9_5_0b11-0.el5.i386

```

Run the rpm -e command to remove the package. Then verify that the package no longer exists.

```

[root@namgdse04 ~]# rpm -e --nodeps TMEA-9_5_0b11-0.el5.i386
warning: /opt/TMEA/properties/devicenumbers.config saved as
/opt/TMEA/properties
/devicenumbers.config.rpm.save
warning: /opt/TMEA/properties/col2anal.config saved as
/opt/TMEA/properties/col2
anal.config.rpm.save
[root@namgdse04 ~]# rpm -qa | grep TMEA
[root@namgdse04 ~]#

```

28.1.14 TMEA Space Cleanup Maintenance (Cleanup Script)

The SA for the TMEA server will need to create a cleanup script to maintain the amount of storage used by TMEA. TMEA does not automatically clean itself up with the current release. The following command should be run on a daily basis to maintain the space estimate of 10GB per array (keeping only 7 days of data locally):

```

[root@se04 ~]# find /HDS_data/TMEA/ -name "*f_*.zip" -mtime +7 -exec rm -rf {} \;

```

This will remove files from the /HDS_data/TMEA/<array serial>/archive/ directory and should not affect other files.

*Note: Currently these archive files are root owned, so pbrun may need to be utilized. There is considerable risk running the command "rm -rf" from the root account, so be sure that the command search criteria (-name "*f_*.zip" -mtime +7) is typed in correctly to prevent unintended deletions.*

28.1.15 Process Impact to RHEL Server

To determine the memory / cpu impact to the REHL server hosting TMEA collections, the following command can be run (not root dependent):

```
[testuser1@namgdse04 ~]$ ps -eo pmem,pcpu,pid,user,args | grep TMEA
0.0 0.0 537 root /bin/sh /opt/TMEA/bin/autotfer VSP
0.1 0.0 11796 root /opt/TMEA/HiPERFC -f /opt/TMEA/properties/config_53112.properties
0.1 0.0 11801 root /opt/TMEA/HiPERFC -f /opt/TMEA/properties/dkcbusy_53112.properties
0.1 0.1 11808 root /opt/TMEA/HiPERFC -f /opt/TMEA/properties/ldevio_53112.properties
0.1 0.0 11814 root /opt/TMEA/HiPERFC -f /opt/TMEA/properties/portio_53112.properties
0.0 0.0 11824 root /opt/TMEA/HiPERFC -f /opt/TMEA/properties/pairstat_53112.properties
0.0 0.0 11833 root /opt/TMEA/HiPERFC -f /opt/TMEA/properties/clprbusy_53112.properties
13.5 0.0 11842 root /opt/TMEA/HiPERFC -f /opt/TMEA/properties/poolstat_53112.properties
0.0 0.0 11886 root /bin/sh /opt/TMEA/bin/autotfer VSP
0.5 0.0 12430 root /opt/TMEA/HiPERFC -f /opt/TMEA/properties/config_10102.properties
0.1 0.0 12435 root /opt/TMEA/HiPERFC -f /opt/TMEA/properties/dkcbusy_10102.properties
0.1 0.2 12440 root /opt/TMEA/HiPERFC -f /opt/TMEA/properties/ldevio_10102.properties
0.1 0.3 12445 root /opt/TMEA/HiPERFC -f /opt/TMEA/properties/portio_10102.properties
0.0 0.0 12450 root /opt/TMEA/HiPERFC -f /opt/TMEA/properties/pairstat_10102.properties
0.0 0.0 12455 root /opt/TMEA/HiPERFC -f /opt/TMEA/properties/clprbusy_10102.properties
2.4 0.0 12460 root /opt/TMEA/HiPERFC -f /opt/TMEA/properties/poolstat_10102.properties
0.0 0.0 12468 root /bin/sh /opt/TMEA/bin/autotfer VSP
0.2 0.0 23500 root /opt/TMEA/HiPERFC -f /opt/TMEA/properties/config_210060.properties
0.0 0.0 23507 root /opt/TMEA/HiPERFC -f /opt/TMEA/properties/dkcbusy_210060.properties
0.1 0.0 23517 root /opt/TMEA/HiPERFC -f /opt/TMEA/properties/ldevio_210060.properties
0.0 0.0 23532 root /opt/TMEA/HiPERFC -f /opt/TMEA/properties/portio_210060.properties
0.0 0.0 23539 root /opt/TMEA/HiPERFC -f /opt/TMEA/properties/pairstat_210060.properties
0.0 0.0 23547 root /opt/TMEA/HiPERFC -f /opt/TMEA/properties/clprbusy_210060.properties
3.5 0.0 23560 root /opt/TMEA/HiPERFC -f /opt/TMEA/properties/poolstat_210060.properties
```

28.1.16 Debug Mode

To troubleshoot problems with particular arrays or sites, it may be useful to enable debug mode for the affected arrays. Use the manage command to enable or disable debug mode. Enabling or disabling debug mode will restart the processes for an array. Enabling debug mode will generate additional log files that may need to be uploaded to HDS Datamart / TMEA Support.

Enable Debug Mode:

```
[root@namgdse04 bin]# ./manage
Enter number to select instance or 'cv' to create VSP instance and 'ca' to
create AMS instance.
 1 53112          running
 2 10102          running
 3 210060         running
? 1

Instance 53112 -- Collector: running, Autotfer: enabled
Last ftp transfer: succeeded at Thu Dec 19 10:22:46 EST 2013

Enter operation code:
stat -- Display detailed status of Collector instance
stop -- Stop the Collector
enddebug -- Enable the debug logs
disdebug -- Disable the debug logs
dis -- Disable auto transfer of data
fc -- Force collection of configuration data
n -- Select new instance
!CMD -- Execute CMD in sh
q -- Quit
? enddebug
TMEA config for 53112 stopped
TMEA dkcbusy for 53112 stopped
TMEA ldevio for 53112 stopped
TMEA portio for 53112 stopped
TMEA pairstat for 53112 stopped
TMEA clprbusy for 53112 stopped
TMEA poolstat for 53112 stopped
TMEA for 53112 stopped
Auto Transfer for 53112 stopped
```

```

TMEA config for 53112 started
TMEA dkcbusy for 53112 started
TMEA ldevio for 53112 started
TMEA portio for 53112 started
TMEA pairstat for 53112 started
TMEA clprbusy for 53112 started
TMEA poolstat for 53112 started
TMEA for 53112 started
/opt/TMEA/bin/autoxfer started at 12/19/13 10:38:15
Auto Transfer for 53112 started

Instance 53112 -- Collector: running, Autoxfer: enabled
Last ftp transfer: succeeded at Thu Dec 19 10:22:46 EST 2013

```

Disable Debug Mode:

```

[root@namgdse04 bin]# ./manage
Enter number to select instance or 'cv' to create VSP instance and 'ca' to
create AMS instance.
 1 53112                running
 2 10102                running
 3 210060               running
? 1

Instance 53112 -- Collector: running, Autoxfer: enabled
Last ftp transfer: succeeded at Thu Dec 19 10:22:46 EST 2013

Enter operation code:
stat -- Display detailed status of Collector instance
stop -- Stop the Collector
endebug -- Enable the debug logs
disdebug -- Disable the debug logs
dis -- Disable auto transfer of data
fc -- Force collection of configuration data
n -- Select new instance
!CMD -- Execute CMD in sh
q -- Quit
? disdebug
TMEA config for 53112 stopped
TMEA dkcbusy for 53112 stopped
TMEA ldevio for 53112 stopped
TMEA portio for 53112 stopped
TMEA pairstat for 53112 stopped
TMEA clprbusy for 53112 stopped
TMEA poolstat for 53112 stopped
TMEA for 53112 stopped
Auto Transfer for 53112 stopped
TMEA config for 53112 started
TMEA dkcbusy for 53112 started
TMEA ldevio for 53112 started
TMEA portio for 53112 started
TMEA pairstat for 53112 started
TMEA clprbusy for 53112 started
TMEA poolstat for 53112 started
TMEA for 53112 started
/opt/TMEA/bin/autoxfer started at 12/19/13 10:43:34
Auto Transfer for 53112 started

```

28.1.17 Automatic Process Management (Keeping TMEA Running)

For large deployments of TMEA, there may be some stability issues with keeping processes running. To keep the processes running automatically, add the autostart script to the crontab (SA should complete this). The changes to the crontab are **highlighted** below. You may have to adjust the mode settings to match other executables in the bin directory (chmod 0755 autostart.sh).

```

[root@namgdse04 bin]# ls -al /opt/TMEA/bin/autostar*
-rwxr-xr-x 1 root root 1332 Dec 19 12:46 /opt/TMEA/bin/autostart.sh

```

```
[root@namgdse04 etc]# crontab -l
# rinfo job starting line
14 02 * * * /opt/rinfo/bin/rinfosend >/var/opt/rinfo/rinfosend.log 2>&1
# rinfo job ending line
0 23 * * * /opt/ecc/exec/check_ecc_processes.sh >/dev/null 2>&1
# SCpacct Entry - DO NOT REMOVE
5 16 * * * /opt/SCpacct/bin/preprocess # SCpacct Entry - DO NOT REMOVE
# SCpacct Entry - DO NOT REMOVE
#
# HDS TMEA Manual Entry 12192013 - checks processes every 30 mins
*/30 * * * * /opt/TMEA/bin/autostart.sh
# HDS TMEA Manual Entry - End
```

When running this option, debug mode should also be enabled.

This script will check to make sure all of the processes are running every thirty minutes.

28.2 TMEA Reporting Website

28.2.1 URL for TMEA

The following URL and account information is provided for regional and global access

<https://tmea-emea.hds.com/tmea/>

This website allows for account maintenance and for report configuration.

28.2.2 Accounts for Accessing TMEA

These accounts are to be used by all Citi SAN staff. The passwords should be changed according to normal account policies.

- Global Account (All Regions)
citiglobal (pass=start123)
- NAM (and LATAM)
citinam (pass=start123)
- EMEA
citiemea (pass=start123)
- APAC
citiapac (pass=start123)

28.2.3 Distribution Lists for Reports

Each account has an email associated with it. This email address is not the email address that is used by the reports. The report distribution list is managed by the TMEA support team. This distribution list would ideally send the reports to the sharepoint sites maintained by each SAN management team/region rather than filling up every SAN Administrators' email accounts with duplicated reports.

28.3 TMEA Support

Normal HDS Escalation procedures should be followed via global support numbers or the hds portal (portal.hds.com). For small changes and/or minor issues the following email address can be used optionally:

datamart.support@hds.com (standard escalation)

For example, it may be easier to send an email to the TMEA team to add a new array to an account and to start processing the data. The traditional method of opening support tickets is still supported and encouraged.

29 Appendix O: SVP Performance Collection

In order to shorten the response time from HDS Support on performance issues, the performance data from the array must be given to HDS as soon as possible. This can be achieved by running a daily export of the performance data and configuring the browser on the SVP to be able to send data to the HDS TUF (Support) website. The following steps should be followed when configuring this on a VSP or HUSVM array:

- Configure “exports” local account on storage navigator
- Post Microcode Update Steps
- Configure Browser

29.1 Configure Exports Account

These instructions cover the necessary steps to implement the performance collection script on the SVP for both VSP and HUS VM arrays. The initial task is to create “exports” account for collecting the data. This account is limited to only this function and cannot be used to make other administrative changes.

Create user group Performance Administrator

The screenshot shows the Hitachi Storage Navigator Administration web interface. The left sidebar contains the Explorer menu with options like Storage Systems, Performance Monitor, Administration, User Groups, Resource Groups, and Encryption Keys. The main content area is titled 'User Groups' and shows a summary of 9 user groups. Below the summary is a table of built-in user groups.

User Group Name	User Group Type	Number of Roles	Number of Resource Groups	Number of Users	All Resource Groups Assigned
Administrator User Group	Built-in	8	2	5	Yes
Audit Log Administrator (View & Modify) User Group	Built-in	2	2	0	Yes
Audit Log Administrator (View Only) User Group	Built-in	2	2	0	Yes
Security Administrator (View & Modify) User Group	Built-in	3	2	0	Yes
Security Administrator (View Only) User Group	Built-in	3	2	0	Yes
Storage Administrator (View & Modify) User Group	Built-in	6	1	0	No
Storage Administrator (View Only) User Group	Built-in	1	1	0	No
Support Personnel Group	Built-in	7	2	1	Yes
System User Group	Built-in	8	2	0	Yes

At the bottom of the interface, there are buttons for 'Create User Group', 'Add Users', 'Edit Resource Group Assignment', and 'More Actions'.

Create User Group

1. Create User Group > 2. Assign Roles > 3. Assign Resource Groups > 4. Confirm

This wizard lets you create a user group. Enter a name of the user group you want to create, and click Check to confirm whether the name is available or not on the external authentication server. Click Next to add roles to this user group.

User Group Name:

(Max. 64 Characters)

Back Next Finish Cancel Help

Add Performance Management role

Create User Group

1. Create User Group > 2. Assign Roles > 3. Assign Resource Groups > 4. Confirm

To assign roles to the user group, select roles from the Unassigned Roles list and click Add. Click Next to assign resource groups to the user group.

Unassigned Roles	
Role	
Audit Log Administrator (View & Modify)	
Audit Log Administrator (View Only)	
Security Administrator (View & Modify)	
Security Administrator (View Only)	
Storage Administrator (Initial Configuration)	
Storage Administrator (Local Copy)	
Storage Administrator (Provisioning)	
Storage Administrator (Remote Copy)	
Storage Administrator (System Resource Management)	
Support Personnel	

Selected: 0 of 10

Add Remove

Assigned Roles	
Role	
Storage Administrator (Performance Management)	
Storage Administrator (View Only)	

Selected: 1 of 2

Back Next Finish Cancel Help

Add resource group and click finish

Create User Group

1.Create User Group > 2.Assign Roles > **3.Assign Resource Groups** > 4.Confirm

When All Resource Groups Assigned is No, select resource groups from the Unassigned Resource Groups list and click Add to assign resource groups to the user group. Click Next to add users to the user group or Finish to confirm.

All Resource Groups Assigned: No

Unassigned Resource Groups					
Resource Group Name (ID)	Number of User Groups	Number of Parity Groups	Number of LDEVs	Number of Ports	Number of Host Groups
testndm (1)	7	0	12	2	1

Selected: 0 of 1

Add Remove

Assigned Resource Groups					
Resource Group Name (ID)	Number of User Groups	Number of Parity Groups	Number of LDEVs	Number of Ports	Number of Host Groups
meta_resource ...	9	19	65267	30	

Selected: 1 of 1

Next Task Option : Continue to Add Users

Back Next Finish Cancel Help

Apply task

Create User Group

1.Create User Group > 2.Assign Roles > 3.Assign Resource Groups > **4.Confirm**

Enter a name for the task. Confirm the settings in the lists and click Apply to add task in Tasks queue for execution.

Task Name: (Max. 32 Characters)

Create User Group					
User Group Name	Number of Roles	Number of Resource Groups	Number of Users	All Resource Groups Assigned	
Performance Administrator	2	1	0	No	

Total: 2

Assigned Roles					
Role					
Storage Administrator (Performance Management)					
Storage Administrator (View Only)					

Total: 2

Assigned Resource Groups					
Resource Group Name (ID)	Number of User Groups	Number of Parity Groups	Number of LDEVs	Number of Ports	Number of Host Groups
meta_resource ...	9	19	65267	30	5855

Total: 1

Go to tasks window for status

Back Next Apply Cancel Help

```
Create user: exports pw: exports
```

[illegible]

Create User

1. Create User > 2. Confirm

Set values for the new user account and click Finish to confirm.

User Name:
(Max 256 Characters)

Account Status: ☒ Enable ☐ Disable

Authentication: ☒ Local ☐ External

Password:
 (6 - 256 Characters)

Re-enter Password:

Back Next Finish Cancel Help

Create User

1. Create User > 2. Confirm

Enter a name for the task. Confirm the settings and click Apply to add task in Tasks queue for execution.

Task Name: (Max. 32 Characters)

Selected User	
Item	Value
User Name	exports
Account Status	Enabled
Authentication	Local
Password	*****
User Group Name	Performance Administrator

☐ Go to tasks window for status

Hitachi Storage Navigator

File Maintenance Actions Reports Settings View Help

Alert Audit Log Resource Locked Logged in as: maintenance

Explorer

- Storage Systems
 - HDS53112(S/N:53112)
 - Tasks
 - Reports
 - Components
 - Parity Groups
 - Logical Devices
 - Pools
 - Ports/Host Groups
 - External Storages
 - Replications
- Performance Monitor
- Administration
 - General Tasks
 - Create Host Groups
 - Create LDEVs
 - Add LUN Paths
 - Create Pools
 - Add External Volumes

more...

Storage Systems

Tasks

HDS53112(S/N:53112) > Tasks

Summary

Completed	In Progress	Waiting	Suspended	Failed
8	1	0	0	0

Tasks

Filter: ON OFF Select All Pages Options 1 / 1


Task Name	Status	Type	User Name	Submission Time	Start Time	End Time	Auto Delete
140116-Cre...	12% In...	Create User	mainten...	2014/01/16 16:24:58	2014/01/16 16:25:07		Enabled
140116-Cre...	Completed	Create User...	mainten...	2014/01/16 16:21:59	2014/01/16 16:22:03	2014/01/16 16:22:56	Enabled
140115-Cre...	Completed	Create Host...	gj06529	2014/01/15 20:06:18	2014/01/15 20:06:22	2014/01/15 20:08:26	Enabled
140115-Cre...	Completed	Create Host...	gj06529	2014/01/15 19:50:05	2014/01/15 19:50:08	2014/01/15 19:52:02	Enabled
131119-Edit...	Completed	Edit Host	jh53740	2013/11/19 16:22:54	2013/11/19 16:22:55	2013/11/19 16:24:18	Enabled
131112-Cre...	Completed	Create User	mainten...	2013/11/12 14:42:58	2013/11/12 14:44:56	2013/11/12 14:47:38	Enabled
131112-Cre...	Completed	Create User	mainten...	2013/11/12 14:42:09	2013/11/12 14:45:06	2013/11/12 14:45:52	Enabled
131112-Cre...	Completed	Create User	mainten...	2013/11/12 14:41:45	2013/11/12 14:43:24	2013/11/12 14:44:06	Enabled
131112-Cre...	Completed	Create User	mainten...	2013/11/12 14:41:05	2013/11/12 14:41:23	2013/11/12 14:42:09	Enabled

Selected: 1 of 9


29.2 Initial Collection and Post Microcode Update Steps

After every microcode upgrade and before the initial performance script collection the HDS CE must validate the latest export.exe is available on the SVP. The following steps should be followed to ensure this. If this step is missed the collection script will fail.


Copy latest export.exe to D:\



Run export.exe from D:\ and confirm the Install Directory is as follows:



Click yes when prompted to overwrite the old .jar files



Two performance collection packages are available one for VSP and one for HUS VM.
Package names are:

HUSVM_export .zip or VSP_export.zip

1. copy pkg to svp under d:\
2. create export folder
3. extract files to d:\export
4. Edit archive.bat and change Device & serial number to match local arrays:


Set FILE_NAME=VSP-53112-%MMDDYYYY%%HHMM%

Set FILE_NAME=HUSVM-210060-%MMDDYYYY%%HHMM%

29.3 Create Daily Scheduled Tasks

A task must be created to run the exports 1x per day. A task must also be created to cleanup and retain only 30 days of exports.

Admin Tools>Task Scheduler>Create basic task



Create Basic Task Wizard

Create a Basic Task

Use this wizard to quickly schedule a common task. For more advanced options or settings such as multiple task actions or triggers, use the Create Task command in the Actions pane.

Trigger
Daily

Action
Start a Program

Finish

Name: run RAIDperf

Description: Performance data collection for HDS arrays

< Back Next > Cancel

Create Basic Task Wizard

Task Trigger

When do you want the task to start?

Trigger
Daily

Action

Finish

One time

When the computer starts

When I log on

When a specific event is logged

< Back Next > Cancel

Create Basic Task Wizard

 Daily

Create a Basic Task

Trigger

Daily

Action

Finish

Start: 1/16/2014 12:00:00 AM ☐ Universal time

Repeat every: 1 days

< Back Next > Cancel

Create Basic Task Wizard

 Action

Create a Basic Task

Trigger

Daily

Action

Finish

What action do you want the task to perform?

☒ Start a program

☐ Send an e-mail

☐ Display a message

< Back Next > Cancel

Create Basic Task Wizard

 Start a Program

Create a Basic Task

Trigger
Daily

Action
Start a Program

Finish

Program/script: D:\export\run RAIDperf.bat Browse...

Add arguments (optional):

Start in (optional): D:\export\

< Back Next > Cancel

Create Basic Task Wizard

 Summary

Create a Basic Task

Trigger
Daily

Action
Start a Program

Finish

Name: run RAIDperf

Description: Performance data collection for HDS arrays



Trigger: Daily; At 12:00 AM every day

Action: Start a program; D:\export\run RAIDperf.bat


☐ Open the Properties dialog for this task when I click Finish

When you click Finish, the new task will be created and added to your Windows schedule.

< Back Finish Cancel



Create daily job to cleanup performance files over 30 days old
Admin Tools>Task Scheduler>Create basic task Setup cleanup program



Create Basic Task Wizard

Create a Basic Task


Use this wizard to quickly schedule a common task. For more advanced options or settings such as multiple task actions or triggers, use the Create Task command in the Actions pane.

Trigger
Action
Finish

Name: cleanup RAID_perf

Description: Cleanup performance data collection files for HDS arrays older than 7 days

< Back Next > Cancel



Create Basic Task Wizard

Task Trigger

When do you want the task to start?

☒ Daily
☐ Weekly
☐ Monthly
☐ One time
☐ When the computer starts
☐ When I log on
☐ When a specific event is logged

< Back Next > Cancel

Create Basic Task Wizard

Daily

Create a Basic Task

Trigger

Start: 1/21/2014 8:00:00 AM ☐ Universal time

Repeat every: 1 days

Action

Finish

< Back Next > Cancel

Create Basic Task Wizard

Action

Create a Basic Task

Trigger

Daily

What action do you want the task to perform?

Action

Finish

☒ Start a program
☐ Send an e-mail
☐ Display a message

< Back Next > Cancel

The image displays two screenshots of the Windows 'Create Basic Task Wizard'.

Top Screenshot: Start a Program

- Title Bar:** Create Basic Task Wizard
- Left Panel:** Trigger (Daily), Action (Start a Program), Finish.
- Program/script:** D:\export\cleanup.bat (with a 'Browse...' button).
- Add arguments (optional):** (Empty text box).
- Start in (optional):** D:\export\
- Buttons:** < Back, Next >, Cancel.

Bottom Screenshot: Summary

- Title Bar:** Create Basic Task Wizard
- Left Panel:** Trigger (Daily), Action (Start a Program), Finish (highlighted).
- Name:** cleanup RAID perf
- Description:** Cleanup performance data collection files for HDS arrays older than 7 days
- Trigger:** Daily; At 2:00 AM every day
- Action:** Start a program; D:\export\cleanup.bat
- Checkbox:** ☐ Open the Properties dialog for this task when I click Finish
- Text:** When you click Finish, the new task will be created and added to your Windows schedule.
- Buttons:** < Back, Finish, Cancel.


29.4 Configuration of Browser

In order to upload the performance collection files quickly to HDS support, the SVP needs the browser to be configured for access to <https://tuf.hds.com>.

Note: On SVPs that have RACC installed, the Citi proxy needs to be added to the same (application) proxy settings as what was used by the RACC application. This is because

RACC uses IE for its proxy settings. This will require an additional firewall rule as RACC and TUF do not use the same IP address at HDS.


The following example shows the webproxy settings and should only be used on non RACC SVPs. SVPs with RACC should use the same application proxy configured on both IE and RACC.




30 Appendix P: Setup of AD for Storage Navigator

To be able to use AD based accounts on Storage Navigator (instead of or in addition to depending on HCS), the following procedures will need to be followed. **This is an optional configuration and is not required.**


1. Insure UDP port 1812 has been opened from the HDS SVP's to the HCS radius server (if ACL needs to be both ways).
2. Add the HDS SVP's as radius clients to existing HCS radius server. Make sure the shared password you enter matches the password used in the radius configuration file created later in this document.




3. You will need a certificate file to give to HDS for use on the SVP. I found the easiest route to accomplish this was to commandeer one already created by Citi. Launch IE and select tool→internet options. Then click the Content tab and then click the certificates button.






4. Select a certificate from the Trusted Root Certification Authorities tab (with a long expiration) from Citi and click export.



5. Follow the screenshots below...



6. Call it Citi.cer (or something similar)



8. Setup the radius configuration file...

Syntax...

```
auth.server.type=radius
auth.server.name=<radius server name>
auth.radius.<radius server name>.protocol=pap
auth.radius.<radius server name>.host=<radius server IP>
auth.radius.<radius server name>.port=1812
auth.radius.<radius server name>.timeout=3
auth.radius.<radius server name>.secret=<radius client password>
auth.radius.<radius server name>.retry.times=3
auth.radius.<radius server name>.attr.NAS-Identifier=<radius server IP>
auth.radius.<radius server name>.domain.name=<radius server domain>
```

Sample of a completed radius file...

```
auth.server.type=radius
auth.server.name=namrhds01
auth.radius.namrhds01.protocol=pap
auth.radius.namrhds01.host=150.110.197.97
auth.radius.namrhds01.port=1812
auth.radius.namrhds01.timeout=3
auth.radius.namrhds01.secret=radiuspass (example)
auth.radius.namrhds01.retry.times=3
auth.radius.namrhds01.attr.NAS-Identifier=150.110.197.97
auth.radius.namrhds01.domain.name=nam.nsroot.net
```

9. Provide the files to the HDS CE and have him/her run the next steps (or you can run them if you have RDP access to the HDS SVP). Just open notepad on the SVP and paste the contents of the radius and certificate files.
10. Copy radius configuration and certificate file to c:\DKC200\TMP on the Primary SVP
11. Execute c:\DKC200\TMP>setexauthconf <radius.conf> <cert.file> from a command prompt

Sample...

```
c:\DKC200\TMP>setexauthconf sn2_ruth_exauth.properties Citi.cer

Reading keystore file...success

Reading certification file...success

Encrypting property file...success

Updating to keystore file...success

All process completed.
Press any key to continue . . .
```

12. Change a user account that's a member of AD to external via storage navigator

Edit User

1. Edit User > 2. Confirm

This wizard lets you edit one or more properties. Check the box in front of the property you want to edit, and then enter the new value.

User Name: kh64535

☐ Account Status: ☒ Enable ☐ Disable


☒ Authentication: ☐ Local ☒ External

Password:
(6 - 256 Characters)

Re-enter Password:

Back Next Finish Cancel Help

13. Verify authentication is external and test it...



14. If successful there will also be a radius message in the Windows Event View under System; IAS; as follows...

Event Type: Information
 Event Source: IAS
 Event Category: None
 Event ID: 1
 Date: 8/19/2011
 Time: 4:11:40 PM
 User: N/A
 Computer: NAMRHDS01
 Description:
 User kh64535 was granted access.
 Fully-Qualified-User-Name = NAM\kh64535
 NAS-IP-Address = <not present>
 NAS-Identifier = 150.110.197.97

Client-Friendly-Name = ORH052N5029-SPA

Client-IP-Address = 172.31.226.167

Calling-Station-Identifier = <not present>

NAS-Port-Type = <not present>

NAS-Port = <not present>

Proxy-Policy-Name = Use Windows authentication for all users

Authentication-Provider = Windows

Authentication-Server = <undetermined>

Policy-Name = Device Manager User Group 1

Authentication-Type = PAP

EAP-Type = <undetermined>

31 Appendix Z: Version Notes

1.0 This is the initial release of the HDS VSP / HUS-VM Solution Guide

1.1 RFP Release of SG for VSP

- Modified HTSM content to remove legacy role [23.1]
- Added list of Citi HTnM Alarms [27.4.2]
- Added Appendix N: Setup of AD for Storage Navigator [28]

1.2 Minor release

- Updated package list [11.2]

1.3 Minor Release for HCS 7.2.1

- Updated LDEV reuse procedures [17.15, ref from best practices]
- Added NDM Seamless Migrations [12.30]
- Added SCSI Reservations on VSP (View / Edit) [12.29]
- Updated Package List and notes [11.3.2, 11.3.5, 11.3.4]

2.0 HUS-VM and 14+2

- Renamed document dropping TagmaStore, adding HUS-VM [Multiple Sections]
- Using HCS instead of HDvM [Multiple Sections]
- Added HUS-VM Hardware [10.3, 12.5]
- Updated Drive and RAID Configurations [12.2, 12.3, 12.4, 12.5]
- Updated WWN Numbering for Host Ports [12.8.1]
- Updated Max Limits for Pools, Array, etc. [10.5, 12.26]
- Updated Package List [11.2]
- Updated Active Directory via Radius to reflect it is a secondary option [13.12]
- Added Active Directory Configuration (direct to AD) [13.11]
- Updated SANScreen to NetApp OnCommand Insight [Multiple Sections]
- Updated RACC and removed VPN dial-in [12.13]
- Added Maximum Ports Per Host [12.11, renumbered all sections 12.12+]
- Updated Power, Dimensions, and Weight [15]
- Added Certificate Imports [13.14]
- Updated RACC and removed VPN-based dial-up [12.14]
- Added HCS Client Settings [13.13]
- Re-added SSL configuration [13.15]

2.1 HUS-VM Host Mode 55 and Other Updates

- Added Shared Memory Requirements [12.33]
- Added Host Mode 55 for HUS-VM for Powerpath [12.2, updated 12.2+]
- Updated Certificate Imports [13.14]
- Corrected SNMP abbreviations [12.6, 12.14.1]

2.2 SOM 729/803 and Modified HMO 55 Added "List Logins" / Login WWNs Function [17.16]

- Updated HDP LUNS for 14+2 RAID Type and SOM917 [12.27]
- Updated package list and details about packages [11.2, 11.3]
- Updated SOM List for 729, 803 [12.27]
- Added Pool Full Procedures [12.28.10]
- Updated Appendix F Firewall ports for SNMP – clarified source/target
- Modified HMO 55 recommendation for all hosts [12.2]

2.3 Standard configurations, latency response rating, and reliability

- Updated RAID configurations: R10 (2+2) – Silver tier, R6 (14+2) – Bronze tier. Removed R6 (6+2) [12.3]
- Drive types: removed 146Gb 15K SAS, old gen SSD, and 1TB, 2TB NL-SAS drive types. Added 300Gb 15K SAS, 1.6TB FMD, and 3TB SAS [12.4]
- Drive rebuild times: removed old gen SSD and 1TB NL-SAS drive types. Added 1.6TB FMD [12.5]
- Standard array configurations: defined latency response and reliability ratings [12.6]
 - o Silver tier R10 (2+2) VSP (including physical layout) and HUS VM
 - o Bronze tier R6 (14+2) HUS VM
 - o Replicated VSP and HUS VM configurations
- DP-Vols: standard LUN sizes [12.28.7]

2.4 Update for TMEA Collector

- Updated HTnM installation for RHEL [27.1, 27.2]
- Removed references to running HTnM on separate server from HCS [27]
- Updated HTnM Perf Reporter Configuration for 60 days data storage [27]
- Added TMEA Collector installation and setup [28]
- Renamed Appendix N: AD for Storage Navigator to Appendix O [29+]

2.4.1 Second Update for TMEA Collector

- Updated RMLIB installation [28.1.4]
- Added Debug Mode [28.1.16]
- Changed command device reference from wwn to scsi [28.1.7]
- Added Automatic Process Management [28.1.17]
- Updated HTnM setup for scsi ids [27]

2.5 All FMD Configuration Options & Performance Collection

- Added DBF drive box [10.3.2.2]
- RAID10 protection [10.4.4, 12.3]
- FMD max drives [10.5]
- FMD spare count [12.7]
- All flash configuration options [12.6.5]
- Added Appendix O – Performance collection script and tuf.hds.com on SVP [29, updated 29+]
- Updated Remote Support RACC for SVP [12.15.3]
- Added SAN “Port Flapping” Alerting [12.34]
- Updated HTnM Reports [27.4]
- Updated Automatic Process Management [28.1.17]