Veritas InfoScale™ Operations Manager 7.2 User Guide

November 2016

VERITAS™

Veritas InfoScale™ Operations Manager 7.2 User Guide

Last updated: 2016-11-02

Document version: 7.2 Rev 0

Legal Notice

Copyright © 2016 Veritas Technologies LLC. All rights reserved.

Veritas, the Veritas Logo, Veritas InfoScale, and NetBackup are trademarks or registered trademarks of Veritas Technologies LLC or its affiliates in the U.S. and other countries. Other names may be trademarks of their respective owners.

This product may contain third party software for which Veritas is required to provide attribution to the third party (“Third Party Programs”). Some of the Third Party Programs are available under open source or free software licenses. The License Agreement accompanying the Software does not alter any rights or obligations you may have under those open source or free software licenses. Refer to the third party legal notices document accompanying this Veritas product or available at:

https://www.veritas.com/about/legal/license-agreements

The product described in this document is distributed under licenses restricting its use, copying, distribution, and decompilation/reverse engineering. No part of this document may be reproduced in any form by any means without prior written authorization of Veritas Technologies LLC and its licensors, if any.

THE DOCUMENTATION IS PROVIDED "AS IS" AND ALL EXPRESS OR IMPLIED CONDITIONS, REPRESENTATIONS AND WARRANTIES, INCLUDING ANY IMPLIED WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE OR NON-INFRINGEMENT, ARE DISCLAIMED, EXCEPT TO THE EXTENT THAT SUCH DISCLAIMERS ARE HELD TO BE LEGALLY INVALID. VERITAS TECHNOLOGIES LLC SHALL NOT BE LIABLE FOR INCIDENTAL OR CONSEQUENTIAL DAMAGES IN CONNECTION WITH THE FURNISHING, PERFORMANCE, OR USE OF THIS DOCUMENTATION. THE INFORMATION CONTAINED IN THIS DOCUMENTATION IS SUBJECT TO CHANGE WITHOUT NOTICE.

The Licensed Software and Documentation are deemed to be commercial computer software as defined in FAR 12.212 and subject to restricted rights as defined in FAR Section 52.227-19 "Commercial Computer Software - Restricted Rights" and DFARS 227.7202, et seq. "Commercial Computer Software and Commercial Computer Software Documentation," as applicable, and any successor regulations, whether delivered by Veritas as on premises or hosted services. Any use, modification, reproduction release, performance, display or disclosure of the Licensed Software and Documentation by the U.S. Government shall be solely in accordance with the terms of this Agreement.

Veritas Technologies LLC

500 E Middlefield Road
Mountain View, CA 94043

http://www.veritas.com

Technical Support

Technical Support maintains support centers globally. All support services will be delivered in accordance with your support agreement and the then-current enterprise technical support policies. For information about our support offerings and how to contact Technical Support, visit our website:

https://www.veritas.com/support

You can manage your Veritas account information at the following URL:

https://my.veritas.com

If you have questions regarding an existing support agreement, please email the support agreement administration team for your region as follows:

Worldwide (except Japan)	CustomerCare@veritas.com
Japan	CustomerCare_Japan@veritas.com

Documentation

Make sure that you have the current version of the documentation. Each document displays the date of the last update on page 2. The document version appears on page 2 of each guide. The latest documentation is available on the Veritas website:

https://sort.veritas.com/documents

Documentation feedback

Your feedback is important to us. Suggest improvements or report errors or omissions to the documentation. Include the document title, document version, chapter title, and section title of the text on which you are reporting. Send feedback to:

doc.feedback@veritas.com

You can also see documentation information or ask a question on the Veritas community site:

http://www.veritas.com/community/

Veritas Services and Operations Readiness Tools (SORT)

Veritas Services and Operations Readiness Tools (SORT) is a website that provides information and tools to automate and simplify certain time-consuming administrative tasks. Depending on the product, SORT helps you prepare for installations and upgrades, identify risks in your datacenters, and improve operational efficiency. To see what services and tools SORT provides for your product, see the data sheet:

https://sort.veritas.com/data/support/SORT_Data_Sheet.pdf

Contents

Section 1 Getting started 25

Chapter 1 Introducing Veritas InfoScale Operations Manager 26

About Veritas InfoScale Operations Manager 26
About increasing efficiency and productivity through storage virtualization 27
Storage Foundation: a complete solution for online storage management 28
Veritas InfoScale Operations Manager: tying it all together 28
About centralized management and monitoring 29
About discovery, monitoring, and management 29

Chapter 2 Using the Management Server console 31

About the Management Server console 31
About the Management Server console Home page 32
About the Global Dashboard 34
About the elements in a perspective view 35
About viewing the summarized information 37
About selecting the objects 38
Moving fields to the properties pane of the Management Server console 38
About drilling down to more information on the selected object 39
About searching for objects 40
Creating new search queries 41
Managing saved search queries 43
Creating and saving searches as Smart Folders 43
Managing Smart Folders 44
About viewing the solutions 46
About viewing the reports 46
About viewing tasks 47
About viewing connectivity graphs 48
About accessing the Veritas InfoScale Operations Manager Help 49

Chapter 3 Examples for using Veritas InfoScale Operations Manager 50

Example: Creating a volume using Veritas InfoScale Operations Manager 51
Example: Creating a service group, adding it to a cluster, and bringing it online using Veritas InfoScale Operations Manager 52
Example: Cluster Server troubleshooting using Veritas InfoScale Operations Manager 55
Service group dependencies – on resources and other service groups 55
Service groups faults 56
Clearing faults for the service group's resources 56
Veritas InfoScale Operations Manager reports for Cluster Server troubleshooting 57
Example: Ensuring the correct level of protection for volumes controlled by Storage Foundation 60
Mirroring: Protection from hardware failures 61
Snapshots: Backing up to prevent the loss of data 62
Example: Improving the availability and the disaster recovery readiness of a service group through fire drills 65
High availability fire drill - Ensuring the high availability capabilities of a service group 66
Disaster recovery fire drill - Ensuring the disaster recovery readiness of a service group 68
Example: Identifying the performance issues of an application using Veritas InfoScale Operations Manager 69
Example: Volume migration using Veritas InfoScale Operations Manager 73
Examples: Identifying and reducing storage waste using Veritas InfoScale Operations Manager 75
Reclaiming thin storage - example 75
Compressing files - example 79
Deduplicating file systems - example 83

Section 2 Managing Veritas InfoScale Operations Manager 88

Chapter 4 Managing user access 89

Creating an Organization 89
Create Organization panel options 91

Create Organization - Select an object panel options 91
Create Organization - Based on a rule panel options 92
Modifying the name of an Organization 94
Modify Organization panel options 94
Deleting an Organization 95
Moving an object to an Organization in a perspective 95
Assigning permissions to user groups on an Organization within a perspective 96
Modifying permissions assigned to user groups on an Organization within a perspective 97
Deleting permissions assigned to user groups on an Organization within a perspective 97
Modifying permissions assigned to user groups on an object within a perspective 98
Verifying a user group in the domain 99
Viewing permissions information 99
Viewing the permissions assigned on a perspective, an Organization, or on an object 100

Chapter 5 Setting up fault monitoring 102

About alerts and rules 102
Creating rules in a perspective 103
Create Rule - Select the type of fault conditions to trigger this rule panel options 105
Create Rule - Select one or more fault topics which will trigger this rule panel options 106
Create Rule - Select organizations panel options 106
Create Rule - Setup notifications panel options 107
Create Rule - Enter name and description panel options 108
Editing rules in a perspective 109
Edit Rule - Select the type of fault condition to trigger this rule panel options 110
Edit Rule - Select one or more fault topics which will trigger this rule panel options 110
Edit Rule - Select organization panel options 111
Edit Rule - Setup notifications panel options 111
Edit Rule - Enter name and description panel options 112
Deleting rules in a perspective 113
Delete Rule panel options 113
Enabling rules in a perspective 114
Enable Rule panel options 114
Disabling rules in a perspective 114

Disable Rule panel options 115
About faults and risks 115
Suppressing faults in a perspective 115
Suppress Faults panel options 116
Restoring a suppressed fault in a perspective 117

Chapter 6 Using reports 118

About reports 118
About using reports 120
Running a report 121
Select scope to run report panel options 121
Saving a report 122
Subscribing for a report 123
Subscribe panel options 123
Editing a report subscription 124
Deleting a report subscription 125
Sending a report through email 125
Email report panel options 126
Viewing my report subscriptions in a perspective 126
Viewing all the report subscriptions in a perspective 127
About the reports available in Veritas InfoScale Operations Manager 127

Section 3 Managing hosts 132

Chapter 7 Overview 133

About performing Storage Foundation and replicator operations 133
About Storage Foundation operations not supported on Windows host 134
Viewing storage summary at the cluster level 135
Viewing faults and risks at the cluster level 136

Chapter 8 Managing File Replicator (VFR) operations 137

About performing File Replicator operations 138
Viewing the VFR option of a host 138
Creating a consistency group 139
Viewing consistency groups 139
Deleting a consistency group 140
Associating a consistency group to a replication job 140
Disassociating a consistency group from a replication job 141
Viewing consistency group properties 141

Creating a replication job 142
Viewing File Replication Jobs 143
Starting a replication job 143
Pausing a replication job 143
Resuming a replication job 144
Stopping a replication job 144
Syncing a replication job 144
Modifying a replication job 145
Deleting a replication job 145
Viewing properties of File Replication Jobs 145

Chapter 9 Managing disk groups and disks 146

About managing disk groups 147
Creating disk groups 148
Create Disk Group - Disk Group Specifications 149
Select Disks panel options 153
Filter Criteria panel options 155
Create Disk Group - Rename disks panel options 155
Recovering disk groups 156
Deporting disk groups 157
Destroying disk groups 157
Importing disk groups 158
Import disk group panel options 159
Adding disks to disk groups 159
Select Disk Group panel options 161
Resizing disks in disk groups 161
Resize Disk panel options 162
Renaming disks in disk groups 163
Rename Disk panel options 165
Upgrading disk groups 166
Splitting disk groups 166
Split Disk Group panel options 167
Moving disk groups 168
Move Disk Group panel options 169
Joining disk groups 170
Join Disk Group panel options 171
About managing disks 172
Initializing disks 172
Initialize Disk panel options 173
Replacing disks 174
Replace Disk panel options 175
Recovering disks 175

Recover Disk panel options 176
Mapping disks 176
Unmapping disks 177
Disconnecting disks 178
Removing disks from disk groups 178
Setting host prefix for disks 179
Bringing disks online 180
Online Disk panel option for making the disks online 181
Taking disks offline 181
Setting disk usage 181
Set Disk Usage panel options 182
Set Disk Usage - Windows host panel options 182
Evacuating disks 184
Evacuate Disk panel options 185
Running or scheduling Trim 186
Trim - Run or Schedule panel options 187
Rescanning disks 188

Chapter 10 Managing volumes 190

About managing Storage Foundation volumes 191
Creating Storage Foundation volumes 192
Create Volume – Select Disk Group and Disk Selection method panel options 194
Volume attributes panel options for creating volumes on UNIX or Linux hosts for specifying values 195
Volume attributes panel options for creating volumes on Windows hosts 198
Add Drive Letter, Path and Create File System details panel options 199
Stopping volumes 200
Recovering volumes 200
Reactivating volumes 201
Deleting volumes 202
Delete Volume panel options 203
Moving volumes 203
Move Volume panel options 203
Renaming volumes 204
Rename Volume panel options 204
Adding mirrors to volumes 205
Add mirror - Options 207
Enable FastResync option panel options 208
Add mirror - Advanced options 209

Removing the mirrors of volumes 212
Remove mirror panel options 212
Creating instant volume snapshots 214
Create Volume Snapshot - disk selection page panel options 216
Create Volume Snapshot - Advanced Options 216
Create Volume Snapshot - Instant Snapshot panel options 217
Creating space optimized snapshots for volumes 218
Create Volume Snapshot - Space Optimized Snapshot panel options 219
Creating mirror break-off snapshots for volumes 220
Create Volume Snapshot - Mirror Break-off Snapshot panel options 221
Dissociating snapshots 222
Dissociate Snapshot panel options 223
Reattaching snapshots 224
Reattach Snapshot panel options 225
Resizing volumes 225
Resize volume panel options 226
Restoring data from the snapshots of volumes 230
Restore Data From Snapshot panel options 231
Refreshing the snapshot of volumes 232
Refresh snapshot panel options 233
Configuring a schedule for volume snapshot refresh 235
Schedule operation panel options 236
Adding snapshot volumes to a refresh schedule 237
Add to existing schedule panel options 237
Removing the schedule for volume snapshot refresh 238
Remove volumes from refresh schedule panel options 239
Setting volume usage 239
Set Volume Usage panel options 239
Splitting snapshots 240
Starting synchronization of snapshots 241
Enabling FastResync on volumes 242
Enable FastResync panel options 242
Disabling FastResync on volumes 243

Chapter 11 Managing file systems 245

About managing file systems 246
Creating file systems 246
Create File System - File System Options 247
Select Volume panel options 250

Advanced Options panel 251
Enabling change logs 254
Disabling change logs 255
Synchronizing change logs 255
Removing change logs 256
Defragmenting file systems 257
Defrag file system panel options 257
Unmounting non clustered file systems from hosts 257
Unmount Confirmation panel 259
Mounting non clustered file systems on hosts 259
Mount Options panel 260
Advanced Mount Options panel options 261
Unmounting clustered file systems 262
Unmount Clustered File System panel options 263
Mounting clustered file systems on hosts 264
Mount File System - Clustered Mount Options 264
Remounting file systems 265
Remount Options panel 266
Checking file systems 268
Check File System panel options 268
Creating file system snapshots 269
Create Snapshot - Create File System Snapshot panel options 269
Snapshot level selection panel options 270
Create File System snapshot - Configure Options 270
Remounting file system snapshot 271
Remount File System Snapshot panel options 272
Mounting file system snapshot 273
Mount File System Snapshot panel options 274
Unmounting file system snapshot 275
Unmount File System Snapshot panel options 275
Removing file system snapshot 276
Remove File System Snapshot panel options 277
Monitoring capacity of file systems 278
Monitor Capacity panel options 278

Chapter 12 Managing SmartIO 280

About managing SmartIO 281
About write-back caching in SmartIO 282
Enabling or disabling SmartIO caching 282
Creating a cache 283
Create Cache panel options 284

Viewing the cache details 285
Viewing the SmartIO Impact analysis chart 286
Changing SmartIO mode 287
Modifying a cache 288
Modify Cache panel options 288
Deleting a cache 291
Loading files to a cache 291
Pinning tablespaces or files to a cache 292
Unpinning tablespaces or files from a cache 292
About using SmartAssist 293
Creating an I/O trace log 293
New I/O Trace Log panel options 294
Viewing I/O trace logs 295
Analyzing an I/O trace log 295
SmartAssist Analysis Options pane 296
Removing an I/O trace log 296

Chapter 13 Managing application IO thresholds 298

About managing application I/O workloads with IOPS settings 298
Managing IO Thresholds 299
Setting IO Thresholds 300
Modifying an App VG 301
Viewing live IOPS charts 302

Chapter 14 Managing replications 303

About managing replications 303
Configuring Storage Foundation replications 304
Configure replication - Create primary panel options 305
Configure replication - Create resources 307
Configure replication - Create primary result panel options 308
Configure replication - Consistency check on secondary host panel options 308
Configure replication - Replication settings for secondary host panel options 310
Configure replication - Add secondary result panel options 313
Adding a Secondary 313
Pausing the replication to a Secondary 314
Pause replication panel options 315
Resuming the replication of a Secondary 315
Resume replication panel options 316
Starting replication to a Secondary 316
Start replication panel options 317

Stopping the replication to a Secondary 318
Stop replication panel options 318
Switching a Primary 319
Switch replication panel options 319
Taking over from an original Primary 320
Takeover Primary panel options 321
Associating a volume 322
Associate Volume - Volume Selection panel options 323
Associate Volume - Advanced Options panel options 323
Resynchronizing a Secondary 324
Removing a Secondary 324
Remove secondary panel options 325
Unconfiguring replication 325
About setting alerts for replication 326
Monitoring replications 326
Monitor replication panel options 327

Chapter 15 Optimizing storage utilization 329

About reclaiming thin storage 329
Performing thin reclamation on file systems or disks 330
Performing thin reclamation on thin pools in enclosures 331
Compressing files 332
About file compression in Veritas InfoScale Operations Manager 333
Selecting directories for compression 334
Setting up compression schedules 335
Starting compression on demand 337
Deduplicating file systems 338
About file system deduplication 338
About deduplication chunk size 340
Implementing deduplication for a file system 341
Starting deduplication for a file system 343
Disabling or removing deduplication for a file system 344

Section 4 Managing high availability and disaster recovery configurations 345

Chapter 16 Overview 347

About high availability and disaster recovery operations 347
Pre-requisites for performing high availability and disaster recovery operations 348

About attributes of Availability objects 348
About Virtual Business Services 349

Chapter 17 Managing clusters 350

About managing clusters 350
Opening a cluster configuration 351
Saving a cluster configuration 351
Closing a cluster configuration 352
Editing attributes of a cluster 352
Edit attribute options 353
Importing a type definition 354

Chapter 18 Managing service groups 356

About managing service groups 357
Creating service groups 358
Create Service Group options 359
Configure System List options 360
Configure Resources options 361
Resource Dependencies options 362
Enabling service groups 363
Disabling service groups 364
Creating Atleast Count dependencies for a resource in a service group 365
About Atleast Count dependency 366
Autoenabling service groups 367
Freezing service groups 367
Unfreezing service groups 368
Flushing service groups 369
Enabling all resources of service groups 370
Disabling all resources of service groups 371
Deleting service groups 371
About linking service groups in a cluster 372
Linking service groups in a cluster 372
Link Service Group options 373
Unlinking service groups 375
About site aware service group operations 376
Prerequisites for using site-related service group operations 377
Limitations of site-related service group operations 377
Bringing service groups online 377
Online service groups panel options 378
Taking service groups offline 379
Offline service groups panel options 380

Switching service groups 381
Switch service groups panel options 382
Clearing faults on service group 383
Clearing the resources in a service group from the Admin Wait state 384
Editing attributes of service groups 385
Modifying the system list for a service group 386
About dependency views 386
Viewing the service group dependency view 387
About modifying a service group 389
About Cluster Server service group alerting and failover reporting 389
Viewing VCS Failover Duration report 390

Chapter 19 Managing systems 391

About managing systems 391
Freezing a system 391
Unfreezing a system 392
Editing attributes of a system 393
Starting the Cluster Server high availability daemon on the hosts in a cluster 393
Stopping the Cluster Server high availability daemon on the systems in a cluster 394

Chapter 20 Managing VSystems 396

About VSystems 396
Starting a virtual machine 396
Stopping a virtual machine 397
Migrating a virtual machine 398

Chapter 21 Managing resources 399

About managing resources 400
Enabling resources 400
Disabling resources 401
Deleting resources 401
Clearing faults on resources 402
Probing resources 403
Taking a resource offline and propagating the state 404
Bringing resources online 405
Taking resources offline 405
Invoking a resource action 406
Invoke Action options 407

Editing attributes of a resource 408
Editing attributes of a resource type 409
Linking resources in a service group 410
Unlinking resources in a service group 410
Adding or modifying resources 411
Marking a resource as critical 412
Marking a resource as non critical 413
Viewing the resource dependency view 413

Chapter 22 Managing global cluster configurations 415

About global clusters 415
About global clusters terminology 416
About creating global clusters 417
Prerequisites for creating global clusters 417
Adding a remote cluster to a local cluster 418
Converting local service groups to global service groups 418
Make Global/Local options 419
Converting global service groups to local service groups 420
About removing a remote cluster from a global cluster setup 421
Taking the wac resource offline 422
Removing the remote cluster from the global cluster setup 422

Chapter 23 Running fire drills 424

About high availability and disaster recovery readiness 424
About high availability fire drills 425
Running the high availability fire drill 426
About disaster recovery fire drills 427
About configuring a fire drill service group 427
Running the disaster recovery fire drill 428
Select remote clusters and global service groups panel options 429
Schedule panel options 430
Summary panel options 431
Editing a fire drill schedule 432
Edit Schedule panel options 432
Deleting fire drill schedules 434
Enabling fire drill schedules 434
Disabling fire drill schedules 435
Viewing fire drill schedules 436

Chapter 24 Using recovery plans 437

About recovery plans 437
Creating recovery plans 438
Create recovery plan panel options 439
Editing recovery plans 442
Edit recovery plan panel options 442
Running recovery plans 444
Run recovery plan panel options 445
Deleting a recovery plan 446
Viewing historical runs of recovery plans 446
Historical Runs panel options 447
Viewing properties of recovery plans 447
Recovery Plan Properties panel options 448
About recovery plan log files 449

Chapter 25 Managing ApplicationHA 450

About ApplicationHA Management 450
Prerequisites for ApplicationHA Management 451
About the ApplicationHA operations 451
Launching ApplicationHA operations from Veritas InfoScale Operations Manager 452
About the ApplicationHA infrastructure 452
Enabling the ApplicationHA infrastructure for a managed host 453
Disabling the ApplicationHA infrastructure for a managed host 453

Chapter 26 Managing application configuration 455

About Application Configuration 455
Prerequisites for application configuration 456
Launching the configure application wizard from Veritas InfoScale Operations Manager 456

Chapter 27 Multi Site Management 458

About Multi-Site Management 458
Features of Multi-Site Management 459
Prerequisites of Multi-Site Management 460
Limitations of Multi-Site Management 460
Setting up a campus cluster 461
Setting up a replicated data cluster 462
Configuring stretch sites 463
Enclosure site assignment panel options 464
Select cluster type panel options 464

Assign site to the systems panel options 465
Site fencing preference panel options 466

Appendix A List of high availability operations 467

Cluster operations 467
System operations 468
Service group operations 469
Resource operations 472

Section 5 Monitoring Storage Foundation HA licenses in the data center 474

Chapter 28 Managing licenses 475

About licenses 475
About Veritas licensing and pricing 476
About the Symantec Performance Value Unit 479
About the Symantec Performance Value Unit for VMware virtual machines 480
About the Symantec Performance Value Unit for Solaris LDOM virtualization server 482
About the Symantec Performance Value Unit for kernel-based virtual machines 483
About the Symantec Performance Value Unit for IBM LPAR 484
About the per-core licensing 484
About assigning price tiers to hosts 485
About license deployment policies 485
Assigning a price tier to a host automatically 486
Assigning a price tier to a host manually 486
Select tier values panel options 487
Select hosts to apply same tier values panel options 488
Creating a license deployment policy 488
Create policy - Details panel options 488
Modifying a license deployment policy 490
Edit thresholds - Details panel options 491
Deleting a license deployment policy 491

Chapter 29 Viewing deployment information 493

Viewing the overview of SFHA licenses in the data center 493
About chargeable deployed licenses 495
Viewing the deployment details 496

Viewing the deployment policy details in the data center 497
Viewing the VOM Deployment Report 498

Chapter 30 Monitoring performance 500

About performance metering statistics 501
About metered resources 501
About space estimation for data logs 505
Enable performance metering for a host 508
Disable performance metering for a host 508
Enable performance metering for a virtualization server 509
Disable performance metering for a virtualization server 509
About Veritas InfoScale Operations Manager performance graphs 510
Pre-requisite commands to view performance graphs for a resource 512
Viewing the performance graphs for a host 513
Viewing the performance graphs for a disk 514
Viewing the performance graphs for volume and file system 515
Viewing the performance graphs for a path 517
Viewing the performance graphs for an initiator 518
Viewing the performance graphs for virtualization server and virtual machines 519
Viewing the performance graphs for a path of a virtualization server 521
Viewing the performance graphs for an enclosure 522
About threshold settings 523
Adding threshold settings for an object 525
Deleting the threshold settings for an object 528
Enabling the threshold settings for an object 531
Disabling the threshold settings for an object 534

Chapter 31 Managing Business Applications 537

About Business Applications in Veritas InfoScale Operations Manager 538
Creating or modifying a Business Application 538
Renaming a Business Application 539
Deleting a Business Application 540
Viewing Business Applications in the data center 540
Viewing the overview of a Business Application 541
Viewing service availability for a Business Application 542
Viewing data availability for a Business Application 543
Viewing SAN connectivity for a Business Application 544

About the makeBE script 545
Objects used in the makeBE script 546
makeBE script CSV file details 546
makeBE script parameters details 547
Limitations of makeBE script 548
makeBE script log files 549
Creating Business Application using the makeBE script 549
Importing Business Application using the makeBE script 550
Exporting Business Application using the makeBE script 551
Updating Business Application using the makeBE script 552
Deleting Business Application using the makeBE script 552

Chapter 32 Managing extended attributes 554

About using extended attributes 554
Setting values to the extended attributes on an object 555
Set Extended Attributes panel options 556
Searching objects to set extended attribute values 556
Set Extended Attributes panel options 557
Modifying the extended attributes value on an object 558

Chapter 33 Managing policy checks 559

About policy checks 559
How signature registration settings work 560
Registering policy signatures 561
Unregistering a signature 562
Setting signature tunables 563
Running a manual policy scan 564
Enabling or disabling policy signatures 564
Viewing policy violation details 565
Viewing or exporting a list of available policy signatures 567
About using custom signatures for policy checks 567
Creating a custom signature script 568
Sample custom signature script 569
Installing a custom signature script 571
Copying the custom signature script 571
Removing a custom signature 572
About using the Distribution Manager Add-on to bundle custom signature scripts 572
Sample setup.pl script for the custom signature 572
Sample unsetup.pl script for a custom signature 573

Chapter 34 Managing Dynamic Multipathing paths 575

About Dynamic Multi-Pathing in Veritas InfoScale Operations Manager 575
Disabling the DMP paths on the initiators of a host 576
Define new case panel options 577
Object selection panel options 579
Path disable panel options 581
Paths disable output summary 581
Disabling the DMP paths on an enclosure 582
Disabling the DMP paths on a virtualization server 583
Re-enabling the DMP paths 584
Path re-enable panel options 585
Paths re-enable output summary 586
DMP Maintenance Result Summary panel 586
Removing a completed DMP maintenance case record 587
Reviewing the output and results of a completed DMP maintenance case 588

Chapter 35 Managing CVM clusters 589

About monitoring and managing CVM clusters in Veritas InfoScale Operations Manager 589
Permissions required for views and operations on CVM cluster objects 591

Chapter 36 Managing Flexible Storage Sharing 593

Implementing Flexible Storage Sharing with Veritas InfoScale Operations Manager 593
About Flexible Storage Sharing 594
Flexible Storage Sharing use cases 595
Flexible Storage Sharing features and support in Veritas InfoScale Operations Manager 597
Exporting and un-exporting disks for Flexible Storage Sharing 598
Enabling or disabling Flexible Storage Sharing on existing shared disk groups 599

Chapter 37 Monitoring the virtualization environment 601

About the virtualization technologies supported 601
About discovering the VMware Infrastructure using Veritas InfoScale Operations Manager 603
How Veritas InfoScale Operations Manager discovers vCenter and ESX servers 603

Information that Veritas InfoScale Operations Manager discovers on the VMware Infrastructure components 604
Viewing the storage mapping information for VMware 605
About the datastores in Veritas InfoScale Operations Manager 606
About the multi-pathing discovery in the VMware environment 607
About near real-time discovery of VMware events 609
About discovering Solaris zones 610
How Veritas InfoScale Operations Manager discovers Solaris zones 611
Information that Veritas InfoScale Operations Manager discovers on Solaris zones 612
Limitations of the discovery of Solaris zones in Veritas InfoScale Operations Manager 613
About discovering logical domains in Veritas InfoScale Operations Manager 614
How Veritas InfoScale Operations Manager discovers Solaris logical domains 615
Information on logical domains that Veritas InfoScale Operations Manager discovers 616
Limitations of the discovery of logical domains in Veritas InfoScale Operations Manager 617
Viewing the storage mapping information for LDoms 617
About discovering LPARs and VIOs in Veritas InfoScale Operations Manager 618
About LPAR storage correlation supported in Veritas InfoScale Operations Manager 620
About Microsoft Hyper-V virtualization discovery 622
Virtual machine discovery in Microsoft Hyper-V 622
Storage mapping discovery in Microsoft Hyper-V 623
Viewing the storage mapping information for Hyper-V 623
About the Kernel-based Virtual Machine (KVM) virtualization discovery in Veritas InfoScale Operations Manager 624
About the reports related to virtualization 625

Chapter 38 Using Web services API 626

About using Veritas InfoScale Operations Manager Web services API 626
Logging in to Veritas InfoScale Operations Manager Web services API 627

Logging out of Veritas InfoScale Operations Manager Web services API 629
About objects supported by Veritas InfoScale Operations Manager Web services API 629
About performing operations using Veritas InfoScale Operations Manager Web services API 633
Examples of performing operations using Veritas InfoScale Operations Manager Web services API 634
Examples of performing operations using XPRTLC and cURL 636
Examples of the output in JSON format 639

Chapter 39 Veritas InfoScale Operations Manager command line interface 641

About the vomadm utility 641
Listing all configured enclosures using the vomadm utility 642
Host management using the vomadm utility 643
Deployment management using the vomadm utility 643
Business Application management using the vomadm utility 644
Service management using the vomadm utility 644
Domain management using the vomadm utility 645
List configured schedules using the vomadm utility 645

Appendix B Command file reference 647

vxlist 648
vomadm 653
xdistc 656

Appendix C Application setup requirements 661

Application setup requirements for Oracle database discovery 661
SQL queries used for the discovery of Oracle database 663
Oracle database discovery in Solaris zones 665
Application setup requirements for Oracle Automatic Storage Management (ASM) discovery 665
SQL queries used for the discovery of Oracle ASM 666
Oracle Automatic Storage Management (ASM) discovery in Solaris zones 667
Application setup requirements for IBM DB2 discovery 667
SQL queries used for the discovery of IBM DB2 database 668
Application setup requirements for Sybase Adaptive Server Enterprise (ASE) discovery 669

SQL queries used for the discovery of Sybase Adaptive Server Enterprise (ASE) 670
Application setup requirements for Microsoft SQL Server discovery 671
SQL queries used for the discovery of Microsoft SQL Server 672
Application setup requirements for Microsoft Exchange Server discovery 672

Glossary 674

Index 675

Section 1

Getting started

- Chapter 1. Introducing Veritas InfoScale Operations Manager
- Chapter 2. Using the Management Server console
- Chapter 3. Examples for using Veritas InfoScale Operations Manager

Chapter 1

Introducing Veritas InfoScale Operations Manager

This chapter includes the following topics:

- About Veritas InfoScale Operations Manager
- About increasing efficiency and productivity through storage virtualization
- Storage Foundation: a complete solution for online storage management
- Veritas InfoScale Operations Manager: tying it all together

About Veritas InfoScale Operations Manager

Veritas InfoScale Operations Manager by Veritas gives you a single, centralized management console for the Storage Foundation High Availability products. You can use it to monitor, visualize, and manage storage and cluster resources, and generate reports about these components in the Management Server domain. Veritas InfoScale Operations Manager helps administrators centrally manage diverse data center environments.

You can also use Veritas InfoScale Operations Manager to visualize and report about the hosts which do not have Storage Foundation High Availability products installed on them.

In Veritas InfoScale Operations Manager, you can establish user credentials such that authorized users can access the product to perform sensitive management tasks. Other users can perform only a basic set of operations, or can only view information.

A typical Veritas InfoScale Operations Manager deployment consists of the following:

- Management Server
- Managed hosts

A Veritas InfoScale Operations Manager deployment may also discover the following:

- Virtualization environment
- SAN/NAS or Unified storage
- SAN fabrics

About increasing efficiency and productivity through storage virtualization

The Storage Foundation products add value to your enterprise by providing storage virtualization—the process of taking multiple physical storage devices and combining them into logical (virtual) storage devices that are allocated to applications and users at will.

Storage virtualization provides a layer of abstraction between applications and the storage they use. This virtualization makes it possible for physical storage in one or more arrays to appear to the application as if it were a single file system on a host. This kind of virtual representation, known as a data container, can take several forms; disk groups and volumes are two of the most common forms.

With storage virtualization you manage data containers rather than individual blocks of storage. This enables you to enjoy the benefits of enterprise-wide server and storage administration.

Storage virtualization offers the following additional benefits:

- Reduced costs: applications have flexibility in terms of the storage they can use. While you store mission-critical data on your best and most reliable (and most expensive) storage arrays, you can store other data that is used by the same applications on less expensive media.
- Better alignment with business needs: storage tiering (or quality of storage service, QoSS) makes it possible for you to provision and use the least expensive storage tier that meets application requirements.
- Reduced risk: you realize better availability and reliability by isolating applications from storage technology decisions.
- Better management of storage: Thin provisioning and reclamation lets you manage your storage in a better manner.

Storage Foundation products continue to add features for enhanced storage utilization and performance, for example SmartIO and Flexible Storage Sharing (FSS). The SmartIO feature of Storage Foundation and High Availability Solutions supports the use of multi-vendor solid-state drives (SSDs) as a read-write cache to improve overall I/O performance. The FSS feature of Storage Foundation Cluster File System High Availability (SFCFSHA) enables network sharing of local storage, cluster wide. The local storage can be in the form of Direct Attached Storage (DAS) or internal disk drives.

See "About Veritas InfoScale Operations Manager" on page 26.

Storage Foundation: a complete solution for online storage management

TheStorage Foundation products combine the industry-leading Veritas Volume Manager, Veritas File System, and other licensed products to provide a complete solution for online storage management. The products include utilities for discovering storage resources throughout the enterprise and for monitoring and managing those resources.

With Storage Foundation products, you can:

- Group physical disks into logical volumes to improve disk utilization and eliminate storage-related downtime
- Move data between different operating systems and storage arrays
- Balance I/O across multiple paths to improve performance
- Replicate data to remote sites for higher availability
- Move unimportant or outdated files to less expensive storage without changing the way users or applications access the files. Defining various levels of tiers for the LUNs lets you manage your data in a better way.
- Manage the Cluster Server and Application HA clusters

See "About Veritas InfoScale Operations Manager" on page 26.

Veritas InfoScale Operations Manager: tying it all together

Veritas InfoScale Operations Manager ties together the various features in the Storage Foundation products in a centralized, standardized way, so your data center can run more efficiently.

Not only can the Veritas InfoScale Operations Manager operators visualize an individual host and the storage behind it; they can also see all instances of Storage Foundation that are running in the data center, across multiple operating system platforms. Having the data at a centralized place helps the operator monitor and manage the entire data center.

See "About centralized management and monitoring" on page 29.

See "About discovery, monitoring, and management" on page 29.

About centralized management and monitoring

In many data centers, the effective management of storage requires a number of different people performing different operations in different locations. For example, a storage administrator monitors and provisions the corresponding array storage, while a server administrator is tasked with ensuring the efficient use of volumes and file systems. An application administrator oversees the applications and databases that consume the storage.

With Veritas InfoScale Operations Manager, all these tasks can be performed at a central point: the Veritas InfoScale Operations Manager console. Using the console, one administrator can easily gather information, monitor and allocate resources, and perform operations on hosts, databases, applications, and storage resources throughout the data center. Thus the disparate roles and tasks of the server administrator, storage administrator, and application administrator converge in the Veritas InfoScale Operations Manager console.

The central administrator can also generate status and inventory reports and distribute the reports to others who need the information.

This centralized operation and reporting is platform and vendor agnostic. Information can be collected, and operations performed, across a variety of operating systems and on storage arrays from multiple hardware vendors. Also, you can monitor and manage the licenses that are installed for Storage Foundation and high availability products in Veritas InfoScale Operations Manager.

The Veritas InfoScale Operations Manager console, running as a thin client, is accessed through any supported Web browser. The console also provides an entry point to the individual Storage Foundation High Availability products.

See "Veritas InfoScale Operations Manager: tying it all together" on page 28.

About discovery, monitoring, and management

Veritas InfoScale Operations Manager includes discovery, administration, and active monitoring of Storage Foundation hosts, DMP, VCS, and ApplicationHA.

System administrators can monitor resources of managed hosts in Veritas InfoScale Operations Manager, select one, and actively manage its attributes in an associated Storage Foundation product—seamlessly within the same browser window. For example, system administrators would want to recover the faulted disks that are listed in Veritas InfoScale Operations Manager. Administrators can select the faulted disk in Veritas InfoScale Operations Manager, drill down to the details of the disk, and manage the recovery of the disk.

See “Veritas InfoScale Operations Manager: tying it all together” on page 28.

Chapter 2

Using the Management Server console

This chapter includes the following topics:

- About the Management Server console
- About the Management Server console Home page
- About the Global Dashboard
- About the elements in a perspective view
- About viewing the summarized information
- About selecting the objects
- About drilling down to more information on the selected object
- About searching for objects
- About viewing the solutions
- About viewing the reports
- About viewing tasks
- About viewing connectivity graphs
- About accessing the Veritas InfoScale Operations Manager Help

About the Management Server console

The Management Server console facilitates perspective driven, role-based access for better manageability of different assets across your data center. Task-oriented

interface and a consistent placement of user interface elements make it easier to navigate across the perspectives and perform the required operations.

See "About the elements in a perspective view" on page 35.

See "About the Management Server console Home page" on page 32.

See "About viewing the solutions" on page 46.

See "About searching for objects" on page 40.

See "About accessing the Veritas InfoScale Operations Manager Help" on page 49.

About the Management Server console Home page

The Figure 2-1 displays the Home page of the Management Server console. Click the icons to perform a specific category of task or view a specific category of information that is related to your data center. For example, click **Availability** to perform the tasks on clusters and service groups.

Figure 2-1 Management Server console



The following options are available:

- **Global Dashboard**: Provides an overall status of the data center (faults, risks, platform distribution, object status), and the information on inventory. You can

also create a customized Dashboard. See "About the Global Dashboard" on page 34.

- **Settings**: Use this to set up the Management Server environment. For example, adding managed hosts, configuring virtualization servers, uploading solutions, configuring SAN switches, configuring storage enclosures, and so on.
  An administrator needs to first perform the setup-related operations, and then use the perspectives for the specific tasks. For example, add the enclosures using the Storage Insight Add-on, and then use the **Storage** perspective to monitor the enclosures.
  Settings is also referred to as the Management Server perspective.
- Perspectives: Depending on the role of the user and the nature of the task to be performed, the tasks are grouped under the following perspectives:
  - **Storage**: Lets you view the information on storage enclosures, associated faults and risks, SAN fabric, switches, and generate reports on storage utilization. You can also perform dynamic multipathing related operations on the enclosures.
  - **Server**: Lets you view the information on hosts and applications (databases, Microsoft Exchange, Virtual Business Service), view associated faults and risks, and generate reports on hosts, volumes, disks, and so on. You can create disk group, volume, file system, configure Application HA, and perform dynamic multi-pathing related operations on the hosts. The Server perspective provides a similar user interface and functionality as that of the VEA Java GUI.
  - **Virtualization**: Lets you view the information on the configured virtualization servers, their associated virtual machines, associated faults and risks, and generate the reports on the virtualization objects and the associated storage. You can also perform dynamic multi-pathing related operations on the VMware ESX server.
  - **Availability**: Lets you perform the operations that are related to clusters, Virtual Business Services, and service groups. For example, you can create service group, Virtual Business Service, and recovery plan. You can also generate reports on High Availability-related objects. The Availability perspective provides a similar user experience and functionality as that of the VCS Java GUI and VCS Management console (VCSMC).
- **SFHA Licensing**: Provides the information on Storage Foundation High Availability licenses, the license deployment details, and reports.

See "About searching for objects" on page 40.

See "About selecting the objects" on page 38.

See "About viewing the summarized information" on page 37.

See "About accessing the Veritas InfoScale Operations Manager Help" on page 49.

About the Global Dashboard

The Global Dashboard provides a quick visual summary of the information in other Veritas InfoScale Operations Manager perspectives. It contains graphic tables and charts that show the status and inventory for all objects that Veritas InfoScale Operations Manager discovers.

The Global Dashboard contains a set of customizable, interactive dashboards, each on a separate tab. Two dashboards are provided by default:

Overview	Includes information on faults, risks, object status, and usage.
Inventory	Includes inventory count and distribution.

Table 2-1 Global Dashboard overview

If you want to...	Do the following...
Navigate (drill down) from a dashboard to a perspective	Click the right-arrow icon on the top right of the chart or table. You are prompted to confirm the navigation.
Remove a chart or table from a dashboard	Click the unpin icon on the top right of the chart or table. You can hover on the icon to view the label **Remove from Dashboard**. You are prompted to confirm the deletion.
Move a chart or table to a different position	Place the cursor on the title bar of the chart or table and drag it.
Add charts or tables to existing dashboards	From any **Overview** tab in a perspective, on the top right of the chart or table that you want to add, click the pin icon. You can hover on the icon to view the label **Add to Dashboard**. In the **Add to Dashboard** window, select the name of the dashboard and click **OK**.

Table 2-1 Global Dashboard overview *(continued)*

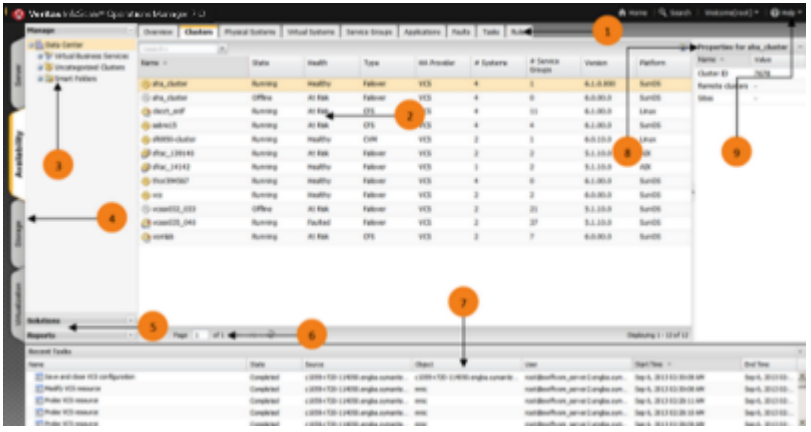
If you want to...	Do the following...
Create a custom dashboard	From any **Overview** tab in a perspective, on the top right of a chart or table, click the pin icon. You can hover on the icon to view the label **Add to Dashboard**. In the **Add to Dashboard** window, type the name of a new dashboard and click **Create**.
Rename a custom dashboard	From the dashboard, right-click the name of the tab and click **Rename**. Enter a new name and click **OK**.
Delete a custom dashboard	From the dashboard, right-click the name of the tab and click **Delete**. You are prompted to confirm the deletion.

See "About the Management Server console" on page 31.

About the elements in a perspective view

The Figure 2-2 displays the console of the **Availability** perspective. Similar user interface elements are used in other views of the Veritas InfoScale Operations Manager Management Server console.

Figure 2-2 User interface elements of the Availability perspective on the Management Server console



The user interface elements are described below:

Table 2-2 User interface elements of the Availability perspective on the Management Server console

Number	User interface element
1	Tabs: Lists the detailed information about a specific object. For example, **Service Groups** tab lists service groups with their properties.
2	Content pane: Based on your selection in the navigation tree, the content pane displays the details that are related the selected object. You can double-click the record to view its details.
3	Navigation tree: Use it to navigate across the Organizations and the objects under the **Data Center** node. For all perspectives, **Data Center** is the parent node. The uncategorized folder contains the objects that are currently not added to any Organization. You can use the **Organize** option to move an object from the uncategorized folder to the desired Organization. **Note:** Depending on the perspective, the uncategorized folder contains different objects. For example, in the storage perspective, the uncategorized folder contains enclosures. For the availability perspective, clusters are placed in the uncategorized folder.
4	Perspectives: Lists the available perspectives.
5	Section: Lists the sections for each perspective. For instance, solutions for the perspectives, and the types of reports that you can generate for the perspective.
6	Table pagination: Displays the records on the specified page.
7	Tasks Pane: Displays the list of recent tasks with the status.
8	Properties pane: Lists the additional properties that are related to the selected record in the content pane. You can move fields from the content pane to the properties pane. See "Moving fields to the properties pane of the Management Server console" on page 38.
9	Menu bar: Provides access to the Help. On the perspectives, the menu bar also provides a **Search** option for searching for objects in the perspective. See "About accessing the Veritas InfoScale Operations Manager Help" on page 49. See "About searching for objects" on page 40.

About viewing the summarized information

The **Overview** tab is present for all objects across all perspectives. It provides a snapshot of overall storage and high availability statistic in your data center including the graphs on faults, alerts, storage allocation, breakup by vendor, and other perspective-specific information.

Figure 2-3 Data Center summary

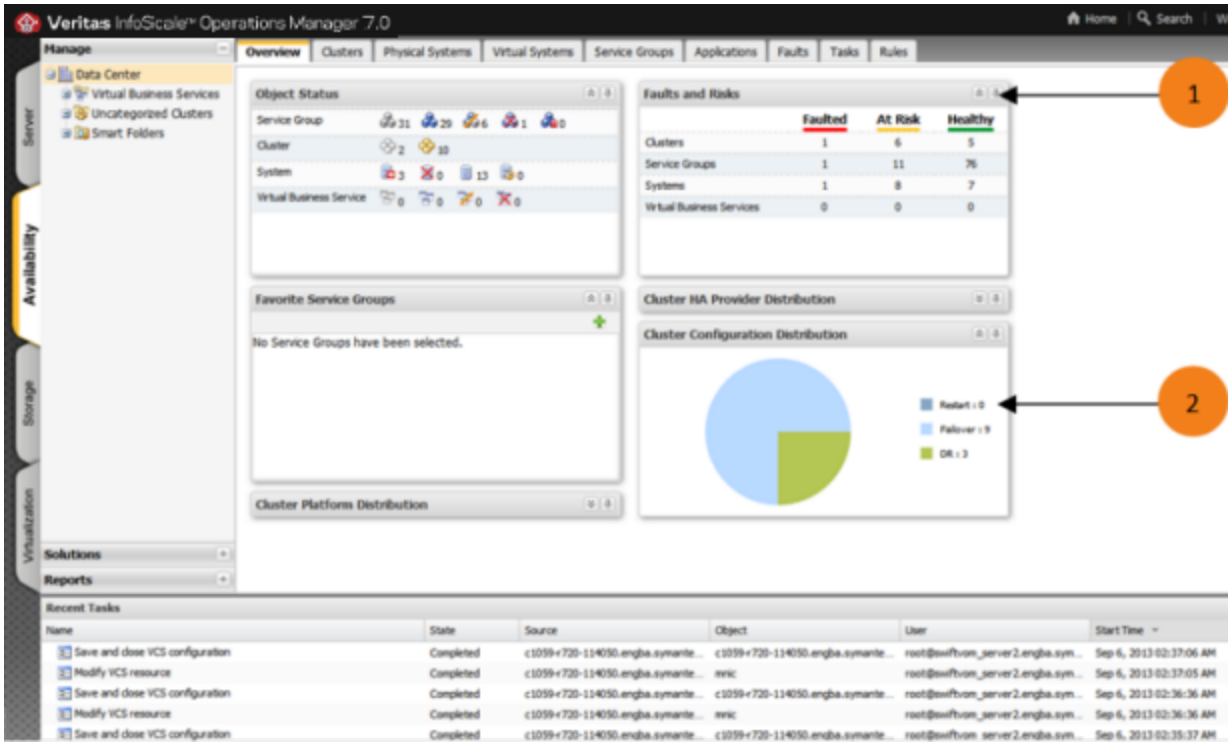


Table 2-3 User interface elements

Number	User interface element
1	Click to add the portlet to the Global Dashboard.
2	Click the legend to turn on or turn off the display of the values in the chart.

About selecting the objects

On the Management Server console, you can select the objects using the following methods. The availability of options depends on your privileges, perspective, and the object.

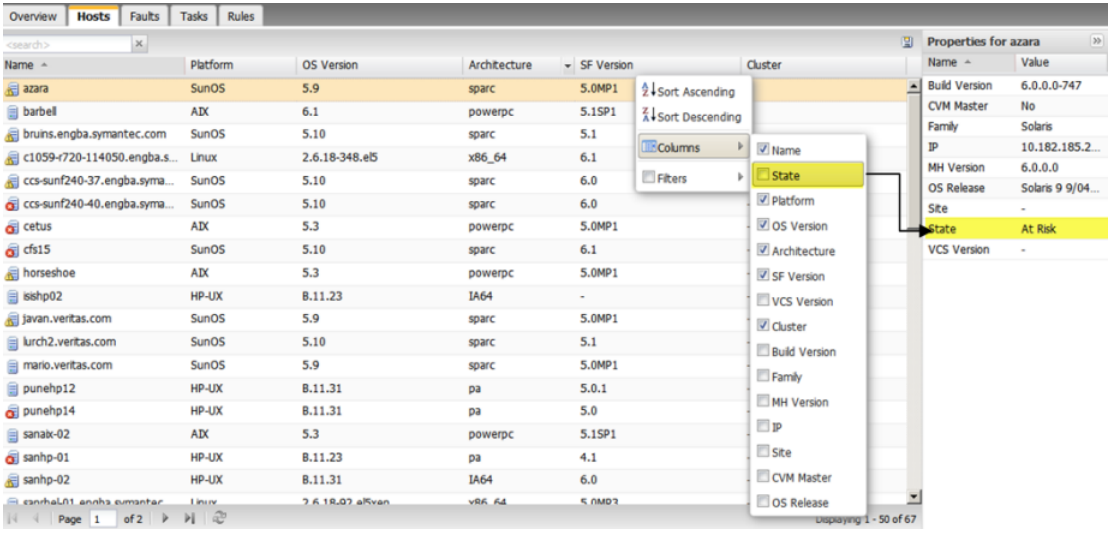
- **Navigation tree**: In the navigation tree, select the required object, and right-click to view its context-sensitive menu.
- **Content pane**: In the content pane, select the required record, and right-click to view the available options. Press Ctrl or Shift for the selection of multiple records.

See "About the Management Server console" on page 31.

Moving fields to the properties pane of the Management Server console

The Management Server console provides you with the flexibility to move fields from the content pane to the properties pane and vice versa. It helps you group most relevant or frequently used fields together for the selected storage object. As illustrated in Figure 2-4, the **State** field is removed from the content pane, and it appears under the **Properties** pane.

Figure 2-4 Selection and removal of fields



See "About the Management Server console" on page 31.

See “About the Management Server console Home page” on page 32.

About drilling down to more information on the selected object

On the Management Server console, double-clicking provides a fast way to navigate from an object record on a table in the content pane to its node in the navigation tree.

This action works only if the type of object that is double-clicked is an object type that appears in the tree. If so, double-clicking the record on the content pane selects the corresponding object in the tree. Otherwise, nothing occurs with the double-click.

For example, if you display the list of hosts for the data center in the content pane, you can filter to a specific host and then double-click the host to go to its node in the tree. However, if you display the list of disks for a host in the content pane and double-click on a record, nothing occurs, since disks do not appear in the tree.

If you double-click on an object that is in the tree but on a different perspective, the console asks you to confirm navigating to the different perspective. The console then opens a browser tab for the new perspective, with the object node selected in the tree.

For example, if you double-click an enclosure for a host on the Server perspective, the action navigates you to the Storage perspective (assuming that you have permission to view that perspective).

In some cases, a tab in the content pane displays subtabs at the bottom. In such a case, selecting a record with a single-click displays more information about that object in the subtabs.

The **Faults and Risks** chart in the data center **Overview** tab for each perspective also lets you drill down to the faulted or at risk objects in the tree. This navigation requires two steps. You first double-click the object type on the chart. The table of faulted and at risk objects of that type is displayed. If the object is of a type that is shown in the tree, you can then double-click the record in the table, which navigates you to the corresponding object in the tree.

See “About selecting the objects” on page 38.

See “About the Global Dashboard” on page 34.

See “About the Management Server console Home page” on page 32.

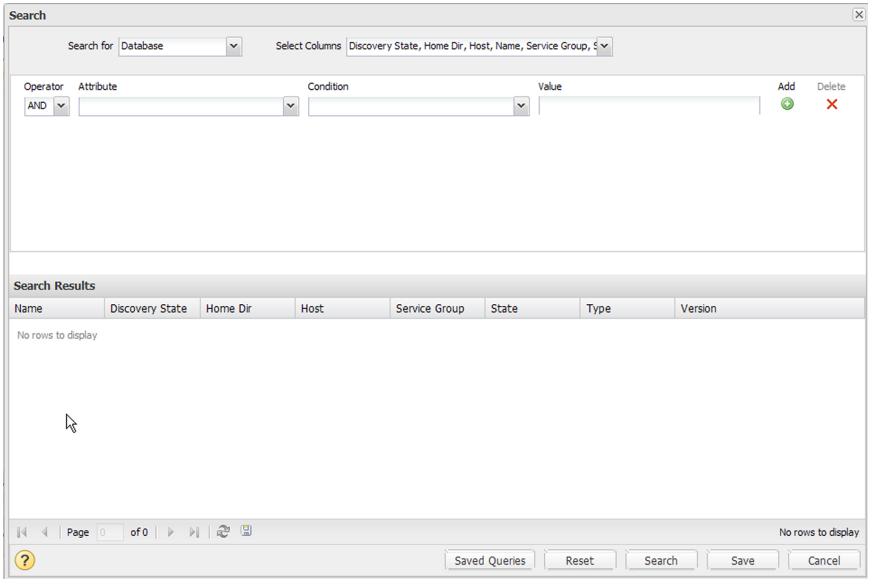
About searching for objects

In the Management Server console, you can use the **Search** button on the menu bar to search for the main objects in the following perspectives:

Server	The objects include Host and related objects such as Disk Group, Disk, Volume, Snapshot, RVG, Enclosure, and Initiator. You can also search for other objects available in this perspective, including Database, Exchange Server, and Package.
Availability	The objects include Cluster and related objects such as Service Group, System, and Resource.
Storage	The objects may include Enclosure, Array Port, LUN/LDEV, RAID Group, Fabric, Switch, Physical Disk, Thin Pool, Storage Rank, Share, Storage Volume, and Qtree. The available objects depend on which add-ons are installed. For example, fabric and switch objects require Fabric Insight Add-on.
Virtualization	The objects include Virtualization Server, Virtual Machine, and Datastore.

For each object search you can specify one or more parameters using an attribute, a condition (operator), and a value. The predefined attributes change depending on the selected object type.

Figure 2-5 Search window



You can create searches for commonly viewed objects and save the searches for later use. For the primary object in a perspective, such as hosts in the **Server** perspective, you also have the choice to save search queries as Smart Folders in the tree.

See "Creating new search queries" on page 41.

See "Managing saved search queries" on page 43.

See "Creating and saving searches as Smart Folders" on page 43.

See "Managing Smart Folders" on page 44.

Creating new search queries

In the Management Server console, you can use the **Search** button on the menu bar to search for a variety of objects in the **Server**, **Availability**, **Storage**, and **Virtualization** perspectives.

You can create searches for commonly viewed objects in a perspective and save the searches for later use.

Note: For the primary object in a perspective, such as hosts in the **Server** perspective, you also have the choice to save search queries as Smart Folders in the tree.

See "Creating and saving searches as Smart Folders" on page 43.

To create a new search

1. In the Management Server console, go to the perspective and click **Search** on the menu bar.
2. If you have previously saved queries, they are displayed in the **Search** window. Click **New Search**.
3. Select from the following options:

Search for	Select the object that you want to search for
Select columns	Choose the columns to appear in the search results. You can also click in the **Custom Name** field and type a custom name for that column.
Operator	You can select AND or OR when adding rules to a search query. You cannot mix operators within the search query. However, if you have selected AND as the operator, and your query includes multiple rules that use the same attribute, Veritas InfoScale Operations Manager will apply the OR operator to the rules with the same attribute.
Attribute	Select the attribute for which you want to specify a value.
Condition	Select the condition, for example, **Starts With**.
Value	Type or select the value, depending on the attribute. Value strings are not case-sensitive.
Add	Click to add another parameter to the search query.
Remove	Click to remove a parameter from the query.

4. Click the **Search** button at the bottom of the window. The results are displayed.
5. To save the search query for future use, click **Save** and specify a name and optionally a description. Click **Save**, then click **OK**. To view the query on the list of saved searches, click **Saved Queries**.

See “About searching for objects” on page 40.

Managing saved search queries

In the Management Server console, you can manage the search queries that you previously created and saved in the **Search** window. You can modify, rename, and delete saved search queries. See the following procedures:

See “To modify a saved search” on page 43.

See “To rename a saved search” on page 43.

See “To delete a saved search” on page 43.

To modify a saved search

1. In the Management Server console, click **Search** on the menu bar.
2. Select the saved search query you want to modify.
3. Click **Modify**.
4. Change the search parameters as desired.
5. To save the changes, click **Save**.

To rename a saved search

1. In the Management Server console, click **Search** on the menu bar.
2. In the list of saved searches, click the icon in the **Rename** column in the row for the search you want to rename.
3. Change the name or description as desired and click **Rename**.

To delete a saved search

1. In the Management Server console, click **Search** on the menu bar.
2. Select the search query you want to delete.
3. Click the icon in the **Delete** column. The search query is removed.

See “About searching for objects” on page 40.

Creating and saving searches as Smart Folders

In the Management Server console, you can use the **Smart Folders** node in the tree to create and save searches regarding the primary object in the current perspective. For example, in the **Server** perspective, you can use **Smart Folders** to search for hosts by various attributes.

Note: To search for objects other than the primary object in the perspective, use the **Search** button on the menu bar.

See "Creating new search queries" on page 41.

To create and save a search as a Smart Folder

1. In the Management Server console, go to the perspective and expand **Manage** in the left pane.
2. Right-click **Smart Folders** and click **New Smart Folder**.
3. In the **Create Smart Folder** window, select from the following options:

Operator	Choose whether to use an AND or OR operator for additional parameters. You cannot mix operators within the search query.
Attribute	Select the attribute for which you want to specify a value.
Condition	Select the condition, for example, **Starts With**.
Value	Type or select the value, depending on the attribute. Value strings are not case-sensitive.
Add	Click to add another parameter to the search query.
Delete	Click to remove a parameter from the query.

4. If you want to preview the query results, click **Preview**.
5. To save the search query for future use as a Smart Folder, click **Save** and specify a name and optionally a description. Click **Save**, then click **OK**.
6. To view the new Smart Folder in the tree, expand **Smart Folders** and select the Smart Folder.

See "About searching for objects" on page 40.

See "Managing Smart Folders" on page 44.

Managing Smart Folders

In the Management Server console, you can manage the search queries that you previously saved as Smart Folders.

You can modify, rename, copy, and delete Smart Folders.

See "To modify a Smart Folder" on page 45.

See "To rename a Smart Folder" on page 45.

See "To copy a Smart Folder" on page 45.

See "To delete a Smart Folder" on page 45.

To modify a Smart Folder

1. In the Management Server console, go to the perspective and expand **Manage** in the left pane.
2. Expand **Smart Folders**.
3. Right-click the Smart Folder you want to modify and select **Modify**.
4. Change the search parameters as desired.
5. Click **Save**.

To rename a Smart Folder

1. In the Management Server console, go to the perspective and expand **Manage** in the left pane.
2. Expand **Smart Folders**.
3. Right-click the Smart Folder you want to rename and select **Rename**.
4. Change the name or description as desired and click **Save**.

To copy a Smart Folder

1. In the Management Server console, go to the perspective and expand **Manage** in the left pane.
2. Expand **Smart Folders**.
3. Right-click the Smart Folder you want to copy and select **Copy**.
4. Edit the name and description of the copy as needed and click **Save**.

To delete a Smart Folder

1. In the Management Server console, go to the perspective and expand **Manage** in the left pane.
2. Expand **Smart Folders**.
3. Right-click the Smart Folder you want to delete and select **Delete**.
4. Confirm the deletion by clicking **Delete**.

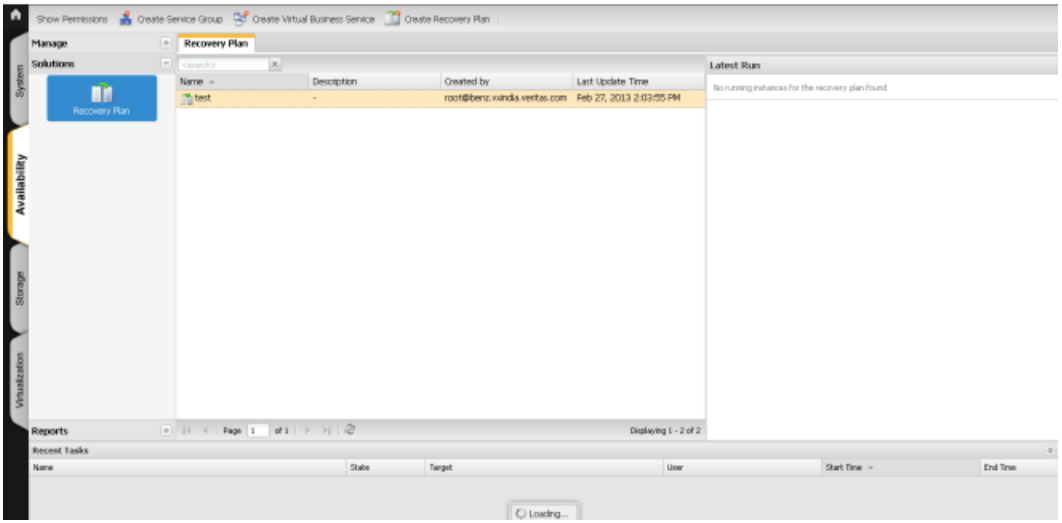
See "About searching for objects" on page 40.

See "Creating and saving searches as Smart Folders" on page 43.

About viewing the solutions

You can use the Management Server console to view the solutions that are available for the selected perspective. For example, in the Figure 2-6, Recovery Plan is shown as the available solution for the Availability perspective.

Figure 2-6 Solutions for the Availability perspective



Veritas InfoScale Operations Manager also provides additional functionality in the form of add-ons. After the successful deployment of the add-on, additional user interface elements (for example, tabs) are displayed on the relevant pages of the Management Server console. For example, after you enable the deep discovery of a storage enclosure using the Storage Insight Add-on, additional tabs are displayed for the storage enclosure. The available add-ons are listed under **Settings** > **Deployment** > **Repository**.

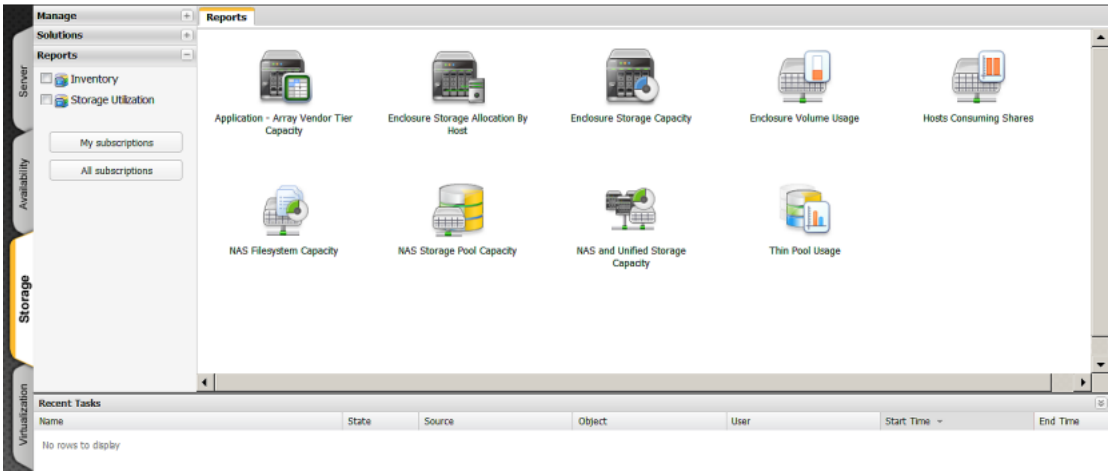
See “About the Management Server console” on page 31.

See “About the Management Server console Home page” on page 32.

About viewing the reports

You can generate and view the reports that are specific to the perspective. The Figure 2-7 displays the reports view of the storage perspective.

Figure 2-7 Management Server reports for the Storage perspective



See "About the Management Server console" on page 31.

See "About the Management Server console Home page" on page 32.

About viewing tasks

A task lets you view the recently executed operations in your data center. Depending on the nature of the operation, you can view the tasks for the following:

- Task and audit related to the general configuration in the data center. For example, adding hosts, refreshing repository, or adding a virtualization server to the Management Server. To view this category of tasks, navigate to **Settings** from the home page of the Management Server console, and select **Audit** and **Tasks** section.
- For each perspective, the **Tasks** tab at the data center-level provides the details of the tasks related to the perspective. For example, in the **Server** perspective, upon deletion of a volume, a task Delete Volume is listed under the **Tasks** tab.

See "About searching for objects" on page 40.

See "About the Management Server console" on page 31.

See "About the Management Server console Home page" on page 32.

About viewing connectivity graphs

The **Connectivity Graph** window shows the connections from an object to its storage components. For example, for a host, you can view the host object linked to HBAs, with the HBAs linked to enclosures.

You can view connectivity graphs for the following objects on the Server and Storage perspectives.

Table 2-4 Objects with connectivity graphs

Perspective	Available connectivity graphs
Server	■ Database: Shows connected host and enclosures. The host can expand to show volumes, paths, and HBAs. Enclosures can expand to show array ports. Enclosures can show connections to additional objects, such as LUNs and physical disks, if Storage Insight Add-on is enabled. ■ Host: Shows connected HBAs and enclosures. ■ Volume: Shows connected paths, HBAs, and enclosures. Enclosures can expand to show array ports. Enclosures can show connections to additional objects, such as LUNs and physical disks, if Storage Insight Add-on is enabled. ■ Disk Group: Shows connected HBAs and enclosures. Enclosures can expand to show array ports.
Storage	Enclosure: Shows connected host. Shows switch and fabric connectivity if Fabric Insight Add-on is enabled.

Following are some tips on using the **Connectivity Graph** window:

- Hover the mouse over an object for a pop-up showing object type and name
- Double-click on an object to expand it (not all objects expand) and click again to close it.
- Click **Zoom** and drag the slider to zoom in or out.
- Click the double-arrow next to **Properties** to open and close the **Properties** pane.

To view a connectivity graph

1 In the perspective, navigate to the object in the tree.

2 Right-click the object and click **Connectivity Graph**.

See “About the Management Server console” on page 31.

About accessing the Veritas InfoScale Operations Manager Help

You can access the product Online Help using the following methods:

- Online mode: In this help mode, the help content in hosted on the web and you need to access the designated server to use the Online Help. It contains the help for the Veritas InfoScale Operations Manager Management Server and its add-ons. For the latest online Help, use the online mode.
- Offline mode: In the offline mode, you need to first download and install the Veritas InfoScale Operations Manager Help add-on on the Management Server. The Veritas InfoScale Operations Manager Help add-on installs the required Help components. After the successful installation of the add-on, you can view the Help for the Veritas InfoScale Operations Manager Management Server and its supported add-ons.

See “About the Management Server console” on page 31.

See “About the Management Server console Home page” on page 32.

Chapter 3

Examples for using Veritas InfoScale Operations Manager

This chapter includes the following topics:

- Example: Creating a volume using Veritas InfoScale Operations Manager
- Example: Creating a service group, adding it to a cluster, and bringing it online using Veritas InfoScale Operations Manager
- Example: Cluster Server troubleshooting using Veritas InfoScale Operations Manager
- Example: Ensuring the correct level of protection for volumes controlled by Storage Foundation
- Example: Improving the availability and the disaster recovery readiness of a service group through fire drills
- Example: Identifying the performance issues of an application using Veritas InfoScale Operations Manager
- Example: Volume migration using Veritas InfoScale Operations Manager
- Examples: Identifying and reducing storage waste using Veritas InfoScale Operations Manager

Example: Creating a volume using Veritas InfoScale Operations Manager

As a user with an Admin role, you are responsible for managing the Storage Foundation product suite that is installed in the data center. Veritas InfoScale Operations Manager provides you a single and a centralized management console for the Storage Foundation and high availability products. You can use Veritas InfoScale Operations Manager to monitor, visualize, and manage storage resources.

Storage Foundation product suite includes multiple Veritas products and optionally licensed features, like Veritas Volume Manager (VxVM), Veritas File System (VxFS), and Volume Replicator (VVR). You can use these products to manage the storage resources in your data center. For example, you can use Veritas Volume Manager to manage physical disks as volumes, which are logical devices.

To perform this task, your user group must be assigned the Admin role on the host or the Server perspective. The permission on the host may be explicitly assigned or inherited from a parent Organization.

The following example describes how you can use Veritas InfoScale Operations Manager to create a volume on a UNIX or Linux managed host.

This example uses the following:

Example_vol	Volume that you create using Veritas Volume Manager.
Example_linux_host	Linux managed host on which you create the Example_vol volume.

To create the Example_vol volume on the Example_linux_host managed host, you do the following:

- In the Management Server console, go to the **Server** perspective and expand **Manage** in the left pane.
- Expand Organization, or **Uncategorized Hosts** to locate and select the host.
- Do one of the following:
  - Right-click on the host and select **Create Volume**.
  - Expand the host and locate the disk group to create the volume. Right-click and select **Create Volume**.
  - Expand the host and select **Volumes**. Right-click and select **Create Volume**.
- In the wizard panels to create a new volume, enter the following information:

Create Volume	Select the disk group on which you want to create the Example_vol volume.
Disks selection option	Specify if you want to select the disks manually.
Volume Attributes	Enter the parameters for the Example_vol volume.
Create cache volume	Enter the snapshot cache volume parameters for the Example_vol volume.
Create File System - File System Options	Select the file system options to create a new file system.
Advanced mount options	Enter the file system options and mount options while creating a new file system.
Create Volume Summary	Click **Finish** to create the volume.

Now, you have successfully created the Example_vol volume on the Linux managed host Example_linux_host.

Example: Creating a service group, adding it to a cluster, and bringing it online using Veritas InfoScale Operations Manager

Veritas InfoScale Operations Manager provides you a single and centralized management console for the Storage Foundation and high availability products. You can use Veritas InfoScale Operations Manager to perform the VCS-related high availability and disaster recovery operations including the following operations:

- Adding a service group to a cluster.
- Bringing a service group online.
- Taking a service group offline.
- Switching the service groups on a remote cluster.
- Making a service group local or global.

The following example describes how you can use Veritas InfoScale Operations Manager to add a service group to a globally configured cluster and bring it online.

This example uses the following:

Example_cluster	Globally configured cluster to which you add a service group.
Example_sg	The **Failover** type service group that you want to add to the **Example_cluster** cluster and bring online.

As part of creating the Example_sg service group, you need to add the following resources and establish their dependencies:

Resource name	Resource type
vol_1	Volume
dg_1	Disk group

When you configure the dependencies of these resources, select vol_1 as the parent resource and dg_1 as the child resource.

To create and add the Example_sg service group to the Example_cluster cluster, you do the following:

- In the Management Server console, go to the **Availability** perspective and select **Manage** in the left pane.
- Expand the Organization or **Uncategorized Clusters** to locate the Example_cluster cluster.
- Right-click on the Example_cluster cluster and select **Create Service Group**.
- In the wizard panels to create and add a service group to a cluster, enter the following information:

Create Service Group	Specify the name of the service group as **Example_sg**. Specify the type of the service group as **Failover**. **Note:** The name that you enter must start with a letter. The name of the service group can contain only letters, numbers, hyphens, or underscores.
Configure System List	Select one or more hosts for the service groups in the **Available Systems** field and move them to the **Systems in Priority Order** field.

Configure Resources	Specify the names and types for the vol_1 and dg_1 resources. Add them to the **Resource List** table.
Resource Dependencies	Select vol_1 as the parent resource and dg_1 as the child resource. Click **Link** to add the selected resources to the **Dependencies** table. Click **Finish**.

To view the Example_sg service group that you have added to the Example_cluster cluster, do the following:

- In the Management Server console, go to the **Availability** perspective and expand **Manage** in the left pane.
- Expand the Organization, or **Uncategorized Clusters** to locate the Example_cluster cluster and then expand the cluster.
- Expand the **Service groups** node and then select the Example_sg service group.
- Click the **Overview** tab.

By default, the Example_sg service group that you have added to the Example_cluster cluster is in offline mode. To put the Example_sg service group into a responsive and functioning state, you must manually bring it online.

To bring the Example_sg service group online, do the following:

- In the Management Server console, go to the **Availability** perspective and select **Manage** in the left pane.
- Expand the Organization or **Uncategorized Clusters** to locate the Example_cluster cluster.
- Expand the cluster and then the **Service Groups** node to locate the Example_sg service group.
- Right-click on the service group and select **Online**.
- In the **Online Service Group** panel, select the host on which you want to bring the service group online and click **OK**.

For more information, refer to the *Veritas InfoScale Operations Manager Add-ons User Guide*.

Example: Cluster Server troubleshooting using Veritas InfoScale Operations Manager

Cluster Server (VCS) is a high availability clustering solution from Storage Foundation High Availability products for UNIX, Linux, and Microsoft Windows operating systems. Veritas InfoScale Operations Manager, being a centralized management console for Storage Foundation HA products, lets you perform various VCS-related operations using Management Server console.

In Veritas InfoScale Operations Manager, clusters, service groups, and replication features provide high availability and disaster recovery of applications. In the Management Server console, Availability perspective provides you a snapshot of high availability and disaster recovery status of your data center. It includes service group states, associated faults, risks, the details about clusters, and replications. This consolidated view lets you perform the following operations:

- View the status of objects that are available in the Availability perspective. These objects include service groups, clusters, systems, and Virtual Business Services. You can view the status under **Object Status** in **Overview** at data center level.
- View top faults and risks under **Faults and Risks** in **Overview** at data center and Organization level. Form this view, you can drill down to the source of fault.
- Monitor the status of selected service groups. You can add service groups that you want to monitor under **Favorite Service Group** in **Overview** at data center level.

See "Service group dependencies – on resources and other service groups" on page 55.

See "Service groups faults" on page 56.

See "Clearing faults for the service group's resources" on page 56.

See "Veritas InfoScale Operations Manager reports for Cluster Server troubleshooting " on page 57.

Service group dependencies – on resources and other service groups

In Veritas InfoScale Operations Manager, a cluster contains one or more service groups, and each service group depends on its constituents resources as well as on other service groups. For each service group, you can view the service group's resources, resources dependencies, and dependency of a service group on other service groups. Use the **Link** option of a service group to establish a service group's relationship with other service groups. However, you can select only those service groups that are part of the same cluster. These are inter-cluster relationships.

Similarly, to edit a service group's resources, use the **Edit** > **Resources** > **Add/Modify Resources** option.

If you want to establish service group dependencies across clusters, and perform controlled service group online and offline operations, you can use Virtual Business Services (VBS). It also includes virtual machine auto start and stop options. These virtual machines are managed by VMware vCenter Server. You can add the virtualization servers using **Settings** > **Virtualization**.

See "Example: Cluster Server troubleshooting using Veritas InfoScale Operations Manager" on page 55.

Service groups faults

A service group is a logical container, which consists of various resources. These resources are categorized in various resources types. Cluster Server uses agents to monitor the status of these resource types. Each agent monitors one resource type. Typically, a service group may fault due to the following reasons:

- Service group resources are in the faulted or unknown state. A fault occurred at the resource level may cause service group failure. You need to first resolve the detected failure (physical removal of any hardware, for instance), and then clear the resource fault using Veritas InfoScale Operations Manager console. An example will be a faulty NIC. In this scenario, first the NIC needs to be replaced, thereafter you need to use the **Clear Fault** option for a selected service group.
- Another scenario is the dependency of a service group's resource on other resources. You can view the resource dependencies under the **Resource Dependency** tab for a service group.
- If a service group is dependent on other service group, the failure of this service group may also cause faults on the former service group. You can view this dependency under the **Service Group Dependency** tab for a service group.

See "Example: Cluster Server troubleshooting using Veritas InfoScale Operations Manager" on page 55.

Clearing faults for the service group's resources

A service group may contain various Cluster Server resources. For example, Disk group, IP, NIC, and so on. Clearing faults on the service group's resources is essential before you bring the faulted service group online. This section explains how to clear fault on a service group's resources. In the same context, note that if the disk is faulted for a particular service group, you might need to perform disk recover and disk replace operations.

You can clear a service group fault by removing the resource faults within that service group. This is essentially a batch operation where you clear faults collectively for all the resources of the selected service group.

To clear faults on service group

1. In the Management Server console, go to the **Availability** perspective and select **Manage** in the left pane.
2. Expand the Organization or **Uncategorized Clusters** to locate the cluster.
3. Expand the cluster and then the **Service Groups** node to locate the service group.
4. Right-click on the service group and select **Clear Fault**.
5. In the **Clear Fault Service Groups** panel, select the system to clear the fault on a specific system. Choose **All Systems** to clear the fault on all the systems. Click **OK**.
6. In the **Result** panel, click **OK**.

If you want to clear the fault for a specific resource, you need to perform the action for individual resource.

To clear faults on resources

1. In the Management Server console, go to the **Availability** perspective and select **Manage** in the left pane.
2. Expand the Organization or **Uncategorized Clusters** to locate the cluster.
3. Expand the cluster, the **Service Groups** node, the service group, and then the resource type to locate the resource.
4. Right-click on the required resource and select **Clear Fault**.
5. In the **Clear Fault Resource** panel, select the system where you want to clear the resource. Choose **All Systems** to clear the resource on all the systems. Click **OK**.
6. In the **Result** panel, click **OK**.

See "Example: Cluster Server troubleshooting using Veritas InfoScale Operations Manager" on page 55.

Veritas InfoScale Operations Manager reports for Cluster Server troubleshooting

You can generate various reports related to trends, activity, and analysis for all HA objects – clusters, service groups, and resources. You cannot run the report for a single service group. You can either run the report on a selected Organization, or

the entire data center. If you select an Organization, all the service groups that are associated with the Organization are scanned.

The following types of reports are generated:

- Activity reports: Cluster Activity, Service Group Activity, and Resource Activity
- Uptime Analysis
- Resource Fault Trending
- Failover Summary

See "Activity reports" on page 58.

See "Uptime Analysis" on page 59.

See "Resource Fault Trending" on page 59.

See "Failover Summary" on page 59.

See "Example: Cluster Server troubleshooting using Veritas InfoScale Operations Manager" on page 55.

Activity reports

These reports capture the change in Cluster Server attributes for high availability objects. For example, the Service Group Activity Report provides information about the change in attributes value for service groups for a specific time duration. Similar reports are generated for resources and clusters. Consider the following example:

Cluster	Service Group	Date	Attribute Name	New Value	Old Value
VCS_Cluster1	Test_SG	Apr 21, 2011 12:49:32 PM	AutoFailOver	True	False

It shows that value of AutoFailOver value has changed from False to True for Test_SG service group.

Some of the attribute values can be set by the user (example - AutoFailOver), and some attribute values are discovered by Cluster Server. The discovered attribute values are read-only to the user (example - State).

For example, if the online resource faults, Cluster Server discovers the change in its State attribute's value. Therefore, when you run the report, this change in the attribute value is captured, and displayed to the user.

Another example will be of AutoFailOver attribute. By defaults, its value is set to True. It specifies that in case of system failure, Cluster Server tries to failover the service group on other available system. However, when you change the value of

AutoFailOver attribute to False, the service group, despite the failure of its current system, will not failover to any other system. The changes are captured in the activity report.

You can generate the following types of activity reports:

- Cluster Activity
- Service Group Activity
- Resource Activity

See "Veritas InfoScale Operations Manager reports for Cluster Server troubleshooting " on page 57.

Uptime Analysis

This report provides information about service group's state over the specified time. Also, the report summarizes the events that affect the online availability of the service groups. This information is useful for troubleshooting.

See "Veritas InfoScale Operations Manager reports for Cluster Server troubleshooting " on page 57.

Resource Fault Trending

This report provides information about resource types that are unstable. You can view the most unstable resource types that generated the maximum faults. From Cluster Server perspective, this information is useful as the Cluster Server uses its agents to monitor the resource types.

See "Veritas InfoScale Operations Manager reports for Cluster Server troubleshooting " on page 57.

Failover Summary

In Veritas InfoScale Operations Manager, whenever a service group restarts on the same node or on another node, it is treated as failover.

This report displays information on the failover that Cluster Server (VCS) supports to ensure high availability. The Failover Summary report displays information on the following types of failovers that are associated with a cluster:

Failover type	Description
Planned failover	The failover that you have planned and executed with the Online, Offline, or the Switch operations.

Failover type	Description
Unplanned automatic failover	The failover that occurred because of faults in the datacenter, which would have affected the high availability adversely. Cluster Server has performed the failover operation without your interference.
Unplanned manual failover	The failover that occurred because of faults in the datacenter. VCS required your interference to clear the faults, or perform the corrective action to complete the failover.

See "Veritas InfoScale Operations Manager reports for Cluster Server troubleshooting " on page 57.

Example: Ensuring the correct level of protection for volumes controlled by Storage Foundation

A volume appears to applications and the operating system as a physical device on which file systems, databases, and other managed data objects can be configured. Storage Foundation is a storage management subsystem that lets you manage the volumes in a better way.

To provide the correct level of protection to volumes is very critical in a data center environment. The method that you use to do this, ensures the increased availability of your data and proper backup mechanisms. In the Management console, Veritas InfoScale Operations Manager lets you configure various features that ensure the protection of the volumes that are controlled by Storage Foundation in your data center.

There are several methods to protect your data using Storage Foundation including snapshots and implementing RAID techniques that use mirroring. Snapshots of volumes let you create point-in-time copies of your data and restore it if the original copy becomes compromised. Mirroring provides additional resilience against hardware failures and potential read performance gains.

Consider that you have a volume by name Datavol that resides on a host that runs AIX, HP-UX, Solaris, Linux, or Windows. To ensure that the data in this volume is protected correctly, you can configure the following features for this volume:

- Create mirrors to ensure enhanced performance of storage and prevent data loss arising out of hardware failures.
- Configure snapshots for backing up data.

See "Mirroring: Protection from hardware failures" on page 61.

See "Snapshots: Backing up to prevent the loss of data" on page 62.

Mirroring: Protection from hardware failures

One of the major threats to the proper retention of data is the hardware failure that can happen to the disks in your data center. This affects the availability of the data in the disks. To prevent this situation, you can mirror the disks or the volumes in your environment. While the mirroring for the volumes is done in the context of the hosts, the mirroring of disks is done from enclosures. This document discusses the mirroring of the volumes that is performed from the context of the hosts.

In Veritas InfoScale Operations Manager, you can configure two types of mirrors for a volume from the host context: a normal mirror, and a snapshot mirror. While the normal mirror acts as a substitute for the volume, the snapshot mirror can be detached from the original volume at a later period so that you can keep this mirror as a snapshot of the original volume. Mirrors for the volumes are created using the plexes available on the unutilized disks that are part of the disk group of the original volume. The disks used for additional mirrors should not be already in use by the volume. When the data is written to the original volume, its mirror is also updated concurrently. When the original volume is not available for data retrieval due to hardware failures, the volume data is still accessible from its unaffected mirror.

To perform this task, your user group must be assigned the Admin role on the host or the Server perspective. The permission on the host may be explicitly assigned or inherited from a parent Organization.

For more information on mirrored volumes, refer to the *Storage Foundation Administrator's Guide*.

To add a datavol mirror to a volume

1. In the Management Server console, go to the **Server** perspective and expand **Manage** in left pane.
2. Expand the Organization or **Uncategorized Hosts** to locate and select the host.
3. Expand the host and select **Volumes**.
4. Right-click on the required volume and select **Mirror** > **Add**.
5. In the **Add mirror - Options** wizard panel, select **Add mirror** and specify the required information. Click **Next**.
6. If you have opted to select the disks manually, select the disks from the **Disk selection** panel. To search for disks using one or more filter criteria, click **Edit**. Click **Next**.

7 In the **Advanced options for Add mirror** panel, specify the required information. Click **Next**.

8 In the **Add mirror summary** verify the details that you have specified for adding mirrors. Click **Finish**.

9 In the **Result** panel, verify that the mirrors have been added successfully.

To add a snapshot mirror to a volume

1 In the Management Server console, go to the **Server** perspective and expand **Manage** in left pane.

2 Expand the Organization or **Uncategorized Hosts** to locate and select the host.

3 Expand the host and select **Volumes**.

4 Right-click on the required volume and select **Mirror** > **Add**.

5 In the **Add mirror - Options** wizard panel, select **Add snap-ready mirror** and specify the required information. Click **Next**.

6 If you have opted to select the disks manually, select the disks from the **Disk selection** panel. To search for disks using one or more filter criteria, click **Edit**. Click **Next**.

7 In the **Add mirror - Summary** verify the details that you have specified for adding mirrors. Click **Finish**.

8 In the **Result** panel, verify that the mirrors have been added successfully.

See "Example: Ensuring the correct level of protection for volumes controlled by Storage Foundation" on page 60.

Snapshots: Backing up to prevent the loss of data

A volume snapshot is the copy of a volume at a specific point in time. Snapshots of volumes let you back up your data and restore it back to the volume at a later time if the original volume becomes compromised

Using the Management Server console, you can configure the following types of snapshots for a volume:

- Full-sized instant snapshots: Full-sized instant snapshot lets you create a snapshot of a volume or a volume set using a compatible volume or volume set in the same disk group. The data in the full-sized snapshot volume is instantly available for the original volume when the original volume becomes compromised. To configure instant snapshot, a target volume must already be present on the same disk group. The volume that you want to keep as a snapshot must not be a snapshot of another volume. This snapshot volume must be of

the same size as of the original volume, and it must be in the Healthy state. Also, the snapshot volume must not have any file system mounted on it. To configure an instant full-sized snapshot for a volume, you must enable the Veritas FlashSnap feature for the original and the snapshot volumes.

- Space-optimized instant snapshot: Space-optimized instant snapshot lets you create snapshots with minimal use of storage resources. As for the full-sized instant snapshots, the data in the space-optimized instant snapshot is instantly available for the original volume when the original volume becomes compromised. When you configure space-optimized instant snapshot for a volume, only the changed data is copied to the snapshot volume. This type of snapshot is ideal for log volumes. The space-optimized instant snapshot configuration does not require you to create a snapshot volume in advance. When you choose to configure a space-optimized instant snapshot, Veritas InfoScale Operations Manager creates a cache object. Only the changed data on the original volume is copied to this cache object. If a cache object is already available on the same disk group out of which the original volume is created, you can use it for configuring the space-optimized instant snapshot. To configure a space-optimized instant snapshot for a volume, you must enable the Veritas FlashSnap feature for the original volume.
- Mirror break-off snapshot: Mirror break-off snapshot is created when one or more mirrored disks in a volume is detached and retained as a different volume. The volume for which you want to configure a mirror break-off snapshot must have one or more plexes in SNAP_READY state. These plexes are used to create a new volume that is a point-in-time copy of the original volume. The Add Mirror operation in Veritas InfoScale Operations Manager provides you with an option to create a plex in the SNAP_READY state.

Although these three types of snapshots serve the same purpose, the mechanism and the configurations for them are different. For Datavol, you can configure all the three types of snapshots. To configure snapshots for volumes, you must have the administrative privilege on the host where you perform this operation.

For more information on configuring snapshots for volumes, refer to the *Storage Foundation Administrator's Guide*.

To configure instant volume snapshot

1. In the Management Server console, go to the **Server** perspective and expand **Manage** in the left pane.
2. Expand the Organization or **Uncategorized Hosts** to locate the host.
3. Expand the host and expand **Volumes**.
4. Right-click on the required volume and select **Snapshot** > **Create**.

5 In the **Create Snapshot - Snapshot Level Selection Page** wizard panel, select **Volume Level Snapshots** and click **Next**.

6 In the **Create Snapshot - Snapshot Type Selection Page** wizard panel, select **Instant** and click **Advance**.

7 In the **Create Snapshot - Advance Options** wizard panel click **Configure** and review the selections. Click **Save** and click **Next**.

8 In the **Create Snapshot - Disk Selection Page** wizard panel, you can select the disk for snapshot automatically or manually.

- If you have selected automatically, click **Next**.
- If you have selected manually, then select the disk in the grid and click **Next**.

9 In the **Configure Snapshot - Summary** panel, verify the configuration information. Click **Finish**.

10 In the **Result** panel, verify that the snapshots have been configured successfully.

To configure space-optimized volume snapshot

1 In the Management Server console, go to the **Server** perspective and expand **Manage** in the left pane.

2 Expand the Organization or **Uncategorized Hosts** to locate the host.

3 Expand the host and expand **Volumes**.

4 Right-click on the required volume and select **Snapshot** > **Create**.

5 In the **Create Snapshot - Snapshot Level Selection Page** wizard panel, select **Volume Level Snapshots** and click **Next**.

6 In the **Create Snapshot - Snapshot Type Selection Page** wizard panel, select **Space optimized** and click **Advance**.

7 In the **Create Snapshot - Advance Options** wizard panel click **Configure** and enter the required information. Click **Save** and click **Next**.

8 In the **Create Snapshot - Disk Selection Page** wizard panel, you can select the disk for snapshot automatically or manually.

- If you have selected automatically, click **Next**.
- If you have selected manually, then select the disk in the grid and click **Next**.

9 In the **Configure Snapshot - Summary** panel, verify the configuration information. Click **Finish**.

10 In the **Result** panel, verify that the snapshots have been configured successfully.

To configure break-off mirror volume snapshot

1 In the Management Server console, go to the **Server** perspective and expand **Manage** in the left pane.

2 Expand the Organization or **Uncategorized Hosts** to locate the host.

3 Expand the host and expand **Volumes**.

4 Right-click on the required volume and select **Snapshot** > **Create**.

5 In the **Create Snapshot - Snapshot Level Selection Page** wizard panel, select **Volume Level Snapshots** and click **Next**.

6 In the **Create Snapshot - Snapshot Type Selection Page** wizard panel, select **Mirror break-off** and click **Advance**.

7 In the **Create Snapshot - Advance Options** wizard panel click **Configure** and enter the required information. Click **Save** and click **Next**.

8 In the **Create Snapshot - Disk Selection Page** wizard panel, you can select the disk for snapshot automatically or manually.

- If you have selected automatically, click **Next**.
- If you have selected manually, then select the disk in the grid and click **Next**.

9 In the **Configure Snapshot - Summary** panel, verify the configuration information. Click **Finish**.

10 In the **Result** panel, verify that the snapshots have been configured successfully.

See "Example: Ensuring the correct level of protection for volumes controlled by Storage Foundation" on page 60.

Example: Improving the availability and the disaster recovery readiness of a service group through fire drills

A service group is a virtual container that contains all the hardware and software resources that are required to run the managed application. Service groups allow VCS to control all the hardware and software resources of the managed application

Introducing Veritas Cluster Server Logical components of VCS 38 as a single unit. When a failover occurs, resources do not fail over individually; the entire service group fails over. If more than one service group is on a system, a group can fail over without affecting the others.

In a datacenter environment, it is important to ensure the availability and the disaster recovery readiness of a service group. Readiness status is the measure of the ability of a service group to successfully failover in its intended or configured way. Readiness takes into account the status of the service group, the system, and the cluster.

Using Veritas InfoScale Operations Manager, you can monitor the following types of readiness:

- The high availability readiness that checks for:
  - The ability of a service group to fail over to a system within the local cluster
  - The ability of a stretch service group to fail over to a system at the local site
- The disaster recovery readiness that checks for:
  - The ability of a global service group to fail over to a system in its target remote cluster at the remote site
  - The ability of a service group in a stretch cluster to fail over to a system in the remote site
  - The correctness of the replication for an application

See “High availability fire drill - Ensuring the high availability capabilities of a service group ” on page 66.

See “Disaster recovery fire drill - Ensuring the disaster recovery readiness of a service group” on page 68.

High availability fire drill - Ensuring the high availability capabilities of a service group

To configure high availability for a database or an application, you may make changes to several infrastructure and configuration settings on multiple systems. To maintain these infrastructure and the configuration settings is difficult because cluster environments can be subject to constant change. Administrators often add disks, create new disk groups and volumes, and add new cluster nodes or new NICs to upgrade and maintain the infrastructure. Updating the Cluster Server configuration to match the changing physical configuration and infrastructure is critical. HA fire drills detect discrepancies between the Cluster Server configuration and the underlying physical configuration and infrastructure on a node. Such discrepancies might prevent bringing up a service group online on a specific node.

Ultimately, the HA fire drill provides the data that is used to update the HA readiness information on the Veritas InfoScale Operations Manager console.

HA fire drill checks the configuration of all the resources of a service group on each system (in the system list for the service group) where the service group is currently offline to ensure that there are no configuration changes or incorrect configurations that may lead to unsuccessful failovers. If the HA fire drill succeeds, there is a higher probability of a successful failover on the target system; there may be some factors outside Cluster Server control that prevent this. However, if the HA fire drill fails, it is quite likely that the service groups will not come online on that system.

The following procedure describes the configuration of the HA fire drill for a service group named Oracle_SG:

To run the high availability fire drill for Oracle_SG service group

1. In the Management Server console, go to the **Availability** perspective and select **Manage** in the left pane.
2. Expand the Organization or **Uncategorized Clusters** to locate the cluster.
3. Expand the cluster and then the **Service groups** node to locate the **Oracle_SG** service group.
4. Right-click on the **Oracle_SG** service group and select **Fire Drill** > **Run HA Fire Drill**.
5. In the **Fire Drill** wizard panel, click **Next**.
6. In the **Schedule** wizard panel, do the following:
   - Select **Schedule for later**.
   - Enter a name for the fire drill schedule and specify the frequency of the fire drill run.
   - Click **Next**.
7. In the **Summary** panel, verify your selections for configuring fire drill for Oracle_SG. Click **Finish**.
8. In the **Result** panel, verify that the HA fire drill for Oracle_SG has been configured successfully.

See “Example: Improving the availability and the disaster recovery readiness of a service group through fire drills” on page 65.

Disaster recovery fire drill - Ensuring the disaster recovery readiness of a service group

The disaster recovery (DR) fire drill feature is a DR solution in Veritas InfoScale Operations Manager, which verifies the ability of a globally configured service group to fail over on a remote cluster, or that of a stretched service group to come online on the remote site in the same campus cluster. A DR fire drill is a zero-downtime test that mimics the configuration, application data, and failover behavior of critical service groups. A successful DR fire drill indicates that it is highly likely for a critical service group to fail over as intended or as configured on to a remote cluster, when it is needed.

The DR fire drill feature lets you do the following:

- Verify that replication for an application works correctly.
- Verify that a DR service group can be brought online successfully.

The objective of the DR fire drill is to bring the fire drill service group online on the remote cluster. The result of this operation verifies the ability of the similarly configured service group to fail over and come online on the remote cluster. When the DR fire drill group comes online, it uses a snapshot of the application data, which is a point-in-time copy of the replicated production data for the application. Fire drill service groups do not interact with outside clients or with other instances of resources. Therefore, they can come online safely even when the service group is online.

Note: The effectiveness of the replication of a service group varies for each replication agent pack that you use. For more information on the limitations of the agent packs, refer to the respective user documentation for the agent pack.

Using Veritas InfoScale Operations Manager, you can run the disaster recovery (DR) fire drill for the specific service groups that you choose in Veritas InfoScale Operations Manager. To perform this operation you must select the service groups that belong to the same cluster.

To perform this operation, you must have any of the following privileges:

- Operator privilege on the service groups for which you want to run the DR fire drill
- Operator privilege on the parent cluster of the service groups for which you run the DR fire drill

Note: To perform a DR fire drill on the service groups, you must create a fire drill service group on the remote cluster. The configuration of the fire drill service group is similar to the configuration of the original service group. For more information, refer to the *Cluster Server Administrator 's Guide*.

The following procedure describes the configuration of DR fire drill for a service group named Oracle_SG.

To run the disaster recovery fire drill for Oracle_SG service group

1. In the Management Server console, go to the **Availability** perspective and select **Manage** in the left pane.
2. Expand the Organization or **Uncategorized Clusters** to locate the cluster.
3. Expand the cluster and then the **Service groups** node to locate the **Oracle_SG** service group.
4. Right-click on the **Oracle_SG** service group and select **Fire Drill** > **Run DR Fire Drill**.
5. In the **Fire Drill** wizard panel, click **Next**.
6. In the **Schedule** wizard panel, do the following:
   - Select **Schedule for later**.
   - Enter a name for the fire drill schedule and specify the frequency of the fire drill run.
   - Click **Next**.
7. In the **Summary** panel, verify your selections for configuring fire drill for Oracle_SG. Click **Finish**.
8. In the **Result** panel, verify that the DR fire drill for Oracle_SG has been configured successfully.

See "Example: Improving the availability and the disaster recovery readiness of a service group through fire drills" on page 65.

Example: Identifying the performance issues of an application using Veritas InfoScale Operations Manager

As an administrator, one of your responsibilities is to ensure the optimum performance of the applications that are running in your data center. Performance issues of your critical applications can have significant negative effect on overall

business. For example, if you fail to address long I/O access times for the storage network supporting your e-commerce transaction system, the seasonal surge in transactions could cause significant slowness and have financial implications on the business.

With the Storage Insight Add-on for Veritas InfoScale Operations Manager, you can get visibility from the application down to the spindle. Veritas InfoScale Operations Manager empowers you to view the detailed storage information of an enclosure. With this added visibility, you can perform storage management operations with confidence.

For detailed information about the storage enclosures in your data center and on using Storage Insight Add-on, refer to the *Veritas InfoScale Operations Manager Add-ons User Guide*.

In Veritas InfoScale Operations Manager, you can view the performance for various objects such as host, volume, disk by using interactive graphs. You can select an object and view graphs for multiple performance parameters.

For the complete list of objects and the types of performance graphs, See "About metered resources" on page 501.

This example explains how you can use the performance graphs on the storage objects in Veritas InfoScale Operations Manager and the ability to view the application to the spindle mapping, to analyze the following problem:

Problem: The Oracle database that is installed on the LUNs, which are part of IBM XIV enclosure, is slow.

In this example, we use the following names:

Table 3-1 Example names

Object	Name
Oracle database	ora_db
Host that you use to access ora_db	lnx_host This host runs on Linux platform.
Volume, which is associated with ora_db	vxvm_vol
File system, which is associated with ora_db	vxfs_fs This File System is a VxFS file system.
Disks that are associated with the IBM XIV enclosure on which ora_db is installed	xiv_disk1 xiv_disk2

Any of the following can result in the poor performance of ora_db:

- The lnx_host that contains ora_db may have many other applications running simultaneously on it. These applications can result in more CPU load on the host or the usage of more memory. If the CPU load or the memory usage exceeds the optimum value, it can affect the performance of ora_db.
  To view the performance graphs for lnx_host associated with ora_db
- The available file system size on vxfs_fs that is associated with ora_db may not be enough for the proper functioning of ora_db. Or the delay in the I/O time on vxvm_vol that is associated with ora_db can affect the performance of ora_db.
  To view the performance graphs for vxfs_fs and vxvm_vol associated with ora_db
- The delay in the I/O time on xiv_disk1 or xiv_disk2 that is associated with ora_db can affect the performance of ora_db.
  To view the performance graphs for xiv_disk1 or xiv_disk2 associated with ora_db
- The delay in the I/O time on an array port that is associated with ora_db can affect the performance of ora_db.
  To view the performance graphs for the array port associated with ora_db

---

Note: The performance graphs for array ports are available only if the enclosure is configured using Storage Insight Add-on.

To view the performance graphs for lnx_host associated with ora_db

1. In the Home page on the Management Server console, go to **Server** perspective, and select **Manage** in the left pane.
2. Expand Organization or **Uncategorized Hosts** to locate **lnx_host**.
3. Click the **Performance** tab.

To view the performance graphs for vxfs_fs and vxvm_vol associated with ora_db

1. In the Home page on the Management Server console, go to **Server** perspective, and select **Manage** in the left pane.
2. Expand **Applications** and expand **Databases** to locate **ora_db**.
3. Click the **Volumes** tab.
4. In the volumes details list, right-click **vxvm_vol**, and select **Performance**.

To view the performance graphs for xiv_disk1 or xiv_disk2 associated with ora_db

1. In the Home page on the Management Server console, go to **Server** perspective, and select **Manage** in the left pane.
2. Expand **Applications** and expand **Databases** to locate **ora_db**.
3. Click the **Disks** tab.
4. In the disks details list, right-click **xiv_disk1** or **xiv_disk2** and select **Performance**.

To view the performance graphs for the array port associated with ora_db

1. In the Home page on the Management Server console, go to **Storage** perspective, and select **Manage** in the left pane.
2. Expand Organization or **Uncategorized Enclosures** to locate the IBM XIV enclosure.
3. Expand the enclosure and expand **Array Ports** to locate the array port associated with **ora_db**
4. Click the **Performance** tab.

Performance charts for xiv_disk1



Performance charts for xiv_disk2



From the performance graphs of xiv_disk1 and xiv_disk2, you notice an issue with the performance of xiv_disk1. The performance graph of xiv_disk1 displays that the average read latency is higher than the optimum value and the bytes read is lower than the optimum value.

From the performance graphs, you have identified the low disk I/O issue of xiv_disk1. This issue can adversely affect the performance of the ora_db. You must understand the specific problem that results in the low disk I/O on xiv_disk1. To understand the specific problem, you must analyze the details of xiv_disk1. Understanding the cause of the problem can help you make appropriate decisions on addressing the performance issue of the ora_db.

As xiv_disk1 and xiv_disk2 are part of IBM XIV enclosure, you can view the detailed information on them with the help of Storage Insight Add-on.

Note: You must configure the deep discovery for IBM XIV storage enclosures to view the details of xiv_disk1 and xiv_disk2.

For information on managing the Storage Insight Add-on, refer to the *Veritas InfoScale Operations Manager Add-ons User Guide*.

Example: Volume migration using Veritas InfoScale Operations Manager

Typical data center storage allocation tasks require you to quickly allocate storage to multiple hosts while maintaining optimum storage utilization. These tasks also involve migrating volumes. The Storage Provisioning and Enclosure Migration Add-on lets you perform the following tasks:

- Template-based storage provisioning: An efficient method to provision storage (with predefined configurations) to the required hosts in your data center. You can create new storage template from volumes or file systems, and use the template to provision storage on the managed hosts.
- Migration of volumes: You can migrate the volumes by host, enclosure or by disk groups. While migrating, you can specify various attributes of the involved storage objects. For example, LUN characteristics (thick or thin), media (SSD or HDD), RAID type, and target layout (concat, striped).

This topic provides an overview of the steps that you need to typically perform to use the volume migration functionality (by host) in Veritas InfoScale Operations Manager. To perform this task, your user group must be assigned the Admin role on the host or the Server perspective. The permission on the host may be explicitly assigned or inherited from a parent Organization.

You also need to verify the following:

- If volumes have been set up for creating snapshots, unprepare those volumes.
- Ensure that disks are provisioned on the new enclosures and are added to the disk group.
- Ensure that the enclosures are added to Veritas InfoScale Operations Manager.

To migrate volumes by host

1. In the Management Server console, go to the **Server** perspective, expand **Solutions** in the left pane, and click **Storage Migration**.
2. In the **Actions** pane, click **Migrate Volumes By Host**.
3. In the **Migrate Volumes By Host** wizard, specify a name and description for the migration task so that you can track the task status. Click **Next**.

4 Select the host and then select the source enclosure. Click **Next**.

5 Select one or more enclosures as the targets to which you want to migrate the volumes.

Under **Select LUN Characteristics**, specify the LUN requirements. It includes selecting enclosure vendor, LUN type (thick, thin, or any), media type (SSD or HDD), LUN classifications, and replication status. Click **Next.**



6 Select the volumes to migrate. Verify the source and target layout. If you want to change the target layout, double-click the volume and select the new layout on the **Change Layout** window.



You can also choose from the following options:

- Keep the same number of paths to the volume (the default) or set a minimum number of paths.
- Optionally, select the check box to delete the original volumes after successfully migrating the volumes to the target enclosure.

Click **Next.**

7 Choose one of the following:

- Schedule the time of the migration.
- To run immediately, select **Run now**.

Warning: You may want to view the Impact Analysis report that is generated by the wizard before you run the migration.

Click **Next**.

8 The wizard generates the Impact Analysis report. It provides the details of the objects that are associated with the volumes being migrated.

View the report and click **Previous** if you want to go back and change anything. Otherwise, click **Finish** and close the confirmation window.

Scheduled migrations are listed on the **Migration Status (By Schedule)** tab. You can check on the status of a specific volume on the **Migration Status (By Volume)** tab.

Examples: Identifying and reducing storage waste using Veritas InfoScale Operations Manager

Primary storage is under increasing pressures in today's IT environment. Data is increasing in size with more complex applications and is stored longer to align to more stringent retention requirements. This data can include anything from spreadsheets from the chief financial officer (CFO), PowerPoints from the sales managers, and archive log files from Oracle databases.

This section introduces you to the following Storage Foundation High Availability features for reducing storage waste and gives examples of how you can use Veritas InfoScale Operations Manager for implementing these features.

Table 3-2 Features for reducing storage waste

Feature	Possible uses
Thin reclamation	Reclaim unused storage space on thin LUNs.
Compression	Reclaim storage space by file system compression. The following are good candidates for compression: ■ Database archive logs ■ Unstructured data
Deduplication	Reclaim storage space by file system deduplication. The following are good candidates for deduplication: ■ Virtual machine boot image files (vmdk files) ■ User home directories ■ File systems with multiple copies of files

Each feature and the required steps using Veritas InfoScale Operations Manager are described under a separate use case.

See "Reclaiming thin storage - example" on page 75.

See "Compressing files - example" on page 79.

See "Deduplicating file systems - example" on page 83.

Reclaiming thin storage - example

With Veritas InfoScale Operations Manager, you can use the Management Server console to gain visibility into and manage thin-reclamation capable devices.

Thin provisioning is a storage array feature that optimizes storage use by allocating and reclaiming the storage on demand. With thin provisioning, the array allocates storage to applications only when the storage is needed, from a pool of free storage. Thin provisioning solves the problem of under-utilization of available array capacity. Administrators do not have to estimate how much storage an application requires. Instead, thin provisioning lets administrators provision large thin or thin reclaim capable LUNs to a host. When the application writes data, the physical storage is allocated from the free pool on the array to the thin-provisioned LUNs.

The two types of thin provisioned LUNs are thin-capable or thin-reclaim capable. Both types of LUNs provide the capability to allocate storage as needed from the free pool. For example, storage is allocated when a file system creates or changes a file. However, this storage is not released to the free pool when files get deleted. Therefore, thin-provisioned LUNs can become 'thick' over time, as the file system starts to include unused free space where the data was deleted. Thin-reclaim capable LUNs address this problem with the ability to release the once-used storage to the pool of free storage. This operation is called thin storage reclamation.

The thin-reclaim capable LUNs do not perform the reclamation automatically. The administrator can initiate a reclamation manually, or with a scheduled reclamation operation.

The Management Server console provides visibility and management for thin reclamation in the context of the following objects that are discovered by Veritas InfoScale Operations Manager:

- Server perspective: Lets you select file systems or disks for a host and perform thin reclamation. The space no longer used is reclaimed from the associated LUNs and made available.
- Storage perspective: Lets you select thin pools from a storage array and either schedule thin reclamation or perform it manually. A thin pool is a collection of devices in the array that are dedicated for use by thin LUNs.

See the requirements for using Veritas InfoScale Operations Manager for thin reclamation.

Requirements for reclaiming thin storage

Example of reclaiming thin storage

Requirements for reclaiming thin storage

To reclaim thin storage requires managed hosts with a supported Storage Foundation version. Thin reclamation is supported on the following Storage Foundation versions:

- UNIX/Linux: 5.0 MP3, or later
- Windows: 5.1 SP1, or later

In addition, the storage must meet the following requirements:

- The storage must be visible to Storage Foundation as thin reclaimable.
- For UNIX: The LUNs must be a part of a Veritas Volume Manager volume which has a mounted VxFS file system.
- For Windows: The LUNs must be a part of a Storage Foundation for Windows dynamic volume which has a mounted NTFS file system.

Reclamation of thin pools requires in addition that the array support the thin-reclamation functionality and that Storage Insight Add-on is configured to discover the enclosure.

Example of reclaiming thin storage

In the following example, a storage administrator reclaims storage from thin pools.

The administrator performs the following procedures to reclaim thin storage.

Run thin reclamation on thin pools

Note: Although not used in this example, you can also select file systems or disks for a selected host and perform thin reclamation.

See "Performing thin reclamation on file systems or disks" on page 330.

Identify thin pools for reclamation

You can run the Top Thin Pools for Reclamation report to show the 10 thin pools with the most reclaimable space.

This report has the following requirements:

- Storage Foundation version 6.0 or later
- Veritas File System (VxFS) disk layout version 9 or later
- Veritas InfoScale Operations Manager managed host (VRTSsfmh) version 5 or later

You can view this information related to the enclosures for which your user group has at least guest role explicitly assigned or inherited from a parent Organization. You can also view the information if your user group has a role assigned on the Storage perspective.

To run the report to identify thin pools for reclamation

1. In the Management Server console, go to the **Storage** perspective and expand **Reports** in the left pane.
2. With the **Storage Utilization** category displayed, click the **Top Thin Pools for Reclamation** report.
3. In the **Select Scope to run report** wizard panel, select the scope of the report and click **Run**.

Run thin reclamation on thin pools

You can select one or more thin pools from a selected storage array and either schedule thin reclamation or perform it manually. Thin pools are available if the array supports them and if Storage Insight Add-on is configured for the selected enclosure. Make sure LUNs from these thin pools are consumed by hosts running Storage Foundation.

To perform this task, your user group must be assigned the Admin role on the enclosure or the Storage perspective. The permission on the enclosure may be explicitly assigned or inherited from a parent Organization.

To schedule thin reclamation on thin pools

1. In the Management Server console, go to the **Storage** perspective and expand **Manage** in the left pane.
2. Click **Data Center** and on the **Enclosures** tab, type all or part of the enclosure name to filter the list of enclosures. Double-click the enclosure name in the table. The enclosure is highlighted and expanded in the tree.
3. In the tree, under the selected enclosure, click **Thin Pools**.
4. In the table, select one or more thin pools, right-click, and select **Schedule Reclamation**.

   Alternatively, to run reclamation manually without scheduling it, you can select **Run Reclamation**.
5. Choose from the options to schedule when thin reclamation runs for the selected thin pools.

Frequency	Select **Once**, **Daily**, **Weekly**, or **Monthly**.
When	The options in the **When** area change depending on the **Frequency** selection. For **Weekly**, you can select **Every weekday** to schedule Monday through Friday or select specific days of the week. For **Monthly**, you can schedule the reclamation to re-occur on a specific day of every month.

See "Examples: Identifying and reducing storage waste using Veritas InfoScale Operations Manager" on page 75.

Compressing files - example

Storage Foundation 6.0 or later enables customers to use host-based compression to optimize existing primary storage. Enabling compression at the file system layer results in storage savings and avoids complex and expensive appliances typically associated with primary compression.

Compression is performed without needing any application changes and with minimal overhead. Compression does not modify the file metadata, nor are inode numbers or file extensions changed. Compression is executed out-of-band, after the write. Once compression is enabled, directories and files begin to have a mix of compressed and uncompressed data blocks. This is managed automatically by the file system, and uncompressed data is compressed during the next sweep.

See "About file compression in Veritas InfoScale Operations Manager" on page 333.

Using the Veritas InfoScale Operations Manager Management Server console you can enable file system compression and view the space savings.

Requirements for compression

Use cases for compression

Example of using compression in Veritas InfoScale Operations Manager

Requirements for compression

Compression has the following requirements:

- Storage Foundation version 6.0 or later
- Veritas File System (VxFS) disk layout version 9 or later
- Veritas InfoScale Operations Manager managed host (VRTSsfmh) version 5 or later

Use cases for compression

The following are some use cases for file compression:

- Unstructured data
  - Studies indicate 90% of user-created data is never accessed after creation.
  - Increased regulatory restrictions leads to longer-term storage.
- Oracle database archive logs
  - Oracle best practices recommend archive logs for database recovery.

- As databases age and are used strictly for reads, log files go stale and unchanged.

How well a system handles the compression and uncompression of files is a key metric in deciding which data types can be compressed and when. Compression is CPU heavy and the CPU load should be considered carefully. Reading from compressed files can also result in performance degradation due to the increased I/O. The total space savings and time to compress or uncompress varies depending on server type, server load, file type, and compression settings. The following table shows some examples of possible savings.

Table 3-3 Compression savings

Data Type	Platform	Original Size	Savings	CPU Usage
Unstructured (80,000 files)	Solaris SPARC 10	5 GB	70%	1 CPU: 6 % 4 CPU: 20%
Oracle archive log	Linux RHES	18 GB	60%	1 CPU: 6 % 4 CPU: 20%

Example of using compression in Veritas InfoScale Operations Manager

In this example, a server administrator has a large set of seldom-used unstructured data in user home directories and will use compression to save storage space.

For examples of using compression with Storage Foundation from the command line, see the *Storage Foundation Administrator's Guide*.

The following object names are used in this example.

Table 3-4 Example names

Object	Name
Host	lnx_host
Volume	vxvm_users
Mount point (file system)	/home

The administrator performs the following procedures to reduce storage waste with file compression.

Locate the volume and mount point for the directories to be compressed

Add a compression schedule

Select the directories for compression

Verify the space saved by compression

To perform compression operations, you must have administrative privileges in Veritas InfoScale Operations Manager for the host on which you invoke the compression operation.

Locate the volume and mount point for the directories to be compressed

To locate the directories in Veritas InfoScale Operations Manager, you need to be able to identify the volume and mount point (file system) by name. The following procedure explains how to use the Management Server console to locate a volume and mount point for a host.

You can also use **Search** on the console menu bar to search for a volume and mount point.

To locate the volume and mount point

1. In the Management Server console, go to the **Server** perspective and expand **Manage** in the left pane.
2. Click **Data Center** and on the **Hosts** tab, type all or part of the host name to filter the list of hosts. For example, type **lnx_host**. Double-click the host name in the table. The host is highlighted and expanded in the tree.
3. In the tree, under the selected host, click **Volumes**.
4. On the **Volumes** tab, filter the list of volumes to locate the volume. For example, type **vxvm_users**.

Add a compression schedule

When selecting a schedule time and duration, keep in mind that compression is a CPU intensive procedure.

To add a compression schedule

1. In the Management Server console, locate the volume and mount point.
2. Right-click the volume on which the file system is mounted and click **Properties**.
3. Click the **Compression** tab.

4 Click **Add Schedule**.

5 In the **Compression Schedule** window, specify the compression options and click **OK**.

Frequency	Select **Daily**, **Weekly**, or **Monthly**.
When	The options in the **When** area change depending on the **Frequency** selection. For **Weekly**, you can select weekday to schedule Monday through Friday or select specific days of the week. For **Monthly**, you can schedule the compression to re-occur on a specific day of every month.
Compression Duration	Specify how long the compression process runs. If all directories are not compressed during the specified duration, at the next scheduled compression run, the process continues with the remaining directories. For example, say that a duration of one hour is set and 10 directories are enabled for compression. After one hour, nine directories are compressed. The compression process stops. At the next scheduled run, the compression process continues with the tenth directory. The process then starts over with the first directory and compresses any new files that were added since the last run. Default: four hours
Number of CPUs to use for compression	Specify how many CPUs to use for the scheduled compression run. Default: 50 percent of the CPUs available for the host, up to 4 CPUs.

Select the directories for compression

You select which directories to compress for the selected file system.

To select the directories for compression

1 In the Management Server console, locate the volume and mount point.

2 Right-click the volume on which the file system is mounted and click **Properties**.

3 Click the **Compression** tab.

4 Select or deselect directories to enable or disable for compression. For example, select **/home** to enable all the user subdirectories under `/home` for compression.

5 Click **Apply**.

The directories are compressed at the next scheduled compression run. If you want to run compression immediately, click **Compress Now**.

Verify the space saved by compression

You can verify the results of file compression by viewing the amount shown as **Space Saved** on the **Compression** tab once the compression run is complete.

You can also run a report to view top savings for file compression.

To view a report on savings by file compression

1 In the Management Server console, go to the **Server** perspective and expand **Reports** in the left pane.

2 With the **Storage Utilization** category displayed, click the **Savings by File System Compression** report.

3 In the **Select Scope to run report** wizard panel, select the scope of the report and click **Run**.

See "Examples: Identifying and reducing storage waste using Veritas InfoScale Operations Manager" on page 75.

Deduplicating file systems - example

Storage Foundation 6.0 and later enables customers to use file system deduplication to optimize existing primary storage. Enabling deduplication at the file system layer results in storage savings and avoids complex and expensive appliances typically associated with file deduplication.

Deduplication is performed without needing any application changes and with minimal overhead. Deduplication does not change the file extension, allowing users and applications to use files normally, without performance effect.

The VxFS deduplication feature works as follows. It eliminates duplicate blocks used by your data by comparing blocks across the file system. When the deduplication feature finds a duplicate block, it removes the space used and instead creates a pointer to the common block. If the duplicate file is changed, thus making the files no longer share the same block, then that changed block is saved to disk instead of the pointer.

For more information on how deduplication works, see the following topic:

See "About file system deduplication" on page 338.

Using the Veritas InfoScale Operations Manager Management Server console you can enable file system deduplication and view the space savings.

Requirements for deduplication

Use cases for deduplication

Example of using deduplication

Requirements for deduplication

Deduplication has the following requirements:

- Storage Foundation version 6.0 or later.
- Veritas File System (VxFS) disk layout version 9 or later.
- Veritas InfoScale Operations Manager managed host (VRTSsfmh) version 6.0 or later.

Use cases for deduplication

The following are good candidates for deduplication:

- Virtual machine boot image files (vmdk files)
- User home directories
- File systems with multiple copies of files

The following might not be the best candidates for deduplication, as they have little or no duplicate data:

- Databases
- Media files, such as JPEG, MP3, and MOV

Example of using deduplication

In the following example, a storage administrator has a volume and file system that are set up to store multiple copies of documents in a source control system.

The following object names are used in this example.

Table 3-5 Example names

Object	Name
Host	lnx_host
Volume	vxvm_source_control
Mount point (file system)	/user_source

The administrator performs the following procedures to reduce storage waste with file deduplication.

Locate the volume and mount point for the directories to undergo deduplication

Implement deduplication

Verify the results of deduplication

To perform deduplication operations, you must have administrative privileges in Veritas InfoScale Operations Manager for the host on which you invoke the deduplication operation.

Locate the volume and mount point for the directories to undergo deduplication

To locate the directories in Veritas InfoScale Operations Manager, you need to be able to identify the volume and mount point (file system) by name. The following procedure explains how to use the Management Server console to locate a volume and mount point for a host.

You can also use **Search** on the console menu bar to search for a volume and mount point.

To locate the volume and mount point

1. In the Management Server console, go to the **Server** perspective and expand **Manage** in the left pane.
2. Click **Data Center** and on the **Hosts** tab, type all or part of the host name to filter the list of hosts. For example, type **lnx_host**. Double-click the host name in the table. The host is highlighted and expanded in the tree.
3. In the tree, under the selected host, click **Volumes**.
4. On the **Volumes** tab, filter the list of volumes to locate the volume. For example, type **vxvm_source_control**.

Implement deduplication

Veritas InfoScale Operations Manager lets you implement deduplication for a selected file system. You configure the deduplication database and optionally set up a schedule.

To implement deduplication for a file system

1. In the Management Server console, locate the volume and mount point. For example, locate **vxvm_source control (/user_source)**.
2. Right-click the volume on which the file system is mounted and click **Properties**.
3. Click the **Deduplication** tab.

4 Click **Configure**.

5 In the **Configure Deduplication** window, you can customize the following options:

Enabled	If you clear the check box, the deduplication operation is disabled. If you want to enable it later, return to this window.
Data Usage	Lets you optimize the database size according to the type of data and the amount of space available for the database. The smaller the chunk size that is selected for data, the more space is required for the database. Once configuration is complete, this parameter cannot be changed except by unconfiguring the database and reconfiguring it. For most data, Veritas recommends the default, **Other (16k)**.

6 To set up a schedule for deduplication, select from the following:

Commit on run number	The deduplication process scans and fingerprints the data before eliminating duplicates. You can schedule the deduplication process to eliminate the duplicates each time it runs (the default value of 1) or every specified number of times. During the times that deduplication does not occur, the deduplication run only updates the fingerprints in the database.
Weekday Schedule	You can select one day of the week or schedule a run every day. Veritas recommends that you schedule deduplication when the system activity is low so as not to interfere with the regular system workload.
Hours	Schedule the hour to begin a deduplication run.

7 Click **Finish**. The deduplication configuration sets up the deduplication database. When a message shows the configuration is complete, click **Close**.

8 If you want to run deduplication now, rather than wait for a scheduled time, click **Scan Now**. Click **Yes** to confirm that you want to begin the deduplication. Once it is begun, you can close the window. The operation runs in the background.

Verify the results of deduplication

You can verify the results of deduplication for a specific file system on the **Deduplication** tab.

The **Space Saved** field shows the most recent savings. You can also view savings over time under **Space Saved History**. To change the time period, select from the **Duration** drop-down list and click **Apply**.

You can also run a report to view top savings for file system deduplication.

To run a report on savings by file system deduplication

1 In the Management Server console, go to the **Server** perspective and expand **Reports** in the left pane.

2 With the **Storage Utilization** category displayed, click the **Savings by File System Deduplication** report.

3 In the **Select Scope to run report** wizard panel, select the scope of the report and click **Run**.

See "Examples: Identifying and reducing storage waste using Veritas InfoScale Operations Manager" on page 75.

Section 2

Managing Veritas InfoScale Operations Manager

- Chapter 4. Managing user access
- Chapter 5. Setting up fault monitoring
- Chapter 6. Using reports

Chapter 4

Managing user access

This chapter includes the following topics:

- Creating an Organization
- Modifying the name of an Organization
- Deleting an Organization
- Moving an object to an Organization in a perspective
- Assigning permissions to user groups on an Organization within a perspective
- Modifying permissions assigned to user groups on an Organization within a perspective
- Deleting permissions assigned to user groups on an Organization within a perspective
- Modifying permissions assigned to user groups on an object within a perspective
- Verifying a user group in the domain
- Viewing permissions information
- Viewing the permissions assigned on a perspective, an Organization, or on an object

Creating an Organization

You can use the Management Server console to group objects within a perspective to form an Organization. An Organization defined in one perspective is not available in another perspective. Table 4-1 lists the objects in each perspective which can be grouped to form an Organization.

Table 4-1 Objects for creating Organizations

Perspective	Object
Server	Hosts
Availability	Clusters
Storage	Enclosures
Virtualization	Virtualization servers

You can create multiple Organizations in the data center. You can also create multiple nested Organizations.

If you create an Organization based on rule, it is applicable for all the new objects added to the domain after the rule is created. For example, if you create an Organization for all Windows hosts, then when a new Windows host is added to the domain, it is included in the Organization.

To perform this task, your user group must be assigned the Admin role on the perspective in which you want to create the Organization.

To create an Organization

1. In the Management Server console, go to the perspective and select **Manage** in the left pane.
2. Right-click **Data Center** and select **Create Organization**. To create a nested Organization, right-click the parent Organization.
3. In the **Create Organization** wizard panel, enter a name for the Organization, and select a method to add objects to the Organization.

   See “Create Organization panel options” on page 91.
4. If you select **Empty Organization**, click **Finish**.
5. If you select **Assign Hosts to Organization**, select the objects, and click **Finish**.

   See “Create Organization - Select an object panel options” on page 91.
6. If you select **Assign Hosts to Organization Based on Rule**, create a rule to select the objects, and click **Finish**.

   See “Create Organization - Based on a rule panel options” on page 92.

See “Modifying the name of an Organization” on page 94.

See “Moving an object to an Organization in a perspective” on page 95.

See “Assigning permissions to user groups on an Organization within a perspective” on page 96.

See “Deleting permissions assigned to user groups on an Organization within a perspective” on page 97.

Create Organization panel options

Use this wizard panel to select a method to add objects to an Organization.

Select **Empty Organization** if you want to create an empty Organization in the perspective.

Or you can select one of the following methods depending on the perspective in which you want to create the Organization.

Table 4-2 Create Organization panel options for selecting a method to add objects to the Organization

Perspective	Selection options
Server	Select either **Assign Hosts to Organization** or **Assign Hosts to Organization Based on Rule**.
Availability	Select either **Assign Clusters to Organization** or **Assign Clusters to Organization Based on Rule**.
Storage	Select either **Assign enclosures to Organization** or **Assign enclosures to Organization Based on Rule**.
Virtualization	Select either **Assign Virtualization Servers to Organization** or **Assign Virtualization Servers to Organization Based on Rule**.

See “Creating an Organization” on page 89.

Create Organization - Select an object panel options

Use this wizard panel to select the objects to be added to the Organization.

Table 4-3 Create Organization panel options for selecting objects in a perspective

Perspective	Action
Server	Select the hosts.
Availability	Select the clusters.

Table 4-3 Create Organization panel options for selecting objects in a perspective *(continued)*

Perspective	Action
Storage	Select the enclosures.
Virtualization	Select the virtualization servers.

See "Creating an Organization" on page 89.

Create Organization - Based on a rule panel options

Use this wizard panel to create a rule for selecting objects to be added to the Organization.

Table 4-4 Create Organization panel options for creating a rule to select objects

Attribute	Displays the list attributes in a perspective.
Condition	Select a condition, for example, **Starts With**.
Value	Enter the value. Value strings are not case-sensitive.
Add	Click to add another search criteria.
Remove	Click to remove the search criteria.
Operator	Choose whether to use an **AND** or **OR** operator for the new search criteria.

Table 4-5 Create Organization panel options for selecting attributes in a perspective

Perspective	Attributes
Server	■ Architecture ■ Cluster ■ IP Address ■ MH Version ■ Name ■ OS Version ■ Platform ■ SF Version ■ State ■ VCS Version
Availability	■ Condition ■ Host Count ■ Name ■ Platform ■ State ■ Sub Type ■ Version
Storage	■ Condition ■ IP Address ■ Name ■ Product ■ Serial ■ Type ■ Vendor
Virtualization	■ Cluster ■ Name ■ Server Type ■ SF Version ■ State

Note: In addition to the predefined attributes, the extended attributes defined for an object are displayed in the **Attributes** drop-down list.

See "Creating an Organization" on page 89.

Modifying the name of an Organization

You can use the Management Server console to modify the name of an existing Organization in a perspective. You can also remove the objects belonging to the Organization.

To perform this task, your user group must be assigned the Admin role on the perspective.

To modify the name of an Organization

1. In the Management Server console, go to the perspective and select **Manage** in the left pane.
2. Right-click the Organization and select **Modify**.
3. In the **Modify Organization** panel, enter a new name, remove the objects, and click **Finish**.

   See "Modify Organization panel options" on page 94.
4. In the **Modify Organization - Result** panel click **Close**.

See "Creating an Organization" on page 89.

See "Deleting an Organization" on page 95.

Modify Organization panel options

Use this wizard panel to modify the name of the Organization.

You can also remove the objects from the Organization.

Table 4-6 Modify Organization panel options for modifying the objects in an Organization

Perspective	Action
Server	Remove the hosts.
Availability	Remove the clusters.
Storage	Remove the enclosures.
Virtualization	Remove the virtualization servers.

See "Modifying the name of an Organization" on page 94.

Deleting an Organization

You can use the Management Server console to delete an existing Organization in a perspective. You can delete an Organization only if it is empty.

To perform this task, your user group must be assigned the Admin role on the perspective.

To delete an Organization

1. In the Management Server console, go to the perspective and select **Manage** in the left pane.
2. Right-click the Organization and select **Delete**.
3. In the **Delete Organization** panel, click **Yes**.
4. In the **Delete Organization - Result** panel, click **Close**.

See "Creating an Organization" on page 89.

See "Modifying the name of an Organization" on page 94.

Moving an object to an Organization in a perspective

You can use the Management Server console to move objects from the **Uncategorized** folder to an Organization in a perspective. Table 4-7 lists the objects in each perspective which can be grouped into an Organization.

Table 4-7 Perspectives and objects

Perspective	Object
Server	Hosts
Availability	Clusters
Storage	Enclosures
Virtualization	Virtualization servers

To perform this task, your user group must be assigned the Admin role on the perspective.

To move an object to an Organization in a perspective

1. In the Management Server console, go to the perspective and select **Manage** in the left pane.
2. Expand **Uncategorized** to select an object.
3. Right-click the object and select **Move To**.
4. In the **Organize** panel select the Organization to which you want to add the object. Click **OK**.
5. In the **Organize - Result** panel, click **Close**.

See "Creating an Organization" on page 89.

See "Modifying the name of an Organization" on page 94.

Assigning permissions to user groups on an Organization within a perspective

Veritas InfoScale Operations Manager makes use the existing user groups which are present in Lightweight Directory Access Protocol (LDAP), or Active Directory (AD), or the authentication mechanism in the native operating system of Windows and UNIX/Linux. Permissions such as Admin or Guest can be assigned to the user groups on an Organization within a perspective. Operator role can be assigned only in the **Availability** perspective. The user groups having the Operator role can perform operations such as taking a service group online or offline, freezing or unfreezing a service group, or running the high availability and disaster recovery fire drill.

User group name is case-sensitive.

To perform this task, your user group must be assigned the Admin role on the perspective.

To assign permissions to user groups on an Organization within a perspective

1. In the Home page on the Management Server console, go to the perspective, and select **Manage** in the left pane.
2. Right-click the Organization and select **Properties**.
3. In the **Permissions** tab, under **Add Permission**, click **Select user group**.
4. In the **Select user group** panel, select the domain, and enter the name of the user group.
5. Click **Validate user group** and click **OK**.
6. Under **Add Permission**, select a role from the drop-down list. Click **Add**.

7 In the **Success** panel click **OK**.

See “Modifying permissions assigned to user groups on an Organization within a perspective” on page 97.

See “Deleting permissions assigned to user groups on an Organization within a perspective” on page 97.

Modifying permissions assigned to user groups on an Organization within a perspective

Veritas InfoScale Operations Manager makes use the existing user groups which are present in the Active Directory or in the native operating system such as Windows or UNIX/Linux.

To perform this task, your user group must be assigned the Admin role on the perspective.

To modify permissions assigned to user groups on an Organization within a perspective

1 In the Home page on the Management Server console, go to the perspective and select **Manage** in the left pane.

2 Right-click the Organization and select **Properties**.

3 In the **Permissions** tab right-click the user group and select **Modify Role**.

4 In the **Modify Role** panel, select a role from the drop-down list, and click **OK**.

5 In the **Modify Role** panel click **Close**.

See “Assigning permissions to user groups on an Organization within a perspective” on page 96.

See “Deleting permissions assigned to user groups on an Organization within a perspective” on page 97.

Deleting permissions assigned to user groups on an Organization within a perspective

Using the Management Server console, you can delete the permissions assigned to user groups on an Organization within a perspective.

To perform this task, your user group must be assigned the Admin role on the perspective.

To delete permissions assigned to user groups on the Organization within a perspective

1. In the Home page on the Management Server console, go to the perspective and select **Manage** in the left pane.
2. Right-click the Organization and select **Properties**.
3. In the **Permissions** tab, right-click the user group and select **Delete**.
4. In the **Delete** panel click **OK**.
5. In the **Delete** panel click **Close**.

See "Assigning permissions to user groups on an Organization within a perspective" on page 96.

See "Modifying permissions assigned to user groups on an Organization within a perspective" on page 97.

Modifying permissions assigned to user groups on an object within a perspective

Objects within an Organization, inherit the permissions that are assigned to the Organization. To exclusively assign permissions to user groups on an object, you need to modify the inherited permissions.

To perform this task, your user group must be assigned the Admin role on the Organization or the perspective.

To modify permissions assigned to user groups on an object within a perspective

1. In the Home page on the Management Server console, go to the perspective and expand **Manage** in the left pane.
2. Expand the Organization to locate the object.
3. Right-click the object and select **Properties**.
4. In the **Permissions** tab right-click the user group and select **Modify Role**.
5. In the **Modify Role** panel, select a role from the drop-down list, and click **OK**.
6. In the **Modify Role** panel click **Close**.

See "Assigning permissions to user groups on an Organization within a perspective" on page 96.

Verifying a user group in the domain

Using the Management Server console, you can verify the user groups in the domain.

The user group names are case-sensitive.

To perform this task, your user group must be assigned the Admin role on the perspective.

To verify a user group in the domain

1. In the Home page on the Management Server console, do one of the following:
   - Click **Settings** and then click **Security**.
   - Click on a perspective and select **Manage** in the left pane.
     Right-click the **Data Center**, the Organization, or an object within an Organization, and select **Properties**.
2. Click the **Permissions** tab.
3. Right-click a user group and select **Verify**.
4. In the **Verify user group** panel, click **Close**.

See "Assigning permissions to user groups on an Organization within a perspective" on page 96.

Viewing permissions information

You can use the Management Server console to view the permissions information. You can view the following details:

Name	Displays the name of the perspective, Organization, and the object on which the permissions are assigned.
Type	Displays the type of the object for example host, cluster, service group, enclosure, or virtualization server. If permissions are assigned on the perspective, the type column displays "Perspective" and in case of an Organization, the type column displays "Organization".
Perspective	Displays the perspective name to which the object or the Organization belongs to. This column is displayed only in the **Management Server** perspective.
User group	Displays the name of the user group.

Role	Displays the type of role that is assigned to the user group.

You can view information for only those objects within the perspective on which you have at least Guest role explicitly assigned or inherited from the parent Organization. You can also view the information if your user group has a role assigned on the perspective.

If multiple user groups have permissions on the same Organization, you can filter the list using **Show only my user groups**. You can only delete the permissions that are assigned to your user group.

You can view this information, if your user group has at least Guest role assigned on the perspective.

To view the permissions information

- In the Home page on the Management Server console, do one of the following:
  - Go to a perspective and select **Manage** in the left pane. Right-click **Data Center** and select **Show Permissions**.
  - Click **Settings** and click **Show Permissions**.

See "Assigning permissions to user groups on an Organization within a perspective" on page 96.

See "Modifying permissions assigned to user groups on an object within a perspective" on page 98.

See "Viewing the permissions assigned on a perspective, an Organization, or on an object" on page 100.

Viewing the permissions assigned on a perspective, an Organization, or on an object

You can use the Management Server console to view the permissions assigned on a perspective, an Organization, or on an object within an Organization.

- You can view the following details:

User group	Displays the name of the user group.
Role	Displays the type of role that is assigned to the user group.

- If multiple user groups have permissions on the same perspective, Organization, or the object, you can filter the list using **Show only my user groups**. You can only delete the permissions that are assigned to your user group.

You can perform the following tasks in this view:

- Assign permissions.
- Modify permissions.
- Delete permissions.
- Verify a user group.

Note: An object within an Organization inherits the permissions assigned on the parent Organization. On an object, you can only modify the permissions.

You can view this information, if your user group has at least Guest role assigned on the perspectives or the Organization.

To view the permissions assigned on a perspective, an Organization, or on an object

- In the Home page on the Management Server console, do one of the following:
  - Click **Settings** and then click **Security**.
    Click the **Permissions** tab and then select a perspective from the drop-down list to view the permissions information for the perspective.
  - Click on a perspective and select **Manage** in the left pane.
    Right-click the **Data Center**, an Organization, or an object within an Organization, select **Properties**, and then click the **Permissions** tab.

See "Assigning permissions to user groups on an Organization within a perspective" on page 96.

See "Modifying permissions assigned to user groups on an object within a perspective" on page 98.

See "Verifying a user group in the domain" on page 99.

See "Deleting permissions assigned to user groups on an Organization within a perspective" on page 97.

See "Viewing permissions information" on page 99.

Chapter 5

Setting up fault monitoring

This chapter includes the following topics:

- About alerts and rules
- Creating rules in a perspective
- Editing rules in a perspective
- Deleting rules in a perspective
- Enabling rules in a perspective
- Disabling rules in a perspective
- About faults and risks
- Suppressing faults in a perspective
- Restoring a suppressed fault in a perspective

About alerts and rules

Data center administrators need to manage the condition of the resources in the data center. Administrators typically define the custom rules that specify what conditions generate an alert, what actions should occur if an alert is detected, and which actions generate which type of alert severity. Using the Management Server console, you can create and maintain rules pertaining to alerts.

You can monitor the faulty status and performance information of your data center by reviewing the alert log on the Management Server console.

You can view the following information on alerts in the data center:

- Information about the alert.
- The source of the alert.

- The time when the alert occurred.

The alert severity levels are:

- Critical
- Warning
- Information

You can create alert rules to receive warnings about events and conditions, such as stopped replication or storage capacity, enabled or disabled I/O paths, faulted clusters and so on.

Using the Management Server console, you can specify to initiate one of the following actions when an alert condition is met:

- Send an email message. For some alert conditions, operators may want to send emails notifying key personnel about the condition. You can specify one or more email addresses to which the alert notification is sent.

  Note: You must provide the details for the SMTP settings before setting the email notification for an alert.

- Send an SNMP trap notification. Some objects are not polled. When events take place, these objects send traps or unsolicited asynchronous SNMP messages to the Server. Some of the rules that Veritas InfoScale Operations Manager uses to monitor objects in the environment rely on SNMP trap-based messages.

  Note: You must configure SNMP trap settings for receiving alert notifications.

- Run a custom script. You can upload a custom script that runs when the alert conditions that are specified by the rule occur.

See “Creating rules in a perspective” on page 103.

Creating rules in a perspective

In the Management Server console, you can create rules to trigger various actions based on alert conditions. Aside from creating rules from the rules option, you can also create rules from selected faults and alerts.

You can choose to create a rule on the **Data Center** or on an Organization.

To create a rule on the **Data Center**, your user group must be assigned the Admin role on the perspective.

To create a rule on an Organization, your user group must be assigned the Admin role on the Organization.

Note: Rules that are created in this manner are applied to one or more chosen Organizations. This can be the data center (the global organization) or one or more Organizations within the data center.

The fault topics listed in the wizard for rule creation are those which are relevant to the perspective. In the Management Server perspective, the fault topics listed for rule creation include all host fault topics and array and switch fault topics. A rule can trigger an email, an SNMP trap and/or a custom script (custom script only available in the Management Server perspective).

For more information on creating rules in the Management Server perspective, see *Veritas InfoScale Operations Manager Installation and Configuration Guide*.

To create a rule in a perspective

1. In the Home page on the Management Server console, select a perspective, and select **Manage** in the left pane.
2. Do one of the following:
   - Right-click **Data Center**, and select **Create Rule**.
   - Right-click an Organization, and select **Create Rule**.
   - Click the **Rules** tab, right-click in the empty tab or an existing rule.
3. In the **Create Rule** wizard panel, do one of the following:
   - Select **This rule will be triggered for all faults of type:**, click **Next** and skip to 5.
   - Select **Enter the fault topics that will trigger the actions for this rule:**, enter the fault definitions separated by a comma (,) or a semicolon (;), and click **Next**. Skip to 5.
   - Select **Choose from a list of fault topics**, click **Next**, and skip to 4.

   See “Create Rule - Select the type of fault conditions to trigger this rule panel options ” on page 105.
4. In the **Create Rule - Select one or more fault topics which will trigger this rule** wizard panel, select the fault topics, and click **Next**.

   See “Create Rule - Select one or more fault topics which will trigger this rule panel options” on page 106.

5 In the **Create Rule - Select organizations** panel, select one or more Organizations, and click **Next**.

See “Create Rule - Select organizations panel options” on page 106.

6 In the **Create Rule - Setup notifications** panel, enter the required information and click **Next**.

See “Create Rule - Setup notifications panel options” on page 107.

7 In the **Create Rule - Enter name and description** panel, enter the required information and click **Finish**.

See “Create Rule - Enter name and description panel options” on page 108.

8 In the **Create Rule - Result** panel, verify that the rule has been successfully created, and click **OK**.

See “About alerts and rules” on page 102.

See “Editing rules in a perspective” on page 109.

See “Deleting rules in a perspective” on page 113.

Create Rule - Select the type of fault conditions to trigger this rule panel options

Use this panel to select a type of fault condition to trigger an alert.

Table 5-1 list the options that you can select to create a rule.

Table 5-1 Create Rule - Select the type of fault conditions to trigger this rule options

Field	Description
This rule will be triggered for all faults of type:	Select this option to trigger a rule for any faults of the selected type. You can select the following types of faults: ■ Fault ■ Risk

Table 5-1 Create Rule - Select the type of fault conditions to trigger this rule options *(continued)*

Field	Description
Enter the fault topics that will trigger the actions for this rule:	Select this option to trigger a rule when the specified fault occurs. You can enter the name of the fault. Use a colon (,) or semicolon (;) to separate multiple entries. Enter **event.alert.vom** to view the list of fault topics. You can choose a fault topic from the list. You can use a wildcard character (*) to select multiple faults. For example, you can enter **event.alert.vom.vm.*** to select all the faults on virtual machines.
Choose from a list of fault topics	Select this option to choose from a list of existing fault definitions. The fault topics listed are those which are relevant to the perspective in which you are creating the rule. If you are creating a rule in the Management Server perspective (Settings), the list includes all host fault topics and array and switch fault topics.

See “Creating rules in a perspective” on page 103.

Create Rule - Select one or more fault topics which will trigger this rule panel options

Use this panel to select the fault topics that will trigger the rule.

The fault topics listed are those which are relevant to the perspective in which you are creating the rule. If you are creating a rule in the Management Server perspective (Settings), the list includes all host fault topics and array and switch fault topics.

See “Creating rules in a perspective” on page 103.

Create Rule - Select organizations panel options

Use this panel to select the Organizations to which the rule applies.

Select the **Data Center** to apply the rule on the data center. Or select an Organization.

When you select an Organization, all the child-Organizations are selected.

When you select **Data Center**, all the Organizations are selected. In this case, your user group must be assigned the Admin role on all the Organizations.

See “Creating rules in a perspective” on page 103.

Create Rule - Setup notifications panel options

Use this wizard panel to set up notifications for the alert.

Table 5-2 lists the options to set up the notification.

Table 5-2 Create Rule - Set up notifications options

Field	Description
Email	Select to set up an email notification when the fault conditions that are specified by the rule occur.
SNMP Trap	Select to send an SNMP trap when the alert conditions that are specified by the rule occur. This option is disabled if SNMP trap settings are not configured. To configure the SNMP trap settings,
Custom script	Select to run a custom script when the alert conditions that are specified by the rule occur. **Note:** You can run a custom script only if you create a rule in the Management Server perspective.

You must set up at least one type of notification for the rule that you create else the rule will not be enabled.

Table 5-3 Notification options

Field	Description
Email: To	Enter the email address of one or more users who want to receive the notification. Separate multiple entries with a comma (,) or a semicolon (;). Example: 123@example.com, 456@example.com

Table 5-3 Notification options *(continued)*

Field	Description
Send email as daily digest	Select to send the email notification as daily digest. All alert notifications are summarized into one email and sent daily to the subscribed users.
Custom script	Browse the custom script file and upload it. You can only upload the following types of scripts: ■ Perl (.pl) ■ Shell (.sh) ■ Batch (.bat)

See "Creating rules in a perspective" on page 103.

Create Rule - Enter name and description panel options

Use this panel to assign a name and description to the alert rule.

Table 5-4 Create Rule - Enter name and description options

Field	Description
Rule Name	Enter the name of the rule. Maximum character limit is 255. Example: Restart stopped ABC program.
Description	Enter a description for this rule. The description should include the purpose of the rule. Maximum character limit is 255. Example: When the ABC program generates a service stopped alert, run the restart program script, and send an alert to the SNMP trap console.
Enable	Clear to disable the rule. An enabled rule monitors alerts for the defined conditions.

See "Creating rules in a perspective" on page 103.

Editing rules in a perspective

Using the Management Server console, you can edit the rules.

To perform this task, your user group must be assigned the Admin role on the perspective or on all the Organizations that are associated with this rule.

To edit a rule in a perspective

1. In the Home page on the Management Server console, select a perspective, and select **Manage** in the left pane.
2. Click the **Rules** tab.
3. In the details pane, right-click a rule and select **Edit**.
4. In the **Edit Rule** wizard panel, do one of the following:
   - Select **This rule will be trigger for all faults of type:**, click **Next** and skip to 5.
   - Select **Enter the fault topics that will trigger the actions for this rule:**, enter the fault definitions separated by a comma (,) or a semicolon (;), and click **Next**. Skip to 5.
   - Select **Choose from a list of fault topics**, click **Next**, and skip to 6.

   See “Edit Rule - Select the type of fault condition to trigger this rule panel options ” on page 110.
5. In the **Edit Rule - Select one or more fault topics which will trigger this rule** wizard panel, select the fault topics, and click **Next**.

   See “Edit Rule - Select one or more fault topics which will trigger this rule panel options” on page 110.
6. In the **Edit Rule - Select organizations** panel, select an organization, and click **Next**.

   See “Edit Rule - Select organization panel options” on page 111.
7. In the **Edit Rule - Setup notifications** panel, enter the required information and click **Next**.

   See “Edit Rule - Setup notifications panel options” on page 111.
8. In the **Edit Rule - Enter name and description** panel, enter the required information, and click **Finish**.

   See “Edit Rule - Enter name and description panel options” on page 112.
9. In the **Edit Rule - Result** panel, verify that the rule has been successfully created, and click **OK**.

See “About alerts and rules” on page 102.

See “Creating rules in a perspective” on page 103.

See “Deleting rules in a perspective” on page 113.

See “Enabling rules in a perspective” on page 114.

Edit Rule - Select the type of fault condition to trigger this rule panel options

Use this panel to select the type of fault conditions to trigger an alert.

Table 5-5 list the options that you can select to create a rule.

Table 5-5 Edit Rule - Select the type of fault condition to trigger this rule

Field	Description
This rule will be triggered for all faults of type:	Select this option to trigger a rule for any faults of the selected type. You can select the following types of faults: ■ Fault ■ Risk
Enter the fault topics that will trigger the actions for this rule separated by (,) or (;):	Select this option to trigger a rule when the specified fault occurs. You can enter the name of the fault. Use a colon (,) or semicolon (;) to separate multiple entries. Enter **event.alert.vom** to view the list of fault topics. You can choose a fault topic from the list. You can use wild character (*) to select multiple faults. For example, you can enter **event.alert.vom.vm.*** to select all the faults on VxVM volumes.
Choose from a list of fault topics	Select this option to choose from a list of existing fault definitions.

See “Editing rules in a perspective” on page 109.

Edit Rule - Select one or more fault topics which will trigger this rule panel options

Use this panel to select the fault topics that will trigger the rule.

See “Editing rules in a perspective” on page 109.

Edit Rule - Select organization panel options

Use this panel to select an Organization to which the alert rule applies.

Select the **Data Center** to apply the rule on the data center. Or select an Organization.

When you select an Organization, all the child-Organizations are selected.

When you select **Data Center**, all the Organizations are selected. In this case, your user group must be assigned the Admin role on all the Organizations.

See "Editing rules in a perspective" on page 109.

Edit Rule - Setup notifications panel options

Use this wizard panel to set up notifications for the alert.

Table 5-6 lists the options to set up the notification.

Table 5-6 Edit Rule panel options to set up notifications

Field	Description
Email	Select to set up an email notification when the fault conditions that are specified by the alert rule occur.
SNMP Trap	Select to send an SNMP trap when the alert conditions that are specified by the alert rule occur. This option is disabled if SNMP trap settings are not configured. To configure the SNMP trap settings,
Custom script	Select to run a custom script when the alert conditions that are specified by the rule occur. **Note:** You can run a custom script only if you edit a rule in the Management Server perspective.

You must set up at least one type of notification for the rule that you create.

Table 5-7 Notification options

Field	Description
Email: To	Enter the email address of one or more users who want to receive the notification. Separate the multiple entries with a comma (,). Example: 123@example.com, 456@example.com
Send email as daily digest	Select to send the email notification as daily digest. All alert notifications are summarized into one email and sent daily to the subscribed users.
Custom script	Browse the custom script file and upload it. You can only upload the following types of scripts: ■ Perl (.pl) ■ Shell (.sh) ■ Batch (.bat)

See "Editing rules in a perspective" on page 109.

Edit Rule - Enter name and description panel options

Use this panel to assign a name and description to the alert rule.

Table 5-8 Edit Rule - Description

Field	Description
Rule Name	Edit the name of the rule. Maximum character limit is 255. Example: Restart stopped ABC program.
Description	Edit the description for this rule. The description should include the purpose of the rule. Maximum character limit is 255. Example: When the ABC program generates a service stopped alert, run the restart program script, and send an alert to the SNMP trap console.

Table 5-8 Edit Rule - Description *(continued)*

Field	Description
Enable	Clear to disable the rule. An enabled rule monitors alerts for the defined conditions.

See "Editing rules in a perspective" on page 109.

Deleting rules in a perspective

Using the Management Server console, you can delete the rules that are no longer required.

To perform this task, your user group must be assigned the Admin role on the perspective.

To delete a rule in a perspective

1. In the Home page on the Management Server console, select a perspective, and select **Manage** in the left pane.
2. Click the **Rules** tab.
3. Right-click a rule and select **Delete**.
4. In the **Delete Rule** wizard panel, review the information, and click **OK**.

   See "Delete Rule panel options" on page 113.
5. In the **Delete Rule - Result** panel, click **OK**.

See "Creating rules in a perspective" on page 103.

See "Enabling rules in a perspective" on page 114.

See "Disabling rules in a perspective" on page 114.

Delete Rule panel options

Use this panel to delete an existing rule. Deleted rules are no longer available for sending emails, generating SNMP traps, or executing custom scripts in response to alerts.

See "Deleting rules in a perspective" on page 113.

Enabling rules in a perspective

Using the Management Server console, you can enable the rules that are in the disabled state.

To perform this task, your user group must be assigned the Admin role on the perspective.

To enable rules in a perspective

1. In the Home page on the Management Server console, select a perspective, and select **Manage** in the left pane.
2. Click the **Rules** tab.
3. Right-click a rule, and select **Enable**.
4. In the **Enable Rule** wizard panel, review the information, and click **OK**.

   See "Enable Rule panel options" on page 114.
5. In the **Enable Rule - Result** panel, click **OK**.

See "Disabling rules in a perspective" on page 114.

See "Creating rules in a perspective" on page 103.

Enable Rule panel options

Use this panel to enable the rule that is in disable state.

See "Enabling rules in a perspective" on page 114.

Disabling rules in a perspective

Using the Management Server console, you can disable the rules that are in the enabled state.

To perform this task, your user group must be assigned the Admin role on the perspective.

To disable rules in a perspective

1. In the Home page on the Management Server console, select a perspective, and select **Manage** in the left pane.
2. Click the **Rules** tab.
3. Right-click a rule, and select **Disable**.

4 In the **Disable Rule** wizard panel, review the information, and click **OK**.

See "Disable Rule panel options" on page 115.

5 In the **Disable Rule - Result** panel, click **OK**.

See "Enabling rules in a perspective" on page 114.

Disable Rule panel options

Use this panel to disable the rule that is in enabled state.

See "Disabling rules in a perspective" on page 114.

About faults and risks

Veritas InfoScale Operations Manager enables you to view all possible problems in the data center that it manages at several levels in the user interface. You can monitor the faulty status and possible risks to the managed resources.

You can view the system identified fault conditions along with their corresponding entities and the affected sources. You can automate error handling by developing the rules that trigger specific actions in response to alert conditions. You can also suppress a fault for a specific duration.

You can view the following information on faults in the data center:

- Conditions of the managed objects (applications, storage enclosures, hosts, clusters and so on) in the data center.
- The source of the fault.
- The time when the fault occurred.

See "Suppressing faults in a perspective" on page 115.

See "Restoring a suppressed fault in a perspective" on page 117.

Suppressing faults in a perspective

Using the Management Server console, you can suppress one or more faults in Veritas InfoScale Operations Manager. To suppress a fault, you can choose one of the following:

- Temporarily hide the fault.
- Disable the fault for the affected fault sources.
- Disable all the faults for the affected fault sources.

For all the options, you can either specify the date and time to keep the faults in the suppressed state, or you can suppress the faults forever.

If a fault is shared in more than one perspective, then it is suppressed in all the perspectives.

To perform this task, your user group must be assigned the Admin role on the Organization or the perspective.

To suppress a fault

1. In the Home page on the Management Server console, select a perspective, and select **Manage** in the left pane.
2. Click **Data Center**.
3. Click the **Faults** tab.
4. Right-click a fault and select **Suppress Faults**.
5. In the **Suppress Faults** wizard panel, enter the required information, and click **OK**.

   See "Suppress Faults panel options" on page 116.

See "Restoring a suppressed fault in a perspective" on page 117.

Suppress Faults panel options

Use this panel to suppress the faults in Veritas InfoScale Operations Manager. You can hide or disable the faults either temporarily or permanently. You can disable a fault for a specific object; however, the fault definition is still considered as active for other objects.

For all these options, you can either specify the date and time to keep the faults in the suppressed state, or suppress the faults forever.

Table 5-9 Suppress Faults panel options

Field	Description
Hide the selected fault(s). Show again if the problem reoccurs.	Select this option to temporarily hide the selected fault. It is essentially hiding the current instance of the fault. The fault is displayed again when it is detected.
Disable the selected fault(s) for the affected fault sources	Select this option to disable the fault for the affected fault source.
Disable all fault(s) for the affected fault sources	Select this option to disable all faults for the affected fault source.

Table 5-9 Suppress Faults panel options *(continued)*

Field	Description
Hide or disable forever	Select this option to hide the fault without specifying any time interval.
Hide or disable until	You can specify the date until which the fault remains suppressed. After this date, the fault is again considered as active in Veritas InfoScale Operations Manager.
Reason for hiding or disabling	Provide the reason why the fault was suppressed. You can enter up to 254 characters for the description.

See "Suppressing faults in a perspective" on page 115.

Restoring a suppressed fault in a perspective

You can restore the fault that is suppressed in Veritas InfoScale Operations Manager.

When you suppress a fault, you set a date until which the fault is suppressed. After the specified date, the fault is again considered as active in the system. However, Veritas InfoScale Operations Manager also provides you with the option to activate the fault before that set date.

To perform this task, your user group must be assigned the Admin role on the Organization or the perspective.

To restore a suppressed fault in a perspective

1. In the Home page on the Management Server console, select a perspective, and select **Manage** in the left pane.
2. Click **Data Center**.
3. Click the **Faults** tab.
4. Right-click the suppressed fault, and select **Restore Faults**.
5. In the **Restore Faults** panel, click **OK**.

See "Suppressing faults in a perspective" on page 115.

Chapter 6

Using reports

This chapter includes the following topics:

- About reports
- About using reports
- Running a report
- Saving a report
- Subscribing for a report
- Editing a report subscription
- Deleting a report subscription
- Sending a report through email
- Viewing my report subscriptions in a perspective
- Viewing all the report subscriptions in a perspective
- About the reports available in Veritas InfoScale Operations Manager

About reports

In Veritas InfoScale Operations Manager, you can generate a variety of reports for various purposes. The following are the broad categories under which the reports are grouped:

- Storage Utilization
  The Storage Utilization category reports provide information on the utilization of the various resources such as file systems, thin pools, enclosures, array volumes, and so on. For example the Array Volume Usage report shows array volume usage across all enclosures. The Underutilized LUNs managed by

Storage Foundation report lists all LUNs for which available capacity is not fully utilized by Storage Foundation, whereas the Underutilized File Systems report lists the file systems for which the available capacity is not fully utilized.

- Trend/Activity
  The Trend or Activity category reports provide information on the activity of a resource during a specified time duration. For example, the Uptime Analysis report displays the total time for which the selected service groups were online. Also, the report summarizes the events that affect the online availability of the service groups.
- Inventory
  The Inventory category reports provide various details about the resources in your data center, such as the mount point of a file system, or type and current state of a cluster, and so on. For example, the All Cluster report lists all the clusters in the data center, whereas the All Virtual Nodes report lists all the virtual machines discovered in your data center

SFHA licensing reports are split into three categories namely, Exception, Inventory, and True up. Exception includes reports such as the Hosts that need attention, Hosts without SFHA Licenses, and Violated Deployment Policies. The inventory category has demo licenses, license lifecycle, and product inventory reports. True up has reports related to the deployment by SPVU, server, processor, and operating system tiers.

Reports can be scoped on the **Data Center** or on an Organization in a perspective. Reports that belong to the Trend/Activity category can be scoped on time. Reports belonging to SFHA licensing cannot be scoped on time or data center.

You can subscribe for a report on a daily, weekly, or monthly basis. You can send the report to multiple recipients by entering the email addresses. You can also save a report as a comma-separated file (CSV) file.

Note: In case of large volume of data in the reports, you need to increase the JVM virtual memory limit to be able to view the reports properly. You can change this value by updating the value for Xmx attribute in the `/var/opt/.VRTSsfmcs/conf/esmweb.cfg` file.

See “About using reports” on page 120.

See “Running a report” on page 121.

See “Saving a report” on page 122.

See “Subscribing for a report” on page 123.

See “About the reports available in Veritas InfoScale Operations Manager” on page 127.

About using reports

In the Management Server console, you can perform the following tasks for reports:

- Run a report.
- Subscribe for a report.
- Save a report as a comma-separated (CSV) file.
- Share the report through email.
- Set the report generation time.

To run a report, you must specify the scope such as **Data Center** or an Organization in the perspective. Reports that belong to the Trend/Activity category can be scoped on time, whereas reports belonging to SFHA licensing cannot be scoped on time or data center.

After you run a report, in the report view, you can subscribe for the report, share the report by email, or save it as a CSV file.

To subscribe for a report, you must specify the following details:

- Frequency at which you want to receive the report.
- The email address at which you want to receive the report.

Note: Ensure that the SMTP settings are configured to receive reports at the specified email address.

You can subscribe for reports at daily, weekly, and monthly frequency. You can also set the report generation time. The default time is 1.00 A.M.

For more information on SMTP settings, and setting the report generation time, refer to the *Veritas InfoScale Operations Manager Installation and Configuration Guide*.

Use **My subscriptions** to view all your report subscriptions in a perspective. Use **All subscriptions** to view the report subscriptions of all the users in the perspective. You can also edit and delete your report subscriptions in this view.

See “About reports” on page 118.

See “Running a report” on page 121.

See “Saving a report” on page 122.

See “Subscribing for a report” on page 123.

See “Viewing my report subscriptions in a perspective” on page 126.

See “Viewing all the report subscriptions in a perspective” on page 127.

Running a report

You can run a report from the **Reports** view in a perspective.

Reports can be scoped on the **Data Center** or on an Organization in a perspective. Additionally, reports that belong to the Trend/Activity category can be scoped on time. Reports belonging to SFHA licensing cannot be scoped on time or data center.

To perform this task, your user group must be assigned the Guest role on at least one object within the perspective. Depending on the permissions you have within the perspective, the data is displayed in the report. For more information on permissions required, refer to the topic for viewing a specific report in the online help.

To run a report

1. In the Management Server console, go to a perspective, and select **Reports** in the left pane.
2. In the **Reports** view, click the report that you want to view.
3. In the **Select scope to run report** wizard panel, select the data range and Organization.

   See “Select scope to run report panel options” on page 121.
4. Click **Run**.

See “Saving a report” on page 122.

See “Subscribing for a report” on page 123.

See “About using reports” on page 120.

Select scope to run report panel options

Use this wizard panel to select the scope for running a report. You can select the data range and the Organization on which you want to run the report.

Table 6-1 Select scope to run report panel options

Field	Description
Date range	You can enter the data range in the **From** and **to** fields to specify the period for which you want to run the report. You can also choose from a combination of number and days, weeks, or months to run a report. For example, you can run the report for last seven days, or last eight weeks, or last nine months. **Note:** The **Date Range** field appears only for the reports that belong to the **Trend/Activity** category.
Organization	Select **Data Center** or an Organization. If you select **Data Center**, the data displayed in the report depends on whether you have permissions on the perspective or on certain Organizations within the data center.

See "Running a report" on page 121.

Saving a report

You can save the contents of a report to view later or to share with other users. You can save the report as a comma-separated (CSV) file on your local computer.

To perform this task, your user group must be assigned the Guest role on at least one object within the perspective. Depending on the permissions you have within the perspective, the data is displayed in the report. For more information on permissions required, refer to the topic for viewing a specific report in the online help.

To save a report

1. In the Management Server console, go to the perspective and select **Reports** in the left pane.
2. Click on a report, select the date range and Organization, and click **Run**.
3. In the report view, click **Save as CSV** to save the report as a CVS file on your local computer.

See "Subscribing for a report" on page 123.

See “Running a report” on page 121.

See “About using reports” on page 120.

Subscribing for a report

You can subscribe for a report and receive it at the email address specified by you at the selected frequency. You can choose from the HTML and CSV delivery format. The default time at which the reports are generated is 1.00 AM.

Ensure that the SMTP settings are configured to receive reports at the email address.

For more information on SMTP and report generation time settings, refer to the *Veritas InfoScale Operations Manager Installation and Configuration Guide*.

After you have subscribed for a report, and there are any changes in the permissions on the selected scope, you may not receive the required information in the report. An email notification informing the change is sent to you.

To perform this task, your user group must be assigned the Guest role on at least one object within the perspective.

To subscribe for a report

1. In the Management Server console, go to the perspective and select **Reports** in the left pane.
2. Click on a report, select the date range and Organization, and click **Run**.
3. In the report view, click **Subscribe**.
4. In the **Subscribe** wizard panel, select the frequency, delivery format, and enter the email address. Click **Subscribe**.

   See “Subscribe panel options” on page 123.
5. In the **Subscribe** panel click **Close**.

See “Editing a report subscription” on page 124.

See “Running a report” on page 121.

See “Saving a report” on page 122.

See “About using reports” on page 120.

Subscribe panel options

Use this wizard panel to subscribe to a report and receive it at the email address specified by you.

Table 6-2 Subscribe panel options

Field	Description
Select frequency	Select the frequency at which you want to receive the report. You can choose to receive the report daily, weekly once on the selected day, or on the selected day of the month.
Delivery format	Select a delivery format, for example, HTML or CSV.
Email address	Enter an email address at which you want to send the report. You can enter multiple email addresses separated by a comma (,).
Selected scope	Displays the scope of the report, such as the **Data Center** or an Organization, and the date range. **Note:** Data range is displayed only for the Trend/Activity report category. Selected scope selection is not available for SFHA licensing reports.

See "Subscribing for a report" on page 123.

Editing a report subscription

You can edit the subscription for a report. You can edit the subscription frequency and the email address at which you want to receive the report.

To perform this task, your user group must be assigned the Guest role on at least one object within the perspective.

To edit a report subscription

1. In the Management Server console, go to the perspective and select **Reports** in the left pane.
2. Click **My subscriptions**.
3. In the **My subscriptions** view, right-click the subscription, and select **Edit**.
4. In the **Subscribe** wizard panel, edit the frequency, delivery format, and the email address. Click **Subscribe**.

   See "Subscribe panel options" on page 123.
5. In the **Subscribe - Result** panel, click **Close**.

See “Running a report” on page 121.

See “Saving a report” on page 122.

See “Subscribing for a report” on page 123.

See “About using reports” on page 120.

Deleting a report subscription

Using the Management Server console, you can delete a report subscription.

To perform this task, your user group must be assigned the Guest role on at least one object within the perspective.

To delete a report subscription

1. In the Management Server console, go to the perspective and select **Reports** in the left pane.
2. Click **My subscriptions**.
3. In the **My subscriptions** view, right-click one or more subscriptions, and select **Delete**.
4. In the **Delete subscription** panel, click **Yes**.
5. In the **Delete subscription - Result** panel, click **Close**.

See “Editing a report subscription” on page 124.

See “Running a report” on page 121.

See “About using reports” on page 120.

Sending a report through email

Using the Management Server console, you can send a report through email to one or more users.

Ensure that the SMTP settings are configured to receive reports through email.

For more information on SMTP settings, refer to the *Veritas InfoScale Operations Manager Installation and Configuration Guide*.

To perform this task, your user group must be assigned the Guest role on at least one object within the perspective.

To delete a report subscription

1. In the Management Server console, go to the perspective and select **Reports** in the left pane.
2. Run a report.
3. In the report view, click **Email** and enter the details.

   See “Email report panel options” on page 126.
4. Click **Finish**.

See “Editing a report subscription” on page 124.

See “Running a report” on page 121.

See “About using reports” on page 120.

Email report panel options

Use this wizard panel to send a report through email to one or more users.

Table 6-3 Email report panel options

Field	Description
Output format	Select a delivery format, for example, HTML or CSV.
Email address	Enter an email address at which you want to send the report. You can enter multiple email addresses separated by a comma (,).

See “Sending a report through email” on page 125.

Viewing my report subscriptions in a perspective

You can view the details of all your report subscriptions in a perspective. You can view details such as the report name, the frequency for which you have subscribed, the email address at which you receive the report, the Organization scope, and the date range of the subscribed report.

In this view, you can perform the following tasks:

- Edit the subscription
- Delete the subscription

To view this information, your user group must be assigned the Admin role on the perspective.

To view my report subscriptions in a perspective

1. In the Home page on the Management Server console, go to the perspective and select **Reports**.
2. Click **My subscriptions**.

See "Editing a report subscription" on page 124.

See "Deleting a report subscription" on page 125.

See "Viewing all the report subscriptions in a perspective" on page 127.

Viewing all the report subscriptions in a perspective

You can view the details of all the report subscriptions in the perspective. You can view details such as the report name, the frequency for which the report is subscribed, the email address at which the report is sent, the Organization scope, the date range of the subscribed report, and the user name.

In this view, you can perform the following tasks:

- Edit the subscription.
- Delete the subscription.

To view this information, your user group must be assigned the Admin role on the Management Server perspective.

To view all the report subscriptions in a perspective

1. In the Home page on the Management Server console, go to the perspective and select **Reports**.
2. Click **All subscriptions**.

See "Editing a report subscription" on page 124.

See "Deleting a report subscription" on page 125.

About the reports available in Veritas InfoScale Operations Manager

The Management Server console includes the following reports available on the Availability, Licensing, Server, Storage, and Virtualization perspectives.

Table 6-4 Availability perspective

Report name	Report description
Cluster Activities	Shows activities for the selected clusters over the specified time range.
Failover Summary	Analyzes failover data of specified clusters.
Policy signature scan summary	Shows the list of cluster systems on which policy signatures have been registered and the last scan status. Scan output can be viewed separately using Show Violations Details.
Resource Fault Trending Report	Analyzes fault trends for resource types in specific clusters over the specified time range.
Resources Activities	Shows activities for the selected Resources over the specified time range.
Service Group Activities	Shows activities for the selected Service Groups over the specified time range.
Uptime Analysis	Calculates the uptime percentage for selected service groups and summarize events affecting the uptime.
VCS Failover Duration	Provides the details on the average failover duration for the service groups in the selected scope.

Table 6-5 Licensing perspective

Report name	Report description
ApplicationHA Deployments	Summarizes ApplicationHA license deployments.
Demo Licenses	Lists all the demo licenses.
Deployments By OS Tier	Lists the total license deployments by OS tier.
Deployments By Processor Tier	Lists the total deployments by processor tier.
Deployments By SPVU	Lists the total license deployments by SPVU.
Deployments By Server Tier	Lists the total license deployments by server tier.
Features Tracking Information	Shows information about the features used on the host and its usage.
Hosts Having Improper Licensing	Lists all the hosts that have improper licensing.

Table 6-5 Licensing perspective *(continued)*

Report name	Report description
Hosts Without SHFA Licenses	Lists all the hosts which do not have SFHA licenses deployed.
License Life Cycle	Lists the history of licenses present on hosts.
Per Host True-Up	Displays the total SPVU required on each host.
Price Tier Sheet	Veritas licenses and price tiering.
Product Inventory	Lists all the products installed.
VOM Deployments	Provides information on hosts connected to Management Server and whether they have Veritas products manageable from the Management Server console.
Violated Deployment Policies	Lists the deployment policies that are violated.

Table 6-6 Server perspective

Report name	Report description
Available Product Updates	Presents consolidated information about the patches applicable to the hosts in the datacenter
Data Protection	Analyzes what redundancy components are in place to protect the storage.
Disks with Single or No Active Paths	Shows all disks with one or no active paths discovered in the data center.
File System Usage	Analyzes how storage is being consumed by file systems.
LUNs Connected to Multiple Hosts	Shows all LUNs that are visible from multiple hosts.
LUNS Not Part of a Disk Group	Shows all LUNs that are not part of a Storage Foundation (SF) disk group.
Policy Signature Scan Summary	Shows the list of hosts on which policy signatures have been registered and the last scan status. Scan output can be viewed separately using Show Violations Details.

Table 6-6 Server perspective *(continued)*

Report name	Report description
Resource Mapping	Shows the mapping between host and disk including path, HBA, array port, enclosure and associated Business Applications.
Savings by File System Compression	Analyzes the vxfs file systems which are benefited by compression.
Savings by File System Deduplication	Analyzes the vxfs file systems which are benefited by deduplication.
SF Product Version	Shows the list of installed products on the host.
Storage Allocation	Analyzes how storage discovered by Storage Foundation (SF) is being allocated and consumed by hosts.
Thin Provisioned LUNs	Shows all LUNs that are thin provisioned.
Top Hosts for Thin Reclamation	Analyzes the storage that can be reclaimed from vxfs file systems. This information is available only for vxfs file systems with SF version 6.0 or later, Layout version 9 or later, File system disks not shared with other file systems and managed host version 5.0 or later.
Underutilized File Systems	Shows all file systems which are less than 80% utilized.
Underutilized LUNS Managed by Storage Foundation	Shows LUNs whose available capacity is not fully utilized (<20% used) by Storage Foundation.
Volumes Not Managed by Storage Foundation	Shows all volumes not managed by Storage Foundation.

Table 6-7 Storage perspective

Report name	Report description
Application - Array Vendor Tier Capacity	Shows capacity distribution of applications on array vendor tiers. This report will show meaningful data only if the enclosure is configured for deep discovery using Storage Insight Add-on.

Table 6-7 Storage perspective *(continued)*

Report name	Report description
Enclosure Storage Allocation by Host	For each host, shows total storage allocated across enclosures.
Enclosure Volume Usage	Shows the array volume usage across all enclosures.
Hosts Consuming Shares	Shows the list of hosts consuming shares.
NAS and Unified Storage Capacity	Shows NAS and Unified Storage Capacity.
NAS File System Capacity	Shows NAS file system capacity across all enclosures.
NAS Pool Allocation Report	Shows report for aggregates/storage pools of NetApp/Celerra arrays.
NAS Storage Pool Capacity	Shows Storage Pool Aggregate capacity across all enclosures.
Thin Pool Usage	For each thin pool, the percentage subscription and consumption along with other thin pool details, including total, subscribed and consumed capacity.
Top Thin Pools for Reclamation	Analyzes the thin pools and enclosures with the most reclaimable space. This information is available only for vxfs file systems with SF version 6.0 or later, Layout version 9 or later, and managed host version 5.0 or later.

Table 6-8 Virtualization perspective

Report name	Report description
Orphaned Virtual Disks	Displays Orphaned Virtual Disks (not used by any virtual machine in the selected reporting scope).
Storage by VM State	Displays Storage Breakup as per the Virtual Machine State.

See “About using reports” on page 120.

Section 3

Managing hosts

- Chapter 7. Overview
- Chapter 8. Managing File Replicator (VFR) operations
- Chapter 9. Managing disk groups and disks
- Chapter 10. Managing volumes
- Chapter 11. Managing file systems
- Chapter 12. Managing SmartIO
- Chapter 13. Managing application IO thresholds
- Chapter 14. Managing replications
- Chapter 15. Optimizing storage utilization

Chapter 7

Overview

This chapter includes the following topics:

- About performing Storage Foundation and replicator operations
- Viewing storage summary at the cluster level
- Viewing faults and risks at the cluster level

About performing Storage Foundation and replicator operations

In the Management Server console, you can perform various Storage Foundation and replicator tasks on a UNIX, Linux, or a Windows host. You can select multiple objects to perform various operations on them simultaneously.

Using the Management Server console you can perform Storage Foundation operations for the following storage objects:

- The volumes and the disk groups that are managed by Storage Foundation.
- The disks that are controlled by Storage Foundation.
- The native file systems and the Veritas File System that are mounted on the Storage Foundation volumes.

Volume replicator is an option of Storage Foundation that works as its fully integrated component. It benefits from the robustness, ease of use, and high performance of Storage Foundation, and at the same time, adds replication capability to Storage Foundation.

Existing Storage Foundation volume configurations, can be replicated and be transparently configured while the application is active.

Using the Management Server console you can perform replicator operations for the following storage objects:

- The disk groups that are managed by Storage Foundation.
- The databases that are controlled by Storage Foundation.
- The RVGs (Replicated Volumes Groups) that are managed by Storage Foundation.

See "About managing disk groups" on page 147.

See "About managing file systems" on page 246.

See "About managing disk groups" on page 147.

See "About managing replications" on page 303.

See "About managing Storage Foundation volumes" on page 191.

See "About Storage Foundation operations not supported on Windows host" on page 134.

See "Permissions required for views and operations on CVM cluster objects" on page 591.

About Storage Foundation operations not supported on Windows host

Following are the Storage Foundation operations that are not supported on Windows host:

Disk and disk group operations:

- Resizing disks in a disk group
- Moving disk groups

Volume operations:

- Moving volumes
- Starting synchronization of a snapshot
- Creating instant volume snapshots
- Creating space optimized snapshots for volumes
- Creating mirror-breakoff snapshots for volumes
- Restoring data from the snapshots of volumes
- Refreshing the snapshot of volumes
- Recovering volumes
- Reattaching snapshots
- Dissociating snapshots

- Splitting snapshots

File system operations:

- Creating file systems
- Defragmenting file systems
- Checking file systems
- Remounting file systems
- Mount file system
- Unmount file system
- Creating file system snapshots
- Mounting file system snapshot
- Unmounting file system snapshot
- Remounting file system snapshot
- Removing file system snapshot
- Enabling change logs
- Disabling change logs
- Synchronizing change logs
- Removing change logs

See "About performing Storage Foundation and replicator operations" on page 133.

Viewing storage summary at the cluster level

Veritas InfoScale Operations Manager enables you to view the storage summary at the cluster level. You can view the following:

- Total storage capacity utilized
  - HDD storage utilized and free space available
  - SSD storage utilized and free space available
- Total free space available

To view the Cluster Storage Summary panel

1. In the Management Server console, go to the **Server** perspective and expand **Manage** in the left pane.
2. Expand **Storage Clusters** and select a CVM cluster.

   The **Cluster Storage Summary** panel appears in the **Overview** tab.

Viewing faults and risks at the cluster level

Veritas InfoScale Operations Manager enables you to view all possible faults and risks at the cluster level. You can view the Faulted, At Risk, and Healthy count for the following:

- Hosts
- Disk Groups
- Volumes
- Disks
- RVGs
- Databases
- Exchange Servers

To view the Faults and Risks panel

1. In the Management Server console, go to the **Server** perspective and expand **Manage** in the left pane.
2. Expand **Storage Clusters** and navigate to a CVM cluster.

   The **Faults and Risks** panel appears in the **Overview** tab.

Chapter 8

Managing File Replicator (VFR) operations

This chapter includes the following topics:

- About performing File Replicator operations
- Viewing the VFR option of a host
- Creating a consistency group
- Viewing consistency groups
- Deleting a consistency group
- Associating a consistency group to a replication job
- Disassociating a consistency group from a replication job
- Viewing consistency group properties
- Creating a replication job
- Viewing File Replication Jobs
- Starting a replication job
- Pausing a replication job
- Resuming a replication job
- Stopping a replication job
- Syncing a replication job
- Modifying a replication job

- Deleting a replication job
- Viewing properties of File Replication Jobs

About performing File Replicator operations

The Management Server console allows you to create a consistency group and replication job.

To perform these tasks, your user group must be assigned the Admin role on the host or the Server perspective. The permission on the host may be explicitly assigned or inherited from a parent Organization.

In VIOM 7.2, a new **VFR** option has been introduced.

See "Viewing the VFR option of a host" on page 138.

Viewing the VFR option of a host

You can view the VFR details from the following paths:

- **Server > Manage > Uncategorized Hosts > Host > VFR**—To view the VFR option
- **Server > Manage > Storage Clusters > VFR > VFR Replication**—To view the VFR Replication option

To view the VFR option

1. In the Management Server console, go to the **Server** perspective and expand **Manage** in the left pane.
2. Expand the Organization or **Uncategorized Hosts** to locate the host.
3. Expand the host.
4. Click **VFR**.

   The VFR option of a host is displayed. It has the following tabs:

   - **File Replication Jobs**—Displays information such as Name, State, Consistency Group, Role, Source, Target, Source Mount Point, Target Mount Point, Frequency, Port, and FCL Enabled.
   - **Consistency Groups**—Displays information such as Consistency Group Name, Mount Point, and Associated Job.

To view the VFR Replication option

1. In the Management Server console, go to the **Server** perspective and expand **Manage** in the left pane.
2. Expand **Storage Clusters**.
3. Select **VFR**.
4. Click the **VFR Replication** tab.

   The VFR Replication tab displays information such as Name, State, Consistency Group, Role, Source, Target, Source Mount Point, Target Mount Point, Frequency, Port, and FCL Enabled.

Creating a consistency group

To create a consistency group

1. In the Management Server console, go to the **Server** perspective and expand **Manage** in the left pane.
2. Expand **Uncategorized Hosts** and select the host.
3. Right-click on the host and select **Replication > File Replicator > Create Consistency Group**.
4. In the **Select file system** wizard panel, select a mount point and click **Next**.

   You must select a mount point to proceed.
5. In the Name field, enter a name for the consistency group.
6. Select a file or directory. Use the **Include** and **Exclude** buttons to include and exclude files or directories.
7. Click **Finish**.

   The consistency group is created.

Viewing consistency groups

To view consistency groups

1. In the Management Server console, go to the **Server** perspective and expand **Manage** in the left pane.
2. Expand **Uncategorized Hosts**, expand the host, and expand **VFR**.
3. Click **Consistency Groups**.

   The **Consistency Groups** panel opens and lists the consistency groups.

Deleting a consistency group

To delete a consistency group

1. In the Management Server console, go to the **Server** perspective and expand **Manage** in the left pane.
2. Expand **Uncategorized Hosts**, expand the host, and expand **VFR**.
3. Click **Consistency Groups**.
4. In the **Consistency Groups** panel, right-click a consistency group and select **Delete Consistency Group**.
5. In the Delete Consistency Group window, click **OK**.

   The consistency group is deleted.

Associating a consistency group to a replication job

To associate a consistency group to a replication job

1. In the Management Server console, go to the **Server** perspective and expand **Manage** in the left pane.
2. Expand **Uncategorized Hosts**, expand the host, and expand **VFR**.
3. Click **Jobs**.
4. Do one of the following:

   - In the **File Replication Jobs** window, right-click the job and select **Associate Consistency Group**.

   Or

   - In the **Jobs** tree, select the job from which you want to disassociate the consistency group. Right-click the job and select **Associate Consistency Group**.

5. In the Associate Consistency Group window, click **OK**.

   The consistency group gets associated to the replication job.

Disassociating a consistency group from a replication job

To dissociate a consistency group from a replication job

1. In the Management Server console, go to the **Server** perspective and expand **Manage** in the left pane.
2. Expand **Uncategorized Hosts**, expand the host, and expand **VFR**.
3. Do one of the following:
   - Click **Consistency Groups**. In the **Consistency Groups** panel, right-click a consistency group and select **Dissociate From Job**.

   Or

   - In the **Jobs** tree, select the job from which you want to disassociate the consistency group. Right-click the job and select **Disassociate Consistency Group**.
4. In the Dissociate From Job window, click **OK**.

   The consistency group is disassociated from the replication job.

Viewing consistency group properties

To view consistency group properties

1. In the Management Server console, go to the **Server** perspective and expand **Manage** in the left pane.
2. Expand **Uncategorized Hosts**, expand the host, and expand **VFR**.
3. Click **Consistency Groups**.
4. In the **Consistency Groups** panel, right-click a consistency group and select **Properties**.

   In the Properties window, you can view the following:
   - General tab—Mount point name and the associated replication job.
   - Include/Exclude list tab—List of files included and excluded.

Creating a replication job

To create a replication job

1. In the Management Server console, go to the **Server** perspective and expand **Manage** in the left pane.
2. Expand **Uncategorized Hosts** and select a host.
3. Right-click on the host and select **Replication > File Replicator > Create Job**.
4. In the File System Selection panel, select the file system and click **Next**.
5. In the Attributes panel, do the following:
   - Job name—Enter the name of the replication job.
   - Source Host—Name, IP, and selected mount point is auto-populated. You can select a different IP from the **Select IP** drop-down list.
   - Target Host—Click **Select** and select a host. The IP is auto-populated after you select the host, but you can change the IP. Click **Select** and select a mount point.
   - Port—The port number is auto-populated.
   - Select Consistency Group—Check the check-box and select a consistency group from the Select Consistency Group window. If the selected consistency group does not exist on the target host for the target mount point, this option creates a consistency group with the same configuration as the source on the target.

     Note: This check-box is enabled only if a consistency groups is available for the file system.

   - Set frequency—Set the frequency for the replication job.
     - Click the **One time** radio button to schedule the job only once.
     - Click the **Periodic** radio button to schedule the job at a periodic interval. You can set it between 15 and 180 minutes.
   - Enable FCL—This check-box is checked by default. You can identify the changes since the last replication iteration using File Change Log (FCL) option.
   - Start Replication—Check this check-box to start the replication job immediately after you finish creating the job.
6. Click **Finish** to finish creating the replication job.

Viewing File Replication Jobs

To view replication jobs that have been created

1. In the Management Server console, go to the **Server** perspective and expand **Manage** in the left pane.
2. Expand **Uncategorized Hosts**, expand the host, and expand **VFR**.
3. Click **Jobs**.

   The Jobs panel opens and lists the replication jobs that have been created.

Starting a replication job

To start a replication job

1. In the Management Server console, go to the **Server** perspective and expand **Manage** in the left pane.
2. Expand **Uncategorized Hosts**, expand the host, and expand **VFR**.
3. Click **Jobs**.
4. In the File Replication Jobs panel, right-click a job and select **Start Job**.
5. In the Start Jobs window, click **OK**.

   The replication job starts.

Pausing a replication job

To pause a replication job

1. In the Management Server console, go to the **Server** perspective and expand **Manage** in the left pane.
2. Expand **Uncategorized Hosts**, expand the host, and expand **VFR**.
3. Click **Jobs**.
4. In the File Replication Jobs panel, right-click a job and select **Pause Job**.
5. In the Pause Jobs window, click **OK**.

   The replication job pauses.

Resuming a replication job

To resume a replication job

1. In the Management Server console, go to the **Server** perspective and expand **Manage** in the left pane.
2. Expand **Uncategorized Hosts**, expand the host, and expand **VFR**.
3. Click **Jobs**.
4. In the File Replication Jobs panel, right-click a job and select **Resume Job**.
5. In the Resume Jobs window, click **OK**.

   The replication job resumes.

Stopping a replication job

To stop a replication job

1. In the Management Server console, go to the **Server** perspective and expand **Manage** in the left pane.
2. Expand **Uncategorized Hosts**, expand the host, and expand **VFR**.
3. Click **Jobs**.
4. In the File Replication Jobs panel, right-click a job and select **Stop Job**.
5. In the Stop Jobs window, click **OK**.

   The replication job stops.

Syncing a replication job

To sync a replication job

1. In the Management Server console, go to the **Server** perspective and expand **Manage** in the left pane.
2. Expand **Uncategorized Hosts**, expand the host, and expand **VFR**.
3. Click **Jobs**.
4. In the File Replication Jobs panel, right-click a job and select **Sync Job**.
5. In the Sync Jobs window, click **OK**.

   The replication job syncs.

Modifying a replication job

To modify a replication job

1. In the Management Server console, go to the **Server** perspective and expand **Manage** in the left pane.
2. Expand **Uncategorized Hosts**, expand the host, and expand **VFR**.
3. Click **Jobs**.
4. In the File Replication Jobs panel, right-click a job and select **Modify Job**.
5. In the Modify Job window, modify the fields as necessary and click **OK**.

Note: You cannot modify the name of a replication job.

Deleting a replication job

To delete a replication job

1. In the Management Server console, go to the **Server** perspective and expand **Manage** in the left pane.
2. Expand **Uncategorized Hosts**, expand the host, and expand **VFR**.
3. Click **Jobs**.
4. In the File Replication Jobs panel, right-click a job and select **Delete Job**.
5. In the Delete Jobs window, click **OK**.

   The replication job is deleted.

Viewing properties of File Replication Jobs

To view properties of a replication job

1. In the Management Server console, go to the **Server** perspective and expand **Manage** in the left pane.
2. Expand **Uncategorized Hosts**, expand the host, and expand **VFR**.
3. Click **Jobs**.
4. In the File Replication Jobs panel, right-click a job and select **Properties**.

   The Properties window displays the properties of the replication job.

Chapter 9

Managing disk groups and disks

This chapter includes the following topics:

- About managing disk groups
- Creating disk groups
- Recovering disk groups
- Deporting disk groups
- Destroying disk groups
- Importing disk groups
- Adding disks to disk groups
- Resizing disks in disk groups
- Renaming disks in disk groups
- Upgrading disk groups
- Splitting disk groups
- Moving disk groups
- Joining disk groups
- About managing disks
- Initializing disks
- Replacing disks

- Recovering disks
- Mapping disks
- Unmapping disks
- Disconnecting disks
- Removing disks from disk groups
- Setting host prefix for disks
- Bringing disks online
- Taking disks offline
- Setting disk usage
- Evacuating disks
- Running or scheduling Trim
- Rescanning disks

About managing disk groups

Following is a list of operations related to disk groups that you can perform in the Management Server console.

See “Creating disk groups” on page 148.

See “Renaming disks in disk groups” on page 163.

See “Resizing disks in disk groups” on page 161.

See “Removing disks from disk groups” on page 178.

See “Splitting disk groups” on page 166.

See “Joining disk groups” on page 170.

See “Moving disk groups” on page 168.

See “Importing disk groups” on page 158.

See “Deporting disk groups” on page 157.

See “Recovering disk groups” on page 156.

See “Destroying disk groups” on page 157.

See “Upgrading disk groups” on page 166.

See "Enabling or disabling Flexible Storage Sharing on existing shared disk groups" on page 599.

See "About performing Storage Foundation and replicator operations" on page 133.

See "Permissions required for views and operations on CVM cluster objects" on page 591.

Creating disk groups

The Management Server console lets you create a disk group using one or more disks that are available on the hosts in your data center. These disk groups can be used to create volumes.

While creating a disk group, you also have the option of enabling encryption for the disk group.

You cannot create a disk group if there are no free disks available on the host.

This operation can be launched from the contexts of hosts and disks.

To perform this task, your user group must be assigned the Admin role on the host or the Server perspective. The permission on the host may be explicitly assigned or inherited from a parent Organization.

You can also create a shared disk group for a CVM cluster from the **Storage Clusters** node. More information is available on permissions required for performing operations from the cluster view.

See "Permissions required for views and operations on CVM cluster objects" on page 591.

To create a disk group

1. In the Management Server console, go to the **Server** perspective and expand **Manage** in the left pane.
2. Expand the Organization or **Uncategorized Hosts** to locate and select the host. Or, select **Storage Clusters** to navigate to a CVM cluster.
3. Do one of the following:
   - Right-click on the host and select **Create Disk Group**.
   - Click on the **Disks** tab to locate the required disk on the host. Right-click on the disk and select **Create Disk Group**.
   - Right-click a CVM cluster and select **Create Disk Group**.

4. In the **Disk Group Specifications** wizard panel, specify the required information, and click **Next**.

   See "Create Disk Group - Disk Group Specifications" on page 149.

5. If you launch the wizard by right-clicking selected disks, and some of the disks are not eligible for the specified disk group type, the wizard lists the not eligible disks. You can continue with the operation for the remaining disks. Click **Next**.

6. If you launch the wizard from a host or from a CVM cluster, in the **Select Disks** wizard panel, select the free disks from the list of disks. To search for disks using one or more filter criteria, click **Edit**. Click **Next**.

   See "Filter Criteria panel options" on page 155.

7. In the **Create Disk Group - Rename disks** wizard panel, specify an option to name the disk. Click **Next**.

   See "Create Disk Group - Rename disks panel options" on page 155.

8. In the **Create Disk Group Summary** wizard panel, verify your selections for creating the disk group. Click **Finish**.

9. In the **Result** panel, verify whether the disk group has been created successfully.

See "About managing disk groups" on page 147.

See "About managing disks" on page 172.

Create Disk Group - Disk Group Specifications

Use this wizard panel to specify the attributes for the disk group.

Table 9-1 lists the attributes that you need to specify for creating disk groups on UNIX or Linux hosts.

Table 9-2 lists the attributes that you need to specify for creating disk groups on Windows hosts.

Table 9-1 Disk group specifications panel options for a UNIX or Linux host

Field	Description
Disk Group Name	Enter a name for the disk group that you want to create. **Note:** If you use special characters [!^<>()|;:'&\\"] while creating the disk group in Command Line Interface (CLI), the navigation for such disk groups fails in the Management Server console. Also some of the views might not function as expected when used in the console.
Enable cross-platform data sharing	Select this check box if you want to enable the sharing of the data of this disk group across other platforms.
Shared (Applicable for clustered host)	Select this check box if you want to use this disk group as a shared disk group. If the wizard is launched from the context of a cluster, this option is required and the check box is selected automatically.
Enable Flexible Storage Sharing	Select this check box to enable the Flexible Storage Sharing feature for a shared disk group in a CVM cluster. See "Implementing Flexible Storage Sharing with Veritas InfoScale Operations Manager" on page 593.
Coordinator (Coordinator disk group is used by VCS)	Select this check box to create a coordinator disk group. A coordinator disk group is exclusively used for VCS I/O fencing.

Table 9-1 Disk group specifications panel options for a UNIX or Linux host *(continued)*

Field	Description
Enable Encryption	Select this check box to enable disk group encryption. This option is available only if the Key Management Server (KMS) client is configured on the host. **Note:** 1. Volumes created under encrypted disk groups are encrypted by default. 2. Encryption is supported on Linux hosts only. 3. Once you enable disk group encryption, it cannot be disabled. For more information about encryption and KMS, see the *Storage Foundation Cluster File System High Availability 7.1 Administrator's Guide - Linux*.
Description	Enter an optional comment.

Table 9-2 Disk group specifications panel options for a Windows host

Field	Description
Disk Group Name	Enter a name for the disk group that you want to create. ■ The name of the disk group can start with a letter or a number. ■ It must not contain dots or spaces. ■ You cannot give the name "BasicGroup" to a disk group that you want to create because this name is reserved by Storage Foundation. ■ The maximum character limit is 18. **Note:** If you use special characters [!^<>()|;:'&\\"] while creating the disk group in Command Line Interface (CLI), the navigation for such disk groups fails in the Management Server console. Also some of the views might not function as expected when used in the console.
Create Cluster Group	Select this check box to create a cluster disk group that Storage Foundation can control. If you select this option, the name of the disk group must be unique across the cluster.
Windows disk management compatible group	Select this check box to create a Windows disk management compatible disk group. For Windows Server 2003, this option creates a disk group that is compatible with the disk groups created with Windows Disk Management and with earlier versions of Storage Foundation for Windows products.
Add private group protection	Select this check box if you want to add protection for the disk group that you create. The private group protection protects the disk group that is located on a shared storage from being accessed by other hosts that are connected to the shared storage pool.

See "Creating disk groups" on page 148.

Select Disks panel options

Use this wizard panel to select disks when you:

- Create a disk group
- Resize a volume
- Add mirrors to a volume
- Remove disks from a disk group
- Add disks to disk groups
- Rename disks in a disk group
- Resize disks in a disk group
- Create a Storage Foundation volume

Note: For CVM clusters: If the wizard is launched from the cluster context, the master node's disk view is shown.

Note: If you are resizing a volume, this wizard panel displays all the disks that belong to the disk group of the volume and all the free disks that are available on the host.

Table 9-3 Select Disks panel options

Field	Description
Disks matching the filter criteria	This section displays the default filter criteria. Click **Edit** to modify the filter criteria. See "Filter Criteria panel options" on page 155.
Reset	Displayed when you modify the default filter criteria. Use it to apply the default filter criteria again.
Name	Displays the name of the free disk.
Condition	Displays the condition of the disk.
Enclosure	Displays the enclosure that holds the disk.
State	Displays the state of the disk.
Total Size	Displays the total available size on the disk.

Table 9-3 Select Disks panel options *(continued)*

Field	Description
Free Size	Displays the unallocated size available in the disk.
Thin	Displays if the disk is a thin reclaimable disk.
# Paths	Displays the number of active paths.
Site	Displays the name of the site for the disk.
FSS State	Displays whether disks are exported or remote. Displays only for disk groups that are enabled for Flexible Storage Sharing (FSS).
Source Host	Displays source host for a disk in a disk group that is enabled for FSS.
Do you want to make the DG Site Consistent?	This check box is available only when you select the disks from different sites. Select the check box to tag the disk group and disks, and make the disk group site consistent.
Selected disk(s):	Displays the number of disks that you have selected.
Total unallocated size in the selected disk(s):	Displays the total free (unallocated) size of the selected disks.

If you are resizing a volume, this wizard panel displays the following additional information:

- The size of the requested volume.

See “Creating disk groups” on page 148.

See “Adding disks to disk groups” on page 159.

See “Resizing volumes” on page 225.

See “Adding mirrors to volumes” on page 205.

See “Removing disks from disk groups” on page 178.

See “Renaming disks in disk groups” on page 163.

See “Resizing disks in disk groups” on page 161.

See “Creating Storage Foundation volumes” on page 192.

Filter Criteria panel options

Use this wizard panel to filter the list of disks available on a host.

Table 9-4 Filter Criteria panel options

Field	Corresponding action in the Edit the Rule Description field
Select Rules	
from enclosure list	Click **Any** and select the enclosure from the list. Click **OK**.
from vendor list	Click **Any** and select the vendor from the **Vendor** field. Click **OK**.
where raid level is	Click **Any** and select the RAID level from the **RAID Level** drop-down field. Click **OK**.
where LUN type is	Click **Any** and select the LUN type from the **LUN Type** drop-down field. Click **OK**.
where Tier is	Click **Any** and select the tier level from the **Tier** drop-down field. Click **OK**.
where active disk path is	Click **Any** and enter the number of active disk paths in the **Active Disk Paths** field. Click **OK**.

See "Creating disk groups" on page 148.

See "Removing disks from disk groups" on page 178.

See "Renaming disks in disk groups" on page 163.

See "Resizing disks in disk groups" on page 161.

See "Creating Storage Foundation volumes" on page 192.

Create Disk Group - Rename disks panel options

Use this wizard panel to provide a Storage Foundation name for the disk that you want to add to a disk group.

Table 9-5 Create Disk Group - Rename disks panel options

Field	Description
Device Name	Select this option to label the disk using the name of the disk. This option is not available if the Enable Flexible Storage Sharing option was selected for the disk group.
Disk Group Name as Prefix	Select this option to use the name of the disk group as the prefix for the disk.
Custom Prefix	Select this option to specify a custom disk prefix. Enter the prefix in the corresponding field. Click on **Populate Names** to assign the prefix as the new name for the disks.
Custom Name	Select this option to enter a new name for the selected disks.

See "Adding disks to disk groups" on page 159.

See "Creating disk groups" on page 148.

Recovering disk groups

In Management Server console, you can recover all the failed disks in a disk group.

You cannot recover a disk group if the selected disk group is not a Storage Foundation disk group or if it is in a deported or disabled state.

To perform this task, your user group must be assigned the Admin role on the host or the Server perspective. The permission on the host may be explicitly assigned or inherited from a parent Organization.

To recover a disk group

1. In the Management Server console, go to the **Server** perspective and expand **Manage** in the left pane.
2. Expand the Organization or **Uncategorized Hosts** to locate the host.
3. Expand the host and then expand **Disk Groups** to locate the disk group to be recovered.
4. Right-click on the selected disk group and select **Recover**.

5 In the **Recover Disk Group** panel, click **OK**.

6 In the **Result** panel, verify that the selected disk group has been recovered successfully. Click **OK**.

See "About managing disk groups" on page 147.

See "About managing disks" on page 172.

Deporting disk groups

The Management Server console lets you deport one or more disk groups.

You cannot deport a disk group if the selected disk group is not a Storage Foundation disk group, or is not in an imported state, or if there is any mounted file system.

To perform this task, your user group must be assigned the Admin role on the host or the Server perspective. The permission on the host may be explicitly assigned or inherited from a parent Organization.

To deport a disk group

1 In the Management Server console, go to the **Server** perspective and expand **Manage** in left pane.

2 Expand the Organization or **Uncategorized Hosts** to locate the host.

3 Expand the host and then expand **Disk Groups** to locate the disk group to be deported.

4 Right-click on the selected disk group and select **Deport**.

5 In the **Deport Disk Group** wizard panel, confirm the disk group you want to deport. Click OK

6 In the **Result** panel, verify that the selected disk group has been deported successfully.

See "About managing disk groups" on page 147.

See "About managing disks" on page 172.

Destroying disk groups

The Management Server console lets you destroy the disk group on a managed host and free the disks in that disk group for re-initialization.

You cannot destroy a disk group if the selected disk group is not a Storage Foundation disk group or if there is any mounted file system.

To perform this task, your user group must be assigned the Admin role on the host or the Server perspective. The permission on the host may be explicitly assigned or inherited from a parent Organization.

To destroy a disk group

1. In the Management Server console, go to the **Server** perspective and expand **Manage** in the left pane.
2. Expand the Organization or **Uncategorized Hosts** to locate the host.
3. Expand the host and then expand **Disk Groups** to locate the disk group to be destroyed.
4. Right-click on the selected disk group and select **Destroy**.
5. In the **Destroy Disk Group** wizard panel, confirm the disk group you want to destroy. Click **OK**.
6. In the **Result** panel, verify that the selected disk group has been destroyed successfully.

See “About managing disk groups” on page 147.

See “About managing disks” on page 172.

Importing disk groups

The Management Server console lets you import a disk group.

You cannot import a disk group if the selected disk group is not a Storage Foundation disk group or if it is not in a deported state.

To perform this task, your user group must be assigned the Admin role on the host or the Server perspective. The permission on the host may be explicitly assigned or inherited from a parent Organization.

To import a disk group

1. In the Management Server console, go to the **Server** perspective and expand **Manage** in the left pane.
2. Expand the Organization or **Uncategorized Hosts** to locate the host.
3. Expand the host and then expand **Disk Groups** to locate the disk group to be imported.
4. Right-click on the selected disk group and select **Import**.

5 In the **Import Disk Group** wizard panel, confirm the disk group you want to import. Click **OK**.

See "Import disk group panel options" on page 159.

6 In the **Result** panel, verify that the selected disk group has been imported successfully.

See "About managing disk groups" on page 147.

See "About managing disks" on page 172.

Import disk group panel options

Use this wizard panel to select the host where you want to import the disk group and the disk group options.

Table 9-6 Import disk group panel options

Field	Description
Disk group import options	
Force	Select to force the disk group import when the host cannot access all disks in the disk group. This option can be used to import a disk group that contains a failed disk, but can lead to disk group inconsistency if all disks are still usable.
Clear Host ID	Select to clear the existing host id stamp (name of the host machine that currently 'owns' the disk group) on all disks in the disk group to be imported. Do not use this option if another host is using any disks in the disk group.
Shared/Clustered	Select to import the disk group as a shared or clustered disk group.

See "Importing disk groups" on page 158.

Adding disks to disk groups

Using the Management Server console, you can add one or more free disks to an existing disk group in your data center.

A disk that is already a part of another disk group cannot be added to a disk group. Also, you cannot add disks to a deported disk group.

This operation can be launched from the contexts of disks and disk groups.

To perform this task, your user group must be assigned the Admin role on the host or the Server perspective. The permission on the host may be explicitly assigned or inherited from a parent Organization.

Note: For Windows operating system, if the disk group site consistency is enabled (by selecting the **Do you want to tag disks with the underlying enclosure site tag**), the new added disk will be tagged with the site tag of the enclosure.

To add a disk to a disk group

1. In the Management Server console, go to the **Server** perspective and expand **Manage** in the left pane.
2. Expand Organization or **Uncategorized Hosts** to locate and select the host.
3. Do one of the following:
   - Expand the host and locate the disk group to add disks. Right-click on the disk group and select **Add Disk**.
   - Click on the **Disks** tab to locate the required disk on the host. Right-click on the disk and select **Add to Disk Group**.
4. Do one of the following:
   - If you launched the operation in the context of disk groups, in the **Disk Selection** wizard panel, select the required disks from the list. To search for disks using one or more filter criteria, click **Edit**. Click **Next**.
     See “Select Disks panel options” on page 153.
     See “Filter Criteria panel options” on page 155.
   - If you launched the operation in the context of disks, in the **Select Disk Group** wizard panel, select the disk group to which you want to add the selected disks. Click **Next**. If some of the disks are not eligible for the specified disk group type, the wizard lists the not eligible disks. You can continue with the operation for the remaining disks. Click **Next**.

   See “Select Disk Group panel options” on page 161.
5. In the **Create Disk Group - Change internal disk name** wizard panel, specify an option to name the disk. Click **Next**.

   See “Create Disk Group - Rename disks panel options” on page 155.

6 In the **Add Disk To Disk Group Summary** wizard panel, verify your selections. Click **Finish**.

7 In the **Result** panel, verify that the disks have been added to the disk group successfully.

See "About managing disk groups" on page 147.

See "About managing disks" on page 172.

Select Disk Group panel options

Use this wizard panel to select a disk group to add the disk that you have selected. You can select only one disk group in this list.

Table 9-7 Select Disk Group panel options

Field	Description
Name	Name of the disk group.
Condition	Condition of the health of the disk group.
State	State of the disk group, whether Imported or Deported.
Type	Type of the disk group, whether Shared or Private.
#Disks	Number of disks available in the disk group.
#Volumes	Number of volumes available in the disk group.
Total Size	Total size of the disk group.
Free Space	Free space available in the disk group that can be used.
FSS	Whether a disk group has a Flexible Storage Sharing (FSS) state of on or off.

See "Adding disks to disk groups" on page 159.

Resizing disks in disk groups

The Management Server console lets you resize a disk in the disk group which is in the control of Storage Foundation. You can resize one or more disks in an existing

disk group on the host. If a disk group contains only one disk, you cannot perform the resize disk operation.

A disk that is a part of a Storage Foundation disk group can be resized, if there are no volumes using the disk. The disks belonging to a deported or a foreign disk group cannot be resized.

To perform this task, your user group must be assigned the Admin role on the host or the Server perspective. The permission on the host may be explicitly assigned or inherited from a parent Organization.

To resize a disk in a disk group

1. In the Management Server console, go to the **Server** perspective and expand **Manage** in the left pane.
2. Expand the Organization or **Uncategorized Hosts** to locate the host.
3. Expand the host and then expand **Disk Groups**.
4. Do one of the following:
   - Right-click on the disk group and select **Resize Disk**.
   - Select the required disk group and click on the **Disks** tab to locate the required disk on the host. Right-click on the disk and select **Resize**.
5. If you launch the operation to resize a disk from the disk group, in the **Disk Selection** wizard panel, select the required disks from the list. To search for disks using one or more filter criteria, click **Edit**. Click **Next**.

   See "Select Disks panel options" on page 153.

   See "Filter Criteria panel options" on page 155.
6. In the **Resize Disk** wizard panel, specify an option to resize the disk. Click **Next**.

   See "Resize Disk panel options" on page 162.
7. In the **Resize Disk Summary** wizard panel, verify that the details of the disk being resized. Click **Finish**.
8. In the **Result** panel, verify that the disk has been resized successfully.

See "About managing disk groups" on page 147.

Resize Disk panel options

Use this wizard panel to enter the details for resizing the disks in the disk group.

Table 9-8 Resize Disk panel options

Field	Description
Force	Force option can be used for the device having the only valid configuration copy for a disk group.
Device Name	Name of the disk that you have selected to resize.
Disk Size	Size of the disk that you have selected to resize.
LUN Size	Actual size of the disk.
Unallocated Size	Amount of free space available for allocation.
New Size	Enter the size to which you want to shrink or grow the disk space. You can specify the size in any of the following units: ■ KB ■ MB ■ GB ■ TB ■ Sectors The disk size cannot exceed the actual LUN size.
Resize Units	Enter the size you want to resize the units to. You can specify the size in any of the following units: ■ KB ■ MB ■ GB ■ TB ■ Sectors

See “Resizing disks in disk groups” on page 161.

Renaming disks in disk groups

The Management Server console lets you rename a disk in the disk group which is in the control of Storage Foundation.

You can rename one or more disks in an existing disk group on the host. The disk belonging to a deported or a foreign disk group cannot be renamed.

This operation can be launched from the contexts of disks and disk groups.

To perform this task, your user group must be assigned the Admin role on the host or the Server perspective. The permission on the host may be explicitly assigned or inherited from a parent Organization.

To rename a disk in a disk group

1. In the Management Server console, go to the **Server** perspective and expand **Manage** in the left pane.
2. Expand the Organization or **Uncategorized Hosts** to locate and select the host.
3. Expand the host and select **Disk Groups**.
4. Do one of the following:
   - Right-click on the disk group and select **Rename Disk**.
   - Select the required disk group and click on the **Disks** tab to locate the required disk on the host. Right-click on the disk and select **Rename**.
5. If you launch the operation to rename a disk from the disk group, in the **Disk Selection** wizard panel, select the disks from the list of disks. To search for disks using one or more filter criteria, click **Edit**. Click **Next**.

   See “Select Disks panel options” on page 153.

   See “Filter Criteria panel options” on page 155.
6. In the **Rename Disk** wizard panel, specify an option to name the disks. Edit the names in the text box if required and click **Next**.

   See “Rename Disk panel options” on page 165.

   See “Rename Disk panel options” on page 165.
7. In the **Rename Disk Summary** wizard panel, verify the details of the disks and the disk group. Click **Finish**.
8. In the **Result** panel, verify that the disks have been renamed in the disk group successfully.

See “About managing disk groups” on page 147.

See “About managing disks” on page 172.

Rename Disk panel options

Use this wizard panel to provide a Storage Foundation name for the disk you want to rename in a disk group.

Select one of the following options to name the disk:

Table 9-9 Rename Disk Details panel options

Field	Description
Device Name	Select this option to label the disk using the name of the disk. This option is not available if the Enable Flexible Storage Sharing option was selected for the disk group.
Disk Group Name as Prefix	Select this option to use the name of the disk group as the prefix for the disk.
Custom Prefix	Select this option to specify a custom disk prefix. Enter the prefix in the corresponding field. Click on **Populate Names** to assign the prefix as the new name for the disks.
Custom Disk Name	Select this option to specify a custom disk name. Enter the disk name in the corresponding field.

Table 9-10 New Name Details panel options

Field	Description
Name	Name of the disk. Lists all the disks available in the disk group.
VxVm Name	Internal name or Device Media (DM) name of the selected disks you want to rename.
New Name	Enter the new name for the selected disk. The name must be less than 32 characters and should not contain the following special characters [!^<>()|;:'&\\"]. It should be alpha-numeric. If one or more disks are being renamed, the same name cannot be used for other disks being renamed.

See “Renaming disks in disk groups” on page 163.

Upgrading disk groups

The Management Server console lets you upgrade the disk groups.

To perform this task, your user group must be assigned the Admin role on the host or the Server perspective. The permission on the host may be explicitly assigned or inherited from a parent Organization.

To upgrade a disk group

1. In the Management Server console, go to the **Server** perspective and expand **Manage** in the left pane.
2. Expand the Organization or **Uncategorized Hosts** to locate and select the host.
3. Expand the host and select **Disk Groups**.
4. Right-click on the selected disk group and select **Upgrade**.
5. In the **Upgrade Disk Group** wizard panel, click **OK**.
6. In the **Result** panel, verify that the selected disk group has been upgraded successfully. Click **OK**.

See “About managing disk groups” on page 147.

See “About managing disks” on page 172.

Splitting disk groups

The Management Server console lets you split a disk group to create a new one which is used to take a backup of the data. To avoid the source disk group from getting removed after the split operation, it should contain at least one disk.

You cannot split a disk group if the selected disk group is not a Storage Foundation disk group or if it is in a deported or disabled state.

To perform this task, your user group must be assigned the Admin role on the host or the Server perspective. The permission on the host may be explicitly assigned or inherited from a parent Organization.

To split a disk group

1. In the Management Server console, go to the **Server** perspective and expand **Manage** in the left pane.
2. Expand the Organization or **Uncategorized Hosts** to locate and select the host.
3. Expand the host and select **Disk Groups**.

4 Right-click on the selected disk group and select **Split**.

5 In the **Split Disk Group** wizard panel, confirm the disk group you want to split. Click **OK**.

See "Split Disk Group panel options" on page 167.

6 In the **Result** panel, verify that the selected disk group has been split successfully.

See "Creating disk groups" on page 148.

See "About managing disk groups" on page 147.

Split Disk Group panel options

Use this wizard panel to split a disk group to create a new one.

Table 9-11 Split Disk Group panel options

Name	Description
Disk Group to Split	Displays the name of the disk group from which the objects have to be split.
New Disk Group	Enter the name of the target disk group to which the split objects have to be moved.
Expand to include disks which provide volume closure sets	Select the **Expand** check box to specify that the objects to be split. It includes all other disks that contain subdisks that are associated with the specified disk group.
Split disk group by	Select one of the following to move the disk groups by: ■ **Disks** ■ **Volumes** ■ **Volume Sets**
Name	Displays the name of the selected disks, volume, or volume sets.
Condition	Displays the condition of the selected disks, volumes, or volume sets.
Size	Displays the size if the selected disks, volumes, or volume sets.

Table 9-11 Split Disk Group panel options *(continued)*

Name	Description
VxVM Name	Displays the internal name or Device Media (DM) name of the selected disks you want to rename.
State	Displays the state of the selected disks.
FSS State	Displays whether disks are exported or remote Displayed only for disk groups that are enabled for Flexible Storage Sharing (FSS).
Source Host	Displays source host for a disk in a disk group that is enabled for FSS.
Mounted	Displays if the volume is mounted or not.

See "Splitting disk groups" on page 166.

Moving disk groups

The Management Server console lets you move the disks within a disk group to another disk group. To avoid the source disk group from getting removed after the move operation, it should contain at least one disk.

You cannot move a disk group if the selected disk group is not a Storage Foundation disk group or if it is in a deported or disabled state. This operation is not supported on Windows host.

To perform this task, your user group must be assigned the Admin role on the host or the Server perspective. The permission on the host may be explicitly assigned or inherited from a parent Organization.

To move a disk group

1. In the Management Server console, go to the **Server** perspective and expand **Manage** in the left pane.
2. Expand the Organization or **Uncategorized Hosts** to locate and select the host.
3. Expand the host and select **Disk Groups**.
4. Right-click on the selected disk group and select **Move**.

5 In the **Move Disk Group** wizard panel, confirm the disk group you want to move. Click **OK**.

See "Move Disk Group panel options" on page 169.

6 In the **Result** panel, verify that the selected disk group has been moved successfully.

See "Creating disk groups" on page 148.

See "About managing disk groups" on page 147.

Move Disk Group panel options

Use this wizard panel to move disks within a disk group to another disk group.

Table 9-12 Move Disk Group panel options

Name	Description
Select Target Disk Group	
Disk Group	Displays the name of the disk group from which the objects have to be moved.
Condition	Displays the condition of the target disk group.
Type	Displays the type of the target disk group.
CDS	Displays the size of the target disk group.
FSS	Displays if Flexible Storage Sharing is on or off for the disk group..
Expand to include disks which provide volume closure sets	Select the **Expand** check box to specify that the objects to be moved. It includes all other disks that contain subdisks that are associated with the specified objects.
Move disk group by	Select one of the following to move the disk groups by: ■ **Disks** ■ **Volumes**
Name	Displays the name of the selected disk or volume.
Condition	Displays the condition of the selected disk or volume.

Table 9-12 Move Disk Group panel options *(continued)*

Name	Description
Size	Displays the size of the selected disk or volume.
VxVM Name	Internal name or Device Media (DM) name of the selected disks you want to rename.
Mounted	Displays if the volume is mounted or not.
FSS State	Displays whether disks are exported or remote Displayed only for disk groups that are enabled for Flexible Storage Sharing (FSS).
Source Host	Displays source host for a disk in a disk group that is enabled for FSS.

See "Moving disk groups" on page 168.

Joining disk groups

The Management Server console lets you join two disk groups by moving all Storage Foundation objects from an imported source disk group to an imported target disk group. The source disk group is removed when the join is complete.

Before performing this operation, all applications that access the volumes should be stopped. Unmount all the file systems that are configured in the volumes.

The reconfiguration must involve an integral number of physical disks. Objects to be joined must not contain open volumes. For a disk group join to succeed, both the source and target disk groups must contain at least one disk that can store copies of the configuration database after the split.

To perform this task, your user group must be assigned the Admin role on the host or the Server perspective. The permission on the host may be explicitly assigned or inherited from a parent Organization.

To join a disk group

1. In the Management Server console, go to the **Server** perspective and expand **Manage** in the left pane.
2. Expand the Organization or **Uncategorized Hosts** to locate and select the host.
3. Expand the host and select **Disk Groups**.

4 Right-click on the selected disk group and select **Join**.

5 In the **Join Disk Group** wizard panel, select the target disk group you want to join. Click **OK**.

See "Join Disk Group panel options" on page 171.

6 In the **Result** panel, verify that the selected disk group has been joined successfully.

See "Creating disk groups" on page 148.

See "About managing disk groups" on page 147.

Join Disk Group panel options

Use this wizard panel to join two disk groups by moving all Storage Foundation objects from an imported source disk group to an imported target disk group. The source disk group is removed when the join is complete.

Table 9-13 Join Disk Group panel options

Name	Description
Source Disk Group	Displays the name of the source disk group from which the objects have to be moved.
Target disk group to join	Select a target disk group from the list displayed.
Name	Displays the name of the target disk group.
Condition	Displays the condition of the target disk group.
Type	Displays whether the target disk group is private or shared.
CDS	Displays whether or not the target disk group is a Cross-Platform Data Sharing (CDS) disk group.
Coordinator	Displays whether or not the target disk group is a coordinator disk group.
FSS	Displays whether or not Flexible Storage Sharing is enabled on the target disk group.

See "Joining disk groups" on page 170.

About managing disks

Following is a list of operations related to disks that you can perform in the Management Server console.

See “Initializing disks” on page 172.

See “Replacing disks” on page 174.

See “Recovering disks” on page 175.

See “Disconnecting disks” on page 178.

See “Removing disks from disk groups” on page 178.

See “Setting host prefix for disks” on page 179.

See “Bringing disks online” on page 180.

See “Taking disks offline” on page 181.

See “Setting disk usage” on page 181.

See “Evacuating disks” on page 184.

See “Running or scheduling Trim ” on page 186.

See “Rescanning disks” on page 188.

See “Exporting and un-exporting disks for Flexible Storage Sharing” on page 598.

See “About performing Storage Foundation and replicator operations” on page 133.

See “Permissions required for views and operations on CVM cluster objects” on page 591.

Initializing disks

The Management Server console lets you initialize a disk and bring it under the control of Storage Foundation.

You can select one or more disks to launch this operation. A disk that is a part of any disk group cannot be initialized. The disks in a foreign or a deported state cannot be initialized. This operation is supported on multiple disks across multiple hosts.

Note: Before you initialize a disk, you need to take a backup of the data.

To perform this task, your user group must be assigned the Admin role on the host or the Server perspective. The permission on the host may be explicitly assigned or inherited from a parent Organization.

Initializing a disk

1. In the Management Server console, go to the **Server** perspective and expand **Manage** in the left pane.
2. Expand the Organization or **Uncategorized Hosts** to locate the host.
3. Expand the host and click on the **Disks** tab.
4. Right-click on the selected disk and select **Initialize**.
5. In the **Initialize Disk** wizard panel, verify the details of the disks selected and also specify the format. Click **Ok**.

   See "Initialize Disk panel options" on page 173.
6. In the **Result** panel, verify that the disks have been initialized successfully.

See "About managing disk groups" on page 147.

See "About managing disks" on page 172.

Initialize Disk panel options

Use this wizard panel to initialize a disk.

Table 9-14 Initialize Disk panel options for UNIX/Linux hosts

Field	Description
Selected Disks	Verify the disk details selected to be initialized.
Specify format	Specify one of the following formats: ■ Cdsdisk ■ Simple ■ Sliced The default format for HP and AIX disk is hpdisk and aixdisk.
Force	Select this option if the initialization is to be forced.

Table 9-15 Initialize Disk panel options for Windows hosts

Field	Description
Selected Disks	Verify the disk details selected to be initialized.

Table 9-15 Initialize Disk panel options for Windows hosts *(continued)*

Field	Description
Specify format	Specify one of the following formats: ■ **MBR**: Master boot record ■ **GPT**: GUID partition table

See "Initializing disks" on page 172.

Replacing disks

The Management Server console lets you replace a faulted disk that is controlled by Storage Foundation and bring the volume that uses this disk back into function. In the **State** column in the list of disks, you can identify a replaceable disk by the label 'In Use - Not accessible'. This disk is highlighted with an error icon in the list of disks.

To perform this task, your user group must be assigned the Admin role on the host or the Server perspective. The permission on the host may be explicitly assigned or inherited from a parent Organization.

Note: Based on the volume of data in the disk, the time taken to complete this operation can be from a few hours to several days.

To replace a disk

1. In the Management Server console, go to the **Server** perspective and expand **Manage** in the left pane.
2. Expand the Organization or **Uncategorized Hosts** to locate the host.
3. Expand the host and select **Disk Groups** and click on the **Disks** tab.
4. Right-click on the selected disk and select **Replace**.
5. In the **Replace Disk** wizard panel, select the new disk. Click **Next**.

   See "Replace Disk panel options" on page 175.
6. In the **Replace Disk Summary** panel, verify the details that you have specified to replace the faulted disk. Click **Finish**.
7. In the **Result** panel, verify that the selected disk has been replaced successfully.

See "Recovering disks" on page 175.

Replace Disk panel options

Use this wizard panel to select a free disk for replacing the faulted disk on a host.

This panel displays the following information on all the free disks on the host where you perform this operation:

Table 9-16 Replace Disk panel options

Field	Description
Disks matching the below filter criteria	This field lists the filter criteria for the disks that are displayed on this panel.
Edit	Click to modify the list of disks by filtering with a different search criteria. See "Filter Criteria panel options" on page 155.
Reset	Click to reset the selections.
Name	Name of the free disk device.
Condition	Condition of the disk.
Enclosure	Enclosure where the disk belongs to.
State	State of the usage of the disk.
Total size	Total size in the disk.
Thin	Whether the disk is a thin disk.
FSS State	Displays whether disks are exported or remote. Displayed only for a disk in a disk group that is enabled for Flexible Storage Sharing (FSS).
Source Host	Displays source host for a disk in a disk group that is enabled for FSS.

See "Replacing disks" on page 174.

Recovering disks

The Management Server console lets you recover the faulted disks that are controlled by Storage Foundation and bring the volumes that use these disks back into function. In the **State** column in the list of disks, you can identify a recoverable

disk by the label 'In Use - Recoverable'. These disks are highlighted with a warning icon on the list of disks.

To perform this task, your user group must be assigned the Admin role on the host or the Server perspective. The permission on the host may be explicitly assigned or inherited from a parent Organization.

Note: Recover operation of the faulted disks can be time consuming. Based on the volume of data in the disks, this operation is completed within a few hours or days.

To recover a disk

1. In the Management Server console, go to the **Server** perspective and expand **Manage** in the left pane.
2. Expand the Organization or **Uncategorized Hosts** to locate the host.
3. Expand the host and select **Disk Groups** and click on the **Disks** tab.
4. Right-click on the selected disk and select **Recover**.
5. In the **Recover Disk** wizard panel, confirm the action.

   See "Recover Disk panel options" on page 176.
6. In the **Result** panel, verify that the selected disks have been recovered successfully.

See "Replacing disks" on page 174.

See "About managing disks" on page 172.

Recover Disk panel options

Use this wizard panel to confirm the action of recovering disconnected disks. Click **Yes** to recover the disks that you have selected.

See "Recovering disks" on page 175.

Mapping disks

The Management Server console lets you map a disk from a remote node in a cluster to a local node. You can use the mapped disk like other local storage on the node. In the Management Server Console, you can view the mapped disk like other disks in the local storage list of the node. You can map a disk to multiple remote nodes.

Note: Disk mapping is not persistent across cluster restarts.

To perform this task, your user group must be assigned the Admin role on the host or the Server perspective. The permission on the host may be explicitly assigned or inherited from a parent Organization.

Note: You cannot map disk groups from the Management Server Console. To map disk groups from the command line, see the *Storage Foundation Cluster File System Administrator's Guide*. Mapping of the disks that are part of a disk group is persistent across cluster restarts.

To map a disk

1. In the Management Server console, go to the **Server** perspective and expand **Manage** in the left pane.
2. Expand the Organization or Uncategorized Hosts to locate and select the host.
3. Right-click the host and select **Map Disks**.
4. In the Map Disks panel, select the remote disk that you want to map to the host, and click **OK**.

See "Unmapping disks" on page 177.

Unmapping disks

The Management Server console lets you unmap a remote disk from a local node in a cluster. When you unmap the remote disk from a node, the disk is removed from the list of local storage on the node.

To perform this task, your user group must be assigned the Admin role on the host or the Server perspective. The permission on the host may be explicitly assigned or inherited from a parent Organization.

To unmap a disk

1. In the Management Server console, go to the **Server** perspective and expand **Manage** in the left pane.
2. Expand the Organization or **Uncategorized Hosts** to locate and select the host.
3. Right-click the host, and select **Unmap Disks**.
4. In the Unmap Disks panel, select the remote disk that you want to unmap, and click **OK**.

See “Mapping disks” on page 176.

Disconnecting disks

The Management Server console lets you disconnect or detach one or more disks. This operation is used to remove a disk from Storage Foundation control.

You cannot disconnect a disk if the selected disk group is not a Storage Foundation disk group or if it is not in an imported state.

To perform this task, your user group must be assigned the Admin role on the host or the Server perspective. The permission on the host may be explicitly assigned or inherited from a parent Organization.

To disconnect a disk

1. In the Management Server console, go to the **Server** perspective and expand **Manage** in the left pane.
2. Expand the Organization or **Uncategorized Hosts** to locate the host.
3. Expand the host and then expand **Disk Groups** and click on the **Disks** tab.
4. Right-click on the selected disk and select **Disconnect Disk**.
5. In the **Disconnect Disk** wizard panel, confirm the disk you want to disconnect. Click **OK**.
6. In the **Result** panel, verify that the selected disk has been disconnected successfully.

See “About managing disk groups” on page 147.

See “About managing disks” on page 172.

Removing disks from disk groups

The Management Server console lets you remove a disk from a disk group which is in the control of Storage Foundation.

You can remove one or more disks from an existing disk group on your host. A disk that is a part of any disk group can be removed if there are no volumes using the disk. The disks cannot be removed from a deported or a foreign disk group.

This operation can be launched from the contexts of disks and disk groups.

Warning: If all the disks within the disk group are selected for removal, the disk group is destroyed.

To perform this task, your user group must be assigned the Admin role on the host or the Server perspective. The permission on the host may be explicitly assigned or inherited from a parent Organization.

To remove a disk from a disk group

1. In the Management Server console, go to the **Server** perspective and expand **Manage** in the left pane.
2. Expand the Organization or **Uncategorized Hosts** to locate and select the host.
3. Expand the host and select **Disk Groups** and click on the **Disks** tab.
4. Do one of the following:
   - Right-click on the disk group and select **Remove Disk**.
   - Click on the **Disks** tab to locate the required disk on the host. Right-click on the disk and select **Remove From Disk Group**.
5. If you launch the operation to remove a disk from the disk group, in the **Disk Selection** wizard panel, select the disks from the list of disks. To search for disks using one or more filter criteria, click **Edit**. If all the disks in the disk group are selected to be removed, the disk group is destroyed. Click **Next**.

   See "Select Disks panel options" on page 153.

   See "Filter Criteria panel options" on page 155.
6. In the **Remove Disk From Disk Group Summary** wizard panel, verify the details of the disks and the disk group. Click **Finish**.
7. To remove the disk from the **Disks** tab, in the **Remove Disk From Disk Group** wizard panel, confirm the disk group from which you want to remove the selected disks. Click **OK**.
8. In the **Result** panel, verify that the disks have been removed successfully from the disk group.

See "Creating disk groups" on page 148.

See "About managing disk groups" on page 147.

See "About managing disks" on page 172.

Setting host prefix for disks

The Management Server console lets you add a host prefix to the disks.

Note: This option is enabled only for Direct-attached storage (DAS).

To perform this task, your user group must be assigned the Admin role on the host or the Server perspective. The permission on the host may be explicitly assigned or inherited from a parent Organization.

To set a host prefix

1. In the Management Server console, go to the **Server** perspective and expand **Manage** in left pane.
2. Expand the Organization or **Uncategorized Hosts** to locate the host.
3. Right-click on the selected host and select **Set Host Prefix**.
4. In the **Set Host Prefix** wizard panel, enter the host prefix. Click **OK**.
5. In the **Result** panel, verify that the host prefix has been applied successfully.

See “About managing disks” on page 172.

Bringing disks online

In Management Server Console, you can restore access to a disk which has been taken offline. The disk is made available to Storage Foundation again.

To perform this task, your user group must be assigned the Admin role on the host or the Server perspective. The permission on the host may be explicitly assigned or inherited from a parent Organization.

To online a disk

1. In the Management Server console, go to the **Server** perspective and expand **Manage** in the left pane.
2. Expand the Organization or **Uncategorized Hosts** to locate and select the host.
3. Expand the host and select **Disk Groups** and click on the **Disks** tab.
4. Right-click on the selected disk and select **Online**.
5. In the **Online Disk** wizard panel, click **OK**.

   See “Online Disk panel option for making the disks online” on page 181.
6. In the **Result** panel, verify that the selected disk has been brought online successfully. Click **OK**.

See “Taking disks offline” on page 181.

See “About managing disks” on page 172.

Online Disk panel option for making the disks online

Use this wizard panel to make the disks online.

See "Bringing disks online" on page 180.

Taking disks offline

In Management Server Console, you can make the disk offline.

To perform this task, your user group must be assigned the Admin role on the host or the Server perspective. The permission on the host may be explicitly assigned or inherited from a parent Organization.

To offline a disk

1. In the Management Server console, go to the **Server** perspective and expand **Manage** in the left pane.
2. Expand the Organization or **Uncategorized Hosts** to locate and select the host.
3. Expand the host and select **Disk Groups** and click on **Disks**.
4. Right-click on the selected disk and select **Offline**.
5. In the **Offline disk** wizard panel, click **OK**.
6. In the **Result** panel, verify that the selected disk has been made offline successfully. Click **OK**.

See "Bringing disks online" on page 180.

See "About managing disks" on page 172.

Setting disk usage

The Management Server console lets you set disk usage to exclude a disk from certain operations by marking it as a spare, excluding it from Hot Relocation, or marking it as a reserved disk, or as a disk reserved for ISP use only on the UNIX host. On the Windows host, you set disk usage to exclude a disk from certain operations by marking it as preferred, secondary, no hot use, and reserved or any one of the three.

To perform this task, your user group must be assigned the Admin role on the host or the Server perspective. The permission on the host may be explicitly assigned or inherited from a parent Organization.

To set disk usage

1. In the Management Server console, go to the **Server** perspective and expand **Manage** in the left pane.
2. Expand the Organization or **Uncategorized Hosts** to locate and select the host.
3. Expand the host, select **Disk Groups**, and click on **Disks**.
4. Right-click on the selected disk and select **Set Disk Usage**.
5. In the **Set Disk Usage - UNIX host** wizard panel, select the option to set the disk usage on a UNIX or Linux host. Click **OK**.

   See "Set Disk Usage panel options" on page 182.

   In the **Set Disk Usage - Windows host** wizard panel, select the option to set the disk usage on a Windows host. Click **OK**.

   See "Set Disk Usage - Windows host panel options" on page 182.
6. In the **Result** panel, verify that the selected disk has been evacuated successfully. Click **OK**.

See "About managing disks" on page 172.

Set Disk Usage panel options

Use this wizard panel to select one of the following options to mark the disk.

Table 9-17 Set Disk Usage panel options

Field	Description
Spare	Select to designate the disk for use by the hot-relocation facility.
No hot use	Select to exclude the disk from use by the hot-relocation facility.
Reserved	Select to not create a subdisk on this disk unless the disk is specified while creating the subdisk.

See "Setting disk usage" on page 181.

Set Disk Usage - Windows host panel options

Use this wizard panel to select one of the following options to mark the disk.

Table 9-18 Set Disk Usage - Windows host panel options

Field	Description
Disk	Displays the name of the disk used.
Hot Relocation Mode	The default for Storage Foundation for Windows is to have automatic hot relocation mode inactive. If an I/O error occurs in a redundant subdisk, the subdisk is not automatically relocated to another disk. Use the Set Disk Usage command to designate preferred disks as targets for hot relocation.
Reserved for manual use	Disks reserved for manual use are not available in automatic selection operations, including hot relocation. In an automatic selection operation, Storage Foundation for Windows chooses the storage where the operation occurs. Generally, the user is given a choice between allowing SFW to "Auto select disks" or "Manually select disks." Examples of commands that allow automatic selection are New Volume and Add Mirror. Reserving a disk for manual use lets you prevent any unwanted volumes or subdisks from being placed on that disk and gives you complete control over the disk.
Hot relocation Usage	Select one of the following Hot Relocation targets: ■ Preferred Hot Relocation Target ■ Secondary Hot Relocation Target ■ Not used as a Hot Relocation Target
Preferred Hot Relocation Target	If there is an I/O failure anywhere in the system, SFW first looks for space on disks that have been marked as preferred hot-relocation targets for redundant subdisks.
Secondary Hot Relocation Target	This option is the default for all disks. During the hot relocation operation, if there are no disks selected as preferred targets or if there is no space available on those disks, Veritas InfoScale Operations Manager chooses space on disks marked as secondary targets.

Table 9-18 Set Disk Usage - Windows host panel options *(continued)*

Field	Description
Not used as a Hot Relocation Target	This option does not allow any hot-relocated subdisks to be moved to the selected disks. It differs from the "Reserved for manual use" option in that the disk remains available for other automatic selection operations.

See "Setting disk usage" on page 181.

Evacuating disks

The Management Server console lets you evacuate a disk by moving the contents of the volumes from one disk to another. If a disk begins to fail, you can attempt to preserve the volumes on that disk by evacuating the disk.

You can evacuate a disk if you plan to remove the disk or use the disk elsewhere. The active volumes on the replacement disk are synchronized automatically.

You cannot recover a disk group if the selected disk group is not a Storage Foundation disk group or if it is in a deported or disabled state.

To perform this task, your user group must be assigned the Admin role on the host or the Server perspective. The permission on the host may be explicitly assigned or inherited from a parent Organization.

To evacuate a disk

1. In the Management Server console, go to the **Server** perspective and expand **Manage** in the left pane.
2. Expand the Organization or **Uncategorized Hosts** to locate and select the host.
3. Expand the host, select **Disk Groups**, and click on **Disks**.
4. Right-click on the selected disk and select **Evacuate**.
5. In the **Evacuate Disk** wizard panel, select a disk selection method. If you choose to manually assign destination disks then select the required disk from the list. Click **OK**.

   See "Evacuate Disk panel options" on page 185.
6. In the **Result** panel, verify that the selected disk has been evacuated successfully. Click **OK**.

See "About managing disks" on page 172.

Evacuate Disk panel options

Use this wizard panel to evacuate disks.

Table 9-19 Evacuate Disk panel options

Field	Description
Auto assign destination disks	Select this option to confirm that the destination disks are automatically assigned for evacuating the disks.
Manually assign destination disks	Select this option to confirm that the destination disks that you specify for evacuating are assigned.
Disks matching the below filter criteria	This section displays the default filter criteria. Click **Edit** to modify the filter criteria. See "Filter Criteria panel options" on page 155.
Reset	Displayed when you modify the default filter criteria. Use it to apply the default filter criteria again.
Name	Displays the name of the free disk device.
Condition	Displays the condition of the disk.
Enclosure	Displays the enclosure that holds the disk.
State	Displays the state of the disk usage.
Total Size	Displays the total available size in the disk.
Thin	Displays if the disk is a thin reclaimable disk.
# Paths	Displays the number of paths.
FSS State	Displays whether disks are exported or remote. Displayed only for disk groups that are enabled for Flexible Storage Sharing (FSS).
Source Host	Displays source host for a disk in disk group that is enabled for FSS.
Required Space	Displays the total disk space required for evacuating the disk.

Table 9-19 Evacuate Disk panel options *(continued)*

Field	Description
Total free space on selected disk(s)	Displays the total available space on the selected disks.

See "Evacuating disks" on page 184.

Running or scheduling Trim

In Management Server Console, you can run or schedule an SSD Trim operation. A Trim operation allows an operating system to inform a solid-state drive (SSD) which blocks of data are no longer considered in use and can be wiped internally. Trim enables the SSD to handle garbage collection overhead, that would otherwise significantly slow down future write operations to the involved blocks. Trim can be used by Storage Foundation (VxFS or VxVM) to clean up the blocks that do not have any valid data in it, making it available for optimum utilization of SSD drives. If a volume has only thin reclaimable LUNs, or only SSDs, or both, Veritas InfoScale Operations Manager performs reclamation on the thin LUNs and Trim operation on the SSD devices.

This operation can be launched from the contexts of disks and volumes.

To perform this task, your user group must be assigned the Admin role on the host or the Server perspective. The permission on the host may be explicitly assigned or inherited from a parent Organization.

To run or schedule a Trim command

1. In the Management Server console, go to the **Server** perspective and expand **Manage** in the left pane.
2. Expand the Organization or **Uncategorized Hosts** to locate and select the host.
3. Do one of the following:
   - Expand the host. Expand **Disk Groups** under the host to locate and select a disk group. Right-click and select **Disks** > **Trim**.
   - Expand the host. Expand **Volume** under the host to locate and select a volume. Right-click and select **File System** > **Trim**.

4 In the panel **Trim - Run/Schedule** wizard panel, select if you want to run or schedule the command. Enter the schedule frequency options and click **OK**.

See "Trim - Run or Schedule panel options" on page 187.

5 In the **Result** panel, verify that the schedule has been set up successfully. Click **OK**.

See "About managing disks" on page 172.

Trim - Run or Schedule panel options

Use this wizard panel to run or schedule a Trim command.

Table 9-20 Trim - Schedule panel options

Field	Description
Run Now	Select to run the Trim command immediately.
Schedule	Select to define the frequency of running the Trim command.
Schedule Name	Enter a name for the schedule.
Schedule Desc	Enter a description of the schedule.
Frequency	Select a frequency for scheduling the SSD Trim command. The values under the **When** column changes with the options that you select here. The available options are daily, weekly, or monthly.

Table 9-20 Trim - Schedule panel options *(continued)*

Field	Description
When	Specify exactly when you want to schedule the SSD Trim command as follows: ■ For **Daily** schedule: Define how often you want to schedule the command on an hourly basis. Also, define the time and start date. ■ For **Weekly** schedule: Select the **Every weekday** option to run the command every week days from Monday to Friday in a week. If you want to run the command on specific days of the week, select the day from the **Recur every week on** field. Also, define the time and start date for the weekly schedule. ■ For **Monthly** schedule: To run the command on a specific day of the month, enter the date in the **Day** field. To run the command on the recurring days of a month, choose the required options from the drop-down list. Also, define the time and start date.

See "Running or scheduling Trim " on page 186.

Rescanning disks

The Management Server console lets you rescan all attached disks for disk configuration changes. It also updates information on removable media.

To perform this task, your user group must be assigned the Admin role on the host or the Server perspective. The permission on the host may be explicitly assigned or inherited from a parent Organization.

To rescan a disk

1. In the Management Server console, go to the **Server** perspective and expand **Manage** in the left pane.
2. Expand the Organization or **Uncategorized Hosts** to locate and select the host.
3. Right-click on the host and select **Rescan Disks**.

4 In the **Rescan Disks** wizard panel, confirm the disks you want to rescan. Click **OK**.

5 In the **Result** panel, verify whether the disks have been rescanned successfully.

See “About managing disks” on page 172.

Chapter 10

Managing volumes

This chapter includes the following topics:

- About managing Storage Foundation volumes
- Creating Storage Foundation volumes
- Stopping volumes
- Recovering volumes
- Reactivating volumes
- Deleting volumes
- Moving volumes
- Renaming volumes
- Adding mirrors to volumes
- Removing the mirrors of volumes
- Creating instant volume snapshots
- Creating space optimized snapshots for volumes
- Creating mirror break-off snapshots for volumes
- Dissociating snapshots
- Reattaching snapshots
- Resizing volumes
- Restoring data from the snapshots of volumes
- Refreshing the snapshot of volumes

- Configuring a schedule for volume snapshot refresh
- Adding snapshot volumes to a refresh schedule
- Removing the schedule for volume snapshot refresh
- Setting volume usage
- Splitting snapshots
- Starting synchronization of snapshots
- Enabling FastResync on volumes
- Disabling FastResync on volumes

About managing Storage Foundation volumes

Following is the list of operations related to volumes that you can perform in the Management Server console.

See “Resizing volumes” on page 225.

See “Stopping volumes” on page 200.

See “Recovering volumes” on page 200.

See “Reactivating volumes ” on page 201.

See “Deleting volumes” on page 202.

See “Moving volumes” on page 203.

See “Renaming volumes” on page 204.

See “Adding mirrors to volumes” on page 205.

See “Removing the mirrors of volumes” on page 212.

See “Creating instant volume snapshots” on page 214.

See “Creating space optimized snapshots for volumes” on page 218.

See “Creating mirror break-off snapshots for volumes” on page 220.

See “Dissociating snapshots” on page 222.

See “Reattaching snapshots” on page 224.

See “Restoring data from the snapshots of volumes” on page 230.

See “Refreshing the snapshot of volumes” on page 232.

See “Configuring a schedule for volume snapshot refresh” on page 235.

See “Adding snapshot volumes to a refresh schedule” on page 237.

See “Removing the schedule for volume snapshot refresh” on page 238.

See “Setting volume usage” on page 239.

See “Splitting snapshots” on page 240.

See “Starting synchronization of snapshots” on page 241.

See “Enabling FastResync on volumes” on page 242.

See “Disabling FastResync on volumes” on page 243.

See “Permissions required for views and operations on CVM cluster objects” on page 591.

Creating Storage Foundation volumes

The Management Server console lets you create a Storage Foundation (VxVM) volume on a host.

To create a Storage Foundation volume, you need to select a Storage Foundation disk group.

Select a host, disk group, or a volume to launch this operation.

To perform this task, your user group must be assigned the Admin role on the host or the Server perspective. The permission on the host may be explicitly assigned or inherited from a parent Organization.

While creating a volume, you also have the option of enabling encryption for the volume.

Note: If the underlying disk group is already encrypted, the volume is encrypted by default.

Encryption is supported only on Linux hosts.

The following capabilities are not supported by volume encryption:

- Encryption of root and swap volumes
- Linked break-off snapshot of encrypted volumes
- Replication of encrypted volumes
- Encryption of existing volumes

To create a Storage Foundation volume on a UNIX/Linux host

1. In the Management Server console, go to the **Server** perspective and expand **Manage** in the left pane.
2. Expand Organization, or **Uncategorized Hosts** to locate and select the host.
3. Do one of the following:
   - Right-click on the host and select **Create Volume**.
   - Expand the host and select **Volumes**. Right-click and select **Create Volume**.
   - Expand the host and locate the disk group to create the volume. Right-click and select **Create Volume**.
4. If you have launched this operation from a disk group, skip this step. If you have launched the operation from a host or from **Volumes**, in the **Create Volume** wizard panel, choose a disk selection method. Click **Next**.

   See "Create Volume – Select Disk Group and Disk Selection method panel options" on page 194.
5. In the **Volume Attributes** wizard panel, specify the attributes of the new volume. Click **Next**.

   See "Volume attributes panel options for creating volumes on UNIX or Linux hosts for specifying values " on page ?.
6. If you manually select the disks from the disk group, click **Next** to view the **Disk Selection** wizard panel. To search for disks using one or more filter criteria, click **Edit**. Select the disks to create the volume. Click **Next**.

   See "Select Disks panel options" on page 153.

   See "Filter Criteria panel options" on page 155.
7. If you want to create a file system on the volume, in the **File System Options** wizard panel, select **Create file system**. Specify the file system and mount options. Click **Next**.

   See "Create File System - File System Options" on page 247.

   See "Advanced Options panel" on page 251.
8. In the **Create Volume Summary** panel, verify the details of the new volume. Click **Finish**.

To create a Storage Foundation volume on a Windows host

1. In the Management Server console, go to the **Server** perspective and expand **Manage** in the left pane.
2. Expand Organization, or **Uncategorized Hosts** to locate the host.

3 Do one of the following:

- Right-click on the host and select **Create Volume**.
- Expand the host and locate the disk group to create the volume. Right-click and select **Create Volume**.
- Expand the host and select **Volumes**. Right-click and select **Create Volume**.

4 In the **Create Volume** wizard panel, choose an option of disk selection. Click **Next**.

5 If you manually select the disks from the disk group, click **Next** to view the **Disk Selection** wizard panel. To search for disks using one or more filter criteria, click **Edit**. Select disks to create the volume. Click **Next**.

See "Select Disks panel options" on page 153.

See "Filter Criteria panel options" on page 155.

6 In the **Volume Attributes** wizard panel, specify the attributes of the new volume. Click **Next**.

See "Add Drive Letter, Path and Create File System details panel options" on page 199.

7 In the **Create Volume Summary** panel, verify the details of the new volume. Click **Finish**.

See "Creating file systems" on page 246.

See "Recovering volumes" on page 200.

See "Deleting volumes" on page 202.

See "About managing Storage Foundation volumes" on page 191.

Create Volume – Select Disk Group and Disk Selection method panel options

Use this wizard panel to choose the disk group for the volume.

Table 10-1 Create Volume – Select Disk Group and Disk Selection method panel options

Field	Description
Select disk group to use for the volume	
Name	Displays the name of the disk group.
Condition	Displays the condition of the disk group.

Table 10-1 Create Volume – Select Disk Group and Disk Selection method panel options *(continued)*

Field	Description
State	Displays the state of the disk group.
Type	Displays the type of the disk group.
#Disks	Displays the number of disks in the disk group.
#Volumes	Displays the number of volumes in the disk group.
Total Size	Displays the total size of the disk group.
Free Space	Displays the free size available in the disk group.

See "Creating Storage Foundation volumes" on page 192.

Volume attributes panel options for creating volumes on UNIX or Linux hosts for specifying values

Use this wizard panel to specify volume attributes for creating a volume on a UNIX or Linux host.

Table 10-2 Create Volume - Volume attributes panel options

Field	Description
Volume Name	Enter a name for the volume. This is a Storage Foundation-specific name that is different from the volume label for the file system. The name must be less than 32 characters and should not contain special characters [!^<>()|;:'&\/"]. **Note:** If you use special characters [!^<>()|;:'&/\"] while creating the volume in Command Line Interface (CLI), certain operations for such volumes fail in the Management Server console. Also some of the views might not function as expected when used in the console.
Size	Specify the volume size or click **Max Size**.

Table 10-2 Create Volume - Volume attributes panel options *(continued)*

Field	Description
Layout	Select a layout type for the volume. Depending on the number of disks available in a disk group or the number of disks selected manually, different layout options are enabled or disabled. Select one of the following layout types: ■ Concat ■ Stripe ■ Mirror-Concat ■ Mirror-Stripe ■ RAID-5 ■ Concat-Mirror ■ Stripe-Mirror
Mirror Info	If you selected Concatenated, Concatenated Mirrored, or Striped Mirror layout, this option is enabled. User can select the total number of mirrors, and set mirror across option.
Stripe Info	If you selected Striped, RAID-5, or Striped Mirror layout, this option is enabled. Users can specify the number of Columns and Stripe Unit Size.. Stripe Across option can also be set.
Advance Options	
Enable Logging	If Mirror Info and Enable FastResync are selected, **Enable Logging** is enabled by default.
Enable FastResync	Select to enable FastResync.
Initialize Zero	Select to clear the volume before enabling it for general use.

Table 10-2 Create Volume - Volume attributes panel options *(continued)*

Field	Description
No layered Volumes	If you are creating a mirrored volume, you can select this option to prevent the creation of a layered volume. In cases where a layered volume layout is appropriate, Storage Foundation can create a layered volume when a non- layered layout is specified. This option ensures that the volume has a non-layered layout. If a layered layout is selected, this option is ignored.
Create cache object	Select this check box to create a cache object.
Enable Encryption	Select this check box to enable encryption for the volume. This option is enabled only if KMS is configured on the host and the underlying disk group is non-encrypted. If the underlying disk group is already encrypted, the volume is encrypted by default. **Note:** 1. Encryption is supported only on Linux hosts. 2. Once you enable volume encryption, it cannot be disabled. For more information about encryption and KMS, see the *Storage Foundation Cluster File System High Availability 7.1 Administrator's Guide - Linux*.
Disk selection method	
Automatic	Select this option to let Storage Foundation automatically decide the disks to use from the disk group.
Manual	Select this option to manually allocate storage from the disk group.

See "Creating Storage Foundation volumes" on page 192.

Volume attributes panel options for creating volumes on Windows hosts

Use this wizard panel to specify volume attributes for creating a volume on a Windows host.

Table 10-3 Volume attributes panel options

Field	Description
Volume Name	Type a name for the volume. This is a Storage Foundation-specific name that is different from the volume label for the file system. The name must be less than 32 characters and should not contain special characters [!^<>()|;:'&\\"]. **Note:** If you use special characters [!^<>()|;:'&/\"] while creating the volume in Command Line Interface (CLI), certain operations for such volumes fail in the Management Server console. Also some of the views might not function as expected when used in the console.
Size	Specify the volume size.
Layout	Select a layout type for the volume. Depending on the number of disks available in a disk group or the number of disks selected manually, different layout options are enabled or disabled. Select one of the following layout types: ■ Concatenated ■ Striped ■ RAID-5
Stripe Info	If you selected Striped or RAID-5, this option is enabled. Stripe Across option can also be set. Users can type the number of Columns and Stripe Unit Size.
Mirror Info	If you selected Concatenated, Concatenated Mirrored, this option is enabled. User can select the total number of mirrors, and set mirror across option.

Table 10-3 Volume attributes panel options *(continued)*

Field	Description
Enable Logging	If Mirror Info is selected, you can select **Enable Logging** to enable logging.

See "Creating Storage Foundation volumes" on page 192.

Add Drive Letter, Path and Create File System details panel options

In this wizard panel, select one of the three options described below:

- Assign a drive letter.
- Do not assign a drive letter.
- Mount as an empty NTFS folder.

Table 10-4 Drive letter, path, and create file system panel options

Field	Description
Format volume	If you assign a drive letter, select this check box to format the volume.
Select a file system	Select one of the following file systems: ■ FAT ■ FAT32 ■ NTFS
Allocation size	Select one of the following allocation sizes: ■ Default ■ 512 ■ 1024 ■ 2048 ■ 4096 ■ 8192 ■ 16K ■ 32K ■ 63K
File system label	Define a label for the file system.
Perform a quick format	Select to perform a quick format of the volume.

Table 10-4 Drive letter, path, and create file system panel options *(continued)*

Field	Description
Enable file and folder compression	Select to enable compression of the file and folder.

See "Creating Storage Foundation volumes" on page 192.

Stopping volumes

The Management Server console lets you stop a volume on a managed host and make it unavailable to an application.

You can stop one or more volumes on a managed host. An error page is displayed if the selected volume is not a Storage Foundation volume. This operation can be performed only on a volume of an imported disk group.

To perform this task, your user group must be assigned the Admin role on the host or the Server perspective. The permission on the host may be explicitly assigned or inherited from a parent Organization.

To stop a volume

1. In the Management Server console, go to the **Server** perspective and expand **Manage** in the left pane.
2. Expand the Organization or **Uncategorized Hosts** to locate the host.
3. Expand the host and then expand **Volumes** to locate the volume to be stopped.
4. Right-click on the required volume and select **Stop**.
5. In the **Stop Volume** wizard panel, confirm the volume you want to stop. Click **OK**.
6. In the **Result** panel, verify that the selected volume has been stopped successfully.

See "Creating Storage Foundation volumes" on page 192.

See "About managing Storage Foundation volumes" on page 191.

Recovering volumes

The Management Server console lets you recover a volume on a managed host. It recovers plexes and volumes after disk replacement.

You can recover one or more volumes on a managed host. An error page is displayed if the selected volume is not a Storage Foundation volume. This operation

can be performed on an imported disk. This operation is not supported on Windows host.

To perform this task, your user group must be assigned the Admin role on the host or the Server perspective. The permission on the host may be explicitly assigned or inherited from a parent Organization.

To recover a volume

1. In the Management Server console, go to the **Server** perspective and expand **Manage** in the left pane.
2. Expand the Organization or **Uncategorized Hosts** to locate and select the host.
3. Expand the host and then expand **Volumes** to locate the volumes to be recovered.
4. Right-click on the required volume and select **Recover**.
5. In the **Recover Volume** wizard panel, confirm the volume you want to recover. Click **OK**.
6. In the **Result** panel, verify that the selected volume has been recovered successfully.

See "Creating Storage Foundation volumes" on page 192.

See "About managing Storage Foundation volumes" on page 191.

Reactivating volumes

The Management Server console lets you reactivate one or more volumes on a managed host.

You can reactivate a volume if the selected volume is a Storage Foundation volume which is in a disabled state. This operation is not supported on UNIX host.

To perform this task, your user group must be assigned the Admin role on the host or the Server perspective. The permission on the host may be explicitly assigned or inherited from a parent Organization.

To reactivate a volume

1. In the Management Server console, go to the **Server** perspective and expand **Manage** in the left pane.
2. Expand the Organization or **Uncategorized Hosts** to locate the host.
3. Expand the host and then expand **Volumes** to locate the volume to be reactivated.

4 Right-click on the required volume and select **Reactivate**.

5 In the **Reactivate Volume** wizard panel, confirm the volume you want to reactivate. Click **OK**.

6 In the **Result** panel, verify that the selected volume has been reactivated successfully.

See "Creating Storage Foundation volumes" on page 192.

See "About managing Storage Foundation volumes" on page 191.

Deleting volumes

The Management Server console lets you remove a volume from your environment if you do not need it anymore. You can delete a volume from the disk groups in Veritas InfoScale Operations Manager. By removing a volume, you can reduce the usage of your disk space. This option also lets you remove a snapshot volume.

In a data center environment, a volume can have other dependencies also. When you delete a volume, the Veritas InfoScale Operations Manager console gives you the detailed information on various dependencies of the selected volume. This report lets you analyze the effects of the delete operation for a volume.

You cannot delete a volume if the volume has applications running on the file system or running on raw volumes or if the volume has data on it. Multiple volumes can be selected for this operation

To perform this task, your user group must be assigned the Admin role on the host or the Server perspective. The permission on the host may be explicitly assigned or inherited from a parent Organization.

To delete a volume

1 In the Management Server console, go to the **Server** perspective and expand **Manage** in the left pane.

2 Expand the Organization or **Uncategorized Hosts** to locate the host.

3 Expand the host and then expand **Volumes** to locate the volume to be deleted.

4 Right-click on the required volume and select **Delete**.

5 In the **Delete Volume** wizard panel, specify the required information. Click **OK**.

See "Delete Volume panel options" on page 203.

6 In the **Result** page, verify that the selected volumes have been deleted successfully.

See "Creating Storage Foundation volumes" on page 192.

See “About managing Storage Foundation volumes” on page 191.

Delete Volume panel options

Use this wizard panel to confirm the volume delete operation.

If you want to delete an enabled volume, select the **Force** check box. If you have an enabled volume in the list of volumes that you want to delete, you must select this check box to continue. Else, the operation fails.

See “Deleting volumes” on page 202.

Moving volumes

The Management Server console lets you move a volume from one set of disks to other disks.

To perform this task, your user group must be assigned the Admin role on the host or the Server perspective. The permission on the host may be explicitly assigned or inherited from a parent Organization.

To move a volume

1. In the Management Server console, go to the **Server** perspective and expand **Manage** in the left pane.
2. Expand the Organization or **Uncategorized Hosts** to locate the host.
3. Expand the host and then expand **Volumes** to locate the volume to be deleted.
4. Right-click on the required volume and select **Move**.
5. In the **Move Volume** wizard panel, specify the required information. Click **OK**.

   See “Move Volume panel options” on page 203.
6. In the **Result** page, verify that the selected volumes have been moved successfully.

See “Creating Storage Foundation volumes” on page 192.

See “About managing Storage Foundation volumes” on page 191.

Move Volume panel options

Use this wizard panel to move a volume from one disk to another disk.

Table 10-5 Move Volumes panel options

Field	Description
Select Source Disks	Select one or more source disks.
Auto assign destination disks	Select this option to confirm that Storage Foundation automatically assigns destination disks for moving the volume.
Manually assign destination disks	Select this option to confirm that Storage Foundation assigns the disks that you specify for moving the volume.

See "Moving volumes" on page 203.

Renaming volumes

The Management Server console lets you rename a volume. You can select one or more volumes to launch this operation.

To perform this task, your user group must be assigned the Admin role on the host or the Server perspective. The permission on the host may be explicitly assigned or inherited from a parent Organization.

To rename a volume

1. In the Management Server console, go to the **Server** perspective and expand **Manage** in the left pane.
2. Expand the Organization or **Uncategorized Hosts** to locate the host.
3. Expand the host and expand **Volumes**.
4. Right-click the selected volume and select **Rename**.
5. In the **Rename Volume** wizard panel, specify the details, and click **Ok**.

   See "Rename Volume panel options" on page 204.
6. In the **Result** panel, verify that the volume is successfully renamed.

See "Creating Storage Foundation volumes" on page 192.

See "About managing Storage Foundation volumes" on page 191.

Rename Volume panel options

Use this wizard panel to rename a volume.

Table 10-6 Rename Volume panel options

Field	Description
Name volumes using:	
Disk Group Name as Prefix	Select this option to use the name of the disk group as the prefix for the volume.
Custom Name	Select this option to specify a custom name. In the **New Name** column, click on the required row and edit the name.
Custom Prefix	Select this option to specify a custom prefix. Enter the prefix in the corresponding field. Click **Populate Names** to assign the prefix as the new name for the volumes.

See "Renaming volumes" on page 204.

Adding mirrors to volumes

Volume mirrors are the exact copies of data of a volume retained in a disk. When you add a mirror to a volume, a copy of the volume is created on the disk which is not already being used by the volume. The Management Server console lets you create the following types of volume mirrors:

- A datavol volume mirror that keeps a backup of the data in a volume
- A volume mirror that can be used as a snapshot volume when it is detached from the volume using the mirror break-off snapshot operation

To perform this task, your user group must be assigned the Admin role on the host or the Server perspective. The permission on the host may be explicitly assigned or inherited from a parent Organization.

To add a datavol mirror to a volume

1. In the Management Server console, go to the **Server** perspective and expand **Manage** in left pane.
2. Expand the Organization or **Uncategorized Hosts** to locate and select the host.
3. Expand the host and select **Volumes**.
4. Right-click on the required volume and select **Add mirror**.

5 In the **Add mirror** wizard panel, select **Add mirror** and specify the required information. Click **Next**.

See “Add mirror - Options” on page 207.

6 If you have opted to select the disks manually, select the disks from the **Disk selection** panel. To search for disks using one or more filter criteria, click **Edit**. Click **Next**.

See “Select Disks panel options” on page 153.

7 In the **Advanced options for Add mirror** panel, specify the required information. Click **Next**.

See “Add mirror - Advanced options” on page 209.

8 In the **Add mirror summary** verify the details that you have specified for adding mirrors. Click **Finish**.

9 In the **Result** panel, verify that the mirrors have been added successfully.

To add a snapshot mirror to a volume

1 In the Management Server console, go to the **Server** perspective and expand **Manage** in left pane.

2 Expand the Organization or **Uncategorized Hosts** to locate and select the host.

3 Expand the host and select **Volumes**.

4 Right-click on the required volume and select **Add mirror**.

5 In the **Add mirror** wizard panel, select **Add snap-ready mirror** and specify the required information. Click **Next**.

See “Add mirror - Options” on page 207.

6 If you have opted to select the disks manually, select the disks from the **Disk selection** panel. To search for disks using one or more filter criteria, click **Edit**. Click **Next**.

See “Select Disks panel options” on page 153.

7 In the **Add mirror summary** verify the details that you have specified for adding mirrors. Click **Finish**.

8 In the **Result** panel, verify that the mirrors have been added successfully.

See “Removing the mirrors of volumes” on page 212.

See “About managing Storage Foundation volumes” on page 191.

Add mirror - Options

Use this wizard panel to add mirrors to an existing volume.

Table 10-7 Add mirror - Options

Field	Description
Operation	
Add mirror	Select this option if you want to add a normal mirror to the volume that you have selected.
Add snap-ready mirror	Select this option if you want to add a snapshot mirror to the volume that you have selected.
Number of mirrors to add	Enter the number of mirrors. The number of mirrors cannot exceed the number of disks available to add mirror. Maximum number of mirrors allowed is 31.
Disk selection	Specify the method of selecting the disks to add the mirror. Select **Manual** to go to the **Disk selection** wizard panel to select the disks. Select **Automatic** to let Storage Foundation choose the disks for mirror.
Disable track alignment	This option is available only for the volumes that reside on a Windows host. Select this check box to disable the track alignment for the mirror. The track alignment feature optimizes the I/O performance by setting the dynamic volumes to store the blocks of data in alignment with the boundaries of the physical track of the disk

Table 10-7 Add mirror - Options *(continued)*

Field	Description
Fast Resync options	This option is available only for adding a snapshot mirror to a volume that resides on a UNIX host. If Fast Mirror Resync (FMR) is disabled for the selected volume, you can enable it by clicking this option. This option is displayed only if FMR is not enabled for the selected volume. Select **Automatic** to enable FMR automatically. Select **Manual** to enable the FMR manually. Click **Enable FastResync** to specify the details. See "Enable FastResync option panel options" on page 208.
Edit	If FMR is enabled for the selected volume, you can edit it by clicking this option. This option is enabled only if the selected volume does not already have a snapshot volume. See "Enable FastResync option panel options" on page 208.

See "Adding mirrors to volumes" on page 205.

Enable FastResync option panel options

Use this wizard panel to enable FastResync for a snapshot volume that you want to create.

Table 10-8 Enable FastResync option panel options

Field	Description
FastResync (DCO) mirror	Select a number from the drop-down list to specify the number of data change object (DCO) mirrors that you want to create. When the data volume is updated, these updates are retained in the form of logs on the DCO mirrors.

Table 10-8 Enable FastResync option panel options *(continued)*

Field	Description
DCO region size (KB)	Enter the size for the DCO mirror that you want to create. The default value is 64 kilobytes.
Enable DRL (Dirty region logging)	Select an option from the drop-down list to enable the Dirty region logging (DRL). DRL expedites the recovery of mirrored volumes after a system crash. The available options are: ■ **On**: Select this option to enable DRL. ■ **Off**: Select this option to disable DRL. ■ **Sequential**: Select this option to configure the sequential DRL. This option is used for volumes such as those that are used for database replay logs for which data are written in a sequential manner. The sequential DRL limits the number of dirty regions in a volume and helps faster recovery of data.
Select DCO disks (By default Volume Manager decides disks to be used)	Select the disks that you want to keep as DCO disks. Select the disks from the list of disks.

See "Add mirror - Options" on page 207.

See "Adding mirrors to volumes" on page 205.

Add mirror - Advanced options

Use this wizard panel to specify advanced options to add a mirror to an existing volume.

This wizard panel displays the following:

Table 10-9 Add mirror - Advanced options

Field	Description
Choose layout	Choose the layout of the mirror that you want to add. Choose **Concatenated** to set the mirror as concatenated mirror. This is the default layout type. If you have selected a single disk for mirroring the volume, this is the only option available for the mirror layout. Choose **Stripe** to set the mirror as a stripe mirror. This option needs minimum two disks for creating the mirror. If the number of selected disks equals the number of mirrors specified then this option is disabled.
Columns	This option is available only for a striped mirror. Enter the number of columns for the striped mirror in the field. A column refers to an area on the disk where all or a portion of the volume resides. Striping is achieved by allocating data alternately and evenly across the columns within a plex.
Stripe unit size	The size of each stripe unit. The default size is 128 kilobytes. The size to be specified should be in multiples of eight.

Table 10-9 Add mirror - Advanced options *(continued)*

Field	Description
Mirror across	Select an option to specify the areas for mirroring. The available options are: ■ Enclosure - Select this option if you want to create the mirrors across enclosures. ■ Controller - This option is available only for the volumes that reside on a UNIX host. Select this option if you want to create the mirrors across controllers. ■ Target - This option is available only for the volumes that reside on a Windows host. Select this option if you want to create the mirrors across targets. ■ Channel - This option is available only for the volumes that reside on a Windows host. Select this option if you want to create the mirrors across channels. ■ Port - This option is available only for the volumes that reside on a Windows host. Select this option if you want to create the mirrors across ports.
Stripe across	Select an option to specify the areas for striping. The available options are: ■ Enclosure - Select this option if you want to stripe the data across enclosures. ■ Controller - This option is available only for the volumes that reside on a UNIX host. Select this option if you want to stripe the data across controllers. ■ Target - This option is available only for the volumes that reside on a Windows host. Select this option if you want to stripe the data across targets. ■ Channel - This option is available only for the volumes that reside on a Windows host. Select this option if you want to stripe the data across channels. ■ Port - This option is available only for the volumes that reside on a Windows host. Select this option if you want to stripe the data across ports.

See “Adding mirrors to volumes” on page 205.

Removing the mirrors of volumes

Volume mirrors are the copies of volumes that are retained in a disk. These disks are updated concurrently when the data in the original volume change. In certain circumstances, such as to free some disk space, you may want to remove the mirrors that you have configured for a volume. The Management Server console lets you add and remove the following types of mirrors:

- A mirror that has been created for a volume.
- A snapshot mirror which is configured to use as a break-off mirror for a volume.

To perform this task, your user group must be assigned the Admin role on the host or the Server perspective. The permission on the host may be explicitly assigned or inherited from a parent Organization.

To remove the mirrors of a volume

1. In the Management Server console, go to the **Server** perspective and expand **Manage** in the left pane.
2. Expand the Organization or **Uncategorized Hosts** to locate the host.
3. Expand the host and select **Volumes**.
4. Right-click on the required volume and select **Mirror** > **Remove**.
5. In the **Remove mirror** wizard panel, specify the required information. Click **Next**.

   See “Remove mirror panel options” on page 212.
6. In the **Remove mirror summary** panel, verify the details that you have specified for removing the volume mirror. Click **Finish**.
7. In the **Result** panel, verify that the selected mirrors have been removed successfully.

See “Adding mirrors to volumes” on page 205.

See “About managing Storage Foundation volumes” on page 191.

Remove mirror panel options

Use this wizard panel to select the normal mirrors or the snapshot mirrors that you want to remove from the volume that you have selected. This view displays the name, size, and the layout of the volume that you have selected.

Note: You cannot remove all the mirrors of a volume. At least one mirror must be retained on the volume that you have selected.

Table 10-10 Remove mirror panel options

Field	Description
Remove mirror	Select this option to remove a data mirror from the volume.
Remove snapshot mirror	Select this option to remove a snapshot mirror from the volume.
Remove by	This option is available only if you remove a volume mirror. Select **Mirror** to choose the specific mirrors to be removed. Select **Disk** to remove the mirrors by choosing the disks from the list.
Mirror selection	This option is available only if you remove a snapshot mirror. Select **Manual** to manually choose the mirrors that you want remove. Select the plexes from the list of mirrors. Select **Automatic** to let Storage Foundation choose the mirrors to be removed.
Convert to data mirror	This option is available only when you remove a snapshot mirror manually. Select this check box to convert the snapshot mirror to a normal data mirror in the volume.
Plex name	Name of the mirror that you want to remove from the volume. This information is displayed only if you: ■ Remove a normal volume mirror by choosing the specific mirrors to be removed. ■ Remove a snapshot mirror by performing a manual selection of mirrors to be removed.

Table 10-10 Remove mirror panel options *(continued)*

Field	Description
Plex state	State of the mirror that you want to remove from the volume. This information is displayed only if you: ■ Remove a normal volume mirror by choosing the specific mirrors to be removed. ■ Remove a snapshot mirror by performing a manual selection of mirrors to be removed.
Plex type	Type of the mirror that you want to remove from the volume. This information is displayed only if you: ■ Remove a normal volume mirror by choosing the specific mirrors to be removed. ■ Remove a snapshot mirror by performing a manual selection of mirrors to be removed.

See "Removing the mirrors of volumes" on page 212.

Creating instant volume snapshots

Instant snapshots let you create snapshots of volumes using compatible volumes in the same disk group or by creating compatible volumes and using them for snapshots. Using the Management Server console you can create instant snapshots for the volumes.

The operation is not supported on Windows host.

To perform this task, your user group must be assigned the Admin role on the host or the Server perspective. The permission on the host may be explicitly assigned or inherited from a parent Organization.

To create an instant volume snapshot

1. In the Management Server console, go to the **Server** perspective and expand **Manage** in the left pane.
2. Expand the Organization, Applications, or **Uncategorized Hosts** to locate the host.

3 Do one of the following:

- Click on the **Databases** tab and locate the required database. Right-click on the database and select **Create Snapshot**.
- Expand the host. Expand **Disk Groups** and select the required disk group. Right-click and select **Create Snapshot**.
- Expand the host. Expand **Volumes** to locate and select the required volume. Right-click and select **Snapshot** > **Create**.

4 In the **Snapshot level selection panel options** wizard panel, select the **Create volume level snapshot** option. Click **Next**.

See "Snapshot level selection panel options" on page 270.

5 In the **Create snapshot - snapshot type selection page panel options** wizard panel, select the snapshot type as **Instant**. Click **Next**.

6 In the **Create snapshot - disk selection page panel options** wizard panel, select the free disks from the list of disks. To search for disks using one or more filter criteria, click **Edit**. Click **Next**.

See "Create Volume Snapshot - disk selection page panel options" on page 216.

7 If you want to create a different snapshot type for each volume, then on the **Create snapshot - snapshot type selection page panel options** click the **Advanced** button. In the **Create snapshot - Advanced Options panel options** wizard panel, click **Configuration** to configure the volume details of the instant snapshot selected.

See "Create Volume Snapshot - Advanced Options" on page 216.

8 In the **Create snapshot - Create instant snapshot** wizard panel, configure the snapshot volume. Click **Save** to go back to the **Create snapshot - Advanced Options** wizard panel. Click **Next**.

See "Create Volume Snapshot - Instant Snapshot panel options" on page 217.

9 In the **Configure snapshot - Configure snapshots summary** panel, verify the configuration information. Click **Finish**.

10 In the **Result** panel, verify that the snapshots have been configured successfully.

See "Creating space optimized snapshots for volumes" on page 218.

See "Creating mirror break-off snapshots for volumes" on page 220.

See "About managing Storage Foundation volumes" on page 191.

Create Volume Snapshot - disk selection page panel options

Table 10-11 Create snapshot - disk selection page panel options

Field	Description
Disk Selection Options	
Select disks for snapshot automatically	Select this option to let Storage Foundation automatically decide the disks to use for the volume.
Select disks for snapshot manually	Select this option to manually allocate storage for the volume.
Disks matching the below filter criteria	This section displays the default filter criteria. Click **Edit** to modify the filter criteria. See "Filter Criteria panel options" on page 155.
Reset	Displayed when you modify the default filter criteria. Use it to apply the default filter criteria again.
Name	Name of the free disk.
Condition	Condition of the disk.
Enclosure	Enclosure that holds the disk.
State	State of the disk.
Total Size	Displays the total available size on the disk.
Thin	Displays if the disk is a thin reclaimable disk.
# Paths	Displays the number of active paths.
FSS State	Displays whether disks are exported or remote; applicable in a CVM cluster that is enabled for Flexible Storage Sharing (FSS).
Source Host	Displays source host for an exported disk in a CVM cluster that is enabled for FSS.

See "Creating instant volume snapshots" on page 214.

Create Volume Snapshot - Advanced Options

Use this wizard panel to create and configure the advanced options of the snapshot.

This panel displays the following information:

Table 10-12 Create snapshot - advanced options panel options

Field	Description
Name	Name of the original volume.
Size	Size of the volume.
State	State of the original volume.
Snapshot type	Select the snapshot type from the drop-down list.
Configuration	Select to configure the type of snapshot selected.

See "Creating instant volume snapshots" on page 214.

Create Volume Snapshot - Instant Snapshot panel options

Use this wizard panel to configure an instant snapshot.

This panel displays the following information:

Table 10-13 Create snapshot - instant snapshot panel options

Field	Description
Create new snapshot volume	Select to create a new snapshot volume.
Use existing volume	Select to use an existing volume.
Available volumes	Select an available volume from the list to use as an existing volume.
Name	Name of the original volume.
Disk Group	Disk group to which the original volume belongs.
State	State of the original volume.
Layout	Layout type of the volume.
Size	Size of the volume.
Synchronize	Select this option to synchronize the data in the snapshot volume.

See "Creating instant volume snapshots" on page 214.

Creating space optimized snapshots for volumes

The Management Server console lets you configure space-optimized snapshots for volumes. The space-optimized snapshot lets you optimally use the space in the snapshot volume. This type of snapshot is ideal for log volumes.

The operation is not supported on Windows host.

To perform this task, your user group must be assigned the Admin role on the host or the Server perspective. The permission on the host may be explicitly assigned or inherited from a parent Organization.

To create a space optimized snapshot

1. In the Management Server console, go to the **Server** perspective and expand **Manage** in the left pane.
2. Expand the Organization, Applications, or **Uncategorized Hosts** to locate the host.
3. Do one of the following:
   - Click on the **Databases** tab and locate the required database. Right-click on the database and select **Create Snapshot**.
   - Expand the host. Expand **Disk Groups** and select the required disk group. Right-click and select **Create Snapshot**.
   - Expand the host. Expand **Volumes** to locate and select the required volume. Right-click and select **Snapshot** > **Create**.
4. In the **Snapshot level selection panel options** wizard panel, select the **Create volume level snapshot** option. Click **Next**.

   See "Snapshot level selection panel options" on page 270.
5. In the **Create snapshot - snapshot type selection page panel options** wizard panel, select the snapshot type as **Space optimized**. Click **Next**.
6. In the **Create snapshot - disk selection page panel options** wizard panel, select the free disks from the list of disks. To search for disks using one or more filter criteria, click **Edit**. Click **Next**.

   See "Create Volume Snapshot - disk selection page panel options" on page 216.

7 If you want to create a different snapshot type for each volume, then on the **Create snapshot - snapshot type selection page panel options**click the **Advanced** button. In the **Create snapshot - Advanced Options panel options** wizard panel, click **Configuration** to configure the volume details of the Space optimized snapshot selected.

See “Create Volume Snapshot - Advanced Options” on page 216.

8 In the **Create snapshot - Create space optimized snapshot** wizard panel, configure the snapshot volume. Click **Save** to go back to the **Create snapshot - Advanced Options** wizard panel. Click **Next**.

See “Create Volume Snapshot - Space Optimized Snapshot panel options” on page 219.

9 In the **Configure snapshot - Configure snapshots summary** panel, verify the configuration information. Click **Finish**.

10 In the **Result** panel, verify that the snapshots have been configured successfully.

See “About managing Storage Foundation volumes” on page 191.

Create Volume Snapshot - Space Optimized Snapshot panel options

Use this wizard panel to configure space-optimized snapshot for a volume. Using this panel, you can create a storage cache object or use a shared cache to store the original data on the data volume.

Table 10-14 Create snapshot - space optimized snapshot panel options

Field	Description
Selected Volume	Name of the volume for which you want to create a space-optimized snapshot
Snapshot name	Enter a name for the space-optimized snapshot. The name can contain alphanumeric characters up to 32 numbers. This field is optional.
Create new cache object	Select this option to create a storage cache object.
Cache volume size	Enter the size of the storage cache that you want to create. Choose the unit of size from the corresponding drop-down list.

Table 10-14 Create snapshot - space optimized snapshot panel options *(continued)*

Field	Description
Mirrors	Select the number of mirrors that you want create on the storage cache.
Autogrow	Select to enable the autogrow feature for the storage cache. If you select this size of the cache increases when the data in the original volume increases.
Use existing cache object	Select this option if you want to use an existing shared cache to store the original data available on the data volume. This option is available only if the storage cache is available on the same disk group out of which the volume is created.

See "Creating space optimized snapshots for volumes" on page 218.

Creating mirror break-off snapshots for volumes

The Management Server console lets you configure mirror break-off snapshots. The mirror break-off snapshot is created when one or more mirrored disks in a volume is detached and retained as a different volume. You can create a mirror break-off snapshot for a volume only if the volume is snap-ready mirror.

The operation is not supported on Windows host.

To perform this task, your user group must be assigned the Admin role on the host or the Server perspective. The permission on the host may be explicitly assigned or inherited from a parent Organization.

To create a mirror break-off snapshot

1. In the Management Server console, go to the **Server** perspective and expand **Manage** in the left pane.
2. Expand the Organization, Applications, or **Uncategorized Hosts** to locate the host.
3. Do one of the following:
   - Click on the **Databases** tab and locate the required database. Right-click on the database and select **Create Snapshot**.

- Expand the host. Expand **Disk Groups** and select the required disk group. Right-click and select **Create Snapshot**.
- Expand the host. Expand **Volumes** to locate and select the required volume. Right-click and select **Snapshot** > **Create**.

4 In the **Snapshot level selection panel options** wizard panel, select the **Create volume level snapshot** option. Click **Next**.

See "Snapshot level selection panel options" on page 270.

5 In the **Create snapshot - snapshot type selection page panel options** wizard panel, select the snapshot type as **Mirror break-off**. Click **Next**.

6 If you want to create a different snapshot type for each volume, then on the **Create snapshot - snapshot type selection page panel options** click the **Advanced** button. In the **Create snapshot - Advanced Options panel options** wizard panel, click **Configuration** to configure the volume details of the Mirror break-off snapshot selected.

See "Create Volume Snapshot - Advanced Options" on page 216.

7 In the **Create snapshot - Create Mirror Break-off Snapshot** wizard panel, configure the snapshot volume. Click **Save** to go back to the **Create snapshot - Advanced Options** wizard panel. Click **Next**.

See "Create Volume Snapshot - Mirror Break-off Snapshot panel options" on page 221.

8 In the **Configure snapshot - Configure snapshots summary** panel, verify the configuration information. Click **Finish**.

9 In the **Result** panel, verify that the snapshots have been configured successfully.

See "About managing Storage Foundation volumes" on page 191.

Create Volume Snapshot - Mirror Break-off Snapshot panel options

Use this wizard panel to configure a snapshot for a volume.

Table 10-15 Create snapshot - break-off snapshot panel options

Field	Description
Selected volume	Name of the volume for which you want to create a Mirror break-off snapshot

Table 10-15 Create snapshot - break-off snapshot panel options *(continued)*

Field	Description
Snapshot volume name	Enter a name for the Mirror break-off snapshot. The name can contain alphanumeric characters up to 32 numbers. This field is optional.
Selection options	
Specify number of mirrors to break off	Select this option to specify the number of mirrors to break off from the volume. If you select this option, Storage Foundation chooses the mirrors to break off.
Mirrors	Select the number of mirrors from the drop-down list.
Select mirror object to break-off	Select this option to choose the mirrors to break off. Choose the plexes from the list.

See "Creating mirror break-off snapshots for volumes" on page 220.

Dissociating snapshots

The Management Server console lets you dissociate a snapshot and turn it into an independent volume.

In the Veritas InfoScale Operations Manager console, you can perform this operation from the following contexts:

- Data volume
- Snapshot volume

When a data volume is selected, an error page is displayed if the selected data volume's disk group is in a deported or foreign state.

When a snapshot volume is selected, an error page is displayed if the selected snapshot volume's disk group is in a deported or foreign state.

The operation is not supported on Windows host.

To perform this task, your user group must be assigned the Admin role on the host or the Server perspective. The permission on the host may be explicitly assigned or inherited from a parent Organization.

To dissociate a snapshot from a volume from the context of data volume

1. In the Management Server console, go to the **Server** perspective and expand **Manage** in the left pane.
2. Expand the Organization or **Uncategorized Hosts** to locate and select the host.
3. Expand the host and select **Volumes**.
4. Right-click on the required volume and select **Snapshot** > **Dissociate**.
5. The **Dissociate Snapshot** wizard panel displays information about all the snapshots available for the selected data volume. You can select the snapshots to be dissociated. Click **Next**.
6. In the **Dissociate Snapshot Summary** wizard panel, verify the details of the data volume and snapshots selected to be dissociated. Click **Finish**.
7. In the **Result** panel, verify that the selected snapshots have been dissociated successfully.

To dissociate a snapshot from the context of snapshot volume

1. In the Management Server console, go to the **Server** perspective and expand **Manage** in the left pane.
2. Expand the Organization or **Uncategorized Hosts** to locate and select the host.
3. Expand the host and select **Volumes**.
4. Right-click on the required volume and select **Snapshot** > **Dissociate**.
5. In the **Dissociate Snapshot** wizard panel, verify the selected snapshots to be dissociated. Click **OK**.
6. In the **Result** panel, verify that the selected snapshots have been dissociated successfully.

See "Reattaching snapshots" on page 224.

See "About managing Storage Foundation volumes" on page 191.

Dissociate Snapshot panel options

Use this wizard panel to select the data volume and view its snapshot.

Table 10-16 Snapshot Table panel options

Field	Description
Name	Displays the name of the snapshot volume.

Table 10-16 Snapshot Table panel options *(continued)*

Field	Description
Snapshot Type	Displays the type of snapshot (fullsize/ space optimized/ instant).
Parent Name	Displays the name of the parent volume.
Time Created	Displays the time it is created.

See "Dissociating snapshots" on page 222.

Reattaching snapshots

The Management Server console lets you reattach a snapshot to the parent snapshot or data volume in the snapshot hierarchy. This operation is commonly used for third mirror break-of or off-host processing.

The operation is not supported on Windows host.

To perform this task, your user group must be assigned the Admin role on the host or the Server perspective. The permission on the host may be explicitly assigned or inherited from a parent Organization.

To reattach a snapshot to the parent snapshot or data volume

1. In the Management Server console, go to the **Server** perspective and expand **Manage** in the left pane.
2. Expand the Organization or **Uncategorized Hosts** to locate and select the host.
3. Expand the host and select **Volumes**.
4. Right-click on the required volume and select **Snapshot** > **Reattach**.
5. In the **Reattach Snapshot** wizard panel, select the parent snapshot or data volume to be reattached. Click **OK**.

   See "Reattach Snapshot panel options" on page 225.

   If the selected snapshot volume has mirrors, the number of mirrors need to be selected. This option allows the user to select mirrors to be reattached to the parent volume.
6. In the **Result** panel, verify that the selected snapshots have been reattached to the parent snapshot or original data volume successfully.

See "Dissociating snapshots" on page 222.

See "About managing Storage Foundation volumes" on page 191.

Reattach Snapshot panel options

Use this wizard panel to select the parent snapshot or data volume that you want to reattach the snapshot to.

Table 10-17 Reattach Snapshot panel options

Field	Description
Name	Name of parent snapshot/ data volume.
Snapshot Type	Type of snapshot
Host Name	Name of host

If the number of mirrors is more than one, you can select the number of mirrors to be reattached from the drop-down list.

See "Reattaching snapshots" on page 224.

Resizing volumes

The applications in your data center reside on the file systems that are mounted on various volumes. You may want to increase or decrease the size of the volume in circumstances such as:

- The higher usage of the file systems that requires you to increase the size of the volume
- The lower usage of the file systems that requires you to decrease the size of the volume without affecting the data

Storage Foundation lets you increase or decrease the size of an individual volume or a volume in a volume set.

To resize a volume, you can use the free disks available on the disk group of the volume or the free disks available on the host where the volume resides. If the free disks that you choose are not part of the disk group of the volume, the Management Server console adds those disks to the disk group of the volume before using them to complete the resize operation.

You can resize a volume only if it is in the healthy state. The resize operations can be performed only on the volumes that are controlled by Storage Foundation. If a volume has mirrors or linked volumes, the sizes of these associated volumes or mirrors also grow when you increase the size of a volume. When you decrease the size of the volume, the file systems that are mounted on the volume are not affected if the volume has enough space to accommodate the file systems. It is also recommended not to reduce the size of the volume if it contains any unmounted

file system because it can result in the loss of data. You cannot grow a volume if the disk groups on which the volume has been created do not have enough space.

If the volume has a VxFS filesystem on it which is full or almost full, before attempting the resize operation, ensure adequate file system space is available.

To perform this task, your user group must be assigned the Admin role on the host or the Server perspective. The permission on the host may be explicitly assigned or inherited from a parent Organization.

To resize a volume

1. In the Management Server console, go to the **Server** perspective and expand **Manage** in the left pane.
2. Expand the Organization or **Uncategorized Hosts** to locate the host.
3. Expand the host and select **Volumes**.
4. Right-click on the required volume and select **Resize**.
5. In the **Resize volume** wizard panel, specify the required information. Click **Next**. If you have selected the **Specify disks for grow** option, select the disks from the **Disk Selection** wizard panel. To search for disks using one or more filter criteria, click **Edit**. Click **Next**.

   See "Resize volume panel options" on page 226.

   See "Select Disks panel options" on page 153.
6. In the **Resize volume summary** panel, verify the details that you have specified for resizing the volume. Click **Finish**.
7. In the **Result** panel, verify that the selected volume has been resized successfully.

See "Creating Storage Foundation volumes" on page 192.

See "About managing Storage Foundation volumes" on page 191.

Resize volume panel options

Use this wizard panel to specify the details for increasing or decreasing the size of a volume.

This wizard panel displays the following information:

Table 10-18 Resize volume panel options

Field	Description
Volume to resize	Displays the information of the volume that you want to resize, such as the total size, the available size, and the number of the mirrored and the linked volumes.
Disk Group	Displays the usage details of the disk group where the selected volume resides.

Table 10-18 Resize volume panel options *(continued)*

Field	Description
Operation	

Table 10-18 Resize volume panel options *(continued)*

Field	Description
	To increase the size of the volume, select **Grow**, and do one of the following: ■ To increase the size of the volume by a specific size or a specific percentage of the current size, select **Grow By** in the drop-down list and do one of the following: ■ To increase the size of the volume by a specific size, enter the size in the corresponding field and select the unit of measure from the corresponding drop-down list. ■ To increase the size of the volume by a specific percentage of the current size, enter the percentage by which you want to increase the size of the volume in the corresponding field. Select **%** from drop-down that lists the units of measure. ■ To increase the size of the volume to a specific size, select **Grow To** from the drop-down list and enter the size in the corresponding field. Select the unit from the drop-down that lists the units of measure. The available units of measure for this option are Kilobytes (KB), Megabytes (MB), Gigabytes (GB), Terabytes (TB) and Percentage (%). You can select **Max Size** to grow the volume to its maximum size. ■ Select **Specify disks for grow** check box if you want to increase the size of specific disks. This option is available only for increasing the size of the volume. If you select this option, you can select the disks from the **Disk Selection** wizard panel. The unallocated space in the disks in this wizard panel is used to increase the size of the volume. You must not specify a size that is higher than the unallocated space in these disks. See "Select Disks panel options" on page 153. ■ Select **Resize file system** check box to resize a VxFS file system along with the volume. This option is enabled only for mounted VxFS file systems. It is disabled for unmounted VxFS and other native file systems. To decrease the size of the volume, select **Shrink**, and do one of the following: ■ To decrease the size of the volume by a specific size or a specific percentage of the current size of the volume, select **Shrink By** in the drop-down list and do one of the following: ■ To decrease the size of the volume by a specific size, enter the size in the corresponding field and select the unit of measure from the corresponding drop-down list. ■ To decrease the size of the volume to a specific percentage

Table 10-18 Resize volume panel options *(continued)*

Field	Description
	of the current size, enter the percentage by which you want to decrease the size of the volume in the corresponding field. Select **%** from drop-down that lists the units of measure. ■ To decrease the size of the volume to a specific size, select **Shrink To** from the drop-down list and enter the size in the corresponding field. Select the unit from the drop-down that lists the units of measure. The available units of measure for this option are Kilobytes (KB), Megabytes (MB), Gigabytes (GB), Terabytes (TB) and Percentage (%). ■ Select **Force** check box if you want to decrease the size of the volume even if the file systems are unmounted.

See "Resizing volumes" on page 225.

Restoring data from the snapshots of volumes

A volume snapshot is the copy of a volume at a specific point in time. Snapshots of volumes let you back up your data and restore it to the original volume at a later time if required. The Management Server console lets you configure various types of snapshots for the volumes that are controlled by Storage Foundation and restore the data to the original volumes in the event of the loss of data.

To restore the data, the original volume must not be a snapshot of another volume. Also, the original volume must not be in use when you restore the backed up data from its snapshot. Before performing this operation, you must stop any application, such as a database, and unmount any file systems that are configured to use the volume.

You cannot perform this operation on a Windows host.

To perform this task, your user group must be assigned the Admin role on the host or the Server perspective. The permission on the host may be explicitly assigned or inherited from a parent Organization.

Note: Before performing this operation, ensure that the data and the snapshot volumes are not open to any application. If any file systems are mounted on the volumes, you must unmount them to prevent loss of data.

To restore data from the snapshot of a volume

1. In the Management Server console, go to the **Server** perspective and expand **Manage** in the left pane.
2. Expand the Organization or **Uncategorized Hosts** to locate and select the host.
3. Expand the host and select **Volumes**.
4. Right-click on the required volume and select **Snapshot** > **Restore**.
5. In the **Restore Data From Snapshot** wizard panel, select snapshots from the snapshot selection list. Click **Next**.

   See "Restore Data From Snapshot panel options" on page 231.
6. In the **Restore Data From Snapshot Summary** panel, verify the details that you have specified for restoring the volume snapshots. Click **Finish**.
7. In the **Result** panel, verify that the data from the selected snapshot volumes have been restored successfully.

See "Creating instant volume snapshots" on page 214.

See "Creating mirror break-off snapshots for volumes" on page 220.

See "About managing Storage Foundation volumes" on page 191.

Restore Data From Snapshot panel options

Use this wizard panel to specify the details for restoring the data from the snapshot of a volume.

Using this wizard panel, you can restore the snapshots of multiple volumes simultaneously.

To do this, click **Save** after configuring the restore operation for one volume and then configure the next one. If you do not want to reset the list, click **Reset**.

This panel displays the following:

Table 10-19 Restore Data From Snapshot panel options

Field	Description
Name	Name of the snapshot volume.
Snapshot Type	Type of the snapshot configuration, such as instant, space-optimized and so on.
Parent name	Displays the name of the parent volume.

Table 10-19 Restore Data From Snapshot panel options *(continued)*

Field	Description
Time created	Time and date when the snapshot was created
Synchronize	Select this option to synchronize the data in the snapshot volume before restoring it back to the original volume.
Attach snapshot plexes to parent on completion of restore	Select this option to attach the plexes in the snapshot volume to the original volume after restoring the data. This option is not available if the snapshot type is space -optimized.

See "Restoring data from the snapshots of volumes" on page 230.

Refreshing the snapshot of volumes

When you refresh a snapshot of a volume, the snapshot volume is replaced with another point-in-time copy of the original volume.

When you refresh the snapshot of a volume, you can optionally create a space-optimized snapshot of the snapshot volume that you refresh. This action ensures proper data backup. The Management Server console lets you refresh the following types of volume snapshots:

- Instant snapshots
- Break-off mirror snapshots

You cannot perform this operation on a Windows host.

To perform this task, your user group must be assigned the Admin role on the host or the Server perspective. The permission on the host may be explicitly assigned or inherited from a parent Organization.

Note: Before refreshing a volume snapshot, you must unmount the file systems that are configured to use the selected volumes to prevent loss of data.

To refresh the snapshot of a volume

1. In the Management Server console, go to the **Server** perspective and expand **Manage** in the left pane.
2. Expand the Organization or **Uncategorized Hosts** to locate the host.

3 Expand the host and select **Volumes**.

4 Right-click on the required snapshot volume and select **Snapshot** > **Refresh**.

5 In the **Refresh snapshot** wizard panel, specify the required information. Click **Next**.

See "Refresh snapshot panel options" on page 233.

6 In the **Refresh snapshots summary** panel, verify the details that you have specified for refreshing the snapshot. Click **Finish**.

7 In the **Result** panel, verify that the snapshot for the selected volume has been refreshed successfully.

See "Configuring a schedule for volume snapshot refresh" on page 235.

See "Adding snapshot volumes to a refresh schedule" on page 237.

See "Removing the schedule for volume snapshot refresh" on page 238.

See "About managing Storage Foundation volumes" on page 191.

Refresh snapshot panel options

Use this wizard panel to specify the options for refreshing the snapshot of a volume.

Table 10-20 Refresh snapshot panel options

Field	Description
Refresh Options	
Synchronize	Select this check box to synchronize the data in the snapshot volume before refreshing it. This option is enabled by default.
Snapshot retention options	

Table 10-20 Refresh snapshot panel options *(continued)*

Field	Description
Preserve snapshot data (creates space optimized snapshot)	Select this check box to create a space-optimized snapshot for the snapshot volume that you refresh. **Note:** The Veritas InfoScale Operations Manager uses the shared cache object available on the disk group of the volume to create the space-optimized snapshot. The shared cache object for a volume is created when the refresh schedule runs for the first time. The Veritas InfoScale Operations Manager uses this cache object to create a space-optimized snapshot on the subsequent runs of the schedule. The format of the cache volume name is <dg_name>_SHARED_CV1 and that of the cache object is <dg_name>_SHARED_CO1.
Snapshot prefix	Enter a prefix that you want to include in the name of the space-optimized snapshot volume. For example, if you enter the prefix as SO, the name of the snapshot volume will be SO_<Name of the original volume>_<Time when the snapshot volume was created>.
Cache volume size	Enter the size of the cache volume for configuring the space-optimized snapshot. Select the unit of measure from the corresponding drop-down list.
Autogrow	Select this check box to enable the automatic growth of the cache volume when data in the original volume increases.
Do not retain more than	Enter the number of previous snapshots that you want to retain.

See “Refreshing the snapshot of volumes” on page 232.

Configuring a schedule for volume snapshot refresh

The refresh operation for a volume snapshot replaces it with a point-in-time copy of the original volume. The Management Server console lets you schedule the refresh operation for the following types of volume snapshots in your data center:

- Instant snapshots
- Mirror break-off snapshots

Note: Before scheduling the refresh operation for volume snapshot, ensure that the volume is not open to any application. To prevent loss of data, you must unmount the file systems that are configured to use the selected volumes before performing this operation.

You cannot perform this operation on a Windows host.

To perform this task, your user group must be assigned the Admin role on the host or the Server perspective. The permission on the host may be explicitly assigned or inherited from a parent Organization.

To configure a schedule for a volume snapshot refresh

1. In the Management Server console, go to the **Server** perspective and expand **Manage** in the left pane.
2. Expand the Organization or **Uncategorized Hosts** to locate the host.
3. Expand the host and select **Volumes**.
4. Right-click on the required snapshot volume and select **Snapshot** > **Schedule** > **Schedule refresh**.
5. In the **Schedule operation** wizard panel, specify the required information.

   Click **Next**.

   See “Schedule operation panel options” on page 236.
6. In the **Refresh snapshot** wizard panel, specify the required information. Click **Next**.

   See “Refresh snapshot panel options” on page 233.
7. In the **Refresh snapshot summary** panel, verify the details that you have specified for refreshing the snapshot. Click **Finish**.
8. In the **Result** panel, verify that the refresh schedule for the selected volume snapshot has been configured successfully.

See “Refreshing the snapshot of volumes” on page 232.

See “Adding snapshot volumes to a refresh schedule” on page 237.

See “Removing the schedule for volume snapshot refresh” on page 238.

See “About managing Storage Foundation volumes” on page 191.

Schedule operation panel options

Use this wizard panel to specify the details for configuring a schedule for a volume snapshot refresh.

This panel displays the following information:

Table 10-21 Schedule operation options

Field	Description
Schedule Name	Enter a name for the schedule for the volume snapshot refresh operation.
Frequency	Select a frequency for running the snapshot refresh operation. The values under the **When** column changes with the options that you choose here. The available options are daily, weekly, and monthly.
When	Specify the details for scheduling the snapshot refresh operation: ■ For daily schedule: Select the time from the **Time** field. ■ For weekly schedule: Select the **Every weekday** option to refresh the snapshot every day from Monday to Friday. If you want to refresh the snapshot on specific days of the week, select the day from the **Recur every week on** field. Select the time from the **Time** field. ■ For monthly schedule: To refresh the snapshot on a specific day of the month, enter the date in the **Day** field. To refresh the snapshot on the recurring days of a month, choose the required options from the drop-down list. Select the time from the **Time** field.

See “Configuring a schedule for volume snapshot refresh” on page 235.

Adding snapshot volumes to a refresh schedule

To keep the data in the snapshot volume up-to-date, you can perform periodic refreshing of the snapshot volume. The Management Server console lets you create schedules for the volume snapshot refresh operation and also add a volume snapshot to the existing schedule that you have configured for volume snapshot refresh

You cannot perform this operation on a Windows host.

To perform this task, your user group must be assigned the Admin role on the host or the Server perspective. The permission on the host may be explicitly assigned or inherited from a parent Organization.

To add a volume snapshot to a refresh schedule

1. In the Management Server console, go to the **Server** perspective and expand **Manage** in the left pane.
2. Expand the Organization or **Uncategorized Hosts** to locate the host.
3. Expand the host and select **Volumes**.
4. Right-click on the required snapshot volume and select **Snapshot** > **Schedule** > **Add to existing schedule**.
5. In the **Add to existing schedule** wizard panel, select the schedule from the list. Click **Finish**.

   See “Add to existing schedule panel options” on page 237.
6. In the **Result** panel, verify that the selected volume snapshot has been added to the schedule successfully.

See “Refreshing the snapshot of volumes” on page 232.

See “Configuring a schedule for volume snapshot refresh” on page 235.

See “Removing the schedule for volume snapshot refresh” on page 238.

See “About managing Storage Foundation volumes” on page 191.

Add to existing schedule panel options

Use this wizard panel to select an existing refresh schedule to add a volume snapshot to it.

Table 10-22 Add to existing schedule panel options

Field	Description
Snapshot refresh schedule name	Name of the schedule for refreshing a volume snapshot.
Schedule details	Details of the schedule
Volumes	Name of the original volume

See "Adding snapshot volumes to a refresh schedule" on page 237.

Removing the schedule for volume snapshot refresh

To keep the data in the snapshot volume up-to-date, you can perform periodic refreshing of the snapshot volume. The Management Server console lets you schedule the refresh operation for the snapshot of a volume in your data center. You can also remove the volume from the refresh schedule. The console deletes the schedule if the volume snapshot that you remove from an existing schedule is the last one in the schedule configuration.

To remove the schedule for a volume snapshot refresh

1. In the Management Server console, go to the **Server** perspective and expand **Manage** in the left pane.
2. Expand the Organization or **Uncategorized Hosts** to locate the host.
3. Expand the host and select **Volumes**.
4. Right-click on the required snapshot volume and select **Snapshot** > **Refresh** > **Remove volume from refresh schedule**.
5. In the **Remove volume from refresh schedule** wizard panel, select the schedules from the list. Click **Finish**.

   See "Remove volumes from refresh schedule panel options" on page 239.
6. In the **Result** panel, verify that the refresh schedules for the selected volume snapshots have been removed successfully.

See "Refreshing the snapshot of volumes" on page 232.

See "Configuring a schedule for volume snapshot refresh" on page 235.

See "Adding snapshot volumes to a refresh schedule" on page 237.

See "About managing Storage Foundation volumes" on page 191.

Remove volumes from refresh schedule panel options

Use this wizard panel to view the refresh schedule that you want to remove for a volume snapshot.

See "Removing the schedule for volume snapshot refresh" on page 238.

Setting volume usage

The Management Server console lets you set the volume usage to set the read policy for a volume. This operation is enabled only of the volume has mirrors on it.

To perform this task, your user group must be assigned the Admin role on the host or the Server perspective. The permission on the host may be explicitly assigned or inherited from a parent Organization.

To set a volume usage

1. In the Management Server console, go to the **Server** perspective and expand **Manage** in the left pane.
2. Expand the Organization or **Uncategorized Hosts** to locate the host.
3. Expand the host and select **Volumes**.
4. Right-click on the required volume and select **Set Usage**.
5. In the **Set Volume Usage** wizard panel, confirm the volume you want to stop. Click **OK**.

   See "Set Volume Usage panel options" on page 239.
6. In the **Result** panel, verify that the selected volume has been stopped successfully.

See "Creating Storage Foundation volumes" on page 192.

See "About managing Storage Foundation volumes" on page 191.

Set Volume Usage panel options

Select one of the following options to set the read policy for a volume on a UNIX, Linux, or Windows host.

Table 10-23 lists the volume usage options that you can set for UNIX or Linux hosts.

Table 10-24 lists the volume usage options that you can set for Windows hosts.

Table 10-23 Set Volume Usage panel options on UNIX or Linux host

Field	Description
Base on layouts	Select this option to let Storage Foundation select the appropriate read policy for the volume mirrors
Round robin	Select this option to read alternates between the volume mirrors.
Site local read	Select this option to read preferentially from plexes at the locally defined site.
Preferred	Select this option to specify a particular mirror from the list, to be used for reads, wherever possible.

Table 10-24 Set Volume Usage panel options on Windows host

Field	Description
Round robin	Select this option to read alternates between the volume mirrors.
Preferred	Select this option to specify a particular mirror from the list, to be used for reads, wherever possible.

See "Setting volume usage" on page 239.

Splitting snapshots

The Management Server console lets you split the snapshot hierarchy into two independent hierarchies.

To perform this task, your user group must be assigned the Admin role on the host or the Server perspective. The permission on the host may be explicitly assigned or inherited from a parent Organization.

To split a snapshot

1. In the Management Server console, go to the **Server** perspective and expand **Manage** in the left pane.
2. Expand the Organization or **Uncategorized Hosts** to locate and select the host.
3. Expand the host and select **Volumes**.

4 Right-click on the required snapshot volume and select **Snapshot** > **Split**.

5 In the **Split Snapshot** wizard panel, confirm the parent snapshot volume to be split. Click **OK**.

6 In the **Result** panel, verify that the selected snapshots have been split from the parent snapshot or data volume.

See "Dissociating snapshots" on page 222.

See "About managing Storage Foundation volumes" on page 191.

Starting synchronization of snapshots

The Management Server console lets you synchronize the contents of an instant snapshot with the contents of the original volume at the point in time the snapshot was taken. If you want to move a snapshot volume to another disk group for export to another computer for off-host processing, you must ensure that the snapshot volume has been completely synchronized.

Synchronization of the contents of a snapshot with its original volume is not possible for space-optimized snapshots.

If synchronization of a snapshot was not started when the snapshot was created, or only synchronization was paused, you can start synchronization.

To perform this task, your user group must be assigned the Admin role on the host or the Server perspective. The permission on the host may be explicitly assigned or inherited from a parent Organization.

To start synchronization

1 In the Management Server console, go to the **Server** perspective and expand **Manage** in the left pane.

2 Expand the Organization or **Uncategorized Hosts** to locate and select the host.

3 Expand the host and select **Volumes**.

4 Right-click on the required snapshot volume and select **Snapshot** > **Start synchronization**.

5 In the **Start synchronization** wizard panel, confirm the snapshot you want to synchronize. Click **OK**.

6 In the **Result** panel, verify that the selected snapshots synchronized successfully.

See "Dissociating snapshots" on page 222.

See "About managing Storage Foundation volumes" on page 191.

Enabling FastResync on volumes

Using the Management Server console you can enable FastResync on volumes.

To perform this task, your user group must be assigned the Admin role on the host or the Server perspective. The permission on the host may be explicitly assigned or inherited from a parent Organization.

To enable FastResync on volumes

1. In the Management Server console, go to the **Server** perspective and expand **Manage** in the left pane.
2. Expand the Organization or **Uncategorized Hosts** to locate and expand the host.
3. Expand **Volumes** to select a volume.
4. Right-click the required volume and select **Enable FastResync**.
5. In the **Enable FastResync** wizard panel, specify the options, and click **OK**.

   See "Enable FastResync panel options" on page 242.
6. In the **Enable FastResync - Result** panel, verify that FastResync is enabled on the volume successfully.

See "Disabling FastResync on volumes" on page 243.

See "About managing Storage Foundation volumes" on page 191.

Enable FastResync panel options

Use this wizard panel to enable FastResync on volumes.

Table 10-25 Enable FastResync panel options

Field	Description
Automatically Configure FastResync Options	Select to automatically enable FastResync.
Manually Configure FastResync Options	Select to manually enable FastResync.

Table 10-25 Enable FastResync panel options *(continued)*

Field	Description
Enable DRL (Dirty region logging)	DRL expedites the recovery of the volumes after a system crash. Select to either enable or disable DRL. Select **sequential** to configure the sequential DRL. This option is used for volumes such as those that are used for database replay logs for which data are written in a sequential manner. The sequential DRL limits the number of dirty regions in a volume and helps faster recovery of data.
Advanced Options	
FastResync (DCO) mirror	Select a number from the drop-down list to specify the number of data change object (DCO) mirrors that you want to create. When the data volume is updated, these updates are retained in the form of logs on the DCO mirrors.
DCO region size (KB)	Enter the size for the DCO mirror that you want to create. The default value is 64 kilobytes.
Select DCO disks	Select the disks from the list that you want to keep as DCO disks.

See "Enabling FastResync on volumes" on page 242.

Disabling FastResync on volumes

Using the Management Server console you can disable FastResync on one or more volumes.

To perform this task, your user group must be assigned the Admin role on the host or the Server perspective. The permission on the host may be explicitly assigned or inherited from a parent Organization.

To disable FastResync on volumes

1. In the Management Server console, go to the **Server** perspective and expand **Manage** in the left pane.
2. Expand the Organization or **Uncategorized Hosts** to locate and expand the host.
3. Expand **Volumes** to select a volume.
4. Right-click the required volume and select **Disable FastResync**.
5. In the **Disable FastResync** wizard panel, verify the details, and click **OK**.
6. In the **Disable FastResync - Result** panel, verify that FastResync is disabled on the volume successfully.

See “Enabling FastResync on volumes” on page 242.

See “About managing Storage Foundation volumes” on page 191.

Chapter 11

Managing file systems

This chapter includes the following topics:

- About managing file systems
- Creating file systems
- Enabling change logs
- Disabling change logs
- Synchronizing change logs
- Removing change logs
- Defragmenting file systems
- Unmounting non clustered file systems from hosts
- Mounting non clustered file systems on hosts
- Unmounting clustered file systems
- Mounting clustered file systems on hosts
- Remounting file systems
- Checking file systems
- Creating file system snapshots
- Remounting file system snapshot
- Mounting file system snapshot
- Unmounting file system snapshot
- Removing file system snapshot

- Monitoring capacity of file systems

About managing file systems

Following is a list of operations related to file systems that you can perform in the Management Server console.

See “Creating file systems” on page 246.

See “Enabling change logs” on page 254.

See “Disabling change logs” on page 255.

See “Synchronizing change logs” on page 255.

See “Removing change logs” on page 256.

See “Defragmenting file systems” on page 257.

See “Unmounting non clustered file systems from hosts” on page 257.

See “Mounting non clustered file systems on hosts” on page 259.

See “Unmounting clustered file systems” on page 262.

See “Mounting clustered file systems on hosts” on page 264.

See “Remounting file systems” on page 265.

See “Checking file systems” on page 268.

See “Creating file system snapshots” on page 269.

See “Remounting file system snapshot” on page 271.

See “Mounting file system snapshot” on page 273.

See “Unmounting file system snapshot” on page 275.

See “Removing file system snapshot” on page 276.

See “About performing Storage Foundation and replicator operations” on page 133.

See “Permissions required for views and operations on CVM cluster objects” on page 591.

Creating file systems

The Management Server console lets you create and mount the file system on a Storage Foundation volume.

To perform this task, your user group must be assigned the Admin role on the host or the Server perspective. The permission on the host may be explicitly assigned or inherited from a parent Organization.

To create a file system

1. In the Management Server console, go to the **Server** perspective and expand **Manage** in the left pane.
2. Expand Organization or **Uncategorized Hosts** to locate and select the host.
3. Do one of the following:
   - Right-click on the host and select **Create File System**.
   - Expand the host and select **Volumes**. Right-click and select **Create File System**.
   - Expand the host. Expand **Volumes** under the host to locate and select a volume. Right-click and select **File System** > **Create File System**.
4. If you did not select a volume to launch the create file system wizard, the **Select Volume** wizard panel is displayed. Select the volume to create the file system. Click **Next**.

   See "Select Volume panel options" on page 250.
5. In the **Create a new File system** panel, select the file system options and the mount options. Click **Next**.

   See "Create File System - File System Options" on page 247.

   See "Advanced Options panel" on page 251.
6. In the **Create File System Summary** panel, verify your selections to create the file system. Click **Finish** to create the file system.

See "Remounting file systems" on page 265.

See "Checking file systems" on page 268.

See "Creating file system snapshots" on page 269.

See "Mounting file system snapshot" on page 273.

See "About managing file systems" on page 246.

Create File System - File System Options

Use this wizard panel to create a new file system.

Select one of the following file systems type to create a new file system. The file system type may vary according to the platform of the host.

- vxfs: available on Unix / Linux if VRTSvxfs package is installed

- ufs: available on SunOS
- ext2: available on Linux
- ext3: available on Linux
- ext4: available on Linux

Table 11-1 File System Options for vxfs, ext2, ext3, ext4 and ufs

Field	Description
Disk Group	Displays disk group name on which selected volume resides.
Volume	Displays the name of the volume used.
File System Type	Select a file system type from the drop-down list.
File System options for vxfs	
Block size	Select the block size from the following options: ■ Default ■ 1024 ■ 2048 ■ 4096 ■ 8192 All units are specified in bytes.
Size	Enter the file system size.
File System options for ext2, ext3, ext4 and ufs	
Block size	Select the block size from the following options: ■ Default ■ 1024 ■ 2048 ■ 4096 ■ 8192 **Note:** For ext file systems, block size 8192 is not allowed. All units are specified in bytes.

Table 11-1 File System Options for vxfs, ext2, ext3, ext4 and ufs *(continued)*

Field	Description
Allocation unit size	Select the allocation unit from the following options: ■ Default ■ 1024 ■ 2048 ■ 4096 ■ 8192 **Note:** For ext file systems, allocation unit size 8192 is not allowed. All units are specified in bytes.

Use the Mount Options wizard panel to set options when mounting a file system.

If the file system has an entry in the file system table, check Mount using options in file system table to use the stored mount options. You cannot use this wizard panel to change the options. To change options, you must remount the file system.

Table 11-2 Mount Options

Field	Description
Mount point	Enter the mount point for the file system. The mount point must be specified, and must be an absolute pathname, that is, it must begin with /.
Read only	Select this option to mount the file system in read-only mode.
Honor setuid	Select this option to allow setuid requests.
Select Caching Mode	For a VxFS cache area, the caching mode determines what kind of caching is performed for the specified mount point. The mode can be avoid cache, read, or writeback. The default mode is read. A VxVM cache area only supports read mode. The panel displays available cache areas, based on the mode you select.

Table 11-2 Mount Options *(continued)*

Field	Description
Add to file system table	Select this option to automatically update the file system table when the file system is mounted.
Mount at boot	This option is available only if you selected **Add to file system table**. Select this option to mount the file system at boot time.
fsck pass	This option is available only if you selected **Add to file system table**. Specify the fsck pass when the file system is mounted at boot time.
Mount Type	Select one of the following mount types which are available for selection only if the disk group is shared: ■ Local ■ Cluster
Select nodes where file system should be mounted	Select one or more nodes where you want to mount the file system. This option is enabled only when the **Mount Type** selected is **Cluster**.
Advanced	Select to specify the advanced options for mounting the file system. See "Advanced Options panel" on page 251.

See "Creating file systems" on page 246.

See "Creating file systems" on page 246.

See "Creating Storage Foundation volumes" on page 192.

Select Volume panel options

Use this wizard panel to select a volume to create a new file system. The panel displays the details of the volumes on which you can create a file system.

Table 11-3 Select volume panel options

Field	Description
Name	Displays the name of the volume.

Table 11-3 Select volume panel options *(continued)*

Field	Description
Condition	Displays the condition of the volume.
Disk Group	Displays the name of the disk group.
Layout	Displays the layout type of the volume.
Size	Displays the size of the volume.
Snapshot Of	Displays the snapshot of the volume.
# Mirrors	Displays the mirrors associated with the volume.
# Volume Snapshots	Displays the count of snapshots.

See "Creating file systems" on page 246.

Advanced Options panel

Use this wizard panel to specify the advanced file system and mount options. The options available depend on the type of file system.

Table 11-4 Advanced options for vxfs, ufs, ext2, ext3, and ext4

Field	Description
Advanced options for vxfs	
Log size (blocks)	Specify the activity logging area size by providing a value for the log size.
Extra Options	Use this area to define other mount options. See the mkfs(1) manual pages. For example: version=8. Separate by comma(,) for specifying multiple values.
Advanced options for ufs	
Label	Define a label for the file system.
Extra options	Use this area to define other mount options. See the mkfs(1) manual pages. For example: mtb=y. Separate by comma(,) for specifying multiple values.

Table 11-4 Advanced options for vxfs, ufs, ext2, ext3, and ext4 *(continued)*

Field	Description
Advanced options for ext2	
Extra options	Use this area to define other mount options. See the mkfs(1) manual pages. For example: `creator-os`. Separate by comma(,) for specifying multiple values.
Advanced options for ext3	
External log device	Enter the name / LABEL of log (journal) block device.
Log size (blocks)	Specify the activity logging area size by providing a value for the log size.
Extra options	Use this area to specify additional parameters to create file systems. See the mkfs(1) manual pages. For example: `creator-os`. Separate by comma(,) for specifying multiple values.
Advanced options for ext4	
Extra options	Use this area to specify additional parameters to create file systems. See the mkfs(1) manual pages. For example: `creator-os`. Separate by comma(,) for specifying multiple values.

Table 11-5 Advanced mount options for vxfs, ufs, ext2, ext3, and ext4

Field	Description
Advanced mount options for vxfs	

Table 11-5 Advanced mount options for vxfs, ufs, ext2, ext3, and ext4 *(continued)*

Field	Description
File system caching policy	Select one of the following: ■ default ■ direct ■ dsync ■ closesync ■ tempcache ■ unbuffered See the mount(1) manual pages.
Convert osync policy	Select one of the following: ■ default ■ direct ■ dsync ■ closesync ■ delay ■ unbuffered See the mount(1) manual pages.
I/O handling policy	Select one of the following: ■ default ■ mwdisable ■ wdisable ■ nodisable ■ mdisable ■ disable See the mount(1) manual pages.
Clear data extents	Select the check box to clear all data extents before file allocation.
Disable access time recording	Select the check box to disable access time recording.
Extra options	Use this area to define other mount options. See the mount(1) manual pages. For example: logiosize=4096. Separate by comma(,) for specifying multiple values.

Table 11-5 Advanced mount options for vxfs, ufs, ext2, ext3, and ext4 *(continued)*

Field	Description
Advanced mount options for ufs, ext2, ext3, and ext4	
Extra options	Use this area to define other mount options. See the mount(1) manual pages. For example: for ufs: quota. For ext2, ext3, and ext4: async. Separate by comma(,) for specifying multiple values.

See "Create File System - File System Options" on page 247.

See "Creating file systems" on page 246.

Enabling change logs

The Management Server console lets you enable a file change log for a mounted VxFS file system on a managed host.

To perform this task, your user group must be assigned the Admin role on the host or the Server perspective. The permission on the host may be explicitly assigned or inherited from a parent Organization.

To enable a file change log

1. In the Management Server console, go to the **Server** perspective and expand **Manage** in the left pane.
2. Expand the Organization or **Uncategorized Hosts** to locate the host.
3. Expand the host and then expand **Volumes**.
4. Right-click on the required volume and select **File System** > **Enable Change Log**.
5. In the **Enable Change Log** wizard panel, confirm the file system for which you want to enable the file change log. Click **OK**.
6. In the **Result** panel, verify that the file change log has been successfully enabled for the selected file system.

See "Disabling change logs" on page 255.

See "About managing file systems" on page 246.

Disabling change logs

The Management Server console lets you disable a file change log for a mounted VxFS file system on a managed host.

To perform this task, your user group must be assigned the Admin role on the host or the Server perspective. The permission on the host may be explicitly assigned or inherited from a parent Organization.

To disable a file change log

1. In the Management Server console, go to the **Server** perspective and expand **Manage** in the left pane.
2. Expand the Organization or **Uncategorized Hosts** to locate the host.
3. Expand the host and then expand **Volumes**.
4. Right-click on the required volume and select **File System** > **Disable Change Log**.
5. In the **Disable Change Log** wizard panel, confirm the file system for which you want to disable the file change log. Click **OK**.
6. In the **Result** panel, verify that the file change log has been successfully disabled for the selected file system.

See "Enabling change logs" on page 254.

See "About managing file systems" on page 246.

Synchronizing change logs

The Management Server console lets you bring the file change log to a stable state by flushing the associated data of a file change log recording interval.

You cannot synchronize a change log if the file system is not a VxFS file system or if it is not mounted.

To perform this task, your user group must be assigned the Admin role on the host or the Server perspective. The permission on the host may be explicitly assigned or inherited from a parent Organization.

To synchronize a file change log

1. In the Management Server console, go to the **Server** perspective and expand **Manage** in the left pane.
2. Expand the Organization or **Uncategorized Hosts** to locate the host.
3. Expand the host and then expand **Volumes**.

4 Right-click on the required volume and select **File System** > **Synchronize Change Log**.

5 In the **Synchronize Change Log** wizard panel, confirm the file system for which you want to synchronize the file change log. Click **OK**.

6 In the **Result** panel, verify that the file change log has been successfully synchronized for the selected file system.

See "Enabling change logs" on page 254.

See "Removing change logs" on page 256.

See "About managing file systems" on page 246.

Removing change logs

The Management Server console lets you remove a file change log for a mounted VxFS file system on a managed host.

The selected file change log must be disabled before removing the file change log file.

To perform this task, your user group must be assigned the Admin role on the host or the Server perspective. The permission on the host may be explicitly assigned or inherited from a parent Organization.

To remove a file change log

1 In the Management Server console, go to the **Server** perspective and expand **Manage** in the left pane.

2 Expand the Organization or **Uncategorized Hosts** to locate the host.

3 Expand the host and select **Volumes**.

4 Right-click on the required volume and select **File System** > **Remove Change Log**.

5 In the **Remove Change Log** wizard panel, confirm the file system for which you want to remove the file change log. Click **OK**.

6 In the **Result** panel, verify that the file change log has been successfully removed for the selected file system.

See "Enabling change logs" on page 254.

See "Synchronizing change logs" on page 255.

See "About managing file systems" on page 246.

Defragmenting file systems

The Management Server console lets you defragment mounted VxFS file systems. When a file system is fragmented, it gets full of un-contiguous locations on the file system which in turn affects performance while reading the data. Defragmenting file systems helps in seeking data quickly. This operation is available only for VxFS file systems.

To perform this task, your user group must be assigned the Admin role on the host or the Server perspective. The permission on the host may be explicitly assigned or inherited from a parent Organization.

To defrag a file system

1. In the Management Server console, go to the **Server** perspective and expand **Manage** in the left pane.
2. Expand the Organization or **Uncategorized Hosts** to locate the host.
3. Expand the host and select **Volumes**.
4. Right-click on the required volume and select **File System** > **Defrag**.
5. In the **Defrag File System** wizard panel, confirm the file system you want to defragment. Click **OK**.

   See "Defrag file system panel options" on page 257.

   This operation may take several minutes, depending on the size of the file system.
6. In the **Result** panel, verify that the command for defragmentation is running successfully and a warning message is displayed.

See "About managing file systems" on page 246.

See "Creating file systems" on page 246.

Defrag file system panel options

Use this wizard panel to confirm and defrag the selected file system. You can define the maximum time you want to run the defrag. If you do specify the maximum time, then the operation is aborted once the time defined is over.

See "Defragmenting file systems" on page 257.

Unmounting non clustered file systems from hosts

The Management Server console lets you unmount a file system that has been mounted on a Storage Foundation volume in your data center. When you unmount

a non-clustered file system from a host, you can optionally remove the entry of the selected file system from the file system table

You cannot unmount a file system if the file system:

- Resides on a host that runs on Windows.
- Has one or more checkpoints that have been mounted. For these types of file systems, you must first unmount the checkpoints before unmounting the file system from the host.

The console lets you unmount the following file systems from a host:

Operating system for the host	Supported File Systems
Solaris	vxfs, ufs
Linux	vxfs, ext2, ext3, and ext4
HP-UX	vxfs
AIX	vxfs

To perform this task, your user group must be assigned the Admin role on the host or the Server perspective. The permission on the host may be explicitly assigned or inherited from a parent Organization.

To unmount a non clustered file system

1. In the Management Server console, go to the **Server** perspective and expand **Manage** in the left pane.
2. Expand the Organization or **Uncategorized Hosts** to locate and select the host.
3. Expand the host and select **Volumes**.
4. Right-click on the required volume and select **File System** > **Unmount**.
5. In the **Unmount Confirmation** wizard panel, click **OK**.

   See “Unmount Confirmation panel” on page 259.
6. In the **Result** page, verify that the selected non-clustered file system has been unmounted successfully.

See “Unmounting clustered file systems” on page 262.

See “Mounting non clustered file systems on hosts” on page 259.

See “About managing file systems” on page 246.

Unmount Confirmation panel

Use this wizard panel to confirm a file system unmount operation on a host.

Clear the **Remove from fstab?** check box to remove the entry for the selected file system in the file system table.

See "Unmounting non clustered file systems from hosts" on page 257.

Mounting non clustered file systems on hosts

For using a file system, you must mount it on a mount point, that is a directory. The Management Server console lets you mount a file system to the mount point on a host.

To mount a file system, the file system must be in an unmounted state.

Operating system for the host	Supported File Systems
Solaris	UNIX file system (UFS)
Linux	Ext2, Ext3, and Ext4

To perform this task, your user group must be assigned the Admin role on the host or the Server perspective. The permission on the host may be explicitly assigned or inherited from a parent Organization.

To mount a non clustered file system

1. In the Management Server console, go to the **Server** perspective and expand **Manage** in the left pane.
2. Expand the Organization or **Uncategorized Hosts** to locate and select the host.
3. Expand the host and select **Volumes**.
4. Right-click on the required volume and select **File System** > **Mount**.
5. In the **Mount Options** wizard panel, specify the required information. Click **OK**.

   See "Mount Options panel" on page 260.
6. In the **Result** page, verify that the selected non-clustered file system has been mounted successfully.

See "Mounting clustered file systems on hosts" on page 264.

See "Unmounting non clustered file systems from hosts" on page 257.

See "About managing file systems" on page 246.

Mount Options panel

Use this wizard panel to specify the options for mounting a non-clustered file system on a host.

Table 11-6 Mount Options

Field	Description
Mount using existing options (from file system table)	Select if the file system is in unmounted state and has an file system table entry.
Mount point	Enter the mount point for the file system. The mount point must be specified, and must be an absolute pathname, that is, it must begin with /.
File System Type	Select a file system type from the drop-down list.
Read only	Select this option to mount the file system in read-only mode.
Honor setuid	Select this option to allow setuid requests.
Select Caching Mode	For a VxFS cache area, the caching mode determines what kind of caching is performed for the specified mount point. The mode can be avoid cache, read, or writeback. The default mode is read. A VxVM cache area only supports read mode. The panel displays available cache areas, based on the mode you select.
Add to file system table	Select this option to automatically update the file system table when the file system is mounted.
Mount at boot	This option is available only if you selected **Add to file system table**. Select this option to mount the file system at boot time.
fsck pass	This option is available only if you selected **Add to file system table**. Specify the fsck pass when the file system is mounted at boot time.
Mount Type	This option is displayed only if the selected volume is a shared volume.

Table 11-6 Mount Options *(continued)*

Field	Description
Advanced	Select to specify the advanced options for mounting the file system. See "Advanced Options panel" on page 251.

See "Mounting non clustered file systems on hosts" on page 259.

Advanced Mount Options panel options

Use this wizard panel to specify new advanced options for mounting a file system.

Click **OK** to save your changes and go back to the **Mount Options** wizard panel.

Table 11-7 Advanced mount options

Field	Description
File system caching policy	Select one of the following: ■ default ■ direct ■ dsync ■ closesync ■ tempcache ■ unbuffered See the mount(1) manual pages.
Convert osync policy	Select one of the following: ■ default ■ direct ■ dsync ■ closesync ■ delay ■ unbuffered See the mount(1) manual pages.

Table 11-7 Advanced mount options *(continued)*

Field	Description
I/O handling policy	Select one of the following: ■ default ■ mwdisable ■ wdisable ■ nodisable ■ mdisable ■ disable See the mount(1) manual pages.
Clear data extents	Select the check box to clear all data extents before file allocation.
Disable access time recording	Select the check box to disable access time recording.
Extra options	Use this area to define other mount options. See the mount(1) manual pages. For example: logiosize=4096. Separate by comma(,) for specifying multiple values.
For ufs, ext2, ext3, and ext4 Not applicable for cluster file system.	
Extra options	Use this area to define other mount options. See the mount(1) manual pages. For example: for ufs: quote. For ext2, ext3, and ext4: async. Separate by comma(,) for specifying multiple values.

See "Mounting non clustered file systems on hosts" on page 259.

See "Mounting clustered file systems on hosts" on page 264.

Unmounting clustered file systems

The Management Server console lets you unmount a clustered file system that has been mounted on one or more nodes of a cluster in your data center.

You cannot unmount a file system if the file system:

- Resides on a host that runs on Windows.

- Has one or more checkpoints that have been mounted. For these types of file systems, you must first unmount the checkpoints before unmounting the file system from the host.

To perform this task, your user group must be assigned the Admin role on the host or the Server perspective. The permission on the host may be explicitly assigned or inherited from a parent Organization.

To unmount a clustered file system

1. In the Management Server console, go to the **Server** perspective and expand **Manage** in the left pane.
2. Expand the Organization or **Uncategorized Hosts** to locate and select the host.
3. Expand the host and select **Volumes**.
4. Right-click on the required volume and select **File System** > **Unmount**.
5. In the **Unmount Cluster File System** wizard panel, specify the required information. Click **OK**.

   See "Unmount Clustered File System panel options" on page 263.
6. In the **Result** page, verify that the selected clustered file system has been unmounted.

See "Unmounting non clustered file systems from hosts" on page 257.

See "Mounting clustered file systems on hosts" on page 264.

See "About managing file systems" on page 246.

Unmount Clustered File System panel options

Use this wizard panel to select the nodes to specify where you want to unmount a clustered file system from.

To unmount the file system from specific nodes, select **Unmount from selected nodes**. Select the required nodes in the **Select nodes to unmount the file system** list.

To unmount the clustered file system from all the nodes and remove the file system from the cluster configuration, select **Remove mount from cluster configuration**. If you select this option, you cannot mount this file system anywhere in the cluster.

See "Unmounting clustered file systems" on page 262.

Mounting clustered file systems on hosts

For using a clustered file system, you must mount it on a mount point of one of the nodes in a cluster. The Management Server console lets you mount a clustered file system to the mount point on a host.

To mount a file system, the file system must be in an unmounted state.

To perform this task, your user group must be assigned the Admin role on the host or the Server perspective. The permission on the host may be explicitly assigned or inherited from a parent Organization.

To mount a clustered file system on a host

1. In the Management Server console, go to the **Server** perspective and expand **Manage** in the left pane.
2. Expand the Organization or **Uncategorized Hosts** to locate the host.
3. Expand the host and select **Volumes**.
4. Right-click on the required volume and select **File System** > **Mount**.
5. In the **Clustered Mount Options** wizard panel, specify the required information. Click **OK**.

   See “Mount File System - Clustered Mount Options” on page 264.
6. In the **Result** page, verify that the selected clustered file system has been mounted successfully.

See “Mounting non clustered file systems on hosts” on page 259.

See “Unmounting clustered file systems” on page 262.

See “About managing file systems” on page 246.

Mount File System - Clustered Mount Options

Use this wizard panel to specify the options for mounting a clustered file system on a host.

The options that you specify in this wizard panel are applied to all the nodes on the cluster.

Table 11-8 Mount File System - Clustered Mount Options

Field	Description
Mount point	Enter the mount point for the file system. The mount point must be specified, and must be an absolute pathname, that is, it must begin with /.
Read only	Select this option to mount the file system in read-only mode.
Honor setuid	Select this option to allow setuid requests.
Select Caching Mode	For a VxFS cache area, the caching mode determines what kind of caching is performed for the specified mount point. The mode can be avoid cache, read, or writeback. The default mode is read. A VxVM cache area only supports read mode. The panel displays available cache areas, based on the mode you select.
Mount Type	Select a cluster.
Select nodes to mount the file system	Select one or more nodes for mounting the file system.
Advanced	Select to specify the advanced options for mounting the file system. See "Advanced Mount Options panel options" on page 261.

See "Mounting clustered file systems on hosts" on page 264.

Remounting file systems

The Management Server console lets you remount an already mounted file systems. This operation is used when the user needs to change mount options while keeping the same mount point.

This operation cannot be performed on a Windows host.

To perform this task, your user group must be assigned the Admin role on the host or the Server perspective. The permission on the host may be explicitly assigned or inherited from a parent Organization.

To remount a file system

1. In the Management Server console, go to the **Server** perspective and expand **Manage** in the left pane.
2. Expand the Organization or **Uncategorized Hosts** to locate the host.
3. Expand the host and select **Volumes**.
4. Right-click on the required volume and select **File System** > **Remount File System**.
5. In the **Remount File System** wizard panel, confirm the file system you want to remount. Click **OK.**

   See "Remount Options panel" on page 266.
6. In the **Result** panel, verify that the file system selected has been remounted successfully.

See "Mounting non clustered file systems on hosts" on page 259.

See "Mounting clustered file systems on hosts" on page 264.

See "About managing file systems" on page 246.

Remount Options panel

Use this wizard panel to remount an already mounted file system.

Table 11-9 Remount Options panel

Field	Description
Read write	Check this check box to change the I/O mode from read only to write mode. **Note:** I/O mode cannot be changed from read write mode to read-only mode during the remount operation.
Disable access time recording	Check this check box to disable file access time updates during remount. This option is available only with VxFS file systems.

Table 11-9 Remount Options panel *(continued)*

Field	Description
Select Caching Mode	For a VxFS cache area, the caching mode determines what kind of caching is performed for the specified mount point. The mode can be avoid cache, read, or writeback. The default mode is read. A VxVM cache area only supports read mode. The panel displays available cache areas, based on the mode you select.
I/O handling policy:	Select one of the following: ■ Default ■ disable ■ nodisable ■ wdisable ■ mwdisable
Convert osync policy	Select one of the following: ■ Default ■ closesync ■ delay ■ direct ■ dsync ■ unbuffered
Extra options	Enter the options that you want to specify in addition to the options that you have specified earlier. For example: logiosize=4096. Separate by comma(,) for specifying multiple values.
Node	Select one or more nodes for remounting the file system. This option is displayed only for clustered file system.

See "Remounting file systems" on page 265.

Checking file systems

The Management Server console lets you check an unmounted file system for consistency.

To perform this task, your user group must be assigned the Admin role on the host or the Server perspective. The permission on the host may be explicitly assigned or inherited from a parent Organization.

To check a file system

1. In the Management Server console, go to the **Server** perspective and expand **Manage** in the left pane.
2. Expand the Organization or **Uncategorized Hosts** to locate the host.
3. Expand the host and select **Volumes**.
4. Right-click on the required volume and select **File System** > **Check**.
5. In the **Check File System** wizard panel, specify the required information. Click **OK.**

   See "Check File System panel options" on page 268.

   This operation may take several minutes, depending on the size of the file system and the number of files in the file system.
6. In the **Result** panel, verify the file system that the file system selected has been checked.

See "Creating file systems" on page 246.

See "About managing file systems" on page 246.

Check File System panel options

Use this wizard panel to check file system.

Table 11-10 Check File System panel options

Field	Description
File System Type	Select a file system type from the drop-down list. This option is not displayed if you select an unmounted file system.
Run full check	Select to run a full check on the file system. This option is available only with VxFS file systems.

See “Checking file systems” on page 268.

Creating file system snapshots

The Management Server console lets you create a file system snapshot.

To perform this task, your user group must be assigned the Admin role on the host or the Server perspective. The permission on the host may be explicitly assigned or inherited from a parent Organization.

To create a file system snapshot

1. In the Management Server console, go to the **Server** perspective and expand **Manage** in the left pane.
2. Do one of the following:
   - Expand the Organization or **Uncategorized Hosts** to select the host. Expand **Volumes**. Right-click required volume and select **Snapshot** > **Create**.
   - Expand the Organization or **Uncategorized Hosts** to select the host. Expand **Disk Groups**. Right-click required disk group and select **Create Snapshot**.
   - Expand **Applications** and expand **Databases**. Right-click the database and select **Create Snapshot**.
3. In the **Snapshot level selection panel options** wizard panel, select the **Create file system snapshot** option. Click **Next**.

   See “Snapshot level selection panel options” on page 270.
4. In the **Create File System snapshot - Advanced Options** wizard panel, specify the mount details. Click **OK**.

   See “Create File System snapshot - Configure Options” on page 270.
5. In the **Result** panel, verify that the file system snapshot has been created successfully.

See “Remounting file system snapshot” on page 271.

See “Unmounting file system snapshot” on page 275.

See “About managing file systems” on page 246.

Create Snapshot - Create File System Snapshot panel options

Use this wizard panel to create and configure file system snapshots.

This panel displays the following information:

Table 11-11 Create snapshot - Create File System Snapshot panel options

Field	Description
File System Name	Displays the name of the selected file system on which the new checkpoint needs to be created.
File system snapshot Name	Specify a file system snapshot name. A valid file system snapshot name must contain only alpha-numeric characters, underscores, dashes, or periods. The length of the name must be between 1 and 19 characters long. It cannot begin with a dash or a period, or end with a period.
Mount	Select this check box if you want to mount the file system snapshot.
Configuration	Select to configure the file system snapshot.

See "Creating file system snapshots" on page 269.

Snapshot level selection panel options

Use this wizard panel to choose the method of creating volume snapshots or file system snapshots with checkpoints.

Table 11-12

Field	Description
Create volume snapshot	Select to create a volume snapshot.
Create file system snapshot	Select to create a file system snapshot.

See "Creating instant volume snapshots" on page 214.

See "Creating space optimized snapshots for volumes" on page 218.

See "Creating mirror break-off snapshots for volumes" on page 220.

See "Creating file system snapshots" on page 269.

Create File System snapshot - Configure Options

Use this wizard panel to create a file system snapshot.

Table 11-13 Create file system snapshot panel options

Field	Description
File system	Displays the name of the selected file system on which the new file system snapshot needs to be created.
File system snapshot name	Specify a checkpoint name. A valid file system snapshot name must contain only alpha-numeric characters, underscores, dashes, or periods. The length of the name must be between 1 and 19 characters long. It cannot begin with a dash or a period, or end with a period.
Removable	Select this check box to create a file system of the type, Removable.
Mount	Select this check box if you want the file system snapshot mounted.
Mount point	If you select **Mount**, then you need to specify a mount point.
Read only	Select this check box to mount the file system snapshot as read-only.
Node	Select one or more nodes where you want to mount the file system snapshot. This option is enabled only for cluster file system.

See "Creating file system snapshots" on page 269.

Remounting file system snapshot

The Management Server console lets you remount an already mounted file system snapshot in a VxFS file system.

To perform this task, your user group must be assigned the Admin role on the host or the Server perspective. The permission on the host may be explicitly assigned or inherited from a parent Organization.

To remount a file system snapshot

1. In the Management Server console, go to the **Server** perspective and expand **Manage** in the left pane.
2. Expand the Organization or **Uncategorized Hosts** to locate and select the host.
3. Expand the host and do one of the following:
   - Expand **Volumes** to locate and select the required volume. Click **Snapshots** tab.
   - Expand **Volumes** to locate and select the required file system snapshot.
4. Right-click on the selected file system snapshot and select **Remount**.
5. In the **Remount File System Snapshot** wizard panel, specify remount details. Click **OK**.

   See "Remount File System Snapshot panel options" on page 272.
6. In the **Result** panel, verify that the checkpoint has been remounted successfully.

See "Creating file system snapshots" on page 269.

See "Mounting file system snapshot" on page 273.

See "Unmounting file system snapshot" on page 275.

See "About managing file systems" on page 246.

Remount File System Snapshot panel options

Use this wizard panel to remount a file system snapshot.

Table 11-14 Remount File System Snapshot panel options

Field	Description
Read Write	Check this check box to change the I/O mode from read only to write mode. **Note:** I/O mode cannot be changed from read write mode to read-only mode during the remount operation.
Disable access time recording	Check this check box to disable file access time updates during remount.

Table 11-14 Remount File System Snapshot panel options *(continued)*

Field	Description
I/O handling policy	Select one of the following: ■ Default ■ disable ■ nodisable ■ wdisable ■ mwdisable
Convert osync policy	Select one of the following: ■ Default ■ closesync ■ delay ■ direct ■ dsync ■ unbuffered
Extra options	Enter the options that you want to specify in addition to the options that you have specified earlier. For example: logiosize=4096. Separate by comma(,) for specifying multiple values.
Node	Select one or more nodes where you want to remount the file system snapshot. This option is enabled only for cluster file system.

See "Remounting file system snapshot" on page 271.

Mounting file system snapshot

The Management Server console lets you mount an existing file system snapshot. Mounted file system snapshot are shown as children of the file system they are based on.

To perform this task, your user group must be assigned the Admin role on the host or the Server perspective. The permission on the host may be explicitly assigned or inherited from a parent Organization.

To mount a file system snapshot

1. In the Management Server console, go to the **Server** perspective and expand **Manage** in the left pane.
2. Expand the Organization or **Uncategorized Hosts** to locate and select the host.
3. Expand the host. Expand **Volumes** to locate and select the required volume. Click on **Snapshots**.
4. Right-click on the selected file system snapshot and select **Mount**.
5. In the **Mount File System Snapshot** wizard panel, specify mount details. Click **OK**.

   See "Mount File System Snapshot panel options" on page 274.
6. In the **Result** panel, verify that the checkpoint has been mounted successfully.

See "Creating file system snapshots" on page 269.

See "Remounting file system snapshot" on page 271.

See "Unmounting file system snapshot" on page 275.

See "About managing file systems" on page 246.

Mount File System Snapshot panel options

Use this wizard panel to mount a file system snapshot.

Table 11-15 Mount File System Snapshot panel options

Field	Description
File system	Displays the name of the selected file system on which the new file system snapshot needs to be mounted.
Snapshot Name	Displays the name of the file system snapshot.
Mount Point	If you select **Mount**, then you need to specify a mount point. A default mount point path is generated that can be edited.
Read only	Select this check box to mount the file system snapshot as read-only.

Table 11-15 Mount File System Snapshot panel options *(continued)*

Field	Description
Node	Select one or more nodes where you want to mount the file system snapshot. This option is enabled only for cluster file system.

See "Mounting file system snapshot" on page 273.

Unmounting file system snapshot

The Management Server console lets you unmount an existing file system snapshot.

To perform this task, your user group must be assigned the Admin role on the host or the Server perspective. The permission on the host may be explicitly assigned or inherited from a parent Organization.

To unmount a file system snapshot

1. In the Management Server console, go to the **Server** perspective and expand **Manage** in the left pane.
2. Expand the Organization or **Uncategorized Hosts** to locate and select the host.
3. Expand the host. Expand **Volumes**, and do one of the following:
   - Select the required volume. Click **Snapshots** tab. Right-click the required mounted file system snapshot and select **Unmount**
   - Locate and right-click the required mounted file system. Select **Unmount**
4. In the **Unmount File System Snapshot** wizard panel, specify unmount details. Click **OK**.

   See "Unmount File System Snapshot panel options" on page 275.
5. In the **Result** panel, verify that the checkpoint has been unmounted successfully.

See "Remounting file system snapshot" on page 271.

See "Creating file system snapshots" on page 269.

See "About managing file systems" on page 246.

Unmount File System Snapshot panel options

Use this wizard panel to unmount a file system snapshot.

Table 11-16 Unmount File System Snapshot panel options

Field	Description
File system	Displays the name of the selected file system on which the file system snapshot needs to be unmounted.
File System Snapshot Name	Displays the file system snapshot name.
Mount Point	Displays the file system snapshot mount point.
Node	Select one or more nodes from where you want to unmount the file system snapshot. This option is enabled only for cluster file system.
Remove cluster mount	Select this check box to remove the cluster file system snapshot mount point. This option is enabled only for a file system snapshot mounted on a clustered file system.
Remove selected snapshot	Select this check box to remove file system snapshot.

See "Unmounting file system snapshot" on page 275.

Removing file system snapshot

The Management Server console lets you remove an existing file system snapshot. A file system snapshot must be unmounted before it can be removed.

To perform this task, your user group must be assigned the Admin role on the host or the Server perspective. The permission on the host may be explicitly assigned or inherited from a parent Organization.

To remove a file system snapshot

1. In the Management Server console, go to the **Server** perspective and expand **Manage** in the left pane.
2. Expand the Organization or **Uncategorized Hosts** to locate and select the host.
3. Do one of the following:

- Expand the host. Expand **Volumes** to locate and select the required volume. Click **Snapshots** tab. Right-click on the selected unmounted file system snapshot and select **Remove**.
- Expand the host. Expand **Volumes** to locate and select the required volume. Right-click required volume and select **Snapshot** > **Remove FS Snapshot**.

4 If you launch the operation to remove the file system snapshot in the context of volumes, then in the **Remove File System Snapshot** wizard panel, specify the details required. Click **OK**.

See "Remove File System Snapshot panel options" on page 277.

5 If you launch the operation in context of snapshots, then in the **Remove File System Snapshot** wizard panel, confirm the file system snapshot you want to remove. Click **OK**.

6 In the **Result** panel, verify that the file system snapshot has been removed successfully.

See "Creating file system snapshots" on page 269.

See "Remounting file system snapshot" on page 271.

See "Unmounting file system snapshot" on page 275.

See "About managing file systems" on page 246.

Remove File System Snapshot panel options

Use this wizard panel to remove a file system snapshot.

Table 11-17 Remove File System Snapshot panel options

Field	Description
File system	Displays the name of the selected file system from which the file system snapshot is to be removed.
Remove last #snapshots	Select this option to select the number of file system snapshots to be removed.
Select snapshot to remove	Select this option to select the file system snapshot to be removed.
Snapshot name	Displays the name of the file system snapshot to be removed.
Creation Time	Displays the time the file system snapshot was created.

See “Removing file system snapshot” on page 276.

Monitoring capacity of file systems

The Management Server console lets you monitor the usage of the file systems in your environment. To monitor the file system usage, you can set the High Usage Warning and the High Usage Risk thresholds. The High Usage Warning threshold triggers risk alert when the file system usage crosses the limit that you have specified. The High Usage Risk triggers the fault alert that needs your urgent attention.

To perform this operation on a Windows host, you need to install Veritas InfoScale Operations Manager managed host version 6.1 or later.

To perform this task, your user group must be assigned the Admin role on the host or the Server perspective. The permission on the host may be explicitly assigned or inherited from a parent Organization.

To monitor capacity of a file system

1. In the Management Server console, go to the **Server** perspective and expand **Manage** in the left pane.
2. Expand the Organization or **Uncategorized Hosts** to locate the host.
3. Expand the host and select **Volumes**.
4. Right-click on the required volume and select **File System** > **Monitor Capacity**.
5. In the **Monitor Capacity** wizard panel, enter the information in the High Usage Warn (%) and in the High Usage Risk (%) fields.

   See “Monitor Capacity panel options” on page 278.
6. Click **Apply all** and click **OK**.
7. In the **Result** panel, verify that you have successfully configured the threshold values for monitoring the file system usage.

See “About managing file systems” on page 246.

Monitor Capacity panel options

Use this wizard panel to set the threshold values for monitoring the file system usage. You can also specify the amount of space that you want to increase on the file system.

Table 11-18 Monitor Capacity panel options

Field	Description
High Usage Warn (%)	Enter a value for the file system size at which a risk alert is triggered on the Management Server.
High Usage Risk (%)	Enter a value for the file system size at which a fault alert is triggered on the Management Server.
Autogrow (%)	Specify the amount of space that you want to increase on the file system when the file system crosses the high usage risk threshold. Autogrow is applicable only for VxFS file systems that are created on Storage Foundation volumes (VxVM).
Apply all	Click to apply the threshold and autogrow values.
Reset all	Click to reset the threshold and autogrow values.
Name	Displays the name of the file system.
Type	Displays the type of the file system.
Size	Displays the total size of the file system.
% Used	Displays the used space on the file system in percentage.
High Usage Warn (%)	Displays the value of the file system size at which the risk alert is triggered on the Management Server.
High Usage Risk (%)	Displays the value of the file system size at which the fault alert is triggered on the Management Server.
Autogrow (%)	Displays the space that is specified to increase on the file system in percentage, when the file system crosses the high usage risk threshold.

See “Monitoring capacity of file systems” on page 278.

Chapter 12

Managing SmartIO

This chapter includes the following topics:

- About managing SmartIO
- Enabling or disabling SmartIO caching
- Creating a cache
- Viewing the cache details
- Viewing the SmartIO Impact analysis chart
- Changing SmartIO mode
- Modifying a cache
- Deleting a cache
- Loading files to a cache
- Pinning tablespaces or files to a cache
- Unpinning tablespaces or files from a cache
- About using SmartAssist
- Creating an I/O trace log
- Viewing I/O trace logs
- Analyzing an I/O trace log
- Removing an I/O trace log

About managing SmartIO

SmartIO supports the use of multi-vendor Solid-State Devices (SSDs) as read-write cache for the high-transaction applications running on a system to improve overall I/O performance. Traditional disks often pose as an I/O bottleneck for high transaction applications. To address this issue, SmartIO facilitates SSD-based cache to drive high performance applications.

When an application issues an I/O request, SmartIO checks to see if the I/O can be serviced from the cache. As applications access data from the underlying volumes or file systems, certain data is moved to the cache based on the internal heuristics. Subsequent I/Os are processed from the cache.

SmartIO supports read and write caching for the VxFS file systems that are mounted on VxVM volumes. It also supports block-level read caching for applications running on VxVM volumes. This type of SmartIO caching supports the applications that run directly over raw volumes. For example, database instances running directly over raw volumes. Volume-level read caching can also be used in cases where VxFS caching cannot be used.

SmartIO requires a managed host with Storage Foundation 6.1 or later for Windows or Linux.

SmartIO requires a managed host with Storage Foundation 6.2 or later for Solaris and AIX.

The volumes to be cached must have the disk group version 190 and the file system layout version 10.

Creating a cache area on shared disks is not recommended. From release 7.1, the creation of multiple cache areas on a single SSD is supported from the Management Server console. SmartIO does not support caching of RAID-5 volumes and DCO volumes.

On a Linux managed host you can create one cache area of VxVM type, and an unlimited number of cache areas of VxFS type.

On a Windows managed host you can create maximum eight cache areas of type VxVM.

For more information on SmartIO, refer to *Storage Foundation Administrator's Guide* for Windows or *Veritas InfoScale SmartIO for Solid State Drives Solutions Guide* for Linux.

See “Creating a cache” on page 283.

See “Deleting a cache” on page 291.

See “Enabling or disabling SmartIO caching” on page 282.

About write-back caching in SmartIO

SmartIO supports write-back caching for applications running on VxFS file systems. In this case, the application reads and writes are satisfied from the cache whenever possible. In write-back mode, the data is written to the SmartIO cache, and the write operation is marked as a success. At a later time, SmartIO transfers the data to the disk and flushes the cache area. Write-back caching expects to improve the latencies of synchronous user data writes. When write-back caching is enabled, read caching is implicitly enabled. Reads are satisfied from the cache if possible, and the file system transparently loads file data into the cache. Both read and write-back caching may be enabled for the same file at the same time.

If a cache device fails, a file that is cached in write-back mode may not be completely present on the disk. SmartIO has a mechanism to flush the data from the cache device when the device comes back online. Storage Foundation Cluster File System High Availability (CFS) provides additional protection from data loss with cache reflection. In the case of CFS, when write-back caching is enabled, SmartIO mirrors the write-back data at the file system level to the other node's SSD cache. This behavior, called cache reflection, prevents loss of write-back data if a node fails.

If a node fails, the other node flushes the mirrored data of the lost node as part of reconfiguration. Cache reflection ensures that write-back data is not lost even if a node fails with pending data.

To enable read, or write-back caching on a VxFS file system, you can select appropriate check boxes while creating a file system and mounting or remounting a clustered or non-clustered file system. You can then also select separate cache areas to enable read and write-back caching.

See “Mounting non clustered file systems on hosts” on page 259.

See “Mounting clustered file systems on hosts” on page 264.

See “Remounting file systems” on page 265.

See “About managing SmartIO” on page 281.

Enabling or disabling SmartIO caching

By default the scope of caching is set to all VxFS mount points or all VxVM volumes. You can disable caching on a specific mount point or volume. You can also enable or disable caching for applications.

If the caching scope is set to selected mount point or volume, caching is not enabled on any mount point or volume by default. You need to explicitly enable caching on the required mount point or volume.

To perform this task, your user group must be assigned the Admin role on the host or the Server perspective. The permission on the host may be explicitly assigned or inherited from a parent organization.

To enable or disable SmartIO caching

1. In the Management Server console, go to the **Server** perspective and expand **Manage** in the left pane.
2. Expand the Applications, Organization, or Uncategorized Hosts.
3. Do one of the following:
   - Expand **Organization** or **Uncategorized Hosts**, and select the required host. Expand **Volumes** to locate the required volume. Right-click the volume, and select **SmartIO** > **Configuration** > **Cache Areas**.
   - Expand **Applications** and expand **Databases** to locate the database and do one of the following.
     - Right-click the database, and select **SmartIO** > **Configuration**.
     - Select **Volumes**. Right-click the required volume, and select **SmartIO** > **Configuration**.
4. In the **Cache Areas** tab of the **SmartIO Overview** panel, do one of the following:
   - Check **SmartIO caching** to enable SmartIO caching.
   - Uncheck **SmartIO caching** to disable SmartIO caching.
5. Select the cache area from the list while enabling caching for applications. For volumes, the cache area is selected by default.

See "Creating a cache" on page 283.

See "Modifying a cache" on page 288.

See "Deleting a cache" on page 291.

Creating a cache

The Management Server console lets you create cache area using the available SSD devices.

You can specify the name of the newly-created cache area.

To perform this task, your user group must be assigned the Admin role on the host or the Server perspective. The permission on the host may be explicitly assigned or inherited from a parent organization.

To create a cache

1. In the Management Server console, go to the **Server** perspective and expand **Manage** in the left pane.
2. Expand the Organization or Uncategorized Hosts to locate and select the host.
3. Right-click the host and select **SmartIO** > **Configuration** > **Cache Areas**.
4. Click **Create Cache**.
5. In the **Create Cache** panel, enter the details, and click **OK**.

   See "Create Cache panel options" on page 284.

See "Modifying a cache" on page 288.

See "Deleting a cache" on page 291.

Create Cache panel options

Use this wizard panel to specify the attributes for the cache area.

Table 12-1 Create Cache panel options

Field	Description
Create cache	Select either Block level caching or a file level caching. For a cache of VxVM type, the Block level caching option is disabled.
Name	Specify a name for the cache area.
Select Caching Mode	For VxFS cache areas, select **Default** mode or **Reserve** mode. Default cache areas are available to all applications except those configured to use a reserved cache areas. You can have only one default cache area online at a time. It is not mandatory to have a default cache area. Reserve cache areas are available specifically to one or more applications. You can configure any number of Reserve cache areas. You can change the caching mode.

Table 12-1 Create Cache panel options *(continued)*

Field	Description
Enable Caching for **Note:** This option is not available for Windows hosts.	Select one of the following: ■ **All VxFS Mount Points**: Select to enable the cache area for all VxFS mount points. ■ **All VxVM Volumes**: Select to enable the cache area for all VxVM volumes.
Select disks	Select the disks to be used in the cache area.
Cache size	The default size of the cache is the total size of the selected disks.

See "Creating a cache" on page 283.

Viewing the cache details

You can use the Management Server console to view the cache details on a host. You can view details such as the cache name, size, state, and the following.

- **Scope**: Displays the scope of the cache. It indicates if the cache is enabled for all the volumes or mount points or selected volumes and mount points.
- **Read Hit Ratio (%)**: Displays the percentage of reads served by the cache.
- **Write Hit Ratio (%)**: Displays the percentage of writes served by the cache. This is displayed only when writeback caching is enabled for VxFS mount points.

For each cached volume/mount point, you can view details such as, Name, Disk Group, and Size. Also, you can view Read Hit Ratio and Write Hit Ratio, and the following:

- **Cache Utilization**: Displays the percentage of the cache that is used by the volume/mount point.

In this view you can perform the create, modify, and delete cache area tasks.

You can view this information, if your user group has at least Guest role assigned on the perspective or the Organization.

To view the cache details

1. In the Management Server console, go to the **Server** perspective, and expand **Manage** in the left pane.
2. Expand the Organization or **Uncategorized Hosts** to locate the host.

3 Right-click the host and click **Smart IO** > **Configuration** > **Cache Areas**.

4 In the Smart IO Overview pane, select a cache to view related volume information.

See "Creating a cache" on page 283.

See "Modifying a cache" on page 288.

See "Deleting a cache" on page 291.

Viewing the SmartIO Impact analysis chart

In the Management Server console, you can view the impact of enabling SmartIO on a host or application with live interactive graphs.

Alternately, you can view the SmartIO impact charts for different durations - six hours, one day, one week, one month, or one year. For Windows hosts, you can review the performance for one day, one month or one year.

You can view four types of charts:

- Total Reads
- Total Writes
- Average Response Time (Read)
- Average Response Time (Write)

On Windows hosts, you can review only the first two charts.

You can view the charts, if your user group has at least Guest role assigned on the perspective or the Organization.

To view the SmartIO impact chart

1 In the Management Server console, go to the **Server** perspective and expand **Manage** in the left pane.

2 Do one of the following:

- Expand the Applications and expand **Databases**.
- Expand the Organization or Uncategorized to locate the host.

3 Right-click the database or the host and select **SmartIO** > **Impact**.

4 Select **Scope** to specify if you want to view the impact charts for selected cache areas or all cache areas.

5 Enter the **Duration** for which you want to plot the SmartIO impact chart.

See "Enabling or disabling SmartIO caching" on page 282.

See “Creating a cache” on page 283.

Changing SmartIO mode

For a VxFS cache area, the caching mode determines what kind of caching is performed for the specified mount point. The mode can be `nocache` (displayed as `avoid cache` on the Management Server Console), `read`, or `writeback`. The default mode is read.

A VxVM cache area only supports read mode.

The caching mode represents the highest level of caching that can be enabled for objects on the mount point. If you specify `avoid cache` mode, SmartIO caching is disabled for the mount point. You cannot enable SmartIO caching for any data objects in that mount point. You must remount the file system to enable caching.

Similarly, if you specify `read` mode, you cannot enable SmartIO `writeback` caching for any data objects in that mount point.

To perform this task, your user group must be assigned the Admin role on the host or the Server perspective. The permission on the host may be explicitly assigned or inherited from a parent organization.

To enable or disable SmartIO caching

1. In the Management Server console, go to the **Server** perspective and expand **Manage** in the left pane.
2. Do one of the following:
   - Expand **Organization** or **Uncategorized Hosts**, and select the required host. Expland **Volumes** to locate the required volume. Right-click the volume, and select **Change SmartIO Mode**.
   - Expand **Applications** and expand **Databases** to locate the database. Do one of the following:
     - Select **Volumes**. Right-click the required volume, and select **Change SmartIO Mode**.
     - Select **Files**. Right-click the required file, and select **Change SmartIO Mode**.
3. In the **Change SmartIO Mode** panel, review the existing SmartIO mode of the selected database, volume, file or filesystem, and then select one of the following SmartIO modes, and then click **OK**:
   - **Read**: To enable caching for only reads.

- **Writeback** To enable caching for both reads and writes, with the writes being periodically written back to the selected volume, file, or filesystem.
- **Avoid Cache**: To disable caching for the the selected database, volume, file, or filesystem.

See "Creating a cache" on page 283.

See "Modifying a cache" on page 288.

See "Deleting a cache" on page 291.

Modifying a cache

The Management Server console lets you modify the attributes for a cache area.

To perform this task, your user group must be assigned the Admin role on the host or the Server perspective. The permission on the host may be explicitly assigned or inherited from a parent organization.

To modify a cache

1. In the Management Server console, go to the **Server** perspective and expand **Manage** in the left pane.
2. Expand the Organization or Uncategorized Hosts to locate and select the host.
3. Right-click the host and select **SmartIO** > **Configuration** > **Cache Areas**.
4. Select the required cache. The **Volumes** tab displays the list of volumes which use the cache.
5. Right-click the cache name and select **Modify Cache**.
6. In the **Modify Cache** panel, specify the required information. Click **OK**.

   See "Modify Cache panel options" on page 288.

See "Creating a cache" on page 283.

See "Deleting a cache" on page 291.

See "Enabling or disabling SmartIO caching" on page 282.

Modify Cache panel options

Use this wizard panel to modify the attributes for a cache area.

Table 12-2 Modify Cache panel options

Field	Description
Change state	Select either **Online** or **Offline**. Online enables the cache area thereby serving the reads and writes from the cache area. Offline disables the cache area. Reads and writes are not served from the cache area.
Select caching mode	Select **Default** mode or **Reserve** mode. Default cache areas are available to all applications except those configured to use a reserved cache areas. You can have only one default cache area online at a time. It is not mandatory to have a default cache area. Reserve cache areas are available specifically to one or more applications. You can configure any number of Reserve cache areas. You can change the caching mode.

Table 12-2 Modify Cache panel options *(continued)*

Field	Description
Enable Caching for **Note:** This option is not available for Windows hosts.	Select one of the following ■ The following two options are displayed if the cache area is created on VxFS file systems. ■ **All Mount Points**: Select this option to enable caching on all mount points. ■ **Selected Mount Points**: Select this option to enable caching on specific mount points. When you select **Selected Mount Points**, caching is not enabled by default. You need to explicitly enable caching on each of the selected mount points. ■ The following two options are displayed if the cache area is created on VxVM volumes. ■ **All Volumes**: Select this option to enable caching for all volumes. ■ **Selected Volumes**: Select this option to enable caching on selected volumes. When you select **Selected Volumes**, then caching is not enabled by default. You need to explicitly enable caching on each of the selected volumes.
Change cache size to	Specify the new cache size.
Add/Remove Disk from Cache	Expand **Add/Remove Disk from Cache** to view and configure the following parameters: **Resize**: To specify disks to add or remove, check the boxes next to the respective disk **Name**, and then select **Add** or **Remove**.

The Modify Cache panel also displays the Current Cache Size and Max Available Cache Size.

See "Modifying a cache" on page 288.

See "Enabling or disabling SmartIO caching" on page 282.

Deleting a cache

The Management Server console lets you delete a cache area.

To perform this task, your user group must be assigned the Admin role on the host or the Server perspective. The permission on the host may be explicitly assigned or inherited from a parent organization.

To delete a cache

1. In the Management Server console, go to the **Server** perspective and expand **Manage** in the left pane.
2. Expand the Organization or Uncategorized Hosts to locate and select the host.
3. Right-click the host and select **SmartIO** > **Configuration** > **Cache Areas**.
4. In the SmartIO Overview panel, right-click the required cache name and select **Delete Cache**.
5. In the **Delete Cache** panel, review the selection, and click **OK**.

See "Creating a cache" on page 283.

See "Modifying a cache" on page 288.

See "Modify Cache panel options" on page 288.

Loading files to a cache

The Management Server console lets you load specific files to a cache.

To perform this task, your user group must be assigned the Admin role on the host or the Server perspective. The permission on the host may be explicitly assigned or inherited from a parent organization.

To load files to a cache

1. In the Management Server console, go to the **Server** perspective and expand **Manage** in the left pane.
2. Expand the **Applications** and expand **Databases** to select a database.
3. Click the **Files** tab.
4. Right-click the required file and select **Load to cache**.
5. In the **Load to cache** panel, review the selection, and click **OK**.

   See "Pinning tablespaces or files to a cache " on page 292.

See "Modifying a cache" on page 288.

See "Deleting a cache" on page 291.

Pinning tablespaces or files to a cache

The Management Server console lets you pin tablespaces and files to a cache.

To perform this task, your user group must be assigned the Admin role on the host or the Server perspective. The permission on the host may be explicitly assigned or inherited from a parent organization.

To pin tablespaces or files to a cache

1. In the Management Server console, go to the **Server** perspective and expand **Manage** in the left pane.
2. Expand the **Applications** and expand **Databases** to select a database.
3. Click the **Tablespaces** or **Files** tab.
4. Right-click on the required tablespace or file and select **Pin to cache**.
5. In the **Pin to cache** panel, review the selection, and click **OK**.

   See “Unpinning tablespaces or files from a cache ” on page 292.

   See “Loading files to a cache” on page 291.

See “Modifying a cache” on page 288.

See “Deleting a cache” on page 291.

Unpinning tablespaces or files from a cache

If you have pinned tablespaces or files to a cache, using the Management Server console you can unpin them from the cache.

To perform this task, your user group must be assigned the Admin role on the host or the Server perspective. The permission on the host may be explicitly assigned or inherited from a parent organization.

To unpin tablespaces or files from a cache

1. In the Management Server console, go to the **Server** perspective and expand **Manage** in the left pane.
2. Expand the **Applications** and expand **Databases** to select a database.
3. Click the **Tablespaces** or **Files** tab.
4. Right-click on the required tablespace or file and select **Unpin from cache**.
5. In the **Unpin from cache** panel, review the selection, and click **OK**.

   See “Pinning tablespaces or files to a cache ” on page 292.

See “Changing SmartIO mode” on page 287.

About using SmartAssist

The SmartAssist tool lets you collect IO trace logs from a set of volumes, mount-points, disks, or diskgroups for a specified duration. The tool then analyzes the log to determine the optimal SmartIO cache size, related latency gain, and performance forecast.

SmartAssist runs in two phases:

1. IO trace log collection
2. IO trace log analysis using SmartIO caching algorithms

You can use the Management Server console to create and analyze I/O trace logs.

For more information, see:

- See “Viewing I/O trace logs” on page 295.
- See “Creating an I/O trace log ” on page 293.
- See “Analyzing an I/O trace log” on page 295.
- See “Removing an I/O trace log” on page 296.

For information on SSD-based caching with SmartIO:

About managing SmartIO (cross reference)

Creating an I/O trace log

The Management Server console lets you create an IO trace log for a set of volumes, mount-points, disks, or disk groups. You can also specify the duration for which SmartAssist must collect the I/O trace log.

For more information on the SmartAssist I/O analysis tool, see:

About SmartAssist (cross reference)

The tool collects I/O traces from the target storage for a specified duration, and saves them at the following location:

```
/var/vx/smartassist
```

To perform this task, your user group must be assigned the Admin role on the host or the Server perspective. The permission on the host may be explicitly assigned or inherited from a parent Organization.

To create an I/O trace log

1. In the Management Server console, go to the **Server** perspective, and expand **Manage** in the left pane.
2. Expand **Organization** or **Uncategorized Hosts** to locate and select the host.
3. Right-click the host and select **SmartIO** > **Configuration**.
4. In the SmartIO overview panel, click **SmartAssist**.
5. Click **New I/O Trace Log** to launch the New I/O Trace Log wizard.
6. In the **New I/O Trace Log** panel, enter the details, and click **OK**.

   For details, see:

   See "New I/O Trace Log panel options " on page 294.

See "Viewing the SmartIO Impact analysis chart" on page 286.

See "Analyzing an I/O trace log" on page 295.

New I/O Trace Log panel options

Use this wizard panel, specify the details for the new I/O trace log.

Table 12-3 New I/O Trace Log panel options

Field	Description
I/O Trace Log Name	Specify a name for the new I/O trace log.
Select Storage Object Type	Specify the type of storage object (or device) for which you want to create the I/O trace log.
List of Storage Objects	Select one or more volumes, mount points, disks, or disk groups for which you want to create an I/O trace log. Select objects of only one type per I/O trace log.
Number of Storage Objects	The number of storage objects that are selected for collecting the I/O trace log.
Trace Log Duration	Specify the time for which SmartAssist must collect the I/O trace log.

See "Creating an I/O trace log " on page 293.

Viewing I/O trace logs

The Management Server Console lets you review details of existing I/O Trace Logs from the SmartAssist tab.

To perform this task, your user group must be assigned the Admin role on the host or the Server perspective. The permission on the host may be explicitly assigned or inherited from a parent Organization.

To view I/O trace logs

1. In the Management Server console, go to the **Server** perspective and expand **Manage** in the left pane.
2. Expand **Organization** or **Uncategorized Hosts** to locate and select the host.
3. Right-click the host and select **SmartIO** > **Configuration**.
4. In the SmartIO overview panel, click **SmartAssist**.

   Review the details of existing I/O trace logs.

See "Creating an I/O trace log " on page 293.

See "Removing an I/O trace log" on page 296.

Analyzing an I/O trace log

The Management Server console lets you analyze the I/O trace log that you collected with SmartAssist. The tool helps you to determine the optimal SmartIO cache size for the volume, mount point, disk, or disk group where you collected the I/O trace log. You can also view related latency gain and performance forecast.

If you plan to allocate a custom cache size, SmartAssist also helps you determine the latency gain and performance forecast for that cache size.

To perform this task, your user group must be assigned the Admin role on the host or the Server perspective. The permission on the host may be explicitly assigned or inherited from a parent Organization.

To analyze an I/O trace log

1. In the Management Server console, go to the **Server** perspective and expand Manage in the left pane.
2. Expand **Organization** or **Uncategorized Hosts** to locate and select the host.
3. Right-click the host, and select **SmartIO** > **Configuration** > **SmartAssist**.
4. In the SmartAssist tab, right-click the required I/O Trace Log, and click **Analyze**.

5 The Analyze I/O Trace Log panel, in the lower pane, SmartAssist by default displays the analysis results for the optimal SmartIO cache size, as follows:

- Total reads: The number of read-hits in the trace log that the optimal cache would serve.
- Cache size: Optimal cache size is displayed by default. If you run the analysis for a custom cache size, that size is displayed in this column
- Read-hit ratio: The percentage of reads that the cache would serve.
- Average latency gain: Response time in milliseconds that the cache would save.
- Speed of reads: The number of the times by which the cache would increase the speed of reads.
- Analysis method: Indicates if recommended or custom cache is used for analysis.

6 Click **View Devices** to view the list of data objects for which you created the I/O Trace Log.

7 To view the same parameters for a different cache size, enter the details in the **SmartAssist Analysis Options** pane, and click **Estimate**.

See "SmartAssist Analysis Options pane" on page 296.

SmartAssist Analysis Options pane

To view SmartAssist analysis for custom cache size, specify the following details

Table 12-4

Field	Description
Analysis with custom size	The cache size for which you want to view latency gain and performance forecast.
Cache latency	Cache latency expected for the given workload.
Data latency	Data latency expected for the given workload.

Removing an I/O trace log

The Management Server console lets you remove an I/O trace log.

To perform this task, your user group must be assigned the Admin role on the host or the Server perspective. The permission on the host may be explicitly assigned or inherited from a parent organization.

To remove an I/O trace log

1. In the Management Server console, go to the **Server** perspective and expand **Manage** in the left pane.
2. Expand Organization or **Uncategorized Hosts** to locate and select the host.
3. Right-click the host and select **SmartIO** > **Configuration** > **SmartAssist**.
4. In the SmartAssist tab, right-click the I/O Trace Log that you want to remove, and click **Remove**.
5. In the Remove I/O Trace Log panel, click **OK**.

Chapter 13

Managing application IO thresholds

This chapter includes the following topics:

- About managing application I/O workloads with IOPS settings
- Managing IO Thresholds
- Setting IO Thresholds
- Modifying an App VG
- Viewing live IOPS charts

About managing application I/O workloads with IOPS settings

When multiple applications use a common storage subsystem, it is important to balance application I/O requests so that the applications can co-exist. You can address this need by setting a maximum threshold on the I/O operations per second (IOPS) for the volumes of an application.

You can use application volume groups (App VGs) to set the thresholds. An application volume group comprises a volume or group of volumes that are associated with an application. Available volumes in an App VG service I/O requests from an application until the application volume group reaches the maximum IOPS threshold. Thereafter Veritas Volume Manager (VxVM) throttles and queues the I/O requests. VxVM services the requests in the next available time interval on priority.

The maximum IOPS threshold determines the maximum number of I/Os processed per second collectively by all the volumes in an application volume group.

In clustered environments, the IOPS threshold for an App VG is propagated to all nodes in the cluster. The threshold applies independently to each node.

From the Management Server console, you can set the maximum IO threshold for an application. VxVM internally sets up an App VG, comprising all volumes that are associated with the application.

See “Managing IO Thresholds” on page 299.

Alternately, Management Server console lets you select specific volumes to construct App VGs as required for an application. You can create multiple App VGs per application. You can specify a different IO threshold for each App VG. You can create App VGs based on the I/O requirements of the application. For example, the redo log volumes of a database application may need a different IOPS threshold, as compared to the archive volumes.

See “Setting IO Thresholds” on page 300.

You can also modify an App VG and its I/O threshold. You can perform add volume or remove volume operations, or you can modify or clear the IOPS threshold. If you do not specify an IO threshold, the App VG is automatically removed.

See “Modifying an App VG” on page 301.

For command line options to manage I/O workloads using maximum IOPS settings, see the *Storage Foundation Administrator's Guide*.

Managing IO Thresholds

You can use the Management Server Console to set maximum IOPS at an application level. When you set this I/O threshold, Veritas Volume Manager (VxVM) internally sets up an application volume group (App VG) of all volumes of that application. VxVM enforces the threshold on the combined IOPS served by the volumes in the App VG.

You can later modify or clear the threshold. If you clear the threshold the App VG is automatically deleted. The volumes are not deleted.

In clustered environments, the IOPS threshold for an App VG is propagated to all nodes in the cluster. The threshold applies independently to each node.

To perform this task, your user group must be assigned the Admin role on the host or the Server perspective. The permission on the host may be explicitly assigned or inherited from a parent Organization.

To manage the IO threshold for an application

1. In the Management Server console, go to the **Server** perspective and expand **Manage** in the left pane.
2. Expand **Applications**, and expand **Databases** to locate and right-click the database.
3. Click **Manage IO Thresholds**.
4. In the Manage IO Thresholds panel, specify a **Maximum IOPS** value, and click **OK**.

See "About managing application I/O workloads with IOPS settings" on page 298.

See "Setting IO Thresholds" on page 300.

See "Viewing live IOPS charts" on page 302.

For command line options and for more information on determining the maximum IOPS value, see the *Storage Foundation Administrator's Guide*.

Setting IO Thresholds

Management Server console lets you select specific volumes to construct App VGs as required for an application. You can create multiple App VGs per application. You can specify a different IO threshold for each App VG. You can create App VGs based on the I/O requirements of the application. For example, the redo log volumes of a database application may need a different IOPS threshold, as compared to the archive volumes.

The maximum IOPS threshold determines the maximum number of I/Os processed per second collectively by all the volumes in an application volume group. If the combined IO requests to the App VG exceed the maximum IOPS value, Veritas Volume Manager throttles the IO requests.

In clustered environments, the IOPS threshold for an App VG is propagated to all nodes in the cluster. The threshold applies independently to each node.

To perform this task, your user group must be assigned the Admin role on the host or the Server perspective. The permission on the host may be explicitly assigned or inherited from a parent Organization.

To set the IO threshold for an App VG

1. In the Management Server console, go to the **Server** perspective and expand **Manage** in the left pane.
2. Expand **Applications**, and expand **Databases** to locate the database. In the **Volumes** tab, select one or more volumes, right-click, and select **Set IO Thresholds**.
3. In the Set IO Threshold panel, specify the **Maximum IOPS value**, specify an **App VG name**, and click **OK**.

See "About managing application I/O workloads with IOPS settings" on page 298.

See "Modifying an App VG" on page 301.

See "Viewing live IOPS charts" on page 302.

For command line options and for more information on forming App VGs and determining the maximum IOPS value, see the *Storage Foundation Administrator's Guide*.

Modifying an App VG

The Management Server Console lets you modify an application volume group (App VG). You can add or remove volumes from the App VG. You can also specify a new maximum IOPS value.

In clustered environments, the IOPS threshold for an App VG is propagated to all nodes in the cluster. The threshold applies independently to each node.

Note: If you clear the maximum IOPS value, the App VG is deleted. Similarly, if you remove all volumes from the App VG, the App VG is deleted. The volumes are not deleted.

To perform this task, your user group must be assigned the Admin role on the host or the Server perspective. The permission on the host may be explicitly assigned or inherited from a parent Organization.

To modify an App VG

1. In the Management Server console, go to the **Server** perspective and expand **Manage** in the left pane.
2. Expand **Applications** and expand **Databases** to locate the database.
3. In the App VGs tab, right-click the appropriate App VG and select **Modify App VG**.

4 In the Modify App VG – Add or remove volumes panel, select volumes to add to the App VG or remove from the App VG, and click **Next**.

5 In the Modify App VG – Manage IO Thresholds panel, specify a new Maximum IOPS value if required.

6 In the Modify App VG – Summary panel, review the summary of proposed App VG modifications, and click **Finish**.

7 In the Modify App VG – Result panel, review the results, and click **OK**.

See "About managing application I/O workloads with IOPS settings" on page 298.

See "Managing IO Thresholds" on page 299.

See "Setting IO Thresholds" on page 300.

See "Viewing live IOPS charts" on page 302.

For command line options to manage I/O workloads using maximum IOPS settings, see the *Storage Foundation Administrator's Guide*.

Viewing live IOPS charts

The Management Server Console displays live charts for I/O operations per second (IOPS) for an application or application volume group (App VG).

To perform this task, your user group must be assigned the Admin role on the host or the Server perspective. The permission on the host may be explicitly assigned or inherited from a parent Organization.

To view IO Performance charts

1 In the Management Server console, go to the **Server** perspective and expand **Manage** in the left pane.

2 Expand **Applications** and expand **Databases** to locate the database.

3 In the Performance tab, do one of the following:

- To view the IOPS for all the volumes that are associated with the application, no further action is required. This chart is displayed by default, as the option 'All' is by default selected in the scope-selection list. The Max IOPS threshold is not displayed.
- To view the IOPS for all volumes in an App VG, select the App VG from the scope-selection list. The chart of the actual IOPS is displayed, along with the Max IOPS threshold.

For more information setting IOPS thresholds for application and App VGs:

See "About managing application I/O workloads with IOPS settings" on page 298.

Chapter 14

Managing replications

This chapter includes the following topics:

- About managing replications
- Configuring Storage Foundation replications
- Adding a Secondary
- Pausing the replication to a Secondary
- Resuming the replication of a Secondary
- Starting replication to a Secondary
- Stopping the replication to a Secondary
- Switching a Primary
- Taking over from an original Primary
- Associating a volume
- Resynchronizing a Secondary
- Removing a Secondary
- Unconfiguring replication
- About setting alerts for replication
- Monitoring replications

About managing replications

Following is a list of replicator operations you can perform in the Management Server console.

Configuring replications

See “Configuring Storage Foundation replications” on page 304.

Adding a secondary

See “Adding a Secondary” on page 313.

Pausing replication to a Secondary

See “Pausing the replication to a Secondary” on page 314.

Resuming replication of a Secondary

See “Resuming the replication of a Secondary” on page 315.

Starting replication to a Secondary

See “Starting replication to a Secondary” on page 316.

Stopping replication to a Secondary

See “Stopping the replication to a Secondary” on page 318.

Switching a Primary

See “Switching a Primary” on page 319.

Taking over from an original Primary

See “Taking over from an original Primary” on page 320.

Removing a secondary

See “Removing a Secondary” on page 324.

See “About performing Storage Foundation and replicator operations” on page 133.

Configuring Storage Foundation replications

In Management Server console, you can configure and set up replication by creating a primary and then adding a secondary.

This operation can be launched from the context of databases or disk groups.

To perform this task, your user group must be assigned the Admin role on the host or the Server perspective. The permission on the host may be explicitly assigned or inherited from a parent Organization.

To configure a replication, create a primary and add a secondary

1. In the Management Server console, go to the **Server** perspective and expand **Manage** in the left pane.
2. Do one of the following:

- Expand **Applications**, expand **Databases**. Right-click the required database and select **Replication** > **Configure**.
- Expand **Uncategorized Hosts** or an **Organization**, expand the host, and expand **Disk Groups**. Right-click the required disk group and select **Configure replication**.

3 In the **Configure replication - Create primary** wizard panel, enter the options to create the primary. Click **Create primary**.

See "Configure replication - Create primary panel options" on page 305.

4 In the **Configure replication - Create primary result** wizard panel, verify that the operation was successful. The successful creation of the primary enables the **Add secondary** button.

See "Configure replication - Create primary result panel options" on page 308.

5 Do one of the following:

- To add the secondary later, click **Close** and skip the remaining steps.
- To immediately add the secondary, click **Add secondary**.

6 In the **Configure replication - Consistency check on secondary host** wizard panel, you can check for configuration consistency result, select the secondary host. Click **Next**.

7 In the **Configure replication - Replication settings for secondary host** wizard panel, select the attributes for the secondary. Click **Finish**.

See "Configure replication - Replication settings for secondary host panel options" on page 310.

8 In the **Configure replication - Add secondary result** wizard panel, verify that the secondary has been added successfully and the replication has started. If you want to add another secondary, click on **Add secondary** tab. If you do not want to add another secondary, click **Close**.

Note: If any operation fails after you have created any object, you can rollback the action of creating the object.

See "Configure replication - Add secondary result panel options" on page 313.

See "Adding a Secondary" on page 313.

Configure replication - Create primary panel options

Use this wizard panel to configure replication and create a primary.

Table 14-1 Configure replication - Create primary panel options

Name	Description
RVG name	Enter a name for the primary replicated volume group (RVG). Rules for RVG names: ■ A valid replicated data set (RDS) name must contain only alpha-numeric characters, underscores, dashes, or periods. ■ The length of the name must be between 1 and 31 characters long. ■ The name cannot begin with a dash or a period, or end with a period.
Volume	Select one or more data volumes or volume sets to be replicated. The selected data volumes are associated to the primary RVG. **Note:** All of the data volumes used by the application MUST be included in the same RVG.
Name	Displays the name of the volume.
File System	Displays the file system name.
State	Displays the health of the volume.
Size	Displays the total size of the volume.
Replicate?	Select the check box if you want to replicate using that particular volume.
Storage Replicator Log (SRL)	Select the volume that is to be configured as the SRL.
New	Select to add a new volume as the SRL.
Use existing	Select to use an existing volume as the SRL.
Volume name	If you are adding a new volume as the SRL, enter a name for the volume.
Size	If you are adding a new volume as the SRL, specify the SRL size. The SRL size must be at least 110MB.

Table 14-1 Configure replication - Create primary panel options *(continued)*

Name	Description
Stripe layout	If you are adding a new volume as the SRL, select the check box if you want use the stripe layout. This option is available only when the disk group has more than one disk.
Manually Select Disk	Click to manually select the disks.
Create VCS Configuration	Select the checkbox if you want to create VCS configuration on the host. This option is enabled only if VCS is installed on the host and you have the privileges to create VCS configurations. If you select this option, you need to create resources for VCS configurations. See "Configure replication - Create resources" on page 307.
Select	Click to select the volume to be associated as the SRL. The volume must meet the following criteria: ■ It should not have any file system mounted on it. ■ It should not be a snapshot volume. ■ It should not be part of another replication. ■ It should not be selected as the data volume to be replicated.

See "Configuring Storage Foundation replications" on page 304.

Configure replication - Create resources

Use this wizard panel to configure replication and create resources for VCS configurations.

Table 14-2 Configure replication - Create resources

Name	Description
Service group Name	Displays the automatically generated name of the service group. You can edit the service group name.

Table 14-2 Configure replication - Create resources *(continued)*

Name	Description
Resources	
Resource Type	Displays the resource types.
Name	Names of the resources. You can edit these names.
Required resources attributes	
Attribute Name	Displays the names of the resource attributes.
Value	Enter or edit the value for the resources.

See "Configuring Storage Foundation replications" on page 304.

Configure replication - Create primary result panel options

Use this wizard panel to view the result of the new primary created and to add a secondary. Once the primary is created, the **Add secondary** tab is enabled. If you chose to create a new SRL, then the result panel also displays the result of a new volume created.

See "Configuring Storage Foundation replications" on page 304.

Configure replication - Consistency check on secondary host panel options

Use this wizard panel to check for consistency before adding a secondary.

Table 14-3 Configure replication - Consistency check on secondary host panel options

Name	Description
Secondary host	Expected value is Host name or IP address. If it is added to Veritas InfoScale Operations Manager, consistency check can be run. If it is not added to Veritas InfoScale Operations Manager, make sure that all shown attributes from primary RVG are present on secondary host. Click **Select host**, to select a host from the existing list.

Table 14-3 Configure replication - Consistency check on secondary host panel options *(continued)*

Name	Description
Configuration consistency check results	Displays the results of the consistency check that is run to match the available objects on the primary RVG with the secondary. If the secondary host is not a part of Veritas InfoScale Operations Manager CMS, you can perform the following checks to ensure consistency: ■ Disk group with the same name is available. ■ Data volumes & SRL with the same name and sizes are available. ■ Ensure that none of the volumes have a mounted file system. ■ In case of Windows host, the data volume selected should have a DCM log created on it.
Primary RVG object	Displays the available objects on the primary RVG.
Match found on secondary	Displays the result of the consistency check that is run to find a match on the primary RVG.
Create Disk Group	Click to create a disk group on secondary, if there is no suitable disk group found on the secondary.
Create missing volume on secondary	If there are no suitable volumes available on the secondary, click to create the volume on secondary. ■ Automatically: Select this option to automatically create the volumes. ■ Manually: Select this option to manually create the volumes on secondary.
Recheck consistency	Click to recheck the consistency after creating a missing disk group or a volume.

Table 14-3 Configure replication - Consistency check on secondary host panel options *(continued)*

Name	Description
Proceed with Add secondary operation	Select this check box if you find any inconsistencies between the primary RVG object and the secondary host, but still want to proceed with the **Add Secondary** operation.

See "Adding a Secondary" on page 313.

See "Configuring Storage Foundation replications" on page 304.

Configure replication - Replication settings for secondary host panel options

Use this wizard panel to select the replication settings for the secondary host.

Table 14-4 Configure replication - Replication settings for secondary host panel options

Name	Description
Primary name/IP	Enter a host name or specify the IP address that can be used for replication. The Secondary host name must be resolvable and reachable from the primary host. For example, london. If you entered the IP address, it must be reachable from the primary host. A valid Host name must contain only alpha-numeric characters, underscores, dashes, or periods. The length of the name must be between 1 and 31 characters long. It cannot begin with a dash or a period, or end with a period.

Table 14-4 Configure replication - Replication settings for secondary host panel options *(continued)*

Name	Description
Secondary name/IP	Enter a host name or specify the IP address that can be used for replication. The Secondary host name must be resolvable and reachable from the primary host. For example, london. If you entered the IP address, it must be reachable from the primary host. A valid Host name must contain only alpha-numeric characters, underscores, dashes, or periods. The length of the name must be between 1 and 31 characters long. It cannot begin with a dash or a period, or end with a period.
Primary RLink	Enter the primary RLink name. A valid RLink name must contain only alpha-numeric characters, underscores, dashes, or periods. The length of the name must be between 1 and 31 characters long. It cannot begin with a dash or a period, or end with a period.
Secondary RLink	Enter the Secondary RLink name.
Start replication	Select the check box to start replication to a Secondary. When you start replication to a Secondary, the data volumes on the Secondary must be synchronized with the data volumes on the primary. Options available for Start Replication operation: ■ Automatic synchronization ■ Synchronization not needed
Advanced options	

Table 14-4 Configure replication - Replication settings for secondary host panel options *(continued)*

Name	Description
Replication modes	For a UNIX/Linux host, select one of the following: ■ Synchronous-Override ■ Asynchronous For a Windows host, select one of the following: ■ Synchronous-Override ■ Synchronous ■ Asynchronous
Protocol	Select one of the following: ■ UDP/IP ■ TCP/IP
SRL Protection	Select one of the following: ■ AutoDCM ■ DCM ■ Override ■ Fail ■ Off
Latency protection	Select one of the following: ■ Off ■ Override
Low Mark	Displays the number of updates in the SRL before the protection becomes inactive and updates succeed.
High mark	Specifies the maximum number of waiting updates in the SRL before the protection becomes active and updates stall or fail.
Set bandwidth limit	Select the check box to specify the bandwidth limit for the replication.

Table 14-4 Configure replication - Replication settings for secondary host panel options *(continued)*

Name	Description
Compress data	Select the check box to compress data before starting the replication to a Secondary. This option is displayed only if the installed Storage Foundation version supports the compression feature.
Create VCS Configuration	Select the checkbox if you want to create VCS configuration on the host. This option is enabled only if VCS is installed on the host and you have the privileges to create VCS configurations. If you select this option, you need to create resources for VCS configurations. See "Configure replication - Create resources" on page 307.

See "Adding a Secondary" on page 313.

See "Configuring Storage Foundation replications" on page 304.

Configure replication - Add secondary result panel options

Use this wizard panel to view the result of the new secondary added. You can also view the status of the replication started. To add another secondary, click **Add secondary**.

See "Adding a Secondary" on page 313.

See "Configuring Storage Foundation replications" on page 304.

Adding a Secondary

In Management Server console, you can add a secondary to the replicated volume group. You can also synchronize the secondary data volumes with the primary data volumes and start the replication.

To perform this task, your user group must be assigned the Admin role on the host or the Server perspective. The permission on the host may be explicitly assigned or inherited from a parent Organization.

To add a Secondary and start replication

1. In the Management Server console, go to the **Server** perspective and expand **Manage** in the left pane.
2. Do one of the following:
   - Expand **Applications**, expand **Databases**. Right-click the required database and select **Replication** > **Add secondary**.
   - Expand **Uncategorized Hosts** or an **Organization**, expand the host, and expand **RVGs**. Right-click the required RVG and select **Add secondary**.
3. In the **Configure replication - Consistency check on secondary host** wizard panel, you can check for configuration consistency result. Click **Close**.

   See "Configure replication - Consistency check on secondary host panel options" on page 308.
4. In the **Configure replication - Replication settings for secondary host** wizard panel, select the accept attributes for the secondary. Click **Finish**.

   See "Configure replication - Replication settings for secondary host panel options" on page 310.
5. In the **Configure replication - Add secondary result** wizard panel, verify that the secondary has been added successfully and the replication has started. If you want to add another secondary, click on **Add secondary** tab. If you do not want to add another secondary, click **Close**.

   Note: If any operation fails after you have created any object, you can rollback the action of creating the object.

   See "Configure replication - Add secondary result panel options" on page 313.

See "Configuring Storage Foundation replications" on page 304.

Pausing the replication to a Secondary

In Management Server console, you can pause the replication to a Secondary. The new and the already queued updates on the Primary are prevented from reaching the Secondary. This operation pauses communication between the Primary and the Secondary. The user should monitor the SRL overflow while performing this operation.

This operation can be launched from the context of databases or hosts.

To perform this task, your user group must be assigned the Admin role on the host or the Server perspective. The permission on the host may be explicitly assigned or inherited from a parent Organization.

To pause replication to a Secondary

1. In the Management Server console, go to the **Server** perspective and expand **Manage** in the left pane.
2. Do one of the following:
   - Expand **Applications**, expand **Databases**. Right-click the required database and select **Replication** > **Pause**.
   - Expand **Uncategorized Hosts** or an **Organization**, expand the host, and expand **RVGs**. Right-click the required RVG and select **Pause**.
3. In the **Pause replication** wizard panel, you can view the Secondary host name, replication status and data status. Click **Finish**.
4. In the **Summary** wizard panel, verify that the selected replication has been paused.

See "Resuming the replication of a Secondary" on page 315.

Pause replication panel options

Use this wizard panel to pause the replication from a primary to a secondary.

Table 14-5 Pause replication panel options

Name	Description
Secondary host name	Displays the name of the secondary host.
Replication status	Displays the replication status of the RVG.
Data status	Displays the data status of the RVG replication.

See "Pausing the replication to a Secondary" on page 314.

Resuming the replication of a Secondary

In Management Server console, you can resume the replication that was paused between the primary and a secondary.

This operation can be launched from the context of databases or hosts.

To perform this task, your user group must be assigned the Admin role on the host or the Server perspective. The permission on the host may be explicitly assigned or inherited from a parent Organization.

To resume replication to a Secondary

1. In the Management Server console, go to the **Server** perspective and expand **Manage** in the left pane.
2. Do one of the following:
   - Expand **Applications**, expand **Databases**. Right-click the required database and select **Replication** > **Resume**.
   - Expand **Uncategorized Hosts** or an **Organization**, expand the host, and expand **RVGs**. Right-click the required RVG and select **Resume**.
3. In the **Resume replication** wizard panel, you can view the secondary host name, replication status and data status. Click **Finish**.
4. In the **Summary** wizard panel, verify that the selected replication has been resumed.

See "Pausing the replication to a Secondary" on page 314.

Resume replication panel options

Use this wizard panel to resume the paused replication between the primary and a secondary.

Table 14-6 Resume replication panel options

Name	Description
Secondary host name	Displays the name of the secondary host.
Replication status	Displays the replication status of the RVG.
Data status	Displays the data status of the RVG replication.

See "Resuming the replication of a Secondary" on page 315.

Starting replication to a Secondary

In Management Server console, you can start the replication of a replicated volume group (RVG).

This operation can be launched from the context of databases or hosts.

To perform this task, your user group must be assigned the Admin role on the host or the Server perspective. The permission on the host may be explicitly assigned or inherited from a parent Organization.

To start replication to a Secondary

1. In the Management Server console, go to the **Server** perspective and expand **Manage** in the left pane.
2. Do one of the following:
   - Expand **Applications**, expand **Databases**. Right-click the required database and select **Replication** > **Start**.
   - Expand **Uncategorized Hosts** or an **Organization**, expand the host, and expand **RVGs**. Right-click the required RVG and select **Start**.
3. In the **Start replication** wizard panel, click **OK**.

   See “Start replication panel options ” on page 317.
4. In the **Summary** wizard panel, verify that the selected replication has been started.

See “Stop replication panel options ” on page 318.

Start replication panel options

Use this wizard panel to start the replication to a secondary replicated data set (RDS).

Table 14-7 Start replication panel options

Name	Description
Using	The replication can be started using one of the following: ■ Automatic synchronization ■ Synchronization not needed
Secondary host name	Displays the name of the secondary host.
Replication status	Displays the replication status of the RVG.
Data status	Displays the data status of the RVG replication.

See “Starting replication to a Secondary” on page 316.

Stopping the replication to a Secondary

In Management Server console, you can stop the replication to a Secondary. This operation fails if the primary and secondary are not up-to-date.

This operation can be launched from the context of databases or hosts.

To perform this task, your user group must be assigned the Admin role on the host or the Server perspective. The permission on the host may be explicitly assigned or inherited from a parent Organization.

To stop replication to a Secondary

1. In the Management Server console, go to the **Server** perspective and expand **Manage** in the left pane.
2. Do one of the following:
   - Expand **Applications**, expand **Databases**. Right-click the required database and select **Replication** > **Stop**.
   - Expand **Uncategorized Hosts** or an **Organization**, expand the host, and expand **RVGs**. Right-click the required RVG and select **Stop**.
3. In the **Stop replication** wizard panel, click **OK**.

   See “Stop replication panel options ” on page 318.
4. In the **Summary** wizard panel, verify that the selected replication has been stopped.

See “Starting replication to a Secondary” on page 316.

Stop replication panel options

Use this wizard panel to stop the replication from a primary to a secondary.

Table 14-8 Stop replication panel options

Name	Description
Secondary host name	Displays the name of the secondary host.
Replication status	Displays the replication status of the RVG.
Data status	Displays the data status of the RVG replication.
Force	Select the **Force** check box to stop the replication even when the primary and secondary RLINKs are not up-to-date.

See "Stopping the replication to a Secondary" on page 318.

Switching a Primary

In Management Server console, switching of a primary to a secondary is useful when the primary must be brought down for maintenance or to make the application active on another node. Migration involves transferring a healthy primary to a secondary when the application that is involved in replication is inactive.

This operation can be launched from the context of databases or hosts.

To perform this task, your user group must be assigned the Admin role on the host or the Server perspective. The permission on the host may be explicitly assigned or inherited from a parent Organization.

To switch a Primary

1. In the Management Server console, go to the **Server** perspective and expand **Manage** in the left pane.
2. Do one of the following:
   - Expand **Applications**, expand **Databases**. Right-click the required database and select **Replication** > **Switch**.
   - Expand **Uncategorized Hosts** or an **Organization**, expand the host, and expand **RVGs**. Right-click the required RVG and select **Switch**.
3. In the **Switch replication** wizard panel, click **Finish**.
4. In the **Summary** wizard panel, verify that the selected replication has been switched.

See "Starting replication to a Secondary" on page 316.

Switch replication panel options

Use this wizard panel to switch the replication from a Primary to a Secondary.

Table 14-9 lists the attributes to switch the replication from a Primary to a Secondary when the RVG is not under VCS control.

Table 14-10 lists the attributes to switch the replication from a Primary to a Secondary when the RVG is under VCS control.

Table 14-9 Switch replication if RVG is not under VCS control panel options

Name	Description
Secondary host name	Displays the name of the Secondary host.

Table 14-9 Switch replication if RVG is not under VCS control panel options *(continued)*

Name	Description
Replication status	Displays the replication status of the RVG.
Data status	Displays the data status of the RVG replication.

Table 14-10 Switch replication if RVG is under VCS control

Name	Description
Secondary host name	Displays the name of the Secondary host.
Replication status	Displays the replication status of the RVG.
Data status	Displays the data status of the RVG replication.
Displays the name of the service group and the remote cluster on which the service group is switched over.	

See "Switching a Primary" on page 319.

Taking over from an original Primary

In Management Server console, you can transfer the Primary role from an original Primary to a Secondary. When the original Primary fails or is destroyed because of a disaster, the takeover procedure enables you to convert a consistent Secondary to a Primary. The takeover of a primary role by a Secondary is useful when the Primary experiences unscheduled downtimes or is destroyed because of a disaster.

Note: The takeover procedure does not guarantee that the new Primary and any additional Secondary RVGs have identical contents. The remaining Secondaries must be completely synchronized with the new Primary.

This operation can be launched from the context of databases or hosts.

To perform this task, your user group must be assigned the Admin role on the host or the Server perspective. The permission on the host may be explicitly assigned or inherited from a parent Organization.

To takeover a Primary

1. In the Management Server console, go to the **Server** perspective and expand **Manage** in the left pane.
2. Do one of the following:
   - Expand **Applications**, expand **Databases**. Right-click the required database and select **Replication** > **Takeover**.
   - Expand **Uncategorized Hosts** or an **Organization**, expand the host, and expand **RVGs**. Right-click the required RVG and select **Takeover**.
3. In the **Takeover** wizard panel, select the failback options and click **OK**.

   See "Takeover Primary panel options " on page 321.
4. In the **Summary** wizard panel, verify that the selected replication has been taken over.

See "Starting replication to a Secondary" on page 316.

Takeover Primary panel options

Table 14-11 lists the attributes to takeover the replication from a Primary to a Secondary when the RVG is not under VCS control.

Table 14-12 lists the attributes to takeover the replication from a Primary to a Secondary when the RVG is under VCS control.

Table 14-11 Takeover replication if RVG is not under VCS control panel options

Name	Description
Secondary host name	Displays the name of the Secondary host.
Replication status	Displays the replication status of the RVG.
Data status	Displays the data status of the RVG replication.
Failback options	
Fast failback	Select this option to incrementally synchronize the original Primary with the new Primary.

Table 14-11 Takeover replication if RVG is not under VCS control panel options *(continued)*

Name	Description
Auto fast failback	Select this option to convert the original Primary to a Secondary after the original Primary becomes available and also to automatically synchronize the data volumes on original primary.
No fast failback	Select this option to synchronize the original primary after it becomes available. Use this option if you are sure that the original primary will not be recovered or if most of the data on the new Primary is going to change while the original Primary is unavailable.

Table 14-12 Takeover replication if RVG is under VCS control

Name	Description
Secondary host name	Displays the name of the Secondary host.
Replication status	Displays the replication status of the RVG.
Data status	Displays the data status of the RVG replication.
Displays the name of the service group and the remote cluster on which the service group is replicated.	

See "Taking over from an original Primary" on page 320.

Associating a volume

In Management Server console, if you want to replicate a volume which is not currently being replicated, you need to associate the volume to the existing RVG.

To perform this task, your user group must be assigned the Admin role on the host or the Server perspective. The permission on the host may be explicitly assigned or inherited from a parent Organization.

To associate a volume

1. In the Management Server console, go to the **Server** perspective and expand **Manage** in the left pane.
2. Expand **Uncategorized Hosts** or an **Organization**, expand the host, and expand **RVGs**. Right-click the required RVG and select **Associate Volume**.
3. In the **Associate Volume - Volume Selection** wizard panel, you can select the volumes that you want to associate to the selected RVG. Click **Next**.

   See "Associate Volume - Volume Selection panel options" on page 323.
4. In the **Associate Volume - Advanced Options** wizard panel, you can create a matching volume on the secondary and also set the synchronization option. Click **Finish**.

   See "Associate Volume - Advanced Options panel options" on page 323.

Associate Volume - Volume Selection panel options

Use this wizard panel to select the volumes that you want to associate to the existing RVG.

See "Associating a volume" on page 322.

Associate Volume - Advanced Options panel options

Use this wizard panel to view if a match for the volume is found on the secondary and set the synchronization option. You can also create the missing volume on the secondary and recheck the consistency.

Table 14-13 Associate Volume - Advanced Options panel options

Name	Description
Auto create volumes	Click to automatically create a volume on the secondary that matches with the selected volume.
Host Name	Displays the name of the host.
Match found on Secondary ?	Displays if a match for the selected volume has been found on secondary.
Synchronization option	Displays the default synchronization option. You can change the synchronization option.
Recheck consistency	Click to recheck the consistency after creating a volume on the secondary.

See “Associating a volume” on page 322.

Resynchronizing a Secondary

In Management Server console, you can resynchronize a secondary. You can perform this operation if secondary is in consistent state but needs a DCM resynchronization

To perform this task, your user group must be assigned the Admin role on the host or the Server perspective. The permission on the host may be explicitly assigned or inherited from a parent Organization.

To resynchronize a Secondary

1. In the Management Server console, go to the **Server** perspective and expand **Manage** in the left pane.
2. Expand **Uncategorized Hosts** or an **Organization**, expand the host, and expand **RVGs**. Right-click the required RVG and select **Resynchronize secondaries**.
3. In the **Resynchronize secondaries** wizard panel, confirm your action. Click **OK**.

Removing a Secondary

In Management Server console, you can remove a secondary replicated volume group (RVG). Before performing this operation, you must stop replication to the specified secondary.

The Remove Secondary operation is irreversible. For the operation to be successful, replication to the secondary must be stopped. This operation does not delete data volumes. It only dissociates the data volumes from the secondary. This operation also removes the VCS configurations, if VCS configurations for replication exists.

This operation can be launched from the context of databases or hosts.

To perform this task, your user group must be assigned the Admin role on the host or the Server perspective. The permission on the host may be explicitly assigned or inherited from a parent Organization.

To remove a Secondary

1. In the Management Server console, go to the **Server** perspective and expand **Manage** in the left pane.
2. Do one of the following:

- Expand **Applications**, expand **Databases**. Right-click the required database and select **Replication** > **Remove secondary**.
- Expand **Uncategorized Hosts** or an **Organization**, expand the host, and expand **RVGs**. Right-click the required RVG and select **Remove secondary**.

3 In the **Remove secondary** wizard panel, click **OK**.

See “Remove secondary panel options ” on page 325.

4 In the **Summary** wizard panel, verify that the selected replication has been taken over.

See “Starting replication to a Secondary” on page 316.

Remove secondary panel options

Use this wizard panel to remove the Secondary from the RVG.

Table 14-14 Remove secondary panel options

Name	Description
Secondary host name	Displays the name of the Secondary host.
Replication status	Displays the replication status of the RVG.
Data status	Displays the data status of the RVG replication.

See “Removing a Secondary” on page 324.

Unconfiguring replication

In Management Server console, you can unconfigure a replication. Invoking this operation first removes all the secondaries, and then removes the primary. This operation also removes the VCS configurations on the host, if you have the required privileges.

To perform this task, your user group must be assigned the Admin role on the host or the Server perspective. The permission on the host may be explicitly assigned or inherited from a parent Organization.

To unconfigure a replication

1. In the Management Server console, go to the **Server** perspective and expand **Manage** in the left pane.
2. Expand **Uncategorized Hosts** or an **Organization**, expand the host, and expand **RVGs**. Right-click the required RVG and select **Unconfigure replication**.
3. In the **Unconfigure replication** wizard panel, confirm your action. Click **OK**.

About setting alerts for replication

During replication, sometimes the secondary storage site lags behind the primary storage site. This can be due to time delay, bandwidth issues, or failure at the secondary storage site.

In Veritas InfoScale Operations Manager, you can monitor any Replicated Volume Group (RVG) for replications that lag behind in terms of log usage (SRL usage) and in terms of data size. You can define high and low threshold levels for raising an alert. You can view the alert in the **Settings > Alerts & Rules** view.

See “Monitoring replications” on page 326.

Monitoring replications

In Veritas InfoScale Operations Manager, you can specify low and high threshold values for raising an alert, when a replication lags behind in terms of size or log usage. An alert is displayed if the data size or log usage reaches the specified thresholds. The severity of the alert differs in cases of low and high thresholds.

You can either specify low and high thresholds individually for each selected replication, or you can specify common threshold values for all the replications.

Note: To disable the replication alerts, set empty values for both the low and the high thresholds.

To perform this task, your user group must be assigned the Admin role on the host or the Server perspective. The permission on the host may be explicitly assigned or inherited from a parent Organization.

To monitor a replication

1. In the Management Server console, go to the **Server** perspective and expand **Manage** in the left pane.
2. Do one of the following:
   - Expand **Applications**, expand **Databases** to select a database. Click the **Replications** tab.
   - Click on an **Organization** and click the **Replications** tab.
   - Expand **Uncategorized Hosts** or an **Organization**, expand the required host, expand **RVGs** to select an RVG. Click the **Links** tab.
3. Right-click on an RVG and select **Monitor replication**.
4. In the **Monitor replication** wizard panel, specify values for low and high thresholds. Click **Finish**.

   See "Monitor replication panel options" on page 327.
5. Result panel displays a message about updating the threshold values successfully in the database. The page displays an error message if the threshold values do not get updated in the database. In the result panel, click **Close**.

Monitor replication panel options

Use this wizard panel to specify low and high threshold values for raising an alert for replications that are lagging.

Table 14-15 Monitor replication panel options

Field	Description
Apply to all	To specify common threshold values for all the selected replications, select **Apply to all**. Selecting this option disables the options to set threshold values for individual replications.
Reset All	To reset all options for setting threshold values for individual replications select **Reset all**.
RVG	Displays the name of the replicated volume group (RVG).
Destination	Displays the name of the destination storage site.

Table 14-15 Monitor replication panel options *(continued)*

Field	Description
Behind By Size Low Threshold (KB)	Specify the low threshold in terms of data size for raising an alert to indicate that the replication is lagging.
Behind By Size High Threshold (KB)	Specify the high threshold in terms of data size for raising an alert to indicate that the replication is lagging.
Log Usage Low Threshold (%)	Specify the low threshold in terms of log usage percentage for raising an alert to indicate that the replication is lagging.
Log Usage High Threshold (%)	Specify the high threshold in terms of log usage percentage for raising an alert to indicate that the replication is lagging.

See “Monitoring replications” on page 326.

Chapter 15

Optimizing storage utilization

This chapter includes the following topics:

- About reclaiming thin storage
- Compressing files
- Deduplicating file systems

About reclaiming thin storage

Thin provisioning and thin reclamation are techniques to achieve efficient storage utilization. Thin provisioning refers to a technique of improving storage utilization by only allocating to a volume the physical storage required to hold its data. With thin provisioning, the storage administrator allocates logical storage to an application. The system releases the physical capacity only when the data is written. However, when the data is deleted from the file system, the physical capacity remains allocated even though it is no longer used. To ensure that your thin storage environment stays thin, you can reclaim this unused space.

You can use Veritas InfoScale Operations Manager to reclaim thin storage and to view the reclaimed physical storage space after the reclamation process is completed.

To reclaim thin storage requires managed hosts with a supported Storage Foundation version. Thin reclamation is supported on the following Storage Foundation versions:

- UNIX/Linux: 5.0 MP3, or later
- Windows: 5.1 SP1, or later

In addition, the storage must meet the following requirements:

- The storage must be visible to Storage Foundation as thin reclaimable.
- For UNIX/Linux: The LUNs must be a part of a Storage Foundation volume which has a mounted VxFS file system.
- For Windows: The LUNs must be a part of a Storage Foundation for Windows dynamic volume which has a mounted NTFS file system.

You can configure thin reclamation in the context of different objects that are discovered by Veritas InfoScale Operations Manager:

- Server perspective: Lets you select multiple file systems or disks for a selected host and perform thin reclamation. The associated LUNs are reclaimed. You can run a report to identify file systems and hosts that are potential candidates for thin reclamation.
- Storage perspective: Lets you select one or more thin pools from a selected storage array and either schedule thin reclamation or perform it manually. The associated LUNs are reclaimed. A thin pool is a collection of devices in the array that are dedicated for use by thin LUNs. Thin pools are available if the array supports them and if Storage Insight Add-on is configured for the selected enclosure. You can run a report to identify thin pools that are potential candidates for reclamation.
  See *Veritas InfoScale Operations Manager Management Server Add-ons User Guide*.

See "Performing thin reclamation on file systems or disks" on page 330.

See "Performing thin reclamation on thin pools in enclosures" on page 331.

Performing thin reclamation on file systems or disks

You can select multiple file systems or disks for a selected host and perform thin reclamation. The associated LUNs are reclaimed.

For requirements for thin reclamation, see the following topic:

See "About reclaiming thin storage" on page 329.

To perform this task, your user group must be assigned the Admin role on the host or the Server perspective. The permission on the host may be explicitly assigned or inherited from a parent Organization.

To run thin reclamation on file systems

1. In the Management Server console, go to the **Server** perspective and expand **Manage** in the left pane.
2. Expand the Organization or **Uncategorized Hosts** to locate the host.

3 Expand the host and expand **Volumes**.

4 Select one or more volumes on which file systems are mounted, right-click and select **File System > Reclaim Thin Storage**.

To run thin reclamation on disks

1 In the Management Server console, go to the **Server** perspective and expand **Manage** in the left pane.

2 Expand the Organization or **Uncategorized Hosts** to locate the host.

3 Select the host and click the **Disks** tab.

4 Select one or more disks, right-click and select **Reclaim thin storage**.

Performing thin reclamation on thin pools in enclosures

You can select one or more thin pools from a selected storage array and either schedule thin reclamation or perform it manually. The associated LUNs are reclaimed.

Note: Thin pools are available if the array supports them and if Storage Insight Add-on is configured for the selected enclosure.

See *Veritas InfoScale Operations Manager Management Server Add-ons User Guide*.

For other requirements for thin reclamation, see the following topic:

See "About reclaiming thin storage" on page 329.

To perform this task, your user group must be assigned the Admin role on the enclosure or the Storage perspective. The permission on the enclosure may be explicitly assigned or inherited from a parent Organization.

To schedule thin reclamation on thin pools

1 In the Management Server console, go to the **Storage** perspective and expand **Manage** in the left pane.

2 Expand the Organization or **Uncategorized Enclosures** to locate the enclosure.

3 Expand the enclosure and select **Thin Pools**.

4 In the table, select one or more thin pools, right-click and select **Schedule Reclamation**.

5 Choose from the options to schedule when thin reclamation runs for the selected thin pools.

Frequency	Select **Once**, **Daily**, **Weekly**, or **Monthly**.
When	The options in the **When** area change depending on the **Frequency** selection. For **Weekly**, you can select **Every weekday** to schedule Monday through Friday or select specific days of the week. For **Monthly**, you can schedule the reclamation to re-occur on a specific day of every month.

To run thin reclamation on thin pools

1 In the Management Server console, go to the **Storage** perspective and expand **Manage** in the left pane.

2 Expand the Organization or **Uncategorized Enclosures** to locate the enclosure.

3 Expand the enclosure and select **Thin Pools**.

4 In the table, select one or more thin pools, right-click and select **Run reclamation**.

Compressing files

Veritas InfoScale Operations Manager lets you compress files by directory on a managed host.

File systems for compression must meet the following requirements:

- Storage Foundation 6.0 or later
- Veritas File System (VxFS) disk layout version 9 or later
- Veritas InfoScale Operations Manager managed host (VRTSsfmh) version 5 or later

Veritas InfoScale Operations Manager does not support file compression on Windows file systems.

See the following topics:

See "About file compression in Veritas InfoScale Operations Manager" on page 333.

See "Setting up compression schedules" on page 335.

See "Selecting directories for compression " on page 334.

See "Starting compression on demand" on page 337.

About file compression in Veritas InfoScale Operations Manager

The compression feature in Storage Foundation enables customers to use host-based compression to optimize existing primary storage. Enabling compression at the file system layer results in storage savings and avoids complex and expensive appliances typically associated with primary compression. Use cases for compression include database archive logs and unstructured data.

Compression is performed without needing any application changes and with minimal overhead. Compression does not modify the file metadata, nor are inode numbers or file extensions changed. Compression is executed out-of-band, after the write.

In Veritas InfoScale Operations Manager, you set up file compression on a host at the file system (mount point) level by selecting directories for compression. You can compress directories on demand, and you can set up a schedule for running the compression process on the host. You can view a report of space saved by file compression. Once compression is enabled, directories and files will begin to have a mix of compressed and uncompressed data blocks. This is managed automatically by the file system, and uncompressed data is compressed during the next scheduled sweep.

Following are more details on how file compression works:

- Only user data is compressible, not VxFS metadata.
- Compression is a property of a file, not a directory. If you compress all files in a directory, for example, any files that you later copy into that directory do not automatically get compressed as a result of being copied into the directory.
- A compressed file is a file with compressed extents. Writes to the compressed file cause the affected extents to get uncompressed; the result can be files with both compressed and uncompressed extents.
- After a file is compressed, the inode number does not change, and file descriptors that are opened before the compression are still valid after the compression.

File compression can have the following interactions with applications:

- In general, applications notice no difference between compressed and uncompressed files, although reads and writes to compressed extents are slower than reads and writes to uncompressed extents. When an application reads a compressed file, the file system does not perform its usual readahead to avoid the CPU load that this can require. However, when reading from the primary fileset, the file system uncompresses an entire compression block (default 1 MB) and leaves these pages in the page cache. Thus, sequential reads of the

file usually only incur an extra cost when crossing a compression block boundary. The situation is different when reading from a file in a Storage Checkpoint; in this case, nothing goes into the page cache beyond the data actually requested. For optimal read performance of a compressed file accessed through a Storage Checkpoint, the application should use a read size that matches the compression block size.

- When writing to compressed extents, ensure that you have sufficient disk space and disk quota limits for the new uncompressed extents since the write uncompresses the extents. If you do not have sufficient disk space, the write can fail with the ENOSPC or EDQUOT error.
- An application that reads data from a compressed file and then copies the file elsewhere, such as `tar`, `cpio`, `cp`, or `vi`, does not preserve compression in the new data. The same is true of some backup programs.
- Backup programs that read file data through the name space do not notice that the file is compressed. The backup program receives uncompressed data, and the compression is lost.

See "Compressing files" on page 332.

Selecting directories for compression

Veritas InfoScale Operations Manager lets you select which directories to compress for the selected file system. Directories that are enabled for compression are compressed during scheduled runs of the compression process or can be compressed on demand.

If a directory was previously compressed, you can uncompress it be deselecting the directory. The uncompression will occur when the compression process runs.

To perform this task, your user group must be assigned the Admin role on the host or the Server perspective. The permission on the host may be explicitly assigned or inherited from a parent Organization.

To select directories for compression

1. In the Management Server console, go to the **Server** perspective and expand **Manage** in the left pane.
2. Expand the Organization or **Uncategorized Hosts** to locate the host.
3. Expand the host and select **Volumes**.
4. Right-click the volume on which the file system is mounted and click **Properties**.
5. Click the **Compression** tab.

6 In the directory tree, select or deselect directories to enable or disable for compression.

7 Click **Apply**.

See "Setting up compression schedules" on page 335.

See "Starting compression on demand" on page 337.

Setting up compression schedules

Veritas InfoScale Operations Manager lets you set up a compression schedule for the selected host. You can add, change, and delete a schedule.

The schedule applies to the directories that are selected for compression for each file system on the host.

To perform this task, your user group must be assigned the Admin role on the host or the Server perspective. The permission on the host may be explicitly assigned or inherited from a parent Organization.

To add a compression schedule

1 In the Management Server console, go to the **Server** perspective and expand **Manage** in the left pane.

2 Expand the Organization or **Uncategorized Hosts** to locate the host.

3 Expand the host and select **Volumes**.

4 Right-click the volume on which the file system is mounted and click **Properties**.

5 Click the **Compression** tab.

6 Click **Add Schedule**.

7 In the **Compression Schedule** window, specify the compression options and click **OK**.

Frequency	Select **Daily**, **Weekly**, or **Monthly**.
When	The options in the **When** area change depending on the **Frequency** selection. For **Weekly**, you can select weekday to schedule Monday through Friday or select specific days of the week. For **Monthly**, you can schedule the compression to re-occur on a specific day of every month.
Compression Duration	Specify how long the compression process runs. If all directories are not compressed during the specified duration, at the next scheduled compression run, the process continues with the remaining directories. For example, say that a duration of one hour is set and 10 directories are enabled for compression. After one hour, nine directories are compressed. The compression process stops. At the next scheduled run, the compression process continues with the tenth directory. The process then starts over with the first directory and compresses any new files that were added since the last run. Default: four hours
Number of CPUs to use for compression	Specify how many CPUs to use for the scheduled compression run. Default: 50 percent of the CPUs available for the host, up to 4 CPUs.

To change a compression schedule

1 On the **Compression** tab, click **Change Schedule**.

2 In the **Compression Schedule** window, update the compression options and click **OK**.

To delete a compression schedule

- On the **Compression** tab, click **Delete Schedule**.

See “Compressing files” on page 332.

Starting compression on demand

Veritas InfoScale Operations Manager lets you start the compression process on a file system immediately if there are no compression-related activities occurring on the host.

To perform this task, your user group must be assigned the Admin role on the host or the Server perspective. The permission on the host may be explicitly assigned or inherited from a parent Organization.

To start compression on demand

1. In the Management Server console, go to the **Server** perspective and expand **Manage** in the left pane.
2. Expand the Organization or **Uncategorized Hosts** to locate the host.
3. Expand the host and select **Volumes**.
4. Right-click the volume on which the file system is mounted and click **Properties**.
5. Click the **Compression** tab.
6. Optionally change the directories that are enabled for compression.
7. Click **Compress Now**.
8. On the **Compress Now** window, specify the options for compression and confirm starting compression.

Compression Duration	Specify how long the compression process runs. All directories may not be processed during the specified duration. For example, a duration of one hour is set and 10 directories are enabled for compression. After one hour, nine directories are compressed. The compression process stops. At the next scheduled run, the compression process continues with the tenth directory, and then starts over with the first directory and compresses any new files added since the last run Default: four hours
Number of CPUs to use for compression	Specify how many CPUs to use for compression. Default: 50 percent of the CPUs available for the host, up to 4 CPUs.

See “Compressing files” on page 332.

Deduplicating file systems

Veritas InfoScale Operations Manager lets you set up data deduplication for a selected file system.

File systems for deduplication must meet the following requirements:

- Storage Foundation 6.0 or later
- Veritas File System (VxFS) disk layout version 9 or later
- Veritas InfoScale Operations Manager managed host (VRTSsfmh) version 6 or later

Veritas InfoScale Operations Manager does not support deduplication on Windows file systems.

See the following topics:

See "About file system deduplication" on page 338.

See "Implementing deduplication for a file system" on page 341.

See "Starting deduplication for a file system" on page 343.

See "Disabling or removing deduplication for a file system" on page 344.

About file system deduplication

The deduplication feature in Storage Foundation enables customers to use file system deduplication to optimize existing primary storage. Enabling deduplication at the file system layer results in storage savings and avoids complex and expensive appliances typically associated with file deduplication.

Deduplication is performed without needing any application changes and with minimal overhead. Deduplication does not change the file extension, allowing users and applications to use files normally, without performance impact.

Before setting up deduplication for a file system, evaluate whether the nature of the data makes it a good candidate for deduplication.

The following are good candidates for deduplication:

- Virtual machine boot image files (vmdk files)
- User home directories
- File systems with multiple copies of files

The following might not be the best candidates for deduplication, as they have little or no duplicate data:

- Databases

- Media files, such as JPEG, MP3, and MOV

The VxFS deduplication feature works as follows. It eliminates duplicate blocks used by your data by comparing blocks across the file system. When the deduplication feature finds a duplicate block, it removes the space used and instead creates a pointer to the common block. If the duplicate file is changed, thus making the files no longer share the same block, then that changed block is saved to disk instead of the pointer.

The deduplication process performs the following tasks:

- Scans the file system for changes
- Fingerprints the data
- Identifies duplicates
- Eliminates duplicates after verifying the duplicates

The space consumed by the deduplication database is a function of the amount of data in the file system and the deduplication chunk size. On Linux or Solaris, Veritas recommends a chunk size of 4k for SFCFSHA, where multiple copies of virtual machine images are accessed over NFS. For all other datasets, Veritas recommends a chunk size of 16k or higher. More information is available on deduplication chunk size.

The deduplication feature has the following limitations:

- A full backup of a deduplicated Veritas File System (VxFS) file system can require as much space in the target as a file system that has not been deduplicated. For example, if you have 2 TB of data that occupies 1 TB worth of disk space in the file system after deduplication, this data requires 2 TB of space on the target to back up the file system, assuming that the backup target does not do any deduplication. Similarly, when you restore such a file system, you must have 2 TB on the file system to restore the complete data. However, this freshly restored file system can be deduplicated again to regain the space savings. After a full file system restore, Veritas recommends that you remove any existing deduplication configuration and that you reconfigure deduplication.
- Deduplication is limited to a volume's primary fileset.
- Deduplication does not support mounted clone and snapshot mounted file system.
- After you restore data from a backup, you must deduplicate the restored data to regain any space savings provided by deduplication.
- If you use the cross-platform data sharing feature to convert data from one platform to another, you must remove the deduplication configuration file and database and re-enable deduplication after the conversion.

- You cannot use the FlashBackup feature of NetBackup in conjunction with the data deduplication feature, because FlashBackup does not support disk layout Version 8 and 9.

About deduplication chunk size

The deduplication chunk size, which is also referred to as deduplication granularity, is the unit at which fingerprints are computed. A valid chunk size is between 4k and 128k and power of two. Once set, the only way to change the chunk size is to remove and re-enable deduplication on the file system.

You should carefully select the chunk size, as the size has significant impact on deduplication as well as resource requirements. The size directly affects the number of fingerprint records in the deduplication database as well as temporary space required for sorting these records. A smaller chunk size results in a large number of fingerprints and hence requires a significant amount of space for the deduplication database.

While the amount of storage that you save after deduplication depends heavily on the dataset and distribution of duplicates within the dataset, the chunk size can also affect the savings significantly. You must understand your dataset to get the best results after deduplication. A general rule of thumb is that a smaller chunk size saves more storage. A smaller chunk size results in more granular fingerprints and in general results in identifying more duplicates. However, smaller chunks have additional costs in terms of database size, deduplication time, and, more importantly, fragmentation. The deduplication database size can be significantly large for small chunk sizes. Higher fragmentation normally results in more file system metadata and hence can require more storage. The space consumed by the deduplication database and the increased file system metadata can reduce the savings achieved via deduplication. Additionally, fragmentation can also have a negative effect on performance. The Veritas File System (VxFS) deduplication algorithms try to reduce fragmentation by coalescing multiple contiguous duplicate chunks.

Larger chunk sizes normally result in a smaller deduplication database size, faster deduplication, and less fragmentation. These benefits sometimes come at the cost of less storage savings. If you have a large number duplicate files that are small in size, you still can choose a chunk size that is larger than the file size. A larger chunk size does not affect the deduplication of files that are smaller than the chunk size. In such cases, the fingerprint is calculated on the whole file, and the files are still deduplicated.

The space consumed by the deduplication database is a function of the amount of data in the file system and the deduplication chunk size. The space consumed by the deduplication database grows with time as new data is added to file system. Additional storage is required for temporary use, such as sorting fingerprints. The

temporary storage may be freed after the work completes. Ensure that sufficient free space is available for deduplication to complete successfully. The deduplication might not start if the file system free space is less than approximately 15%. The deduplication sometimes needs more than 15% free space for smaller chunk sizes. In general, the space consumed reduces significantly with larger chunk sizes. Veritas recommends that you have approximately 20% free space for 4k chunks.

Implementing deduplication for a file system

Veritas InfoScale Operations Manager lets you implement deduplication for a selected file system. You configure the deduplication database and optionally set up a schedule.

To perform this task, your user group must be assigned the Admin role on the host or the Server perspective. The permission on the host may be explicitly assigned or inherited from a parent Organization.

To implement deduplication for a file system

1. In the Management Server console, go to the **Server** perspective and expand **Manage** in the left pane.
2. Expand the Organization or **Uncategorized Hosts** to locate the host.
3. Expand the host and select **Volumes**.
4. Right-click the volume on which the file system is mounted and click **Properties**.

5 Click the **Deduplication** tab. The tab includes the following:

Configure	Lets you configure the file system for deduplication. Configuration sets up a small database. The **Configure Deduplication** window includes several customizable parameters and also lets you set up a schedule for running deduplication.
Unconfigure	Lets you remove an existing configuration. If you want to disable deduplication temporarily but leave the configuration in place, you do so on the **Configure Deduplication** window. See "Disabling or removing deduplication for a file system" on page 344.
Scan Now	Lets you start the deduplication operation immediately if there is an existing configuration.
Space Saved History	Tracks the results of deduplication. To change the time period, select from the **Duration** drop-down list and click **Apply**.

6 Click **Configure**.

7 In the **Configure Deduplication** window, you can customize the following options:

Enabled	If you clear the checkbox, the deduplication operation is disabled. If you want to enable it later, return to this window.
Data Usage	Lets you optimize the database size according to the type of data and the amount of space available for the database. The smaller the chunk size selected for data, the more space is required for the database. Once configuration is complete, this parameter cannot be changed except by unconfiguring the database and reconfiguring it. For most data, Veritas recommends the default, **Other (16k)**. Options include **VMDK files (4k)** and **Home directories (8k)**.

8 To set up a schedule for deduplication, select from the following:

Commit on run number	The deduplication process scans and fingerprints the data before eliminating duplicates. You can schedule the deduplication process to eliminate the duplicates each time it runs (the default value of 1) or every specified number of times. During times that deduplication does not occur, the deduplication run only updates the fingerprints in the database.
Weekday Schedule	You can select one day of the week or schedule a run every day. Veritas recommends that you schedule deduplication when the system activity is low so as not to interfere with the regular system workload.
Hours	Schedule the hour to begin a deduplication run.

9 Click **Finish**. The deduplication configuration sets up the deduplication database. When a message shows the configuration is complete, click **Close**.

10 If you want to run deduplication now, rather than wait for a scheduled time, click **Scan Now**. Click **Yes** to confirm that you want to begin the deduplication. Once it is begun, you can close the window. The operation runs in the background.

See "Deduplicating file systems" on page 338.

Starting deduplication for a file system

After you configure deduplication for a file system, you can start the deduplication operation on demand.

To perform this task, your user group must be assigned the Admin role on the host or the Server perspective. The permission on the host may be explicitly assigned or inherited from a parent Organization.

To start deduplication for a file system

1 In the Management Server console, go to the **Server** perspective and expand **Manage** in the left pane.

2 Expand the Organization or **Uncategorized Hosts** to locate the host.

3 Expand the host and select **Volumes**.

4 Right-click the volume on which the file system is mounted and click **Properties**.

5 Click the **Deduplication** tab.

6 Click **Scan Now**.

See “Deduplicating file systems” on page 338.

Disabling or removing deduplication for a file system

If you have set up deduplication for a file system, you can later either temporarily disable or remove deduplication for that file system.

To perform this task, your user group must be assigned the Admin role on the host or the Server perspective. The permission on the host may be explicitly assigned or inherited from a parent Organization.

To disable or remove deduplication for a file system

1 In the Management Server console, go to the **Server** perspective and expand **Manage** in the left pane.

2 Expand the Organization or **Uncategorized Hosts** to locate the host.

3 Expand the host and select **Volumes**.

4 Right-click the volume on which the file system is mounted and click **Properties**.

5 Click the **Deduplication** tab.

6 Choose one of the following:

To disable deduplication	■ Click **Configure**. ■ In the **Configure Duplication** window, clear the check box for **Enabled** and click **Finish**.
To remove deduplication	■ Click **Unconfigure**. ■ Confirm that you want to remove the configuration.

See “Deduplicating file systems” on page 338.

Section 4

Managing high availability and disaster recovery configurations

- Chapter 16. Overview
- Chapter 17. Managing clusters
- Chapter 18. Managing service groups
- Chapter 19. Managing systems
- Chapter 20. Managing VSystems
- Chapter 21. Managing resources
- Chapter 22. Managing global cluster configurations
- Chapter 23. Running fire drills
- Chapter 24. Using recovery plans
- Chapter 25. Managing ApplicationHA
- Chapter 26. Managing application configuration
- Chapter 27. Multi Site Management

- Appendix A. List of high availability operations

Chapter 16

Overview

This chapter includes the following topics:

- About high availability and disaster recovery operations
- Pre-requisites for performing high availability and disaster recovery operations
- About attributes of Availability objects
- About Virtual Business Services

About high availability and disaster recovery operations

Using Veritas InfoScale Operations Manager, you can manage the high availability application clusters based on Cluster Server. The Management Server console helps you perform all the VCS-related high availability and disaster recovery operations such as setting up global cluster (GCO), or bringing a service group or resource online.

You must note the following when you perform the high availability (HA) and disaster recovery (DR) operations using Veritas InfoScale Operations Manager:

- To perform DR operations such as remote online, remote offline, and switch, the remote cluster need not be managed by the same Management Server as the local cluster. However, at least one cluster in the global cluster option (GCO) setup needs to be managed by Management Server. To use Veritas InfoScale Operations Manager for optimal management of clusters in an existing GCO setup, it is recommended that you include the clusters within a single Management Server domain.
- You are prompted to open a closed configuration to perform an operation that requires configuration changes. If you choose to open the configuration

automatically, Veritas InfoScale Operations Manager provides an option to close the configuration after the operation is completed.

- When an operation that involves configuration changes is performed, Veritas InfoScale Operations Manager automatically dumps the configuration at the end of the operation.
- For performing batch operations, you must have the appropriate privileges on all the objects that you have selected. If you do not have appropriate privileges on any of the selected objects, the entire operation does not succeed. In such cases, you need to deselect the objects where you do not have the privileges, and launch the wizard again to perform the operation.

See "About managing clusters" on page 350.

See "About managing resources" on page 400.

See "About high availability and disaster recovery readiness" on page 424.

Pre-requisites for performing high availability and disaster recovery operations

You must ensure the following before performing high availability and disaster recovery operations :

- The Veritas InfoScale Operations Manager 3.0, or later managed host package must be installed. It is recommended to install Veritas InfoScale Operations Manager 7.2 managed host package to get the latest features.
- The managed host must be managed by a Management Server.
- Cluster Server must be configured and running on the managed host.

For best results, it is recommended that all the nodes in the cluster must be managed by Management Server. However, this is not mandatory.

See "About high availability and disaster recovery operations" on page 347.

About attributes of Availability objects

Cluster Server components are configured using the attributes of these components. Attributes contain data about the cluster, systems, service groups, resources, resource types, agent, and heartbeats. For example, the value of a service group's SystemList attribute specifies on which systems the group is configured and the priority of each system within the group. Each attribute has a definition and a value. Attributes also have default values assigned when a value is not specified.

In the Management Server console, you can view the attributes of the Availability objects such as clusters, service groups, systems, resources, and resource types. You can edit the values of the attributes that have an edit icon against the attribute-name.

Table 16-1 Attribute type and description

Attribute type	Description
Attribute name in Bold	Indicates that the attribute is a must-configure attribute. You must set the value of such an attribute.
Attribute name in Italics	Indicates that the attribute is an important attribute. An important attribute may not be a must-configure attribute.
Attribute name in Bold and Italics	Indicates that the attribute is an important and must-configure attribute.
Attribute name with an edit icon	Indicates that the attribute is an editable attribute.

See "Editing attributes of a cluster" on page 352.

See "Editing attributes of service groups" on page 385.

See "Editing attributes of a system" on page 393.

See "Editing attributes of a resource" on page 408.

See "Editing attributes of a resource type" on page 409.

About Virtual Business Services

The Virtual Business Services feature provides visualization, orchestration, and reduced frequency and duration of service disruptions for multi-tier business applications running on heterogeneous operating systems and virtualization technologies. A virtual business service represents the multi-tier application as a consolidated entity that helps you manage operations for a business service. It builds on the high availability and disaster recovery provided for the individual tiers by Veritas InfoScale products such as Cluster Server.

Application components that are managed by Cluster Server or Microsoft Failover Clustering can be actively managed through a virtual business service.

You can use the Veritas InfoScale Operations Manager Management Server console to create, configure, and manage virtual business services.

For more information on Virtual Business Services see the *Virtual Business Service–Availability User's Guide*.

Chapter 17

Managing clusters

This chapter includes the following topics:

- About managing clusters
- Opening a cluster configuration
- Saving a cluster configuration
- Closing a cluster configuration
- Editing attributes of a cluster
- Importing a type definition

About managing clusters

Clusters manage your highly available applications as redundant, identically-configured service groups. These service groups are distributed among the cluster nodes and can fail over from one node to another.

The cluster management views of the Veritas InfoScale Operations Manager console enable you to manage individual clusters through a single console. Currently, you can perform tasks such as open configuration, save configuration and close configuration on individual clusters.

See “Opening a cluster configuration” on page 351.

See “Saving a cluster configuration” on page 351.

See “Closing a cluster configuration” on page 352.

See “Starting the Cluster Server high availability daemon on the hosts in a cluster” on page 393.

See “Stopping the Cluster Server high availability daemon on the systems in a cluster” on page 394.

See “Importing a type definition” on page 354.

Opening a cluster configuration

In the Management Server console, you can make changes to a cluster configuration file by opening the configuration. You can perform this task only when the cluster configuration is closed.

To perform this task, your user group must be assigned the Admin role on the cluster or the Availability perspective. The permission on the cluster may be explicitly assigned or inherited from a parent Organization.

To open a cluster configuration

1. In the Management Server console, go to the **Availability** perspective and select **Manage** in the left pane.
2. Expand the Organization, or **Uncategorized Clusters** to locate the cluster.
3. Right-click on the required cluster and select **Configuration** > **Open**.
4. In the **Open configuration** panel, click **OK**.
5. In the result panel, click **OK**.

See “About managing clusters” on page 350.

See “Saving a cluster configuration” on page 351.

See “Closing a cluster configuration” on page 352.

Saving a cluster configuration

In the Management Server console, you can make changes to a cluster configuration and save the updated configuration to disk.

To perform this task, your user group must be assigned the Admin role on the cluster or the Availability perspective. The permission on the cluster may be explicitly assigned or inherited from a parent Organization.

To save a cluster configuration

1. In the Management Server console, go to the **Availability** perspective and select **Manage** in the left pane.
2. Expand the Organization, or **Uncategorized Clusters** to locate the cluster.
3. Right-click on the required cluster and select **Configuration** > **Save**.
4. In the **Save configuration** panel, click **OK**.
5. In the result panel, click **OK**.

See “About managing clusters” on page 350.

See “Opening a cluster configuration” on page 351.

See “Closing a cluster configuration” on page 352.

Closing a cluster configuration

In the Management Server console, you can close a cluster configuration to make sure that no changes are made to it.

To perform this task, your user group must be assigned the Admin role on the cluster or the Availability perspective. The permission on the cluster may be explicitly assigned or inherited from a parent Organization.

To close a cluster configuration

1. In the Management Server console, go to the **Availability** perspective and select **Manage** in the left pane.
2. Expand the Organization, or **Uncategorized Clusters** to locate the cluster.
3. Right-click on the required cluster and select **Configuration** > **Close**.
4. In the **Close configuration** panel, click **OK**.
5. In the result panel, click **OK**.

See “About managing clusters” on page 350.

See “Opening a cluster configuration” on page 351.

See “Saving a cluster configuration” on page 351.

Editing attributes of a cluster

The Edit attribute operation lets you configure the policy behavior of various objects in Cluster Server. This operation also helps you define the attribute values that enable Cluster Server monitor different resources for high availability.

To perform this task, your user group must be assigned the Admin role on the cluster or the Availability perspective. The permission on the cluster may be explicitly assigned or inherited from a parent Organization.

Note: You can edit only the editable attributes of the cluster. The editable attributes have an edit icon along with the attribute name in the list of attributes.

To edit attributes of a cluster

1. In the Management Server console, go to the **Availability** perspective and select **Manage** in the left pane.
2. Expand the Organization or **Uncategorized Clusters** to locate the cluster.
3. Right-click the cluster and select **Properties**.
4. Select the **Attributes** tab.
5. Right-click the attribute that you want to edit, and select **Edit Attribute**.
6. Edit the values in the fields. Click **OK**.

   See "Edit attribute options" on page 353.
7. In the **Result** page, click **OK**.

See "About managing clusters" on page 350.

See "Saving a cluster configuration" on page 351.

Edit attribute options

Use this wizard panel to edit attributes of a cluster, host, service group, resources, or resource types.

Table 17-1 Edit attribute panel options

Field	Description
Attribute	Displays the attribute that you want to edit. This field is not editable.
Apply to	Determines the scope of the attribute. Typically, this option helps you assign the same attribute values for all systems, or different values to different systems. You cannot toggle between the scopes for the service group attributes.

Table 17-1 Edit attribute panel options *(continued)*

Field	Description
Attribute Value	Enter the new value for the attribute based on the dimension and click **OK**. ■ For scalar (Single) attributes, enter the value in the field. ■ For vector or key-list attributes, enter the values in the fields. Click **Add** to add the attribute value. To remove a value, select the value from the value list and click **Remove**. ■ For the attributes for associations, enter the key and values in the fields and click **Add**. To delete an attribute value, select the value from the value table and click the delete option.

See "Editing attributes of a cluster" on page 352.

See "Editing attributes of a system" on page 393.

See "Editing attributes of service groups" on page 385.

See "Editing attributes of a resource" on page 408.

See "Editing attributes of a resource type" on page 409.

See "Configure Resources options" on page 361.

Importing a type definition

In the Management Server console, you can import custom resource types definition into the configuration. It is commonly used when any ISV agent is installed on the VCS cluster to import the agents type definition. On Windows, this type definition is imported by default. On Linux, the administrator needs to perform this action.

To import a type definition

1. In the Management Server console, go to the **Availability** perspective and select **Manage** in the left pane.
2. Expand the Organization, or **Uncategorized Clusters** to locate the cluster.
3. Right-click on the required cluster and select **Import Type Definition**.
4. In the **Import Type Definition** wizard panel, do the following:
   - Select the cluster node to import the type definition.

- Provide the agent installation path on the cluster node where the type definition file is present. It is expected that the agent is already installed on the cluster node and the type definition file is available.

5 Click **Finish**.

See “About managing clusters” on page 350.

Chapter 18

Managing service groups

This chapter includes the following topics:

- About managing service groups
- Creating service groups
- Enabling service groups
- Disabling service groups
- Creating Atleast Count dependencies for a resource in a service group
- About Atleast Count dependency
- Autoenabling service groups
- Freezing service groups
- Unfreezing service groups
- Flushing service groups
- Enabling all resources of service groups
- Disabling all resources of service groups
- Deleting service groups
- About linking service groups in a cluster
- Linking service groups in a cluster
- Unlinking service groups
- About site aware service group operations
- Prerequisites for using site-related service group operations

- Limitations of site-related service group operations
- Bringing service groups online
- Taking service groups offline
- Switching service groups
- Clearing faults on service group
- Clearing the resources in a service group from the Admin Wait state
- Editing attributes of service groups
- Modifying the system list for a service group
- About dependency views
- Viewing the service group dependency view
- About modifying a service group
- About Cluster Server service group alerting and failover reporting
- Viewing VCS Failover Duration report

About managing service groups

A service group is a virtual container that contains all the hardware and the software resources that are required to run the managed application. Service groups allow VCS to control all the hardware and the software resources of the managed application as a single unit. When a failover occurs, resources do not fail over individually; the entire service group fails over. If more than one service group is on a system, a group can fail over without affecting the others.

In the Veritas InfoScale Operations Manager console, the **Service Groups** tab in a cluster view, or the **Service Groups** view enables you to do the following:

- Collectively manage multiple service groups
- Individually manage a single service group

See “Enabling service groups” on page 363.

See “Freezing service groups” on page 367.

See “Enabling all resources of service groups” on page 370.

See “Flushing service groups” on page 369.

See “Clearing faults on service group” on page 383.

See “Linking service groups in a cluster” on page 372.

See “Autoenabling service groups” on page 367.

See “Editing attributes of service groups” on page 385.

See “Creating service groups” on page 358.

See “Converting local service groups to global service groups” on page 418.

See “Converting global service groups to local service groups” on page 420.

See “Modifying the system list for a service group” on page 386.

Creating service groups

In the Management Server console, you can create a service group that is associated with a cluster.

To perform this task, your user group must be assigned the Admin role on the cluster or the Availability perspective. The permission on the service group may be explicitly assigned or inherited from a parent Organization or cluster.

To create a service group

1. In the Management Server console, go to the **Availability** perspective and select **Manage** in the left pane.
2. Expand the Organization or **Uncategorized Clusters** to locate the cluster.
3. Right-click on the required cluster and select **Create Service Group**.
4. In the **Create Service Group** panel, do one of the following:
   - To add a service group and configure it, enter the required information and click **Next**.
   - To add a service group and configure it at a later time, enter the required information and click **Finish**.

   See “Create Service Group options” on page 359.
5. In the **Configure System List** panel, do one of the following:
   - To configure the system list and configure the resources, enter the required information and click **Next**.
   - To configure the system list and to configure resources at a later time, enter the required information and click **Finish**.

   See “Configure System List options” on page 360.
6. In the **Configure Resources** panel, do one of the following:

- To configure the resources and resource dependencies, enter the required information and click **Next**.
- To configure the resources and to configure resource dependencies at a later time, enter the required information and click **Finish.**

See “Configure Resources options” on page 361.

7 In the **Resource Dependencies** panel, enter the required information and click **Finish**.

- To configure the resource dependencies, enter the required information and click **Next**.
- To configure the resource dependencies at a later time, click **Finish**.

8 In the **Result** panel, click **OK**.

See “About managing service groups” on page 357.

See “Deleting service groups” on page 371.

See “About modifying a service group” on page 389.

See “Adding or modifying resources” on page 411.

Create Service Group options

Use this wizard panel to create a service group in a cluster.

Table 18-1 Create Service Group wizard panel options

Field	Description
Name	Enter a name for the service group. The name must be unique among the service groups in the cluster. The name that you enter must start with a letter. The name of the service group can contain only letters, numbers, hyphens, or underscores.

Table 18-1 Create Service Group wizard panel options *(continued)*

Field	Description
Type	Select the type of the service group. The available options are : ■ **Failover** - Runs on one system at a time. Failover groups are used for most applications that do not support multiple systems to simultaneously access the application's data. ■ **Parallel** - Runs simultaneously on more than one system in the cluster. Parallel service groups are appropriate for the applications that manage multiple application instances running simultaneously without data corruption. ■ **Hybrid** - A combination of the failover and parallel service groups. It behaves as a failover group within a system zone and a parallel group across system zones. The hybrid service groups are used for replicated data clusters.

See "Creating service groups" on page 358.

Configure System List options

Use this wizard panel to specify the systems where you want to run the service group. You can use the arrow options to move systems between the columns, or change the priority for the systems.

Table 18-2 Configure System List wizard panel options

Column	Description
Available Systems	Lists the available systems that you can use for running the service group.

Table 18-2 Configure System List wizard panel options *(continued)*

Column	Description
Systems in Priority Order	Lists the systems that you have selected, in the order of your priority. The following option are available: ■ **Autostart**: On selecting the Autostart check box, a service group automatically restarts after a faulted persistent resource becomes online. ■ **Priority**: Enables you to prioritize each system within a service group. The systems have a default value assigned when they are added to the **Systems in Priority Order** option. You can also edit the assigned default value.

See "Creating service groups" on page 358.

See "Modifying the system list for a service group" on page 386.

Configure Resources options

Use this wizard panel to configure the resources for a service group. You can add multiple resources to a service group. Click the **Add** option to add the resource to the service group.

Table 18-3 Configure Resources wizard panel options

Field	Description
Name	Enter a name for the resource. A resource name should be unique in a cluster. The name that you enter must start with a letter. The name of the resource can contain only letters, numbers, hyphens, or underscores.
Type	Select a type from the drop-down field.
Enabled	Select this check box if you want to enable the resource.
Critical	Select this check box if you want to mark the resource as critical.

Table 18-3 Configure Resources wizard panel options *(continued)*

Field	Description
Edit resource attribute	Click this option to edit the resource attributes. In the **Edit Attributes** page, you can see the important and must-configure attributes in the list. Select the attribute that you want to edit, and enter the required information in the fields. See "Edit attribute options" on page 353.
Delete resource	Click this option to delete the resource.

See "Creating service groups" on page 358.

See "Adding or modifying resources" on page 411.

Resource Dependencies options

Use this wizard panel to create resource dependencies by linking the resources. Click **Link** to link the parent resource to the child resource.

Use the **Configure Atleast Dependencies** checkbox to configure Atleast count dependency for the child resources.

Note: Multiple selection of child resources for Atleast dependency is allowed only when you check the **Configure Atleast Dependencies** checkbox.

Table 18-4 Resource Dependencies panel options

Field	Description
Select Parent	Select the parent resource from the drop-down list.
Select Child	Select the child resource from the list.
Delete	Click this option to delete the resource dependency.

Table 18-4 Resource Dependencies panel options *(continued)*

Field	Description
Atleast Count	Select the dependency count from the drop-down list to set a minimum number of child resources to be online before the parent comes online. For example, if an application depends on three IPs and if this application has to be brought online or has to remain online, at least two IPs must be online to enable the application to be online.

See "Creating service groups" on page 358.

See "Linking resources in a service group" on page 410.

See "Unlinking resources in a service group" on page 410.

See "Adding or modifying resources" on page 411.

Enabling service groups

In the Management Server console, you can enable the disabled service groups so that they can be brought online. You can enable a service group only if it is in a disabled state on a cluster host.

To perform this task, your user group must be assigned the Admin role on the service group or the Availability perspective. The permission on the service group may be explicitly assigned or inherited from a parent Organization or cluster.

Note: Before you perform this task, you need to open the cluster configuration if the configuration is not already open.

To enable service groups

1. In the Management Server console, go to the **Availability** perspective and select **Manage** in the left pane.
2. Expand the Organization or **Uncategorized Clusters** to locate the cluster.
3. Expand the cluster and then the **Service Groups** node to locate the service group.
4. Right-click on the service group and select **Edit** > **Enable**.
5. In the **Enable Service Group** panel, do the following:

- Select the system to enable the service group on. Choose **All Systems** to enable the service group on all systems.
- Click **OK**.

6 In the **Result** panel, click **OK**.

See "About managing service groups" on page 357.

See "Autoenabling service groups" on page 367.

See "Disabling service groups" on page 364.

See "Enabling resources" on page 400.

See "Enabling all resources of service groups" on page 370.

Disabling service groups

In the Management Server console, you can disable the service groups to prevent them from being brought online. You can use this feature to temporarily prevent agents from monitoring service groups on a host while they undergo maintenance operations. You can disable a service group only if it is in an enabled state on a cluster host.

To perform this task, your user group must be assigned the Admin role on the service group or the Availability perspective. The permission on the service group may be explicitly assigned or inherited from a parent Organization or cluster.

Note: Before you perform this task, you need to open the cluster configuration if the configuration is not already open.

To disable service groups

1 In the Management Server console, go to the **Availability** perspective and select **Manage** in the left pane.

2 Expand the Organization or **Uncategorized Clusters** to locate the cluster.

3 Expand the cluster and then the **Service Groups** node to locate the service group.

4 Right-click on the service group and select **Edit** > **Disable**.

5 In the **Disable Service Group** panel, do the following:

- To disable a single service group, select the system to disable the service group on. Choose **All Systems** to disable the service group on all the systems.

- Click **OK**.

6 In the **Result** panel, click **OK**.

See "About managing service groups" on page 357.

See "Enabling service groups" on page 363.

See "Autoenabling service groups" on page 367.

See "Disabling resources" on page 401.

See "Disabling all resources of service groups" on page 371.

Creating Atleast Count dependencies for a resource in a service group

In the Management Server console, you can create dependencies (link) between two resources of a service group. When you create a link between two resources, you need to designate one resource as the parent resource and the other one as a child resource. The child resource must come online before the parent resource.

In an Atleast Count dependency, the parent resource depends on a set of child resources. You can configure a minimum number of child resources for the parent resource. The configured child resources must be online for the parent to be brought online or to remain online.

The parent resource of one dependency can be the child resource of another dependency. In a service group, several parent-child dependencies exist to support a single application resource. In a dependency diagram, this application resource occupies the apex of the diagram.

To perform this task, your user group must be assigned the Admin role on the service group or the Availability perspective. The permission on the service group may be explicitly assigned or inherited from a parent Organization or cluster.

To create Atleast Count dependency for a resource in a service group

1 In the Management Server console, go to the **Availability** perspective and select **Manage** in the left pane.

2 Expand the Organization or **Uncategorized Clusters** to locate the cluster.

3 Right-click on the required cluster and select **Create Service Group**.

4 In the **Create Service Group** panel, do one of the following:

- To add a service group and configure it, enter the required information and click **Next**.

- To add a service group and configure it at a later time, enter the required information and click **Finish**.

See “Create Service Group options” on page 359.

5 In the **Configure System List** panel, do one of the following:

- To configure the system list and configure the resources, enter the required information and click **Next**.
- To configure the system list and to configure resources at a later time, enter the required information and click **Finish**.

See “Configure System List options” on page 360.

6 In the **Configure Resources** panel, do one of the following:

- To configure the resources and resource dependencies, enter the required information and click **Next**.
- To configure the resources and to configure resource dependencies at a later time, enter the required information and click **Finish.**

See “Configure Resources options” on page 361.

7 In the **Resource Dependencies** panel, select the parent and the child resources.

8 Check the **Configure Atleast Dependencies** check box for selecting the child resources that need to be configured.

Note: The configured child resources need to be online for the parent to be online.

---

9 Select the Child resources, and click **Link**.

10 Select the number of child resources that needs to be set for Atleast Count dependency from the **Atleast Count** drop-down option.

11 In the **Result** panel, click **OK**.

About Atleast Count dependency

In an Atleast Count dependency, the parent resource depends on a set of child resources. The parent resource is brought online or can remain online only if the minimum number of child resources in this resource set is online.

For example, the parent resource (res1) depends on the child resources (res2, res3, res4, res5, and res6). When Atleast dependency is set at two, the parent resource can be brought online only when two or more resources come online and

the parent resource can remain online only until two or more child resources are online.

Autoenabling service groups

In the Management Server console, a service group is autodisabled until Veritas Cluster Server (VCS) probes all resources and checks that they are ready to come online. You can autoenable a service group in situations where the VCS engine is not running on one of the systems in the cluster. In this case, you must change the disabled state of the service group to enable the group on another system in the cluster.

To perform this task, your user group must be assigned the Admin role or the Operator role on the service group or the Availability perspective. The permission on the service group may be explicitly assigned or inherited from a parent Organization or cluster.

To autoenable a service group

1. In the Management Server console, go to the **Availability** perspective and select **Manage** in the left pane.
2. Expand the Organization or **Uncategorized Clusters** to locate the cluster.
3. Expand the cluster and then the **Service Groups** node to locate the service group.
4. Right-click on the service group and select **Edit** > **Autoenable**.
5. In the **Autoenable Service Group** panel, select the system to autoenable the service group on. Click **OK**.
6. In the **Result** panel, click **OK**.

See “About managing service groups” on page 357.

See “Enabling service groups” on page 363.

See “Disabling service groups” on page 364.

See “Enabling resources” on page 400.

Freezing service groups

In the Management Server console, you can freeze the service groups. Freezing is useful for performing maintenance tasks on service groups. You can freeze a service group either temporarily or persistently. The temporarily frozen service group automatically unfreezes when Cluster Server gets restarted. The persistently frozen service group needs to be explicitly unfrozen.

To perform this task, your user group must be assigned the Admin role or Operator role on the service group or the Availability perspective. The permission on the service group may be explicitly assigned or inherited from a parent Organization or cluster.

Note: To freeze the service groups persistently, your user group must be assigned the Admin role.

To freeze service groups

1. In the Management Server console, go to the **Availability** perspective and select **Manage** in the left pane.
2. Expand the Organization or **Uncategorized Clusters** to locate the cluster.
3. Expand the cluster and then the **Service Groups**

   node to locate the service group.
4. Right-click on the service group and select **Freeze**.
5. In the **Freeze Service Group** panel, do the following:
   - Select **Freeze Persistently** to ensure that the service groups remain frozen unless explicitly unfrozen.
   - Click **OK**.
6. In the **Result** panel, click **OK**.

See "About managing service groups" on page 357.

See "Unfreezing service groups" on page 368.

Unfreezing service groups

In the Management Server console, you can unfreeze the frozen service groups.

To perform this task, your user group must be assigned the Admin role or Operator role on the service group or the Availability perspective. The permission on the service group may be explicitly assigned or inherited from a parent Organization or cluster.

Note: To unfreeze the service groups that have been frozen persistently, your user group must be assigned the Admin role.

To unfreeze service groups

1. In the Management Server console, go to the **Availability** perspective and select **Manage** in the left pane.
2. Expand the Organization or **Uncategorized Clusters** to locate the cluster.
3. Expand the cluster and then the **Service Groups** node to locate the service group.
4. Right-click on the service group and select **Unfreeze**.
5. In the **Unfreeze Service Group** panel, click **OK**.
6. In the **Result** panel, click **OK**.

See "About managing service groups" on page 357.

See "Freezing service groups" on page 367.

Flushing service groups

In the Management Server console, you can address resource-related problems if they occur while a service group comes online or goes offline. As a service group comes online or goes offline, the resources in the group also come online or go offline. If the online or the offline operation hangs on a particular resource, flush the service group to clear the Waiting to go online or Waiting to go offline states from its resources. Flushing the service group typically leaves the cluster in a partial state. After you complete this process, resolve the issue with the particular resource and then proceed with bringing the service group online or taking it offline.

Note: The flush operation does not halt the resource operations (such as online or offline) that are already running. If a running operation succeeds after a flush command is executed, the resource state might change depending on the operation.

To perform this task, your user group must be assigned the Admin role or the Operator role on the service group or the Availability perspective. The permission on the service group may be explicitly assigned or inherited from a parent Organization or cluster.

To flush a service group

1. In the Management Server console, go to the **Availability** perspective and select **Manage** in the left pane.
2. Expand the Organization or **Uncategorized Clusters** to locate the cluster.
3. Expand the cluster and then the **Service Groups** node to locate the service group.

4 Right-click on the service group and select **Flush**.

5 In the **Flush Service Group** panel, select the system to flush the service group on. Click **OK**.

6 In the **Result** panel, click **OK**.

See "About managing service groups" on page 357.

Enabling all resources of service groups

In the Management Server console, you can enable all the resources of a service group. You need to enable the resources of a service group before bringing them online.

To perform this task, your user group must be assigned the Admin role on the service group or the Availability perspective. The permission on the service group may be explicitly assigned or inherited from a parent Organization or cluster.

Note: You must have resources configured for a service group before you enable them. Before you enable the resources of a service group, you must open the configuration of the cluster that it belongs to.

To enable all resources of service groups

1 In the Management Server console, go to the **Availability** perspective and select **Manage** in the left pane.

2 Expand the Organization or **Uncategorized Clusters** to locate the cluster.

3 Expand the cluster and then the **Service Groups** node to locate the service group.

4 Right-click on the service group and select **Edit** > **Resources** > **Enable All**.

5 In the **Enable All Resources** panel, click **OK**.

6 In the **Result** panel, click **OK**.

See "About managing service groups" on page 357.

See "Disabling all resources of service groups" on page 371.

See "Enabling service groups" on page 363.

See "Disabling service groups" on page 364.

See "Enabling resources" on page 400.

See "Disabling resources" on page 401.

Disabling all resources of service groups

In the Management Server console, you can prevent the resources of a service group from being brought online by disabling them.

To perform this task, your user group must be assigned the Admin role on the service group or the Availability perspective. The permission on the service group may be explicitly assigned or inherited from a parent Organization or cluster.

Note: All the resources of a service group must be taken offline before they are disabled. Before you disable the resources of a service group, you must open the configuration of the cluster that it belongs to.

To disable all resources of service groups

1. In the Management Server console, go to the **Availability** perspective and select **Manage** in the left pane.
2. Expand the Organization or **Uncategorized Clusters** to locate the cluster.
3. Expand the cluster and then the **Service Groups** node to locate the service group.
4. Right-click on the service group and select **Edit** > **Resources** > **Disable All**.
5. In the **Disable All Resources** panel, click **OK**.
6. In the **Result** panel, click **OK**.

See "About managing service groups" on page 357.

See "Enabling all resources of service groups" on page 370.

See "Enabling service groups" on page 363.

See "Disabling service groups" on page 364.

See "Enabling resources" on page 400.

See "Disabling resources" on page 401.

Deleting service groups

In the Management Server console, you can delete one or more service groups that are no longer required. Before you delete any service group, you need to unlink the dependencies for the service group.

To perform this task, your user group must be assigned the Admin role on the service group or the Availability perspective. The permission on the service group may be explicitly assigned or inherited from a parent Organization or cluster.

To delete service groups

1. In the Management Server console, go to the **Availability** perspective and select **Manage** in the left pane.
2. Expand the Organization or **Uncategorized Clusters** to locate the cluster.
3. Expand the cluster and then the **Service Groups** node to locate the service group.
4. Right-click on the service group and select **Delete**.
5. In the **Delete Service Group** panel, click **OK**.
6. In the **Result** panel, click **OK**.

See "About managing service groups" on page 357.

See "Creating service groups" on page 358.

See "Linking service groups in a cluster" on page 372.

See "Modifying the system list for a service group" on page 386.

About linking service groups in a cluster

Service groups are linked in a cluster to ensure that they go online in the correct order. Service groups are linked two at a time. A service group often depends upon the resources and the run status of another service group.

The service group that must come online first is called the child service group. The service group that must come online after the child is called the parent service group.

A service group can be the child of one service group and the parent of another service group at the same time. In a complex service group, several parent-child dependencies exist to ultimately support a single application service group.

See "Linking service groups in a cluster" on page 372.

Linking service groups in a cluster

In the Management Server console, you can link a dependent pair of service groups in a cluster.

To perform this task, your user group must be assigned the Admin role on the service group or the Availability perspective. The permission on the service group may be explicitly assigned or inherited from a parent Organization or cluster.

Note: Before you perform this task, you need to open the cluster configuration if the configuration is not already open.

To link two service groups in a cluster

1. In the Management Server console, go to the **Availability** perspective and select **Manage** in the left pane.
2. Expand the Organization or **Uncategorized Clusters** to locate the cluster.
3. Expand the cluster and then the **Service Groups** node to locate the service group.
4. Right-click on the service group and select **Edit** > **Link**.
5. In the **Link Service Group** panel, specify the details of the dependencies and click **Finish**.

   See "Link Service Group options" on page 373.
6. In the **Result** panel, click **OK**.

See "About managing service groups" on page 357.

See "About linking service groups in a cluster" on page 372.

See "Unlinking service groups" on page 375.

See "Deleting service groups" on page 371.

Link Service Group options

Use this wizard panel to add the selected service group as a child group to a parent service group.

Table 18-5 Link Service Group panel options

Field	Description
Child Group	Select the service group that you intend to define as the child service group.

Table 18-5 Link Service Group panel options *(continued)*

Field	Description
Relationship	Select the relationship that defines the relationship of the parent group with the state of the child group. **Online local** -The parent group must wait for the child group to be brought online before it can start. **Offline local** - The parent group can be started only if the child group is offline, and vice versa. This relationship prevents conflicting applications from running on the same system. You cannot specify any dependency type if you are using this relationship. **Online Remote** - An instance of parent group depends on one or more instances of the child group being online on any system other than the system on which the parent is online. In this relationship, you cannot specify the dependency type as Hard. **Online Global**- The child group must be online somewhere in the cluster before the parent group can come online. In this relationship, you cannot specify the dependency type as Hard.

Table 18-5 Link Service Group panel options *(continued)*

Field	Description
Dependency Type	Select the type of dependency that defines the rigidity of the link between the parent and the child groups. The available options are: **Firm** - The child group must be online before the parent group is brought online; the location of the dependency determines where the child group must be online. **Soft** - The child group must be online before the parent group being brought online; the location of the dependency determines where the child group must be online. For example, in an online-local- soft dependency, an instance of the child group must be online on the same system before the parent group can come online. **Hard** - The child and the parent groups fail over to the same system together when either the child or the parent faults.

See "Linking service groups in a cluster" on page 372.

Unlinking service groups

In the Management Server console, you can unlink a dependent pair of service groups in a cluster.

To perform this task, your user group must be assigned the Admin role on the service group or the Availability perspective. The permission on the service group may be explicitly assigned or inherited from a parent Organization or cluster.

Note: Before you perform this task, you need to open the cluster configuration if the configuration is not already open.

To unlink two service groups in a cluster

1. In the Management Server console, go to the **Availability** perspective and select **Manage** in the left pane.
2. Expand the Organization or **Uncategorized Clusters** to locate the cluster.

3 Expand the cluster and then the **Service Groups** node to locate the service group.

4 Right-click on the service group and select **Edit** > **Unlink**.

5 In the **Unlink Service Group** panel, select the parent group from which you want to unlink the service group. Click **OK**.

6 In the **Result** panel, click **OK**.

See "About managing service groups" on page 357.

See "About linking service groups in a cluster" on page 372.

See "Linking service groups in a cluster" on page 372.

See "Deleting service groups" on page 371.

About site aware service group operations

While performing service group operations in a campus cluster, you can provide site-related information. With site awareness, the administrator can fail over a service group to another site in an event of failure at the local site. The feature is supported for the following service group operations:

- Online
- Offline
- Switch

Use the **Availability** perspective of the Management Server console to specify the site or the system to fail over the selected service group. For example, you can bring a selected service group online on any system of site A, or take a service group offline on all the systems of site B.

You can select single or multiple service groups for these operations. However, note that the site-related operations are applicable only when the selected service groups are configured on the same stretch cluster; service groups from stretch and non-stretch clusters will not be supported. The operation is supported for parallel, failover, and hybrid service groups.

Note: The site-related operation uses the multi-site management feature. It does not use the `SystemZones` attribute of Cluster Server (VCS) service group.

See "Prerequisites for using site-related service group operations" on page 377.

See "Limitations of site-related service group operations" on page 377.

Prerequisites for using site-related service group operations

The site-related operations for the configured service groups are available only when:

- The selected service group is a part of a campus cluster.
- You have already provided the site name (referred to as site tagging) to the cluster nodes where the service groups are configured. You need to use the **Configure Stretch Sites** wizard of the multi-site management feature to assign the site tag to the cluster nodes.
  See "Configuring stretch sites" on page 463.
- You use the 6.1 version of Cluster Server, Veritas InfoScale Operations Manager Management Server and the managed hosts.

See "About site aware service group operations" on page 376.

Limitations of site-related service group operations

This section lists the limitations of site-related service group operations:

- When multiple service groups are selected from stretch and non-stretch clusters, the site option is not available to the user.
- Propagate and Force options are not available for multiple service group operations.

See "About site aware service group operations" on page 376.

Bringing service groups online

Use the Management Server console to manually put one or more service groups in a functioning state (referred to as online) on a specific system, any system of a selected site, or any available system in the cluster. The site awareness lets you fail over the service groups to another site in an event of failure at the local site. The online operation is applicable to failover, parallel, and hybrid service groups, or their combination (applicable to multiple service groups selection).

To perform this task, your user group must be assigned the Admin role or Operator role on the service group or the Availability perspective. The permission on the service group may be explicitly assigned or inherited from a parent Organization or cluster.

To bring service groups online

1. In the Management Server console, go to the **Availability** perspective and select **Manage** in the left pane.
2. Expand the Organization or **Uncategorized Clusters** to locate the cluster.
3. Expand the cluster and then the **Service Groups** node to locate the service group.
4. Under the **Service Groups** tab, select the required service group. Press Ctrl or Shift for the selection of multiple service groups.
5. Right-click and select **Online**.

   See "Online service groups panel options" on page 378.
6. In the **Result** panel, click **OK**.

See "About site aware service group operations" on page 376.

Online service groups panel options

Use this panel to provide site and systems (cluster nodes) related information to bring one or more service groups online.

For single service group:Table 18-6

For multiple service groups:Table 18-7

Table 18-6 Online single service group panel options

Field	Description
Select the system you want to online service group on:	Select the required system to bring the selected service group online. A parallel service group can fail over to any specific system, all systems in the cluster, or all systems from the selected site. A hybrid service group can fail over to any specific system, all systems in the cluster, or any system from the selected site. A failover service group can fail over to any specific system, any system in the cluster, or any system from the selected site.

Table 18-6 Online single service group panel options *(continued)*

Field	Description
Force (enforce a takeover if the cluster holding authority is down)	Select the check box to enforce that the global service group fails over on another available cluster. If this cluster does not have the required authority to bring the global service group online, the **Force** option assigns the required authority to that cluster; thereby ensuring successful failover of the global service group. Authority is a persistent service group attribute, and it designates which cluster has the right to bring a global service group online.
Propagate (All of its required child groups are also brought online)	Select the check box to bring all child service groups online along with the parent service group.

Table 18-7 Online multiple service groups panel options

Field	Description
Do you want to online the following service group(s)?	For multiple service groups operation (with any combination of failover, parallel or hybrid service groups), you can fail over the service groups to any of the available systems in the cluster or any available site. Select the appropriate option (**Any System** or the site name) to bring the service groups online.

Note: For the online operation on multiple service groups, though Veritas InfoScale Operations Manager displays all sites for the selected service groups, it is important to note that some sites may not be applicable to a service group or groups. The online operation will be successful only for those sites which are applicable to the selected service groups. The online operation is ignored for the remaining service groups.

See "Bringing service groups online" on page 377.

Taking service groups offline

Use the Management Server console to take one or more service groups offline on a specific system, any system of a selected site, or any available system in the cluster.

To perform this task, your user group must be assigned the Admin role or Operator role on the service group or the Availability perspective. The permission on the service group may be explicitly assigned or inherited from a parent Organization or cluster.

To take service groups offline

1. In the Management Server console, go to the **Availability** perspective and select **Manage** in the left pane.
2. Expand the Organization or **Uncategorized Clusters** to locate the cluster.
3. Expand the cluster and then the **Service Groups** node to locate the service group.
4. Under the **Service Groups** tab, select the required service group. Press Ctrl or Shift for the selection of multiple service groups.
5. Right-click and select **Offline**.

   See "Offline service groups panel options" on page 380.
6. In the **Result** panel, click **OK**.

See "About site aware service group operations" on page 376.

Offline service groups panel options

Use this panel to provide site and systems (cluster nodes) related information to take one or more service groups offline.

For single service group:Table 18-8

For multiple service groups:Table 18-9

Table 18-8 Offline single service group panel options

Field	Description
Select the system you want to offline service group	Select the required cluster node (system) where you want to take the selected service group offline. A parallel service group can fail over to any specific system, all systems in the cluster, or all systems from the selected site. A hybrid service group can fail over to any specific system, or all systems in the cluster. A failover service group can fail over to any specific system.

Table 18-8 Offline single service group panel options *(continued)*

Field	Description
Force	Select the check box to force the service group to go offline even if its resource is waiting to complete the probe operation.
If probed	Select the check box to take the service group offline only if the probe operation for the service group's resources is already completed.
Propagate	Select the check box to take all required parent service groups (that have hard or firm dependency) offline along with the selected service group.

Table 18-9 Offline multiple service groups panel options

Field	Description
Select the system you want to offline service groups	Select the required cluster node (system) where you want to take the service groups offline. Depending on the combination of the service group types (failover, parallel, or hybrid), you can select a specific system or a site.

Note: For the offline operation on multiple service groups, though Veritas InfoScale Operations Manager displays all sites for the selected service groups, it is important to note that some sites may not be applicable to a service group or groups. The offline operation will be successful only for those sites which are applicable to the selected service groups. The offline operation is ignored for the remaining service groups.

See "Taking service groups offline" on page 379.

Switching service groups

Use the Management Server console to switch one or more service groups to a specific system, any system of a selected site, or any available system in the cluster.

Note: The switch operation is not supported for the parallel service groups. It is available only for failover and hybrid service groups with the site option being available only for the failover service groups. Site option is not available for the hybrid service group (for single and collective service groups operations).

To perform this task, your user group must be assigned the Admin role or the Operator role on the service group or the Availability perspective. The permission on the service group may be explicitly assigned or inherited from a parent Organization or cluster.

To switch service groups

1. In the Management Server console, go to the **Availability** perspective and select **Manage** in the left pane.
2. Expand the Organization or **Uncategorized Clusters** to locate the cluster.
3. Expand the cluster and then the **Service Groups** node to locate the service group.
4. Under the **Service Groups** tab, select the required service group. Press Ctrl or Shift for the selection of multiple service groups.
5. Right-click and select **Switch**.

   See "Switch service groups panel options" on page 382.
6. In the **Result** panel, click **OK**.

See "About site aware service group operations" on page 376.

Switch service groups panel options

Use this panel to provide site and systems (cluster nodes) related information to switch one or more service groups to another system.

For single service group: Table 18-10

For multiple service groups: Table 18-11

Table 18-10 Switch single service group panel options

Field	Description
Select the system you want to switch service group to:	Select the required system to switch the selected service group to. A failover service group can switch to any specific system, any system in the cluster, or any system from the selected site. A hybrid service group can switch to any specific system. The switch operation is not supported for the parallel service group.

Table 18-11 Switch multiple service groups panel options

Field	Description
Do you want to switch the following service group(s)?	Select **Any System** or the site name to switch the selected service groups to any available system or the selected site.

Note: For the switch operation on multiple service groups, though Veritas InfoScale Operations Manager displays all sites for the selected service groups, it is important to note that some sites may not be applicable to a service group or groups. The switch operation will be successful only for those sites which are applicable to the selected service groups. The switch operation is ignored for the remaining service groups.

See "Switching service groups" on page 381.

Clearing faults on service group

In the Management Server console, you can clear a service group by removing the resource faults within that service group. After you clear resource faults, you can bring the service groups online. A resource fault in a service group may occur in several situations, such as a power failure or a faulty configuration.

To perform this task, your user group must be assigned the Admin role or the Operator role on the service group or the Availability perspective. The permission on the service group may be explicitly assigned or inherited from a parent Organization or cluster.

To clear the faults on service group

1. In the Management Server console, go to the **Availability** perspective and select **Manage** in the left pane.
2. Expand the Organization or **Uncategorized Clusters** to locate the cluster.
3. Expand the cluster and then the **Service Groups** node to locate the service group.
4. Right-click on the service group and select **Clear Fault**.
5. In the **Clear Fault Service Groups** panel, do the following:
   - To clear the fault on a specific system, select the system. Choose **All Systems** to clear the fault on all the systems.

- Click **OK**.

6 In the **Result** panel, click **OK**.

See "About managing service groups" on page 357.

Clearing the resources in a service group from the Admin Wait state

In the Management Server console, you can clear the resources in a service group from the Admin Wait state. The resources may enter an Admin Wait state in a service group due to any of the following reasons:

- The offline function did not complete within the expected time.
- The offline function was ineffective.
- The online function did not complete within the expected time.
- The online function was ineffective.
- The resource was taken offline unexpectedly.
- The monitor function consistently failed to complete within the expected time.

Before clearing the resources in the admin wait state, it is recommended that you take manual action to solve the problems.

To perform this task, your user group must be assigned the Admin role or the Operator role on the service group or the Availability perspective. The permission on the service group may be explicitly assigned or inherited from a parent Organization or cluster.

To clear the resources in a service group from the Admin Wait state

1 In the Management Server console, go to the **Availability** perspective and select **Manage** in the left pane.

2 Expand the Organization or **Uncategorized Clusters** to locate the cluster.

3 Expand the cluster, and then the **Service Groups** node to locate the service group.

4 Right-click on the required service group and select **Clear Admin Wait**.

5 In the **Clear Admin Wait** panel, do the following:

- Select the system where you want to clear the resource from the Admin Wait state.
- Select **Fault Service Group** to mark the service group as faulted.

- Click **OK**.

6 In the **Result** panel, click **OK**.

See "Bringing resources online" on page 405.

See "Taking resources offline" on page 405.

See "Taking a resource offline and propagating the state" on page 404.

See "Probing resources" on page 403.

Editing attributes of service groups

In the Management Server console, you can edit some of the attributes of the service groups.

To perform this task, your user group must be assigned the Admin role on the service group or the Availability perspective. The permission on the service group may be explicitly assigned or inherited from a parent Organization or cluster.

Note: You can edit only the editable attributes of the service groups. The editable attributes have an edit icon along with the attribute name in the list of attributes.

To edit attributes of a service group

1 In the Management Server console, go to the **Availability** perspective and expand **Manage** in the left pane.

2 Expand the Organization or **Uncategorized Clusters** to locate the cluster.

3 Expand the cluster and then the **Service Groups** node to locate the service group.

4 Right-click the service group and select **Properties**.

5 Select the **Attributes** tab.

6 Right-click the attribute that you want to edit, and select **Edit Attribute**.

7 Edit the values in the fields. Click **OK**.

See "Edit attribute options" on page 353.

8 In the **Result** page, click **OK**.

See "About managing service groups" on page 357.

See "About modifying a service group" on page 389.

Modifying the system list for a service group

In the Management Server console, you can modify the list of the systems where a service group can go online.

To perform this task, your user group must be assigned the Admin role on the service group or the Availability perspective. The permission on the service group may be explicitly assigned or inherited from a parent Organization or cluster.

Note: Before you perform this task, you need to open the cluster configuration if the configuration is not already open.

To modify the system list for a service group

1. In the Management Server console, go to the **Availability** perspective and select **Manage** in the left pane.
2. Expand the Organization or **Uncategorized Clusters** to locate the cluster.
3. Expand the cluster and then the **Service Groups** node to locate the service group.
4. Right-click on the service group and select **Edit** > **Modify System List**.
5. In the **Configure System List** panel, specify the details and click **OK**.

   See "Configure System List options" on page 360.
6. In the **Result** panel, click **OK**.

See "About managing service groups" on page 357.

See "About modifying a service group" on page 389.

About dependency views

The dependency view is a graphical representation of a service group's dependency with other service groups or the dependency of the resources for a selected service group.

The Management Server console provides two types of dependency views:

- Service group dependency view - launched from cluster view and service group view
- Resource dependency view - launched from service group view and resource view

The dependency views have the following components:

Component	Description
Main section	Located in the center of the dependency view window. To pan or move the view in any direction, click and drag the mouse in the desired direction. Using the main section, you can perform the following actions: Service group dependency view: ■ Right-click the service group icon and select the required operation from the shortcut menu to perform a service group-level operation. Resource dependency view: ■ Right-click the resource icon and select the required operation from the shortcut menu to perform a resource-level operation.
Horizontal zoom bar	Located at the top of the dependency view window. To zoom in or zoom out, click on the zoom bar or drag the zoom sliders. Click **Reset Zoom** to reset the zoom to 100%.
System tabs	Located at the bottom of the dependency view window. Tabs represent the systems present in the system list of the selected service group.

Note: The dependency view is automatically refreshed every 20 seconds. You can also refresh the view by refreshing the browser.

See "Viewing the service group dependency view" on page 387.

See "Viewing the resource dependency view" on page 413.

Viewing the service group dependency view

In the Management Server console, you can view the graphical representation of the dependencies between the service groups in a cluster.

You can view this information related to the service groups for which your user group has at least Guest role explicitly assigned or inherited from a parent Organization. You can also view information if your user group has at least Guest role assigned on the Availability perspective.

You can perform all the operations that are related to a service group by right-clicking on the required service group.

Note: The service group dependency view that is launched from the cluster view shows the dependency for all the service groups in that cluster. The service group dependency view that is launched from the service group view displays only the dependencies for the selected service group.

To view service groups dependencies from cluster view

1. In the Management Server console, go to the **Availability** perspective and select **Manage** in the left pane.
2. Expand the organization Entity, or **Uncategorized Clusters** to locate the cluster.
3. Select the required cluster.
4. Select **Service Group Dependency** tab.

   This view displays the dependencies for all the service groups in the selected cluster.

To view service group dependencies from service group view

1. In the Management Server console, go to the **Availability** perspective and select **Manage** in the left pane.
2. Expand the organization Entity, or **Uncategorized Clusters** to locate the cluster.
3. Expand the cluster and then select the service group.
4. Select **Service Group Dependency** tab.

   This view displays only the dependencies for the service group that you have chosen.

See “About dependency views” on page 386.

See “Viewing the resource dependency view” on page 413.

About modifying a service group

You can modify a service group that you have created. Modifying a service group involves performing one or more tasks in the following list:

- Modifying the resources in the service group
  See "Adding or modifying resources" on page 411.
- Modifying the System List for the service group
  See "Modifying the system list for a service group" on page 386.
- Editing attributes for the service group

About Cluster Server service group alerting and failover reporting

The Cluster Server (VCS) service group failover alerting and reporting provides the following features:

- It calculates the average failover duration for the automatic unplanned failover of VCS service groups in the data center. The storage administrator can view the service groups that take more time to come online on a system. The alerting feature also helps you understand the overall performance of Cluster Server in the data center. The failover duration of VCS service groups is provided by the **VCS Failover Duration** report.
- It provides the ability to set the threshold (in seconds) for the failover duration alert. Veritas InfoScale Operations Manager raises an alert if the current failover duration is more than the specified threshold. You can set the failover duration threshold for the following objects:
  - Cluster
  - Service group

  The value of the failover duration set at the cluster level is applicable to all its constituent service groups unless explicitly set at a particular service group-level.

Note: To use the reporting and alerting feature, make sure you use 6.1, or later versions of Cluster Server, Veritas InfoScale Operations Manager Management Server and the managed hosts.

See "Viewing VCS Failover Duration report" on page 390.

See "About threshold settings" on page 523.

See "Adding threshold settings for an object" on page 525.

Viewing VCS Failover Duration report

In Veritas InfoScale Operations Manager, the VCS Failover Duration report provides the details on the failover duration for the service groups in the selected scope. You can add all types of service groups (failover, parallel, and hybrid) for the failover duration reporting. The report provides the minimum, maximum, and the average failover duration for the service groups, thereby providing an overview of high availability (HA) status of the applications in the data center. All users (administrator, operator, and guest) can run this report. The reporting and alerting is supported only for Cluster Server service groups. Other third-party clustering products are not supported. Also ensure that VCS nodes (managed hosts) are running Veritas Operations Manger 6.1 or later.

When a service group is failed over across a global site or between different clusters, the failover report displays the host name and cluster name of the source and target systems in the following format *<Host-Name>::<Cluster-Name>*.

For example, when the host (host1) in the cluster (clus2) is failed over to the host (host5) in another cluster (clus7), the report indicates the source as *host1::clus2* and target as *host5::clus7*.

You can perform the following tasks in this view:

- Subscribe for the report.
- Save the report as a CSV file.
- Email the report.

You can view this information related to the clusters and service groups for which your user group has at least Guest role explicitly assigned or inherited from a parent Organization. You can also view information if your user group has at least Guest role assigned on the Availability perspective.

To view the VCS Failover Duration report

1. In the Management Server console, go to the **Availability** perspective and expand **Reports** in the left pane.
2. Click the **VCS Failover Duration** report.
3. In the **Select scope for VCS Failover Duration report** wizard panel, select the scope of the report and click **Run**.

   See “Select scope to run report panel options” on page 121.

Chapter 19

Managing systems

This chapter includes the following topics:

- About managing systems
- Freezing a system
- Unfreezing a system
- Editing attributes of a system
- Starting the Cluster Server high availability daemon on the hosts in a cluster
- Stopping the Cluster Server high availability daemon on the systems in a cluster

About managing systems

In the Management Server console, you can manage the systems that are associated with a cluster. You can perform operations such as freeze or unfreeze on any system.

See “Freezing a system” on page 391.

See “Unfreezing a system” on page 392.

Freezing a system

In the Management Server console, you can freeze a system to prevent the service groups on it from failing over to another system. This task is particularly useful when you perform a system upgrade. A frozen system automatically unfreezes when Cluster Server is restarted.

To perform this task, your user group must be assigned the Admin or Operator role on the cluster or the Availability perspective. The permission on the cluster may be explicitly assigned or inherited from a parent Organization.

To freeze the system persistently, or to failover active service groups to another system, you must have administrator privilege.

To freeze a system

1. In the Management Server console, go to the **Availability** perspective and select **Manage** in the left pane.
2. Expand the Organization or **Uncategorized Clusters** to locate the cluster.
3. Expand the cluster and then expand the **Systems** node to locate the system.
4. Right-click on the system and select **Freeze**.
5. In the **Freeze System** panel, do the following:
   - Select **Freeze Persistently** to ensure that the systems remain frozen unless explicitly unfrozen.
   - Select **Evacuate service groups running on the System** to fail over the active service groups on the systems to another system in the cluster.
6. Click **OK**.
7. In the **Result** panel, click **OK**.

See "Unfreezing a system" on page 392.

Unfreezing a system

In the Management Server console, you can unfreeze a system that has been frozen. You can also unfreeze the systems that were persistently frozen.

To perform this task, your user group must be assigned the Admin or Operator role on the cluster or the Availability perspective. The permission on the cluster may be explicitly assigned or inherited from a parent Organization.

To unfreeze a system that has been frozen persistently, you must have administrator privilege.

To unfreeze a system

1. In the Management Server console, go to the **Availability** perspective and select **Manage** in the left pane.
2. Expand the Organization or **Uncategorized Clusters** to locate the cluster.
3. Expand the cluster and then expand the **Systems** node to locate the system.
4. Right-click on the system and select **Unfreeze**.
5. In the **Unfreeze System** panel, click **OK**.
6. In the **Result** panel, click **OK**.

See “Freezing a system” on page 391.

Editing attributes of a system

In the Management Server console, you can edit some of the attributes of the system.

To perform this task, your user group must be assigned the Admin role on the cluster or the Availability perspective. The permission on the cluster may be explicitly assigned or inherited from a parent Organization.

Note: You can edit only the editable attributes of the system. The editable attributes have an edit icon along with the attribute name in the list of attributes.

To edit attributes of a system

1. In the Management Server console, go to the **Availability** perspective and select **Manage** in the left pane.
2. Expand the Organization or **Uncategorized Clusters** to locate the cluster.
3. Expand the cluster and then expand the **Systems** node to locate the system.
4. Right-click on the system and select **Properties**.
5. Select the **Attributes** tab.
6. Right-click the attribute that you want to edit, and select **Edit Attribute**.
7. Edit the values in the fields. Click **OK**.

   See “Edit attribute options” on page 353.
8. In the **Result** page, click **OK**.

See “About managing systems” on page 391.

Starting the Cluster Server high availability daemon on the hosts in a cluster

In the Management Server console, you can start the Cluster Server high availability daemon (HAD) on the hosts in a cluster. You can start the Cluster Server HAD on the selected hosts, or on all hosts in a cluster.

To perform this task, your user group must be assigned the Admin role on the cluster or the Availability perspective. The permission on the cluster may be explicitly assigned or inherited from a parent Organization.

To start the Cluster Server HAD on the hosts in a cluster

1. In the Management Server console, go to the **Availability** perspective and select **Manage** in the left pane.
2. Expand the Organization or **Uncategorized Clusters** to locate the cluster.
3. Right-click on the cluster and select **Start VCS**.
4. In the **Start VCS** wizard panel, do the following:
   - To start the Cluster Server HAD on a specific host, select the host. Choose **All Systems** to start the Cluster Server HAD on all hosts in the cluster.
   - Click **OK**.
5. In the **Result** panel, click **OK**.

See "Stopping the Cluster Server high availability daemon on the systems in a cluster" on page 394.

Stopping the Cluster Server high availability daemon on the systems in a cluster

In the Management Server console, you can stop the Cluster Server high availability daemon (HAD) on the selected systems, or on all systems in a cluster. You may want to stop the Cluster Server high availability daemon (HAD) on the systems in case of a maintenance operation for the devices.

To perform this task, your user group must be assigned the Admin role on the cluster or the Availability perspective. The permission on the cluster may be explicitly assigned or inherited from a parent Organization.

To stop the Cluster Server HAD on the hosts in a cluster

1. In the Management Server console, go to the **Availability** perspective and select **Manage** in the left pane.
2. Expand the Organization or **Uncategorized Clusters** to locate the cluster.
3. Right-click on the cluster and select **Stop VCS**.
4. In the **Stop VCS** wizard panel, do the following:
   - To stop the Cluster Server HAD on a specific system, select the system. Choose **All Systems** to stop the Cluster Server HAD on all systems in the cluster.
   - To force the Cluster Server HAD to stop without stopping the applications or service groups that are running on the systems, select **Force**.

- Click **OK**.

5 In the **Result** panel, click **OK**.

See “Starting the Cluster Server high availability daemon on the hosts in a cluster” on page 393.

Chapter 20

Managing VSystems

This chapter includes the following topics:

- About VSystems
- Starting a virtual machine
- Stopping a virtual machine
- Migrating a virtual machine

About VSystems

In the Management Server console, you can view the virtual machines (VM) that are monitored. This view is available only for the clusters having virtual machines that are monitored. In Veritas InfoScale Operations Manager, the virtual machines that are managed under Veritas Cluster Server (VCS) as resources are called VSystems.

You can perform operations such as start VM, stop VM, and migrate on VSystems.

See “Starting a virtual machine” on page 396.

See “Stopping a virtual machine” on page 397.

See “Migrating a virtual machine” on page 398.

Starting a virtual machine

In the Management Server console, you can start the virtual machine (VM) by starting the service group that monitors the virtual machine.

To perform this task, your user group must be assigned the Admin role on the service group or the Availability perspective. The permission on the service group may be explicitly assigned or inherited from a parent Organization or cluster.

To start a virtual machine

1. In the Management Server console, go to the **Availability** perspective and select **Manage** in the left pane.
2. Expand the Organization or **Uncategorized Clusters** to locate the cluster.
3. Select the required cluster in the navigation tree and select the **VSystems** tab.
4. Right-click on the required virtual machine and select **Start VM**.
5. In the **Start Virtual Machine** panel, do the following:
   - Select the physical server where you want to start the VM.
   - Click **OK**.
6. In the **Result** panel, click **OK**.

See "About VSystems" on page 396.

See "Stopping a virtual machine" on page 397.

See "Migrating a virtual machine" on page 398.

Stopping a virtual machine

In the Management Server console, you can stop the virtual machine (VM) by stopping the service group that monitors the virtual machine.

To perform this task, your user group must be assigned the Admin role on the service group or the Availability perspective. The permission on the service group may be explicitly assigned or inherited from a parent Organization or cluster.

To stop a virtual machine

1. In the Management Server console, go to the **Availability** perspective and select **Manage** in the left pane.
2. Expand the Organization or **Uncategorized Clusters** to locate the cluster.
3. Select the required cluster in the navigation tree and select the **VSystems** tab.
4. Right-click on the required virtual machine and select **Stop VM**.
5. In the **Stop Virtual Machine** panel, do the following:
   - Select the system to enable the service group on. Choose **All Systems** to enable the service group on all systems.
   - Click **OK**.
6. In the **Result** panel, click **OK**.

See "About VSystems" on page 396.

See “Starting a virtual machine” on page 396.

See “Migrating a virtual machine” on page 398.

Migrating a virtual machine

Using Veritas InfoScale Operations Manager 7.2, you can perform a live migration for the virtual machines (VM) that are configured as resource in the VCS service groups. In a service group, you can configure only one virtual machine as a resource. During migration, you move a running virtual machine between different physical servers without disconnecting the clients or applications. Unlike switch operation, live migration feature involves minimal downtime.

For more information on migrating Virtual machines, refer to *Cluster Server 6.1 Administrator's Guide*.

Note: To be able to migrate a VM, You need the managed host package version 6.1 on the host.

To perform this task, your user group must be assigned the Admin role on the service group or the Availability perspective. The permission on the service group may be explicitly assigned or inherited from a parent Organization or cluster.

To migrate a virtual machine

1. In the Management Server console, go to the **Availability** perspective and select **Manage** in the left pane.
2. Expand the Organization or **Uncategorized Clusters** to locate the cluster.
3. Select the required cluster in the navigation tree and select the **VSystems** tab.
4. Right-click on the required virtual machine and select **Migrate**.
5. In the **Migrate Virtual Machine** panel, do the following:
   - Select the system where you want to migrate the virtual system.
   - Click **OK**.
6. In the **Result** panel, click **OK**.

See “About VSystems” on page 396.

See “Starting a virtual machine” on page 396.

See “Stopping a virtual machine” on page 397.

Chapter 21

Managing resources

This chapter includes the following topics:

- About managing resources
- Enabling resources
- Disabling resources
- Deleting resources
- Clearing faults on resources
- Probing resources
- Taking a resource offline and propagating the state
- Bringing resources online
- Taking resources offline
- Invoking a resource action
- Editing attributes of a resource
- Editing attributes of a resource type
- Linking resources in a service group
- Unlinking resources in a service group
- Adding or modifying resources
- Marking a resource as critical
- Marking a resource as non critical
- Viewing the resource dependency view

About managing resources

Resources are the most basic elements of a service group. You can use the Management Server console to manage single or multiple resources.

See “Enabling resources” on page 400.

See “Disabling resources” on page 401.

See “Probing resources” on page 403.

See “Editing attributes of a resource” on page 408.

See “Editing attributes of a resource type” on page 409.

See “Clearing faults on resources” on page 402.

See “Clearing the resources in a service group from the Admin Wait state” on page 384.

See “Deleting resources” on page 401.

See “Linking resources in a service group” on page 410.

See “Bringing resources online” on page 405.

See “Taking resources offline” on page 405.

See “Invoking a resource action” on page 406.

See “Taking a resource offline and propagating the state” on page 404.

See “Adding or modifying resources” on page 411.

Enabling resources

In the Management Server console, you can enable one or more resources. You need to first enable the resources before bringing them online.

To perform this task, your user group must be assigned the Admin role on the service group or the Availability perspective. The permission on the service group may be explicitly assigned or inherited from a parent Organization or cluster.

To enable resources

1. In the Management Server console, go to the **Availability** perspective and select **Manage** in the left pane.
2. Expand the Organization or **Uncategorized Clusters** to locate the cluster.
3. Expand the cluster, the **Service Groups** node, the service group, and then the resource type to locate the resource.
4. Right-click on the required resource and select **Enable**.

5 In the **Enable resource** panel, click **OK**.

6 In the **Result** panel, click **OK**.

See "About managing resources" on page 400.

See "Disabling resources" on page 401.

See "Enabling service groups" on page 363.

See "Autoenabling service groups" on page 367.

See "Enabling all resources of service groups" on page 370.

Disabling resources

In the Management Server console, you can prevent a resource from being brought online by disabling the resource. You can disable one or more resources at a time.

To perform this task, your user group must be assigned the Admin role on the service group or the Availability perspective. The permission on the service group may be explicitly assigned or inherited from a parent Organization or cluster.

To disable resources

1 In the Management Server console, go to the **Availability** perspective and select **Manage** in the left pane.

2 Expand the Organization or **Uncategorized Clusters** to locate the cluster.

3 Expand the cluster, the **Service Groups** node, the service group, and then the resource type to locate the resource.

4 Right-click on the required resource and select **Disable**.

5 In the **Disable resource** panel, click **OK**.

6 In the **Result** panel, click **OK**.

See "About managing resources" on page 400.

See "Enabling resources" on page 400.

See "Disabling service groups" on page 364.

See "Disabling all resources of service groups" on page 371.

Deleting resources

In the Management Server console, you can delete the resources in a service group. You can delete multiple resources simultaneously.

To perform this task, your user group must be assigned the Admin role on the service group or the Availability perspective. The permission on the service group may be explicitly assigned or inherited from a parent Organization or cluster.

To delete resources

1. In the Management Server console, go to the **Availability** perspective and select **Manage** in the left pane.
2. Expand the Organization or **Uncategorized Clusters** to locate the cluster.
3. Expand the cluster, the **Service Groups** node, the service group, and then the resource type to locate the resource.
4. Right-click on the required resource and select **Delete**.
5. In the **Delete resource** panel, click **OK**.
6. In the **Result** panel, click **OK**.

See "About managing resources" on page 400.

See "Disabling all resources of service groups" on page 371.

See "Enabling resources" on page 400.

See "Enabling all resources of service groups" on page 370.

Clearing faults on resources

In the Management Server console, you can clear the faults on a resource.

To perform this task, your user group must be assigned the Admin role or the Operator role on the service group or the Availability perspective. The permission on the service group may be explicitly assigned or inherited from a parent Organization or cluster.

To clear faults on resources

1. In the Management Server console, go to the **Availability** perspective and select **Manage** in the left pane.
2. Expand the Organization or **Uncategorized Clusters** to locate the cluster.
3. Expand the cluster, the **Service Groups** node, the service group, and then the resource type to locate the resource.
4. Right-click on the required resource and select **Clear Fault**.
5. In the **Clear Fault Resource** panel, select the system where you want to clear the resource. Choose **All Systems** to clear the resource on all the systems. Click **OK**.
6. In the **Result** panel, click **OK**.

See “About managing resources” on page 400.

See “Clearing faults on service group” on page 383.

Probing resources

In the Management Server console, you can probe a resource to confirm that it is properly configured before you bring it online. You can use the Management Server console to probe one or more resources.

Note: A resource must be enabled before you probe it.

To perform this task, your user group must be assigned the Admin role or the Operator role on the service group or the Availability perspective. The permission on the service group may be explicitly assigned or inherited from a parent Organization or cluster.

To probe resources

1. In the Management Server console, go to the **Availability** perspective and select **Manage** in the left pane.
2. Expand the Organization or **Uncategorized Clusters** to locate the cluster.
3. Expand the cluster, the **Service Groups** node, the service group, and then the resource type to locate the resource.
4. Right-click on the required resource and select **Probe**.
5. In the **Probe resource(s)** panel, select the systems on which you want to probe the resource. Select **All systems** to probe the resource on all the systems. Click **OK**.
6. In the **Result** panel, click **OK**.

See “About managing resources” on page 400.

See “Bringing resources online” on page 405.

See “Taking resources offline” on page 405.

See “Taking a resource offline and propagating the state” on page 404.

See “Clearing the resources in a service group from the Admin Wait state” on page 384.

Taking a resource offline and propagating the state

In the Management Server console, you can take a resource and all of its dependents offline. When you take the parent resource offline, the offline state propagates to the child resources. You can bring multiple resources of a service group offline simultaneously.

You cannot perform this task if any of the following conditions exist:

- The resource does not depend on any other resource
- The resource does not depend on any other resource
- The resource is offline

To perform this task, your user group must be assigned the Admin role or the Operator role on the service group or the Availability perspective. The permission on the service group may be explicitly assigned or inherited from a parent Organization or cluster.

To take a resource offline and propagate the state

1. In the Management Server console, go to the **Availability** perspective and select **Manage** in the left pane.
2. Expand the Organization or **Uncategorized Clusters** to locate the cluster.
3. Expand the cluster, the **Service Groups** node, the service group, and then the resource type to locate the resource.
4. Right-click on the required resource and select **offline Propagate**.
5. In the **Offline Propagate for Resource** panel, do the following:
   - To take a resource offline and propagate the offline state to the child resources on a specific system, select the system.
   - To take the resource offline without considering the state of the parent resource, select **Ignore parent**.
   - Click **OK**.
6. In the **Result** panel, click **OK**.

See “About managing resources” on page 400.

See “Bringing resources online” on page 405.

See “Taking resources offline” on page 405.

See “Clearing the resources in a service group from the Admin Wait state” on page 384.

See “Probing resources” on page 403.

Bringing resources online

In the Management Server console, you can manually bring resources online. You can bring multiple resources of a service group online simultaneously. To bring a resource online, you must first enable that resource.

To perform this task, your user group must be assigned the Admin role or the Operator role on the service group or the Availability perspective. The permission on the service group may be explicitly assigned or inherited from a parent Organization or cluster.

To bring a resource online

1. In the Management Server console, go to the **Availability** perspective and select **Manage** in the left pane.
2. Expand the Organization or **Uncategorized Clusters** to locate the cluster.
3. Expand the cluster, the **Service Groups** node, the service group, and then the resource type to locate the resource.
4. Right-click on the required resource and select **Online**.
5. In the **Online Resource** panel, do the following:
   - To bring a resource online on a specific system, select the system.
   - To let a global resource (that has authority on a remote cluster) have the authority on the local cluster on which it is brought online, select **Force**.
   - Click **OK**.
6. In the **Result** panel, click **OK**.

See “About managing resources” on page 400.

See “Taking resources offline” on page 405.

See “Clearing the resources in a service group from the Admin Wait state” on page 384.

See “Probing resources” on page 403.

Taking resources offline

In the Management Server console, you can manually take resources offline. You can take multiple resources of a service group offline simultaneously.

To perform this task, your user group must be assigned the Admin role or the Operator role on the service group or the Availability perspective. The permission on the service group may be explicitly assigned or inherited from a parent Organization or cluster.

To take a resource offline

1. In the Management Server console, go to the **Availability** perspective and select **Manage** in the left pane.
2. Expand the Organization or **Uncategorized Clusters** to locate the cluster.
3. Expand the cluster, the **Service Groups** node, the service group, and then the resource type to locate the resource.
4. Right-click on the required resource and select **offline**.
5. In the **Offline Resource** panel, do the following:
   - To take a resource offline on a specific system, select the system.
   - To take the resource offline without considering the state of the parent resource, select **Ignore parent**.
   - Click **OK**.
6. In the **Result** panel, click **OK**.

See "About managing resources" on page 400.

See "Bringing resources online" on page 405.

See "Taking a resource offline and propagating the state" on page 404.

See "Probing resources" on page 403.

See "Clearing the resources in a service group from the Admin Wait state" on page 384.

Invoking a resource action

In the Management Server console, you can run a predefined script that performs an action on a resource. A few examples of predefined resource actions are splitting disk groups and joining disk groups.

To perform this task, your user group must be assigned the Admin role or the Operator role on the service group or the Availability perspective. The permission on the service group may be explicitly assigned or inherited from a parent Organization or cluster.

To invoke a resource action

1. In the Management Server console, go to the **Availability** perspective and select **Manage** in the left pane.
2. Expand the Organization or **Uncategorized Clusters** to locate the cluster.
3. Expand the cluster, the **Service Groups** node, the service group, and then the resource type to locate the resource.
4. Right-click on the required resource and select **Invoke Action**.
5. In the **Invoke Action** panel, enter the required information. Click **OK**.

   See "Invoke Action options" on page 407.
6. In the **Result** panel, click **OK**.

See "About managing resources" on page 400.

See "Adding or modifying resources" on page 411.

See "Linking resources in a service group" on page 410.

See "Unlinking resources in a service group" on page 410.

See "Deleting resources" on page 401.

Invoke Action options

Use this wizard panel to provide the input to invoke a resource action.

Table 21-1 Invoke resource action wizard panel options

Field	Description
Select the action that you want to invoke	Select the action that you want to invoke from the drop-down list. This list contains only the Supported Actions that are supported on the resource that you have chosen.
Select the system you want to invoke action on	Select the system you want to invoke action on from the drop-down list.

Table 21-1 Invoke resource action wizard panel options *(continued)*

Field	Description
Action Arguments	Enter the arguments for the action that you want to invoke and click **Add**. For example, you can use arguments for the Application and the Mount actions as follows: ■ Application - Select **getcksum** as the action and enter the path to the application. Click **Add** and click **OK** to invoke the action on the selected system. ■ Mount point - Select **mountpoint.vfd** and enter **fix** as the action argument. Click **Add**, and then click **OK** to create the mount point on the selected system.

See "Invoking a resource action" on page 406.

Editing attributes of a resource

In the Management Server console, you can edit some of the attributes of the resources.

To perform this task, your user group must be assigned the Admin role on the service group or the Availability perspective. The permission on the service group may be explicitly assigned or inherited from a parent Organization or cluster.

Note: You can edit only the editable attributes of the resource. The editable attributes have an edit icon along with the attribute name in the list of attributes.

To edit attributes of a resource

1. In the Management Server console, go to the **Availability** perspective and select **Manage** in the left pane.
2. Expand the Organization or **Uncategorized Clusters** to locate the cluster.
3. Expand the cluster, the **Service Groups** node, the service group, and then the resource type to locate the resource.
4. Right-click the resource and select **Properties**.
5. Select the **Attributes** tab.
6. Right-click the attribute that you want to edit, and select **Edit Attribute**.

7 Edit the values in the fields. Click **OK**.

See "Edit attribute options" on page 353.

8 In the **Result** page, click **OK**.

See "About managing resources" on page 400.

See "Adding or modifying resources" on page 411.

Editing attributes of a resource type

In the Management Server console, you can edit some of the attributes of the resource types.

To perform this task, your user group must be assigned the Admin role on the service group or the Availability perspective. The permission on the service group may be explicitly assigned or inherited from a parent Organization or cluster.

Note: You can edit only the editable attributes of a resource type. The editable attributes have an edit icon along with the attribute name in the list of attributes.

To edit attributes of a resource type

1 In the Management Server console, go to the **Availability** perspective and expand **Manage** in the left pane.

2 Expand the Organization or **Uncategorized Clusters** to locate the cluster.

3 Expand the cluster, the **Service Groups** node, and the service group to locate the resource type.

4 Right-click the resource type and select **Properties**.

5 Select the **Attributes** tab.

6 Right-click the attribute that you want to edit, and select **Edit Attribute**.

7 Edit the values in the fields. Click **OK**.

See "Edit attribute options" on page 353.

8 In the **Result** page, click **OK**.

See "About managing clusters" on page 350.

See "Saving a cluster configuration" on page 351.

Linking resources in a service group

In the Management Server console, you can create links (dependencies) between two resources of a service group. When you create a link between two resources, you need to designate one resource as the parent resource and the other one as a child resource. The child resource must come online before the parent resource.

The parent resource of one dependency can be the child resource of another dependency. In a service group, several parent-child dependencies exist to support a single application resource. In a dependency diagram, this application resource occupies the apex of the diagram.

To perform this task, your user group must be assigned the Admin role on the service group or the Availability perspective. The permission on the service group may be explicitly assigned or inherited from a parent Organization or cluster.

To link resources with a parent-child dependency

1. In the Management Server console, go to the **Availability** perspective and select **Manage** in the left pane.
2. Expand the Organization or **Uncategorized Clusters** to locate the cluster.
3. Expand the cluster, the **Service Groups** node, the service group, and then the resource type to locate the resource.
4. Right-click on the required resource and select **Link/Unlink Resources**.
5. In the **Resource Dependencies** panel, Select the parent and the child resources, and select **Link**. Click **OK**.
6. In the **Result** panel, click **OK**.

See “About managing resources” on page 400.

See “Adding or modifying resources” on page 411.

See “Invoking a resource action” on page 406.

See “Unlinking resources in a service group” on page 410.

Unlinking resources in a service group

In the Management Server console, you can unlink two resources in a service group to remove their parent-child dependency.

To perform this task, your user group must be assigned the Admin role on the service group or the Availability perspective. The permission on the service group may be explicitly assigned or inherited from a parent Organization or cluster

To unlink resources in a service group

1. In the Management Server console, go to the **Availability** perspective and select **Manage** in the left pane.
2. Expand the Organization or **Uncategorized Clusters** to locate the cluster.
3. Expand the cluster, the **Service Groups** node, the service group, and then the resource type to locate the resource.
4. Right-click on the required resource and select **Link/Unlink Resources**.
5. In the **Resource Dependencies** panel, do the following:
   - Click the **Delete** option that corresponds to the resource dependency that you want to remove.
   - Click **OK**.
6. In the **Result** panel, click **OK**.

See "About managing resources" on page 400.

See "Linking resources in a service group" on page 410.

See "Adding or modifying resources" on page 411.

See "Invoking a resource action" on page 406.

See "Deleting resources" on page 401.

Adding or modifying resources

In the Management Server console, you can add resources to a service group or modify the existing service groups.

To perform this task, your user group must be assigned the Admin role on the service group or the Availability perspective. The permission on the service group may be explicitly assigned or inherited from a parent Organization or cluster.

To add or modify resources in a service group

1. In the Management Server console, go to the **Availability** perspective and select **Manage** in the left pane.
2. Expand the Organization or **Uncategorized Clusters** to locate the cluster.
3. Expand the cluster, and the **Service Groups** node to locate the service group.
4. Right-click on the required service group and select **Edit** > **Resources** > **Add/Modify Resources**.
5. In the **Configure Resources** panel, do the following:

- Enter the required information to add a resource. To modify a resource, use the options that are available in the **Resource List** table.
  See "Configure Resources options" on page 361.
- To configure the resource dependencies, click **Next**.
- To add or modify the resource and configure the resource dependencies at a later time, click **Finish**.

6 In the **Resource Dependencies** panel, do the following:

- Enter the required information to link the resources.
- Click **Finish**.

7 In the **Result** panel, click **OK**.

See "About managing resources" on page 400.

See "Linking resources in a service group" on page 410.

See "Unlinking resources in a service group" on page 410.

See "Invoking a resource action" on page 406.

See "Deleting resources" on page 401.

Marking a resource as critical

In the Management Server console, you can mark a resource of a service group as critical. When a resource that is marked as critical or its dependent resource faults, the service group where these resources exist also moves to the faulted state. Cluster Server takes the failed resource offline and updates the service group status to ONLINE|PARTIAL. This operation also ensures that the service group does not come online, as part of its online process, when a critical resource fails to come online

To perform this task, your user group must be assigned the Admin role or the Operator role on the service group or the Availability perspective. The permission on the service group may be explicitly assigned or inherited from a parent Organization or cluster.

To mark a resource as critical

1 In the Management Server console, go to the **Availability** perspective and select **Manage** in the left pane.

2 Expand the Organization or **Uncategorized Clusters** to locate the cluster.

3 Expand the cluster, the **Service Groups** node, the service group, and then the resource type to locate the resource.

4 Right-click on the required resource and select **Mark Critical**.

5 In the **Mark Resource as Critical** panel, click **OK**.

6 In the **Result** panel, click **OK**.

See "Marking a resource as non critical" on page 413.

See "About managing resources" on page 400.

Marking a resource as non critical

In the Management Server console, you can mark a critical resource in a service group as non-critical. When you mark a critical resource as non-critical, the service group does not move to the faulted state when the resources that are marked non-critical, or its dependent resources are faulted.

To perform this task, your user group must be assigned the Admin role or the Operator role on the service group or the Availability perspective. The permission on the service group may be explicitly assigned or inherited from a parent Organization or cluster.

To mark a resource as non critical

1 In the Management Server console, go to the **Availability** perspective and select **Manage** in the left pane.

2 Expand the Organization or **Uncategorized Clusters** to locate the cluster.

3 Expand the cluster, the **Service Groups** node, the service group, and then the resource type to locate the resource.

4 Right-click on the required resource and select **Mark Non Critical**.

5 In the **Mark Resource as Non Critical** panel, click **OK**.

6 In the **Result** panel, click **OK**.

See "Marking a resource as critical" on page 412.

See "About managing resources" on page 400.

Viewing the resource dependency view

In the Management Server console, you can view the graphical representation of the dependencies between the resources in a service group.

You can view this information related to the service groups for which your user group has at least Guest role explicitly assigned or inherited from a parent Organization or cluster. You can also view information if your user group has at least Guest role assigned on the Availability perspective.

You can perform all the operations that are related to a resource by right-clicking on the required resource.

To view resource dependencies from service group view

1. In the Management Server console, go to the **Availability** perspective and select **Manage** in the left pane.
2. Expand the Organization, or **Uncategorized Clusters** to locate the cluster.
3. Expand the cluster and then select the service group.
4. Select **Resource Dependency** tab.

   This view displays the dependencies for all the resources in the selected service group.

To view resource dependencies from resource view

1. In the Management Server console, go to the **Availability** perspective and select **Manage** in the left pane.
2. Expand the Organization, or **Uncategorized Clusters** to locate the cluster.
3. Expand the cluster, the **Service groups** node, the service group, and then the resource type to select the resource.
4. Select **Resource Dependency** tab.

   This view displays only the dependencies for the resource that you have chosen.

See “About dependency views” on page 386.

See “Viewing the service group dependency view” on page 387.

Chapter 22

Managing global cluster configurations

This chapter includes the following topics:

- About global clusters
- About global clusters terminology
- About creating global clusters
- Converting local service groups to global service groups
- Converting global service groups to local service groups
- About removing a remote cluster from a global cluster setup

About global clusters

Global cluster technology is a central feature of high availability disaster recovery (HA/DR). Part of HA/DR uses objects called global clusters and global service groups to achieve remote, cross-cluster failover. Global clusters and global service groups have corresponding remote clusters and remote service groups configured at a distant geographic location.

Table 22-1 describes the objects in a global cluster.

Table 22-1 Global cluster objects

Object	Description
Global cluster	A global cluster contains at least one remote cluster.

Table 22-1 Global cluster objects *(continued)*

Object	Description
Global service group	A global service group is a service group that has been configured on one or more remote clusters.
Remote cluster	A remote cluster is a cluster that is specifically configured as a failover target for a global service group. The local cluster heartbeats with this cluster.

If failover within the local cluster is not possible, a global service group fails over to its preconfigured remote cluster.

Note: Remote clusters are usually configured far away from their corresponding global clusters.

See "About global clusters terminology" on page 416.

See "About creating global clusters" on page 417.

See "About removing a remote cluster from a global cluster setup" on page 421.

About global clusters terminology

In a global cluster configuration, a global cluster is at the local site and its corresponding remote cluster is at a remote site. At least one global service group is configured on both the local cluster and on the remote cluster.

In the context of the global cluster management task descriptions, global and remote have the following definitions:

- Global indicates the following:
  - The cluster or the service group is configured at the local site.
  - The cluster or the service group is also configured with a specific failover target cluster or failover target group at a distant remote site.

  Remote indicates the following:
  - The cluster or the service group is configured at a distant site, far away from the local site.
  - The cluster or the service group is configured as a failover target for its global counterpart.

The terms are switched if the perspective is from the remote site.

See "About global clusters" on page 415.

About creating global clusters

The process of creating a global cluster involves the following tasks:

- Configure global cluster heartbeats (inter-cluster heartbeats) to monitor the health of the failover target clusters.
- Create a common service group on a cluster at the local site and on a cluster at a remote site.
- Convert the service group that is common to both the local and the remote clusters into a global service group.

See "About global clusters" on page 415.

See "Prerequisites for creating global clusters" on page 417.

See "Adding a remote cluster to a local cluster" on page 418.

See "Converting local service groups to global service groups" on page 418.

Prerequisites for creating global clusters

Ensure that the following prerequisites are met when you create global clusters:

- The `ClusterService` service group for all clusters must be properly configured for Cluster Server (VCS) global cluster operations.
- The `csgnic` resource or `gconic` resource (NIC resource), `gcoip` (IP resource), and the `wac` resources must come online.
- The service group that is intended to serve as the global service group must have the same name on the global and the remote cluster.
- The global cluster and the remote cluster must have unique names.
- The global cluster and the remote cluster use the same version of VCS.
- The global cluster and the remote cluster must use the same operating system.
- When you use Veritas InfoScale Operations Manager to create the global cluster, the remote cluster should be part of the same Management Server domain as the local cluster.

See "About creating global clusters" on page 417.

Adding a remote cluster to a local cluster

In the Management Server console, you can designate a remote cluster as the failover target for the service groups that are configured on a local cluster. You must begin with two standalone clusters. This task can be performed from either cluster.

To perform this task, your user group must be assigned the Admin role on the cluster or the Availability perspective. The permission on the cluster may be explicitly assigned or inherited from a parent Organization.

To add a remote cluster to a local cluster

1. In the Management Server console, go to the **Availability** perspective and expand **Manage** in the left pane.
2. Expand the Organization, or **Uncategorized Clusters** to locate the cluster.
3. Right-click on the required cluster and select **Disaster Recovery** > **Setup GCO**.
4. In the **Setup GCO between selected clusters** panel, select the local cluster and the cluster that is to be configured as the remote cluster.
5. Click **Finish** to confirm the action.
6. In the **Result** page, click **OK**.

See "About creating global clusters" on page 417.

See "Prerequisites for creating global clusters" on page 417.

Converting local service groups to global service groups

In the Management Server console, you can convert a local service group to global service group.

After connecting the intended global (local) and the intended remote cluster, you must convert the local service group that is common to the remote cluster into a global group.

For each service group on the local cluster that you intend to make global, configure an identical service group on the remote cluster. A service group that is intended to serve as the global group must have the same name on the local cluster and on the remote cluster.

To perform this task, your user group must be assigned the Admin role on the service group or the Availability perspective. The permission on the service group may be explicitly assigned or inherited from a parent Organization or cluster.

To convert a local service group to a global service group

1. In the Management Server console, go to the **Availability** perspective and select **Manage** in the left pane.
2. Expand the Organization or **Uncategorized Clusters** to locate the cluster.
3. Expand the cluster and then the **Service Groups** node to locate the service group.
4. In the service groups list, right-click on the required service groups and select **Edit** > **Make Global/Local**.
5. In the **Make Global/Local** panel, select the clusters on which you want to make the service group global. Move the selected clusters from the **Available Clusters** list to the **Clusters in Priority Order** list.

   See "Make Global/Local options" on page 419.
6. To make the service group global, click **OK**.
7. In the **Result** panel, click **OK**.

See "About creating global clusters" on page 417.

Make Global/Local options

Use this wizard panel to select the clusters on which you want to make the service group local or global, and to select the failover policy.

Table 22-2 Cluster Selection panel options

Option	Description
Available Clusters	Lists the clusters on which you can make the service group global. To make the service group global, you need to select the clusters in this column and move them under the **Clusters in Priority Order** column. To make the service group local, you need to select the clusters in the **Clusters in Priority Order** column and move them under this column .

Table 22-2 Cluster Selection panel options *(continued)*

Option	Description
Clusters in Priority Order	Lists the clusters on which the service group is currently global. To make the service group local, you need to select the clusters in this column and move them under the **Available Clusters** column. To make the service group global, you need to select the clusters in the **Available Clusters** column and move them under this column .
Select Cluster Failover Policy	Lists the policies for cluster failover. To prevent a group from automatically failing over to another cluster, you need to select **Manual**. To enable a group to automatically fail over to another cluster if it is unable to fail over within the cluster, or if the entire cluster faults, you need to select **Auto**. To enable a group to automatically fail over to another cluster if it is unable to fail over within the cluster, you need to select **Connected**. **Note:** Failover of a global service group should not be automatic as the data present at the secondary site may not be updated. For this reason, the default policy for global service groups is **Manual** instead of **Auto**.

See "Converting local service groups to global service groups" on page 418.

See "Converting global service groups to local service groups" on page 420.

Converting global service groups to local service groups

You can convert a global service group back to a local service group. The conversion does the following:

- Removes the remote cluster from the cluster list of the service group on the global cluster.
- Removes the global cluster from the cluster list of the service group on the remote cluster.

To perform this task, your user group must be assigned the Admin role on the service group or the Availability perspective. The permission on the service group may be explicitly assigned or inherited from a parent Organization or cluster.

To convert a global service group to a local service group

1. In the Management Server console, go to the **Availability** perspective and select **Manage** in the left pane.
2. Expand the Organization or **Uncategorized Clusters** to locate the cluster.
3. Expand the cluster and then the **Service Groups** node to locate the service group.
4. In the service groups list, right-click on the required service groups and select **Edit** > **Make Global/Local**.
5. In the **Make Global/Local** panel, select the clusters on which you want to make the service group local. Move the selected clusters from the **Available Clusters** list to the **Clusters in Priority Order** list.

   See "Make Global/Local options" on page 419.
6. To make the service group local, click **OK**.
7. In the **Result** panel, click **OK**.

See "About removing a remote cluster from a global cluster setup" on page 421.

About removing a remote cluster from a global cluster setup

To remove a remote cluster from a global cluster setup, perform the following tasks:

- On the remote cluster, take the `wac` resource in the `ClusterService` group offline.
  See "Taking the wac resource offline " on page 422.
- The second step in removing a remote cluster is to prevent it from being considered as a failover target. Remove the remote cluster from the cluster lists of each global service group and remote service group that participates in this global relationship. You need to convert the global service groups to local service groups.
  See "Converting global service groups to local service groups" on page 420.

You must repeat this procedure for all global service groups that are configured in the cluster.

- Remove the remote cluster from the global cluster setup.
  See "Removing the remote cluster from the global cluster setup" on page 422.

You cannot remove a remote cluster if any of the following apply:

- The remote cluster is part of a cluster list for a global service group
- The cluster is part of a global heartbeats list for a global service group

You must have one of the following privileges to perform the tasks:

- Administrator privilege on the remote cluster that is to be removed
- Administrator privilege on the global service groups that are configured on the remote cluster

Taking the wac resource offline

To remove a remote cluster from a global cluster, the first step is to take the `wac` resource offline on the remote cluster.

To perform this task, your user group must be assigned the Admin role on the remote cluster or the Availability perspective. The permission on the cluster may be explicitly assigned or inherited from a parent Organization.

To take the `wac` resource offline

1. In the Management Server console, go to the **Availability** perspective and expand **Manage** in the left pane.
2. Expand the Organization, or **Uncategorized Clusters** to locate the remote cluster.
3. Select the required cluster and then select **Resources** tab.
4. Right-click on the `wac` resource and select **offline**.

See "About removing a remote cluster from a global cluster setup" on page 421.

Removing the remote cluster from the global cluster setup

The last step in removing a remote cluster is to remove, or unlink, that cluster from the global cluster.

Note: Before you perform this task, you must take the `wac` resource offline on the remote cluster, and remove the remote cluster from the cluster lists of all global service groups.

To remove the remote cluster from the global cluster setup

1. In the Management Server console, go to the **Availability** perspective and expand **Manage** in the left pane.
2. Expand the Organization, or **Uncategorized Clusters** to locate the cluster.
3. Right-click on the required cluster and select **Disaster recovery** > **Remove GCO**.
4. In the **Remove GCO between selected clusters** panel, select the local cluster and the remote cluster that is to be removed from the global cluster setup. Click **Next**.
5. Click **Finish** to confirm the action.
6. In the **Result** page, click **OK**.

See “About removing a remote cluster from a global cluster setup” on page 421.

Chapter 23

Running fire drills

This chapter includes the following topics:

- About high availability and disaster recovery readiness
- About high availability fire drills
- Running the high availability fire drill
- About disaster recovery fire drills
- About configuring a fire drill service group
- Running the disaster recovery fire drill
- Editing a fire drill schedule
- Deleting fire drill schedules
- Enabling fire drill schedules
- Disabling fire drill schedules
- Viewing fire drill schedules

About high availability and disaster recovery readiness

Readiness status is the measure of the ability of a service group to fail over in its intended or the configured fashion. Readiness takes into account the status of the service group, the system, and that of the cluster.

Veritas InfoScale Operations Manager lets you monitor the following types of readiness:

- The high availability readiness that checks for:

- The ability of a service group to fail over to a system within the local cluster
- The ability of a stretch service group to fail over to a system at the local site

- The disaster recovery readiness that checks for:
  - The ability of a Cluster Server global service group to fail over to a system in its target remote cluster at the remote site
  - The ability of a service group in a stretch cluster to fail over to a system in the remote site

See "About high availability and disaster recovery operations" on page 347.

About high availability fire drills

The high availability (HA) fire drill checks whether the service groups in a cluster are ready to go online on a specific system in the same cluster.

To configure high availability for a database or an application, you may make changes to several infrastructure and configuration settings on multiple systems. To maintain these infrastructure and configuration settings is difficult because cluster environments can be subject to constant change. Administrators often add disks, create new disk groups and volumes, and add new cluster nodes or new NICs to upgrade and maintain the infrastructure. Updating the Cluster Server configuration to match the changing physical configuration and infrastructure is critical. HA fire drills detect discrepancies between the Cluster Server configuration and the underlying physical configuration and infrastructure on a node. Such discrepancies might prevent bringing up a service group online on a specific node. Ultimately, the HA fire drill provides the data that is used to update the HA readiness information on the Veritas InfoScale Operations Manager console.

The HA fire drill checks all the resources in a service group and is run on each system (in the system list for the service group) where the service group is offline.

To run an HA fire drill, the service group must be both of the following:

- Online in the cluster.
- Composed or partially composed of the resource types that are under the management of the agents that support the HA fire drill. These resources must have VFDs to perform a successful HA fire drill.

You can schedule an HA fire drill on a service group, or manually run it whenever you want.

Note: If the HA fire drill passes, it does not guarantee that the service group will be online on the target system; there may be some factors outside Cluster Server control that prevent it from being online. However, if the HA fire drill fails, it is quite likely that the service groups will not come online on that system.

See "Running the high availability fire drill" on page 426.

Running the high availability fire drill

In the Management Server console, you can run the high availability (HA) fire drill on a service group.

To perform this task, your user group must be assigned the Admin or Operator role on the service group or the Availability perspective. The permission on the service group may be explicitly assigned or inherited from a parent Organization or cluster.

To run the high availability fire drill

1. In the Management Server console, go to the **Availability** perspective and select **Manage** in the left pane.
2. Expand the Organization or **Uncategorized Clusters** to locate the cluster.
3. Expand the cluster and then the **Service groups** node to locate the service group.
4. Right-click on the service group and select **Fire Drill** > **Run HA Fire Drill**.
5. In the **Fire Drill** wizard panel, click **Next**.
6. In the **Schedule** wizard panel, specify the schedule to run the HA fire drill and click **Next**.

   See "Schedule panel options" on page 430.
7. In the **Summary** panel, click **Finish**.

   See "Summary panel options" on page 431.
8. In the **Result** panel, click **OK**.

See "About high availability fire drills" on page 425.

See "Running the disaster recovery fire drill" on page 428.

See "Editing a fire drill schedule" on page 432.

See "Deleting fire drill schedules" on page 434.

See "Enabling fire drill schedules" on page 434.

See "Disabling fire drill schedules" on page 435.

About disaster recovery fire drills

The disaster recovery (DR) fire drill feature is a DR solution in Veritas InfoScale Operations Manager. The DR fire drill verifies the ability of a globally configured service group to fail over on a remote cluster, or that of a stretched service group to come online on the remote site in the same campus cluster. A DR fire drill is a zero-downtime test that mimics the configuration, application data, and failover behavior of critical service groups. A successful DR fire drill indicates that it is highly likely for a critical service group to fail over as intended or as configured on to a remote cluster, when it is needed.

The DR fire drill feature lets you do the following:

- Verify that replication for an application works correctly.
- Verify that a DR service group can be brought online successfully.

To perform a DR fire drill on the service groups, you must create a fire drill service group on the remote cluster. The configuration of the fire drill service group is similar to the configuration of the original service group.

The objective of the DR fire drill is to bring the fire drill service group online on the remote cluster. The result of this operation verifies the ability of the similarly-configured service group to fail over and come online on the remote cluster. When the DR fire drill group comes online, it uses a snapshot of the application data, which is a point-in-time copy of the replicated production data for the application. Fire drill service groups do not interact with outside clients or with other instances of resources. Therefore, they can come online safely even when the service group is online.

To ensure the failover, it is recommended that you disable the DR fire drill when they are not in use.

See “Running the disaster recovery fire drill” on page 428.

About configuring a fire drill service group

If you want to run the disaster recovery (DR) fire drill on a service group in your data center, it is important that you create a fire drill service group using Cluster Server. The configuration of the fire drill service group is similar to the configuration of the original service group.

When you use the fire drill service group in Veritas InfoScale Operations Manager for running DR fire drill, ensure the following:

- The service group on which you want to run the DR fire drill is configured on a remote cluster

- The fire drill service group has the same name as the global service group with a suffix, "_fd". For the service groups belong to the Windows clusters earlier to the 5.1SP1 release, the name of the fire drill group must be in the following format: "FDxx_ sg". Here, "FD" identifies the name of the fire drill service group, "xx" identifies any number from 00 to 99, and "sg" identifies the name of the service groups on which you want to run the DR fire drill. For the later Windows releases, the name of the fire drill service group must have the same name as the global service group with a suffix, "_fd".
- The `UserStrGlobal` attribute of the fire drill group contains the string “FD:app_group_name”. The “FD:app_group_name” identifies the name of the service group on the remote cluster.
- The fire drill service group has an offline-local dependency with the service group on the remote cluster.

See “About disaster recovery fire drills” on page 427.

Running the disaster recovery fire drill

In Veritas InfoScale Operations Manager, you can perform disaster recovery (DR) fire drills on the service groups.

To perform this task, your user group must be assigned the Admin or Operator role on the service group or the Availability perspective. The permission on the service group may be explicitly assigned or inherited from a parent Organization or cluster.

To run the disaster recovery fire drill

1. In the Management Server console, go to the **Availability** perspective and select **Manage** in the left pane.
2. Expand the Organization or **Uncategorized Clusters** to locate the cluster.
3. Expand the cluster and then the **Service Groups** node to locate the service group.
4. Right-click on the service group and select **Fire Drill** > **Run DR Fire Drill**.
5. In the **Select Remote Clusters and Global Service Groups** wizard panel, specify the required details and click **Next**.

   See “Select remote clusters and global service groups panel options” on page 429.
6. In the **Schedule** wizard panel, specify the schedule to run the DR fire drill and click **Next**.

   See “Schedule panel options” on page 430.

7 In the **Summary** panel, verify the selections that you have made and click **Finish**.

See "Summary panel options" on page 431.

8 In the **Result** panel, click **OK**.

See "About high availability fire drills" on page 425.

See "Running the high availability fire drill" on page 426.

See "Editing a fire drill schedule" on page 432.

See "Deleting fire drill schedules" on page 434.

See "Enabling fire drill schedules" on page 434.

See "Disabling fire drill schedules" on page 435.

Select remote clusters and global service groups panel options

Use this wizard panel to select the remote clusters and the global service groups where you want to run the disaster recovery (DR) fire drill.

Table 23-1 Select remote clusters and global service groups panel options

Field	Description
Remote Clusters of the cluster [cluster name]	Lists the remote clusters that are available with Veritas InfoScale Operations Manager. Select the check box corresponding to the name of the cluster and select the name of the host from the drop-down in the **Run fire drill on host** column. You can select **Any System** from the drop-down list to run the DR fire drill on any of the systems on the cluster.
Global service groups	Lists the global service groups available on the selected clusters. This table displays the name, state, and type of the service groups, and also the primary and the remote clusters where the service groups are configured. The DR fire drill is performed on the remote cluster. Select the check boxes to select the service groups for which you want to run the DR fire drill.

Table 23-1 Select remote clusters and global service groups panel options *(continued)*

Field	Description
Reset the fire drill test after successful run	Select this check box to take the fire drill service group offline immediately after a successful fire drill.

See "Running the disaster recovery fire drill" on page 428.

See "About configuring a fire drill service group" on page 427.

Schedule panel options

Use this wizard panel to create a fire drill schedule or run a fire drill on the selected service group in the cluster.

Select **Run Now** to run the fire drill immediately after you set it up. Select **Schedule for later** to specify a schedule for running the fire drill later.

Table 23-2 Schedule panel options

Field	Description
Schedule Name	Enter a name for the schedule for running the fire drill.
Schedule Desc	Enter the schedule description for running the fire drill.
Frequency	Select a frequency for running the fire drill. The values under the **When** column changes with the options that you choose here. The available options are once, daily, weekly, and monthly.

Table 23-2 Schedule panel options *(continued)*

Field	Description
When	Specify exactly when you want to run the fire drill as follows: ■ For daily schedule: Select the frequency of the fire drill if you want to run the fire drill multiple times in a day. Select the time from the **Time** field and start date from the **Start Date** field. ■ For weekly schedule: Select the **Every weekday** option to run the fire drill every week days from Monday to Friday in a week. If you want to run the fire drill on specific days of the week, select the day from the **Recur every week on** field. Select the time from the **Time** field and start date from the **Start Date** field. ■ For monthly schedule: To run the fire drill on a specific day of the month, enter the date in the **Day** field. To run the fire drill on the recurring days of a month, choose the required options from the drop-down list. Select the time from the **Time** field and start date from the **Start Date** field.

See “Running the high availability fire drill” on page 426.

See “Running the disaster recovery fire drill” on page 428.

Summary panel options

Use this wizard panel to verify your selection for running the fire drill on the selected service groups. Click **Finish**.

See “Editing a fire drill schedule” on page 432.

See “Running the high availability fire drill” on page 426.

See “Running the disaster recovery fire drill” on page 428.

Editing a fire drill schedule

In the Management Server console, you can edit the schedules that you have created for running a high availability (HA) or a disaster recovery (DR) fire drill. You can identify the HA and the DR fire drills from the **Category** column in this page.

To perform this task, your user group must be assigned the Admin or Operator role on the service group or the Availability perspective. The permission on the service group may be explicitly assigned or inherited from a parent Organization or cluster.

To modify a fire drill schedule

1. In the Management Server console, go to the **Availability** perspective and select **Manage** in the left pane.
2. Expand the Organization or **Uncategorized Clusters** to locate the cluster.
3. Expand the cluster and select the **Service groups** node.
4. Select **Fire Drill Schedules** tab.
5. In the schedules list, right-click on the required fire drill schedule. Select **Edit**.
6. In the **Edit Schedule** wizard panel, edit the schedule to run the fire drill. Click **OK**.

   See "Edit Schedule panel options" on page 432.
7. In the **Summary** panel, verify the selections that you have made. Click **Finish**.

   See "Summary panel options" on page 431.
8. In the **Result** panel, click **OK**.

See "Deleting fire drill schedules" on page 434.

See "Enabling fire drill schedules" on page 434.

See "Disabling fire drill schedules" on page 435.

See "Running the disaster recovery fire drill" on page 428.

See "Running the disaster recovery fire drill" on page 428.

Edit Schedule panel options

Use this wizard panel to edit a fire drill schedule that you have created earlier.

Table 23-3 Edit Schedule panel options

Field	Description
Schedule Name	Displays the name of the fire drill schedule.

Table 23-3 Edit Schedule panel options *(continued)*

Field	Description
Schedule Desc	Edit the schedule description for running the fire drill.
Frequency	Select a frequency for running the fire drill. The values under the **When** column changes with the options that you choose here. The available options are daily, weekly, and monthly.
When	Specify exactly when you want to run the fire drill as follows: ■ For daily schedule: Select the frequency of the fire drill if you want to run the fire drill multiple times in a day. Select the time from the **Time** field and start date from the **Start Date** field. ■ For weekly schedule: Select the **Every weekday** option to run the fire drill every week days from Monday to Friday in a week. If you want to run the fire drill on specific days of the week, select the day from the **Recur every week on** field. Select the time from the **Time** field and start date from the **Start Date** field. ■ For monthly schedule: To run the fire drill on a specific day of the month, enter the date in the **Day** field. To run the fire drill on the recurring days of a month, choose the required options from the drop-down list. Select the time from the **Time** field and start date from the **Start Date** field.

See “Editing a fire drill schedule” on page 432.

See “Viewing fire drill schedules” on page 436.

See “Deleting fire drill schedules” on page 434.

See “Running the high availability fire drill” on page 426.

See “Running the disaster recovery fire drill” on page 428.

Deleting fire drill schedules

In the Management Server console, you can delete the schedules that you have created for running a high availability (HA) or a disaster recovery (DR) fire drill.

To perform this task, your user group must be assigned the Admin or Operator role on the service group or the Availability perspective. The permission on the service group may be explicitly assigned or inherited from a parent Organization or cluster.

To delete a fire drill schedule

1. In the Management Server console, go to the **Availability** perspective and select **Manage** in the left pane.
2. Expand the Organization or **Uncategorized Clusters** to locate the cluster.
3. Expand the cluster and select the **Service groups** node.
4. Select **Fire Drill Schedules** tab.
5. In the schedules list, right-click on the required fire drill schedule. Select **Delete**.
6. In the **Delete Schedule** wizard panel, click **OK**.
7. In the **Result** panel, click **OK**.

See “Editing a fire drill schedule” on page 432.

See “Enabling fire drill schedules” on page 434.

See “Disabling fire drill schedules” on page 435.

See “Running the disaster recovery fire drill” on page 428.

See “Running the disaster recovery fire drill” on page 428.

Enabling fire drill schedules

In the Management Server console, you can enable the disabled schedules that you have created for running a high availability (HA) or a disaster recovery (DR) fire drill.

To perform this task, your user group must be assigned the Admin or Operator role on the service group or the Availability perspective. The permission on the service group may be explicitly assigned or inherited from a parent Organization or cluster.

To enable a fire drill schedule

1. In the Management Server console, go to the **Availability** perspective and select **Manage** in the left pane.
2. Expand the Organization or **Uncategorized Clusters** to locate the cluster.

3 Expand the cluster and select the **Service groups** node.

4 Select **Fire Drill Schedules** tab.

5 In the schedules list, right-click on the required fire drill schedule. Select **Enable**.

6 In the **Enable Schedule(s)** panel, click **OK**.

7 In the **Result** panel, click **OK**.

See "Editing a fire drill schedule" on page 432.

See "Deleting fire drill schedules" on page 434.

See "Disabling fire drill schedules" on page 435.

See "Running the disaster recovery fire drill" on page 428.

See "Running the disaster recovery fire drill" on page 428.

Disabling fire drill schedules

In the Management Server console, you can delete the schedules that you have created for running a high availability (HA) or a disaster recovery (DR) fire drill.

To perform this task, your user group must be assigned the Admin or Operator role on the service group or the Availability perspective. The permission on the service group may be explicitly assigned or inherited from a parent Organization or cluster.

To disable a fire drill schedule

1 In the Management Server console, go to the **Availability** perspective and select **Manage** in the left pane.

2 Expand the Organization or **Uncategorized Clusters** to locate the cluster.

3 Expand the cluster and select the **Service groups** node.

4 Select **Fire Drill Schedules** tab.

5 In the schedules list, right-click on the required fire drill schedule. Select **Disable**.

6 In the **Disable Schedule(s)** wizard panel, click **OK**.

7 In the **Result** panel, click **OK**.

See "Editing a fire drill schedule" on page 432.

See "Deleting fire drill schedules" on page 434.

See "Enabling fire drill schedules" on page 434.

See "Running the disaster recovery fire drill" on page 428.

See "Running the disaster recovery fire drill" on page 428.

Viewing fire drill schedules

In the Management Server console, you can view the schedules that you have created for the high availability and the disaster recovery fire drills.

From this view, you can edit, delete, enable, or disable a fire drill schedule.

This page displays the following information for a fire drill schedule:

Name	Name of the fire drill schedule.
Recurrence	Schedule of the fire drill run.
Service Group	Name of the service group that is associated with the fire drill.
Start Date	Date when the fire drill is scheduled to start.
End Date	Date when the fire drill is scheduled to end.
Status	The current status of the fire drill. Indicates whether the fire drill is enabled, disabled, or invalid.
Category	The category of the fire drill. Indicates whether the schedule is for a high availability fire drill or a disaster recovery fire drill.

To view the fire drill schedules and details

1. In the Management Server console, go to the **Availability** perspective and select **Manage** in the left pane.
2. Expand the Organization or **Uncategorized Clusters** to locate the cluster.
3. Expand the cluster and select the **Service groups** node.
4. Select **Fire Drill Schedules** tab.

See “Editing a fire drill schedule” on page 432.

See “Deleting fire drill schedules” on page 434.

See “Enabling fire drill schedules” on page 434.

See “Disabling fire drill schedules” on page 435.

See “Running the disaster recovery fire drill” on page 428.

See “Running the disaster recovery fire drill” on page 428.

Chapter 24

Using recovery plans

This chapter includes the following topics:

- About recovery plans
- Creating recovery plans
- Editing recovery plans
- Running recovery plans
- Deleting a recovery plan
- Viewing historical runs of recovery plans
- Viewing properties of recovery plans
- About recovery plan log files

About recovery plans

You can use recovery plans to enhance active disaster recovery capabilities for the objects in your data center. This automation lets you group multiple Virtual Business Services (VBS) and service groups and run the predefined tasks on these entities in the desired sequence.

You can specify the following tasks in a recovery plan:

- Start VBS or stop VBS
- Bring a service group online or take a service group offline
- Run custom scripts

You can run the custom scripts on the managed hosts that are a part of the recovery plan. You can add multiple such custom scripts to your recovery plan. Every time

a script is run, it returns a value. Table 24-1 lists the various return values with corresponding platforms and descriptions.

Table 24-1

Return value	Platform	Description
0	Windows, Unix/Linux	Script executed successfully when no explicit return code is specified in the script.
-1	Windows	Unable to fork the script.
-2	Windows, Unix/Linux	Script did not complete within the timeout specified.
-3	Unix/Linux	Unable to fork the script.

A default value of 5 minutes is considered as timeout for any script. This timeout can be configured while running the recovery plan. If the script does not return a value within five minutes or takes more than five minutes to run, then the task is marked as a failure.

Using the Veritas InfoScale Operations Manager console, you can create, edit, delete, and run the recovery plans. The Veritas InfoScale Operations Manager console provides the progress and the status of a recovery plan.

See “Creating recovery plans” on page 438.

Creating recovery plans

The Management Server lets you create a recovery plan and view the list of Virtual Business Services, service groups, or custom scripts. You can start or stop Virtual Business Services and bring service groups online or offline. To run a custom script, the name of the custom script and on which host it is run, needs to be specified. You can also specify whether a task is critical or not.

To perform this task, your user group must be assigned the Admin role on the cluster or the Availability perspective. The permission on the service group may be explicitly assigned or inherited from a parent Organization or cluster.

If your user group is assigned the Operator or Guest role on the cluster, you can only view the Virtual Business Services in the tree node, but not use them to perform this operation.

To create a recovery plan

1. In the Management Server console, go to the **Availability** perspective and expand **Solutions** in the left pane.
2. Do one of the following:
   - Click on **Recovery Plan** to view a list of all recovery plans. Right-click on a recovery plan and select **Create**.
   - Click on **Recovery Plan**. If there are no existing recovery plans listed, right-click on the pane and click **Create**.
3. In the **Create Recovery Plan** wizard panel, specify the name, description, and other task details of the recovery plan. Click **Finish**.

   See "Create recovery plan panel options" on page 439.
4. On the **Result** panel, verify that the recovery plan has been successfully created. Click **OK**.

See "Editing recovery plans" on page 442.

See "Deleting a recovery plan" on page 446.

See "Running recovery plans" on page 444.

Create recovery plan panel options

Use this wizard panel to specify a name, description, and tasks for the recovery plan. Table 24-2 describes the options for creating a recovery plan. Table 24-3 describes the properties of the tasks that are defined as part of the recovery plan.

Table 24-2 Create recovery plan

Field	Description
Name	Enter a name for the recovery plan.
Description	Enter the description of the recovery plan.
Add	Click to add a Virtual Business Service (VBS), service group, or a script to the recovery plan. An empty VBS does not get listed for selection when you try to create a recovery plan.

Table 24-2 Create recovery plan *(continued)*

Field	Description
Import	Click to import an existing recovery plan. The tasks of the imported recovery plan along with the associated objects, are appended to the list of tasks. If you import a script task, the script is also copied for the new recovery plan. Following are some exceptions where only the template (tasks without objects) of the recovery plan is imported: ■ If you do not have access privileges on any of the objects (Virtual Business Service, service group, or host) that are associated with a task of the imported recovery plan. ■ If any of the objects that are associated with a task of the imported recovery plan have been deleted. In these conditions, you can select any other object of the same object type for that particular task.
Delete	Click to delete a Virtual Business Service, service group, or a script task from the recovery plan.
Move Up	Click to change the order of the tasks within the recovery plan upwards.
Move Down	Click to change the order of the tasks within the recovery plan downwards.
Type	Select the type of recovery plan object from the drop-down list.
Name	Specify the name of the recovery plan object selected.

Table 24-2 Create recovery plan *(continued)*

Field	Description
Action	Select an action from the drop-down list. This drop-down list displays the actions that can be performed on the objects with in recovery plan. For a site-aware service group, you can select the site on which the online/offline operation needs to be performed. For a non site-aware service group, you can choose to online/offline on any system. For VBS and script, available actions are start/stop and execute respectively. Site-aware service group operations are available only if the systems that belong to the cluster have the managed host version 6.1 installed on them.

Table 24-3 Task Properties

Description	You can add the description of the task that is selected.
Critical	If a task marked as critical fails, recovery plan execution is aborted and remaining tasks are not executed. If a task marked as non-critical fails, recovery plan execution continues with the remaining tasks.
Timeout (mins)	Default value for **Timeout** is **5 minutes**. You can modify the timeout value. When you specify the timeout, you need to ensure that the task gets executed within the specified time. The task may or may not complete successfully, but recovery plan does not track any changes after the expiry of specified timeout.

See “Creating recovery plans” on page 438.

Editing recovery plans

The Management Server console lets you edit a recovery plan and view the list of Virtual Business Services, service groups, or custom scripts. You can start or stop Virtual Business Services and bring service groups online or offline. You can specify whether a task is critical or not and also run custom scripts. To run a custom script, you must specify the name of the custom script and on which host it is to be run. You can also add, delete, or change the order of the tasks displayed in the recovery plan from the list displayed.

To perform this task, your user group must be assigned the Admin role on the cluster or the Availability perspective. The permission on the service group may be explicitly assigned or inherited from a parent Organization or cluster.

To edit a recovery plan

1. In the Management Server console, go to the **Availability** perspective and expand **Solutions** in the left pane.
2. Click on **Recovery plan** to view a list of all recovery plans.
3. Right-click on the required recovery plan and select **Edit**.
4. In the **Edit Recovery Plan** wizard panel, specify the name, description, and tasks of the recovery plan. Click **Finish**.

   See "Edit recovery plan panel options" on page 442.
5. On the **Summary** panel verify that the recovery plan has been successfully edited, click **OK**.

See "Creating recovery plans" on page 438.

See "Deleting a recovery plan" on page 446.

See "Running recovery plans" on page 444.

Edit recovery plan panel options

Use this wizard panel to modify a name, description, tasks, and task properties for the recovery plan. Table 24-4 describes the options for editing a recovery plan. Table 24-5 describes the properties of the tasks that are defined as part of the recovery plan.

Table 24-4 Edit recovery plan

Field	Description
Name	Modify the name for the recovery plan.

Table 24-4 Edit recovery plan *(continued)*

Field	Description
Description	Modify the description of the recovery plan.
Add	Click to add a Virtual Business Service, service group, or a script to the recovery plan. An empty VBS does not get listed for selection when you try to edit a recovery plan.
Import	Click to import an existing recovery plan. The tasks of the existing recovery plan are appended to the list of tasks.
Delete	Click to delete a Virtual Business Service, service group, or a script from the recovery plan.
Move Up	Click to change the order of the tasks within the recovery plan upwards.
Move Down	Click to change the order of the tasks within the recovery plan downwards.
Type	Select the type of recovery plan object from the drop-down list.
Name	Specify the name of the recovery plan object selected.
Action	Select an action from the drop-down list. This drop-down list displays the actions that can be performed on the objects with in recovery plan. For a site-aware service group, you can select the site on which the online/offline operation needs to be performed. For a non site-aware service group, you can choose to online/offline on any system. For VBS and script, available actions are start/stop and execute respectively.

Table 24-5 Task Properties

Description	You can modify the description of the task that is selected.

Table 24-5 Task Properties *(continued)*

Critical	Default value for **Critical** is **No**. you can modify the default value. If a task is marked as critical and it fails, then the recovery plan is aborted and the remaining tasks are not run. If the task is marked as not critical, then even if the task fails, the recovery plan execution is not stopped.
Timeout (mins)	Default value for **Timeout** is **5 minutes**. You can modify the timeout value. When you specify the timeout, you need to ensure that the task gets executed out within the specified time. The task gets executed in the background and the recovery plan does not track any changes that may happen to the task from the backend. Recovery plan does not track even if the forked process that is executed from the recovery plan for script execution gets killed in the backend.

See “Editing recovery plans” on page 442.

Running recovery plans

The Management Server console lets you run a recovery plan. If you do not have permission to run a task specified in the recovery plan, then the recovery plan fails to run. Execution of recovery plan is blocked if you do not have access to any of the objects that are part of the recovery plan, or the objects have been deleted. Appropriate error messages are logged. After resolving the errors, you can run the recovery plan again and choose to skip the completed tasks. The operation can be aborted.

To perform this task, your user group must be assigned the Admin or Operator role on the cluster or the Availability perspective. The permission on the cluster may be explicitly assigned or inherited from a parent organization.

To run a recovery plan

1. In the Management Server console, go to the **Availability** perspective and expand **Solutions** in the left pane.
2. Click on **Recovery Plan** to view a list of all recovery plans.
3. Right-click on the required recovery plan and select **Execute**.

4 In the **Execute Recovery Plan** wizard panel, confirm the recovery plan that you want to run. Also, indicate if any of the tasks need to be skipped. Click **Finish**.

See "Run recovery plan panel options" on page 445.

5 On the **Result** panel, verify that the recovery plan was run successfully. Click **OK**.

See "Creating recovery plans" on page 438.

See "Editing recovery plans" on page 442.

See "Deleting a recovery plan" on page 446.

Run recovery plan panel options

Use this wizard panel to review tasks for the recovery plan. You can also indicate if any of the tasks need to be skipped.

Table 24-6 Running recovery plans

Field	Description
Type	View the type of task selected.
Name	View the name of the selected object.
Action	View the action selected.
Critical	View the criticality of the task.
Timeout (min)	Specify the duration you want to wait while the script, Virtual Business Service, or service group task runs. Recovery plan stops after the expiry of the timeout that is specified. Veritas InfoScale Operations Manager does not track the status of the tasks after the timeout. Change in timeout while running a recovery plan is applicable only for the current run/execution.
Skip	Select to skip the task execution from the current run.
Offline shared service group(s) while stopping any VBS	Select the check box to make the shared service groups offline while stopping any Virtual Business Service.

See “Running recovery plans” on page 444.

Deleting a recovery plan

The Management Server console lets you delete an existing recovery plan if it is not required. You can delete only those recovery plans that are currently not in execution state. You can either delete a single recovery plan or multiple recovery plans.

To perform this task, your user group must be assigned the Admin role on the cluster or the Availability perspective. The permission on the service group may be explicitly assigned or inherited from a parent Organization or cluster.

To delete a recovery plan

1. In the Management Server console, go to the **Availability** perspective and expand **Solutions** in the left pane.
2. Click on **Recovery Plan** to view a list of all recovery plans.
3. Right-click on the required recovery plan and select **Delete**.
4. In the **Delete Recovery Plan** wizard panel, confirm that you want to delete the recovery plan. Click **OK**.
5. On the **Result** panel, verify that the recovery plan has been successfully deleted. Click **OK**.

See “Creating recovery plans” on page 438.

See “Editing recovery plans” on page 442.

See “Running recovery plans” on page 444.

Viewing historical runs of recovery plans

The Management Server console lets you view the historical runs of a recovery plan.

To perform this task, your user group must be assigned at least Guest role on the cluster or the Availability perspective. The permission on the cluster may be explicitly assigned or inherited from a parent organization.

To view historical runs of a recovery plan

1. In the Management Server console, go to the **Availability** perspective and expand **Solutions** in the left pane.
2. Click on **Recovery Plan** to view a list of all recovery plans.

3 Right-click on the required recovery plan and select **Historical Runs**.

4 In the **Historical Runs** wizard panel, view the historical details of the recovery plan. Click **Close**.

See "Historical Runs panel options" on page 447.

See "Creating recovery plans" on page 438.

See "Editing recovery plans" on page 442.

See "Deleting a recovery plan" on page 446.

Historical Runs panel options

Use this wizard panel to view the historical details of the recovery plan.

The **Historical Runs** view of the recovery plan displays the following details:

Table 24-7 Historical Runs panel options

Field	Description
State	The state of the recovery plan for this particular execution instance.
Source	The source host details of the recovery plan task.
User	The user name of the user who initiated the recovery plan execution.
Start Time	The date and time when the recovery plan execution is started.
End Time	The date and time when the recovery plan execution is completed.

See "Creating recovery plans" on page 438.

Viewing properties of recovery plans

The Management Server console lets you view the properties of a recovery plan.

To perform this task, your user group must be assigned at least Guest role on the cluster or the Availability perspective. The permission on the cluster may be explicitly assigned or inherited from a parent organization.

To view properties of a recovery plan

1. In the Management Server console, go to the **Availability** perspective and expand **Solutions** in the left pane.
2. Click on **Recovery Plan** to view a list of all recovery plans.
3. Right-click on the required recovery plan and select **Properties**.
4. In the **Recovery Plan Properties** wizard panel, view the properties of the recovery plan. Click **Close**.

   See "Recovery Plan Properties panel options" on page 448.

See "Creating recovery plans" on page 438.

See "Editing recovery plans" on page 442.

See "Deleting a recovery plan" on page 446.

Recovery Plan Properties panel options

Use this wizard panel to view the properties of the recovery plan. The fields that are displayed in this panel cannot be edited or selected to perform any task. Table 24-8 describes the properties of a recovery plan while Table 24-9 describes the properties of a task that belongs to a recovery plan.

Table 24-8 Recovery Plan Properties panel options

Field	Description
Type	View the type of task selected.
Name	View the name of the selected object. An icon is displayed along with the name to display the current status of the selected object.
Action	View the action for the task. If the object is a service group that belongs to a stretched cluster, then you can see the selected action as well as the site, where the task is performed.

Table 24-9 Task Properties

Description	View the description for the task that is selected.
Critical	View the criticality for the task that is selected.

Table 24-9 Task Properties *(continued)*

Timeout (mins)	View the timeout value for the task that is selected.

See “Creating recovery plans” on page 438.

About recovery plan log files

Veritas InfoScale Operations Manager maintains several log files that operators can use for troubleshooting. Following is a list

Table 24-10 Log files for recovery plan

Log file category	Log file name
GUI logs	/var/opt/VRTSsfmcs/logs/WebDebugLog.txt
Any unhandled exception logs	/var/opt/VRTSsfmcs/logs/tomcat.log
Script execution logs	/var/opt/VRTSsfmh/logs/rplan_script_execute.log
Push file logs	/var/opt/VRTSsfmh/logs/push_file.log

See “About recovery plans” on page 437.

Chapter 25

Managing ApplicationHA

This chapter includes the following topics:

- About ApplicationHA Management
- Prerequisites for ApplicationHA Management
- About the ApplicationHA operations
- Launching ApplicationHA operations from Veritas InfoScale Operations Manager
- About the ApplicationHA infrastructure
- Enabling the ApplicationHA infrastructure for a managed host
- Disabling the ApplicationHA infrastructure for a managed host

About ApplicationHA Management

ApplicationHA is a high availability product that provides monitoring capabilities for the applications that run inside a virtual machine. ApplicationHA uses a persistent heartbeat mechanism to communicate the state of the application to the high availability provider in the virtualization infrastructure layer.

When an application that runs inside a virtual machine fails, ApplicationHA tries to restart the application. If ApplicationHA cannot restart the application, it attempts to gracefully restart the virtual machine. If the internal restart does not restore the application, ApplicationHA sends a message to the high availability provider to externally restart the virtual machine. If that fails, the high availability provider may fail over the virtual machine to another virtualization server.

Using the Management Server console, you can perform the ApplicationHA operations in the following virtualization environments:

- VMware ESX/ESXi

- Linux Kernel Virtual Machine (KVM)
- IBM AIX Logical Partitions (LPAR)
- Oracle VM Server (OVM) for SPARC (Solaris LDOM)

The high availability provider in the VMware environment is VMwareHA. The high availability provider in the KVM, OVM Server for SPARC, and LPAR environments is Cluster Server.

To configure Veritas InfoScale Operations Manager to discover virtual machines in these virtualization environments, you must perform the following steps:

- Add the VMware or the KVM virtual machines as managed hosts to Management Server.
- Add the LPARs as managed hosts to Management Server.
- Add the logical domains in an OVM Server for SPARC as managed hosts to Management Server.

See "Prerequisites for ApplicationHA Management " on page 451.

See "About the ApplicationHA operations" on page 451.

Prerequisites for ApplicationHA Management

The LPAR agent blocks LPM functionality for the management LPAR that hosts VCS when it manages and monitors LPAR resources. When VCS stops managing LPARs, LPM functionality is available for the VCS system. The LPM functionality to migrate the management LPAR remains blocked if the LPAR agent crashes or is terminated. For more information, refer to the *Veritas™ Cluster Server Bundled Agents Reference Guide*.

See "About ApplicationHA Management" on page 450.

See "About the ApplicationHA operations" on page 451.

About the ApplicationHA operations

The Management Server console provides visibility into different applications that are discovered from physical and virtual machines. It provides the information based on the following:

- Virtualization technology: VMware, Kernel-based Virtual Machine (KVM), logical partitions (LPARs), and logical domains (LDoms)
- High availability provider: ApplicationHA and Cluster Server
- Scope: Entire domain, Organization, or at the host-level

The Management Server console provides an interface to perform the ApplicationHA operations. You can perform the operations on the applications that run on a virtual machine, logical domain, or a logical partition.

You must have Admin role assigned on the host or the Server perspective for performing the ApplicationHA operations in Veritas InfoScale Operations Manager. The permission on the host may be explicitly assigned or inherited from a parent Organization.

See "About ApplicationHA Management" on page 450.

See "Prerequisites for ApplicationHA Management " on page 451.

Launching ApplicationHA operations from Veritas InfoScale Operations Manager

Using Veritas InfoScale Operations Manager, you can perform tasks to monitor the applications on ApplicationHA guest on the virtual machines, the logical partitions, and the logical domains in your data center.

To launch ApplicationHA operations from Veritas InfoScale Operations Manager

1. In the Management Server console, go to the **Server** perspective and expand **Manage** in the left pane.
2. Expand the Organization, or **Uncategorized Hosts** to locate the host.
3. Right-click on the required host and select **Manage ApplicationHA**.
4. Click on the required operation on the left side of the wizard panel.

See "About the ApplicationHA operations" on page 451.

About the ApplicationHA infrastructure

The ApplicationHA infrastructure feature lets you create a private network between the Cluster Server (VCS) cluster and the ApplicationHA hosts in Veritas InfoScale Operations Manager. Using this private network, an ApplicationHA host communicates with the VCS cluster to inform about an application fault. Using the Veritas InfoScale Operations Manager console, you can enable or disable ApplicationHA infrastructure for a selected managed host. The nodes of the VCS that manage the ApplicationHA guest is referred to as infrastructure hosts.

The supported operating systems for this feature are Linux, AIX, and Solaris. The supported servers are KVM Server, LDom Server, and LPAR virtual machine. Also,

ensure that the managed hosts are configured as agent in the Veritas InfoScale Operations Manager.

See "Enabling the ApplicationHA infrastructure for a managed host" on page 453.

See "Disabling the ApplicationHA infrastructure for a managed host" on page 453.

Enabling the ApplicationHA infrastructure for a managed host

Using the Management Server console, you can enable the ApplicationHA infrastructure for a selected managed host. This operation creates a private network between the Cluster Server (VCS) cluster and the ApplicationHA host. Using this private network, the ApplicationHA host communicates with the VCS cluster to inform about an application fault.

You can also use this feature to configure the auto-registration option of the VCS cluster. The auto-registration feature enables automated communication over the private network between the virtual machines and the infrastructure host, provided that the ApplicationHA is configured on the virtual machines.

To perform this task, your user group must be assigned the Admin role or the Operator role on the cluster or the Availability perspective. The permission on the cluster may be explicitly assigned or inherited from a parent Organization.

This operation can be performed only on Kernel-based Virtual Machine (KVM), Logical Domains (LDOM), and logical partition (LPAR) virtual machines.

To enable the ApplicationHA infrastructure for a managed host

1. In the Management Server console, go to the **Availability** perspective and select **Manage** in the left pane.
2. Expand the Organization or **Uncategorized Clusters** to locate the cluster.
3. Expand the cluster and then expand the **Systems** node to locate the system.
4. Right-click on the system and select **ApplicationHA Infrastructure** > **Enable**.
5. In the **Enable ApplicationHA Infrastructure** panel, click **OK** to confirm.

See "Disabling the ApplicationHA infrastructure for a managed host" on page 453.

Disabling the ApplicationHA infrastructure for a managed host

Using the Management Server console, you can disable the ApplicationHA infrastructure for a selected managed host.

To perform this task, your user group must be assigned the Admin role or the Operator role on the cluster or the Availability perspective. The permission on the cluster may be explicitly assigned or inherited from a parent Organization.

This operation can be performed only on Kernel-based Virtual Machine (KVM), Logical Domains (LDOM), and logical partition (LPAR) virtual machines.

To disable the ApplicationHA infrastructure for a managed host

1. In the Management Server console, go to the **Availability** perspective and select **Manage** in the left pane.
2. Expand the Organization or **Uncategorized Clusters** to locate the cluster.
3. Expand the cluster and then expand the **Systems** node to locate the system.
4. Right-click on the system and select **ApplicationHA Infrastructure** > **Disable**.
5. In the **Disable ApplicationHA Infrastructure** panel, click **OK** to confirm.

See "Enabling the ApplicationHA infrastructure for a managed host" on page 453.

Chapter 26

Managing application configuration

This chapter includes the following topics:

- About Application Configuration
- Prerequisites for application configuration
- Launching the configure application wizard from Veritas InfoScale Operations Manager

About Application Configuration

Veritas InfoScale Operations Manager uses the VCS application configuration wizard to configure application monitoring under VCS control running in virtual or physical environments.

Note: This feature is enabled only with VCS versions 6.0.1 (Windows), 6.0.2 (Linux) and an agent pack that is installed on VCS versions 6.1.0 (Solaris, AIX) onwards. In case of Linux and UNIX, ensure that VRTSvcswiz package is installed.

You can configure application monitoring either in a start/stop mode on a single system, or a failover mode on multiple systems.

See "Prerequisites for application configuration" on page 456.

See "Launching the configure application wizard from Veritas InfoScale Operations Manager" on page 456.

Prerequisites for application configuration

Following are the prerequisites for using the application configuration feature in Veritas InfoScale Operations Manager:

- The cluster should already be configured between the systems where you want to configure application for monitoring.
- The required ports are not blocked by a firewall.
- The application is installed identically on all the systems that will be part of the VCS cluster.
- The disk drive containing the application data files and the registry replication information is present on the local system.

For more information on the prerequisites for application configuration, see the application-specific VCS agent guide. To verify that VCS supports wizard-based configuration for a particular application in your physical or virtual environment, see the *Veritas Cluster Server Release Notes*.

See "About Application Configuration" on page 455.

See "Launching the configure application wizard from Veritas InfoScale Operations Manager" on page 456.

Launching the configure application wizard from Veritas InfoScale Operations Manager

Using Veritas InfoScale Operations Manager, you can configure application monitoring running in virtual or physical environment. You can launch the wizard from a cluster or from a system.

To launch the configure application wizard from a cluster

1. In the Management Server console, go to the **Availability** perspective and expand **Manage** in the left pane.
2. Expand the Organization, or **Uncategorized Clusters** to locate the cluster.
3. Right-click on the required cluster and select **Configure Application**.
4. In the **System selection** wizard panel, select the system where the application is running.

To launch the configure application wizard from a system

1. In the Management Server console, go to the **Availability** perspective and expand **Manage** in the left pane.
2. Expand the Organization, or **Uncategorized Clusters**, and then expand **Systems** to locate the system.
3. Right-click on the required system and select **Configure Application** to launch the wizard.

See "About Application Configuration" on page 455.

See "Prerequisites for application configuration" on page 456.

Chapter 27

Multi Site Management

This chapter includes the following topics:

- About Multi-Site Management
- Features of Multi-Site Management
- Prerequisites of Multi-Site Management
- Limitations of Multi-Site Management
- Setting up a campus cluster
- Setting up a replicated data cluster
- Configuring stretch sites

About Multi-Site Management

Multi-Site Management facilitates uniform site configuration and management across Storage Foundation and Cluster Server objects. From the Veritas InfoScale Operations Manager Management Server console, you can assign site names to the enclosures and to the cluster's hosts that consume storage from the enclosure. You can also set the site fencing preferences for these hosts.

Multi-Site Management is supported on the following operating systems:

Platforms	Operating systems	Level
Linux	Red Hat Enterprise Linux 6	Update 3, 4
	Red Hat Enterprise Linux 5	Update 5, 6, 7, 8, 9
	SUSE Linux Enterprise 11	SP2
	Oracle Linux 6	Update 3, 4
	Oracle Linux 5	Update 5, 6, 7, 8. 9

For Linux, only 64-bit operating systems are supported on the AMD Opteron or the Intel Xeon EM64T (x86_64) processor line.

Solaris	Solaris 10	Update 8. 9, 10, 11
	Solaris 11	SRU1, Update 1
AIX	AIX 6.1	TL5 or later (Power 5, Power 6, or Power 7)
	AIX 7.1	TL0 or later (Any chipset that the operating system supports)
Microsoft Windows	Windows Server 2008	R2
	Windows Server 2012	R2
	Windows Server 2012	

See “Features of Multi-Site Management” on page 459.

See “Prerequisites of Multi-Site Management ” on page 460.

See “Limitations of Multi-Site Management ” on page 460.

Features of Multi-Site Management

The key features of Multi-Site Management in Veritas InfoScale Operations Manager are as follows:

- Ability to assign site names to the enclosures and the hosts of the selected cluster.
- Provision to select the type of cluster configuration - campus cluster, replicated data cluster (RDC), and mixed mode.
- Ability to add and edit site fencing preference.

- Configuration of service group dependency.
- Auto-population of site tagging during discovery of managed hosts: After a host is added to the Management Server, it discovers the existing site tags defined for the disk group and disks on each host, and populates the information in the Veritas InfoScale Operations Manager database.

See "Setting up a campus cluster" on page 461.

See "Setting up a replicated data cluster" on page 462.

Prerequisites of Multi-Site Management

Before you start using the Multi-Site Management feature in Veritas InfoScale Operations Manager, ensure that the following prerequisites are met:

- All hosts of the cluster must be using Storage Foundation High Availability version 6.0.1, or later.
- For non-Windows platforms, the Management Server and the managed hosts must be using Veritas InfoScale Operations Manager 6.0, or later. For Windows platform, the Management Server and the managed hosts must be using Veritas InfoScale Operations Manager 6.1, or later.

See "About Multi-Site Management " on page 458.

Limitations of Multi-Site Management

The limitations of Multi-Site Management are as follows:

- Site tagging of groups of entities (for example, Virtual Business Service) is not supported.
- Rule-based site tagging is not supported.
- Auto site tagging (propagation) of entities between the storage and the hosts, and vice versa is not supported.
- Site-level operations (for example, evacuate) are not supported.
- Clusters with system zones defined on service groups cannot be site tagged. You need to first clear the system zones attributes, and then tag them.
- For Windows operating system, when the user adds a disk to a disk group using the command line, the disk will not be site tagged.

See "About Multi-Site Management " on page 458.

Setting up a campus cluster

For campus cluster configuration, the steps are as follows:

Table 27-1 Setting up a campus cluster

Steps	Action	Description
Step 1	Configure Stretch Sites using the Veritas InfoScale Operations Manager Management Server console.	This step includes assigning the site tag to the storage enclosure from the **Storage** perspective and tagging systems of the required cluster from the **Availability** perspective of the Veritas InfoScale Operations Manager Management Server console. See "Configuring stretch sites" on page 463. You can create the disk group (explained in step 2) before or after tagging the enclosure and hosts. Depending on whether the disk group is already configured, step 2 presents two scenarios.
Step 2	When the disk group is already configured (with disks from different enclosures) before you perform step 1.	In this scenario, when you tag the enclosure and hosts (as explained in step 1), Veritas InfoScale Operations Manager automatically makes the disk group site consistent. **Note:** Site consistency is also configured through the Configure Stretch Cluster wizard.
	When the disk group is not configured before you perform step 1.	After tagging the enclosure and hosts (as explained in step 1), you can create a new disk group using the Veritas InfoScale Operations Manager console. While creating the disk group, an option will be provided to make the disk group site-aware. See "Configuring stretch sites" on page 463.
Step 3	Create mirrored volume on the disk group.	Create mirrored volumes on the disk group created above. It is required for end-to-end setup of campus cluster.

Note: If the site tags are removed using VxVM or VCS command line, Veritas InfoScale Operations Manager still shows the site tags on the Veritas InfoScale Operations Manager console. To get the updated status, wait for 30 minutes, or refresh the host on the Veritas InfoScale Operations Manager console. For more information and instructions on site tag removal, refer to *Storage Foundation High Availability Solutions Disaster Recovery Implementation Guide*.

See "Features of Multi-Site Management" on page 459.

See "Setting up a replicated data cluster" on page 462.

Setting up a replicated data cluster

For replicated data cluster configuration, the steps are as follows:

Table 27-2 Setting up a replicated data cluster

Steps	Action	Description
Step 1	Set up replication.	It can be done using: ■ Array-based replication. For example, EMC Symmetrix Remote Data Facility (SRDF) configuration. Refer to the vendor documentation for more information. ■ Replication using Volume Replicator. Refer to *Cluster Server Administrator's Guide* for the information on setting-up a replicated data cluster configuration.
Step 2	Configuring site tag using the Veritas InfoScale Operations Manager console.	It includes site tagging the cluster systems, and setting their preferences from the **Availability** perspective. See "Configuring stretch sites" on page 463.

For more information, refer to *Storage Foundation High Availability Solutions Disaster Recovery Implementation Guide*.

See "Features of Multi-Site Management" on page 459.

See "Setting up a campus cluster" on page 461.

Configuring stretch sites

You can use the Veritas InfoScale Operations Manager Management server console to configure stretch sites. The following three configurations are supported:

- For campus cluster, first tag the enclosure with the site name and thereafter, assign the same site tag to the cluster's hosts. Set the required site fencing preferences. Storage Foundation and Cluster Server commands are run for the campus cluster configuration.
- For replicated data cluster (RDC) configuration, assign site tags for the nodes of the cluster and set the site fencing preferences. Enclosure tagging is not required for RDC configuration. Cluster Server commands are run for the RDC configuration.
- For mixed mode (a combination of Storage Foundation volume mirroring and array or host-based replication), tag the enclosure with the site name and assign the same tag to the cluster's hosts. Set the required site fencing preferences. Like campus cluster, Storage Foundation and Cluster Server commands are run for the mixed mode.

Note: The remove site tag operation is not supported through the Veritas InfoScale Operations Manager console. Refer to the *Storage Foundation High Availability Solutions Disaster Recovery Implementation Guide* for more information.

To configure stretch sites using Veritas InfoScale Operations Manager

---

Note: For campus cluster and mixed configurations, run all the following steps. For RDC configuration (since enclosure tagging is not required), go to the step 3.

---

1. In the Management Server console, go to the **Storage** perspective, and locate the required enclosures. Enclosures from both sites in the Veritas InfoScale Operations Manager Management Server domain should be tagged. The disks of these enclosures are or will be a part of the disk groups that are used for the campus cluster.
2. Right-click the enclosure, and select **Assign Sites**.

   See "Enclosure site assignment panel options" on page 464.
3. Go to the **Availability** perspective, and locate the cluster that uses storage from the enclosure. Right-click the cluster and then select **Configure Stretch Sites**.

4 In the **Specify Cluster Type** panel, select the required cluster configuration. Click **Next**.

See "Select cluster type panel options" on page 464.

5 In the **Assign Sites to Systems** panel, select the site tag for each system of the cluster. Click **Next**.

See "Assign site to the systems panel options" on page 465.

6 In the **Specify Site Fencing Preference** panel, set the preferences for the sites.

See "Site fencing preference panel options" on page 466.

7 Click **Finish** to complete the site assignment for the systems. In the **Result** panel click **Show Commands** to view all commands that are run for the site configuration operation. If any command fails to run, the wizard does not run subsequent commands.

See "Features of Multi-Site Management" on page 459.

See "Setting up a campus cluster" on page 461.

See "Setting up a replicated data cluster" on page 462.

Enclosure site assignment panel options

Use this panel to assign the site name to the enclosure which is used by the cluster. You can select the site name from the **Site** drop-down list, or you can provide a new site name.

See "Configuring stretch sites" on page 463.

Select cluster type panel options

Use this panel to specify the type of cluster for the site assignment operation.

Table 27-3 Select cluster type panel options

Field	Description
Campus cluster	For campus cluster configuration, the data is copied between the sites using the Storage Foundation volume mirroring feature. For a campus cluster, you need to tag the enclosure and systems. Storage Foundation and Cluster Server commands are run for the campus cluster configuration. Ensure that the enclosures that are associated with the selected campus cluster are already tagged.
Replicated data cluster (RDC)	For RDC, only host based or array-level replication happens. Cluster Server commands are run (that is, only VCS objects are tagged).
Mixed mode	For mixed mode, the data is copied between sites using the combination of Storage Foundation volume mirroring and host or array-based replication technologies.

See "Configuring stretch sites" on page 463.

See "Setting up a campus cluster" on page 461.

See "Setting up a replicated data cluster" on page 462.

Assign site to the systems panel options

Use this panel to assign the site tag for the systems of the selected cluster.

Table 27-4 Assign site to the systems panel options

Field	Description
System	Displays the systems from the selected cluster.

Table 27-4 Assign site to the systems panel options *(continued)*

Field	Description
Site	Lists the tags at the enclosure level. For example, if the system gets storage from enclosure E1 (with tag T1) and E2 (with tag T2), the drop-down list shows T1 and T2. **Note:** If you want to provide a new tag for the system, you can edit the tag name.

See "Configuring stretch sites" on page 463.

See "Setting up a campus cluster" on page 461.

See "Setting up a replicated data cluster" on page 462.

Site fencing preference panel options

Use this panel to set the site preferences for the systems of the selected cluster.

Table 27-5 Site fencing preference panel options

Field	Description
Configure Site Fencing Preference	Select the check box to enable the site fencing preference. If the `UseFence` cluster attribute is set to SCSCI3, the **Configure Site Fencing Preference** check box is available. If the `UseFence` cluster attribute is set to None, the check box is not available, and an appropriate message is displayed.
Site	The name of the site tag that is set for the system.
Preference	Select the site preference for the site - highest, high, medium, and low.

Note: The host site fencing preference option is not supported for Microsoft Windows operating system. It will be supported in the upcoming releases of Veritas InfoScale Operations Manager. You can still set these preferences. However, they will be effective only when Cluster Server (VCS) supports them for Windows operating system.

See "Configuring stretch sites" on page 463.

See "Setting up a campus cluster" on page 461.

See "Setting up a replicated data cluster" on page 462.

Appendix A

List of high availability operations

This appendix includes the following topics:

- Cluster operations
- System operations
- Service group operations
- Resource operations

Cluster operations

Table A-1 lists the operations that you can perform on clusters, and indicates whether the operation can be performed on a single cluster, or on multiple clusters.

Table A-1 Cluster operations

Operation	Supported for single or multiple clusters?
Opening a cluster configuration	Single
Saving a cluster configuration See "Saving a cluster configuration" on page 351.	Single
Closing a cluster configuration See "Closing a cluster configuration" on page 352.	Single

Table A-1 Cluster operations *(continued)*

Operation	Supported for single or multiple clusters?
Editing attributes of a cluster See "Editing attributes of a cluster" on page 352.	Single
Setting up GCO between clusters See "Adding a remote cluster to a local cluster" on page 418.	Multiple
Importing type definition See "Importing a type definition" on page 354.	Single
Removing GCO between clusters See "About removing a remote cluster from a global cluster setup" on page 421.	Multiple

System operations

Table A-2 lists the operations that you can perform on systems, and indicates whether the operation can be performed on a single system, or on multiple systems.

Table A-2 System operations

Operation	Supported for single or multiple cluster hosts?
Freezing systems See "Freezing a system" on page 391.	Single and multiple
Unfreezing systems See "Unfreezing a system" on page 392.	Single and multiple
Editing attributes of a system See "Editing attributes of a system" on page 393.	Single

Table A-2 System operations *(continued)*

Operation	Supported for single or multiple cluster hosts?
Starting the Cluster Server high availability daemon on the systems in a cluster See "Starting the Cluster Server high availability daemon on the hosts in a cluster" on page 393.	Single
Stopping the Cluster Server high availability daemon on the systems in a cluster See "Stopping the Cluster Server high availability daemon on the systems in a cluster" on page 394.	Single

Service group operations

Table A-3 lists the operations that you can perform on service groups, and indicates whether the operation can be performed on a single service group, or on multiple service groups.

Table A-3 Service group operations

Operation	Supported for single or multiple service groups?
Adding service groups to a cluster See "Creating service groups" on page 358.	Single
Modifying service groups See "About modifying a service group" on page 389.	Single
Enabling service groups See "Enabling service groups" on page 363.	Single and multiple
Disabling service groups See "Disabling service groups" on page 364.	Single and multiple
Modifying the System List for a service group See "Modifying the system list for a service group" on page 386.	Single

Table A-3 Service group operations *(continued)*

Operation	Supported for single or multiple service groups?
Enabling all resources in a service group See "Enabling all resources of service groups" on page 370.	Single and multiple
Disabling all resources in a service group See "Disabling all resources of service groups" on page 371.	Single and multiple
Bringing service groups online See "Bringing service groups online" on page 377.	Single and multiple
Taking service groups offline See "Taking service groups offline" on page 379.	Single and multiple
Freezing service groups See "Freezing service groups" on page 367.	Single and multiple
Unfreezing service groups See "Unfreezing service groups" on page 368.	Single and multiple
Linking service groups See "Linking service groups in a cluster" on page 372.	Single
Unlinking service groups See "Unlinking service groups" on page 375.	Single
Switching service groups See "Switching service groups" on page 381.	Single and multiple
Flushing a service group See "Flushing service groups" on page 369.	Single
Autoenabling a service group See "Autoenabling service groups" on page 367.	Single

Table A-3 Service group operations *(continued)*

Operation	Supported for single or multiple service groups?
Clearing faults on a service group See "Clearing faults on service group" on page 383.	Single and multiple
Clearing resources in a service group from Admin Wait state See "Clearing the resources in a service group from the Admin Wait state" on page 384.	Single and multiple
Editing attributes of service groups See "Editing attributes of service groups" on page 385.	Single
Converting local service groups to global service groups See "Converting local service groups to global service groups" on page 418.	Single
Converting global service groups to local service groups See "Converting global service groups to local service groups" on page 420.	Single
Running the high availability fire drill for the selected service groups See "Running the high availability fire drill" on page 426.	Single and multiple
Running the disaster recovery fire drill for the selected service groups See "Running the disaster recovery fire drill" on page 428.	Single and multiple
Deleting service groups See "Deleting service groups" on page 371.	Single and multiple

Resource operations

Table A-4 lists the operations that you can perform on resources, and indicates whether the operation can be performed on a single resource, or on multiple resources.

Table A-4 Resource operations

Operation	Supported for single or multiple resources?
Adding or modifying resources See "Adding or modifying resources" on page 411.	Single
Enabling resources See "Enabling resources" on page 400.	Single and multiple
Disabling resources See "Disabling resources" on page 401.	Single and multiple
Editing attributes of a resource See "Editing attributes of a resource" on page 408.	Single
Editing attributes of a resource type See "Editing attributes of a resource type" on page 409.	Single
Bringing resources online See "Bringing resources online" on page 405.	Single and multiple
Taking resources offline See "Taking resources offline" on page 405.	Single and multiple
Taking resources offline and propagating the state See "Taking a resource offline and propagating the state" on page 404.	Single and multiple
Probing the resources See "Probing resources" on page 403.	Single and multiple

Table A-4 Resource operations *(continued)*

Operation	Supported for single or multiple resources?
Clearing faults on a resource See "Clearing faults on resources" on page 402.	Single and multiple
Linking resources in a service group See "Linking resources in a service group" on page 410.	Single
Unlinking resources in a service group See "Unlinking resources in a service group" on page 410.	Single
Invoking a resource action See "Invoking a resource action" on page 406.	Single
Marking a resource as critical See "Marking a resource as critical" on page 412.	Single and multiple
Marking a resource as non critical See "Marking a resource as non critical" on page 413.	Single and multiple
Deleting resources See "Deleting resources" on page 401.	Single and multiple

Section 5

Monitoring Storage Foundation HA licenses in the data center

- Chapter 28. Managing licenses
- Chapter 29. Viewing deployment information

Chapter 28

Managing licenses

This chapter includes the following topics:

- About licenses
- About Veritas licensing and pricing
- About assigning price tiers to hosts
- About license deployment policies
- Assigning a price tier to a host automatically
- Assigning a price tier to a host manually
- Creating a license deployment policy
- Modifying a license deployment policy
- Deleting a license deployment policy

About licenses

Typically, if you want to check the status of your Storage Foundation licenses, you must examine each server individually. A data center may have many types of hosts running a variety of software which can be a time-consuming process. Using the Veritas InfoScale Operations Manager console, you can view the real-time status of the Storage Foundation product licenses deployed on the managed host.

You can view the total number of licenses deployed, the chargeable and non-chargeable licenses, as well as the expired licenses. You can view the total number of hosts and virtual hosts with deployed licenses, as well as hosts with licenses that need to be consolidated. You can also view the product-wise deployment distribution in an interactive graph.

You can run reports that provide details on demo licenses, deployments by operating system tier and server tier. You can also run reports for hosts without Storage Foundation High Availability licenses, and to get details about the deployment policies that are violated.

Licenses deployed on a managed host are discovered at an interval of 24 hours.

See "About Veritas licensing and pricing" on page 476.

See "About assigning price tiers to hosts" on page 485.

See "Viewing the overview of SFHA licenses in the data center" on page 493.

See "Viewing the deployment details" on page 496.

See "Viewing the deployment policy details in the data center" on page 497.

About Veritas licensing and pricing

The Veritas storage and server management high availability products which operate on UNIX (Solaris, HP-UX, AIX), Linux and Windows must be certified to work on servers. The product certification differs from platform to platform. Certification is how storage and server management high availability products are licensed by platform. The price tier information is available on the Veritas Services and Operations Readiness Tools (SORT) website. You can configure Management Server to update the price tier information automatically or you can do it manually.

Table 28-1 Types of Veritas price tiers

Tier Type	Platforms supported	Description
Server Price Tier	UNIX: Solaris, HP-UX, and IBM AIX	The Veritas storage and server management, high availability products on UNIX are generally licensed and priced per server. The pricing depends on the server price tier to which a server is assigned. Veritas has established 12 server tiers (Tier A through Tier N), with Tier A representing servers with less processing capability and Tier N representing the highest-performing servers. Customers must pay the per-server price for each server that runs storage and server management, high availability products. This pricing meter applies to all versions of Veritas storage and server management, high availability products on UNIX (Solaris, HP-UX, AIX). The available price tiers for the server tier are A, B, C, D, E, F, G, H, I, J, K, L, M, N Server tier is applicable for license versions prior to 7.x as well as for license versions 7.x.

Table 28-1 Types of Veritas price tiers *(continued)*

Tier Type	Platforms supported	Description
Processor Price Tier	All platforms	The Veritas storage and server management, high availability products can be licensed and priced per tiered processor. The price depends on the tier to which a processor is assigned. A processor can be assigned to three different processor tiers from Processor Tier 1 through Processor Tier 3 based on the performance characteristics of the processor, such as number of processing cores and the architecture. Customers must pay the per-tiered-processor price for every processor that runs storage and server management, high availability products. The available price tiers for the processor tier are Tier 1, Tier 2, and Tier 3. Processor Price Tier is applicable for license versions prior to 7.x.
Operation System Price Tier	Microsoft Windows	The Veritas storage and server automation, high availability products which operate on the Microsoft Windows operating system are priced per server. The price also depends on the edition of Windows that is installed; for example, standard, enterprise, or data center. The Storage Foundation and high availability products on Windows do not follow a per-processor pricing meter. The available price tiers for the operating system tier are Level 1, Level 2, and Level 3. Operation System Price Tier is applicable for license versions prior to 7.x.

Table 28-1 Types of Veritas price tiers *(continued)*

Tier Type	Platforms supported	Description
Symantec Performance Value Unit (SPVU)	All platforms	The following are considered for calculating SPVU for the Veritas storage and server management, high availability products on a host: ■ The operating system on the host ■ The make and the model of the processor ■ Number of cores per processor Symantec Performance Value Unit (SPVU) is applicable for license versions prior to 7.x.
Per-core licenses	All platforms	The value is calculated by accumulating the total number of cores on the host used by the product. Per-core licences are applicable only for license versions 7.x onwards.

See "About the Symantec Performance Value Unit" on page 479.

See "About licenses" on page 475.

See "Viewing the overview of SFHA licenses in the data center" on page 493.

About the Symantec Performance Value Unit

From the Storage Foundation and high availability 6.0 release, Veritas uses the Symantec Performance Value Unit (SPVU) pricing method for Storage Foundation and high availability products. The total number of SPVUs required for a specific product is calculated by multiplying the number of cores used by the product with the SPVU value per core.

The SPVU per core is calculated based on the following:

- The operating system on the host
- The make and the model of the processor
- The number of cores per processor

Note: Version 4.1 is the minimum version of the Veritas InfoScale Operations Manager managed host to be installed on the host to calculate the SPVU.

To calculate the SPVU for the products that run on IBM LPAR and kernel-based virtual machines, you need to install Veritas InfoScale Operations Manager managed host version 5.0 or later.

Based on these criteria, a host is assigned a specific SPVU value per core. The number of cores used by the physical and the virtual hosts is different. On the products that are deployed on a physical host, the SPVU is calculated for the actual cores used by the products. The number of cores used by all the products is the same on a physical host. For example, you have a physical host that has eight used cores, and you have installed Storage Foundation and Cluster Server on this host. To calculate the SPVU for this host, all the eight used cores are considered for both the products. Essentially, on a physical host, the number of used cores is the same for all the products that run on the same operating system.

For the virtualization servers, the SPVUs for the products are calculated differently for various types of virtualization servers. Veritas InfoScale Operations Manager can calculate the SPVUs required for the products that are installed on VMware, Solaris LDOM, the Solaris Zones virtualization servers, kernel-based virtual machines, and IBM LPAR.

The SPVU for Application HA is calculated only when it is deployed on Linux and Windows virtualization platforms.

See "About Veritas licensing and pricing" on page 476.

See "About the Symantec Performance Value Unit for VMware virtual machines" on page 480.

See "About the Symantec Performance Value Unit for Solaris LDOM virtualization server" on page 482.

See "About the Symantec Performance Value Unit for kernel-based virtual machines" on page 483.

See "About the Symantec Performance Value Unit for IBM LPAR" on page 484.

About the Symantec Performance Value Unit for VMware virtual machines

Veritas InfoScale Operations Manager cannot calculate the Symantec Performance Value Units (SPVU) for the products installed on the VMware virtual machines if it cannot recognize the mapping between the VMware ESX virtualization server and the virtual machines. To enable Veritas InfoScale Operations Manager to recognize

the mapping between the VMware ESX virtualization server and the virtual machines, you must perform the following tasks:

- Add the virtualization server to Veritas InfoScale Operations Manager using VMware vCenter by navigating to **Settings** > **Virtualization**.
- Add the virtual machines separately to Veritas InfoScale Operations Manager as managed hosts.

Note: If you add the virtual machines to Veritas InfoScale Operations Manager as managed hosts before you add the VMware ESX virtualization server, it can take up to 24 hours for Veritas InfoScale Operations Manager to calculate the SPVUs. Similarly, if you add only the virtual machines to Veritas InfoScale Operations Manager, and do not add the virtualization server, Veritas InfoScale Operations Manager cannot calculate the SPVUs for the products that are installed on the virtualization server. However, Veritas InfoScale Operations Manager discovers all the licenses for the products in this scenario.

When the SPVU is calculated for the products that run on the virtual machine of a VMware ESX server, the total number of cores on the server is compared with the total number of Virtual CPUs (vCPUs) that are allocated to the virtual machines that have the same guest operating system and the same product. The smaller number among these two is the total number of cores that need to be licensed (core-to-license) for the server. The SPVU is calculated by multiplying this number with the SPVU value per core.

For example, consider a VMware ESX Server that has 3 virtual machines and a total number of 8 cores. Two of the virtual machines use 6 vCPUs each and run Storage Foundation for Windows. The total number of the vCPUs on the virtual machine is calculated by adding 6 to 6, which is 12. To calculate the core-to-license for the server, 12 is compared with 8 (the number of cores on the physical VMware ESX Server). Because 8 is the smaller number, the core to license for the products on this VMware ESX server is 8. To find the required SPVU for the products on this virtual machine, 8 is multiplied with the actual SPVU value per core.

See “About the Symantec Performance Value Unit” on page 479.

See “About the Symantec Performance Value Unit for Solaris LDOM virtualization server” on page 482.

See “About the Symantec Performance Value Unit for kernel-based virtual machines” on page 483.

See “About the Symantec Performance Value Unit for IBM LPAR” on page 484.

About the Symantec Performance Value Unit for Solaris LDOM virtualization server

To enable Veritas InfoScale Operations Manager to calculate the Symantec Performance Value Units (SPVU) for the products that are installed on an LDOM virtualization server, you must perform the following tasks

- Add the virtualization server to Veritas InfoScale Operations Manager as a managed host.
- Add the guest domains to Veritas InfoScale Operations Manager as managed hosts.

Note: If you add the domains to Veritas InfoScale Operations Manager as managed hosts before you add the LDOM virtualization server, it can take up to 24 hours for Veritas InfoScale Operations Manager to calculate the SPVUs. Similarly, if you add only the domains to Veritas InfoScale Operations Manager, and do not add the LDOM virtualization server, Veritas InfoScale Operations Manager cannot calculate the SPVUs for the products that are installed on the virtualization server. However, Veritas InfoScale Operations Manager discovers all the licenses on the products in this scenario.

To calculate the SPVU for the products on a Solaris LDOM virtualization server, Veritas InfoScale Operations Manager considers the total number of threads that are allocated to the logical domains (the control and the guest domains) where the same product is installed. For example, consider a Solaris LDOM server with 4 cores, with 8 threads in each core. This server has a control domain with 10 threads and two guest domains with 6 and 4 threads each. On all the three domains, Storage Foundation is installed. The total number of threads is 20. The total number of core-to-license is calculated by dividing 20 (total threads) by 8 (thread per core), which is rounded to 3. The SPVU is calculated by multiplying 3 with the SPVU value per core.

The SPVUs for the products on Solaris Zones are calculated based on the cores on the global zones. The cores on the local zones are not considered for calculating SPVUs for the products on Solaris Zones.

See “About the Symantec Performance Value Unit” on page 479.

See “About the Symantec Performance Value Unit for VMware virtual machines” on page 480.

See “About the Symantec Performance Value Unit for kernel-based virtual machines” on page 483.

See “About the Symantec Performance Value Unit for IBM LPAR” on page 484.

About the Symantec Performance Value Unit for kernel-based virtual machines

To calculate the Symantec Performance Value Unit (SPVU) for the products that run on kernel-based virtual machines (KVM), you need to install Veritas InfoScale Operations Manager managed host version 5.0.

When the SPVU is calculated for the products that run on the virtual machine, the total number of cores on the server is compared with the total number of Virtual CPUs (vCPUs) that are allocated to the virtual machines that have the same guest operating system and the same product. The smaller number among these two is the total number of cores that need to be licensed (core-to-license) for the server. The SPVU is calculated by multiplying this number with the SPVU value per core.

For example, consider a virtual server that has 3 virtual machines and a total number of 8 cores. Two of the virtual machines use 6 vCPUs each and run Storage Foundation for Windows. The total number of vCPUs on the virtual machine is calculated by adding 6 to 6, which is 12. To calculate the core-to-license for the virtual server, 12 is compared with 8 (the number of physical cores on the server). Because 8 is the smaller number, the core to license for the products on this server is 8. To find the required SPVU for the products on this virtual machine, 8 is multiplied with the actual SPVU value per core.

When the SPVU is calculated for the products that run in the KVM cluster environment, the total number of vCPUs of all the virtual machines in the cluster is compared with the total number of cores on the server. The smaller number among these two is the total number of cores that need to be licensed (core-to-license) for the server. The SPVU is calculated by multiplying this number with the SPVU value per core.

For example in a KVM cluster environment, the total number of vCPUs on all the virtual machines in the cluster is 24 and the total number of cores is 35. To calculate the core-to-license, 24 is compared with 35. Because 24 is the smaller number, the core-to-license for the products on this server is 24. To find the required SPVU for the products in this cluster environment, 24 is multiplied with the actual SPVU value per core.

For the products that run on KVM, and are deployed on the server as well as the virtual machines, the total number of cores on the server is used to calculate the SPVU.

See “About the Symantec Performance Value Unit” on page 479.

See “About the Symantec Performance Value Unit for VMware virtual machines” on page 480.

See “About the Symantec Performance Value Unit for Solaris LDOM virtualization server” on page 482.

See "About the Symantec Performance Value Unit for IBM LPAR" on page 484.

About the Symantec Performance Value Unit for IBM LPAR

To calculate the Symantec Performance Value Unit (SPVU) for the products that run on IBM LPAR, you need to install Veritas InfoScale Operations Manager managed host version 5.0.

When the SPVU is calculated for the products that run on IBM LPAR, the sum of Virtual CPUs (vCPUs) is compared with the number of cores on the physical LPAR server. The smaller number among these two is the total number of cores that need to be licensed (core-to-license) for the server. The SPVU is calculated by multiplying this number with the SPVU value per core.

Note: The vCPUs are allocated cores in case of dedicated LPAR. They are entitled capacity in case of capped LPAR and reserved capacity in case of uncapped LPAR.

When the SPVU is calculated for the products that run on IBM LPAR in a cluster environment, the sum of all vCPUs of all the virtual machines in the cluster is compared with the sum of physical cores on the LPAR server in the cluster. The smaller number among these two is used to calculate the license.

See "About the Symantec Performance Value Unit" on page 479.

See "About the Symantec Performance Value Unit for VMware virtual machines" on page 480.

See "About the Symantec Performance Value Unit for Solaris LDOM virtualization server" on page 482.

See "About the Symantec Performance Value Unit for kernel-based virtual machines" on page 483.

About the per-core licensing

From the Storage Foundation High Availability 7.0 release, Veritas uses the per-core licensing pricing method for Storage Foundation products. The value is calculated by accumulating the total number of cores on the host used by the product.

The calculation for number of cores used by the physical hosts and the virtual hosts is different. On the products that are deployed on a physical host, the cores-to-license value is calculated for the actual cores used by the products. The number of cores used by all the products is the same on a physical host. For example, you have a physical host that has eight used cores, and you have installed Storage Foundation and Cluster Server on this host. To calculate the per-core value for this host, all the eight used cores are considered for both the products.

Essentially, on a physical host, the number of used cores is the same for all the products that run on the same operating system.

For a virtualization server, the cores-to-license value for the products is calculated as minimum of the following two:

- Virtual CPUs assigned to the virtual machines associated with that particular virtualization server.
- Total cores present on the virtualization server.

See "About Veritas licensing and pricing" on page 476.

See "About the Symantec Performance Value Unit for VMware virtual machines" on page 480.

See "About the Symantec Performance Value Unit for Solaris LDOM virtualization server" on page 482.

See "About the Symantec Performance Value Unit for kernel-based virtual machines" on page 483.

See "About the Symantec Performance Value Unit for IBM LPAR" on page 484.

About assigning price tiers to hosts

You can use operating system-specific commands to find host characteristics. This includes the make and model of the host, processor type, and processor count. However, although you can discover hardware information for most hosts, you may not have all the characteristics of a host. In that case, it is called an "unknown tier".

The assign price tier feature lets you assign price tiers to an unknown host. It eliminates the need to find host characteristics manually.

You can assign a price tier to a host by selecting the server price tier, processors price tier, operating system price tier, or the Symantec Performance Value Unit (SPVU) price tier.

See "Assigning a price tier to a host manually" on page 486.

About license deployment policies

Veritas InfoScale Operations Manager lets you create license deployment policies to manage the license deployment in the data center.

You can define a risk threshold value and a fault threshold value for a specific license. If the number of licenses that you have deployed in the data center exceeds the risk threshold value, Veritas InfoScale Operations Manager generates a warning. The alert is also generated if the number of licenses that are deployed exceeds the

fault threshold value. You can view alerts in the Veritas InfoScale Operations Manager console.

See "Creating a license deployment policy" on page 488.

Assigning a price tier to a host automatically

You can assign price tiers to a single host or multiple hosts automatically in Veritas InfoScale Operations Manager using the uploaded price tier information.

To perform this task, your user group must be assigned the Admin role on the Server perspective.

To assign a price tier to a host automatically

1. In the Home page on the Management Server console, click **Licensing**.
2. Click the **Deployment details** tab.
3. In the deployment details list, right-click the host, and select **Assign price tier > Automatically**.
4. In the **Assign price tier automatically** wizard panel, verify the host that you have selected. Click **OK**.
5. In the **Assign price tier automatically - Result** panel, click **OK**.

See "About assigning price tiers to hosts" on page 485.

Assigning a price tier to a host manually

You can manually assign a price tier to a host by selecting the server, processor type, operating system, or the Symantec Price Value Unit (SPVU) price tier.

To perform this task, your user group must be assigned the Admin role on the Server perspective.

To assign a price tier to a host manually

1. In the Home page on the Management Server console, click **SFHA Licensing**.
2. Click the **Deployment details** tab.
3. In the deployment details list, right-click the host, and select **Assign price tier > Manually**.
4. In the **Assign price tier manually - Select tier values** panel, specify the required information, and click **Next**.

   See "Select tier values panel options" on page 487.

5 In the **Assign price tier manually - Select hosts to apply same tier values** panel, select the hosts and click **Finish**.

See "Select hosts to apply same tier values panel options" on page 488.

6 In the **Assign price tier manually - Result** panel, click **OK**.

See "About assigning price tiers to hosts" on page 485.

Select tier values panel options

Use this wizard panel to manually assign the latest price tier to the host on which the licenses are deployed.

You can view the host name, platform, CPU model, and server model details for the selected host.

Note: All the information may not apply to a host that you have selected. You can specify the values from the drop-down boxes for the information that is applicable for the host.

Table 28-2 Select tier values panel options

Field	Description
Server tier	Select the appropriate server price tier. The server price tiers are represented using alphabets A through N. The server price tier applies to the Solaris, the HP-UX, and the AIX platforms.
Processor tier	Select the appropriate processor type price tier. The processor price tiers are represented as Tier 1, Tier 2, Tier 3, and Tier 4. The processor price tier applies to all the supported platforms.
OS tier	Select the appropriate operating system price tier. The operating system price tiers are represented as Level 1, Level 2, and Level 3. The operating system price tier applies to the Windows platforms.
SPVU	Enter the value for the Symantec Performance Value Unit (SPVU) per processor core of the host.
Show hosts having similar configuration	This check box is displayed only if the selected host has similar hosts based on make, model, processor, or operating system in the data center.

See “Assigning a price tier to a host manually” on page 486.

See “Select hosts to apply same tier values panel options” on page 488.

Select hosts to apply same tier values panel options

Use this wizard panel to select similar hosts based on make, model, processor, or operating system.

You can view details such as the host name, platform, server tier, processor tier, operating system tier, and the SPVU value for each host.

Select the hosts to which you want to apply same tier values.

See “Assigning a price tier to a host manually” on page 486.

See “Select tier values panel options” on page 487.

Creating a license deployment policy

Using the Management Server console, you can create a license deployment policy that lets you receive various types of alerts on the deployment of licenses in the data center.

To perform this task, your user group must be assigned the Admin role on the Server perspective.

To create a license deployment policy

1. In the Home page on the Management Server console, click **Licensing**.
2. Click the **Deployment policy** tab.
3. Click **Create policy**.
4. In the **Create policy - Details** wizard panel, enter the details, and click **Finish**.
5. In the **Create policy - Result** panel, click **OK**.

See “Modifying a license deployment policy” on page 490.

See “Deleting a license deployment policy” on page 491.

Create policy - Details panel options

Use this wizard panel to create a license deployment policy.

Table 28-3 Create policy - Details panel options

Field	Description
Policy name	Specify the name of the deployment policy. You can use the name to identify the deployment policy later. For example, the name of the policy helps you identify the policy in an alert message.
Platform	Select the platform based on which the licenses are deployed.
Tier type	Select one of the following price tier types based on the platform on which the licenses are deployed: ■ SPVU ■ Server tier ■ OS tier ■ Per Core ■ Processor tier Following is the list of platforms and their tier types: ■ Solaris, HP-UX, and AIX - ■ For license versions prior to 7.x - SPVU, Server tier, and Processor price tier ■ For license versions 7.x onwards - Server tier and Per -core licensing ■ Linux - ■ For license versions prior to 7.x - SPVU and Processor price tiers ■ For license versions 7.x onwards - Per -core licensing ■ Windows - ■ For license versions prior to 7.x - SPVU, Processor tier, and OS price tier ■ For license versions 7.x onwards - Per -core licensing
Tier value	Select from the following price tiers based on the platform that you have chosen. ■ **Processor** tier - Tier 1, Tier 2, Tier 3, and Tier 4. ■ **Server** tier - Alphabets A-N. ■ **OS** tier - Level 1, Level 2, and Level 3. **Note:** This field does not apply to the SPVU tier type.
Product name	Select the Storage Foundation product for which you want to create a deployment policy.

Table 28-3 Create policy - Details panel options *(continued)*

Field	Description
Product edition	Select the edition of the product for which you want to create the deployment policy.
Product version	Select the version of the product for which you want to create the deployment policy.
Risk threshold	Enter the number of license deployments for which Veritas InfoScale Operations Manager must generate a warning. If you have selected SPVU as the tier type, you must specify the number of SPVUs in this field.
Fault threshold	Enter the number of license deployments for which Veritas InfoScale Operations Manager must generate an alert. If you have selected SPVU as the tier type, you must specify the number of SPVUs in this field. If you have selected Per-core licensing as the tier type, you need to specify the number of cores-to-licence in this field.

See "Creating a license deployment policy" on page 488.

Modifying a license deployment policy

You can only modify the risk and fault threshold values for a license deployment policy that you have already created.

To perform this task, your user group must be assigned the Admin role on the Server perspective.

To modify a license deployment policy

1. In the Home page on the Management Server console, select **Licensing**.
2. Click the **Deployment policy** tab.
3. Right-click the policy and select **Edit thresholds**.
4. In the **Edit thresholds - Details** wizard panel, modify the values in the **Risk threshold** and **Fault threshold** fields, and click **Finish**.

   See "Edit thresholds - Details panel options" on page 491.
5. In the **Edit thresholds - Result** panel, click **OK**.

See "Creating a license deployment policy" on page 488.

See "Deleting a license deployment policy" on page 491.

Edit thresholds - Details panel options

Use this wizard panel to modify the **Risk threshold** and the **Fault threshold** values that you have defined for a license deployment policy.

If the policy that you want to edit has no name, you can provide a name to the policy in this panel.

Table 28-4 Edit thresholds - Details panel options

Field	Description
Policy name	Displays the name of the policy that you have provided. If the policy that you want to edit has no name, you can provide a name to the policy in this field.
Platform	Displays the platform that you have selected when you created the policy.
Tier type	Displays the tier type that you have selected when you created the policy.
Tier value	Displays the tier value that you have selected when you created the policy.
Product name	Displays the Storage Foundation product that you have selected when you created the policy.
Product edition	Displays the edition of the Storage Foundation product that you have selected when you created the policy.
Product version	Displays the version of the Storage Foundation product that you have selected when you created the policy.
Risk threshold	Modify the number of license deployments for which Veritas InfoScale Operations Manager must generate a warning.
Fault threshold	Modify the number of license deployments for which Veritas InfoScale Operations Manager must generate an alert. You must verify the license deployments in your data center.

See “Modifying a license deployment policy” on page 490.

Deleting a license deployment policy

Using the Management Server console, you can delete a license deployment policy that you no longer require.

To perform this task, your user group must be assigned the Admin role on the Server perspective.

To delete a license deployment policy

1. In the Home page on the Management Server console, click **Licensing**.
2. Click the **Deployment policy** tab.
3. Right-click the policy to display the shortcut menu, click **Delete policy**.
4. In the **Delete policy** wizard panel, click **OK**.
5. In the **Delete policy - Result** panel, click **OK**.

See “Creating a license deployment policy” on page 488.

See “Modifying a license deployment policy” on page 490.

Chapter 29

Viewing deployment information

This chapter includes the following topics:

- Viewing the overview of SFHA licenses in the data center
- Viewing the deployment details
- Viewing the deployment policy details in the data center
- Viewing the VOM Deployment Report

Viewing the overview of SFHA licenses in the data center

In the Management Server console, this view provides you the overview of the licenses deployed on your data center. You can view the following tables and chart:

- **Host deployment summary** table
  Table 29-1
- **License deployment summary** table
  Table 29-2
- **Product deployment summary** chart
  In this chart, you can view the version-wise deployments of products. Rest the mouse pointer on the chart to view the product versions and the number of deployments.

Table 29-1 Host deployment summary

Field	Description
Total hosts having licenses	Displays the number of hosts in the data center on which the product licenses are deployed.
Virtual hosts having licenses	Displays the number of virtual hosts in the data center on which the product licenses are deployed.
Hosts having improper licensing	Displays the number of hosts having deployed licenses which have overlapping features. These licenses can be consolidated. For example, if Volume Replicator Standard edition and Storage Foundation High Availability Enterprise edition is installed on a host. The license for Volume Replicator can be consolidated with Storage Foundation High Availability with Volume Replicator license. Also displays the hosts that have multiple licenses of the same product and the hosts that have the Distributed Site Management (DSM) feature enabled without having the license for Storage Foundation with HA/DR.

Table 29-2 License deployment summary

Field	Description
Total licenses	Displays the total number of licenses deployed in the data center.
Chargeable licenses	Displays the total number of deployed licenses that are chargeable. See "About chargeable deployed licenses" on page 495.
Non-chargeable licenses	Displays the total number of deployed licenses that are non-chargeable.
Expired licenses	Displays the number of expired licenses in the data center.

Table 29-2 License deployment summary *(continued)*

Field	Description
Violated policies	Displays the number of deployment policies that are violated.

You can view this information related to the hosts, if your user group has at least Guest role assigned on the Server perspective.

To view the overview of SFHA licenses in the data center

- In the Home page on the Management Server console, click **SFHA Licensing**.

See "About licenses" on page 475.

See "Viewing the deployment details" on page 496.

See "Viewing the deployment policy details in the data center" on page 497.

About chargeable deployed licenses

A deployed license is chargeable based on factors such as the type of the license, the edition, the version, and the feature set of the license.

When a single license is deployed on a host, that license is always marked as chargeable.

Consider the following situations in which a host has multiple redundant licenses for the same product, but either the version or the edition is different:

- If a host has multiple licenses for the same product with different editions, but same versions, the license with the highest edition is marked as chargeable. For example, if Storage Foundation version 5.0 Standard edition and Storage Foundation version 5.0 Enterprise edition are deployed on the host, then Storage Foundation Enterprise edition is chargeable.
- If a host has multiple licenses for the same product with different versions, but same editions, the version of the corresponding package is considered to mark the license as chargeable. For example, if you have deployed Storage Foundation 5.0 Standard and Storage Foundation 5.1 Standard on one of your hosts, Storage Foundation 5.0 Standard license is chargeable if the VRTSvxvm 5.0 package is installed on the host.
- If a host has multiple licenses for the same product with different versions and different editions, then the license with the highest edition is chargeable. For example, if Storage Foundation version 4.1 Enterprise and Storage Foundation version 5.0 Standard are deployed, then Storage Foundation 4.1 Enterprise is chargeable.

When a host has multiple products deployed, with different versions and different editions, then the license for which product evidence is found is marked as chargeable. Else the product with the highest license is chargeable.

The overlapping of the feature set of the licenses also determines license accountability. For example, if Volume Replicator Standard edition and Storage Foundation High Availability Enterprise edition are installed on a host, then both the licenses are marked as chargeable, and the host is marked as "Host having improper licensing".

All point product licenses, for example ApplicationHA licenses or Database licenses are always chargeable.

See "Viewing the overview of SFHA licenses in the data center" on page 493.

Viewing the deployment details

In the Management Server console, this view provides you an overview of the licenses that are installed on a host in your data center. You can view the edition, version, SPVU, and check if the license is chargeable. You can view the details of license such as the key number, serial number, expiry date, the server price tier, processor price tier, and operating system tier information

You can perform the following tasks in this view:

- Assign price tier automatically.
- Assign price tier manually.
- View SPVU.
- View child licenses.
  This view displays the list of child licenses that are installed through the parent product license.
- View enabled features.
  This view displays the list of products and features that are enabled by the parent product license.

You can view this information related to the hosts, if your user group has at least Guest role assigned on the Server perspective.

To view the deployment details in the data center

1. In the Home page on the Management Server console, click **SFHA Licensing**.
2. Click the **Deployment Details** tab.

See "About chargeable deployed licenses" on page 495.

See "About assigning price tiers to hosts" on page 485.

See “Assigning a price tier to a host manually” on page 486.

See “About the Symantec Performance Value Unit” on page 479.

Viewing the deployment policy details in the data center

In the Management Server console, this view provides you an overview of the deployed policies in your data center. You can view the details of the policies such as the name of the policy, the product name, edition, version, tier type, tier, risk and fault thresholds, and the total number of license deployments.

You can perform the following tasks in this view:

- Create a policy.
- Edit the risk and fault thresholds of a policy.
- Delete a policy.

A deployment policy for which the licenses have exceeded the fault threshold value is highlighted in red. A deployment policy for which the licenses have exceeded the risk threshold value is highlighted in yellow.

You can enter the policy name or product name in the **Search** text box at the top to filter the **License deployment policies** table.

You can also filter the table based on the following:

- Policies that are at risk
- Policies that are faulted

You can view this information related to the hosts, if your user group has at least Guest role assigned on the Server perspective.

To view the deployment policy details

1. In the Home page on the Management Server console, click **SFHA Licensing**.
2. Click the **Deployment Policy** tab.
3. In the **Deployment Policy** tab, review the details of the license deployment policies.

See “Creating a license deployment policy” on page 488.

See “Modifying a license deployment policy” on page 490.

See “Deleting a license deployment policy” on page 491.

See “About license deployment policies” on page 485.

See "Viewing the overview of SFHA licenses in the data center" on page 493.

See "Viewing the deployment details" on page 496.

Viewing the VOM Deployment Report

This report provides information on hosts that are connected to Management Server and whether they have Veritas products manageable from the Management Server console.

The overview section displays the following information:

- The total number of hosts that report to Management Server.
- The number of hosts that run Storage Foundation, Cluster Server, Application HA, or Dynamic Multipathing.
- The number of hosts that have Storage Foundation Enterprise licenses.
- The number of physical switches discovered.
- The number of physical Fibre Channel ports discovered.
- The number of enclosures exporting NAS shares.
- The number of hosts that do not run Storage Foundation, Cluster Server, AppHA, or Dynamic Multipathing.
- The number of hosts that have Storage Foundation Standard or Storage Foundation Basic licenses.
- The number of virtual switches discovered.
- The number of virtual Fibre Channel ports discovered.

In the table you can view the name of the host, operating system on which the host runs, and the type of the host. You can view the configuration type such as agentless host, agent host, or a host with no configuration as well as the following details:

- Whether Storage Foundation, Cluster Server, Dynamic Multipathing, or Application HA is installed on the host.
- The edition of the Storage Foundation, Cluster Server, Dynamic Multipathing, or Application HA installed on the host. The value can be Enterprise, Standard, or Basic.

You can perform the following tasks in this view:

- Subscribe for the report.
- Save the report as a CSV file.
- Email the report.

You can view this information related to the hosts, if your user group has at least Guest role assigned on the Server perspective.

To view the VOM Deployment Report

1. In the Home page on the Management Server console, click **SFHA Licensing**.
2. Click the **Reports** tab.
3. Click **VOM Deployment Report**.

See "About reports" on page 118.

See "Subscribing for a report" on page 123.

See "Saving a report" on page 122.

See "Sending a report through email" on page 125.

Chapter 30

Monitoring performance

This chapter includes the following topics:

- About performance metering statistics
- About metered resources
- About space estimation for data logs
- Enable performance metering for a host
- Disable performance metering for a host
- Enable performance metering for a virtualization server
- Disable performance metering for a virtualization server
- About Veritas InfoScale Operations Manager performance graphs
- Viewing the performance graphs for a host
- Viewing the performance graphs for a disk
- Viewing the performance graphs for volume and file system
- Viewing the performance graphs for a path
- Viewing the performance graphs for an initiator
- Viewing the performance graphs for virtualization server and virtual machines
- Viewing the performance graphs for a path of a virtualization server
- Viewing the performance graphs for an enclosure
- About threshold settings
- Adding threshold settings for an object

- Deleting the threshold settings for an object
- Enabling the threshold settings for an object
- Disabling the threshold settings for an object

About performance metering statistics

Historical performance data of various resources is collected in a fixed-size binary file. The older data is overwritten as new data arrives in a circular round robin array. The number of metrics, frequency of data insertion, number of objects, and the roll-up databases affect the size of binary file. The higher resolution data is compressed to a lower resolution data. For example, data is collected every 5-minutes for the last 24 hours for daily performance analysis. For monthly performance analysis, average of the 5-minutes data for every two hours is used. For the yearly performance log, average of the 5-minutes data for every 24 hours is used.

See "About metered resources" on page 501.

See "About space estimation for data logs" on page 505.

About metered resources

Table 30-1 lists the resources that are metered and the difference in the log configurations of previous versions of Veritas InfoScale Operations Manager and Veritas InfoScale Operations Manager 5.0 and onwards. For example, to understand the percentage of CPU utilization of a system, the metering is done every 5 minutes for 24 hours in Veritas InfoScale Operations Manager 5.0. Whereas the metering in the previous versions of Veritas InfoScale Operations Manager was 5 minutes for one week.

Hosts metering on Windows platform is available in Veritas InfoScale Operations Manager version 6.0 and onwards. Table 30-2 lists the metered resources and the log configurations on Windows platform.

The data that is collected during metering is used to generate the performance charts of the resource. In Veritas InfoScale Operations Manager version 6.1 and onwards, this data is also used to evaluate the threshold settings defined on the resources.

Control Host Add-on version 6.1and onwards is required to enable performance metering for the VMware ESX server and Virtual Machine. Storage Insight Add-on version 6.1 and onwards is required to enable performance metering for storage

array port, adapter, and enclosure. To view the performance graphs for volume and disk, Storage Foundation HA products must be installed on the host.

Table 30-1 Log configurations for UNIX hosts

Resource	Chart name	Log configuration in the previous versions	Log configuration in version 5.0 and onwards
Host	Available Memory (in KB) Average CPU Load	5 minutes / 1 week	5 minutes / 1 day 2 hours / 1 month 1 day / 1 year
Host	CPU Utilization (in percentage) Swap in Rate (in KB/sec) Used Swap (in KB)	Not available	5 minutes / 1 day 2 hours / 1 month 1 day / 1 year
File system	Size/Used (in KB)	6 hours / 1 month 1 day / 1 year	6 hours / 1 month 1 day / 1 year
Volume	Average Read/Write Latency Bytes Read/Written (in bytes)	1 minute / 6 hours 5 minutes / 1 day 2 hours / 1 month 1 day / 1 year	1 minute / 6 hours 5 minutes / 1 day 2 hours / 1 month 1 day / 1 year
Disk	Average Read/Write Latency Bytes Read/Written (in bytes)	1 minutes / 6 hours 5 minutes / 1 day 2 hours / 1 month 1 day / 1 year	1 minutes / 6 hours 5 minutes / 1 day 2 hours / 1 month 1 day / 1 year
Host Initiator	Average Read/Write Latency Bytes Read/Written (in bytes) Read/Write Errors Read/Write Queue Lengths	1 minute / 6 hours 5 minutes / 1 day 2 hours / 1 month 1 day / 1 year	5 minutes / 1 day 2 hours / 1 month 1 day / 1 year

Table 30-1 Log configurations for UNIX hosts *(continued)*

Resource	Chart name	Log configuration in the previous versions	Log configuration in version 5.0 and onwards
Host Enclosure	Average Read/Write Latency for Host Average Bytes Read/Written for Host (in bytes)	Not available	5 minutes / 1 day 2 hours / 1 month 1 day / 1 year
Host Path	Average Read/Write Latency Bytes Read/ Written(in bytes)	Not available	5 minutes / 1 day 2 hours / 1 month
VMware ESX server and Virtual Machine	Available Memory (in KB) CPU Utilization (in percentage) Swap in Rate (in KB/sec) Used Swap (in KB)	Not available	5 minutes / 1 day 2 hours / 1 month 1 day / 1 year
VMware ESX server Initiator	Average Read/Write Latency Bytes Read/Written (in bytes)	Not available	5 minutes / 1 day 2 hours / 1 month 1 day / 1 year
VMware ESX server Enclosure	Average Read/Write Latency for Host Bytes Read/Written (in bytes)	Not available	5 minutes / 1 day 2 hours / 1 month 1 day / 1 year
VMware ESX server Path	Average Read/Write Latency Bytes Read/ Written(in bytes)	Not available	5 minutes / 1 day 2 hours / 1 month 1 day / 1 year
Storage array - Port	IO Operations per second IO Throughput per second	Not available	30 minutes / 1 day 2 hours / 1 month 1 day / 1 year

Table 30-1 Log configurations for UNIX hosts *(continued)*

Resource	Chart name	Log configuration in the previous versions	Log configuration in version 5.0 and onwards
Storage array - Adapter	IO Operations per second IO Throughput per second	Not available	30 minutes / 1 day 2 hours / 1 month 1 day / 1 year
Storage array - Enclosure	Average Read/Write Latency for Host Average Bytes Read/Written for Host (in bytes) IO Operations per second	Not available	30 minutes / 1 day 2 hours / 1 month 1 day / 1 year

Table 30-2 Log configurations for Windows hosts

Resource	Chart name	Log configuration in version 6.0 and onwards
Host	Available Memory (in KB) Average CPU Load CPU Utilization (in percentage) Swap in Rate (in KB/sec) Used Swap (in KB)	5 minutes / 1 day 2 hours / 1 month 1 day / 1 year
File system	Average Read/Write Latency Bytes Read/Write (in bytes)	5 minutes / 1 day 2 hours / 1 month 1 day / 1 year
Volume	Average Read/Write Latency Bytes Read/Written (in bytes)	5 minutes / 1 day 2 hours / 1 month 1 day / 1 year
Disk	Average Read/Write Latency Bytes Read/Written (in bytes)	5 minutes / 1 day 2 hours / 1 month 1 day / 1 year

See “About performance metering statistics” on page 501.

See “About space estimation for data logs” on page 505.

See “About threshold settings” on page 523.

About space estimation for data logs

Table 30-3 describes the space estimation for data logs for the various resources. For estimation purposes, the data in the Number of resources column is according to the standard environment. The metrics collected column represents the number of metrics collected for each resource. For example, in case of DMP paths, the total number of metrics collected is four: bytes read, bytes written, read average, and write average.

Data logs for host, volume, disk, file system, path, and initiator are stored on the managed host. The data logs for virtualization server, virtual machine, path, and initiator are stored on the Control Host. For storage array (port, adapter, and enclosure), data log for 1 day is stored on the discovery host, where as all the other logs are stored on Management Server.

Note: If Veritas InfoScale Operations Manager is configured in high availability environment, storage array port, adapter, and enclosure logs are saved on a shared disk. VMware ESX server and virtual machines logs are also saved on a shared disk.

Table 30-4 lists the space estimation for data logs for host, file system, volume, and disk on Windows platform.

Table 30-3 Space estimation for data logs

Name of resource	Number of resources	Number of metrics collected	Interval of collection	Duration of collection	Size in KB	Size in KB for a single object
Host, VMware ESX server, and Virtual Machine	1	5	5 minutes	1 day	24	24
	1	5	2 hours	1 month	29	29
	1	5	1 day	1 year	30	30
Multipathing paths	1000	4	5 minutes	1 day	18967	19
	1000	4	2 hours	1 month	23477	24

Table 30-3 Space estimation for data logs *(continued)*

Name of resource	Number of resources	Number of metrics collected	Interval of collection	Duration of collection	Size in KB	Size in KB for a single object
Initiator	4	9	5 minutes	1 day	171	43
	4	18	2 hours	1 month	423	106
	4	18	1 day	1 year	428	107
Enclosure	4	4	5 minutes	1 day	76	19
	4	8	2 hours	1 month	8	2
	4	8	1 day	1 year	190	46
File system	100	3	5 minutes	1 day	1423	14
	100	3	1 day	1 year	1784	18
Volume	100	4	1 minute	6 hours	2348	23
	100	4	5 minutes	1 day	1898	19
	100	4	2 hours	1 month	2348	23
	100	4	1 day	1 year	2379	24
Disk	100	4	1 minute	6 hours	2348	23
	100	4	5 minutes	1 day	1898	19
	100	4	2 hours	1 month	2347	23
	100	4	1 day	1 year	2379	23
Storage array - Array port	32	2	30 minutes	1 day	304	9
	32	4	2 hours	1 month	751	23
	32	4	1 day	1 year	761	24
Storage array - Adapter	8	2	30 minutes	1 day	76	9
	8	4	2 hours	1 month	188	23
	8	4	1 day	1 year	190	24

Table 30-3 Space estimation for data logs *(continued)*

Name of resource	Number of resources	Number of metrics collected	Interval of collection	Duration of collection	Size in KB	Size in KB for a single object
Storage array -Enclosure	1	1	30 minutes	1 day	5	5
	1	2	2 hours	1 month	12	12
	1	2	1 day	1 year	12	12

Table 30-4 Space estimation for data logs for Windows hosts

Name of resource	Number of resources	Metrics collected	Interval of collection	Duration of collection	Size in KB	Size in KB for a single object
Host	1	5	5 mins	1 day	24	24
	1	5	2 hours	1 month	29	29
	1	5	1 day	1 year	30	30
File system	100	4	5 minutes	1 day	1898	19
	100	4	2 hours	1 month	2348	23
	100	4	1 day	1 year	2379	24
Volume	100	4	5 minutes	1 day	1898	19
	100	4	2 hours	1 month	2348	23
	100	4	1 day	1 year	2379	24
Disk	100	4	5 minutes	1 day	1898	19
	100	4	2 hours	1 month	2347	23
	100	4	1 day	1 year	2379	23

See “About performance metering statistics” on page 501.

See “About metered resources” on page 501.

Enable performance metering for a host

Use this option to enable performance metering for a host.

You cannot enable performance metering if the configuration type of the host is agentless or no-configuration. You can view the configuration type of the host in the Management Server perspective.

To perform this task, your user group must be assigned the Admin role on the host or the Server perspective. The permissions on the host may be explicitly assigned or inherited from a parent Organization.

To enable performance metering for a host

1. In the Management Server console, go to the **Server** perspective and select **Manage** in the left pane.
2. Expand the Organization or **Uncategorized Hosts** to locate the host.
3. Right-click the host and select **Properties**.
4. Click the **Performance** tab.
5. Select **Enable performance metering** and click **OK**.

See “Disable performance metering for a host” on page 508.

See “About performance metering statistics” on page 501.

See “About metered resources” on page 501.

See “About space estimation for data logs” on page 505.

Disable performance metering for a host

Use this option to disable performance metering for a host.

To perform this task, your user group must be assigned the Admin role on the host or the Server perspective. The permission on the host may be explicitly assigned or inherited from a parent Organization.

To disable performance metering for a host

1. In the Management Server console, go to the **Server** perspective and select **Manage** in the left pane.
2. Expand the Organization or **Uncategorized Hosts** to locate the host.
3. Right-click the host and select **Properties**.
4. Click the **Performance** tab.
5. Clear the **Enable performance metering** check box, and click **OK**.

See “Enable performance metering for a host” on page 508.

See “About performance metering statistics” on page 501.

See “About metered resources” on page 501.

See “About space estimation for data logs” on page 505.

Enable performance metering for a virtualization server

Use this option to enable performance metering for VMware, LDOM, KVM, and Hyper-V virtualization servers.

To enable performance metering for VMware virtualization servers, Control Host Add-on version 6.1 or later is required.

To perform this task, your user group must be assigned the Admin role on the virtualization server or the Virtualization perspective. The permissions on the virtualization server may be explicitly assigned or inherited from a parent Organization.

To enable performance metering for a virtualization server

1. In the Management Server console, go to the **Virtualization** perspective and select **Manage** in the left pane.
2. Expand the Organization or **Uncategorized Servers** to locate the virtualization server.
3. Right-click the virtualization server and select **Properties**.
4. Click the **Performance** tab.
5. Select **Enable performance metering**, and click **OK**.

See “Disable performance metering for a virtualization server” on page 509.

See “About performance metering statistics” on page 501.

See “About metered resources” on page 501.

See “About space estimation for data logs” on page 505.

Disable performance metering for a virtualization server

Use this option to disable performance metering for a VMware, LDOM, KVM, and Hyper-V virtualization servers.

To perform this task, your user group must be assigned the Admin role on the virtualization server or the Virtualization perspective. The permission on the virtualization server may be explicitly assigned or inherited from a parent Organization.

To disable performance metering for a virtualization server

1. In the Management Server console, go to the **Virtualization** perspective, and select **Manage** in the left pane.
2. Expand the Organization or **Uncategorized Servers** to locate the virtualization server.
3. Right-click the virtualization server and select **Properties**.
4. Click the **Performance** tab.
5. Clear the **Enable performance metering** check box, and click **OK**.

See "Enable performance metering for a virtualization server" on page 509.

See "About performance metering statistics" on page 501.

See "About metered resources" on page 501.

See "About space estimation for data logs" on page 505.

About Veritas InfoScale Operations Manager performance graphs

You can view the performance of the following objects by using interactive graphs.

Table 30-5

Perspective	Objects
Server	Host Disk Volume and file system Path Initiators
Virtualization	Virtualization server and virtual machine Path

Table 30-5 *(continued)*

Perspective	Objects
Storage	Enclosure Array port Adapter

Install Control Host Add-on version 6.1 and above on a managed host to view the performance graphs for the above objects. Install Storage Insight Add-on version 6.1 and above to view the performance graphs for an array port and adapter in the Storage perspective.

You can select an object and view graphs for multiple performance parameters. These graphs are line charts. The lines that represent the performance parameters are rendered in different colors. The X-axis represents the time duration and the Y-axis represents the performance parameter. Linear trend line is displayed for every performance parameter. The linear trend line is plotted by fitting a line on the historical data. Linear trend is not available for live charts.

Performance graphs for volume and disk are available if they are managed by VxVM. Performance charts for path (physical host), initiator, and enclosure are available only if they are managed by physical hosts running DMP. Performance charts of a path for a virtualization server are available only if the path is managed by the virtualization servers running VxDMP.

You can perform the following actions on the performance graphs:

- Specify the duration for which you want to view the performance of the object that you have selected.
- Move the mouse pointer over the line chart to view the performance data. A tool tip is displayed corresponding to the position of the mouse pointer on the graph.
- Click the desired legend to temporarily hide or display the line graph.

See "Pre-requisite commands to view performance graphs for a resource" on page 512.

See "Viewing the performance graphs for a host" on page 513.

See "Viewing the performance graphs for a disk" on page 514.

See "Viewing the performance graphs for volume and file system" on page 515.

See "Viewing the performance graphs for a path" on page 517.

See "Viewing the performance graphs for an initiator" on page 518.

See "Viewing the performance graphs for virtualization server and virtual machines" on page 519.

See "Viewing the performance graphs for a path of a virtualization server" on page 521.

See "Viewing the performance graphs for an enclosure" on page 522.

Pre-requisite commands to view performance graphs for a resource

Table 30-6 lists the pre-requisite commands that are required to view performance graphs for a host.

Table 30-6 Pre-requisite commands for a host

Chart name	Pre-requisite commands
Available Memory	Solaris/Linux: `vmstat` AIX: `svmon` HP-UX: `swapinfo`
Swap in Rate	VMware Guest/ESX: `SWPR/S` Solaris: `sar` AIX: `vmstat` HP-UX: `vmstat` Linux: `sar`
Used Swap	Solaris/AIX: `swap` Linux: `cat /proc/swaps` HP-UX: `swapinfo` VMware: `vSphere APIs`
Size/Used	`vxlist`

To view the performance graphs for enclosure, multipathing path, and initiator for an ESX server running VxDMP, the user account that is used to configure the vCenter discovery must have the following privileges:

```
Host\CIM\CIM Interaction
```

To view the **Available Memory**, **CPU Utilization**, **Swap in Rate**, and the **Used Swap** performance graphs for VMware ESX servers and virtual machines, the user account that is used to configure the vCenter discovery must have the following privileges:

```
Host\Configuration\Change Settings
```

See "About Veritas InfoScale Operations Manager performance graphs" on page 510.

Viewing the performance graphs for a host

In the Management Server console, you can view the performance of a host in an interactive graph. You can review the performance of a host for different durations - 6 hours, 24 hours, 1 week, 1 month, 1 year, or based on live data.

Table 30-7 lists the performance graphs for a host.

Table 30-7 Performance graphs for a host

Performance graph name	Description
Available Memory	Displays the free memory in the selected host and the trend for the specified duration.
Average CPU Load	Displays the CPU queue length on the selected host and the trend for the specified duration. For Windows hosts, this graph displays the threads waiting for CPU on the selected host and the trend for the specified duration.
CPU Utilization	Displays the CPU utilization (in percentage) on the selected host and the trend for the specified duration.
Swap in Rate	Displays the swap in rate (in KB per second) for the selected host and the trend for the specified duration.
Used Swap	Displays the swap that is used by the selected host and the trend for the specified duration.

You can view these performance graphs for the hosts, for which your user group has at least Guest role explicitly assigned or inherited from a parent Organization. You can also view the graphs if your user group has at least Guest role assigned on the Server perspective.

To view the performance graphs for a host

1. In the Management Server console, go to the **Server** perspective and select **Manage** in the left pane.
2. Expand the Organization or **Uncategorized Hosts** to locate the host.
3. Select the host, and click the **Performance** tab. To change the duration, use the drop-down list.

See "About Veritas InfoScale Operations Manager performance graphs" on page 510.

See "Enable performance metering for a host" on page 508.

Viewing the performance graphs for a disk

In the Management Server console, you can view the performance of a disk in an interactive graph. You can review the performance of a disk for different durations - 6 hours, 24 hours, 1 week, 1 month, 1 year, or based on live data.

You can view the performance graphs only for a Storage Foundation disk and if the disk belongs to a disk group. Performance charts are not displayed for virtual machines having virtual initiator.

Table 30-8 lists the performance graphs for a disk.

Table 30-8 Performance graphs for a disk

Performance graph name	Description
Average Read/Write Latency	Displays the average read and write latency for the disk and the trend for the specified duration.
Bytes Read/Written	Displays the number of bytes read and written (in bytes) on the disk and the trend for the specified duration.

You can view these performance graphs for the hosts, for which your user group has at least Guest role explicitly assigned or inherited from a parent Organization. You can also view the graphs if your user group has at least Guest role assigned on the Server perspective.

To view the performance graphs for a disk associated with a host

1. In the Management Server console, go to the **Server** perspective and select **Manage** in the left pane.
2. Expand the Organization or **Uncategorized Hosts** to locate the host.
3. Expand the host and then expand **Disk Groups** to locate the disk group.
4. Select the disk group and click the **Disks** tab.
5. In the disks list, right-click on a disk and select **Performance**. To change the duration, use the drop-down list.
6. Click the ellipses to select another disk. The ellipses are displayed only if the disk is a shared disk.

To view the performance graphs for a disk associated with an application

1. In the Management Server console, go to the **Server** perspective and select **Manage** in the left pane.
2. Expand **Applications** and then expand **Databases**.

3 Select the database and click the **Disks** tab.

4 In the disks list, right-click on a disk and select **Performance**. To change the duration, use the drop-down list.

5 Click the ellipses to select another disk. The ellipses are displayed only if the disk is a shared disk.

See "About Veritas InfoScale Operations Manager performance graphs" on page 510.

See "Enable performance metering for a host" on page 508.

Viewing the performance graphs for volume and file system

In the Management Server console, you can view the performance of a volume and file system in an interactive graph. You can review the performance of a volume for different durations - 6 hours, 24 hours, 1 week, 1 month, 1 year, or based on live data. You can review the performance of a file system for a duration of one month or one year. In case of Windows hosts, you can review the performance of a file system for a duration of six hours, 24 hours, one week, one month, one year, or based on live data.

If a file system is mounted on a volume, then you can view the file system graphs in the volume performance graphs view. If there is no file system mounted on the volume, then only volume performance graphs are displayed.

You can view the performance graphs only for a Storage Foundation volume and if the volume belongs to a disk group.

Table 30-9 lists the performance graphs for a volume.

Table 30-9 Performance graphs for a volume

Performance graph name	Description
Average Read/Write Latency	Displays the average read and write latency for the volume and the trend for the specified duration.
Bytes Read/Written	Displays the number of bytes read and written (in bytes) on the volume and the trend for the specified duration.

Table 30-10 lists the performance graphs for a file system.

Table 30-10 Performance graphs for a file system

Performance graph name	Description
Size/Used	Displays the size and the used space of the file system and the trend for the specified duration. This graph is displayed only for UNIX hosts.
FileSystem - Average Read/Write Latency	Displays the average read and write latency for the file system and the trend for the specified duration. This graph is displayed only for native Windows file systems that is FAT, FAT32, NTFS, and ReFS.
FileSystem - Bytes Read/Write	Displays the number of bytes read and written on the file system and the trend for the specified duration. This graph is displayed only for native Windows file systems that is FAT, FAT32, NTFS, and ReFS.

You can view these performance graphs for the hosts, for which your user group has at least Guest role explicitly assigned or inherited from a parent Organization. You can also view the graphs if your user group has at least Guest role assigned on the Server perspective.

To view the performance graphs for volume and file system associated with a host

1. In the Management Server console, go to the **Server** perspective and select **Manage** in the left pane.
2. Expand the Organization or **Uncategorized Hosts** to locate the host.
3. Expand the host and then expand **Volumes** to locate the volume to view its performance graph.
4. Select the volume, and click the **Performance** tab. To change the duration, use the drop-down list.
5. Click the ellipses in the file system graph to select another file system. The ellipses are displayed only if the file system is a shared file system.

To view the performance graphs for volume and file system associated with an application

1. In the Management Server console, go to the **Server** perspective and select **Manage** in the left pane.
2. Expand **Applications** and then expand **Databases**.
3. Select the database and click the **Volumes** tab.

4 In the volumes list, right-click on a volume and select **Performance**. To change the duration, use the drop-down list.

5 Click the ellipses in the file system graph to select another file system. The ellipses are displayed only if the file system is a shared file system.

See "About Veritas InfoScale Operations Manager performance graphs" on page 510.

See "Enable performance metering for a host" on page 508.

Viewing the performance graphs for a path

In the Management Server console, you can view the performance of a path in an interactive graph. You can review the performance of a path for different durations - 6 hours, 24 hours, 1 week, 1 month, or based on live data.

You can view the performance graphs for a path only if the path is managed by DMP or VxDMP. Performance charts are not displayed for virtual machines having virtual initiator.

Table 30-11 lists the performance graphs for a path.

Table 30-11 Performance graphs for a path

Performance graph name	Description
Average Read/Write Latency	Displays the average read and write latency of the selected path and the trend for the specified duration.
Bytes Read/Written	Displays the number of bytes read and written (in bytes) in the selected path and the trend for the specified duration.

You can view these performance graphs for the hosts, for which your user group has at least Guest role explicitly assigned or inherited from a parent Organization. You can also view the graphs if your user group has at least Guest role assigned on the Server perspective.

To view the performance graph for a path associated with a host

1 In the Management Server console, go to the **Server** perspective and select **Manage** in the left pane.

2 Expand the Organization or **Uncategorized Hosts** to locate the host.

3 Expand the host and then expand the **Disk Groups**.

4 Click on a disk.

5 In the **Paths** tab, right-click the path and select **Performance**. To change the duration, use the drop-down list.

To view the performance graph for a path associated with an application

1 In the Management Server console, go to the **Server** perspective and select **Manage** in the left pane.

2 Expand **Applications** and then expand **Databases**.

3 Select the database and click the **Disks** tab.

4 Click on a disk

5 In the **Paths** tab, right-click the path and select **Performance**. To change the duration, use the drop-down list.

See "About Veritas InfoScale Operations Manager performance graphs" on page 510.

Viewing the performance graphs for an initiator

In the Management Server console, you can view the performance of an initiator in an interactive graph. You can review the performance of an initiator for different durations - 6 hours, 24 hours, 1 week, 1 month, or based on live data.

You can view the performance graphs for an initiator only if the initiator is managed by DMP or VxDMP. Performance charts are not displayed for virtual machines having virtual initiator.

Table 30-12 lists the performance graphs for an initiator.

Table 30-12 Performance graphs for an initiator

Performance graph name	Description
Average Read/Write Latency	Displays the average read and write latency for the initiator and the trend for the specified duration.
Bytes Read/Written	Displays the number of bytes read and written (in bytes) on the initiator and the trend for the specified duration.
Read/Write Errors	Displays the number of read and write errors on the initiator and the trend for the specified duration.

Table 30-12 Performance graphs for an initiator *(continued)*

Performance graph name	Description
Read/Write Queue Lengths	Displays the length of read and write queue on the initiator and the trend for the specified duration. **Note:** Read/Write Errors and Read/Write Queue Lengths graphs for an initiator are displayed only for Storage Foundation version 5.1 and above. These charts are available only for non-virtualized hosts.

You can view these performance graphs for the hosts, for which your user group has at least Guest role explicitly assigned or inherited from a parent Organization. You can also view the graphs if your user group has at least Guest role assigned on the Server perspective.

To view the performance graphs for an initiator

1. In the Management Server console, go to the **Server** perspective and select **Manage** in the left pane.
2. Expand the Organization or **Uncategorized Hosts** to locate the host.
3. Select the host, and click the **Initiators** tab.
4. In the initiators list, right-click an initiator and select **Performance**. To change the duration, use the drop-down list.

See "About Veritas InfoScale Operations Manager performance graphs" on page 510.

Viewing the performance graphs for virtualization server and virtual machines

In the Management Server console, you can view the performance of virtualization servers and virtual machines in an interactive graph. You can review the performance for different durations - 6 hours, 24 hours, 1 week, 1 month, 1 year, or based on live data.

You can view the performance graphs for VMware and LDOM virtualization servers, and VMware virtual machines. You can view the performance graphs for LDOM and LPAR virtual machines only if the VRTSsfmh package is installed on them.

Table 30-13 lists the performance graphs for virtualization server and virtual machine.

Table 30-13 Performance graphs for virtualization server and virtual machine

Performance graph name	Description
Available Memory	Displays the free memory (in GB) in the selected virtualization server and virtual machine and the trend for the specified duration.
CPU Utilization	Displays the CPU utilization (in percentage) on the selected virtualization server and virtual machine and the trend for the specified duration.
Swap in Rate	Displays the swap in rate (in KB per second) for the selected virtualization server and virtual machine and the trend for the specified duration.
Used Swap	Displays the swap (in GB) that is used by the selected virtualization server and virtual machine and the trend for the specified duration.

You can view these performance graphs for the virtualization servers, for which your user group has at least Guest role explicitly assigned or inherited from a parent Organization. You can also view the graphs if your user group has at least Guest role assigned on the Virtualization perspective.

To view the performance graphs for an ESX Server

1. In the Management Server console, go to the **Virtualization** perspective and select **Manage** in the left pane.
2. Expand the Organization or **Uncategorized Servers** to locate the virtualization server.
3. Select the virtualization server, and click the **Performance** tab. To change the duration, use the drop-down list.

To view the performance graphs for a virtual machine

1. In the Management Server console, go to the **Virtualization** perspective and select **Manage** in the left pane.
2. Expand the Organization or **Uncategorized Servers** to locate the virtualization server.
3. Expand the virtualization server, and expand **Virtual Machines** to locate the virtual machine.
4. Select a virtual machine, and click the **Performance** tab. To change the duration, use the drop-down list.

See "About Veritas InfoScale Operations Manager performance graphs" on page 510.

See "Enable performance metering for a host" on page 508.

Viewing the performance graphs for a path of a virtualization server

In the Management Server console, you can view the performance for a path of a virtualization server in an interactive graph. You can review the performance of a path for different durations - 6 hours, 24 hours, 1 week, 1 month, or based on live data.

You can view the performance graphs for a path only if the path is managed by DMP or VxDMP.

Table 30-14 lists the performance graphs for a path of a virtualization server.

Table 30-14 Performance graphs for a path of a virtualization server

Performance graph name	Description
Average Read/Write Latency	Displays the average read and write latency of the selected path and the trend for the specified duration.
Bytes Read/Written	Displays the number of bytes read and written (in bytes) in the path and the trend for the specified duration.

You can view these performance graphs for the virtualization servers, for which your user group has at least Guest role explicitly assigned or inherited from a parent Organization. You can also view the graphs if your user group has at least Guest role assigned on the Virtualization perspective.

To view the performance graph for a path of a virtualization server

1. In the Management Server console, go to the **Virtualization** perspective and select **Manage** in the left pane.
2. Expand the Organization or **Uncategorized Servers** to locate the virtualization server.
3. Select the virtualization server and click the **Disks** tab to select a disk
4. In the **Paths** tab, right-click the path, and select **Performance**. To change the duration, use the drop-down list.

See "About Veritas InfoScale Operations Manager performance graphs" on page 510.

See "Viewing the performance graphs for virtualization server and virtual machines" on page 519.

Viewing the performance graphs for an enclosure

In the Management Server console, you can view the performance of an enclosure in an interactive graph. You can review the performance of an enclosure for different durations - 6 hours, 24 hours, 1 week, 1 month, 1 year, or based on live data.

You can view the performance graphs for an enclosure only if the storage provisioned to a host from the enclosure is managed by DMP or VxDMP. You can also view the performance charts if the enclosure is discovered by Storage Insight Add-on version 6.1 or later.

Average Read/Write Latency for Host and **Bytes Read/Written for Host** charts are not displayed for virtual machines having virtual initiator.

Table 30-15 lists the performance graphs for an enclosure.

Table 30-15 Performance graphs for an enclosure

Performance graph name	Description
Average Read/Write Latency for Host	Displays the average read and write latency and the trend for the selected host for the specified duration.
Bytes Read/Written for Host	Displays the bytes read and written (in KB) and the trend for the selected host for the specified duration
IO Operations per second	Displays the number of IO operations per second and the linear trend for the selected host for the specified duration. **Note:** You can view this graph only if Storage Insight Add-on version 6.1 or later is installed This graph is displayed only for EMC Symmetrix, EMC CLARiiON, EMC VNX (Block), NetApp and IBM XIV arrays. This graph cannot be rendered for **Live** data.

You can view these performance graphs for the enclosures, for which your user group has at least Guest role explicitly assigned or inherited from a parent Organization. You can also view the graphs if your user group has at least Guest role assigned on the Storage perspective.

To view the performance graphs for an enclosure

1. In the Management Server console, go to the **Storage** perspective and select **Manage** in the left pane.
2. Expand Organization or **Uncategorized Enclosures** to locate the enclosure.
3. Do one of the following:
   - If the enclosure is EMC VNX (Block), expand **Components** to select the **Block**.
   - For other enclosures, skip to step 4.
4. Click the **Performance** tab. To change the duration, use the drop-down list.
5. Click the ellipses to select a host.

See "About Veritas InfoScale Operations Manager performance graphs" on page 510.

About threshold settings

Using Veritas InfoScale Operations Manager 7.2, you can set the threshold values on the respective objects for the metrics that are described in Table 30-16. A Risk or Fault is raised when the threshold condition is violated.

Managed host version 6.1 or later is required to enable threshold settings on the host and is supported only from one Management Server. Threshold settings are not supported on an agentless host. You can view the configuration type of the host in the Management Server perspective.

Managed host versions 5.x and 6.x have the following default threshold values which are not displayed on the Management Server console.

- CPU utilization = 98 %
- Available Memory = 102400 KB

An alert of the risk severity type is raised if these default threshold values are violated. However if you define new threshold values on a managed host version 6.1, then the default values are overwritten.

The following table lists the objects and the metrics on which you can set the threshold value. It also lists the evaluation intervals for the objects.

Table 30-16 Objects, metrics, and evaluation intervals

Object	Metrics	Evaluation intervals
Host	CPU Utilization Available Memory Average CPU Load	5 minutes
Disk, volume, and path associated with a host.	Average Read Latency Average Write Latency	5 minutes On a UNIX/Linux host, the performance statistics for disk and volume are collected every one minute. Last five samples are considered for threshold evaluation and if any of these samples violate the threshold value, an alert is raised.
Host Initiator	Average Read Latency Average Write Latency Read Queue Length Write Queue Length Read Errors Write Errors	5 minutes **Note:** Threshold setting for host initiator on a Windows host is not supported.
Cluster and service group	Failover Duration	Threshold evaluation happens on failover event.

Sample values are collected for the performance metrics for an object at the interval mentioned in Table 30-16. If any of the collected sample values violate any of the threshold values defined for that metric, an alert is raised with severity specified in threshold definition.

The alert is cleared only when the collected performance sample value does not violate any of the threshold values defined for that performance metric.

See “Adding threshold settings for an object” on page 525.

See “Deleting the threshold settings for an object” on page 528.

See “Enabling the threshold settings for an object” on page 531.

See “Disabling the threshold settings for an object” on page 534.

Adding threshold settings for an object

Using the Management Server console, you can set the threshold values for an object. An appropriate fault or risk alert is raised if the threshold value is violated.

You can also select multiple objects of the same type, such as hosts, volumes, or disks. When you select multiple objects, the previously set values are overwritten. Based on the number of objects, the operation might take a couple of minutes to complete. You can view the status in the **Recent Tasks** pane.

You cannot select multiple clusters in the Availability perspective, to set the threshold value.

For a cluster, the set threshold value is applicable for all the service groups within the cluster, unless it is explicitly defined on a service group.

Metric, operator, threshold, and severity type of raised alert (fault or risk) values should not be identical while setting multiple thresholds on an object.

Note: Ensure that performance metering is enabled on the host to receive the threshold violation alerts (fault or risk).

- To add the threshold settings for a host
- To add the threshold settings for a disk
- To add the threshold settings for a volume
- To add the threshold settings for a path associated with a host
- To add the threshold settings for a host initiator
- To add the threshold settings for a cluster
- To add the threshold settings for a service group

To perform this task, your user group must be assigned the Admin role on the host or the Server perspective. The permissions on the host may be explicitly assigned or inherited from a parent Organization.

To add the threshold settings for a host

1. In the Management Server console, go to the **Server** perspective and select **Manage** in the left pane.
2. Expand the **Organization** or **Uncategorized Hosts** to locate the host.
3. Right-click the host and select **Properties**.
4. Click the **Performance** tab.
5. Click **Add** to create a blank row.

6 Select from the following options.

Metrics	Select the metrics from the drop-down list.
Operator	Select either '>' or '<'.
Threshold	Enter a threshold value.
Raise	Select either **Fault** or **Risk** alert severity type.
Status	Displays the status of the threshold setting. When you add new metrics the status is **Enabled** by default.

7 Click **Apply** and click **OK**.

To add the threshold settings for a disk

1 In the Management Server console, go to the **Server** perspective and select **Manage** in the left pane.

2 Expand the **Organization** or **Uncategorized Hosts** to locate the host.

3 Click the **Disk** tab.

4 Right-click a disk and select **Properties**.

5 Click the **Performance** tab.

6 Click **Add** to create a blank row.

7 Select from the options, click **Apply**, and then click **OK**.

To add the threshold settings for a volume

1 In the Management Server console, go to the **Server** perspective and select **Manage** in the left pane.

2 Expand the **Organization** or **Uncategorized Hosts** to locate the host.

3 Expand the host and then expand Volumes to locate the volume.

4 Right-click on the volume and select **Properties**.

5 Click the **Performance** tab.

6 Click **Add** to create a blank row.

7 Select from the options, click **Apply**, and then click **OK**.

To add the threshold settings for a path associated with a host

1. In the Management Server console, go to the **Server** perspective and select **Manage** in the left pane.
2. Expand the **Organization** or **Uncategorized Hosts** to locate the host.
3. Expand the host and then expand **Disk Groups** to select a disk group.
4. Click on the **Disks** tab.
5. In the **Paths** tab, right-click the path and select **Properties**.
6. Click the **Performance** tab.
7. Click **Add** to create a blank row.
8. Select from the options, click **Apply**, and then click **OK**.

To add the threshold settings for a host initiator

1. In the Management Server console, go to the **Server** perspective and select **Manage** in the left pane.
2. Expand the **Organization** or **Uncategorized Hosts** to locate the host.
3. Select the host, and click the **Initiators** tab.
4. In the initiators list, right-click an initiator and select **Properties**.
5. Click the **Performance** tab.
6. Click **Add** to create a blank row.
7. Select from the options, click **Apply**, and then click **OK**.

To add the threshold settings for a cluster

1. In the Management Server console, go to the **Availability** perspective and select **Manage** in the left pane.
2. Expand the **Organization** or **Uncategorized Clusters** to locate the cluster.
3. Right-click the cluster, and select **Properties**.
4. Click the **Performance** tab.
5. Click **Add** to create a blank row.
6. Select from the options, click **Apply**, and then click **OK**.

Note: To perform this task on a cluster, your user group must be assigned the Admin role on the cluster or the Availability perspective. The permission on the cluster may be explicitly assigned or inherited from a parent Organization.

To add the threshold settings for a service group

1. In the Management Server console, go to the **Availability** perspective and select **Manage** in the left pane.
2. Expand the **Organization** or **Uncategorized Clusters** to locate the cluster.
3. Expand the cluster, and expand **Service Groups**.
4. Right-click the service group and select **Properties**.
5. Click the **Performance** tab, click **Add** to create a blank row.
6. Select from the options, click **Apply**, and then click **OK**.

Note: To perform this task on a service group, your user group must be assigned the Admin or Operator role on the service group or the Availability perspective. The permission on the service group may be explicitly assigned or inherited from a parent Organization or cluster.

See "About threshold settings" on page 523.

See "Deleting the threshold settings for an object" on page 528.

See "Enabling the threshold settings for an object" on page 531.

See "Disabling the threshold settings for an object" on page 534.

Deleting the threshold settings for an object

Using the Management Server console, you can delete the threshold values set for an object.

- To delete the threshold settings for a host
- To delete the threshold settings for a disk
- To delete the threshold settings for a volume
- To delete the threshold settings for a path associated with a host
- To delete the threshold settings for an initiator
- To delete the threshold settings for a cluster
- To delete the threshold settings for a service group

To perform this task, your user group must be assigned the Admin role on the host or the Server perspective. The permissions on the host may be explicitly assigned or inherited from a parent Organization.

To delete the threshold settings for a host

1. In the Management Server console, go to the **Server** perspective and select **Manage** in the left pane.
2. Expand the **Organization** or **Uncategorized Hosts** to locate the host.
3. Right-click the host and select **Properties**.
4. Click the **Performance** tab.
5. Select the threshold settings which you want to delete and click **Delete**.
6. Click **Apply** and click **OK**.

To delete the threshold settings for a disk

1. In the Management Server console, go to the **Server** perspective and select **Manage** in the left pane.
2. Expand the **Organization** or **Uncategorized Hosts** to locate the host.
3. Click the **Disk** tab.
4. Right-click a disk and select **Properties**.
5. Click the **Performance** tab.
6. Select the threshold settings which you want to delete and click **Delete**.
7. Click **Apply** and click **OK**.

To delete the threshold settings for a volume

1. In the Management Server console, go to the **Server** perspective and select **Manage** in the left pane.
2. Expand the **Organization** or **Uncategorized Hosts** to locate the host.
3. Expand the host and then expand Volumes to locate the volume.
4. Right-click on the volume and select **Properties**.
5. Click the **Performance** tab.
6. Select the threshold settings which you want to delete and click **Delete**.
7. Click **Apply** and click **OK**.

To delete the threshold settings for a path associated with a host

1. In the Management Server console, go to the **Server** perspective and select **Manage** in the left pane.
2. Expand the **Organization** or **Uncategorized Hosts** to locate the host.
3. Expand the host and then expand **Disk Groups**.
4. Click on a disk.

5 In the **Paths** tab, right-click the path and select **Properties**.

6 Click the **Performance** tab.

7 Select the threshold settings which you want to delete and click **Delete**.

8 Click **Apply** and click **OK**.

To delete the threshold settings for an initiator

1 In the Management Server console, go to the **Server** perspective and select **Manage** in the left pane.

2 Expand the **Organization** or **Uncategorized Hosts** to locate the host.

3 Select the host, and click the **Initiators** tab.

4 In the initiators list, right-click an initiator and select **Properties**.

5 Click the **Performance** tab.

6 Select the threshold settings which you want to delete and click **Delete**.

7 Click **Apply** and click **OK**.

To delete the threshold settings for a cluster

1 In the Management Server console, go to the **Availability** perspective and select **Manage** in the left pane.

2 Expand the **Organization** or **Uncategorized Clusters** to locate the cluster.

3 Right-click the cluster, and select **Properties**.

4 Click the **Performance** tab.

5 Select the threshold settings which you want to delete and click **Delete**.

6 Click **Apply** and click **OK**.

Note: To perform this task on a cluster, your user group must be assigned the Admin role on the cluster or the Availability perspective. The permission on the cluster may be explicitly assigned or inherited from a parent Organization.

To delete the threshold settings for a service group

1 In the Management Server console, go to the **Availability** perspective and select **Manage** in the left pane.

2 Expand the **Organization** or **Uncategorized Clusters** to locate the cluster.

3 Expand the cluster, and expand **Service Groups**.

4 Right-click the service group and select **Properties**.

5 Click the **Performance** tab.

6 Select the threshold settings which you want to delete and click **Delete**.

7 Click **Apply** and click **OK**.

Note: To perform this task on a service group, your user group must be assigned the Admin or Operator role on the service group or the Availability perspective. The permission on the service group may be explicitly assigned or inherited from a parent Organization or cluster.

See "About threshold settings" on page 523.

See "Adding threshold settings for an object" on page 525.

Enabling the threshold settings for an object

Using the Management Server console, you can enable the threshold values for an object.

- To enable the threshold settings for a host
- To enable the threshold settings for a disk
- To enable the threshold settings for a volume
- To enable the threshold settings for a path associated with a host
- To enable the threshold settings for an initiator
- To enable the threshold settings for a cluster
- To enable the threshold settings for a service group

To perform this task, your user group must be assigned the Admin role on the host or the Server perspective. The permissions on the host may be explicitly assigned or inherited from a parent Organization.

To enable the threshold settings for a host

1 In the Management Server console, go to the **Server** perspective and select **Manage** in the left pane.

2 Expand the **Organization** or **Uncategorized Hosts** to locate the host.

3 Right-click the host and select **Properties**.

4 Click the **Performance** tab.

5 Select the threshold settings which you want to enable and click **Enable**.

6 Click **Apply** and click **OK**.

To enable the threshold settings for a disk

1. In the Management Server console, go to the **Server** perspective and select **Manage** in the left pane.
2. Expand the **Organization** or **Uncategorized Hosts** to locate the host.
3. Click the **Disk** tab.
4. Right-click a disk and select **Properties**.
5. Click the **Performance** tab.
6. Select the metrics threshold settings you want to enable and click **Enable**.
7. Click **Apply** and click **OK**.

To enable the threshold settings for a volume

1. In the Management Server console, go to the **Server** perspective and select **Manage** in the left pane.
2. Expand the **Organization** or **Uncategorized Hosts** to locate the host.
3. Expand the host and then expand Volumes to locate the volume.
4. Right-click on the volume and select **Properties**.
5. Click the **Performance** tab.
6. Select the threshold settings which you want to enable and click **Enable**.
7. Click **Apply** and click **OK**.

To enable the threshold settings for a path associated with a host

1. In the Management Server console, go to the **Server** perspective and select **Manage** in the left pane.
2. Expand the **Organization** or **Uncategorized Hosts** to locate the host.
3. Expand the host and then expand **Disk Groups**.
4. Click on a disk.
5. In the **Paths** tab, right-click the path and select **Properties**.
6. Click the **Performance** tab.
7. Select the threshold settings which you want to enable and click **Enable**.
8. Click **Apply** and click **OK**.

To enable the threshold settings for an initiator

1. In the Management Server console, go to the **Server** perspective and select **Manage** in the left pane.
2. Expand the **Organization** or **Uncategorized Hosts** to locate the host.

3 Select the host, and click the **Initiators** tab.

4 In the initiators list, right-click an initiator and select **Properties**.

5 Click the **Performance** tab.

6 Select the threshold settings which you want to enable and click **Enable**.

7 Click **Apply** and click **OK**.

To enable the threshold settings for a cluster

1 In the Management Server console, go to the **Availability** perspective and select **Manage** in the left pane.

2 Expand the **Organization** or **Uncategorized Clusters** to locate the cluster.

3 Right-click the cluster, and select **Properties**.

4 Click the **Performance** tab.

5 Select the threshold settings which you want to enable and click **Enable**.

6 Click **Apply** and click **OK**.

Note: To perform this task on a cluster, your user group must be assigned the Admin role on the cluster or the Availability perspective. The permission on the cluster may be explicitly assigned or inherited from a parent Organization.

To enable the threshold settings for a service group

1 In the Management Server console, go to the **Availability** perspective and select **Manage** in the left pane.

2 Expand the **Organization** or **Uncategorized Clusters** to locate the cluster.

3 Expand the cluster, and expand **Service Groups**.

4 Right-click the service group and select **Properties**.

5 Click the **Performance** tab.

6 Select the threshold settings which you want to enable and click **Enable**.

7 Click **Apply** and click **OK**.

Note: To perform this task on a service group, your user group must be assigned the Admin or Operator role on the service group or the Availability perspective. The permission on the service group may be explicitly assigned or inherited from a parent Organization or cluster.

See “About threshold settings” on page 523.

See “Disabling the threshold settings for an object” on page 534.

Disabling the threshold settings for an object

Using the Management Server console, you can disable the threshold values for an object.

- To disable the threshold settings for a host
- To disable the threshold settings for a disk
- To disable the threshold settings for a volume
- To disable the threshold settings for a path associated with a host
- To disable the threshold settings for an initiator
- To disable the threshold settings for a cluster
- To disable the threshold settings for a service group

To perform this task, your user group must be assigned the Admin role on the host or the Server perspective. The permissions on the host may be explicitly assigned or inherited from a parent Organization.

To disable the threshold settings for a host

1. In the Management Server console, go to the **Server** perspective and select **Manage** in the left pane.
2. Expand the **Organization** or **Uncategorized Hosts** to locate the host.
3. Right-click the host and select **Properties**.
4. Click the **Performance** tab.
5. Select the threshold settings which you want to disable and click **Disable**.
6. Click **Apply** and click **OK**.

To disable the threshold settings for a disk

1. In the Management Server console, go to the **Server** perspective and select **Manage** in the left pane.
2. Expand the **Organization** or **Uncategorized Hosts** to locate the host.
3. Click the **Disk** tab.
4. Right-click a disk and select **Properties**.
5. Click the **Performance** tab.
6. Select the threshold settings which you want to disable and click **Disable**.
7. Click **Apply** and click **OK**.

To disable the threshold settings for a volume

1. In the Management Server console, go to the **Server** perspective and select **Manage** in the left pane.
2. Expand the **Organization** or **Uncategorized Hosts** to locate the host.
3. Expand the host and then expand Volumes to locate the volume.
4. Right-click on the volume and select **Properties**.
5. Click the **Performance** tab.
6. Select the threshold settings which you want to disable and click **Disable**.
7. Click **Apply** and click **OK**.

To disable the threshold settings for a path associated with a host

1. In the Management Server console, go to the **Server** perspective and select **Manage** in the left pane.
2. Expand the **Organization** or **Uncategorized Hosts** to locate the host.
3. Expand the host and then expand **Disk Groups**.
4. Click on a disk.
5. In the **Paths** tab, right-click the path and select **Properties**.
6. Click the **Performance** tab.
7. Select the threshold settings which you want to disable and click **Disable**.
8. Click **Apply** and click **OK**.

To disable the threshold settings for an initiator

1. In the Management Server console, go to the **Server** perspective and select **Manage** in the left pane.
2. Expand the **Organization** or **Uncategorized Hosts** to locate the host.
3. Select the host, and click the **Initiators** tab.
4. In the initiators list, right-click an initiator and select **Properties**.
5. Click the **Performance** tab.
6. Select the threshold settings which you want to disable and click **Disable**.
7. Click **Apply** and click **OK**.

To disable the threshold settings for a cluster

1. In the Management Server console, go to the **Availability** perspective and select **Manage** in the left pane.
2. Expand the **Organization** or **Uncategorized Clusters** to locate the cluster.

3 Right-click the cluster, and select **Properties**.

4 Click the **Performance** tab.

5 Select the threshold settings which you want to disable and click **Disable**.

6 Click **Apply** and click **OK**.

Note: To perform this task on a cluster, your user group must be assigned the Admin role on the cluster or the Availability perspective. The permission on the cluster may be explicitly assigned or inherited from a parent Organization.

To disable the threshold settings for a service group

1 In the Management Server console, go to the **Availability** perspective and select **Manage** in the left pane.

2 Expand the **Organization** or **Uncategorized Clusters** to locate the cluster.

3 Expand the cluster, and expand **Service Groups**.

4 Right-click the service group and select **Properties**.

5 Click the **Performance** tab.

6 Select the threshold settings which you want to disable and click **Disable**.

7 Click **Apply** and click **OK**.

Note: To perform this task on a service group, your user group must be assigned the Admin or Operator role on the service group or the Availability perspective. The permission on the service group may be explicitly assigned or inherited from a parent Organization or cluster.

See “About threshold settings” on page 523.

See “Enabling the threshold settings for an object” on page 531.

Chapter 31

Managing Business Applications

This chapter includes the following topics:

- About Business Applications in Veritas InfoScale Operations Manager
- Creating or modifying a Business Application
- Renaming a Business Application
- Deleting a Business Application
- Viewing Business Applications in the data center
- Viewing the overview of a Business Application
- Viewing service availability for a Business Application
- Viewing data availability for a Business Application
- Viewing SAN connectivity for a Business Application
- About the makeBE script
- Creating Business Application using the makeBE script
- Importing Business Application using the makeBE script
- Exporting Business Application using the makeBE script
- Updating Business Application using the makeBE script
- Deleting Business Application using the makeBE script

About Business Applications in Veritas InfoScale Operations Manager

A Business Application is a collection of objects within Veritas InfoScale Operations Manager that constitutes a unit of management. The objects can be of different types.

One use case for Business Applications is as follows: Veritas InfoScale Operations Manager can discover some applications, such as some database applications, and display information about their related objects. In cases where Veritas InfoScale Operations Manager is unable to discover an application, you can create a Business Application. For this purpose, you can create a representation of the application using a base object. For example, if the application MyApp is the only one that uses a specific disk group, MyAppDG, then the MyAppDG disk group can be used as the base object for the MyApp Business Application. Veritas InfoScale Operations Manager is then able to determine all associated objects for that disk group and collect that data for monitoring in the Business Application.

You can select one or more of the following objects when creating Business Applications:

- Hosts
- Volumes
- Disk Groups
- Service Groups
- Databases
- Exchange Servers

You can monitor the status of the set of associated objects, both in overview and in detail.

See "Creating or modifying a Business Application" on page 538.

See "Renaming a Business Application" on page 539.

See "Deleting a Business Application" on page 540.

See "Viewing Business Applications in the data center" on page 540.

See "Viewing the overview of a Business Application" on page 541.

Creating or modifying a Business Application

You can use the Management Server console to create or modify a Business Application. A Business Application is a collection of objects within Veritas InfoScale

Operations Manager that constitutes a unit of management. The objects can be of different types.

You can create multiple Business Applications in the data center.

To perform this task, your user group must be assigned the Admin role on the host or the Server perspective. The permission on the host may be explicitly assigned or inherited from a parent Organization.

To create a Business Application

1. In the Management Server console, go to the **Server** perspective and expand **Manage** in the left pane.
2. Right-click **Business Applications** and select **Create Business Application**.
3. In the **Create Business Application** wizard panel, enter a name for the Business Application and optionally enter a description. Select one or more of the available object types: hosts, volumes, disk groups, service groups, databases, and Exchange Servers. Click **Next**.
4. On the subsequent panels, select the objects to add to the Business Application.
5. Review the information on the **Summary** panel and if satisfied, click **Next**.

   The Business Application is listed on the table. To view details, double-click the table row.

To modify a Business Application

1. In the Management Server console, go to the **Server** perspective and expand **Manage** in the left pane.
2. Expand **Business Applications**.
3. Right-click a Business Application and select **Modify**.

   If you want to modify only the name of the Business Application, it is easier to use the Rename option.

4. In the **Modify Business Application** wizard panel, change any information desired, including the name. Click **Next**.
5. On the subsequent panels, select the objects to add to the Business Application.
6. Review the information on the **Summary** panel and if satisfied, click **Next**.

See "About Business Applications in Veritas InfoScale Operations Manager" on page 538.

Renaming a Business Application

In the Management Server console you can rename existing Business Applications.

To perform this task, your user group must be assigned the Admin role on the host or the Server perspective. The permission on the host may be explicitly assigned or inherited from a parent Organization.

To rename a Business Application

1. In the Management Server console, go to the **Server** perspective and expand **Manage** in the left pane.
2. Expand **Business Applications**.
3. Right-click a Business Application and select **Rename**.
4. Type the new name and optionally a description and click **Next**.

See "About Business Applications in Veritas InfoScale Operations Manager" on page 538.

Deleting a Business Application

In the Management Server console you can delete existing Business Applications.

To perform this task, your user group must be assigned the Admin role on the host or the Server perspective. The permission on the host may be explicitly assigned or inherited from a parent Organization.

To delete a Business Application

1. In the Management Server console, go to the **Server** perspective and expand **Manage** in the left pane.
2. Expand **Business Applications**.
3. Right-click a Business Application and select **Delete**.
4. Click **Yes** to proceed with the deletion.

See "About Business Applications in Veritas InfoScale Operations Manager" on page 538.

Viewing Business Applications in the data center

You can use the Management Server console to view a list of the Business Applications in the data center.

For each Business Application you can view a graphical display showing status of service availability, data availability, and SAN connectivity. By hovering on the status icon you can view more details.

You can view this information related to the hosts for which your user group has at least Guest role explicitly assigned or inherited from a parent Organization. You

can also view the information if your user group has at least Guest role assigned on the Server perspective.

To view the Business Applications in the data center

1. In the Management Server console, go to the **Server** perspective and expand **Manage** in the left pane.
2. In the tree, click **Business Applications**.
3. For a status overview, hover on any status icon. You can double-click a business application for more information.

See "About Business Applications in Veritas InfoScale Operations Manager" on page 538.

See "Viewing the overview of a Business Application" on page 541.

Viewing the overview of a Business Application

You can use the Management Server console to view a summary of the status of a Business Application.

The **Overview** tab includes status panels for the following:

- **Service Availability:** Shows information about whether high availability clusters are configured, whether disaster recovery is configured, and the status of clusters and related objects, as follows:
  - **Health**: Determined by the state of service groups and clusters that provide high availability.
  - **Hidden Risks**: Determined by service groups in a frozen state or critical resources that are not enabled.
  - **Stability**: Determined by the number of unplanned failovers within the last seven days.
- **Data Availability**: Shows information about replication and snapshots and shows the status of storage assets.
  - **Health**: Determined by the status of volumes and file systems, disks, and Replicated Volume Groups (RVGs).
  - **Hidden Risks**: Determined by amount of data redundancy using volume snapshots and file system copies.
  - **Stability**: Applies only if VVR is configured. Stability is determined by the amount of the Storage Replicator Log (SRL) lag.
- **SAN Connectivity**: Shows information on whether redundancy is configured for the storage I/O paths. Redundancy includes multi-pathing at the host level

and single point of failure at the switch or fabric level. The availability of switch and fabric information depends on whether Fabric Insight Add-on is configured.

- **Health**: Determined by switch level or fabric level failures or by disks in a multipathing degrade state.
- **Hidden Risks**: Determined by the number of disks without I/O redundancy configured.

The **Overview** tab also includes charts, including Storage Distribution (by vendor) and Storage Usage (Raw, Usable, Consumed), and an activity log.

You can view this information related to the hosts for which your user group has at least Guest role explicitly assigned or inherited from a parent Organization. You can also view the information if your user group has at least Guest role assigned on the Server perspective.

To view the overview of a Business Application

1. In the Management Server console, go to the **Server** perspective and expand **Manage** in the left pane.
2. Expand **Business Applications**.
3. Select a Business Application and click the **Overview** tab.

   For details about the metrics and computations that are used in a status panel, click the status panel. To view the computations used, hover on an information icon on the details panel.
4. More information is available on other tabs.

   See "Viewing service availability for a Business Application" on page 542.

   See "Viewing data availability for a Business Application" on page 543.

   See "Viewing SAN connectivity for a Business Application" on page 544.

See "About Business Applications in Veritas InfoScale Operations Manager" on page 538.

Viewing service availability for a Business Application

You can use the Management Server console to view information about the Business Application objects related to service availability.

The objects listed can include the following:

- Clusters
- Service groups

- Resources

You can view this information related to the hosts for which your user group has at least Guest role explicitly assigned or inherited from a parent Organization. You can also view the information if your user group has at least Guest role assigned on the Server perspective.

To view service availability for a Business Application

1. In the Management Server console, go to the **Server** perspective and expand **Manage** in the left pane.
2. Expand **Business Applications**.
3. Select a Business Application and click the **Service Availablity** tab.

   You can double-click a cluster to drill down to the **Availability** perspective, if you have access to that perspective.

See "About Business Applications in Veritas InfoScale Operations Manager" on page 538.

See "Viewing the overview of a Business Application" on page 541.

Viewing data availability for a Business Application

You can use the Management Server console to view information about the Business Application objects related to data availability.

The objects listed can include the following:

- Databases
- Exchange Servers
- Hosts
- Initiators
- Disk groups
- Volumes
- Disks
- Replications

Depending on the object type, graphs show performance information for the selected object. For example, for the selected host, a graph shows average CPU load and load average trend. The same information is available if you drill down to the host and view its performance graph.

You can view this information related to the hosts for which your user group has at least Guest role explicitly assigned or inherited from a parent Organization. You can also view the information if your user group has at least Guest role assigned on the Server perspective.

To view data availability for a Business Application

1. In the Management Server console, go to the **Server** perspective and expand **Manage** in the left pane.
2. Expand **Business Applications**.
3. Select a Business Application and click the **Data Availability** tab.

   You can drill down to objects that are available as top-level objects in the tree, including hosts, databases, or Exchange Servers. To drill down to an object, double-click on it.

See "About Business Applications in Veritas InfoScale Operations Manager" on page 538.

See "Viewing the overview of a Business Application" on page 541.

Viewing SAN connectivity for a Business Application

You can use the Management Server console to view information about the Business Application objects related to SAN connectivity.

The objects listed can include the following, depending on whether Fabric Insight Add-on is configured:

- Paths
- Enclosures
- RAID groups
- Thin pools
- Adapters
- Array ports
- Switches

You can view this information related to the hosts for which your user group has at least Guest role explicitly assigned or inherited from a parent Organization. You can also view the information if your user group has at least Guest role assigned on the Server perspective.

To view SAN connectivity for a Business Application

1. In the Management Server console, go to the **Server** perspective and expand **Manage** in the left pane.
2. Expand **Business Applications**.
3. Select a Business Application and click the **SAN Connectivity** tab.

See "About Business Applications in Veritas InfoScale Operations Manager" on page 538.

See "Viewing the overview of a Business Application" on page 541.

About the makeBE script

Using makeBE command-line script, Veritas InfoScale Operations Manager users can perform various operations related to Business Applications using command line. The makeBE script is packaged with the `VRTSsfmcs` package, and you need to use the `vomadm` utility to run makeBE script. The `vomadm` utility lets you run various commands on Veritas InfoScale Operations Manager Management Server. The makeBE script is one such command.

The makeBE script lets you perform the following operations that are related to Business Application:

- Create a Business Application
- Update a Business Application
- Delete a Business Application
- Import Business Applications from a comma separated value (CSV) file. You need to ensure that the CSV file already contains the information about the objects that you want to add to the Business Application.
- Export the objects of a Business Application to a comma separated value (CSV) file.

In the high availability (HA) environment, the makeBE commands are not supported on non-active cluster nodes (slave). In non-high availability environment, makeBE commands are not supported on the managed host.

See "Objects used in the makeBE script" on page 546.

See "makeBE script CSV file details" on page 546.

See "makeBE script parameters details" on page 547.

See "Limitations of makeBE script" on page 548.

Objects used in the makeBE script

The makeBE script uses the following objects for its Business Application-related operations:

- **Base objects**: These objects are typically all the objects that you use when you create a Business Application using the Veritas InfoScale Operations Manager console. The objects are host, database, disk group, File System, service group, enclosure, cluster, and Exchange server.
- **makeBE CSV file**: The `makeBE` CSV file contains one or many occurrences of the following pattern in each line:

  `<BusinessApps_name>,<obj_type>,<obj_id>,<oper_type>,<optional information column>`

  Each line instructs the `makeBE` command how to manipulate the Business Application. Each `makeBE` CSV file may contain multiple lines to create the Business Application. You can use a `makeBE` CSV file to manipulate multiple Business Applications.

See "Creating Business Application using the makeBE script" on page 549.

See "Importing Business Application using the makeBE script" on page 550.

See "Exporting Business Application using the makeBE script" on page 551.

See "Updating Business Application using the makeBE script" on page 552.

makeBE script CSV file details

The makeBE script CSV file contains one or many occurrences of the following pattern in each line:

`<BusinessApps_name>,<obj_type>,<obj_id>,<oper_type>,<optional information column>`

For example

`'BusinessApp1','host','{00120050-5680-5af4-0000-000094d5b73c}','I','system1.domain.com'`

The detailed information about the makeBE CSV file parameters is provided in the following table:

Table 31-1 `makeBE` script CSV file details

`<BusinessApps_name>`	You can use `BusinessApps_name` to refer to any existing or a new Business Application. This parameter is common to import, export, and user-defined import operations.

Table 31-1 `makeBE` script CSV file details *(continued)*

`<oper_type>`	Indicates the type of operation performed on the Business Application. ■ I: This option is used to add an object to the Business Application, or update if the object already exists. ■ D: This option is used to remove an object from the Business Application. These parameters are common to import, export, and user-defined import operations.
`<obj_type>` and `<obj_id>`	Provides the description for the Business Application object. ■ <obj_type>='AE_description', <obj_id>=detail description of the Business Application The above-mentioned parameter is common to import, export, and user-defined import operations. However, the following parameters are specific to import and export operations: ■ <obj_type>='host',<obj_id>=host_id in the host table ■ <obj_type>='database',<obj_id>=db_id in the db table ■ <obj_type>='diskgroup',<obj_id>=dg_id in the diskgroup table ■ <obj_type>='filesystem',<obj_id>=fs_id in the filesystem table ■ <obj_type>='servicegroup',<obj_id>=sg_id in the servicegroup table ■ <obj_type>='enclosure',<obj_id>=enc_id in the array_subsystem table ■ <obj_type>='cluster',<obj_id>=hc_id in the host_container table ■ <obj_type>='exchange_server',<obj_id>=exch_id in the exchange_server table The following parameter is specific to the user-defined import operation: ■ <obj_type>='host_name_pattern',<obj_id>=regular expression pattern
`<optional information column>`	This column is populated with the object name when the export option is given. For the other operations, this column is ignored. This option is common to import, export, and user-defined import operations.

See "About the makeBE script " on page 545.

makeBE script parameters details

You can view the `makeBE` script-related parameters by typing the following command on the Veritas InfoScale Operations Manager Management Server:

```
perl vomadm makeBE --help
```

The command output displays three operations that you can perform to manipulate the Business Applications. The supported operations are **import**, **export**, and **user_defined_import**. With these commands, you also need to specify the `infile` (a CSV file) that contains the information about the Business Application, its objects, and the type of operation you want to perform on the Business Application.

```
vomadm makeBE [--import <infile> | --export <outfile> |
--user_defined_import <infile>]
```

`infile`: full path to the input file

`outfile`: full path to the output file

- import: Use import option to create a Business Applications, add objects to a Business Applications, and remove the objects from the selected Business Applications. Import option also lets you update and delete an existing Business Applications. The operation type parameter (`<oper_type>`) mentioned in the CSV file determines the type of operation. To add an object to the Business Application, the value of `<oper_type>` is set to I, and for the delete operation, it is set to D.
- export: Use export option to export a Business Application's content to a CSV file. The file can be used for the import, or verification purpose.
- user_defined_import: Like the import option, you can use this option to add or remove Business Application objects. However, this option also provides additional flexibility for the Business Application operations. For instance, you can specify `<obj_type>` as 'host_name_pattern' and `<obj_id>` as the 'regular expression pattern'.
  For example, the following expression adds all the hosts starting with 'vcs' to the TestBA Business Application: In this example, the value of 'regular expression pattern' `is vcs.*`.
  ```
  'TestBA','host_name_pattern',vcs.*,'I',
  ```

See “About the makeBE script ” on page 545.

Limitations of makeBE script

This section lists the limitations of makeBE script:

- You cannot export a single Business Application object using the `makeBE` script from Management Server.
- Virtual Business Services operations are not supported.

See “About the makeBE script ” on page 545.

makeBE script log files

The generated log files of the makeBE operations are listed below:

- UNIX
  - `/var/opt/VRTSsfmh/logs/vomadm.log`
  - `/var/VRTSsfmcs/SFMdb3.dblog`
- Windows
  - `<appdirmh>/logs/vomadm.log`
  - `<appdircs>/logs/SFMdb3.dblog`

See "About the makeBE script " on page 545.

Creating Business Application using the makeBE script

You can create a Business Application using the makeBE script. To create a Business Application, you need to use the import option of the makeBE script. The import option uses an existing CSV file to create a Business Application. The content of the CSV file determines the Business Application's objects.

To create a Business Application using the makeBE script

1. Log on to the Management Server as the administrator or the root user.
2. Run the following commands on Management Server:
   - For UNIX Management Server: `/opt/VRTSsfmh/bin> perl vomadm makeBE --import /tmp/Samplefile.csv`
   - For Windows Management Server: `C:\Program Files\VERITAS\VRTSsfmh\bin>perl.exe vomadm makeBE --import C:\tmp\Samplefile.csv`
3. After the successful import, the Business Application is displayed on the Management Server console.

A sample CSV file used for the Business Application create operation is listed below. The Business Application is TestBA that contains a host, enclosure, database, disk group, and file system.

```
'TestBA','host','{422e85ef-33e8-9257-2be9-704ff07a9d38}','I',

'TestBA','enclosure','EMC_CLARiiON_CK200082401064','I',
'EMC_CLARiiON_CK200082401064'
```

```
'TestBA','database','MSSQL@win2k3-6\\SQL32BIT@
{0012001a-a03b-5952-0000-000038caa494}@-1107251074','I','TestDB'

'TestBA','diskgroup','{b932c8c8-b9d0-11e0-a90c-001aa03b5950}','I',
'win2k3-6-Dg0'

'TestBA','filesystem','\\?\\Volume{58ca12bd-11a6-11e0-9946-001aa03b5950}
\\','I','\\?\\Volume{58ca12bd-11a6-11e0-9946-001aa03b5950}\\'
```

See "About the makeBE script " on page 545.

Importing Business Application using the makeBE script

You can import the content of an existing CSV file using the `makeBE` script, and create a Business Application. The content of the imported CSV file determines the Business Application's objects.

To import a Business Application using the makeBE script

1. Log on to the Management Server as the administrator.
2. Run the following commands on Management Server:
   - For UNIX Management Server: `/opt/VRTSsfmh/bin> perl vomadm makeBE --import /temp/Samplefile.csv`
   - For Windows Management Server: `C:\Program Files\VERITAS\VRTSsfmh\bin>perl.exe vomadm makeBE --import C:\temp\Samplefile.csv`

A sample CSV file used for the Business Application import operation is listed below. The Business Application is TestBA that contains a database, host, service group, file system, and a storage enclosure.

```
'TestBA','database','Sybase@nbk1@sybaseha@

{00020003-bae2-42d7-0000-000083e242d7}','I','nbk1'

'TestBA','host','{564dd88a-0ae2-a9bb-7854-b50e3e0376ee}','I',

'TestBA','servicegroup','VCSAppMonHBSG@ApplicationHA_pl8dom5-',

'I','VCSAppMonHBSG'

'TestBA','filesystem','{00120015-c5f4-2b83-0000-00004a2cf9a3}

\\Device\\Harddisk0\\Partition2','I',

'{00120015-c5f4-2b83-0000-00004a2cf9a3}\\Device\\Harddisk0\\Partition2'
```

```
'TestBA','enclosure','EMC_CLARiiON_CK200082401064',
'I','EMC_CLARiiON_CK200082401064'
'TestBA','BusinessApps_description','Full description of the Business Application','I',
```

After the successful import, the Business Application is displayed on Management server console containing the above listed objects.

See "About the makeBE script " on page 545.

Exporting Business Application using the makeBE script

If you have created a Business Application using the Management Server console, you can export the Business Application-related information to a CSV file using the export option of makeBE script.

To export a Business Application using the makeBE script

1. Log on to the Management Server as the administrator.
2. Run the following commands on Management Server:
   - For UNIX Management Server: `/opt/VRTSsfmh/bin> perl vomadm makeBE --export /tmp/Samplefile.out`
   - For Windows Management Server: `C:\Program Files\VERITAS\VRTSsfmh\bin>perl.exe vomadm makeBE --export C:\temp\Samplefile.out`

A sample CSV file generated after the Business Application export operation is listed below. The Business Application is TestBA that contains three hosts and two storage enclosures. The last line of the CSV file includes Business Application description.

Each object entry also includes its object ID. For example, each host entry includes its host_id. For the first host entry, the object ID is 422e85ef-33e8-9257-2be9-704ff07a9d38.

```
'TestBA','host','{422e85ef-33e8-9257-2be9-704ff07a9d38}','I',
'TestBA','host','{4207f2c1-8538-115f-d3b0-80638d0b0f2d}','I',
'TestBA','host','{42070e20-1106-3042-99ea-428e38dea523}','I',
'TestBA','enclosure','EMC000287971357','I','EMC000287971357'
'TestBA','enclosure','EMC_CLARiiON_CK200082401064','I','EMC_CLARiiON_CK200082401064'
```

```
'TestBA','BA_description','Business Apps description.','I',
```

See “About the makeBE script ” on page 545.

Updating Business Application using the makeBE script

Using the makeBE script, you can update an existing Business Application. For example, adding or removing host, service groups, databases, clusters and so on. The `--user_defined_import` option of `makeBE` script lets you specify the CSV file that contains the necessary information to update the Business Application.

To update a Business Application using the makeBE script

1. Log on to the Management Server as the administrator.
2. Run the following commands on Management Server:
   - For UNIX Management Server: `/opt/VRTSsfmh/bin> perl vomadm makeBE --user_defined_import /temp/Samplefile.out`
   - For Windows Management Server: `C:\Program Files\VERITAS\VRTSsfmh\bin>perl.exe vomadm makeBE --user_defined_import C:\temp\Samplefile.csv`

A sample CSV file used for the user_defined_import operation is listed below:

```
'ALL_Hosts','host_name_pattern',vom-.*,'I',
'ALL_Hosts','AE_description',' ','I',
```

Note: `--user_defined_import` option provides you with the flexibility to enter host name pattern. For example, all the host name starting with vom-.*.

See “About the makeBE script ” on page 545.

Deleting Business Application using the makeBE script

You can delete an existing Business Application using the `makeBE` script. To delete a Business Application, you need to use the import option of the `makeBE` script. The import option uses an existing CSV file to delete a Business Application. The <oper_type> parameter (set to D) in the CSV file indicates the type of operation (delete).

To delete a Business Application using the makeBE script

1. Log on to the Management Server as the administrator.
2. Run the following commands on Management Server:
   - For UNIX Management Server: `/opt/VRTSsfmh/bin> perl vomadm makeBE --import /tmp/Samplefile`
   - For Windows Management Server: `C:\Program Files\VERITAS\VRTSsfmh\bin>perl.exe vomadm makeBE --import C:\tmp\Samplefile.csv`

A sample CSV file used for the delete operation is listed below. Note that the value of `<oper_type>` defined in CSV file is D. It indicates that delete operation is performed on the Business Application.

```
'TestBA','BA_description','Business Application description','D',
```

See “About the makeBE script ” on page 545.

Chapter 32

Managing extended attributes

This chapter includes the following topics:

- About using extended attributes
- Setting values to the extended attributes on an object
- Searching objects to set extended attribute values
- Modifying the extended attributes value on an object

About using extended attributes

An extended attribute is an additional user-defined attribute that provides additional details about an object in Veritas InfoScale Operations Manager. This information about an object is in addition to the details that Veritas InfoScale Operations Manager discovers for that object. You can define multiple extended attributes on the objects using the Management Server console. You can use the extended attributes to search, filter, and sort the objects in the Management Server console. You can also manage the extended attributes using the Veritas InfoScale Operations Manager Web services API.

You can define an extended attribute and associate it with the relevant objects. You need to set the value for the extended attribute when you associate it with the object.

Table 32-1 lists the object types supporting extended attributes and the perspective to which these objects belong to.

Table 32-1 Object types supporting extended attributes

Objects	Perspective
Host, disk, disk group, volume, snapshot, exchange server, database	Server
Cluster, service group	Availability
Enclosure, switch, fabric, fabric zone	Storage
Virtualization server, virtual machine	Virtualization

You can set the value for the extended attribute on the objects in one of the following ways:

- By selecting an object in a perspective
- By searching and filtering the objects in the data center in a perspective

Note: For more information on adding an extended attribute, refer to the *Veritas InfoScale Operations Manager Installation and Configuration Guide*.

Setting values to the extended attributes on an object

Using the Management Server console, you can set the values for the extended attributes that are defined for an object in the relevant perspective.

Table 32-2 lists the object types supporting extended attributes and the perspective to which these objects belong to.

Table 32-2 Object types supporting extended attributes

Objects	Perspective
Host, file system, disk, disk group, volume, exchange server	Server
Cluster, service group	Availability
Enclosure, switch, fabric, fabric zone	Storage
Virtualization server, virtual machine	Virtualization

For more information on extended attributes, refer to the *Veritas InfoScale Operations Manager Installation and Configuration Guide*.

To perform this task, your user group must be assigned the Admin role on the object or on the relevant perspective. The permission on the object may be explicitly assigned or inherited from a parent Organization.

To set values to the extended attributes on one or more objects

1. In the Home page on the Management Server console, go to the perspective and select **Manage** in the left pane.
2. Expand the Organization or the Uncategorized folder to locate the object.
3. Right-click the object and select **Set Extended Attributes**.
4. In the **Set Extended Attributes** panel, enter the value for the extended attribute. Click **OK**.

   See "Set Extended Attributes panel options" on page 556.
5. In the **Set Extended Attributes - Result** panel, click **Close**.

Set Extended Attributes panel options

Use this panel to specify the values for the extended attributes.

Select the extended attribute check box for which you want to enter a value. The maximum length of the value can be 256 characters.

Searching objects to set extended attribute values

Using the Management Server console, you can search and filter the objects in a perspective, and set the extended attribute values.

Table 32-3 lists the object types supporting extended attributes and the perspective to which these objects belong to.

Table 32-3 Object types supporting extended attributes

Objects	Perspective
Host, file system, disk, disk group, volume, exchange server	Server
Cluster, service group	Availability
Enclosure, switch, fabric, fabric zone	Storage
Virtualization server, virtual machine	Virtualization

For more information on extended attributes, refer to the *Veritas InfoScale Operations Manager Installation and Configuration Guide*.

To perform this task, your user group must be assigned the Admin role on the relevant perspective.

To search objects to set the extended attribute values

1. In the Management Server console, go to the perspective and select **Manage** in the left pane.
2. Right-click **Data Center** and select **Set Extended Attributes**.
3. In the **Set Extended Attributes** wizard panel, enter the search criteria, and click **Preview** if you want to view the search results.

   See "Set Extended Attributes panel options" on page 557.
4. Click **Set Extended Attributes**.
5. In the **Set Extended Attributes** panel, enter the value for the extended attribute, and click **OK**.

   See "Set Extended Attributes panel options" on page 556.
6. In the **Set Extended Attributes - Result** panel click **Close**.

Set Extended Attributes panel options

Use this wizard panel to search for objects in a perspective.

Table 32-4 Set Extended Attributes panel options

Field	Description
Search for	Select the object that you want to search for
Select columns	Choose the columns to appear in the search results.
Reset	Click to clear the selections.
Update	Click to update the columns that appear in the search results.
Attribute	Select the attribute for which you want to specify a value.
Condition	Select the condition, for example, **Starts With**.
Value	Type the value. Value strings are not case-sensitive.

Table 32-4 Set Extended Attributes panel options *(continued)*

Field	Description
Add	Click to add another parameter to the search query.
Remove	Click to remove from the query.
Operator	Choose whether to use an AND or OR operator for the new parameter.

Modifying the extended attributes value on an object

Using the Management Server console you can modify the values that you have set for the extended attributes defined for a specific object. You can modify the values for the extended attributes from the details view of the relevant object.

When you modify the value of an extended attribute for a single object, the existing values are displayed. You need to overwrite the existing value. When you select multiple objects, the existing extended attribute values are not displayed.

For more information on extended attributes, refer to the *Veritas InfoScale Operations Manager Installation and Configuration Guide*.

To perform this task, your user group must be assigned the Admin role on the object or on the relevant perspective. The permission on the object may be explicitly assigned or inherited from a parent Organization.

To modify the values that are set for the extended attributes on an object

1. In the Home page on the Management Server console, go to the perspective and select **Manage** in the left pane.
2. Expand the Organization or the Uncategorized folder to locate the object.
3. In the details view of the object, select one or more objects.
4. Right-click the objects and select **Set Extended Attributes**.
5. In the **Set Extended Attributes** panel, modify the values, and click **OK**.

   See "Set Extended Attributes panel options" on page 556.
6. In the **Set Extended Attributes - Result** panel click **Close**.

Chapter 33

Managing policy checks

This chapter includes the following topics:

- About policy checks
- How signature registration settings work
- Registering policy signatures
- Unregistering a signature
- Setting signature tunables
- Running a manual policy scan
- Enabling or disabling policy signatures
- Viewing policy violation details
- Viewing or exporting a list of available policy signatures
- About using custom signatures for policy checks

About policy checks

The Veritas InfoScale Operations Manager policy check uses individual rules, called policy signatures, to validate whether the data center configuration conforms to a predefined standard. Veritas InfoScale Operations Manager contains a set of predefined signatures. You can also create customized signatures.

Policy check scans identify the resources in the Veritas InfoScale Operations Manager environment that do not meet the specified standards. Violations can be tracked over time. The policy check helps you better identify and address risks in your environment by providing notification of error and risk conditions on a proactive basis, before they evolve into something more severe.

You select the signatures to be scanned and a schedule for when to run the scan. This process is called registering signatures. You can also run scans on demand. Once the scan is performed, you can view the violation results. You can enable or disable the signatures temporarily.

The time that is required for the scan depends on the number of hosts against which the policy checks are run and the number of signatures run.

You cannot register signatures on ESX servers and non-global zones because the VRTSsfmh package is not installed on them.

See "How signature registration settings work" on page 560.

See "Registering policy signatures" on page 561.

See "Unregistering a signature" on page 562.

See "Setting signature tunables" on page 563.

See "Running a manual policy scan" on page 564.

See "Enabling or disabling policy signatures" on page 564.

See "Viewing policy violation details" on page 565.

See "Viewing or exporting a list of available policy signatures" on page 567.

See "About using custom signatures for policy checks" on page 567.

How signature registration settings work

Signature registration is the process of selecting which policy signatures are to be scanned and optionally specifying a schedule for when to run the scan.

You can determine the scope of the signature scan, that is, which hosts are scheduled to be scanned for the selected signatures, depending on where you register the signatures, as follows:

- From the **Server** perspective, you can register signatures by Organization and by individual host.
- From the **Availability** perspective, you can register signatures by Organization and by individual cluster.

Registering signatures by Organization helps ensure policy compliance within the Organization and is also more convenient.

It is possible to register signatures using different settings for a host or a cluster than for the rest of the Organization to which it belongs. The most recent registration settings take precedence.

Therefore, the recommended sequence is to first register the signatures at the Organization level. Then, as needed, you can modify signature settings at the lower (host or cluster) level. Note that if you later modify the signature settings at the Organization level, the Organization settings will override any settings that were modified earlier at the lower level.

Some signatures have tunable parameters. After registering signatures, you can view the list of registered signatures to identify any that have tunable parameters, and then set the values for the tunables. As with signature registration, you can apply tunable parameters at the Organization level or at the host or cluster level, and the most recent settings take precedence.

Registered signatures are enabled by default. You can temporarily disable signatures so as to exclude them from a scan. As with other settings, you can enable or disable signatures at the Organization level or at the host or cluster level, and the most recent settings take precedence. With the appropriate permission to modify settings for Veritas InfoScale Operations Manager, you can also enable or disable policy signatures for the data center, and again, the most recent settings take precedence.

See “Registering policy signatures” on page 561.

See “About policy checks” on page 559.

Registering policy signatures

The Management Server console lets you select which policy signatures are to be scanned and optionally specify a schedule for when to run the scan. This process is called signature registration.

You determine the scope of the signature scan, that is, which hosts are scheduled to be scanned for the selected signatures, depending on where you register the signatures, as follows:

- From the **Server** perspective, you can register signatures by Organization and by individual host.
- From the **Availability** perspective, you can register signatures by Organization and by individual cluster.

To perform this task on the Server perspective, your user group must be assigned the Admin role on the host or the Server perspective. The permission on the host may be explicitly assigned or inherited from a parent Organization. To perform this task on the Availability perspective, your user group must be assigned the Admin role on the cluster or the Availability perspective. The permission on the cluster may be explicitly assigned or inherited from a parent Organization.

To register policy signatures

1 In the Management Server console, choose from the following:

Server perspective	Select an Organization or a host
Availability perspective	Select an Organization or a cluster

2 Right-click the selected object and click **Policy Checks > Register Policy Signatures**.

This menu option is also available from the **Registered Signatures** tab.

3 In the **Select Signatures** panel, select the check boxes next to the available signatures that you want to register and click **Next**.

4 To specify a schedule for the scan, select **Select schedule** and specify the schedule details. Click **Next**.

5 On the **Summary** panel, review the information. If you do not want to run the signature check immediately, deselect that option. Click **Finish**.

The signature is listed on the **Registered Signatures** tab as enabled.

If a signature has tunable parameters (default settings that you can modify), they are listed on the table and you can modify them.

See "Setting signature tunables" on page 563.

See "About policy checks" on page 559.

See "How signature registration settings work" on page 560.

Unregistering a signature

In the Management Server console, you can unregister a policy signature. After you unregister a signature, it is not included in the future policy scans for the selected object.

To perform this task on the Server perspective, your user group must be assigned the Admin role on the host or the Server perspective. The permission on the host may be explicitly assigned or inherited from a parent Organization. To perform this task on the Availability perspective, your user group must be assigned the Admin role on the cluster or the Availability perspective. The permission on the cluster may be explicitly assigned or inherited from a parent Organization.

To unregister a signature

1 In the Management Server console, choose from the following:

Server perspective	Select an Organization or a host
Availability perspective	Select an Organization or a cluster

2 Click the **Registered Signatures** tab.

3 Right-click the signature that you want to remove, and then click **Unregister Signatures**.

4 Confirm that you want to unregister the selected signature.

See "About policy checks" on page 559.

Setting signature tunables

In the Management Server console, some policy check signatures may have tunable parameters, which are default settings that you can modify. You can modify tunables after you register a signature for use in policy check scans. The **Tunable Defaults** column on the **Registered Signatures** tab shows the default tunable settings if a signature has any associated tunables.

You can set signature tunables at the Organization level or at the host or cluster level. If you set tunables for a host or cluster and then later set them for the Organization, the most recent settings take precedence.

To perform this task on the Server perspective, your user group must be assigned the Admin role on the host or the Server perspective. The permission on the host may be explicitly assigned or inherited from a parent Organization. To perform this task on the Availability perspective, your user group must be assigned the Admin role on the cluster or the Availability perspective. The permission on the cluster may be explicitly assigned or inherited from a parent Organization.

To set signature tunables

1 In the Management Server console, choose from the following:

Server perspective	Select an Organization or a host
Availability perspective	Select an Organization or a cluster

2 Click the **Registered Signatures** tab.

3 Right-click the row for a signature that has tunables that you want to modify and click **Set Tunables**.

4 In the **Set Policy Signature Tunables** panel, enter the new tunable values and click **OK**.

5 In the **Result** panel, verify that the operation was successful.

See "About policy checks" on page 559.

Running a manual policy scan

In the Management Server console, you can run a policy scan manually.

The policy scan applies to the selected host on the **Server** perspective or to the selected cluster on the **Availability** perspective. You can also run a scan for an Organization.

To perform this task on the Server perspective, your user group must be assigned the Admin role on the host or the Server perspective. The permission on the host may be explicitly assigned or inherited from a parent Organization. To perform this task on the Availability perspective, your user group must be assigned the Admin role on the cluster or the Availability perspective. The permission on the cluster may be explicitly assigned or inherited from a parent Organization.

To run a manual policy scan

1 In the Management Server console, choose from the following:

Server perspective	Select an Organization or a host
Availability perspective	Select an Organization or a cluster

2 Right-click the selected object and click **Policy Checks > Scan Signatures Now**.

3 Confirm starting the scan.

See "About policy checks" on page 559.

Enabling or disabling policy signatures

In the Management Server console, you can enable or disable policy signatures. When you register policy signatures, they are enabled by default.

You can enable or disable individual signatures for individual hosts or clusters or by Organization. The most recent settings take precedence.

To perform this task on the Server perspective, your user group must be assigned the Admin role on the host or the Server perspective. The permission on the host may be explicitly assigned or inherited from a parent Organization. To perform this

task on the Availability perspective, your user group must be assigned the Admin role on the cluster or the Availability perspective. The permission on the cluster may be explicitly assigned or inherited from a parent Organization.

You can also enable or disable policy signatures for the data center in the Management Server perspective.

See the *Veritas InfoScale Operations Manager Management Server Installation and Configuration Guide*.

To enable policy signatures

1. In the Management Server console, choose from the following:

Server perspective	Select an Organization or a host
Availability perspective	Select an Organization or a cluster

2. Click the **Registered Signatures** tab.
3. Select one or more signatures to enable.
4. Right-click the signatures and click **Enable Signatures**.
5. Confirm whether to enable the signatures.

To disable policy signatures

1. In the Management Server console, choose from the following:

Server perspective	Select an Organization or a host
Availability perspective	Select an Organization or a cluster

2. Click the **Registered Signatures** tab.
3. Select one or more signatures to disable.
4. Right-click the signatures and click **Disable Signatures**.
5. Confirm whether to disable the signatures.

See “About policy checks” on page 559.

Viewing policy violation details

A policy violation occurs when a configuration option on a managed host does not meet the condition that a registered policy signature sets for it. The violation is discovered and reported on by a policy signature scan.

The Management Server console displays detailed information about the violations that are discovered by policy signature scans. You can view the policy violations for a selected host on the **Server** perspective or for a selected cluster on the **Availability** perspective. You can also view policy violations for an Organization.

The violations are grouped in the following categories.

Table 33-1 Categories of policy violations

Category	Description
All	A list that shows both the new and recurring violations, based on the previous five scans.
New	New policy violations in the current scan that were not part of the previous five scans.
Recurring	Violations that occurred in one or more of the previous five scans and are continuing to happen in the current scan.
Resolved	Policy violations that were found in the previous five scans that are no longer occurring.

The **Overview** tab includes information on the signature names, the signature categories (for example, the platform), the number of errors, the number of warnings, and the total number of violations.

The details include the signature that was violated, the object on which the violation occurred, and the details of the violation.

Note: A Policy Signature Scan Summary report is available in the Reports view.

Policy violations also generate a fault for the host. The data center **Faults** tab provides a central view to manage and monitor all faults of the corresponding perspective, including policy violation faults. For the faults that are generated on Management Server, or on the managed host versions lower than 6.1, the email notifications and faults may be displayed in the console within 30 minutes.

You can also set up a rule to receive offline notification of policy violations.

See “Creating rules in a perspective” on page 103.

You can view this information related to the hosts (on the Server perspective) or the clusters (on the Availability perspective) for which your user group has at least Guest role explicitly assigned or inherited from a parent Organization. You can also view information if your user group has at least Guest role assigned on the perspective.

To view policy violation details

1 Choose one of the following:

Server perspective	Select an Organization or a host
Availability perspective	Select an Organization or a cluster

2 Right-click the selected object and click **Policy Checks > Show Violation Details**.

3 Review the overview information on the **Overview** tab. Choose from the other tabs for details.

See "About policy checks" on page 559.

Viewing or exporting a list of available policy signatures

In the Management Server console, you can view all policy signatures and export them to a file. You must have the appropriate permissions for the Veritas InfoScale Operations Manager **Settings** perspective.

You can view this information, if your user group has Admin role assigned on the Management Server perspective.

To view or export a list of available policy signatures

1 In the Home page on the Management Server console, click **Settings**.

2 Click **Policy Signatures**.

3 If you want to export the signatures to a file, click the **Save** icon in the upper right of the window.

See "About policy checks" on page 559.

About using custom signatures for policy checks

Veritas InfoScale Operations Manager lets you add custom signatures for policy checks. You can use the custom signatures the same way that you use the predefined signatures that are installed with Veritas InfoScale Operations Manager. You can schedule the custom signatures for scanning on specified sets of hosts.

You must copy the custom signature script to Management Server and run the `pcreg` utility to install it. Then, you must copy the custom signature script to the managed hosts where you want to run it. Also, you can use the Distribution Manager

Add-on to bundle one or more custom signature scripts in a custom solution, for upload to the Management Server database and deployment to the managed hosts.

See “About policy checks” on page 559.

See “Creating a custom signature script” on page 568.

See “Installing a custom signature script” on page 571.

See “Copying the custom signature script” on page 571.

See “Removing a custom signature” on page 572.

Creating a custom signature script

You can create a custom signature script for policy checks in Veritas InfoScale Operations Manager. When creating the script you can use any programming language. Also, the custom signature script can have any type of code.

You must ensure the following when you create the custom signature script:

- Script must return 0 for success.
- Script must return a value other than 0 for errors.

In case of error, the script must print as follows:

Table 33-2 Script output when an error occurs

Item	Value
RESULT_RC	Any of the following exit codes: ■ **0** - Check passed. ■ **1** - Check fails with warning. ■ **2** - Check fails with error. ■ **100** - Not supported or cannot run check.
RESULT_TXT	Displays the failure message.
RESULT_REMEDY	Displays the message that helps resolve the problem.

The custom signature script must support the --what command line argument, which ensures the following output when you use the `custom signature script.pl --what` command:

Table 33-3 Output of the command: `custom signature script.pl --what`

Item	Value
CHECK_ID	Displays a unique value. This value can be a number or a string. The CHECK_ID and the name of the script file need to be the same for the custom signature to work properly. Example: PC_VXVM_UNUSED_VOLUMES
CHECK_NAME	Displays the name of the custom signature.
CHECK_DESCRIPTION	Displays the description for the signature.
CHECK_KEYWORDS	Displays the category to which the custom signature belongs. Example: UTILIZATION,VXVM,PERFORMANCE,VCS

Note: The custom signature must return the code 0 on successful installation on Management Server using the `pcreg` utility. It must not give any other output.

See "Sample custom signature script" on page 569.
See "About using custom signatures for policy checks" on page 567.
See "Installing a custom signature script" on page 571.
See "Copying the custom signature script" on page 571.
See "Removing a custom signature" on page 572.
See "About policy checks" on page 559.

Sample custom signature script

The following is an example of a custom signature script for policy checks in Veritas InfoScale Operations Manager.

Note: The following sample script is created using PERL.

```
#!/opt/VRTSsfmh/bin/perl
BEGIN { @INC = ("/opt/VRTSsfmh/lib/modules"); }
use strict;
```

```
if ($ARGV[0] eq "--what")
{
    print "CHECK_ID: PC_VXVM_UNUSED_VOLUMES\n";
    print "CHECK_NAME: VxVM Unused Volumes\n";
    print "CHECK_DESCRIPTION: Check to see if any VxVM volumes are unused.\n";
    print "CHECK_KEYWORDS: UTILIZATION,UNIX,VXVM\n";
    exit (0);
}

if($^O =~ /Win32/i)
{
    print "RESULT_RC=100\n";
    print "RESULT_TXT=Windows host detected\n";
    print "RESULT_REMEDY=This check is intended only for unix hosts\n";\
    exit(100);
}

my $exitcode = 0;
my @items = `/opt/VRTSsfmh/bin/vxlist -t stats vol`;
my $iocount = {};

foreach (@items) {
  if(/^vol\s+(\S+)\s+(\S+)\s+(read|write)\s+(\S+)\s+/) {
    my $name = $2."/".$1;
    $iocount->{$name} += $4;
  }
}
foreach (keys %$iocount) {
  if($iocount->{$_} == 0) {
     /(\S+)\/(\S+)/;
     print "RESULT_RC=2\n";
     print "RESULT_TXT=VxVM volume $2 in disk group $1 may be unused.
            It has performed no I/O since last reboot.\n";
     print "RESULT_REMEDY=There may be data on the volume, such as an unmounted file system.
            If not, space can be reclaimed by deleting the volume,\n";
     $exitcode = 2;
  }
}

exit($exitcode);
```

See “Creating a custom signature script” on page 568.

Installing a custom signature script

After copying a custom signature script to Management Server, you must install the custom signature. Ensure that the script is in the executable format.

To install a custom signature on Management Server, you can use the `pcreg` utility, which is available in the following locations:

- On UNIX/Linux-based Management Server - /opt/VRTSsfmh/bin
- On Windows-based Management Server - C:\Program Files\VERITAS\VRTSsfmh\bin

Use the following syntax for running the utility to install a custom signature:

- `pcreg --script` *`custom signature script`*

Note: Provide the full path to the custom signature script.

After you install the custom signature, you must copy the custom signature script to the managed hosts where you want to register it.

See "About using custom signatures for policy checks" on page 567.

Copying the custom signature script

You must ensure that the custom policy check signature exists on the managed host where you want it to run it.

You must copy the custom script to the following locations:

- On UNIX/Linux-based managed host - /opt/VRTSsfmh/lib/signatures
- On Windows-based managed host - C:\Program Files\VERITAS\VRTSsfmh\lib\signatures

You can copy the custom signature to the managed hosts manually. Also, you can use the Distribution Manager Add-on to push the custom signatures to the required managed hosts.

See "About using custom signatures for policy checks" on page 567.

See "Creating a custom signature script" on page 568.

See "Installing a custom signature script" on page 571.

See "Removing a custom signature" on page 572.

Removing a custom signature

To remove a custom policy check signature from Veritas InfoScale Operations Manager, you can run the following utility:

```
pcreg --id policy check id
```

See "About using custom signatures for policy checks" on page 567.

About using the Distribution Manager Add-on to bundle custom signature scripts

Using the Distribution Manager Add-on, you can bundle one or more custom policy check signature scripts, along with the set up and the unsetup scripts, to an `.sfa` archive. The setup script can upload the custom signature scripts at a specific location on Management Server and copy the custom signature scripts at a specific location on the managed host. Once the add-on is created for the custom signature, you can deploy it using the Deployment Management feature. The unsetup script of the add-on can remove the custom signature from Management Server and from managed hosts.

For information on using the Distribution Manager Add-on, see *Veritas InfoScale Operations Manager Add-ons User Guide*.

See "Sample setup.pl script for the custom signature" on page 572.

See "Sample unsetup.pl script for a custom signature" on page 573.

See "About using custom signatures for policy checks" on page 567.

Sample setup.pl script for the custom signature

The following is the sample `setup.pl` script for the custom signature:

```
#!/opt/VRTSsfmh/bin/perl

BEGIN { @INC = ("/opt/VRTSsfmh/lib/modules"); }

use VRTS::AddOnInfo;
use File::Copy;
require File::Spec;

my $store_dir = $ENV{'Store'};
my $install_dir = VRTS::Paths::get_path("InstallDir");
my $pcregutil = File::Spec->catfile($install_dir,"bin","pcreg");
my $sigdir = File::Spec->catfile($install_dir,"lib","signatures");
```

```
# Creating array of signature files name
@signature_files = ("UnUsedVolumes.pl");

# Copy file under signatures directory
foreach $sigfile (@signature_files)
{
        print "Copying script $sigfile to $sigdir\n";
        copy($sigfile,$sigdir);

        print "Making script executable\n";
  my $pc = File::Spec->catfile($sigdir,$sigfile);
        system("chmod +x $pc");

        # Upload signature only on CS
  if ( VRTS::AddOnInfo::is_CS() )
        {
   print "Uploading script\n";
   my $cmd1 = VRTS::Util::make_command($pcregutil,"--script",$pc);
   system($cmd1);
        }
}
exit(0);
```

Sample unsetup.pl script for a custom signature

The following is the sample `unsetup.pl` script for a custom signature:

```
#!/opt/VRTSsfmh/bin/perl

BEGIN { @INC = ("/opt/VRTSsfmh/lib/modules"); }

use VRTS::AddOnInfo;
use File::Copy;
require File::Spec;

my $install_dir = VRTS::Paths::get_path("InstallDir");
my $sigdir = File::Spec->catfile($install_dir,"lib","signatures");
my $pcregutil = File::Spec->catfile($install_dir,"bin","pcreg");

# Creating array of signature files name and signature name
@signature_files = ("UnUsedVolumes.pl");
@signature_names = ("PC_VXVM_UNUSED_VOLUMES");
```

```
# Remove script on CS
if ( VRTS::AddOnInfo::is_CS() )
{
 foreach $signame (@signature_names)
 {
  print "Remove signature $signame\n";
  system("$pcregutil --id $signame");
 }
}

# Remove signature from signature directory
foreach $sigfilename (@signature_files)
{
 print "Remove script $sigfilename\n";
 my $sigfile = File::Spec->catfile($sigdir,$sigfilename);
 File::Path::rmtree("$sigfile");
}
exit(0);
```

Chapter 34

Managing Dynamic Multipathing paths

This chapter includes the following topics:

- About Dynamic Multi-Pathing in Veritas InfoScale Operations Manager
- Disabling the DMP paths on the initiators of a host
- Disabling the DMP paths on an enclosure
- Disabling the DMP paths on a virtualization server
- Re-enabling the DMP paths
- Removing a completed DMP maintenance case record
- Reviewing the output and results of a completed DMP maintenance case

About Dynamic Multi-Pathing in Veritas InfoScale Operations Manager

Dynamic Multi-Pathing lets you direct the communication to the storage through different sets of nodes to achieve failover and load balancing. When a connection fails, Dynamic Multi-Pathing spreads the communication to alternate paths automatically.

In certain circumstances, you may want to perform maintenance operations on the array ports or on the initiators in your data center. For example, upgrade the firmware that has been installed on one or more initiators. You need to temporarily disable the multipathing paths that connect these array ports and the initiators to the storage, and re-enable these paths after the maintenance operations are completed. Veritas

InfoScale Operations Manager provides you the options of disabling the multipathing paths by choosing an array, virtualization server, or initiators in your data center.

In the Management Server console, dynamic multipathing operations and discovery of multipathing paths are supported on the following:

- Physical hosts that have Dynamic Multi-Pathing installed.
- VMware ESX servers that have Dynamic Multi-Pathing for VMware installed.

Veritas InfoScale Operations Manager also supports the discovery of EMC PowerPath/Virtual Edition (VE) of ESX Server, Microsoft MPIO, and HP Native multipathing. For these technologies, Veritas InfoScale Operations Manager discovers the multipathing information that includes the paths under the metanodes, path count, active_path count, and the state of each path.

Before you disable the paths, you must ensure that the connected resources have alternate paths to the underlying physical storage. Veritas InfoScale Operations Manager provides you the options to verify the associated volumes, disk groups, hosts, and the applications that currently use the storage that is connected through the selected initiators or array ports. You need to have administrative privileges on host, enclosure, or on virtualization server to perform these path maintenance operations.

Note: You cannot use Veritas InfoScale Operations Manager to manage the paths for Windows managed hosts.

See "Disabling the DMP paths on the initiators of a host" on page 576.

See "Disabling the DMP paths on an enclosure" on page 582.

See "Disabling the DMP paths on a virtualization server" on page 583.

Disabling the DMP paths on the initiators of a host

You can use the Management Server console to select host and disable all paths of one or more initiators of the host.

To disable the DMP paths on the initiators of a host

1 In the Management Server console, do one of the following:

- In the **Server** perspective select the required host, right-click and then select **DMP Path Management**.
- In the **Server** perspective select the host, and under the **Initiators** tab, select the required initiator. Right-click the initiator and then select **DMP Path Management**.

- In the **Server** perspective, go to the **Solutions** section, click **DMP Path Management**, and under **Actions**, click **Start DMP Path Management**.

2 In the **Define New Case** panel, enter the required information. Click **Next**.

See "Define new case panel options" on page 577.

3 In the next panel, select one or more initiators that contain the DMP paths that you want to disable.

See "Object selection panel options" on page 579.

4 In the **Path Disable Confirmation** panel, review the details of the DMP paths. To view the information on the relationship of the DMP paths with other storage resources in your data center, click **View Associations Report**. Click **Disable** to disable all the listed DMP paths.

See "Path disable panel options" on page 581.

5 In the **Paths Disabled** panel, review the details of the disabled paths and the commands that were run to complete the operation. Do one of the following:

- To save and close the configuration and re-enable the DMP paths later, click **Save and Close**.
  See "Paths disable output summary" on page 581.
- To re-enable the DMP paths after the maintenance operations are completed, click **Next**. The wizard prompts you to re-enable the DMP paths.
  See "Re-enabling the DMP paths " on page 584.

See "About Dynamic Multi-Pathing in Veritas InfoScale Operations Manager" on page 575.

Define new case panel options

Use this panel to start a new dynamic multipathing maintenance case. The panel lets you specify a name for the new DMP maintenance case that you want to start.

Table 34-1

Lists the options related to creating a new DMP case for host.

Table 34-2

Lists the options related to creating a new DMP case for enclosure.

Table 34-3

Lists the options related to creating a new DMP case for virtualization server.

Table 34-1 Define new case panel options for the Host

Field	Description
Define New Case	
Case Name	Name of the new DMP maintenance case that you want to start.
Description	The additional information that you want to include for the new DMP maintenance case. This field is optional.
Select the hosts that contain the initiators on which you want to perform maintenance. If you have selected an individual host, the following options are not applicable. These options are displayed only when the DMP maintenance case is started from the **Solutions** section of the Server perspective.	
Name	The name of the host.
State	The state of the host.
IP address	The IP address of the host.
Architecture	The architecture of the host, such as PowerPC and SPARC.
Platform	The operating system that the host uses.

See "Disabling the DMP paths on the initiators of a host" on page 576.

Table 34-2 Define new case panel options for the enclosure

Field	Description
New Case options	
Case Name	Name of the new DMP maintenance case that you want to start.
Description	The additional information that you want to include for the new DMP maintenance case. This field is optional.
Select the enclosure that contains one or more array ports on which you want to perform maintenance. If you have selected an individual enclosure, the following options are not applicable. These options are displayed only when the DMP maintenance case is started from the **Solutions** section of Storage perspective.	
Name	The name of the enclosure.

Table 34-2 Define new case panel options for the enclosure *(continued)*

Field	Description
Serial	The serial number of the enclosure.
Vendor	The name of the manufacturer of the enclosure.
Product	The name of the array model.
Type	Type of array.

See "Disabling the DMP paths on an enclosure" on page 582.

Table 34-3 Define new case panel options for the virtualization server

Field	Description
New Case options	
Case Name	Name of the new DMP maintenance case that you want to start.
Description	The additional information that you want to include for the new DMP maintenance case. This field is optional.
Select the hosts that contain the initiators on which you want to perform maintenance. If you have selected an individual host, the following options are not applicable. These options are displayed only when the DMP maintenance case is started from the **Solutions** section of Virtualization perspective.	
Name	The name of the host.
State	The state of the host.
IP address	The IP address of the host.
Architecture	The architecture of the host, such as PowerPC and SPARC.
Platform	The operating system that the host uses.

See "Disabling the DMP paths on a virtualization server" on page 583.

Object selection panel options

This panel lets you select the initiators (for host and virtualization server), and array ports or adapters (for enclosure) to perform DMP path maintenance operations.

Note: For enclosure, if Storage Insight Add-on is installed, array ports are displayed. In absence of Storage Insight Add-on, array port World Wide Name (WWN) is listed.

Table 34-4

Lists the initiators for DMP path management operation.

Table 34-5

Lists the array ports for DMP path management operation.

Table 34-6

Lists the initiators of a virtualization server for DMP path management operation.

Table 34-4 Select Initiators panel options

Field	Description
Name	The name of the initiator.
Host	The host to which the initiator is associated.
Driver Version	The version of the driver for the initiator.
Firmware Version	The version of the firmware that is used on the initiator.
Vendor	The manufacturer of the initiator.

See "Disabling the DMP paths on the initiators of a host" on page 576.

Table 34-5 Select array port panel options

Field	Description
Name	The name of the array port
Array Port WWN	The World Wide Name for the array port.

See "Disabling the DMP paths on an enclosure" on page 582.

Table 34-6 Select Initiators panel options for virtualization server

Field	Description
Name	The name of the initiator.
Host	The host to which the initiator is associated.
Driver Version	The version of the driver for the initiator.

Table 34-6 Select Initiators panel options for virtualization server *(continued)*

Field	Description
Firmware Version	The version of the firmware that is used on the initiator.
Vendor	The manufacturer of the initiator.

See "Disabling the DMP paths on a virtualization server" on page 583.

Path disable panel options

Use this panel to view the details of the DMP paths that you have selected to disable. You can click the **View Association Report** to view the information on the relationship of the DMP paths that you want to disable with other storage resources. If one or more paths are not supported for the operation, they are listed separately. Click **View Paths** to view the details of such paths.

Table 34-7 Path disable panel options

Field	Description
Name	The names of the DMP paths that you have selected for disabling.
Status	Current status of the DMP Paths that you have selected for disabling.
Devices	The name of the disk where the path is located.
Active Paths	The number of the active DMP paths.
Host	Host to which the DMP Paths that you have selected to disable, is connected.
Initiator	Initiator to which the DMP path is associated.

See "Disabling the DMP paths on the initiators of a host" on page 576.

See "Disabling the DMP paths on an enclosure" on page 582.

See "Disabling the DMP paths on a virtualization server" on page 583.

Paths disable output summary

Use this panel to view the output summary of the DMP path disabling operation.

Table 34-8 Paths disable output summary panel options

Field	Description
Path Disable Operation Output Summary	■ Total number of successful DMP commands. ■ Total number of failed DMP commands.
Command Details	■ Hosts on which the successful DMP commands were run. ■ Details of the DMP paths that are disabled. ■ Details of the failed DMP commands and the reason for the failure.

See "Disabling the DMP paths on the initiators of a host" on page 576.

See "Disabling the DMP paths on an enclosure" on page 582.

See "Disabling the DMP paths on a virtualization server" on page 583.

Disabling the DMP paths on an enclosure

You can use the Management Server console to select an enclosure and disable all paths of one or more array ports of the enclosure.

To disable the DMP paths on an enclosure

1 In the Management Server console, do one of the following:

- In the **Storage** perspective select the required enclosure, right-click and then select **DMP Path Management**.
- In the **Storage** perspective, go to the **Solutions** section, click **DMP Path Management**, and under **Actions**, click **Start DMP Path Management.**

2 In the **Define New Case** panel, enter the required information. Click **Next**.

See "Define new case panel options" on page 577.

3 In the **Specify Array Ports** panel, select one or more array ports to perform DMP maintenance.

See "Object selection panel options" on page 579.

4 In the **Path Disable Confirmation** panel, review the details of the DMP paths. To view the information on the relationship of the DMP paths with other storage resources in your data center, click **View Associations Report**. Click **Disable** to disable all the listed DMP paths.

See "Path disable panel options" on page 581.

5 In the **Paths Disabled** panel, review the details of the disabled paths and the commands that were run to complete the operation. Do one of the following:

- To save and close the configuration and re-enable the DMP paths later, click **Save and Close**.
  See "Paths disable output summary" on page 581.
- To re-enable the DMP paths after the maintenance operations are completed, click **Next**. The wizard prompts you to re-enable DMP paths.
  See "Re-enabling the DMP paths " on page 584.

See "About Dynamic Multi-Pathing in Veritas InfoScale Operations Manager" on page 575.

Disabling the DMP paths on a virtualization server

You can use the Management Server console to select a virtualization server and disable all paths of one or more of its initiators.

To disable the DMP paths on a virtualization server

1 In the Management Server console, do one of the following:

- In the **Virtualization** perspective select the required host, right-click and then select **DMP Path Management**.
- In the **Virtualization** perspective, go to the **Solutions** section, click **DMP Path Management**, and under **Actions**, click **Start DMP Path Management**.

2 In the **Define New Case** panel, enter the required information. Click **Next**.

See "Define new case panel options" on page 577.

3 In the next panel, select one or more initiators that contain the DMP paths that you want to disable.

See "Object selection panel options" on page 579.

4 In the **Path Disable Confirmation** panel, review the details of the DMP paths. To view the information on the relationship of the DMP paths with other storage resources in your data center, click **View Associations Report**. Click **Disable** to disable all the listed DMP paths.

See "Path disable panel options" on page 581.

5 In the Paths Disabled panel, review the details of the disabled paths and the commands that were run to complete the operation. Do one of the following:

- To save and close the configuration and re-enable the DMP paths later, click **Save and Close**.
  See "Paths disable output summary" on page 581.
- To re-enable the DMP paths after the maintenance operations are completed, click **Next**. The wizard prompts you to re-enable DMP paths.
  See "Re-enabling the DMP paths " on page 584.

See "About Dynamic Multi-Pathing in Veritas InfoScale Operations Manager" on page 575.

Re-enabling the DMP paths

You can re-enable the DMP paths for the initiator, enclosure, or the virtualization server that you have disabled as a part of the maintenance operation. You can re-enable the DMP paths in the following ways:

- Immediately after disabling the DMP paths. Use the same wizard panel to re-enable the disabled paths.
- Using the **Resume Case and Re-Enable Paths** options for the DMP case under the **Waiting Cases** section on the **Solution** section. For initiator, use **Solution** section under the Server perspective, for enclosure, use **Solution** section in the Storage perspective, and for virtualization server, use **Solution** section in the Virtualization perspective.

To re-enable the DMP paths immediately after disabling them

1 After you disable the DMP paths on an array, initiator, or a virtualization server, click **Next**.

See "Paths disable output summary" on page 581.

2 In the Paths Re-Enable Confirmation panel, click **Re-Enable**.

See "Path re-enable panel options" on page 585.

3 In the **Path Enable Operation Output Summary** panel, review the details of the re-enabled DMP paths.

See "Paths re-enable output summary" on page 586.

4 In the **DMP Maintenance Result Summary** page, view the details of the current paths and the commands run to complete the operation. Click **Close**.

See "DMP Maintenance Result Summary panel" on page 586.

To re-enable the DMP paths using the Solutions page

1 In the Management Server console, do the following:

- For enclosure, go to the **Solutions** page of the Storage perspective.
- For initiator, go to the **Solutions** page of the Server perspective.
- For virtualization server, go to the **Solutions** page of the Virtualization perspective.

2 Under the **Waiting Cases** section, select the required case. Right-click and then select **Resume Case and Re-Enable Paths**.

3 In the Paths Re-Enable Confirmation panel, click **Re-Enable**.

See "Path re-enable panel options" on page 585.

4 In the **Path Enable Operation Output Summary** panel, review the details of the re-enabled DMP paths.

See "Paths re-enable output summary" on page 586.

5 In the **DMP Maintenance Result Summary** page, view the details of the current paths and the commands run to complete the operation. Click **Close**.

See "DMP Maintenance Result Summary panel" on page 586.

See "About Dynamic Multi-Pathing in Veritas InfoScale Operations Manager" on page 575.

Path re-enable panel options

Use this wizard panel to view the details of the DMP paths that you have selected to re-enable.

Table 34-9 Path re-enable confirmation panel options

Field	Description
Name	Names of the DMP paths that you have selected.
Devices	Name of the disk where the path is located.

Table 34-9 Path re-enable confirmation panel options *(continued)*

Field	Description
Active Paths	The number of the active DMP paths.
Host	Host to which the DMP Paths that you have selected for enabling is connected.
Initiator	Initiator to which the DMP path is associated.

See "Re-enabling the DMP paths " on page 584.

Paths re-enable output summary

Use this wizard panel to view the output summary of the DMP path re-enabling operation. This wizard panel displays the following information:

Table 34-10 Paths re-enable output summary panel options

Field	Description
Path Enable Operation Output Summary	■ Total number of successful commands. ■ Total number of failed DMP commands.
Command Details	■ Hosts on which the successful DMP commands were run ■ Details of the DMP paths that are re-enabled ■ Details of the failed DMP commands and the reason for the failure.

See "Re-enabling the DMP paths " on page 584.

DMP Maintenance Result Summary panel

Use this wizard panel to view the overall summary of the Dynamic Multipathing (DMP) maintenance case that you have completed.

This panel lists the following details:

- DMP paths that were managed and their current statuses.
- All commands that are run during the DMP maintenance operation and their results.

Table 34-11 DMP Maintenance Result Summary panel options

Field	Description
Current Path Status	
Name	Name of the DMP path on which you have completed the state management.
Status	The current status of the DMP path.
Device	The disk device to which the DMP path is associated.
Active Paths	Number of active paths.
Host	Host to which the DMP path is associated.
Initiator	The initiator to which the DMP path is associated.
Array port	Array port to which the DMP path is associated.
All Commands Executed and Results	
Disable Path Commands Executed	Displays the names of the commands that are run for disabling the DMP paths.
Enable Path Commands Executed	Displays the name of the commands that are run for enabling the DMP paths.

See “Re-enabling the DMP paths ” on page 584.

Removing a completed DMP maintenance case record

The Veritas InfoScale Operations Manager database keeps the records of each of the completed DMP maintenance cases. If you do not want to keep this record for a longer period of time, you can remove them from the database using the Veritas InfoScale Operations Manager console.

To remove a completed DMP maintenance case record

1 In the Management Server console, do the following:

- For enclosure, go to the **Solutions** page of the Storage perspective.
- For host, go to the **Solutions** page of the Server perspective.

- For virtualization server, go to the **Solutions** page of the Virtualization perspective.

2 Under the **Completed Cases** section, select the required case. Right-click and then select **Delete Case**.

3 On the **Delete Completed Case** panel, click **Yes** to confirm.

See "About Dynamic Multi-Pathing in Veritas InfoScale Operations Manager" on page 575.

Reviewing the output and results of a completed DMP maintenance case

Using the Veritas InfoScale Operations Manager console, you can view the status of the DMP path and the list of commands that were run as part of the DMP maintenance process.

To review the output and results of a completed DMP maintenance case

1 In the Management Server console, do the following:

- For enclosure, go to the **Solutions** page of the Storage perspective.
- For host, go to the **Solutions** page of the Server perspective.
- For virtualization server, go to the **Solutions** page of the Virtualization perspective.

2 Under the **Completed Cases** section, select the required case. Right-click and then select **Show Case Output and Results**.

3 In the **DMP Maintenance Result Summary** page, view the details of the current paths and the commands run to complete the operation.

See "DMP Maintenance Result Summary panel" on page 586.

See "About Dynamic Multi-Pathing in Veritas InfoScale Operations Manager" on page 575.

Chapter 35

Managing CVM clusters

This chapter includes the following topics:

- About monitoring and managing CVM clusters in Veritas InfoScale Operations Manager
- Permissions required for views and operations on CVM cluster objects

About monitoring and managing CVM clusters in Veritas InfoScale Operations Manager

Storage Foundation Cluster File System High Availability (SFCFSHA) includes the Cluster Volume Manager (CVM) as a component. CVM expands the functionality of the Storage Foundation volume manager (VxVM) to add support for a clustered environment. CVM enables the cluster nodes to simultaneously access and manage a set of disks or LUNs under Storage Foundation control.

The Server perspective of the Management Server console lets you monitor information about CVM clusters and supports Storage Foundation (SF) operations on cluster objects such as disk groups, volumes, and disks.

Note: Operations on SF disk groups of the cluster-shared type and operations on volumes on such disk groups are not available on a Windows CVM cluster.

Similar views are provided in the Veritas InfoScale Operations Manager Web services API.

Table 35-1 Monitoring and managing CVM clusters in the Server perspective

Operation or view	Support in the Server perspective
View all CVM clusters	From the storage clusters node in the tree you can display a list of all CVM clusters. The table columns provide top-level information such as whether the cluster supports Flexible Storage Sharing (FSS).
Search for CVM clusters	The global search feature includes clusters as an object. You can create and save search queries based on cluster attributes.
View detailed information about objects in each cluster	You can view detailed information about objects in the cluster including: ■ Hosts ■ All disks of all hosts in the cluster ■ Disk groups ■ Volumes and file systems ■ Applications
Use local or partially shared storage across all hosts in the cluster	In FSS-capable clusters and hosts, you can export or un-export disks from the disks view of the cluster. This operation is also available from the disks view of a host.
View the correlation of shared and exported storage across hosts in a cluster.	In the disks view of the cluster, you can group the information by disk to show which disks on different hosts are the same. For example, the view shows the exported disks and corresponding remote disks that are shared using the FSS export feature. The disks view of a host shows similar information on a sub-tab.
Create shared disk groups or perform disk group operations in a cluster.	You can create shared disk groups from either the storage clusters node in the tree, a selected cluster, or a selected cluster host. If the cluster supports FSS, you can enable FSS for the shared disk group and add exported disks. All disk group operations that are supported from a host view are also available from the cluster view.
Perform volume and file system operations in a cluster.	All volume and file system operations that are supported from a host view are also available from the cluster view.

See “Permissions required for views and operations on CVM cluster objects” on page 591.

See “About Flexible Storage Sharing” on page 594.

Permissions required for views and operations on CVM cluster objects

In the Server perspective of the Management Server console, the primary object for assigning roles and permissions is the host. Permissions can be assigned to the perspective, to Organizations of hosts, or to individual hosts.

The Server perspective includes separate views for CVM clusters. If your user group has at least Guest role assigned on at least one of the hosts in the cluster, you can view the information for the cluster and the hosts in it. The permission on the host may be explicitly assigned or inherited from a parent Organization. You can also view the information if your user group has at least Guest role assigned on the Server perspective.

You can perform operations by right-clicking CVM cluster objects (for example, disks and volumes) in the context of the cluster view, as well as from a host context.

To perform most Storage Foundation operations from the cluster context requires Admin rights to the host that is the master node of the selected CVM cluster. The permission on the host may be explicitly assigned or inherited from a parent Organization. The affected operations include disk group, disk, volume, and file system operations.

Note: It is recommended to assign Admin rights on all the cluster nodes since the master role can change.

- To perform the following operations from a host in the cluster view requires Admin rights on the concerned host:
  - Create disk group
  - Create volume
  - Create file system
  - Rescan disks
  - Refresh
  - Move to (Organization)
  - DMP path management

    - Create extended attributes
- To perform the following operations on disks in the cluster view requires Admin rights on the concerned host:
    - Reclaim thin storage
    - Disconnect
    - Offline
    - Online
    - Export
    - Unexport
    - Resize
    - Rename
    - Evacuate
    - Replace
    - Recover
    - Initialize
    - Set disk usage
    - Trim
    - Set extended attributes

See “About monitoring and managing CVM clusters in Veritas InfoScale Operations Manager ” on page 589.

Chapter 36

Managing Flexible Storage Sharing

This chapter includes the following topics:

- Implementing Flexible Storage Sharing with Veritas InfoScale Operations Manager
- About Flexible Storage Sharing
- Flexible Storage Sharing use cases
- Flexible Storage Sharing features and support in Veritas InfoScale Operations Manager
- Exporting and un-exporting disks for Flexible Storage Sharing
- Enabling or disabling Flexible Storage Sharing on existing shared disk groups

Implementing Flexible Storage Sharing with Veritas InfoScale Operations Manager

To implement Flexible Storage Sharing (FSS) with Veritas InfoScale Operations Manager use the following process.

Table 36-1 Implementing Flexible Storage Sharing with Veritas InfoScale Operations Manager

Task	For more information
Understand requirements and use cases for FSS	See "About Flexible Storage Sharing" on page 594. See "Flexible Storage Sharing use cases" on page 595. See "Flexible Storage Sharing features and support in Veritas InfoScale Operations Manager " on page 597.
Export disks for use in an FSS-capable cluster	See "Exporting and un-exporting disks for Flexible Storage Sharing" on page 598.
Create a shared disk group with FSS enabled and add exported disks Or Enable FSS on an existing shared disk group and add exported disks	See "Enabling or disabling Flexible Storage Sharing on existing shared disk groups" on page 599. See "Adding disks to disk groups" on page 159.

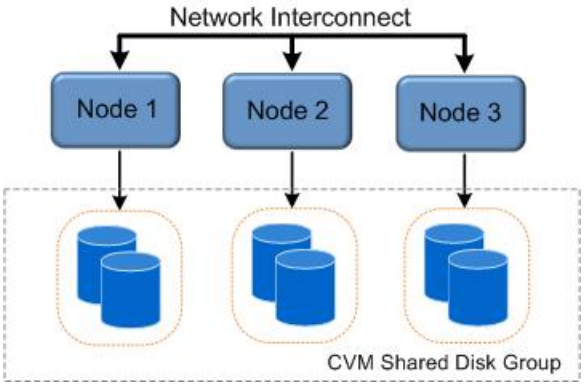
About Flexible Storage Sharing

Flexible Storage Sharing (FSS) is a feature of Storage Foundation Cluster File System High Availability (SFCFSHA) that enables network sharing of local storage, cluster wide. The local storage can be in the form of Direct Attached Storage (DAS) or internal disk drives. Network shared storage is enabled by using a network interconnect between the nodes of a cluster.

FSS allows network shared storage to co-exist with physically shared storage, and logical volumes can be created using both types of storage creating a common storage namespace. Logical volumes using network shared storage provide data redundancy, high availability, and disaster recovery capabilities, without requiring physically shared storage, transparently to file systems and applications.

Refer to the Storage Foundation and High Availability Solutions Software Compatibility List for information on FSS support by your SFCFSHA version and platform.

Figure 36-1 Flexible Storage Sharing Environment



See "Flexible Storage Sharing use cases" on page 595.

See "Flexible Storage Sharing features and support in Veritas InfoScale Operations Manager " on page 597.

See "Implementing Flexible Storage Sharing with Veritas InfoScale Operations Manager" on page 593.

Flexible Storage Sharing use cases

The following list includes several use cases for which you would want to use the FSS feature:

Use of local storage in current use cases	The FSS feature supports all current use cases of the Storage Foundation Cluster File System High Availability (SFCFSHA) stack without requiring SAN-based storage.

Off-host processing	Data Migration: ■ From shared (SAN) storage to network shared storage ■ From network shared storage to SAN storage ■ From storage connected to one node (DAS)/cluster to the storage connected to a different node (DAS)/cluster, that do not share the storage Back-up/Snapshots: An additional node can take a back-up by joining the cluster and reading from volumes/snapshots that are hosted on the DAS/shared storage, which is connected to one or more nodes of the cluster, but not the host taking the back-up.
DAS SSD benefits leveraged with existing SFCFSHA features	■ Mirroring across DAS SSDs connected to individual nodes of the cluster. DAS SSDs provides better performance than SAN storage (including SSDs). FSS provides a way to share these SSDs across cluster. ■ Keeping one mirror on the SSD and another on the SAN storage provides faster read access due to the SSDs, and also provide high availability of data due to the SAN storage. ■ There are several best practices for using SSDs with Storage Foundation. All the use-cases are possible with SAN attached SSDs in clustered environment. With FSS, DAS SSDs can also be used for similar purposes.
FSS with SmartIO for file system caching	If the nodes in the cluster have internal SSDs as well as HDDs, the HDDs can be shared over the network using FSS. You can use SmartIO to set up a read cache using the SSDs. The read cache can service volumes created using the network-shared HDDs.
Campus cluster configuration	Campus clusters can be set up without the need for Fibre Channel (FC) SAN connectivity between sites.

See “About Flexible Storage Sharing” on page 594.

Flexible Storage Sharing features and support in Veritas InfoScale Operations Manager

Flexible Storage Sharing (FSS) is a feature of Storage Foundation Cluster File System High Availability (SFCFSHA). FSS is supported in Veritas InfoScale Operations Manager as follows:

- Veritas InfoScale Operations Manager checks if Cluster Volume Manager (CVM) clusters are FSS capable and if the cluster nodes have the required managed host version (Veritas InfoScale Operations Manager 6.1 or later) in order to perform FSS operations from Veritas InfoScale Operations Manager.
- Veritas InfoScale Operations Manager discovers whether FSS is enabled for a disk group and the information is shown in the Management Server console for the disk group. Disk groups must be enabled for FSS before adding network shared disks.
- Veritas InfoScale Operations Manager discovers whether a disk is exported (network shared disk), remote, or neither and will show the FSS state for a disk in the console.

  Disk export is an FSS-specific operation performed on a disk that is locally visible to a host to make it visible to all hosts in the CVM cluster. A remote disk is a disk that has been exported from another host in the cluster.
- The Management Server console contains wizards for support of the FSS operations of enabling/disabling FSS for an existing disk group and exporting/unexporting disks for shared use in the CVM cluster. Wizards for existing SF operations support FSS operations as needed.
- CVM clusters (both FSS capable and non-FSS capable) are listed in the tree in the Server perspective. The cluster view in the Server perspective provides storage visualization views, including how each host in the cluster is using storage, how much is local, how much is remote, and so on. The Create Disk Group operation is available at the cluster level. Other SF operations are available from objects in the selected cluster.
- You can create queries for a Cluster object in the Search feature on the Server perspective. For example, you can search for clusters with the attribute FSS Capable.
- When you create a volume for a FSS-enabled disk group, you can specify to mirror the volume across hosts. To help validate that mirrored volumes are not based off local storage from a single host, you can register the Cross Enclosure/Host Mirroring policy signature for the hosts.

For more information on FSS requirements and limitations, see the SFCFSHA documentation.

See "About Flexible Storage Sharing" on page 594.

See "About monitoring and managing CVM clusters in Veritas InfoScale Operations Manager " on page 589.

Exporting and un-exporting disks for Flexible Storage Sharing

To use disks that are not physically shared across all nodes of a Cluster Volume Manager cluster, the disks must first be exported for network sharing. Exporting a device makes the device available to all nodes in the cluster. The export is required for SAN disks as well as local disks to make them fully shared with any new nodes that may join the cluster. Exported disks can be added only to a disk group that is enabled for Flexible Storage Sharing (FSS).

Boot disks, opaque disks, disks that are part of imported or deported disk groups, and non-Storage Foundation (VxVM) disks cannot be exported.

To perform this task, your user group must be assigned the Admin role on the host or the Server perspective. The permission on the host may be explicitly assigned or inherited from a parent Organization.

To export disks for Flexible Storage Sharing

1. In the Management Server console, go to the **Server** perspective and expand **Manage** in the left pane.
2. Expand **Clusters** and click the **Disks** tab.

   You can also perform this operation when viewing disks on a host.
3. Right-click the disk and select **Export**.

   You can also select multiple disks at a time for export.
4. Confirm that you want to export the disk. Click **OK**.

   Once a disk is exported, you can view the associated remote disks on the **Disks** tab.

   You can add the exported disk to a shared disk group that has FSS enabled. You can enable FSS while creating a shared disk group or using the Set/Unset FSS operation on an existing disk group.

   See "Enabling or disabling Flexible Storage Sharing on existing shared disk groups" on page 599.

To un-export disks previously exported for Flexible Storage Sharing

1. In the Management Server console, go to the **Server** perspective and expand **Manage** in the left pane.
2. Expand **Clusters** and click the **Disks** tab.

   You can also perform this operation when from a host context.
3. Right-click the disk and select **Unexport**.
4. Confirm that you want to un-export the disk. Click **OK**.

   Once a disk is un-exported, it is no longer visible to other nodes in the cluster as available storage.

See "Implementing Flexible Storage Sharing with Veritas InfoScale Operations Manager" on page 593.

See "Permissions required for views and operations on CVM cluster objects" on page 591.

Enabling or disabling Flexible Storage Sharing on existing shared disk groups

The Management Server console lets you enable or disable Flexible Storage Sharing (FSS) on existing shared disk groups in a Cluster Volume Manager (CVM) cluster that supports FSS

You can also enable FSS when you create a new disk group in a CVM cluster that supports FSS.

To perform this task, your user group must be assigned the Admin role on the current host, if in the host context, or on the master node of the cluster, if in the cluster context, or the Server perspective. The permission on the host may be explicitly assigned or inherited from a parent Organization.

To enable or disable FSS on an existing shared disk group

1. In the Management Server console, go to the **Server** perspective and expand **Manage** in the left pane.
2. Expand **Clusters** and click the **Shared Disk Groups** tab.

   You can also perform this operation from a selected disk group in a host context.
3. Right-click the disk group and select **Set/Unset FSS**.

   The current state of FSS (on or off) is displayed.
4. Confirm that you want to set FSS on or off for the disk group.

See “Implementing Flexible Storage Sharing with Veritas InfoScale Operations Manager” on page 593.

See “Permissions required for views and operations on CVM cluster objects” on page 591.

Chapter 37

Monitoring the virtualization environment

This chapter includes the following topics:

- About the virtualization technologies supported
- About discovering the VMware Infrastructure using Veritas InfoScale Operations Manager
- About discovering Solaris zones
- About discovering logical domains in Veritas InfoScale Operations Manager
- About discovering LPARs and VIOs in Veritas InfoScale Operations Manager
- About Microsoft Hyper-V virtualization discovery
- About the Kernel-based Virtual Machine (KVM) virtualization discovery in Veritas InfoScale Operations Manager
- About the reports related to virtualization

About the virtualization technologies supported

Veritas InfoScale Operations Manager supports the following virtualization technologies:

- VMware ESX
- Solaris Zones
- Oracle VM Server for SPARC (previously called Sun Logical Domains - LDoms)
- Logical partition (LPAR)

- Microsoft Hyper-V
- Kernel-based Virtual Machine (KVM): Red Hat Enterprise Linux as the KVM Server

For VMware ESX discovery, a designated Control Host discovers the VMware vCenter Server in the data center. This discovery displays those ESX servers that VMware vCenter Server manages and the virtual machines that are configured on the ESX servers.

For Solaris Zones discovery, the zone agentlet that is present in the `VRTSsfmh` package, which is installed on a Solaris managed host, discovers the Global zones that are configured on the host. This discovery displays the non-Global zones that are configured on the Global Zone.

For Sun LDoms discovery, the LDom agentlet that is present in the `VRTSsfmh` package, which is installed on a Solaris managed host, discovers the LDom Server that is configured on the host. This discovery displays the LDom clients that are configured on the LDom Server.

For logical partition (LPAR) discovery, Veritas InfoScale Operations Manager can use Hardware Management Console (HMC), a `VRTSsfmh` package that is installed on the LPAR client, or a `VRTSsfmh` package installed as a part of DMP on the VIO server. Control Host is required for the HMC discovery.

For Microsoft Hyper-V discovery, Veritas InfoScale Operations Manager discovers Hyper-V server (with `VRTSsfmh` package on it), and its correlation with the Hyper-V virtual machines. It also discovers the storage that is provisioned to the guests and its correlation with the virtual machine and the Hyper-V server. The Hyper-V guest, when added (using agent or agentless option) to Veritas InfoScale Operations Manager Management Server domain, provides storage mapping discovery.

For Kernel-based Virtual Machine (KVM) discovery, Veritas InfoScale Operations Manager discovers KVM virtual machines on the Linux host if the KVM modules are installed, and configured on the virtualization server (KVM Server). Veritas InfoScale Operations Manager discovers basic information about KVM virtual machines. For example, virtual machine name, CPU, and so on.

See “About discovering the VMware Infrastructure using Veritas InfoScale Operations Manager” on page 603.

See “About discovering Solaris zones” on page 610.

See “About discovering logical domains in Veritas InfoScale Operations Manager” on page 614.

See “About discovering LPARs and VIOs in Veritas InfoScale Operations Manager” on page 618.

See “About Microsoft Hyper-V virtualization discovery” on page 622.

About discovering the VMware Infrastructure using Veritas InfoScale Operations Manager

In Veritas InfoScale Operations Manager, a managed host that is designated as a Control Host can discover the VMware Infrastructure.

The managed host on which you have installed the Control Host Add-on discovers the information on the following VMware infrastructure components in your data center:

- vCenter servers which manage one or more ESX servers.
- ESX servers on which the individual virtual machines are configured.

Ensure that the Control Hosts can ping the vCenter servers or the ESX servers from which they can discover the information on VMware Infrastructure. A Control Host and a vCenter server or an ESX server from which the Control Host wants to discover the information on VMware Infrastructure must be on the same subnet.

The Control Host Add-on that you install on the designated Control Hosts contains the VMware Infrastructure SDK (VI SDK), which provides a standard interface for the VMware servers and the Control Hosts to access the VMware Infrastructure. A Control Host reports the information on the VMware Infrastructure that it discovers to Management Server. The Management Server coalesces the data that it receives from the Control Host and populates the relevant views.

For more information on the supported versions of vCenter servers and ESX servers, refer to the *Veritas InfoScale Operations Manager Hardware and Software Compatibility List (HSCL)*.

See "About the virtualization technologies supported" on page 601.

See "How Veritas InfoScale Operations Manager discovers vCenter and ESX servers" on page 603.

How Veritas InfoScale Operations Manager discovers vCenter and ESX servers

Veritas InfoScale Operations Manager uses designated Control Hosts to discover the information on the virtual machines. You must install the Control Host Add-on on the managed hosts that you want to designate as Control Hosts. Control Host uses the VMware Infrastructure SDK (VI SDK) to access the vCenter and the ESX servers.

When you configure the virtualization discovery in Veritas InfoScale Operations Manager, you must ensure that you have appropriate privileges to access the vCenter or the ESX servers. Also, you must ensure that you have Browse datastore

privileges on the vCenter or the ESX servers from which you want to discover the VMware Infrastructure information.

From Veritas InfoScale Operations Manager 7.0 onwards, the default method of VMware discovery skips the datastore browsing. Due to this change, you can see the following differences in the data discovered using the previous versions of Veritas InfoScale Operations Manager and the data discovered using Veritas InfoScale Operations Manager 7.0:

- Only those virtual disks are discovered that are attached to some VMware virtual machine.
- Some of the virtual disk attributes are not discovered such as Physical Allocation, Used Capacity, % Utilization, and thin/non-thin.

If you require these details, the datastore browsing can be enabled by setting the `datastore_browse` flag to 1 in the `virtualization.conf` file.

For more information on configuration settings using the `virtualization.conf` file, see the *Veritas InfoScale Operations Manager Installation and Configuration Guide*.

The vCenter server contains a web server, which is an Apache Tomcat server. The web services that are hosted on the web server communicate with the VMware Infrastructure. After you configure a virtualization discovery, the Control Host uses VI SDK to communicate with the web services that are hosted on the web server. For this communication, the Control Host uses the HTTPS protocol.

The URL for the VMware SDK web services is as follows:

https://*traditional host name of the vCenter or the ESX servers*/sdk

After the discovery of VMware Infrastructure, the Control Host reports the discovered data to Management Server.

See "About discovering the VMware Infrastructure using Veritas InfoScale Operations Manager" on page 603.

See "Information that Veritas InfoScale Operations Manager discovers on the VMware Infrastructure components" on page 604.

See "About the datastores in Veritas InfoScale Operations Manager" on page 606.

Information that Veritas InfoScale Operations Manager discovers on the VMware Infrastructure components

Discovery of the VMware Infrastructure components provides the following information:

- Host name and IP address of the VMware Infrastructure components that Veritas InfoScale Operations Manager discovers.
- Operating system handles of the VMware Infrastructure components that Veritas InfoScale Operations Manager discovers.
- Correlation of operating system handles to the virtual disks that are associated with the virtual machines configured on the ESX servers.

See "About discovering the VMware Infrastructure using Veritas InfoScale Operations Manager" on page 603.

See "How Veritas InfoScale Operations Manager discovers vCenter and ESX servers" on page 603.

See "About the datastores in Veritas InfoScale Operations Manager" on page 606.

Viewing the storage mapping information for VMware

The storage mapping view provides information on how storage is allocated to VMware ESX server and its associated virtual machines, the type of virtual disk used, actual physical storage allocation, source and guest devices, and other relevant information.

- Source Name: The path of the `vmdk` file in the datastore used for storage allocation.
- Source Device: The name of the device within the VMware virtualization server.
- Datastore name: The name of the datastore that contains the `vmdk`. A datastore is a virtual representation of combined underlying physical storage resources in the data center.
- vDisk Type: The type of virtual disk - thin, flat, or raw. This information is available only if datastore_browse is set to 1.
- Size: Storage allocation for the virtual disk.
- Physical Allocation: Storage that is currently consumed This information is available only if datastore_browse is set to 1.
- Virtual Disk: The name of the virtual disk that is using the `vmdk`. This information is available only at the virtualization server level.
- Guest device name: The name of the device within the guest machine. The information is available only when the `VRTSsfmh` package is installed on the virtual machine.
- Target type: The type of destination device. For example, disk.

You can view this information related to the virtualization servers for which your user group has at least guest role explicitly assigned or inherited from a parent Organization. You can also view the information if your user group has at least guest role assigned on the Virtualization perspective.

To view the storage mapping information for VMware at the virtualization server level

1. In the Management Server console, go to the **Virtualization** perspective and expand **Manage** in the left pane.
2. Select **Data Center** and expand the Organization.
3. Select the required virtualization server.
4. Click the **Storage Mapping** tab.

To view the storage mapping information for VMware at the virtual machine level

1. In the Management Server console, go to the **Virtualization** perspective and expand **Manage** in the left pane.
2. Select **Data Center** and expand the Organization.
3. Select the required virtualization server.
4. Expand the **Virtual Machines** node.
5. Click the required virtual machine.
6. Click the **Storage Mapping** tab.

See "About near real-time discovery of VMware events" on page 609.

See "About discovering the VMware Infrastructure using Veritas InfoScale Operations Manager" on page 603.

See "About the virtualization technologies supported" on page 601.

About the datastores in Veritas InfoScale Operations Manager

A datastore is a collection of physical storage that is created based on the disks and LUNs that are attached to an ESX server. The storage is assigned to VMware virtual machines from the datastore. Veritas InfoScale Operations Manager lets you view the storage mapping between the datastore and the virtual machines.

A datastore is a storage location for virtual machine files. This storage location can be a VMFS volume or a directory on an NFS file system.

Also, VMware can assign storage to the virtual machines directly from the physical disks that are available with the ESX servers. This storage assignment is known as Raw Device Mapping.

See "About discovering the VMware Infrastructure using Veritas InfoScale Operations Manager" on page 603.

About the multi-pathing discovery in the VMware environment

The support for multi-pathing discovery in the VMware environment lets you discover information that is related to the disk multi-pathing for the ESX servers in your data center. It lets you manage disk paths under VxDMP on the ESX servers. It is enabled using Control Host Add-on. For each disk, you can discover the following information:

- Total number of paths, and number of active paths.
- The details about the path.
- The information about the multi-pathing applications. The multi-pathing applications can be VMware native multi-pathing (NMP), EMC PowerPath/Virtual Edition, or Dynamic Multi-Pathing (DMP) for VMware.

When DMP for VMware is the multi-pathing solution, you can enable or disable the disk paths using the DMP maintenance wizard using the Veritas InfoScale Operations Manager console.

Note: You can also generate the report containing the details about the DMP licenses that are installed on the VMware ESX servers. The license discovery is also enabled using Control Host Add-on.

See "About the user privileges for multi-pathing discovery in the VMware environment" on page 607.

About the user privileges for multi-pathing discovery in the VMware environment

The following table lists the privileges that are required for the multi-pathing discovery on the ESX server:

Table 37-1

Privilege	Is it optional	What if the privilege is not given
Datastore\Browse Datastore	Yes	This flag indicates if datastores need to be browsed for discovering the details of virtual disks. Skipping the datastore browsing causes following differences in VMware discovery as compared to the previous versions of Veritas InfoScale Operations Manager: ■ Only those virtual disks are discovered that are attached to some VMware virtual machine. ■ Some of the virtual disks attributes such as Physical Allocation, Used Capacity, % Utilization, and thin/non-thin are not discovered.
Host\CIM\CIM Interaction	No. However, it is required only for the ESX servers that run VxDMP. It is not required if you don't have any ESX running VxDMP.	■ VxDMP tunables discovery does not happen. ■ The discovery of VxDMP licenses does not happen. ■ DMP path maintenance operations in Veritas InfoScale Operations Manager will not be able to operate on the disk paths in the ESX servers that are managed by VxDMP.

Table 37-1 *(continued)*

Privilege	Is it optional	What if the privilege is not given
Host\Configuration\Change Settings	Yes	Veritas InfoScale Operations Manager cannot discover the version of VxDMP running inside the VMware ESX server. There is no effect on the ESX servers that are not running VxDMP.

See "About the multi-pathing discovery in the VMware environment" on page 607.

About near real-time discovery of VMware events

With near real-time discovery of VMware events, any change in the state of a virtual machine (for example, VM powered on) and changes occurring at the vCenter Server infrastructure-level (for example, VM created) in the Management Server domain are updated in the Veritas InfoScale Operations Manager database in near real-time.

The near real-time discovery of VMware infrastructure enables the partial discovery of ESX servers managed under a vCenter Server. For example, if an SNMP trap is received for a virtual machine (VM1) hosted on ESX1, Veritas InfoScale Operations Manager runs the discovery cycle only for ESX1. Other ESX servers under that vCenter Server are not re-discovered. This discovery is triggered by the event notification from the VMware vCenter Server to the Management Server using SNMP traps. :

For near real-time discovery, ensure to configure the VMware vCenter Server and the Management Server in the same domain. This discovery is supported for the following events occurring at a VMware vCenter Server-level:

Table 37-2 Supported events for near-real time discovery

Discovered state	Event as shown in VMware vCenter Server	Applicable with the Management Server version
Virtual machine powered on	VM powered on	6.0, or later
Virtual machine powered off	VM powered off	6.0, or later

Table 37-2 Supported events for near-real time discovery *(continued)*

Discovered state	Event as shown in VMware vCenter Server	Applicable with the Management Server version
Virtual machine Distributed Resource Scheduler (DRS) powered on	DRS VM powered on	6.0, or later
Virtual machine suspended	VM suspended	6.0, or later
Virtual machine created	VM created	6.1, or later
Virtual machine migrated Hot migration: A powered-on virtual machine is migrated from one ESX server to another ESX server.	VM migrated	6.1, or later
Virtual machine relocated from one ESX server to another Cold migration: A powered-off virtual machine is migrated from one ESX server to another ESX server.	VM relocating	6.1, or later
Virtual machine renamed	VM renamed	6.1, or later
Virtual machine migrated to another host by VMware DRS (Distributed Resource Scheduler)	DRS VM migrated	6.1, or later

Note: The near real-time update of virtual machines is supported on VMware vCenter Server 4.x, 5.x and 6.0.

For more information on setting-up near real-time (NRT) discovery, refer to the *Veritas InfoScale Operations Manager Installation and Configuration Guide*.

About discovering Solaris zones

Veritas InfoScale Operations Manager discovers the zones that are created on Solaris 10 or 11 host that it manages. The Zone Agentlet that is installed along with

the VRTSsfmh package on a Solaris 10 or 11 host lets Veritas InfoScale Operations Manager Management Server discover the Global Zones and the non-Global Zones on the host.

Through the discovery of Solaris zones, Veritas InfoScale Operations Manager discovers the following information:

- The non-Global Zones that are associated with a Global Zone.
- The storage that is exported from the Global Zone to non-Global Zones.
- Databases that run in non-Global Zones of a Solaris 10 host.

Note: Veritas InfoScale Operations Manager discovers the Oracle databases on a non-Global Zone only if the non-Global Zone is in the Running state.

Note: Veritas InfoScale Operations Manager does not support the discovery of applications other than Oracle databases on non-Global Zones.

See "About the virtualization technologies supported" on page 601.

How Veritas InfoScale Operations Manager discovers Solaris zones

In Veritas InfoScale Operations Manager, the zone agentlet that is part of the VRTSsfmh package facilitates the discovery of Solaris zones. Through the zone agentlet, Veritas InfoScale Operations Manager discovers the Global Zone and the associated non-Global Zones.

Note: Discovery of zones is performed only in the Global Zones.

Veritas InfoScale Operations Manager uses the following Solaris utilities to discover global and non-Global Zones:

Table 37-3 Solaris utilities to discover global and non-Global Zones

Utility	Purpose
zoneadm	Lists the non-Global Zones that are configured on the Global Zone on the Solaris 10 managed host
zonecfg	Displays the details of each non-global configurations

Table 37-3 Solaris utilities to discover global and non-Global Zones *(continued)*

Utility	Purpose
`zlogin`	Logs on to a non-Global Zone from a Global Zone

See "About discovering Solaris zones" on page 610.

See "Information that Veritas InfoScale Operations Manager discovers on Solaris zones" on page 612.

See "Limitations of the discovery of Solaris zones in Veritas InfoScale Operations Manager" on page 613.

Information that Veritas InfoScale Operations Manager discovers on Solaris zones

Discovery of Solaris zones provides the following information:

Table 37-4 Solaris zones information that Veritas InfoScale Operations Manager discovers

Type of information	Components
Non-Global Zones , which are configured on Global Zones	Veritas InfoScale Operations Manager discovers the following: ■ The devices that are exported from the Global Zone to the non-Global Zones ■ The file systems that are mounted in the non-Global Zones
Storage that is exported from the Global Zone to the non-Global Zones	Veritas InfoScale Operations Manager discovers the following: ■ Full operating system handles (not slices) ■ Veritas Volume Manager volumes ■ ZFS volumes

Table 37-4 Solaris zones information that Veritas InfoScale Operations Manager discovers *(continued)*

Type of information	Components
Oracle databases inside non-Global Zones	Veritas InfoScale Operations Manager discovers the following: ■ Oracle databases ■ Oracle RAC databases ■ Oracle database on solaris9 branded zone ■ Oracle database on Solaris native branded zone **Note:** Veritas InfoScale Operations Manager does not support the discovery of secure Oracle databases on the non-Global Zones. A secure Oracle database has a password set on 'sysdba' account to secure them.

See "About discovering Solaris zones" on page 610.

See "How Veritas InfoScale Operations Manager discovers Solaris zones" on page 611.

See "Limitations of the discovery of Solaris zones in Veritas InfoScale Operations Manager" on page 613.

Limitations of the discovery of Solaris zones in Veritas InfoScale Operations Manager

The following limitations apply to the discovery of Solaris zones in Veritas InfoScale Operations Manager:

- Veritas InfoScale Operations Manager does not support the discovery of Solaris zones on Solaris managed hosts where both the zones and the LDoms are configured.
- Veritas InfoScale Operations Manager does not recommend the installation of the `VRTSsfmh` package in non-Global Zones.

Note: Conversely, you can install Veritas InfoScale Operations Manager Management Server in a non-Global Zone and add the corresponding Global Zone as a managed host to Management Server.

- Veritas InfoScale Operations Manager does not discover the devices that are exported to non-Global Zones if the `zlogin` utility is not allowed on them.
- Veritas InfoScale Operations Manager discovers only the native and the Solaris-branded zones.

See "About discovering Solaris zones" on page 610.

See "How Veritas InfoScale Operations Manager discovers Solaris zones" on page 611.

See "Information that Veritas InfoScale Operations Manager discovers on Solaris zones" on page 612.

About discovering logical domains in Veritas InfoScale Operations Manager

Logical domains (LDoms) is the paravirtualization technology from Oracle Sun. An LDom provides a separate virtualized operating system environment and a virtualized CPU that are created within a Solaris operating system instance. Each LDom uses an independent kernel. Each LDom contains a dedicated, virtualized operating system, and a virtualized CPU. You can start, stop, and restart the operating system that runs inside an LDom. Each LDom functions as a full virtual machine with a subset of hardware resources that you can configure as required. You can run your applications on the LDoms.

The physical server in your data center on which the LDoms are created is known as LDom Server. Individual Guest LDoms that are created on an LDom Server can have several different roles, which are based on the context and usage of the LDoms.

The following are the four major types of roles of the LDoms:

Table 37-5 LDom roles

LDom role	Description
Control domain	Creates and manages other LDoms and services by communicating with the hypervisor that is present in the LDom Server.
Service domain	Provides the services to other LDoms that are created on the LDom Server. For example, the Service domain provides a virtual network switch or a virtual disk service.

Table 37-5 LDom roles *(continued)*

LDom role	Description
I/O domain	Accesses the input or output devices directly. For example, the I/O domain can access a network device.
Guest domain	Uses the services from the Service and the I/O domains. The Control domain manages the Guest domains.

The LDom Agentlet that is installed along with the `VRTSsfmh` package on a Solaris traditional host lets Veritas InfoScale Operations Manager Management Server discover the LDom Server and Guest LDoms that are configured on the LDom server.

See "About the virtualization technologies supported" on page 601.

See "How Veritas InfoScale Operations Manager discovers Solaris logical domains" on page 615.

See "Information on logical domains that Veritas InfoScale Operations Manager discovers" on page 616.

See "Limitations of the discovery of logical domains in Veritas InfoScale Operations Manager" on page 617.

How Veritas InfoScale Operations Manager discovers Solaris logical domains

In Veritas InfoScale Operations Manager, the LDom agentlet that is part of the `VRTSsfmh` package facilitates the discovery of LDoms on a Solaris managed host. Through the LDom agentlet, Veritas InfoScale Operations Manager discovers the LDom Server and the associated guest LDoms.

Veritas InfoScale Operations Manager uses the `ldm` command to discover the details of Solaris Logical domains.

The following are the major options that Veritas InfoScale Operations Manager uses with the `ldm` command to discover the details of Solaris Logical domains:

Table 37-6 Options that Veritas InfoScale Operations Manager uses with the `ldm` command

`ldm` command and option	Purpose
`ldm list`	Lists all the LDoms that are configured on the LDom Server and their details
`ldm devices`	Discovers the information on the CPU and memory of the LDoms that Veritas InfoScale Operations Manager discovers

See "About discovering logical domains in Veritas InfoScale Operations Manager" on page 614.

See "Information on logical domains that Veritas InfoScale Operations Manager discovers" on page 616.

See "Limitations of the discovery of logical domains in Veritas InfoScale Operations Manager" on page 617.

Information on logical domains that Veritas InfoScale Operations Manager discovers

Discovery of Solaris LDoms provides the following information:

Table 37-7 LDom information that Veritas InfoScale Operations Manager discovers

Type of Information	Components
Guest LDoms that are configured on LDom Servers	Veritas InfoScale Operations Manager discovers the following: ■ Total number of Virtual CPUs ■ Number of Virtual CPUs by core ■ Number of Virtual CPUs that are assigned to each LDom ■ Total system memory ■ Available system memory ■ Memory, which is assigned to each LDom

See "About discovering logical domains in Veritas InfoScale Operations Manager" on page 614.

See "How Veritas InfoScale Operations Manager discovers Solaris logical domains" on page 615.

See "Limitations of the discovery of logical domains in Veritas InfoScale Operations Manager" on page 617.

Limitations of the discovery of logical domains in Veritas InfoScale Operations Manager

The following limitations apply to the discovery of Solaris LDoms in Veritas InfoScale Operations Manager:

- Veritas InfoScale Operations Manager does not discover Solaris LDoms that are in the inactive state.
- Veritas InfoScale Operations Manager does not perform the discovery of Solaris LDoms along with the discovery of Solaris zones.

See "About discovering logical domains in Veritas InfoScale Operations Manager" on page 614.

See "How Veritas InfoScale Operations Manager discovers Solaris logical domains" on page 615.

See "Information on logical domains that Veritas InfoScale Operations Manager discovers" on page 616.

Viewing the storage mapping information for LDoms

You can use the Management Server console to view the storage mapping information for LDoms. Review the following information on this page:

- Source Name: The source providing storage to the LDoms.
- Source Device: The name of the source device providing storage to the LDoms.
- Source Type: The source type that provides the storage. For example, file, disk, or volume.
- Guest device name: The name of the device within LDoms.
- Target type: The type of destination device. For example, disk.
- Storage Container Name: The name of the storage container providing storage to the LDoms.
- Storage Container Type: The type of the storage container providing storage to the LDoms. For example, file system.
- Virtual Machine: The name of the virtual machine consuming the storage.
- Physical allocation: The storage capacity of the devices that are exported to LDom guest. For example Disk, Volume, and File.

You can view this information related to the virtualization servers for which your user group has at least guest role explicitly assigned or inherited from a parent Organization. You can also view the information if your user group has at least guest role assigned on the Virtualization perspective.

To view the storage mapping information for LDoms at the virtualization server level

1. In the Management Server console, go to the **Virtualization** perspective and expand **Manage** in the left pane.
2. Select **Data Center** and expand the Organization.
3. Select the required virtualization server.
4. Click the **Storage Mapping** tab.

To view the storage mapping information for LDoms at the virtual machine level

1. In the Management Server console, go to the **Virtualization** perspective and expand **Manage** in the left pane.
2. Select **Data Center** and expand the Organization.
3. Select the required virtualization server.
4. Expand the **Virtual Machines** node.
5. Click the required LDom.
6. Click the **Storage Mapping** tab.

See "About discovering logical domains in Veritas InfoScale Operations Manager" on page 614.

See "About the virtualization technologies supported" on page 601.

About discovering LPARs and VIOs in Veritas InfoScale Operations Manager

You can use Veritas InfoScale Operations Manager to configure LPAR server, and discover the information that is related to LPARs, VIO clients, and VIO servers in your data center. Agentless discovery of client LPARs and VIO servers is not supported.

Note: The Veritas InfoScale Operations Manager supports only legitimate filename characters in an LPAR profile name. The special characters reserved for Operating System usage (for example space, "\", "$", "!", "&") are not supported. It is recommended to use upper and lower case alphabets, numeric values (0-9), "_" and "-" for the LPAR profile name.

LPAR discovery mechanisms can be grouped into the following categories:

- Discovery using the Hardware Management Console (HMC): The HMC server manages LPAR servers and lets you discover information related to VIO servers and VIO clients. You can use the virtualization management option on the Veritas InfoScale Operations Manager console to add the HMC server to Management Server.

  To add the HMC server to Veritas InfoScale Operations Manager, you need to install the Control Host add-on on the host where the HMC server should be added. Virtual SCSI disks on LPAR client are supported. However, NPIV, or virtual Fibre Channel disks are not supported. Currently, only Virtual SCSI disks backed by native or DMP devices are supported. By configuring HMC server only (without the `VRTSsfmh` package), you can discover information about the exported storage from the VIO server to the VIO clients and the devices that are given to the VIO server from the storage area network (SAN).

- Discovery using the `VRTSsfmh` package that is installed on the LPAR client: The presence of the `VRTSsfmh` package on LPAR client provides additional information about them. This information is correlated with the information that is discovered using the HMC server. Virtual SCSI device discovery, and Virtual SCSI device correlation with the source device in VIO server is also supported.

  Note: Veritas InfoScale Operations Manager supports only native disks as the back-end devices for the VIO server. These disks can be controlled by Microsoft Multipath I/O (MPIO) and Dynamic Multi-Pathing (DMP). Disks that are controlled by any third-party multipathing software (or the logical volumes), when used as the backing devices, do not have end-to-end correlation available.

- Discovery using the `VRTSsfmh` package that is installed as a part of DMP on the VIO server: When a VIO server having DMP version 6.0 or later is added, it provides the discovery of DMP backed exported storage along with the normal managed host discovery. For end-to-end correlation, DMP version 6.0 or later on the VIO server is required. Storage mapping for DMP backed devices is available only if the VIO server (with DMP installed) is added to Veritas InfoScale Operations Manager Management Server.

- Storage Insight Add-on lets you discover complete information about arrays and LUNs from the SAN, which are allocated to the VIO server.

Note: When an array (consumed by the VIO server) is configured, or a VIO server (with DMP) is added to Veritas InfoScale Operations Manager Management Server, refreshing the corresponding HMC discovery is recommended to view the end-to-end correlation immediately in the Veritas InfoScale Operations Manager console.

See "About the virtualization technologies supported" on page 601.

About LPAR storage correlation supported in Veritas InfoScale Operations Manager

Veritas InfoScale Operations Manager provides the support for storage correlation of VIO servers and clients. The storage correlation support for VIO servers and clients provides the information that is related to VIO servers and clients storage consumption at each layer. The following VIO server and clients-related information is provided:

- Information about the assigned storage to the VIO client; whether the storage in the VIO client is directly assigned from the storage area network (SAN), or through a VIO server
- Discovery of VIO servers and correlation of VIO server storage with the SAN.
- Detail of the storage that is exported from the VIO server to the VIO client, and the mapping between the VIO server source device and the VIO client target device
- Information about which VIO server participates in storage allocation to the VIO clients, and how much storage is allocated.
- Information about how much storage is allocated to the VIO server, and how much storage is allocated to the VIO clients from that VIO server.
- Information about how much of the allocated storage is consumed by the VIO client for various applications, and file systems.

See "About discovering LPARs and VIOs in Veritas InfoScale Operations Manager" on page 618.

See "About the virtualization technologies supported" on page 601.

Viewing storage mapping information for LPARs

You can use the Management Server console to view the storage mapping information for LPARs and VIOs. Review the following information on this page:

- Source Name: The source providing storage to LPARs. For example, LUN.
- Source Device: The name of the device within VIO server.
- Source Type: The source type that provides the storage. For example, disk.
- Physical Allocation: The storage that is currently consumed.
- Virtual Disk: The name of the LPAR using the storage.
- Guest device name: The name of the device within the LPAR. The information is available only when `VRTSsfmh` package is installed on the virtual machine.
- Target type: The type of destination device. For example, disk.
- Source VIOs: The name of the VIO server providing the virtualized I/O services to LPARs.

You can view this information related to the virtualization servers for which your user group has at least guest role explicitly assigned or inherited from a parent Organization. You can also view the information if your user group has at least guest role assigned on the Virtualization perspective.

To view the storage mapping information for LPAR at the virtualization server level

1. In the Management Server console, go to the **Virtualization** perspective and expand **Manage** in the left pane.
2. Select **Data Center** and expand the Organization.
3. Select the required virtualization server.
4. Click the **Storage Mapping** tab.

To view the storage mapping information for LPAR at the virtual machine level

1. In the Management Server console, go to the **Virtualization** perspective and expand **Manage** in the left pane.
2. Select **Data Center** and expand the Organization.
3. Select the required virtualization server.
4. Expand the **Virtual Machines** node.
5. Click the required LPAR.
6. Click the **Storage Mapping** tab.

See "About discovering LPARs and VIOs in Veritas InfoScale Operations Manager" on page 618.

See "About the virtualization technologies supported" on page 601.

About Microsoft Hyper-V virtualization discovery

Hyper-V is a hypervisor-based virtualization technology from Microsoft for x86-64 systems. You can use Veritas InfoScale Operations Manager to discover Hyper-V host and virtual machine-related information if the Hyper-V role is enabled on the managed host. Veritas InfoScale Operations Manager uses the Hyper-V WMI API for the discovery.

Hyper-V discovery can be grouped into the following categories:

- Virtual machine discovery: Hyper-V virtual machine discovery by Veritas InfoScale Operations Manager and its correlation with the Hyper-V server.
- Exported storage discovery: Discovery of storage that is provisioned to the guests and its correlation with the virtual machine and Hyper-V server.

See "Virtual machine discovery in Microsoft Hyper-V " on page 622.

See "Storage mapping discovery in Microsoft Hyper-V" on page 623.

Virtual machine discovery in Microsoft Hyper-V

Veritas InfoScale Operations Manager lets you discover information about Hyper-V virtual machines. For example, the name of the virtual machine, allocated memory, CPU, state, and the storage exported (virtual hard disks and pass through disks) from Hyper-V server to Hyper-V guest. Veritas InfoScale Operations Manager discovers all virtual machines including the virtual machines without the guest operating system installed.

Agent and agentless discoveries of Hyper-V virtual machines are supported. However, for the agentless method, the discovered information is limited. To discover more information about the configured virtual machines, the agent discovery method should be used. It provides detailed information about the virtual machines.

Virtual machine discovery prerequisites are as follows:

- The `VRTSsfmh` package should be installed on the Hyper-V server (parent partition).
- The Hyper-V role should be enabled.
- The Windows Management Instrumentation (WMI) service should be running.

A limitation of virtual machine discovery is listed below:

- Hyper-V discovery is not supported on an agentless Hyper-V Server (parent partition) to which the Hyper-V virtual machines are associated.

See "About Microsoft Hyper-V virtualization discovery" on page 622.

Storage mapping discovery in Microsoft Hyper-V

Veritas InfoScale Operations Manager discovers the storage provisioned to the guests from the host's local storage, or storage area network (SAN). The Hyper-V guest (with or without `VRTSsfmh` package), when added to the Veritas InfoScale Operations Manager Management Server domain, provides storage mapping discovery.

Additional storage attributes are also displayed on the page. For example, size, type of the storage (VHD or passthrough disk), and the storage container (volume on the host where virtual storage is provisioned). The storage device handles on the guest will be mapped to the corresponding VHD or passthrough disk provisioned from host. Veritas InfoScale Operations Manager also discovers the snapshot disks provisioned to the VMS.

The storage mapping discovery prerequisites are as follows:

- The Hyper-V server must be running Microsoft Windows 2008 R2 or later operating system.
- Windows Management Instrumentation (WMI) should be running on the guest.

The storage mapping discovery limitation is as follows:

- Storage correlation is not supported for Linux guests.

See "About Microsoft Hyper-V virtualization discovery" on page 622.

Viewing the storage mapping information for Hyper-V

You can use the Management Server console to view the storage mapping information for Hyper-V. Review the following information on this page:

- Source Name: The name of the source providing storage (.vhd file) to the Hyper-V virtual machines.
- Source Device: The name of the source device providing storage to the Hyper-V virtual machines.
- Storage Container Name: The name of the storage container that stores the .vhd files.
- Storage Container Type: The storage container type.
- vDisk Type: Type of virtual disk type. For example, dynamic and differencing.

- Virtual Machine: The name of the virtual machine consuming the provided storage.
- Guest device name: The name of the device within Hyper-V virtual machine.
- Target type: The type of destination device. For example, disk.

You can view this information related to the virtualization servers for which your user group has at least guest role explicitly assigned or inherited from a parent Organization. You can also view the information if your user group has a role assigned on the Virtualization perspective.

To view the storage mapping information for Hyper-V at the virtualization server level

1. In the Management Server console, go to the **Virtualization** perspective and expand **Manage** in the left pane.
2. Select **Data Center** and expand the Organization.
3. Select the required virtualization server.
4. Click the **Storage Mapping** tab.

To view the storage mapping information for Hyper-V at the virtual machine level

1. In the Management Server console, go to the **Virtualization** perspective and expand **Manage** in the left pane.
2. Select **Data Center** and expand the Organization.
3. Select the required virtualization server.
4. Expand the **Virtual Machines** node.
5. Click the required Hyper-V virtual machine.
6. Click the **Storage Mapping** tab.

See "About the virtualization technologies supported" on page 601.

About the Kernel-based Virtual Machine (KVM) virtualization discovery in Veritas InfoScale Operations Manager

Kernel-based Virtual Machine (KVM) is a full virtualization solution for Linux on x86 hardware containing virtualization extensions (Intel VT or AMD-V). Veritas InfoScale Operations Manager discovers KVM virtual machines on the Linux host if the KVM modules are installed, and configured. Veritas InfoScale Operations Manager

discovers basic information about only running virtual machines. For example, virtual machine name, CPU, and so on. Veritas InfoScale Operations Manager uses `virsh` commands to discover KVM-related information.

Kernel-based Virtual Machine (KVM) discovery pre-requisites are as follows:

- `VRTSsfmh` package must be present on the Linux host.
- KVM modules must be installed and configured.

Kernel-based Virtual Machine (KVM) discovery limitation is as follows:

- Exported storage discovery, and storage correlation are not supported.

See "About the virtualization technologies supported" on page 601.

About the reports related to virtualization

In Veritas InfoScale Operations Manager, you can generate the following reports that display the details of the virtualization environment that Veritas InfoScale Operations Manager discovers:

Orphaned Virtual Disks report	Displays the details on the virtual disks that are not used by any virtual machine in the selected reporting scope. The information includes the name of the virtual disk, its type, physical storage allocation for the virtual disk, and other properties.
Storage Breakup by VM State report	Displays the information on storage breakup and virtual machine states. The report provides the information on the storage being consumed by powered ON virtual machines, powered OFF virtual machines and suspended virtual machines.

See "About the virtualization technologies supported" on page 601.

Chapter 38

Using Web services API

This chapter includes the following topics:

- About using Veritas InfoScale Operations Manager Web services API
- Logging in to Veritas InfoScale Operations Manager Web services API
- Logging out of Veritas InfoScale Operations Manager Web services API
- About objects supported by Veritas InfoScale Operations Manager Web services API
- About performing operations using Veritas InfoScale Operations Manager Web services API
- Examples of performing operations using Veritas InfoScale Operations Manager Web services API
- Examples of performing operations using XPRTLC and cURL
- Examples of the output in JSON format

About using Veritas InfoScale Operations Manager Web services API

Veritas InfoScale Operations Manager provides an API that can be accessed over the HTTPS protocol using any standard HTTPS client. The interface provides the ability to query Veritas InfoScale Operations Manager discovered data and to manage user defined attributes for certain object types. The API can be used for searching the objects, listing their properties, and setting the extended attributes on them. Using Veritas InfoScale Operations Manager 6.1 you can also perform operations on some objects. These APIs can be invoked using the XPRTLC client or any other HTTPS client like cURL.

An extended attribute is an user-defined attribute that provides additional details about an object in Veritas InfoScale Operations Manager. These extended attributes can be managed using the Web services API provided by Veritas InfoScale Operations Manager.

Note: Web services API can also be invoked from a host that is not a part of the Management Server domain but has connectivity to the Management Server.

The base URL to access Web services API:

`https://`*`ManagementServer_hostname`*`:14161/vom/api`

where *ManagementServer_hostname* is the host name, fully-qualified host name, or IP address of Management Server.

See "Logging in to Veritas InfoScale Operations Manager Web services API" on page 627.

See "About objects supported by Veritas InfoScale Operations Manager Web services API" on page 629.

See "About performing operations using Veritas InfoScale Operations Manager Web services API" on page 633.

Logging in to Veritas InfoScale Operations Manager Web services API

The Veritas InfoScale Operations Manager Web services API requires users to log in to perform the operations. The authentication information is sent as a part of the HTTPS request body.

You can use one of the following methods to log in:

- With your user credentials.
- Generate a certificate using HTTPS client (XPRTLC or cURL) and use the same to log in.

Use one of the following methods to log in to Web services API using your user credentials:

- Launch the following URL for the log in interface:
  `https://`*`ManagementServer_hostname`*`:14161/vom/api/login`
- Run the following XPRTLC command:

```
xprtlc -l https://ManagementServer_hostname:14161/vom/api/login
-d user=user -d password=password -d
domain=ManagementServer_hostname
```

- Run the following cURL command:

```
curl -g -k -d user=user -d password=password -d
domain=ManagementServer_hostname
https://ManagementServer_hostname:14161/vom/api/login
```

Use the following commands to generate a certificate using XPRTLC or cURL:

- For XPRTLC:

```
xprtlc -l https://ManagementServer_hostname:14161/vom/api/gencert
-d user=user -d password=password -d
domain=ManagementServer_hostname > cert.txt
```

- For cURL:

```
curl -g -k -d user=user -d password=password -d
domain=ManagementServer_hostname
https://ManagementServer_hostname:14161/vom/api/gencert >
/root/cert.txt
```

Use the following commands to log in using the certificate:

- For XPRTLC:

```
xprtlc -l https://ManagementServer_hostname:14161/vom/api/login
-f certfile=@/root/cert.txt
```

- For cURL:

```
curl -g -k -F certfile=@/root/cert.txt
https://ManagementServer_hostname:14161/vom/api/login
```

After you login, a session ID is created using which you can access the Web services API. The session ID is valid for 30 minutes. After the session expires, you need to log in again for a new session ID.

Sample session ID:

{"cookie":"**JSESSIONID=38B752A5DCF210717F5C99D867A17379;**","current_server_time":"Mon Jan 28 03:48:03 PST 2013","max_active_interval":30,"expires_at":"Mon Jan 28 04:18:03 PST 2013"}

See "About using Veritas InfoScale Operations Manager Web services API" on page 626.

See "Logging out of Veritas InfoScale Operations Manager Web services API" on page 629.

Logging out of Veritas InfoScale Operations Manager Web services API

Use one of the following methods to log out:

- Launch the following URL for the log out interface:
  `https://ManagementServer_hostname:14161/vom/api/logout`
- Run the following XPRTLC command:
  `xprtlc -b "session id" -l`
  `https://ManagementServer_hostname:14161/vom/api/logout`
- Run the following cURL command:
  `curl -g -k -b "session id"`
  `https://ManagementServer_hostname:14161/vom/api/logout`

See "Logging in to Veritas InfoScale Operations Manager Web services API" on page 627.

See "About using Veritas InfoScale Operations Manager Web services API" on page 626.

About objects supported by Veritas InfoScale Operations Manager Web services API

The Veritas InfoScale Operations Manager Web services API can be used for searching the objects, setting the extended attributes on them, and also perform certain operations.

Table 38-1 lists some of the supported objects and the corresponding object name to be used in the URL.

Table 38-1 Object name used in the URL

Object	Object name used in the URL
Host	host
Disk	disk
Disk group	diskgroup
Volume	volume
HBA	hba
Virtual Business Services-dependency	vbs-dependency

Table 38-1 Object name used in the URL *(continued)*

Object	Object name used in the URL
Virtual Business Services	vbs
Database	db
Applications database file	db/file
Database - Volume Manager	dbvm
Database - Volume Manager disk	dbvmdisk
Database volume disk group	dbvmdiskgroup
Tablespaces	tablespace
Instances	instance
Exchange servers database file	dbfile
Exchange Servers	exch
Licenses	license
Packages	package
Product	product
Cluster Volume Manager (CVM)	cluster
Service group	sg
System	system
Enclosure	enclosure
Fabric	fabric
Switch	switch
Zone	zone
Switch chassis	switch-chassis
Virtualization server	vserver
Virtual Machine	vm
Storage Mapping	storagemapping
Storage Pool	storagepool

Table 38-1 Object name used in the URL *(continued)*

Object	Object name used in the URL
Logical devices	ldev
RAID group	raidgroup
Thin pool	thinpool
vfiler (NetApp) or Virtual Data Mover (VNX or Celerra)	vfiler
NAS Filesystem Replication	fsrep
Physical devices	pdev
Storage Volume	storagevolume
Storage provisioning templates	template/storage
Replicated volume group	rvg
Hosts participating in RVG replication	rvghost
Replication link between primary and secondary RVGs	rvglink
Recovery plan	recoveryplan
Recovery plan tasks	recoveryplan/task
SmartIO cache area	iocachearea
View all the tasks in a perspective, except for the Management Server perspective	task
View all the sub-tasks in a perspective, except for the Management Server perspective	subtask

You can view HBA from the **Server** perspective as well as the **Storage** perspective. The HBA from **Server** perspective lists the HBAs available on the given host. HBA from **Storage** perspective lists the HBAs to which LUNs are exported from the given enclosure.

To manage an object using Web services API, your user group must be assigned the Admin role on the perspective within the Management Server console in which the object is typically managed.

For example, to manage a host your user group must be assigned the Admin role on the **Server** perspective.

Table 38-2 lists the objects and the perspectives on which you require the Admin role, and the perspective name that you need to use in the URL.

Table 38-2 Objects and perspectives

Object	Perspective	Perspective name used in the URL
Host, disk, disk group, volume, system HBA, exchange server, Virtual Business Services, Virtual Business Services - dependency, databases, database - volume manager, database - volume manager disk, database - volume manager diskgroup, table space, instances, database - file, exchange server, exchange server database file, application databases, replicated volume group (RVG), RVG links, RVG hosts, storage provisioning templates, Cluster Volume Manager (CVM), SmartIO cache area, license, package, and product.	Server	server
Cluster, service group, resource, system, Virtual Business Services -dependency, Virtual Business Services, and recovery plan tasks.	Availability	hadr
Enclosure, logical device (LDEVs), port, HBA, replication, RAID group, thin pool, rank, policy, share, vfiler (NetApp) or Virtual Data Mover (VNX or Celerra), NAS Filesystem Replication, physical device (PDEVs), storage volume, fabric, switch, switch chassis, switch port, zone, zone member, and connectivity.	Storage	storage
Virtualization server, virtual machine, storage mapping, disk, storage pool, cluster, and system HBA.	Virtualization	virt

See "About using Veritas InfoScale Operations Manager Web services API" on page 626.

See "Logging in to Veritas InfoScale Operations Manager Web services API" on page 627.

About performing operations using Veritas InfoScale Operations Manager Web services API

Table 38-3 lists the URLs you can use to view the objects and manage their extended attributes. The supported output format is JSON. The extended attributes are reported in the metadata with editable value set to true. All successful operations return `HTTP 200 OK` response. An unsuccessful operation has an error element in the output which describes the error condition.

Table 38-3 Operations and URLs

Operation	HTTP method	URL
Get metadata for object	GET	*base_url*/meta/*perspective/object type*
Define extended attribute	POST	*base_url*/meta/*perspective/object type*/add?name=**ea_name**
Delete extended attribute	POST	*base_url*/meta/*perspective/object type*/delete?name=**ea_name**
Modify extended attribute	POST	*base_url*/meta/*perspective/object type*/modify?name=**ea_name**&new_name=**ea_new_name**
Set extended attribute	POST	*base_url*/update/*perspective/object type/object id*?**ea_name**=**value**
Get object attributes	GET	*base_url*/query/*perspective/object type/object id*

Where base_url is `https://ManagementServer_hostname:14161/vom/api`. Object type is the objects supported by Veritas InfoScale Operations Manager Web services API, for example host, cluster, LDEV, and virtualization server. **ea_name** is the name of the extended attribute and **value** is the value you assign to the extended attribute.

Using the Veritas InfoScale Operations Manager Web services API you can also perform the following operations:

- Start and stop Virtual Business Services.
  This operation can be performed in the **Server** and **Availability** perspective.
- Start and stop VVR replication.
- Run a recovery plan.
- Provision storage using a storage template.

To perform this operation you need to install Storage Provisioning and Enclosure Migration Add-on version 6.1.

- Perform thin reclamation on disks, volumes, and thin pools of an enclosure.
  To perform this operation on an enclosure, you need to install Storage Insight Add-on version 6.1.
  See "About reclaiming thin storage" on page 329.

Use the following URL to view the above listed operations.

base_url/op/

The GET request on an operation URL shows a sample URL for executing the operation along with the payload information.

For example to start the replication of a replicated volume group (RVG), the start_replication operation URL would include the host ID along with the RVG ID.

base_url/op/server/*host/{host_id}/rvg/{rvg_id}*/startreplication

Enter the operation URL and the payload information in any HTTPS client to execute the operation. On successful completion of the operation, the Task URL is displayed. Click on the Task URL to view complete information about the operation.

operation_urls are displayed when *base_url*/query is executed for an object. This is the list of operations that can be performed on that object.

Note: operation_urls are displayed only when there are operations which are supported from Web API. If there are no operations available **operation_urls** is not displayed.

See "About using Veritas InfoScale Operations Manager Web services API" on page 626.

See "Logging in to Veritas InfoScale Operations Manager Web services API" on page 627.

Examples of performing operations using Veritas InfoScale Operations Manager Web services API

Following are some examples on performing operations using Web services API:

- To get metadata for a host:

  ```
  https://symantecdomain.example.com:14161/vom/api/meta/server/host
  ```

- To define an extended attribute named **department** on object type cluster:

```
https://symantecdomain.example.com:14161/vom/api/meta/hadr/cluster
/add?name=department
```

Adds and displays the extended attribute **department** for object type cluster in the Management Server console.

- To modify an extended attribute named **department** to **Dept_new** on object type cluster:

```
https://symantecdomain.example.com:14161/vom/api/meta/hadr/cluster
/modify?name=department&new_name=Dept_new
```

- To delete an extended attribute named **Dept_new** on object type cluster:

```
https://symantecdomain.example.com:14161/vom/api/meta/hadr/cluster
/delete?name=dept_new
```

When you delete an extended attribute, it is not displayed in the Management Server console.

- To update the extended attribute value for a host where the extended attribute is location:

```
https://symantecdomain.example.com:14161/vom/api/update/server/host
/myhost_id?location=1st floor
```

- To filter disk group objects whose display type is private:

```
https://symantecdomain.example.com:14161/vom/api/query/server/host
/myhost_id/diskgroup?display_type=private
```

- To run a recovery plan

```
https://symantecdomain.example.com:14161/vom/api/op/hadr/
recoveryplan/rplan_id/execute/
```

Payload information required to run a recovery plan is task order and time out duration. If a task is to be skipped during run-time, IS_SKIP should be true. If shared service groups are to be brought offline then OfflineSharedSg should be set to true. Use the above URL along with the payload information in an HTTPS client to run the recovery plan operation.

- To provision storage using a storage template, enter the following URL to view the payload information.

```
https://symantecdomain.example.com:14161/vom/api/op/server/
template/storage/storage_template_id/provision
```

Payload information required to execute the storage template operation is host ID, disk group ID, and volume size. Use the above URL along with the payload information in a HTTPS client to execute the operation.

Note: If "Ask reason for all operations" is enabled under **Advanced Authorization** in the **Management Server** perspective, then it is mandatory to to provide a reason for performing the operation.

See "About performing operations using Veritas InfoScale Operations Manager Web services API" on page 633.

Examples of performing operations using XPRTLC and cURL

Following are some examples on performing operations using XPRTLC and cURL:

To get metadata for a host:

- `xprtlc -m POST -b "session id" -l https://symantecdomain.example.com:14161/vom/api/meta/server/host`

- `curl -g -k -X POST -b "session id" https://symantecdomain.example.com:14161/vom/api/meta/server/host`

To define an extended attribute named **department** on object type cluster:

- `xprtlc -m POST -b "session id" -l https://symantecdomain.example.com:14161/vom/api/meta/hadr/cluster/add?name=department`

- `curl -g -k -X POST -b "session id" https://symantecdomain.example.com:14161/vom/api/meta/hadr/cluster/add?name=department`

  Adds and displays the extended attribute **department** for object type cluster in the Management Server console.

To modify an extended attribute named **department** to **Dept_new** on object type cluster:

- `xprtlc -m POST -b "session id" -l https://symantecdomain.example.com:14161/vom/api/meta/hadr/cluster/modify?name=department&new_name=Dept_new`

- curl -g -k -X POST -b "*session id*" https://symantecdomain.example.com:14161/vom/api/meta/hadr/cluster/modify?name=department&new_name=Dept_new

To delete an extended attribute named **Dept_new** on object type cluster:

- xprtlc -m POST -b "*session id*" -l https://symantecdomain.example.com:14161/vom/api/meta/hadr/cluster/delete?name=dept_new

- curl -g -k -X POST -b "*session id*" https://symantecdomain.example.com:14161/vom/api/meta/hadr/cluster/delete?name=dept_new

  When you delete an extended attribute, it is not displayed in the Management Server console.

To update the extended attribute value for a host where the extended attribute is location:

- xprtlc -m POST -b "*session id*" -l https://symantecdomain.example.com:14161/vom/api/update/server/host/**myhost_id**?location=1st floor

- curl -g -k -X POST -b "*session id*" https://symantecdomain.example.com:14161/vom/api/update/server/host/**myhost_id**?location=1st floor

To filter object disk group whose display type is private:

- xprtlc -b "*session id*" -l https://symantecdomain.example.com:14161/vom/api/query/server/host/**myhost_id**/diskgroup?display_type=private

- curl -g -k -b "*session id*" https://symantecdomain.example.com:14161/vom/api/query/server/host/**myhost_id**/diskgroup?display_type=private

To run a recovery plan:

- xprtlc -m POST -b "*session id*" -l 'https://symantecdomain.example.com:14161/vom/api/op/hadr/recoveryplan/**myrplan_id**/execute' -d 'payload={"Tasks": [{"task_order":"1", "IS_SKIP":false}, {"task_order":"1", "TIMEOUT":"5"}, {"task_order":"2", "IS_SKIP":false,

```
"TIMEOUT":2}], "OfflineSharedSg":true}' -d
'reason=recovery plan execution for ticket no. 123'
```

- ```
  curl -g -k -X POST -b "session id" -l
  'https://symantecdomain.example.com:14161/vom/api/op/hadr/recoveryplan/
  myrplan_id/execute' -d
  'payload={"Tasks": [{"task_order":"1", "IS_SKIP":false},
   {"task_order":"1","TIMEOUT":5},
   {"task_order":"2", "IS_SKIP":false,
   "TIMEOUT":2}], "OfflineSharedSg":true}' -d
  'reason=recovery plan execution for ticket no. 123'
  ```

To provision storage using a storage template

- ```
  xprtlc -b "session id" -m POST -d
  'payload={"HostId":"my_host_id",
   "DiskGroupId":"myDG_id", "VolSize":"vol_size"}' -l
  'https://symantecdomain.example.com:14161/vom/api/op/server/
   template/storage/storage_template_ID/provision'
  ```

- ```
  curl -g -k -X POST -b "session id" -l
  'https://symantecdomain.example.com:14161/vom/api/op/server/
   template/storage/storage_template_ID/provision' -d
  'payload={"HostId": {"my_host_id",
   "DiskGroupId":"myDG_id", "VolSize":"vol_size"}'
  ```

Where, **symantecdomain.example.com** is the name of your Management Server, server is the Server perspective, hadr is the Availability perspective, **myhost_id** is the encoded ID of the host, **myDG_id** is the encoded ID of the disk group, and **myrplan_id** is the recovery plan ID.

If "Ask reason for all operations" is enabled under **Advanced Authorization** in the **Management Server** perspective, then you need to provide a reason for performing the operation.

Note: The output of the `query` URL for an object displays the **encoded_id** value. You need to use this encoded ID value to update the extended attribute for that object.

Similarly the GET request on an operation URL shows a sample URL for executing the operation.

See "About performing operations using Veritas InfoScale Operations Manager Web services API" on page 633.
  ```

Examples of the output in JSON format

Following are some examples of the output in JSON format:

Output for metadata:

```
[
  {
    "name":"id"
    "perspective":"SERVER"
    "editable":false
  },

  {
    "name":"name"
    "perspective":"SERVER"
    "editable":false
  },
  {
    "name":"user defined ea"
    "perspective":"SERVER"
    "editable":true
  },
  .
  .
  .
  {
   "name":"associated_urls",
   "urls":[
            {"name":"volumes",
             "template":[
                          "https://symantecdomain.example.com/vom/api/query
                           /server/host/{host-id}/volume"
                        ],
            }
            {"name":"disks",
             "template":[
                          "https://symantecdomain.example.com/vom/api/query
                           /server/host/{host-id}/disk"
                        ],
             }
             .
             .
```

```
            ]
]
```

Output for query:

```
[
  {
    "user defined ea":"1st floor",
    "encoded_id":"{00120050-56a6-15f3-0000-0000f684d3ab}",
    "id":"{00120050-56a6-15f3-0000-0000f684d3ab}",
    "associated_urls":{
    "diskgroups":"https://symantecdomain.example.com/vom/api/query/server
                /host/{00120050-56a6-15f3-0000-0000f684d3ab}/diskgroup",
    "hbas":"https://symantecdomain.example.com/vom/api/query/server/host/
           {00120050-56a6-15f3-0000-0000f684d3ab}/hba",
    "volumes":"https://symantecdomain.example.com/vom/api/query/server/host/
              {00120050-56a6-15f3-0000-0000f684d3ab}/volume",
    "disks":"https://symantecdomain.example.com/vom/api/query/server/host/
             {00120050-56a6-15f3-0000-0000f684d3ab}/disk"
  },
    "name":"10.255.255.255",
    "url":"https://symantecdomain.example.com/vom/api/query/server/host/
          {00120050-56a6-15f3-0000-0000f684d3ab}",
  .
  .
  }
]
```

See “About performing operations using Veritas InfoScale Operations Manager Web services API” on page 633.

Chapter 39

Veritas InfoScale Operations Manager command line interface

This chapter includes the following topics:

- About the vomadm utility
- Listing all configured enclosures using the vomadm utility
- Host management using the vomadm utility
- Deployment management using the vomadm utility
- Business Application management using the vomadm utility
- Service management using the vomadm utility
- Domain management using the vomadm utility
- List configured schedules using the vomadm utility

About the vomadm utility

The `vomadm` utility lets you perform several operations falling into following categories:

Option	Category	Operation
arrayinfo	config-list	Listing all the enclosures that are configured through Storage Insight Add-on

Option	Category	Operation
host-mgmt	HostManagement	Host management in Veritas InfoScale Operations Manager
hotfix	deployment	Deployment of hotfix in Veritas InfoScale Operations Manager
makeBE	BEManagement	Business Application management in Veritas InfoScale Operations Manager
service	Service-Management	Service management in Veritas InfoScale Operations Manager
domain-mgmt	Domain-Management	Domain management in Veritas InfoScale Operations Manager
schedule		Listing all the schedules that are configured on the managed host.

Note: You need to have administrative privileges to run these commands.

The `host-mgmt` and `makeBE` options can be used only on the Veritas InfoScale Operations Manager Management Server. These commands are supported on Linux and Windows operating system.

The `vomadm` command syntax is described below:

- For Windows Management Server: `c:\Program Files\Veritas\VRTSsfmh\bin>perl.exe vomadm < command>`
- For UNIX Management Server: `/opt/VRTSsfmh/bin/vomadm <command>`

Listing all configured enclosures using the vomadm utility

Use the following command to list all the enclosures configured through Storage Insight Add-on:

`vomadm arrayinfo --config-list <type> [--output-format <json>]`

Where

`config-list`: Method used to configure the enclosure under Storage Insight Add-on , default value is 'all'.

`output-format`: Type of output that the user wants, default is tabular format.

See “About the vomadm utility” on page 641.

Host management using the vomadm utility

Use the following commands to manage the configured hosts in Veritas InfoScale Operations Manager.

`vomadm host-mgmt [--remove --host <hostname> | --remove --hostfile <file> | --list ]`

Where,

`list`: It lists all hosts configured as agents to Veritas InfoScale Operations Manager Management Server.

`remove`: It removes specified host(s) from Veritas InfoScale Operations Manager. To remove single host, use the `--host` option. To remove multiple hosts use the `--hostfile` option.

`host`: The name for the managed host as known by Veritas InfoScale Operations Manager Management Server as listed in the `--list` option.

`hostfile`: The path to a file containing list of host names. Hosts should be listed in the file with each host on a new line.

See “About the vomadm utility” on page 641.

Deployment management using the vomadm utility

Use the following commands to deploy hotfix using the `vomadm` utility:

`vomadm hotfix [--install <hf> | --uninstall <hf-id> | --list | --get-status <task-id>]`

Where

`install`: Installs hotfix specified by the hotfix file.

`uninstall`: Uninstalls the hotfix specified by the hotfix id.

`get-status`: Gets the status of the task specified by the task id.

`hf`: full path to the hotfix file.

`hf-id`: ID of the hotfix.

`task-id`: ID of the task generated when operation is fired.

See “About the vomadm utility” on page 641.

Business Application management using the vomadm utility

The makeBE command option in `vomadm` utility lets you perform various business application-related operations using the command line. The makeBE command syntax is as follows:

```
vomadm makeBE [--import <infile> | --export <outfile> |
--user_defined_import <infile>]
```

Where

`infile`: Absolute path to the input file.

`outfile`: Absolute path to the output file.

See "About the makeBE script " on page 545.

See "Creating Business Application using the makeBE script" on page 549.

See "Updating Business Application using the makeBE script" on page 552.

See "About the vomadm utility" on page 641.

Service management using the vomadm utility

Use the following command to manage the various services in Veritas InfoScale Operations Manager:

```
vomadm service { --start | --stop | --restart | --status | --version
| --help } [process]
```

Where

`start`: Starts the specified service, or all the services that the script manages.

`stop`: Stops the specified service, or all the services that the script manages.

`restart`: Restarts the specified service, or all the services that the script manages.

`status`: Displays the status of the specified service, or all the services that the script manages.

`version`: Displays the version of the `VRTSsfmcs` package or `RTSsfmh` package that is installed on the Management Server host.

`help`: Displays help to use the `service` command and various options available with the command.

`process`: Name of any of the following processes:

- `web` - Veritas InfoScale Operations Manager Web Server
- `at` - Veritas InfoScale Operations Manager Authentication Service
- `xprtld` - Veritas InfoScale Operations Manager Messaging Service
- `db` - Veritas InfoScale Operations Manager Database Service
- `dcli` - Veritas InfoScale Operations Manager Distributed Command Line Daemon
- `sd` - Veritas InfoScale Operations Manager watchdog
- `xtrapd` - SNMP trap service
- `ALL` - All the processes that the script manages

See "About the vomadm utility" on page 641.

Domain management using the vomadm utility

Use the following commands to manage the configured hosts in Veritas InfoScale Operations Manager.

`vomadm domain-mgmt [--remove-all | --remove <domain> | --show-tasks [--latest <n>]]`

Where,

`remove-all`: Deregisters all the hosts from all the Management Server domains, except from the domain where the command is run.

`remove <domain>`: Deregisters all the hosts from Management Server that is specified in the `<domain>` argument. For the `<domain>` argument, use the name that is used while configuring the Management Server.

`show-tasks [--latest <n>]`: Shows the status of all tasks that have been performed or currently running for this command. The `latest` argument when used with `show-tasks` option shows the latest *n* tasks specified, where *n* is a number.

See "About the vomadm utility" on page 641.

List configured schedules using the vomadm utility

Use the following command to list the configured schedules on the managed host.

`vomadm schedule[--help | --list [--<class>]`

Where,

`list --<class>`: displays all the schedules that belong to the class specified.

See “About the vomadm utility” on page 641.

Appendix B

Command file reference

This appendix includes the following topics:

- vxlist
- vomadm
- xdistc

vxlist

`vxlist` – displays records of the Storage Foundation configuration.

SYNOPSIS

`vxlist [-option] [keyword] [arguments] [storage_object_names ...]`

DESCRIPTION

`vxlist` lists Storage Foundation objects.

To display the `vxlist` command output, the `vxdclid` daemon must be running. If `vxdclid` is not running, run `/opt/VRTSsfmh/adm/dclisetup.sh` as a root user.

KEYWORDS

`alert`
: Lists the Veritas Volume Manager alerts.

`cache`
: Lists the Volume Manager cache objects.

`disk`
: Lists the disks.

`diskgroup|dg`
: Lists the Volume Manager disk groups.

`dmp`
: Lists the supported Array Support Libraries (ASL).

`enclosure|enclr|array`
: Lists the enclosures.

`filesystem|fs`
: Lists the mounted file systems.

`getfield`
: Lists the specified fields from Veritas Volume Manager (VxVM) records. Used with the -F option.

`hba|hostport|controller|ctrl`
: Lists the controllers.

`lun`

Lists the Storage Insight Add-on information of LUNs. The information is only available if the host has been added to a Management Server domain, and the associated array has been enabled for Storage Insight.

The first use of the `lun` keyword fetches the array information from Management Server. Any `vxlist lun` command that is run within the next one hour uses the `vxdclid` cache. The `vxdclid` cache may be stale if array parameters have been modified. The `vxlist lun` command that is run after one hour of the last `vxlist lun` command fetches the latest array information from Management Server again.

To force the `vxlist lun` command to fetch the updated array information, use the `rescan` keyword.

`nodeinfo`

Lists the Cluster Volume Manager node information.

`path`

Lists the paths.

`plex|pl`

Lists the plexes.

`rescan`

Fetches updated Storage Foundation information.

See the `lun` keyword.

`snapshot|snap`

Lists the Volume Manager snapshots.

`subdisk|sd`

Lists the subdisks.

`tag|tags`

Lists the volume tags.

`targetport|tca`

Lists the target ports.

`task`

Lists the Volume Manager tasks that are running.

`umfilesystem|umfs`

Lists the unmounted file systems that are referenced in the file system table file.

`volume|vol`

Lists the Volume Manager volumes.

`vset`

Lists the Volume Manager volume sets.

OPTIONS

`-a|--all`

Displays all fields or sections including those that have no data.

`-d|--deliminator` *string*

Uses the specified *string* instead of spaces to delimit fields in tabular display.

`-e|--exact`

Displays all size-related numbers in sectors.

See the `-u` option.

`-F|--format` *"objtype:field1[,field2...] [objtype:field1[,field2...]]"*

Displays the fields specified for one or more object types. For each object type, specify the object type, a colon, and a comma-separated list of field names. A field name can be any field listed in the output of 'vxprint -m' for that object type. The following object types are supported:

- disk
- diskgroup|dg
- enclosure|enclr|array
- filesystem|fs
- hba|hostport|controller|ctrl
- path
- targetport|tca
- volume|vol
- vset

`-k|--kilobyte`

Displays all size-related numbers in kilobytes.

`-g|--diskgroup` *dg*

Lists storage objects in the specified disk group.

`-H|--help [`*objtype*`]`

Displays usage information.

`-l|--long`

Displays in long format.

`-B|--bare` *`field1[,field2...] objtype`*

Displays bare format. Displays only the specified long format fields

`-O|--output [csv|long|table]`

Displays information in the selected format. The default is the 'table' format.

`-p|--property` *`object_name1 object_name2...`*

Displays the property pages of the specified LUNs. The sections without data are not displayed unless the `--all` option is used.

`-q|--suppress`

Suppresses headers in tabular output format.

`-s|--sections` *`sectionname,... object_name`*

Displays only the specified sections in the property page. Sections with no data are not displayed unless the `--all` option is used.

`-t|--table [default|lun|stats]`

Displays LUN information in the specified table format. The default format is 'default'.

`-u|--unit [p|t|g|m|k|blocks|bytes|scaled]`

Displays all size-related numbers in the specified unit. The default is 'scaled'.

EXAMPLES

This section provides usage examples for `vxlist`.

EXAMPLE 1:

To display `vxlist` usage for viewing information on disks.

```
vxlist -H disk
```

EXAMPLE 2:

To display the fields Device, Status, Log Info, and VDID for disks, in bare format, delimited by the '++' string.

```
vxlist -B "Device,Status,Log Info,VDID" -d ++ disk
```

EXAMPLE 3:

To display only the disks section in the property page for the volume named `vol_1`.

```
vxlist -s disks vol vol_1
```

EXAMPLE 4:

To display the property pages with the disks section for the volumes named `vol_1` and `vol_2`.

```
vxlist -p -s disks vol vol_1 vol_2
```

EXAMPLE 5:

To display the fields device_tag, guid, and mediatype for the disk `disk_1`.

```
vxlist -g dg_1 -F disk:device_tag,guid,mediatype getfield disk_1
```

EXAMPLE 6

To display the specified disk and volume fields for the `dg_1` disk group:

```
vxlist -g dg_1 -F "disk:device_tag,guid,mediatype volume:state"
getfield
```

FILES

`/etc/vx/dcli/sfm/conf/dcli_conf.ini`
: The `vxlist` and `vxdclid` configuration file

`/etc/vx/dcli/log/server_A`
: The `vxdclid` log file

NOTES

The default location of `vxlist` is `/opt/VRTSsfmh/bin/vxlist`. There is also a `vxlist` link named `/etc/vx/bin/vxlist`.

Windows-based Management Server does not support the `vxlist` command.

vomadm

`vomadm` – enables you to perform Veritas InfoScale Operations Manager related operations using the command line. The available options are listing configured enclosures, host management, hot fix deployment management, business application management, service management, and domain management.

SYNOPSIS

```
vomadm arrayinfo

vomadm host-mgmt

vomadm hotfix

vomadm makeBE

vomadm service

vomadm domain-mgmt
```

DESCRIPTION

Use `arrayinfo` option for listing configured enclosures, `host-mgmt` option for host management, `hotfix` option for hot fix deployment, `makeBE` option for business applications operations, `service` option for service management, and `domain-mgmt` option for Veritas InfoScale Operations Manager domain management.

KEYWORDS

`arrayinfo`
: Used to list all the enclosures configured through Storage Insight Add-on.

`host-mgmt`
: Used to manage the configured hosts in Veritas InfoScale Operations Manager.

`hotfix`
: Used to install and uninstall hot fix.

`makeBE`
: Used to perform various business application-related operations. For example, create a business application, import a business application.

`service`
: Used to start, stop, or restart one or more Veritas InfoScale Operations Manager services on a Management Server or to get help, version of `VRTSsfmcs` package

and the VRTSsfmh package, or status of one or more Veritas InfoScale Operations Manager services.

domain-mgmt

Used to manage Veritas InfoScale Operations Manager domains configured to managed hosts.

OPTIONS

For arrayinfo

```
[--config-list <type> [--output-format<json>] ]
```

For host-mgmt

```
[--remove --host <hostname> | --remove --hostfile <file> | --list
]
```

For hotfix

```
[--install <hf> | --uninstall <hf-id> | --list ]
```

For makeBE

```
[--import <infile> | --export <outfile> | --user_defined_import
<infile>]
```

For service

```
[{ --start | --stop | --restart | --status | --version | --help
} <process> ]
```

For domain-mgmt

```
[--remove-all | --remove <domain> | --show-tasks {--latest<n>}]
```

FILES

Log files

```
/var/opt/VRTSsfmh/logs/vomadm.log
```

EXAMPLES

This section provides the usage example for vomadm.

EXAMPLE 1:

To list all the enclosures configured through Storage Insight Add-on in json format.

```
/opt/VRTSsfmh/bin/vomadm arrayinfo --config-list all --output-format
json
```

EXAMPLE 2:

To list all the hosts configured as agents toVeritas InfoScale Operations Manager Management Server.

```
/opt/VRTSsfmh/bin/vomadm host-mgmt --list
```

EXAMPLE 3:

To uninstall a hot fix.

```
/opt/VRTSsfmh/bin/vomadm hotfix --unistall VRTSsfmch-4.1.142.0.sfa
```

EXAMPLE 4:

To create a business application using the `makeBE` command.

```
/opt/VRTSsfmh/bin/vomadm makeBE --import /tmp/Samplefile
```

EXAMPLE 5:

To display the status of all the services.

```
/opt/VRTSsfmh/bin/vomadm service --status All
```

EXAMPLE 6:

To start all the services.

```
/opt/VRTSsfmh/bin/vomadm service --start All
```

EXAMPLE 7:

To stop all the services.

```
/opt/VRTSsfmh/bin/vomadm service --stop All
```

EXAMPLE 8:

To deregister hosts from all the Veritas InfoScale Operations Manager domains, except the domain from which the command is run.

```
/opt/VRTSsfmh/bin/vomadm domain-mgmt --remove-all
```

NOTES

The default location of `vomadm` command is `/opt/VRTSsfmh/bin/vomadm`

COPYRIGHTS

Copyright ©) 2016 Veritas Technologies. All rights reserved.

xdistc

`xdistc` – command-line interface to VRTSsfmh distributor

SYNOPSIS

`xdistc [OPTIONS] --push` *`localfile remotefile`*

`xdistc [OPTIONS] --run --command` *`arg1 arg2..`*

`xdistc [OPTIONS] --push` *`localfile remotefile`* `--run --command` *`arg1 arg2..`*

`xdistc --results --id` *`requestid`* `[--wait` *`duration`*`]`

DESCRIPTION

`xdistc` is the command-line interface to the `VRTSsfmh` distributor. You can use `xdistc` to copy files or run commands across all managed hosts in a centrally managed domain. You can perform these tasks on the available hosts that are currently running, and on the unavailable hosts when they are started. When you run `xdistc`, the task that is specified with it continues to run in the background even when xdistc has stopped running.

You can run `xdistc` only from a Management Server host that has the `xprtld` daemon running. You must be logged on as root to run `xdistc`.

OPTIONS

`--push` *`localfile remotefile`*

Copies a file to multiple managed hosts. The symbolic names `$TMPDIR`, `$VARDIR`, and `$TMPFILE` can be used as destination file paths. `$TMPDIR` typically points to the `/tmp` directory, but it may vary on Windows managed hosts. `$VARDIR` points to the `/VRTSsfmh/var` directory. To use, append the file name after the symbolic name. For example, `$TMPDIR/myfile.txt`, or `$VARDIR/myfile2.txt`. `$TMPFILE` creates a temporary file name ensuring no collision with other files in `/tmp`. This is useful with the `--run` option.

You can specify only a few designated directories as the destination.

To overwrite existing files while copying, use the `--force` option.

`--run --command` *`arg1 arg2..`*

Runs a command on multiple managed hosts. In this form, a command is used from commands previously whitelisted on each destination host. The arguments

after the double dash are passed directly to the command. You can specify a request ID with the `--id` option. If the request ID is not specified, a random ID is internally created.

`--push` *`localfile remotefile`* `--run --command` *`arg1 arg2..`*

Used to copy files and run a command on multiple hosts simultaneously. This option is useful when you want to copy an executable file to multiple managed hosts, and run it.

`--results`

Retrieves the `stdout` and `stderr` results from a run request. You can specify this option with the `--run` option to view the results of the command that is executed. You must specify the request ID if you use the `--results` option without the `--run` option. You can use the `--wait` option to specify the time the `xdistc` script should wait to obtain the results.

`--os` *`osname`*

Specifies the operating system. The task that is specified with `xdistc` is run on the managed hosts that have the specified operating system running. The *osname* must be specified as one of the following: **SunOS**, **Linux**, **HP-UX**, **AIX**, **Windows**. You can also specify multiple operating systems. For example, to specify AIX and Linux, use **'/AIX|Linux/'**, including the single quotes.

`--cpu` *`cputype`*

Specifies the processor. The task that is specified with `xdistc` is run on the managed hosts that have the specified processor. The *cputype* must be specified as one of the following: **sparc**, **x86_64**, **powerpc**, **x86**, **x64**, **i386**, **i686**.

`--host` *`hostname`*

Specifies the host. The task that is specified with `xdistc` is run on the specified managed host. You can specify the option multiple times to specify multiple hosts.

`--hostfile` *`filename`*

Specifies a file that contains the names of managed hosts. The file must be whitespace delimited. The task that is specified with `xdistc` is run on the managed hosts that are specified in the file.

`--when` *`spec`*

Specifies the state of the managed hosts. The possible values for *spec* are '**now**', '**up**', or '**now,up**', without the single quotes. If you specify '**now**', the task that you specify with `xdistc` is run on the managed hosts that are already started. If you specify '**up**', the task that you specify with `xdistc` is run on the managed hosts that are being started or restarted, or the managed hosts that

are being added to the centrally managed domain. The default value of *spec* is '**now,up**'.

`--id` *requestid*

Specifies the request ID. You can use the `--id` option with the `--push` or `--run` options to assign a request ID. It also collects the output when it is used with the `--results` option. If not specified, an ID is generated internally.

`--ttl` *timespec*

Specifies the time that `xdistc` should remember the request and the output results of the request. You can specify the time for which `xdistc` should attempt to send the request to the managed hosts. The value of *timespec* can be in days, hours, or minutes. For example, you can use any one of the following values to specify that `xdistc` should remember the request and the output results for a day: **1d**, **24h**, or **1440m**, where **d** stands for days, **h** stands for hours, and **m** stands for minutes. You can also specify one of the following values for *timespec*: '**complete**', or '**forever**', without the single quotes. If you specify *timespec* as '**complete**', `xdistc` deletes the request after it has run the task on the specified hosts. If you specify '**forever**', the request is not deleted automatically.

`--wait` *seconds*

Used with the `--results` option to specify the time the `xdistc` script should wait to obtain the results. You must specify the time in seconds. The default value is **0** seconds, which causes the `xdistc` script to wait indefinitely for the results.

`--force`

Specifies that existing files should be overwritten when files are copied to the managed host.

`--delete`

Specifies that the copied file should be deleted from the managed host after the command that is specified with the `--run` option has completed. The `--delete` option is used only when the `--push` option and the `--run` option are used together.

`--permission` *p*

Specifies the access permissions for the file that is copied to the managed host. You can specify the access permission as an octal number. For example, **500**, or **444**.

`--whitename` *name*

Specifies that after a file is copied, it should be included in the whitelist to be run later.

`--user` *username*

Specifies the user name to be used when the task is run from `xdistc` on each managed host. The default is user name is **vxss:///sfm_admin//**.

`--uri` *uri*

Specifies a raw URI to call. This is a lower-level interface above the `--push` and the `--run` options that is used to invoke arbitrary URLs.

`--d` *option=value*

Specifies the additional values that `xdistc` passes as form data to remote URLs.

NOTES

The default location of `xdistc` is `/opt/VRTSsfmh/bin/xdistc`.

The default log file for `xdistc` is `/var/opt/VRTSsfmh/logs/xdist.log`.

Windows-based Management Server does not support the `xdistc` command.

EXAMPLES

This section provides usage examples for `xdistc`.

EXAMPLE 1:

To copy the `script.sh` file to all Linux managed hosts, run it, and display the results. The file will be deleted from the managed host after it is run. The request will remain active on Management Server for a day. The command will be sent to any new hosts joining the domain during this time, and results can be collected until the request is automatically cleaned up on Management Server.

```
xdistc --ttl 1d --os Linux --push script.sh /tmp/remote.sh --run
--delete --results
```

EXAMPLE 2:

To copy the `/root/script.sh` file to all managed hosts, set its access permissions, and whitelist it. The request will remain active on Management Server for a day. The command will be sent to any new hosts joining the domain during this time, and results can be collected until the request is automatically cleaned up on Management Server.

```
xdistc --ttl 1d --push /root/script.sh /var/opt/VRTSsfmh/script99.sh
--permission 755 --whitename script99
```

EXAMPLE 3:

To run the whitelisted `script99` command with arguments. The request will remain active on Management Server for a day. The command will be sent to any new hosts joining the domain during this time, and results can be collected until the request is automatically cleaned up on Management Server.

```
xdistc --ttl 1d --run -- script99 arg1 arg2
```

EXAMPLE 4:

To run `script99` using the whitelist and the lower-level URI interface. The request will remain active on Management Server for a day. The command will be sent to any new hosts joining the domain during this time, and results can be collected until the request is automatically cleaned up on Management Server.

```
xdistc --ttl 1d --uri admin/whitelist.pl/run --d
argv=["script99","arg1","arg2"]
```

Appendix C

Application setup requirements

This appendix includes the following topics:

- Application setup requirements for Oracle database discovery
- Application setup requirements for Oracle Automatic Storage Management (ASM) discovery
- Application setup requirements for IBM DB2 discovery
- Application setup requirements for Sybase Adaptive Server Enterprise (ASE) discovery
- Application setup requirements for Microsoft SQL Server discovery
- Application setup requirements for Microsoft Exchange Server discovery

Application setup requirements for Oracle database discovery

For Oracle database discovery, ensure that Veritas InfoScale Operations Manager Management Server environment and Oracle database are properly configured. For the supported versions of Oracle database, refer to the *Veritas InfoScale Operations Manager Hardware and Software Compatibility List (HSCL)*.

Oracle database discovery prerequisites

For Veritas InfoScale Operations Manager to discover Oracle database, the home directory of the Oracle database must be accessible. To access the home directory of the database, one of the following prerequisites must be met. If none of the two

prerequisite are met, no detail about the Oracle database is discovered. It is also applicable to the discovery of Oracle Real Application Clusters (RAC).

- The Oracle database is clustered using Cluster Server.
- For Oracle database discovery on Unix/Linux operating systems, the `oratab` file must have the Oracle database listed along with its proper home directory. For Windows operating system, Veritas InfoScale Operations Manager agentlet reads the Windows registry to get the Oracle database configuration details.

Permissions to system tables

When you perform discovery of Oracle databases, ensure that you have the right permission to the following system tables:

ALL_ARGUMENTS	ALL_TAB_COMMENTS	DBA_TS_QUOTAS
ALL_CATALOG	ALL_TRIGGERS	DBA_USERS
ALL_COL_COMMENTS	ALL_TRIGGER_COLS	DBA_VIEWS
ALL_CONSTRAINTS ALL_CONS_COLUMNS	ALL_TYPES	DICTIONARY
ALL_DB_LINKS ALL_ERRORS	ALL_UPDATABLE_COLUMNS	DICT_COLUMNS
ALL_INDEXES	ALL_USERS	GLOBAL_NAME
ALL_IND_COLUMNS	ALL_VIEWS	NLS_DATABASE_PARAMETERS
ALL_LOBS	DATABASE_COMPATIBLE_LEVEL	NLS_INSTANCE_PARAMETERS
ALL_OBJECTS	DBA_DB_LINKS	NLS_SESSION_PARAMETERS
ALL_OBJECT_TABLES	DBA_ERRORS	PRODUCT_COMPONENT_VERSION
ALL_SEQUENCES	DBA_OBJECTS	ROLE_TAB_PRIVS
ALL_SNAPSHOTS	DBA_ROLES	SESSION_PRIVS
ALL_SOURCE	DBA_ROLE_PRIVS	SESSION_ROLES
ALL_SYNONYMS	DBA_SOURCE	SYSTEM_PRIVILEGE_MAP
ALL_TABLES	DBA_TABLESPACES	TABLE_PRIVILEGES
ALL_TAB_COLUMNS	DBA_TAB_PRIVS	TABLE_PRIVILEGE_MAP
ALL_TAB_COL_STATISTICS	DBA_TRIGGERS	

Oracle database discovery

Veritas InfoScale Operations Manager discovers the Oracle database using the following methods:

- **Automatic discovery**: It is the default discovery method. After you install the `VRTSsfmh` package, and add the host to the Veritas InfoScale Operations Manager Management Server domain, the database family of Veritas InfoScale Operations Manager automatically gathers Oracle configuration information. The agentlet uses `sysdba` user account for the discovery. If the password exists for `sysdba` user account, Veritas InfoScale Operations Manager agentlet fails to perform the automatic discovery. If required, you can stop the automatic discovery of Oracle database, as described in the following procedure.
- **Manual Discovery**: If the automatic discovery is not possible (since the password is required), the database is listed in the **Server** perspective of the Management Server console with the discovery state as **Partial (needs credentials)**. In this scenario, right-click the database, and select **Set Credentials**. Enter the user name and password for the Oracle database connection.

You can stop the automatic discovery of Oracle database if you do not want the Veritas InfoScale Operations Manager database agentlet to automatically discover the database.

To stop the automatic discovery of Oracle database

1. In the **Home** page on the Management Server console, click **Settings**.
2. Click **Host**.
3. In the host listing page, select the host for which you want to stop the database discovery.
4. Right-click the host and then select **Properties**.
5. Click the **Discovery families** tab.
6. For the DB family, click the pause button. It will stop the subsequent discovery of the database family on the host. Once you pause the family, the **Frequency** column shows the status as **Paused**. Click the play button to restart the database discovery.

See "SQL queries used for the discovery of Oracle database" on page 663.

See "Oracle database discovery in Solaris zones" on page 665.

SQL queries used for the discovery of Oracle database

The Veritas InfoScale Operations Manager agentlet uses the following SQL queries to discover Oracle database:

- select 'db_name=',value from v$parameter where name = 'db_name';
- select 'total_tablespaces=',count(*) from v$tablespace;
- select 'total_files=',SUM(c) from (select count (*) c from v$datafile union all select count (*) c from v$tempfile);
- select 'total_logfiles=',count(*) from v$logfile;
- select 'status=',status from v$instance;
- select 'parallel=',parallel from v$instance;
- select 'datsize=',sum(bytes) from dba_data_files;
- select 'logsize=',sum(bytes) from sys.v_$log;
- select 'tmpsize=',sum(bytes) from dba_temp_files;
- select 'free_space=',sum(bytes) from dba_free_space;
- select '--archive_log_files--' from dual;
- select a.RECID,'++',a.NAME,'++',a.BLOCKS,'++',a.BLOCK_SIZE,'++',a.STATUS from v$archived_log a;
- select '--redo_log_files--' from dual;
- select [a.member], '++',nvl(b.status, 'null'),'++', b.group#, '++',b.bytes from v$logfile a, v$log b where a.group#=b.group#;
- set colsep;
- select '--tablespace_size--' from dual;
- select t.tablespace_name tablespace,nvl(tsf.bytes,0) free FROM sys.dba_tablespaces t,sys.sm$ts_avail tsa,sys.sm$ts_free tsf WHERE t.tablespace_name = tsa.tablespace_name (+) AND t.tablespace_name = tsf.tablespace_name (+) AND t.tablespace_name not in (select tablespace_name from dba_tablespaces where extent_management = 'LOCAL' and contents = 'TEMPORARY') UNION SELECT t.tablespace_name tablespace,nvl(lt.free,0) FROM sys.dba_tablespaces t,(SELECT h.tablespace_name,sum ((h.bytes_free + h.bytes_used) - p.bytes_used) free FROM sys.v_$temp_space_header h,sys.v_$Temp_extent_pool p wHERE p.file_id(+) = h.file_id group by h.tablespace_name) lt WHERE t.tablespace_name = lt.tablespace_name ORDER BY 1;
- select '--tablespace_status--' from dual;

- `select tablespace_name, status from sys.dba_tablespaces order by tablespace_name;`
- `select '--tablespace_info--' from dual;`
- `select * from (select ts.name name, nvl(f.name, 'null'), nvl(f.file#, 0), nvl(f.bytes, 0), nvl(f.status, 'null') from v$tablespace ts, v$datafile f where ts.ts#= f.ts# union all select ts.name name, nvl(f.name, 'null'), nvl(f.file#, 0), nvl(f.bytes, 0), nvl(f.status, 'null') from v$tablespace ts, v$tempfile f where ts.ts#= f.ts#) order by name;`

See "Application setup requirements for Oracle database discovery " on page 661.

Oracle database discovery in Solaris zones

Veritas InfoScale Operations Manager supports Oracle database discovery in Solaris zones. Veritas InfoScale Operations Manager managed host (agent) must be installed on the Global zone. The discovery happens from the Global zone, by connecting to all the zones.

Note: The automatic discovery of Oracle database cannot be paused for individual zones. The discovery must be paused on the Global zone, which has the agent installed, and it will stop the discovery on the local zones.

See "Application setup requirements for Oracle database discovery " on page 661.

See "About the virtualization technologies supported" on page 601.

Application setup requirements for Oracle Automatic Storage Management (ASM) discovery

For Oracle Automatic Storage Management (ASM) discovery, ensure that Veritas InfoScale Operations Manager Management Server environment and Oracle ASM are properly configured. For the supported versions of Oracle ASM, refer to the *Veritas InfoScale Operations Manager Hardware and Software Compatibility List (HSCL).*

Oracle ASM discovery prerequisites

For Veritas InfoScale Operations Manager to discover Oracle ASM, the home directory of Oracle ASM is obtained from the `oratab` file. If the entry of ASM is missing from `oratab` file, Veritas InfoScale Operations Manager fails to discover Oracle ASM details.

Oracle ASM discovery

Veritas InfoScale Operations Manager discovers the Oracle ASM using the following method:

- **Automatic discovery**: It is the default discovery method. After you install `VRTSsfmh` package, the database (DB) family of Veritas InfoScale Operations Manager automatically gathers Oracle ASM configuration information. The agentlet uses the `sysdba` user for the discovery of Oracle ASM. If password exists for `sysdba` user account, the Veritas InfoScale Operations Manager agentlet fails to perform the automatic discovery of Oracle ASM.

You can stop the automatic discovery of Oracle ASM if you do not want the Veritas InfoScale Operations Manager database agentlet to automatically discover the Oracle ASM.

To stop the automatic discovery of Oracle ASM

1. In the **Home** page on the Management Server console, click **Settings**.
2. Click **Host**.
3. In the host listing page, select the host for which you want to stop the Oracle ASM discovery.
4. Right-click the host and then select **Properties**.
5. Click the **Discovery families** tab.
6. For the DB family, click the pause button. It will stop the subsequent discovery of the Oracle ASM family on the host. Once you pause the family, the **Frequency** column shows the status as **Paused**. Click the play button to restart the Oracle ASM discovery.

See "SQL queries used for the discovery of Oracle ASM" on page 666.

See "Oracle Automatic Storage Management (ASM) discovery in Solaris zones" on page 667.

SQL queries used for the discovery of Oracle ASM

The Veritas InfoScale Operations Manager agentlet uses the following SQL queries to discover Oracle ASM:

- To discover Oracle ASM diskgroup in 10gR1 release:

```
select GROUP_NUMBER, NAME, STATE, TYPE, TOTAL_MB, FREE_MB from
v$asm_diskgroup;
```

- To discover Oracle ASM diskgroup in 10gR2 or later release:

```
select GROUP_NUMBER, NAME, STATE, TYPE, TOTAL_MB, FREE_MB,
REQUIRED_MIRROR_FREE_MB, USABLE_FILE_MB, DATABASE_COMPATIBILITY
from v$asm_diskgroup;
```

- To discover Oracle ASM disk:

```
select GROUP_NUMBER, NAME, PATH, STATE, TOTAL_MB, FREE_MB from
v$asm_disk;
```

See "Application setup requirements for Oracle Automatic Storage Management (ASM) discovery" on page 665.

See "Oracle Automatic Storage Management (ASM) discovery in Solaris zones" on page 667.

Oracle Automatic Storage Management (ASM) discovery in Solaris zones

Veritas InfoScale Operations Manager supports Oracle ASM discovery in Solaris zones. For the discovery, the managed host (agent) must be installed on the Global zone. The discovery happens from the Global zone, by connecting to all the zones.

Note: The automatic discovery of Oracle ASM cannot be paused for individual zones. The discovery must be paused on the Global zone, which has the agent installed, and it will stop the discovery on the local zones.

See "Application setup requirements for Oracle Automatic Storage Management (ASM) discovery" on page 665.

See "SQL queries used for the discovery of Oracle ASM" on page 666.

Application setup requirements for IBM DB2 discovery

For IBM DB2 (Extended Edition and Enterprise-Extended Edition) discovery, ensure that Veritas InfoScale Operations Manager Management Server environment and IBM DB2 are properly configured. For the supported versions of IBM DB2, refer to the *Veritas InfoScale Operations Manager Hardware and Software Compatibility List (HSCL)*.

Veritas InfoScale Operations Manager agentlet attempts to gather DB2 configuration information automatically. Automatic information gathering allows DB agentlet to discover and return information about DB2 instances without intervention.

You can stop the automatic discovery of IBM DB2 if you do not want the Veritas InfoScale Operations Manager agentlet to automatically discover the IBM DB2 database.

To stop the automatic discovery of IBM DB2 database

1. In the **Home** page on the Management Server console, click **Settings**.
2. Click **Host**.
3. In the host listing page, select the host for which you want to stop the IBM DB2 database discovery.
4. Right-click the host and then select **Properties**.
5. Click the **Discovery families** tab.
6. For the DB family, click the pause button. It will stop the subsequent discovery of IBM DB2 database family on the host. Once you pause the family, the **Frequency** column shows the status as **Paused**. Click the play button to restart the IBM DB2 database discovery.

See "SQL queries used for the discovery of IBM DB2 database" on page 668.

SQL queries used for the discovery of IBM DB2 database

The Veritas InfoScale Operations Manager agentlet uses the following SQL queries to discover IBM DB2 database:

- *db2_home_dir*/sqllib/adm/db2licm -l
- *db2_home_dir*/sqllib/bin/db2 get dbm configuration
- *db2_home_dir*/sqllib/bin/db2 list database directory
- *db2_home_dir*/sqllib/bin/db2 connect to <database_name>
- *db2_home_dir*/sqllib/bin/db2 get db configuration for <database_name>
- *db2_home_dir*/sqllib/bin/db2 list tablespaces show detail
- *db2_home_dir*/sqllib/bin/db2 list tablespace containers for <tablespace_name> show detail
- For IBM DB2 versions earlier than 9.7:

```
print J "su $sid -c \" $home_dir/sqllib/bin/db2 connect to
$db_name\; $home_dir/sqllib/bin/db2 \\\"select
TABLESPACE_NAME,total_pages as TBSPC_Pages, USED_PAGES, PAGE_SIZE,
TABLESPACE_STATE, TBS_CONTENTS_TYPE, FREE_PAGES, EXTENT_SIZE,
PREFETCH_SIZE, NUM_CONTAINERS, USABLE_PAGES from table
```

```
(snapshot_tbs_cfg (\'$db_name\',-2)) as snapshot_tbs_cfg\\\";
$home_dir/sqllib/bin/db2 disconnect $db_name\"";
```

- For IBM DB2 version 9.7 or later:

```
print J "su $sid -c \" $home_dir/sqllib/bin/db2 connect to
$db_name\; $home_dir/sqllib/bin/db2 \\\"SELECT TBSP_NAME as
TABLESPACE_NAME, TBSP_TOTAL_PAGES as TBSPC_Pages, TBSP_USED_PAGES
as USED_PAGES, TBSP_PAGE_SIZE as PAGE_SIZE, TBSP_STATE as
TABLESPACE_STATE, TBSP_CONTENT_TYPE as TBS_CONTENTS_TYPE ,
TBSP_FREE_PAGES as FREE_PAGES, TBSP_EXTENT_SIZE as EXTENT_SIZE,
TBSP_PREFETCH_SIZE as PREFETCH_SIZE, TBSP_NUM_CONTAINERS as
NUM_CONTAINERS, TBSP_USABLE_PAGES as USABLE_PAGES FROM
TABLE(MON_GET_TABLESPACE('',-1)) AS snapshot_tbs_cfg\\\";
$home_dir/sqllib/bin/db2 disconnect $db_name\"";
```

- *db2_home_dir*/sqllib/bin/db2

See "Application setup requirements for IBM DB2 discovery" on page 667.

Application setup requirements for Sybase Adaptive Server Enterprise (ASE) discovery

For Sybase Adaptive Server Enterprise (ASE) discovery, ensure that Veritas InfoScale Operations Manager Management Server environment and Sybase database are properly configured. For the supported versions of Sybase database, refer to the *Veritas InfoScale Operations Manager Hardware and Software Compatibility List (HSCL)*.

Sybase ASE discovery prerequisites

Assign a password to the system administrator (SA) user. For more information on application setup, refer to Sybase ASE documentation.

Sybase ASE discovery

Veritas InfoScale Operations Manager does not provide automatic discovery of Sybase ASE. To enable Sybase ASE discovery, you must configure Veritas InfoScale Operations Manager to discover Sybase instances.

To enable Sybase ASE discovery in Veritas InfoScale Operations Manager

1. In the **Home** page on the Management Server console, click **Server** perspective.
2. Expand **Applications** under **Manage**.

3 Under **Databases**, select the desired Sybase instance.

4 Right-click the database and then select **Set Credentials**. Enter the username and password that can be used for the Sybase ASE connection.

You can stop the automatic discovery of Sybase ASE if you do not want the Veritas InfoScale Operations Manager agentlet to automatically discover the database.

To stop the automatic discovery of Sybase ASE

1 In the **Home** page on the Management Server console, click **Settings**.

2 Click **Host**.

3 In the host listing page, select the host for which you want to stop the Sybase ASE discovery.

4 Right-click the host and then select **Properties**.

5 Click the **Discovery families** tab.

6 For the DB family, click the pause button. It will stop the subsequent discovery of Sybase ASE family on the host. Once you pause the family, the **Frequency** column shows the status as **Paused**. Click the play button to restart the Sybase ASE discovery.

See "SQL queries used for the discovery of Sybase Adaptive Server Enterprise (ASE)" on page 670.

SQL queries used for the discovery of Sybase Adaptive Server Enterprise (ASE)

The Veritas InfoScale Operations Manager agentlet uses the following SQL queries to discover Sybase Adaptive Server Enterprise (ASE):

- `sp_helpdb`
- `sp_helpdevice`
- `sp_helpsegment`
- `select @@pagesize`
- `use <database_name>`
- `sp_helpdb <database_name>`
- `select count(*) from sysobjects where type = 'U' or type = 'S'`
- `select name from syssegments`
- `sp_helpsegment <segment_name>`

See “Application setup requirements for Sybase Adaptive Server Enterprise (ASE) discovery” on page 669.

Application setup requirements for Microsoft SQL Server discovery

For Microsoft SQL Server discovery, ensure that Veritas InfoScale Operations Manager Management Server environment and Microsoft SQL Server are properly configured. For the supported versions of Microsoft SQL Server, refer to the *Veritas InfoScale Operations Manager Hardware and Software Compatibility List (HSCL)*.

Microsoft SQL Server discovery

Veritas InfoScale Operations Manager discovers the Microsoft SQL Server using the following methods:

- **Automatic discovery**: Automatic discovery is the default discovery method. After you install `VRTSsfmh` package, and add the host to a Veritas InfoScale Operations Manager Management Server, the database (DB) family of Veritas InfoScale Operations Manager automatically gathers MS SQL Server configuration information. The agentlet runs as local system user and connects to MS SQL Server without credentials. If local system user does not have access to the MS SQL Server, automatic discovery collects limited information.
- **Manual Discovery**: If automatic discovery is not possible because the local system user does not have access to the MS SQL Server, the database is listed in the Server perspective of the Management Server console, with the discovery state as **Partial (needs credentials)**. Right-click on the database and select **Set Credentials**. Enter the username and password for the MS SQL Server connection.

You can stop the automatic discovery of Microsoft SQL Server if you do not want the Veritas InfoScale Operations Manager agentlet to automatically discover the Microsoft SQL Server.

To stop the automatic discovery of Microsoft SQL Server

1. In the **Home** page on the Management Server console, click **Settings**.
2. Click **Host**.
3. In the host listing page, select the host for which you want to stop the MS SQL Server discovery.
4. Right-click the host and then select **Properties**.

5 Click the **Discovery families** tab.

6 For the DB family, click the pause button. It will stop the subsequent discovery of the Microsoft SQL Server family on the host. Once you pause the family, the **Frequency** column shows the status as **Paused**. Click the play button to restart the Microsoft SQL Server discovery.

See "SQL queries used for the discovery of Microsoft SQL Server" on page 672.

SQL queries used for the discovery of Microsoft SQL Server

The Veritas InfoScale Operations Manager agentlet uses the following SQL queries to discover Microsoft SQL Server:

- `select * from sysdatabases`
- `DBCC SQLPERF(LOGSPACE)`
- `use database_name; exec sp_helpfilegroup`
- `use database_name; DBCC SHOWFILESTATS`
- `use database_name; select * from dbo.sysfiles where groupid=group_id`

Note: The agentlet runs the first query for every SQL Server instance. This query returns the list of the databases in an instance. For the remaining set of queries, the *database_name* is the value that is returned from the first query.

See "Application setup requirements for Microsoft SQL Server discovery" on page 671.

Application setup requirements for Microsoft Exchange Server discovery

For Microsoft Exchange Server discovery, ensure that Veritas InfoScale Operations Manager Management Server environment and Microsoft Exchange Server are properly configured. For the supported versions of Microsoft Exchange Server, refer to the *Veritas InfoScale Operations Manager Hardware and Software Compatibility List (HSCL)*.

Veritas InfoScale Operations Manager automatically discovers the Microsoft Exchange Server details and no additional configuration is required.

See "About Business Applications in Veritas InfoScale Operations Manager" on page 538.

See “About Veritas InfoScale Operations Manager” on page 26.

Glossary

Authentication Service	See Symantec Product Authentication Service.
cluster	A set of hosts (each termed a node) that share a set of disks and are connected by a set of redundant heartbeat networks.
cluster communication	Communication between clusters using either of the two core communication protocols defined by Veritas Cluster Server: GAB and LLT. The communication takes place by means of heartbeat signals sent between systems or fast kernel-to-kernel broadcasts.
DMP (Dynamic Multipathing)	A feature of Veritas Volume Manager that provides greater reliability and better performance by using path failover and load balancing for multiported disk arrays connected to host systems through multiple paths. DMP detects the various paths to a disk using a mechanism that is specific to each supported array type. DMP can also differentiate between different enclosures of a supported array type that are connected to the same host system.
file system	A means of organizing the addressable (LUN) storage of one or more physical or virtual disks to give users and applications a convenient way of organizing files. File systems appear to users and applications as directories arranged in a hierarchy.
logical unit number	See LUN.
LUN (logical unit number)	A unique and discrete addressable unit or logical volume that may reside inside one or more simple or array storage devices. LUNs are exposed to the outside world through an addressing scheme presented to the host as SCSI LUN numbers. Each LUN has a unique device handle and represents a logical volume.
node	An object in a network. In Veritas Cluster Server, node refers specifically to one of any number of hosts in a cluster. See also object.

Index

A

adding
- disks to a disk group 159
- mirrors to volumes 205
- resources 411
- service groups 358

alerts and rules 102
App VG
- modify 301

application configuration
- about 455
- launching 456
- prerequisites 456

ApplicationHA 450
ApplicationHA infrastructure
- about 452
- disable 453
- enable 453

ApplicationHA Management
- overview 450
- prerequisites 451
- supported virtual machines 450

ApplicationHA operations
- about 451
- disabling ApplicationHA infrastructure 453
- enabling ApplicationHA infrastructure 453
- launching 452

arrays
- disabling DMP paths 582

assigning price tier
- automatically 486
- manually 486
- processor price tier 485
- server price tier 485

Atleast Count dependency
- child resources 366
- create 365

attributes
- editing 385

Availability
- edit attributes 348

B

break-off mirror volume snapshot
- configure 220

Business Applications
- creating 538
- deleting 540
- modifying 538
- overview 538
- renaming 539

C

change log
- synchronizing 255

checking file system 268
checkpoint
- creating 269

clearing faults on resources 402
cluster hosts
- freezing 391
- managing 391
- starting Cluster Server HAD 393
- stopping Cluster Server HAD 394
- unfreezing 392

Cluster Server HAD
- starting 393
- stopping 394

clustered file systems
- mounting 264
- unmounting 262

clusters
- close configuration 352
- editing attributes 352
- FSS 594
- managing 350
- open configuration 351
- save configuration 351

compression 333
- adding directories 334
- scheduling 335
- starting on demand 337

configure
fire drill service groups 427
instant volume snapshot 214
creating
checkpoint 269
file systems 246
recovery plan 438
space-optimized volume snapshot 218
volumes 192
custom signature
copying to the managed hosts 571
creating script 568
registering on Management Server 571
removing 572
sample script 569
using Distribution Manager Add-on 572

D

datastore 606
deduplication
configuring 341
disabling 344
overview 338
removing 344
starting 343
deep array discovery
ESX servers 603
defragmenting file system 257
deleting
fire drill schedules 434
resources 401
service groups 371
dependency views 386
disabling
all resources in a service group 371
fire drill schedules 435
service groups 364
discovering VMware Infrastructure 603
disk group
adding disks 159
creating 148
deporting 157
destroying 157
importing 158
joining 170
moving 168
recovering 156
removing disks 178
renaming disks 163
disk group *(continued)*
resizing disks 161
splitting 166
upgrading 166
disks
disconnecting 178
evacuating 184
exporting for FSS 598
host prefix 179
initializing 172
mapping 176
offlinig 181
online 180
recover disks 175
recovering 175
replace disks 174
rescanning 188
thin reclamation 330
unmapping 177
usage 181
DMP maintenance case
removing a record 587
view results and output 588
DMP paths
disabling on a host 576
disabling on a virtualization server 583
disabling on an array 582
Dynamic Multi-Pathing 575

E

editing attributes 348
host 393
resource 408
resource type 409
enabling
all resources in a service group 370
fire drill schedules 434
resources 400–401
service groups 363
ESX servers
datastore 606
deep array discovery 603
extended attribute
modifying value 558
searching and setting values 556
setting values 555
using 554

F

faults
    restoring 117
    suppressing 115
faults and risks 115
file change log
    disabling 255
    enabling 254
    removing 256
file compression 333
    adding directories 334
    scheduling 335
    starting on demand 337
file system snapshot
    mount 273
    remounting 271
    removing 276
    unmount 275
file systems
    checking 268
    creating 246
    deduplication 338, 341, 343–344
    defragmenting 257
    remounting 265
    thin reclamation 330
files
    pinning to cache 292
    unpinning from cache 292
fire drill
    configure fire drill service groups 427
    deleting schedules 434
    disabling schedules 435
    enabling schedules 434
    modifying schedule 432
    running DR fire drill 428
    running HA fire drill 426
    viewing schedules 436
FSS
    enabling or disabling on existing disk group 599
    exporting disks 598
    functionality 594
    use cases 595

G

global clusters 415
    adding remote clusters 418
    creating 417
    objects 415
    prerequisites for creating 417
global clusters *(continued)*
    removing remote clusters 421–422
    terminology 416
global service groups
    converting from local 418
    converting to local 420

H

HA-DR
    pre-requisites 348
Hardware Management Console 618
host
    disabling DMP paths 576
    editing attributes 393
Hyper-V virtualization discovery 622

I

I/O Trace Log
    panel options 294
I/O trace log
    analyzing 295
    viewing 295
importing
    disk group 158
    type definition 354
instant volume snapshot
    configure 214
IO Thresholds
    setting 300

L

LDom discovery
    limitations 617
    method 615
LDoms
    discovery 614
    discovery information 616
    roles 614
license deployment policy
    creating 488
    deleting 491
    modifying 490
licenses
    about 475
    accountable deployed license 495
    child licenses 496
    deployed licenses 475
    deployment details 496

licenses *(continued)*
deployment policy details 497
discovery interval 475
features enabled 496
host deployment summary 493
license deployment summary 493
product deployment summary 493
SPVU details 496
licensing and pricing
OS tier 476
per-core 476, 484
processor tier 476
server tier 476
SPVU 476, 479
IBM LPAR 484
kernel-based virtual machines 483
Solaris LDOM virtualization server 482
VMware virtual machine 480
linking
resources 410
service groups 372, 375
LPAR
about discovering 618

M

makeBE script
about 545
creating business application 549
CSV file 546
deleting business application 552
export 547
exporting business application 551
import 547
importing business application 550
limitations 548
log files 549
objects 546
updating business application 552
user defined import 547
Management Server console
about 31
home page 32
user interface element 35
migrating
VM 398
modifying
fire drill schedule 432
resources 411
system list 386
monitoring capacity 278
mounting
clustered file systems 264
file system snapshot 273
non clustered file systems 259
Multi Site Management
about 458
features 459
limitations 460
prerequisites 460
setting-up campus cluster 461
setting-up Replicated Data Cluster 462
multi-pathing discovery
user privileges 607
VMware environment 607

N

non clustered file systems
mounting 259
unmounting 257

O

object
modifying permissions 98
offline and propagate
resources 404
online help
about 49
Organization
assigning permissions 96
creating 89
deleting 95
deleting permissions 97
modifying name 94
modifying permissions 97

P

performance graphs
about 510
disk 514
enclosure 522
file system 515
host 513
initiator 518
path 517
virtual machine 519
VMware ESX server 519
VMware ESX Server path 521

performance graphs *(continued)*
volume 515
performance metering
about 501
disable for host 508
disable for virtualization server 509
enable for host 508
enable for virtualization server 509
metered resources 501
permissions
assigning on Organization 96
deleting on Organization 97
modifying on object 98
modifying on Organization 97
policy checks
about 559
bundling custom signatures using Distribution Manager Add-on 572
enabling or disabling signatures 564
exporting signatures 567
registering signatures 560–561
running manual scans 564
unregistering signatures 562
viewing policy violation details 565
policy violations
viewing scan details 565

R

reactivating volume 201
reattaching
snapshot 224
recover disks 175
recovering disk group 156
recovering volume 200
recovery plan
creating 438
deleting 446
editing 442
executing 444
log files 449
overview 437
viewing historical runs 446
viewing properties 447
refreshing
break-off link snapshot 232
break-off mirror snapshot 232
instant snapshot 232
space-optimized snapshot 232
remote clusters 418
removing 421–422
wac resource 422
remounting
file system 265
file system snapshot 271
removing
disks from a disk group 178
DMP maintenance record 587
file change log 256
file system snapshot 276
volume mirrors 212
renaming
disks in a disk group 163
replace disks 174
replications
adding secondary 313
associating volume 322
configuring 304
monitoring 326
pausing 314
removing secondary 324
resuming 315
resynchronizing secondary 324
setting alerts 326
starting 316
stopping 318
switching primary 319
taking over primary 320
unconfiguring replication 325
report subscription
all subscriptions 127
delete 125
edit 124
my subscriptions 126
reports
about 118
all subscriptions 127
delete subscription 125
edit subscription 124
email 125
my subscriptions 126
running 121
saving 122
subscribing 123
using 120
virtualization 625
VOM deployment 498

resize
disk groups 161
volumes 225
resource action
invoking 406
resource dependency
viewing 413
resource type
editing attributes 409
resources
adding 411
bringing online 405
clearing faults 402
clearing from admin wait state 384
deleting 401
editing attributes 408
enabling 400–401
linking 410
managing 400
mark as critical 412
mark as non critical 413
modifying 411
offline and propagate 404
probing 403
taking offline 405
unlinking 410
restore data from snapshot 230
rules
about 102
creating 103
deleting 113
disabling 114
editing 109
enabling 114

S

scheduling file compression 335
scripts
vxlist 648
xdistc 656
Search feature 41, 43
service group dependency
viewing 387
service groups
adding 358
bringing online 377
clearing faults 383
deleting 371
disabling 364
service groups *(continued)*
disabling all resources 371
enabling 363
enabling all resources 370
flushing 369
freezing 367
linking 372, 375
managing 357
switch 381
taking offline 379
unfreezing 368
setting volume usage 239
signatures
custom 567, 572
enabling or disabling 564
exporting 567
registering 560–561
setting tunables 563
unregistering 562
Smart Folders 43–44
SmartAssist
about 293
SmartIO caching
about 281
creating 283
deleting 291
enabling or disabling 282, 287
impact analysis 286
modifying 288
viewing details 285
write-back 282
snapshot
add to a refresh schedule 237
configure break-off mirror snapshot 220
configure instant snapshot 214
dissociating 222
reattaching 224
refresh 232
remove from refresh schedule 238
restore data from snapshot 230
schedule refresh 235
splitting 240
solaris zones 610
Global Zones 610
non-Global Zones 610
solaris zones discovery 610
information 612
limitations 613
method 611

solaris zones discovery *(continued)*
utilities 611
space estimation data logs 505
splitting
disk group 166
snapshot 240
starting
Cluster Server HAD 393
synchronization of snapshots 241
VM 396
stopping
Cluster Server HAD 394
VM 397
stopping volume 200
stretch sites
site fencing preference 463
synchronizing change log 255
system list
modify 386

T

tablespaces
pinning to cache 292
unpinning from cache 292
tablespaces or files
pinning to cache 291
taking offline
resources 405
service groups 379
thin reclamation
on file systems or disks 330
on thin pools 331
overview 329
threshold settings
about 523
configuring 525
deleting 528
disabling 534
enabling 531
Thresholds
managing 299
Trim
running 186
scheduling 186
type definition
importing 354

U

unlinking
resources 410
unmounting
clustered file systems 262
file system snapshot 275
non clustered file systems 257
upgrading disk group 166
user group permissions
assigning on Organization 96
deleting on Organization 97
modifying on object 98
modifying on Organization 97

V

Veritas InfoScale Operations Manager
about 26
viewing
fire drill schedules 436
resource dependency 413
service group dependency 387
Virtual Business Services
overview 349
virtualization management 450
virtualization reports 625
virtualization server
disabling DMP paths 583
virtualization technology
kernel-based virtual machine 601
LPAR 601
Microsoft Hyper-V 601
Solaris LDom 601
Solaris zones 601
VMware 601
VMware discovery 603
datastore 606
ESX servers 603
information 604
near real-time update of virtual machine states 609
vCenter servers 603
VMware Infrastructure
discovery 603
VMware SDK 603
volume
renaming 204
usage 239
volume mirrors
remove 212

volumes
add snapshot to a refresh schedule 237
configure break-off mirror volume snapshot 220
configure instant snapshot 214
creating 192
deleting 202
disable fastresync 243
enable fastresync 242
moving 203
reactivate 201
recover 200
refresh snapshot 232
remove snapshot from refresh schedule 238
resize 225
restore data from snapshot 230
stopping 200
vomadm
about 641
Business application management 644
deployment 643
domain management 645
host management 643
listing enclosures 642
service management 644
VSystems
managing 396
migrating 398
starting 396
stopping 397
vxlist 648

W

wac resource
taking offline 422
Web services API
base URL 626
extended attributes 626
logging in
cURL 627
XPRTLC 627
logging out
cURL 629
XPRTLC 629
operation examples 634
XPRTLC cURL 636
operations 633
output examples 639
supported objects 629

X

xdistc 656

Z

Zone Agentlet 610