InfoScale™ Operations Manager 9.1 User's Guide

arctera™

Last updated: 2025-12-12

Legal Notice

Copyright ©2025 Arctera US LLC. All rights reserved.

Arctera and the Arctera Logo are trademarks or registered trademarks of Arctera US 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 Arctera 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 Arctera product or available at:

https://www.arctera.io/license-agreements

The product described in this document is distributed under licenses restricting its use, copying, distribution, and de-compilation/reverse engineering. No part of this document may be reproduced in any form by any means without prior written authorization of Arctera US 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. ARCTERA US 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 Arctera 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.

Arctera US LLC | www.arctera.io

Legal Notice for Artificial Intelligence (AI) Technology

Arctera Incorporates artificial intelligence (AI) technology to implement certain features that use data submitted by a customer using the `vxvm-stats-collector.service` service (“Customer data”). These features are disabled by default and must be enabled for use in "Anomaly Detection using AI" functionality. More details may be found in the Arctera Product InfoScale documentation (the “Documentation”) in the "Anomaly detection and monitoring" section of the *Arctera InfoScale™ Operations Manager User's Guide*.

Your organization ("Customer") agrees to the following terms and conditions when enabling the use of AI in this service:

AI-generated output may contain errors and unexpected results. Customers must use caution when relying upon responses and validate all output before use.

The Customer must not do the following:

- Include any Personal Data in its prompts.
- Use AI to create content that is illegal, confidential, harmful, misleading, or violates third-party rights or privacy.
- Use AI in Arctera products/s to make decisions that have a call for human judgment, including uses that may have health or safety consequences, without intervention by a human operator who will make the final decision.

Customer must use AI feature in Arctera products/s solely as integrated with Arctera products/s and only for purposes consistent with the intended functionality of the AI features. Arctera makes no representations and provides no warranties about the completeness, reliability, or accuracy of AI-generated output.

Technical Support

Technical Support maintains support centers globally. Technical Support's primary role is to respond to specific queries about product features and functionality. The Technical Support group also creates content for our online Knowledge Base. The Technical Support group works collaboratively with the other functional areas within the company to answer your questions in a timely fashion.

Our support offerings include the following:

- A range of support options that give you the flexibility to select the right amount of service for any size organization
- Telephone and/or Web-based support that provides rapid response and up-to-the-minute information
- Upgrade assurance that delivers software upgrades
- Global support purchased on a regional business hours or 24 hours a day, 7 days a week basis
- Premium service offerings that include Account Management Services

For information about our support offerings, you can visit our website at the following URL:

www.arctera.io/support

All support services will be delivered in accordance with your support agreement and the then-current enterprise technical support policy.

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 latest documentation is available here:
https://sort.veritas.com/arctera/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:
productdocs@arctera.io

You can also see documentation information or ask a question on the Arctera community site:
https://vox.veritas.com/category/arctera-discussions/discussions/infoscale

Services and Operations Readiness Tools (SORT)

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 26

Chapter 1 Introducing InfoScale Operations Manager 27

About InfoScale Operations Manager 27
Connecting to the InfoScale Operations Manager console 28
About increasing efficiency and productivity through storage virtualization 29
Storage Foundation: a complete solution for online storage management 30
InfoScale Operations Manager: tying it all together 31
About centralized management and monitoring 31
About discovery, monitoring, and management 32

Chapter 2 Using the Management Server console 33

About the Management Server console 33
About the Management Server console Home page 34
About the Global Dashboard 36
About the elements in a perspective view 37
About viewing the summarized information 39
About selecting the objects 40
Moving fields to the properties pane of the Management Server console 40
About drilling down to more information on the selected object 41
About searching for objects 42
Creating new search queries 43
Managing saved search queries 45
Creating and saving searches as Smart Folders 46
Managing Smart Folders 47
About viewing the solutions 48
About viewing the reports 48
About viewing tasks 49
About viewing connectivity graphs 50
About accessing the InfoScale Operations Manager Help 51

Chapter 3 Examples for using InfoScale Operations Manager 52

Example: Creating a volume using InfoScale Operations Manager 53
Example: Creating a service group, adding it to a cluster, and bringing it online using InfoScale Operations Manager 54
Example: Cluster Server troubleshooting using InfoScale Operations Manager 57
Service group dependencies – on resources and other service groups 57
Service groups faults 58
Clearing faults for the service group's resources 58
InfoScale Operations Manager reports for Cluster Server troubleshooting 59
Example: Ensuring the correct level of protection for volumes controlled by Storage Foundation 62
Mirroring: Protection from hardware failures 63
Snapshots: Backing up to prevent the loss of data 64
Example: Improving the availability and the disaster recovery readiness of a service group through fire drills 68
High availability fire drill - Ensuring the high availability capabilities of a service group 69
Disaster recovery fire drill - Ensuring the disaster recovery readiness of a service group 70
Example: Identifying the performance issues of an application using InfoScale Operations Manager 72
Example: Volume migration using InfoScale Operations Manager 75
Examples: Identifying and reducing storage waste using InfoScale Operations Manager 78
Reclaiming thin storage - example 79
Compressing files - example 81
Deduplicating file systems - example 85

Section 2 Managing InfoScale Operations Manager 90

Chapter 4 Managing user access 91

Creating an Organization 91
Create Organization panel options 93

Create Organization - Select an object panel options 93
Create Organization - Based on a rule panel options 94
Modifying the name of an Organization 95
Modify Organization panel options 96
Deleting an Organization 96
Moving an object to an Organization in a perspective 97
Assigning permissions to user groups on an Organization within a perspective 97
Modifying permissions assigned to user groups on an Organization within a perspective 98
Deleting permissions assigned to user groups on an Organization within a perspective 99
Modifying permissions assigned to user groups on an object within a perspective 99
Verifying a user group in the domain 100
Viewing permissions information 101
Viewing the permissions assigned on a perspective, an Organization, or on an object 102

Chapter 5 Setting up fault monitoring 104

About alerts and rules 104
Creating rules in a perspective 105
Create Rule - Select the type of fault conditions to trigger this rule panel options 107
Create Rule - Select one or more fault topics which will trigger this rule panel options 108
Create Rule - Select organizations panel options 108
Create Rule - Setup notifications panel options 109
Create Rule - Enter name and description panel options 110
Editing rules in a perspective 111
Edit Rule - Select the type of fault condition to trigger this rule panel options 112
Edit Rule - Select one or more fault topics which will trigger this rule panel options 112
Edit Rule - Select organization panel options 113
Edit Rule - Setup notifications panel options 113
Edit Rule - Enter name and description panel options 114
Deleting rules in a perspective 115
Delete Rule panel options 115
Enabling rules in a perspective 116
Enable Rule panel options 116
Disabling rules in a perspective 116

Disable Rule panel options 117
About faults and risks 117
Suppressing faults in a perspective 117
Suppress Faults panel options 118
Restoring a suppressed fault in a perspective 119

Chapter 6 Using reports 120

About reports 120
About using reports 122
Running a report 123
Select scope to run report panel options 123
Saving a report 124
Subscribing for a report 125
Subscribe panel options 125
Editing a report subscription 126
Deleting a report subscription 127
Sending a report through email 127
Email report panel options 128
Viewing my report subscriptions in a perspective 128
Viewing all the report subscriptions in a perspective 129
About the reports available in InfoScale Operations Manager 129
About global reports 133
Adding management servers in the Global Reports perspective 134
Creating a report 136

Section 3 Managing hosts 138

Chapter 7 Overview 139

About performing Storage Foundation and replicator operations 139
About Storage Foundation operations not supported on Windows host 140
Viewing storage summary at the cluster level 141
Viewing faults and risks at the cluster level 142

Chapter 8 Working with the uncategorized hosts 143

Working with the uncategorized hosts 143

Chapter 9 Managing File Replicator (VFR) operations 145

About performing File Replicator operations 146
Viewing the VFR option of a host 146
Creating a consistency group 147
Viewing consistency groups 147
Deleting a consistency group 148
Associating a consistency group to a replication job 148
Disassociating a consistency group from a replication job 149
Viewing consistency group properties 149
Creating a replication job 150
Viewing File Replication Jobs 151
Starting a replication job 151
Pausing a replication job 151
Resuming a replication job 152
Stopping a replication job 152
Syncing a replication job 152
Modifying a replication job 153
Deleting a replication job 153
Viewing properties of File Replication Jobs 153

Chapter 10 Managing disk groups and disks 154

About managing disk groups 155
Creating disk groups 156
    Create Disk Group - Disk Group Specifications 157
    Select Disks panel options 160
    Filter Criteria panel options 162
    Create Disk Group - Rename disks panel options 162
Recovering disk groups 163
Deporting disk groups 164
Destroying disk groups 164
Importing disk groups 165
    Import disk group panel options 166
Adding disks to disk groups 166
    Select Disk Group panel options 168
Resizing disks in disk groups 168
    Resize Disk panel options 170
Renaming disks in disk groups 171
    Rename Disk panel options 172
Upgrading disk groups 173
Splitting disk groups 174
    Split Disk Group panel options 174
Moving disk groups 176

Move Disk Group panel options 176
Joining disk groups 178
Join Disk Group panel options 178
About managing disks 179
Initializing disks 180
Initialize Disk panel options 181
Replacing disks 181
Replace Disk panel options 182
Recovering disks 183
Recover Disk panel options 184
Mapping disks 184
Unmapping disks 185
Disconnecting disks 185
Removing disks from disk groups 186
Setting host prefix for disks 187
Bringing disks online 188
Online Disk panel option for making the disks online 188
Taking disks offline 188
Setting disk usage 189
Set Disk Usage panel options 190
Set Disk Usage - Windows host panel options 190
Evacuating disks 192
Evacuate Disk panel options 193
Running or scheduling Trim 194
Trim - Run or Schedule panel options 195
Rescanning disks 196

Chapter 11 Managing volumes 198

About managing Storage Foundation volumes 199
Creating Storage Foundation volumes 200
Create Volume – Select Disk Group and Disk Selection method panel options 202
Volume attributes panel options for creating volumes on UNIX or Linux hosts for specifying values 203
Volume attributes panel options for creating volumes on Windows hosts 205
Add Drive Letter, Path and Create File System details panel options 207
Encrypting existing volumes 208
Starting online migration to encrypt a volume 210
Switching plexes to verify the copied data 211
Aborting online migration 212

Committing the background changes to encrypt a volume 213
Stopping volumes 213
Recovering volumes 214
Reactivating volumes 215
Deleting volumes 215
Delete Volume panel options 216
Moving volumes 217
Move Volume panel options 217
Renaming volumes 218
Rename Volume panel options 218
Adding mirrors to volumes 219
Add mirror - Options 220
Enable FastResync option panel options 222
Add mirror - Advanced options 223
Removing the mirrors of volumes 226
Remove mirror panel options 227
Creating instant volume snapshots 229
Create Volume Snapshot - disk selection page panel options 230
Create Volume Snapshot - Advanced Options 231
Create Volume Snapshot - Instant Snapshot panel options 232
Creating space optimized snapshots for volumes 233
Create Volume Snapshot - Space Optimized Snapshot panel options 234
Creating mirror break-off snapshots for volumes 235
Create Volume Snapshot - Mirror Break-off Snapshot panel options 236
Dissociating snapshots 237
Dissociate Snapshot panel options 238
Reattaching snapshots 239
Reattach Snapshot panel options 239
Resizing volumes 240
Resize volume panel options 241
Restoring data from the snapshots of volumes 244
Restore Data From Snapshot panel options 245
Refreshing the snapshot of volumes 246
Refresh snapshot panel options 247
Configuring a schedule for volume snapshot refresh 249
Schedule operation panel options 250
Adding snapshot volumes to a refresh schedule 251
Add to existing schedule panel options 251
Removing the schedule for volume snapshot refresh 252
Remove volumes from refresh schedule panel options 253

Setting volume usage 253
Set Volume Usage panel options 253
Splitting snapshots 254
Starting synchronization of snapshots 255
Enabling FastResync on volumes 256
Enable FastResync panel options 256
Disabling FastResync on volumes 257

Chapter 12 Managing file systems 259

About managing file systems 260
Creating file systems 260
Select Volume panel options 261
Create File System - File System Options 262
Configure SecureFS immutable checkpoints 265
Validate SecureFS checkpoints 271
Recover SecureFS checkpoints 273
Advanced Options panel 276
Enabling change logs 278
Disabling change logs 279
Synchronizing change logs 280
Removing change logs 280
Defragmenting file systems 281
Defrag file system panel options 282
Unmounting non clustered file systems from hosts 282
Unmount Confirmation panel 283
Mounting non clustered file systems on hosts 283
Mount Options panel 284
Advanced Mount Options panel options 285
Unmounting clustered file systems 287
Unmount Clustered File System panel options 288
Mounting clustered file systems on hosts 288
Mount File System - Clustered Mount Options 289
Remounting file systems 290
Remount Options panel 291
Checking file systems 292
Check File System panel options 293
Creating file system snapshots 293
Create Snapshot - Create File System Snapshot panel options 294
Snapshot level selection panel options 294
Create File System snapshot - Configure Options 295
Remounting file system snapshot 296

Remount File System Snapshot panel options 297
Mounting file system snapshot 298
Mount File System Snapshot panel options 298
Unmounting file system snapshot 299
Unmount File System Snapshot panel options 300
Removing file system snapshot 300
Remove File System Snapshot panel options 301
Monitoring capacity of file systems 302
Monitor Capacity panel options 303

Chapter 13 Managing SmartIO 305

About managing SmartIO 306
About write-back caching in SmartIO 307
Enabling or disabling SmartIO caching 307
Creating a cache 308
Create Cache panel options 309
Viewing the cache details 310
Viewing the SmartIO Impact analysis chart 311
Changing SmartIO mode 312
Modifying a cache 313
Modify Cache panel options 313
Deleting a cache 316
Loading files to a cache 316
Pinning tablespaces or files to a cache 317
Unpinning tablespaces or files from a cache 317
About using SmartAssist 318
Creating an I/O trace log 318
New I/O Trace Log panel options 319
Viewing I/O trace logs 320
Analyzing an I/O trace log 320
SmartAssist Analysis Options pane 321
Removing an I/O trace log 321

Chapter 14 Managing application IO thresholds 323

About managing application I/O workloads with IOPS settings 323
Managing IO Thresholds 324
Setting IO Thresholds 325
Modifying an App VG 326
Viewing live IOPS charts 327

Chapter 15 Managing replications 328

About managing replications 329
Configuring Storage Foundation replications 329
Configure replication panel options 330
Configure replication - Create resources 334
Configure replication - Create primary result panel options 335
Configure replication - Consistency check on secondary host panel options 335
Configure replication - Replication settings for secondary host panel options 336
Configure replication - Add secondary result panel options 339
Adding a Secondary 340
Pausing the replication to a Secondary 341
Pause replication panel options 341
Resuming the replication of a Secondary 342
Resume replication panel options 342
Starting replication to a Secondary 343
Start replication panel options 343
Stopping the replication to a Secondary 344
Stop replication panel options 344
Switching a Primary 345
Switch replication panel options 346
Taking over from an original Primary 346
Takeover Primary panel options 347
Associating a volume 348
Associate Volume - Volume Selection panel options 349
Associate Volume - Advanced Options panel options 349
Resynchronizing a Secondary 350
Removing a Secondary 350
Remove secondary panel options 351
Unconfiguring replication 351
About setting alerts for replication 352
Monitoring replications 352
Monitor replication panel options 353
Implementation of new faults 354
Support for global VVR replication monitoring thresholds 355

Chapter 16 Optimizing storage utilization 358

About reclaiming thin storage 358
Performing thin reclamation on file systems or disks 359
Compressing files 360
About file compression in InfoScale Operations Manager 360

Selecting directories for compression 362
Setting up compression schedules 362
Starting compression on demand 364
Deduplicating file systems 365
About file system deduplication 366
About deduplication chunk size 367
Implementing deduplication for a file system 368
Starting deduplication for a file system 371
Disabling or removing deduplication for a file system 371

Section 4 Managing high availability and disaster recovery configurations 373

Chapter 17 Overview 374

About high availability and disaster recovery operations 374
Pre-requisites for performing high availability and disaster recovery operations 375
About attributes of Availability objects 375
About Virtual Business Services 376

Chapter 18 Managing clusters 377

About managing clusters 377
Opening a cluster configuration 378
Saving a cluster configuration 378
Closing a cluster configuration 379
Editing attributes of a cluster 379
Edit attribute options 380
Importing a type definition 381

Chapter 19 Managing service groups 383

About managing service groups 384
Creating service groups 385
Create Service Group options 386
Configure System List options 387
Configure Resources options 388
Resource Dependencies options 389
Enabling service groups 390
Disabling service groups 391
Creating Atleast Count dependencies for a resource in a service group 392

About Atleast Count dependency 393
Autoenabling service groups 394
Freezing service groups 394
Unfreezing service groups 395
Flushing service groups 396
Enabling all resources of service groups 397
Disabling all resources of service groups 398
Deleting service groups 398
About linking service groups in a cluster 399
Linking service groups in a cluster 399
Link Service Group options 400
Unlinking service groups 402
About site aware service group operations 403
Prerequisites for using site-related service group operations 404
Limitations of site-related service group operations 404
Bringing service groups online 404
Online service groups panel options 405
Taking service groups offline 406
Offline service groups panel options 407
Switching service groups 408
Switch service groups panel options 409
Clearing faults on service group 410
Clearing the resources in a service group from the Admin Wait state 411
Editing attributes of service groups 412
Modifying the system list for a service group 413
About dependency views 414
Viewing the service group dependency view 415
About modifying a service group 416
About Cluster Server service group alerting and failover reporting 416
Viewing VCS Failover Duration report 417

Chapter 20 Managing systems 419

About managing systems 419
Freezing a system 419
Unfreezing a system 420
Editing attributes of a system 421
Starting the Cluster Server high availability daemon on the hosts in a cluster 421
Stopping the Cluster Server high availability daemon on the systems in a cluster 422

Chapter 21 Managing resources 424

About managing resources 425
Enabling resources 425
Disabling resources 426
Deleting resources 426
Clearing faults on resources 427
Probing resources 428
Taking a resource offline and propagating the state 429
Bringing resources online 430
Taking resources offline 430
Invoking a resource action 431
    Invoke Action options 432
Editing attributes of a resource 433
Editing attributes of a resource type 434
Linking resources in a service group 435
Unlinking resources in a service group 435
Adding or modifying resources 436
Marking a resource as critical 437
Marking a resource as non critical 438
Viewing the resource dependency view 438

Chapter 22 Managing global cluster configurations 440

About global clusters 440
About global clusters terminology 441
About creating global clusters 442
    Prerequisites for creating global clusters 442
    Adding a remote cluster to a local cluster 443
Converting local service groups to global service groups 443
    Make Global/Local options 444
Converting global service groups to local service groups 446
About removing a remote cluster from a global cluster setup 447
    Taking the wac resource offline 447
    Removing the remote cluster from the global cluster setup 448

Chapter 23 Running fire drills 449

About high availability and disaster recovery readiness 449
About high availability fire drills 450
Running the high availability fire drill 451
About disaster recovery fire drills 452
About configuring a fire drill service group 452
Running the disaster recovery fire drill 453

Select remote clusters and global service groups panel options 454
Schedule panel options 455
Summary panel options 456
Editing a fire drill schedule 456
Edit Schedule panel options 457
Deleting fire drill schedules 459
Enabling fire drill schedules 459
Disabling fire drill schedules 460
Viewing fire drill schedules 461

Chapter 24 Using recovery plans 463

About recovery plans 463
Creating recovery plans 464
Create recovery plan panel options 465
Editing recovery plans 468
Edit recovery plan panel options 468
Running recovery plans 470
Run recovery plan panel options 471
Deleting a recovery plan 472
Viewing historical runs of recovery plans 473
Historical Runs panel options 473
Viewing properties of recovery plans 474
Recovery Plan Properties panel options 475
About recovery plan log files 475

Chapter 25 Managing application configuration 477

About Application Configuration 477
Prerequisites for application configuration 478
Launching the configure application wizard from InfoScale Operations Manager 478

Chapter 26 Multi Site Management 480

About Multi-Site Management 480
Features of Multi-Site Management 481
Prerequisites of Multi-Site Management 482
Limitations of Multi-Site Management 482
Setting up a campus cluster 482
Setting up a replicated data cluster 484
Configuring stretch sites 484
Enclosure site assignment panel options 486

Select cluster type panel options 486
Assign site to the systems panel options 487
Site fencing preference panel options 487

Appendix A List of high availability operations 489

Cluster operations 489
System operations 490
Service group operations 491
Resource operations 494

Section 5 Monitoring Storage Foundation HA licenses in the data center 496

Chapter 27 Managing licenses 497

About licenses 497
About Arctera licensing and pricing 498
About the Symantec Performance Value Unit 501
About the Symantec Performance Value Unit for VMware virtual machines 502
About the Symantec Performance Value Unit for Solaris LDOM virtualization server 504
About the Symantec Performance Value Unit for kernel-based virtual machines 505
About the Symantec Performance Value Unit for IBM LPAR 506
About the per-core licensing 506
About assigning price tiers to hosts 507
About license deployment policies 508
Assigning a price tier to a host automatically 508
Assigning a price tier to a host manually 508
Select tier values panel options 509
Select hosts to apply same tier values panel options 510
Creating a license deployment policy 510
Create policy - Details panel options 511
Modifying a license deployment policy 513
Edit thresholds - Details panel options 513
Deleting a license deployment policy 514
Notification about a corrupted VxFS file system 515
Certificate expiry detection and renewal feature 519

Chapter 28 Viewing deployment information 522

Viewing the overview of SFHA licenses in the data center 522
About chargeable deployed licenses 524
Viewing the deployment details 525
Viewing the deployment policy details in the data center 526
Viewing the InfoScale Operations Manager Deployment Report 527

Chapter 29 Monitoring performance 529

About performance metering statistics 530
About metered resources 530
About space estimation for data logs 534
Enable performance metering for a host 536
Disable performance metering for a host 537
Enable performance metering for a virtualization server 537
Disable performance metering for a virtualization server 538
About InfoScale Operations Manager performance graphs 539
Pre-requisite commands to view performance graphs for a resource 540
Viewing the performance graphs for a host 541
Viewing the performance graphs for a disk 542
Viewing the performance graphs for volume and file system 544
Viewing the performance graphs for a path 546
Viewing the performance graphs for an initiator 547
Viewing the performance graphs for virtualization server and virtual machines 548
Viewing the performance graphs for a path of a virtualization server 549
Viewing the performance graphs for an enclosure 550
About threshold settings 552
Adding threshold settings for an object 553
Deleting the threshold settings for an object 557
Enabling the threshold settings for an object 560
Disabling the threshold settings for an object 563

Chapter 30 Anomaly detection and monitoring 566

About anomaly detection 566
Requirements 567
Enabling anomaly detection 567
Training the ML engine 568
Viewing graphs and faults for anomalies 568

Chapter 31 Managing Business Applications 571

About Business Applications in InfoScale Operations Manager 572
Creating or modifying a Business Application 572
Renaming a Business Application 574
Deleting a Business Application 574
Viewing Business Applications in the data center 574
Viewing the overview of a Business Application 575
Viewing service availability for a Business Application 576
Viewing data availability for a Business Application 577
Viewing SAN connectivity for a Business Application 578
About the makeBE script 579
Objects used in the makeBE script 580
makeBE script CSV file details 580
makeBE script parameters details 582
Limitations of makeBE script 582
makeBE script log files 583
Creating Business Application using the makeBE script 583
Importing Business Application using the makeBE script 584
Exporting Business Application using the makeBE script 585
Updating Business Application using the makeBE script 586
Deleting Business Application using the makeBE script 586

Chapter 32 Managing extended attributes 588

About using extended attributes 588
Setting values to the extended attributes on an object 589
Set Extended Attributes panel options 590
Searching objects to set extended attribute values 590
Set Extended Attributes panel options 591
Modifying the extended attributes value on an object 592

Chapter 33 Managing policy checks 593

About policy checks 593
How signature registration settings work 594
Registering policy signatures 595
Unregistering a signature 596
Setting signature tunables 597
Running a manual policy scan 598
Enabling or disabling policy signatures 598
Viewing policy violation details 599
Viewing or exporting a list of available policy signatures 601
About using custom signatures for policy checks 601

Creating a custom signature script 602
Sample custom signature script 603
Installing a custom signature script 605
Copying the custom signature script 605
Removing a custom signature 606
About using the Distribution Manager Add-on to bundle custom signature scripts 606
Sample setup.pl script for the custom signature 606
Sample unsetup.pl script for a custom signature 607

Chapter 34 Managing Dynamic Multipathing paths 609

About Dynamic Multi-Pathing in InfoScale Operations Manager 609
Disabling the DMP paths on the initiators of a host 610
Define new case panel options 611
Object selection panel options 614
Path disable panel options 615
Paths disable output summary 616
Disabling the DMP paths on an enclosure 616
Disabling the DMP paths on a virtualization server 617
Re-enabling the DMP paths 618
Path re-enable panel options 619
Paths re-enable output summary 620
DMP Maintenance Result Summary panel 620
Removing a completed DMP maintenance case record 621
Reviewing the output and results of a completed DMP maintenance case 622

Chapter 35 Managing CVM clusters 623

About monitoring and managing CVM clusters in InfoScale Operations Manager 623
Permissions required for views and operations on CVM cluster objects 625

Chapter 36 Managing Flexible Storage Sharing 627

Implementing Flexible Storage Sharing with InfoScale Operations Manager 627
About Flexible Storage Sharing 628
Flexible Storage Sharing use cases 629
Flexible Storage Sharing features and support in InfoScale Operations Manager 631

Exporting and un-exporting disks for Flexible Storage Sharing 632
Enabling or disabling Flexible Storage Sharing on existing shared disk groups 633

Chapter 37 Monitoring the virtualization environment 635

About the virtualization technologies supported 635
About discovering the VMware Infrastructure using InfoScale Operations Manager 637
How InfoScale Operations Manager discovers vCenter and ESX servers 637
Information that InfoScale Operations Manager discovers on the VMware Infrastructure components 638
Viewing the storage mapping information for VMware 639
About the datastores in InfoScale Operations Manager 640
About the multi-pathing discovery in the VMware environment 641
About near real-time discovery of VMware events 643
About discovering Solaris zones 645
How InfoScale Operations Manager discovers Solaris zones 645
Information that InfoScale Operations Manager discovers on Solaris zones 646
Limitations of the discovery of Solaris zones in InfoScale Operations Manager 647
About discovering logical domains in InfoScale Operations Manager 648
How InfoScale Operations Manager discovers Solaris logical domains 649
Information on logical domains that InfoScale Operations Manager discovers 650
Limitations of the discovery of logical domains in InfoScale Operations Manager 650
Viewing the storage mapping information for LDoms 651
About discovering LPARs and VIOs in InfoScale Operations Manager 652
About LPAR storage correlation supported in InfoScale Operations Manager 654
About Microsoft Hyper-V virtualization discovery 655
Virtual machine discovery in Microsoft Hyper-V 656
Storage mapping discovery in Microsoft Hyper-V 656
Viewing the storage mapping information for Hyper-V 657

About the Kernel-based Virtual Machine (KVM) virtualization discovery in InfoScale Operations Manager 658
About the reports related to virtualization 659

Chapter 38 Using Web services API 660

About using InfoScale Operations Manager Web services API 660
Logging in to InfoScale Operations Manager Web services API 661
Logging out of InfoScale Operations Manager Web services API 662
About objects supported by InfoScale Operations Manager Web services API 663
About performing operations using InfoScale Operations Manager Web services API 666
Examples of performing operations using InfoScale Operations Manager Web services API 669
Examples of performing operations using XPRTLC and cURL 674
Examples of the output in JSON format 681

Chapter 39 InfoScale Operations Manager command line interface 684

About the vomadm utility 684
Host management using the vomadm utility 685
Deployment management using the vomadm utility 686
Business Application management using the vomadm utility 686
Service management using the vomadm utility 687
Domain management using the vomadm utility 687
List configured schedules using the vomadm utility 688

Appendix B Command file reference 689

vxlist 690
vomadm 695
xdistc 698

Appendix C Application setup requirements 703

Application setup requirements for Oracle database discovery 703
SQL queries used for the discovery of Oracle database 706
Oracle database discovery in Solaris zones 707
Application setup requirements for Oracle Automatic Storage Management (ASM) discovery 707
SQL queries used for the discovery of Oracle ASM 708

Oracle Automatic Storage Management (ASM) discovery in Solaris zones 709
Application setup requirements for IBM DB2 discovery 709
SQL queries used for the discovery of IBM DB2 database 710
Application setup requirements for Sybase Adaptive Server Enterprise (ASE) discovery 711
SQL queries used for the discovery of Sybase Adaptive Server Enterprise (ASE) 712
Application setup requirements for Microsoft SQL Server discovery 713
SQL queries used for the discovery of Microsoft SQL Server 714
Application setup requirements for Microsoft Exchange Server discovery 714

Glossary 716

Index 721

Section 1

Getting started

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

Chapter 1

Introducing InfoScale Operations Manager

This chapter includes the following topics:

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

About InfoScale Operations Manager

InfoScale Operations Manager by Arctera 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. InfoScale Operations Manager helps administrators centrally manage diverse data center environments.

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

InfoScale Operations Manager uses two-factor authentication (2FA) to protect user accounts from getting compromised. A combination of your login credentials and a six-digit PIN is used to authenticate every user account. On gaining access to the management console, 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 InfoScale Operations Manager deployment consists of the following:

- Management Server
- Managed hosts

A InfoScale Operations Manager deployment may also discover the following:

- Virtualization environment

Connecting to the InfoScale Operations Manager console

1. On a client system that has a network connection to the host, open a Web browser.

   Your browser must be configured to accept cookies. If you are using pop-up blockers, either disable them or configure them to accept pop-ups from the host.

2. In the browser's address field, type the following URL and press Enter:

   https://hostname:14161/

   where hostname is the host name, fully-qualified host name, or IP address of Management Server.

   Example: https://myhost.example.com:14161/

3. In the **username** and **password** fields, type credentials for a valid operating system network domain account.

   The Authentication Service automatically recognizes users in the domain—for example, unixpwd or NT—on which the Authentication Broker host is a member.

4. Click **Login**.

5. On the screen that appears, set the six-digit pin for two-factor authentication:

   a. Enter a six-digit pin of your choice in the first text box.

   b. Re-enter the same pin in the second box to confirm the pin.

   c. Enter your email address in the third box. The OTP is sent to this email address if you forget your PIN.

      Provide an accurate email address so that you are not locked out of your account if you forget your PIN.

   d. Confirm your email address by re-entering it in the fourth box.

   e. Click **Validate**. The InfoScale Operations Manager dashboard is displayed.

Resetting the six-digit two-factor authentication PIN

If you forget your six-digit PIN that you had set for accessing the InfoScale Operations Manager dashboard, you can reset your PIN by using the following steps:

1. Access the InfoScale Operations Manager login screen.
2. Enter your user credentials and click **Login**.
3. When prompted for the six-digit PIN, click **Reset**.

   An OTP is sent to the registered email address.
4. On the screen that appears, set the six-digit Pin for two-factor authentication:

   a. Enter a six-digit pin of your choice in the first text box.

   b. Re-enter the same pin in the second box to confirm the pin.

   c. Enter your email address in the third box. The OTP is sent to this email address if you forget your PIN.

      Provide an accurate email address so that you are not locked out of your account if you forget your PIN.

   d. Confirm your email address by re-entering it in the fourth box.

   e. Enter the OTP that you received on your registered email address in the last box.

   f. Click **Validate**. The InfoScale Operations Manager dashboard is displayed.

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 InfoScale Operations Manager" on page 27.

Storage Foundation: a complete solution for online storage management

TheStorage Foundation products combine the industry-leading Volume Manager (VxVM), File System (VxFS), 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 InfoScale Operations Manager" on page 27.

InfoScale Operations Manager: tying it all together

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 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 31.

See "About discovery, monitoring, and management" on page 32.

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 InfoScale Operations Manager, all these tasks can be performed at a central point: the 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 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 InfoScale Operations Manager.

The 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 "InfoScale Operations Manager: tying it all together" on page 31.

About discovery, monitoring, and management

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 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 InfoScale Operations Manager. Administrators can select the faulted disk in InfoScale Operations Manager, drill down to the details of the disk, and manage the recovery of the disk.

See "InfoScale Operations Manager: tying it all together" on page 31.

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 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 37.

See "About the Management Server console Home page" on page 34.

See "About viewing the solutions" on page 48.

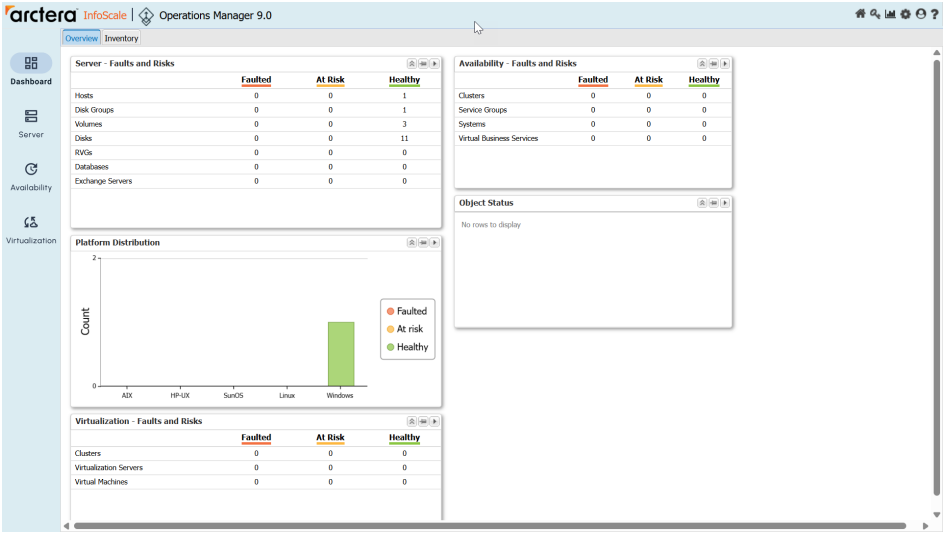
See "About searching for objects" on page 42.

See "About accessing the InfoScale Operations Manager Help" on page 51.

About the Management Server console Home page

The following figure 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:

- **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 36.

- **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.
  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:
  - **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.
  - **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).
  - **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.
- **Licensing**: Provides the information on Storage Foundation High Availability licenses, the license deployment details, and reports.
- **Global Reports**: Lets you view the collection of InfoScale deployment data from multiple management servers. You can access the Global Reports perspective from the top right corner of the menu bar. The Global Reports perspective lets you manage subscriptions for the available reports.

See "About searching for objects" on page 42.

See "About selecting the objects" on page 40.

See "About viewing the summarized information" on page 39.

See "About accessing the InfoScale Operations Manager Help" on page 51.

About the Global Dashboard

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

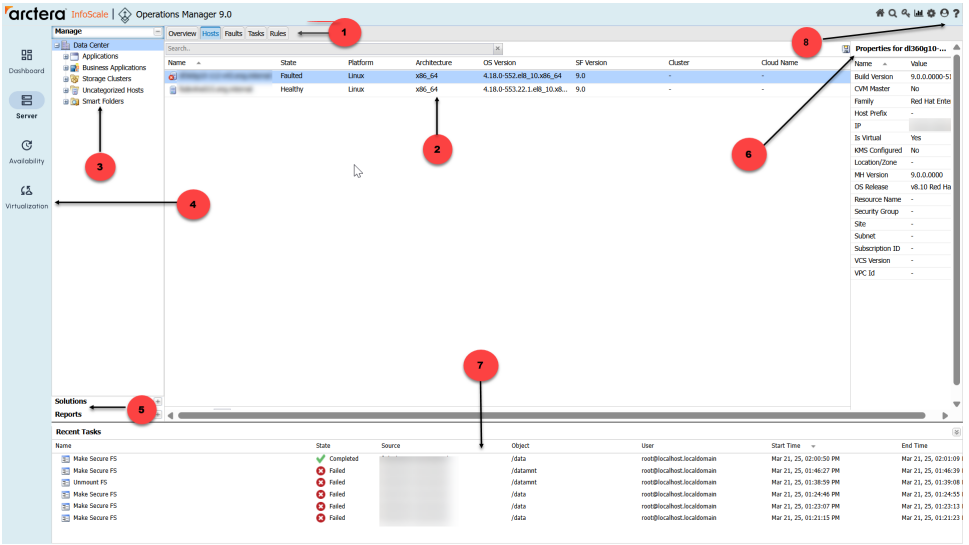
If you want to...	Do the following...
Create a custom dashboard	From any **Overview** tab in a perspective, on the top right of a chart or table, click the pin icon. You can hover on the icon to view the label **Add to Dashboard**. In the **Add to Dashboard** window, type the name of a new dashboard 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 33.

About the elements in a perspective view

The following graphic displays the console of the **Availability** perspective. Similar user interface elements are used in other views of the InfoScale Operations Manager Management Server console.

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



The user interface elements are described below:

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

Number	User interface element
1	Tabs: Lists the detailed information about a specific object.
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, for the Availability perspective, clusters are placed in the uncategorized folder.
4	Perspectives: Lists the available perspectives.

Table 2-2 User interface elements of the perspective on the Management Server console *(continued)*

Number	User interface element
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	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 40.
7	Tasks Pane: Displays the list of recent tasks with the status.
8	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 InfoScale Operations Manager Help" on page 51. See "About searching for objects" on page 42.

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

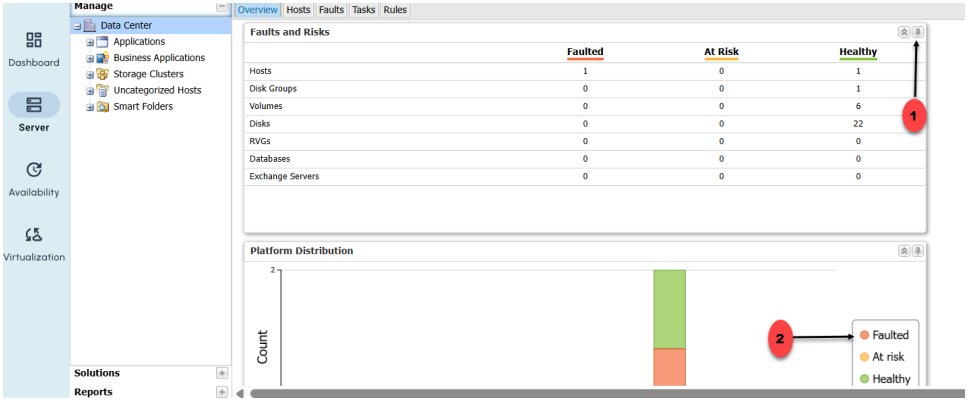


Table 2-3 User interface elements

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.

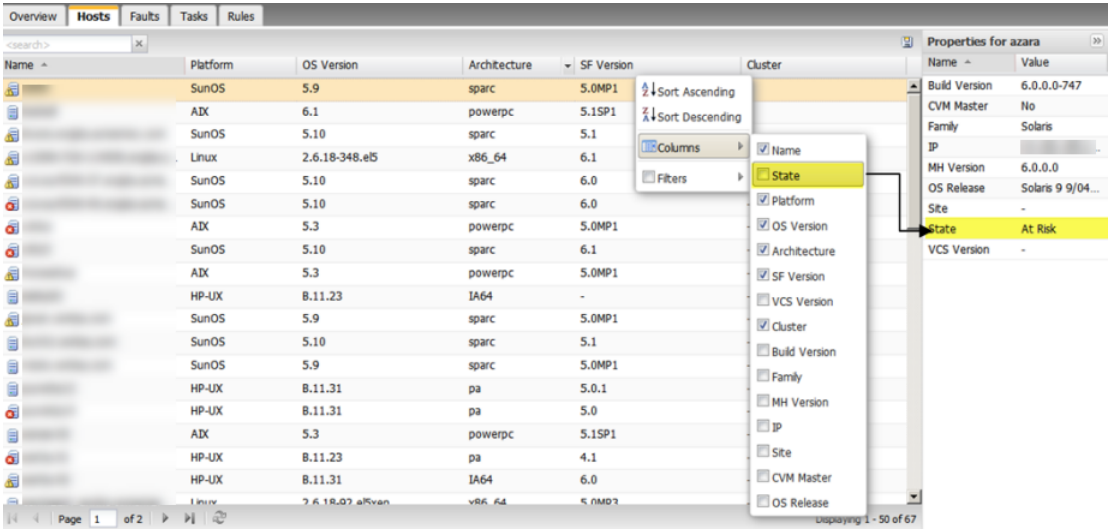
- **Navigation tree**: In the navigation tree, select the required object, and right-click to view its context-sensitive menu.
- **Content pane**: In the content pane, select the required record, and right-click to view the available options. Press Ctrl or Shift for the selection of multiple records.

See "About the Management Server console" on page 33.

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 33.

See "About the Management Server console Home page" on page 34.

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.

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 40.

See "About the Global Dashboard" on page 36.

See "About the Management Server console Home page" on page 34.

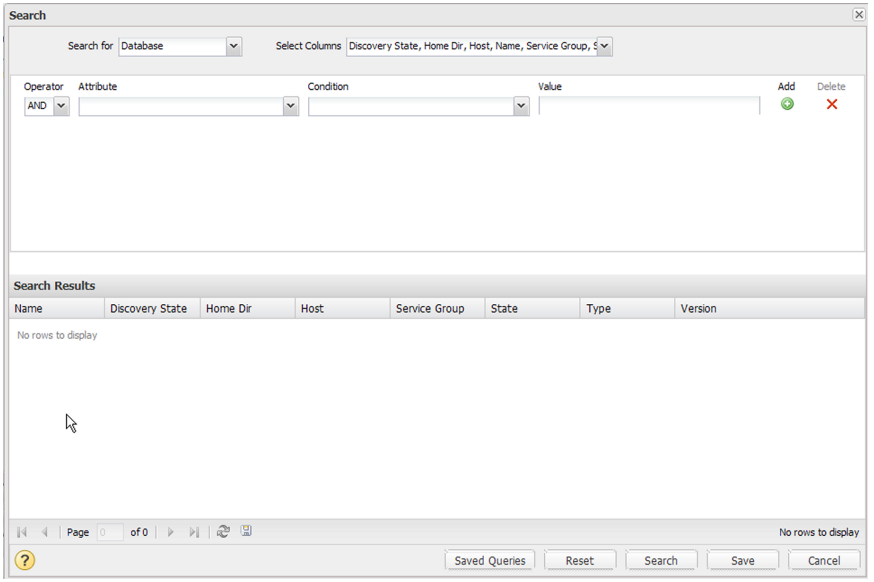
About searching for objects

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.
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 43.

See "Managing saved search queries" on page 45.

See "Creating and saving searches as Smart Folders" on page 46.

See "Managing Smart Folders" on page 47.

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 46.

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, 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 42.

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 45.

See "To rename a saved search" on page 45.

See "To delete a saved search" on page 45.

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 42.

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 43.

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 42.

See "Managing Smart Folders" on page 47.

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 47.

See "To rename a Smart Folder" on page 47.

See "To copy a Smart Folder" on page 47.

See "To delete a Smart Folder" on page 47.

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**.

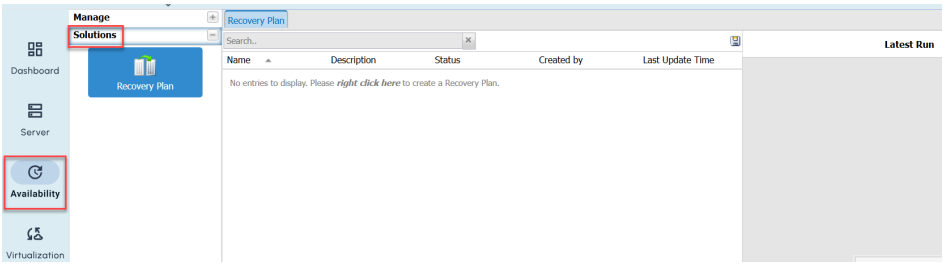
See "About searching for objects" on page 42.

See "Creating and saving searches as Smart Folders" on page 46.

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 following graphic, Recovery Plan is shown as the available solution for the Availability perspective.

Figure 2-6 Solutions for the Availability perspective



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**.

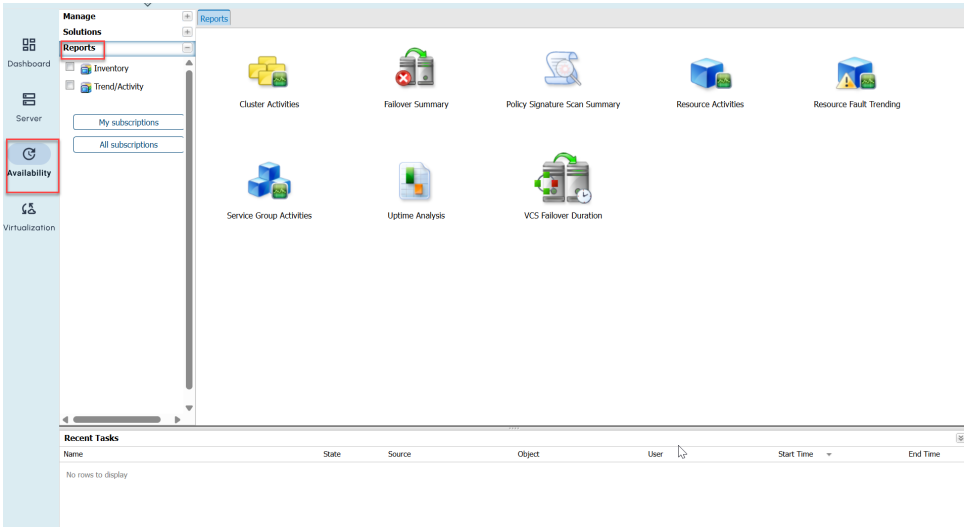
See "About the Management Server console" on page 33.

See "About the Management Server console Home page" on page 34.

About viewing the reports

You can generate and view the reports that are specific to the perspective. The following graphic displays the reports view of the Availability perspective.

Figure 2-7 Management Server reports for the Availability perspective



See "About the Management Server console" on page 33.

See "About the Management Server console Home page" on page 34.

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 42.

See "About the Management Server console" on page 33.

See "About the Management Server console Home page" on page 34.

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 perspective.

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.

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 33.

About accessing the 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 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 InfoScale Operations Manager Help add-on on the Management Server. The 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 InfoScale Operations Manager Management Server and its supported add-ons.

See "About the Management Server console" on page 33.

See "About the Management Server console Home page" on page 34.

Chapter 3

Examples for using InfoScale Operations Manager

This chapter includes the following topics:

- Example: Creating a volume using InfoScale Operations Manager
- Example: Creating a service group, adding it to a cluster, and bringing it online using InfoScale Operations Manager
- Example: Cluster Server troubleshooting using 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 InfoScale Operations Manager
- Example: Volume migration using InfoScale Operations Manager
- Examples: Identifying and reducing storage waste using InfoScale Operations Manager

Example: Creating a volume using 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. InfoScale Operations Manager provides you a single and a centralized management console for the Storage Foundation and high availability products. You can use InfoScale Operations Manager to monitor, visualize, and manage storage resources.

Storage Foundation product suite includes multiple Arctera products and optionally licensed features, like Volume Manager (VxVM), File System (VxFS), and Veritas Volume Replicator (VVR). You can use these products to manage the storage resources in your data center. For example, you can use 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 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 InfoScale Operations Manager

InfoScale Operations Manager provides you a single and centralized management console for the Storage Foundation and high availability products. You can use 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 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 *InfoScale Operations Manager Add-ons User Guide*.

Example: Cluster Server troubleshooting using 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. 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 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 57.

See "Service groups faults" on page 58.

See "Clearing faults for the service group's resources" on page 58.

See "InfoScale Operations Manager reports for Cluster Server troubleshooting " on page 59.

Service group dependencies – on resources and other service groups

In 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 InfoScale Operations Manager" on page 57.

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 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 InfoScale Operations Manager" on page 57.

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 InfoScale Operations Manager" on page 57.

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 60.

See "Uptime Analysis" on page 61.

See "Resource Fault Trending" on page 61.

See "Failover Summary" on page 61.

See "Example: Cluster Server troubleshooting using InfoScale Operations Manager" on page 57.

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 "InfoScale Operations Manager reports for Cluster Server troubleshooting " on page 59.

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 "InfoScale Operations Manager reports for Cluster Server troubleshooting " on page 59.

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 "InfoScale Operations Manager reports for Cluster Server troubleshooting " on page 59.

Failover Summary

In 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.
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 "InfoScale Operations Manager reports for Cluster Server troubleshooting " on page 59.

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, 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, 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 63.

See "Snapshots: Backing up to prevent the loss of data" on page 64.

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 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 62.

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 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, 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 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 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 62.

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 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 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 69.

See "Disaster recovery fire drill - Ensuring the disaster recovery readiness of a service group" on page 70.

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 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 68.

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 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 InfoScale Operations Manager, you can run the disaster recovery (DR) fire drill for the specific service groups that you choose in 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 68.

Example: Identifying the performance issues of an application using 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 InfoScale Operations Manager, you can get visibility from the application down to the spindle. 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 *Arctera InfoScale Operations Manager Add-ons User Guide*.

In 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.

A complete list of objects and the types of performance graphs is provided.

See "About metered resources" on page 530.

This example explains how you can use the performance graphs on the storage objects in 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

Table 3-1 Example names *(continued)*

Object	Name
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

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.

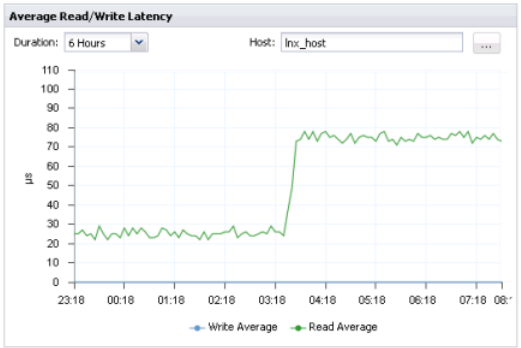
To view the performance graphs for vxfs_fs and vxvm_vol associated with ora_db

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

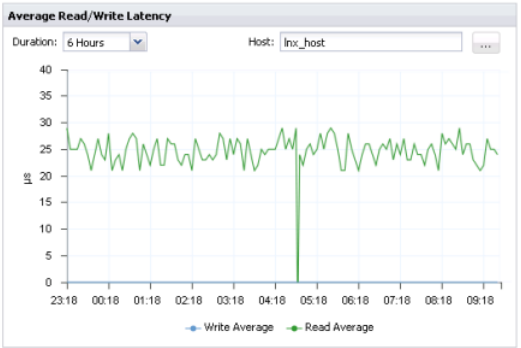
1. In the Home page on the Management Server console, go to **Server** perspective, and select **Manage** in the left pane.
2. Expand **Applications** and expand **Databases** to locate **ora_db**.
3. Click the **Disks** tab.
4. In the disks details list, right-click **xiv_disk1** or **xiv_disk2** and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

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 *Arctera InfoScale Operations Manager Add-ons User Guide*.

Example: Volume migration using 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 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 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**.

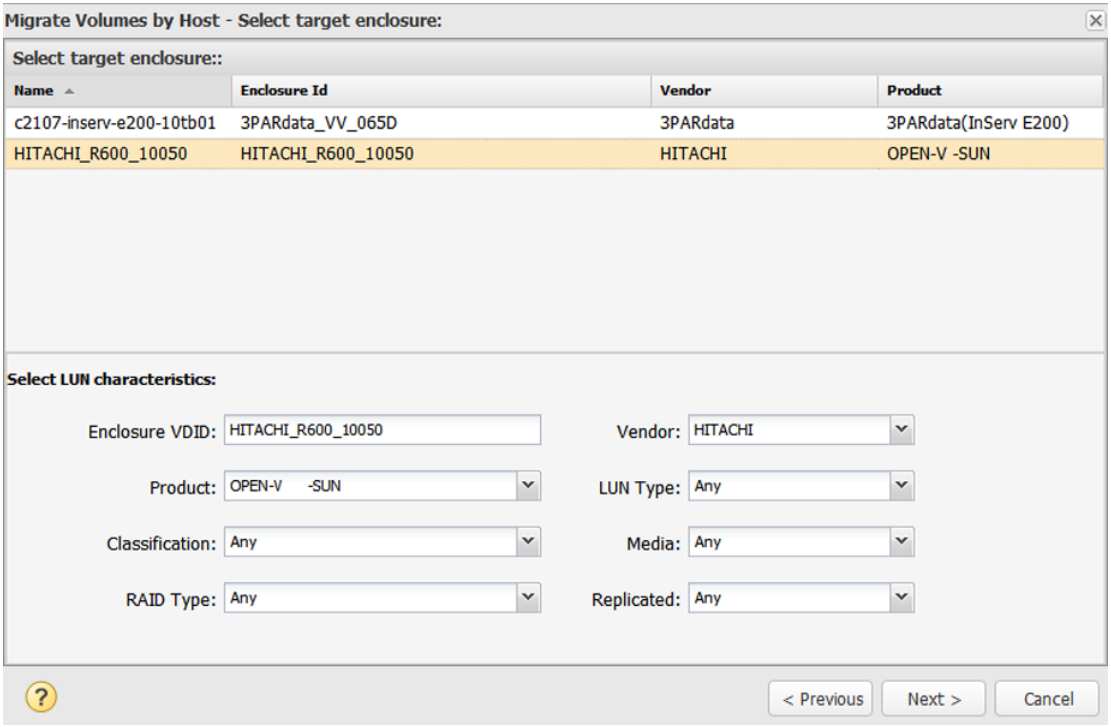
2 In the **Actions** pane, click **Migrate Volumes By Host**.

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**.

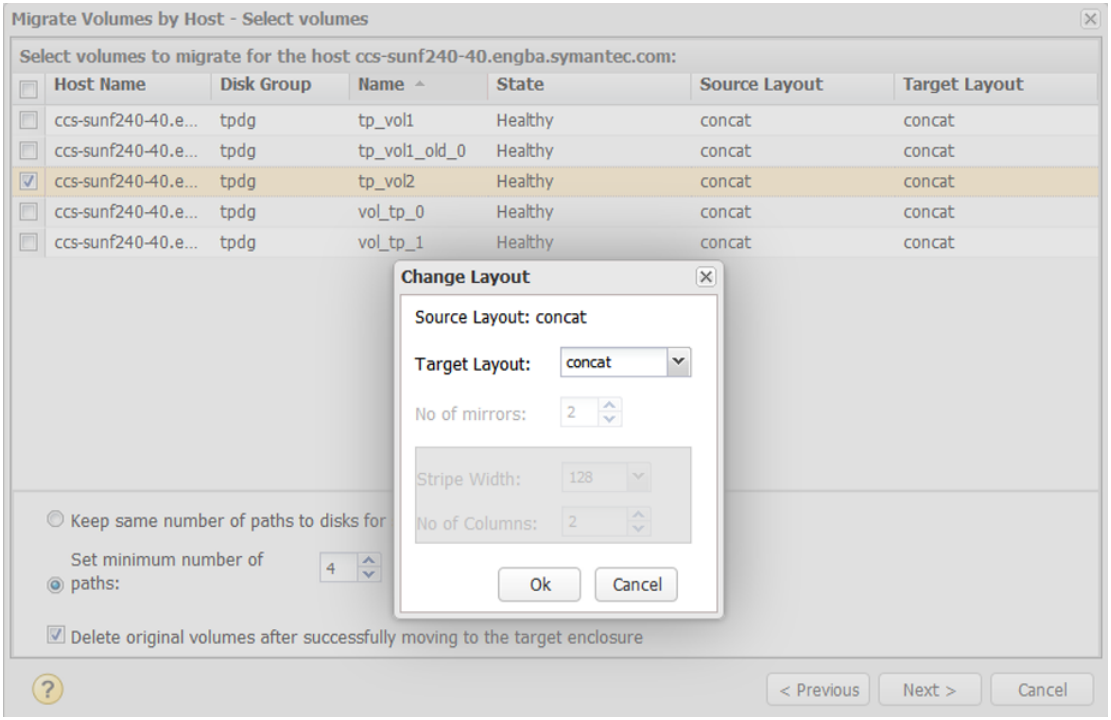
4 Select the host and then select the source enclosure. Click **Next**.

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 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 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 InfoScale Operations Manager are described under a separate use case.

See "Reclaiming thin storage - example" on page 79.

See "Compressing files - example" on page 81.

See "Deduplicating file systems - example" on page 85.

Reclaiming thin storage - example

With 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 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.

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 359.

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:

- InfoScale 8.0 or later
- File System (VxFS) disk layout version 9 or later
- InfoScale Operations Manager managed host (VRTSsfmh) version 7.4.2 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.

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. The permission on the enclosure may be explicitly assigned or inherited from a parent Organization.

See "Examples: Identifying and reducing storage waste using InfoScale Operations Manager" on page 78.

Compressing files - example

Storage Foundation 7.4.2 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 InfoScale Operations Manager" on page 360.

Using the 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 InfoScale Operations Manager

Requirements for compression

Compression has the following requirements:

- InfoScale 7.4.2 or later
- File System (VxFS) disk layout version 9 or later
- InfoScale Operations Manager managed host (VRTSsfmh) version 7.4.2 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 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 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 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 InfoScale Operations Manager" on page 78.

Deduplicating file systems - example

Storage Foundation 7.4.2 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 366.

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

Requirements for deduplication

Deduplication has the following requirements:

- InfoScale 7.4.2 or later.
- File System (VxFS) disk layout version 9 or later.
- InfoScale Operations Manager managed host (VRTSsfmh) version 7.4.2 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 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 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

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, Arctera 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. Arctera 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 InfoScale Operations Manager" on page 78.

Section 2

Managing 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
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 93.

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 93.

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 94.

See "Modifying the name of an Organization" on page 95.

See "Moving an object to an Organization in a perspective" on page 97.

See "Assigning permissions 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 99.

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 91.

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

See "Creating an Organization" on page 91.

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

Table 4-5 Create Organization panel options for selecting attributes in a perspective *(continued)*

Perspective	Attributes
Availability	■ Condition ■ Host Count ■ Name ■ Platform ■ State ■ Sub Type ■ Version
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 91.

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 96.

4 In the **Modify Organization - Result** panel click **Close**.

See "Creating an Organization" on page 91.

See "Deleting an Organization" on page 96.

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 95.

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 91.

See "Modifying the name of an Organization" on page 95.

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 91.

See "Modifying the name of an Organization" on page 95.

Assigning permissions to user groups on an Organization within a perspective

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 98.

See "Deleting permissions assigned to user groups on an Organization within a perspective" on page 99.

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

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 97.

See "Deleting permissions assigned to user groups on an Organization within a perspective" on page 99.

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 97.

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

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 97.

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 97.

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 97.

See "Modifying permissions assigned to user groups on an object within a perspective" on page 99.

See "Viewing the permissions assigned on a perspective, an Organization, or on an object" on page 102.

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 97.

See "Modifying permissions assigned to user groups on an object within a perspective" on page 99.

See "Verifying a user group in the domain" on page 100.

See "Deleting permissions assigned to user groups on an Organization within a perspective" on page 99.

See "Viewing permissions information" on page 101.

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 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 105.

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 *Arctera 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 107.

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 108.

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 108.

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

See "Create Rule - Setup notifications panel options" on page 109.

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 110.

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 104.

See "Editing rules in a perspective" on page 111.

See "Deleting rules in a perspective" on page 115.

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 105.

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 105.

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 105.

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 105.

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 105.

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 112.
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 112.
6. In the **Edit Rule - Select organizations** panel, select an organization, and click **Next**.

   See "Edit Rule - Select organization panel options" on page 113.
7. In the **Edit Rule - Setup notifications** panel, enter the required information and click **Next**.

   See "Edit Rule - Setup notifications panel options" on page 113.
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 114.
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 104.

See "Creating rules in a perspective" on page 105.

See "Deleting rules in a perspective" on page 115.

See "Enabling rules in a perspective" on page 116.

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 111.

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 111.

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 111.

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 111.

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 111.

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 115.

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

See "Creating rules in a perspective" on page 105.

See "Enabling rules in a perspective" on page 116.

See "Disabling rules in a perspective" on page 116.

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 115.

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 116.

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

See "Disabling rules in a perspective" on page 116.

See "Creating rules in a perspective" on page 105.

Enable Rule panel options

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

See "Enabling rules in a perspective" on page 116.

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 117.

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

See "Enabling rules in a perspective" on page 116.

Disable Rule panel options

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

See "Disabling rules in a perspective" on page 116.

About faults and risks

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 117.

See "Restoring a suppressed fault in a perspective" on page 119.

Suppressing faults in a perspective

Using the Management Server console, you can suppress one or more faults in 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 118.

See "Restoring a suppressed fault in a perspective" on page 119.

Suppress Faults panel options

Use this panel to suppress the faults in 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 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 117.

Restoring a suppressed fault in a perspective

You can restore the fault that is suppressed in 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, 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 117.

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 InfoScale Operations Manager
- About global reports
- Adding management servers in the Global Reports perspective
- Creating a report

About reports

In 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 122.

See "Running a report" on page 123.

See "Saving a report" on page 124.

See "Subscribing for a report" on page 125.

See "About the reports available in InfoScale Operations Manager" on page 129.

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 *Arctera 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 120.

See "Running a report" on page 123.

See "Saving a report" on page 124.

See "Subscribing for a report" on page 125.

See "Viewing my report subscriptions in a perspective" on page 128.

See "Viewing all the report subscriptions in a perspective" on page 129.

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 123.
4. Click **Run**.

See "Saving a report" on page 124.

See "Subscribing for a report" on page 125.

See "About using reports" on page 122.

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 123.

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 125.

See "Running a report" on page 123.

See "About using reports" on page 122.

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 *Arctera 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 125.
5. In the **Subscribe** panel click **Close**.

See "Editing a report subscription" on page 126.

See "Running a report" on page 123.

See "Saving a report" on page 124.

See "About using reports" on page 122.

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 125.

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 125.

5 In the **Subscribe - Result** panel, click **Close**.

See "Running a report" on page 123.

See "Saving a report" on page 124.

See "Subscribing for a report" on page 125.

See "About using reports" on page 122.

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 126.

See "Running a report" on page 123.

See "About using reports" on page 122.

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 *Arctera 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 128.

4 Click **Finish**.

See "Editing a report subscription" on page 126.

See "Running a report" on page 123.

See "About using reports" on page 122.

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 127.

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 126.

See "Deleting a report subscription" on page 127.

See "Viewing all the report subscriptions in a perspective" on page 129.

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 126.

See "Deleting a report subscription" on page 127.

About the reports available in 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	Arctera 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 Arctera 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 7.4.2 or later, Layout version 9 or later, File system disks not shared with other file systems and managed host version 7.4.2 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.
VVR status of data replicated	Provides details on the replication status of a Replicated Volume Group (RVG) at the time of report generation.
VVR Data Transferred-Weekly	Provides details on the amount of the data that is transferred using the VVR configuration over a period of one week.
VVR Data Transferred-Monthly	Provide details on the amount of data that is transferred using a VVR configuration over a period of one month.

Table 6-7 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 122.

About global reports

This InfoScale Operations Manager patch introduces a global reporting feature, which facilitates collection of InfoScale deployment data from multiple management servers.

Managing global reports

To use the global reporting feature, you access the new Global Reports perspective from the menu bar and first add the required management servers. Thereafter, you use customized SQL queries to create and run reports, which comprise data that is collected from those management servers. The Global Reports perspective also lets you manage subscriptions for the available reports.

Prerequisite

Ensure that port 5636 is open on any management server that you want to add to the Global Reports perspective. This requirement also applies to the main Central Management Server (CMS), which is the management server from which you plan to generate the global reports.

Adding management servers in the Global Reports perspective

To allow connections on a management server

1 Navigate to the database configuration directory on the management server that you want to add to the Global Reports perspective.

On Windows `systemDrive:\ProgramData\Symantec\VRTSsfmcs\db\data`

Replace *systemDrive* with the appropriate value.

On UNIX `/var/opt/VRTSsfmcs/db/data`

2 Open the `pg_hba.conf` file and add the **host SFMdb3 habguest *mainCMS_FQDN* password** line for the main InfoScale Operations Manager CMS:

```
# TYPE  DATABASE        USER    ADDRESS          METHOD

# "local" is for Unix domain socket connections only
local   all             all                      trust
# IPv4 local connections:
host    all             all     127.0.0.1/32     trust
host SFMdb3 habguest mainCMS_FQDN password
# IPv6 local connections:
host    all             all     ::1/128          trust
# Allow replication connections from localhost, by
# a user with the replication privilege.
local   replication     all                      trust
host    replication     all     127.0.0.1/32     trust
host    replication     all     ::1/128          trust
```

Note: In the **`host SFMdb3 habguest mainCMS_FQDN password`** line:

- Single spaces must be used as separators, instead of tab characters.

- The `password` string is a literal string, not a variable to be replaced with a user password.

- Use management server FQDN instead of ***mainCMS_FQDN***

---

3 Open the `postgresql.conf` file and add or update the line for ports as follows:

```
listen_addresses = '*'
```

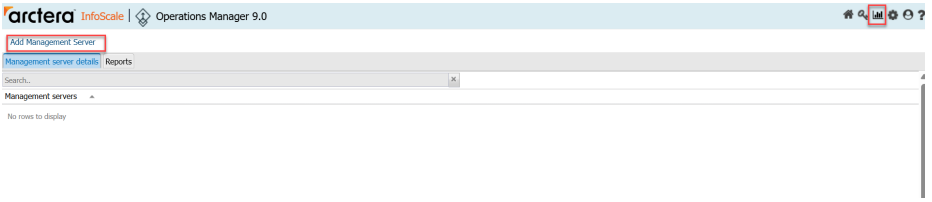
4 Restart the database.

On Windows
```
systemDrive:\Program Files\VERITAS\VRTSsfmcs\bin\vomsc --restart db
```
Replace *systemDrive* with the appropriate value.

On UNIX
```
/opt/VRTSsfmcs/bin/vomsc --restart db
```

To work with management servers in the Global Reports perspective

1 Click the **Global Reports** menu on the InfoScale Operations Manager console.



2 In the Global Reports perspective, you can use the Management server details tab to perform the following operations:

- Click **Add Management Server** above the tab title. On the Add management server dialog, enter the name or the IP address of the management server and its database password.
  To identify the `habguest` password of a management server, run the following command on that management server:

On Windows
```
systemDrive:\Program Files\Veritas\VRTSsfmh\bin\xdbadm.exe -g -u habguest -c systemDrive:\ProgramData\Symantec\VRTSsfmcs\conf
```
Replace *systemDrive* with the appropriate value.

On UNIX
```
/opt/VRTSsfmh/bin/xdbadm -g -u habguest -c /var/opt/VRTSsfmcs/conf
```

A new row for the corresponding management server then appears on the tab.

Note: Ensure that you do not add multiple entries for the same management server. Otherwise, reports that retrieve data from such a management server will contain duplicate records.

- To update the database password of a management server, click **Edit management server** details from the context menu of the corresponding row. Use the appropriate fields on the Edit management server details dialog.
- To remove a management server from this perspective, click **Delete management server** from the context menu of the corresponding row. Provide your confirmation on the dialog box that appears. After you delete a management server, the corresponding row no longer appears on the tab.

Creating a report

To create a report

1 Open the Global Reports perspective and click the **Reports** tab.

In the **Reports** tab, by default, you can see few reports from different perspectives. You can run report after assigning required management server.

2 You can perform the following operations from the left pane:

- Specify the management servers from which you want to retrieve data for the report. Servers that are added to the Global Reports perspective are listed at the bottom, select from this list.
- To filter the reports that appear in the content pane, you can select or deselect the management servers. By default, none of the management servers are selected, and the content pane displays all the available reports . When you select one or more management servers, the content pane displays only those reports that retrieve data from the corresponding servers.

3 You can perform the following operations from the contents pane:

- To run a report, click the kebab menu on the desired report in the content pane and then click **Run**. The report appears in a new window. You can open or save the report as a CSV file, email it in the HTML or the CSV format, or subscribe to receive it. Subscribed reports are scheduled, so users can receive them at specific times.
- To modify the criteria of a report, click the kebab menu on the desired report in the content pane and then click **Modify**. You can modify its

description, specify different criteria by using the appropriate SQL query values, and select or deselect the management servers on which to run the report.

- To remove a report, click the kebab menu on the desired report in the content pane and then click **Remove**. The report no longer appears in the content pane, and any subscriptions to the report are also deleted.

Section 3

Managing hosts

- Chapter 7. Overview
- Chapter 8. Working with the uncategorized hosts
- Chapter 9. Managing File Replicator (VFR) operations
- Chapter 10. Managing disk groups and disks
- Chapter 11. Managing volumes
- Chapter 12. Managing file systems
- Chapter 13. Managing SmartIO
- Chapter 14. Managing application IO thresholds
- Chapter 15. Managing replications
- Chapter 16. 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 File System (VxFS) 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 155.

See "About managing file systems" on page 260.

See "About managing disk groups" on page 155.

See "About managing replications" on page 329.

See "About managing Storage Foundation volumes" on page 199.

See "About Storage Foundation operations not supported on Windows host" on page 140.

See "Permissions required for views and operations on CVM cluster objects" on page 625.

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 139.

Viewing storage summary at the cluster level

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

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

Working with the uncategorized hosts

This chapter includes the following topics:

- Working with the uncategorized hosts

Working with the uncategorized hosts

InfoScale Operations Manager discovers the hosts and their associations to storage resources and network devices.

When a new host is added to the InfoScale Operations Manager CMS, the host is discovered and listed as an uncategorized host in the Server perspective. If the host belongs to a cluster, the cluster is displayed under **Uncategorized Cluster** in the Availability perspective.

When you click on **Uncategorized Hosts**, host details like state, platform, architecture, SF version, build version, IP address, VCS version, and so on are displayed. InfoScale Operations Manager can also discover hosts deployed in an Azure environment and accordingly display the cloud name, subscription Id, resource name, subnet Id, location details.

Note: InfoScale Operations Manager supports the discovery of InfoScale hosts in an Azure cloud environment only if the hosts are running on Linux or Windows.

Some of these attributes are displayed as columns for each uncategorized host, whereas the other attributes are displayed in the Properties area in the perspective view. You can right-click each of the attributes and select Show as column to display the selected attribute as a column in the table.

You can view attributes related to uncategorized cluster in the Availability perspective.

Chapter 9

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 InfoScale Operations Manager 7.2 or later versions, a **VFR** option has been introduced.

See "Viewing the VFR option of a host" on page 146.

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 10

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 156.

See "Renaming disks in disk groups" on page 171.

See "Resizing disks in disk groups" on page 168.

See "Removing disks from disk groups" on page 186.

See "Splitting disk groups" on page 174.

See "Joining disk groups" on page 178.

See "Moving disk groups" on page 176.

See "Importing disk groups" on page 165.

See "Deporting disk groups" on page 164.

See "Recovering disk groups" on page 163.

See "Destroying disk groups" on page 164.

See "Upgrading disk groups" on page 173.

See "Enabling or disabling Flexible Storage Sharing on existing shared disk groups" on page 633.

See "About performing Storage Foundation and replicator operations" on page 139.

See "Permissions required for views and operations on CVM cluster objects" on page 625.

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 625.

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 157.

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 162.

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 162.

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 155.

See "About managing disks" on page 179.

Create Disk Group - Disk Group Specifications

Use this wizard panel to specify the attributes for the disk group.

Table 10-1 lists the attributes that you need to specify for creating disk groups on UNIX or Linux hosts.

Table 10-2 lists the attributes that you need to specify for creating disk groups on Windows hosts.

Table 10-1 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 InfoScale Operations Manager" on page 627.
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.
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: Volumes created under encrypted disk groups are encrypted by default. Once you enable disk group encryption, it cannot be disabled. The only workaround is to create a copy of the unencrypted data onto another volume with the same layout and then replace the encrypted volume. For more information about encryption and KMS, refer to the platform-specific *Storage Foundation Administrator's Guide*.
Description	Enter an optional comment.

Table 10-2 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 156.

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 10-3 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 162.
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.

Table 10-3 Select Disks panel options *(continued)*

Field	Description
Total Size	Displays the total available size on the disk.
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 156.

See "Adding disks to disk groups" on page 166.

See "Resizing volumes" on page 240.

See "Adding mirrors to volumes" on page 219.

See "Removing disks from disk groups" on page 186.

See "Renaming disks in disk groups" on page 171.

See "Resizing disks in disk groups" on page 168.

See "Creating Storage Foundation volumes" on page 200.

Filter Criteria panel options

Use this wizard panel to filter the list of disks available on a host.

Table 10-4 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 156.

See "Removing disks from disk groups" on page 186.

See "Renaming disks in disk groups" on page 171.

See "Resizing disks in disk groups" on page 168.

See "Creating Storage Foundation volumes" on page 200.

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 10-5 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 166.

See "Creating disk groups" on page 156.

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 155.

See "About managing disks" on page 179.

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 155.

See "About managing disks" on page 179.

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 155.

See "About managing disks" on page 179.

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 166.

6 In the **Result** panel, verify that the selected disk group has been imported successfully.

See "About managing disk groups" on page 155.

See "About managing disks" on page 179.

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 10-6 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 165.

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 160.
  See "Filter Criteria panel options" on page 162.
- 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 168.

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 162.

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 155.

See "About managing disks" on page 179.

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 10-7 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 166.

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.

You can also resize disks within disk groups that are part of FSS configurations, in which disks are connected with each cluster node and accessible to other nodes through network interconnect. In this case, the disk resize operation internally invokes the `vxdisk resize` command with the appropriate options.

See "About Flexible Storage Sharing" on page 628.

For details on growing the existing storage by using the `vxdisk resize` command, refer to the platform-specific *Storage Foundation 9.0 Administrator's Guide*.

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 **More Operations** > **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 160.

   See "Filter Criteria panel options" on page 162.
6. In the **Resize Disk** wizard panel, specify an option to resize the disk. Click **Next**.

   See "Resize Disk panel options" on page 170.

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 155.

Resize Disk panel options

Use this wizard panel to enter the details for resizing the disks in the disk group.

Table 10-8 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.

Table 10-8 Resize Disk panel options *(continued)*

Field	Description
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 168.

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 160.

See "Filter Criteria panel options" on page 162.

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 172.

See "Rename Disk panel options" on page 172.

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 155.

See "About managing disks" on page 179.

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 10-9 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.

Table 10-9 Rename Disk Details panel options *(continued)*

Field	Description
Custom Disk Name	Select this option to specify a custom disk name. Enter the disk name in the corresponding field.

Table 10-10 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 171.

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 155.

See "About managing disks" on page 179.

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 174.

6 In the **Result** panel, verify that the selected disk group has been split successfully.

See "Creating disk groups" on page 156.

See "About managing disk groups" on page 155.

Split Disk Group panel options

Use this wizard panel to split a disk group to create a new one.

Table 10-11 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.
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 174.

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 176.
6. In the **Result** panel, verify that the selected disk group has been moved successfully.

See "Creating disk groups" on page 156.

See "About managing disk groups" on page 155.

Move Disk Group panel options

Use this wizard panel to move disks within a disk group to another disk group.

Table 10-12 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.

Table 10-12 Move Disk Group panel options *(continued)*

Name	Description
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.
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 176.

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 178.
6. In the **Result** panel, verify that the selected disk group has been joined successfully.

See "Creating disk groups" on page 156.

See "About managing disk groups" on page 155.

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 10-13 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 178.

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 180.

See "Replacing disks" on page 181.

See "Recovering disks" on page 183.

See "Disconnecting disks" on page 185.

See "Removing disks from disk groups" on page 186.

See "Setting host prefix for disks" on page 187.

See "Bringing disks online" on page 188.

See "Taking disks offline" on page 188.

See "Setting disk usage" on page 189.

See "Evacuating disks" on page 192.

See "Running or scheduling Trim " on page 194.

See "Rescanning disks" on page 196.

See "Exporting and un-exporting disks for Flexible Storage Sharing" on page 632.

See "About performing Storage Foundation and replicator operations" on page 139.

See "Permissions required for views and operations on CVM cluster objects" on page 625.

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 181.
6. In the **Result** panel, verify that the disks have been initialized successfully.

See "About managing disk groups" on page 155.

See "About managing disks" on page 179.

Initialize Disk panel options

Use this wizard panel to initialize a disk.

Table 10-14 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 10-15 Initialize Disk panel options for Windows hosts

Field	Description
Selected Disks	Verify the disk details selected to be initialized.
Specify format	Specify one of the following formats: ■ **MBR**: Master boot record ■ **GPT**: GUID partition table

See "Initializing disks" on page 180.

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 182.

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 183.

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 10-16 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 162.
Reset	Click to reset the selections.
Name	Name of the free disk device.
Condition	Condition of the disk.

Table 10-16 Replace Disk panel options *(continued)*

Field	Description
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 181.

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 184.

6 In the **Result** panel, verify that the selected disks have been recovered successfully.

See "Replacing disks" on page 181.

See "About managing disks" on page 179.

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 183.

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 185.

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 184.

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 155.

See "About managing disks" on page 179.

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 160.

See "Filter Criteria panel options" on page 162.

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 156.

See "About managing disk groups" on page 155.

See "About managing disks" on page 179.

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 179.

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 188.
6. In the **Result** panel, verify that the selected disk has been brought online successfully. Click **OK**.

See "Taking disks offline" on page 188.

See "About managing disks" on page 179.

Online Disk panel option for making the disks online

Use this wizard panel to make the disks online.

See "Bringing disks online" on page 188.

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 188.

See "About managing disks" on page 179.

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 190.

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 190.

6 In the **Result** panel, verify that the selected disk has been evacuated successfully. Click **OK**.

See "About managing disks" on page 179.

Set Disk Usage panel options

Use this wizard panel to select one of the following options to mark the disk.

Table 10-17 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 189.

Set Disk Usage - Windows host panel options

Use this wizard panel to select one of the following options to mark the disk.

Table 10-18 Set Disk Usage - Windows host panel options

Field	Description
Disk	Displays the name of the disk used.

Table 10-18 Set Disk Usage - Windows host panel options *(continued)*

Field	Description
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, InfoScale Operations Manager chooses space on disks marked as secondary targets.

Table 10-18 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 189.

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 193.
6. In the **Result** panel, verify that the selected disk has been evacuated successfully. Click **OK**.

See "About managing disks" on page 179.

Evacuate Disk panel options

Use this wizard panel to evacuate disks.

Table 10-19 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 162.
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 10-19 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 192.

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, 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 195.

5 In the **Result** panel, verify that the schedule has been set up successfully. Click **OK**.

See "About managing disks" on page 179.

Trim - Run or Schedule panel options

Use this wizard panel to run or schedule a Trim command.

Table 10-20 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 10-20 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 194.

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 179.

Chapter 11

Managing volumes

This chapter includes the following topics:

- About managing Storage Foundation volumes
- Creating Storage Foundation volumes
- Encrypting existing 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 "Encrypting existing volumes" on page 208.

See "Resizing volumes" on page 240.

See "Stopping volumes" on page 213.

See "Recovering volumes" on page 214.

See "Reactivating volumes " on page 215.

See "Deleting volumes" on page 215.

See "Moving volumes" on page 217.

See "Renaming volumes" on page 218.

See "Adding mirrors to volumes" on page 219.

See "Removing the mirrors of volumes" on page 226.

See "Creating instant volume snapshots" on page 229.

See "Creating space optimized snapshots for volumes" on page 233.

See "Creating mirror break-off snapshots for volumes" on page 235.

See "Dissociating snapshots" on page 237.

See "Reattaching snapshots" on page 239.

See "Restoring data from the snapshots of volumes" on page 244.

See "Refreshing the snapshot of volumes" on page 246.

See "Configuring a schedule for volume snapshot refresh" on page 249.

See "Adding snapshot volumes to a refresh schedule" on page 251.

See "Removing the schedule for volume snapshot refresh" on page 252.

See "Setting volume usage" on page 253.

See "Splitting snapshots" on page 254.

See "Starting synchronization of snapshots" on page 255.

See "Enabling FastResync on volumes" on page 256.

See "Disabling FastResync on volumes" on page 257.

See "Permissions required for views and operations on CVM cluster objects" on page 625.

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

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 202.

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 160.

See "Filter Criteria panel options" on page 162.

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 262.

See "Advanced Options panel" on page 276.

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 160.

See "Filter Criteria panel options" on page 162.

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 207.

7 In the **Create Volume Summary** panel, verify the details of the new volume. Click **Finish**.

See "Creating file systems" on page 260.

See "Recovering volumes" on page 214.

See "Deleting volumes" on page 215.

See "About managing Storage Foundation volumes" on page 199.

Create Volume – Select Disk Group and Disk Selection method panel options

Use this wizard panel to choose the disk group for the volume.

Table 11-1 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.

Table 11-1 Create Volume – Select Disk Group and Disk Selection method panel options *(continued)*

Field	Description
Condition	Displays the condition of the disk group.
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 200.

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 11-2 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 11-2 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.
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.

Table 11-2 Create Volume - Volume attributes panel options *(continued)*

Field	Description
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: Once volume encryption is enabled, it cannot be disabled. The only workaround is to create a copy of the unencrypted data onto another volume with the same layout and mirror settings and then replace the encrypted volume For more information about encryption and KMS, refer to the platform-specific *Storage Foundation Administrator's Guide*.
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 200.

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 11-3 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.
Enable Logging	If Mirror Info is selected, you can select **Enable Logging** to enable logging.

See "Creating Storage Foundation volumes" on page 200.

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 11-4 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.
Enable file and folder compression	Select to enable compression of the file and folder.

See "Creating Storage Foundation volumes" on page 200.

Encrypting existing volumes

The online volume encryption feature lets you migrate unencrypted volumes at rest to encrypted ones. Using this feature a volume is migrated without application downtime, that is, while the file system is mounted and while the I/Os are running. The online migration operation involves mirroring the existing storage configured under a volume, which requires an equal amount of additional storage that gets used in the background.

This feature is available only with:

- InfoScale 9.0 and later on:
  - Linux managed hosts
  - Linux and Windows for Central Management Server (CMS)
- InfoScale Operations Manager Management Server 9.0 and later

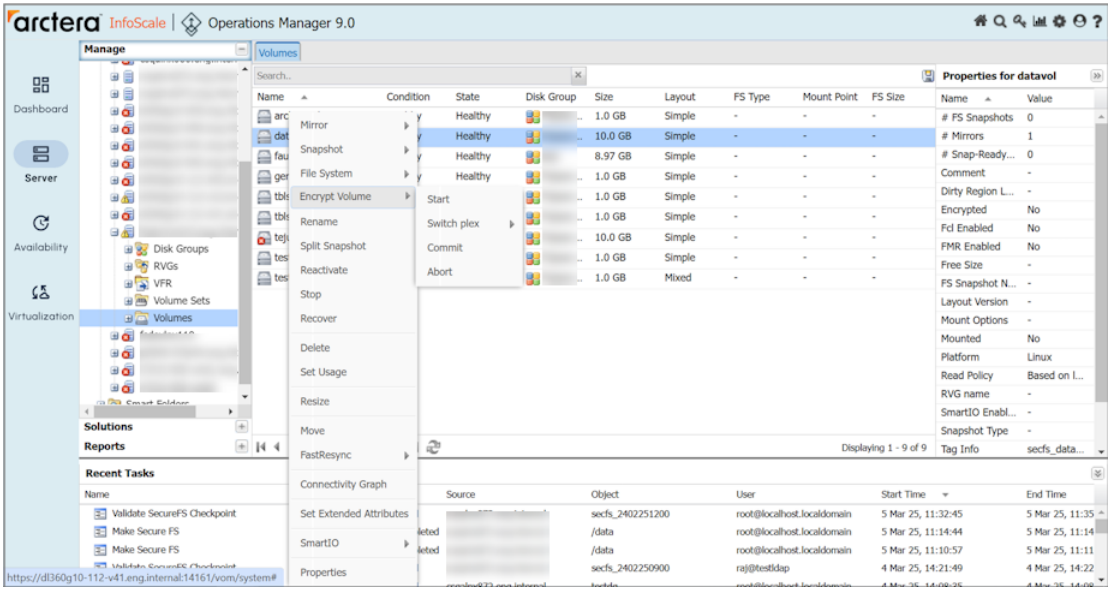
Online migration involves the Start phase, in which the process is initiated, and the Commit phase, in which the background changes made to the volume are finalized. The unencrypted volume is migrated to an encrypted one when both these phases are completed successfully. After the Start phase is complete and before you can initiate the Commit phase, you can abort the migration or switch plexes. The switching of reads between the source (unencrypted) plex and the target (unencrypted) plex helps verify the data copied during the Start phase. Meanwhile, the writes continue to happen on both the plexes.

Limitations:

- Encryption of RAID 5 and erasure coded (EC) volumes is not supported on Linux hosts.
- Only one online migration can be performed on an unencrypted volume at a time.
- Only one top-level mirror plex can be migrated at a time.

The volume encryption operations can be launched from the context of an existing unencrypted volume. To perform these operations, your user group must be assigned the admin role on the host or in the Server perspective. The permission on the host may be explicitly assigned or inherited from a parent Organization.

Figure 11-1 Menus for encrypting an existing unencrypted volume



Online migration involves the following operations:

- See "Starting online migration to encrypt a volume" on page 210.
- See "Switching plexes to verify the copied data" on page 211.
- See "Aborting online migration" on page 212.
- See "Committing the background changes to encrypt a volume" on page 213.

Each of these operations invokes an InfoScale command in the background, and you can view the status of each operation in the **Recent Tasks** pane.

For details on the volume encryption feature and the related commands, refer to the platform-specific *Storage Foundation Administrator's Guide*.

To troubleshoot any issues that you may encounter while performing these operations, you can review the information logged in the following InfoScale Operations Manager files:

- Management Server logs: `WebDebugLog.txt` and `tomcat.log`
- Agent log: `mhrun.log`

Starting online migration to encrypt a volume

The Server perspective of the InfoScale Operations Manager is used to initiate the online migration of an unencrypted volume to an encrypted one.

To start online migration

1. In the Management Server console, go to the **Server** perspective and expand **Manage** in the left pane.
2. Perform one of the following actions:
   - Expand **Uncategorized Hosts** or an Organization, expand the host, and expand **Volumes**. Right-click the unencrypted volume that you want to encrypt and select **Encrypt Volume** > **Start**.
   - Select the host and locate the unencrypted volume under it that you want to encrypt. Right-click the volume and select **Encrypt Volume** > **Start**.

   The **Start Encryption** wizard appears.
3. On the **Volume Attributes** panel, specify the following details:

   Target layout:
   - (Default) To use the same layout as the source volume, select the **Same as Source** option.
   - To use a different layout, select the required option and specify any related details like **Mirror Info** or **Stripe Info**, if applicable.

   Disk selection method:
   - **Automatic**: Indicates that all the disks across which the volume spans are to be encrypted. This option is selected by default.
   - **Manual**: Select this option if you want to select a subset of the disks to be encrypted.
4. Click **Next**.

5 On the **Start Encryption - Summary** panel, verify all the specifications that you have provided for the volume encryption. If you need to make any changes to the specifications on the previous panels, you can click **Previous**; otherwise click **Finish** to proceed.

The Start phase runs as a background process. In the volumes list, the **Encrypted** column displays the status as **In Progress**.

This operation invokes the `vxassist migrate` command with the `-o start` option and the relevant parameters based on your selections on the previous wizard panels.

6 On the **Start Encryption - Result** panel, verify the status of the operation and click **OK** to close the wizard.

When the process completes successfully, a new plex is added to the top-level volume. The newly added plex has a layered volume underneath, which in turn contains the same number of plexes as the top-level volume. In the volumes list, the **Encrypted** column displays the status as **Commit Pending**.

If the operation fails, an error message is displayed; accordingly take the appropriate action.

See "Encrypting existing volumes" on page 208.

Switching plexes to verify the copied data

When the encrypted (target) plexes are first created, the data reads continue to happen from the unencrypted (source) plexes. You can switch the reads to the encrypted plexes to verify that all the original data has been copied.

To read data from the target plex

1 Right-click the top-level volume for which you have started online migration and click **Encrypt Volume** > **Switch plex** > **Target plex**.

2 On the **Target plex** dialog box, verify the name of the volume and the target plex, and click **OK**.

This operation invokes the `vxassist migrate` command with the `-o switchtgt` option and the relevant parameters.

3 On the **Target plex - Result** panel, verify the status of the operation and click **OK** to close the dialog box.

The encrypted plexes are now set to read-write, and data is read from these plexes. All other plexes are set to write-only.

To read data from the source plex

1 Right-click the top-level volume for which you have started online migration and click **Encrypt Volume** > **Switch plex** > **Source plex**.

2 On the **Source plex** dialog box, verify the name of the volume and the source plex, and click **OK**.

This operation invokes the `vxassist migrate` command with the `-o switchsrc` option and the relevant parameters.

3 On the **Source plex - Result** panel, verify the status of the operation and click **OK** to close the dialog box.

The unencrypted plexes are now set to read-write, and data is read from these plexes. All other plexes are set to write-only.

Optionally, you can verify the switching of the data reads between source and the target plexes by using the `vxprint` command on the InfoScale hosts.

See "Encrypting existing volumes" on page 208.

Aborting online migration

If you find any inconsistencies in the copied data or for any other reason, you may choose to not enable the volume for encryption right away. In that case, you can abort the online migration to reinstate the volume to its original layout, but only if the Commit phase has not yet been initiated.

To abort online migration

1 Right-click the top-level volume for which you have started online migration and click **Encrypt Volume** > **Abort**.

2 On the **Abort Volume** dialog box, verify the name of the volume and click **OK**.

This operation invokes the `vxassist migrate` command with the `-o abort` option and the relevant parameters.

3 On the **Abort Volume - Result** panel, verify the status of the operation and click **OK** to close the dialog box.

The unencrypted volume layout in now restored to its original configuration.

In the volumes list, the **Encrypted** column displays the status as **No**.

See "Encrypting existing volumes" on page 208.

Committing the background changes to encrypt a volume

The Server perspective of the InfoScale Operations Manager is used to finalize the online migration of an unencrypted volume to an encrypted one.

To commit online migration

1. Right-click the top-level volume for which you have started online migration and click **Encrypt Volume** > **Commit**.
2. On the **Commit Volume** dialog box, verify the name of the volume and click **OK**.

   This operation invokes the `vxassist migrate` command with the `-o commit` option and the relevant parameters.
3. On the **Commit Volume - Result** panel, verify the status of the operation and click **OK** to close the dialog box.

   When the process completes successfully, the encrypted plex is attached to the original volume. A new backup volume is created on the additional storage with the same volume name and the `-old` suffix and its encryption state as OFF. The unencrypted (source) plex is attached to the backup volume.

The encryption state of the original volume is now set to ON. In the volumes list, the **Encrypted** column displays its status as **Yes**.

Note: After the online migration is committed successfully, the volume encryption cannot be disabled. The only workaround is to create a copy of the unencrypted data onto another volume with the same layout and mirror settings and then replace the encrypted volume.

See "Encrypting existing volumes" on page 208.

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 200.

See "About managing Storage Foundation volumes" on page 199.

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 200.

See "About managing Storage Foundation volumes" on page 199.

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 200.

See "About managing Storage Foundation volumes" on page 199.

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 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 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 216.
6. In the **Result** page, verify that the selected volumes have been deleted successfully.

See "Creating Storage Foundation volumes" on page 200.

See "About managing Storage Foundation volumes" on page 199.

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 215.

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 217.
6. In the **Result** page, verify that the selected volumes have been moved successfully.

See "Creating Storage Foundation volumes" on page 200.

See "About managing Storage Foundation volumes" on page 199.

Move Volume panel options

Use this wizard panel to move a volume from one disk to another disk.

Table 11-5 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 217.

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 218.
6. In the **Result** panel, verify that the volume is successfully renamed.

See "Creating Storage Foundation volumes" on page 200.

See "About managing Storage Foundation volumes" on page 199.

Rename Volume panel options

Use this wizard panel to rename a volume.

Table 11-6 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 218.

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 220.
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 160.
7. In the **Advanced options for Add mirror** panel, specify the required information. Click **Next**.

   See "Add mirror - Advanced options" on page 223.
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 220.

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 160.

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 226.

See "About managing Storage Foundation volumes" on page 199.

Add mirror - Options

Use this wizard panel to add mirrors to an existing volume.

Table 11-7 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.

Table 11-7 Add mirror - Options *(continued)*

Field	Description
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
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 222.

Table 11-7 Add mirror - Options *(continued)*

Field	Description
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 222.

See "Adding mirrors to volumes" on page 219.

Enable FastResync option panel options

Use this wizard panel to enable FastResync for a snapshot volume that you want to create.

Table 11-8 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.
DCO region size (KB)	Enter the size for the DCO mirror that you want to create. The default value is 64 kilobytes.

Table 11-8 Enable FastResync option panel options *(continued)*

Field	Description
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 220.

See "Adding mirrors to volumes" on page 219.

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 11-9 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 11-9 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.

Table 11-9 Add mirror - Advanced options *(continued)*

Field	Description
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 219.

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 227.
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 219.

See "About managing Storage Foundation volumes" on page 199.

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 11-10 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.

Table 11-10 Remove mirror panel options *(continued)*

Field	Description
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.
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.

Table 11-10 Remove mirror panel options *(continued)*

Field	Description
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 226.

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 294.

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 230.

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 231.

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 232.

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 233.

See "Creating mirror break-off snapshots for volumes" on page 235.

See "About managing Storage Foundation volumes" on page 199.

Create Volume Snapshot - disk selection page panel options

Table 11-11 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.

Table 11-11 Create snapshot - disk selection page panel options *(continued)*

Field	Description
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 162.
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 229.

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 11-12 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 229.

Create Volume Snapshot - Instant Snapshot panel options

Use this wizard panel to configure an instant snapshot.

This panel displays the following information:

Table 11-13 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 229.

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 294.
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 230.
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 231.

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 234.

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 199.

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 11-14 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.
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.

Table 11-14 Create snapshot - space optimized snapshot panel options *(continued)*

Field	Description
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 233.

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 294.

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 231.

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 236.

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 199.

Create Volume Snapshot - Mirror Break-off Snapshot panel options

Use this wizard panel to configure a snapshot for a volume.

Table 11-15 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
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.

Table 11-15 Create snapshot - break-off snapshot panel options *(continued)*

Field	Description
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 235.

Dissociating snapshots

The Management Server console lets you dissociate a snapshot and turn it into an independent volume.

In the 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 239.

See "About managing Storage Foundation volumes" on page 199.

Dissociate Snapshot panel options

Use this wizard panel to select the data volume and view its snapshot.

Table 11-16 Snapshot Table panel options

Field	Description
Name	Displays the name of the snapshot volume.
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 237.

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 239.

   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 237.

See "About managing Storage Foundation volumes" on page 199.

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 11-17 Reattach Snapshot panel options

Field	Description
Name	Name of parent snapshot/ data volume.
Snapshot Type	Type of snapshot

Table 11-17 Reattach Snapshot panel options *(continued)*

Field	Description
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 239.

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 241.

See "Select Disks panel options" on page 160.

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 200.

See "About managing Storage Foundation volumes" on page 199.

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 11-18 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 11-18 Resize volume panel options *(continued)*

Field	Description
Operation	

Table 11-18 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 160. ■ 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.

Table 11-18 Resize volume panel options *(continued)*

Field	Description
	■ To decrease the size of the volume to a specific percentage 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 240.

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 245.

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 229.

See "Creating mirror break-off snapshots for volumes" on page 235.

See "About managing Storage Foundation volumes" on page 199.

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 11-19 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 11-19 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 244.

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 247.

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 249.

See "Adding snapshot volumes to a refresh schedule" on page 251.

See "Removing the schedule for volume snapshot refresh" on page 252.

See "About managing Storage Foundation volumes" on page 199.

Refresh snapshot panel options

Use this wizard panel to specify the options for refreshing the snapshot of a volume.

Table 11-20 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 11-20 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 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 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 246.

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 250.
6. In the **Refresh snapshot** wizard panel, specify the required information. Click **Next**.

   See "Refresh snapshot panel options" on page 247.
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 246.

See "Adding snapshot volumes to a refresh schedule" on page 251.

See "Removing the schedule for volume snapshot refresh" on page 252.

See "About managing Storage Foundation volumes" on page 199.

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 11-21 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 249.

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 251.
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 246.

See "Configuring a schedule for volume snapshot refresh" on page 249.

See "Removing the schedule for volume snapshot refresh" on page 252.

See "About managing Storage Foundation volumes" on page 199.

Add to existing schedule panel options

Use this wizard panel to select an existing refresh schedule to add a volume snapshot to it.

Table 11-22 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 251.

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 253.

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 246.

See "Configuring a schedule for volume snapshot refresh" on page 249.

See "Adding snapshot volumes to a refresh schedule" on page 251.

See "About managing Storage Foundation volumes" on page 199.

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 252.

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 253.
6. In the **Result** panel, verify that the selected volume has been stopped successfully.

See "Creating Storage Foundation volumes" on page 200.

See "About managing Storage Foundation volumes" on page 199.

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 11-23 lists the volume usage options that you can set for UNIX or Linux hosts.

Table 11-24 lists the volume usage options that you can set for Windows hosts.

Table 11-23 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 11-24 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 253.

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 237.

See "About managing Storage Foundation volumes" on page 199.

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 237.

See "About managing Storage Foundation volumes" on page 199.

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 256.
6. In the **Enable FastResync - Result** panel, verify that FastResync is enabled on the volume successfully.

See "Disabling FastResync on volumes" on page 257.

See "About managing Storage Foundation volumes" on page 199.

Enable FastResync panel options

Use this wizard panel to enable FastResync on volumes.

Table 11-25 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 11-25 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 256.

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 256.

See "About managing Storage Foundation volumes" on page 199.

Chapter 12

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 260.

See "Enabling change logs" on page 278.

See "Disabling change logs" on page 279.

See "Synchronizing change logs" on page 280.

See "Removing change logs" on page 280.

See "Defragmenting file systems" on page 281.

See "Unmounting non clustered file systems from hosts" on page 282.

See "Mounting non clustered file systems on hosts" on page 283.

See "Unmounting clustered file systems" on page 287.

See "Mounting clustered file systems on hosts" on page 288.

See "Remounting file systems" on page 290.

See "Checking file systems" on page 292.

See "Creating file system snapshots" on page 293.

See "Remounting file system snapshot" on page 296.

See "Mounting file system snapshot" on page 298.

See "Unmounting file system snapshot" on page 299.

See "Removing file system snapshot" on page 300.

See "About performing Storage Foundation and replicator operations" on page 139.

See "Permissions required for views and operations on CVM cluster objects" on page 625.

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 261.
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 262.

   See "Advanced Options panel" on page 276.
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 290.

See "Checking file systems" on page 292.

See "Creating file system snapshots" on page 293.

See "Mounting file system snapshot" on page 298.

See "About managing file systems" on page 260.

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 12-1 Select volume panel options

Field	Description
Name	Displays the name of the volume.
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 260.

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 12-2 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.

Table 12-2 File System Options for vxfs, ext2, ext3, ext4 and ufs *(continued)*

Field	Description
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.
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 12-3 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.
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

Table 12-3 Mount Options *(continued)*

Field	Description
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 276.

Configure SecureFS immutable checkpoints

InfoScale offers the ability to configure secure file systems (SecureFS). SecureFS provides a recovery mechanism in case of a data corruption or deletion, and a ransomware event. The SecureFS feature helps protect your data through the use of snapshots and non-modifiable checkpoints.

The secure file system feature is enhanced to identify the underlying application (generic application, Oracle, or PostgreSQL) and then secure the file system accordingly. In earlier versions SecureFS feature was available only for general applications and for Oracle. With 9.0 release, SecureFS supports PostgreSQL databases. You can configure immutable checkpoint on a file system that is mounted on a InfoScale volume. Immutable checkpoints are supported only on VxFS. Hence all directories including the subdirectories should be VxFS mount points.

Prerequisites

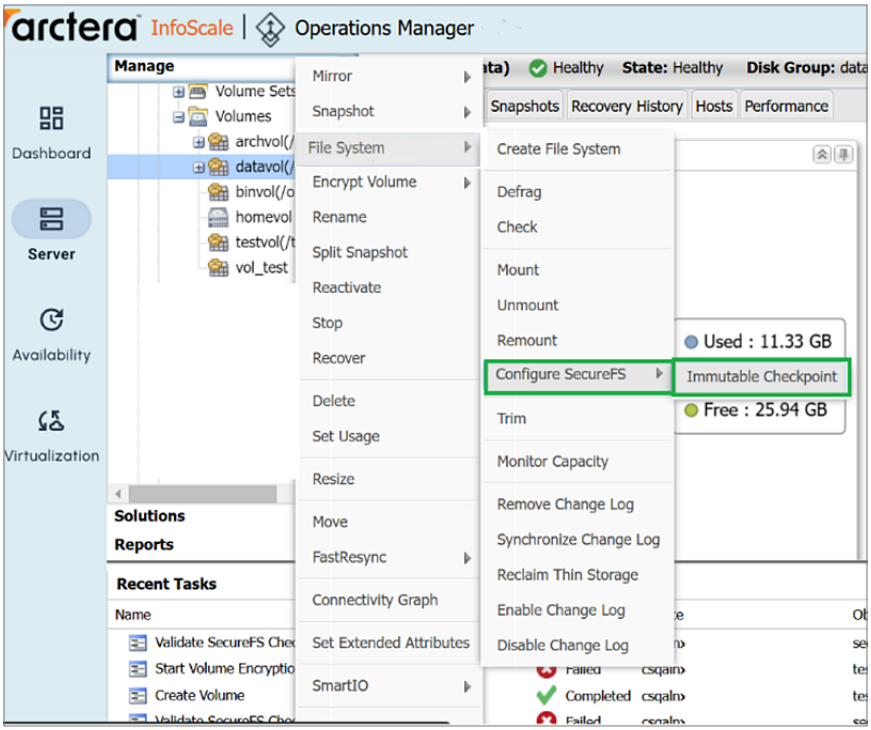
- Before you configure SecureFS immutable checkpoints, refer to the *InfoScale Storage Foundation Administrator's Guide - Linux* for more details about the SecureFS feature.
- To configure immutable checkpoint for an application, ensure that you select the host and the volume on which the application is installed or configured.
- 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 an immutable checkpoint

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 Expand the host and **Volumes** to locate the volume on which you want to configure immutable checkpoints.

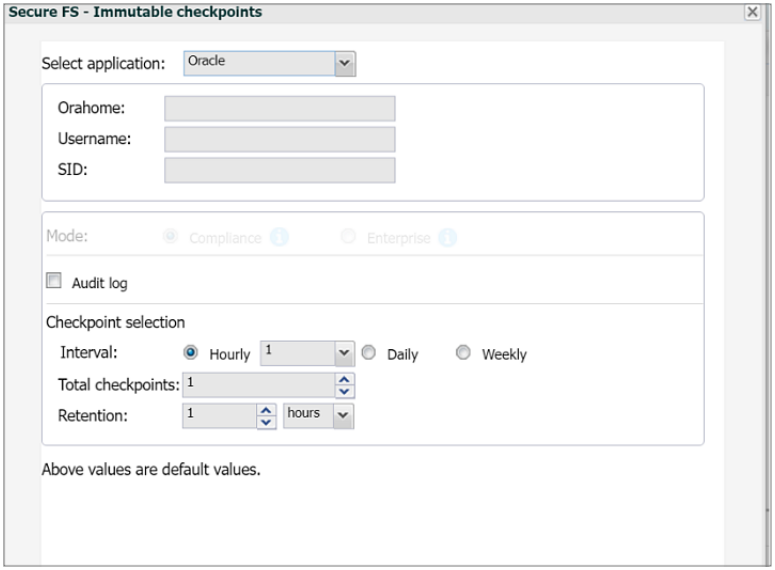
4 Right-click the volume and select **File System** > **Configure SecureFS** > **Immutable Checkpoint**.



5 In the **Secure FS - Immutable checkpoints** panel, from the **Select application** dropdown field, select the application, and then enter the required information.

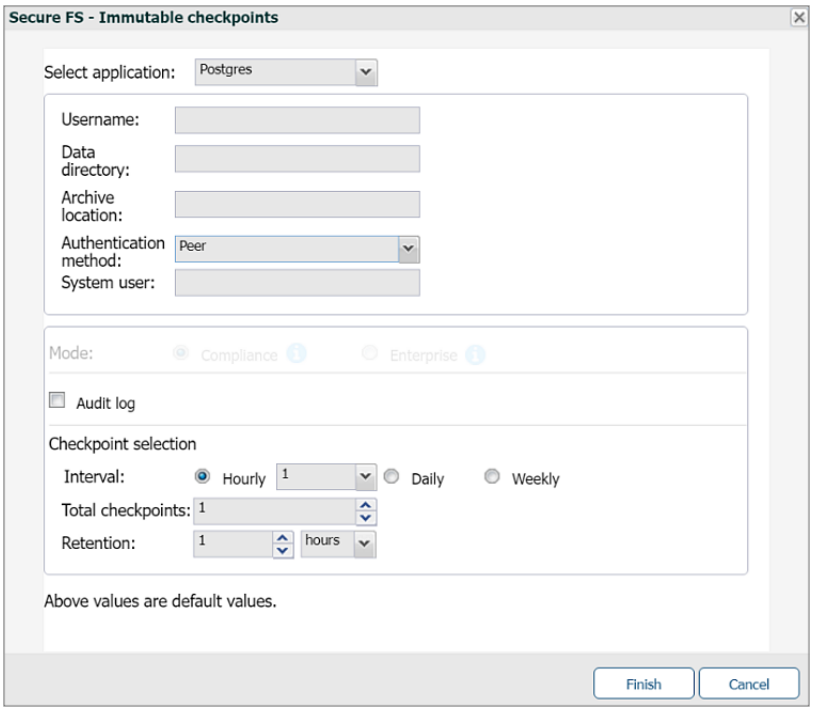
6 For a generic application, choose **General** and go to step 9.

7 For Oracle database, choose **Oracle** and then specify the following information and then go to step 9:



Orahome	Enter the Oracle home directory path that is set for the `ORACLE_HOME` environment variable. It is the path to the Oracle binaries and configuration files.
Username	Enter the Oracle user who has the privileges to start and stop the Oracle database instance.
SID	Enter the Oracle instance name that is specified for the $ORACLE_SID variable.

8 For PostgreSQL database, choose **Postgres** and then specify the following information:



Username	Enter the dedicated OS user name that is created when PostgreSQL is installed. This user has the privileges to start, stop, and monitor the PostgreSQL database server operations.
Data directory	Enter the absolute path of the directory that contains the database that the PostgreSQL instance manages.
Archive location	Enter the archive location that is configured to store the PostgreSQL Write-Ahead Log (WAL) files. To prevent SecureFS failures, this location must be on a different VxFS file system than where the PostgreSQL instance is deployed.

Authentication method	Choose the authentication method from the dropdown list. The following options are available: ■ Trust This method assumes that any user who can connect to the PostgreSQL server is authorized to access the database with whatever database user name they specify. ■ Peer This method relies on the operating system user to connect to the PostgreSQL database.
System user	Enter the operating system user that can connect to the PostgreSQL database. Specify this parameter only if the Authentication method is selected as Peer.

9 Specify the following details for the checkpoint creation:

Mode	Specifies the mode of the file system. While configuring SecureFS on an existing file system, the mode cannot be switched to compliance mode (Compliance mode refers to soft WORM) if the file system has WORM enabled. On the other hand, if WORM (WORM refers to Enterprise mode) is not enabled, the configuration process automatically converts the file system to Enterprise mode. Compliance mode is a default mode. You can change the mode later if required.
Audit log	Specifies the logging of certain events to the persistent storage. You can choose to enable or disable it. By default, it is enabled. Once enabled it cannot be disabled.
Interval	Specify the time interval frequency at which the SecureFS checkpoint creation must trigger. You can specify the time duration in Hours, Days, and Weeks. The minimum supported time duration is one hour, which means that the checkpoint creation gets triggered every hour.
Total checkpoints	Enter the maximum number of checkpoints to create and store.
Retention	Specify the time duration for which you want to retain the checkpoint. You can specify the retention time duration in hours, days, and weeks. The minimum recommended retention period is one hour.

10 Click **Finish** to trigger the SecureFS checkpoints creation process.

11 The **Secure FS - Result** pane displays the progress of the operation and the commands that InfoScale Operations Manager uses to configure the checkpoints. After the process completes successfully, click **OK**.

Depending on the checkpoint interval configuration, InfoScale Operations Manager runs the commands to create the SecureFS checkpoints. Once the retention period is reached, the oldest checkpoint is automatically deleted and a new checkpoint is created. The process repeats for the defined time interval.

Alerts for SecureFS checkpoints

The following alerts enhance monitoring and fault detection for the SecureFS checkpoint:

- The following alert is generated if the `vxfstaskd` daemon is not running on the host:
  The vxfstaskd process is not running on this host
  The fault is resolved once the `vxfstaskd` service is up and running.
- The following alert is generated if it fails to create a SecureFS checkpoint for a file system.
  Failed to create securefs checkpoint for filesystem
  The fault is cleared once the underlying issue with the specified file system is resolved.

You can view the faults in the **Faults and Risks** panel of the host in the **Server** perspective or in the **Server - Faults and Risks** panel in the **Overview** tab of the Dashboard by drilling down to the hosts on which the fault occurred.

Validate SecureFS checkpoints

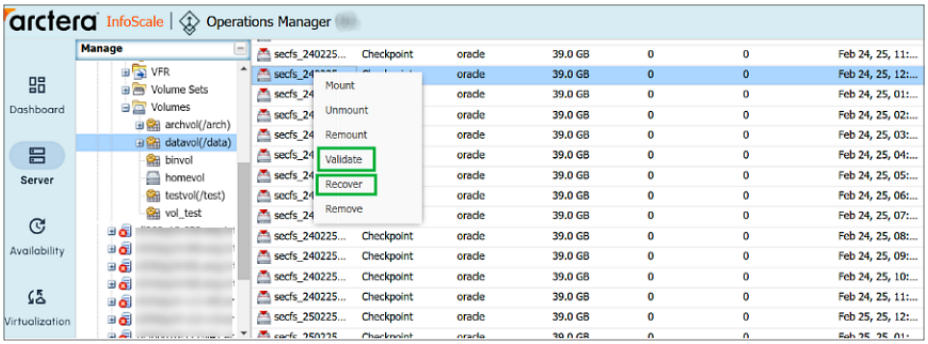
The secure file system (SecureFS) feature is enhanced to help perform several validations on the SecureFS checkpoints. These validations determine whether the selected checkpoint is viable and usable for a recovery.

To view and validate the immutable checkpoints for an application, ensure that you select the host and the volume on which the application is installed or configured. 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 validate a SecureFS checkpoint

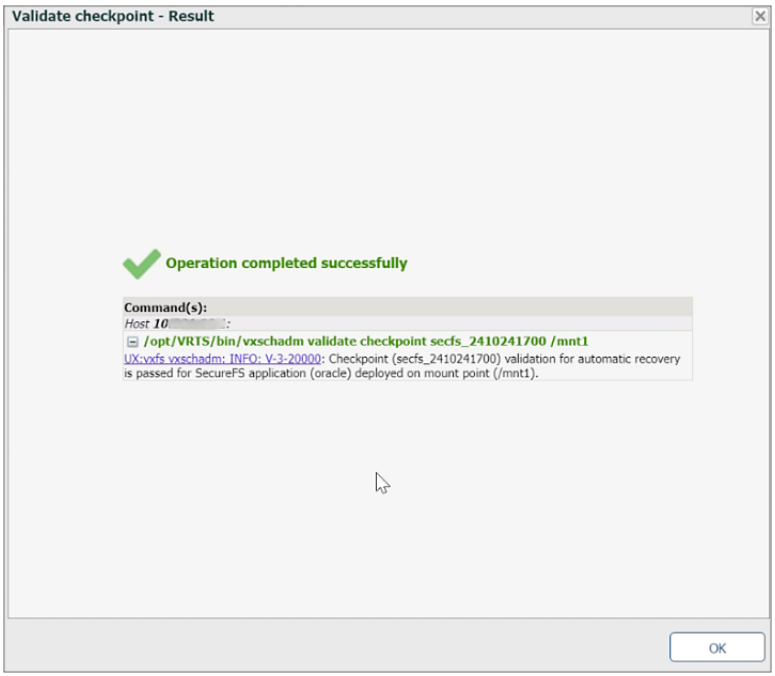
1. In the Management Server console, go to the **Server** perspective and expand **Manage** in the left pane.
2. Expand **Organization** or **Uncategorized Hosts** to locate and select the host.
3. Expand the host and **Volumes** to locate the volume for which you have configured immutable checkpoints.
4. Click the **Snapshots** tab in the right pane.

5 Right-click on a checkpoint and select **Validate**.



6 On the **Validate checkpoint** panel, verify that the desired checkpoint is selected and then click **OK** to begin the validation.

7 The **Validate checkpoint - Result** pane displays the checkpoint validation status. InfoScale Operations Manager runs the InfoScale commands to validate the checkpoint. You may see a message similar to the following:



After the validation completes successfully, click **OK**. The validated checkpoint is ready to use for a recovery.

Recover SecureFS checkpoints

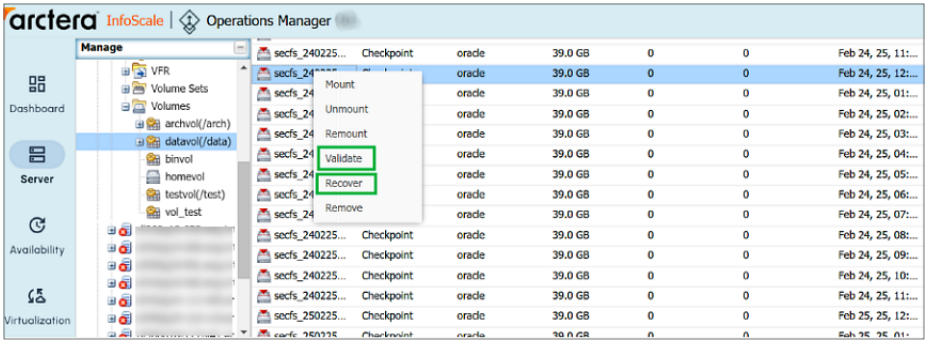
The secure file system (SecureFS) feature helps protect your data against ransomware attacks and accidental deletions through the use of snapshots and non-modifiable checkpoints. In case of a data loss event, you can use these storage checkpoints to perform data recovery using the SecureFS snapshots. In this release, the SecureFS snapshot recovery process is enhanced to be more efficient and user-friendly. Termed as "Single Click Recovery" the process focuses on ease of use and improving the recovery time objective (RTO) in SecureFS environments.

Prerequisites

- To view and validate the immutable checkpoints for an application, ensure that you select the host and the volume on which the application is installed or configured. 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.
- If the application is not configured with InfoScale Cluster Server (VCS), you must stop the application manually before you initiate the recovery process. If VCS manages the application, InfoScale Operations Manager freezes the VCS application service group and pauses the SecureFS configuration for the application, until the recovery is completed.
- Ensure that you have allocated sufficient free disk space on the main file system used by the application. For example, if the checkpoint size is 500 GB, you must have at least 500 GB or more free space available on the file system where the checkpoint is being restored.

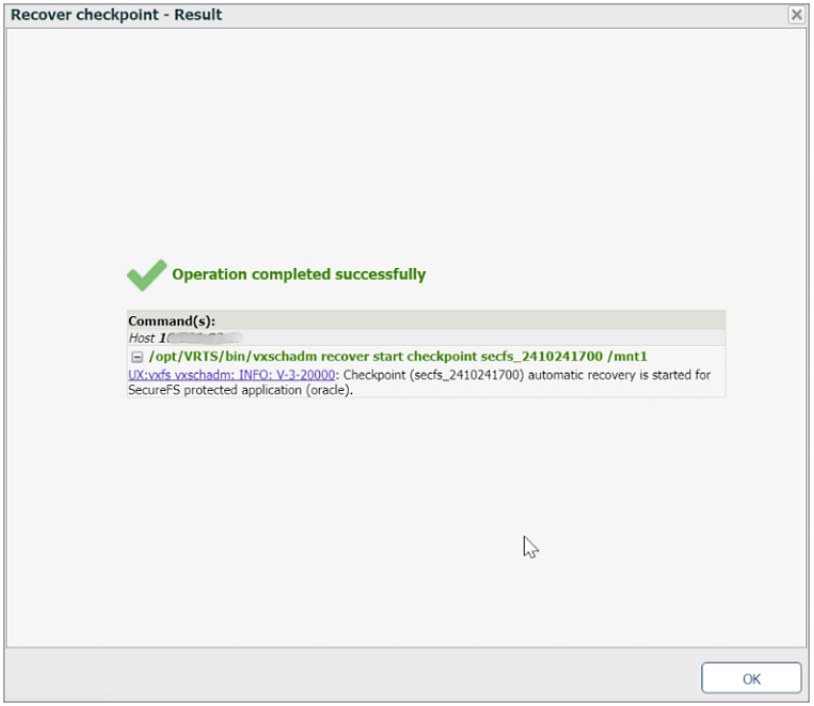
To recover SecureFS checkpoints

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. Expand the host and **Volumes** to locate the volume for which you have configured immutable checkpoints.
4. Click the **Snapshots** tab in the right pane.
5. Right-click on a checkpoint and select **Recover**.



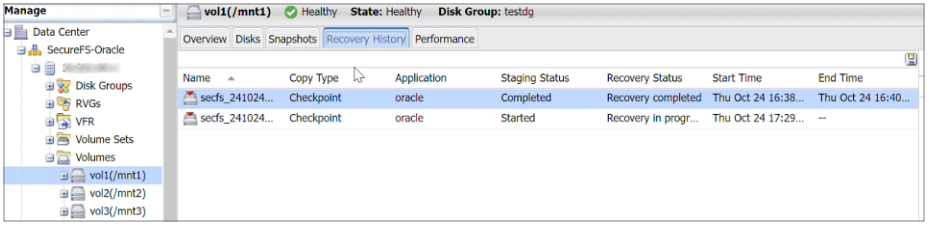
6. On the **Recover checkpoint** panel, verify that the desired checkpoint is selected and then click **Next** to begin the recovery.

7 The **Recover checkpoint - Result** pane displays the checkpoint recovery status. InfoScale Operations Manager runs the InfoScale commands to recover the checkpoint. You may see a message similar to the following:



After the recovery process is started successfully, click **OK**.

8 Click the **Recovery History** tab to view the progress of the checkpoint recovery. InfoScale Operations Manager displays the status in real time.



9 After the recovery is completed successfully, you can perform any additional steps that are required to start the application and make it functional again.

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 12-4 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.
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.

Table 12-4 Advanced options for vxfs, ufs, ext2, ext3, and ext4 *(continued)*

Field	Description
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 12-5 Advanced mount options for vxfs, ufs, ext2, ext3, and ext4

Field	Description
Advanced mount options for vxfs	
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 12-5 Advanced mount options for vxfs, ufs, ext2, ext3, and ext4 *(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.
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 262.

See "Creating file systems" on page 260.

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 279.

See "About managing file systems" on page 260.

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 278.

See "About managing file systems" on page 260.

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 278.

See "Removing change logs" on page 280.

See "About managing file systems" on page 260.

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 278.

See "Synchronizing change logs" on page 280.

See "About managing file systems" on page 260.

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 282.

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 260.

See "Creating file systems" on page 260.

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 281.

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
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 283.
6. In the **Result** page, verify that the selected non-clustered file system has been unmounted successfully.

See "Unmounting clustered file systems" on page 287.

See "Mounting non clustered file systems on hosts" on page 283.

See "About managing file systems" on page 260.

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 282.

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 284.
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 288.

See "Unmounting non clustered file systems from hosts" on page 282.

See "About managing file systems" on page 260.

Mount Options panel

Use this wizard panel to specify the options for mounting a non-clustered file system on a host.

Table 12-6 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.

Table 12-6 Mount Options *(continued)*

Field	Description
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.
Advanced	Select to specify the advanced options for mounting the file system. See "Advanced Options panel" on page 276.

See "Mounting non clustered file systems on hosts" on page 283.

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 12-7 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.
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 12-7 Advanced mount options *(continued)*

Field	Description
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 283.

See "Mounting clustered file systems on hosts" on page 288.

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 288.

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 282.

See "Mounting clustered file systems on hosts" on page 288.

See "About managing file systems" on page 260.

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 287.

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 289.

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 283.

See "Unmounting clustered file systems" on page 287.

See "About managing file systems" on page 260.

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 12-8 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.

Table 12-8 Mount File System - Clustered Mount Options *(continued)*

Field	Description
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 285.

See "Mounting clustered file systems on hosts" on page 288.

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 291.

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 283.

See "Mounting clustered file systems on hosts" on page 288.

See "About managing file systems" on page 260.

Remount Options panel

Use this wizard panel to remount an already mounted file system.

Table 12-9 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.
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

Table 12-9 Remount Options panel *(continued)*

Field	Description
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 290.

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 293.

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 260.

See "About managing file systems" on page 260.

Check File System panel options

Use this wizard panel to check file system.

Table 12-10 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 292.

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 294.

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 295.

5 In the **Result** panel, verify that the file system snapshot has been created successfully.

See "Remounting file system snapshot" on page 296.

See "Unmounting file system snapshot" on page 299.

See "About managing file systems" on page 260.

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 12-11 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 293.

Snapshot level selection panel options

Use this wizard panel to choose the method of creating volume snapshots or file system snapshots with checkpoints.

Table 12-12

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 229.

See "Creating space optimized snapshots for volumes" on page 233.

See "Creating mirror break-off snapshots for volumes" on page 235.

See "Creating file system snapshots" on page 293.

Create File System snapshot - Configure Options

Use this wizard panel to create a file system snapshot.

Table 12-13 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.

Table 12-13 Create 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 "Creating file system snapshots" on page 293.

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 297.
6. In the **Result** panel, verify that the checkpoint has been remounted successfully.

See "Creating file system snapshots" on page 293.

See "Mounting file system snapshot" on page 298.

See "Unmounting file system snapshot" on page 299.

See "About managing file systems" on page 260.

Remount File System Snapshot panel options

Use this wizard panel to remount a file system snapshot.

Table 12-14 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.
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 296.

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 298.
6. In the **Result** panel, verify that the checkpoint has been mounted successfully.

See "Creating file system snapshots" on page 293.

See "Remounting file system snapshot" on page 296.

See "Unmounting file system snapshot" on page 299.

See "About managing file systems" on page 260.

Mount File System Snapshot panel options

Use this wizard panel to mount a file system snapshot.

Table 12-15 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.

Table 12-15 Mount File System Snapshot panel options *(continued)*

Field	Description
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.
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 298.

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 300.

5 In the **Result** panel, verify that the checkpoint has been unmounted successfully.

See "Remounting file system snapshot" on page 296.

See "Creating file system snapshots" on page 293.

See "About managing file systems" on page 260.

Unmount File System Snapshot panel options

Use this wizard panel to unmount a file system snapshot.

Table 12-16 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 299.

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 301.

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 293.

See "Remounting file system snapshot" on page 296.

See "Unmounting file system snapshot" on page 299.

See "About managing file systems" on page 260.

Remove File System Snapshot panel options

Use this wizard panel to remove a file system snapshot.

Table 12-17 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.

Table 12-17 Remove File System Snapshot panel options *(continued)*

Field	Description
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 300.

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 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 303.

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.

Note: On Windows, InfoScale Operations Manager supports FS monitoring capacity only when the FS is under VxVM.

See "About managing file systems" on page 260.

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 12-18 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.

Table 12-18 Monitor Capacity panel options *(continued)*

Field	Description
% 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 302.

Chapter 13

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 7.4.2 or later for Windows or Linux.

SmartIO requires a managed host with Storage Foundation 7.4.2 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 *Arctera InfoScale SmartIO for Solid State Drives Solutions Guide* for Linux.

See "Creating a cache" on page 308.

See "Deleting a cache" on page 316.

See "Enabling or disabling SmartIO caching" on page 307.

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 283.

See "Mounting clustered file systems on hosts" on page 288.

See "Remounting file systems" on page 290.

See "About managing SmartIO" on page 306.

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 308.

See "Modifying a cache" on page 313.

See "Deleting a cache" on page 316.

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 309.

See "Modifying a cache" on page 313.

See "Deleting a cache" on page 316.

Create Cache panel options

Use this wizard panel to specify the attributes for the cache area.

Table 13-1 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 13-1 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 308.

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 308.

See "Modifying a cache" on page 313.

See "Deleting a cache" on page 316.

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 307.

See "Creating a cache" on page 308.

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 308.

See "Modifying a cache" on page 313.

See "Deleting a cache" on page 316.

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 313.

See "Creating a cache" on page 308.

See "Deleting a cache" on page 316.

See "Enabling or disabling SmartIO caching" on page 307.

Modify Cache panel options

Use this wizard panel to modify the attributes for a cache area.

Table 13-2 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 13-2 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 313.

See "Enabling or disabling SmartIO caching" on page 307.

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 308.

See "Modifying a cache" on page 313.

See "Modify Cache panel options" on page 313.

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 317.

See "Modifying a cache" on page 313.

See "Deleting a cache" on page 316.

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 317.

   See "Loading files to a cache" on page 316.

See "Modifying a cache" on page 313.

See "Deleting a cache" on page 316.

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 317.

See "Changing SmartIO mode" on page 312.

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 320.
- See "Creating an I/O trace log " on page 318.
- See "Analyzing an I/O trace log" on page 320.
- See "Removing an I/O trace log" on page 321.

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 319.

See "Viewing the SmartIO Impact analysis chart" on page 311.

See "Analyzing an I/O trace log" on page 320.

New I/O Trace Log panel options

Use this wizard panel, specify the details for the new I/O trace log.

Table 13-3 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 318.

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 318.

See "Removing an I/O trace log" on page 321.

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 321.

SmartAssist Analysis Options pane

To view SmartAssist analysis for custom cache size, specify the following details

Table 13-4

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 14

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 324.

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 325.

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 326.

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 323.

See "Setting IO Thresholds" on page 325.

See "Viewing live IOPS charts" on page 327.

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 323.

See "Modifying an App VG" on page 326.

See "Viewing live IOPS charts" on page 327.

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 323.

See "Managing IO Thresholds" on page 324.

See "Setting IO Thresholds" on page 325.

See "Viewing live IOPS charts" on page 327.

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 323.

Chapter 15

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
- Implementation of new faults
- Support for global VVR replication monitoring thresholds

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 329.

Adding a secondary

See "Adding a Secondary" on page 340.

Pausing replication to a Secondary

See "Pausing the replication to a Secondary" on page 341.

Resuming replication of a Secondary

See "Resuming the replication of a Secondary" on page 342.

Starting replication to a Secondary

See "Starting replication to a Secondary" on page 343.

Stopping replication to a Secondary

See "Stopping the replication to a Secondary" on page 344.

Switching a Primary

See "Switching a Primary" on page 345.

Taking over from an original Primary

See "Taking over from an original Primary" on page 346.

Removing a secondary

See "Removing a Secondary" on page 350.

See "About performing Storage Foundation and replicator operations" on page 139.

Configuring Storage Foundation replications

In Management Server console, you can configure and set up replication by using Configure Replication Operation.

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 VVR replication from the Management Server console

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** wizard panel, enter the options to configure the replication and click **Start replication** to start the VVR replication.

See "Configure replication panel options" on page 330.

Note: If any operation fails after you have created any object, you can rollback the action of creating the object.

Configure replication panel options

Use this wizard panel to configure replication.

Table 15-1 Configure replication 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.

Table 15-1 Configure replication panel options *(continued)*

Name	Description
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.
Secondary host	Click **Select host**, to select a host from the existing list.
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. Disk group with the same name is available.
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 matching volume on Disk group	Select this option to automatically create the volumes on secondary disk group.
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.

Table 15-1 Configure replication panel options *(continued)*

Name	Description
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.
Size	If you are adding a new volume as the SRL, specify the SRL size. The SRL size must be at least 110MB.
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.
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.
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

Table 15-1 Configure replication panel options *(continued)*

Name	Description
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.
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.

Table 15-1 Configure replication panel options *(continued)*

Name	Description
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.
Start replication	Starts the replication process to configure replication between primary and secondary. See "Configure replication - Create resources" on page 334.

See "Configuring Storage Foundation replications" on page 329.

Configure replication - Create resources

Use this wizard panel to configure replication and create resources for VCS configurations.

Table 15-2 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.
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 329.

Configure replication - Create primary result panel options

Use this wizard panel to view the result of the new primary and secondary created and you can view the status of the replication started.

See "Configuring Storage Foundation replications" on page 329.

Configure replication - Consistency check on secondary host panel options

Use this wizard panel to check for consistency before adding a secondary.

Table 15-3 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 InfoScale Operations Manager, consistency check can be run. If it is not added to 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.
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 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.

Table 15-3 Configure replication - Consistency check on secondary host panel options *(continued)*

Name	Description
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.
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 340.

See "Configuring Storage Foundation replications" on page 329.

Configure replication - Replication settings for secondary host panel options

Use this wizard panel to select the replication settings for the secondary host.

Table 15-4 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.
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.

Table 15-4 Configure replication - Replication settings for secondary host panel options *(continued)*

Name	Description
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	
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

Table 15-4 Configure replication - Replication settings for secondary host panel options *(continued)*

Name	Description
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.
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 334.

See "Adding a Secondary" on page 340.

See "Configuring Storage Foundation replications" on page 329.

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 340.

See "Configuring Storage Foundation replications" on page 329.

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 335.
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 336.
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 339.

See "Configuring Storage Foundation replications" on page 329.

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 342.

Pause replication panel options

Use this wizard panel to pause the replication from a primary to a secondary.

Table 15-5 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 341.

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 341.

Resume replication panel options

Use this wizard panel to resume the paused replication between the primary and a secondary.

Table 15-6 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 342.

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 343.

4 In the **Summary** wizard panel, verify that the selected replication has been started.

See "Stop replication panel options " on page 344.

Start replication panel options

Use this wizard panel to start the replication to a secondary replicated data set (RDS).

Table 15-7 Start replication panel options

Name	Description
Using	The replication can be started using one of the following: ■ Automatic synchronization ■ Synchronization not needed

Table 15-7 Start replication panel options *(continued)*

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 "Starting replication to a Secondary" on page 343.

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 344.
4. In the **Summary** wizard panel, verify that the selected replication has been stopped.

See "Starting replication to a Secondary" on page 343.

Stop replication panel options

Use this wizard panel to stop the replication from a primary to a secondary.

Table 15-8 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 344.

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 343.

Switch replication panel options

Use this wizard panel to switch the replication from a Primary to a Secondary.

Table 15-9 lists the attributes to switch the replication from a Primary to a Secondary when the RVG is not under VCS control.

Table 15-10 lists the attributes to switch the replication from a Primary to a Secondary when the RVG is under VCS control.

Table 15-9 Switch 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.

Table 15-10 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 345.

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 347.

4 In the **Summary** wizard panel, verify that the selected replication has been taken over.

See "Starting replication to a Secondary" on page 343.

Takeover Primary panel options

Table 15-11 lists the attributes to takeover the replication from a Primary to a Secondary when the RVG is not under VCS control.

Table 15-12 lists the attributes to takeover the replication from a Primary to a Secondary when the RVG is under VCS control.

Table 15-11 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.

Table 15-11 Takeover replication if RVG is not under VCS control panel options *(continued)*

Name	Description
Failback options	
Fast failback	Select this option to incrementally synchronize the original Primary with the new Primary.
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 15-12 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 346.

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 349.

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 349.

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 348.

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 15-13 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.

Table 15-13 Associate Volume - Advanced Options panel options *(continued)*

Name	Description
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 348.

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 351.

4 In the **Summary** wizard panel, verify that the selected replication has been taken over.

See "Starting replication to a Secondary" on page 343.

Remove secondary panel options

Use this wizard panel to remove the Secondary from the RVG.

Table 15-14 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 350.

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 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 352.

Monitoring replications

In 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 353.

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 15-15 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 15-15 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 352.

Implementation of new faults

The following new faults were introduced:

Volume Replication

- **SRL Disabled**
  If RVG volume is SRL and state is 'DISABLED' then **SRL for RVG %obj_name% is disabled** fault gets raised. This is a part of VVR discovery.
- **RLINK detached**
  If the Rlink state for RVG is 'DETACHED' then **Rlink for RVG %obj_name% is detached** fault gets raised. This is a part of VVR discovery.
- A new fault of type warning is implemented for Volume Replicator to identify when SRL and data volumes are on the same disk. The fault is raised on the RVG consisting of such SRL and data volume, and is then propagated to the respective host.
  Fault message is 'SRL and data volumes are on the same disk' and topic ID is 'event.alert.vom.vvr.datavol.srlvol.same.disk' . The topic ID can be used to create rule in the Server perspective.

Cluster Server

A new fault of type error is implemented for VCS. The fault is raised on the remote cluster which shut down abruptly and is not able to update its resources or service groups status to the reporting Operations Manager Management Server.

The Management Server reports the status of the remote cluster as faulted. Fault message is 'Cluster is down. Administrative action is required.' and topic ID is 'event.alert.vom.vcs.cluster.down'. The topic ID can be used to create rule in the Availability perspective.

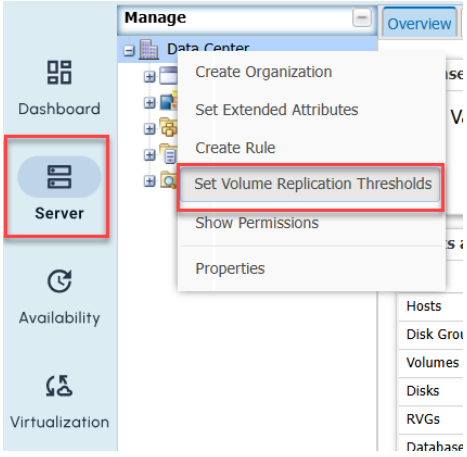
Support for global VVR replication monitoring thresholds

Using **Set Volume Replication Thresholds** option you can specify a common low or high replication threshold values to raise an alert for an Organization.

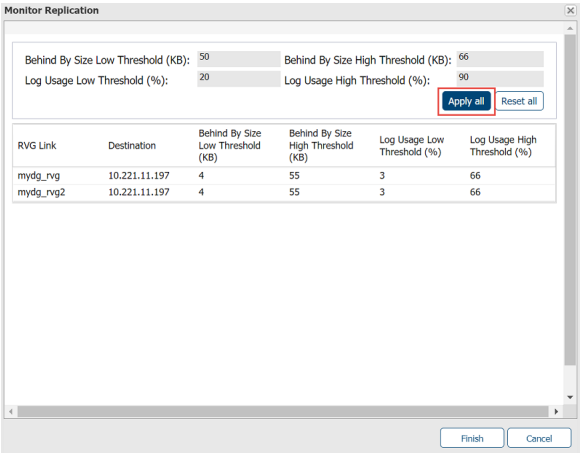
To perform this task, your user group must be assigned the Admin role on the host or the Server perspective. The permission on the host may be explicitly assigned or inherited from a parent Organization.

Perform the following procedure to monitor VVR replication threshold globally

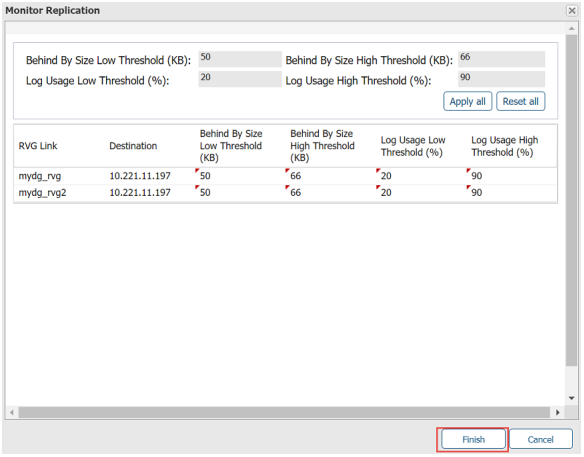
1. Right – click the **Data Center** menu to use **Set Volume Replication Threshold** option. The Monitor Replication window appears.

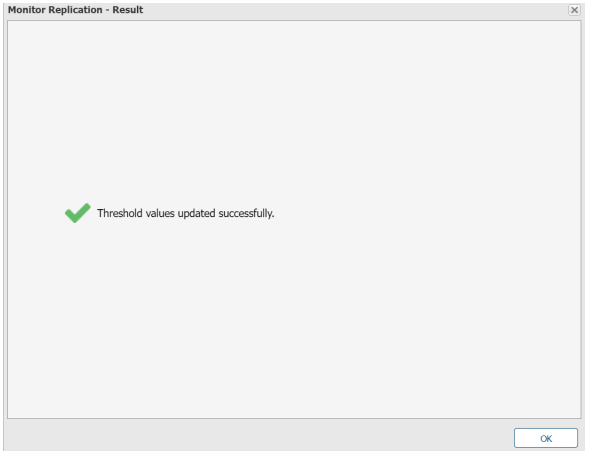


2. Set the VVR replication thresold values for all the RVG links for an organization and click the **Apply all** button.



3. Click the **Finish** button, to completed the process. A success message appears.




Monitor Replication - Result
Threshold values updated successfully.
OK

Chapter 16

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 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: 7.4.2 or later
- Windows: 7.4.2 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 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.

See "Performing thin reclamation on file systems or disks" on page 359.

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 358.

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**.

Compressing files

InfoScale Operations Manager lets you compress files by directory on a managed host.

File systems for compression must meet the following requirements:

- Storage Foundation 7.4.2 or later
- File System (VxFS) disk layout version 9 or later
- InfoScale Operations Manager managed host (VRTSsfmh) version 7.4.2 or later

InfoScale Operations Manager does not support file compression on Windows file systems.

See the following topics:

See "About file compression in InfoScale Operations Manager" on page 360.

See "Setting up compression schedules" on page 362.

See "Selecting directories for compression " on page 362.

See "Starting compression on demand" on page 364.

About file compression in 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 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 360.

Selecting directories for compression

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 362.

See "Starting compression on demand" on page 364.

Setting up compression schedules

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 360.

Starting compression on demand

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 360.

Deduplicating file systems

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 7.4.2 or later
- File System (VxFS) disk layout version 9 or later
- InfoScale Operations Manager managed host (VRTSsfmh) version 7.4.2 or later

InfoScale Operations Manager does not support deduplication on Windows file systems.

See the following topics:

See "About file system deduplication" on page 366.

See "Implementing deduplication for a file system" on page 368.

See "Starting deduplication for a file system" on page 371.

See "Disabling or removing deduplication for a file system" on page 371.

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, Arctera recommends a chunk size of 4k for SFCFSHA, where multiple copies of virtual machine images are accessed over NFS. For all other datasets, Arctera 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 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, Arctera 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 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. Arctera recommends that you have approximately 20% free space for 4k chunks.

Implementing deduplication for a file system

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 371.
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, Arctera 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. Arctera 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 365.

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 365.

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 365.

Section 4

Managing high availability and disaster recovery configurations

- Chapter 17. Overview
- Chapter 18. Managing clusters
- Chapter 19. Managing service groups
- Chapter 20. Managing systems
- Chapter 21. Managing resources
- Chapter 22. Managing global cluster configurations
- Chapter 23. Running fire drills
- Chapter 24. Using recovery plans
- Chapter 25. Managing application configuration
- Chapter 26. Multi Site Management
- Appendix A. List of high availability operations

Chapter 17

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 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 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 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, InfoScale Operations Manager provides an option to close the configuration after the operation is completed.

- When an operation that involves configuration changes is performed, 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 377.

See "About managing resources" on page 425.

See "About high availability and disaster recovery readiness" on page 449.

Pre-requisites for performing high availability and disaster recovery operations

You must ensure the following before performing high availability and disaster recovery operations :

- The InfoScale Operations Manager 3.0, or later managed host package must be installed. It is recommended to install InfoScale Operations Manager 9.1 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 374.

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 17-1 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 379.

See "Editing attributes of service groups" on page 412.

See "Editing attributes of a system" on page 421.

See "Editing attributes of a resource" on page 433.

See "Editing attributes of a resource type" on page 434.

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 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 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 18

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 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 378.

See "Opening a cluster configuration" on page 378.

See "Saving a cluster configuration" on page 378.

See "Saving a cluster configuration" on page 378.

See "Closing a cluster configuration" on page 379.

See "Closing a cluster configuration" on page 379.

See "Starting the Cluster Server high availability daemon on the hosts in a cluster" on page 421.

See "Stopping the Cluster Server high availability daemon on the systems in a cluster" on page 422.

See "Importing a type definition" on page 381.

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 377.

See "Saving a cluster configuration" on page 378.

See "Closing a cluster configuration" on page 379.

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 377.

See "Opening a cluster configuration" on page 378.

See "Closing a cluster configuration" on page 379.

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 377.

See "Opening a cluster configuration" on page 378.

See "Saving a cluster configuration" on page 378.

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 380.
7. In the **Result** page, click **OK**.

See "About managing clusters" on page 377.

See "Saving a cluster configuration" on page 378.

Edit attribute options

Use this wizard panel to edit attributes of a cluster, host, service group, resources, or resource types.

Table 18-1 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 18-1 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 379.

See "Editing attributes of a system" on page 421.

See "Editing attributes of service groups" on page 412.

See "Editing attributes of a resource" on page 433.

See "Editing attributes of a resource type" on page 434.

See "Configure Resources options" on page 388.

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 377.

Chapter 19

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 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 390.

See "Freezing service groups" on page 394.

See "Enabling all resources of service groups" on page 397.

See "Flushing service groups" on page 396.

See "Clearing faults on service group" on page 410.

See "Linking service groups in a cluster" on page 399.

See "Autoenabling service groups" on page 394.

See "Editing attributes of service groups" on page 412.

See "Creating service groups" on page 385.

See "Converting local service groups to global service groups" on page 443.

See "Converting global service groups to local service groups" on page 446.

See "Modifying the system list for a service group" on page 413.

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 386.
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 387.
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 388.

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 384.

See "Deleting service groups" on page 398.

See "About modifying a service group" on page 416.

See "Adding or modifying resources" on page 436.

Create Service Group options

Use this wizard panel to create a service group in a cluster.

Table 19-1 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 19-1 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 385.

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 19-2 Configure System List wizard panel options

Column	Description
Available Systems	Lists the available systems that you can use for running the service group.

Table 19-2 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 385.

See "Modifying the system list for a service group" on page 413.

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 19-3 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 19-3 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 380.
Delete resource	Click this option to delete the resource.

See "Creating service groups" on page 385.

See "Adding or modifying resources" on page 436.

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 19-4 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 19-4 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 385.

See "Linking resources in a service group" on page 435.

See "Unlinking resources in a service group" on page 435.

See "Adding or modifying resources" on page 436.

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 384.

See "Autoenabling service groups" on page 394.

See "Disabling service groups" on page 391.

See "Enabling resources" on page 425.

See "Enabling all resources of service groups" on page 397.

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 384.

See "Enabling service groups" on page 390.

See "Autoenabling service groups" on page 394.

See "Disabling resources" on page 426.

See "Disabling all resources of service groups" on page 398.

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 386.

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 387.

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 388.

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 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 384.

See "Enabling service groups" on page 390.

See "Disabling service groups" on page 391.

See "Enabling resources" on page 425.

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 384.

See "Unfreezing service groups" on page 395.

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 384.

See "Freezing service groups" on page 394.

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 384.

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 384.

See "Disabling all resources of service groups" on page 398.

See "Enabling service groups" on page 390.

See "Disabling service groups" on page 391.

See "Enabling resources" on page 425.

See "Disabling resources" on page 426.

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 384.

See "Enabling all resources of service groups" on page 397.

See "Enabling service groups" on page 390.

See "Disabling service groups" on page 391.

See "Enabling resources" on page 425.

See "Disabling resources" on page 426.

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 384.

See "Creating service groups" on page 385.

See "Linking service groups in a cluster" on page 399.

See "Modifying the system list for a service group" on page 413.

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 399.

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 400.

6 In the **Result** panel, click **OK**.

See "About managing service groups" on page 384.

See "About linking service groups in a cluster" on page 399.

See "Unlinking service groups" on page 402.

See "Deleting service groups" on page 398.

Link Service Group options

Use this wizard panel to add the selected service group as a child group to a parent service group.

Table 19-5 Link Service Group panel options

Field	Description
Child Group	Select the service group that you intend to define as the child service group.

Table 19-5 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 19-5 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 399.

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 384.

See "About linking service groups in a cluster" on page 399.

See "Linking service groups in a cluster" on page 399.

See "Deleting service groups" on page 398.

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 404.

See "Limitations of site-related service group operations" on page 404.

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 484.
- You use the 7.4.2 or later version of Cluster Server, InfoScale Operations Manager Management Server and the managed hosts.

See "About site aware service group operations" on page 403.

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 403.

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 405.

6 In the **Result** panel, click **OK**.

See "About site aware service group operations" on page 403.

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 19-6

For multiple service groups:Table 19-7

Table 19-6 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 19-6 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 19-7 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 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 404.

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 407.

6 In the **Result** panel, click **OK**.

See "About site aware service group operations" on page 403.

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 19-8

For multiple service groups:Table 19-9

Table 19-8 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 19-8 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 19-9 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 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 406.

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 409.
6. In the **Result** panel, click **OK**.

See "About site aware service group operations" on page 403.

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 19-10

For multiple service groups: Table 19-11

Table 19-10 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 19-11 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 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 408.

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 384.

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 430.

See "Taking resources offline" on page 430.

See "Taking a resource offline and propagating the state" on page 429.

See "Probing resources" on page 428.

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 380.

8 In the **Result** page, click **OK**.

See "About managing service groups" on page 384.

See "About modifying a service group" on page 416.

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 387.

6 In the **Result** panel, click **OK**.

See "About managing service groups" on page 384.

See "About modifying a service group" on page 416.

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%.

Component	Description
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 415.

See "Viewing the resource dependency view" on page 438.

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 414.

See "Viewing the resource dependency view" on page 438.

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 436.
- Modifying the System List for the service group
  See "Modifying the system list for a service group" on page 413.
- 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. 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, InfoScale Operations Manager Management Server and the managed hosts.

See "Viewing VCS Failover Duration report" on page 417.

See "About threshold settings" on page 552.

See "Adding threshold settings for an object" on page 553.

Viewing VCS Failover Duration report

In 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 Arctera Operations Manger 7.4.2 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 123.

Chapter 20

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 419.

See "Unfreezing a system" on page 420.

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 420.

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 419.

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 380.

8 In the **Result** page, click **OK**.

See "About managing systems" on page 419.

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 422.

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 421.

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 425.

See "Disabling resources" on page 426.

See "Probing resources" on page 428.

See "Editing attributes of a resource" on page 433.

See "Editing attributes of a resource type" on page 434.

See "Clearing faults on resources" on page 427.

See "Clearing the resources in a service group from the Admin Wait state" on page 411.

See "Deleting resources" on page 426.

See "Linking resources in a service group" on page 435.

See "Bringing resources online" on page 430.

See "Taking resources offline" on page 430.

See "Invoking a resource action" on page 431.

See "Taking a resource offline and propagating the state" on page 429.

See "Adding or modifying resources" on page 436.

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 425.

See "Disabling resources" on page 426.

See "Enabling service groups" on page 390.

See "Autoenabling service groups" on page 394.

See "Enabling all resources of service groups" on page 397.

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 425.

See "Enabling resources" on page 425.

See "Disabling service groups" on page 391.

See "Disabling all resources of service groups" on page 398.

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 425.

See "Disabling all resources of service groups" on page 398.

See "Enabling resources" on page 425.

See "Enabling all resources of service groups" on page 397.

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 425.

See "Clearing faults on service group" on page 410.

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 425.

See "Bringing resources online" on page 430.

See "Taking resources offline" on page 430.

See "Taking a resource offline and propagating the state" on page 429.

See "Clearing the resources in a service group from the Admin Wait state" on page 411.

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 425.

See "Bringing resources online" on page 430.

See "Taking resources offline" on page 430.

See "Clearing the resources in a service group from the Admin Wait state" on page 411.

See "Probing resources" on page 428.

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 425.

See "Taking resources offline" on page 430.

See "Clearing the resources in a service group from the Admin Wait state" on page 411.

See "Probing resources" on page 428.

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 425.

See "Bringing resources online" on page 430.

See "Taking a resource offline and propagating the state" on page 429.

See "Probing resources" on page 428.

See "Clearing the resources in a service group from the Admin Wait state" on page 411.

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 432.

6 In the **Result** panel, click **OK**.

See "About managing resources" on page 425.

See "Adding or modifying resources" on page 436.

See "Linking resources in a service group" on page 435.

See "Unlinking resources in a service group" on page 435.

See "Deleting resources" on page 426.

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 431.

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 380.

8 In the **Result** page, click **OK**.

See "About managing resources" on page 425.

See "Adding or modifying resources" on page 436.

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 380.

8 In the **Result** page, click **OK**.

See "About managing clusters" on page 377.

See "Saving a cluster configuration" on page 378.

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 425.

See "Adding or modifying resources" on page 436.

See "Invoking a resource action" on page 431.

See "Unlinking resources in a service group" on page 435.

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 425.

See "Linking resources in a service group" on page 435.

See "Adding or modifying resources" on page 436.

See "Invoking a resource action" on page 431.

See "Deleting resources" on page 426.

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 388.
- 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 425.

See "Linking resources in a service group" on page 435.

See "Unlinking resources in a service group" on page 435.

See "Invoking a resource action" on page 431.

See "Deleting resources" on page 426.

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 438.

See "About managing resources" on page 425.

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 437.

See "About managing resources" on page 425.

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 414.

See "Viewing the service group dependency view" on page 415.

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 441.

See "About creating global clusters" on page 442.

See "About removing a remote cluster from a global cluster setup" on page 447.

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 440.

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 440.

See "Prerequisites for creating global clusters" on page 442.

See "Adding a remote cluster to a local cluster" on page 443.

See "Adding a remote cluster to a local cluster" on page 443.

See "Converting local service groups to global service groups" on page 443.

See "Converting local service groups to global service groups" on page 443.

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 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 442.

See "Adding a remote cluster to a local cluster" on page 443.

See "Converting global service groups to local service groups" on page 446.

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 442.

See "Prerequisites for creating global clusters" on page 442.

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 444.

6 To make the service group global, click **OK**.

7 In the **Result** panel, click **OK**.

See "About creating global clusters" on page 442.

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 .
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 443.

See "Converting global service groups to local service groups" on page 446.

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 444.
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 447.

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 447.
- 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 446.
  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 448.

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 447.

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 447.

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.

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 374.

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 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 451.

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 455.
7. In the **Summary** panel, click **Finish**.

   See "Summary panel options" on page 456.
8. In the **Result** panel, click **OK**.

See "About high availability fire drills" on page 450.

See "Running the disaster recovery fire drill" on page 453.

See "Editing a fire drill schedule" on page 456.

See "Deleting fire drill schedules" on page 459.

See "Enabling fire drill schedules" on page 459.

See "Disabling fire drill schedules" on page 460.

About disaster recovery fire drills

The disaster recovery (DR) fire drill feature is a DR solution in 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 453.

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 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 452.

Running the disaster recovery fire drill

In 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 454.

6 In the **Schedule** wizard panel, specify the schedule to run the DR fire drill and click **Next**.

See "Schedule panel options" on page 455.

7 In the **Summary** panel, verify the selections that you have made and click **Finish**.

See "Summary panel options" on page 456.

8 In the **Result** panel, click **OK**.

See "About high availability fire drills" on page 450.

See "Running the high availability fire drill" on page 451.

See "Editing a fire drill schedule" on page 456.

See "Deleting fire drill schedules" on page 459.

See "Enabling fire drill schedules" on page 459.

See "Disabling fire drill schedules" on page 460.

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 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.

Table 23-1 Select remote clusters and global service groups panel options *(continued)*

Field	Description
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.
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 453.

See "About configuring a fire drill service group" on page 452.

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 451.

See "Running the disaster recovery fire drill" on page 453.

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 456.

See "Running the high availability fire drill" on page 451.

See "Running the disaster recovery fire drill" on page 453.

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 457.

7 In the **Summary** panel, verify the selections that you have made. Click **Finish**.

See "Summary panel options" on page 456.

8 In the **Result** panel, click **OK**.

See "Edit Schedule panel options" on page 457.

See "Deleting fire drill schedules" on page 459.

See "Enabling fire drill schedules" on page 459.

See "Disabling fire drill schedules" on page 460.

See "Running the disaster recovery fire drill" on page 453.

See "Running the disaster recovery fire drill" on page 453.

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 456.

See "Viewing fire drill schedules" on page 461.

See "Deleting fire drill schedules" on page 459.

See "Running the high availability fire drill" on page 451.

See "Running the disaster recovery fire drill" on page 453.

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 "Viewing fire drill schedules" on page 461.

See "Editing a fire drill schedule" on page 456.

See "Enabling fire drill schedules" on page 459.

See "Disabling fire drill schedules" on page 460.

See "Running the disaster recovery fire drill" on page 453.

See "Running the disaster recovery fire drill" on page 453.

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 "Viewing fire drill schedules" on page 461.

See "Editing a fire drill schedule" on page 456.

See "Deleting fire drill schedules" on page 459.

See "Disabling fire drill schedules" on page 460.

See "Running the disaster recovery fire drill" on page 453.

See "Running the disaster recovery fire drill" on page 453.

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 "Viewing fire drill schedules" on page 461.

See "Editing a fire drill schedule" on page 456.

See "Deleting fire drill schedules" on page 459.

See "Enabling fire drill schedules" on page 459.

See "Running the disaster recovery fire drill" on page 453.

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 456.

See "Deleting fire drill schedules" on page 459.

See "Enabling fire drill schedules" on page 459.

See "Disabling fire drill schedules" on page 460.

See "Running the disaster recovery fire drill" on page 453.

See "Running the disaster recovery fire drill" on page 453.

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 InfoScale Operations Manager console, you can create, edit, delete, and run the recovery plans. The InfoScale Operations Manager console provides the progress and the status of a recovery plan.

See "Creating recovery plans" on page 464.

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 465.

4 On the **Result** panel, verify that the recovery plan has been successfully created. Click **OK**.

See "Editing recovery plans" on page 468.

See "Deleting a recovery plan" on page 472.

See "Running recovery plans" on page 470.

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 464.

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 468.
5. On the **Summary** panel verify that the recovery plan has been successfully edited, click **OK**.

See "Creating recovery plans" on page 464.

See "Deleting a recovery plan" on page 472.

See "Running recovery plans" on page 470.

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.
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 468.

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 471.

5 On the **Result** panel, verify that the recovery plan was run successfully. Click **OK**.

See "Creating recovery plans" on page 464.

See "Editing recovery plans" on page 468.

See "Deleting a recovery plan" on page 472.

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.

Table 24-6 Running recovery plans *(continued)*

Field	Description
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. 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 470.

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 464.

See "Editing recovery plans" on page 468.

See "Running recovery plans" on page 470.

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 473.

See "Creating recovery plans" on page 464.

See "Editing recovery plans" on page 468.

See "Deleting a recovery plan" on page 472.

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 464.

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 475.

See "Creating recovery plans" on page 464.

See "Editing recovery plans" on page 468.

See "Deleting a recovery plan" on page 472.

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.
Timeout (mins)	View the timeout value for the task that is selected.

See "Creating recovery plans" on page 464.

About recovery plan log files

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 463.

Chapter 25

Managing application configuration

This chapter includes the following topics:

- About Application Configuration
- Prerequisites for application configuration
- Launching the configure application wizard from InfoScale Operations Manager

About Application Configuration

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 478.

See "Launching the configure application wizard from InfoScale Operations Manager" on page 478.

Prerequisites for application configuration

Following are the prerequisites for using the application configuration feature in 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 *InfoScale Cluster Server Release Notes*.

See "About Application Configuration" on page 477.

See "Launching the configure application wizard from InfoScale Operations Manager" on page 478.

Launching the configure application wizard from InfoScale Operations Manager

Using 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 477.

See "Prerequisites for application configuration" on page 478.

Chapter 26

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 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	Update 8.6	
	SUSE Linux Enterprise 12	SP5	
	SUSE Linux Enterprise 15	SP3, SP4	
	Oracle Linux	Update 8.6, 9.0	
	For Linux, only 64-bit operating systems are supported on the AMD Opteron or the Intel Xeon EM64T (x86_64) processor line.		
	AIX	AIX 7.2	TL4, TL5
		AIX 7.3	TL0, TL1
Microsoft Windows	Windows Server	2019, 2022	

See "Features of Multi-Site Management" on page 481.

See "Prerequisites of Multi-Site Management " on page 482.

See "Limitations of Multi-Site Management " on page 482.

Features of Multi-Site Management

The key features of Multi-Site Management in 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 InfoScale Operations Manager database.

See "Setting up a campus cluster" on page 482.

See "Setting up a replicated data cluster" on page 484.

Prerequisites of Multi-Site Management

Before you start using the Multi-Site Management feature in InfoScale Operations Manager, ensure that the following prerequisites are met:

- All hosts of the cluster must be using Storage Foundation High Availability version 8.0, or later.
- For non-Windows platforms, the Management Server and the managed hosts must be using InfoScale Operations Manager 8.0.2, or later. For Windows platform, the Management Server and the managed hosts must be using InfoScale Operations Manager 8.0, or later.

See "About Multi-Site Management " on page 480.

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 480.

Setting up a campus cluster

For campus cluster configuration, the steps are as follows:

Table 26-1 Setting up a campus cluster

Steps	Action	Description
Step 1	Configure Stretch Sites using the 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 InfoScale Operations Manager Management Server console. See "Configuring stretch sites" on page 484. 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), 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 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 484.
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, InfoScale Operations Manager still shows the site tags on the InfoScale Operations Manager console. To get the updated status, wait for 30 minutes, or refresh the host on the 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 481.

See "Setting up a replicated data cluster" on page 484.

Setting up a replicated data cluster

For replicated data cluster configuration, the steps are as follows:

Table 26-2 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 Veritas 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 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 484.

For more information, refer to *Storage Foundation High Availability Solutions Disaster Recovery Implementation Guide*.

See "Features of Multi-Site Management" on page 481.

See "Setting up a campus cluster" on page 482.

Configuring stretch sites

You can use the 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 InfoScale Operations Manager console. Refer to the *Storage Foundation High Availability Solutions Disaster Recovery Implementation Guide* for more information.

To configure stretch sites using 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 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 486.
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 486.
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 487.

6 In the **Specify Site Fencing Preference** panel, set the preferences for the sites.

See "Site fencing preference panel options" on page 487.

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 481.

See "Setting up a campus cluster" on page 482.

See "Setting up a replicated data cluster" on page 484.

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 484.

Select cluster type panel options

Use this panel to specify the type of cluster for the site assignment operation.

Table 26-3 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).

Table 26-3 Select cluster type panel options *(continued)*

Field	Description
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 484.

See "Setting up a campus cluster" on page 482.

See "Setting up a replicated data cluster" on page 484.

Assign site to the systems panel options

Use this panel to assign the site tag for the systems of the selected cluster.

Table 26-4 Assign site to the systems panel options

Field	Description
System	Displays the systems from the selected cluster.
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 484.

See "Setting up a campus cluster" on page 482.

See "Setting up a replicated data cluster" on page 484.

Site fencing preference panel options

Use this panel to set the site preferences for the systems of the selected cluster.

Table 26-5 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 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 484.

See "Setting up a campus cluster" on page 482.

See "Setting up a replicated data cluster" on page 484.

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 378.	Single
Closing a cluster configuration See "Closing a cluster configuration" on page 379.	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 379.	Single
Setting up GCO between clusters See "Adding a remote cluster to a local cluster" on page 443.	Multiple
Importing type definition See "Importing a type definition" on page 381.	Single
Removing GCO between clusters See "About removing a remote cluster from a global cluster setup" on page 447.	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 419.	Single and multiple
Unfreezing systems See "Unfreezing a system" on page 420.	Single and multiple
Editing attributes of a system See "Editing attributes of a system" on page 421.	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 421.	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 422.	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 385.	Single
Modifying service groups See "About modifying a service group" on page 416.	Single
Enabling service groups See "Enabling service groups" on page 390.	Single and multiple
Disabling service groups See "Disabling service groups" on page 391.	Single and multiple

Table A-3 Service group operations *(continued)*

Operation	Supported for single or multiple service groups?
Modifying the System List for a service group See "Modifying the system list for a service group" on page 413.	Single
Enabling all resources in a service group See "Enabling all resources of service groups" on page 397.	Single and multiple
Disabling all resources in a service group See "Disabling all resources of service groups" on page 398.	Single and multiple
Bringing service groups online See "Bringing service groups online" on page 404.	Single and multiple
Taking service groups offline See "Taking service groups offline" on page 406.	Single and multiple
Freezing service groups See "Freezing service groups" on page 394.	Single and multiple
Unfreezing service groups See "Unfreezing service groups" on page 395.	Single and multiple
Linking service groups See "Linking service groups in a cluster" on page 399.	Single
Unlinking service groups See "Unlinking service groups" on page 402.	Single
Switching service groups See "Switching service groups" on page 408.	Single and multiple
Flushing a service group See "Flushing service groups" on page 396.	Single

Table A-3 Service group operations *(continued)*

Operation	Supported for single or multiple service groups?
Autoenabling a service group See "Autoenabling service groups" on page 394.	Single
Clearing faults on a service group See "Clearing faults on service group" on page 410.	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 411.	Single and multiple
Editing attributes of service groups See "Editing attributes of service groups" on page 412.	Single
Converting local service groups to global service groups See "Converting local service groups to global service groups" on page 443.	Single
Converting global service groups to local service groups See "Converting global service groups to local service groups" on page 446.	Single
Running the high availability fire drill for the selected service groups See "Running the high availability fire drill" on page 451.	Single and multiple
Running the disaster recovery fire drill for the selected service groups See "Running the disaster recovery fire drill" on page 453.	Single and multiple
Deleting service groups See "Deleting service groups" on page 398.	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 436.	Single
Enabling resources See "Enabling resources" on page 425.	Single and multiple
Disabling resources See "Disabling resources" on page 426.	Single and multiple
Editing attributes of a resource See "Editing attributes of a resource" on page 433.	Single
Editing attributes of a resource type See "Editing attributes of a resource type" on page 434.	Single
Bringing resources online See "Bringing resources online" on page 430.	Single and multiple
Taking resources offline See "Taking resources offline" on page 430.	Single and multiple
Taking resources offline and propagating the state See "Taking a resource offline and propagating the state" on page 429.	Single and multiple
Probing the resources See "Probing resources" on page 428.	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 427.	Single and multiple
Linking resources in a service group See "Linking resources in a service group" on page 435.	Single
Unlinking resources in a service group See "Unlinking resources in a service group" on page 435.	Single
Invoking a resource action See "Invoking a resource action" on page 431.	Single
Marking a resource as critical See "Marking a resource as critical" on page 437.	Single and multiple
Marking a resource as non critical See "Marking a resource as non critical" on page 438.	Single and multiple
Deleting resources See "Deleting resources" on page 426.	Single and multiple

Section 5

Monitoring Storage Foundation HA licenses in the data center

- Chapter 27. Managing licenses
- Chapter 28. Viewing deployment information

Chapter 27

Managing licenses

This chapter includes the following topics:

- About licenses
- About Arctera 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
- Notification about a corrupted VxFS file system
- Certificate expiry detection and renewal feature

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 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 Arctera licensing and pricing" on page 498.

See "About assigning price tiers to hosts" on page 507.

See "Viewing the overview of SFHA licenses in the data center" on page 522.

See "Viewing the deployment details" on page 525.

See "Viewing the deployment policy details in the data center" on page 526.

About Arctera licensing and pricing

The Arctera storage and server management high availability products which operate on UNIX (Solaris, 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 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 27-1 Types of Arctera price tiers

Tier Type	Platforms supported	Description
Server Price Tier	UNIX: Solaris, and IBM AIX	The Arctera 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. Arctera 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 Arctera storage and server management, high availability products on UNIX (Solaris and 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 27-1 Types of Arctera price tiers *(continued)*

Tier Type	Platforms supported	Description
Processor Price Tier	All platforms	The Arctera 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 Arctera 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 27-1 Types of Arctera price tiers *(continued)*

Tier Type	Platforms supported	Description
Symantec Performance Value Unit (SPVU)	All platforms	The following are considered for calculating SPVU for the Arctera 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 501.

See "About licenses" on page 497.

See "Viewing the overview of SFHA licenses in the data center" on page 522.

About the Symantec Performance Value Unit

From the Storage Foundation and high availability 6.0 release, Arctera 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 7.2 is the minimum version of the 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 InfoScale Operations Manager managed host version 7.2 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. 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 Arctera licensing and pricing" on page 498.

See "About the Symantec Performance Value Unit for VMware virtual machines" on page 502.

See "About the Symantec Performance Value Unit for Solaris LDOM virtualization server" on page 504.

See "About the Symantec Performance Value Unit for kernel-based virtual machines" on page 505.

See "About the Symantec Performance Value Unit for IBM LPAR" on page 506.

About the Symantec Performance Value Unit for VMware virtual machines

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 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 InfoScale Operations Manager using VMware vCenter by navigating to **Settings** > **Virtualization**.
- Add the virtual machines separately to InfoScale Operations Manager as managed hosts.

Note: If you add the virtual machines to InfoScale Operations Manager as managed hosts before you add the VMware ESX virtualization server, it can take up to 24 hours for InfoScale Operations Manager to calculate the SPVUs. Similarly, if you add only the virtual machines to InfoScale Operations Manager, and do not add the virtualization server, InfoScale Operations Manager cannot calculate the SPVUs for the products that are installed on the virtualization server. However, 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 501.

See "About the Symantec Performance Value Unit for Solaris LDOM virtualization server" on page 504.

See "About the Symantec Performance Value Unit for kernel-based virtual machines" on page 505.

See "About the Symantec Performance Value Unit for IBM LPAR" on page 506.

About the Symantec Performance Value Unit for Solaris LDOM virtualization server

To enable 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 InfoScale Operations Manager as a managed host.
- Add the guest domains to InfoScale Operations Manager as managed hosts.

Note: If you add the domains to InfoScale Operations Manager as managed hosts before you add the LDOM virtualization server, it can take up to 24 hours for InfoScale Operations Manager to calculate the SPVUs. Similarly, if you add only the domains to InfoScale Operations Manager, and do not add the LDOM virtualization server, InfoScale Operations Manager cannot calculate the SPVUs for the products that are installed on the virtualization server. However, 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, 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 501.

See "About the Symantec Performance Value Unit for VMware virtual machines" on page 502.

See "About the Symantec Performance Value Unit for kernel-based virtual machines" on page 505.

See "About the Symantec Performance Value Unit for IBM LPAR" on page 506.

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 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 501.

See "About the Symantec Performance Value Unit for VMware virtual machines" on page 502.

See "About the Symantec Performance Value Unit for Solaris LDOM virtualization server" on page 504.

See "About the Symantec Performance Value Unit for IBM LPAR" on page 506.

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 InfoScale Operations Manager managed host version 7.2.

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 501.

See "About the Symantec Performance Value Unit for VMware virtual machines" on page 502.

See "About the Symantec Performance Value Unit for Solaris LDOM virtualization server" on page 504.

See "About the Symantec Performance Value Unit for kernel-based virtual machines" on page 505.

About the per-core licensing

From the Storage Foundation High Availability 7.0 release, Arctera 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 Arctera licensing and pricing" on page 498.

See "About the Symantec Performance Value Unit for VMware virtual machines" on page 502.

See "About the Symantec Performance Value Unit for Solaris LDOM virtualization server" on page 504.

See "About the Symantec Performance Value Unit for kernel-based virtual machines" on page 505.

See "About the Symantec Performance Value Unit for IBM LPAR" on page 506.

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 508.

About license deployment policies

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, 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 InfoScale Operations Manager console.

See "Creating a license deployment policy" on page 510.

Assigning a price tier to a host automatically

You can assign price tiers to a single host or multiple hosts automatically in 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 507.

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 509.

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 510.

6 In the **Assign price tier manually - Result** panel, click **OK**.

See "About assigning price tiers to hosts" on page 507.

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 27-2 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 and 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.

Table 27-2 Select tier values panel options *(continued)*

Field	Description
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 508.

See "Select hosts to apply same tier values panel options" on page 510.

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 508.

See "Select tier values panel options" on page 509.

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 513.

See "Deleting a license deployment policy" on page 514.

Create policy - Details panel options

Use this wizard panel to create a license deployment policy.

Table 27-3 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.

Table 27-3 Create policy - Details panel options *(continued)*

Field	Description
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 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.
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.

Table 27-3 Create policy - Details panel options *(continued)*

Field	Description
Risk threshold	Enter the number of license deployments for which 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 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 510.

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 513.

5 In the **Edit thresholds - Result** panel, click **OK**.

See "Creating a license deployment policy" on page 510.

See "Deleting a license deployment policy" on page 514.

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 27-4 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 InfoScale Operations Manager must generate a warning.
Fault threshold	Modify the number of license deployments for which 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 513.

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**.

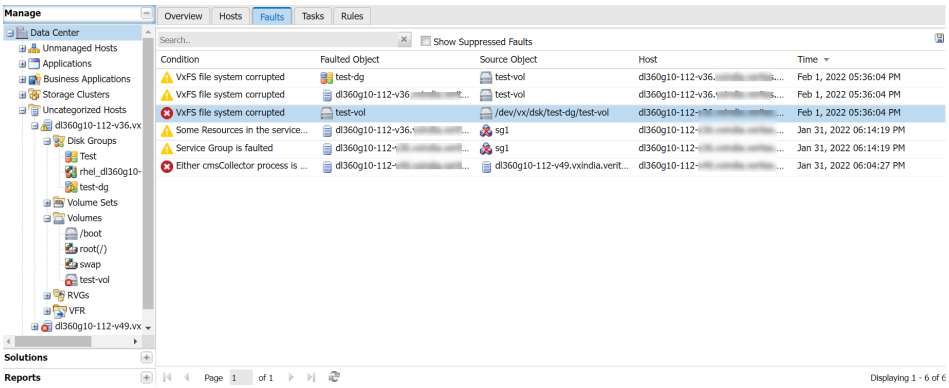
5 In the **Delete policy - Result** panel, click **OK**.

See "Creating a license deployment policy" on page 510.

See "Modifying a license deployment policy" on page 513.

Notification about a corrupted VxFS file system

Arctera InfoScale Operations Manager discovers FULLFSCK flag on a VxFS file system every 24 hours. If the file system is corrupted, a fault 'SF_FILESYSTEM_CORRUPTED' is raised.

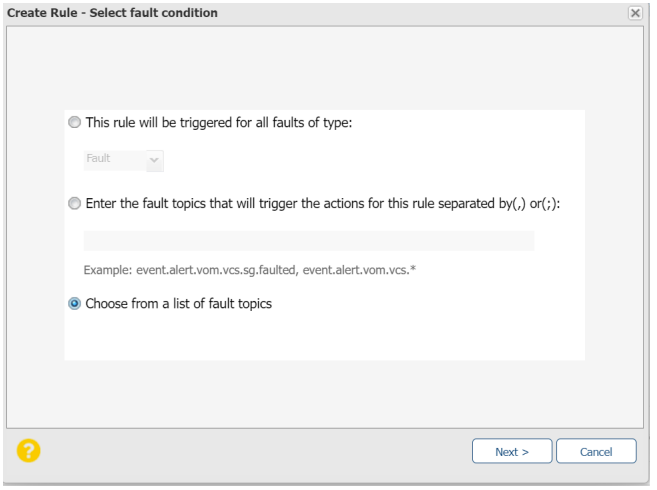


You can create a rule to get an email/SNMP notification about this fault.

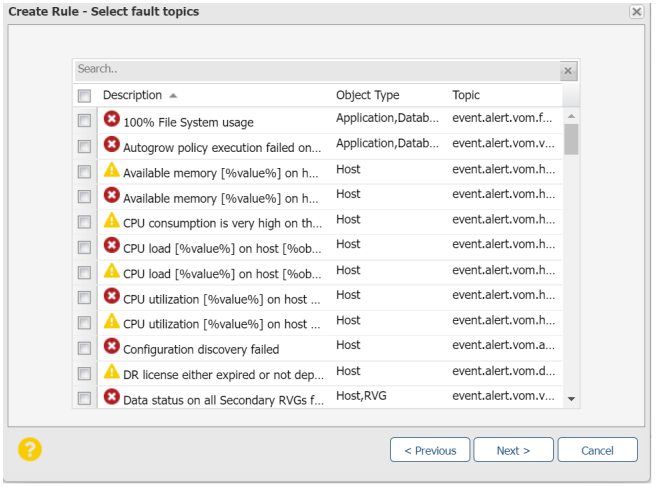
Complete the following steps

1 In the Server perspective, click **Rules** and right-click to create a new rule.

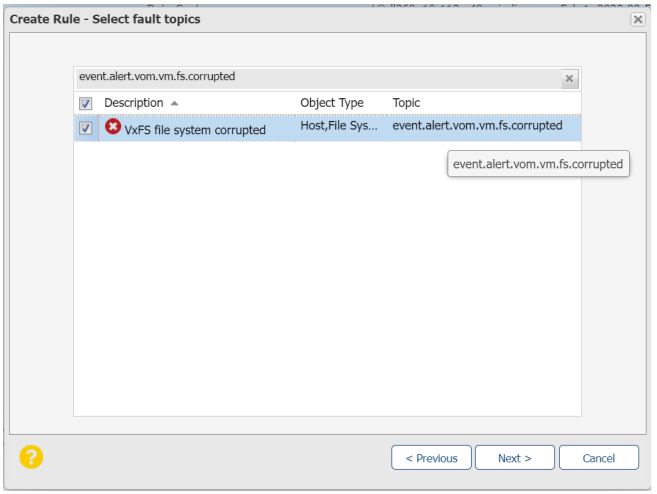
2 Select **Choose from a list of fault topics** in the following screen. Click **Next**.



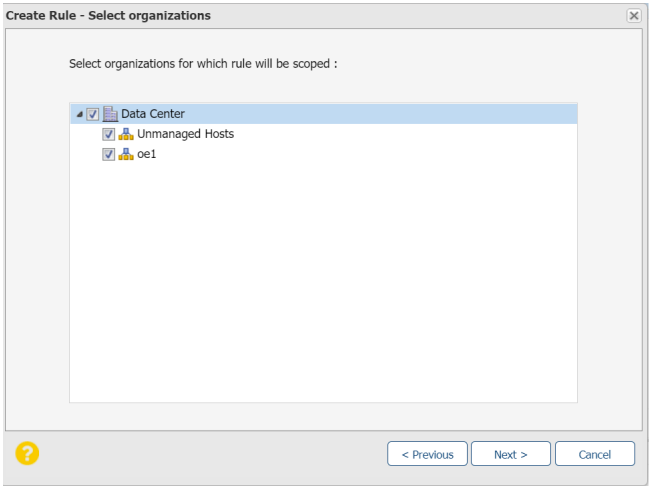
3 Enter **VxFS file system corrupted** to search this topic or select **event.alert.vom.vm.fs.corrupted** in **Topic** .



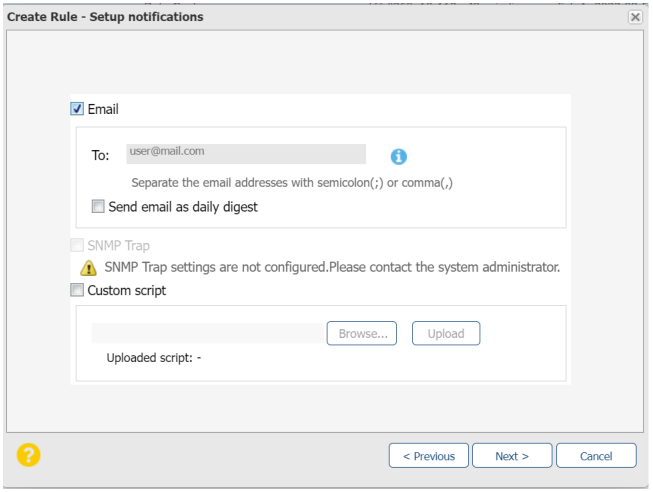
4 Select the rule as indicated in the following screen. Click **Next**.



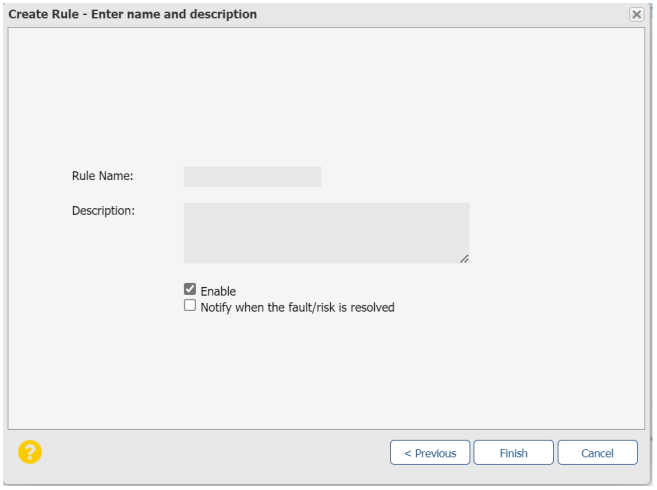
5 Select Organizations for which you want to send this notification. Click **Next**.



6 Select **Email** and enter the email addresses of those who want to receive this notification. You can enter multiple email addresses as indicated. Click **Next**.



7 Enter a **Name** and **Description** for this rule. Optionally, you can click the check box next to **Notify when the fault/risk is resolved**. If you select this option, an email notification is sent when the file system is back. Click **Finish**.



You have successfully created a rule for email notifications to be sent when a file system is corrupted. Email notification contains the corrupted file system path.

Certificate expiry detection and renewal feature

InfoScale Operations Manager provides a provision to detect and renew the Management Server and the InfoScale Operations Manager Agents' certificates.

During the HOST discovery, InfoScale Operations Manager performs the certificate expiry detection. If the expiry date of the certificate is found 120 days or less, then below fault appears on the InfoScale Operations Manager GUI with number of remaining expiry days.

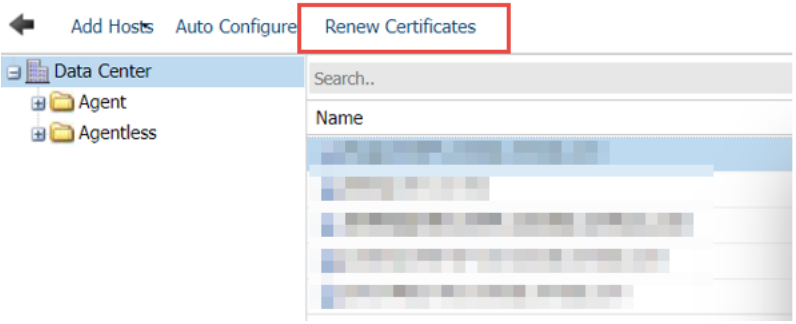
"InfoScale Operations Manager certificates will be expired within %days_remaining% days. Renew certificates"

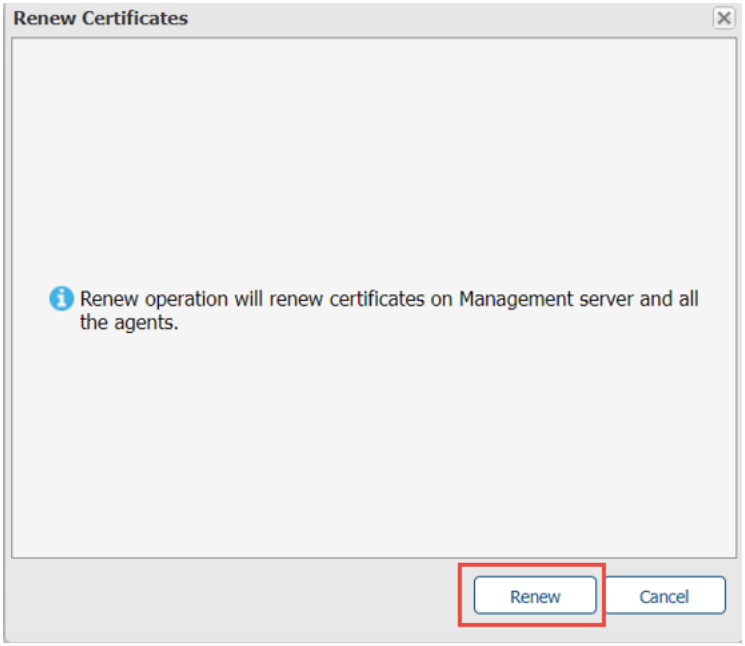
Whenever you see the fault in InfoScale Operations Manager GUI, you need to perform renew certificates operation to renew the certificates. If you fail to renew the certificates in 120 days, then InfoScale Operations Manager Management Server and Agents may stop working.

To renew certificates

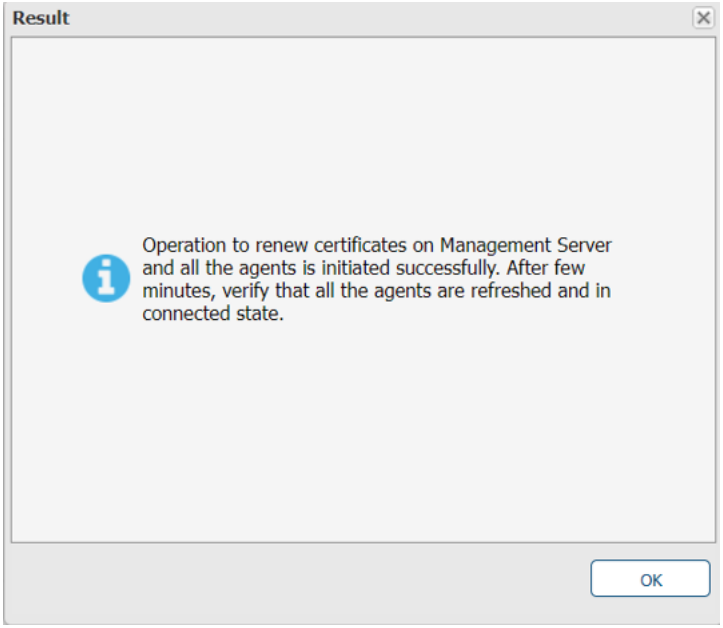
You can start the certificate renewal process when your certificates are about to expire by performing the following steps:

1. On the InfoScale Operations Manager Console, click **Settings** > **Host**.
2. Click the **Renew Certificate** option at the top of the menu bar. The **Renew Certificates** pop-up window opens.





3. Click the **Renew** button to renew the certificates.



Once Management Server and Agents certificates are successfully renewed, the fault disappears.

Note: While performing certificate renew operation, InfoScale Operations Manager validates whether all the agents are in the connected and reachable state. In case if any Agent is not connected or not reachable then you must need to fix the communication issue to renew the certificates.

---

4. If you are using third party certificates for `sfm_agent` then InfoScale Operations Manager will detect it. In this case, renew certificate operation will renew certificates on Management Server only and you have to manually run the `import_sfm_agent_certificate.pl` script to update `sfm_agent` certificates for InfoScale Operations Manager Agents.

Chapter 28

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 InfoScale Operations Manager 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 28-1
- **License deployment summary** table
  Table 28-2
- **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 28-1 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 28-2 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 524.
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 28-2 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 497.

See "Viewing the deployment details" on page 525.

See "Viewing the deployment policy details in the data center" on page 526.

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 522.

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 524.

See "About assigning price tiers to hosts" on page 507.

See "Assigning a price tier to a host manually" on page 508.

See "About the Symantec Performance Value Unit" on page 501.

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 510.

See "Modifying a license deployment policy" on page 513.

See "Deleting a license deployment policy" on page 514.

See "About license deployment policies" on page 508.

See "Viewing the overview of SFHA licenses in the data center" on page 522.

See "Viewing the deployment details" on page 525.

Viewing the InfoScale Operations Manager Deployment Report

This report provides information on hosts that are connected to Management Server and whether they have Arctera 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, 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, 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, or Dynamic Multipathing is installed on the host.
- The edition of the Storage Foundation, Cluster Server, or Dynamic Multipathing 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 Deployment Report

1 In the Home page on the Management Server console, click **SFHA Licensing**.

2 Click the **Reports** tab.

3 Click **Deployment Report**.

See "About reports" on page 120.

See "Subscribing for a report" on page 125.

See "Saving a report" on page 124.

See "Sending a report through email" on page 127.

Chapter 29

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 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 530.

See "About space estimation for data logs" on page 534.

About metered resources

Table 29-1 lists the resources that are metered and the difference in the log configurations of previous versions of InfoScale Operations Manager and 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 InfoScale Operations Manager 5.0. Whereas the metering in the previous versions of InfoScale Operations Manager was 5 minutes for one week.

Hosts metering on Windows platform is available in InfoScale Operations Manager version 6.0 and onwards. Table 29-2 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 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 29-1 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
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

Table 29-1 Log configurations for UNIX hosts *(continued)*

Resource	Chart name	Log configuration in the previous versions	Log configuration in version 5.0 and onwards
Host 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
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

Table 29-1 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 - 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 29-2 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 530.

See "About space estimation for data logs" on page 534.

See "About threshold settings" on page 552.

About space estimation for data logs

Table 29-3 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 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 29-4 lists the space estimation for data logs for host, file system, volume, and disk on Windows platform.

Table 29-3 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 29-3 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
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

Table 29-4 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

Table 29-4 Space estimation for data logs for Windows hosts *(continued)*

Name of resource	Number of resources	Metrics collected	Interval of collection	Duration of collection	Size in KB	Size in KB for a single object
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 530.

See "About metered resources" on page 530.

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 537.

See "About performance metering statistics" on page 530.

See "About metered resources" on page 530.

See "About space estimation for data logs" on page 534.

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 536.

See "About performance metering statistics" on page 530.

See "About metered resources" on page 530.

See "About space estimation for data logs" on page 534.

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 538.

See "About performance metering statistics" on page 530.

See "About metered resources" on page 530.

See "About space estimation for data logs" on page 534.

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 537.

See "About performance metering statistics" on page 530.

See "About metered resources" on page 530.

See "About space estimation for data logs" on page 534.

About InfoScale Operations Manager performance graphs

You can view the performance of the following objects by using interactive graphs.

Table 29-5

Perspective	Objects
Server	Host Disk Volume and file system Path Initiators
Virtualization	Virtualization server and virtual machine Path
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 540.

See "Viewing the performance graphs for a host" on page 541.

See "Viewing the performance graphs for a disk" on page 542.

See "Viewing the performance graphs for volume and file system" on page 544.

See "Viewing the performance graphs for a path" on page 546.

See "Viewing the performance graphs for an initiator" on page 547.

See "Viewing the performance graphs for virtualization server and virtual machines" on page 548.

See "Viewing the performance graphs for a path of a virtualization server" on page 549.

See "Viewing the performance graphs for an enclosure" on page 550.

Pre-requisite commands to view performance graphs for a resource

Table 29-6 lists the pre-requisite commands that are required to view performance graphs for a host.

Table 29-6 Pre-requisite commands for a host

Chart name	Pre-requisite commands
Available Memory	Solaris/Linux: `vmstat` AIX: `svmon`
Swap in Rate	VMware Guest/ESX: `SWPR/S` Solaris: `sar` AIX: `vmstat` Linux: `sar`

Table 29-6 Pre-requisite commands for a host *(continued)*

Chart name	Pre-requisite commands
Used Swap	Solaris/AIX: `swap` Linux: `cat /proc/swaps` 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 InfoScale Operations Manager performance graphs" on page 539.

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 29-7 lists the performance graphs for a host.

Table 29-7 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.

Table 29-7 Performance graphs for a host *(continued)*

Performance graph name	Description
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 InfoScale Operations Manager performance graphs" on page 539.

See "Enable performance metering for a host" on page 536.

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 29-8 lists the performance graphs for a disk.

Table 29-8 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.

Table 29-8 Performance graphs for a disk *(continued)*

Performance graph name	Description
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 InfoScale Operations Manager performance graphs" on page 539.

See "Enable performance metering for a host" on page 536.

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 29-9 lists the performance graphs for a volume.

Table 29-9 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 29-10 lists the performance graphs for a file system.

Table 29-10 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.

Table 29-10 Performance graphs for a file system *(continued)*

Performance graph name	Description
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 InfoScale Operations Manager performance graphs" on page 539.

See "Enable performance metering for a host" on page 536.

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 29-11 lists the performance graphs for a path.

Table 29-11 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 InfoScale Operations Manager performance graphs" on page 539.

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 29-12 lists the performance graphs for an initiator.

Table 29-12 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.
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 InfoScale Operations Manager performance graphs" on page 539.

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 29-13 lists the performance graphs for virtualization server and virtual machine.

Table 29-13 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.

Table 29-13 Performance graphs for virtualization server and virtual machine *(continued)*

Performance graph name	Description
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 InfoScale Operations Manager performance graphs" on page 539.

See "Enable performance metering for a host" on page 536.

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 29-14 lists the performance graphs for a path of a virtualization server.

Table 29-14 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 InfoScale Operations Manager performance graphs" on page 539.

See "Viewing the performance graphs for virtualization server and virtual machines" on page 548.

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 8.0.2 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 29-15 lists the performance graphs for an enclosure.

Table 29-15 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 InfoScale Operations Manager performance graphs" on page 539.

About threshold settings

Using InfoScale Operations Manager 9.1, you can set the threshold values on the respective objects for the metrics that are described in Table 29-16. A Risk or Fault is raised when the threshold condition is violated.

Managed host version 7.2 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 7.2 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 29-16 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.

Table 29-16 Objects, metrics, and evaluation intervals *(continued)*

Object	Metrics	Evaluation intervals
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 29-16. 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 553.

See "Deleting the threshold settings for an object" on page 557.

See "Enabling the threshold settings for an object" on page 560.

See "Disabling the threshold settings for an object" on page 563.

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 552.

See "Deleting the threshold settings for an object" on page 557.

See "Enabling the threshold settings for an object" on page 560.

See "Disabling the threshold settings for an object" on page 563.

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 552.

See "Adding threshold settings for an object" on page 553.

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 552.

See "Disabling the threshold settings for an object" on page 563.

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 552.

See "Enabling the threshold settings for an object" on page 560.

Chapter 30

Anomaly detection and monitoring

This chapter includes the following topics:

- About anomaly detection
- Requirements
- Enabling anomaly detection
- Training the ML engine
- Viewing graphs and faults for anomalies

About anomaly detection

InfoScale provides an AI-powered anomaly detection system that employs advanced machine learning algorithms to provide comprehensive threat identification and system monitoring capabilities. The system continuously monitors file system statistics and entropy patterns to establish baseline profiles, and then applies machine learning models to detect deviations from established trends. The system analyzes data points over time to detect established patterns of behavior and identifies outliers as anomalies.

The following file system attributes are monitored:

- Rename count
- Read count
- Write count

Any deviations in these attributes are considered as an anomaly and reported in the InfoScale Operations Manager UI as trend graphs and real-time alerts are generated.

Entropy is calculated by sampling a percentage of data for unusual patterns. By default 10% of the data is sampled for deviation in data patterns. Changes in entropy values indicate potential malicious activities. These changes are flagged as anomalies and displayed in the UI as entropy graphs.

AI-based anomaly detection is available with InfoScale version 9.0.2 or later.

At a high level, AI-based anomaly detection includes the following stages:

- **Stage 1: Enable AI-based anomaly detection for file systems.**
  See "Enabling anomaly detection" on page 567.
- S**Stage 2: Train the ML engine on representative workloads.**
  See "Training the ML engine" on page 568.
- **Stage 3: View trend graphs and entropy variation graphs.**
  See "Viewing graphs and faults for anomalies " on page 568.

Requirements

AI-based anomaly detection is available with:

- InfoScale version 9.0.2 or later
- On a Linux Management Server
- On a Linux host

Enabling anomaly detection

You can enable anomaly detection for a file system created on a volume. When you enable anomaly detection, the following file system attributes are monitored:

- Rename count
- Read count
- Write count

Any deviations detected in these attributes are identified as anomalies and represented visually as graphs in the UI. Alerts are generated and shown in the UI.

To enable anomaly detection:

1. In the Management Server console, go to the **Server** perspective and expand **Manage** in the left pane.
2. Search or locate the host in the left pane.
3. Expand the host and expand **Volumes**.
4. Click the volume that has the file system that you want to monitor for anomalies.
5. Click **Anomaly Watch**. Select **Anomaly Detection** and when prompted for confirmation click **Yes**.

   Anomaly detection is enabled and a successful notification is displayed on the screen.

   After enabling anomaly detection, it takes the ML engine 30 days to learn baseline data patterns. During this training period, no graphs are shown in the UI. After the training duration is over, trend graphs for the monitored attributes are shown in the UI.

Training the ML engine

After you enable AI-based anomaly detection in the InfoScale Operations Manager UI, the ML engine must be calibrated with representative workloads that run in the customer environment. The training duration is 30 days during which the engine learns the data patterns seen in typical workloads that run in the customer's environment and establishes baseline profiles. After the training period, the subsequent data is compared against this baseline for detecting variations and identifying these as anomalies.

After the model is trained, it starts predicting on new incoming data in fixed-size batches to optimize computational efficiency and throughput.

Note: Predictions are processed in discrete batches, which may lead to minor inaccuracies, especially when anomalies span across multiple batches. Such cross-boundary anomalies might not be fully captured within the context of individual batches, potentially affecting detection accuracy.

Viewing graphs and faults for anomalies

You can view the trend graphs and faults in real time in the UI and receive email notifications for the detected anomalies.

Viewing graphs

You can view the graphs for entropy changes and for the anomalies that are detected during the read, write, and rename operations performed on a file system.

To view graphs:

1. In the Management Server console, go to the **Server** perspective and expand **Manage** in the left pane.
2. Search or locate the host in the left pane.
3. Expand the host and expand **Volumes**.
4. Click the volume that has the file system on which anomaly detection is enabled.
5. Click **Anomaly Watch**.

   The following trend graphs are shown:

   - **Rename Count**: Shows measure of rename operations analyzed over a time series.
     - Green contours: Represent the count of actual rename operations.
     - Blue contours: Represent the count of rename operations predicted by the ML engine.
     - Blue range: Represents the accepted range of rename operations predicted by the ML engine. Any count beyond this threshold is considered a deviation and shown as an anomaly indicated in red.
   - **Read Count**: Shows measure of read operations analyzed over a time series.
     - Green contours: Represent the count of actual read operations.
     - Blue contours: Represent the count of read operations predicted by the ML engine.
     - Blue range: Represents the accepted range of read operations predicted by the ML engine. Any count beyond this threshold is considered a deviation and shown as an anomaly indicated in red.
   - **Write count**: Shows measure of write operations analyzed over a time series.
     - Green contours: Represent the count of actual write operations.
     - Blue contours: Represent the count of write operations predicted by the ML engine.

- Blue range: Represents the accepted range of write operations predicted by the ML engine. Any count beyond this threshold is considered a deviation and shown as an anomaly indicated in red.

- **High Entropy**: Shows the entropy levels analyzed over a time series.
  - Green contour: Represents the entropy level for the actual data.
  - Blue range: Represents the accepted entropy range that is predicted by the ML engine.

Viewing faults and receiving email notifications

Faults are generated if the count for file system attributes deviates from the ML engine predicted range. If a fault occurs and rule is set with the fault topic **event.alert.vom.fs.anomaly**, email notification is sent to the recipients for the generated faults. Ensure that you provide the details for the SMTP settings before setting the email notification for a fault.

The following fault is generated if an anomaly is detected:

The *monitored-attributes* patterns observed on the *file-system-name* file system indicate a few anomalies. Verify the data on the file system. Note: This alert originates from an early-stage ML model and may produce false positives if it has not yet fully learned the relevant patterns.

Chapter 31

Managing Business Applications

This chapter includes the following topics:

- About Business Applications in 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 InfoScale Operations Manager

A Business Application is a collection of objects within 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: InfoScale Operations Manager can discover some applications, such as some database applications, and display information about their related objects. In cases where 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. 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 572.

See "Renaming a Business Application" on page 574.

See "Deleting a Business Application" on page 574.

See "Viewing Business Applications in the data center" on page 574.

See "Viewing the overview of a Business Application" on page 575.

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 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 InfoScale Operations Manager" on page 572.

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 InfoScale Operations Manager" on page 572.

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 InfoScale Operations Manager" on page 572.

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 InfoScale Operations Manager" on page 572.

See "Viewing the overview of a Business Application" on page 575.

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 576.

   See "Viewing data availability for a Business Application" on page 577.

   See "Viewing SAN connectivity for a Business Application" on page 578.

See "About Business Applications in InfoScale Operations Manager" on page 572.

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 InfoScale Operations Manager" on page 572.

See "Viewing the overview of a Business Application" on page 575.

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 InfoScale Operations Manager" on page 572.

See "Viewing the overview of a Business Application" on page 575.

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 InfoScale Operations Manager" on page 572.

See "Viewing the overview of a Business Application" on page 575.

About the makeBE script

Using makeBE command-line script, 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 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 580.

See "makeBE script CSV file details" on page 580.

See "makeBE script parameters details" on page 582.

See "Limitations of makeBE script" on page 582.

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 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 583.

See "Importing Business Application using the makeBE script" on page 584.

See "Exporting Business Application using the makeBE script" on page 585.

See "Updating Business Application using the makeBE script" on page 586.

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 579.

makeBE script parameters details

You can view the `makeBE` script-related parameters by typing the following command on the 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 579.

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 579.

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 579.

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 579.

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 579.

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 579.

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 579.

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 579.

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 InfoScale Operations Manager. This information about an object is in addition to the details that 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 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 *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 *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 590.

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 *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 591.

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 590.

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 *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 590.
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 InfoScale Operations Manager policy check uses individual rules, called policy signatures, to validate whether the data center configuration conforms to a predefined standard. InfoScale Operations Manager contains a set of predefined signatures. You can also create customized signatures.

Policy check scans identify the resources in the 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 594.

See "Registering policy signatures" on page 595.

See "Unregistering a signature" on page 596.

See "Setting signature tunables" on page 597.

See "Running a manual policy scan" on page 598.

See "Enabling or disabling policy signatures" on page 598.

See "Viewing policy violation details" on page 599.

See "Viewing or exporting a list of available policy signatures" on page 601.

See "About using custom signatures for policy checks" on page 601.

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 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 595.

See "About policy checks" on page 593.

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 597.

See "About policy checks" on page 593.

See "How signature registration settings work" on page 594.

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 593.

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 593.

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 593.

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 *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 593.

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 105.

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 593.

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 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 593.

About using custom signatures for policy checks

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 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 593.

See "Creating a custom signature script" on page 602.

See "Installing a custom signature script" on page 605.

See "Copying the custom signature script" on page 605.

See "Removing a custom signature" on page 606.

Creating a custom signature script

You can create a custom signature script for policy checks in 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 603.
See "About using custom signatures for policy checks" on page 601.
See "Installing a custom signature script" on page 605.
See "Copying the custom signature script" on page 605.
See "Removing a custom signature" on page 606.
See "About policy checks" on page 593.

Sample custom signature script

The following is an example of a custom signature script for policy checks in 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 602.

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 601.

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 601.

See "Creating a custom signature script" on page 602.

See "Installing a custom signature script" on page 605.

See "Removing a custom signature" on page 606.

Removing a custom signature

To remove a custom policy check signature from InfoScale Operations Manager, you can run the following utility:

```
pcreg --id policy check id
```

See "About using custom signatures for policy checks" on page 601.

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 *InfoScale Operations Manager Add-ons User Guide*.

See "Sample setup.pl script for the custom signature" on page 606.

See "Sample unsetup.pl script for a custom signature" on page 607.

See "About using custom signatures for policy checks" on page 601.

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 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 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. 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.

InfoScale Operations Manager also supports the discovery of EMC PowerPath/Virtual Edition (VE) of ESX Server and Microsoft MPIO . For these technologies, 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. 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 InfoScale Operations Manager to manage the paths for Windows managed hosts.

See "Disabling the DMP paths on the initiators of a host" on page 610.

See "Disabling the DMP paths on an enclosure" on page 616.

See "Disabling the DMP paths on a virtualization server" on page 617.

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 611.

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 614.

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 615.

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 616.
- 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 618.

See "About Dynamic Multi-Pathing in InfoScale Operations Manager" on page 609.

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 610.

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.	

Table 34-2 Define new case panel options for the enclosure *(continued)*

Field	Description
Name	The name of the enclosure.
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 616.

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 617.

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 610.

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 616.

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.
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 617.

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 610.

See "Disabling the DMP paths on an enclosure" on page 616.

See "Disabling the DMP paths on a virtualization server" on page 617.

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 610.

See "Disabling the DMP paths on an enclosure" on page 616.

See "Disabling the DMP paths on a virtualization server" on page 617.

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 611.

3 In the **Specify Array Ports** panel, select one or more array ports to perform DMP maintenance.

See "Object selection panel options" on page 614.

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 615.

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 616.
- 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 618.

See "About Dynamic Multi-Pathing in InfoScale Operations Manager" on page 609.

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 611.

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 614.

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 615.

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 616.
- 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 618.

See "About Dynamic Multi-Pathing in InfoScale Operations Manager" on page 609.

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 616.

2 In the Paths Re-Enable Confirmation panel, click **Re-Enable**.

See "Path re-enable panel options" on page 619.

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 620.

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 620.

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 619.

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 620.

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 620.

See "About Dynamic Multi-Pathing in InfoScale Operations Manager" on page 609.

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.
Active Paths	The number of the active DMP paths.

Table 34-9 Path re-enable confirmation panel options *(continued)*

Field	Description
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 618.

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 618.

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 618.

Removing a completed DMP maintenance case record

The 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 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 InfoScale Operations Manager" on page 609.

Reviewing the output and results of a completed DMP maintenance case

Using the 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 620.

See "About Dynamic Multi-Pathing in InfoScale Operations Manager" on page 609.

Chapter 35

Managing CVM clusters

This chapter includes the following topics:

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

About monitoring and managing CVM clusters in 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 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.

Table 35-1 Monitoring and managing CVM clusters in the Server perspective *(continued)*

Operation or view	Support in the Server perspective
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 625.

See "About Flexible Storage Sharing" on page 628.

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 InfoScale Operations Manager " on page 623.

Chapter 36

Managing Flexible Storage Sharing

This chapter includes the following topics:

- Implementing Flexible Storage Sharing with InfoScale Operations Manager
- About Flexible Storage Sharing
- Flexible Storage Sharing use cases
- Flexible Storage Sharing features and support in 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 InfoScale Operations Manager

To implement Flexible Storage Sharing (FSS) with InfoScale Operations Manager use the following process.

Table 36-1 Implementing Flexible Storage Sharing with InfoScale Operations Manager

Task	For more information
Understand requirements and use cases for FSS	See "About Flexible Storage Sharing" on page 628. See "Flexible Storage Sharing use cases" on page 629. See "Flexible Storage Sharing features and support in InfoScale Operations Manager " on page 631.
Export disks for use in an FSS-capable cluster	See "Exporting and un-exporting disks for Flexible Storage Sharing" on page 632.
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 633. See "Adding disks to disk groups" on page 166.

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.

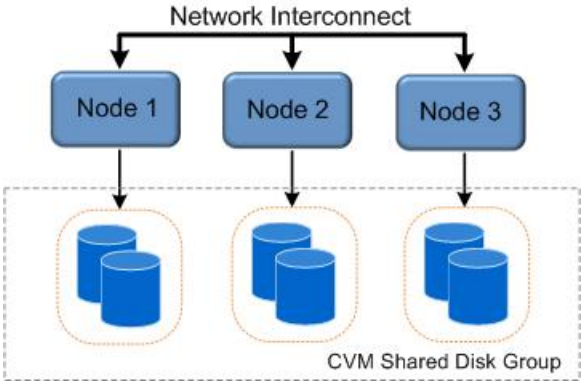
FSS allows growing the existing storage by growing the LUN, which is termed as Dynamic LUN Expansion (DLE). To use the DLE feature, you resize the appropriate disks using the Management Server console, which internally invokes the InfoScale `disk resize` command. The command intelligently detects the remote disks and gets the required protocol executed to complete the LUN expansion; you don't need to specify any additional options. No explicit master switching is required.

You can leverage DLE to grow the storage on Linux hosts in cloud environments.

See "Resizing disks in disk groups" on page 168.

For information on FSS support for you platform and InfoScale configuration, refer to the platform-specific *InfoScale 9.1 Software Compatibility List*.

Figure 36-1 Flexible Storage Sharing Environment



See "Flexible Storage Sharing use cases" on page 629.

See "Flexible Storage Sharing features and support in InfoScale Operations Manager " on page 631.

See "Implementing Flexible Storage Sharing with InfoScale Operations Manager" on page 627.

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 628.

Flexible Storage Sharing features and support in InfoScale Operations Manager

Flexible Storage Sharing (FSS) is a feature of Storage Foundation Cluster File System High Availability (SFCFSHA). FSS is supported in InfoScale Operations Manager as follows:

- InfoScale Operations Manager checks if Cluster Volume Manager (CVM) clusters are FSS capable and if the cluster nodes have the required managed host version (InfoScale Operations Manager 8.0 or later) in order to perform FSS operations from InfoScale Operations Manager.
- 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.
- 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 628.

See "About monitoring and managing CVM clusters in InfoScale Operations Manager " on page 623.

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 633.

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 InfoScale Operations Manager" on page 627.

See "Permissions required for views and operations on CVM cluster objects" on page 625.

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 InfoScale Operations Manager" on page 627.

See "Permissions required for views and operations on CVM cluster objects" on page 625.

Chapter 37

Monitoring the virtualization environment

This chapter includes the following topics:

- About the virtualization technologies supported
- About discovering the VMware Infrastructure using InfoScale Operations Manager
- About discovering Solaris zones
- About discovering logical domains in InfoScale Operations Manager
- About discovering LPARs and VIOs in InfoScale Operations Manager
- About Microsoft Hyper-V virtualization discovery
- About the Kernel-based Virtual Machine (KVM) virtualization discovery in InfoScale Operations Manager
- About the reports related to virtualization

About the virtualization technologies supported

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, 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, 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 InfoScale Operations Manager Management Server domain, provides storage mapping discovery.

For Kernel-based Virtual Machine (KVM) discovery, 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). 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 InfoScale Operations Manager" on page 637.

See "About discovering Solaris zones" on page 645.

See "About discovering logical domains in InfoScale Operations Manager" on page 648.

See "About discovering LPARs and VIOs in InfoScale Operations Manager" on page 652.

See "About Microsoft Hyper-V virtualization discovery" on page 655.

About discovering the VMware Infrastructure using InfoScale Operations Manager

In 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 *InfoScale Operations Manager Hardware and Software Compatibility List (HSCL)*.

See "About the virtualization technologies supported" on page 635.

See "How InfoScale Operations Manager discovers vCenter and ESX servers" on page 637.

How InfoScale Operations Manager discovers vCenter and ESX servers

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 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 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 InfoScale Operations Manager and the data discovered using 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 *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 InfoScale Operations Manager" on page 637.

See "Information that InfoScale Operations Manager discovers on the VMware Infrastructure components" on page 638.

See "About the datastores in InfoScale Operations Manager" on page 640.

Information that 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 InfoScale Operations Manager discovers.
- Operating system handles of the VMware Infrastructure components that 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 InfoScale Operations Manager" on page 637.

See "How InfoScale Operations Manager discovers vCenter and ESX servers" on page 637.

See "About the datastores in InfoScale Operations Manager" on page 640.

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 643.

See "About discovering the VMware Infrastructure using InfoScale Operations Manager" on page 637.

See "About the virtualization technologies supported" on page 635.

About the datastores in 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. 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 InfoScale Operations Manager" on page 637.

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 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 641.

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 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 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	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 641.

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 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, 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	8.0.2, or later
Virtual machine powered off	VM powered off	8.0.2, 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	8.0.2, or later
Virtual machine suspended	VM suspended	8.0.2, or later
Virtual machine created	VM created	8.0.2, or later
Virtual machine migrated Hot migration: A powered-on virtual machine is migrated from one ESX server to another ESX server.	VM migrated	8.0.2, 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	8.0.2, or later
Virtual machine renamed	VM renamed	8.0.2, or later
Virtual machine migrated to another host by VMware DRS (Distributed Resource Scheduler)	DRS VM migrated	8.0.2, or later

Note: The near real-time update of virtual machines is supported on VMware vCenter Server 7.x.

For more information on setting-up near real-time (NRT) discovery, refer to the *Arctera InfoScale Operations Manager Installation and Configuration Guide.*

About discovering Solaris zones

InfoScale Operations Manager discovers the zones that are created on Solaris 11 host that it manages. The Zone Agentlet that is installed along with the `VRTSsfmh` package on a Solaris 11 host lets InfoScale Operations Manager Management Server discover the Global Zones and the non-Global Zones on the host.

Through the discovery of Solaris zones, 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 11 host.

Note: InfoScale Operations Manager discovers the Oracle databases on a non-Global Zone only if the non-Global Zone is in the Running state.

Note: 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 635.

How InfoScale Operations Manager discovers Solaris zones

In InfoScale Operations Manager, the zone agentlet that is part of the `VRTSsfmh` package facilitates the discovery of Solaris zones. Through the zone agentlet, InfoScale Operations Manager discovers the Global Zone and the associated non-Global Zones.

Note: Discovery of zones is performed only in the Global Zones.

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

Table 37-3 Solaris utilities to discover global and non-Global Zones *(continued)*

Utility	Purpose
`zonecfg`	Displays the details of each non-global configurations
`zlogin`	Logs on to a non-Global Zone from a Global Zone

See "About discovering Solaris zones" on page 645.

See "Information that InfoScale Operations Manager discovers on Solaris zones" on page 646.

See "Limitations of the discovery of Solaris zones in InfoScale Operations Manager" on page 647.

Information that InfoScale Operations Manager discovers on Solaris zones

Discovery of Solaris zones provides the following information:

Table 37-4 Solaris zones information that InfoScale Operations Manager discovers

Type of information	Components
Non-Global Zones , which are configured on Global Zones	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	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 InfoScale Operations Manager discovers *(continued)*

Type of information	Components
Oracle databases inside non-Global Zones	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**: 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 645.

See "How InfoScale Operations Manager discovers Solaris zones" on page 645.

See "Limitations of the discovery of Solaris zones in InfoScale Operations Manager" on page 647.

Limitations of the discovery of Solaris zones in InfoScale Operations Manager

The following limitations apply to the discovery of Solaris zones in InfoScale Operations Manager:

- InfoScale Operations Manager does not support the discovery of Solaris zones on Solaris managed hosts where both the zones and the LDoms are configured.
- InfoScale Operations Manager does not recommend the installation of the `VRTSsfmh` package in non-Global Zones.

  Note: Conversely, you can install InfoScale Operations Manager Management Server in a non-Global Zone and add the corresponding Global Zone as a managed host to Management Server.

- InfoScale Operations Manager does not discover the devices that are exported to non-Global Zones if the `zlogin` utility is not allowed on them.

- InfoScale Operations Manager discovers only the native and the Solaris-branded zones.

See "About discovering Solaris zones" on page 645.

See "How InfoScale Operations Manager discovers Solaris zones" on page 645.

See "Information that InfoScale Operations Manager discovers on Solaris zones" on page 646.

About discovering logical domains in 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.
I/O domain	Accesses the input or output devices directly. For example, the I/O domain can access a network device.

Table 37-5 LDom roles *(continued)*

LDom role	Description
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 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 635.

See "How InfoScale Operations Manager discovers Solaris logical domains" on page 649.

See "Information on logical domains that InfoScale Operations Manager discovers" on page 650.

See "Limitations of the discovery of logical domains in InfoScale Operations Manager" on page 650.

How InfoScale Operations Manager discovers Solaris logical domains

In 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, InfoScale Operations Manager discovers the LDom Server and the associated guest LDoms.

InfoScale Operations Manager uses the `ldm` command to discover the details of Solaris Logical domains.

The following are the major options that InfoScale Operations Manager uses with the `ldm` command to discover the details of Solaris Logical domains:

Table 37-6 Options that 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 InfoScale Operations Manager discovers

See "About discovering logical domains in InfoScale Operations Manager" on page 648.

See "Information on logical domains that InfoScale Operations Manager discovers" on page 650.

See "Limitations of the discovery of logical domains in InfoScale Operations Manager" on page 650.

Information on logical domains that InfoScale Operations Manager discovers

Discovery of Solaris LDoms provides the following information:

Table 37-7 LDom information that InfoScale Operations Manager discovers

Type of Information	Components
Guest LDoms that are configured on LDom Servers	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 InfoScale Operations Manager" on page 648.

See "How InfoScale Operations Manager discovers Solaris logical domains" on page 649.

See "Limitations of the discovery of logical domains in InfoScale Operations Manager" on page 650.

Limitations of the discovery of logical domains in InfoScale Operations Manager

The following limitations apply to the discovery of Solaris LDoms in InfoScale Operations Manager:

- InfoScale Operations Manager does not discover Solaris LDoms that are in the inactive state.

- InfoScale Operations Manager does not perform the discovery of Solaris LDoms along with the discovery of Solaris zones.

See "About discovering logical domains in InfoScale Operations Manager" on page 648.

See "How InfoScale Operations Manager discovers Solaris logical domains" on page 649.

See "Information on logical domains that InfoScale Operations Manager discovers" on page 650.

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 InfoScale Operations Manager" on page 648.

See "About the virtualization technologies supported" on page 635.

About discovering LPARs and VIOs in InfoScale Operations Manager

You can use 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 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 InfoScale Operations Manager console to add the HMC server to Management Server.

  To add the HMC server to 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: 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 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 InfoScale Operations Manager Management Server, refreshing the corresponding HMC discovery is recommended to view the end-to-end correlation immediately in the InfoScale Operations Manager console.

See "About the virtualization technologies supported" on page 635.

About LPAR storage correlation supported in InfoScale Operations Manager

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 InfoScale Operations Manager" on page 652.

See "About the virtualization technologies supported" on page 635.

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 InfoScale Operations Manager" on page 652.

See "About the virtualization technologies supported" on page 635.

About Microsoft Hyper-V virtualization discovery

Hyper-V is a hypervisor-based virtualization technology from Microsoft for x86-64 systems. You can use InfoScale Operations Manager to discover Hyper-V host and virtual machine-related information if the Hyper-V role is enabled on the managed host. 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 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 656.

See "Storage mapping discovery in Microsoft Hyper-V" on page 656.

Virtual machine discovery in Microsoft Hyper-V

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. 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 655.

Storage mapping discovery in Microsoft Hyper-V

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 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. 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 2019 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 655.

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 635.

About the Kernel-based Virtual Machine (KVM) virtualization discovery in 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). InfoScale Operations Manager discovers KVM virtual machines on the Linux host if the KVM modules are installed, and configured. InfoScale Operations Manager discovers basic information about only running virtual machines. For example, virtual machine name, CPU, and so on. 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 635.

About the reports related to virtualization

In InfoScale Operations Manager, you can generate the following reports that display the details of the virtualization environment that 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 635.

Chapter 38

Using Web services API

This chapter includes the following topics:

- About using InfoScale Operations Manager Web services API
- Logging in to InfoScale Operations Manager Web services API
- Logging out of InfoScale Operations Manager Web services API
- About objects supported by InfoScale Operations Manager Web services API
- About performing operations using InfoScale Operations Manager Web services API
- Examples of performing operations using InfoScale Operations Manager Web services API
- Examples of performing operations using XPRTLC and cURL
- Examples of the output in JSON format

About using InfoScale Operations Manager Web services API

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 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 InfoScale Operations Manager 8.0.2 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 InfoScale Operations Manager. These extended attributes can be managed using the Web services API provided by 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 InfoScale Operations Manager Web services API" on page 661.

See "About objects supported by InfoScale Operations Manager Web services API" on page 663.

See "About performing operations using InfoScale Operations Manager Web services API" on page 666.

Logging in to InfoScale Operations Manager Web services API

The 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 -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 InfoScale Operations Manager Web services API" on page 660.

See "Logging out of InfoScale Operations Manager Web services API" on page 662.

Logging out of 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 InfoScale Operations Manager Web services API" on page 661.

See "About using InfoScale Operations Manager Web services API" on page 660.

About objects supported by InfoScale Operations Manager Web services API

The 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
Virtual Business Services	vbs
Database	db
Applications database file	db/file
Database - Volume Manager	dbvm

Table 38-1 Object name used in the URL *(continued)*

Object	Object name used in the URL
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
Virtualization server	vserver
Virtual Machine	vm
Storage Mapping	storagemapping
Storage Pool	storagepool
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

Table 38-1 Object name used in the URL *(continued)*

Object	Object name used in the URL
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
Data Center	datacenter

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, product, Data Center and Cluster.	Server	server
Cluster, service group, resource, system, Virtual Business Services -dependency, Virtual Business Services, and recovery plan tasks.	Availability	hadr
Virtualization server, virtual machine, storage mapping, disk, storage pool, cluster, and system HBA.	Virtualization	virt

See "About using InfoScale Operations Manager Web services API" on page 660.

See "Logging in to InfoScale Operations Manager Web services API" on page 661.

About performing operations using 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 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 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 8.0.2
- 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 358.
- Move managed host in the Data Center.

You can move managed hosts in the Data Center. This operation can be performed in the server perspective.

Note: Managed host can be moved only in the Data Center.

Note: If the rule is set on the parent Organization, then the managed host will move to the defined Organization.

- Move managed hosts in the Organization/child-Organization.
  Provide the parent Organization name first, followed by the subsequent child-Organizations names (s).
- Move cluster in the Data Center.

Note: The cluster can be moved only in the Data Center.

Note: If the rule is set on the parent Organization, then the cluster will move to defined Organization.

- Move cluster in the Organization/child-Organization.
  Provide the parent Organization name first, followed by the subsequent child-Organization name(s).
- Bring service group online with the parameters as force or propogate.
- Switch service group.
- Offline service group with the parameters as force, probe, or propogate.
- Freeze service group with the parameter as persistent.
- Unfreeze service group.
- Clear fault on service group.
- Clear the **clearadminwait** state for service group.
- Flush service group.
- Remove or unconfigure a managed host

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 InfoScale Operations Manager Web services API" on page 660.

See "Logging in to InfoScale Operations Manager Web services API" on page 661.

Examples of performing operations using InfoScale Operations Manager Web services API

Following are some examples on performing operations using Web services API:

- To get metadata for a host:

  ```
  https://domain.example.com:14161/vom/api/meta/server/host
  ```

- To define an extended attribute named **department** on object type cluster:

  ```
  https://domain.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://domain.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://domain.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://domain.example.com:14161/vom/api/update/server/host
  /myhost_id?location=1st floor
  ```

- To filter disk group objects whose display type is private:

  ```
  https://domain.example.com:14161/vom/api/query/server/host
  /myhost_id/diskgroup?display_type=private
  ```

- To run a recovery plan

  ```
  https://domain.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://domain.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.

- To online service group with the parameters as **force** or **propogate**.

  ```
  https://domain.example.com:14161/vom/api/op/hadr/servicegroup/
  service_group_id/host/{hostname}/{force}/{propogate}/online
  ```

Example: Service group online with the parameter as **force**

```
https://domain.example.com:14161/vom/api/op/hadr/servicegroup/
ABC_12/host/example.com/true/false/online
```

Example: Service group online with the parameter as **propogate**

```
https://domain.example.com:14161/vom/api/op/hadr/servicegroup/
ABC_12/host/example.com/false/true/online
```

- To offline service group with the parameters as **force**, **probe**, or **propogate**.

```
https://domain.example.com:14161/vom/api/op/hadr/servicegroup/
service_group_id/host/{hostname}/{force}/{probe}/{propogate}/offline
```

Example: Service group offline with the parameter as **force**

```
https://domain.example.com:14161/vom/api/op/hadr/servicegroup/
x1@linux_clus-00000/host/example.com/true/false/false/offline
```

Example: Service group offline with the parameter as **propogate**

```
https://domain.example.com:14161/vom/api/op/hadr/servicegroup/
x1@linux_clus-00000/host/example.com/false/false/true/offline
```

Example: Service group offline with the parameters as **force** and **probe**

```
https://domain.example.com:14161/vom/api/op/hadr/servicegroup/
x1@linux_clus-00000/host/example.com/true/true/false/offline
```

- To Switch service group.

```
https://domain.example.com:14161/vom/api/op/hadr/servicegroup/
service_group_id/host/{hostname}/switch/
```

Example

```
https://domain.example.com:14161/vom/api/op/hadr/servicegroup/
x1@linux_clus-00000/host/example.com/switch/
```

- To move managed host in the Data Center.

```
https://domain.example.com:14161/vom/api/op/server/
{host_id}/moveto/datacenter/
```

Example

```
https://domain.example.com:14161/vom/api/op/server/
34001234-1234-4dxx-0000-00000000000/moveto/datacenter/
```

- To move managed host from Organization OE1 to child-Organization OE3 of Organization OE2.

  ```
  https://domain.example.com:14161/vom/api/
  op/server/{host_id}/moveto_oe/OE2,OE3/
  ```

  Example

  ```
  https://domain.example.com:14161/vom/api/op/server/
  0001234-1234-4dzz-0000-00000000011/moveto_oe/OE2,OE3/
  ```

- To move cluster in the Data Center.

  ```
  https://domain.example.com:14161/vom/api/
  op/hadr/{cluster_id}/moveto/datacenter/
  ```

  Note: Cluster ID is the encoded_id.

  Example

  ```
  https://domain.example.com:14161/vom/api/op/
  hadr/rhelclust-10098/moveto/datacenter/
  ```

- To move cluster from an organization OE1 to child-Organization OE3 of Organization OE2.

  ```
  https://domain.example.com:14161/vom/api/
  op/hadr/{cluster_id}/moveto_oe/OE2,OE3/
  ```

  Example

  ```
  https://domain.example.com:14161/vom/api/op/
  hadr/rhelclust-10098/moveto_oe/OE2,OE3/
  ```

- To freeze service group with the parameter as **persistent**.

  ```
  https://domain.example.com:14161/vom/api/
  op/hadr/servicegroup/{service_group_id}/{persistent}/freeze/
  ```

  Example: Freeze service group with the parameter as persistent

```
https://domain.example.com:14161/vom/api/
op/hadr/servicegroup/x1@linux_clus-00000/true/freeze/
```

Example: Freeze a service group without any parameter

```
https://domain.example.com:14161/vom/api/
op/hadr/servicegroup/x1@linux_clus-00000/false/freeze/
```

- To unfreeze service group.

```
https://domain.example.com:14161/vom/api/
op/hadr/servicegroup/{service_group_id}/unfreeze/
```

Example

```
https://domain.example.com:14161/vom/api/
op/hadr/servicegroup/x1@linux_clus-00000/unfreeze/
```

- To clear fault on service group.

```
https://domain.example.com:14161/vom/api/
op/hadr/servicegroup/{service_group_id}/host/{hostname}/clear/
```

Example

```
https://domain.example.com:14161/vom/api/op/hadr/servicegroup/
x1@linux_clus-00000/host/example.com/clear
```

- To clear the **clearadminwait** state for service group.

```
https://domain.example.com:14161/vom/api/
op/hadr/servicegroup/{service_group_id}/host/{hostname}/{fault}/clearadmin
```

Example: Clear the clearadminwait state and the fault.

```
https://domain.example.com:14161/vom/api/op/hadr/servicegroup/
x1@linux_clus-00000/host/example.com/true/clearadminwait/
```

Example: Clear the clearadminwait state without clearing the fault.

```
https://domain.example.com:14161/vom/api/op/hadr/servicegroup/
x1@linux_clus-00000/host/example.com/false/clearadminwait/
```

- To flush service group.

```
https://domain.example.com:14161/vom/api/
op/hadr/servicegroup/{service_group_id}/host/{hostname}/flush/
```

Example

```
https://domain.example.com:14161/vom/api/op/hadr/servicegroup/
ABC_12/host/example.com/flush/
```

- To remove or unconfigure a managed host
  Enter the following URL in a browser:

  ```
  https://managementServerID:14161/vom/api/op/infra/host/hostID/remove
  ```

  You can identify the value of `hostID` for the managed host that you want to remove or unconfigure by using the `host` API that is described earlier.

Note: If "Ask reason for all operations" is enabled under **Advanced Authorization** in the **Management Server** perspective, then it is mandatory to provide a reason for performing the operation.

See "About performing operations using InfoScale Operations Manager Web services API" on page 666.

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://domain.example.com:14161/vom/api/meta/server/host`

- `curl -g -k -X POST -b "session id" https://domain.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://domain.example.com:14161/vom/api/meta/hadr/cluster/add?name=department`

- `curl -g -k -X POST -b "session id" https://domain.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://domain.example.com:14161/vom/api/meta/hadr/cluster/modify?name=department&new_name=Dept_new`

- `curl -g -k -X POST -b "session id" https://domain.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://domain.example.com:14161/vom/api/meta/hadr/cluster/delete?name=dept_new`

- `curl -g -k -X POST -b "session id" https://domain.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://domain.example.com:14161/vom/api/update/server/host/`**`myhost_id`**`?location=1st floor`

- `curl -g -k -X POST -b "session id" https://domain.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://domain.example.com:14161/vom/api/query/server/host/`**`myhost_id`**`/diskgroup?display_type=private`

- curl -g -k -b "*session id*"
  https://domain.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://domain.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://domain.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://domain.example.com:14161/vom/api/op/server/
   template/storage/**storage_template_ID**/provision'

- curl -g -k -X POST -b "*session id*" -l
  'https://domain.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**"}'

- To move managed host in the Data Center.

  curl -g -k -X POST -b "*session id*"
  https://domain.example.com:14161/vom/api/
  op/server/{**host_id**}/moveto/datacenter/

Example

```
curl -g -k -X POST -b "JSESSIONID=13C8F4E4203BFE96DD6D1LFC69FC7B65;"
https://domain.example.com:14161/vom/api/
op/server/%7B0002345-5432-4dff-0000-00000000000%7D/moveto/datacenter/
```

- To move managed host from an Organization OE1 to child Organization OE3 of Organization OE2.

```
curl -g -k -X POST -b "session id"
https://domain.example.com:14161/vom/
api/op/server/{host_id}/moveto_oe/OE2,OE3/
```

Example

```
curl -g -k -X POST -b "JSESSIONID=13C8F4E4203BFE96DD6D1LFC69FC7B65;"
https://domain.example.com:14161/vom/api/
op/server/%7B0002345-5432-4dff-0000-00000000000%7D/moveto_oe/OE2,OE3/
```

- To move cluster in the Data Center.

```
curl -g -k -X POST -b "session id"
https://domain.example.com:14161/vom/
api/op/hadr/{cluster_id}/moveto/datacenter/
```

Example

```
curl -g -k -X POST -b "JSESSIONID=13C8F4E4203BFE96DD6D1LFC69FC7B65;"
https://domain.example.com:14161/vom/api/op/
hadr/rhelclust-42058/moveto/datacenter/
```

- To move cluster from an Organization OE1 to child Organization OE3 of Organization OE2.

```
curl -g -k -X POST -b "session id"
https://domain.example.com:14161/vom/
api/op/hadr/{cluster_id}/moveto_oe/OE2,OE3/
```

Example

```
curl -g -k -X POST -b "JSESSIONID=13C8F4E4203BFE96DD6D1LFC69FC7B65;"
https://domain.example.com:14161/vom/api/op/
hadr/rhelclust-42058/moveto_oe/OE2,OE3/
```

- To online service group with the parameters **force** or **propogate**.

```
curl -g -k -X POST -b "session id"
https://domain.example.com:14161/vom/api/op/
hadr/servicegroup/{service_group_id}/host/{hostname}/{force}/{propogate}/o
```

Example: Service group online with parameter force

```
curl -g -k -X POST -b "JSESSIONID=13C8F4E4203BFE96DD6D1LFC69FC7B65;"
https://domain.example.com:14161/vom/api/op/hadr/servicegroup/
x1@linux_clus-00000/host/example.com/true/false/online
```

Example: Service group online with parameter propogate

```
curl -g -k -X POST -b "JSESSIONID=13C8F4E4203BFE96DD6D1LFC69FC7B65;"
https://domain.example.com:14161/vom/api/op/hadr/servicegroup/
x1@linux_clus-00000/host/example.com/false/true/online
```

- To offline service group with the parameters **force**, **propogate** or **probe**.

```
curl -g -k -X POST -b "session id"
https://domain.example.com:14161/vom/api/op/
hadr/servicegroup/{service_group_id}/host/{hostname}/{force}/{probe}/{prop
```

Example: Service group offline with parameter force

```
curl -g -k -X POST -b "JSESSIONID=13C8F4E4203BFE96DD6D1LFC69FC7B65;"
https://domain.example.com:14161/vom/api/op/hadr/servicegroup/
x1@linux_clus-00000/host/example.com/true/false/false/offline
```

Example: Service group offline operation with the parameters force and probe

```
curl -g -k -X POST -b "JSESSIONID=13C8F4E4203BFE96DD6D1LFC69FC7B65;" https
x1@linux_clus-00000/host/example.com/true/true/false/offline
```

- To switch service group.

```
curl -g -k -X POST -b "session id"
https://domain.example.com:14161/vom/api/op/
hadr/servicegroup/{service_group_id}/host/{hostname}/switch/
```

Example

```
curl -g -k -X POST -b "JSESSIONID=13C8F4E4203BFE96DD6D1LFC69FC7B65;"
https://domain.example.com:14161/vom/api/op/hadr/servicegroup/
x1@linux_clus-00000/host/example.com/switch/
```

- To freeze service group with parameter **persistent**.

```
curl -g -k -X POST -b "session id"
https://domain.example.com:14161/vom/api/op/
hadr/servicegroup/{service_group_id}/{persistent}/freeze/
```

Example

```
curl -g -k -X POST -b "JSESSIONID=13C8F4E4203BFE96DD6D1LFC69FC7B65;"
https://domain.example.com:14161/vom/api/
op/hadr/servicegroup/x1@linux_clus-00000/true/freeze/
```

Example: Freeze service group without any parameter

```
curl -g -k -X POST -b "JSESSIONID=13C8F4E4203BFE96DD6D1LFC69FC7B65;"
https://domain.example.com:14161/vom/api/
op/hadr/servicegroup/x1@linux_clus-00000/false/freeze/
```

- To unfreeze service group.

```
curl -g -k -X POST -b "session id"
https://domain.example.com:14161/vom/api/op/
hadr/servicegroup/{service_group_id}/unfreeze/
```

Example

```
curl -g -k -X POST -b "JSESSIONID=13C8F4E4203BFE96DD6D1LFC69FC7B65;"
https://domain.example.com:14161/vom/api/
op/hadr/servicegroup/x1@linux_clus-00000/unfreeze/
```

- To clear fault on service group.

```
curl -g -k -X POST -b "session id"
https://domain.example.com:14161/vom/api/op/
hadr/servicegroup/{service_group_id}/host/{hostname}/clear/
```

Example

```
curl -g -k -X POST -b "JSESSIONID=13C8F4E4203BFE96DD6D1LFC69FC7B65;"
https://domain.example.com:14161/vom/api/op/hadr/servicegroup/
x1@linux_clus-00000/host/example.com/clear
```

- To clear the **clearadminwait** state for service group.

```
curl -g -k -X POST -b "session id"
https://domain.example.com:14161/vom/api/op/
hadr/servicegroup/{service_group_id}/host/{hostname}/{fault}/clearadminwai
```

Example: Clear the clearadminwait state and the fault

```
curl -g -k -X POST -b "JSESSIONID=13C8F4E4203BFE96DD6D1LFC69FC7B65;"
https://domain.example.com:14161/vom/api/op/hadr/servicegroup/
x1@linux_clus-00000/host/example.com/true/clearadminwait
```

Example: Clear the clearadminwait state without clearing the fault

```
curl -g -k -X POST -b "JSESSIONID=13C8F4E4203BFE96DD6D1LFC69FC7B65;"
https://domain.example.com:14161/vom/api/op/hadr/servicegroup/
x1@linux_clus-00000/host/example.com/false/clearadminwait
```

- To flush service group.

```
curl -g -k -X POST -b "session id"
https://domain.example.com:14161/vom/api/op/
hadr/servicegroup/{service_group_id}/host/{hostname}/flush/
```

Example

```
curl -g -k -X POST -b "JSESSIONID=13C8F4E4203BFE96DD6D1LFC69FC7B65;"
https://domain.example.com:14161/vom/api/op/hadr/servicegroup/
x1@linux_clus-00000/host/example.com/flush
```

- To fetch the details of a managed host
  1. Before you can use the `query` API to fetch the details of a managed host, identify the session ID for the secure API login by using the following command:

```
curl -g -k -d user="userName" -d password="userPassword" -d domain="manage
https://managementServerID:14161/vom/api/login
```

Specify the appropriate values of the username and password of the CMS.
A sample session ID value is as follows:

```
{"cookie":"JSESSIONID=2886C4EE175075D72C0AB91A1D9CF846;","current_server_t
26 10:57:36 GMT-12:00 2020","max_active_interval":30,"expires_at":"Mon Oct
GMT-12:00 2020"}
```

2. Fetch the host details by using the following command:

```
curl -g -k -X POST -b "sessionID"
https://managementServerID.com:14161/vom/api/query/server/host/
```

- To remove or unconfigure a managed host
  Run the following command to use the API to remove or unconfigure a managed host:

```
curl -g -k -X POST -b "sessionID"
https://managementServerID.com:14161/vom/api/op/infra/host/%7bhostID%7d/re
```

Note: The curl command does not support the use of curly brackets, { and }. Replace { with %7b and } with %7d as the corresponding markup text. To be able to use this API with Google Chrome, install the RestMan addon. To be able to use this API with Mozilla FireFox, install the REStClient addon.

Where, **domain.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, **cluster_id** is the encoded ID of the cluster, 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 InfoScale Operations Manager Web services API" on page 666.

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://domain.example.com/vom/api/query
                         /server/host/{host-id}/volume"
                       ],
           }
           {"name":"disks",
            "template":[
                        "https://domain.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://domain.example.com/vom/api/query/server
                 /host/{00120050-56a6-15f3-0000-0000f684d3ab}/diskgroup",
    "hbas":"https://domain.example.com/vom/api/query/server/host/
           {00120050-56a6-15f3-0000-0000f684d3ab}/hba",
    "volumes":"https://domain.example.com/vom/api/query/server/host/
              {00120050-56a6-15f3-0000-0000f684d3ab}/volume",
```

```
    "disks":"https://domain.example.com/vom/api/query/server/host/
            {00120050-56a6-15f3-0000-0000f684d3ab}/disk"
  },
    "name":"10.255.255.255",
    "url":"https://domain.example.com/vom/api/query/server/host/
          {00120050-56a6-15f3-0000-0000f684d3ab}",
  .
  .
  }
]
```

See "About performing operations using InfoScale Operations Manager Web services API" on page 666.

Chapter 39

InfoScale Operations Manager command line interface

This chapter includes the following topics:

- About 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
host-mgmt	HostManagement	Host management in InfoScale Operations Manager
hotfix	deployment	Deployment of hotfix in InfoScale Operations Manager

Option	Category	Operation
makeBE	BEManagement	Business Application management in InfoScale Operations Manager
service	Service-Management	Service management in InfoScale Operations Manager
domain-mgmt	Domain-Management	Domain management in 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 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>`

Host management using the vomadm utility

Use the following commands to manage the configured hosts in InfoScale Operations Manager.

`vomadm host-mgmt [--remove --host <hostname> | --remove --hostfile <file> | --list ]`

Where,

`list`: It lists all hosts configured as agents to InfoScale Operations Manager Management Server.

`remove`: It removes specified host(s) from 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 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 684.

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 684.

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 579.

See "Creating Business Application using the makeBE script" on page 583.

See "Updating Business Application using the makeBE script" on page 586.

See "About the vomadm utility" on page 684.

Service management using the vomadm utility

Use the following command to manage the various services in 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` - InfoScale Operations Manager Web Server
- `at` - InfoScale Operations Manager Authentication Service
- `xprtld` - InfoScale Operations Manager Messaging Service
- `db` - InfoScale Operations Manager Database Service
- `dcli` - InfoScale Operations Manager Distributed Command Line Daemon
- `sd` - InfoScale Operations Manager watchdog
- `xtrapd` - SNMP trap service
- `ALL` - All the processes that the script manages

See "About the vomadm utility" on page 684.

Domain management using the vomadm utility

Use the following commands to manage the configured hosts in 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 684.

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 684.

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 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 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 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 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 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 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 InfoScale Operations Manager services.

domain-mgmt

Used to manage 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 toInfoScale 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 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 ©) 2025 Arctera US LLC. 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**, **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 InfoScale Operations Manager Management Server environment and Oracle database are properly configured. For the supported versions of Oracle database, refer to the *InfoScale Operations Manager Hardware and Software Compatibility List (HSCL).*

Oracle database discovery prerequisites

For 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, 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

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 InfoScale Operations Manager Management Server domain, the database family of 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, 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 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 706.

See "Oracle database discovery in Solaris zones" on page 707.

SQL queries used for the discovery of Oracle database

The 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 703.

Oracle database discovery in Solaris zones

InfoScale Operations Manager supports Oracle database discovery in Solaris zones. 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 703.

See "About the virtualization technologies supported" on page 635.

Application setup requirements for Oracle Automatic Storage Management (ASM) discovery

For Oracle Automatic Storage Management (ASM) discovery, ensure that InfoScale Operations Manager Management Server environment and Oracle ASM are properly configured. For the supported versions of Oracle ASM, refer to the *InfoScale Operations Manager Hardware and Software Compatibility List (HSCL).*

Oracle ASM discovery prerequisites

For 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, InfoScale Operations Manager fails to discover Oracle ASM details.

Oracle ASM discovery

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 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 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 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 708.

See "Oracle Automatic Storage Management (ASM) discovery in Solaris zones" on page 709.

SQL queries used for the discovery of Oracle ASM

The 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 707.

See "Oracle Automatic Storage Management (ASM) discovery in Solaris zones" on page 709.

Oracle Automatic Storage Management (ASM) discovery in Solaris zones

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 707.

See "SQL queries used for the discovery of Oracle ASM" on page 708.

Application setup requirements for IBM DB2 discovery

For IBM DB2 (Extended Edition and Enterprise-Extended Edition) discovery, ensure that InfoScale Operations Manager Management Server environment and IBM DB2 are properly configured. For the supported versions of IBM DB2, refer to the *InfoScale Operations Manager Hardware and Software Compatibility List (HSCL).*

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 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 710.

SQL queries used for the discovery of IBM DB2 database

The 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 709.

Application setup requirements for Sybase Adaptive Server Enterprise (ASE) discovery

For Sybase Adaptive Server Enterprise (ASE) discovery, ensure that InfoScale Operations Manager Management Server environment and Sybase database are properly configured. For the supported versions of Sybase database, refer to the *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

InfoScale Operations Manager does not provide automatic discovery of Sybase ASE. To enable Sybase ASE discovery, you must configure InfoScale Operations Manager to discover Sybase instances.

To enable Sybase ASE discovery in 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 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 712.

SQL queries used for the discovery of Sybase Adaptive Server Enterprise (ASE)

The 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 711.

Application setup requirements for Microsoft SQL Server discovery

For Microsoft SQL Server discovery, ensure that InfoScale Operations Manager Management Server environment and Microsoft SQL Server are properly configured. For the supported versions of Microsoft SQL Server, refer to the *InfoScale Operations Manager Hardware and Software Compatibility List (HSCL)*.

Microsoft SQL Server discovery

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 InfoScale Operations Manager Management Server, the database (DB) family of 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 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 714.

SQL queries used for the discovery of Microsoft SQL Server

The 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 713.

Application setup requirements for Microsoft Exchange Server discovery

For Microsoft Exchange Server discovery, ensure that InfoScale Operations Manager Management Server environment and Microsoft Exchange Server are properly configured. For the supported versions of Microsoft Exchange Server, refer to the *InfoScale Operations Manager Hardware and Software Compatibility List (HSCL)*.

InfoScale Operations Manager automatically discovers the Microsoft Exchange Server details and no additional configuration is required.

See "About Business Applications in InfoScale Operations Manager" on page 572.

See "About InfoScale Operations Manager" on page 27.

Glossary

Active/active configuration	A failover configuration where each system runs a service group. If either system fails, the other one takes over and runs both service groups. Also symmetric configuration.
Active/passive configuration	A failover configuration consisting of one service group on a primary system, and one dedicated backup system. Also asymmetric configuration.
addressable unit	Any storage resource in the network that is ready to be allocated for use by hosts and applications. Also AddrUnit or AU. See also LUN
allocated storage	The total amount of addressable storage in LUNs that is designated for use by specific hosts. A LUN is considered allocated when a host operating system has written a device handle for the LUN (in other words, claimed the LUN) or when the array has masked the LUN to a specific target. Contrast with unallocated storage
application	A program or group of programs designed to perform a specific task. Oracle Database and Arctera NetBackup are examples of applications.
Authentication Service	See Symantec Product Authentication Service.
bridge	A device that connects and passes packets between two segments of a storage network that use the same communications protocol. See also router
capacity	The amount of storage an object can allocate or use.
claimed storage	Storage for which at least one host's operating system has created a device handle. Contrast with unclaimed storage
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 Cluster Server (VCS): GAB and LLT. The communication takes place by means of heartbeat signals sent between systems or fast kernel-to-kernel broadcasts.
configured storage	Physical storage that has been formatted and is ready to be apportioned into RAID groups.

	Contrast with unconfigured storage
device handle	The name the operating system uses to identify a storage resource (known as an addressable unit or LUN), and the correct means (driver, system call) to access it. Also OS handle.
disk group	A collection of disks that share a common configuration. A disk group configuration is a set of records containing detailed information on existing Veritas Volume Manager objects (such as disk and volume attributes) and their relationships. Each disk group has an administrator-assigned name and an internally defined unique ID. The root disk group (rootdg) is a special private disk group that always exists.
DMP (Dynamic Multipathing)	A feature of Volume Manager (VxVM) 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.
event	A notification that indicates when an action, such as an alert or a change in state, has occurred for one or more objects on the storage network.
failover	A backup operation that automatically switches to a standby database, server, or network if the primary system fails or is temporarily shut down for servicing.
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.
firmware	A set of software instructions set permanently in a device's memory.
GBIC	Gigabit interface converter. A widely used transceiver module for Fibre Channel. A GBIC is modular and hot-swappable and can be either copper or optical.
Global Service Group	A VCS service group that spans across two or more clusters. The ClusterList attribute for this group contains the list of clusters over which the group spans.
Group Atomic Broadcast (GAB)	A communication mechanism of the VCS engine that manages cluster membership, monitors heartbeat communication, and distributes information throughout the cluster.
hub	A common connection point for devices in the storage network. The hub may be unmanaged, IP-managed, or FC-managed. An unmanaged hub is passive in the sense that it serves simply as a conduit for data, moving the data from one storage resource to another. IP-managed and FC-managed hubs are intelligent, containing features an administrator can use to monitor the traffic passing through the hub and configure each port in the hub.

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 Cluster Server (VCS), node refers specifically to one of any number of hosts in a cluster. See also object.
object	A single, unique addressable entity on a storage network. It is possible for objects to be present within objects. For example, while a tape array is an object, each individual tape drive within the array is also an object. A host is an object, and the HBA inside the host is also an object. Each object has one or more attributes and can be a member of one or more zones.
Object Reference or OID (Object ID)	A key which uniquely identifies an object in the discovery data store. OIDs are represented in XML files as hexadecimal strings with a maximum length of 128 characters.
physical fabric	The physical components of a fabric, including all switches and all other SAN objects. You can configure one or more virtual fabrics—each one isolated from the others—based on the hardware components in the physical fabric.
policy	A set of rules, or configuration settings, that are applied across a number of objects in the storage network. You establish policies to help you monitor and manage the network. Each policy associates certain sets of conditions with storage resources and defines actions to be taken when these conditions are detected.
RAID	Redundant Array of Independent Disks. A set of techniques for managing multiple disks for cost, data availability, and performance. See also mirroringstriping
resource	Any of the individual components that work together to provide services on a network. A resource may be a physical component such as a storage array or a switch, a software component such as Oracle8i or a Web server, or a configuration component such as an IP address or mounted file system.
SAN	Acronym for "storage area network." A network linking servers or workstations to devices, typically over Fibre Channel, a versatile, high-speed transport. The storage area network (SAN) model places storage on its own dedicated network, removing data storage from both the server-to-disk SCSI bus and the main user network. The SAN includes one or more hosts that provide a point of interface with LAN users, as well as (in the case of large SANs) one or more fabric switches and SAN hubs to accommodate a large number of storage devices.
SCSI	Small Computer Systems Interface. A hardware interface that allows for the connection of multiple peripheral devices to a single expansion board that plugs

into the computer. The interface is widely used to connect personal computers to peripheral devices such as disk and media drives.

seeding A technique used to protect a cluster from a preexisting network partition. By default, when a system comes up, it is not seeded. Systems can be seeded automatically or manually. Only systems that have been seeded can run VCS. Systems are seeded automatically only when an unseeded system communicates with a seeded system or when all systems in the cluster are unseeded and able to communicate with each other.

See network partition

service group A collection of resources working together to provide application services to clients. It typically includes multiple resources, hardware- and software-based, working together to provide a single service.

SMTP Simple Mail Transfer Protocol, a commonly used protocol for sending email messages between servers.

SnapMirror A method of mirroring volumes and qtrees on NetApp unified storage devices. With SnapMirror, a user can schedule or initiate data transfers, request information about transfers, update a mirror, and manage mirrors.

See mirroring

snapshot A point-in-time image of a volume or file system that can be used as a backup.

SNMP The Simple Network Management Protocol for Internet network management and communications used to promote interoperability. SNMP depends on cooperating systems that must adhere to a common framework and a common language or protocol.

striping A layout technique that spreads data across several physical disks by mapping the data to successive media, known as stripes, in a cyclic pattern. Also RAID Level 0.

switch A network device to which nodes attach and which provides high-speed switching of node connections via link-level addressing.

system The physical hardware on which data and applications reside, and the connections between them.

topology The physical or logical arrangement of resources on the storage network and the connections between them.

unused storage Storage to which data has not been written.

Contrast with used storage

virtual IP address A unique IP address associated with a VCS cluster. This address can be used on any system in the cluster, along with other resources in the VCS cluster service

group. A virtual IP address is different from a system's base IP address, which corresponds to the system's host name.

See also IP address

virtualization

Representing one or more objects, services, or functions as a single abstract entity so that they can be managed or acted on collectively. An example of virtualization is the creation of a virtual fabric from a switch and associated storage resources as a means of controlling access and increasing scalability in the storage network.

Index

A

adding
- disks to a disk group 166
- mirrors to volumes 219
- resources 436
- service groups 385

alerts and rules 104
App VG
- modify 326

application configuration
- about 477
- launching 478
- prerequisites 478

arrays
- disabling DMP paths 616

assigning price tier
- automatically 508
- manually 508
- processor price tier 507
- server price tier 507

Atleast Count dependency
- child resources 393
- create 392

attributes
- editing 412

Availability
- edit attributes 375

B

break-off mirror volume snapshot
- configure 235

Business Applications
- creating 572
- deleting 574
- modifying 572
- overview 572
- renaming 574

C

change log
- synchronizing 280

checking file system 292
checkpoint
- creating 293

clearing faults on resources 427
cluster hosts
- freezing 419
- managing 419
- starting Cluster Server HAD 421
- stopping Cluster Server HAD 422
- unfreezing 420

Cluster Server HAD
- starting 421
- stopping 422

clustered file systems
- mounting 288
- unmounting 287

clusters
- close configuration 379
- editing attributes 379
- FSS 628
- managing 377
- open configuration 378
- save configuration 378

compression 360
- adding directories 362
- scheduling 362
- starting on demand 364

configure
- fire drill service groups 452
- instant volume snapshot 229

creating
- checkpoint 293
- file systems 260
- recovery plan 464
- space-optimized volume snapshot 233
- volumes 200

custom signature
- copying to the managed hosts 605

custom signature *(continued)*
creating script 602
registering on Management Server 605
removing 606
sample script 603
using Distribution Manager Add-on 606

D

datastore 640
deduplication
configuring 368
disabling 371
overview 366
removing 371
starting 371
deep array discovery
ESX servers 637
defragmenting file system 281
deleting
fire drill schedules 459
resources 426
service groups 398
dependency views 414
disabling
all resources in a service group 398
fire drill schedules 460
service groups 391
discovering VMware Infrastructure 637
disk group
adding disks 166
creating 156
deporting 164
destroying 164
importing 165
joining 178
moving 176
recovering 163
removing disks 186
renaming disks 171
resizing disks 168
splitting 174
upgrading 173
disks
disconnecting 185
evacuating 192
exporting for FSS 632
host prefix 187
initializing 180
mapping 184
disks *(continued)*
offlinig 188
online 188
recover disks 183
recovering 183
replace disks 181
rescanning 196
thin reclamation 359
unmapping 185
usage 189
DMP maintenance case
removing a record 621
view results and output 622
DMP paths
disabling on a host 610
disabling on a virtualization server 617
disabling on an array 616
Dynamic Multi-Pathing 609

E

editing attributes 375
host 421
resource 433
resource type 434
enabling
all resources in a service group 397
fire drill schedules 459
resources 425–426
service groups 390
ESX servers
datastore 640
deep array discovery 637
extended attribute
modifying value 592
searching and setting values 590
setting values 589
using 588

F

faults
restoring 119
suppressing 117
faults and risks 117
file change log
disabling 279
enabling 278
removing 280

file compression 360
adding directories 362
scheduling 362
starting on demand 364
file system snapshot
mount 298
remounting 296
removing 300
unmount 299
file systems
checking 292
creating 260
deduplication 366, 368, 371
defragmenting 281
remounting 290
thin reclamation 359
files
pinning to cache 317
unpinning from cache 317
fire drill
configure fire drill service groups 452
deleting schedules 459
disabling schedules 460
enabling schedules 459
modifying schedule 456
running DR fire drill 453
running HA fire drill 451
viewing schedules 461
FSS
enabling or disabling on existing disk group 633
exporting disks 632
functionality 628
use cases 629

G

global clusters 440
adding remote clusters 443
creating 442
objects 440
prerequisites for creating 442
removing remote clusters 447–448
terminology 441
global service groups
converting from local 443
converting to local 446

H

HA-DR
pre-requisites 375
Hardware Management Console 652
host
disabling DMP paths 610
editing attributes 421
Hyper-V virtualization discovery 655

I

I/O Trace Log
panel options 319
I/O trace log
analyzing 320
viewing 320
importing
disk group 165
type definition 381
InfoScale Operations Manager
about 27
instant volume snapshot
configure 229
IO Thresholds
setting 325

L

LDom discovery
limitations 650
method 649
LDoms
discovery 648
discovery information 650
roles 648
license deployment policy
creating 510
deleting 514
modifying 513
licenses
about 497
accountable deployed license 524
child licenses 525
deployed licenses 497
deployment details 525
deployment policy details 526
discovery interval 497
features enabled 525
host deployment summary 522
license deployment summary 522

licenses *(continued)*
product deployment summary 522
SPVU details 525
licensing and pricing
OS tier 498
per-core 498, 506
processor tier 498
server tier 498
SPVU 498, 501
IBM LPAR 506
kernel-based virtual machines 505
Solaris LDOM virtualization server 504
VMware virtual machine 502
linking
resources 435
service groups 399, 402
LPAR
about discovering 652

M

makeBE script
about 579
creating business application 583
CSV file 580
deleting business application 586
export 582
exporting business application 585
import 582
importing business application 584
limitations 582
log files 583
objects 580
updating business application 586
user defined import 582
Management Server console
about 33
home page 34
user interface element 37
modifying
fire drill schedule 456
resources 436
system list 413
monitoring capacity 302
mounting
clustered file systems 288
file system snapshot 298
non clustered file systems 283
Multi Site Management
about 480
Multi Site Management *(continued)*
features 481
limitations 482
prerequisites 482
setting-up campus cluster 482
setting-up Replicated Data Cluster 484
multi-pathing discovery
user privileges 641
VMware environment 641

N

non clustered file systems
mounting 283
unmounting 282

O

object
modifying permissions 99
offline and propagate
resources 429
online help
about 51
Organization
assigning permissions 97
creating 91
deleting 96–97
deleting permissions 99
modifying name 95
modifying permissions 98

P

performance graphs
about 539
disk 542
enclosure 550
file system 544
host 541
initiator 547
path 546
virtual machine 548
VMware ESX server 548
VMware ESX Server path 549
volume 544
performance metering
about 530
disable for host 537
disable for virtualization server 538
enable for host 536

performance metering *(continued)*
enable for virtualization server 537
metered resources 530
permissions
assigning on Organization 97
deleting on Organization 99
modifying on object 99
modifying on Organization 98
policy checks
about 593
bundling custom signatures using Distribution Manager Add-on 606
enabling or disabling signatures 598
exporting signatures 601
registering signatures 594–595
running manual scans 598
unregistering signatures 596
viewing policy violation details 599
policy violations
viewing scan details 599

R

reactivating volume 215
reattaching
snapshot 239
recover disks 183
recovering disk group 163
recovering volume 214
recovery plan
creating 464
deleting 472
editing 468
executing 470
log files 475
overview 463
viewing historical runs 473
viewing properties 474
refreshing
break-off link snapshot 246
break-off mirror snapshot 246
instant snapshot 246
space-optimized snapshot 246
remote clusters 443
removing 447–448
wac resource 447
remounting
file system 290
file system snapshot 296
removing
disks from a disk group 186
DMP maintenance record 621
file change log 280
file system snapshot 300
volume mirrors 226
renaming
disks in a disk group 171
replace disks 181
replications
adding secondary 340
associating volume 348
configuring 329
monitoring 352
pausing 341
removing secondary 350
resuming 342
resynchronizing secondary 350
setting alerts 352
starting 343
stopping 344
switching primary 345
taking over primary 346
unconfiguring replication 351
report subscription
all subscriptions 129
delete 127
edit 126
my subscriptions 128
reports
about 120
all subscriptions 129
delete subscription 127
edit subscription 126
email 127
IOM deployment 527
my subscriptions 128
running 123
saving 124
subscribing 125
using 122
virtualization 659
resize
disk groups 168
volumes 240
resource action
invoking 431
resource dependency
viewing 438

resource type
editing attributes 434
resources
adding 436
bringing online 430
clearing faults 427
clearing from admin wait state 411
deleting 426
editing attributes 433
enabling 425–426
linking 435
managing 425
mark as critical 437
mark as non critical 438
modifying 436
offline and propagate 429
probing 428
taking offline 430
unlinking 435
restore data from snapshot 244
rules
about 104
creating 105
deleting 115
disabling 116
editing 111
enabling 116

S

scheduling file compression 362
scripts
vxlist 690
xdistc 698
Search feature 43, 45
service group dependency
viewing 415
service groups
adding 385
bringing online 404
clearing faults 410
deleting 398
disabling 391
disabling all resources 398
enabling 390
enabling all resources 397
flushing 396
freezing 394
linking 399, 402
managing 384
service groups *(continued)*
switch 408
taking offline 406
unfreezing 395
setting volume usage 253
signatures
custom 601, 606
enabling or disabling 598
exporting 601
registering 594–595
setting tunables 597
unregistering 596
Smart Folders 46–47
SmartAssist
about 318
SmartIO caching
about 306
creating 308
deleting 316
enabling or disabling 307, 312
impact analysis 311
modifying 313
viewing details 310
write-back 307
snapshot
add to a refresh schedule 251
configure break-off mirror snapshot 235
configure instant snapshot 229
dissociating 237
reattaching 239
refresh 246
remove from refresh schedule 252
restore data from snapshot 244
schedule refresh 249
splitting 254
solaris zones 645
Global Zones 645
non-Global Zones 645
solaris zones discovery 645
information 646
limitations 647
method 645
utilities 645
space estimation data logs 534
splitting
disk group 174
snapshot 254
starting
Cluster Server HAD 421

starting *(continued)*
synchronization of snapshots 255
stopping
Cluster Server HAD 422
stopping volume 213
stretch sites
site fencing preference 484
synchronizing change log 280
system list
modify 413

T

tablespaces
pinning to cache 317
unpinning from cache 317
tablespaces or files
pinning to cache 316
taking offline
resources 430
service groups 406
thin reclamation
on file systems or disks 359
overview 358
threshold settings
about 552
configuring 553
deleting 557
disabling 563
enabling 560
Thresholds
managing 324
Trim
running 194
scheduling 194
type definition
importing 381

U

unlinking
resources 435
unmounting
clustered file systems 287
file system snapshot 299
non clustered file systems 282
upgrading disk group 173
user group permissions
assigning on Organization 97
deleting on Organization 99
user group permissions *(continued)*
modifying on object 99
modifying on Organization 98

V

viewing
fire drill schedules 461
resource dependency 438
service group dependency 415
Virtual Business Services
overview 376
virtualization reports 659
virtualization server
disabling DMP paths 617
virtualization technology
kernel-based virtual machine 635
LPAR 635
Microsoft Hyper-V 635
Solaris LDom 635
Solaris zones 635
VMware 635
VMware discovery 637
datastore 640
ESX servers 637
information 638
near real-time update of virtual machine states 643
vCenter servers 637
VMware Infrastructure
discovery 637
VMware SDK 637
volume
renaming 218
usage 253
volume mirrors
remove 226
volumes
add snapshot to a refresh schedule 251
configure break-off mirror volume snapshot 235
configure instant snapshot 229
creating 200
deleting 215
disable fastresync 257
enable fastresync 256
moving 217
reactivate 215
recover 214
refresh snapshot 246

volumes *(continued)*
remove snapshot from refresh schedule 252
resize 240
restore data from snapshot 244
stopping 213
vomadm
about 684
Business application management 686
deployment 686
domain management 687–688
host management 685
service management 687
vxlist 690

W

wac resource
taking offline 447
Web services API
base URL 660
extended attributes 660
logging in
cURL 661
XPRTLC 661
logging out
cURL 662
XPRTLC 662
operation examples 669
XPRTLC cURL 674
operations 666
output examples 681
supported objects 663

X

xdistc 698

Z

Zone Agentlet 645