



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

HDS TagmaStore USP/USPV Solutions Guide

GDSE-05015-SG

Solutions Guide

2.9.5

1 Version Control

Version	Date	Author	Reviewed	Comments
2.0	14-Sep-07	Gary Johnson	Chris Clayton	USPV, HTSM, & VPN
2.1	26-Sep-07	Gary Johnson	Global Storage Community	Disk & LDEV Timing
2.2	02-Nov-07	Gary Johnson	Minor Release	Package Names
2.3	11-Feb-07	Gary Johnson	Minor Release	Minor Edits
2.4	28-Feb-08	Gary Johnson	Parul Jain	SNMP & Misc Edits
2.5	14-Jul-08	Gary Johnson	Global Storage Community	HDP, Tuning Mgr, Syslog
2.6	27-Oct-08	Gary Johnson Chris Clayton	Global Storage Community	HDP, HDvM, HTnM, HTSM
2.7	30-Dec-08	Gary Johnson	Kevin Hobaugh	USPVM
2.8	23-Jan-09	Gary Johnson	Kevin Hobaugh	Zero Data Removal Tool
2.9	25-Jun-09	Gary Johnson	Teo, Kee-Huat	Lower cost tiering for replication
2.9.1	17-Jul-09	Gary Johnson	Minor Release	Updated HTnM AFR Solaris
2.9.2	18-Dec-09	Gary Johnson	Global Storage Community	Updated HDvM & HTnM
2.9.3	30-Mar-10	Gary Johnson	Global Storage Community	Additional HDvM & HTnM Updates
2.9.4	25-Jul-12	Gary Johnson	Minor Release	SCSI Reservations
2.9.5	29-Apr-13	Jeff Hernandez	Gary Johnson	SOM 803 & DRU procedure

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	Reviewed By
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 “*GT Global Distributed Storage and Backup Engineering” distribution list or alternatively to gdse@citigroup.com.

5 Latest Versions

It is recommended that the latest version of the document be used from the GDSE website <http://gdse.ny.ssmc.com>. Newer versions may include expanded scope, or comments 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.....	12
8	RELEASE STATUS	12
9	PREQUISITES FOR A NEW HDS ENVIRONMENT	13
9.1	REQUIRED LICENSING FOR ARRAY (BASED ON RAW ARRAY STORAGE).....	13
9.2	REQUIRED LICENSING PER SITE (BASED ON TOTAL RAW STORAGE)	13
10	INSTALLATION.....	14
10.1	REQUIRED INSTALLATION STEPS.....	14
10.2	SOFTWARE PACKAGES	14
10.3	SOFTWARE PACKAGE DETAILS / NOTES	15
11	BEST PRACTICES	17
11.1	STORAGE HOST SETUP (HOST MODES)	17
11.2	RAID TYPES	18
11.3	DRIVE TYPES.....	19
11.3.1	Using Slower Drives on a USP.V.....	19
11.4	ARRAY TYPES	20
11.5	ARRAY CACHE	22
11.6	SPARE DRIVES.....	23
11.7	LOGICAL DEVICE (LDEV) SIZING AND CAPACITY	24
11.8	LOGICAL DEVICE (LDEV) NUMBERING	25
11.8.1	LDEVs when using TSM.....	26
11.9	LUN EXPANSION	27
11.9.1	Expansion with A New LUN	27
11.9.2	Expansion with Host Volume Management	27
11.9.3	Expansion with LUSE Volumes.....	27
11.10	LUSE VOLUMES	27
11.11	DRIVE REBUILD TIMES.....	28
11.12	LDEV CREATE/DESTROY TIMING	28
11.13	HOST PORT BALANCING AND REDUNDANCY	29
11.13.1	USP Host Cards.....	29
11.13.2	USPV Host Cards	31
11.13.3	Maximum Host Port Performance (USPV) Configuration	33
11.13.4	WWN Numbering for USP/USPV Host Ports	34
11.14	TOPOLOGY	34
11.15	MANAGEMENT SOFTWARE (DM vs. SN)	34
11.16	SAN MONITORING.....	35
11.17	ARRAY MONITORING: ALERTS AND REMOTE CONNECTIONS	35
11.17.1	SMNP	35
11.17.2	Vendor Notification (Dial-Out / HiTrack)	35
11.17.3	Vendor Remote Access (VPN).....	35
11.18	HOST PATH REDUNDANCY WITH DMP	36
11.19	HOST PATH MONITORING (WITH DMP)	37

11.20	RMI TIME-OUT.....	38
11.21	ARRAY CONFIGURATION BACKUP, REVIEW, AND LARGE-SCALE CHANGES	38
11.22	ORDERING (STANDARD CONFIGURATIONS)	39
11.22.1	USPV Configurations.....	39
11.22.2	USP Configurations (Legacy).....	41
11.23	NETWORK TIME PROTOCOL (SNTP OR NTP) – TIME SYNCHRONIZATION	42
11.24	MULTIPLE ARRAY MANAGEMENT	43
11.25	POLLING INTERVALS FOR DEVICE MANAGER.....	43
11.26	ENTERPRISE ALERTING (SNMP).....	44
11.27	VMWARE	46
11.28	HDS SOFTWARE SERVER SIZING	46
11.29	RUNNING MULTIPLE VERSIONS OF SERVER SOFTWARE SIMULTANEOUSLY	46
11.30	PHYSICAL VIEW (RUNNING SN FROM HDvM).....	46
11.31	USING MULTIPLE STORAGE TIERS (DISK TYPES) WITH REPLICATION.....	47
11.32	BACKUP, RESTORE, AND MIGRATION OF DEVICE MANAGER AND RELATED DATABASES.....	47
11.32.1	Backup HDvM Databases.....	47
11.32.2	Restore HDvM Databases.....	48
11.32.3	Restrictions to Backup and Restores.....	48
11.32.4	Migrations (Export and Import of HDvM databases).....	48
11.33	THIN PROVISIONING (HDP) ON USPv ARRAYS	50
11.34	MAXIMUM LUN SIZE	50
11.35	INCREASED HEAP SIZE (HDvM MEMORY EXPANSION)	50
11.36	SYSTEM OPTIONS AND SYSTEM DISKS	51
11.37	SCSI RESERVATIONS ON USPv (VIEW / MODIFY).....	53
11.37.1	Viewing Host SCSI Reservations	53
11.37.2	Modifying Host SCSI Reservations on USPv	55
12	SECURITY PRACTICES.....	58
12.1	INTERFACES TO THE USP/USPV	58
12.2	DEFAULT USERIDS	58
12.3	SPECIAL USER ACCOUNTS (FUNCTIONAL IDS).....	58
12.4	POWERBROKER CONFIGURATION (CLI AND CCI ONLY)	59
12.5	AUDIT LOGGING (eTRUST).....	63
12.5.1	Device Manager Auditing.....	63
12.6	TURNING ON SSL AND SECURING DEVICE MANAGER (WINDOWS).....	63
12.6.1	Creating a Keypair	63
12.6.2	Restarting the Device Manager Server	64
12.6.3	Enabling TLS/SSL Server Security.....	65
12.6.4	Generating a Private Key Using SSLC Utility.....	66
12.6.5	Creating a Certificate Signing Request	66
12.6.6	Creating a Self-Signed Certificate	66
12.6.7	Activating the SSL Configuration	67
12.6.8	Setting the SSL Host Name	67
12.6.9	Limiting Client IPs to Known Management Systems (ACL for HDvM).....	67
12.6.10	Using CLI with HTTPS	68
12.7	USER ENTITLEMENTS REPORTING	69
12.7.1	Storage Navigator User Entitlements	69
12.7.2	Device Manager User Entitlements	70
12.8	DECOMMISSIONING OF STORAGE	70
12.9	SETTING SECURITY POLICY IN DEVICE MANAGER.....	70
12.10	CCI SECURITY	70
12.11	HDvM CLI SECURITY	71
12.12	HDvM ARRAY NAME.....	71
12.13	ACTIVE DIRECTORY AUTHENTICATION WITH HDvM	71
12.13.1	IAS Installation and Setup	71
12.13.2	HDvM Configuration of Radius.....	80

12.13.3	<i>HDvM User Setup</i>	83
12.14	ICQ COMPENSATING CONTROLS AND FAILURES.....	84
12.14.1	<i>Recommended Level-3 Procedures</i>	85
12.14.2	<i>ICQ Section 2.8, 2.9 (Password Mgmt)</i>	86
12.14.3	<i>ICQ Section 2.11 (Functional ID Password)</i>	86
12.14.4	<i>ICQ Section 3.1 (Entitlement Review)</i>	86
12.14.5	<i>ICQ Section 3.2 (Audit Logs)</i>	86
12.14.6	<i>ICQ Section 3.4 (Audit Logs)</i>	86
12.14.7	<i>ICQ Section 5.3 (Audit Logs)</i>	86
12.14.8	<i>ICQ Section 6.4 (RBAC)</i>	86
12.14.9	<i>ICQ Section 6.5 (Maker/Checker)</i>	86
12.14.10	<i>ICQ Section 6.6 (Default Passwords)</i>	86
13	REFERENCES	87
14	APPENDIX A: MANAGEMENT STATION	88
14.1.1	<i>Number of Servers</i>	88
14.1.2	<i>Device Manager Client for Wintel</i>	89
14.1.3	<i>Device Manager Client for Solaris</i>	89
14.1.4	<i>Device Manager Command Line Client for Solaris (DM CLI)</i>	89
14.1.5	<i>Storage Navigator Client for Wintel</i>	89
14.1.6	<i>Storage Navigator Client for Solaris</i>	89
14.1.7	<i>Storage Navigator Server</i>	89
14.1.8	<i>Device Manager Windows 2k3/2k Server</i>	89
15	APPENDIX B: CONFIGURATION OF ARRAY VIA STORAGE NAVIGATOR	90
15.1	ACCESSING STORAGE NAVIGATOR	90
15.2	INFORMATION SCREEN	90
15.3	MODIFY MODE.....	92
15.4	SECURITY IN STORAGE NAVIGATOR.....	92
15.5	SVP PANEL (ARRAY LISTING) – ADDING NEW ARRAYS.....	95
15.6	ENVIRONMENTAL PARAMETERS.....	97
15.7	DOWNLOAD CONFIGURATION, USER ACCOUNTS, LOG FILES IN CONTROL PANEL.....	98
15.8	TRACE FILES DOWNLOAD (USPV ONLY).....	101
15.9	SYSLOG BASED AUDIT LOG	102
15.10	SYSLOG AUDIT RECORD FORMAT	103
15.10.1	<i>Syslog Server Setup</i>	106
15.11	LEGACY AUDIT LOG (USP ONLY)	107
15.12	PORT STATUS	108
15.13	SYSTEM EVENTS	109
15.14	PORT DEFINITION	110
15.15	RAID (PARITY) GROUP DEFINITION	110
15.16	VOLUME (LDEV) CREATION / DEFINITION	111
15.16.1	<i>Install Custom Volume (CV)</i>	111
15.16.2	<i>Make Volume</i>	115
15.17	VOLUME (LDEV) DELETION.....	116
15.17.1	<i>Volume to Space</i>	116
15.17.2	<i>Make Volume</i>	117
15.17.3	<i>Initialize Volume</i>	117
15.18	SPECIAL DISPLAY OF RAID 10 (4+4) PARITY GROUPS	117
15.19	LUSE VOLUMES	119
15.20	ASSIGNMENT LDEVs TO HOST GROUPS.....	124
15.21	CHANGING THE HBAS ON HOST GROUPS (HBA SWAP).....	129
15.22	REMOVING LUNS FROM HOST GROUPS (DELETING PATHS)	130
15.23	SETTING UP COMMAND DEVICES (FOR SHADOWIMAGE CCI)	131
15.24	DISK SHREDDING	132

15.25	SYSTEM STATUS.....	138
15.26	LICENSE KEYS.....	139
15.27	CONFIGURATION FILE LOADER (LARGE-SCALE CHANGES)	141
15.28	CONFIGURING PERFORMANCE DATA COLLECTION	142
16	APPENDIX C: HICOMMAND DEVICE MANAGER NAVIGATION AND USAGE	145
16.1	VERSION REQUIREMENTS.....	145
16.2	LOGGING INTO HDVM SERVER	145
16.3	USERS AND PERMISSIONS	147
16.3.1	Users	147
16.3.2	Permissions.....	149
16.3.3	Resource Groups.....	150
16.4	SECURITY	152
16.4.1	Password Policy	153
16.4.2	Failed Login Attempt Account Lock.....	153
16.4.3	Banner Page	154
16.5	SUBSYSTEMS	154
16.5.1	Add Subsystem	155
16.5.2	LUN Scan.....	157
16.6	HOSTS	158
16.6.1	Add Host	158
16.6.2	Modify Host.....	159
16.6.3	Add Storage.....	159
16.6.4	Remove Host.....	159
16.7	LOGICAL GROUPS	159
16.8	DEVICE MANGER REPORTS	161
16.9	ALERTS	162
16.10	TASKS	164
17	APPENDIX D: COMMAND LINE INTERFACE FOR DEVICE MANAGER.....	165
17.1	HiCOMMAND CLI SETUP (SOLARIS)	165
17.2	HiCOMMAND CLI SETUP (WINDOWS)	165
17.3	SECURE OPTION	165
17.4	SECURE PASSWORD PROMPTING	165
17.5	PROPERTIES FILE.....	166
17.5.1	Using properties file to store read-only username/password	166
17.6	CLI STRUCTURE.....	167
17.7	HOST MODES	168
17.8	ARRAY DISCOVERY.....	169
17.9	ADDING LUNS.....	170
17.10	CLI EXAMPLE USE CASE	173
18	APPENDIX E: COMMAND CONTROL INTERFACE / SHADOWIMAGE	180
18.1	CCI SECURITY	180
18.2	HORCM FILES.....	180
18.2.1	HORCM_MON	180
18.2.2	HORCM_CMD	180
18.2.3	HORCM_DEV	181
18.2.4	HORCM_INST.....	181
18.3	CCI CONFIGURATION FILES	181
18.4	GETTING THE LUN INFORMATION.....	183
18.5	STARTING THE INSTANCES	183
18.6	STARTING REPLICATION	184
18.7	RESYNCHRONIZING VOLUMES	185
18.8	DELETING PAIR RELATIONSHIPS	185

19	APPENDIX F: FIREWALL PORTS	187
20	APPENDIX G: DATACENTER PHYSICAL REQUIREMENTS	189
20.1	USP/USPV POWER (3-PHASE USA)	189
20.2	USP/USPV POWER (3- PHASE EU).....	189
20.3	USPVM POWER.....	190
20.4	USP/USPV DIMENSIONS AND SPACING	191
20.5	USPVM DIMENSIONS AND SPACING.....	191
20.6	WEIGHT	192
20.7	HVAC.....	193
20.8	NETWORK CONNECTIVITY	193
20.9	PHONE LINE CONNECTIVITY	193
20.10	FIBRE CONNECTIVITY	193
20.11	MANAGEMENT SERVER 1	193
20.12	MANAGEMENT SERVER 2 (CLI)	193
21	APPENDIX H: DEVICE MANAGER SAMPLE REPORTS.....	195
21.1	REPORT: PHYSICAL CONFIGURATION OF STORAGE SUBSYSTEM (HTML)	195
21.2	REPORT: STORAGE UTILIZATION BY HOST (HTML)	196
21.3	REPORT: STORAGE UTILIZATION BY LOGICAL GROUP (HTML)	197
21.4	REPORT: USERS AND RESOURCE GROUPS (HTML).....	198
22	APPENDIX I: HDS DATA MIGRATION TOOLS.....	199
22.1.1	<i>Methodology</i>	199
22.2	VOLUME MIGRATION	199
22.3	CROSS-SYSTEM COPY	200
22.4	RECOMMENDATIONS	200
23	APPENDIX J: CITI MCAFEE VIRUS SCAN AND EPO AGENT.....	201
23.1	INSTALLATION OF ANTI-VIRUS SOFTWARE	201
23.1.1	<i>Prior To Beginning Installation</i>	201
23.1.2	<i>Installation Steps</i>	201
23.2	GLOBAL EPO POLICY EXCLUSIONS	205
23.3	UNINSTALLING ANTI-VIRUS SOFTWARE	206
23.4	REPORTING AND TREE STRUCTURE.....	207
24	APPENDIX K: SETUP OF HITRACK SERVICE FOR HTTPS	208
24.1	SETUP OF HI-TRACK SOFTWARE	208
25	APPENDIX L: HDS TIERED STORAGE MANAGER	212
25.1	WHAT IS TSM?	212
25.2	DOCUMENTATION	213
25.3	OVERVIEW OF MIGRATION STEPS	214
25.3.1	<i>Steps in the actual migration process</i>	215
25.4	TSM SERVER INSTALLATION AND CONFIGURATION	217
25.4.1	<i>Licensing</i>	217
	<i>Note that once licenses have been applied, it is necessary to first refresh Device Manager and then refresh TSM before TSM will see them as activated.</i>	217
25.4.2	<i>TSM Server Requirements</i>	217
25.4.3	<i>Installing the TSM Server</i>	217
25.4.4	<i>Connecting with the web client</i>	217
25.4.5	<i>User Management setup</i>	218
25.4.6	<i>Configuration of server.properties files</i>	218
25.4.7	<i>Setting up integration with HiCommand Tuning Manager</i>	219
25.4.8	<i>Initial refreshes and creating Storage Domains.</i>	220
25.4.9	<i>Building Storage Tiers</i>	220

25.4.10	SANscreen integration	221
25.4.11	eTrust Audit setup	222
25.4.12	Audit Log Configuration - event log	222
25.5	CITIGROUP TSM BEST PRACTICE AND OPERATIONAL NOTES	223
25.5.1	I/O loading during migrations	223
25.5.2	Refreshing the TSM database	223
25.5.3	Delayed execution of a task	224
25.5.4	Reporting on storage usage by Tier	224
25.5.5	Disk shredding	224
25.5.6	Backup/recovery of TSM database	225
25.5.7	Task Completion Notification	225
25.5.8	Dealing with failures.....	226
25.5.9	Usage of migration policies.....	226
25.6	CLIENT CLI INSTALLATION, CONFIGURATION AND OPERATIONS.....	227
25.6.1	Introduction	227
25.6.2	Installing the CLI Client	227
25.6.3	Configuration of properties file	227
25.6.4	Environment Variables	228
25.6.5	PowerBroker configuration	229
25.6.6	CLI Operations	230
25.6.7	CLI migration example	232
25.6.8	CLI shredding tasks	232
25.6.9	CLI shredding example.....	233
25.7	USING HTSM WITH REMOTELY-REPLICATED STORAGE	234
25.7.1	General	234
25.7.2	TrueCopy Sync (TC) specific comments	235
25.7.3	Universal Replicator (UR) specific comments.....	235
25.8	USING HTSM WITH DYNAMICALLY PROVISIONED STORAGE.....	237
25.8.1	Introduction	237
25.8.2	Migrations of DP volumes to DP pools	238
25.8.3	Migration of DP volumes to DP volumes	243
25.8.4	Migration of a DP volume to a non-DP volume	243
25.8.5	Migration of non-DP volume to a DP pool.....	243
25.8.6	Migration of a non-DP volume to a DP volume	243
26	APPENDIX M: SNMP MIB DEFINITIONS	245
27	APPENDIX N: HITACHI DYNAMIC - THIN PROVISIONING (HDP).....	246
27.1	WHAT IS HDP?	246
27.2	MIGRATION TO HDP	246
27.2.1	Block based Migrations	246
27.2.2	Host Based Migrations	246
27.3	HDP CONFIGURATION VIA STORAGE NAVIGATOR.....	246
27.3.1	Create a Pool.....	247
27.3.2	Delete a Pool	248
27.3.3	Expanding a Pool.....	251
27.3.4	Create DP-VOL (V-VOL)	252
27.3.5	Delete DP-VOL(V-VOL).....	257
27.3.6	Change DP-VOL(V-VOL) Threshold.....	258
27.3.7	Resetting Thresholds for HDP	259
27.4	HDP CONFIGURATION VIA DEVICE MANAGER (HDVM)	259
27.4.1	Create Pool	259
27.4.2	Delete Pool	262
27.4.3	Expanding a Pool.....	265
27.4.4	Create DP-VOL	266
27.4.5	Delete DP-VOL.....	271

27.5	SNMP BASED ALERTS FOR HDP	273
27.5.1	Setting up SNMP for Threshold alerts	273
27.5.2	Setting up HDvM for Threshold alerts.....	273
27.6	HDP REPORTING.....	275
27.6.1	Reporting on HDP Volumes via Device Manager	275
27.6.2	Reporting on HDP Volumes via Tuning Manager (HTnM)	276
27.7	CITI HDP BEST PRACTICES	277
27.7.1	Virtual to Physical Capacity (Thin Provisioning Rules).....	277
27.7.2	Pool Composition	277
27.7.3	Pool ID Numbering.....	278
27.7.4	Pool Creation.....	278
27.7.5	Pool Growth.....	278
27.7.6	Migrations.....	279
27.7.7	DP-VOL Sizing	279
27.7.8	Thresholds.....	279
27.7.9	HDP Alerts.....	279
27.7.10	Pool Full Condition	282
27.7.11	Application Migrations	283
27.7.12	Application and Host Operating System Considerations.....	283
27.8	ZERO DATA REMOVAL TOOL	285
27.8.1	Caveats and Restrictions.....	285
27.8.2	Using the Tool.....	285
27.9	SYSTEM OPTIONS FOR HDP (ARRAY SETTINGS).....	287
28	APPENDIX O: PERFORMANCE MONITOR DATA EXTRACT TOOL	288
28.1	WHAT IS PERFORMANCE MONITOR DATA EXTRACT TOOL?	288
28.2	PREPARING TO USE THE EXTRACT TOOL.....	288
28.2.1	Preparing the Array.....	288
28.2.2	Choosing a Server.....	288
28.2.3	Temporary User	288
28.2.4	Software	288
28.3	USING THE EXTRACT TOOL.....	288
28.3.1	Creating the Command File.....	288
28.3.2	Optional Replication Commands.....	289
28.3.3	Optional Data Ranges (Longrange and Shorrange Settings)	289
28.3.4	Creating a Batch File.....	289
28.3.5	Running the Export Tool.....	290
28.3.6	Timing the Export Tool	290
28.3.7	Error Codes from the Export Tool	290
28.3.8	Optional Data Ranges (Longrange and Shorrange Settings)	291
28.4	TRANSMITTING THE DATA TO HDS.....	291
29	APPENDIX P: SERVER PRIORITY MANAGER (QOS)	292
29.1	CONSIDERATIONS AND RISKS	292
29.2	SETUP AND USE OF SPM	292
30	APPENDIX Q: TUNING MANAGER (HTNM).....	294
30.1	INSTALLATION OF SOFTWARE	294
30.2	TUNING MANAGER CONFIGURATION	294
30.2.1	Separate Tuning Manager and Device Manager Servers.....	294
30.2.2	Creating Array Agents (Called RAID Agents).....	296
30.2.3	Agent Processes	304
30.2.4	Verify Agents.....	304
30.2.5	Polling Times	305
30.2.6	Manual Polling	307
30.3	PERFORMANCE REPORTING METHODOLOGY.....	307

30.3.1	Basic Terminology	307
30.3.2	Review of a Typical Application Transaction	308
30.3.3	Response Time Factors	308
30.3.4	General Performance Monitoring	308
30.3.5	Performance Analysis Methodology	308
30.4	CITI STANDARD REPORTS	309
30.4.1	Setup of Citi Reports	309
30.4.2	Subsystem Flow Analysis Reports.....	310
30.4.3	Subsystem Trend Reports.....	311
30.4.4	Host (Server) Level Reports.....	311
30.5	CITI STANDARD ALARMS	312
30.5.1	Setup of Citi Alarms.....	312
30.5.2	Email Alerting Setup.....	312
30.5.3	RAID Alarms.....	315
30.6	PROCESSOR TO PORT MAPPINGS	317
30.7	THIN POOL (HDP) TREND REPORTING.....	318
30.8	VIRTUAL VOLUME TREND REPORTING (DETAIL OF POOL TREND REPORT)	322
30.9	GENERAL ARRAY REPORTING.....	324
31	APPENDIX R: SUMMARY OF SETUP FILES FOR HDVM (AND HITACHI STORAGE COMMAND SUITE).....	329
31.1	DISPATCHER (POLLING PERIOD).....	329
31.2	HTNM (EMAIL ALERTS).....	329
31.3	HTSM.....	329
31.4	SSL FOR HDVM.....	329
31.4.1	Create SSL keys for secure CLI access.....	329
31.4.2	Generate private key.....	331
31.4.3	Create a Certificate Signing Request.....	332
31.4.4	Creating a Self-Signed Certificate	332
31.4.5	Copy new key files to prod SSL directory	332
31.4.6	Modify conf file to use HTTPS and find the key files	333
31.4.7	Change SSL URL to HTTPS	333
31.4.8	Setup GUI/CLI on other Windows Hosts	333
32	APPENDIX Z: VERSION NOTES	334

7 Introduction

In order to provide Citi with more choices for Enterprise Class storage, Global Distributed Storage Engineering is certifying the latest storage array product from Hitachi Data Systems. The HDS TagmaStore Universal Store Platform (USP / USPV) is a recent generation of enterprise-class storage that will be certified for use within Citi Technology Infrastructure (CTI).

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

8 Release Status

There is an approved standard build for this product.

9 Prerequisites for a new HDS Environment

The following items are required in order to implement a TagmaStore USP/USPV according to GDSE standards:

- A Storage Management Workstation (Windows 2k3/2k). This server will run all of the GUI interfaces for Storage Navigator, Device Manager, and Tiered Storage Manager. This server will also run the Device Manager and the Tiered Storage Manager servers. Since DM/TSM has a built-in web server, then care should be taken to make sure there are no conflicts with other applications.
- If command line functionality is desired, then an additional management SOE system (Solaris only) will be required to run Device Manager CLI, and optionally the TSM CLI and ShadowImage CCI. PowerBroker infrastructure is also a pre-requisite.
- A Citi Extranet Platform (CEP) or firewall that encompasses and selectively isolates the HDS SVP devices (IP connections to array).
- A VPN Secure Vendor Dial-up environment for HDS dial-in (i.e. Safeword Cards and CitiAnywhere VPN SW, Axeda Gateway Service).
- An eTrust Audit logging infrastructure environment.
- An Onaro SANSscreen SAN auditing environment is required (via Hitachi Device Manager).
- Appropriate physical floor space, cooling, power, and redundant storage SAN fabric connectivity.

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

- Resource Manager (a basic software package)
 - o Storage Navigator
 - o LUN Manager
 - o LUSE/LVI (Custom volume sizing tool)
 - o SANTinel (LUN port sharing)
 - o DiskShredding (Volume based security overwrite)
- ShadowImage (Internal array replication – based on replication raw size)
- Performance Monitor
- Hitachi Dynamic Provisioning*

**Hitachi Dynamic Provisioning is not yet certified/supported for use at Citi.*

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

- HiCommand Device Manager (Reporting, CLI, and SANSscreen)
- HiCommand Tiered Storage Manager

10 Installation

10.1 Required Installation Steps

Major Steps:

- Physical Installation of Array into Enterprise Class Data Facility
- Network Installation of SVPs
- Installation of Management Server 1 (Windows)
- (Optional) Installation of Management Server 2 for CLI (Solaris)
- Redundant management server(s)
- Firewalled Network Segment (CEP) Isolation of SVPs
- Configuration of Array via Storage Navigator (via Mgmt Server)
- Configuration of HTTPS based dial-home for alerting to HDS
- VPN access and configuration (Vendor Remote Access)
- Setup of Device Manager (necessary for ECC, CLI and SANScreen)
- eTrust and SANScreen setup and configuration (Auditing)
- Allocation of storage to storage hosts

10.2 Software Packages

Software	Platform	Package Name
HDS Command Control Interface CLI (CCI) version 01-19-03	Solaris 10	<i>HORCM (directory HORCM_01-19-03)</i>
HDS Command Control Interface CLI (CCI) version 01-22-03	Solaris 10	<i>HORCM (directory HORCM_01-22-03)</i>
HDS Device Manager HDvM 6.0 Server	Windows 2k3/2k	<i>hdvm6000</i>
HDS Device Manager HDvM 6.0.0.5 Patch Only	Windows 2k3/2k	<i>hdvm6500</i>
HDS Device Manager HDvM 6.0 CLI	Solaris 10	<i>HDvMCLI (directory HDvMCLI_60)</i>
HDS Device Manager HDvM 6.0 CLI	Windows 2k3	<i>hdcl6010</i>
HDS Device Manager HDvM 6.0 Agent (for HRpM)	Solaris 10	<i>HDvM_AGT_60.tar.gz (directory HDvM_AGT_60)</i>
HDS Device Manager HDvM 6.2 Server	Windows 2k3	<i>hdvm6210</i>
HDS Device Manager HDvM 6.2 CLI	Solaris 10	<Not Finalized>
HDS Device Manager HDvM 6.2 CLI	Windows 2k3	<Not Finalized>
HDS Device Manager HDvM 6.2 CLI	RH Linux 5.4	<Not Finalized>
HDS Device Manager HDvM 6.2 Agent (for HRpM)	Solaris 10	<Not Finalized>
HDS Replication Manager HRpM 6.0 Server	Windows 2k3	<i>hrpm6000</i>
HDS Replication Manager HRpM 6.2 Server	Windows 2k3	<Not Finalized>
HDS Tiered Storage Manager HTSM 6.0 Server	Windows 2k3	<i>htsm6000</i>
HDS Tiered Storage Manager HTSM 6.0 CLI	Solaris 10	<i>PRhtsmcli_6.0.tar</i>
HDS Tiered Storage Manager HTSM 6.2 Server	Windows 2k3	<i>hdvm6210 (same as HDvM 6.2)</i>
HDS Tiered Storage Manager HTSM 6.2 CLI	Solaris 10	<Not Finalized>
HDS Tiered Storage Manager HTSM 6.2 CLI	RH Linux 5.4	<Not Finalized>
HDS Tuning Manager HTnM 6.0 Server	Windows 2k3	<i>htnm6000</i>
HDS Tuning Manager HTnM 6.0 Agent for RAID	Windows 2k3	<i>htar6000</i>
HDS Tuning Manager HTnM 6.0 Agent for RAID and RAID Manager Library (RMLIB)	Solaris 10	<i>HTnM_AFR_60.A.tar.gz (directory HTnM_AFR_60.A)</i>
HDS Tuning Manager HTnM 6.2 Server	Windows 2k3	<i>htnm6200</i>
HDS Tuning Manager HTnM 6.2 Agent for RAID	Windows 2k3	<Not Finalized>
HDS Tuning Manager HTnM 6.2 Agent for RAID and RAID	Solaris 10	<Not Finalized>

Manager Library (RMLIB)		
HDS Tuning Manager HTnM 6.2 Agent for RAID	RH Linux 5.4	<Not Finalized>
HDS Tuning Manager HTnM 6.2 Citi Reports and Alarms	Windows 2k3	<Not Finalized>
Sun Java JRE 1.5.0_11 (with JWS)	Windows 2k3/2k	jrjw1510
Sun Java JRE 1.6.0_16	Solaris 10	jre16016 (directory jre160161.6.0_16)
Sun Java JRE 1.6.0_18 (with JWS)	Windows 2k3	<Not Finalized>
Performance Monitor Extract Tool USP 50-09-71	Windows 2K3	hpet5090
Performance Monitor Extract Tool USPV 60-04-04	Windows 2K3	hpet6040
Performance Monitor Extract Tool USPV 60-05-15	Windows 2k3	hpet1510
Performance Monitor Extract Tool USPV 60-06-21	Windows 2k3	hpet2110

10.3 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.

HDvMCLI

If you are going to upgrade your HDvM CLI package, then you will want to preserve a copy the HiCommandCLI.properties file. Once you have made a copy, you should remove the old package and add the new package. Using JRE 1.6 or later will enable the password hiding.

HORCM

You will need to preserve any HORCM configuration files before uninstalling the old package and installing the new package.

HTnM_AFR_60

This package is not a package but a tar/gzip file that contains an installation script. You will want to uncompress and expand the tar file locally and run the installation script. That script will add the package to your system. See [Setup of Agent for RAID for Solaris 10](#) for detailed instructions on using this package.

HDvM_AGT_60

This package is not a package but a tar/gzip file that contains an installation script. You will want to uncompress and expand the tar file locally and run the installation script. That script will add the package to your system.

hdvm6500

This spac only copies an exe file to the local package directory "c:\apps_srv\instpkgs\hdvm6500". Run the executable and select to install or expand files (then run setup). If you are running HDvM 6.0 then the patch (6.0.0-05) will be applied. Otherwise it will not apply the patch.

hpet6040

This spac only copies the needed files and subdirectories to the package directory "c:\apps_srv\instpkgs\hpet6040". The contents can be copied elsewhere or used where they are.

hpet5090

This spac only copies the needed files and subdirectories to the package directory "c:\apps_srv\instpkgs\hpet5090". The contents can be copied elsewhere or used where they are.

jre16016

This Solaris package will also 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.

11 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 Storage Navigator](#)
- [Appendix C: HiCommand Device Manager Navigation and Usage](#)
- [Appendix D: Command Line Interface for Device Manager](#)
- [Appendix L: Tiered Storage Manager](#)

11.1 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: [GDSE Supported MultiPath Configuration Matrix](#).

Each supported host OS has its own stack information relating to the configuration and use with SAN-attached storage. Please refer to the following link for each host type:

<http://engineering.citigroup.net/gdse/products/matrices/>

See the section [Host Path Monitoring \(with DMP\)](#) for additional information about host monitoring using DMP.

For Host modes in setting up Host Groups, use the following chart:

Operating System Type	Host Mode	Option	Array
AIX	OF		USPV / USP
AIX with VCS	OF	2, 22	USPV / USP
AIX with Oracle AMS	OF		USPV / USP
AIX with Oracle RAC & AMS	OF		USPV / USP
AIX with HACMP	OF	15	USPV / USP
AIX with VIO vSCSI	OF		USPV / USP
ESX (VMWare)	21*		USPV
ESX (VMWare)	OA	19	USP
HP-UX	3	12	USPV / USP
HP-UX with VCS	3	2,12,22	USPV / USP
HP-UX with Oracle AMS	3	12	USPV / USP
Linux (RH)	0		USPV / USP
Linux (RH) with RHCS (Red Hat Linux Cluster)	0		USPV / USP
Linux (RH) with VCS	0	2, 22	USPV / USP
Linux (RH) with VDBE/AC for Oracle RAC	0	2	USPV / USP
Solaris with VM	9	2	USPV / USP
Solaris with VCS	9	2, 22	USPV / USP
Solaris with VDBE/AC for Oracle RAC	9	2	USPV / USP
Solaris with Oracle AMS	9	2	USPV / USP
Solaris with ZFS	9	2	USPV / USP
Solaris with Sun Cluster	9		USPV / USP

Windows	2C		USPV / USP
Windows with VCS	2C	2, 22	USPV / USP
Windows with VDBE/AC for Oracle RAC	2C	2	USPV / USP
Windows Cluster	2C		USPV / USP
Windows Cluster with Veritas Volume Mgr	2C		USPV / USP

*Note: VMWare is also supported with Host Mode 01, but since some volumes may be LUSE, 21 is recommended. LUSE volumes are legacy and no longer promoted.

For more information regarding how to set the correct host mode, see [Assignment LDEVS to Host Groups](#).

11.2 RAID Types

The following supported RAID Groups (also called Parity Groups) have the best performance characteristics on the HDS USP (assuming the same 15k speed):

Disk Configuration	Relative Capacity of Array	Relative Performance Level
<i>RAID 10 (2+2)</i>	Lowest	Highest
RAID 10 73GB (4+4)	...	...
RAID 10 146GB (4+4)	...	...
RAID 10 300GB (4+4)	...	...
<i>RAID 5 (3+1)</i>	...	...
RAID 5 73GB (7+1)	...	...
RAID 5 146GB (7+1)	...	...
RAID 5 300GB (7+1)	...	...
RAID 6 1000GB (6+2)	Highest	Lowest

The other RAID configurations did not have acceptable performance and/or cost characteristics (RAID 5 3+1, and RAID 10 2+2) in comparison to the supported configurations. The only supported RAID 5 configuration is 7+1. The only supported RAID 10 configuration is 4+4.

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 configuration.

Special Exception for USPVM:

Due to a very rare but possible vulnerability, there can be data loss with a R5 7+1 configuration if both back end loops fail on a given BED. No other raid types are affected by this vulnerability. This vulnerability has a low probability but can result in data loss (2-disk failure on raid group). Due to this, GDSE is promoting the use of RAID 5 in a 3+1 configuration and discouraging the use of a 7+1 configuration. However, this will reduce the cost effectiveness of the array.

On HDS, RAID10 has been implemented as striping before mirroring rather than mirroring before striping. When we refer to RAID10 on HDS, we actually are referring to RAID 0+1.

11.3 Drive Types

The following drive types are supported on the USP:

Disk Type	Back-End Speed
73GB (15k)	2GB
146GB (15k)	2GB
300 GB (10k)	2GB

The following drive types are supported on the USPV:

Disk Type	Back-End Speed
73GB (15k) FC	4Gb
146GB (15k) FC	4Gb
300 GB (10k)* FC	2Gb*
300 GB (15k) FC	4Gb
450 GB (15k) FC	4Gb
1000 GB (7k) SATA	4Gb

***Warning:** Due to the slower back-end speed, any 300GB 10k drives need to be on dedicated BED (DKA) Pairs to avoid lowering the speed of any 4Gb drives to 2Gb. This is a critical design issue and must be considered when using the slower 300GB 10k drives. The use of 15k drives is preferred for performance reasons. This drive is no longer recommended and is supported only in legacy environments.

11.3.1 Using Slower Drives on a USPV

This configuration is supported only by exception and is not a desirable configuration.

The use of both slower (2Gb) and faster (4Gb) drives in the same USPV array requires the use of dedicated BED (DKA) pairs. This will limit the number of drives of each type and the location as it compartmentalizes the array. This does not apply to the USP as there is only a single drive speed (2Gb). The following diagram shows recommended 2Gb placement in the Citi USPV models. This shows the recommended placement of slower interface speed (2Gb) drives. The red box indicates placement of the drives on a USPV array.



Note: The above diagram does not account for spare drives. An 8-disk parity group must be used for spares that is controlled by the same BED (DKA) pair, thus reducing the total number of useable 2Gb drives by 8. These disks also follow the same rules for similar parity groups (8-disk parity groups in R5 7+1 configuration).

11.4 Array Types

There are three models of the USP used at Citi:

Component	Small (R0)	Medium (R0,R1)	Large (R0,R1,L1)
Number of Back End Directors (BEDs)	2	4	8
Total fibre channel loops per DA pair	16	16,32	16,16,16,32
Minimum number of disk drives	16	128	392
Maximum number of disk drives	120	376	632
Maximum raw capacity			165.0 TB
Maximum usable capacity*	27.4 TB	86.0 TB	144.5 TB

Min/Max Cache Memory (non-replicated) + Shared Memory	24/40 GB +16GB	64/128 GB +16GB	96/200 GB +16GB
Min/Max Cache Memory (replicated) + Shared Memory	32/64 GB +20GB	64/128 GB +20GB	96/200 GB +20GB
Recommended connectivity (FA port pairs)	8	16	24

*Based on 146GB R5

There are five models of the USPV used at Citi:

Component	USPVM (R0)	Small (R0)	Medium (R0,R1)	Large (R0,R1,L1)	Extra Large (R1,R2,L1,L2)
Number of Back End Directors (BEDs)	2	2	4	8	8
Total fibre channel loops per DA pair	8	16	16,32	16,16,16,32	32,32,32,32 No R0 or R2L
Minimum number of disk drives	16	16	128	392	640
Maximum number of disk drives	116	120	376	632	1016
Maximum raw capacity*				165.0 TB	265.1 TB
Maximum usable capacity	27.4 TB	27.4 TB	86.0 TB	144.5 TB	232.0 TB
Min/Max Cache Memory (non-replicated) + Shared Memory	24/40 GB +16GB	24/40 GB +16GB	64/128 GB +16GB	96/200 GB +16GB	128/256 GB +16GB
Min/Max Cache Memory (replicated) + Shared Memory	32/64 GB +16GB	32/64 GB +20GB	64/128 GB +20GB	96/200 GB +20GB	160/512 GB +20GB
Recommended connectivity (FA port pairs)	8	8	16	24	32

*Based on 300GB R5

There may be variation of orders that do not exactly fit these models. Cache and BEDs should be adjusted according such that the following rules are never exceeded:

Maximum Drive Loop Count: 32

11.5 Array Cache

There are three models of arrays that are used by Citi on the USP and five on the USPV:

Citi USP Configuration	HDS USP Model	Number BED Pairs	Recommended Cache (including Shrd Mem)
Small	100	1	32GB
Medium	600	2	32GB
Large	1100	4	64GB

Citi USPV Configuration	HDS USPV Model	Number BED Pairs	Cache** Normal / HA*	Shared Memory Normal / HA
USPVM	USPVM	1	24-40/ 32-64GB	16 / 20GB
Small	USPV	2	24-40/ 32-64GB	16 / 20GB
Medium	USPV	Up to 4	64-120 / 96-192GB	16 / 20GB
Large	USPV	Up to 8	96-200 / 160-256GB	16 / 20GB
Extra Large	USPV	8	128-256 / 200GB+	16/ 20GB

The Small Citi model (on USP only) should not be used for replicated environments or for high-intensity I/O applications. This is due to the limited number of BEDs or Back-End Director that control all internal disk activity. This does not apply to the USPV.

**On the USPVM models, there is a range of cache for the Normal and HA variations. There is a range because the amount of cache is based on the total raw capacity of the array. For example, on the medium USPVM model, the standard array cache ranges from 64-120GB. If the array is fully populated with 146GB drives, then 64GB is appropriate. If the array is fully populated with 300GB drives, then 120GB is appropriate.

*HA is for High Availability. This array can be configured for use with HUR (Hitachi Universal Replicator) and or True Copy Synchronous. HUR and True Copy Synchronous are not yet certified / supported for use at Citi.

Thin Provisioning (HDP) is not expected to require additional shared memory. HDP is not yet certified for use at Citi.

Note: If there are very strong business requirements and special conditions that dictate a total storage capacity that is between the standard configurations, then it is advised that the larger configuration be used and the desired number of disks be removed in increments of 8. The cache, control memory, spare drives, BEDs, host fiber cards, array frames, and all other physical aspects should remain the same. Software licensing may need to be re-sized appropriately. These special configurations are discouraged.

The USPVM cannot be upgraded to a USPVM. The drives can be reused but all data will be lost.

11.6 Spare Drives

There should be at least one spare drive per BED or Back-End Director. The following charts are a summary of the number of BEDS and recommended spare drives (1 spare per bed the minimum recommended):

Citi USP Configuration	HDS USP Model	Number BED Pairs	Drives (Cabinets)	Number Spare Drives
Small	100	1	120(DKC)	2
Medium	600	2	376(DKC,R1)	2
Large	1100	4	632(DKC,R1,L1)	4

Citi USPVM Configuration	HDS USPVM Model	Number BED Pairs	Drives (Cabinets)	Number Spare Drives
USPVM	USPVM	1	112 (DKC)	2
Small	USPVM	2	120 (DKC)	2
Medium	USPVM	4	376 (DKC,R1)	4
Large	USPVM	8	632 (DKC,R1,L1)	6
Extra Large	USPVM	8	1016(DKC,R1,R2,L1,L2)*	8

*Note: The Extra Large configuration will have 128 empty drive slots. These empty slots will be in the DKC if configured initially (called R0). If the array is upgraded to this size, then R2 Lower should be empty.

If there are two drive types, then the spares should be split between the two types. If there are three drive types, then 2 additional spares should be added. There should be 1 spare per BED and at least 2 per drive type. *For example, on an array with 4 BEDs and 3 drive types (146,300, 300-10k) there should be 6 total spares (2 each). If there were only two drive types with 4 BEDs, then 4 total spares would be appropriate.*

Having more spare than required is not a problem. However each set of 8 drives takes away a potential parity group. The spare drives reside in the 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 146GB and 300GB drives may benefit from having 6x or 8x 300GB drives rather than 2 or 3 of each drive as spares.

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.

11.7 Logical Device (LDEV) Sizing and Capacity

The USP supports fixed or custom-size volume definitions. The current standard is to use OPEN-V CVS emulation for custom volume sizing of all volumes. This gives us maximum flexibility (including LUSE functionality without a performance penalty). To keep this storage consistent with all other storage platforms and for operational efficiency, the following LDEV sizes should be used on a standard configuration. This has been optimized for a 146GB R5(7+1) configuration.

Role	Volumes
Default Standard Volumes (STDs) for all data storage types	34.69GB (35,522.812 MB)
LUSE Volumes	69.38GB, 104.07GB, 138.76GB (Or any 34.69GB combination <=277.52GB from separate Parity Groups)
Custom Volumes (CVOs)	Any size from 4GB to 277.52GB created from destroyed LDEVs / freespace. These are typical of luns migrating from other storage arrays with different default lun sizes.
Command Devices (CMDs) for use with ShadowImage (BCVs)	50.625 MB
Clustering/Quorum Volumes (QUOs) for use with clusters	1000.625 MB

USPVM Configuration (116 Drives – Max 26.9 TB Usable w/ R5 7+1 300GB) USPVM Only

LDEV / Config	300GB R5(7+1)	300GB R10(4+4)	146GB R5(7+1)	146GB R10(4+4)
CMDs – 50.625MB	14 (1x14pg)	28 (2x14pg)	14 (1x14pg)	14 (2x14pg)
QUO – 1000.625MB	14 (1x14pg)	28 (2x14pg)	14 (1x14pg)	14 (2x14pg)
STDs – 35522.812 MB	756 (54x14pg)	420(30x14pg)	378 (27x14pg)	196 (14x14pg)
Freespace	5570 MB x 14	16845 MB x 28	0.0 MB x 14	24996 MB x 14

Note: The use of R5 7+1 on the USPVM array is discouraged. The figures are here for comparison purposes. For more information see [RAID Types](#).

USPVM Configuration (continued) USPVM Only

LDEV / Config	300GB R5(3+1)	146GB R5(3+1)
CMDs – 50.625MB	28 (1x28pg)	28 (1x28pg)
QUO – 1000.625MB	28 (1x28pg)	28 (1x28pg)
STDs – 35522.812 MB	644 (23x28pg)	308 (11x28pg)
Freespace	6948 MB x 28	19524 MB x 28

Standard Small Configuration (120 Drives – Max 27.4 TB Usable w/ R5 300GB) USP/USPV

LDEV / Config	300GB R5(7+1)	300GB R10(4+4)	146GB R5(7+1)	146GB R10(4+4)
<i>CMDs – 50.625MB</i>	15 (1x15pg)	30 (2x15pg)	15 (1x15pg)	30 (2x15pg)
<i>QUO – 1000.625MB</i>	15 (1x15pg)	30 (2x15pg)	15 (1x15pg)	30 (2x15pg)
<i>STDs – 35522.812 MB</i>	810 (54x15pg)	450(30x15pg)	405 (27x15pg)	210 (14x15pg)
<i>Freespace</i>	5570 MB x 14	16845 MB x 30	0.0 MB x 15	24996 MB x 30

Standard Medium Configuration (376 Drives – Max 86.0 TB Usable w/R5 300GB) USP/USPV

LDEV / Config	300GB R5(7+1)	300GB R10(4+4)	146GB R5(7+1)	146GB R10(4+4)
<i>CMDs – 50.625MB</i>	47 (1x47pg)	94 (2x47pg)	47 (1x47pg)	94 (2x47pg)
<i>QUO – 1000.625MB</i>	47 (1x47pg)	94 (2x47pg)	47 (1x47pg)	94 (2x47pg)
<i>STDs – 35522.812 MB</i>	2538(54x47pg)	1410(30x47pg)	1269 (27x47pg)	658 (14x47pg)
<i>Freespace</i>	5570 MB x 47	16845MB x 94	0.0 MB x 47	24996 MB x 94

Standard Large Configuration (632 Drives – Max 144.5 TB Usable w/ R5 300GB) USP/USPV

LDEV / Config	300GB R5(7+1)	300GB R10(4+4)	146GB R5(7+1)	146GB R10(4+4)
<i>CMDs – 50.625MB</i>	79 (1x79pg)	158 (2x79pg)	79 (1x79pg)	158 (2x79pg)
<i>QUO – 1000.625MB</i>	79 (1x79pg)	158 (2x79pg)	79 (1x79pg)	158 (2x79pg)
<i>STDs – 35522.812 MB</i>	4266(54x79pg)	2370 (30x79pg)	2133 (27x79pg)	1106 (14x79pg)
<i>Freespace</i>	5570 MB x 79	16845MB x 158	0.0 MB x 79	24996 MB x 158

Extra Large Configuration (1016 Drives – Max 232 TB Usable w/ R5 300GB) USP/USPV Only

LDEV / Config	300GB R5(7+1)	300GB R10(4+4)	146GB R5(7+1)	146GB R10(4+4)
<i>CMDs – 50.625MB</i>	127 (1x127pg)	254(2x127pg)	127 (1x127pg)	254(2x127pg)
<i>QUO – 1000.625MB</i>	127 (1x127pg)	254 (2x127pg)	127 (1x127pg)	254(2x127pg)
<i>STDs – 35522.812 MB</i>	6858(54x127pg)	3810(30x127pg)	3429(27x127pg)	1778 (14x127pg)
<i>Freespace</i>	5570 MB x 127	16845MB x 254	0.0 MB x 127	24996 MB x 254

CMDs and QUOs should be at the beginning/top of the Parity Group (created first) to allow for the destruction and reuse of LDEVs as Custom Volumes. STDs should be allocated from the top of the list to further maximize this potential reuse. Freespace can be combined to create larger volumes that span parity groups (LUSE) or can be used to create smaller custom volumes.

11.8 Logical Device (LDEV) Numbering

In order to have consistent configurations worldwide, use the following CU - Logical Device (CU:LDEV) numbering scheme.

Format: CU:LDEV (USPV has an additional octet that precedes the address i.e. 00:CU:LDEV). For example, 01:00 is the first address for CU01. This would be a RAID 5 (7+1) volume (LDEV). On the USPV, this would be preceded by "00" (00:01:00). This preceding octet is reserved for future use.

CU Prefix (Hex)	Address Range (Hex)	Type	RAID or DeviceType	Drive Type
CU 00 – 1F	00 – FF	Physical	RAID 5 (7+1)	146,300,450GB FC
CU 10 – 1F	00 – FF	Physical	RAID 5 (7+1)	146,300,450GB FC
CU 20 – 2F	00 – FF	Physical	RAID 6 (6+2)	1000GB SATA
CU 30 – 39	00 – FF	Physical	RAID 6 (6+2)	1000GB SATA
CU 3A – 3B	00 – FF	<i>Future Use</i>	<i>Future Use</i>	<i>Future Use</i>
CU 3C	00 – FF	Physical	<i>QUO Devices</i>	146,300,450GB FC
CU 3D	00 – FF	Physical	<i>QUO Devices</i>	146,300,450GB FC
CU 3E	00 – FF	Physical	<i>CMD Devices</i>	Any
CU 3F	00 – FF	Physical	<i>CMD Devices</i>	Any
CU 40 – 4F	00 – FF	Physical	<i>POOL Devices</i>	146,300,450GB FC
CU 50 – 5F	00 – FF	Physical	<i>POOL Devices</i>	1000GB SATA
CU 60 – 6F	00 – FF	Virtual	<i>DP-VOL Devices</i>	<i>146,300,450GB FC Pools</i>
CU 70 – 7F	00 – FF	Virtual	<i>DP-VOL Devices</i>	<i>146,300,450GB FC Pools</i>
CU 80 – 8F	00 – FF	Virtual	<i>DP-VOL Devices</i>	<i>146,300,450GB FC Pools</i>
CU 90 – 9F	00 – FF	Virtual	<i>DP-VOL Devices</i>	<i>146,300,450GB FC Pools</i>
CU A0 – AF	00 – FF	Virtual	<i>DP-VOL Devices</i>	<i>1000GB SATA Pools</i>
CU B0 – BF	00 – FF	Virtual	<i>DP-VOL Devices</i>	<i>1000GB SATA Pools</i>
CU C0 – CF	00 – FF	<i>Future Use</i>	<i>Future Use</i>	<i>Future Use</i>
CU D0 – DF	00 – FF	<i>Future Use</i>	<i>Future Use</i>	<i>Future Use</i>
CU E0 – EF	00 – FF	<i>Future Use</i>	<i>Future Use</i>	<i>Future Use</i>
CUF0 – FE	00 – FF	<i>Future Use</i>	<i>Future Use</i>	<i>Future Use</i>

Notes:

- RAID 10 (2+2) is not recommended due to expense/performance.
- RAID 6 (6+2) is recommended for use with FC and SATA drives.
- QUO Devices are for clustering.
- CMD Devices are used for communication via fiber to the array. This is typically used in CCI (command line interface for ShadowImage and TrueCopy/HUR). It is also used by Tuning Manager.
- POOL Devices are dedicated RAID groups that comprise a HDP Pool.
- DP-VOL Devices are virtual (thin) devices that are from a HDP Pool.

To achieve proper load balancing for sequentially numbered LDEVs, the array should be configured initially using DISPERSED mode. This address schema (performed by HDS/SUN staff at array set-up) will balance the LDEVs across all of the RAID groups (parity groups) assigned to a CU range. The result will be balanced addresses that can be allocated to hosts sequentially without impact to performance.

When manually assigning the addresses for multiple devices, the interval function should be used. The interval function is used to disperse (or skip addresses) between the physically sequential volumes. To use this function, all of the volumes need to be selected from the upper list and then click on the first (unused) address box. The interval value is the number total number of RAID/Parity groups to spread the addresses over (usually the total number of parity groups). This is recommended for creating multiple volumes and preventing assignment of physically sequential volumes to the same host.

11.8.1 LDEVs when using TSM

It is important to understand that use of Tiered Storage Manager will disrupt the LDEV layout. This is due to the LDEV address swap that occurs when data is moved between 2 devices. The new device takes on the LDEV address of the old device and the old device

takes on the LDEV address of the new device. Therefore, if TSM is being used, the administrator will need to pay close attention to the Parity Group column when allocating LUNs to ensure that the same host is not allocated LUNs from the same physical parity group. When TSM is not in use, sequential addresses will be from different parity groups (in dispersed mode). This may not be true in device addresses have been swapped by TSM.

11.9 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:

11.9.1 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.

11.9.2 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.

11.9.3 Expansion with LUSE Volumes

As described in the following section ([LUSE Volumes](#)), LUSE volumes can be used to add space to an existing LUN. The original LUN will need to be the primary LUN in the LUSE volume. The LUSE volume will be addressed by the original LUN once the expansion is complete. Note: this will result in a concatenated (not stripped) volume. **It is also recommended that a backup of the data be made prior to creating the LUSE volume, no matter what OS is being used.** This is the least host-intensive method, though systems that are not AIX or Windows based will lose the data, defeating this advantage.

11.10 LUSE Volumes

Use of LUSE volumes is not recommended (see below). 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. A LUSE can be made from LDEVs or other LUSE volumes. Only the original, first LDEV can have any paths associated with it.

Warning: Except for systems running Windows 2k3/2k or AIX, creating and releasing LUSE volumes is a destructive process. **Be sure to back up your data before preceding no matter what OS is being used.**

Because the volumes are concatenated sets and the additional logic lookups, they will generally not perform as well as a fully striped volume if the LUSE components (LDEVs) came from the same parity group (RAID group). Therefore, to maximize the performance of the LUSE volumes, they should be made with LDEVs that are from different parity groups. *This should be easy to configure by selecting sequential LDEV addresses, as sequential LDEVs should be on different parity groups (if TSM is not in use).*

Warning: Block-level migration (i.e. TSM or UVM) is not possible from LUSE volumes without significant professional services. Therefore the use of LUSE volumes is no longer recommended due to performance and migration limitations. Host-based copies typically require significant application downtime and must be considered before using LUSE volumes. LUSE volumes are now supported on legacy environments only.

11.11 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 will be longer on a more active array.

USPV Drive Type	Rebuild Time for RAID 1 (2+2)*	Rebuild Time for RAID5 (3+1)**	Rebuild Time for RAID5 (7+1)	Rebuild Time for RAID6 (6+2)***
146GB FC(15k)	0h 55m	1h 35m	2h 20m	2h 15m
300GB FC(10k)	2h 10m	3h 05m	4h 20m	4h 20m
300GB FC(15k)	1h 45m	2h 30m	3h 30m	3h 30m
450GB FC(15k)	1h 40m	1hr 40m	4h 00m	5h 50m
600GB FC(15k)	1h 10m	5h 00m	7h 00m	7h 00m
1000GB SATA(7.2k)	12h 00m	11h 40m	14h 50m	14h 50m
2000GB SATA(7.2k)	24h 15m	24h 30m	31h 10m	31h 10m

USP Drive Type	Rebuild Time for RAID 1 (2+2)*	Rebuild Time for RAID5 (7+1)	Rebuild Time for RAID6 (6+2)
73GB (15k)	0h 30m	1h 05m	1h 15m
146GB (15k)	0h 55m	2h 20m	2h 15m
300GB (10k)	2h 10m	4h 20m	4h 20m

Note: *RAID 1 (2+2) is not a supported RAID type at Citi. However, it does give a good estimate for the RAID 1 (4+4) that is supported.

** RAID 5 (3+1) is the recommended RAID for all drive types on USPVM only

*** RAID 6 (6+2) is the recommended RAID for all drive types on USPVs

600GB FC and 2000GB SATA are not yet certified for use at Citi.

11.12 LDEV Create/Destroy Timing

Creation of LDEVs

The creating of LDEVs varies depending on the size of the LDEV and if there is disk contention on the Parity Group. If many LDEVs are being created within the same Parity Group, then the timing will not be as bad as creating the volumes sequentially. The following are estimated creation times on an inactive USP and USPV array:

Size of LDEV	Est. Time for Create on USP	Est. Time for Create on USPV
1 GB	0h 5m	0h 3m
10 GB	0h 5m	0h 3m

50 GB	0h 13m	0h 5m
100 GB	0h 15m	0h 7m
200 GB	0h 24m	0h 13m

These times were based on 146 GB 15K drives in a RAID 5(7+1) parity group. It is noted that the first LDEV in a parity group may take significantly longer (+45 minutes).

Creation of LUSE

The timing of the creation of one or more LUSE volumes is not size dependent and is fixed at approximately 5 minutes (for an inactive array).

Destruction of LDEVs (Volume to Space) and LUSE Volumes

Unless LDEVs are destroyed via a Make Volume operation, the timing is approximately 5 minutes for an inactive array to complete the operation.

11.13 Host Port Balancing and Redundancy


On the USP and 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 (very 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. The hosts cards for the USP and USPV are different.


11.13.1 USP Host Cards

The primary host port card slots are CHA-1E (Cluster 1) and CHA-2Q (Cluster 2). Additional cards are added to slots CHA-1F (Cluster 1) and CHA-2R (Cluster 2). Cards and ports can be viewed by selecting the *Port* tab of the *LUN Manager* screen.



There are three types of host port cards, 8-port, 16-port, and 32-port cards. The 8-port cards have 4 processors (called MPs) per card. Since the cards are purchased in pairs to be split across the cluster barrier, HDS describes the pair of cards as a 16 Port Feature. The 16-port cards have 8 MPs per card and are added as a 32 Port Feature. The 32-port cards have 16 MPs per card and are added as a 64 Port Feature.

Citi uses 8-port cards. Two-HBA hosts (one HBA per redundant fabric) should have one HBA connected to each member of a card pair. For example, if a host using Fabric A HBA is connected to port 1A, then the Fabric B HBA should be connected to port 2A. The following diagram displays the location of each port on the first 4 host cards.



Having four HBAs, one needs to consider the processors on the host port cards. Each processor controls two ports (every other port on each card is controlled by the same MP). For example, the same MP controls 1A & 5A. Therefore, a quad host HBA configuration

would have fabric A connections to ports 1A and 3A to use different MPs. The fabric B connections would be 2A and 4A to span the different MPs.

Maximum Host Port Performance Configuration (USP)

To maximize performance, a single port per MP configuration should be used. If 32 ports are purchased (4x 8-port cards) and only 16 connections to the SAN (8 per fabric) are made, then it would be prudent to follow this method. Under these conditions, the following ports should be configured (based on 8-port cards):

	USP Fabric A	USP Fabric B
1	1A (Slot CHA-1E, Port 1) WWN: 50060e8004xxxx00	2A (Slot CHA-2Q, Port 1) WWN: 50060e8004xxxx10
2	3A (Slot CHA-1E, Port 2) WWN: 50060e8004xxxx20	4A (Slot CHA-2Q, Port 2) WWN: 50060e8004xxxx30
3	1C (Slot CHA-1E, Port 5) WWN: 50060e8004xxxx02	2C (Slot CHA-2Q, Port 5) WWN: 50060e8004xxxx12
4	3C (Slot CHA-1E, Port 6) WWN: 50060e8004xxxx22	4C (Slot CHA-2Q, Port 6) WWN: 50060e8004xxxx32
5	1E (Slot CHA-1F, Port 1) WWN: 50060e8004xxxx04	2E (Slot CHA-2R, Port 1) WWN: 50060e8004xxxx14
6	3E (Slot CHA-1F, Port 2) WWN: 50060e8004xxxx24	4E (Slot CHA-2R, Port 3) WWN: 50060e8004xxxx34
7	1G (Slot CHA-1F, Port 5) WWN: 50060e8004xxxx06	2G (Slot CHA-2R, Port 5) WWN: 50060e8004xxxx16
8	3G (Slot CHA-1F, Port 6) WWN: 50060e8004xxxx26	4G (Slot CHA-2R, Port 6) WWN: 50060e8004xxxx36

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.

The remaining ports (i.e. 5A,7A,6A,8A,...) can be used for external storage (migrations, etc.). The remaining ports can be used for external storage (migrations, etc.).

11.13.2 USPV Host Cards

The USPV uses half-height cards. In each slot, there are 2 subslots (upper and lower). Cards in the upper and lower subslots can be replaced without affecting the other card in the same slot. There are 2 types of fiber channel based cards that are used, 4-port cards and 8-port cards. The standard card for Citi is the 8-port card. Since the cards are purchased in pairs to be split across the cluster barrier, HDS describes the pair of cards as a 16 Port Feature.


The primary host port card slots are CHA-1EU (Cluster 1, Slot E, Upper) and CHA-2QU (Cluster 2, Slot Q, Upper). Additional cards are added to slots CHA-1FU (CL1, F-Upper) and CHA-2RU (CL 2, R-Upper). Then cards are added to the lower slots CHA-1EL (CL1,E-Lower) and CHA-2QL (CL2,Q-Lower).

Card Slot Order (Order of adding host cards)

CHA-1EU, CHA-2QU
CHA-1FU, CHA-2RU
CHA-1EL, CHA-2QL
CHA-1FL, CHA-2RL
CHA-1GU, CHA-2TU
CHA-1HU, CHA-2UU


Note: On the USPVM, the slot numbers are slightly different. The front-end card slots in order of use for Fabric A would be on Cluster 1 Slots 1G, 1H, and 1B. The front-end card slots in order of use for Fabric A would be on Cluster 2 Slots 2L, 2K, and 2E.

Cards and ports can be viewed by selecting the *Port* tab of the *LUN Manager* screen.



As described above, Citi uses 8-port cards. Two-HBA hosts (one HBA per redundant fabric) should have one HBA connected to each member of a card pair. For example, if a host using Fabric A HBA is connected to port 1A, then the Fabric B HBA should be connected to port 2A.

The remaining ports (i.e. 5A,7A,6A,8A,...) can be used for external storage (migrations, etc.). The following diagram displays the location of each port on the first 4 host cards.



Having four HBAs, one needs to consider the processors on the host port cards. Each processor controls two ports (every other port on each card is controlled by the same MP). For example, the same MP controls 1A & 5A. Therefore, a quad host HBA configuration would have fabric A connections to ports 1A and 3A to use different MPs. The fabric B connections would be 2A and 4A to span the different MPs.

11.13.3 Maximum Host Port Performance (USPV) Configuration

To maximize performance, a single port per MP configuration should be used. If 32 ports are purchased (4x 8-port cards, or 2x "16-port options") and only 16 connections to the SAN (8 per fabric) are made, then it would be prudent to follow this method. Under these conditions, the following ports should be configured (based on 8-port cards):

	Fabric A	Fabric B
1	1A (Slot CHA-1EU, Port 1) WWN: 50060e8005xxxx00	2A (Slot CHA-2QU, Port 1) WWN: 50060e8005xxxx10
2	3A (Slot CHA-1EU, Port 2) WWN: 50060e8005xxxx20	4A (Slot CHA-2QU, Port 2) WWN: 50060e8005xxxx30
3	1B (Slot CHA-1EU, Port 5) WWN: 50060e8005xxxx01	2B (Slot CHA-2QU, Port 5) WWN: 50060e8005xxxx11
4	3B (Slot CHA-1EU, Port 6) WWN: 50060e8005xxxx21	4B (Slot CHA-2QU, Port 6) WWN: 50060e8005xxxx31
5	1E (Slot CHA-1FU, Port 1) WWN: 50060e8005xxxx04	2E (Slot CHA-2RU, Port 1) WWN: 50060e8005xxxx14
6	3E (Slot CHA-1FU, Port 2) WWN: 50060e8005xxxx24	4E (Slot CHA-2RU, Port 3) WWN: 50060e8005xxxx34

7	1F (Slot CHA-1FU, Port 5) WWN: 50060e8005xxxx05	2F (Slot CHA-2RU, Port 5) WWN: 50060e8005xxxx15
8	3F (Slot CHA-1FU, Port 6) WWN: 50060e8005xxxx25	4F (Slot CHA-2RU, Port 6) WWN: 50060e8005xxxx35

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.

Note: On the USPVM, the slot numbers are slightly different. The front-end card slots in order of use for Fabric A would be on Cluster 1 Slots 1G, 1H, and 1B. The front-end card slots in order of use for Fabric A would be on Cluster 2 Slots 2L, 2K, and 2E.

11.13.4 WWN Numbering for USP/USPV Host Ports

The WWN of each port is shown above. The “xxxx” is the four hexadecimal digit equivalent of the serial number. For example, the serial number 10192 converts to “27d0”. The following is the USP and USPV WWN pattern:

WWN Component	Value
Prefix (HDS TagmaStore)	50060e80+(04-USP or 05-USPV)
Serial Number (Convert 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

For example, a USP SN#10192, the port CL2A would have a WWN of 50060e8004**27d010** and port CL5-C would have a WWN of 50060e8004**27d042**. A USPV with the same serial number would have a WWN for the port CL2A of 50060e80**05**27d010.

Note: The CLI for Device Manager uses a different port numbering scheme than the CLX-X described above. See the CLI section for more details.

11.14 Topology

The USP should be connected to the core directors from each redundant fabric. There should be at least eight 2Gb (or 4Gb) fibre connections from each of the fabrics. Additional ports may be connected if these ports become overloaded.

A maximum of a “7-to-1” fan-in ratio is recommended on the USP.

A maximum of a “10-to-1” fan-in ratio is recommended on the USPV.

11.15 Management Software (DM vs. SN)

Storage Navigator is designed to manage one or more arrays (USP/USPV or 9980 Lightning Series). If there are other types of HDS arrays to be managed, then Hitachi Device Manager (DM) will be needed. In addition, if command line functionality is desired, then the DM environment along with the CLI and CCI components (if appropriate) will need to be installed on a system. At Citi we limit the server to be installed on Windows. The CCI and CLI are best suited for UNIX.

If an environment runs both SN and DM, then it is important to understand that DM has a different conceptual view of devices. In addition, there must be a translation of decimal and hexadecimal numbers to accommodate this heterogeneous environment.

Note: To keep the DM and SN configurations synchronized, the best way to launch SN is from DM (using the Physical View button).

There are a few functions (for example DiskShredding) that reside in SN and are not available from DM.

DM is required for most third-party toolsets to monitor, report, or manage the USP/USPV. This includes Onaro SANScreen and EMC Control Center.

The USPV array requires the use of DM version 5.7 or later. The USP array does not have a version restriction.

11.16 SAN Monitoring

Onaro SANScreen should be the primary SAN monitoring tool with end-to-end path change monitoring. SANScreen requires the use of Hitachi Device Manager.

11.17 Array Monitoring: Alerts and Remote Connections

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

11.17.1 SNMP

The USP array does support SNMP event broadcasting. SNMP alerts can be configured to send traps to a HDvM (Device Manager) server and SNMP Trap Servers (ESS). For more information see [Enterprise Alerting \(SNMP\)](#).

11.17.2 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](#).

11.17.3 Vendor Remote Access (VPN)


Remote access by HDS is governed by the Vendor Remote Access Policy. Under this policy:

1. The vendor will call a specified Citi help desk requesting dial-in access. The help desk will provide the username / temporary (Safeword) password (and the dial-in number if necessary).
2. The vendor will connect to the Citi network via a temporary VPN connection. Upon successful connection, the vendor's system will be assigned an IP address

and will be attached to the Citi network. The vendor should have IP access only to SVPs that are located within a CEP / firewall.

3. The vendor will then use Microsoft Remote Desktop client to connect and control the SVP to obtain information /correct problem.
4. The vendor will disconnect the Remote Desktop session and the VPN connection.

The following diagram is a typical configuration:



The CEP / firewall will be configured such that only defined ports will be open between the SVPs and the pool of IPs reserved for VPN. This will ensure that only the VPN connection will have remote desktop access to the SVPs. If HDS personnel need to perform remote desktop on-site, then they will have hands-on access to the internal networks on the array.

11.18 Host Path Redundancy with DMP

On a Solaris 8 host, DMP should be enabled natively. Powerpath is likely already set up and configured for any existing storage. Powerpath prior to 4.5 should not recognize any HDS LUNs. The LUNs will be ignored by PowerPath. The paths should be recognized by DMP. If the path is not enabled, then the path will need to be enabled using vxddmpadm.

Powerpath version 4.5 and later should be disabled for HDS LUNs, allowing for DMP to retain control. Refer to the appropriate host stack solution guide for details on how to perform this.

On storage hosts with Solaris 8 with Veritas Volume Manager, the Array Support Library (ASL) for the Hitachi USP array should be installed on the host system. Improper behavior of devices may occur without the ASL.

- o For Veritas Volume Manager 3.X: package HTC-USP
Additional information:
<http://seer.support.veritas.com/docs/271477.htm>
- o For Veritas Volume Manager 4.X: package VRTS-USP
Additional information:
<http://seer.support.veritas.com/docs/272025.htm>

The ASL is used to make Veritas aware of the storage array characteristics. This information is usually handled by Powerpath.

11.19 Host Path Monitoring (with DMP)

With Solaris 8 or later, DMP is used as the primary path failover and load-balancing scheme. This information does not apply to any other storage host type.

As described in [SAN Monitoring](#), Onaro SANscreen should be implemented to monitor end-to path changes.

The Veritas DMP command line interface will show an alert status when a path failure occurs. SANmon currently handles the path management monitoring process in conjunction with PowerPath. SANmon is a tool written to periodically check HBA and path status on the Solaris environment. It may need modification in order to handle DMP without PowerPath.

As an alternative to SANmon for systems that no longer use it, a script could be created to be run frequently on the local host.

Note: Until changes to the monitoring environment can be made corporate- wide, it is suggested that scripts be used to manually parse errors from /var/adm/messages and to use some of the suggested scripting.

The following are some suggested scripting to be used on the host to further monitor host path issues:

1. To get a list of the controllers under dmp control, use the following listctlr command:

```
vxddmpadm listctlr all
```

```
root@sangs6(/)#vxddmpadm listctlr all
CTLR-NAME    ENCLR-TYPE    STATE    ENCLR-NAME
=====
c1           Disk         ENABLED    Disk
c3           EMC          ENABLED    EMC2
c2           EMC          ENABLED    EMC2
c3           GENESIS      ENABLED    GENESIS0
c2           GENESIS      ENABLED    GENESIS0
```

2. To check the status on all the disk attached to that controller, run the getsubpaths command:

```
vxddmpadm getsubpaths ctrl=c#
```

```
root@sangs6(/)#vxddmpadm getsubpaths ctrl=c2
```

NAME	STATE	PATH-TYPE[M]	DMPNODENAME	ENCLR-TYPE	ENCLR-NAME	ATTRS
c2t50d4s2	ENABLED	-	c2t50d4s2	EMC	EMC2	-
c2t50d2s2	ENABLED	-	c2t50d2s2	EMC	EMC2	-
c2t40d11s2	ENABLED	-	c2t40d11s2	GENESIS	GENESIS0	-
c2t40d10s2	ENABLED	-	c2t40d10s2	GENESIS	GENESIS0	-
c2t40d5s2	ENABLED	-	c2t40d5s2	GENESIS	GENESIS0	-
c2t40d4s2	ENABLED	-	c2t40d4s2	GENESIS	GENESIS0	-
c2t40d3s2	ENABLED	-	c2t40d3s2	GENESIS	GENESIS0	-
c2t40d2s2	ENABLED	-	c2t40d2s2	GENESIS	GENESIS0	-
c2t40d1s2	ENABLED	-	c2t40d1s2	GENESIS	GENESIS0	-
c2t40d0s2	ENABLED	-	c2t40d0s2	GENESIS	GENESIS0	-

11.20 RMI Time-Out

RMI Time-Out Period is the amount of time that a user cannot log back into the system after failing to log on. The default value is 1 minute. This value should be changed to 15 minutes.

11.21 Array Configuration Backup, Review, and Large-Scale Changes

There are two areas that deal with the different tasks of the configuration files. The two areas are the Control Panel Download/Restore and the Configuration File Loader. The Control Panel is used to backup, review, or restore current configuration files. The Configuration File Loader is used to make (or review) large-scale configuration changes.

Download (Backup/Review) from Control Panel

The following items can be downloaded for purposes of backup/restore from Storage Navigator:

- Storage Device List
- User Account List
- Environment Parameter List

Regular backups should be made as frequently as array changes are made to ensure a restorable configuration in case it is lost. Loss of this configuration is unlikely, but the possibility exists. The file format created here is a compressed file that contains a .csv file (comma separated spreadsheet).

For instructions go to [Downloading Configuration, User Accounts and Log Files in Control Panel](#).

Restore (with Backup) from Control Panel

The following configuration files can be restored via the Control Panel:

- Storage Device List
- User Account List
- Environment Parameter List

For instructions go to [Download Configuration, User Accounts, Log Files in Control Panel](#).

Note: A Restore is potentially dangerous and should be done with HDS supervision.

Configuration File Loader (Large Scale Changes)

Configuration File Loader is used to make large-scale or “batch” changes to the LUN Manager configuration of the TagmaStore™ USP subsystem, using either a text editor or spreadsheet software to edit configuration definition files (spreadsheet files) and apply the changes. You can apply the same configuration definition file to different subsystems, and you can define the configuration when you are offline. The file format created here is a .spd file (spreadsheet).

This should be the primary form of configuration review for an array. This file should be submitted to HDS/Sun Microsystems as the approved configuration of an array before it is deployed. This should be similar to a “Bin File Review” process.

For instructions go to the [Configuration File Loader \(Large-Scale Changes\)](#).

11.22 Ordering (Standard Configurations)

The following is the standard Citi configurations. Part numbers have been removed to prevent this document from frequent updates. For a detailed LDEV layout under the various drive and raid configurations, see the [Logical Device \(LDEV\) Sizing](#) section under Best Practices.

The USPV is the hardware platform of choice between the USP and the USPV. Unless there are strong business reasons to purchase a USP, the USPV should be the default choice.

11.22.1 USPV Configurations

USPVM Configuration – Max 26.9TB (RAID 5 300GB) Usable

- HDS USPV
- 1 DKC
- 116 Drives (146GB 15k, 300GB 15k, 1000GB 7.2k SATA)
- 2-4 Spare Drives (at least 2 per drive type)
- 1 BED (DKA) Pairs
- 24-40GB or 32-64GB Cache (Normal or HA)*
- 16GB or 20GB Control Memory (Normal or HA)*
- 6x 8-port host fiber cards (called 3x “16-Port Option”)
- Dual SVPs
- Software
 - o Resource Manager
 - o HiCommand Device Manager
 - o HiCommand Tiered Storage Manager
 - o Performance Monitor
 - o ShadowImage with CCI
 - o Dynamic Provisioning

*HA is for High Availability. This array can be configured for use with HUR (Hitachi Universal Replicator) and or True Copy Synchronous.

Standard Small Configuration – Max 27.4TB (RAID 5 300GB) Usable

- HDS USPV
- 1 DKC
- 120 Drives (146GB 15k, 300GB 15k, 1000GB 7.2k SATA)

- 4 Spare Drives (at least 2 per drive type)
- 2 BED (DKA) Pairs
- 24-40GB or 32-64GB Cache (Normal or HA)*
- 16GB or 20GB Control Memory (Normal or HA)*
- 6x 8-port host fiber cards (called 3x “16-Port Option”)
- Dual SVPs
- Software
 - o Resource Manager
 - o HiCommand Device Manager
 - o HiCommand Tiered Storage Manager
 - o Performance Monitor
 - o ShadowImage with CCI
 - o Dynamic Provisioning

*HA is for High Availability. This array can be configured for use with HUR (Hitachi Universal Replicator) and or True Copy Synchronous.

Standard Medium Configuration – Max 86.0TB (RAID 5 300GB) Usable

- HDS USPV
- 1 DKC
- 1 DKU (R1)
- 376 Drives (146GB 15k, 300GB 15k, 1000GB 7.2k SATA)
- 4 Spare Drives (at least 2 per drive type)
- 4 BED (DKA) Pairs
- 64-120GB or 96-192GB Cache (Normal or HA)*
- 16GB or 20GB Control Memory (Normal or HA)*
- 8x 8-port host fiber cards (called 4x “16-Port Option”)
- Dual SVPs
- Software
 - o Resource Manager
 - o HiCommand Device Manager
 - o HiCommand Tiered Storage Manager
 - o Performance Monitor
 - o ShadowImage with CCI
 - o Dynamic Provisioning

*HA is for High Availability. This array can be configured for use with HUR (Hitachi Universal Replicator) and or True Copy Synchronous.

Standard Large Configuration – Max 144.5TB (RAID 5 300GB) Usable

- HDS USPV
- 1 DKC
- 2 DKUs (R1 & L1)
- 632 Drives (146GB 15k, 300GB 15k, 1000GB 7.2k SATA)
- 6 Spare Drives (at least 2 per drive type)
- 8 BED (DKA) Pairs
- 96-200GB or 160-256GB Cache (Normal or HA)*
- 16GB or 20GB Control Memory (Normal or HA)*
- 8x 8-port host fiber cards (called 4x “16-Port Option”)
- Dual SVPs
- Software
 - o Resource Manager
 - o HiCommand Device Manager

- o HiCommand Tiered Storage Manager
- o Performance Monitor
- o ShadowImage with CCI
- o Dynamic Provisioning

*HA is for High Availability. This array can be configured for use with HUR (Hitachi Universal Replicator) and or True Copy Synchronous.

Standard XLarge Configuration – Max 144.5TB (RAID 5 300GB) Usable

- HDS USPV
- 1 DKC
- 4 DKUs (R1, R2* & L1, L2)
- 1016 Drives (146GB 15k, 300GB 15k, 1000GB 7.2k SATA)
- 6-8 Spare Drives (at least 2 per drive type)
- 8 BED (DKA) Pairs
- 128-256GB or 160-512GB Cache (Normal or HA)**
- 16GB or 20GB Control Memory (Normal or HA)**
- 10x 8-port host fiber cards (called 4x “16-Port Option”)
- Dual SVPs
- Software
 - o Resource Manager
 - o HiCommand Device Manager
 - o HiCommand Tiered Storage Manager
 - o Performance Monitor
 - o ShadowImage with CCI
 - o Dynamic Provisioning

*1016 Drive configuration requires that R0 (DKU) or lower half of R2 be empty (no drives)

**HA is for High Availability. This array can be configured for use with HUR (Hitachi Universal Replicator) and or True Copy Synchronous.

Note: If there are very strong business requirements and special conditions that dictate a total storage capacity that is between the standard configurations, then it is advised that the larger configuration be used and the desired number of disks be removed in increments of 8. The cache, control memory, spare drives, BEDs, host fiber cards, array frames, and all other physical aspects should remain the same. Software licensing may need to be re-sized appropriately. These special configurations are discouraged.

11.22.2 USP Configurations (Legacy)

Standard Small Configuration – Max 27.4TB (RAID 5 300GB) Usable

- HDS USP 100
- 1 DKC
- 120 Drives (146GB 15k, 73GB 15k, 300GB 10k)
- 2 Spare Drives (per drive type)
- 1 BED (DKA) Pairs
- 28GB Cache
- 6GB Control Memory
- 4x 8-port host fiber cards (called 2x “16-Port Option”)
- Dual SVPs
- Software
 - o Resource Manager – 18TB (Raw) License Key
 - o HiCommand Device Manager – 18TB (Raw) License Key
 - o Performance Monitor – 18TB (Raw) License Key

- o ShadowImage with CCI – 7TB (Usable) License Key*

*Note: The Small configuration of the USP is not appropriate for intensive I/O applications including frequent use of ShadowImage or TrueCopy replication software. A larger configuration is recommended.

Standard Medium Configuration – Max 86.0TB (RAID 5 300GB) Usable

- HDS USP 600
- 1 DKC
- 1 DKU (R1)
- 376 Drives (146GB 15k, 73GB 15k, 300GB 10k)
- 2 Spare Drives (per drive type)
- 2 BED (DKA) Pairs
- 64GB Cache
- 6GB Control Memory
- 4x 8-port host fiber cards (called 2x “16-Port Option”)
- Dual SVPs
- Software
 - o Resource Manager – 55TB (Raw) License Key
 - o HiCommand Device Manager – 55TB (Raw) License Key
 - o Performance Monitor – 55TB (Raw) License Key
 - o ShadowImage with CCI – 23TB (Usable) License Key

Standard Large Configuration – Max 144.5TB (RAID 5 300GB) Usable

- HDS USP 1100
- 1 DKC
- 2 DKUs (R1 & L1)
- 632 Drives (146GB 15k, 73GB 15k, 300GB 10k)
- 4 Spare Drives (at least 2 per drive type)
- 4 BED (DKA) Pairs
- 120GB Cache
- 6GB Control Memory
- 6x 8-port host fiber cards (called 3x “16-Port Option”)
- Dual SVPs
- Software
 - o Resource Manager – 91TB (Raw) License Key
 - o HiCommand Device Manager – 91TB (Raw) License Key
 - o Performance Monitor – 91TB (Raw) License Key
 - o ShadowImage with CCI – 39TB (Usable) License Key

Note: If there are very strong business requirements and special conditions that dictate a total storage capacity that is between the standard configurations, then it is advised that the larger configuration be used and the desired number of disks be removed in increments of 8. The cache, control memory, spare drives, BEDs, host fiber cards, array frames, and all other physical aspects should remain the same. Software licensing may need to be re-sized appropriately. These special configurations are discouraged.

11.23 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.

11.24 Multiple Array Management

For sites that have multiple USP/USPV arrays, it would be prudent to keep SVP panel list up to date on more than one SVP. This could be a primary SVP for another array or it could be a secondary SVP for the same array. Secondary SVP panel screen can be accessed by their IP address. The SVPs do not share these lists naturally. When a secondary SVP takes over the functions of a primary SVP, it uses its own array list, which must be maintained separately.

To maintain a single list between multiple SVPs, the Tool/Control panel can be used to download and upload the user and SVP list. For example, if the (Storage Navigator) user list on one SVP is up to date. This user list can be exported and then uploaded to all other SVPs. This will be easier than maintaining separate user lists. For details on how to access the download/upload functions see [Download Configuration, User Accounts, Log Files in Control Panel](#).

Alternatively, 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. This would also prevent having to maintain lists on multiple SVPs. Desktop icons/shortcuts could be copied from one management server to another on a manual or automated basis.

For more information on how to maintain the SVP panel list see [SVP Panel \(Array Listing\) – Adding New Arrays](#).

11.25 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.

11.26 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 (On the USP this screen is located as a tab under License Key Information).

Once in modify mode, add the new destination ip address to the *Manager* and the *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).

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

Click Apply to make the changes effective.

The SNMP trap ports (62, 64 TCP) 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 following are sample traps as seen from a Trap Monitoring system:

Date	Time	Priority	Hostname	Message
12-13-2007	11:31:31	Local7.Debug	10.40.53.25	community=public, enterprise=1.3.6.1.4.1.116.3.11.4.1.1, uptime=23890462, 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=BF4A21, 1.3.6.1.4.1.116.5.11.4.2.4=1.3.6.1.4.1.116.5.11.4.1.1.7.1.2, 1.3.6.1.4.1.116.5.11.4.2.5=2007/12/13, 1.3.6.1.4.1.116.5.11.4.2.6=10:48:12
12-13-2007	10:53:42	Local7.Debug	10.40.53.13	community=public, enterprise=1.3.6.1.4.1.116.3.11.4.1.1, uptime=23673151, agent_ip=10.40.53.13, generic_num=6, specificTrap_num=3, specificTrap_name=eventTrapREFCODE, version=Ver1, 1.3.6.1.4.1.116.5.11.4.2.1=10192, 1.3.6.1.4.1.116.5.11.4.2.2=RAID500, 1.3.6.1.4.1.116.5.11.4.2.3=BF5010, 1.3.6.1.4.1.116.5.11.4.2.4=1.3.6.1.4.1.116.5.11.4.1.1.6.1.7, 1.3.6.1.4.1.116.5.11.4.2.5=2007/12/13, 1.3.6.1.4.1.116.5.11.4.2.6=10:55:48
12-13-2007	10:26:31	Local7.Debug	10.40.53.13	community=public, enterprise=1.3.6.1.4.1.116.3.11.4.1.1, uptime=23510088, agent_ip=10.40.53.13, generic_num=6, specificTrap_num=3, specificTrap_name=eventTrapREFCODE, version=Ver1, 1.3.6.1.4.1.116.5.11.4.2.1=10192, 1.3.6.1.4.1.116.5.11.4.2.2=RAID500, 1.3.6.1.4.1.116.5.11.4.2.3=BF7082, 1.3.6.1.4.1.116.5.11.4.2.4=1.3.6.1.4.1.116.5.11.4.1.1.7.1.3, 1.3.6.1.4.1.116.5.11.4.2.5=2007/12/13, 1.3.6.1.4.1.116.5.11.4.2.6=10:28:39
12-13-2007	10:07:43	Local7.Debug	10.40.53.25	community=public, enterprise=1.3.6.1.4.1.116.3.11.4.1.1, uptime=23387614, 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=BF4411, 1.3.6.1.4.1.116.5.11.4.2.4=1.3.6.1.4.1.116.5.11.4.1.1.6.1.6, 1.3.6.1.4.1.116.5.11.4.2.5=2007/12/13, 1.3.6.1.4.1.116.5.11.4.2.6=09:59:53
12-13-2007	10:02:41	Local7.Debug	10.40.53.13	community=public, enterprise=1.3.6.1.4.1.116.3.11.4.1.1, uptime=23367070, agent_ip=10.40.53.13, generic_num=6, specificTrap_num=3, specificTrap_name=eventTrapREFCODE, version=Ver1, 1.3.6.1.4.1.116.5.11.4.2.1=10192, 1.3.6.1.4.1.116.5.11.4.2.2=RAID500, 1.3.6.1.4.1.116.5.11.4.2.3=BF4310, 1.3.6.1.4.1.116.5.11.4.2.4=1.3.6.1.4.1.116.5.11.4.1.1.6.1.6, 1.3.6.1.4.1.116.5.11.4.2.5=2007/12/13, 1.3.6.1.4.1.116.5.11.4.2.6=10:04:38
12-19-2007	15:57:41	Local7.Debug	10.40.53.25	community=public, enterprise=1.3.6.1.4.1.116.3.11.4.1.1, uptime=43749416, 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=BF78A2, 1.3.6.1.4.1.116.5.11.4.2.4=1.3.6.1.4.1.116.5.11.4.1.1.6.1.8, 1.3.6.1.4.1.116.5.11.4.2.5=2007/12/19, 1.3.6.1.4.1.116.5.11.4.2.6=15:58:05, 1.3.6.1.4.1.116.5.11.4.2.7="FAN ABNORMALITY"

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).
 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.

11.27 VMWare

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

- HDvM (Device Manager) 6.0
- HTnM (Tuning Manager) 6.0 (including agent for RAID)
- HTSM (Tiered Storage Manager) 6.0

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 restricted for use with Storage Management applications only. This is required for Tuning Manager Agent for RAID.

11.28 HDS Software 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 15 X-LARGE USPV Frames (with all 300GB drives):
 - Four Quad-Core Intel® Xeon™ X7350 Processor (2.93 GHz)
 - 8 GBytes of Memory
 - 35GB Disk space

11.29 Running Multiple Versions of Server Software Simultaneously

It may be necessary to run two versions of server software at the same time. 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. HDS and GDSE support running two versions of HDvM only if once instance is used for making changes. The other version must be used for read-only purposes.

11.30 Physical View (Running SN from HDvM)

Running Physical View from Device Manager (HDvM) launches an instance of the Storage Navigator (SN) client. This is the preferred way to run Storage Navigator because it synchronizes the HDvM database with the changes made under SN.

Note: the same username and password must be configured on both HDvM and SN.

To be able to run 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 \(ACL for HDvM\)](#).

11.31 Using Multiple Storage Tiers (Disk Types) with Replication

HDS Replication software is flexible in that the source (primary) and destination (secondary) volumes can be from different drive types. In order to bring down costs associated with internal and external replication, GDSE supports the use of lower-cost tiers as the destination (secondary) storage on USPV arrays.

With HUR, the destination should not jump by more than one drive class as compared to the source. For example, if the source is 146GB, then the destination should be no larger than the 300GB FC 15k. Using a 146GB → 450GB is not recommended for performance reasons.

The following chart summarizes the supported configurations:

USPV Replication Type	Source (PVOL)	Destination (SVOL)
Internal (ShadowImage)	146GB FC 15k 300GB FC 15k 450GB FC 15k	Any of: 146GB FC 15k 300GB FC 15k 450GB FC 15k 1000GB SATA 7.2k
External Synchronous (TrueCopy)	146GB FC 15k 300GB FC 15k 450GB FC 15k	<i>Should be same due to performance dependency.</i>
External Asynchronous (HUR)	146GB FC 15k 300GB FC 15k 450GB FC 15k	Any of (but no more than one drive type slower): 146GB FC 15k 300GB FC 15k 450GB FC 15k

Notes: If the source is a 146GB FC drive, then the destination should be a 146GB FC or 300GB FC.

TrueCopy (Synchronous Replication) in most configurations must write to the destination before completing the source write so the use of a lower-cost tier at the destination is not recommended. It should be the same as the source.

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

11.32.1 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.

11.32.2 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

11.32.3 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.

11.32.4 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>hcmdsdbtrans /export /workpath C:\temp\trans_work /file
c:\hdvm_trans\db_arc /auto

```

```

KAPM05900-I The hcmdsdbtrans 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 hcmdsdbtrans 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>hcmdsdbtrans /import /workpath C:\temp\trans
_exp /file C:\hdvm_exp\db_arc /type DeviceManager /auto
KAPM05900-I The hcmdsdbtrans 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 hcmsdbtrans 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 hcmsdbtrans 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
```

11.33 Thin Provisioning (HDP) on USPV Arrays

Thin Provisioning (also called HDP) is the default configuration for USPV 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. For details on how to setup and manage an array with HDP see [Citi HDP Best Practices](#). The USP does not support thin provisioning.

11.34 Maximum LUN size

The LUN size limit for thin provisioning is now 2TB (may be increased in future). The maximum traditional thick LUN size (assignable a host) under Citi standards remains 277.52 GB

11.35 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

11.36 System Options (SOM) and System Disks

There are several system options 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 USPv systems:

System Option	Description	Notes
729	HDP: Protect on write when HDP full	This will prevent data corruption when pool is full.
803	HDP: Protect on write when HDP encounters pool volume error (Data Retention Utility enabled)	This will prevent data corruption when pool volume error occurs.
460	Control information backed up to SVP disk drive with PS poweroff	This will save a copy of critical cache-based data.
727	Control information backed up to system disk with PS poweroff	This will save a copy of critical cache-based data.
734	HDP: Repeated SNMP Threshold Exceeded Traps every 8hrs	This is for microcode 60-05-15 and above.


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 the Data Retention Utility recovery procedures.


As a part of the above setting (option 727), a system disk must be created. A system disk is used by the array to maintain certain cache data (from shared memory) during a powerdown or unexpected system failure. It is also used for quick formats. The system disk cannot be a thin LUN (DP-VOL), but it can reside on a parity group that is dedicated to thin storage. This disk will be a dedicated LUN (10GB) that cannot be used for any other storage. If it is not possible to configure a system disk from existing storage, then the system disk should be configured during the next physical disk upgrade (or reconfiguration).

To create a system disk, follow these steps:


1. From Storage Navigator (HDvM → Resources → Arrays → Physical View), go to Lun Expansion/VLL and select the VLL tab. Select Install CV on the right-click menu for a Parity Group that has at least 10GB free storage.



2. Select 1x 10240MB (10GB) LDEV then hit Set.



3. Click Next. Choose an appropriate LDEV address. Click Next.
4. System Disk Setting. When you select the disk and click Set, you will get this confirmation dialog. Click OK.



- Continue with the remainder of the create LDEV process (hitting Apply to make the changes effective). Do not select the quick format option as the system disk cannot be quick formatted. If you inadvertently selected quick format, you will need to manually format the LDEV (right click for option).


11.37 SCSI Reservations on USPV (View / Modify)

In order to view or release Host SCSI reservations, you will need to access the SVP via remote desktop. This will be limited to the HDS RACC system and possibly the HCS system(s). You will be prompted for a password. You will need to login with Administrator/raid-login. You will probably have to fill out an informational screen about who is logging into the SVP.


WARNING: Do not log off the SVP. The applications (including dial-home capabilities) that run on the SVP are not services and will not run if you log off the system. Disconnect from the remote desktop session instead of a logging off when session is no longer needed.

11.37.1 Viewing Host SCSI Reservations

- Select Maintenance on the SVP Main Panel




2. Select **LUNM** (LUN Management from Display Menu). The current configuration load may take a few minutes.



3. Select the port that contains the host group you want to view.

Note: There can be different reservations on different paths to the same device. You could have no reservation on one path and an active reservation on another path.

4. Change the *Display* from Group to LUN.
5. Select the host group you want to view.
6. Under View, select LUN Status. The status column (far right) will show if there is a SCSI Reservation (i.e. PGR / Key) or no reservation (i.e. Normal).



The following are the current SCSI Reservation Types:

Normal	Normal device (no reservation)
BLK	It is not ready due to blockade.
OPR	It is reserved by the normal Open Reserve command.
KEY	Persistent Group Reserve key is set.
PGR	It is reserved by the Persistent Group Reserve command.
MFR	It is reserved by Mainframe.
H35R	It is reserved from the H3500 server.
ACA	It is in the ACA ACTIVE status.

7. You can select the host group from another port to view reservations from another path.
8. Close any of the screens opened during this process. Disconnect the remote desktop session when no longer needed.

WARNING: Do not log off the SVP. The applications (including dial-home capabilities) that run on the SVP are not services and will not run if you log off the system. Disconnect from the remote desktop session instead of a logging off when session is no longer needed.

11.37.2 Modifying Host SCSI Reservations on USPV

To clear SCSI reservations on a USPV, you will need to follow a slightly different path as viewing. That is because the Install screen listed below is not selectable unless you are in Modify Mode.

In addition to having a lock on the array (while in Modify Mode), you will need to obtain a password from HDS support to be able to clear reservations. This is a time-based password that lasts for an hour (based on the time of the SVP). You should consider two or more passwords (one for each hour) in case you are delayed or run into complications.


1. Set Modify Mode.

Note: Modify Mode can be selected before you go to the maintenance screen from the main SVP Application. Once in Modify Mode, a lock on the array is present and no other changes can be made until the the lock is released (or SVP is rebooted). This lock does not expire.

2. Select Install from the SVP main panel.
3. Select Chanhe Configuration on the Install Panel.
4. Select LUN Configuration.
5. Select the port that contains the host group you want to view.

Note: There can be different reservations on different paths to the same device. You could have no reservation on one path and an active reservation on another path.

6. Change the *Display* from Group to LUN.
7. Select the host group you want to view.
8. Under View, select LUN Status. The status column (far right) will show if there is a SCSI Reservation (i.e. PGR / Key) or no reservation (i.e. Normal).




The following are the current SCSI Reservation Types:

Normal
BLK


Normal device (no reservation)
It is not ready due to blockade.

OPR	It is reserved by the normal Open Reserve command.
KEY	Persistent Group Reserve key is set.
PGR	It is reserved by the Persistent Group Reserve command.
MFR	It is reserved by Mainframe.
H35R	It is reserved from the H3500 server.
ACA	It is in the ACA ACTIVE status.

9. Select the LUN(s) in LUN list that you wish to clear the reservations. All reservations for the LUN on this path/host group will be removed.
10. Right-click on selected LUN, or click LUN Management screen LUN command.
11. Select Force Reset on shortcut menu.
12. Click Yes for the message SEC3777i



13. Enter the "one-time" password provided by HDS Support. This password is time sensitive (lasts for an hour based on SVP time) and case sensitive (any letters should be uppercase).



14. The status should change to Normal on the selected LUNs.

Note: There can be different reservations on different paths to the same device. You could have no reservation on one path and an active reservation on another path.

15. Repeat steps 7-14 for any other ports / host groups that are mapped to the LUN.
16. Close any of the screens opened during this process. Turn off Modify Mode. Disconnect the remote desktop session when when no longer needed.

WARNING: Do not log off the SVP. The applications (including dial-home capabilities) that run on the SVP are not services and will not run if you log off the system. Disconnect from the remote desktop session instead of a logging off when session is no longer needed.

12 Security Practices

12.1 Interfaces to the USP/USPV

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

Interface	Abbreviation	Type
Storage Navigator	SN	HTTP
Device Manager GUI	DM	HTTP/HTTPS
Device Manager CLI	CLI	TCP
Host access using Veritas DMP	HBA	FC Only
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
Tiered Storage Manager GUI	TSM	HTTP/HTTPS
Tiered Storage Manager CLI	TSMCLI	RMI

12.2 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 userids. For more information see [User Management](#).

For information regarding the setup of other user accounts, see [Special User Accounts](#).

The default userid and password for the SVP is “Administrator/raid-login”. This password should be changed immediately as all HDS customers would have this in common with Citi.

12.3 Special User Accounts (Functional IDs)

There are several special user accounts that need to be created to tie in external applications to the HDS arrays. 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.
- The passwords must follow all of the same security content and failed login attempt rules as do all userids/passwords.
- 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 [IT Risk Management Technical Requirements for Functional Accounts](#).

The following user accounts are functional IDs and will need at least annual password renewals (more frequent changes are desirable):

Onaro SANScreen User in HDvM

On Device Manager an ID needs to be created for Onaro SANScreen. To make global auditing consistent, create a user “*Sanscreen*” with view permissions only (with appropriate storage resource allocations). On SANScreen, this userid needs to be associated with each array object. This is a non-privileged, interactive functional id account. For more information regarding Onaro SANScreen see [SANScreen Solutions Guide](#).

Device Manager User in SN

On Storage Navigator for each array, an ID needs to be created for Device Manager. To make global auditing consistent, create a user “*DeviceMgr*” with modify permission for all functions and the storage administrator role enabled. The account administrator role for this user should be disabled. On Device Manager, this userid needs to be associated with each array object. This is a non-privileged, interactive functional id account.

Tiered Storage Manager User in HDvM

On Device Manager, an ID needs to be created for Tiered Storage Manager. To make global auditing consistent, create a user “*tsm_admin*” with admin privileges except for user management (with appropriate storage resource allocations). This is a non-privileged, interactive functional id account. For more information regarding TSM, see [Appendix L: Tiered Storage Manager](#).

EMC Control Center (ECC) User in HDvM

On Device Manager an ID needs to be created for EMC Control Center (ECC). To make global auditing consistent, create a user “*ECCUser*” with view permissions only (with appropriate storage resource allocations). This is a non-privileged, interactive functional id account. For more information regarding EMC Control Center see [EMC Control Center Solutions Guide](#).

12.4 PowerBroker Configuration (CLI and CCI Only)

In order to secure the CLI and CCI environments, PowerBroker needs to be configured to include commands for the HDS command line applications. The CLI package(s) are installed with root-only access. PowerBroker will allow specified users to access these commands and will log it. The Solaris DM command line client will refer to a PowerBroker server that contains a configuration file that lists specific users to have access to specified commands.

The CLI package for Solaris has been provided with an alternative abbreviated name (HCLI.sh) for the HiCommand CLI shell script (normally HiCommandCLI.sh). PowerBroker (using pbrun) must be invoked in order to run the HDS command line. The following are example commands, response, and the corresponding log entry for PowerBroker.

Sample Commands (Non-functional)

```
# HiCommandCLI.sh GetServerInfo -u server
ksh: HiCommandCLI.sh: cannot execute
```

```
# HCLI.sh GetServerInfo -u system
ksh: HCLI.sh: cannot execute
```

Sample Commands (Functional)

```
# pbrun HCLI.sh GetStorageArray -u system
Password required; Password: manager
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
...
PowerBroker Log entry:
adminprogs[] 'symdg' 'symdev' 'symmir' 'symbcv' 'symmask' 'symmaskdb' 'syminq'
'symacl' 'symcfg' 'symlmf' 'symrdf' 'symcli' 'symconfigure' 'symevent' 'symgat
e' 'symld' 'symoptmz' 'sympd' 'symaudit' 'navicli' 'HC' 'HCLI.sh' 'hc' 'HiComman
dCLI.sh'
adminusers[] 'ludwig' 'am37267' 'mb27571' 'sk34158' 'gj06529'
argc 4
argv[] 'HCLI.sh' 'GetStorageArray' '-u' 'system'
bkgd 0
command 'HCLI.sh'
cwd '/export/opt/HiCommandCLI'
date '2005/11/08'
day 8
dayname 'Tue'
env[] '_=/opt/pb/bin/pbrun' 'MANPATH=/usr/man:/export/opt/emc/EMCpower/man:/us
r/local/man:/usr/local/qt/doc/man' '_INIT_UTC_RELEASE=5.8' 'HZ=100' '_INIT_UTC_M
ACHINE=sun4u' 'PATH=/usr/sbin:/usr/bin:/usr/ucb:/etc:/usr/openwin/bin:/usr/local
/bin:/usr/ccs/bin:/usr/local/microsoft/bin:/export/home/ludwig/bin:/usr/local/sa
mba/bin:/usr/lib:/export/opt/netscape:/work:/usr/local/bin:/usr/local/qt/bin:/us
r/dt/bin:/export/opt/emc/WideSky/V5.2.0/bin:/export/opt/emc/SYMCLI/V5.5.0/bin:/v
ar/tmp/VolumeLogix:/opt/pb/bin:.' '_INIT_UTC_VERSION=Generic_117350-20'
'EDITOR=
emacs' '_INIT_RUN_NPREV=0' 'HISTFILE=/tmp/ksh.log4032'
'_INIT_NET_STRATEGY=none'
'LOGNAME=gj06529' '_INIT_UTC_NODENAME=sangs10' '_INIT_UTC_ISA=sparc'
'MAIL=/var
/mail/gj06529' '_INIT_PREV_LEVEL=S' 'VCMDBDEVICE=/dev/rdisk/c2t16d0s2'
'DISPLAY=1
0.98.128.139:0.0' 'SHELL=/bin/ksh'
'LD_LIBRARY_PATH_64=/export/opt/emc/EMCpower/
lib/sparcv9' 'HOME=/usr/gary' '_INIT_UTC_SYSNAME=SunOS'
'RAD_LOG_FACILITY=LOG_LO
CAL0'
'LD_LIBRARY_PATH=/export/opt/emc/EMCpower/lib:/export/opt/emc/VRTSvc/EMC/
lib:/usr/lib:/usr/local/lib:/usr/openwin/lib:/usr/ucblib:/usr/local/qt/lib:.' 'T
ERM=ansi' 'QTDIR=/usr/local/qt' 'PWD=/export/opt/HiCommandCLI' 'TZ=US/Eastern'
'_INIT_RUN_LEVEL=3' '_INIT_UTC_PLATFORM=SUNW,Ultra-80'
event 'Accept'
eventlog '/var/opt/pb/pb.eventlog'
exitstatus 'undefined'
false 0
group 'staff'
groups[] 'staff'
host 'sangs10'
hosts[] 'sangs10' 'sangs7'
hour 15
iolog '/var/log/pb/pb.gj06529.HCLI.sh.agaqZ3'
iologtemplate '/var/log/pb/pb.gj06529.HCLI.sh.XXXXXX'

```



```

lineinfile '/etc/opt/pb/pb.conf'
linenum '23'
localmode 0
lognopassword 1
logport 12347
logservers[] 'sundried'
logstderr 1
logstdin 1
logstdout 1
masterhost 'sundried'
minute 34
month 11
nice 0
pbclientname 'pbrun'
pblocaldmachine 'sun4u'
pblocaldnodename 'sangs10'
pblocaldrelease '5.8'
pblocaldsysname 'SunOS'
pblocaldversion 'Generic_117350-20'
pblogdmachine 'sun4u'
pblogdnodename 'sundried'
pblogdrelease '5.8'
pblogdsysname 'SunOS'
pblogdversion 'Generic_108528-13'
pbmasterdmachine 'sun4u'
pbmasterdnodename 'sundried'
pbmasterdrelease '5.8'
pbmasterdsysname 'SunOS'
pbmasterdversion 'Generic_108528-13'
pbrunmachine 'sun4u'
pbrunnodename 'sangs10'
pbrunrelease '5.8'
pbrunsysname 'SunOS'
pbrunversion 'Generic_117350-20'
pbversion '3.5.2-01'
pid 28568
ptyflags 7
rejectnullpasswords 0
requestuser 'gj06529'
runargv[] 'HCLI.sh' 'GetStorageArray' '-u' 'system'
runbkgd 0
runcommand 'HCLI.sh'
runcwd '/export/opt/HiCommandCLI'
runenv[] '_=/opt/pb/bin/pbrun' 'MANPATH=/usr/man:/export/opt/emc/EMCpower
/man:/usr/local/man:/usr/local/qt/doc/man' '_INIT_UTS_RELEASE=5.8' 'HZ=100' '_IN
IT_UTS_MACHINE=sun4u' 'PATH=/bin:/usr/bin:/usr/sbin:/export/opt/emc/SYMCLI/V5.5.
0/bin:/export/opt/Navisphere/bin:/export/opt/HiCommandCLI'
'_INIT_UTS_VERSION=Ge
neric_117350-20' 'EDITOR=emacs' '_INIT_RUN_NPREV=0'
'HISTFILE=/tmp/ksh.log4032'
'_INIT_NET_STRATEGY=none' 'LOGNAME=gj06529'
'_INIT_UTS_NODENAME=sangs10' '_INIT_
UTS_ISA=sparc' 'MAIL=/var/mail/gj06529' '_INIT_PREV_LEVEL=S'
'VCMDBDEVICE=/dev/r
dsk/c2t16d0s2' 'DISPLAY=10.98.128.139:0.0' 'SHELL=/bin/ksh'
'_LD_LIBRARY_PATH_64=

```

```

/export/opt/emc/EMCpower/lib/sparcv9' 'HOME=/usr/gary'
'_INIT_UTS_SYSNAME=SunOS'
'_RAD_LOG_FACILITY=LOG_LOCAL0'
'_LD_LIBRARY_PATH=/export/opt/emc/EMCpower/lib:/ex
port/opt/emc/VRTSvcs/EMC/lib:/usr/lib:/usr/local/lib:/usr/openwin/lib:/usr/ucbli
b:/usr/local/qt/lib:.' 'TERM=ansi' 'QTDIR=/usr/local/qt' 'PWD=/export/opt/HiComm
andCLI' 'TZ=US/Eastern' '_INIT_RUN_LEVEL=3'
'_INIT_UTS_PLATFORM=SUNW,Ultra-80'
rungroup      '!g!'
rungroups[]    '!G!'
runhost 'sangs10'
runlocalmode 0
runnice 0
runptyflags 7
runtimeout 0
runumask 18
runuser 'root'
status 0
submithost 'sangs10'
submithostip '146.128.134.203'
subprocuser 'root'
time '15:34:54'
timezone 'EST5EDT'
true 1
ttyname '/dev/pts/2'
umask 18
uniqueid '9280016f43710bee6F98'
user 'gj06529'
year 2005
event 'Finish'
exitdate '2005/11/08'
exitstatus 'Command finished with exit status 0'
exittime '15:35:02'
uniqueid '9280016f43710bee6F98'

```

The username and password should not be defined in the HiCommandCLI.properties file. The username should be specified on the command line entry (using the `-u username` option). The user will be prompted for the password. The password (clear text entry) will not be logged, but the rest of the command will be logged with PowerBroker.

To avoid transmission of the clear text password, use the `-s` (or `-secure`) option in the command line. This must be configured after SSL has been configured for the server. For more information, see [Using CLI with HTTPS](#).

For further information on PowerBroker and its configuration for SAN, see [PowerBroker Implementation for SAN](#).

For further information regarding the CLI, see [Appendix D: Command Line Interface for Device Manager](#) and the CCI, see [Appendix E: Command Control Interface / ShadowImage](#).

For details of PowerBroker configuration for TSM, see [Appendix L: HDS Tiered Storage Manager](#).

12.5 Audit Logging (eTrust)

In order to integrate into the eTrust logging system, log files must be provided to an eTrust agent running on the SVP or on a management station. As new events are added to the log file they are processed into the eTrust master logs.

Note: Solaris and Wintel systems may already be monitored via CA eTrust. Modifications will need to be made to support both audit logging and standard Solaris logging to different destinations.

In order to automatically send audit records to an eTrust 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 an eTrust repository for reporting and record keeping. Refer to the latest version of the [eTrust Storage Infrastructure Monitoring](#) document for the latest procedures. 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 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).

12.5.1 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 HiCommand Suite 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 HiCommand Common Component to enable the new settings.

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

TSM auditing also requires this edit. For more information see [Audit Log Configuration - event log](#).

12.6 Turning on SSL and Securing Device Manager (Windows)

In order for Device Manager to support SSL sessions, the following steps must be followed to turn on SSL functionality. It is important that all steps be followed to make the Device Manager web server more secure and thus to limit security exposure.

12.6.1 Creating a Keypair

1. To create a keypair, run the following batch file from the command line:

```
C:\Program Files\HiCommand\DeviceManager\HiCommandServer\HiKeyTool.bat
```

The following menu appears:

```
=====
Copyright(C) 2001, 2005, Hitachi, Ltd.
HiKeytool v4.1
=====
```

- 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) Display contents of Device Manager Server TrustStore
- 11) Display verbose contents of Device Manager Server TrustStore
- 12) Delete an entry from the Device Manager Server TrustStore
- 13) Change Device Manager Server TrustStore Password
- 14) Exit

2. Select option 1 Make KeyPair/Self-Signed Certificate and use the following values except for location and passwords.

```
Enter Server Name [default=emcsrv07]:
Enter Organizational Unit [default=HiCommand Device Manager Administration]:
Enter Organization Name [default=emcsrv07]:Citi
Enter your City or Locality:NYC
Enter your State or Province:NY
Enter your two-character country-code [default=US]:
Enter Key Alias [default=emcsrv07]:
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]:999
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 X500Name for
  emcsrv07...
Creating the Device Manager Server KeyPair for emcsrv07 at:
  C:\Program Files\HiCommand\DeviceManager\HiCommandServer\keystore
  <this can take up to a minute>
All done.
(A)nother command or E(x)it?x
Exiting...
```

12.6.2 Restarting the Device Manager Server

The server needs to be restarted after creating a Keypair.

To Stop the Server:

```
C:\Program Files\HiCommand\DeviceManager\hicommand.bat stop
The HiCommandServer service is stopping.
The HiCommandServer service was stopped successfully.

Press any key to continue . . .
```

To Start the Server:

```
C:\Program Files\HiCommand\DeviceManager\hicommand.bat start
The HiCommandServer service is starting.
The HiCommandServer service was started successfully.

Press any key to continue . . .
```

To check the status of the Server:

```
C:\Program Files\HiCommand\DeviceManager\hicommand.bat status
HiCommand Server is running.
Press any key to continue . . .
```

To stop the component services for Device Manager:

```
C:\Program Files\HiCommand\base\bin\hcmdssrv.exe /stop
KAPM05009-I Already stopped service. service-name=HiCommand Suite Single Sign On
Service
KAPM05009-I Already stopped service. service-name=HiCommand Suite Common Web Service
KAPM06439-I The HiRDB service has stopped.
```

To check the status of the component services for Device Manager:

```
C:\Program Files\HiCommand\base\bin\hcmdssrv.exe /status
KAPM06441-I The HiRDB service has already stopped.
KAPM05009-I Already stopped service. service-name=HiCommand Suite Common Web Service
KAPM05009-I Already stopped service. service-name=HiCommand Suite Single Sign On
Service
```

12.6.3 Enabling TLS/SSL Server Security

1. To enable SSL, run the HiKeyTool from the command line again:

```
C:\Program Files\HiCommand\DeviceManager\HiCommandServer\HiKeyTool.bat
```

The following menu appears:

```
=====
Copyright(C) 2001, 2005, Hitachi, Ltd.
HiKeytool v4.1
=====
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) Display contents of Device Manager Server TrustStore
11) Display verbose contents of Device Manager Server TrustStore
12) Delete an entry from the Device Manager Server TrustStore
13) Change Device Manager Server TrustStore Password
14) Exit
```

2. Select option 2 Set Device Manager Security Level and use the

```
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
```

3. You will need to restart the server to affect this change. See [Restarting the Device Manager Server](#).

12.6.4 Generating a Private Key Using SSLC Utility

To generate a private key, run the SSLC utility from the command line using the following format:

```
C:\Program Files\HiCommand\Base\httpsd\sslc\bin\sslc.exe genrsa -out key-file [512|1024|2048]

For example,
C:\Program Files\HiCommand\Base\httpsd\sslc\bin>sslc genrsa -out demoCA\httpsdkey
y.pem 2048
Loading 'entropy' into random state - 1060 semi-random bytes loaded
Generating 2 prime RSA private key, 2048 bit long modulus
.....+++++
.....+++++

e is 65537 (0x10001)
```

If you encounter any error messages about not being able to load random state, then restart the entire system and retry the command. Be sure to use 2048 bit length for maximum encryption.

12.6.5 Creating a Certificate Signing Request

After creating a Certificate Signing Request (CSR), it will need to be sent to the Certificate Authority (CA). The CA will then return a signed certificate. The exact format of the CSR will vary depending on which CA is chosen.

If the sslc.cnf file is configured in advance, then you will not need to enter information such as Country Name and Locality Name on the command line. Entries in the cnf file will be the default answers to questions. This is a recommended, but not a required step.

The format of the command is:

```
C:\Program Files\HiCommand\Base\httpsd\sslc\bin\sslc.exe req -config configuration-file -new -
key key-file -out CSR-file

For example:
sslc req -config demoCA\sslc.cnf -new -key demoCA\httpsdkey.pem -out demoCA\httpsd.csr
```

12.6.6 Creating a Self-Signed Certificate

If you do not have a signed and trusted certificate, you can create a self-signed certificate using the x509 utility. This is also useful for testing.

The format of the command is:

```
C:\Program Files\HiCommand\Base\httpsd\sslc\bin\sslc.exe x509 -in CSR-file -out certificate-file -
req -signkey key-file -days #days

For example,
sslc x509 -in demoCA\httpsd.csr -out demoCA\selfcert.pem -req -signkey
demoCA\httpsdkey.pem -days 999

Signature OK
Subject=/C=US/ST=NewYork/L=NYC/O=Citi/OU=Website/CN=www.citigroup.com
```

12.6.7 Activating the SSL Configuration

1. Modify the file "httpsd.conf" in C:\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 SSLCertificateFile remains commented out.

```
...
Listen 23015
SSLDisable

Listen 23016
<VirtualHost emcsrv07:23016>
  ServerName emcsrv07
  SSLEnable
  SSLRequireSSL
  SSLCertificateFile "C:/Program Files/HiCommand/Base/httpsd/conf/ssl/server/selfcert.pem"
  SSLCertificateKeyFile "C:/Program Files/HiCommand/Base/httpsd/conf/ssl/server/httpsdkey.pem"
  # SSLCertificateFile "C:/Program Files/HiCommand/Base/httpsd/conf/ssl/server/selfcert.pem"
  SSLSessionCacheTimeout 3600
</VirtualHost>

Include "C:/Program Files/HiCommand/Base/hwc/redirector/mod_jk.conf"
Alias /manual/DeviceManager/ "C:/Program
Files/HiCommand/DeviceManager/HiCommandServer/docroot/HiCommand/"
```

2. Verify that the locations of the certificate files are the same as the directory locations listed in the httpsd.conf file. (Copy the 2 .pem from the demoCA subdirectory to the server subdirectory named in the .conf file above)

3. You will need to restart the server to affect this change. See [Restarting the Device Manager Server](#).

12.6.8 Setting the SSL Host Name

When you set SSL, you will need to specify a logical host name for the cluster manager IP address. That IP address must be enabled and accessible. If that IP address was specified during Device Manager installation, execute the following command:

```
C:\Program Files\HiCommand\Base\bin\hcmdschgurl.bat /change http:// <cluster manager
IPaddress>:port number http:// <host name for the cluster manager IP address>:port number

To find out the current setting, use:
hcmdschgurl.bat /print

For example:
hcmdschgurl.bat /print
http://146.128.135.116:23015 HiCommand Device Manager

hcmdschgurl.bat /change http://146.128.135.116:23015 https://emcsrv07:23016
KAPM06111-I The URL was changed from "http://146.128.135.116:23015" to "https://
emcsrv07:23016"
```

Have users log into **https://server:23016/DeviceManager/** instead of the http address (note the port number change and https)

You will need to restart the server to affect this change — a full system restart is recommended.

12.6.9 Limiting Client IPs to Known Management Systems (ACL for HDvM)

There is an access control list that will limit client (GUI or CLI/API) access to the HDvM server. Management stations (CLI/ GUI), Onaro servers, ECC servers, etc. should be listed here to

further secure the system by limiting access to the Storage Management team(s). In order to limit the use of Device Manager to known management systems, you will need to modify the following file:

```
c:\Program Files\HiCommand\DeviceManager\HiCommandServer\config\server.properties
```

Default Value

```
Server.http.security.clientIP=*. *.*.*
```

Example1

```
server.http.security.clinetIP=127.0.0.1,192.168.93.*
```

Example2

```
server.http.security.clinetIP=127.0.0.1,192.168.93.123, 192.168.93.101, 192.168.93.111
```

In Example1 above, the value will allow connections from the local server that Device Manager is installed and any server in the 93 subnet.

In Example2 above, the value will allow connections from the local server that Device Manager is installed and several known management servers.

You will need to restart the server to affect this change. See [Restarting the Device Manager Server](#).

If this setting has been made to the server and an additional client attaches without being listed, an unable to connect error will result on the client. For example, on a CLI client, the error message will appear:

*Error performing command: Error reading response (check server availability and url):
java.net.SocketException: Connection reset by peer: Connection reset by peer*

12.6.10 Using CLI with HTTPS

To implement CLI while running HTTPS on the Device Manager server, after SSL has been configured:

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

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

You will need to provide a Device Manager username and password to access this file. Since Solaris SOE systems do not have supported GUI capabilities, it is suggested that you use the SAN Management system (Windows) to download the file.

2. Store the downloaded file in the CLI installation directory on the Solaris CLI system. (/export/opt/HiCommandCLI). Use the filename HiCommandCerts.

3. Set the path name of the HiCommandCerts file in the
HDVM_CLI_CERTS_PATH=/export/opt/HiCommandCLI/HiCommandCerts

4. Execute the desired command. For example, type as shown below to execute the GetServerInfo command:
HiCommandCLI.sh -s https://<devicemanagerserver>:2443/service GetServerInfo

Note: Observe the following guidelines:

- Use the https protocol for the Device Manager server URL.
- Use the https port for the Device Manager server URL (2443, unless configured differently in the Server's configuration file).
- Use the -s (or --secure) switch.

Modify the properties file to use the https address rather than the http address to prevent having to specify it in every command.

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

# set the server url
#HiCommandCLI.serverurl=http://localhost:2001/service
HiCommandCLI.serverurl=https://bnehdvm:2443/service
#HiCommandCLI.serverurl=http://10.129.148.30:2001/service
```

Once this configuration has been established, you will need to add the -s option (or --secure) to all commands. If you leave this option out, you will probably get an error such as:

Error performing command: Error reading response (check server availability and url): java.net.SocketException: SSL implementation not available


This error can also be caused by not finding the HiCommandCerts file.

12.7 User Entitlements Reporting

The user entitlements reporting for the USP is a weak area for the product. Limited RBAC and separation of duties further complicates this environment.

12.7.1 Storage Navigator User Entitlements

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](#).



There is a user list download function, but it contains very limited information. For details and sample downloads, see [Download Configuration, User Accounts, Log Files in Control Panel](#).

12.7.2 Device Manager User Entitlements

The best reporting for users 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 \(HTML\)](#).

12.8 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](#).

12.9 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, and the banner text.

The following password policy changes should be made:

- The password length should be set to 6.
- The password content should be set to alphanumeric.
- The password and username cannot be the same.

For more detailed information on how to make these changes, see [Password Policy](#).

The following failed login lockout changes should be made:

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

For more information on how to change this setting, see [Failed Login Attempt Account Lock](#).

In addition there should be a banner with information about the limited use of the software. For more information on how to change this setting, see [Banner Page](#).

12.10 CCI Security

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.

12.11 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 (i.e. copy \HiCommandCLI to a different directory and rename the executable to select the different paths).

In addition the secure (-s) option should be used.

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

12.12 HDvM Array Name


HDvM uses a default name of "USP@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 Onaro (NetApp) 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.

12.13 Active Directory Authentication with HDvM


Active Directory can be used to authenticate users for HDvM. Users are still 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 because it achieves some improvement in the overall security of the system and reduces the workload associated with entitlement reviews. In order to configure HDvM with Active Directory, the following steps must be completed.

12.13.1 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 HDvM 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 "Device Manager 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 Aadd. 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.

12.13.2 HDvM Configuration of Radius

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

Note: Save a copy of the file; delete all lines with "ldap" in them; edit the word "ServerName" to the name of the server you installed IAS on. Also change "localhost" to the same name as servername; change 127.0.0.1 to IP of the IAS server. Uncomment all lines (remove #) 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

2. Define the IAS server and secret password (entered on the IAS client panel) to HDvM. 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.
```

3. 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.
```

4. Validate your server is defined

```
hcmdsradiussecret /list
```


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

5. 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>


```

12.13.3 HDvM User Setup

1. Next change the accounts defined in HDvM 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.

12.14 ICQ Compensating Controls and Failures

An ICQ (Infrastructure Compliance Questionnaire) has been completed for the HDS TagmaStore USP/USPV Storage Array product.

This product is dependent on a SOE Management server (Windows), SOE CLI Management server (Solaris), a CEP/firewall, eTrust Audit, VPN (Remote Access), and Onaro SANSscreen. To meet the security requirements the following compensating controls are to be implemented:

Security Issue	Compensating Control
Password Management (CISS 4.0 2.1)	All users will be authenticated by the SOE [management] server. See section ICQ Section 2.8, 2.9 (Password Mgmt) below.
Password Management?(CISS 4.0 2.2, 2.2.1, 2.6, 2.8)	Handled via the SOE user account login process on a SOE server. See section ICQ Section 2.8, 2.9 (Password Mgmt) below.
Clear Text Passwords (CISS 4.0 2.5.2)	CLI displays passwords in clear text, however the authentication that is enforced on the CLI is done by the SOE server and PowerBroker that does not display passwords in clear text. Only authorized users (limited to storage admin group and system admin group) can use the specific management (SOE) server with the software. Users must log into the management server (logged who was on server at any given time). Users must have detailed and privileged knowledge of the management software operations (limited to storage admin group). Users must have detailed and privileged knowledge of the storage environment (limited to storage admin group). These conditions significantly lower the risk of the compensating control, making it fully effective.
Password Management (CISS 4.0 2.6, 2.8)	Handled via the SOE user account login process on a SOE server. See section ICQ Section 2.8, 2.9 (Password Mgmt) below.
Functional ID Passwords (CISS 4.0 2.12)	See section ICQ Section 2.11 (Functional ID Password) below.
Audit Logs (CISS 4.0 11.1.1, 11.1.2)	Will be handled by the SOE server security logging and with admin/security logs being sent to eTrust Audit logging system.
Maker/Checker: CISS 4.0 6.4.2	See section ICQ Section 6.5 (Maker/Checker) below.
Default Access: CISS 4.0 6.6	See section ICQ Section 6.6 (Default Passwords) below.
Encryption (CISS 4.0 x.x)	Passwords can be stored using MD5 hashing, however it is for secondary authentication. The primary authentication (compliant) is the SOE server. AES and 3DES are used for SSL. Only authorized users (limited to storage admin group and system admin group) can use the specific management (SOE) server with the software. Users must log into the management server (logged who was on server at any given time). Users must have detailed and privileged knowledge of the management software operations (limited to storage admin group). Users must have detailed and privileged knowledge of the storage environment (limited to storage admin group). These conditions significantly lower the risk of the compensating control, making it fully effective.

12.14.1 Recommended Level-3 Procedures

To meet the security requirements as defined in the ICQ, the following Level-3 procedures should be developed and implemented by the Regional Operations Teams:

12.14.2 ICQ Section 2.8, 2.9 (Password Mgmt)

A Level-3 password management process should be implemented such that all passwords are changed every 60 days and that UserIDs are disabled after 60 days of inactivity. Components that need to be included in this process are Device Manager and Storage Navigator. The SOE server forces password complexity. Passwords and userids should not be shared.

12.14.3 ICQ Section 2.11 (Functional ID Password)

A Level-3 password management process should be implemented such that all functional passwords are changed every 365 days and that UserIDs are disabled after 60 days of inactivity. Functional IDs should not be used by any users other than system processes.

12.14.4 ICQ Section 3.1 (Entitlement Review)

A Level-3 security review process should be implemented such that all UserIDs are reviewed to verify that they have the appropriate privileges for their job responsibility. Components that need to be included in this process are Device Manager and Storage Navigator.

12.14.5 ICQ Section 3.2 (Audit Logs)

A Level-3 audit review process should be implemented such that audit logs are reviewed on at least a quarterly basis.

12.14.6 ICQ Section 3.4 (Audit Logs)

A Level-3 audit review process should be implemented such that audit logs are reconciled to individual users on at least a quarterly basis.

12.14.7 ICQ Section 5.3 (Audit Logs)

A Level-3 audit review process should be implemented such that audit logs are reviewed for any suspicious activity or security violations. Suspicious activities may include (but not limited to), repeated failed login attempts, inconsistencies between change management and actual system changes, and frequent security changes. If any discrepancies or suspicious activity is identified in the audit trails, obtain evidence of escalation and closure.

12.14.8 ICQ Section 6.4 (RBAC)

A Level-3 ISA process should be implemented such that a security user role is created and used only for assigned management staff performing ISA functions. Only ISA functions will be used by this role (add/delete/modify users). System administrators making changes to the array configuration will use a storage administrator role.

12.14.9 ICQ Section 6.5 (Maker/Checker)

A Level-3 maker/checker process should be implemented such that all SAN access changes to the array are documented in a maker/checker methodology using Onaro SanScreen and/or eTrust. This requires that the SanScreen system be given administrative privileges to the array via Device Manager. See section [Special User Accounts](#) of this document for a description the userids required. For more information about eTrust, see [Audit Logging \(eTrust\)](#).

In addition to SanScreen and/or eTrust, there should be a comprehensive change management and verification process in place such that all changes are reviewed and approved prior to any implementation and verified afterwards.

12.14.10 ICQ Section 6.6 (Default Passwords)

A Level-3 password management process should be implemented such that all default passwords are changed prior to implementation of the array and passwords for all users on a 60-day basis thereafter. Components that need to be included in this process are Device Manager, Storage Navigator, Tiered Storage Manager, and the SVP.


13 References

Document	Location
GDSE HDS USP Product Page	http://engineering.citigroup.net/gdse/products/platforms/tagmstore/
GDSE Supported MultiPath Configuration Matrix	http://engineering.citigroup.net/gdse/products/matrices/supported_multipath_config.pdf
GUE Website	http://unixeng.ssmb.com/
eTrust Storage Infrastructure Monitoring	http://engineering.ti.citigroup.net/gdse/products/services/etrustaudit/
Veritas Array Support Library for HDS TagmaStore USP on Veritas 4.X	http://seer.support.veritas.com/docs/272025.htm
Veritas Array Support Library for HDS TagmaStore USP on Veritas 3.X	http://seer.support.veritas.com/docs/271477.htm
PowerBroker Implementation for SAN	http://engineering.citigroup.net/gdse/products/platforms/symmetrix/stdsdocs/05006-s.pdf

14 Appendix A: Management Station

It is important to understand that the SOE management servers are on the corporate network and the SVPs must be behind a firewall or CEP. The Management servers communicate with the array via the SVPs through the firewall / CEP. The management servers should be restricted for use by only designated Storage Administrators. This is a corrective action to give the environment the necessary security and audit features.

The following is a logical depiction of a typical environment:



A Management Station will be used to hold the Hitachi Device Manager Server/Client software and Storage Navigator client software for managing the IPS. This server should be SOE built and have a remote desktop server running on it. The Management Station should be treated as the only point of management for the HDS USPs at each site. Administrators needing to manage remotely can remote desktop to this management station and then use Device Manager to manage the USP.

14.1.1 Number of Servers

There should be no more than 8-10 arrays managed by a single DM instance (assuming 1-2 processors). Additional memory and processors may be able to compensate for this limitation. There should be at least one wintel server for DM and GUI (DM and SN)

clients. There should be at least one Solaris server for the CLI and syslog. Redundant systems (1 Wintel, 1 Solaris if possible) are also recommended.

14.1.2 Device Manager Client for Wintel

Internet Explorer 6.0 or higher

Requires SUN JRE 1.4.2 or higher (including Java Web Start)

JRE 1.5.0 (including Java Web Start) can also be used. .

14.1.3 Device Manager Client for Solaris

This is not recommended due to lack of internal support of a web browser on a Solaris server.

14.1.4 Device Manager Command Line Client for Solaris (DM CLI)

Requires PowerBroker to be installed and configured to log any use of the command line interface (CLI) to DM. PB should be configured to send its logs to a centralized log server for periodic review. This server should also be configured to receive syslog information from the SVP. Syslog is a native service for Solaris, but can also be configured on other systems.

14.1.5 Storage Navigator Client for Wintel

Internet Explorer 6.0 or higher

Requires SUN JRE 1.4.2 or higher (including Java Web Start)

14.1.6 Storage Navigator Client for Solaris

This is not recommended due to lack of internal support of a web browser on a Solaris server.

14.1.7 Storage Navigator Server

This software is hosted by the built-in SVPs on the array.

14.1.8 Device Manager Windows 2k3/2k Server

This configuration is the preferred method of installation. It is also the most efficient version of Device Manager server.

Note: The IP address that is entered as a part of the installation (*Enter the IP address or host name of the server that the client accesses from a Web browser:*) is the IP address of the web server that will be installed and configured for accessing Device Manager. It is not the IP address of the array.

15 Appendix B: Configuration of Array via Storage Navigator

15.1 Accessing Storage Navigator


The best way to access Storage Navigator is to launch it from Device Manager. This is the “physical view” button available for each storage array that is defined in the DM database.

To obtain access to Storage Navigator (SN) directly, the following URLs are to be used:

<http://www.xxx.yyy.zzz/cgi-bin/utility/sjc0000.cgi> (USP)
<http://www.xxx.yyy.zzz/> (USPV)

Where www.xxx.yyy.zzz is the IP address of the primary SVP.

There should be at least one array listed on the main screen. Since the SN from one SVP can manage multiple arrays, it may be desirable to consolidate the arrays into a single list (other USPs, USPVs, and 9980's). *However, maintaining a list of arrays on all arrays would be an unnecessary complexity. It would be easier to have a shortcut or desktop icon for each SVP instead.*





To edit the list, select the modify mode button  at the top right of the screen (and enter the username and password of an administrator). Additional arrays may be added or modified. Once the list is satisfactory, then the array of interest should be selected by its number or nickname hyperlinks. Do not use commas or other punctuation in the location field.

The first time you load SN, it takes several minutes to load the Java Applet (10.6MB on USPV). This is normal for the first time only.

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.

15.2 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.



On the USP, installed options will be navigatable 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.

15.3 Modify Mode

To make a change to any modifiable fields throughout Storage Navigator, 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 . While in modify mode, the session will automatically time-out if there is no activity for more than 30 minutes (for changing the default value see section 11.2). This can be manually reset (even if the time has expired) by selecting the *Reset* button.

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.

15.4 Security in Storage Navigator

The default user id and password for SN is root/root. However this password should be changed immediately and other users should be added to Storage User Account List. Root should not be used after the initial configuration of userids.

Note: with USP microcode releases prior to 50-07-69, the user access is controlled via the control panel. To access this utility, (while logged in as a superuser) press the user account icon  in the far upper right corner of Storage Navigator. Starting with 50-07-69, these functions have been removed from the Control Panel.

The Account page (USP: SN → Information Option → Account Tab, USPV: Go → Security → Account) shows user information and what entitlements they have. The Account screen allows the user to create, delete, and modify user accounts.



The account information outline is on the left side of the page and lists users under the Storage Administrator (SA) folder or under specific logical partitions (01-31). Your userid will have a specific open-mouth smile icon(☺). Other users will have another icon with a closed-mouth smile(😊). To view the attributes of a specific user, select the user from the SA list. Immediately to the right of this list will be the User Information List.

Note: HDS provides for security based on partitions. Citi does not currently use this feature so partition-specific features are ignored. Partitions are defined as SLPR00 (ALL) and SLPR01-SLPR31.

SLPR ID indicates the account type. This should be SA (Storage Administrator) for all users.

User Roles

Each storage administrator can have one or more of the following roles assigned:

- **Account Administrator Role.** This role should be assigned to ISA personnel. The account administrator role can include the ability to create, modify, or delete user accounts, as follows:
 - o **Disable:** The user can neither view nor modify the settings of other user accounts.
 - o **View:** The user can view the operation authority of other user accounts.
 - o **Modify:** The user can add and delete user accounts, and can change the settings of other user accounts.
- **Audit Log Administrator Role.** The audit log administrator role can include the ability to download the audit log and access the syslog panel, as follows:
 - o **Disable:** The user cannot download the Audit log file or access the Syslog panel. Storage Partition administrators have audit log functions disabled.

- o **View:** The user can download the audit log file, view but not change the settings in the Syslog panel, and download the syslog information file.
- o **Modify:** The user can download the audit log file and use the Syslog panel to set the syslog servers and download the syslog information file.
- **Storage Administrator Role.** The storage administrator role can include the ability to perform storage device list and set environment operations.
 - o **Disable:** The user cannot perform Storage Device List operations, cannot use optional products, and cannot access the Set Environment panel.
 - o **Enable:** The user can perform the Storage Device List operations, and can make settings in the Set Environment panel.


Specific Entitlement Access

The permission information list is on the lower right of the Account panel. The following list is for the USPV:

Function	Function Detail	Permission
System Information	Information	Modify
LUN Manager	LUN Manager	Modify
LUN Manager	Port	Modify
LUN Manager	Authentication	Modify
LUN Expansion / VLL	LUN Expansion	Modify
LUN Expansion / VLL	VLL	Modify
LUN Expansion / VLL	Pool	Modify
LUN Expansion / VLL	V-VOL	Modify
Cache Residency Manager	Cache Residency	Modify
Performance Manager	Performance Monitor	Modify
Performance Manager	Volume Migration	Modify
TrueCopy	TrueCopy	Modify
Universal Replicator	Universal Replicator	Modify
ShadowImage	ShadowImage	Modify
TrueCopy for z/OS(R)	TrueCopy for z/OS(R)	Modify
Universal Replicator for z/OS(R)	Universal Replicator for z/OS(R)	Modify
ShadowImage for z/OS(R)	ShadowImage for z/OS(R)	Modify
Mainframe Connection	Compatible PAV	Modify
Mainframe Connection	Volume Security	Modify
Mainframe Connection	Volume Retention Manager	Modify
Data Retention Utility	Data Retention	Modify
Universal Volume Manager	Universal Volume Manager	Modify
Configuration File Loader	Configuration File Loader	Modify
Environmental Settings	License Key	Modify
Environmental Settings	SNMP Information	Modify
Environmental Settings	Partition Definition	Modify
Environmental Settings	License Key Partition Definition	Modify

- Permission indicates the access level for each function.
 - o Modify allows the user to perform the selected function.
 - o View allows the user to view the panel, but not perform the selected function.

If changes to the user areas are made, then the following screen may appear:



This indicates that a configuration file download (backup) needs to be made. See [Download Configuration, User Accounts, Log Files in Control Panel](#) for details on how to download these files.

Modifying Users

While in modify mode, you can select one or more rows and right click the permission field on one row to change the selection to Modify or View. You can also change the role permissions described previously. After your changes have been made, select the *Apply* button at the bottom of the screen to make the changes effective.

Modifying User Passwords

While in modify mode, to delete a specific user, select the id from the Account Information outline. Right-click the selection and choose *Change Password* from the menu. After your changes have been made, select the *Apply* button at the bottom of the screen to make the changes effective.

Deleting Users

While in modify mode, to delete a specific user (other than root or the id you are using), select the id from the Account Information outline. Right-click the selection and choose *Delete User* from the menu. After your changes have been made, select the *Apply* button at the bottom of the screen to make the changes effective.

Creating Users


While in modify mode, to create a specific user, select the desired outline folder from the Account Information outline. This should be the SA folder, not a logical partition folder. Right-click the selection and choose *New User* from the menu. A new user screen appears.

 A Windows-style dialog box titled "Hitachi TagmaStore USP" with a tab labeled "New User". It contains three input fields: "User ID" (with a hint "(Number:1-16)"), "Password" (with a hint "(Number:6-16)"), and "Re-enter Password". At the bottom right are "Set" and "Cancel" buttons.

You will need to enter a unique userid and a password (min length of 6 and should be alphanumeric). Select the *Set* button to confirm the entry or *Cancel* to cancel the entry. After your changes have been made, select the *Apply* button at the bottom of the screen to make the changes effective.

15.5 SVP Panel (Array Listing) – Adding New Arrays


When first connecting to Storage Navigator, there is a list of the arrays that the SVP is aware of. This is a manually configured list.



To edit the array list, select the edit icon  in the upper right hand corner. To access the edit screen, the user must have the Storage Administrator Role. For more information on this role, see [Security in Storage Navigator](#). A username and password will be requested (based on the security defined within Storage Navigator from the SVP being accessed).



Once authorized, the array list will change colors to pink and you will have the option to delete, edit, or create (new entry) arrays on the list.



Deleting an existing entry will require a confirmation before the change is made.

Editing existing entries is similar to adding new entries. The following detail screen appears and allows the user to make changes to all fields. The most important fields are the IP address and a unique nickname. All of the remaining fields are informational. Do not check the HTTP block option.

Change Storage Entry - Microsoft Internet Explorer

Storage Entry

No. 1

Nickname	388GX10HU
Product Name	Tagma USP600
Serial No.	10192
Host Name / IP Address	☐ Host Name ☒ IP Address 10 . 40 . 53 . 13
Location	Warren Lab
HTTP Block	☐ SVP which has blocked out HTTP

Submit Cancel

For information on managing multiple arrays, see [Multiple Array Management](#).

15.6 Environmental Parameters

The environmental parameters, such as user time-outs, can be accessed via the control panel.

To access this utility on the USP, press the control panel icon  in the far upper right corner of Storage Navigator.

To access this utility on the USPV, access the following URL:

<http://www.xxx.yyy.zzz/cgi-bin/utility/toolpanel.cgi>

Where www.xxx.yyy.zzz is the IP address of the primary SVP.

Tool Panel - Microsoft Internet Explorer

File Edit View Favorites Tools Help

Back Forward Stop Search Favorites

Address <http://10.40.53.25/cgi-bin/utility/toolpanel.cgi> Go

Links Windows Marketplace Tomcat Server Administration EMC Policy Manager

Tool Panel


Control Panel
Download Trace Files
Update Certificate Files
Set Up HTTP Blocking
Release HTTP Blocking

Login Message

All information transmitted hereby is intended only for the use of the addressee (s) named above and may contain confidential and privileged information. Any unauthorized review, use, disclosure or distribution of confidential and privileged information is prohibited.

Done Local intranet

Select the control panel button. After entering an administrator userid and password, the control panel should now be visible:



User modifiable fields:

- RMI time-out period – This is the length of time that a user cannot log back into SN if the password is not entered correctly three times **(It is recommended that this value be changed to 15 min)**.
- RMI time-out period for Modify – This is the automatic time-out length of time for a user that has placed SN into modify mode without any keyboard activity.
- Polling Period for checking exclusive lock – This value synchronizes SN with CE activities.
- Polling period for checking obstacles – This value synchronizes SN with other changes to the array.
- Status Panel hidden – will hide status panel

15.7 Download Configuration, User Accounts, Log Files in Control Panel

Go to the control panel screen.

To access this utility on the USP, press the control panel icon  in the far upper right corner of Storage Navigator.

To access this utility on the USPV, access the following URL:

<http://www.xxx.yyy.zzz/cgi-bin/utility/toolpanel.cgi>

Where www.xxx.yyy.zzz is the IP address of the primary SVP.

Select the control panel button. After entering an administrator userid and password, the control panel should now be visible.

The default *Download* tab is selected to download backup copies of the configuration files:


- Storage Device List
- User Account List
- Environment Parameter List
- Syslog version of Auditlog (This does not download the log – only some parameters)

Note: The preferred method of Audit Log retrieval is via Syslog. This Syslog version of the audit log has a different format and is described in [Syslog Based Audit Log](#)

This is a very useful tool if you would like to maintain the same user list and/or SVP list on all of the SVPs that are under control. To accomplish this, update the users in a single instance of Storage Navigator. Use the following screens to download the User Account List. This User account list can then be copied to every SVP and then uploaded. This not recommended for frequent updates, but can be managed periodically. This can also be done with the Storage Device list that is the list of SVPs for each array. For more information on multiple array management, see [Multiple Array Management](#).



To download these files, select the checkbox next to the file desired and then select the *Submit* button. A verification screen and then a file transfer dialog will appear.



Note: The downloaded files will be in a compressed tar format. The download message box will prompt for a filename if you choose to save the file. The filename suggested should have a .tar extension instead of the suggested .tgz extension. In order for WinZip to be able to read the file, it must have a .tar extension.

The following are some sample outputs:

```
Userlist.csv
root,607d687c8ec1c54c94b1a54e0de4c4ec,FFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFF
FFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFF
FFFFFFFFFFFFFFFFFFFFFFFF,1,FFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFF
FFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFF
FFF,65535,0,1900/01/01,0,0,1900/01/01,0,0
st49663,576e85cd17c412c099621287a9ff1a5f,FFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFF
FFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFF
F0FEFEFFFFFFFFFFFFFFFF,1,FFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFF
FFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFF
FFFFF,65535,0,1900/01/01,0,0,1900/01/01,0,0
ry92530,97345801fa62e89d0b345d7c3d40b0f5,FFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFF
FFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFF
F0FEFEFFFFFFFFFFFFFFFF,1,FFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFF
FFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFF
FFFFF,65535,0,1900/01/01,0,0,1900/01/01,0,0
ld59486,d7e4c004cc35ac2f7af6e241bf137b18,00000000000000000000000000000000
000000000000000000000000000000000000000000000000000000000000000000
030F0F81,0,FFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFF
FFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFBF,65535,
0,1900/01/01,0,0,1900/01/01,0,0
DeviceMgr,07a420dda82c9e8535425db8c9ba156d,00000000000000000000000000000000
000000000000000000000000000000000000000000000000000000000000000000
010703030B01,0,FFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFF
FFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFF3F,65
535,0,1900/01/01,0,0,1900/01/01,0,0
kh53086,eb9a9b414def9370acf4b09b6251b191,FFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFF
FFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFF
FF0FEFEFFFFFFFFFFFFFFFF,1,FFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFF
FFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFF
FFFFF,65535,0,1900/01/01,0,0,1900/01/01,0,0
gj06529,26ed838aa3b3c1455e4e3078082af2c2,ffffffffffffffffffffffffffffffffffff
ffffffffffffffffffffffffffffffffffffffffffffffffffffffffffffffffffffffffffff
ffffffffffffffffffff,65535,0,1900/01/01,0,0,1900/01/01,0,0
gdse,f4dcaca3032538217008cab8ae90b299,ffffffffffffffffffffffffffffffffffff
ffffffffffffffffffffffffffffffffffffffffffffffffffffffffffffffffffffffffffff
ffffffffffff,65535,0,1900/01/01,0,0,1900/01/01,0,0
cd15948,4e958a3e8a513b6e755a73980f0cd524,FFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFF
FFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFF
FF0FEFEFFFFFFFFFFFFFFFF,1,FFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFF
FFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFF
FFFFF,65535,0,1900/01/01,0,0,1900/01/01,0,0

Env.csv
500,32,30,35,110,10000,01-02-03,1,20,30,JRE1.4.2,0

Strlist.csv
1,283GX11HU,102-10,10102,10.40.53.25,Warren NJ,,0
2,388GX10HU,101-11,10102,10.40.53.13,Warren NJ,,0

Syslog Parameters:
Adtprm.txt
[AdtPrm]
ServerAddr0=0000:0000:0000:0000:000A:0028:002A:001E
ServerAddr1=0000:0000:0000:0000:00C0:00A8:0000:0002
```



```

ServerPort0=00514
ServerPort1=00514
ServerIpVer0=0
ServerIpVer1=0
OutFlg0=1
OutFlg1=0
OutDetail=1
LocationInfo=Local_USP@Warren_Lab10

```

md5_summary.txt (USPV Only)
203c593ef0c3e134d97fe14f09383310

Banner Text (USPV Only):

Wm.txt

All information transmitted hereby is intended only for the use of the addressee(s) named above and may contain confidential and privileged information. Any unauthorized review, use, disclosure or distribution of confidential and privileged information is prohibited.

Downloaded configuration files can be restored (if necessary). To display the restore screen select the *Restore* tab.

Archive files (from Download tab) or CSV files (from Configuration File Utility) can be uploaded here.

Warning: This will overwrite the existing configuration and could change access to the storage.

Use the *Browse* button to point the correct file location.


15.8 Trace Files Download (USPV only)

In addition to the control panel and download screens described above, the USPVP allows the user to download the array trace files. This is usually requested by HDS customer support.

To access the tools utility on the USPVP, go the following URL:

<http://www.xxx.yyy.zzz/cgi-bin/utility/toolpanel.cgi>

Where www.xxx.yyy.zzz is the IP address of the primary SVP.



Select the Download Trace Files button. After entering an administrator username and password, a confirmation dialog will appear. You will then be prompted for a local filename and location. This produces a tgz file can be a large file (> 20MB). The tgz file extracts to be a tar file. The tar file contains approximately 2,810 files.

15.9 Syslog Based Audit Log

The USP or USPV system can be configured to automatically send all audit records from the SVP to one or two syslog servers. This eliminates the need for a manual download of the audit records described in the previous section. If syslog has been configured and is sending records subsequently to eTrust, then warning messages about the audit log (wraparound) can be ignored.

The Syslog screen (USP: Information Option → Syslog tab, USPV: Go Menu → Security → Syslog) shows the enabled/disabled status of two syslog server destinations. 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.

The screenshot shows the 'Storage Navigator' interface with the 'Syslog' tab selected. The 'Syslog Server Setting' section includes the following fields and options:

- Output to Primary Server:** ☒ Enable ☐ Disable
- Primary Server Setting:** IP Address: 10 . 40 . 42 . 30 Port Number: 514
- Output to Secondary Server:** ☐ Enable ☒ Disable
- Secondary Server Setting:** IP Address: 192 . 168 . 0 . 2 Port Number: 514
- Location Identification Name:** JSPV@WarrenLab10
- Output Detailed Information:** ☒ Enable ☐ Disable
- Download Syslog:** [Download button]

At the bottom right, there are 'Apply' and 'Cancel' buttons.

For more information about eTrust Audit logging, see [Audit Logging \(eTrust\)](#).

15.10 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 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,1.1,268330,,2006-10-27T11:23:37.5Z,Storage,SVP,ConfigurationAccess,Success,uid=gj06529,R500:10192,,Local_USP@Warren_Lab10,,,,,,268062,BasicLog,,,99,[BASE],ChangeMode ,Modify,Normal end,Seq.=516290
Oct 27 12:30:46 [10.40.53.13.20.88] Oct 27 11:30:46 SVP Storage: CELFSS,1.1,268331,,2006-10-27T11:30:46.2Z,Storage,SVP,ConfigurationAccess,Success,uid=gj06529,

[illegible]

Oct 27 12:38:03 [10.40.53.13.20.140] Oct 27 11:38:03 SVP Storage: CELFSS,1.1,26
8353,,2006-10-27T11:38:03.2Z,Storage,SVP,ConfigurationAccess,Success,uid=gj06529
,R500:10192,,Local_USP@Warren_Lab10,,,,,,,,268073,BasicLog,,,99,[Account],Set Us
er Account,,Normal end,Seq.=516301

Oct 27 12:38:03 [10.40.53.13.20.141] Oct 27 11:38:03 SVP Storage: CELFSS,1.1,26
8354,,2006-10-27T11:38:03.2Z,Storage,SVP,ConfigurationAccess,Success,uid=gj06529
,R500:10192,,Local_USP@Warren_Lab10,,,,,,,,268073,DetailLog,,,+Operation=Delete
Oct 27 12:38:03 [10.40.53.13.20.142] Oct 27 11:38:03 SVP Storage: CELFSS,1.1,26
8355,,2006-10-27T11:38:03.2Z,Storage,SVP,ConfigurationAccess,Success,uid=gj06529
,R500:10192,,Local_USP@Warren_Lab10,,,,,,,,268073,DetailLog,,,+{UserID,Authority
,SLPRNo,WritePermission,ReadPermission}=[{TestUser,-,-,-,-}] Num of Users=1
Oct 27 12:38:11 [10.40.53.13.20.145] Oct 27 11:38:11 SVP Storage: CELFSS,1.1,26
8356,,2006-10-27T11:38:11.1Z,Storage,SVP,ConfigurationAccess,Success,uid=gj06529
,R500:10192,,Local_USP@Warren_Lab10,,,,,,,,268074,BasicLog,,,99,[BASE],ChangeMod
e.View.Normal end,Seq.=516302

```
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,1.1,204
082,,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 \(NTP\) – Time Synchronization](#). The time reported in eTrust will be based on the time set on the array SVP(s).

15.10.1 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.

15.10.1.1 Syslog on Solaris 8

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

Solaris 8

```

/etc/init.d/syslog
Change call to syslogd from /usr/sbin/syslogd -t to /usr/sbin/syslog
to allow for remote messages to be sent to syslog.

```



```
/usr/sbin/syslogd >/dev/msglog 2>&1 &
```

```
/etc/syslog.conf
```

Add a line for “user.info” to point to a “HDS_messages” file. This file requires the use of tab characters. Note this will require additional scripting to handle file management (prevent file from growing too large).

```
user.info                /var/adm/HDS_messages
```

15.10.1.2 Syslog on Solaris 10

There are two configuration files that need to be modified (on Solaris 10) to enable it for remote logging. In addition, there are SVC area values to modify.

Solaris 10

```
/etc/default/syslogd
```

```
LOG_FROM_REMOTE=YES (all other lines should be commented out)
```

```
/etc/syslog.conf
```

Add a line for “user.info” to point to a “HDS_messages” file. This file requires the use of tab characters. Note this will require additional scripting to handle file management (prevent file from growing too large).

```
user.info                /var/adm/HDS_messages
```

To modify the SVC area values and restart the syslog daemon in Solaris 10, use the following commands:

```
svccfg -s svc:/system/system-log setprop config/log_from_remote = true
svcadm restart svc:/system/system-log
```

For more information about configuring eTrust Audit logging to read the syslog messages, see [Audit Logging \(eTrust\)](#)

15.11 Legacy Audit Log (USP Only)


The preferred method of Audit Log retrieval is via Syslog. This Syslog version of the audit log has a different format and is described in [Syslog Based Audit Log](#). Both the USP and the USPV use the Syslog version of Audit Log.

Contents of the Legacy Audit Log records are:


- Time
- User name
- Setting contents
- Result

System configuration changes, logins/logouts, failed logins, and physical system events are represented in the audit (audit.log) and login (login.log) files.

If there has not been a recent download of audit events an error message may appear (at startup) stating that the audit file is running out of space or that it has run out of space and it has wrapped around. This indicated that audit records have been overwritten. The *Audit* button will also go to a red color with appropriate warning or wraparound message.



Selecting the *Audit* button produces the download screen, where the audit log can be stored locally on the machine you are accessing Storage Navigator. The suggested filename is a function of the current date.



15.12 Port Status

The System screen (USP: Information Option → System tab, USPV: Go Menu → System Information → Port Status) shows port status and some of the basic software and revision information. The port status shows which ports are installed (bold) and which ports are fabric-connected (gold round button).

Storage Navigator - IP:10.40.53.25 - HITACHI DATA SYSTEMS

File Go Help 0 minute(s) remaining in session. Reset User ID : gj06529


Information Port Status Status

Port Status

Cluster-1				Cluster-2			
16FS Fibre 8ch-4mp				16FS Fibre 8ch-4mp			
1A	3A	5E	7E	1J	3J	5N	7N
5A	7A	5E	7E	5J	7J	5N	7N
1B	3B	1F	3F	1K	3K	1P	3P
5B	7B	5F	7F	5K	7K	5P	7P
1C	3C	1G	3G	1L	3L	1Q	3Q
5C	7C	5G	7G	5L	7L	5Q	7Q
1D	3D	1H	3H	1M	3M	1R	3R
5D	7D	5H	7H	5M	7M	5R	7R
1E	3E	1I	3I	1O	3O	1S	3S
5E	7E	5I	7I	5O	7O	5S	7S
1F	3F	1J	3J	1T	3T	1U	3U
5F	7F	5J	7J	5T	7T	5U	7U
1G	3G	1K	3K	1V	3V	1W	3W
5G	7G	5K	7K	5V	7V	5W	7W
1H	3H	1L	3L	1X	3X	1Y	3Y
5H	7H	5L	7L	5X	7X	5Y	7Y
1I	3I	1M	3M	1Z	3Z	1AA	3AA
5I	7I	5M	7M	5Z	7Z	5AA	7AA
1J	3J	1N	3N	1AB	3AB	1AC	3AC
5J	7J	5N	7N	5AB	7AB	5AC	7AC
1K	3K	1O	3O	1AD	3AD	1AE	3AE
5K	7K	5O	7O	5AD	7AD	5AE	7AE
1L	3L	1P	3P	1AF	3AF	1AG	3AG
5L	7L	5P	7P	5AF	7AF	5AG	7AG
1M	3M	1Q	3Q	1AH	3AH	1AI	3AI
5M	7M	5P	7P	5AH	7AH	5AI	7AI
1N	3N	1R	3R	1AJ	3AJ	1AK	3AK
5N	7N	5R	7R	5AJ	7AJ	5AK	7AK
1O	3O	1S	3S	1AL	3AL	1AM	3AM
5O	7O	5S	7S	5AL	7AL	5AM	7AM
1P	3P	1T	3T	1AN	3AN	1AO	3AO
5P	7P	5T	7T	5AN	7AN	5AO	7AO
1Q	3Q	1U	3U	1AP	3AP	1AQ	3AQ
5Q	7Q	5U	7U	5AP	7AP	5AQ	7AQ
1R	3R	1V	3V	1AR	3AR	1AS	3AS
5R	7R	5V	7V	5AR	7AR	5AS	7AS
1S	3S	1W	3W	1AT	3AT	1AU	3AU
5S	7S	5V	7V	5AT	7AT	5AU	7AU
1T	3T	1X	3X	1AV	3AV	1AW	3AW
5T	7T	5X	7X	5AT	7AT	5AW	7AW
1U	3U	1Y	3Y	1AX	3AX	1AY	3AY
5U	7U	5Y	7Y	5AX	7AX	5AY	7AY
1V	3V	1Z	3Z	1BA	3BA	1BB	3BB
5V	7V	5Z	7Z	5AX	7AX	5BB	7BB
1W	3W	1AA	3AA	1BC	3BC	1BD	3BD
5W	7W	5AA	7AA	5AX	7AX	5BB	7BB
1X	3X	1AB	3AB	1BE	3BE	1BF	3BF
5X	7X	5AB	7AB	5AX	7AX	5BB	7BB
1Y	3Y	1AC	3AC	1BG	3BG	1BH	3BH
5Y	7Y	5AB	7AB	5AX	7AX	5BB	7BB
1Z	3Z	1AD	3AD	1BI	3BI	1BJ	3BJ
5Z	7Z	5AB	7AB	5AX	7AX	5BB	7BB
1AA	3AA	1AE	3AE	1BK	3BK	1BL	3BL
5AA	7AA	5AB	7AB	5AX	7AX	5BB	7BB
1AB	3AB	1AF	3AF	1BM	3BM	1BN	3BN
5AB	7AB	5AB	7AB	5AX	7AX	5BB	7BB
1AC	3AC	1AG	3AG	1BO	3BO	1BP	3BP
5AB	7AB	5AB	7AB	5AX	7AX	5BB	7BB
1AD	3AD	1AH	3AH	1BQ	3BQ	1BR	3BR
5AB	7AB	5AB	7AB	5AX	7AX	5BB	7BB
1AE	3AE	1AI	3AI	1BS	3BS	1BT	3BT
5AB	7AB	5AB	7AB	5AX	7AX	5BB	7BB
1AF	3AF	1AJ	3AJ	1BU	3BU	1BV	3BV
5AB	7AB	5AB	7AB	5AX	7AX	5BB	7BB
1AG	3AG	1AK	3AK	1BW	3BW	1BX	3BX
5AB	7AB	5AB	7AB	5AX	7AX	5BB	7BB
1AH	3AH	1AL	3AL	1BY	3BY	1BZ	3BZ
5AB	7AB	5AB	7AB	5AX	7AX	5BB	7BB
1AI	3AI	1AM	3AM	1C0	3C0	1C1	3C1
5AB	7AB	5AB	7AB	5AX	7AX	5BB	7BB
1AJ	3AJ	1AN	3AN	1C2	3C2	1C3	3C3
5AB	7AB	5AB	7AB	5AX	7AX	5BB	7BB
1AK	3AK	1AO	3AO	1C4	3C4	1C5	3C5
5AB	7AB	5AB	7AB	5AX	7AX	5BB	7BB
1AL	3AL	1AP	3AP	1C6	3C6	1C7	3C7
5AB	7AB	5AB	7AB	5AX	7AX	5BB	7BB
1AM	3AM	1AQ	3AQ	1C8	3C8	1C9	3C9
5AB	7AB	5AB	7AB	5AX	7AX	5BB	7BB
1AN	3AN	1AR	3AR	1CA	3CA	1CB	3CB
5AB	7AB	5AB	7AB	5AX	7AX	5BB	7BB
1AO	3AO	1AS	3AS	1CC	3CC	1CD	3CD
5AB	7AB	5AB	7AB	5AX	7AX	5BB	7BB
1AP	3AP	1AT	3AT	1CE	3CE	1CF	3CF
5AB	7AB	5AB	7AB	5AX	7AX	5BB	7BB
1AQ	3AQ	1AU	3AU	1CG	3CG	1CH	3CH
5AB	7AB	5AB	7AB	5AX	7AX	5BB	7BB
1AR	3AR	1AV	3AV	1C0	3C0	1C1	3C1
5AB	7AB	5AB	7AB	5AX	7AX	5BB	7BB
1AS	3AS	1AW	3AW	1C2	3C2	1C3	3C3
5AB	7AB	5AB	7AB	5AX	7AX	5BB	7BB
1AT	3AT	1AX	3AX	1C4	3C4	1C5	3C5
5AB	7AB	5AB	7AB	5AX	7AX	5BB	7BB
1AU	3AU	1AY	3AY	1C6	3C6	1C7	3C7
5AB	7AB	5AB	7AB	5AX	7AX	5BB	7BB
1AV	3AV	1AZ	3AZ	1C8	3C8	1C9	3C9
5AB	7AB	5AB	7AB	5AX	7AX	5BB	7BB
1AW	3AW	1BA	3BA	1C0	3C0	1C1	3C1
5AB	7AB	5AB	7AB	5AX	7AX	5BB	7BB
1AX	3AX	1BB	3BB	1C2	3C2	1C3	3C3
5AB	7AB	5AB	7AB	5AX	7AX	5BB	7BB
1AY	3AY	1BC	3BC	1C4	3C4	1C5	3C5
5AB	7AB	5AB	7AB	5AX	7AX	5BB	7BB
1AZ	3AZ	1BD	3BD	1C6	3C6	1C7	3C7
5AB	7AB	5AB	7AB	5AX	7AX	5BB	7BB
1BA	3BA	1BE	3BE	1C8	3C8	1C9	3C9
5AB	7AB	5AB	7AB	5AX	7AX	5BB	7BB
1BB	3BB	1BF	3BF	1CA	3CA	1CB	3CB
5AB	7AB	5AB	7AB	5AX	7AX	5BB	7BB
1BC	3BC	1BG	3BG	1CC	3CC	1CD	3CD
5AB	7AB	5AB	7AB	5AX	7AX	5BB	7BB
1BD	3BD	1BH	3BH	1CE	3CE	1CF	3CF
5AB	7AB	5AB	7AB	5AX	7AX	5BB	7BB
1BE	3BE	1BI	3BI	1CG	3CG	1CH	3CH
5AB	7AB	5AB	7AB	5AX	7AX	5BB	7BB
1BF	3BF	1BJ	3BJ	1C0	3C0	1C1	3C1
5AB	7AB	5AB	7AB	5AX	7AX	5BB	7BB
1BG	3BG	1BK	3BK	1C2	3C2	1C3	3C3
5AB	7AB	5AB	7AB	5AX	7AX	5BB	7BB
1BH	3BH	1BL	3BL	1C4	3C4	1C5	3C5
5AB	7AB	5AB	7AB	5AX	7AX	5BB	7BB
1BI	3BI	1BM	3BM	1C6	3C6	1C7	3C7
5AB	7AB	5AB	7AB	5AX	7AX	5BB	7BB
1BJ	3BJ	1BN	3BN	1C8	3C8	1C9	3C9
5AB	7AB	5AB	7AB	5AX	7AX	5BB	7BB
1BK	3BK	1BO	3BO	1CA	3CA	1CB	3CB
5AB	7AB	5AB	7AB	5AX	7AX	5BB	7BB
1BL	3BL	1BP	3BP	1CC	3CC	1CD	3CD
5AB	7AB	5AB	7AB	5AX	7AX	5BB	7BB
1BM	3BM	1BQ	3BQ	1CE	3CE	1CF	3CF
5AB	7AB	5AB	7AB	5AX	7AX	5BB	7BB
1BN	3BN	1BR	3BR	1CG	3CG	1CH	3CH
5AB	7AB	5AB	7AB	5AX	7AX	5BB	7BB
1BO	3BO	1BS	3BS	1C0	3C0	1C1	3C1
5AB	7AB	5AB	7AB	5AX	7AX	5BB	7BB
1BP	3BP	1BT	3BT	1C2	3C2	1C3	3C3
5AB	7AB	5AB	7AB	5AX	7AX	5BB	7BB
1BQ	3BQ	1BU	3BU	1C4	3C4	1C5	3C5
5AB	7AB	5AB	7AB	5AX	7AX	5BB	7BB
1BR	3BR	1BV	3BV	1C6	3C6	1C7	3C7
5AB	7AB	5AB	7AB	5AX	7AX	5BB	7BB
1BS	3BS	1BW	3BW	1C8	3C8	1C9	3C9
5AB	7AB	5AB	7AB	5AX	7AX	5BB	7BB
1BT	3BT	1BX	3BX	1CA	3CA	1CB	3CB
5AB	7AB	5AB	7AB	5AX	7AX	5BB	7BB
1BU	3BU	1BY	3BY	1CC	3CC	1CD	3CD
5AB	7AB	5AB	7AB	5AX	7AX	5BB	7BB
1BV	3BV	1BZ	3BZ	1CE	3CE	1CF	3CF
5AB	7AB	5AB	7AB	5AX	7AX	5BB	7BB
1BW	3BW	1C0	3C0	1CG	3CG	1CH	3CH
5AB	7AB	5AB	7AB	5AX	7AX	5BB	7BB
1BX	3BX	1C1	3C1	1C0	3C0	1C1	3C1
5AB	7AB	5AB	7AB	5AX	7AX	5BB	7BB
1BY	3BY	1C2	3C2	1C2	3C2	1C3	3C3
5AB	7AB	5AB	7AB	5AX	7AX	5BB	7BB
1BZ	3BZ	1C3	3C3	1C4	3C4	1C5	3C5
5AB	7AB	5AB	7AB	5AX	7AX	5BB	7BB
1C0	3C0	1C4	3C4	1C6	3C6	1C7	3C7
5AB	7AB	5AB	7AB	5AX	7AX	5BB	7BB
1C1	3C1	1C5	3C5	1C8	3C8	1C9	3C9
5AB	7AB	5AB	7AB	5AX	7AX	5BB	7BB
1C2	3C2	1C6	3C6	1CA	3CA	1CB	3CB
5AB	7AB	5AB	7AB	5AX	7AX	5BB	7BB
1C3	3C3	1C7	3C7	1CC	3CC	1CD	3CD
5AB	7AB	5AB	7AB	5AX	7AX	5BB	7BB
1C4	3C4	1C8	3C8	1CE	3CE	1CF	3CF
5AB	7AB	5AB	7AB	5AX	7AX	5BB	7BB
1C5	3C5	1C9	3C9	1CG	3CG	1CH	3CH
5AB	7AB	5AB	7AB	5AX	7AX	5BB	7BB
1C6	3C6	1CA	3CA	1C0	3C0	1C1	3C1
5AB	7AB	5AB	7AB	5AX	7AX	5BB	7BB
1C7	3C7	1CB	3CB	1C2	3C2	1C3	3C3
5AB	7AB	5AB	7AB	5AX	7AX	5BB	7BB
1C8	3C8	1CC	3CC	1C4	3C4	1C5	3C5
5AB	7AB	5AB	7AB	5AX	7AX	5BB	7BB
1C9	3C9	1CD	3CD	1C6	3C6	1C7	3C7
5AB	7AB	5AB	7AB	5AX	7AX	5BB	7BB
1CA	3CA	1CE	3CE	1C8	3C8	1C9	3C9
5AB	7AB	5AB	7AB	5AX	7AX	5BB	7BB
1CB	3CB	1CF	3CF	1CA	3CA	1CB	3CB
5AB	7AB	5AB	7AB	5AX	7AX	5BB	7BB
1CC	3CC	1CG	3CG	1CC	3CC	1CD	3CD
5AB	7AB	5AB	7AB	5AX	7AX	5BB	7BB
1CD	3CD	1CH	3CH	1CE	3CE	1CF	3CF
5AB	7AB	5AB	7AB	5AX	7AX	5BB	7BB

- 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.



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.


15.14 Port Definition

1. Select the Port Definition screen (LUN Manager Option → Port tab)
2. Make sure that all ports to be connected to the SAN (for host connectivity) are P-to-P (Point to Point) Connection, Fabric enabled, Auto (for 2GB and 4GB SANs), and Target Attribute. If the Target attribute is set incorrectly (i.e. it was used for another option) then refer to:

External:	Universal Volume Manger
Initiator or RCU Target:	True Copy

15.15 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. For example, RAID Group 1-2 is the second RAID group from Box 1. The following shows how boxes are organized:



Note: A “box” may also be referred to as a “chassis” in the CLI.

Citi uses 8-disk raid groups for maximum performance. A USP can have the DKC (primary controller and disk cabinet) and up to four additional DKU (disk cabinets). Typical configuration will consist of a DKC and 1-2 DKUs.

At a minimum, all of the raid groups on a new array should be configured and available for volume creation as free space.

15.16 Volume (LDEV) Creation / Definition


You can create volumes out of freespace or out of existing volumes in 2 different ways. They are very similar in procedure, but they differ in the destruction of any existing volumes:

To create volumes out of freespace (non-destructive), follow the instructions below for [Install Custom Volume \(CV\)](#).


To create volumes from existing volumes (destructive), follow the instructions below for [Make Volume](#).

15.16.1 Install Custom Volume (CV)

1. Select the Volume Definition screen (LUN Expansion Option → VLL tab).
2. Volumes can be created out of free space by selecting a BOX then a Parity Group. Right click on the free space and select *Install CV*.



3. A dialog will appear allowing you to define the size of the volumes.



4. Enter size (in MB) of the volumes and quantity. For example, for 34.69GB volumes, enter 35522. Once the size is entered, the dialog will calculate how many of the volumes will fit evenly into the freespace. Then click *Set*.
5. The dialog will then display the volumes to be created on the left hand side. You may create additional volumes or click *Next* and go to next step.

Storage Navigator

Install CV

Parity Group: 3-8-(1)

How to Create Volumes

No.	Emulation	Capacity
1	OPEN-V	35522 MB
2	OPEN-V	35522 MB
3	OPEN-V	35522 MB
4	OPEN-V	35522 MB
5	OPEN-V	35522 MB
6	OPEN-V	35522 MB
7	OPEN-V	35522 MB
8	OPEN-V	35522 MB
9	OPEN-V	35522 MB
10	OPEN-V	35522 MB
11	OPEN-V	35522 MB
12	OPEN-V	35522 MB
13	OPEN-V	35522 MB
14	OPEN-V	35522 MB
15	OPEN-V	35522 MB
16	OPEN-V	35522 MB
17	OPEN-V	35522 MB
18	OPEN-V	35522 MB

available LDEVs : 1994

Emulation: OPEN-V

Capacity Unit: MB

Setting Mode: ☒ Specify capacity and number
☐ Divide free space by number
☐ Divide free space by capacity
☐ Set remaining space as volume

Capacity: 35522 MB (46 - 5567)

No.: 54 (0)

Buttons: Set, Delete, Clear, Next, Cancel

In order to have consistent configurations worldwide, use the CU (Logical Device) Numbering scheme as defined in the Best Practices.

- Assign a CU number and select an address for the volume(s). There are up to 255 addresses per CU number. You will need to select the volume on the upper list and then click on an unused volume address (non-grey box). The interval function is used to disperse (or skip addresses) between the physically sequential volumes. To use this function, all of the volumes need to be selected from the upper list and then click on the first (unused) address box. The interval value is the number of RAID (parity) groups that will have designated addresses within the CU address range. This is recommended for creating multiple volumes to prevent assigning physically sequential volumes to the same host.

Storage Navigator

Install CV

Parity Group: 3-8-(1)

Select LDEV No.: Please click the following cell

Select LDKC No.: 00

Select CU No.: 01

Interval: 0

No.	Emulation	Capacity	LDKC:CU:LDEV
34	OPEN-V	35522 MB	00:01:80
35	OPEN-V	35522 MB	00:01:8C
36	OPEN-V	35522 MB	00:01:90
37	OPEN-V	35522 MB	00:01:8E
38	OPEN-V	35522 MB	00:01:8F
39	OPEN-V	35522 MB	00:01:90
40	OPEN-V	35522 MB	00:01:91
41	OPEN-V	35522 MB	00:01:92
42	OPEN-V	35522 MB	00:01:93
43	OPEN-V	35522 MB	00:01:94
44	OPEN-V	35522 MB	00:01:95
45	OPEN-V	35522 MB	00:01:96
46	OPEN-V	35522 MB	00:01:97
47	OPEN-V	35522 MB	00:01:98
48	OPEN-V	35522 MB	00:01:99
49	OPEN-V	35522 MB	00:01:9A
50	OPEN-V	35522 MB	00:01:9B
51	OPEN-V	35522 MB	00:01:9C
52	OPEN-V	35522 MB	00:01:9D
53	OPEN-V	35522 MB	00:01:9E
54	OPEN-V	35522 MB	00:01:9F

Buttons: Back, Next, Cancel

- If the wrong address was selected, right click the volume in the volume listing and select **Back**. This will remove the address selection and goes to the previous screen. Click the **Next** button to continue.
- The System Disk Setting is not used at this time. It is used for creating/releasing special devices. Select **Next** to continue.

Storage Navigator

Install CV

Parity Group: 3-8-(1)

System Disk Setting

No.	LDKC:CU:LDEV	Emulation	Capacity	Option
1	00:01:6A	OPEN-V	35522 MB	-
2	00:01:6B	OPEN-V	35522 MB	-
3	00:01:6C	OPEN-V	35522 MB	-
4	00:01:6D	OPEN-V	35522 MB	-
5	00:01:6E	OPEN-V	35522 MB	-
6	00:01:6F	OPEN-V	35522 MB	-
7	00:01:70	OPEN-V	35522 MB	-
8	00:01:71	OPEN-V	35522 MB	-
9	00:01:72	OPEN-V	35522 MB	-
10	00:01:73	OPEN-V	35522 MB	-
11	00:01:74	OPEN-V	35522 MB	-
12	00:01:75	OPEN-V	35522 MB	-
13	00:01:76	OPEN-V	35522 MB	-

☐ Hide this dialog box
(If you select this option, this dialog box will not reappear until you switch from the VLL window to another window.)

Set Release

Back Next Cancel

9. You may also get an SSID setting screen. This is used for mainframe and should be set as the default value.
10. A confirmation dialog screen will appear. Select OK to confirm the volume addressing to be created.

Storage Navigator


Install CV

Do you want to set the following custom volume?(3010 9010)

VDEV	Emulation	Capacity	LDKC:CU:LDEV	Option
3-8-(1)	OPEN-V	35522 MB	00:01:6A	-
3-8-(1)	OPEN-V	35522 MB	00:01:6B	-
3-8-(1)	OPEN-V	35522 MB	00:01:6C	-
3-8-(1)	OPEN-V	35522 MB	00:01:6D	-
3-8-(1)	OPEN-V	35522 MB	00:01:6E	-
3-8-(1)	OPEN-V	35522 MB	00:01:6F	-
3-8-(1)	OPEN-V	35522 MB	00:01:70	-
3-8-(1)	OPEN-V	35522 MB	00:01:71	-
3-8-(1)	OPEN-V	35522 MB	00:01:72	-
3-8-(1)	OPEN-V	35522 MB	00:01:73	-
3-8-(1)	OPEN-V	35522 MB	00:01:74	-
3-8-(1)	OPEN-V	35522 MB	00:01:75	-
3-8-(1)	OPEN-V	35522 MB	00:01:76	-
3-8-(1)	OPEN-V	35522 MB	00:01:77	-
3-8-(1)	OPEN-V	35522 MB	00:01:78	-
3-8-(1)	OPEN-V	35522 MB	00:01:79	-
3-8-(1)	OPEN-V	35522 MB	00:01:7A	-
3-8-(1)	OPEN-V	35522 MB	00:01:7B	-
3-8-(1)	OPEN-V	35522 MB	00:01:7C	-
3-8-(1)	OPEN-V	35522 MB	00:01:7D	-

Back OK Cancel

11. Review of volumes to be created. The bold blue (in italics) volumes and resulting free space are to be created.




12. If all volumes to be added have been addressed, then click the APPLY button to start the disk formatting process. This process will take a few minutes, depending on the number of volumes to be created and the size of the parity group(s).

15.16.2 Make Volume

6. You cannot delete the last LUN on a Parity Group. If you want to create new volumes out of the last volume, then Right Click on existing volume out of the box listings to select Make Volume (**WARNING: this will destroy volume**). Whole Parity Group Recreate (Make Volume) creates new LDEVs out of the parity group without having to delete the existing LDEVs first. This step is only skipped from the user perspective. All existing LDEVs will be destroyed and replaced with the new LDEVs. Within the VLL tab of the LUN Expansion option, select the appropriate (Box then) Parity Group. Right-click on the parity group to select *Make Volume* from the pop-up menu. You will get a screen very similar to the Install CV screen that allows you to create LDEVs and assign addresses within the party group.

Warning: This will destroy any existing volumes on the parity group.




Follow the steps under [Install Custom Volume](#) (starting from step 3) to complete the Make Volume process. Once the LDEVs have been created logically, the new settings appear on the panel in blue italics but are not yet implemented. If you want to make the changes to the subsystem, then select *Apply* on the VLL tab and then select *OK* on the confirmation message. If you want to cancel the changes to the subsystem, select *Cancel* on the VLL tab and then select *OK* on the confirmation message.

15.17 Volume (LDEV) Deletion


There are two methods of deleting LDEVs from a parity group: Volume to Space and Make Volume.

15.17.1 Volume to Space

Volume to Space creates free space from selected LDEVs. The new and any existing freespace will not be coalesced with any other freespace unless they reside sequentially within the parity group. Within the VLL tab of the LUN Expansion option, select the appropriate (Box then) Parity Group and the appropriate LDEVs to be deleted. While in modify mode, right-click on one of the selected LDEVs and select Volume to Space from the pop-up menu.



Do not select the freespace entry as an LDEV to be deleted, or you will receive an error message that not all of the paths have been deleted. A confirmation message will appear, showing the LDEVs to be deleted. Click OK to continue with the operation.



The new settings appear on the panel in blue italics but are not yet implemented. If you want to make the changes to the subsystem, then select *Apply* on the VLL tab and then select *OK* on the confirmation message. If you want to cancel the changes to the subsystem, select *Cancel* on the VLL tab and then select *OK* on the confirmation message.

15.17.2 Make Volume

See [Make Volume](#) under Volume (LDEV) Creation / Definition.


15.17.3 Initialize Volume

This method is not recommended as it creates a single volume for a parity group. This method is typically used by a CE who is adding storage, but not setting up any useful volumes.


15.18 Special Display of RAID 10 (4+4) Parity Groups

When creating LDEVs (see previous topic) from RAID 10 (4+4) Parity Groups, it is important to know that they may appear in the configuration as two concatenated RAID 10 (2+2) Parity Groups. HDS states that is only a display issue – the R10 is fully striped over all of the disk members. Half of the total Parity Group storage (approximately 136GB on a 73GB R10 4+4) will be under each of the parity groups. These Parity Groups will be displayed with a green icon

instead of the normal blue icon . To show the relationship between the parity groups, right click the parity group to display the Concatenation List.



This will display the members of this group.




When assigning LDEVs to Host Groups (see next section), LDEVs from 4+4 parity groups will also list the concatenation list as an option on a right-click of the LDEV list.

15.19 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. To create a LUSE volume, select the LUN Expansion/VLL screen from the main Storage Navigator screen. The LUSE tab is the default view.



The LDEV Information outline is on the upper left of the LUN Expansion tab and displays an outline view of the CU and LDEV numbers installed on the subsystem (for example, CU 01 has a number of LDEVs, including 01:00).




The format of the LDEV number changes if the LDEV is an expanded (LUSE) volume (for example, if LDEV volume 01:00 consists of five concatenated LDEVs, its number is 01:00*5). The top LDEV in the LUSE volume defines the name of the LDEV.

The LDEV information outline will display some or all of the following icons:

- Normal Variable-sized volume (CV)
- Expanded (LUSE) volume

The LDEV detail information is displayed in the LDEV Operational table on the lower half of the LUN Expansion tab.



If an existing LUSE extends over two or more parity groups, the PG column displays the smaller parity group number.

Numbering LUSE Volumes

The head of the LUSE becomes LUN number. All of the LUSE components must be in the same CU. So if you had a LUSE volume that consisted of 00:10 00:12 and 00:1F, the LUSE LUN would be 00:10.

Creating LUSE Volumes

Warning: Except for systems running Windows 2k3/2k or AIX, creating and releasing LUSE volumes is a destructive process. Be sure to back up your data before proceeding.

1. While in Modify mode, select a CU number (address range) from the LDEV outline. The LDEV Detail table shows all LDEVs in the selected CU.

- In the LDEV detail table, select and then right-click the LDEVs that you want to include in the LUSE volume. This will display the Set LUSE Volume pop-up menu.

LDEV	Emulation	Capacity	RAID	PG	Paths
00:02	OPEN-V CVS	3.05 GB	1(2D+2D)	1 - 1	
00:03	OPEN-V CVS	3.05 GB	1(2D+2D)	1 - 1	
00:04	OPEN-V CVS	3.05 GB	1(2D+2D)	1 - 1	
00:05	OPEN-V CVS	3.05 GB	1(2D+2D)	1 - 1	
00:06	OPEN-V CVS	3.05 GB	1(2D+2D)	1 - 1	1
00:07	OPEN-V CVS			1 - 1	
00:08	OPEN-V			1 - 1	
00:09	OPEN-V			1 - 1	1
00:0A	OPEN-V			1 - 1	

Set LUSE Volume
Release LUSE Volume
Concatenation List

Alternatively, you could select a CU number (address range) and then select the arrow button in the Select an LDEV drop down box (on the lower right of the panel) then select the first LDEV for the LUSE from the Free LDEVs list. Select one or more LDEVs for the LUSE volume, then select the <<Add button to move the selected LDEVs to the Expanded LDEVs list. Deleting an LDEV can be accomplished with a selected LDEV and then selecting >>Delete. Select the Set button to continue.

Expanded LDEVs

00:00 [Top]

Selected LDEVs: 0

<< Add

Delete >>


Set

Free LDEVs


LDEV	RAID	PG
00:13	1(2D+2D)	1 - 1
00:17	1(2D+2D)	1 - 1
00:18	1(2D+2D)	1 - 1
00:1A	1(2D+2D)	1 - 1
00:1B	1(2D+2D)	1 - 1
00:1C	1(2D+2D)	1 - 1
00:1D	1(2D+2D)	1 - 1
00:1E	1(2D+2D)	1 - 1
00:22	1(2D+2D)	1 - 1

Selected LDEVs: 1


- If you want to use a path-defined LDEV as the top volume in the LUSE volume see the LUSE restrictions with existing paths on LDEVs (next sub-section).
- When you start a LUSE operation that would combine a path-defined LDEV or LUSE volume with another path-defined LDEV or LUSE volume, a panel will appear with a message asking if you want to continue the operation. Click the OK button if you want to continue or the Cancel button if you do not.



- The set LUSE Confirmation panel displays. Select OK to continue, or select Cancel to cancel the operation.



6. If the top LDEV in the LUSE volume has no paths to it, the Data Initialize Warning panel displays. Select OK to continue, or select Cancel to cancel the operation.



7. The new settings appear on the panel in blue italics but are not yet implemented. If you want to make the changes to the subsystem, then select *Apply* on the LUN Expansion tab and then select *OK* on the confirmation message.
8. If you want to cancel the changes to the subsystem, select *Cancel* on the LUN Expansion tab and then select *OK* on the confirmation message.

LUSE Restrictions with Existing Paths on LDEVs

When you create a LUSE volume, the top LUSE volume can be either an LDEV or a LUSE volume that has one or more paths defined to it. Creating LUSE volumes using a path-defined LDEV has the following restrictions:

- Only the top volume in the LUSE volume to be created can have any paths. The other volumes in the LUSE volume to be created must not have any paths. The maximum number of paths to that top LDEV is the maximum group number (128) multiplied by the maximum number of ports (64). There is one path per group.
- The path(s) to the top volume must be defined from either a Windows® 2000 host, a Windows® 2003 host, or an AIX® 5.2 host.

Caution: If you want to perform a LUSE operation that includes a Windows® LDEV that has a defined path, the host mode must be 2C (Windows® Extension). If it is not 2C, you must change it before the LUSE operation.

Resetting an Unregistered LUSE Volume

After you have made the settings for a LUSE volume, but before you select Apply to implement those settings, the prospective LUSE volume is shown on the LUN Expansion tab in blue bold italics. At the point in the process you can opt to reset (right-click menu) the LUSE volume and release the individual LDEVs by using the reset LUSE volume function. This is different from merely canceling all of the stored actions, because you can choose to reset some but not all of the unregistered LUSE volumes.


Releasing a LUSE Volume

Warning: Except for systems running Windows 2k3/2k or AIX, creating and releasing LUSE volumes is a destructive process. Be sure to back up your data before proceeding.

1. Change to modify mode and remove all paths to the LUSE volume that is to be released (see [Removing LUNs from Host Groups \(Deleting Paths\)](#))
2. Select the CU number (address range) from the LDEV Information outline. The LDEV Detail table (upper right-hand corner of the LUN Expansion tab) shows all of the LDEVs in the selected CU.
3. Select on or more LUSE volumes in the LDEV Detail table. Right-click the LUSE volume to select *Release LUSE Volume* on the pop-up menu.

LDEV	Emulation	Capacity	RAID	PG	Paths
00:00	OPEN-3	0.04 GB	1(2D+2D)	1 - 1	
00:01	OPEN-9	0.05 GB	1(2D+2D)	1 - 1	
00:02	OPEN-3 *3	6.87 GB	1(2D+2D)	1 - 1	
00:05	OPEN-3 *3			1 - 1	
00:06	OPEN-3			1 - 1	1
00:10	OPEN-E			1 - 1	
00:11	OPEN-E			1 - 1	
00:12	OPEN-E	0.04 GB	1(2D+2D)	1 - 1	
00:13	OPEN-3	0.04 GB	1(2D+2D)	1 - 1	

4. The Release LUSE Volume Confirmation panel will display.



- Verify the LUSE volumes listed are the correct volumes to be released then click OK. The new settings will appear in blue italics, but are not yet implemented.
5. If you want to make the changes to the subsystem, then select *Apply* on the LUN Expansion tab and then select *OK* on the confirmation message.
 6. If you want to cancel the changes to the subsystem, select *Cancel* on the LUN Expansion tab and then select *OK* on the confirmation message.


Changing LUSE Capacities

To expand the capacity of a LUSE volume, either select a LUSE volume that you want to expand, and then add LDEVs or LUSE volume(s), or first select LDEVs or LUSE volume(s) that you want to add, and then select a LUSE volume to be expanded.


You may not reduce the capacity of an existing LUSE volume. If you want a LUSE volume to have a smaller capacity, you must first release the LUSE volume (refer to previous section) and then re-define the LUSE volume.

15.20 Assignment LDEVs to Host Groups


1. Select LUN Manager Option.
2. Select a port and then right click to select enable LUN security (if not already enabled).




3. Confirm that you are turning on LUN security.
4. Repeat for all ports listed.
5. When all ports are highlighted in blue, then hit the Apply button.
6. All ports should now indicate that LUN Security is enabled.
7. Select Host Group 00 under the port that needs to be configured for host access. If Host Group 00 is already in use for another host, then select Add Host Group.
8. Right click the group and select Change Host Group




9. Assign the group name to the hostname being added to the port (max 16 characters).
10. Select the Appropriate Host Mode. For a list of the valid supported host modes, see [Storage Host Setup \(Host Modes\)](#).




For HP-UX, Host mode 03 should be used. In addition, to prevent ghost LUNs from appearing Host Mode Option 12 should be used (select the *Option* button).



11. Click OK.
12. Select LDEVs to be added to this port either by drag/drop from CU/LUN list or by highlighting and clicking on the Add LU Path button. The Add LU Path function associates the LDEV with the highlighted HOST LUN number (typically 0000 as the first LUN). Selecting a different CU number from the CU drop down list will display LDEVs for that CU address range.






Note: The LDEVs listed may already be allocated to other servers. All LDEVs are visible in the list (grouped by CU). Check the Path column to see if it already has an access path defined. LDEVs that have been already assigned will have a value in the path column. A value of two paths is normal. Four paths is usually for clusters.


Use the following key to differentiate the different status of LDEVs:

Symbol	Meaning
	Indicates a logical volume to which an LU path is defined.
	Indicates an expanded LU.
	Indicates a command device (required for CCI)
	Indicates a secured command device.
	Indicates a LUN to which no LU path is defined.
	Indicates a volume that has not been allocated.
	Indicates a volume to which Virtual LVI/LUN (CVS) has been applied (allocated).

13. You can create the second path (or any other path for this or other hosts) before hitting *Apply*. This is accomplished by right clicking the storage group name and selecting *Copy*. If you copy the storage group 00 to another port, it will not replace the storage group 00 on the second port. Once you have enough host groups beyond storage group 00, then this method is preferred. There is a maximum of 1024 host groups per port.




14. Alternatively, the LUN Paths can be copied to a group already created on another port. This second method is recommended for the first hosts on each port (using host group 00) or for adding LUNs to existing host groups. To copy the LUN Paths, right click the selected LUNs at the source location, and select copy Paths. Then go to the target port and right click to select paste Paths.



15. Verify that all LUN Paths are on all of the desired ports with the same Host Group number. The same numbering and naming of the host groups is to reduce complexity in the configuration.


16. Add the WWN of the Host HBA to each of the appropriate host groups. The array should give a list of connected HBAs. If the HBA is not listed, then it may have to be entered manually. Enter a description with the following format: "hostname_fabric "




17. Hit the *Apply* button to make all of the changes effective.

15.21 Changing the HBAs on Host Groups (HBA Swap)

1. While in Modify mode, select and expand the host group that is of interest. For most hosts there will be two host groups, one for the A-fabric connection and another one for the B-fabric connection.
2. Right-click the HBA listed for the host group and select Change WWN & Name.



1. Update the HBA WWN and the name (if appropriate). Select OK when modifications are complete.




Note that selecting the “*Apply to change to other ports, too*” option will update other host groups. However, most hosts will have only a single host group for a given HBA.

2. Hit the *Apply* button to make all of the changes effective.

15.22 Removing LUNs from Host Groups (Deleting Paths)

1. While in Modify mode, select and expand the host group that is of interest.
2. In the Host Group detail box, select the LUNs that will be released. Multiple LUNs can be selected individually (using CTRL key) or with a range (using SHFT key).
3. Once the LUNs have been selected, right-click one of the LUNs and select Release LU Path.




4. Since there are usually at least two host groups for every LUN for a given host, this should be repeated for the same LUNs on the other host group(s).


5. If you want to make the changes to the subsystem, then select *Apply* on the LUN Manager tab and then select *OK* on the confirmation message.
6. If you want to cancel the changes to the subsystem, select *Cancel* on the LUN Manager tab and then select *OK* on the confirmation message.

15.23 Setting Up Command Devices (For ShadowImage CCI)

1. Select an existing LUN/device from the LU Path & Security Menu




2. Right Click the device and select *Convert To Command Device*
3. A warning message will appear.



4. Click Yes and then hit Apply to make the changes effective.
5. It is not recommended that different hosts share the same command device due to possible sharing violations.
6. 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.

7. **WARNING:** Never send standard i/o to the device once it has been converted to a command device (i.e. format disk, write i/o, etc).



15.24 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. Volume shredder is available for both open and mainframe systems. Its API can be called by other functions such as Tiered Storage Manager CLI. For more information on TSM CLI see [Appendix L: HDS Tiered Storage Manager](#).

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 (pool-VOLs)
- Journal volumes
- Virtual volumes (V-VOLs) – This refers to Copy-on-Write Snapshot volumes not externalized volumes.
- Volumes containing extents

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)

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

Blocking I/O Activity in the Target Volume

1. Change to Modify mode. Select the LUN Expansion (LUSE)/Virtual LVI/LUN (VLL) button to open the LUSE/VLL panel. The LUN Expansion tab is the default view. Select the VLL tab.




2. You will first need to block I/O activity in the target volume.


- (a) In the LDEV Information Table select the LDEV Status button
- (b) Select a target volume(s) to be shredded.
- (c) Right click the target volume to display the Blockade popup menu, then select Blockade.




The target volume information will be displayed in blue italic letters.



- (d) Select Apply, then select OK on the confirmation message.




The status column will reflect the change from Normal to Blocked. To see the status column, the windows will need to be scrolled to the right.



(e) If you want to reverse the blocking process (unblock), right click on the volume to display the Restore pop-up menu select Restore, select Apply, then select OK on the confirmation message

Defining the Shredding Settings

1. Using the default shredding parameters will write three passes: 00, FF, then 00. If you want to use these default settings, then skip to Shredding the Data in the Target Volume.
2. Select the *Shredding Param* button on the VLL main panel to open the Shredding Param. Panel.



3. In the Write Data Type section, do one of the following:
 - (a) If you want to define the shredding settings using the default settings, select Default and go to step 4.

Shredding Param.

Write Data Type
☒ Default (00-FF-00)
☐ User

Write Data
☒ Define (Max: 4 hex digits)
☐ Random

Format List

NO.	Data
1	00
2	FF
3	00

Buttons: Set, Delete, OK, Cancel

(b) If you want to define the shredding settings yourself, select *User*. Otherwise, using *Default* will write three passes: 00, FF, then 00. If you want to use the default settings, then skip to Shredding the Data in the Target Volume.


4. In the *Write Data* section, do one of the following:

(a) If you want to select the overwrite data randomly:

- Select *Random*.
- Select *Set*. A row will be added to the *Format List*.
- Continue adding rows by selecting *Random* and then *Set* until there are at least three rows in the *Format List*, including the *Last Data* row. (This is important to ensure that the data is completely erased.)
- To delete a row, select that row and then select *Delete*. You cannot delete the *Last Data* row.

(b) If you want to define the settings:

- Select *Define*.
- Enter the dummy data in the text box. This text box is enabled only if you have previously selected *User* as the *Write Data Type*.
- Select *Set*, and a row will be added to the *Format List*. To delete a row, select that row and then select *Delete*. Do not delete the *Last Data* row.
- Continue adding rows by selecting dummy data and then selecting *Set* until there are at least three rows in the *Format List*, including the *Last Data* row (This is important to ensure that the data is completely erased).
- To delete a row, select that row and then select *Delete*. You cannot delete the *Last Data* row.



Shredding Param.

Write Data Type

☐ Default (00-FF-00)

☒ User

Write Data

☒ Define (Max. 16 hex digits)

☐ Random

Format List

NO.	Data
1	55
2	55
Last Data	00

Set

Delete

OK **Cancel**

5. Select *OK* to complete the action, or select *Cancel* to cancel the action. The Shredding Param. panel closes, and you are returned to the VLL main panel.


Shredding the Data in the Target Volume

1. Select and right-click the target volume(s) in the LDEV Information table to display the Shredding popup menu.



2. Select *Shredding*, and the target volume will be displayed in blue italic letters.

3. Select *Apply*. The Shredding Data Output panel displays.



Shredding data output

non-object Volume

CULDEV

object Volume

CULDEV
00 B6

Add


Delete

OK

Volumes in the non-Object volume list will *not* have their shredding data saved to a file. Volumes in the Object volume list will have their shredding data saved to a file. The volume selected for shredding will be in the non-Object volume list.

- If you want to add a volume to Object volume list, select the volume(s) in the non-Object volume list and then select the *Add* button. You can add up to three volumes in this list.
- If you want to remove a volume from the Object volume list, select the volume, then select the *Delete* button.


4. Select *OK* and then *Apply* changes on the next dialog. A confirmation message appears. Select *OK* again, and you are returned to the VLL main panel and shredding operation begins. During the shredding operation, a progress panel displays.




Caution: If you select Cancel during the shredding operation, you cannot restart the operation, and the status of the data in the volume is unknown.

On an inactive array, three 34GB volumes from the same 146GB disk RAID 5 (7+1) parity group took approximately 39 minutes to complete the default 3-pass pattern.

5. When the shredding operation is finished, the Shredding Complete confirmation panel displays.




If the process ends in an error, an error message is displayed.



6. If you want to download the shredding results* immediately after the shredding operation, do the following:

- (a) Select *Download* on the Shredding Complete Confirmation panel to display the download dialogue box.
- (b) Select the destination folder, and then select *OK*.



7. If you do not want to download the shredding results* immediately after the operation, select *OK* to close the confirmation panel and return to the VLL main panel. The volume status will revert to Normal.

8. To download the results* of completed shredding operations, select the *Download* button in the VLL main panel to display the *Download* dialogue box. Specify the file destination, then select *Save*.

*Note: The term "results" is used loosely here. There are 3 binary files generated as "results" (one for each pass). If the contents of bin files using a hex editor are all zero's (0x00 hex), it indicates the shredding operation was successful. If not all zero's, it indicates that the shredding operation ended with error. A more comprehensive and useful results file has been requested from HDS.

15.25 System Status

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. This is especially true when an HDS/Sun CE has control (modify mode) of the array via the SVP Application.

Ref Code	Error Level	Status	Date
BF4808	Moderate	Not Complete	2005/03/30 04:00:47
BF480E	Moderate	Not Complete	2005/03/30 04:00:47
BF4806	Moderate	Not Complete	2005/03/30 04:00:47
BF480C	Moderate	Not Complete	2005/03/30 04:00:47
BF4804	Moderate	Not Complete	2005/03/30 04:00:47
BF480A	Moderate	Not Complete	2005/03/30 04:00:47
BF4802	Moderate	Not Complete	2005/03/30 04:00:47
BF4800	Moderate	Not Complete	2005/03/30 04:00:46
BF60D1	Moderate	Not Complete	2005/01/27 13:00:20
BF4A23	Moderate	Not Complete	2005/01/27 04:00:40
BF5A26	Moderate	Not Complete	2005/01/27 04:00:40
BF5A22	Moderate	Not Complete	2005/01/27 04:00:40
BF4A21	Moderate	Not Complete	2005/01/27 04:00:39
BF4A23	Moderate	Complete	2004/11/20 03:59:48
BF4A21	Moderate	Complete	2004/11/20 03:59:47
BF60D1	Moderate	Complete	2004/11/24 13:00:12
BF4A23	Moderate	Complete	2004/11/24 04:00:32
BF4A21	Moderate	Complete	2004/11/24 04:00:32
BF4A23	Moderate	Complete	2004/11/20 04:00:33
BF4A21	Moderate	Complete	2004/11/20 04:00:33
BF60D1	Moderate	Complete	2004/11/19 05:55:47
BF4A23	Moderate	Complete	2004/11/19 04:00:23
BF4A21	Moderate	Complete	2004/11/19 04:00:23
452000	Service	Complete	2004/11/19 03:58:26
451000	Service	Complete	2004/11/19 03:53:24
EF1000	Serious	Complete	2004/11/19 03:47:16
463000	Moderate	Complete	2004/11/19 03:47:16
461000	Service	Complete	2004/11/19 03:47:14
460004	Moderate	Complete	2004/11/19 03:47:13

- 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

For additional information, an event can be selected and a right-click will reveal a detail option. Select *Detail* and a dialog screen will appear.

Detail	
Ref.Code	BF4A23
Error Level	Moderate
Status	Not Complete
Date	2005/01/27 04:00:40
Error Section	ENVIRONMENTAL ERROR
Error Detail	DKUPS WARNING

15.26 License Keys

To enter new license keys or to view exiting license keys, select the *License Key Option* from the main menu (USPV: Go → Environmental Settings → License Keys)

Storage Navigator - IP:10.40.53.25 -

File Go Help 24 minute(s) remaining in session. Reset User ID: gj06529

License Key SNMP Information Partition Definition License Key Partition Definition

License Key

Product Name	Key Type	Permitted Volumes	Term(days)	Status
Cache Residency Manager	Permanent	Unlimited		Installed
Cache Residency Manager for z/OS(R)	Permanent	Unlimited		Installed
Compatible PAV	Permanent	Unlimited		Installed
Data Retention Utility	Permanent	Unlimited		Installed
Volume Retention Manager	Permanent	Unlimited		Installed
Database Validator	Permanent	Unlimited		Installed
Dynamic Provisioning	Permanent	Unlimited		Installed
Open Volume Management	Permanent	Unlimited		Installed
LUN Manager	Permanent	Unlimited		Installed

License Install

Mode: ☒ KeyCode ☐ File ☐ Available

KeyCode:

FileName: Browse

License Status

☒ Enable ☐ Disable Set


53.03 TB (MF 0.00 TB)

Total Storage


Install Uninstall Apply Cancel

To edit a license key entry (adding new license keys), select the product name on the drop down list. For a key-code entry (select the KeyCode mode if not already selected) enter the value of the key code into the KeyCode field and select *Install*. For a file-based entry, select the File mode, use the Refer button to open a file browser to find the license key file (xxx.plk). File-based entry is recommended.


After the entry is made or the file is selected, a confirmation screen will appear:



If an incorrect entry is made, then an error screen will appear:




If the license key is verified, then the license will be highited in blue italics.



Press the *Apply* button to make the changes effective.

15.27 Configuration File Loader (Large-Scale Changes)


Configuration File Loader is used to make large-scale or “batch” changes to the LUN Manager configuration of the TagmaStore USP/USPV subsystem, using either a text editor or spreadsheet software to edit configuration definition files (spreadsheet files) and apply the changes. You can apply the same configuration definition file to different subsystems, and you can define the configuration when you are offline. To import or export the configuration into a single file, select the *Configuration File Loader* option from the main menu (USPV: Go → Configuration File Loader).



If no files have been recently uploaded, then this screen will be blank. It normally contains the errors encountered during an upload.


There is Select Sheet option on the version for the USPV. Use *LUN Manager* only at this time. The *Volume Retention* option is not supported at this time.

To export a configuration select the *Export* button and the following dialog will appear:



The whole filename can be entered (must end in .spd) or the *Reference* button can be selected to browse for a specific location. To export the configuration after the filename has been specified, select the *Export* button on this screen. Select *Close* to exit the screen.

To import a configuration select the *Import* button and the following dialog will appear:



15.28 Configuring Performance Data Collection

The array should be configured to collect performance data on a frequent basis. This data can be used to troubleshoot specific problems or to analyze overall performance of the array over periods of time. HDS states that the collection of this data should not impact the array's performance (< 0.1%).

To turn on performance data collection, go to the Performance Management Options Screen (Go → Performance Management → Monitoring Options). While in modify mode, select enable for current status and change the gathering interval to 5 Minutes. You will then need to select every possible address space on the array so that the entire array information is stored. Select Apply when your changes are complete.

Storage Navigator - IP:10.40.53.25 -

Storage Navigator

File Go Help 30 minute(s) remaining in session. [Reset] [Icons] User ID : gj06529

Performance Management TC Monitor TCz Monitor UR Monitor URz Monitor Monitoring Options

Monitoring Options

Monitoring Switch

Current Status: ☒ Enable ☐ Disable Gathering Interval: 5 Minute(s)

Monitoring Target

PG	LDKC#	CU	+0	+1	+2	+3	+4	+5	+6	+7	+8	+9	+A	+B	+C	+D	+E	+F		
	00	00	S	S	S	S	N	N	N	N	S	N	S	S	N	N	N	N	N	Select Release
		10	N	N	N	N	N	N	N	N	N	N	N	N	N	N	N	N	N	
		20	N	N	N	N	N	N	N	N	N	N	N	N	N	N	N	N	N	
		30	N	N	N	N	N	N	N	N	N	N	N	N	N	S	N	S	N	
		40	N	N	N	N	N	N	N	N	N	N	N	N	N	N	N	N	N	
		50	N	N	N	N	N	N	N	N	N	N	N	N	N	N	N	N	N	
		60	N	N	N	N	N	N	N	N	N	N	N	N	N	N	N	N	N	
		70	N	N	N	N	N	N	N	N	N	N	N	N	N	N	N	N	N	
		80	N	N	N	N	N	N	N	N	N	N	N	N	N	N	N	N	N	
		90	N	N	N	N	N	N	N	N	N	N	N	N	N	N	N	N	N	
		A0	N	N	N	N	N	N	N	N	N	N	N	N	N	N	N	N	N	
		B0	N	N	N	N	N	N	N	N	N	N	N	N	N	N	N	N	N	
		C0	N	N	N	N	N	N	N	N	N	N	N	N	N	N	N	N	N	
		D0	N	N	N	N	N	N	N	N	N	N	N	N	N	N	N	N	N	
		E0	N	N	N	N	N	N	N	N	N	N	N	N	N	N	N	N	N	
		F0	N	N	N	N	N	N	N	N	N	N	N	N	N	N	N	S	N	
Monitoring Target CUs : 255																				

PG Selection Support [Apply] [Cancel]

PP Status : Normal

The whole filename can be entered (must end in .spd) or the *Reference* button can be selected to browse for a specific location. To import the configuration after the filename has been specified, select the *Import* button on this screen. Select *Close* to exit the screen. Any errors or status associated with the import will be displayed on the Configuration File Loader screen.



Warning: This will overwrite the existing configuration, which may corrupt / disable a portion or the entire array. Do not upload a file without being absolutely sure about the changes being made.

16 Appendix C: HiCommand Device Manager Navigation and Usage


16.1 Version Requirements

The USPV array requires the use of HiCommand Device Manager to be version 5.7 or later. The USP array does not have a version restriction.

16.2 Logging Into HDvM Server

The web page for accessing HDS Device Manager is:


`http://<HDMserver>:23015/DeviceManager/`




After the server has been configured for SSL (see [Turning on SSL for Device Manager \(Windows\)](#)), the following address should be used:

`https://<HDMserver>:23016/DeviceManager/`


When the initial screen comes up, it has a button for the login screen. Select the *Go To Login Screen* button if the login screen does not appear automatically.



Before any login can occur for the first time, the license key must be entered. The license key can be accessed by selecting the *License* button. A key can be entered manually or a license file (for example “HDS9900V SN10192 20041025 225929.plk”) may be used. Using a file is recommended. Select *Save* after one of these methods is used. A valid license key will update the screen with the serial number of an array and the raw storage limit of the license. Multiple files may be needed for multiple arrays if the capacity is exceeded. The license key for previous versions of the software work with later versions.




Once the license key has been validated, select *Close* to close the screen and go back to the login screen. The default username/password is system/manager. This should be changed as soon as possible. The following is the main screen visible once the license key has been entered:



16.3 Users and Permissions

Select the *Users and Permissions* link on the Object Tree under Administration. It will display Users and Permissions.




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 usersids.

For information regarding special user accounts to be created, see [Special User Accounts](#).

16.3.1 Users

Select the Users list the current users. Click on an existing user to list the properties of that user. You are able to Add, Modify, Delete, Lock, Unlock, and Change Permissions of users.



Modify User

To modify an existing user once the user has been selected, click the *Edit Profile* button. To change the password for a user, click the *Change Password* button.

Add User

To add a user, go to the main user screen (Administration → User and Permissions → Users). Select the Add User button.

The screenshot shows the 'Add User' form in the 'HiCommand Device Manager' web application. The form is titled 'Add User' and includes a 'User Profile' section with the instruction 'Please enter a value into each field.' The fields are:

- *User ID:
- *Password:
- *Verify Password:
- Full Name:
- E-mail:
- Description:

At the bottom of the form, there is a legend: '*Required field'. The form has 'OK' and 'Cancel' buttons at the bottom right.


Fill out the appropriate fields and click Ok when done. The password field must meet the password content requirements as defined in the password policy. For more information on the password policy, see [Password Policy](#). After a user has been entered it is necessary to assign them the appropriate resources (see [Resource Groups](#)).

Delete User

To delete a user, select the user from the user list (select checkbox next to user name) then click on the *Delete Users* button. Alternatively, you could have the user detail information displayed and select the *Delete User* button. You will be asked to confirm the deletion.

Lock User


You can lock out a user from having access to Device Manager. A locked user cannot log in (invalid user id or password error message). If a user logs in and is then locked, the user's session becomes invalid. To lock a user or more than one user, select the user (select checkbox next to user) then click the *Lock Users* button. You will get a confirmation screen:



The status of the user will change to Locked (from Active) and the user icon  will change to . To unlock a user, select the user (checkbox next to username) and then click on the *Unlock Users* button.

16.3.2 Permissions

In order for the user to log into the system and have access to certain areas, the user's permissions must be established. Otherwise, a user will not be able to login (invalid userid or password error message). To assign a user's permissions, select the user from the user list. The user information detail will be displayed. To modify the permissions, click the *Change Permission* button. A dialog screen will be displayed that will allow you to check each permission for that user:



The following are each available permission:

Application

Displays the application names. For each application, you can select a check box to grant permission to the user. Select the All Applications check box to grant permissions for all applications. For example, HDvM is Device Manager and HTSM is Tiered Storage Manager.

Admin

A user with Admin permission can manage resources and set up applications.

Modify

A user with Modify permission can set and update allocated resources.

Execute

A user with Execute permission can execute tasks such as batch processing defined by a user with Modify permission, but cannot define, edit, or delete a task.

View

A user with View permission can view the information about allocated resources.

Peer

Peer permission can only be set for the Device Manager agent. *Citi does not use the Device Manager agent.*


To save the changes to the permission, click on the OK button.

16.3.3 Resource Groups


Device Manager uses Resource Groups to assign certain storage resources to each user. At Citi, we do not use different Resource Groups. However, we have to assign at least "All Resources" in order for a userid to have the appropriate access.

16.3.3.1 Resource Group Allocation


Go to the Resource Group Allocation screen (Administration → Resource Groups → Resource Group Allocation). You will get a security screen verifying the download of a certificate in order to load the Allocation screen.




Click Yes to Proceed. The screen will need to download a Java Applet. If you do not have JRE WebStart, then you may get a download screen instead of the Java application screens.



Once security access to the Java applet is allowed, then the user management screen is displayed.



If a user does not have “All Resource” under the Resource Group column, then it needs to be modified. To modify this value, select the user and then the *Change* button. You will not be able to modify the user “System”.



Change the user’s Resource Group from “<No Group>” to All Resources”. Press Ok when done.

16.3.3.2 Resource Group Administration

Resource Group Administration is used to create different Resource Groups. Citi does not use this functionality at this time.

16.3.3.3 Troubleshooting Resource Group Allocation and Administration

1. If you are unable to initiate this screen, confirm that you are running a JRE package that has JWS (Java Web Start) enabled. If you use the JRE (Java Runtime Engine) package listed at the beginning of this document (see [Software Packages](#)), then you should not have an issue. The USPV (and USPVM) Storage Navigator require the same JWS to function properly.
2. When running the resource group allocation or administration screens, the following error can appear (after JRE with JWS is confirmed and is working):



The resource screen will not load. Other java webstart pages (such as the USPV Storage Navigator) will still load without error. To resolve this issue, there is a setting on internet explorer that will cause these symptoms.

Internet Options → Advanced → Security → Do not save encrypted pages to disk
(The check box/value should be off to resolve issue)

Note: On Internet Explorer 7, the Internet Options are under Tools.

16.4 Security

The security screens define the password policy (length, content, etc.), failed login attempt lockouts, and the banner text.

16.4.1 Password Policy

Go to the Password Policy Page (Administration → Security → Password).

HDvM **HITACHI DATA SYSTEMS**

Password [Help](#)

Password Policy

Enter the minimum number of characters (uppercase, lowercase, numeric, and symbolic), to add up to or less than the minimum password length.

Minimum length(1 - 256): characters

Minimum number of characters:

uppercase(0 - 256) characters

lowercase(0 - 256) characters

numeric(0 - 256) characters

symbolic(0 - 256) characters

Allow password to be the same as the User ID: ☐ Yes ☒ No

[OK](#) [Cancel](#)

Done Local intranet

The following changes should be made to meet Citi Security Standards:

- Change minimum length to 6
- Change numeric to 1
- Allow username and password to be the same to: No

Click OK when done.

16.4.2 Failed Login Attempt Account Lock

Go to the Account Lock screen (Administration → Security → Account Lock)

Select *Edit Settings* to modify the default settings.

HDvM **HITACHI DATA SYSTEMS**

Account Lock [Help](#)

Account Lock Policy

☒ Lockout a user account after number of failed attempts

Maximum number of login attempts (1 - 10) : times

[OK](#) [Cancel](#)

Done Local intranet

The following changes should be made to meet Citi Security Standards:

- Check the box for Lockout a user account after number of failed attempts

- Change the number of login attempts to 6.
Click **OK** when done.

16.4.3 Banner Page

Go to the Banner Page (Administration → Security → Banner)

To change the banner text, select the **Edit Message** button.

To meet the Citi Security standards, banner text (typically provided by BISO) must be added to the text field. Click **OK** when done.

Preview is used to see how the text will appear on the login screen.

16.5 SubSystems


To list existing subsystems select **Resources** and then **Subsystems** from the Object Tree. Click on the **Add Subsystem** link to add arrays into the list.

16.5.1 Add Subsystem


- For a USP, the type is USP. For a USPV, the type is USPV.
- The IP Address is the IP address of the primary SVP.
- The UserID is an administrator ID that has been defined within Storage Navigator. For information regarding special user accounts to be used, see [Special User Accounts](#).
- The password is the password that is associated with that ID.

Click **OK** to save the array definition.


After entering the IP address, UserID and password, a discovery screen will appear. Depending on the configuration, the discovery may take several minutes.




If the discovery process is not successful due to an incorrect userid and/or password, the following screen will be displayed. Select **CLOSE** to cancel the storage system addition.




If the IP address is incorrect, or if the firewall is blocking the ports, then you will get the following error message. Select **CLOSE** to cancel the storage system addition.




If a user has their Storage Navigator session in Modify mode, then the following error will appear:




If the discovery process is successful, then a confirmation screen will appear. Select **CLOSE** to complete the changes.



To list the existing arrays, select *List Objects*.



To display detailed array information, select a listed array off the list. *Show Properties* refreshes the display.



Display Options


- o Physical View – Loads Storage Navigator (using the Device Manager ID). This works only if there is a TCP 2001 Port allowed between the SVP and HDvM. For more information see [Appendix F: Firewall Ports](#).
- o Modify Properties—To change the username, password or IP address of an array
- o LUN Scan – Scans the array for currently configured LDEVs.
- o Remove Subsystem – Removes the subsystem from the database

16.5.2 LUN Scan

LUN Scans are used to create logical objects (based on host ports) for each physical volume within the array. This is how Device Manager keeps track of allocated and unallocated storage.

Note: Do not perform a LUN Scan until hostnames and HBA WWNs have been entered manually. Otherwise, each HBA will be assumed to be an independent server in the database.


A warning screen will be displayed and must be confirmed before proceeding.



16.6 Hosts

16.6.1 Add Host

Hosts need to be added manually to Device Manager to present an accurate database since Citi does not use the HDS Host agent. To create a host entry, select *Add Host*.




Enter the server name in the name field. To add HBAs to the World Wide Names List, select the *Add* button. Enter the HBA WWN in the dialog box that is displayed (using a

space(s) as a separator, a period as a separator, or with no separators between the hexadecimal values). Select OK to save the changes. A successful addition dialog will be displayed. If an HBA is entered incorrectly then storage for only the correct HBA will be displayed. You cannot enter an HBA that has already been entered for another host.

16.6.2 Modify Host

To change the name of a host or to change the HBAs associated with a host, select Modify Properties after a host has been selected from the Object Tree. If an HBA is changing and there is already storage associated with the hba, then device manager will confirm which Host Group (with LDEVs) will need to be updated with the new HBA information.



To repair an incorrect entry for an hba, the old entry should be removed, saved and then a new HBA entry created. If there is storage change confirmation, do not select LUNs (under Host Groups) that are for other HBAs .

16.6.3 Add Storage

Storage can be added to a host by selecting *Add Storage*. Adding storage to a host is termed adding storage to a Storage Domain. In Storage Navigator, the equivalent to a Storage Domain is a Host Group.

16.6.4 Remove Host

To remove a host from the database and to remove its HBAs from accessing any ports (Host Groups), select Remove Host after a host has been selected from the Object Tree. There will be a confirmation of removal along with a display of what ports the HBAs are authorized to use.


16.7 Logical Groups

Select My Groups then select the *Logical Groups* link on the Object Tree. A LUN Scan will automatically create logical groups based on the currently allocated LDEVs by port.


Note: Logical Groups are a category not present in Storage Navigator.

Logical Groups are used to organize users and LDEV allocation. Logical Groups are typically used to organize by server/application type or functional role. For example Oracle Database servers may be distinguished from Application servers. This level of abstraction will be visible in the reporting from Device Manager. Logical Groups can be created with parent and child relationships.

Logical Groups can be initially setup using the LUN Scan function (under *All Storage*). Logical Groups can be renamed and have the icon changed by selecting the group and then selecting the *Modify Properties* link. Parent relationships can also be defined on this screen. The highest level group has no parent (None).



Deleting Logical groups can have an effect on storage allocation. If you try to delete a Logical group, then you will get a warning message:




This warning also lists all of the current allocation that is associated with the logical group. Keeping the access path will delete the group without effecting allocation. Otherwise allocation will be removed with the logical group deletion.

Note: If you do not check the "Keep Access Path" box, then all of the allocation will be removed and all of the servers listed will no longer have access to their storage.

16.8 Device Manger Reports

To see the list of available reports, select the *Reports* link on the Object Tree.



The available reports are:

- [Physical Configuration of Storage Subsystem](#)
- [Storage Utilization by Host](#)
- [Storage Utilization by Logical Group](#)
- [Users and Permissions](#)

The Detailed Array Reports is not a working function in the current microcode.


Samples of these reports are in the [Device Manager Reporting Appendix](#).

There are no facilities for automated reporting from Device Manager at this time.


16.9 Alerts

To see all of the active alerts on defined arrays, select the *Alerts* link under Administration. The screen will display the alerts that have occurred on configured arrays since the Device Manager has been active.


Note: These events do not correlate with the status screen on Storage Navigator.



To see additional detail select the description link and a more detailed screen will be displayed.



In order to obtain an up-to date list of events, select the Monitor Events link. Choosing this item displays the Monitor - Alerts dialog box. In the Monitor - Alerts dialog box, you can actively monitor the alerts. If an alert occurs, the system automatically detects it and reports the alert on the dialog box.



16.10 Tasks

To see the currently configured tasks, select the *Task* link under Administration. These tasks are generated by making changes to the configuration or allocation of any array.

- Operations in progress – Displays a list of tasks.
- Operation – Displays the details of user operation.
- Subsystem – Displays the name of the storage subsystem on which the operation is to be performed.
- Logical Group – Displays the logical group name to be operated.
- Start Time – Displays the start time for the operation.

17 Appendix D: Command Line Interface for Device Manager

CLI (Command Line Interface) is the component of 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.

17.1 HiCommand CLI Setup (Solaris)

Install the Solaris package for HDS Device Manager CLI

The software will be placed in:

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

The package also provides for PowerBroker by installing the scripts with root privileges only. Power Broker must also be configured. For more information see [PowerBroker Configuration \(CLI and CCI Only\)](#).

17.2 HiCommand CLI Setup (Windows)

Install the Windows SPAC for HDS Device Manager CLI

The software will be placed in:

C:\HiCommandCLI

17.3 Secure Option

To avoid transmission of the clear text password, use the `-s` (or `-secure`) option in the command line. This must be configured after SSL has been configured for the server. For more information, see [Using CLI with HTTPS](#).

17.4 Secure Password Prompting

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's JRE version does not affect this feature, only the client.

with JRE = 1.5_11

```
C:\HiCommandCLI>hicommandcli getstoragearray -u system
Password required; Password: manager
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
...
```

JRE = 1.6_7

```
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
```

17.5 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.

Note: The password entries are not encrypted. Do not use a password file or clear text entry in the properties file. You will be prompted for the password. This password entry will not be logged by PowerBroker if you use the prompted password.

Set the model number and serial number of the array in this file as well.

17.5.1 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 Response: "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.

GDSE 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 to a different directory and rename the executable to select the different paths).

17.6 CLI Structure

The structure of CLI is as follows:

HiCommandCLI {server-url} {server-command} [options] [parameters]

For simplification, a copy of HiCommandCLI.sh called HCLI.sh has been provided in the package.

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

17.7 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"

17.8 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.

17.9 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
```

17.10 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.USP.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=USB port=5 domain=2 lun=000 devnum=4096
```

List of 1 Path elements:

An instance of Path

objectID=PATH.USB.14053.5.2.4096
devNum=4,096
portID=5
domainID=2
scsilID=15
LUN=0
wwnSecurityValidity=true

hc.bat addlun serialnum=14053 model=USB port=133 domain=1 lun=000 devnum=4096

List of 1 Path elements:

An instance of Path

objectID=PATH.USB.14053.133.1.4096
devNum=4,096
portID=133
domainID=1
scsilID=15
LUN=0
wwnSecurityValidity=true

- Create Logical Groups and Assign Host/Paths to Logical Groups

Create Logical Group Hierarchy

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.USB.14053.5.2.4096,PATH.USB.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.USB.14053.5.2.4096
    devNum=4,096
    portID=5
    domainID=2
    scsiID=15
    LUN=0
    wwnSecurityValidity=true
  An instance of Path
    objectID=PATH.USB.14053.133.1.4096
    devNum=4,096
    portID=133
    domainID=1
    scsiID=15
    LUN=0
    wwnSecurityValidity=true
```

18 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.

Note: The CCI package also provides for PowerBroker by installing the scripts with root privileges only. Power Broker must also be configured. For more information see [PowerBroker Configuration \(CLI and CCI Only\)](#).

18.1 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.

18.2 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.

18.2.1 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.

18.2.2 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.

18.2.3 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.

18.2.4 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.

18.3 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/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 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.

18.4 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.

18.5 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.

18.6 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

18.7 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 -
```

18.8 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.

19 Appendix F: Firewall Ports


The recommended configuration for the USP is to have the SVPs located behind a firewall or CEP. Device Manager (server) should reside on an SOE Wintel server on the corporate network. The CLI client for Device Manager should reside on an SOE Solaris server with Power Broker logging for all HDS CLI commands. The Device Manager GUI and Storage Navigator GUI should be run from the SOE Wintel server. The DM server and the SN client will go through the firewall to communicate with the SVPs. The following **one-way** 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	**NEW**
	TCP	51100	**NEW**
Device Manager CLI			
	TCP	2001	Does not go through firewall
CCI for ShadowImage			
		User Config	Does not go through firewall
Storage Navigator to SVP			
Standard web interface (client)	TCP	80	HTTP
	TCP	443	HTTPS
	TCP	21	FTP
	TCP	1099	
	TCP	51099	
	TCP	51100	
HDS Remote Connect			
Windows Remote Desktop Client	TCP	3389	MS-RDP
HDS Dial-Home (HiTrack)			
Webproxy (HTTPS via webproxy)	TCP	8080	
DNS (For lookup of webproxy server)*	UDP	53	DNS
SVP to Syslog			

Audit Records to Syslog Server	UDP	514	Syslog
SVP to SNTP (Time Server)			
Time Synchronization	UDP	123	SNTP
SVP to ePO (Virus Scan)			
	TCP	18080	
	TCP	18083	
	TCP	18085	
	TCP	8801	
ePO to SVP (Virus Scan)			
	TCP	18081	
	TCP	18082	
SVP to ESS/SNMP Server			
	UDP	161	SNMP
	UDP	162	SNMP


Figure 1 – Firewall Port tables

*Note: The DNS firewall entry is optional. This is used to lookup the ip address for the webproxy server as many sites implement webproxy as a DNS pool. If the webproxy server is a known, fixed ip address, then the dns lookup is not necessary.




20.3 USPVM Power

The USPVM array variant of the USPV requires single-phase 220volt 30A service. The connector is a NEMA L6-30P(USA) or IEC 309(International). There are four of these connections for the DKU unit




20.4 USP/USPV Dimensions and Spacing

The USP and USPV arrays enclosure are comprised of a DKC (CPU/primary cabinet) and zero or more DKUs (additional secondary disk cabinets). The DKC is 762mm by 925mm. There is an additional service clearance requirement of 727mm on the front and back of the DKC. The DKU is 660mm by 925mm. There is an additional service clearance requirement of 627mm on the front and back of the DKU. Both units are 1860mm in height. The units will be installed with the DKU to the right of the DKC if facing the front of the DKC (called R1). If additional drive space is needed, a secondary DKU would be installed to the left of the DKC (called L1). A third DKU would be installed to the right of R1 and would be called R2. The last DKU would be installed to the left of L1 and would be called L2.



20.5 USPVM Dimensions and Spacing

The USPVM uses a standard 19-inch rack (40U) to house the DKC (36U is used). The secondary rack is not a supported configuration.



In addition to the standard 19-inch rack, one can use the HDS rack for the USPVM. The dimensions are as follows:

Width: 601mm
 Depth: 1050mm*
 Height: 2001mm

* Note: An additional 575mm of depth is required at the rear of the DKC to allow for the access door.

20.6 Weight

The following table assumes 288 (+ 4 Spare) 146GB disks, 4 BEDs. This weight will vary somewhat based on different numbers of drives.

USP/USPV Subsystem	Weight
Approximate weight of DKC:	955 kg
Approximate weight of DKU:	600 kg
<i>Total approximate weight:</i>	<i>1,555kg</i>

The following table assumes 120 300GB disks with 2 BEDs in a single DKC.

USPVM Subsystem	Weight
Approximate weight of DKC:	400 kg
<i>Total approximate weight:</i>	<i>400 kg</i>

20.7 HVAC

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

The USP and USPV are air-cooled. Air must enter the subsystem through the airflow intakes at the sides and bottoms of the frames and must be exhausted out of the top, 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 120 146GB disks in the DKC and 150 146GB disks in the DKU.

USP/USPV Subsystem	Heat Output (kW)	Heat Output (BTUs)	Air Flow (m3/min)	Air Flow (ft3/min)
DKC	6.889	23,506	44	1553.84
DKU	4.395	14,672	31	1094.75

The following table assumes 120 146GB disks in the DKC

USPVM Subsystem	Heat Output (kW)	Heat Output (BTUs)	Air Flow (m3/min)	Air Flow (ft3/min)
DKC	3.0	10,236	10.6	363.73

20.8 Network Connectivity

2 SVPs (1x 10/100/1000 each) within CEP / firewall
1-2 Management Servers (2x 10/100/1000 each)

20.9 Phone Line Connectivity

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

20.10 Fibre Connectivity

LC Connector to (Patch Panel Connection) 50 micron multimode (8x to 32x connections)
This may be regional specific.

20.11 Management Server 1

- Windows OS (Dedicated or under VMWare)
- IP address on corporate network with access to SVP via firewall.
- All GUIs will be run from this system.
- Only designated storage admins will have access to this system.
- Redundant fiber connectivity to each array will be needed for performance (Tuning Manager) and replication (Replication Manager)

Note: A web browser (running JRE 1.5 or later with JWS) is necessary to run the SN or DM clients. If a CLI client is used, the client must have JRE 1.6 or later installed.

20.12 Management Server 2 (CLI)

- Solaris 10
- Using Power Broker to log all command line activity
- IP address on corporate network with access to Management Server 1.

- Redundant fiber connectivity to each array will be needed for replication (HUR,TC ShadowImage) functions.
- No GUIs will be run from this system.
- This system may be appropriate to run syslog services.
- To secure the CLI, the client must have JRE 1.6 or later installed.

21 Appendix H: Device Manager Sample Reports

21.1 Report: Physical Configuration of Storage Subsystem (HTML)

Physical Configuration of Storage Subsystem - Microsoft Internet Explorer

Address: <https://ecctest01ms:2443/service/PhysicalConfiguration>

Links: Windows Marketplace, Tomcat Server Administration, EMC Policy Manager

Physical Configuration of Storage Subsystem

Subsystem	Model	Serial	Configured Capacity	Last Refreshed
USP_V@10.40.53.25	USP_V	10102	16.56 TB	2007/08/28 17:31:38
IP Address	Microcode_SVP	Microcode_DKC	Cache Memory	
10.40.53.25	60-01-30/00	60-01-31-00/00	32.00 GB	

ACP 1: 23 AG / 15.75 TB

CL 1

CHA-1EU Fibre	CHA-1FU Fibre
1A 1B	1E 1F
3A 3B	3E 3F
5A 5B	5E 5F
7A 7B	7E 7F

CL 2

CHA-2QU Fibre	CHA-2RU Fibre
2A 2B	2E 2F
4A 4B	4E 4F
6A 6B	6E 6F
8A 8B	8E 8F

Chassis 2

ACP 1
0 AG
60 Disks

Chassis 1

ACP 1
15 AGs
60 Disks

Chassis 4

<Empty>

ACP 1
0 AG
32 Disks

Frame R0

<Empty>

Frame R1

<Empty>

Chassis 3

ACP 1
8 AGs
32 Disks

Chassis 1

1-1-1	1-2-1	1-3-1	1-4-1	1-5-1	1-6-1	1-7-1	1-8-1	1-9-1	1-10-1
1-11-1	1-12-1	1-13-1	1-14-1	1-15-1					

Chassis 2

Chassis 3

3-1-1	3-2-1	3-3-1	3-4-1	3-5-1	3-6-1	3-7-1	3-8-1
-------	-------	-------	-------	-------	-------	-------	-------

Chassis 4

Done

Local intranet

21.2 Report: Storage Utilization by Host (HTML)

Storage Utilization by Host Report - Microsoft Internet Explorer									
File Edit View Favorites Tools Help									
Back Forward Stop Home Search Favorites									
Address https://ecctest01ms:2443/service/StorageUtilizationByHosts#Anchorad16-p520-02-l2 Go									
Links Windows Marketplace Tomcat Server Administration EMC Policy Manager									
host_5								69.38 GB	0 %
hpl25p								104.07 GB	0 %
namgdse12								173.45 GB	0 %
sanaix1_A								69.38 GB	0 %
sangs1								450.97 GB	0 %
sangs3								136.81 GB	0 %
sangs4								312.21 GB	0 %
sangs5								69.38 GB	0 %
sangs6								277.52 GB	0 %
sangs8								34.69 GB	0 %
t2000								69.38 GB	0 %
uehp10n0								312.21 GB	0 %
uelabhp7								138.76 GB	0 %
uevm1								34.69 GB	0 %
uvm2								69.38 GB	0 %
Total								5.27 TB	0 %
Host: ad16-p520-02-l2									
CL3-A ad16-p520-02-l2/c10-p520-02									
Home / LUN SCAN / USP@10.40.53.13									
Subsystem									
USP@10.40.53.13									
LDEV	Type	SLPR	CLPR	Port/HSD/LUN	Emu Type	Copy Type	Copy Role	Copy Status	Volume A
Host		Assigned LDEV		External Subsystem	External LDEV	File System		Capacity	% Used
10:01	-	0	0	CL3-A/AIX_p520-02_cluster/00	OPEN-V	Unpaired			CVS
ad16-p520-02-l2 : 10.00.00.00.C9.42.F8.EB									34.69 GB
c10-p520-02 : 10.00.00.00.C9.42.FA.F1									
10:1B	-	0	0	CL3-A/AIX_p520-02_cluster/01	OPEN-V	Unpaired			CVS
ad16-p520-02-l2 : 10.00.00.00.C9.42.F8.EB									34.69 GB
c10-p520-02 : 10.00.00.00.C9.42.FA.F1									
10:23	-	0	0	CL3-A/AIX_p520-02_cluster/05	OPEN-V	Unpaired			CVS
ad16-p520-02-l2 : 10.00.00.00.C9.42.F8.EB									34.69 GB
c10-p520-02 : 10.00.00.00.C9.42.FA.F1									
10:24	-	0	0	CL3-A/AIX_p520-02_cluster/06	OPEN-V	Unpaired			CVS
ad16-p520-02-l2 : 10.00.00.00.C9.42.F8.EB									34.69 GB
c10-p520-02 : 10.00.00.00.C9.42.FA.F1									
3C:30	-	0	0	CL3-A/AIX_p520-02_cluster/02	OPEN-V	Unpaired			CVS
ad16-p520-02-l2 : 10.00.00.00.C9.42.F8.EB									500 MB

21.3 Report: Storage Utilization by Logical Group (HTML)

Storage Utilization							
Subsystem	Model	Serial No.	IP Address	Configured Capacity	Open-Allocated	Open-Unallocated	Mainframe-Unspecified
USP@10.40.53.13	USP	10192	10.40.53.13	10.58 TB	2.95 TB	7.63 TB	0 MB
USP_V@10.40.53.25	USP_V	10102	10.40.53.25	16.56 TB	1.01 TB	15.54 TB	0 MB
Total				27.15 TB	3.97 TB	23.17 TB	0 MB

Groups	Number of LUNs
LUN SCAN	297
USP@10.40.53.13	227
CL1-A	6
CL1-A c11-dl585-01/c11-dl585-02	6
CL1-A c15-dl380-01	3
CL1-A d8-v240-03	5
CL1-A d8-v240-04	5
CL1-A ecctest01ms	3
CL1-A q15-dl585-01	3
CL1-A q16-dl585-01	2
CL1-A q16-dl858-02	2
CL1-A q6-dl585-02	1
CL1-A q6-dl585-03	1
CL1-A namqds12	5
CL1-A sangs1	9
CL1-A sangs3	3
CL1-A sangs6	4
CL1-A sangs8	1
CL1-A t2000	2
CL1-A uehp10n0	9
CL1-A uelabhp7	3
CL1-A uvem2	1
CL2-A c11-dl585-01/c11-dl585-02	6
CL2-A c15-dl380-01	3
CL2-A d8-v240-03	5
CL2-A d8-v240-04	5
CL2-A ecctest01ms	3
CL2-A q15-dl585-01	3
CL2-A q16-dl585-01	2
CL2-A q16-dl858-02	2
CL2-A q6-dl585-02	1
CL2-A q6-dl585-03	1
CL2-A namqds12	5
CL2-A sangs1	9
CL2-A sangs3	3

21.4 Report: Users and Resource Groups (HTML)

Users Report - Microsoft Internet Explorer

File Edit View Favorites Tools Help

Address <https://ecctest01ms:2443/service/UsersAndPermissions> Go

Links Windows Marketplace Tomcat Server Administration EMC Policy Manager

Users and Resource Groups

Resource Groups	Number of Users
All Resources	8
<No Group>	1

All Resources

User ID	Full Name	User Description
audit_test	Audit Test (No modify permission)	
cc83851		
ecc	EMC Control Center	
gj06529	Gary Johnson	
HaUser		Peer connections
onaro	Onaro SANscreen	
System		Built-in account
tsm_admin	TSM Manager Account for DM	

<No Group>

User ID	Full Name	User Description
test01		

Done Local intranet

22 Appendix I: HDS Data Migration Tools

Hitachi **Universal Volume Manager** (UVM) is a program product that enables the virtualization of a storage subsystem. When an external storage array fibre channel port(s) are connected to the TagmaStore USP, the UVM enables you to operate the external subsystem as if it is a part of the local TagmaStore USP. Although UVM enables virtualization, HDS' UVM is not a suitable virtualization strategy for Citi. **The use of UVM is restricted to data migrations only.**

In order to provide Citi with the necessary tools and best practices for migrating SAN attached data off of EMC Symmetrix (only) storage arrays to HDS storage arrays, Global Distributed Storage Engineering evaluated two HDS migration tools. Both are part of the HDS Universal Store Platform (USP) or TagmaStore and utilize Universal Volume Manager. The tools evaluated were “**Volume Migration**”¹ and “**Cross-System Copy**”.

22.1.1 Methodology

- Obtain license keys for the following HDS software products
 - o Universal Volume Manager
 - o Cross System Copy
 - o Volume Migration
 - o Performance Monitor
- Configure HDS Front End ports as external ports
 - o Use dedicated Front End ports or add additional ports (in pairs of two)
 - o Maintain redundancy and balance per best practices (one for each fabric)
 - o GDSE recommends pairing two redundant HDS external ports to two redundant EMC FA's
- Configure External Storage
 - o Limited to EMC Symmetrix storage arrays only
 - o Zone the EMC Symmetrix storage array FA's, containing the devices that need to be migrated, to the HDS external ports (one to one zoning)
- Assign EMC devices to HDS
 - o Provide the HDS external ports (WWNs) access to the EMC devices via symmask.
 - o HDS limits external LUNs per path to 255
 - o Prior to provisioning EMC devices to HDS, stop I/O to the EMC devices and export volume groups or device groups on the host, if applicable.
 - o In addition, remove EMC provisioning for the host from the VCMDB and remove zoning
- Discover and configure the external LUNs as a Virtual LUN (UVM)²
- Create target LUN(s) on the HDS array
 - o The source (external array) and target (HDS array) are required to be the exact same size.
 - o Before creating the HDS target LUN, retrieve the size of the source LUN, by viewing the HDS Virtual LUN. This will guarantee that the source and target are exactly the same size.

22.2 Volume Migration

Volume Migration has an automated feature, which enables you to balance workloads among parity groups to improve system performance. It also allows for migrations to be performed manually. The manual process is best suited for Citi migration plans.

¹ See References section for links to document: “Universal Volume Manager User's Guide”.

² See References section for links to document: “Configure Externally Attached Storage Systems – Cookbook”.

During migration, host devices are available to the host via the HDS virtual LUN. Host access to data via the old paths is not allowed. Follow the best practices detailed earlier in this document to add the host to a host group and give the host access to the virtual LUN³. This method is advantageous when the host has a limited outage window. *There is a possible performance impact to the host data, while the data is being migrated.*

Upon completion, this method automatically changes the LDEV number for the source and target. The host would be accessing the same LDEV number, but it would be residing on an actual HDS device.

For details on how to perform this type of migration, please reference the Volume Migration manual.

22.3 Cross-System Copy

Cross-System Copy⁴ is part of the Universal Volume Manager. Much like Volume Migration, Cross-System Copy enables you to copy data from an external storage array to the TagmaStore USP. Unlike Volume Migration, Cross-System Copy does not allow for the data to be accessed from the host while the migration is in process. This method of migration requires an outage for the duration of the migration.

For details on how to perform this type of migration, please reference the Cross-System Copy manual.

Upon completion, follow the best practices detailed earlier in this document to add the host to a host group and give the host access to the HDS LUN.

22.4 Recommendations

- GDSE recommends that migrations be performed during non-peak hours when I/O patterns are low or when the host can tolerate an outage.
- Multiple HDS Front End ports (redundant) can be configured as External Ports. External Ports should be paired with redundant EMC FA ports. If you're migrating data from multiple EMC Symmetrix arrays, then you'll need a pair of HDS External Ports for each array. This should help with hosts that have data spread across multiple Symmetrix arrays.
- Zone one HDS External Port to one EMC Symmetrix FA.
- Both migration processes do a block level copy of the data at the LU level. Migration times tested on a 16.8GB device took 15 minutes. Based on that, it's estimated that a 250GB device would take approximately 3.5 hours to migrate.
- The most simplistic of the two migration methods is Cross-System Copy. Using this method does require a host outage. In instances where a host can tolerate downtime, GDSE recommends using Cross-System Copy.
- If downtime is not an option, then Volume Migration should be used, allowing the host to access the HDS Virtual LUN during the migration.

³ See References section for links to document: "Universal Volume Manager User's Guide" and "Performance Manager User's Guide - Performance Monitor, Volume Migration, and Server Priority Manager"

⁴ See References section for links to document: "Cross-System Copy User's Guide"

23 Appendix J: Citi McAfee Virus Scan and ePO Agent

23.1 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.

23.1.1 Prior To Beginning Installation

3. 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 and 18082 to communicate to the SVPs. For a complete listing of firewall ports see [Appendix F: Firewall Ports](#).
4. Obtain a SiteList.xml file that describes the local ePO server locations from regional ePO / Virus Scan administrators. The ePO installation will require this file. 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 3.X ePO agent (CMA3601.zip or 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.

5. It is critical that you contact the regional ePO administrator to make sure the exclusions (see [Global ePO Policy 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.

23.1.2 Installation Steps

6. 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.
7. Copy both zip files and "sitelist.xml" to the SVP. The "CMA3601.zip" (or more recent version) file is for the ePO agent and the "VSE85_Non-SOE.zip" (or more recent version) file is for Virus Scan.
8. Unzip the "VSE85_Non-SOE.zip" file into an installation directory if not uncompressed already.

Note: Winzip is not required to uncompress zip files on the WindowsXP (or Windows Vista) that is on the SVP.

VSE Non-interactive Install (default and recommended)

This is the default behavior of the package. To change to interactive, modify setup.ini:

bSilent=1 (change to 0 for interactive)

Run setupvse.exe. The installation will complete in 1-2 minutes. Proceed to Step 6.


VSE Interactive Install (optional)


This is not the default behavior of the package. To change the installation to non-interactive, modify setup.ini:

bSilent=0 (change to 1 for non-interactive)


Run setupvse.exe.

Select the custom install. Under custom install, un-select "Auto-Update", "Lotus Notes...", and "Microsoft Outlook...". Also disable the buffer Overflow and Script Scan to avoid any performance issues. Keep On-Access Scan and On-Demand Scan. Your options should look like the following screen:







Click Next. Un-select “Install Alert Manager Server” and “Install Desktop Firewall”



Click Next. Un select “Import AutoUpdate repository”. Select the “Enable On-Access Scanner at the end of installation”



Click Next. Select "Show system tray icon with all menu options"



Click Next. The installation will begin.

9. Start the Virus Scan application (if not already started by the install program).
Once started, there should be a small icon  or  in the system tray (near the clock) in the task bar.
10. ePO Install. Unzip the "CMA3601.zip" into an installation directory if not uncompressed already(i.e. FrmInst.exe or FramePkg.exe). Go to the directory that contains FrmInst.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:


```
FrmInst.exe /install=agent /siteinfo=<path to>sitelist.xml
```

```
FramePkg.exe /install=agent /siteinfo=<path to>sitelist.xml
```


For example:

FramePkg.exe /install=agent /siteinfo=c:\temp\virus\sitelist.xml

A dialog box should appear and show the installation status.



11. Repeat steps 4-8 for the secondary SVP.

Once both products have been successfully installed, confirm that the regional ePO administrator can configure the SVPs from the ePO server. **This is important to keep the virus dat files up to date and to allow for any other updates to the virus software.**


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.

23.2 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:


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.

23.3 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 Frmlnst.exe. Run the following command (from the command line) to install the product:

```
Frmlnst.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.

23.4 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.

24 Appendix K: Setup of HiTrack Service for HTTPS

Outbound configuration and Alert Information 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). 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. HiTrack Version B.8 05/30/2007 or later is required. The information being transmitted to HDS is non-confidential and is encrypted. This applies to USP and USPV arrays.

24.1 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.



Hi-Track Service Manager

General | Site | Modem | **Center** | Subsystems | Notification | Utilities

Hi-Track Center

Contact Attempt Sequence

Order	Method	Location
1	HTTPS	ushtinet01.hds...
2	HTTPS	ushtinet02.hds...

Maximum sequence passes: 4


Call Time

Daily call time: 11:27 AM Weekly day: Saturday Randomize

Copyright © 1995-2007 Hitachi Data Systems Version B.8 - 5/30/2007

OK Apply Help

4. Primary FTP Server. Select Add to create a new entry.



Contact Method Profile

Method

Type: ☐ Dial-up ☐ FTP ☒ HTTPS Times to attempt: 4 (0 = disabled)

Properties

Phone number:

HDS host: ushtinet01.hds.com

Username: ht_http_row

Password: ☒ Hide password

☒ Passive mode ☐ Use SSL

Port Range: Minimum Maximum

Proxy Server

Type: HTTP Tunnel An HTTP proxy that tunnels for FTP and HTTPS

Host or IP: webproxy.ssmb.com

Port: 8080

Username: (if applicable)

Password: (if applicable) ☒ Hide password

Test OK Close

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.ssmb.com
Port: 8080

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.ssmb.com
Port: 8080

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.*

25 Appendix L: HDS Tiered Storage Manager

25.1 What is TSM?

HDS Tiered Storage Manager (TSM) 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 Storage Domain.

The driver for certification of this technology is to enable simple migration between tiers within a storage array, primarily to allow applications to be introduced on the lowest cost tier and only migrated to a higher tier if there are performance issues. It can also be used to migrate from a higher tier to a lower tier within a storage array to reduce chargeback costs.

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.**
- **User-definable migration groups** allow all LDEVs associated with an application to be migrated as a single task. A migration group is a fixed list of source volumes for the migration.
- **Storage Tiers defined in terms of RAID type and disk size.** TSM defines manages Storage Tiers in terms of LDEV characteristics, rather than lists of devices. Once these definitions are setup, TSM 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 TSM 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 TSM control. Tasks created with the Web Client can be executed via the CLI and vice versa. The CLI is supported on Solaris 10 only.

TSM has its own server which should reside on the same Windows 2003 host as the Device Manager server. TSM depends on Device Manager (DM) for all communication with the actual storage array. The web and CLI clients communicate with the TSM server. The TSM server in turn interfaces with Device Manager for all information on and control of the storage arrays.

25.2 Documentation

HDS provide particularly good documentation for TSM and no attempt is made to duplicate that information here, other than to provide an overview of the migration steps. Anyone implementing TSM will need to consult the first two documents. Users of TSM must be familiar with the User's Guide as a minimum.

1. **HiCommand Tiered Storage Manager User's Guide** – This is the key document. Covers general overview of TSM usage as well as typical operations using the Web client. Sections on locking and unlocking tasks should be ignored as this functionality is not being used by Citi. Shredding tasks will usually not be used; shredding will usually be undertaken as part of a migration task.

<https://catecollaboration.citigroup.net/domains/platstor/gdse/products/platforms/Tagmastore/hc0909.pdf>

2. **HiCommand Tiered Storage Manager Server Installation and Configuration Guide** – Detailed description of the server component of TSM, including backing up and restoring repositories.

<https://catecollaboration.citigroup.net/domains/platstor/gdse/products/platforms/Tagmastore/hc1502.pdf>

3. **HiCommand Tiered Storage Manager CLI Reference Guide** – Detailed description of TSM CLI. Assumes knowledge of TSM usage covered in the User Guide.

<https://catecollaboration.citigroup.net/domains/platstor/gdse/products/platforms/Tagmastore/hc0919.pdf>

4. **HiCommand Tiered Storage Manager Messages Operations Guide** – Classified descriptions of error messages and log file entries.

<https://catecollaboration.citigroup.net/domains/platstor/gdse/products/platforms/Tagmastore/hc0929.pdf>

25.3 Overview of Migration Steps

Before migrations can take place, there are two main setup tasks

- 1) **Creation of Storage Domain.** This is a one-off task to bring a USP storage controller array under TSM control. This can only be done by a user who has Admin permission.
- 2) **Creation of Storage Tiers.** These are pools of free storage to which migrations take place. These will usually be set up once when TSM 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.


Then for each specific migration there are the following steps:

- 1) **Creation of a Migration Group.** The Migration Group (MG) defines the source volumes for the migration. MGs 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.
- 2) **Creation of a Migration Task.** This involves selecting the Storage Tier and devices within that tier to which the LDEVs in the MG will be migrated to. It is at this time that the user also selects if automatic erasure of the released storage is to be undertaken (recommended) and if the task is to be executed immediately after creation. Note that when using the CLI, one must first create a migration plan, and then create a migration task as a separate activity. This is to allow one to modify the migration plan (e.g. select different migration target devices to those suggested by TSM) before creating the migration task. Modification to a CLI migration plan is undertaken using a text editor.
- 3) **Execution of the Migration Task.** Tasks created with the Web Client can be executed via the CLI and viceversa. Separation of Migration Task creation and execution allows
 - a. Migrations to be created in advance and executed during a green zone
 - b. Task creation and execution to be undertaken by different people. There is a TSM execute permission that allows a user to execute but not create tasks.

Note that after the data is copied from source to target (100% progress on the task progress status), there can be several minutes wait whilst the swaps take place and each volume reaches “Success status”. It is normal for each volume to reach “Success status” in turn, not simultaneously.

Erasure of released storage automatically takes place at the end of migration, if selected at task creation time. Separate erasure tasks can also be created (CLI only) to allow erasure to be deferred if necessary (not recommended).

The following diagram summarizes the migration task execution differences and commonalities between the web and CLI clients.



For a detailed description of how to set up Migration tasks see the HDS document **HiCommand Tiered Storage Manager User's Guide**. In addition, if the CLI is to be used for migration tasks, read the **HiCommand Tiered Storage Manager CLI Reference Guide**.


25.3.1 Steps in the actual migration process

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 Array Group A-A-A but after migration is part of Array Group B-B-B. Hence, sequential LDEV numbers may be in different Array Groups when TSM is used in an array.
- Although the Migration Source volume is 1:01 and the Migration Target Volume is 11:01, erase of data 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 erasure.



25.4 TSM Server Installation and Configuration

25.4.1 Licensing

The following licenses key need to be installed on the USP Domain Controller

- Tiered Storage Manager (apply when first start up TSM)
- Volume Migrator2 (apply via Storage Navigator)
- Volume Shredder (apply via Storage Navigator)

As TSM depends on Device Manager, this of course must also be installed and licensed.

Keys must be obtained from HDS as they are locked to the USP serial number. The licenses can either be applied by HDS engineer (via the SVP) or by end user using Storage Navigator.

Note that once licenses have been applied, it is necessary to first refresh Device Manager and then refresh TSM before TSM will see them as activated.

25.4.2 TSM Server Requirements

The TSM server should be installed on the same server as Device Manager (DM). Windows 2003 is the standard platform for DM and TSM.

No more than 5 Storage Domains can be managed by a single server.

Note that the HTSM CLI should be installed on a Solaris 10 server (see later section).

25.4.3 Installing the TSM Server

The latest Windows TSM Server wrap name is identified in the TSM release notes.

During the installation:

- For "IP address or host name" enter the IP address of the server on which you are installing the TSM server software.
- Take the defaults for Port numbers.

25.4.4 Connecting with the web client

TSM web client requires:

- Internet Explorer 6.0 or higher
- Sun JRE 1.4.2 or higher

The usual way to launch the web client is from the DM web client dashboard.

Alternatively, you can connect directly with a web browser using one of the following URLs:

If SSL has not yet been set up in the HiCommand Web Server

`http://ServerHost:23015/TieredStorageManager/`

Once SSL has been set up (as described in the Security Practices Section) then use URL

`https://ServerHost:23016/TieredStorageManager/`

[Note use of https and port 23016 in this case where SSL is used]

25.4.5 User Management setup

DM and TSM both share the same authorization domain (single sign-on) and each tool can be launched from the other without additional authorization. However, as permissions are used to restrict operations performed by individual users, it is necessary to grant additional permissions to DM users who need to use TSM.

When you first log in after installing TSM, log in with the system account and, if necessary, set up a user account that has the User Management permission and a user account that has the TSM Admin permission.

User management is common between HiCommand Suite products and hence between DM and TSM. However, four permissions are provided for performing TSM operations. User management permission can be combined with any other permission. 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 TSM should have their own account with the appropriate permissions for their role.

TSM Permission	Description
Admin	Dedicated account permission for managing the storage domain. Permits user to manage storage domains and view all TSM resources
Modify	Permits user to use all functions except creating and deleting storage domains. Needed for daily operation of TSM.
Execute	Permits user to execute existing tasks and view all TSM resources.
View	Permits user to view all TSM 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.

User entitlement reporting with TSM is the same process as for Device Manager as both products share the same HiCommand administration database.

25.4.6 Configuration of server.properties files

The following changes should be made to the TSM server properties file

`C:\Program Files\HiCommand\TieredStorageManager\conf\server.properties`

`server.repository.autoRefresh.pollingIntervalInMinute=360`

This ensures that TSM checks every 360 minutes to see if the Device Manager (DM) repository (for which it depends for all device information) has been refreshed. If TSM detects a refresh, then TSM storage domain associated with the storage subsystem in DM is refreshed. The exact

value that this parameter is set to is not critical; that fact that it is set is important otherwise automatic refreshing will not take place.

```
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.mail.smtp.host=<your SMTP host>
```

```
server.mail.from=<e-mail-address>
```

These allow e-mail notification when migration tasks complete. In order to send e-mail notification via the Citigroup mail routers the “from” address for all e-mails must be changed to a Citigroup valid address. This address can be created for you by Messaging Systems Support group. See <http://gemsbb.nj.ssmb.com/mss/> (select “policies” then read “Applications that send Email”) for details of the process.

```
server.rmi.server=2
```

```
server.securty.port=24500
```

These changes allow the TSM CLI client to connect using SSL communication.

Note that setting `server.rmi.server=2` enables SSL communications but does not disable non-SSL communications via port 20352.

25.4.7 Setting up integration with HiCommand Tuning Manager

The following setup will allow one to view average and maximum usage of an array group from within Tiered Storage Manager. This can be used to check acquired usage information from Tuning Manager, and then migrate volumes from array groups that have high busy rates to array groups with low busy rates.

First the “HiCommand Suite TuningService” service must be started on the Tuning Manager server. This can be achieved using the Windows Services program. Also, change “Startup Type” to “Automatic”.

Then, on the TSM server edit the file

```
C:\Program Files\HiCommand\TieredStorageManager\conf\tuningmanager.properties
```

as follows:

```
htnm.servers=1
```

```
htnm.server.0.host=<Tuning-Manager-Hostname-or-IP>
```

```
htnm.server.0.port=23015
```

Use port 23015 for non-SSL communications, and 23016 for SSL communications.

Stop and restart the TSM server. Then log into the TSM and refresh the storage domains.

This is sufficient to allow Array Group busy and max busy % information to be available to TSM. For further information on configuring performance information acquisition with TSM, see the Hitachi Tiered Storage Manager User’s Guide.

25.4.8 Initial refreshes and creating Storage Domains.

Before using TSM for the first time with a new Storage Domain, it is strongly recommended that a “Refresh All” is done in Storage Navigator for the Domain Controller, followed by a Device Manager and then a TSM refresh for the said Domain Controller. Failure to do so may cause migration task setup to fail due to configuration inconsistencies. Testing has shown that the “Refresh All” within Storage Navigator is necessary.

When creating the name of a Storage Domain, do not use special characters (e.g. @) as these can cause issue when used with the CLI (see CLI section).

25.4.9 Building Storage Tiers

LDEVs in all storage tiers should be of identical size to permit migration to and from each tier. The standard LDEV size is 36375360 KB.

The RAID types and disk types supported within Citi change as a function of time. See Section 11 for currently supported configurations. At the present time, it is recommended that only two tiers be created. Some possible tiers are indicated below.

Suggested Storage Tier name*	Raid Type	Physical Disk Size
RAID5_146GB	RAID5 (7D+1P)	146GB
RAID5_300GB	RAID5 (7D+1P)	300GB
RAID6_1TB	RAID6 (6D+2P)	1TB
Thin_RAID6_146GB	RAID6 (6D+2P)	146GB
Thin_RAID6_300GB	RAID6 (6D+2P)	300GB
Thin_RAID6_1TB	RAID6 (6D+2P)	1TB

* Tier name should not include spaces if GDSE CLI TSM reporting script is to be used. Alternatively, the Storage class name (Gold, Silver) could also be included in the Storage Tier name.

The exact ratio of disk types to install is difficult to judge and may be determined by:

- % of each disk type included in a particular procurement deal
- Anticipated customer requirements for each tier

The following example gives several different options for installed disk ratios, showing the approximate relative amount of available usage disk space in each pool. Other ratios of disks are of course possible.

Configuration Example	Physical 300GB disks*	Physical 146GB disks*	Usable storage in tier RAID5 300GB	Usable storage in tier RAID5 146GB
A	50%	50%	66%	33%
B	60%	40%	75%	25%
C	80%	20%	90%	10%

* % by spindle count.

The TSM storage tiers should be set up when a Storage Array is first brought under TSM control. It is essential that tiers are not created with overlapping definitions so that an LDEV meets the criteria of more than one tier. If this happens, the LDEV will appear in both tiers and double-counting will occur when viewing the capacity (total, free and used) of each tier.

The following tier definitions are suggested:

Tier Name	TSM definition
RAID5_300GB	Using Volume Filter: RAIDLevel = 'RAID5(7D+1P)' AND DiskCapacityInGB > 150 AND DiskCapacityInGB < 350
RAID5_146GB	Using Volume Filter: RAIDLevel = 'RAID5(7D+1P)' AND DiskCapacityInGB > 100 AND DiskCapacityInGB < 150
RAID6_1TB	Using Volume Filter: RAIDLevel = 'RAID6(6D+2P)' AND DiskCapacityInGB > 800 AND DiskCapacityInGB < 1200
Thin_RAID6_*	Using Pool Filter: PoolID = 'DP x' OR PoolID = 'DP y' OR PoolID = 'DP z' Where DP x, DP y and DP z are the DP pools that map to this tier (see separate section on using TSM with DP volume later)

25.4.10 SANscreen integration

TSM migrations do not cause any changes to logical configuration within the SAN (LDEVs allocated to a host does not change) and hence such changes cannot automatically be detected by SANscreen. eTrust audit should be used for auditing changes executed with TSM.

Selecting an HDS storage array in SANscreen (v3.5), and then the “Volumes” icon at the bottom of the screen brings up a list of Volumes (LDEVs), which does show both RAID type and Disk Size. Hence, it is possible to determine the tier that an LDEV is currently part of using SANscreen.

25.4.11 eTrust Audit setup

eTrust Audit configuration for TSM is covered in the document “**eTrust Audit Storage Monitor Solutions Guide**”. The latest version can be found:

<http://catecollaboration.citigroup.net/domains/platstor/gdse/products/techstdshmpgs/etruststorageaudit.aspx>

25.4.12 Audit Log Configuration - event log

The following changes should be made to the TSM 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 HTSMserverMessagesN.log files.

```
C:\Program Files\HiCommand\TieredStorageManager\conf\logger.properties

logger.sysloglevel=20
```

TSM information will be logged as part of the HiCommand integrated logging, and hence the source in the event log for TSM messages will be “HBase Storage Mgmt Log”.

Event log audit logging to the event log is also controlled at the HiCommand Common component level and hence may already be enabled on any system already running other HiCommand products, such as Device Manager.

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 TSM audit log information, one must check the environment settings file for the HiCommand Suite Common Component is as follows:

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

Log.Event.Category StartStop,Authentication,AccessControl,ConfigurationAccess
Log.Level 6
```

This will log all information that TSM can generate. If you need to edit this file, it will be necessary to restart HiCommand Common Component to enable the new settings.

TSM information will be logged as part of the HiCommand integrated logging, and hence the source in the event log for TSM messages will be “HBase Storage Mgmt Log”.

Device Manager also requires this same edit. For more information see .

Note also that CLI Auditing is also recorded in the HTSMServerMessageN.log file and Web Client Authentication is also recorded in the HTSMGuiMessageN.log file.

25.5 Citigroup TSM Best Practice and Operational Notes

25.5.1 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.

25.5.2 Refreshing the TSM database

TSM has its own database. Information to populate the TSM database is obtained from Device Manager. Device Manager in turn has its own database with array configuration and status information which needs to be updated from the storage arrays.

A mismatch can occur between the DM and TSM databases if configuration of a storage subsystem is changed, including changes to pair state made with Replication Manager. If changes are made to the subsystem with Storage Navigator, it will be necessary to ensure that the DM database is first refreshed and then the TSM database before any TSM work is undertaken. If changes are to the subsystem are made via DM, then only the TSM database needs to be refreshed.

The consequence of the TSM database not being up to date when doing a migration is very serious - to quote the User Manual "data might be migrated to an unexpected area". The following Best Practice is recommended to avoid this:

- 1) Set up autoRefresh as discussed in the TSM server installation and configuration section. This will ensure that TSM will check periodically if the DM database has been refreshed, and if so refresh the TSM database. This will reduce the chance of the TSM database being inconsistent with the DM database.
- 2) Always refresh the TSM database before setting up a migration task in case any changes have taken place in the DM database since the TSM database was last refreshed. If you are sure that no changes can have taken place (e.g. no other staff have access to the system since the TSM database was last refreshed) then this is not necessary. The time of last TSM database refresh can be found on the "Detailed Information" tab for the Storage Domain.
- 3) If Storage Navigator is used at your site, it is necessary to first refresh DM database and then refresh TSM database if there is any possibility of changes having been made to the array configuration since these databases were last refreshed.

If in any doubt, refresh the DM database and then the TSM database to ensure that they are both consistent with the current configuration of the array.

25.5.3 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.

TSM contains no in-built scheduler for task execution. If automated delayed execution of a migration task is required, an external scheduler will need to be used.

25.5.4 Reporting on storage usage by Tier

The eventual goal is to provide utilization and chargeback reports via the Strategic Storage reporting infrastructure. However, this cannot be done at this time as:

- Not all regions are using ECC for chargeback reporting.
- ECC is unable to determine all the underlying disk characteristics for an LDEV that define which Citi Storage Class the LDEV belongs to. Hence the basic information needed to chargeback by tiers is not available.
- Customized reporting would be required with ECC to provide the chargeback reports required, even if the missing information were available.

Engineering will:

- a) continue to work with EMC and HDS to try to resolve the missing information issues
- b) examine the capabilities of new versions of ECC to see if customized reporting for HDS can be avoided

In the meantime, local regional allocation reporting scripts can be used. Alternatively, the latest TSM Release Notes provide a sample script **tsm_report** can provide reports in excel-compatible format (tab delimited) on usage by Storage Tier for each host in a Storage Domain.

<https://catercollaboration.citigroup.net/domains/platstor/gdse/products/techstdshmpgs/tagmastore.aspx>

25.5.5 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 using the CreateMigrationTask with the `erasedata=yes` qualifier, when using the CLI, and 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.

25.5.6 Backup/recovery of TSM database

The TSM database can be backed up as part of the integrated backup of all HiCommand databases. Also the TSM database can be restored separately. The backup and restore process is as follows:

To backup HiCommand databases

1) Backup databases

```
cd C:\Program Files\HiCommand\Base\bin
hcmdsbackup /dir C:\<directory>\HiCommand_backup_<date>
```

To restore just the TSM database

1) Stop TSM server

2) Stop DM server and common services

3) Restore TSM database only

```
cd C:\Program Files\HiCommand\Base\bin
hcmdsdb /restore C:\<directory>\HiCommand_backup_<date> \database\backup_hdb /type
TieredStorageManager
```

4) Restart DM and then TSM servers

25.5.7 Task Completion Notification

The web client shows the current status of a migration and this will be the usual place where status can be checked. The non-interactive notification method is e-mail.

An example of a notification mail is shown below. Note that it does show if data erasure was undertaken as part of the migration.

The following migration task has finished:

```
Task ID: TK1if3iy2s
Task type: Migration
Status: Success
Task owner: cc83851
Execution request time: 2007/07/16 09:13:36 End time: 2007/07/16 09:34:15
Storage domain: USPV10102 Migration group: TestMG
Previous target storage tier:
```

```
Target storage tier: RAID6 300GB
Erase data: Yes
Migrated volumes: "00:00:02"
```

This message was sent automatically by the Tiered Storage Manager server.

The CLI command `GetTasks` can also be used to get the status of a task. This could be used to create a script that can trigger an alternative notification process, if desired. For example:

- a) `htsmcli GetTasks -u <username> -p <password> migrationgroupname="MyApp"`
- b) `htsmcli GetTasks -u <username> -p <password> id=<task-id>`

will return information, including status (e.g. Success), for a) all groups associated with this migration group and b) the specified task-id respectively, which can be tested against.

CLI scripts can be produced once the actual operational requirement is clearly defined.

Finally, Task completion is also logged in the `HTSMServerMessageN` file.

25.5.8 Dealing with failures

As previously described, migration failures can occur if a migration is attempted when high I/O rates are occurring to the source volumes. This can leave the TSM database in an inconsistent state regarding the status of source and target volumes involved in the migration. A refresh of the TSM database usually resolves any issues.

If you are unsure if the swap took place during a failure, check which the tier the LDEVs in the migration group are now in. Make sure you refresh the TSM database first as otherwise volumes may be marked as used (when they are not) and the volumes may not appear in the correct tier.

25.5.9 Usage of migration policies

The target volumes from a Storage tier that gets auto-selected by TSM for a migration can be influenced using array group selection rules (rule tab on Create MG page of web client). The default is "BalancedCapacity", which balances the number of candidate LDEV migration targets from with the array group of the storage tier (i.e. selects targets from the Array Groups with the greatest capacity) . If a particular application is particularly I/O intensive, the rule can be changed to "MaximumCoverage" which will distribute the migration target LDEVs across as many array groups as possible. Conversely, the rule "Minimum Coverage", attempts to select all targets from the same Array Group. See the User's Manual for further details on LDEV selection rules.

Finally, 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 the following entry in the `server.properties` file for the TSM server.

```
server.migrationPlan.candidateVolumeCountLimit=false
```

However, under normal circumstance, there is not reason not to accept the auto-selected recommendations for target devices.

25.6 Client CLI installation, configuration and operations

25.6.1 Introduction

The web client is well suited to the operational model of TSM. It is particularly strong compared with the CLI in areas such as Migration Group volume selection, and migration pair selection, which can often be an interactive process. Monitoring and management of tasks is also much easier with the web client.

However, there is also an important role for the CLI in terms of automation, and hence is supported within Citigroup.

25.6.2 Installing the CLI Client

The HTSM CLI (package name: PRhtsmcli on the GDS Repository) should be installed on a Solaris 10 server. See latest Release Notes for exact location on the GDS Repository of the current version.

Install with the command:

```
# pkgadd -d . PRhtsmcli
```

The package differs from the tarball delivered by Hitachi in the following respects:

1. The postinstall will run the setup script automatically.
2. There is an additional softlink “current” which points to the current version tree. This is provided to avoid the need to put the current version number into either scripts or into the PowerBroker configuration (see below for more details)
3. The htsmcli script permissions have been changed to only allow it to be run by root, or from user accounts via PowerBroker (which run as root on behalf of the user)

25.6.3 Configuration of properties file

Note that the GDSE package of the TSM CLI adds a softlink “current” which points to the current version tree. For example,

```
/export/opt/TieredStorageManager/current → /export/opt/TieredStorageManager/0620
```

Hence, it is possible to refer to the TSM CLI installation path in a version independent manner for scripts and PowerBroker installation. Use of the “current” softlink is assumed for the rest of this document.

The following changes are recommended for the CLI properties file

```
/export/opt/TieredStorageManager/current/CLI/htsmcli.properties
```

The following enable the use of secure SSL communications between the TSM CLI client and the TSM server. As well as forcing the use of secure communications, this entry removes the need to use the `-s` qualifier on command lines when one wants to invoke communications via SSL).

```
option.secure=true
```

The following entries avoid you needing to specify these parameters everytime a CLI command is issued. Substitute your own server IP address and storagedomainname in. Note that if the storagedomainname is to be defined in this properties file, then it must not contain an @ character. Hence, ensure your storagedomainname does not include the "@" symbol (it will if you accept the default name when creating the storage domain). Storage domains can be renamed in TSM. If you have to manage multiple storage domains, then obviously you will not be able to use the properties file to set this parameter.

```
# set the server location
htmsserver.location=rmi://10.40.42.19:24500/HTSMserver #(for SSL communications)

*or*

htmsserver.location=rmi://10.40.42.19:20352/HTSMserver #(for non-SSL communications)

# set the Storage domain name, for commands that use this parameter
parameter.storagedomainname=USP10192
```

The final step to enable SSL communications is to download the electronic certificate from the TSM server from URL:

<http://<TSM-server-IP>:23015/TieredStorageManager/TieredStorageManagerCerts>

This file (TieredStorageManagerCerts) should be placed onto the Solaris TSM CLI client

/export/opt/TieredStorageManager/current/CLI/TieredStorageManagerCerts

When using SSL communications, all CLI scripts will now need the following added at their start to ensure that this certificate can be found:

```
HTSM_CLI_CERTS_PATH=/opt/TieredStorageManager/current/CLI/TieredStorageManagerCerts
export HTSM_CLI_CERTS_PATH
```

Similarly the above environment variable will need to be set when running TSM CLI commands interactively.

Do not use the parameter option.password to define the password directly, or to define a password file. When using the CLI, use the arguments -u <username> -p @<password file path> to pass the password to the TSM server. Do not specify the actual password else this will be stored in the PowerBroker log file.

With HTSM 6.2 once can now use the new htsmaccount CLI command to encode a password in the htsmcli.properties file. This allows one to omit entering a password when executing CLI commands. However, this ties usage of the CLI to a single account. For further details of the htsmaccount command, see the HTSM CLI Reference Guide (p2-13).

<https://catecollaboration.citigroup.net/domains/platstor/gdse/products/platforms/Tagmastore/hc0919.pdf>

25.6.4 Environment Variables

The following two environment variables should be set before running any CLI commands:

```
HTSM_CLI_DIR=/opt/TieredStorageManager/current/CLI
export HTSM_CLI_DIR
HTSM_CLI_HOME=/opt/TieredStorageManager/current/CLI
```



```
export HTSM_CLI_HOME
HTSM_CLI_CERTS_PATH=/opt/TieredStorageManager/current/CLI/TieredStorageManagerCerts
export HTSM_CLI_CERTS_PATH
```

HTSM_CLI_DIR defines the location of the CLI directory

HTSM_CLI_HOME is an environment variable used to report the location of the TSM htsmcli.properties file. If this is not set, settings in the properties file will not be used.

HTSM_CLI_CERTS_PATH defined the locaton of the SSL Certificate for TSM

25.6.5 PowerBroker configuration

The CLI must be run from user accounts via PowerBroker. The CLI package has been installed with root-only access. PowerBroker will allow specified users to access these commands and log it.

Powerbroker configuration is covered in the Section entitled PowerBroker Configuration. However, a summary of the TSM specific points are highlighted below for clarity and convenience:

Changes required on pbmaster server:

- Add TSM CLI client to /etc/hosts, if necessary
- Modify pb.conf on pbmster to support TSM
 - Add TSM CLI server to hosts list
 - Add htsmcli to adminprogs list
 - Add /export/opt/TieredStorageManager/current/CLI to PATH environment variable definition. Use of the current softlink is preferred as it makes the installation version independent and changes to PowerBroker set up will not be needed when TSM version is upgraded.

Changes required on TSM CLI client host:

- Add pbmaster to /etc/hosts, if necessary
- Install PowerBroker client software, if not already installed
- Copy /etc/opt/pb/pb.key from PBmaster to same directory on TSM CLI client host
- Change permissions of pb.key to 600 on TSM CLI client host

Set TSM environment variables before running htsmcli commands via pbrun. For example:

```
#!/bin/sh
HTSM_CLI_DIR=/opt/TieredStorageManager/current/CLI
export HTSM_CLI_DIR
HTSM_CLI_HOME=/opt/TieredStorageManager/current/CLI
export HTSM_CLI_HOME
#
echo $1 > /tmp/123
/opt/pb/bin/pbrun htsmcli CreateMigrationTask -u cc83851 -p @/tmp/123
migrationplan="hds_luns1" erasedata=no
```

```
rm /tmp/123
```

Note the use of the syntax `-p @/tmp/123` (where `/tmp/123` temporarily contains the user's password) to provide the password to the TSM server. TSM does not prompt for a password if it is not supplied (unlike DM), but will instead fail. Hence, it is necessary to use the `-p @<password file path>` syntax to provide the password, as defining this in the `server.properties` file is not recommended for security reasons). The syntax `-p <actual-password>` should not be used as this would log the password in the PowerBroker log file. In the above example, the password is passed as an argument to the script, written to a temporary password file on-the-fly, and then deleted again after the command is run. In this way, the password is not persistent on disk, and yet is also not recorded in the PowerBroker log.

An edited (for length) example of the entry in `pb.eventlog` is shown below

```
requestuser      'cc83851'
runargv[]        'htsmcli' 'CreateMigrationTask' '-u' 'cc83851' '-p' '@/tmp/cac'
'migrationplan=hds_luns1' 'erasedata=no'
runbkgd 0
runcommand      'htsmcli'
runcwd          '/opt/TieredStorageManager/current/CLI'
runenv[]        'DISPLAY=tantort:0.0' 'EDITOR=vi' 'HISTFILE=//.history.cc83851'
'HISTSIZE=256000' 'HOME=/' 'HTSM_CLI_DIR=/opt/TieredStorageManager/current/CLI'
'HTSM_CLI_HOME=/opt/TieredStorageManager/current/CLI' 'HZ=100'
'LD_LIBRARY_PATH=/usr/lib:/usr/local/lib:/usr/openwin/lib:/usr/ucblib:/lib:/usr/storapi/
shlib:.' 'LOGNAME=cc83851' 'MAIL=/var/mail/cc83851'
'MANPATH=/usr/local/man:/doc/man/usr/storapi/man:/usr/storapi/storman:/usr/share/man'
'PATH=/bin:/usr/bin:/usr/sbin:/export/opt/emc/SYMCLI/V6.0.3/bin:/export/opt/Navisphere/bi
n:/export/opt/HiCommandCLI:/export/opt/HORCM:/export/opt/TieredStorageManager/current/CLI
' 'PRINTER=lp' 'PS1=cc8385#' 'TERM=vt100' 'TMPDIR=/dev/vx/rndmp' 'TZ=US/Eastern'
rungroup        '!g!'
rungroups[]     '!G!'
runhost         'f5-v440-01'
runlocalmode    0
runnice 0
runptyflags     7
runtimeout      0
runumask        18
runuser         'root'
status 0
submithost      'sangs11'
submithostip    '10.40.42.49'
subprocuser     'root'
time            '08:07:49'
timezone        'EST5EDT'
true 1
ttyname         '/dev/pts/2'
umask 18
uniqueid        '0a282a214667f51549F3'
user            'cc83851'
year 2007
'*/event        'Finish'
exitdate        '2007/06/07'
exitstatus      'Command finished with exit status 0'
exittime        '08:08:28'
uniqueid        '0a282a214667f51549F3'
```

25.6.6 CLI Operations

The most commonly used CLI commands relating to managing migrations are list below (non-exhaustive list). The names are on the whole self-explanatory.

- Refresh (refreshes the storage domain)

- CreateMigrationGroup
- DeleteMigratonGroup
- GetMigrationGroups
- ModifyMigrationGroup
- AddVolumeToMigrationGroup
- RemoveVolumeFromMigrationGroup

- CreateStorageTier
- DeleteStorageTier
- GetStorageTiers
- ModifyStorageTiers

- CreateMigrationPlan
- CreateMigrationTask
- GetTasks
- ExecuteTask
- DeleteTasks

There are also additional commands relating to shredding tasks, managing storage domains and other task management operations.

Full descriptions, with examples can be found in the excellent **HiCommand Tiered Storage Manager CLI Reference Guide** and little further detail is given here, bar a few simple examples.

<https://catecollaboration.citigroup.net/domains/platstor/gdse/products/platforms/Tagmastore/hc0919.pdf>

It is important to understand the following key points relating to migration task management via the CLI.

1. With the CLI it is necessary to create a migration plan before a migration task can be created. The migration plan is a text file that defines the characteristics of the migration. There is no such concept of a migration plan with the web client; a migration task is created directly.
2. It is not possible to use "CreateMigrationPlan" if the user is already logged on via the web client.
3. With the web client, editing of the source-target device pairs is done via an interactive GUI. With the CLI, it is necessary to edit the migration plan file with a text editor. For example, if you do not like the default migration targets recommended by TSM, you will need to comment out the recommended targets and uncomment the ones you want before creating the migration task. Targets in the migration plan begin with "targetControllerDeviceNumber". Changing source-target pairings is much easier with the web client as this can be done by simply selecting new target devices.
4. With the CLI, a migration task is created with the CreateMigrationTask Command, using the Migration Plan file as an input.
5. Tasks created with the CLI can be managed with the web client and viceversa. For example, it is possible to create a task with the web client and execute via the CLI (and viceversa).
6. If you do not execute a CLI generated task immediately (using the --execute option for CreateMigration Task), you will need to check (or parse if generating tasks in a script) the

output of “CreateMigrationPlan” to capture the task-id, as this is needed by “ExecuteTask”. Alternatively, one can later use “GetTasks” to get the id of the required task, using suitable filters such as “migrationgroupname”.

25.6.7 CLI migration example

The following show example commands for common migration tasks. In practice, all would be run via PowerBroker (preface each command with /opt/pb/bin/pbrun). Username and password arguments are not shown for clarity. It is assumed storagedomainname and htmsserver location are set in the htmscli.properties file and hence also are not on command line.

1) Create and populate migration group

```
# htmscli CreateMigrationGroup name="Test04" description="CLI generated MG"
# htmscli AddVolumeToMigrationGroup name="Test04" ControllerDeviceNumber=0:07
# htmscli AddVolumeToMigrationGroup name="Test04" ControllerDeviceNumber=0:08
```

2) Create migration plan

```
# htmscli CreateMigrationPlan -o MigPlanTest04 migrationgroupname="Test04"
targetstoragetiername="RAID5 146GB"
```

3) Review MigPlanTest04 with editor to confirm target volumes (uncommented targetControllerDeviceNumber= entries) are acceptable, and edit to taste

4a) To create migration task from the migration plan and execute immediately:

```
# htmscli CreateMigrationPlan --execute migrationplan="MigPlanTest04"
```

4b) Or to create migration task for delayed execution:

```
# htmscli CreateMigrationPlan migrationplan="MigPlanTest04"
```

5) To execute a previously created task

```
# htmscli ExecuteTask id=<task-id>
```

25.6.8 CLI shredding tasks

If for some reason shredding cannot be done at migration time (or more likely you forgot to select the erase option), or DoD shredding is required for some exceptional reason, then this can be done via a separate CLI shredding task. However, be aware that after a migration, the addresses of the source and target volume are swapped and hence you would need to shred the LDEVs with the addresses of the TARGET devices that were selected from the Storage Tier for the migration. For example, if you migrate from

01:00 → 02:1F

after the migration, the old data to be erased will be on 02:1F as the final stage of a migration task is to swap the addresses of the source and target so that the address of the LDEVs seen by the host never changes (e.g. 01:00 in this case).

Hence, if you do a migration with Migration Group XYZ and then do a shredding with of Migration group XYZ, you will be attempting to erase the production data, as the source and target LDEV

numbers will have been swapped (in practice, TSM will not allow you to create a shredding task for disks with paths). Instead you need to create a Migration Group that contains the LDEV numbers that were originally the target devices of the migration. If you erase as part of the migration task, this confusion never comes up as the erasure is handled transparently.

25.6.9 CLI shredding example

The following show example commands for common migration tasks. In practice, all would be run via PowerBroker (preface each command with /opt/pb/bin/pbrun). Username and password arguments are not shown for clarity. It is assumed storagedomainname and htmsserver location are set in the htmscli.properties file and hence also are not on command line.

1) Create and populate migration group. These volumes must not have a path to a host.

```
# htmscli CreateMigrationGroup name="Shred" description="Shredding MG"
# htmscli AddVolumeToMigrationGroup name="Shred" ControllerDeviceNumber=2:01
# htmscli AddVolumeToMigrationGroup name="Shred" ControllerDeviceNumber=2:02
```

2) Create shredding task. --force is required if you do not want the CLI to prompt for confirmation when used in a script.

```
# htmscli CreateShreddingTask --force migrationgroupname="shred" shreddingmethod=ZERO-ONCE
```

3) Execute shredding task

```
# htmscli ExecuteTask id=<task-id>
```

25.7 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.

25.7.1 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.
- ZPR cannot be run after a normal to DP pool migration if a replication relationship is in place. If ZPR needs to be run, replication relationship must be removed first.
- 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 UR?
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>GDSE 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

25.7.2 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 HTSM 6.2 and 60-05-11 microcode or later.
- 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 its 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 **can not be** a thick volume.

25.7.3 Universal Replicator (UR) specific comments

- UR 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 6.2.
- Split off the ShadowImage copy of the S-VOL after replication is suspended but before migration is undertaken.

After suspending replication, it is necessary to refresh first Device Manager (DM) and then Tiered Storage Manager (TSM) to update flags in TSM that the volumes are no longer being replicated. This will mark the volumes as available for migration. The refreshes must be executed for both the replication source and target arrays. Note that, if Replication Manager is used for the suspension, then it is still necessary to refresh the DM and then the TSM databases.

- Set up the migration in TSM. Note that it is not possible to set up a migration task before replication is suspended and the DM and TSM 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 TSM for P-VOL array
 3. refresh DM and then TSM for S-VOL array
 4. once the replication state is PAIR, resync your ShadowImage copy of the S-VOL

Note that, if Replication Manager is used for the suspension, then it is still necessary to refresh the DM and then the TSM databases.

25.8 Using HTSM with Dynamically Provisioned storage

25.8.1 Introduction

Use of TSM is supported with Dynamically Provisioned (DP) storage. There are five migration scenarios involving DP volumes, of which only three are recommended (see the previous section for support matrix for using TSM with replicated DP-volumes).

Source is a DP Volume		When migrating <u>from</u> a DP volume, and if the “Erase remaining data on source volumes” is selected, then the target LDEV (or the LDEV created out of the target DP pool by the migration process) is deleted at the end of the migration, returning space to the DP pool for reuse.
Source	Target	Comments
DP Volume	DP Pool	<i>Migration to DP Pools is the preferred DP → DP migration process.</i> The <u>temporary</u> LDEV created by the migration is deleted at the end of the migration, resulting in a “clean” migration.
DP Volume	DP Volume	Testing has found a number of issues with migration of DP volumes to existing DP volumes (see Release Notes for further details). <i>It is recommended that DP volumes be migrated to DP pools rather than DP volumes.</i>
DP Volume	Non-DP Volume	Recommended method when target is non-DP volume. The target LDEV is deleted after the migration has occurred, leaving a gap in the LDEV address range. This can only be avoided by <u>not</u> selecting the “Erase remaining data on source volumes”, in which case the data will need to be separately erased from the target LDEV (former source storage) using the CLI (see Section 23.6.8).

Source is a non-DP Volume		When migrating <u>from</u> a non-DP volume, if the “Erase remaining data on source volumes” is selected, then the target LDEV (or the LDEV created out of the target DP pool by the migration process) is <u>not</u> deleted at the end of the migration.
Source	Target	Comments
Non-DP volume	DP Pool	Recommended method when source is non-DP volume. Consumed capacity of the DP volume after the migration will be 100%, and it will be necessary to run the Zero Data Removal Tool to reclaim zero-page data from the volume back into the DP pool. Migration from a Non-DP volume to a DP pool will result in the creation of a permanent new (non-DP) LDEV.
Non-DP volume	DP Volume	Alternative method (not recommended) Consumed capacity of the DP volume after the migration will be 100%, and it will be necessary to run the Zero Data Removal Tool to reclaim zero-page data from the volume back into the DP pool. Migration from a Non-DP volume to a DP volume will <u>not</u> result in the creation of a permanent new (non-DP) LDEV. However, such a migration requires that a DP-volume-based Storage Tier be built for each DP pool, and hence is more complex than just migrating to a DP-pool-based Storage Tier.

25.8.2 Migrations of DP volumes to DP pools

25.8.2.1 Introduction

Migrations of DP volumes to DP pools (rather than DP volumes) are recommended to ensure that storage from the source volume is returned to the storage pool after the migration.

25.8.2.2 Building Storage Tiers with DP pools

Tiers should be built using the **Pool Filter** in the “Create Tier” dialog box. It will be necessary to first identify with Device Manager (DM) which DP pools need to be associated with which tiers. For example, when building a “thin” storage tiers built from RAID 5 300GB disks, one first needs to identify which DP pools are built from 300GB disks configured as RAID5.

If you do not already this documented, one way to find this information from within TSM, is to begin to create a temporary storage tier using the volume filter and the following filter conditions (for the above example):

DynamicProvisioning = 'DP-Pool-VOL' AND DiskCapacityInGB > 250 AND DiskCapacityInGB < 350 AND RAIDLevel = 'RAID5(7D+1P)'

Click on “Search” and all DP-Pool_VOLs that are built from 300GB RAID5 will be shown in the volume list. This listing also shows the Pool ID for each of the DP-Pool_VOLs, and hence it is possible to identify all the DP pools that need to be included in a given storage tier. Once the DP pools associated with each proposed storage tier have been noted (by varying the filter condition), the creation of this temporary storage tier should be cancelled.

Create Tier - Microsoft Internet Explorer

HTSM HITACHI

Create Tier Help

Properties

Use this form to create your Storage Tier.

*Name:

Description:

***Filter Conditions**

Edit preferred filter condition below.

Formula:

Select filter condition:

Show entries matching of the conditions below:

Attribute: is

Disk Capacity > GB

Disk Capacity < GB

RAID Level

* Required field

Volume List (Total:38)

Rows/Page: 25 Page: 1 of 2

LDEV	VOLSER	Dynamic Provisioning	Pool ID	Cons. Capacity
00:02:11		DP-Pool-VOL	DP 47	
00:02:12		DP-Pool-VOL	DP 47	
00:02:13		DP-Pool-VOL	DP 47	
00:02:14		DP-Pool-VOL	DP 47	
00:02:15		DP-Pool-VOL	DP 3	
00:02:16		DP-Pool-VOL	DP 3	
00:02:17		DP-Pool-VOL	DP 3	
00:02:18		DP-Pool-VOL	DP 3	
00:02:19		DP-Pool-VOL	DP 3	
00:02:1A		DP-Pool-VOL	DP 3	
00:02:1B		DP-Pool-VOL	DP 3	
00:02:1C		DP-Pool-VOL	DP 3	
00:02:1F		DP-Pool-VOL	DP 0	
00:02:20		DP-Pool-VOL	DP 0	
00:02:21		DP-Pool-VOL	DP 0	

Identifying which DP pools are associated with each disk size and RAID type

Single tiers that include ALL DPs that comprise a tier (or storage class) should then be built using:

Filter Condition:	Pool Filter
Show entries matching:	Any
Filter condition formula:	PoolID = 'DP 0' OR PoolID = 'DP 3' OR PoolID = 'DP 47'

In the above example, this Storage Tier comprises of the three DP pools “DP 0”, “DP 3” and “DP 47”. Note the gap between the “DP” and the pool number (e.g. DP 47). Note that you will probably have only one DP pool per tier, but thus documentation needs to cover potential cases where there are more.

Create Tier

Use this form to create your Storage Tier.

*Name:

Description:

*Filter Conditions

Edit preferred filter condition below.

Formula:

Select filter condition:

Show entries matching of the conditions below:

Attribute: Value:

is

is

is

Reset Search

* Required field

Pool List (Total:3)

Pool ID	Status	Capacity		Threshold1
		Total	Free	
DP 0	Normal	272.58 GB	254.41 GB	70%
DP 3	Over Threshold	825.80 GB	94.37 GB	70%
DP 47	Normal	134.28 GB	116.11 GB	70%

Creating a DP pool based Tier

25.8.2.3 Migrating to DP pools

Migrations should be configured from the HTSM GUI using the Migrate MG dialog box.

When executing a migration, it will always initially give the following “error” in the “Profile for Migration” pane when you click on “Check Compatibility”

“The migration plan cannot be created because none of the volume pairs in the migration plan are volume subject to migration”

and you will note that there will be no pool identified under the “Target Pool ID” column for any devices in the MG (see figure below).

Storage Domain: USP10102

Migration Group(Source):
 eliot2
 HUR_vol_0413
 multi
 TestUSPV
 TestUSPV_2

Storage Tier(Target):
 RAID5_146
 RAID5_300
 Thin_RAID1_146GB
 Thin_RAID5_146GB
 Thin_RAID5_300GB

Information :
 multi:

Information :
 Thin_RAID5_300GB:

Check Compatibility

Migration Group(Source): multi → Storage Tier(Target): Thin_RAID5_300GB
 Time Estimate 0:00:00 (+ 0:00:00)
 Compatibilities:

Profile for Migration

The migration plan cannot be created because none of the volume pairs in the migration plan are volumes subject to migration.

Configuration	No. of Free Volumes in Tier		Difference
	Needed	Available	
OPEN-V, 34.69 GB (72750720 Blocks)	-	-	-

Migration Pairs:

Source	Vol. ID	LDEV	Target Pool ID	Status	Subsystem		Ctrl. Array Group		Array Group		AG B
					Source	Target	Source	Target	Source	Target	
	00:FE:10	00:FE:10	-		HDS10102		X4-1		X4-1		
	00:FE:11	00:FE:11	-		HDS10102		X4-1		X4-1		

Edit Pairs Auto-match Pairs

Setting up a migration to a DP pool

This is easily addressed by clicking on the “Edit Pairs” button (marked on figure above). Indeed, until you go through this procedure to select a pool, it is not possible to click on “OK” in the “Migrate MG” window to submit the migration plan.

Clicking on the “Edit Pairs” button will bring up the “Edit Migration Pairs” dialog box (see figure below). You then need to identify the pairs to be edited (all devices in the MG), pick a DP pool from this Storage Tier (in the lower panel), and then clicking on “Replace Pair” button. All devices in the migration group can be paired with a target DP in a single operation. Note that one still has to go through this process even if the Tier only had a single DP pool.

In the example below, all volumes in the MG (devices FE:00 and FE:01) have been selected (left-most checkbox ticked) and “DP 47” selected in the lower panel. The user now needs to click on “Replace Pair” to define the migration as being to DP47 for all volumes in the MG, and then click on “OK”

Edit Migration Pairs - Microsoft Internet Explorer

HTSM HITACHI

Edit Migration Pairs

Migration Group(Source): multi Storage Tier(Target): Thin_RAID5_300GB

Current Pair(s): select one or more pairs to replace or uncouple.

	Source		Target	Status	Subsystem		Ctrl. Array Group		Array Group		AG Bus	
	Vol. ID	LDEV	Pool ID		Source	Target	Source	Target	Source	Target	Source	Target
☒	00:FE:10	00:FE:10	-		HDS10102		X4-1		X4-1			
☒	00:FE:11	00:FE:11	-		HDS10102		X4-1		X4-1			

Do Not Migrate

Available Pool(s): select a pool to replace the current pair(s) Replace Pair

	Pool ID	Status	Capacity		Threshold1	Threshold2	Usage Rate	CLPR	Volumes
			Total	Free					
☐	DP 0	Normal	272.58 GB	254.41 GB	70%	80%	6%	0	15
☐	DP 3	Over Threshold	825.80 GB	94.37 GB	70%	80%	88%	0	3
☒	DP 47	Normal	134.28 GB	116.11 GB	70%	80%	13%	0	1

OK Cancel

Selecting a DP pool as the target for a migration

It is recommended that the “**Erase remaining data on source volumes**” option always be selected when migrating from DP volumes to ensure former source storage is returned to the DP pool after it has been erased.

Migration Task Settings:

Task Description:

Notification Email:

☒ Execute immediately after confirmation.

☒ **Erase remaining data on source volumes.**
(Not choosing this option will leave the source volume capacity consumed.)

If the “Erase remaining data on source volumes” option is selected, then:

- a) The temporary target LDEV created by the migration process is removed once the migration is completed.

- b) The original source storage (with address of target device) is scrubbed.
- c) The original source is returned to its DP pool for reuse (if migration is from a thin device).

These are all considered highly desirable actions.

25.8.3 Migration of DP volumes to DP volumes

Testing has found a number of issues with migration of DP volumes to existing DP volumes (see release notes for further details). As a consequence, it is recommended that DP volumes always are migrated to DP pool.

25.8.4 Migration of a DP volume to a non-DP volume

The process for migrating a DP volume to a non-DP volume is the same as for migrating a non-DP volume to a non-DP volume, as described earlier in this Appendix. However, note that if the “Erase remaining data on source volumes” is selected, then the target LDEV **is deleted** at the end of the migration. This is because the erase takes place AFTER the LDEV number swap, but which time the “target” LDEV number is now assigned to DP storage, and hence is returnable to its DP pool. This deletion of the LDEV can only be avoided by not selecting the “Erase remaining data on source volumes”, in which case the data will need to be separately erased from the target LDEV (former source storage) using the CLI (see Section 23.6.8).

After a migration from a DP volume to a non-DP volume, you may be required to refresh DM and TSM.

25.8.5 Migration of non-DP volume to a DP pool

The process for migration from a non-DP volume is the same as described above for the migration of a DP volume to a SP pool-based Storage Tier (created as defined above using the Pool filter). However, note that if the “Erase remaining data on source volumes” is selected, then the LDEV created out of the target DP pool by the migration process **is not deleted** at the end of the migration. This is because the erase takes place AFTER the LDEV number swap, but which time the LDEV number created by the migration process is now assigned to non-DP storage, and hence is not returnable to any DP pool.

Note that consumed capacity of the DP volume after the migration will be 100%, and it will be necessary to run the Zero Data Removal Tool to reclaim zero-page data from the volume back into the DP pool.

After a migration from a non-DP volume to a DP pool, you may be required to refresh DM and TSM.

25.8.6 Migration of a non-DP volume to a DP volume

The migration method from a non-DP volume to a DP pool has the advantage that no new LDEVs are created, but does require that a Storage Tier be created that contains DP Volumes (as opposed to DP Pools).

The process for migrating a non-DP volume to a volume-based DP Storage Tier (created as defined below using the Volume filter) is the same as for migrating a non-DP volume to a non-DP volume, as described earlier in this Appendix. However, note that if the “Erase remaining data on source volumes” is selected, then the target LDEV is **not deleted** at the end of the migration. This is because the erase takes place AFTER the LDEV number swap, but which time the LDEV number created by the migration process is now assigned to non-DP storage, and hence is not returnable to any DP pool.

Note that consumed capacity of the DP volume after the migration will be 100%, and it will be necessary to run the Zero Data Removal Tool to reclaim zero-page data from the volume back into the DP pool.

After a migration from a non-DP volume to a DP volume, you may be required to refresh DM and TSM.

25.8.6.1 Building Storage Tiers with DP volumes

For migration of a non-DP volume to a DP volume, a tier that contain a single DP needs to be built using:

Filter Condition:	Volume Filter
Show entries matching:	All
Filter condition formula:	(PoolID = 'DP 1') AND (DynamicProvisioning is DP- Vol)

In the above example, this Storage Tier comprises of the one DP pool “DP 1”. Note the gap between the “DP” and the pool number (i.e. DP 1).

Note that this only supports only one DP pool per DP-Volume-based Storage Tier definition and requires us to treat each DP pool as a separate Storage Tier (TSM does not support the syntax required to built DP-Volume-based Storage Tiers from multiple DP pools).

26 Appendix M: 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

27 Appendix N: Hitachi Dynamic - Thin Provisioning (HDP)

27.1 What is HDP?

Hitachi Dynamic Provisioning (HDP) allows for the virtualization of traditional block storage such that multiple applications can be pooled together to recapture unused storage blocks.

The driver for certification of this technology is to save upfront capital expenses by avoiding or delaying the purchase of additional storage.

Thin provisioning is now the default storage configuration standard for all HDS systems. Business justification is required for deviation from this directive.

27.2 Migration to HDP

27.2.1 Block based Migrations

Applications can be migrated to thin storage using block based tools such as UVM (and TSM). This is the standard and recommended migration method. Once the LUNs have been migrated, the host will then need to be re-pointed to the new LUNs (this is automatic in same-array migrations). Once a block-based migration is complete, all of the pages will be allocated (100% utilized). To reclaim (and return to the pool) the unused portion of the LUNs, the ZPR function will need to be run against the LDEV.

27.2.2 Host Based Migrations

To migrate to HDP on a host level, an existing application must perform a host-based LUN to LUN migration to obtain the thin storage advantages. This will require significant application downtime in some cases. The applications being migrated will:

1. Shutdown updates to the application data. Typically this means the application will be turned off.
2. Once the new LUNs are visible to the host, data should be manually copied from the legacy LUNs. Applications that perform file level copies (such as RoboCopy) can be used. Block level copies should not be used.
3. Once the data has been copied and verified, the application should be configured to use the new LUNs and then returned to normal status. The old LUNs should be returned once the application has fully verified.

Applications that are being initially established do not need to perform this migration.

27.3 HDP Configuration via Storage Navigator


It is noted that Storage Navigator describes the thin volumes that are allocated to hosts as Virtual Volumes (or V-VOLs). Hitachi Device Manager (HDvM) refers to these volumes as Dynamic Provisioning Volumes (or DP-VOLs). Device Manager refers to the volumes that comprise the pool storage as Dynamic Provisioning Pool Volumes (or DP Pool Vols). For instructions on using Device Manager with HDP see [HDP Configuration via Device Manager \(HDvM\)](#).

- Must first install SM for SYS area and PP keys [HDS CE]
- Turn on required system modes (see [System Options and System Disks](#)) [HDS CE]
- Create an HDP pool including one or more LDEVs (pool volumes).
- When creating the HDP pool, set one threshold for free capacity management.


- HDP pool has two thresholds; one is the value mentioned above, the other is fixed value (80%).
- Create Pool with minimum of 4 LDEVs, 1 LDEVs each consuming all of 4 RAID groups (each RG behind different BED pair). Alternatively, you could divide each parity group by 4 LDEVs.
- Create virtual volumes (DP-VOLs) and associate each to a single HDP pool

27.3.1 Create a Pool


1. Go to the Pool tab under LUN Expansion (Storage Navigator Go Menu → LUN Expansion/VLL → Pool). Select Modify Mode.



2. Right click the Dynamic Provisioning folder and select "New Pool" option.



3. Enter pool data and select Set.




The 'Storage Navigator' dialog box, titled 'New Pool', contains the following fields and controls:

- Pool ID:** A text input field containing '001' with a range indicator '(0-127)' to its right.
- Threshold:** A dropdown menu showing '55' with a range indicator '(%)' to its right.
- Attribute:** A dropdown menu showing 'Dynamic Provisioni...'.
- Buttons:** 'Set' and 'Cancel' buttons at the bottom right.

The Pool ID is a unique numeric value to identify the pool (0-127). The threshold is the alerting threshold (default of 70%) for the pool. This threshold value can be changed at a later time. There is a system level threshold of 80% that is not modifiable. Therefore choose an appropriate threshold value other than 80%.

27.3.2 Delete a Pool

In order to delete a pool, you must not have any DP-VOLs(V-VOLs) associated with the pool. Right click on the Dynamic Provisioning folder and select Delete Pools.



The 'Storage Navigator' window, titled 'Pool', shows the following components:

- Top Tabs:** LUN Expansion, Virtual LVI/LUN (VLL), Pool, V-VOL.
- Pool Management Summary (Top Right):**
 - Pool Mng. Capacity : used / total
 - C.O.W. Snapshot : 0.00 GB / 0.00 GB
 - Dynamic Provisioning : 2744.76 GB / 2714717.57 GB
- Pool List (Left):** A tree view showing 'Subsystem' > 'C.O.W. Snapshot' > 'Dynamic Provis' > a list of pools: Pool:0, Pool:1, Pool:47, Pool:101, Pool:102, Pool:120, Pool:121. A right-click context menu is open over 'Dynamic Provis' with options: 'New Pool', 'Delete Pool(s)', and 'Restore Pool(s)'. 'Delete Pool(s)' is highlighted.
- Pool Information (Bottom Left):** A table showing details for the selected pool:

Pool ID	:	
Status	:	
Total Capacity	:	
Used Capacity	:	
Usage Rate	:	
Threshold1	:	
Threshold2	:	
Attribute	:	
CLPR	:	
- Pool-VOL (Right):** A table with columns: LDKC:CU:LDEV, Capacity, Emulation, RAID, PG, Int. VOL Info. Below the table are fields for 'Total Capacity:', 'Used LDEVs:', and 'Remaining LDEVs:'. An 'Add Pool-VOL' button is below these fields.
- Free LDEVs (Bottom Right):** A section with 'Free LDEVs' and a table with columns: LDKC:CU:LDEV, Capacity, Emulation, RAID, PG, Int. VOL Info, Ext. Below the table is a 'Selected LDEVs: 0' field.
- Buttons:** 'Optimize', 'SIM Complete Request' (checkbox), 'Apply', and 'Cancel' at the bottom.


Any pools that meet the criteria will be marked as deleted (blue italics). Unfortunately you cannot select individual pools to be deleted.

27.3.2.1 Associate volumes with Pools


1. To associate a volume with an HDP Pool (i.e. add physical storage to a pool), select the desired Pool-ID. Select the LDKC (00 is only current valid choice) and then CU

address range from the “free LDEVs” list. You will see only volumes that have not been allocated to hosts or used against other pools (or C.O.W. pools).

For future reference, the attribute column on the VLL screen shows if a volume has been assigned to a specific pool. You cannot jump to another screen at this point without losing the changes entered above.




2. Select specific volumes from “free LDEVs” list then select “Add Pool-VOL”. *This is similar to adding LDEVs to a host group.*
3. Select OK on the confirmation screen.
4. Select Apply to make the changes effective.




- When ALL LDEVs have been added click on Optimize. Optimization should take a few seconds up to a max of 20 minutes depending on the number of LDEVs being added to the pool.

27.3.2.2 Changing the Threshold of a Pool

- Go to the Pool tab under LUN Expansion (Storage Navigator Go Menu → LUN Expansion/VLL → Pool). Select Modify Mode.
- Select the Pool to be modified. Right click the Pool ID and select Change Pool Information.




- Change the threshold percentage to the new value. Do not change the value to 80% as that threshold has already been set by the system. Click OK.




4. Select Apply to make the changes effective.

27.3.3 Expanding a Pool

1. Go to the Pool tab under LUN Expansion (Storage Navigator Go Menu → LUN Expansion/VLL → Pool). Select Modify Mode.
2. Select the Pool to be modified.
3. Select the LDKC (00 is only current valid choice) and then CU address range from the “free LDEVs” list. You will see only volumes that have not been allocated to hosts or used against other pools (or C.O.W. pools).



4. Select specific volumes from “free LDEVs” list then select “Add Pool-VOL”. *This is similar to adding LDEVs to a host group.*
5. Select OK on the confirmation screen.
6. Select Apply to make the changes effective.



6. When ALL LDEVs have been added click on Optimize. Optimization should take a few seconds up to a max of 20 minutes depending on the number of LDEVs being added to the pool.


Warning: Optimize should be run only one time after adding additional LDEVs to the pool. Selecting Optimize at a time other than after adding LDEVs will result in an unbalanced page allocation that will no longer spread I/O over all of the LDEVs correctly.

27.3.4 Create DP-VOL (V-VOL)

In order to create a DP-VOL (V-VOL), you have to create a group and then associate it with a pool.

27.3.4.1 Create an HDP DP-VOL Group (V-VOL Group)

1. Go to the V-VOL tab under LUN Expansion (Storage Navigator Go Menu → LUN Expansion/VLL → V-VOL). Select Modify Mode.
2. Right Click of Dynamic Provisioning Folder and select 'New V-Vol Group' Option.



3. Enter V-VOL Group data. Each V-VOL group is a logical container for up to 3,072 GB (or 88x 34.69GB standard Citi volumes) of storage. However, the current methodology is to put only a single V-VOL per V-VOL group. The group number is the next available group number. Emulation type should be Open-V. The CLPR should be 00. The "Copy of V-VOL" field is the number of additional V-VOLs to create. For example, if you want to create 10x V-Vols, then you would put a value of 9 here.

HDS recommends placing a single DP-VOL per V-VOL group to take advantage of LUN expansion after the LUN has been assigned. Citi has not used dynamically sized LUNs in the past as many operating systems do not support them. However, Citi has adopted the methodology to create 1 V-VOL per V-VOL group.

4. Select Next for display of Create V-VOL box


5. Select Set to add to V-VOL list. There will be only one LDEV listed here.

No.	Emulation	Capacity
1	OPEN-V	35522 MB


6. Select Next to create the addresses for the DP-VOLs. This is very similar to creating addresses for standard LDEVs. Using the interval option is not recommended as the pool is already being created from different Parity Groups. There is also a different range of CU:LDEVs for thin volumes. You should see all of the V-Vols being created in the top box.

No.	V-VOL Group	Emulation	Capacity	LDKC:CU:LDEV
7	X 32 - 1	OPEN-V	35522 MB	00:50:06 X
8	X 33 - 1	OPEN-V	35522 MB	00:50:07 X
9	X 34 - 1	OPEN-V	35522 MB	00:50:08 X
10	X 35 - 1	OPEN-V	35522 MB	00:50:09 X

7. Select Next to go to the confirmation screen. Select OK. You will only see the first V-VOL on this list.




8. You will be able to navigate through the V-VOL groups and see all of the new groups and V-VOLs in blue. To make the changes effective, click on the apply button.




Note: The DP-VOLs (V-VOLs) will have a "X" designation after their LDKC:CU:LDEV address. For example "00:FE:50 X"

27.3.4.2 Associate DP-VOL(V-VOL) with a HDP Pool


1. Go to the V-VOL tab (Storage Navigator Go Menu → VLL → V-VOL)
2. Right click on Dynamic Provisioning and select "Associate V-VOL Groups with Pool" from the pop-up window to display "Connect Pool" dialog.




3. Select the V-VOL groups you wish to associate with a pool. Then select Set.



4. Select appropriate pool from the list and then select "NEXT" to change threshold.



5. Select threshold (5%-300% (5% inc)). This threshold is for each DP-VOL(V-VOL). Use the default of 5% as this is not a heavily used threshold.



6. Click "SET".
7. Select Apply to make the changes effective.

The DP-VOLs can now be allocated to a host as a LUN like any LDEV. The DP-VOLs (V-VOLs) will have a "X" designation after their LDKC:CU:LDEV address. For example "00:FE:50 X"

27.3.5 Delete DP-VOL(V-VOL)

To delete the DP-VOL (V-VOL), you have to disassociate the volume from the pool. This will take the volume out of the pool. You then can delete the DP-VOL (V-VOL). Unfortunately, you have to delete the DP-VOLs(V-VOLs) at the group level. Therefore any DP-VOLs (V-VOLs) in the same group will also be deleted.

27.3.5.1 Release DP-VOL(V-VOL) from a HDP Pool

1. Go to the V-VOL tab (Storage Navigator Go Menu → VLL → V-VOL). Right click on Dynamic Provisioning and select Release V-VOL Groups from Pool.

V-VOL Group	CLPR	V-VOL Cnt	LDEV	Pool ID	Threshold
X 21-1	00 : CLPR0	1	00:E0:01 X	Pool:120	5%
X 22-1	00 : CLPR0	1	00:E0:02 X	Pool:120	5%
X 23-1	00 : CLPR0	1	00:E0:03 X	Pool:120	5%
X 24-1	00 : CLPR0	1	00:E0:04 X	Pool:120	5%
X 25-1	00 : CLPR0	1	00:E1:00 X	Pool:121	5%
X 26-1	00 : CLPR0	1	00:E1:01 X	Pool:121	5%
X 27-1	00 : CLPR0	1	00:E1:02 X	Pool:121	5%
X 28-1	00 : CLPR0	1	00:E1:03 X	Pool:121	5%
X 29-1	00 : CLPR0	1	00:E1:04 X	Pool:121	5%
X 30-1	00 : CLPR0	1	00:50:04 X	Pool:0	5%
X 31-1	00 : CLPR0	1	00:50:05 X	Pool:0	5%
X 32-1	00 : CLPR0	1	00:50:06 X	Pool:0	5%
X 33-1	00 : CLPR0	1	00:50:07 X	Pool:0	5%
X 34-1	00 : CLPR0	1	00:50:08 X	Pool:0	5%
X 35-1	00 : CLPR0	1	00:50:09 X	Pool:0	5%

2. Select the V-VOL groups you want to remove from the pool (this will destroy any data on the V-VOLs).
3. You will need to acknowledge that you are deleting the data from the V-VOLs.
4. Select Apply to make the changes effective.

Warning: This will delete any data on that V-VOL and will make it unusable until it is associated with a pool in the future.

27.3.6 Change DP-VOL(V-VOL) Threshold


1. Go to the V-VOL tab (Storage Navigator Go Menu → VLL → V-VOL)
2. Select the V-VOL Group that contains the V-VOL to be associated with an HDP pool.
3. Right click one or more DP-VOLs (V-VOLs) to be associated with the HDP pool, and select "Change threshold" from the pop-up window to display the confirmation dialog.
4. Select threshold (5%-300% (5% inc)). This threshold is for each DP-VOL(V-VOL). Use the default of 5% as this is not a heavily used threshold.

- Click "SET" and then Apply to make the changes effective.

27.3.7 Resetting Thresholds for HDP

To change the status of the Pool to NORMAL:

- Go to the Pool tab under LUN Expansion (Storage Navigator Go Menu → LUN Expansion/VLL → Pool). Select Modify Mode.
- Select the "SIM Complete Request" check box



- Select the Apply button. Select OK to verify the changes.


27.4 HDP Configuration via Device Manager (HDvM)

Using Device Manager to setup and modify HDP is simpler than Storage Navigator and is encouraged.


Note: With HDvM you cannot disassociate DP-VOLs from a Pool (Storage Navigator Only). You can delete DP-VOLs, but this is performed on the array group level. HDvM creates each DP-VOL in its own array (DP-Vol) group.

27.4.1 Create Pool

Go to the Dynamic Provisioning screen (All Storage → select array → Pools → Dynamic Provisioning)



Select the Create DP Pool button to start the Create DP Pool wizard.



Select the Pool ID (0-127) and the user threshold (default of 70%). Select Next to continue.

1. Introduction ▶ 2. Create DP Pool ▶ 3. Select DP Pool VOLS ▶ 4. Confirm ▶ 5. Finish

1. From the LDEV list, select the LDEVs to be registered into the Dynamic Provisioning pool.
2. Click Next to confirm the settings.

☐	LDEV	Label	Capacity	Type	Array Group	Volume Attr	SLPR	CLPR
☐	00:00:5A		34.69 GB	-	1-3-1	CVS	0	0
☐	00:01:05		34.69 GB	-	3-1-1	CVS	0	0
☒	00:01:13		34.69 GB	-	3-2-1	CVS	0	0
☒	00:01:14		34.69 GB	-	3-2-1	CVS	0	0
☒	00:01:15		34.69 GB	-	3-2-1	CVS	0	0
☒	00:01:16		34.69 GB	-	3-2-1	CVS	0	0
☐	00:02:06		34.69 GB	-	1-1-1	CVS	0	0

Filter Filter Off

< Back Next > Cancel Help

Select the LDEVs that will comprise the pool. Select Next to continue.

Note: The LDEVs are normally very large (total Parity Group capacity or divided by 4) and should come from dedicated parity groups.

1. Introduction ▶ 2. Create DP Pool ▶ 3. Select DP Pool VOLS ▶ 4. Confirm ▶ 5. Finish

Please check the values below, select the confirmation checkbox, and then click Confirm.


Properties	
Pool ID	0
Threshold1	70%

LDEV	Label	Capacity	Type	Array Group	Volume Attr	SLPR	CLPR	EMU Type
00:01:13		34.69 GB	-	3-2-1	CVS	0	0	OPEN-V
00:01:14		34.69 GB	-	3-2-1	CVS	0	0	OPEN-V
00:01:15		34.69 GB	-	3-2-1	CVS	0	0	OPEN-V


☒ Yes, I've confirmed information above.

< Back Confirm Cancel Help

Confirm the information and then select the confirmation check box. Select Confirm to create the new DP Pool.




Once the Pools have been created, you will get a successful message like:




Select the finish button to complete the wizard. The new pool should now be in the pool list.

27.4.2 Delete Pool


Go to the Dynamic Provisioning screen (All Storage → select array → Pools → Dynamic Provisioning). Select the pool you wish to create.




Select the Delete DP Pool button.




Once you select OK, the Pool Delete function will proceed.



Once completed, you should see a successful message like:




If you try to delete a DP pool that has DP-VOLs associated with it, then you will get the following error:




You will need to delete the DP-VOLs (or disassociate them from the pool) before attempting to delete the DP Pool.

27.4.3 Expanding a Pool


To expand a DP Pool, go to the desired DP screen (All Storage → desired array → Pools → Dynamic Provisioning → desired pool).



Select the Expand DP Pool button to start the wizard.




Select the volumes that are to be used on expanding the pool (these volumes should be from dedicated parity groups and should be very large). Select OK to continue.



Once you have confirmed the choice of LDEVs to be added to the pool, select the OK button.

Note: Adding volumes to a pool cannot be undone without deleting the entire pool.


Once completed, you will get an acknowledgement screen:




27.4.4 Create DP-VOL

To create a DP-VOL, go to the desired DP screen (All Storage → desired array → Pools → Dynamic Provisioning → desired pool).


When you create DP-VOLs, it does not prompt you to specify the address of the DP-VOL (it goes to an Auto address that is the next available address). You must select the Auto address to manually address each DP-VOL. Once you have selected the DP Pool that you want to use to create DP-VOLs, go to the DP-VOL Tab.




Selecting the DP-VOLs button will bring up a wizard dialog box that will step you through the process. After the introduction you will be prompted for the size (MB,GB,TB), volume threshold (use default 5%), quantity, and emulation type (always OPEN-V). You can change the size of each DP-VOL on a later step.



Select Next to continue.



To change the size or address of any of the DP-VOLs, select the blue Auto address.




The address can be changed to a known address or left at Auto. Auto will allow the system to select the next available address. The Capacity, Threshold, and Emulation can be changed to a different value if desired. Select OK to make the changes effective or Cancel to remove your changes. Select Next to continue to the confirmation screen.

LDEV	Capacity	Threshold	EMU Type
00:FD:00	35522 MB	5%	OPEN-V
Auto	35522 MB	5%	OPEN-V
00:FD:01	35522 MB	5%	OPEN-V

Check the Yes box and select Confirm to continue. There will be a processing dialog box while the volumes are being created.


Command submitted for processing. This operation may take several minutes.

You will get a summary screen if successful:



HDvM creates each DP-VOL in its own DP-VOL Group. This will allow you to delete these DP-VOLs at a later date without affecting any other DP-VOLs. This also allows you to take advantage of dynamic expansion of the DP-VOL (not all OS support this).

If you specified addresses that were already in use, then you will get the following error screen:




You can go back to the previous screens and change the address(es) that are in conflict.


27.4.5 Delete DP-VOL

To delete DP-VOLs, select the DP Pool (All Storage > Array > Pools > Dynamic Provisioning > Pool) and then the DP-VOL Tab.

Note: When deleting DP-VOLs, the DP-VOL group must be deleted. You cannot disassociate a DP-VOL from the Pool (like you can in Storage Navigator). This is not an issue if you have 1 DP-Vol per array group.




Check the box next to the DP-VOLs you wish to delete and select the Delete DP-VOLs button.




Click OK to delete the volumes.


You will get a dialog box during the delete and finally an operation complete dialog box if successful.



If there are other DP-VOLs in the same DP-VOL group, you will get a warning message that those volumes will be deleted as well.



If you have allocated any of the selected DP-VOLs to hosts, you will not be able to delete the volumes.



27.5 SNMP Based Alerts for HDP

To generate alerts for HDP thresholds, you will configure SNMP to send traps to HDvM. The thresholds will trigger a SNMP trap. The SNMP trap (when received by HDvM) will generate an email to an individual or distribution list with a predefined message.

27.5.1 Setting up SNMP for Threshold alerts

SNMP is used to send trap messages to an SNMP infrastructure. SNMP must be configured to send traps to the HDvM server in order for HDvM to generate emails. For instructions on using SNMP alerts from the SVP to an enterprise alerting system see Section [Enterprise Alerting \(SNMP\)](#). For details about the SNMP traps, see [HDP Alerts via SNMP Traps](#).

27.5.2 Setting up HDvM for Threshold alerts

HDvM can be set up for sending email alerts to an SMTP email system for array SIMs (including HDP Thresholds). Using email for alerts is less reliable than SNMP traps coming from the array (SVP) and should only be used as a secondary (or backup) means of alerting. HDvM is alerted via SNMP from the array, and then sends out an email based on that alert.

Note: For email alerts to work, you must have SNMP traps set up to be sent to the HDvM server.

The following are sample settings that need to be changed to enable email based alerts:

```
C:\Program Files\HiCommand\DeviceManager\HiCommandServer\config\server.properties
server.mail.enabled=true
server.mail.from=GDSE@citigroup.com           <= use CAPS on address
server.mail.smtp.host=199.1.2.3               <= use your authorized SMTP server
server.mail.smtp.port=25
server.mail.smtp.auth=false
server.mail.alert.type=Trap
server.mail.alert.status=Moderate
```

The email address should be in all CAPS to be SMTP compatible

```
C:\Program Files\HiCommand\DeviceManager\HiCommandServer\config\dispatcher.properties
server.dispatcher.daemon.receiveTrap=true
server.dispatcher.traps.purgePeriod=0
```

```
C:\Program Files\HiCommand\DeviceManager\HiCommandServer\config
\customizedsnmptrap.properties
customizedsnmptrap.customizedSNMPTrapEnable = true
```

The Device Manager Server (with Common Services) needs to be restarted after making the above changes.

The following is a sample of the email (default format) that is generated by HDvM when it receives a SNMP trap from the USPv:

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.

Note: the Source field is the name of the array designated in Device Manager. To get a better name (i.e. with serial number) update the name of the array in Device Manager.

To modify the contents of the email message, there is a configuration file that can be changed. The following is a sample of how you could modify the email message.

```
C:\Program Files\HiCommand\DeviceManager\HiCommandServer\config\mail-alert-
detection.txt
Subject:[DVM] Alert Notification
```

The following alert occurred.

MessageID: \${messageID}
Alert Type: \${alertType}
Source: \${source}
Status: \${status}
Component: \${component}
Description: \${description}
Recommended Action: \${recommendedAction}
Additional Info: \${additionalInfo}
Occurrence Time: \${occurrenceTime}

This message was sent automatically by the Device Manager server.

```
*****
**                                     **
**   This automated message has been modified by GDSE.   **
**   HDVM                                                 **
**                                     **
*****
```

27.6 HDP Reporting

There are two good engines for reporting on HDP-Volumes. For an instant snapshot view of HDP, Device Manager is an easy view of the data. For more extensive reporting, Tuning Manager should be used. However, Tuning Manager software will need to be installed and configured.

27.6.1 Reporting on HDP Volumes via Device Manager

Device Manager provides some basic level reporting on the HDP environment. However, the reporting is only screen based and is not automated at this time.

Note: Storage Navigator describes the thin volumes that are allocated to hosts as Virtual Volumes (or V-VOLs). Hitachi Device Manager (HDvM) refers to these volumes as Dynamic Provisioning Volumes (or DP-VOLs). Device Manager refers to the volumes that comprise the pool storage as Dynamic Provisioning Pool Volumes (or DP Pool Vols).


Log into HDvM. Go to All Storage and drill down to the appropriate USP array. Under a USP array that has Dynamic Provisioning configured, there should be a Pool storage area. Select the Pool storage to obtain a view of all of the Pools.

Pool ID	Capacity	Free Capacity	% Used	Threshold1	Threshold2	Num DP Pools
DP 1	238.01 GB	174.80 GB	26%	25%	80%	
DP 101	134.28 GB	126.08 GB	6%	10%	80%	
DP 102	99.70 GB	99.70 GB	0%	70%	80%	

This screen-only information shows the following information about each pool:

- Capacity (in GB)
- Free Capacity (in GB)
- % Used
- Threshold 1 Percentile (User-Based)
- Threshold 2 Percentile (Fixed System at 80%)
- Number of DP Pool Vols (Volumes that compose Pool)
- Number of DP-VOLs (Virtual Volumes created from Pool)
- Capacity of DP-VOLs (Virtual Capacity)
- Current Status (i.e. Normal, Over Threshold)

Selecting one of the pools will drill down to the next level, showing the volumes that make up the pool (DP Pool Vols) and the virtual volumes that are allocated to hosts (DP-VOLs).



This screen-only information shows the following information about the pool:

DP Pool Vol Tab:

- LDEV
- Capacity
- Array Group

DP-VOL Tab:

- LDEV
- Host (with WWN)
- Port
- LUN (host LUN ID)
- Virtual Array Group
- Capacity

27.6.2 Reporting on HDP Volumes via Tuning Manager (HTnM)

Tuning Manager provides more comprehensive reporting, but requires additional software to be installed and configured. With Version 6.0 and later, this software can reside on the same server as Device Manager. You will also need SAN access to a command device visible to Tuning Manager (directly or indirectly).

For more details on the setup and use of Tuning Manager see [Appendix Q: Tuning Manager \(HTnM\)](#).

27.7 Citi HDP Best Practices

27.7.1 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.

GDSE Thin Provisioning Rule #1: The maximum supported amount of virtual storage that GDSE 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.

GDSE 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.

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).

27.7.2 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

- 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.
- The LDEVs that should be used as POOL Volumes are in the CU:LDEV range of 40:00 – 5F:FF. For more details on LDEV ranges see [Logical Device \(LDEV\) Numbering](#).
- The LDEVs should be created as DP-VOLs that are in the CU range of 60:00 – 9F:FF. For more details on LDEV ranges see [Logical Device \(LDEV\) Numbering](#).
- 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:

	<u>60-03-06</u>
Maximum number of LDEVs comprising a Pool	1024
Maximum size of LDEV comprising a Pool	4TB
Maximum number of DP-VOLs from a Pool	8192

Maximum size of a DP-VOL	4TB*
Maximum number of Pools per array	127

**Note: The max LUN size for Citi approved configurations is 2TB*

27.7.3 Pool ID Numbering

With the use of Tiered Storage Management, DP-VOLs can come from any pool on the array and can change pool assignments with a migration. However, a given DP-VOL will show what Pool ID it is assigned to. This allows us the opportunity to associate a Pool ID with a classification of storage.

Use the following chart to assign Pool IDs:

Pool ID	Drive Type	RAID Type
0-9	4Gb FC 146GB (15k)	R5
10-19	4Gb FC 300GB (15k)	R5
20-29	(Future Use - FC)	
30-39	(Future Use - FC)	
40-49	(Future Use - FC)	
50-59	4Gb SATA 1000GB (7.2k)	R6
60-69	SATA (Future Use – SATA)	
70-79	SATA (Future Use – SATA)	
80-89	SATA (Future Use – SATA)	
90-99	SATA (Future Use – SATA)	
100-109	SAS (Future Use – SAS)	
110-119	SAS (Future Use – SAS)	
120-128	SAS (Future Use – SAS)	

27.7.4 Pool Creation

- Create Pool with 4-8 PGs (based on 1 LDEVs per PG). Older pools may have been created with 4 LDEVs per PG.
- Optimize after all growth is complete – do not optimize at any other time (this will resort free-list)
- Create DP-Vols (V-Vols). DP-Vols will have “X” designation at end of the displayed CU:LDEV address. Standard LUN (DP-VOL) sizes should be used where possible. There may be a need for larger DP-VOL sizes, but they present additional complexities with future migrations and backups. Custom volume sizes may be appropriate for hosts that would need more than 250GB or that cannot effectively manage a large number of LUNs. The LDEVs should be created as DP-VOLs that are in the CU range of 60:00 – 9F:FF. For more details on LDEV ranges see [Logical Device \(LDEV\) Numbering](#)

HDS recommends placing a single DP-VOL per V-VOL Group to take advantage of LUN expansion after the LUN has been assigned. Citi has not used dynamically sized LUNs in the past as many operating systems do not support them. Citi does not support their use at this time.

27.7.5 Pool Growth

- Upgrade should be multiples of original size of pool (n): n, n*2, n*3, ...

- Upgrade pool by 4 RAID Groups (or multiples of 4) at a time (for best BED/DKA performance). This is a minimum of 32 drives.
- 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 8 BED/DKAs
- If Array back-end was upgraded (4 → 8 BEDs) then 2 pools should be used (each spread over 4 BEDs).
- Optimize only once after growth.

Restriping function needs to be run after any growth (function expected in 2008)

27.7.6 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 by HDS, but not yet certified by GDSE.

27.7.7 DP-VOL Sizing

Standard volume sizes can be used (34.69GB) 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).

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 2TB. For more information on this restriction, see [Maximum LUN size](#).

27.7.8 Thresholds

V-VOL threshold – set to 5% and ignore. The V-VOL threshold is a reverse threshold.

User definable (Pool Threshold 1)

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

Fixed (Pool Threshold 2) – 80%

27.7.9 HDP Alerts

27.7.9.1 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 \(SNMP\)](#).

27.7.9.2 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 \(SNMP\)](#).

27.7.9.3 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 threshold1 (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.

27.7.10 Pool Full Condition

In order to prevent any potential data corruption, the 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 Go > Data Retention Utility > Data Retention and locate the LDEV with the Protect Attribute enabled.

Hitachi Storage Navigator - IP:10.40.53.13 - S/N:53112 - D/N:HDS53112 -

Hitachi Storage Navigator

File Go Help

Logged in as: jh53740

Data Retention

Data Retention

LDKC: 00

CU Group: 00 - 3F

Storage System

- CU 00
- CU 01
- CU 02
- CU 03
- CU 04
- CU 05
- CU 06
- CU 08
- CU 0A
- CU 1E
- CU 28
- CU 2B
- CU 2E
- CU 30
- CU 31
- CU 32
- CU 3E


LDEV	Attribute	Emulation	Capacity	S-VO
04 X	Protect	OPEN-V	20.00 GB	Enabl
07	Read/Write	OPEN-V	34.69 GB	Enabl
11	Read/Write	OPEN-V	34.69 GB	Enabl
19	Read/Write	OPEN-V	34.69 GB	Enabl
1E X	Read/Write	OPEN-V	34.69 GB	Enabl
1F	Read/Write	OPEN-V	34.69 GB	Enabl
24	Read/Write	OPEN-V	34.69 GB	Enabl
29	Read/Write	OPEN-V	34.69 GB	Enabl
2C	Read/Write	OPEN-V	34.69 GB	Enabl
2E	Read/Write	OPEN-V	1.00 GB	Enabl
57 X	Read/Write	OPEN-V	1.00 GB	Enabl
58 X	Read/Write	OPEN-V	1.00 GB	Enabl
59 X	Read/Write	OPEN-V	1.00 GB	Enabl
5A X	Read/Write	OPEN-V	1.00 GB	Enabl
DF X	Read/Write	OPEN-V	2.00 GB	Enabl
E0 X	Read/Write	OPEN-V	2.00 GB	Enabl
E1 X	Read/Write	OPEN-V	2.00 GB	Enabl
E2 X	Read/Write	OPEN-V	2.00 GB	Enabl

Operation : No operation

Read/Write : 17 Read Only : 0 Protect : 1

Expiration Lock : Disable -> Enable Apply Cancel

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](#).

27.7.11 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](#).

27.7.12 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 overallocation 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 JFS2 VxFS	Low (write every 64MB – 65%) High 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.

27.8 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.


27.8.1 Caveats and Restrictions

You cannot run this against volumes that are currently being replicated to (via TrueCopy, Universal Replicator, ShadowImage, or Volume Migrator). This does not apply to volumes that were used at one time by these products.


In addition, you cannot use this tool on volumes that are being used by volume shredding.

27.8.2 Using the Tool

You will need to find out the V-VOL group of the volume you want to run this against. The easiest way is to look in Device Manager or to go to the Basic Information screen on Storage Navigator (File → Basic Information → Dynamic Provisioning Tab).



While in Modify Mode, go to the Lun Expansion screen (Lun Expansion → V-VOL). Find the V-VOL group that has your volume in it. Select your volume and then right-click. Select Discard Zero Data.



Select Yes on the confirmation dialog box. You can select any other volumes you want to run this against as well. Select Apply (and OK to confirmation) to start the process.

It should take approximately 5-10 minutes per volume depending on the background processing available on the array. The volume can be used as normal (read/write) during this process. This is a sequential process (one volume is zero-page discarded at a time). You can also make other changes to the array while this process is underway.

The V-VOL screen can be refreshed. Volumes that are currently being processed have an operation status of “Discarding Zero Data”. You can also see the status on the Basic Information Screen for a more concise view. The operation status column will return to “Ready” once the process has been completed on the volume (and after a screen refresh).

Basic Information Display

Port | LUN | LDEV | LDEV Status | Dynamic Provisioning | Connected Host

Dynamic Provisioning CU or Pool: LDKC:CU Pool: Pool:1

DP-VOL

Capacity	Parity Group	Usage Rate	Threshold	Paths	Operation	Pool ID
34.69 GB	X4-1 - (1)	0 %	5 %	0	Ready	1
34.69 GB	X4-1 - (1)	0 %	5 %	0	Ready	1
34.69 GB	X4-1 - (1)	85 %	5 %	2	Discarding Zero Data	1
34.69 GB	X4-1 - (1)	100 %	5 %	2	Discarding Zero Data	1
34.69 GB	X4-1 - (1)	0 %	5 %	0	Ready	1
34.69 GB	X4-1 - (1)	0 %	5 %	0	Ready	1
34.69 GB	X4-1 - (1)	0 %	5 %	0	Ready	1
34.69 GB	X4-1 - (1)	0 %	5 %	0	Ready	1
34.69 GB	X4-1 - (1)	0 %	5 %	0	Ready	1
34.69 GB	X4-1 - (1)	0 %	5 %	0	Ready	1

Pool-VOL

LDKC:CU:LDEV	Emulation	Capacity	Parity Group

Export Close

Volumes that have not been processed yet will also have an operation status of “Discarding Zero Data”. While in Modify Mode, you can cancel or interrupt the process on queued or actively discarding volumes by right clicking the volume on the V-VOL screen and selecting Stop Discarding Zero Data (and selecting Apply with OK confirmation).

27.9 System Options for HDP (Array Settings)

There are several system options that are needed for HDP 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).

For a detailed list of all of the recommended system options, see [System Options and System Disks](#).

28 Appendix O: Performance Monitor Data Extract Tool

28.1 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.

28.2 Preparing to Use the Extract Tool

28.2.1 Preparing the Array

The array must be configured to save performance data in order for the extract tool to function.

28.2.2 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.

28.2.3 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.

28.2.4 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\instpkgs\ <package name>

You can move the software with the subdirectories to any location desired.

28.3 Using the Extract Tool

28.3.1 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	

28.3.2 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

28.3.3 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.

28.3.4 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.getmondat.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.

28.3.5 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.

28.3.6 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

28.3.7 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
```

28.3.8 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.

28.4 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.

29 Appendix P: 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

29.1 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.

29.2 Setup and Use of SPM

The following are the steps required to setup and use SPM for the TagamaStore Series Arrays:

1. Log into the array via Storage Navigator (HDvM → Resources → Arrays → Physical View)
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

30 Appendix Q: 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.

30.1 Installation of Software

Tuning Manager should be installed on the same server as HDvM if there is enough CPU and memory resources available. Otherwise, Tuning Manager can reside on a separate server. The recommended configuration at Citi is to install HDvM, HTnM, HTSM (and HRpM for replicated sites) on the same server. This server can be a physical server or a VMWare host running Windows 2003.

30.2 Tuning Manager 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.

30.2.1 Separate Tuning Manager and Device Manager Servers

For the case where there is not enough memory or resources to run both HDvM and HTnM on the same server, this section describes the setup procedures. These steps should not be followed on a unified (all HDS server software installed on same) system. Please note that the recommended configuration is to run HDvM, (HTSM,) and HTnM on the same server.

30.2.1.1 Linking Tuning Manager and Device Manager

The following are the example command line steps to link Tuning Manager and Device Manager via a SSO connection:

```
Example configuration:
Device Manager is running on a windows server "ecctest01ms"
Tuning Manager is running on a windows server "ecctest05as"
-----

C:\Program Files\HiCommand\Base\bin>hcmdssrv /status
KAPM06440-I The HiRDB service has already started.
KAPM05007-I Already started service. service-name=HBase Storage Mgmt Web
Service

KAPM05007-I Already started service. service-name=HBase Storage Mgmt Common
Serv
ice
KAPM05007-I Already started service. service-name=HiCommand Suite
TuningManager
KAPM05009-I Already stopped service. service-name=HiCommand Suite
TuningService
KAPM05007-I Already started service. service-name=HiCommand Performance
Reporter

C:\Program Files\HiCommand\Base\bin>hcmdssrv /stop /server TuningManager
KAPM05017-I Succeeded in stopping of service. service-name=HiCommand Suite
Tunin
gManager

C:\Program Files\HiCommand\Base\bin>hcmdssup /host ecctest01 /port 23015
```

```

C:\Program Files\HiCommand\Base\bin>hcndsrep /add /type TuningManager /url
http:
//ecctest05as | ecctest05as:23015/TuningManager/login.do /user tnm_admin /pass
G0ldfinger
The filename, directory name, or volume label syntax is incorrect.

C:\Program Files\HiCommand\Base\bin>hcndsrep
KAPM04515-E Option is invalid. option = /user

usage: hcndsrep{
    /add { /type ProductName [/displayname DisplayName] /url LaunchUrl }
|
    /delete{ /type ProductName [/displayname DisplayName] }
|
    /print { [/type ProductName] }
} [/nolog] /user USERNAME /pass PASSWORD

C:\Program Files\HiCommand\Base\bin>hcndsrep /add /type TuningManager /url
http:
//ecctest05as:23015/TuningManager/login.do /user tnm_admin /pass G0ldfinger
KAPM04556-E Failed to connect to HiCommand Suite Common Web Service.

C:\Program Files\HiCommand\Base\bin>hcmdssup
KAPM04521-E Option is invalid.
Usage: hcmdssup [/nolog] /host hostname /port portnumber [/sslport portnumber]
    option
    /nolog                                Restrain output message to commandline
    /host HOSTNAME                        Hostname or IP-address which
repository
exist
    /port PORTNUMBER                      Portnumber of repository service
    /sslport SSLPORTNUMBER                Portnumber of ssl

C:\Program Files\HiCommand\Base\bin>hcmdssup /host ecctest01ms /port 23015

C:\Program Files\HiCommand\Base\bin>hcndsrep /add /type TuningManager /url
http:
//ecctest05as:23015/TuningManager/login.do /user tnm_admin /pass xxxxxxxx
KAPM04560-I hcndsrep add has succeeded. TuningManager

C:\Program Files\HiCommand>cd ..\..\TuningManager\bin

C:\Program Files\HiCommand\TuningManager\bin>htm-switch-mode -sso
2007-10-18 15:15:47,500 api-INFO-8002: The switch to 'SSO' mode was performed,
p
lease restart HiCommand Tuning Manager Service.

C:\Program Files\HiCommand\TuningManager\bin>cd ..\..\Base\bin

C:\Program Files\HiCommand\Base\bin>hcmdssrv /start /server TuningManager
KAPM06440-I The HiRDB service has already started.
KAPM05016-I Succeeded in starting of service. service-name=HiCommand Suite
Tunin
gManager

```

30.2.1.2 Configuring HDvM to be ready for HTnM

There is a small batch file that needs to be run after the HTnM installation that will prepare HDvM for remote connections made by Tuning Manager (when HTnM and HDvM are not on same server).

```

C:\Program Files\HiCommand\DeviceManager\HicommandServer\tools\htmsetup.bat
=====
==
htmsetup

```

```

=====
==

This tool configures the Device Manager settings for connections from Tuning
Man
ager.
1) Configure the connection settings for Tuning Manager installed on the
remote
machine.
2) Display the current settings.
3) Exit.
Enter a value[1 - 3]: 2

The current settings will now be displayed.

The current connection settings for Tuning Manager installed on the remote
machi
ne are as follows:
DBMS remote connection: disabled
DBMS service port number: not set

This tool configures the Device Manager settings for connections from Tuning
Man
ager.
1) Configure the connection settings for Tuning Manager installed on the
remote
machine.
2) Display the current settings.
3) Exit.
Enter a value[1 - 3]: 1

Configure the connection settings for Tuning Manager installed on the remote
mac
hine. Specify the operation:
1) Enable the remote connection.
2) Disable the remote connection.
3) Cancel.
Enter a value[1 - 3]: 1

Processing to enable the connection from Tuning Manager installed on the
remote
machine will now start.
Specify the service port number that the DBMS uses for remote connection.
Enter the port number [5001 - 65535, default = 24220]:

Now processing...
This might take several minutes.
Addition of Windows Firewall inbound exceptions: finished
Addition of Windows Firewall outbound exceptions: not installed
Enabling of the DBMS remote connection: finished

Note: If there was a Tuning Manager instance running locally, then you may get
the following error:

Enabling of the DBMS remote connection: failed
KAIC00859-E The DBMS cannot be stopped because suite products are connected to
i
t.
KAIC30916-E Configuration of the settings to enable the connection from Tuning
M
anager installed on a remote machine has failed. Some errors occurred.

```

30.2.2 Creating Array Agents (Called RAID Agents)

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).

30.2.2.1 Setup of Agent for RAID for 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 `jpcinssetup` 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>jpcinssetup 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 `jpcinlist` command. You must specify the agent type (RAID agent – `agtd`)

```

C:\Program Files\HiCommand\TuningManager\jp1pc\tools>jpcinslist 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 `jpctrl` command.

```

Correctly functioning processes:
C:\Program Files\HiCommand\TuningManager\jp1pc\tools>jpctrl 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
Inactive*
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)

```

30.2.2.2 Setup of Agent for RAID for Solaris 10

Typically this setup is performed on a Solaris 10 host when the HTnM server does not have SAN connectivity to a given array.

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 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).

To install the Agent For RAID:

```

You must run this script by executing the full path. You must also give it a
temporary directory to use for unpacking.

#/var/opt/HTnM_AFR_60.A/install.sh /temp
Begin installation of Hitachi Tuning Manager - Agent for RAID - 6.0.0(6.0.0-
00)
Fri Jul 10 18:56:33 SGT 2009

Hitachi Tuning Manager cannot be installed after you install this product. If
you want to install Hitachi Tuning Manager and this product on the same host,
you must install Hitachi Tuning Manager first.

```


If you should want to uninstall this product later on, please note that the product versions of all other Agents on the same host must be 5.7 or later. To uninstall this product, upgrade the Agents to version 5.7 or later, or uninstall the Agents whose versions are 5.5 or earlier before uninstalling this product.

```
Do you want to continue the installation?
To continue, enter Y. (Default Y) Y/N > y
> Install version ok [ 060000 ]
> /opt/jp1pc/tools/jpcctrl: No such file exists or permission was denied.
> Stop NNM services.
> /opt/OV/bin/ovstop: No such file exists or permission was denied.
> /usr/bin/mkdir /opt/jp1pc: Completed.
Installing The Common Files.
```

{ Some details of installation removed from example}

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 Jul 10 19:00:28 SGT 2009

Done installing Hitachi Tuning Manager - Agent for RAID.

The package name installed should be HTMRAID.

To install the RAID Manager Library (RMLIB):

Run RLinst.sh /opt/HTnM (you have to specify the directory where to place the RMLIB directory. This directory must already exist (you may have to provide the directory again if prompted). Run rmlibinstall.sh after the setup to make sure the symbolic link is in place.

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. After zoning and allocating the command device, you should be able to see the device using the cfgadm command:

```
#cfgadm -al -o show_SCSI_LUN
c1::50060e800563af00,1      disk      connected   configured   unknown
c2::50060e800563af10,1      disk      connected   configured   unknown
```

If the devices do not show up associated with the correct WWN (of the CL port), then your zoning and/or WWN of the host on the array. You should be able to see the raw device in format after running `devfsadm`. Do not add the raw devices to Veritas VxVm or a ZFS pool.

Use `jpctdlistraid` to list the command devices visible to the server.

```
[/opt/jp1pc/tools]
sangs9# jpctdlistraid
KAVF18700-I The detection of the monitorable storage subsystem has begun.
"PRODUCT"          ,"SERIAL"    ,"LDEV"      ,"SLPR" ,"PORT" ,"DEVICE_FILE"
"USP_VM"            ,"25519"    ,"00:3E:0C" ,"0"    ,"CL1-
A", "/dev/rdisk/c1t50060E800563AF00d1s2"
"USP_VM"            ,"25519"    ,"00:3E:0C" ,"0"    ,"CL2-
A", "/dev/rdisk/c2t50060E800563AF10d1s2"
```

Note: If this host is also used as a solutions enabler server or has other LUNs, you may also get some errors:

```
KAVF18742-W An error occurred during system call or OS command execution.
(name=open, rc=-1, error=2)
KAVF18743-W An attempt to get the device information has failed. (device
name=/dev/rdisk/emcpower4as2)
KAVF18742-W An error occurred during system call or OS command execution.
(name=open, rc=-1, error=2)
```

```
KAVF18743-W An attempt to get the device information has failed. (device
name=/dev/rdisk/emcpower5as2)
```

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 of the array. You will then have a series of questions that will be asked interactively (Storage Model=2, Command Device= /dev/rdisk/..., Unassigned Open Volume Monitoring =N, Mainframe Volume Monitoring=N, Store Version =2.0).

```
[/opt/jp1pc/tools]
sangs9# jpcinssetup agtd -inst 99999
Storage Model (1:Thunder/AMS or 2:Lightning/USP/USP V) : 2
Command Device File Name [/dev/rdisk/c0t0d0s2] :
/dev/rdisk/c1t50060E800563AF99d1s2
Unassigned Open Volume Monitoring (Y/N) [N] :
Mainframe Volume Monitoring (Y/N) [N] :
Store Version [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 `jpcinlist` command. You must specify the agent type (RAID agent – `agtd`)

```
[/opt/jp1pc/tools]
sangs9# jpcinlist agtd
25519
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:
[/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]
sangs9# jpcinssetup agtd -inst 99999
Storage Model (1:Thunder/AMS or 2:Lightning/USP/USP V) : 2
Command Device File Name [/dev/rdisk/c0t0d0s2] :
/dev/rdisk/c1t50060E800563AF99d1s2
Unassigned Open Volume Monitoring (Y/N) [N] :
Mainframe Volume Monitoring (Y/N) [N] :
Store Version [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)
```

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.

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.
```

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 forcedly terminated. Do you want to
continue? (Y/N)Y
KAVE05090-I Service is stopped.
KAVE05087-I The connecting host has changed. (sangs9 -> ecctest01ms)
```

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”.

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)
```

30.2.3 Agent Processes


The following is a list of processes that are created by the agent and Tuning Manager server.

30.2.4 Verify Agents

To view the agents from Tuning Manager, 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 2 hours). The agents may show up under the related agents column. Refer to the Tuning Manager Installation Guide for additional troubleshooting.




Alternatively 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.



30.2.5 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 interval and at least 2 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 single agent. Select the agent 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.

30.2.6 Manual 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.



30.3 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.

30.3.1 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.

30.3.2 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.

30.3.3 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)
- Priority can be managed using Server Priority Manager (see [Appendix P: Server Priority Manager](#)).

Port Utilization and Throughput

Cache Utilization

Impact of Cache Size on Performance

How Application Utilize Cache

Cache Management

LDEV Response Time

Array Group Read/Write IOPS

30.3.4 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

30.3.5 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.

30.3.5.1 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

30.4 Citi Standard Reports

There are several standard reports that are part of the Citi HTnM Reporting package. The reports are organized by Subsystem Flow Analysis and Subsystem Trend Analysis. There is also a host level report, but that will not be frequently used as there are very few servers with HTnM host agents.


30.4.1 Setup of Citi Reports

Creating (or updating) the Citi HTnM reports in a HTnM database is a simple process. Once the package has been installed, the following steps should be followed:

1. Open a command prompt (Start → Run → cmd)
2. Set PR_HOME environment variable to the location of the Performance Reporter software.
Set PR_HOME="c:\program~1\HiCommand\TuningManager\PerformanceReporter"
3. Verify the PR_HOME environment variable by running a SET command to list all current variables. If the value is not correct then, repeat step 2.
4. Run the setup script to put the new reports into the system. This script will delete any previous versions of the reports.

30.4.2 Subsystem Flow Analysis Reports

To find these reports start Performance Reporter (HTnM → Go → Performance Reporter). Select Reports, User Defined Reports, and then Citi TagmaStore Reports.



These reports are named in sequence to correspond to the order in which they should be run when troubleshooting a performance problem.

30.4.2.1 Inclusive and Exclude Report Versions


On several reports, there are two versions of the same report. There is an inclusive version that shows all values and a second version that excludes all data below the criteria. For example, there is a report "02 Host Port CPU (MP of CHA) Util" and "02 Host Port CPU (MP of CHA) Util >50% Only". The second version of the report only shows data where the processor utilization is > 50%. If the CPU is running at 48%, then it will not show up in the report. Therefore both versions are included.

30.4.2.2 Response Time Reports

The response time reports (read and write) are broken up by the CU Address range of the LDEV. This is a reflection of the HTnM databases and cannot be combined into a single report.

30.4.3 Subsystem Trend Reports

To find these reports start Performance Reporter (HTnM → Go → Performance Reporter). Select Reports, User Defined Reports, and then Citi TagmaStore Reports.



These reports are intended to show the health of the array over time. These reports complement the analysis reports and cover much of the same data.

30.4.3.1 LDEV Reports

The LDEV reports (read response, write response, characteristics) are broken up by the CU Address range of the LDEV. This is a reflection of the HTnM databases and cannot be combined into a single report.

30.4.3.2 Thin Pool Reports

There is a Thin Pool Utilization Trend report and a Virtual Volume Configuration Report. The Utilization Trend report is very powerful in that it shows the utilization trend of the pool to help you predict when the pool will run out of storage. The Virtual Volume Configuration Report is a detail report showing all of the virtual volumes for a given pool. It can be run from the Thin Pool Utilization Trend report (by clicking on the Pool ID) or run directly.

30.4.4 Host (Server) Level Reports

This report directory contains only a single report that displays real-time statistics about the host. There should only be a single instance of the Host Agent (Host Agent for Windows) in

the environment. This host agent will be installed on the HTnM server to make sure that the server is not running out of CPU, Memory, or storage for the HDS software.

If more host agents are deployed at some time in the future, these reports may be expanded.

30.5 Citi Standard Alarms

There are several predefined alarms that correspond to the predefined reports discussed in the previous section. Alarms are used by HTnM to define criteria from which warning and abnormal conditions are met and then appropriate alerting is sent out.

30.5.1 Setup of Citi Alarms

Creating (or updating) the Citi HTnM alarms in a HTnM database is a simple process. Once the package has been installed, the following steps should be followed:

1. Open a command prompt (Start → Run → cmd)
2. Set PR_HOME environment variable to the location of the Performance Reporter software.
`Set PR_HOME=c:\program~1\HiCommand\TuningManager\PerformanceReporter`
3. Verify the PR_HOME environment variable by running a SET command to list all current variables. If the value is not correct then, repeat step 2.
4. Open the alarm script file with an ASCII editor like notepad. Search and replace "ecctest01ms" with the name of your server running HTnM and save the file.
5. Run the setup script to put the new alarms into the system. This script will delete any previous versions of the alarms.
6. 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 .
7. 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 a

30.5.2 Email Alerting Setup

30.5.2.1 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.

30.5.2.2 System Alerts Page

The System Alerts page should also be updated (Explorer → Administration → System Reports). Select the Edit System Alerts button.


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.

30.5.2.3 SNMP Alerting Setup


This functionality has not been tested with the ESS team. Implementing SNMP traps for alarms is not recommended at this time.

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.

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.



30.5.3 RAID Alarms

To find these alarms start Performance Reporter (HTnM → Go → Performance Reporter). Select Alarms, RAID, and then USP.




Note: The alarms listed above show that they have been activated as the icon has changed to .

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%.

Triggering Alarms

Once the criteria 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.

30.6 Processor to Port Mappings


In order to know what processor to report on for a given host port, the following chart lists the host ports and processor information. This pattern repeats and is easy to follow.

Port	Board Location	Processor (MP)
CL1-A	CHA-1EU	1EU_MP0
CL1-B	CHA-1EU	1EU_MP2
CL1-C	CHA-1EL	1EL_MP0
CL1-D	CHA-1EL	1EL_MP2
CL1-E	CHA-1FU	1FU_MP0
CL1-F	CHA-1FU	1FU_MP2
CL1-G	CHA-1FL	1FL_MP0
CL1-H	CHA-1FL	1FL_MP2
CL1-J	CHA-1GU	1GU_MP0
CL1-K	CHA-1GU	1GU_MP2
CL1-L	CHA-1GL	1GL_MP0
CL1-M	CHA-1GL	1GL_MP2
CL1-N	CHA-1HU	1HU_MP0
CL1-P	CHA-1HU	1HU_MP2
CL1-Q	CHA-1HL	1HL_MP0
CL1-R	CHA-1HL	1HL_MP2
CL2-A	CHA-2QU	2QU_MP0
CL2-B	CHA-2QU	2QU_MP2
CL2-C	CHA-2QL	2QL_MP0
CL2-D	CHA-2QL	2QL_MP2
CL2-E	CHA-2RU	2RU_MP0
CL2-F	CHA-2RU	2RU_MP2
CL2-G	CHA-2RL	2RL_MP0


CL2-H	CHA-2RL	2RL_MP2
CL2-J	CHA-2TU	2TU_MP0
CL2-K	CHA-2TU	2TU_MP2
CL2-L	CHA-2TL	2TL_MP0
CL2-M	CHA-2TL	2TL_MP2
CL2-N	CHA-2UU	2UU_MP0
CL2-P	CHA-2UU	2UU_MP2
CL2-Q	CHA-2UL	2UL_MP0
CL2-R	CHA-2UL	2UL_MP2
CL3-A	CHA-1EU	1EU_MP1
CL3-B	CHA-1EU	1EU_MP3
CL3-C	CHA-1EL	1EL_MP1
CL3-D	CHA-1EL	1EL_MP3
CL3-E	CHA-1FU	1FU_MP1
CL3-F	CHA-1FU	1FU_MP3
CL3-G	CHA-1FL	1FL_MP1
CL3-H	CHA-1FL	1FL_MP3
CL3-J	CHA-1GU	1GU_MP1
CL3-K	CHA-1GU	1GU_MP3
CL3-L	CHA-1GL	1GL_MP1
CL3-M	CHA-1GL	1GL_MP3
CL3-N	CHA-1HU	1HU_MP1
CL3-P	CHA-1HU	1HU_MP3
CL3-Q	CHA-1HL	1HL_MP1
CL3-R	CHA-1HL	1HL_MP3
CL4-A	CHA-2QU	2QU_MP1
CL4-B	CHA-2QU	2QU_MP3
CL4-C	CHA-2QL	2QL_MP1
CL4-D	CHA-2QL	2QL_MP3
CL4-E	CHA-2RU	2RU_MP1
CL4-F	CHA-2RU	2RU_MP3
CL4-G	CHA-2RL	2RL_MP1
CL4-H	CHA-2RL	2RL_MP3
CL4-J	CHA-2TU	2TU_MP1
CL4-K	CHA-2TU	2TU_MP3
CL4-L	CHA-2TL	2TL_MP1
CL4-M	CHA-2TL	2TL_MP3
CL4-N	CHA-2UU	2UU_MP1
CL4-P	CHA-2UU	2UU_MP3
CL4-Q	CHA-2UL	2UL_MP1
CL4-R	CHA-2UL	2UL_MP3

30.7 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:




30.8 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.






30.9 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):



31 Appendix R: Summary of Setup Files for HDvM (and Hitachi Storage Command Suite)

The following are needed changes to the ini files for the HDvM and related servers. They are provided in order to expedite implementation and setup of a new site.

31.1 Dispatcher (polling Period)

x:\Program Files\HiCommand\DeviceManager\HiCommandServer\config\dispatcher.properties
 server.dispatcher.daemon.pollingPeriod=60 (change from 5 to 60)

31.2 HTnM (email Alerts)

X:\Program Files\HiCommand\TuningManager\conf\user.properties
 #pr.incontextlaunch.mode=
 alert.email.smtpport=25
 alert.email.from=gds@citi.com NOTE: This needs to be a valid citi email address
 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

31.3 HTSM

Server Properties

Update server.properties in x:\Program Files\HiCommand\TieredStorageManager\conf

```
server.rmi.port=20352
server.rmi.secure=2
server.rmi.security.port=24500
server.base.initialsynchro=false
server.base.ipaddress=127.0.0.1
server.base.port=23015
server.mail.smtp.host=mailhub-vip.nj.ssmb.com
server.mail.from=HTSM_NAMMHDS01@citi.com
server.mail.errorsTo=kery.hines@citi.com
server.mail.smtp.port=25
server.mail.smtp.auth=false
server.eventNotification.mail.to=kery.hines@citi.com
server.eventMonitoringIntervalInMinute=180
server.migration.multiExecution=8
server.checkOutVolumeRange=true
server.repository.autoRefresh.pollingIntervalInMinute=360
server.repository.autoRefresh.serverStarted=false
server.repository.autoRefresh.recurringLocalTime=
server.migration.dataErase.defaultValue=true
server.migrationPlan.candidateVolumeCountLimit=true
server.repository.dvmModifyCheck.pollingIntervalInMinute=60
server.migrationPlan.candidateCapacityGroupDisplayMaxCount=4
```

31.4 SSL for HDvM

31.4.1 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 X500Name 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...

31.4.2 Generate private key

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

Program Files\HiCommand\Base\httpsd\sslc\bin

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)

31.4.3 Create a Certificate Signing Request

Program Files\HiCommand\Base\httpsd\ssl\bin

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

Using configuration from demoCA\ssl.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
```

31.4.4 Creating a Self-Signed Certificate

Program Files\HiCommand\Base\httpsd\ssl\bin

```
D:\Program Files\HiCommand\Base\httpsd\ssl\bin>ssl 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

31.4.5 Copy new key files to prod SSL directory

Copy **httpsdkey.pem** and **selfcert.pem** from
 x:\Program Files\HiCommand\Base\httpsd\ssl\bin\demoCA to
 x:\Program Files\HiCommand\Base\httpsd\conf\ssl\server

31.4.6 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 SSLCACertificateFile 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"
# SSLCACertificateFile "D:/Program
Files/HiCommand/Base/httpsd/conf/ssl/cacert/anycert.pem"
    SSLSessionCacheTimeout 3600
</VirtualHost>
```

31.4.7 Change SSL URL to HTTPS

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

Program Files\HiCommand\Base\bin

```
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 servers to make these changes effective.

31.4.8 Setup GUI/CLI on other Windows Hosts

1. Download the HiCommandCerts file from the Device Manager server from the following (http) URL:
<http://<devicemanagerserver>:2001/service/HiCommandCerts> and save in directory x:\Program Files\HiCommand\DeviceManager\HiCommandCLI

32 Appendix Z: Version Notes

1.0 – This is the initial release of the HDS USP TagmaStore Solutions Guide

1.1 – Minor edits to ICQ and RA references

1.2 – Updated noticed about RA. Engineering Management canceled the RA due to low residual risk of product. Restriction of required CAP for RA was removed. Restriction of required CAP for Standard Build still exists. Also updated the following sections:

- Appendix F: Firewall ports – Simplified port lists for Device Manager
- LUSE Volumes – Added a description for the numbering of LUSE volumes
- Security Practices – Added a section on Decommissioning of Storage
- Appendix B: Configuration of the Array... -- Added a section on Changing the HBAs on Host Groups (HBA Swap)
- Appendix B: Configuration of the Array – updated Host Group Mode list for VCS
- Best Practices – Added LUNs per port guidelines.
- Best Practices – Added Cache recommendations for Small, Medium, and Large configurations.
- Best Practices – Added timing estimates and restrictions for disk shredding.
- The following row entries were removed from the [Version Control Table](#):

Version	Date	Author	Reviewed	Comments
0.9.2	22-Nov-05	Gary Johnson	Global Storage Community	RFC
0.1	12-Jul-05	Gary Johnson	Parul Jain Kevin Hobaugh	Peer Review
0.2	14-Jul-05	Gary Johnson	Global Storage Community	RFC
0.3	22-Jul-05	Gary Johnson	Global Storage Community	
0.4	04-Aug-05	Gary Johnson	Global Storage Community	Updates to LUSE and std configs
0.5	19-Aug-05	Gary Johnson	Global Storage Community	Added small config and DM network info
0.6	30-Sep-05	Gary Johnson	Global Storage Community	Std Configs, Vol layouts, WWDS, eTrust, ICQ
0.7	14-Oct-05	Gary Johnson	Global Storage Community	ICQ Updates
0.8	04-Nov-05	Gary Johnson	Global Storage Community	SN & DM Security Controls PBRUN
0.9	08-Nov-05	Gary Johnson	Global Storage Community	CCI, PBRUN, CLI w/ https
0.9.1	21-Nov-05	Gary Johnson	Parul Jain Eliot Wilson Chris Clayton	Peer Review, Security, Reorg

1.3 – This is a minor update. An appendix was added to provide a best practice for the “HDS Migration Tools”.

1.4 – This is a minor update. Added a section to Best Practices regarding the LUN Expansion. Changed file properties of file to reflect new GDSE template standards. Updated the Best Practices for Host setups by referring to each Host OS support matrix.

1.5 – Added ECC User account to Special Accounts. Added heterogeneous raid support to RAID types section of Best Practices. Updated eTrust references to include syslog functionality and added syslog firewall port. Added syslog configuration information to Appendix B. Added variable passing for username and password on CLI scripting.

Updated the references table. The address spaces (CU) for R5 and R10 have been swapped per request because of the dominant use of RAID 5. References to clustering not supported were removed (HDS has qualified several Citi clusters.) Added description to Host Mode Option 2 – DBEdition and VCS.

- 1.5.1 -- Corrected Syslog examples
- 1.6 -- Section 11.6.9 Added connection error message when limiting IP addresses. Added more syslog samples.
- 1.6.1 -- Added new security features for Storage Navigator in section 13.2.2 Security in Storage Navigator. Updated default password recommendations in 11.2.9, 11.3, 13.2.1.1, and 13.2.2. Updated functional ID recommendations in section 10.14 Special User Accounts.
- 1.6.2 Updated Storage Navigator User Entitlements (13.2.4). Added Time Synchronization (SNTP). Updated Firewall list to include SNTP (13.6).
- 1.6.3 Added SVP Panel security and setup (13.2.5). Added recommended encryption be set to RSA 2048 (11.6.4). Changed Citigroup logo to Citi logo and replaced all Citigroup with Citi references (many sections). Updated references to reflect new GDSE website.
- 1.6.4 Added Appendix J (13.10) on McAfee Virus Scan and ePO agent installation instructions. Added ePO/Virus Scan rules to Appendix F--Firewall Rules (13.6).
- 1.6.5 Updated Device Manager Polling period to include filename of configuration file (10.23). Removed statement regarding "CLI does not work with CLI." (13.4.2). Added Appendix K on Setup of HiTrack Service for HTTPS. Updated sections to remove description of modem and replaced with HTTPS alerts (10.15, 9.1, 13.1)
- 1.6.6 Appendix L (14) added to provide Best Practice and other Solutions Guide material for Hitachi Tiered Storage Manager.
- 2.0 USPVP and TSM updates.
 - Added support for 300GB Drives
 - Updated entire document for coverage of the USPVP series array (8, 8.1, 8.2, 9.1, 9.2, 10, 10.1, 10.2, 10.3, 10.4, 10.5, 10.6, 10.7, 10.9, 10.10, 10.11, 10.12, 10.15, 10.16, 10.18, 10.21, 11.1, 11.3, 13.1, 14.9, 14.16, 14.24, 15.1, 15.5, 15.7, 17, 19, 21, 25).
 - Added Tiered Storage Manager references as new component (Appendix L: 24).
 - Added VPN to replace WWDS (10.15, 13.7.7).
 - Updated document links (10.1).
 - Restructured Appendix numbering such that each appendix has its own chapter number (13.1-13.11 → 13-25)
 - Added SSL for Storage Navigator
 - The following row entries were removed from the [Version Control Table](#):

1.0	21-Dec-05	Gary Johnson	Global Storage Community	Pkg Info, Sec Roles
1.1	23-Dec-05	Gary Johnson	Minor Release	ICQ / RA
1.2	29-Jun-06	Gary Johnson	Minor Release	See Release Notes Section
1.3	07-Jul-06	Kevin Hobaugh	Gary Johnson	Added HDS Migration Tool Information
1.4	17-Jul-06	Gary Johnson	Kevin Hobaugh	Added Host LUN Expansion

 - Added Security Policy (password, banner, failed login) for Device Manager
 - Added supported drive types (10.3)
 - Removed ICQ answer details in Security Best Practices due to redundancy. ICQ can be accessed by interested parties.
 - Added additional examples to CLI section (16)
- 2.1 Additional USPVP changes.
 - Updated LDEV destroy and create timing (10.11)

- Added Disk Rebuild Times (10.10)
- Updated GDSE website address from "<http://engineering.ti.citigroup.net/gdse/>" to "<http://engineering.citigroup.net/gdse/>"
- Removed references to RAID 6 (14+2) – does not exist on USPV or USP
- Updated cache charts based on HDS recommendations (10.4)
- Updated BED Pairs recommendations for Small USPV (10.4,10.5,10.22)
- Updated 2Gb Drive restrictions on USPV (10.3, 10.5)
- Added DM version 5.7 requirement for USPV (10.14, 15.1)
- Added confirmation of internet access step to HTTPS dial-home (23.1)
- Added 300GB R10 4+4 (10.2, 10.6)

2.2 Additional USPV updates

- Changed BED numbers for Large USPV (10.3, 10.4, 10.5, 10.22)
- Updated maintaining multiple user and multiple SVPs lists (10.24, 14.7)
- Corrected cache values in ordering section (10.22)
- Updated Unix packages for CLI & CCI (9.2)

2.3 Additional Minor updates

- Updated host port configuration on large model [10.22.1]
- Re-added (removed in version 2.0) ICQ detail information that corresponds to updated ICQ. Updated ICQ descriptions [11.10]
- Added restriction for Functional ID use [10.16]
- Updated eTrust Solutions Guide document link [11.4, 12]

2.4 SNMP & Additional Minor updates

- Added SNMP Alerting Sections [10.x, Appendix M]
- Updated access control list for HDvM [11.5.9]
- Added 2001 Firewall port for Physical View (running Storage Navigator) from HDvM [18,15.5.1]
- Updated firewall ports for HTTPS Dial-home [23, 18, 10.17]
- Updated McAfee and ePO installation notes to reflect new versions of software [22]
- The following 2006 row entries were removed from the [Version Control Table](#):

1.5	07-Aug-06	Gary Johnson	Kevin Hobaugh	eTrust Syslog
1.5.1	09-Aug-06	Gary Johnson	Minor Release	Corrected Syslog samples
1.6	27-Oct-06	Gary Johnson	Kevin Hobaugh	IP restriction, Syslog server setup
1.6.1	03-Nov-06	Gary Johnson	Kevin Hobaugh	Updated SN security
1.6.2	07-Nov-06	Gary Johnson	Global Storage Community	Added SNTP Updated FW

- Removed old WWDS reference [12]

2.5 Hitachi Dynamic / Thin Provisioning (HDP) & Tuning Manager updates

- Added HDP Best Practices [Appendix N]
- Added HDP Configuration [Appendix N]
- Added HDP Reporting via Device Manager [Appendix N]
- Added HDP Reporting via Tuning Manager (with Tuning Manager Setup) [Appendix N]
- Added Performance Data Gathering [14.28]
- Added HDP Application Considerations [25.5.9]
- Updated Syslog info to include Solaris 10 [13.10.1]
- Restated supported RAID configurations [9.4]

- Corrected password listing [25.4.2]
- Updated VSE and ePO instructions for new website and newer packages [21]
- Updated SNMP list descriptions [9.28]
- Added extra-large array configuration [9]
- Added CU:LDEV ranges for Pool-Vols and DP-Vols [9.10]
- Updated Pool ID Numbering [25.5.3]
- Added email content modification on HDP email alerts [25.3.9]
- Updated Host Mode table for all hosts [9.3,13.20,15.4]

2.6 HDP, Device Manager, Tuning Manager updates

- Added Regional Reviewers Chart renumbering all other sections [2]
- Changed HDP Thin Volume methodology to use 1 DP-Vol / Array Group [27]
- Updated pool growth guidelines [27]
- Updated migrations wording [27]
- Updated Reporting for HDP [27]
- Added 1000GB SATA drive to USPV [11.5, 11.6, 11]
- Updated CU:LDEV address space [11.8]
- Updated spare drive recommendations [11.6]
- Updated HDP best practices [27]
- The following 2007 row entries were removed from the [Version Control Table](#):

1.6.3	22-Mar-07	Gary Johnson	Kevin Hobaugh	Added Panel Security
1.6.4	04-Apr-07	Gary Johnson	Richard Sopp	Added McAfee Virus Scan
1.6.5	03-Aug-07	Gary Johnson	Chris Clayton	HTTPS DialHome & Misc Edits
1.6.6	07-Aug-07	Chris Clayton	Gary Johnson	Appendix L (HTSM)

- Renamed Release Notes to Version Notes [28]
- Updated packages chart[10.2] and added package notes [10.3]
- Moved Special User Accounts to Security Practices [12.3]
- Added USPV XL Sample Configuration [11.22.1]
- Added sample general reports for Tuning Manager [27.5.2.8]
- Added statement regarding shared host ports for all OS types [11.13]
- Added Security Practices for HDvM CLI [12.11]
- Updated HDvM CLI procedures and security [17]
- Added VMWare Server Info [11.27]
- Added Server Sizing Info [11.28]
- Added Running Multiple Versions Simultaneously [11.29]
- Added HDvM Array Name Recommendation [12.12]
- Added 1000GB Drives [11.2, 11.3]
- Added updates for HTSM 6.0 [10.2, 25]
- Updated Special Accounts (Functional IDs) [12.3]
- Reorganized HDP SN Setup to Match HDP HDvM Setup [27]

2.7 USPVM updates

- Updated standard build statement [8]
- Added USPVM configurations [11.2,11.4,11.5,11.6,11.7,11.13.2,11.13.3,11.22.1]
- Added USPVM Power, Sizing, and HVAC [20.3,20.5,20.6,20.7]
- Updated Management Station Descriptions [20.10, 20.11, 20.12]
- Updated firewall ports for new HDvM ports [19]
- Added Performance Monitor Extract Tool [Appendix O: 28, 10.2]

2.8 Zero Data Removal Tool

- Added Zero Data Removal Tool Section [27.8]
- 2.9 Misc Updates and New Lower-Cost Tiering for Replication
- Updated HDP SIM Code Chart [27.7.9.3]
 - Added SNMP trap info (from Tivoli alerts) [27.7.9.1]
 - Added lower-cost replication tiering statement [11.31]
 - Added Host Modes for Windows Cluster with VxVm, AIX w/HACMP [11.1]
 - Updated AV section to reflect Windows XP and Vista SVP installations [23.1, 23.1.2]
 - Added troubleshooting with Resource Allocation and Administration screens [16.3.3.3]
 - New Quorum size of 1000MB [11.7]
 - Updated max cache size for USPVM to 16GB [11.4]
 - Added database backup/restore/migration for HDvM [11.32]
 - Updated package list [10.2]
 - Updated package directory installation folder [28.2.4]
 - Removed duplicate description of HDvM server sizing [11.28]
 - Updated package name for HTnM AFR for Solaris [10.2]
 - Added HTnM Solaris AFR setup instructions [27.6.2.3]
 - Updated HTnM Verify Agents [27.6.2.4]
 - Updated HDP Pairty Group Sourcing from 4x to allow for 1x [27.3, 27.7.2]
- 2.9.1 Updates for HTnM AFR
- Added Solaris agent package instructions and examples
- 2.9.2 Updates for HTnM and HTSM 6.2
- Updated Thin Provisioning directive (default storage configuration) [27.1]
 - Updated AFR for Windows Setup [27.6.2.2]
 - Added package notes section for package instructions/notes [10.3]
 - Updated JRE Instructions for Solaris "jre16016" [10.2, 10.3]
 - Added Server Priority Manager (Appendix P) [29] – renumbered 29+ sections
 - Added new SIM alerts for Pools [27.7.9.3]
 - Updated Host Options Table for RHCS and Oracle AMS RAC on AIX [11.1]
 - Updated 450GB drive type and R6(6+2) recommendations [11.8]
 - Updated LUSE as not recommended [11.10]
 - Added USPVM HDP requirement for new deployments [11.32]
 - Updated Appendix L: Tiered Storage Manager for use with thin pools and updated HTSM 6.2 software [25]
 - Added System Options Table and System Disk instructions [27.9, 11.36]
 - Removed Array Settings – replaced with Pool Full behavior description [27.7.7]
 - Added system disk instructions [11.36]
 - Updated Tuning Manager Appendix to include new methodology [30]
 - Add port / port CPU mapping [30.6]
 - Add in new java heap info for expanded HDvM memory settings [11.35]
 - Changed max LUN size for thin volumes to be 2TB [11.34, 27.7.2]
 - Created Appendix Q for Tuning Manager (moved from appendix N) [27,30]
 - Added Active Directory Instructions for HDvM Authentication [12.13]
 - Added Appendix R for HDvM Summary of Changes [31]
- 2.9.3 Additional Updates for HTnM and HTSM 6.2
- Updated application migration recommendations for HDP [27.7.1]
 - Updated Host Mode chart, added AIX entries [11.1]
 - Updated statement about single server [30.2.1]
 - Removed reference to host mode 19 and old microcode minimum [11.1]

- Removed 5.x packages from table, added 6.2 packages to table [10.2]
- Minor Update on AFR on Solaris [10.3.3.3]
- Added warnings for not using a command device as a standard storage device [15.23, 18]
- Updated SNTP to add description of CE screen location and to refer to NTP [11.23]
- Updates SYSLOG server setup [15.10.3]

2.9.4 Minor Update

- Added SCSI Reservations on USPV [11.37]

2.9.5 SOM 803 and DRU Procedure

- Added SOM 803 to table and updated Notes section [11.36]
- Added DRU procedure [27.7.10]