



Citrix XenServer® 7.0 Administrator's Guide

Published October 2017
1.0 Edition



Citrix XenServer® 7.0 Administrator's Guide

Copyright © 2016 Citrix Systems, Inc. All Rights Reserved.
Version: 7.0

Citrix, Inc.
851 West Cypress Creek Road
Fort Lauderdale, FL 33309
United States of America

Disclaimers

This document is furnished "AS IS." Citrix, Inc. disclaims all warranties regarding the contents of this document, including, but not limited to, implied warranties of merchantability and fitness for any particular purpose. This document may contain technical or other inaccuracies or typographical errors. Citrix, Inc. reserves the right to revise the information in this document at any time without notice. This document and the software described in this document constitute confidential information of Citrix, Inc. and its licensors, and are furnished under a license from Citrix, Inc.

Citrix Systems, Inc., the Citrix logo, Citrix XenServer and Citrix XenCenter, are trademarks of Citrix Systems, Inc. and/or one or more of its subsidiaries, and may be registered in the United States Patent and Trademark Office and in other countries. All other trademarks and registered trademarks are property of their respective owners.

Trademarks

Citrix®
XenServer®
XenCenter®

Contents

1. Document Overview	1
1.1. Introducing XenServer	1
1.1.1. Benefits of Using XenServer	1
1.1.2. Administering XenServer	2
1.2. XenServer Documentation	2
2. Managing Users	3
2.1. Authenticating Users With Active Directory (AD)	3
2.1.1. Configuring Active Directory Authentication	4
2.1.2. User Authentication	6
2.1.3. Removing Access for a User	7
2.1.4. Leaving an AD Domain	8
2.2. Role Based Access Control	8
2.2.1. Roles	9
2.2.2. Definitions of RBAC Roles and Permissions	10
2.2.3. Using RBAC with the CLI	16
2.2.3.1. To List All the Available Defined Roles in XenServer	16
2.2.3.2. To Display a List of Current Subjects:	16
2.2.3.3. To Add a Subject to RBAC	17
2.2.3.4. To Assign an RBAC Role to a Created subject	17
2.2.3.5. To Change a Subject's RBAC Role:	17
2.2.4. Auditing	18
2.2.4.1. Audit Log xe CLI Commands	18
2.2.4.2. To Obtain All Audit Records From the Pool	18
2.2.4.3. To Obtain Audit Records of the Pool Since a Precise Millisecond Timestamp	18
2.2.4.4. To Obtain Audit Records of the Pool Since a Precise Minute Timestamp	18
2.2.5. How Does XenServer Compute the Roles for the Session?	18
3. XenServer Hosts and Resource Pools	20
3.1. Hosts and Resource Pools Overview	20

3.2. Requirements for Creating Resource Pools	20
3.3. Creating a Resource Pool	21
3.4. Creating Heterogeneous Resource Pools	22
3.5. Adding Shared Storage	22
3.6. Removing a XenServer Host from a Resource Pool	23
3.7. Preparing a Pool of XenServer Hosts for Maintenance	23
3.8. Export Resource Pool Data	24
3.8.1. To Export Resource Data	25
3.9. High Availability	26
3.9.1. HA Overview	26
3.9.1.1. Overcommitting	27
3.9.1.2. Overcommitment Warning	27
3.9.1.3. Host Fencing	27
3.9.2. Configuration Requirements	27
3.9.3. Restart configuration settings	28
3.9.3.1. Start order	29
3.10. Enabling HA on your XenServer Pool	29
3.10.1. Enabling HA Using the CLI	29
3.10.2. Removing HA Protection from a VM using the CLI	30
3.10.3. Recovering an Unreachable Host	30
3.10.4. Shutting Down a host When HA is Enabled	30
3.10.5. Shutting Down a VM When it is Protected by HA	31
3.11. Host Power On	31
3.11.1. Powering on Hosts Remotely	31
3.11.2. Using the CLI to Manage Host Power On	31
3.11.2.1. To Enable Host Power On Using the CLI	31
3.11.2.2. To Turn on Hosts Remotely Using the CLI	32
3.11.3. Configuring a Custom Script for XenServer's Host Power On Feature	32
3.11.3.1. Key/Value Pairs	32
3.11.3.1.1. host.power_on_mode	32
3.11.3.1.2. host.power_on_config	32

3.11.3.2. Sample Script	33
3.12. Communicating with XenServer Hosts and Resource Pools	33
4. Networking	34
4.1. Networking Support	34
4.2. vSwitch Networks	34
4.3. XenServer Networking Overview	35
4.3.1. Network Objects	36
4.3.2. Networks	36
4.3.3. VLANs	36
4.3.3.1. Using VLANs with Management Interfaces	36
4.3.3.2. Using VLANs with Virtual Machines	36
4.3.3.3. Using VLANs with Dedicated Storage NICs	36
4.3.3.4. Combining Management Interfaces and Guest VLANs on a Single Host NIC	37
4.3.4. Jumbo frames	37
4.3.5. NIC Bonds	37
4.3.5.1. Active-Active Bonding	39
4.3.5.2. Active-Passive Bonding	40
4.3.5.3. LACP Link Aggregation Control Protocol Bonding	41
4.3.5.4. Switch Configuration	43
4.3.5.4.1. Switch Configuration for LACP Bonds	44
4.3.6. Initial Networking Configuration after Setup	44
4.3.7. Changing Networking Configuration	45
4.3.7.1. Modifying the Network Object	45
4.3.8. Changing the Up Delay for Bonds	45
4.4. Managing Networking Configuration	46
4.4.1. Cross-Server Private Networks	46
4.4.2. Creating Networks in a Standalone Server	47
4.4.3. Creating Networks in Resource Pools	47
4.4.4. Creating VLANs	47
4.4.5. Creating NIC Bonds on a Standalone Host	48

4.4.5.1. Creating a NIC Bond	48
4.4.5.2. Controlling the MAC Address of the Bond	48
4.4.5.3. Reverting NIC Bonds	49
4.4.6. Creating NIC Bonds in Resource Pools	49
4.4.6.1. Adding NIC Bonds to New Resource Pools	49
4.4.6.2. Adding NIC Bonds to an Existing Pool	50
4.4.7. Configuring a Dedicated Storage NIC	50
4.4.8. Using SR-IOV Enabled NICs	51
4.4.9. Controlling the Rate of Outgoing Data (QoS)	52
4.4.10. Changing Networking Configuration Options	53
4.4.10.1. Hostname	53
4.4.10.2. DNS Servers	53
4.4.10.3. Changing IP Address Configuration for a Standalone Host	53
4.4.10.4. Changing IP Address Configuration in Resource Pools	53
4.4.10.5. Management Interface	54
4.4.10.6. Disabling Management Access	55
4.4.10.7. Adding a New Physical NIC	55
4.4.10.8. Using Switch Port Locking	55
4.4.10.8.1. Requirements	55
4.4.10.8.2. Notes	55
4.4.10.8.3. Implementation Notes	56
4.4.10.8.4. Examples	56
4.4.10.8.5. How Switch-port Locking Works	57
4.4.10.8.6. VIF Locking-Mode States	57
4.4.10.8.7. Configuring Switch Port Locking	58
4.4.10.8.8. Preventing a Virtual Machine from Sending or Receiving Traffic from a Specific Network	59
4.4.10.8.9. Removing a VIF's Restriction to an IP Address	60
4.4.10.8.10. Simplifying VIF Locking Mode Configuration in the Cloud	60
4.4.10.8.11. Using Network Settings for VIF Traffic Filtering	61
4.5. Networking Troubleshooting	61
4.5.1. Diagnosing Network Corruption	61

4.5.2. Emergency Network Reset	62
4.5.2.1. Verifying the Network Reset	62
4.5.2.2. Using the CLI for Network Reset	62
4.5.2.2.1. Pool Master Command Line Examples	64
4.5.2.2.2. Pool Member Command Line Examples	64
5. Storage	66
5.1. Storage Overview	66
5.1.1. Storage Repositories (SRs)	66
5.1.2. Virtual Disk Image (VDI)	66
5.1.3. Physical Block Devices (PBDs)	66
5.1.4. Virtual Block Devices (VBDs)	67
5.1.5. Summary of Storage objects	67
5.1.6. Virtual Disk Data Formats	67
5.1.6.1. VDI Types	67
5.1.6.2. Creating a Raw Virtual Disk Using the xe CLI	67
5.1.6.3. Converting Between VDI Formats	68
5.1.6.4. VHD-based VDIs	68
5.2. Storage Repository Formats	69
5.2.1. Local LVM	69
5.2.1.1. LVM Performance Considerations	69
5.2.1.2. Creating a Local LVM SR (lvm)	70
5.2.2. Local EXT3	70
5.2.2.1. Creating a Local EXT3 SR (ext)	70
5.2.3. udev	70
5.2.4. ISO	71
5.2.5. Software iSCSI Support	71
5.2.5.1. XenServer Host iSCSI configuration	71
5.2.6. Software FCoE Storage	72
5.2.6.1. Creating a Software FCoE SR	72
5.2.7. Hardware Host Bus Adapters (HBAs)	72
5.2.7.1. Sample QLogic iSCSI HBA setup	72

5.2.7.2. Removing HBA-based SAS, FC or iSCSI Device Entries	73
5.2.8. Shared LVM Storage	73
5.2.8.1. Creating a Shared LVM Over iSCSI SR Using the Software iSCSI Initiator (lvmoiscsi)	73
5.2.8.2. Creating a Shared LVM over Fibre Channel / Fibre Channel over Ethernet / iSCSI HBA or SAS SR (lvmohba)	74
5.2.9. NFS and SMB	76
5.2.9.1. Creating a Shared NFS SR (NFS)	76
5.2.9.2. Creating a Shared SMB SR (SMB)	77
5.2.10. LVM over Hardware HBA	78
5.3. Storage Configuration	78
5.3.1. Creating Storage Repositories	78
5.3.2. Probing an SR	79
5.4. Storage Multipathing	81
5.5. XenServer and IntelliCache	82
5.5.1. IntelliCache Deployment	82
5.5.1.1. Enabling on Host Installation	83
5.5.1.2. Converting an Existing Host to Use Thin Provisioning	83
5.5.1.3. VM Boot Behavior	84
5.5.1.3.1. VM Caching Behavior Settings	84
5.5.1.4. Implementation Details and Troubleshooting	84
5.6. Storage Read Caching	85
5.6.1. Enabling and Disabling	85
5.6.2. Limitations	86
5.6.3. Comparison with IntelliCache	86
5.6.4. To Set the Read Cache Size	86
5.6.4.1. How to View the Current dom0 Memory Allocation	86
5.6.4.2. XenCenter Display Notes	87
5.7. Managing Storage Repositories	87
5.7.1. Removing SRs	87
5.7.2. Introducing an SR	88
5.7.3. Live LUN Expansion	88

5.7.4. Live VDI Migration	89
5.7.4.1. Limitations and Caveats	89
5.7.4.2. To Move Virtual Disks using XenCenter	89
5.7.5. Cold VDI Migration between SRs (Offline Migration)	89
5.7.5.1. Copying All of a VMs VDIs to a Different SR	89
5.7.5.2. Copying Individual VDIs to a Different SR	89
5.7.6. Converting Local Fibre Channel SRs to Shared SRs	90
5.7.7. Reclaiming Space for Block-Based Storage on the Backing Array Using Discard	90
5.7.8. Automatically Reclaiming Space When Deleting Snapshots	91
5.7.8.1. Reclaiming Space Using the Off Line Coalesce Tool	91
5.7.9. Adjusting the Disk IO Scheduler	92
5.7.10. Virtual Disk QoS Settings	92
6. Configuring VM Memory	94
6.1. What is Dynamic Memory Control (DMC)?	94
6.1.1. The Concept of Dynamic Range	94
6.1.2. The Concept of Static Range	94
6.1.3. DMC Behavior	95
6.1.4. How Does DMC Work?	95
6.1.5. Memory Constraints	95
6.2. xe CLI Commands	96
6.2.1. Display the Static Memory Properties of a VM	96
6.2.2. Display the Dynamic Memory Properties of a VM	96
6.2.3. Updating Memory Properties	97
6.2.4. Update Individual Memory Properties	97
6.3. Upgrade Issues	98
7. XenServer Memory Usage	99
7.1. Control Domain Memory	99
7.1.1. Changing the Amount of Memory Allocated to the Control Domain	99
7.1.2. How Much Memory is Available to VMs	100
8. Disaster Recovery and Backup	101

8.1. Understanding XenServer DR	101
8.2. DR Infrastructure Requirements	101
8.3. Deployment Considerations	103
8.3.1. Steps to Take Before a Disaster	103
8.3.2. Steps to Take After a Disaster	103
8.3.3. Steps to Take After a Recovery	103
8.4. Enabling Disaster Recovery in XenCenter	103
8.5. Recovering VMs and vApps in the Event of Disaster (Failover)	104
8.6. Restoring VMs and vApps to the Primary Site After Disaster (Failback)	105
8.7. Test Failover	105
8.8. vApps	107
8.8.1. Using the Manage vApps dialog box in XenCenter	107
8.9. Backing Up and Restoring XenServer Hosts and VMs	107
8.9.1. Backing up Virtual Machine metadata	109
8.9.1.1. Backing up single host installations	109
8.9.1.2. Backing up pooled installations	109
8.9.2. Backing up XenServer hosts	110
8.9.3. Backing up VMs	111
8.10. VM Snapshots	111
8.10.1. Regular Snapshots	111
8.10.2. Quiesced Snapshots	111
8.10.3. Snapshots with memory	112
8.10.4. Creating a VM Snapshot	112
8.10.5. Creating a snapshot with memory	112
8.10.6. To list all of the snapshots on a XenServer pool	113
8.10.7. To list the snapshots on a particular VM	113
8.10.8. Restoring a VM to its previous state	114
8.10.8.1. Deleting a snapshot	114
8.10.9. Snapshot Templates	115
8.10.9.1. Creating a template from a snapshot	115
8.10.9.2. Exporting a snapshot to a template	115

8.10.9.3. Advanced Notes for Quiesced Snapshots	116
8.11. Coping with machine failures	116
8.11.1. Member failures	116
8.11.2. Master failures	117
8.11.3. Pool failures	118
8.11.4. Coping with Failure due to Configuration Errors	118
8.11.5. Physical Machine failure	118
9. Monitoring and Managing XenServer	120
9.1. Monitoring XenServer Performance	120
9.1.1. Available Host Metrics	120
9.1.2. Available VM Metrics	123
9.1.3. Analyzing and Visualizing Metrics in XenCenter	125
9.1.3.1. Configuring Performance Graphs	125
9.1.3.1.1. Configuring the Graph Type	126
9.1.4. Configuring Metrics	126
9.1.5. Using RRDs	127
9.1.5.1. Analysing RRDs using HTTP	128
9.1.5.2. Analysing RRDs using rrd2csv	128
9.2. Alerts	128
9.2.1. Using XenCenter to View Alerts	128
9.2.1.1. XenCenter Performance Alerts	128
9.2.1.1.1. To Configure Performance Alerts	129
9.2.1.2. XenCenter Alerts	130
9.2.1.3. XenCenter Software Update Alerts	131
9.2.2. Configuring Performance Alerts Using the xe CLI	131
9.2.2.1. Generic Example Configuration	133
9.3. Configuring Email Alerts	133
9.3.1. Enabling Email Alerts Using XenCenter	134
9.3.2. Enabling Email Alerts using the xe CLI	134
9.3.2.1. Sending Email Alerts Through Authenticated SMTP Servers	134
9.3.2.1.1. Additional Configuration Options	134

9.4. Custom Fields and Tags	135
9.5. Custom Searches	135
9.6. Determining throughput of physical bus adapters	136
10. Troubleshooting	137
10.1. Support	137
10.2. Health Check	137
10.3. XenServer host logs	138
10.3.1. Sending host log messages to a central server	138
10.4. XenCenter logs	138
10.5. Troubleshooting connections between XenCenter and the XenServer host	138
A. Command Line Interface	140
A.1. Basic xe Syntax	140
A.2. Special Characters and Syntax	141
A.3. Command Types	141
A.3.1. Parameter Types	142
A.3.2. Low-level Parameter Commands	143
A.3.3. Low-level List Commands	143
A.4. xe Command Reference	144
A.4.1. Appliance Commands	144
A.4.1.1. Appliance Parameters	144
A.4.1.2. appliance-assert-can-be-recovered	144
A.4.1.3. appliance-create	144
A.4.1.4. appliance-destroy	145
A.4.1.5. appliance-recover	145
A.4.1.6. appliance-shutdown	145
A.4.1.7. appliance-start	145
A.4.2. Audit Commands	145
A.4.2.1. audit-log-get parameters	145
A.4.2.2. audit-log-get	145
A.4.3. Bonding Commands	146
A.4.3.1. Bond Parameters	146

A.4.3.2. bond-create	146
A.4.3.3. bond-destroy	146
A.4.4. CD Commands	146
A.4.4.1. CD Parameters	146
A.4.4.2. cd-list	147
A.4.5. Console Commands	148
A.4.5.1. Console Parameters	148
A.4.6. Disaster Recovery (DR) Commands	148
A.4.6.1. drtask-create	148
A.4.6.2. drtask-destroy	149
A.4.6.3. vm-assert-can-be-recovered	149
A.4.6.4. appliance-assert-can-be-recovered	149
A.4.6.5. appliance-recover	149
A.4.6.6. vm-recover	149
A.4.6.7. sr-enable-database-replication	149
A.4.6.8. sr-disable-database-replication	149
A.4.6.9. Example Usage	149
A.4.7. Event Commands	150
A.4.7.1. Event Classes	150
A.4.7.2. event-wait	150
A.4.8. GPU Commands	151
A.4.8.1. Physical GPU (pGPU) Parameters	151
A.4.8.2. pGPU Operations	152
A.4.8.2.1. pgpu-param-set	152
A.4.8.2.2. pgpu-param-get-uuid	152
A.4.8.2.3. pgpu-param-set-uuid	152
A.4.8.2.4. pgpu-param-add-uuid	152
A.4.8.3. GPU Group Parameters	152
A.4.8.3.1. GPU Group Operations	153
A.4.8.4. Virtual GPU Parameters	153
A.4.8.5. Virtual GPU Type Parameters	154

A.4.8.6. Virtual GPU Operations	155
A.4.8.6.1. vgpu-create	155
A.4.8.6.2. vgpu-destroy	155
A.4.8.6.3. Disabling VNC for VMs with virtual GPU	155
A.4.9. Host Commands	155
A.4.9.1. Host Selectors	155
A.4.9.2. Host Parameters	156
A.4.9.3. host-backup	159
A.4.9.4. host-bugreport-upload	159
A.4.9.5. host-crashdump-destroy	159
A.4.9.6. host-crashdump-upload	159
A.4.9.7. host-disable	159
A.4.9.8. host-dmesg	159
A.4.9.9. host-emergency-management-reconfigure	160
A.4.9.10. host-enable	160
A.4.9.11. host-evacuate	160
A.4.9.12. host-forget	160
A.4.9.13. host-get-system-status	160
A.4.9.14. host-get-system-status-capabilities	161
A.4.9.15. host-is-in-emergency-mode	162
A.4.9.16. host-apply-edition	162
A.4.9.17. host-license-add	162
A.4.9.18. host-license-view	162
A.4.9.19. host-logs-download	162
A.4.9.20. host-management-disable	162
A.4.9.21. host-management-reconfigure	163
A.4.9.22. host-power-on	163
A.4.9.23. host-get-cpu-features	163
A.4.9.24. host-set-cpu-features	163
A.4.9.25. host-set-power-on	163
A.4.9.26. host-reboot	164

A.4.9.27. host-restore	164
A.4.9.28. host-set-hostname-live	164
A.4.9.29. host-shutdown	164
A.4.9.30. host-syslog-reconfigure	165
A.4.9.31. host-data-source-list	165
A.4.9.32. host-data-source-record	165
A.4.9.33. host-data-source-forget	165
A.4.9.34. host-data-source-query	166
A.4.10. Message Commands	166
A.4.10.1. Message Parameters	166
A.4.10.2. message-create	166
A.4.10.3. message-destroy	166
A.4.10.4. message-list	166
A.4.11. Network Commands	167
A.4.11.1. Network Parameters	167
A.4.11.2. network-create	168
A.4.11.3. network-destroy	168
A.4.12. Patch (Update) Commands	168
A.4.12.1. Patch Parameters	168
A.4.12.2. patch-apply	169
A.4.12.3. patch-pool-clean	169
A.4.12.4. patch-pool-apply	169
A.4.12.5. patch-precheck	169
A.4.12.6. patch-upload	169
A.4.13. PBD Commands	169
A.4.13.1. PBD Parameters	169
A.4.13.2. pbd-create	170
A.4.13.3. pbd-destroy	170
A.4.13.4. pbd-plug	170
A.4.13.5. pbd-unplug	170
A.4.14. PIF Commands	170

A.4.14.1. PIF Parameters	171
A.4.14.2. pif-forget	173
A.4.14.3. pif-introduce	173
A.4.14.4. pif-plug	173
A.4.14.5. pif-reconfigure-ip	173
A.4.14.6. pif-scan	174
A.4.14.7. pif-unplug	174
A.4.15. Pool Commands	174
A.4.15.1. Pool Parameters	174
A.4.15.2. pool-designate-new-master	175
A.4.15.3. pool-dump-database	175
A.4.15.4. pool-eject	176
A.4.15.5. pool-emergency-reset-master	176
A.4.15.6. pool-emergency-transition-to-master	176
A.4.15.7. pool-ha-enable	176
A.4.15.8. pool-ha-disable	176
A.4.15.9. pool-join	176
A.4.15.10. pool-recover-slaves	176
A.4.15.11. pool-restore-database	176
A.4.15.12. pool-sync-database	177
A.4.16. Storage Manager Commands	177
A.4.16.1. SM Parameters	177
A.4.17. SR Commands	177
A.4.17.1. SR Parameters	178
A.4.17.2. sr-create	179
A.4.17.3. sr-destroy	179
A.4.17.4. sr-enable-database-replication	179
A.4.17.5. sr-disable-database-replication	179
A.4.17.6. sr-forget	179
A.4.17.7. sr-introduce	179
A.4.17.8. sr-probe	180

A.4.17.9. sr-scan	180
A.4.18. Task Commands	180
A.4.18.1. Task Parameters	180
A.4.18.2. task-cancel	181
A.4.19. Template Commands	181
A.4.19.1. VM Template Parameters	182
A.4.19.2. template-export	188
A.4.20. Update Commands	188
A.4.20.1. update-upload	188
A.4.21. User Commands	188
A.4.21.1. user-password-change	188
A.4.22. VBD Commands	188
A.4.22.1. VBD Parameters	188
A.4.22.2. vbd-create	190
A.4.22.3. vbd-destroy	190
A.4.22.4. vbd-eject	190
A.4.22.5. vbd-insert	190
A.4.22.6. vbd-plug	191
A.4.22.7. vbd-unplug	191
A.4.23. VDI Commands	191
A.4.23.1. VDI Parameters	191
A.4.23.2. vdi-clone	192
A.4.23.3. vdi-copy	193
A.4.23.4. vdi-create	193
A.4.23.5. vdi-destroy	193
A.4.23.6. vdi-forget	193
A.4.23.7. vdi-import	193
A.4.23.8. vdi-introduce	194
A.4.23.9. vdi-pool-migrate	194
A.4.23.10. vdi-resize	194
A.4.23.11. vdi-snapshot	194

A.4.23.12. vdi-unlock	194
A.4.24. VIF Commands	195
A.4.24.1. VIF Parameters	195
A.4.24.2. vif-create	196
A.4.24.3. vif-destroy	197
A.4.24.4. vif-plug	197
A.4.24.5. vif-unplug	197
A.4.24.6. vif-configure-ipv4	197
A.4.24.7. vif-configure-ipv6	197
A.4.25. VLAN Commands	197
A.4.25.1. vlan-create	197
A.4.25.2. pool-vlan-create	197
A.4.25.3. vlan-destroy	198
A.4.26. VM Commands	198
A.4.26.1. VM Selectors	198
A.4.26.2. VM Parameters	198
A.4.26.3. vm-assert-can-be-recovered	204
A.4.26.4. vm-cd-add	205
A.4.26.5. vm-cd-eject	205
A.4.26.6. vm-cd-insert	205
A.4.26.7. vm-cd-list	205
A.4.26.8. vm-cd-remove	205
A.4.26.9. vm-clone	206
A.4.26.10. vm-compute-maximum-memory	206
A.4.26.11. vm-copy	206
A.4.26.12. vm-crashdump-list	206
A.4.26.13. vm-data-source-list	207
A.4.26.14. vm-data-source-record	207
A.4.26.15. vm-data-source-forget	207
A.4.26.16. vm-data-source-query	207
A.4.26.17. vm-destroy	208

A.4.26.18. vm-disk-add	208
A.4.26.19. vm-disk-list	208
A.4.26.20. vm-disk-remove	208
A.4.26.21. vm-export	208
A.4.26.22. vm-import	209
A.4.26.23. vm-install	209
A.4.26.24. vm-memory-shadow-multiplier-set	209
A.4.26.25. vm-migrate	210
A.4.26.26. vm-reboot	211
A.4.26.27. vm-recover	211
A.4.26.28. vm-reset-powerstate	211
A.4.26.29. vm-resume	211
A.4.26.30. vm-shutdown	211
A.4.26.31. vm-start	212
A.4.26.32. vm-suspend	212
A.4.26.33. vm-uninstall	212
A.4.26.34. vm-vcpu-hotplug	212
A.4.26.35. vm-vif-list	212

Chapter 1. Document Overview

This document is a system administrator's guide for XenServer, the complete server virtualization platform from Citrix. It contains procedures to guide you through configuring a XenServer deployment. In particular, it focuses on setting up storage, networking and resource pools, and how to administer XenServer hosts using the `xe` command line interface.

This document covers the following topics:

- Managing Users with Active Directory and Role Based Access Control
- Creating Resource Pools and setting up High Availability
- Configuring and Managing Storage Repositories
- Configuring Virtual Machine Memory using Dynamic Memory Control
- Setting Control Domain Memory on a XenServer host
- Configuring Networking
- Recovering Virtual Machines using Disaster Recovery and Backing Up Data
- Monitoring XenServer Performance Metrics and Configuring Alerts
- Troubleshooting XenServer
- Using the XenServer `xe` command line interface

1.1. Introducing XenServer

XenServer is the complete server virtualization platform from Citrix. The XenServer package contains all you need to create and manage a deployment of virtual x86 computers running on Xen, the open-source paravirtualizing hypervisor with near-native performance. XenServer is optimized for both Windows and Linux virtual servers.

XenServer runs directly on server hardware without requiring an underlying operating system, which results in an efficient and scalable system. XenServer works by abstracting elements from the physical machine (such as hard drives, resources and ports) and allocating them to the *virtual machines* running on it.

A virtual machine (VM) is a computer composed entirely of software that can run its own operating system and applications as if it were a physical computer. A VM behaves exactly like a physical computer and contains its own virtual (software-based) CPU, RAM, hard disk and network interface card (NIC).

XenServer lets you create VMs, take VM disk snapshots and manage VM workloads. For a comprehensive list of major XenServer features, visit www.citrix.com/xenserver.

1.1.1. Benefits of Using XenServer

Using XenServer reduces costs by:

- Consolidating multiple VMs onto physical servers
- Reducing the number of separate disk images that need to be managed
- Allowing for easy integration with existing networking and storage infrastructures

Using XenServer increases flexibility by:

- Allowing you to schedule zero downtime maintenance by using XenMotion to live migrate VMs between XenServer hosts
- Increasing availability of VMs by using High Availability to configure policies that restart VMs on another XenServer host if one fails
- Increasing portability of VM images, as one VM image will work on a range of deployment infrastructures



1.1.2. Administering XenServer

There are two methods by which to administer XenServer: XenCenter and the XenServer Command-Line Interface (CLI).

XenCenter is a graphical, Windows-based user interface. XenCenter allows you to manage XenServer hosts, pools and shared storage, and to deploy, manage and monitor VMs from your Windows desktop machine.

The XenCenter on-line *Help* is a useful resource for getting started with XenCenter and for context-sensitive assistance.

The **XenServer Command-line Interface (CLI)** allows you to administer XenServer using the Linux-based `xe` commands.

1.2. XenServer Documentation

XenServer documentation shipped with this release includes:

- *XenServer Release Notes* cover new features in XenServer 7.0 and any advisories and known issues that affect this release.
- *XenServer Quick Start Guide* provides an introduction for new users to the XenServer environment and components. This guide steps through the installation and configuration essentials to get XenServer and the XenCenter management console up and running quickly. After installation, it demonstrates how to create a Windows VM, VM template and pool of XenServer hosts. It introduces basic administrative tasks and advanced features, such as shared storage, VM snapshots and XenMotion live migration.
- *XenServer Installation Guide* steps through the installation, configuration and initial operation of XenServer and the XenCenter management console.
- *XenServer Virtual Machine User's Guide* describes how to install Windows and Linux VMs within a XenServer environment. This guide explains how to create new VMs from installation media, from VM templates included in the XenServer package and from existing physical machines (P2V). It explains how to import disk images and how to import and export appliances.
- *XenServer Administrator's Guide* gives an in-depth description of the tasks involved in configuring a XenServer deployment, including setting up storage, networking and pools. It describes how to administer XenServer using the `xe` Command Line Interface.
- *vSwitch Controller User's Guide* is a comprehensive user guide to the vSwitch Controller for XenServer.
- *Supplemental Packs and the DDK* introduces the XenServer Driver Development Kit, which can be used to modify and extend the functionality of XenServer.
- *XenServer Software Development Kit Guide* presents an overview of the XenServer SDK. It includes code samples that demonstrate how to write applications that interface with XenServer hosts.
- *XenAPI Specification* is a reference guide for programmers to the XenServer API.

For additional resources, visit the [Citrix Knowledge Center](#).

Chapter 2. Managing Users

Defining users, groups, roles and permissions allows you to control who has access to your XenServer hosts and pools and what actions they can perform.

When you first install XenServer, a user account is added to XenServer automatically. This account is the local super user (LSU), or root, which is authenticated locally by the XenServer computer.

The LSU, or root, is a special user account intended for system administration and has all rights or permissions. In XenServer, the LSU is the default account at installation. The LSU is only authenticated by XenServer and does not require any external authentication service. If an external authentication service fails, the LSU can still log in and manage the system. The LSU can always access the XenServer physical server through SSH.

You can create additional users by adding the Active Directory accounts through either XenCenter's **Users** tab or the xe CLI. If your environment does not use Active Directory, you are limited to the LSU account.

Note:

When you create new users, XenServer does not assign newly created user accounts RBAC roles automatically. As a result, these accounts do not have any access to the XenServer pool until you assign them a role.

These permissions are granted through roles, as discussed in [Section 2.1, "Authenticating Users With Active Directory \(AD\)"](#).

2.1. Authenticating Users With Active Directory (AD)

If you want to have multiple user accounts on a server or a pool, you must use Active Directory user accounts for authentication. This lets XenServer users log in to a pool using their Windows domain credentials.

The only way you can configure varying levels of access for specific users is by enabling Active Directory authentication, adding user accounts, and assign roles to those accounts.

Active Directory users can use the xe CLI (passing appropriate `-u` and `-pw` arguments) and also connect to the host using XenCenter. Authentication is done on a per-resource pool basis.

Access is controlled by the use of *subjects*. A subject in XenServer maps to an entity on your directory server (either a user or a group). When external authentication is enabled, the credentials used to create a session are first checked against the local root credentials (in case your directory server is unavailable) and then against the subject list. To permit access, you must create a subject entry for the person or group you wish to grant access to. This can be done using XenCenter or the xe CLI.

If you are familiar with XenCenter, note that the XenServer CLI uses slightly different terminology to refer to Active Directory and user account features:

XenCenter Term	XenServer CLI Term
Users	Subjects
Add users	Add subjects

Understanding Active Directory Authentication in the XenServer Environment

Even though XenServer is Linux-based, XenServer lets you use Active Directory accounts for XenServer user accounts. To do so, it passes Active Directory credentials to the Active Directory domain controller.

When added to XenServer, Active Directory users and groups become XenServer subjects, generally referred to as simply users in XenCenter. When a subject is registered with XenServer, users/groups are authenticated with Active Directory on login and do not need to qualify their user name with a domain name.

Note:

By default, if you did not qualify the user name (for example, enter either mydomain\myuser or myser@mydomain.com), XenCenter always attempts to log users in to Active Directory authentication servers using the domain to which it is currently joined. The exception to this is the LSU account, which XenCenter always authenticates locally (that is, on the XenServer) first.

The external authentication process works as follows:

1. The credentials supplied when connecting to a server are passed to the Active Directory domain controller for authentication.
2. The domain controller checks the credentials. If they are invalid, the authentication fails immediately.
3. If the credentials are valid, the Active Directory controller is queried to get the subject identifier and group membership associated with the credentials.
4. If the subject identifier matches the one stored in the XenServer, the authentication is completed successfully.

When you join a domain, you enable Active Directory authentication for the pool. However, when a pool is joined to a domain, only users in that domain (or a domain with which it has trust relationships) can connect to the pool.

Note:

Manually updating the DNS configuration of a DHCP-configured network PIF is unsupported and might cause Active Directory integration, and consequently user authentication, to fail or stop working.

2.1.1. Configuring Active Directory Authentication

XenServer supports use of Active Directory servers using Windows 2003 or later.

Active Directory authentication for a XenServer host requires that the same DNS servers are used for both the Active Directory server (configured to allow for interoperability) and the XenServer host. In some configurations, the active directory server may provide the DNS itself. This can be achieved either using DHCP to provide the IP address and a list of DNS servers to the XenServer host, or by setting values in the PIF objects or using the installer if a manual static configuration is used.

Citrix recommends enabling DHCP to broadcast host names. In particular, the host names `localhost` or `linux` should not be assigned to hosts.

Warning:

XenServer hostnames should be unique throughout the XenServer deployment.

Note the following:

- XenServer labels its AD entry on the AD database using its hostname. Therefore, if two XenServer hosts have the same hostname and are joined to the same AD domain, the second XenServer will overwrite the AD entry of the first XenServer, regardless of if they are in the same or in different pools, causing the AD authentication on the first XenServer to stop working.

It is possible to use the same hostname in two XenServer hosts, as long as they join different AD domains.

- The XenServer hosts can be in different time-zones, as it is the UTC time that is compared. To ensure synchronization is correct, you may choose to use the same NTP servers for your XenServer pool and the Active Directory server.
- Mixed-authentication pools are not supported (that is, you cannot have a pool where some servers in the pool are configured to use Active Directory and some are not).
- The XenServer Active Directory integration uses the Kerberos protocol to communicate with the Active Directory servers. Consequently, XenServer does not support communicating with Active Directory servers that do not utilize Kerberos.
- For external authentication using Active Directory to be successful, it is important that the clocks on your XenServer hosts are synchronized with those on your Active Directory server. When XenServer joins the Active Directory domain, this will be checked and authentication will fail if there is too much skew between the servers.

Warning:

Host names must consist solely of no more than 63 alphanumeric characters, and must not be purely numeric.

Once you have Active Directory authentication enabled, if you subsequently add a server to that pool, you are prompted to configure Active Directory on the server joining the pool. When you are prompted for credentials on the joining server, enter Active Directory credentials with sufficient privileges to add servers to that domain.

Active Directory integration

Make sure that the following firewall ports are open for outbound traffic in order for XenServer to access the domain controllers.

Port	Protocol	Use
53	UDP/TCP	DNS
88	UDP/TCP	Kerberos 5
123	UDP	NTP
137	UDP	NetBIOS Name Service
139	TCP	NetBIOS Session (SMB)
389	UDP/TCP	LDAP
445	TCP	SMB over TCP
464	UDP/TCP	Machine password changes
3268	TCP	Global Catalog Search

Note:

To view the firewall rules on a Linux computer using *iptables*, run the following command:
`iptables -nL`

Note:

XenServer uses PowerBroker Identity Services (PBIS) to authenticate the AD user in the AD server, and to encrypt communications with the AD server.

How does XenServer manage the machine account password for AD integration?

Similarly to Windows client machines, PBIS automatically updates the machine account password, renewing it once every 30 days, or as specified in the machine account password renewal policy in the AD server.

Enabling external authentication on a pool

- External authentication using Active Directory can be configured using either XenCenter or the CLI using the command below.

```
xe pool-enable-external-auth auth-type=AD \
  service-name=<full-qualified-domain> \
  config:user=<username> \
  config:pass=<password>
```

The user specified needs to have Add/remove computer objects or workstations privileges, which is the default for domain administrators.

Note:

If you are not using DHCP on the network used by Active Directory and your XenServer hosts, you can use these two approaches to setup your DNS:

- Set up your domain DNS suffix search order for resolving non-FQDNs:

```
xe pif-param-set uuid=<pif-uuid_in_the_dns_subnetwork> \
  "other-config:domain=suffix1.com suffix2.com suffix3.com"
```

- Configure the DNS server to use on your XenServer hosts:

```
xe pif-reconfigure-ip mode=static dns=<dnshost> ip=<ip> \
  gateway=<gateway> netmask=<netmask> uuid=<uuid>
```

- Manually set the management interface to use a PIF that is on the same network as your DNS server:

```
xe host-management-reconfigure pif-uuid=<pif_in_the_dns_subnetwork>
```

Note:

External authentication is a per-host property. However, Citrix advises that you enable and disable this on a per-pool basis – in this case XenServer will deal with any failures that occur when enabling authentication on a particular host and perform any roll-back of changes that may be required, ensuring that a consistent configuration is used across the pool. Use the **host-param-list** command to inspect properties of a host and to determine the status of external authentication by checking the values of the relevant fields.

Disabling external authentication

- Use XenCenter to disable Active Directory authentication, or the following xe command:

```
xe pool-disable-external-auth
```

2.1.2. User Authentication

To allow a user access to your XenServer host, you must add a subject for that user or a group that they are in. (Transitive group memberships are also checked in the normal way, for example: adding a subject for group A, where group A contains group B and user 1 is a member of group B would permit access to user 1.) If you wish to manage user permissions in Active Directory, you could create a single group that you then add and

remove users to/from; alternatively, you can add and remove individual users from XenServer, or a combination of users and groups as your would be appropriate for your authentication requirements. The subject list can be managed from XenCenter or using the CLI as described below.

When authenticating a user, the credentials are first checked against the local root account, allowing you to recover a system whose AD server has failed. If the credentials (i.e. username then password) do not match/authenticate, then an authentication request is made to the AD server – if this is successful the user's information will be retrieved and validated against the local subject list, otherwise access will be denied. Validation against the subject list will succeed if the user or a group in the transitive group membership of the user is in the subject list.

Note:

When using Active Directory groups to grant access for Pool Administrator users who will require host ssh access, the number of users in the Active Directory group must not exceed 500.

Allowing a user access to XenServer using the CLI

- To add an AD subject to XenServer:

```
xe subject-add subject-name=<entity name>
```

The entity name should be the name of the user or group to which you want to grant access. You may optionally include the domain of the entity (for example, '<xend\user1>' as opposed to '<user1>') although the behavior will be the same unless disambiguation is required.

Using the CLI to Revoke User Access

- Find the user's subject identifier. This is the user or the group containing the user (removing a group would remove access to all users in that group, providing they are not also specified in the subject list). To do this use the `subject-list` command:

```
xe subject-list
```

This returns a list of all users.

You may wish to apply a filter to the list, for example to find the subject identifier for a user named `user1` in the `testad` domain, you could use the following command:

```
xe subject-list other-config:subject-name='<testad\user1>'
```

- Remove the user using the **subject-remove** command, passing in the subject identifier you learned in the previous step:

```
xe subject-remove subject-uuid=<subject-uuid>
```

- You may wish to terminate any current session this user has already authenticated. See [Terminating all authenticated sessions using xe](#) and [Terminating individual user sessions using xe](#) for more information about terminating sessions. If you do not terminate sessions the users whose permissions have been revoked may be able to continue to access the system until they log out.

Listing subjects with access

- To identify the list of users and groups with permission to access your XenServer host or pool, use the following command:

```
xe subject-list
```

2.1.3. Removing Access for a User

Once a user is authenticated, they will have access to the server until they end their session, or another user terminates their session. Removing a user from the subject list, or removing them from a group that is in the

subject list, will not automatically revoke any already-authenticated sessions that the user has; this means that they may be able to continue to access the pool using XenCenter or other API sessions that they have already created. In order to terminate these sessions forcefully, XenCenter and the CLI provide facilities to terminate individual sessions, or all currently active sessions. See the XenCenter help for more information on procedures using XenCenter, or below for procedures using the CLI.

Terminating all authenticated sessions using xe

- Execute the following CLI command:

```
xe session-subject-identifier-logout-all
```

Terminating individual user sessions using xe

1. Determine the subject identifier whose session you wish to log out. Use either the **session-subject-identifier-list** or **subject-list** xe commands to find this (the first shows users who have sessions, the second shows all users but can be filtered, for example, using a command like **xe subject-list other-config:subject-name=xendt\\user1** – depending on your shell you may need a double-backslash as shown).
2. Use the **session-subject-logout** command, passing the subject identifier you have determined in the previous step as a parameter, for example:

```
xe session-subject-identifier-logout subject-identifier=<subject-id>
```

2.1.4. Leaving an AD Domain

Warning:

When you leave the domain (that is, disable Active Directory authentication and disconnect a pool or server from its domain), any users who authenticated to the pool or server with Active Directory credentials are disconnected.

Use XenCenter to leave an AD domain. See the XenCenter help for more information. Alternately run the **pool-disable-external-auth** command, specifying the pool uuid if required.

Note:

Leaving the domain will not cause the host objects to be removed from the AD database. See [this](#) knowledge base article for more information about this and how to remove the disabled host entries.


2.2. Role Based Access Control

XenServer's Role Based Access Control (RBAC) allows you to assign users, roles, and permissions to control who has access to your XenServer and what actions they can perform. The XenServer RBAC system maps a user (or a group of users) to defined roles (a named set of permissions), which in turn have associated XenServer permissions (the ability to perform certain operations).

As users are not assigned permissions directly, but acquire them through their assigned role, management of individual user permissions becomes a matter of simply assigning the user to the appropriate role; this simplifies common operations. XenServer maintains a list of authorized users and their roles.

RBAC allows you to easily restrict which operations different groups of users can perform- thus reducing the probability of an accident by an inexperienced user.

To facilitate compliance and auditing, RBAC also provides an Audit Log feature.



RBAC depends on Active Directory for authentication services. Specifically, XenServer keeps a list of authorized users based on Active Directory user and group accounts. As a result, you must join the pool to the domain and add Active Directory accounts before you can assign roles.

The local super user (LSU), or root, is a special user account used for system administration and has all rights or permissions. In XenServer, the local super user is the default account at installation. The LSU is authenticated via XenServer and not external authentication service, so if the external authentication service fails, the LSU can still log in and manage the system. The LSU can always access the XenServer physical host via SSH.

RBAC process

This is the standard process for implementing RBAC and assigning a user or group a role:

1. Join the domain. See [Enabling external authentication on a pool](#)
2. Add an Active Directory user or group to the pool. This becomes a subject. See [Section 2.2.3.3, "To Add a Subject to RBAC"](#).
3. Assign (or modify) the subject's RBAC role. See [Section 2.2.3.4, "To Assign an RBAC Role to a Created subject"](#).

2.2.1. Roles

XenServer is shipped with the following six, pre-established roles:

- *Pool Administrator* (Pool Admin) – the same as being the local root. Can perform all operations.

Note:

The local super user (root) will always have the "Pool Admin" role. The Pool Admin role has the same permissions as the local root.

- *Pool Operator* (Pool Operator) – can do everything apart from adding/removing users and modifying their roles. This role is focused mainly on host and pool management (i.e. creating storage, making pools, managing the hosts etc.)
- *Virtual Machine Power Administrator* (VM Power Admin) – creates and manages Virtual Machines. This role is focused on provisioning VMs for use by a VM operator.
- *Virtual Machine Administrator* (VM Admin) – similar to a VM Power Admin, but cannot migrate VMs or perform snapshots.
- *Virtual Machine Operator* (VM Operator) – similar to VM Admin, but cannot create/destroy VMs – but can perform start/stop lifecycle operations.
- *Read-only* (Read Only) – can view resource pool and performance data.

Note:

You cannot add, remove or modify roles in this version of XenServer.

Warning:

You can not assign the role of *pool-admin* to an AD group which has more than 500 members, if you want users of the AD group to have SSH access.

For a summary of the permissions available for each role and more detailed information on the operations available for each permission, see [Section 2.2.2, “Definitions of RBAC Roles and Permissions”](#).

All XenServer users need to be allocated to an appropriate role. By default, all new users will be allocated to the Pool Administrator role. It is possible for a user to be assigned to multiple roles; in that scenario, the user will have the union of all the permissions of all their assigned roles.

A user's role can be changed in two ways:

1. Modify the subject -> role mapping (this requires the assign/modify role permission, only available to a Pool Administrator.)
2. Modify the user's containing group membership in Active Directory.

2.2.2. Definitions of RBAC Roles and Permissions

The following table summarizes which permissions are available for each role. For details on the operations available for each permission, see Definitions of permissions.

Table 2.1. Permissions available for each role

Role permissions	Pool Admin	Pool Operator	VM Power Admin	VM Admin	VM Operator	Read Only
Assign/modify roles	X					
Log in to (physical) server consoles (through SSH and XenCenter)	X					
Server backup/restore	X					
Import/export OVF/OVA packages and disk images	X					
Set cores per socket	X					
Convert virtual machines using XenServer Conversion Manager	X					
Switch-port locking	X	X				
Log out active user connections	X	X				
Create and dismiss alerts	X	X				
Cancel task of any user	X	X				
Pool management	X	X				
Storage XenMotion	X	X	X			
VM advanced operations	X	X	X			
VM create/destroy operations	X	X	X	X		
VM change CD media	X	X	X	X	X	

Role permissions	Pool Admin	Pool Operator	VM Power Admin	VM Admin	VM Operator	Read Only
VM change power state	X	X	X	X	X	
View VM consoles	X	X	X	X	X	
XenCenter view management operations	X	X	X	X	X	
Cancel own tasks	X	X	X	X	X	X
Read audit logs	X	X	X	X	X	X
Connect to pool and read all pool metadata	X	X	X	X	X	X
Configure virtual GPU	X	X				
View virtual GPU configuration	X	X	X	X	X	X
Access the config drive (CoreOS VMs only)	X					
Container management	X					
Health Check configuration	X	X				
View Health Check results and settings	X	X	X	X	X	X

Definitions of Permissions

The following table provides additional details about permissions:

Table 2.2. Definitions of permissions

Permission	Allows Assignee To	Rationale/Comments
Assign/modify roles	- Add/remove users - Add/remove roles from users - Enable and disable Active Directory integration (being joined to the domain)	This permission lets the user grant himself or herself any permission or perform any task. Warning: This role lets the user disable the Active Directory integration and all subjects added from Active Directory.
Log in to server consoles	- Server console access through ssh - Server console access through XenCenter	Warning: With access to a root shell, the assignee could arbitrarily reconfigure the entire system, including RBAC.
Server backup/restore VM create/destroy operations	- Back up and restore servers - Back up and restore pool metadata	The ability to restore a backup lets the assignee revert RBAC configuration changes.

Permission	Allows Assignee To	Rationale/Comments
Import/export OVF/OVA packages and disk images	- Import OVF and OVA packages - Import disk images - Export VMs as OVF/OVA packages	
Set cores-per-socket	Set the number of cores per socket for the VM's virtual CPUs	This permission enables the user to specify the topology for the VM's virtual CPUs.
Convert VMs using XenServer Conversion Manager	Convert VMware VMs to XenServer VMs	This permission lets the user convert workloads from VMware to XenServer by copying batches of VMware VMs to XenServer environment.
Switch-port locking	Control traffic on a network	This permission lets the user to block all traffic on a network by default, or define specific IP addresses from which a VM is allowed to send traffic.
Log out active user connections	Ability to disconnect logged in users	
Create/dismiss alerts		Warning: A user with this permission can dismiss alerts for the entire pool. Note: The ability to view alerts is part of the Connect to Pool and read all pool metadata permission.
Cancel task of any user	Cancel any user's running task	This permission lets the user request XenServer cancel an in-progress task initiated by any user.

Permission	Allows Assignee To	Rationale/Comments
Pool management	- Set pool properties (naming, default SRs) - Enable, disable, and configure HA - Set per-VM HA restart priorities - Configure DR and perform DR failover, failback and test failover operations - Enable, disable, and configure Workload Balancing (WLB) - Add and remove server from pool - Emergency transition to master - Emergency master address - Emergency recover slaves - Designate new master - Manage pool and server certificates - Patching - Set server properties - Configure server logging - Enable and disable servers - Shut down, reboot, and power-on servers - Restart toolstack - System status reports - Apply license - Live migration of all other VMs on a server to another server, due to either Maintenance Mode, or HA - Configure server management interface and secondary interfaces - Disable server management - Delete crashdumps - Add, edit, and remove networks - Add, edit, and remove PBDs/PIFs/VLANs/Bonds/SRs - Add, remove, and retrieve secrets	This permission includes all the actions required to maintain a pool. Note: If the management interface is not functioning, no logins can authenticate except local root logins.

Permission	Allows Assignee To	Rationale/Comments
Storage XenMotion	- Migrate VMs from one host to another host - Move Virtual Disk (VDIs) from one SR to another SR	
VM advanced operations	- Adjust VM memory (through Dynamic Memory Control) - Create a VM snapshot with memory, take VM snapshots, and roll-back VMs - Migrate VMs - Start VMs, including specifying physical server - Resume VMs	This permission provides the assignee with enough privileges to start a VM on a different server if they are not satisfied with the server XenServer selected.
VM create/destroy operations	- Install or delete - Clone/copy VMs - Add, remove, and configure virtual disk/CD devices - Add, remove, and configure virtual network devices - Import/export XVA files - VM configuration change	Server backup/restore VM create/destroy operations
VM change CD media	- Eject current CD - Insert new CD	Import/export OVF/OVA packages; import disk images
VM change power state	- Start VMs (automatic placement) - Shut down VMs - Reboot VMs - Suspend VMs - Resume VMs (automatic placement)	This permission does not include start_on, resume_on, and migrate, which are part of the VM advanced operations permission.
View VM consoles	See and interact with VM consoles	This permission does not let the user view server consoles.
XenCenter view management operations	- Create and modify global XenCenter folders - Create and modify global XenCenter custom fields - Create and modify global XenCenter searches	Folders, custom fields, and searches are shared between all users accessing the pool
Cancel own tasks	Lets a user cancel their own tasks	

Permission	Allows Assignee To	Rationale/Comments
Read audit log	Download the XenServer audit log	
Connect to pool and read all pool metadata	- Log in to pool - View pool metadata - View historical performance data - View logged in users - View users and roles - View messages - Register for and receive events	
Configure virtual GPU	- Specify a pool-wide placement policy - Assign a virtual GPU to a VM - Remove a virtual GPU from a VM - Modify allowed virtual GPU types - Create, destroy, or assign a GPU group	
View virtual GPU configuration	View GPUs, GPU placement policies, and virtual GPU assignments	
Access the config drive (CoreOS VMs only)	- Access the config driver of the VM - Modify the cloud-config parameters	
Container management	- Start - Stop - Pause - Resume - Access information about the container	
Health Check configuration	- Enable Health Check - Disable Health Check - Update Health Check settings - Manually upload a server status report	
View Health Check results and settings	- View the results of a Health Check upload - View Health Check enrollment settings	

Note:

In some cases, a Read Only user cannot move a resource into a folder in XenCenter, even after receiving an elevation prompt and supplying the credentials of a more privileged user. In this case, log on to XenCenter as the more privileged user and retry the action.

2.2.3. Using RBAC with the CLI

2.2.3.1. To List All the Available Defined Roles in XenServer

- Run the command: `xe role-list`

This command returns a list of the currently defined roles, for example:

```

uuid( R0): 0165f154-ba3e-034e-6b27-5d271af109ba
name ( R0): pool-admin
description ( R0): The Pool Administrator role has full access to all
features and settings, including accessing Dom0 and managing subjects,
roles and external authentication

uuid ( R0): b9ce9791-0604-50cd-0649-09b3284c7dfd
name ( R0): pool-operator
description ( R0): The Pool Operator role manages host- and pool-wide resources,
including setting up storage, creating resource pools and managing patches, and
high availability (HA).

uuid( R0): 7955168d-7bec-10ed-105f-c6a7e6e63249
name ( R0): vm-power-admin
description ( R0): The VM Power Administrator role has full access to VM and
template management and can choose where to start VMs and use the dynamic memory
control and VM snapshot features

uuid ( R0): aaa00ab5-7340-bfbc-0d1b-7cf342639a6e
name ( R0): vm-admin
description ( R0): The VM Administrator role can manage VMs and templates

uuid ( R0): fb8d4ff9-310c-a959-0613-54101535d3d5
name ( R0): vm-operator
description ( R0): The VM Operator role can use VMs and interact with VM consoles

uuid ( R0): 7233b8e3-eacb-d7da-2c95-f2e581cdbf4e
name ( R0): read-only
description ( R0): The Read-Only role can log in with basic read-only access

```

Note:

This list of roles is static; it is not possible to add, remove, or modify roles.

2.2.3.2. To Display a List of Current Subjects:

- Run the command `xe subject-list`

This will return a list of XenServer users, their uuid, and the roles they are associated with:



```

uuid ( R0): bb6dd239-1fa9-a06b-a497-3be28b8dca44
subject-identifier ( R0): S-1-5-21-1539997073-1618981536-2562117463-2244
other-config (MRO): subject-name: example01\user_vm_admin; subject-upn: \
  user_vm_admin@XENDT.NET; subject-uid: 1823475908; subject-gid: 1823474177; \
  subject-sid: S-1-5-21-1539997073-1618981536-2562117463-2244; subject-gecos: \
  user_vm_admin; subject-displayname: user_vm_admin; subject-is-group: false; \
  subject-account-disabled: false; subject-account-expired: false; \
  subject-account-locked: false; subject-password-expired: false
roles (SR0): vm-admin

```

```

uuid ( R0): 4fe89a50-6a1a-d9dd-afb9-b554cd00c01a
subject-identifier ( R0): S-1-5-21-1539997073-1618981536-2562117463-2245
other-config (MRO): subject-name: example02\user_vm_op; subject-upn: \
  user_vm_op@XENDT.NET; subject-uid: 1823475909; subject-gid: 1823474177; \
  subject-sid: S-1-5-21-1539997073-1618981536-2562117463-2245; \
  subject-gecos: user_vm_op; subject-displayname: user_vm_op; \
  subject-is-group: false; subject-account-disabled: false; \
  subject-account-expired: false; subject-account-locked: \
  false; subject-password-expired: false
roles (SR0): vm-operator

```

```

uuid ( R0): 8a63fbf0-9ef4-4fef-b4a5-b42984c27267
subject-identifier ( R0): S-1-5-21-1539997073-1618981536-2562117463-2242
other-config (MRO): subject-name: example03\user_pool_op; \
  subject-upn: user_pool_op@XENDT.NET; subject-uid: 1823475906; \
  subject-gid: 1823474177; subject-s id:
  S-1-5-21-1539997073-1618981536-2562117463-2242; \
  subject-gecos: user_pool_op; subject-displayname: user_pool_op; \
  subject-is-group: false; subject-account-disabled: false; \
  subject-account-expired: false; subject-account-locked: \
  false; subject-password-expired: false
roles (SR0): pool-operator

```

2.2.3.3. To Add a Subject to RBAC

In order to enable existing AD users to use RBAC, you will need to create a subject instance within XenServer, either for the AD user directly, or for one of their containing groups:

1. Run the command `xe subject-add subject-name=<AD user/group>`

This adds a new subject instance.

2.2.3.4. To Assign an RBAC Role to a Created subject

Once you have added a subject, you can assign it to an RBAC role. You can refer to the role by either its uuid or name:

1. Run the command:

```
xe subject-role-add uuid=<subject uuid> role-uuid=<role_uuid>
```

or

```
xe subject-role-add uuid=<subject uuid> role-name=<role_name>
```

For example, the following command adds a subject with the uuid b9b3d03b-3d10-79d3-8ed7-a782c5ea13b4 to the Pool Administrator role:

```
xe subject-role-add uuid=b9b3d03b-3d10-79d3-8ed7-a782c5ea13b4 role-name=pool-admin
```

2.2.3.5. To Change a Subject's RBAC Role:

To change a user's role it is necessary to remove them from their existing role, and add them to a new role:

1. Run the commands:

```
xe subject-role-remove uuid=<subject uuid> role-name= \
    <role_name_to_remove>
xe subject-role-add uuid=<subject uuid > role-name= \
    <role_name_to_add>
```

To ensure that the new role takes effect, the user should be logged out and logged back in again (this requires the "Logout Active User Connections" permission - available to a Pool Administrator or Pool Operator).

Warning:

Once you have added or removed a pool-admin subject, there can be a delay of a few seconds for ssh sessions associated to this subject to be accepted by all hosts of the pool.

2.2.4. Auditing

The RBAC audit log will record any operation taken by a logged-in user.

- the message will explicitly record the Subject ID and user name associated with the session that invoked the operation.
- if an operation is invoked for which the subject does not have authorization, this will be logged.
- if the operation succeeded then this is recorded; if the operation failed then the error code is logged.

2.2.4.1. Audit Log xe CLI Commands

```
xe audit-log-get [since=<timestamp>] filename=<output filename>
```

This command downloads to a file all the available records of the RBAC audit file in the pool. If the optional parameter 'since' is present, then it only downloads the records from that specific point in time.

2.2.4.2. To Obtain All Audit Records From the Pool

Run the following command:

```
xe audit-log-get filename=/tmp/auditlog-pool-actions.out
```

2.2.4.3. To Obtain Audit Records of the Pool Since a Precise Millisecond Timestamp

Run the following command:

```
xe audit-log-get since=2009-09-24T17:56:20.530Z \
filename=/tmp/auditlog-pool-actions.out
```


2.2.4.4. To Obtain Audit Records of the Pool Since a Precise Minute Timestamp

Run the following command:

```
xe audit-log-get since=2009-09-24T17:56Z \
filename=/tmp/auditlog-pool-actions.out
```

2.2.5. How Does XenServer Compute the Roles for the Session?

1. The subject is authenticated via the Active Directory server to verify which containing groups the subject may also belong to.
2. XenServer then verifies which roles have been assigned both to the subject, and to its containing groups.
3. As subjects can be members of multiple Active Directory groups, they will inherit all of the permissions of the associated roles.



In this illustration, since Subject 2 (Group 2) is the Pool Operator and User 1 is a member of Group 2, when Subject 3 (User 1) tries to log in, he or she inherits both Subject 3 (VM Operator) and Group 2 (Pool Operator) roles. Since the Pool Operator role is higher, the resulting role for Subject 3 (User 1) is Pool Operator and not VM Operator.

Chapter 3. XenServer Hosts and Resource Pools

This chapter describes how resource pools can be created through a series of examples using the `xe` command line interface (CLI). A simple NFS-based shared storage configuration is presented and a number of simple VM management examples are discussed. Procedures for dealing with physical node failures are also described.

3.1. Hosts and Resource Pools Overview

A *resource pool* comprises multiple XenServer host installations, bound together into a single managed entity which can host Virtual Machines. When combined with shared storage, a resource pool enables VMs to be started on *any* XenServer host which has sufficient memory and then dynamically moved between XenServer hosts while running with minimal downtime (XenMotion). If an individual XenServer host suffers a hardware failure, then the administrator can restart the failed VMs on another XenServer host in the same resource pool. If high availability (HA) is enabled on the resource pool, VMs will automatically be moved if their host fails. Up to 16 hosts are supported per resource pool, although this restriction is not enforced.

A pool always has at least one physical node, known as the *master*. Only the master node exposes an administration interface (used by XenCenter and the XenServer Command Line Interface, known as the `xe` CLI); the master forwards commands to individual members as necessary.

Note:

If the pool's master fails, master re-election will only take place if High Availability is enabled.

3.2. Requirements for Creating Resource Pools

A resource pool is a homogeneous (or heterogeneous with restrictions, see [Section 3.4, “Creating Heterogeneous Resource Pools”](#)) aggregate of one or more XenServer hosts, up to a maximum of 16. The definition of homogeneous is:

- the CPUs on the server joining the pool are the same (in terms of vendor, model, and features) as the CPUs on servers already in the pool.
- the server joining the pool is running the same version of XenServer software, at the same patch level, as servers already in the pool

The software will enforce additional constraints when joining a server to a pool – in particular:

- it is not a member of an existing resource pool
- it has no shared storage configured
- there are no running or suspended VMs on the XenServer host which is joining
- there are no active operations on the VMs in progress, such as one shutting down

You must also check that the clock of the host joining the pool is synchronized to the same time as the pool master (for example, by using NTP), that its management interface is not bonded (you can configure this once the host has successfully joined the pool), and that its management IP address is static (either configured on the host itself or by using an appropriate configuration on your DHCP server).

XenServer hosts in resource pools may contain different numbers of physical network interfaces and have local storage repositories of varying size. In practice, it is often difficult to obtain multiple servers with the exact same CPUs, and so minor variations are permitted. If you are sure that it is acceptable in your environment for hosts with varying CPUs to be part of the same resource pool, then the pool joining operation can be forced by passing a `--force` parameter.

Note:

The requirement for a XenServer host to have a static IP address to be part of a resource pool also applies to servers providing shared NFS or iSCSI storage for the pool.

Although not a strict technical requirement for creating a resource pool, the advantages of pools (for example, the ability to dynamically choose on which XenServer host to run a VM and to dynamically move a VM between XenServer hosts) are only available if the pool has one or more shared storage repositories. If possible, postpone creating a pool of XenServer hosts until shared storage is available. Once shared storage has been added, Citrix recommends that you move existing VMs whose disks are in local storage into shared storage. This can be done using the **xe vm-copy** command or using XenCenter.

3.3. Creating a Resource Pool

Resource pools can be created using either the XenCenter management console or the CLI. When you join a new host to a resource pool, the joining host synchronizes its local database with the pool-wide one, and inherits some settings from the pool:

- VM, local, and remote storage configuration is added to the pool-wide database. All of these will still be tied to the joining host in the pool unless you explicitly take action to make the resources shared after the join has completed.
- The joining host inherits existing shared storage repositories in the pool and appropriate PBD records are created so that the new host can access existing shared storage automatically.
- Networking information is partially inherited to the joining host: the *structural* details of NICs, VLANs and bonded interfaces are all inherited, but *policy* information is not. This policy information, which must be re-configured, includes:
 - the IP addresses of management NICs, which are preserved from the original configuration
 - the location of the management interface, which remains the same as the original configuration. For example, if the other pool hosts have their management interfaces on a bonded interface, then the joining host must be explicitly migrated to the bond once it has joined.
 - Dedicated storage NICs, which must be re-assigned to the joining host from XenCenter or the CLI, and the PBDs re-plugged to route the traffic accordingly. This is because IP addresses are not assigned as part of the pool join operation, and the storage NIC is not useful without this configured correctly. See [Section 4.4.7, “Configuring a Dedicated Storage NIC”](#) for details on how to dedicate a storage NIC from the CLI.

To join XenServer hosts *host1* and *host2* into a resource pool using the CLI

1. Open a console on XenServer host *host2*.
2. Command XenServer host *host2* to join the pool on XenServer host *host1* by issuing the command:

```
xe pool-join master-address=<host1> master-username=<administrators_username> \
master-password=<password>
```

The *master-address* must be set to the fully-qualified domain name of XenServer host *host1* and the *password* must be the administrator password set when XenServer host *host1* was installed.

Naming a resource pool

- XenServer hosts belong to an unnamed pool by default. To create your first resource pool, rename the existing nameless pool. Use tab-complete to find the `pool_uuid`:

```
xe pool-param-set name-label=<"New Pool"> uuid=<pool_uuid>
```


3.4. Creating Heterogeneous Resource Pools

XenServer simplifies expanding deployments over time by allowing disparate host hardware to be joined into a resource pool, known as heterogeneous resource pools. Heterogeneous resource pools are made possible by leveraging technologies in Intel (FlexMigration) and AMD (Extended Migration) CPUs that provide CPU "masking" or "leveling". These features allow a CPU to be configured to *appear* as providing a different make, model, or functionality than it actually does. This enables you to create pools of hosts with disparate CPUs but still safely support live migrations.

Note:

The CPUs of XenServer hosts joining a heterogeneous pool must be of the same vendor (that is, AMD, Intel) as the CPUs on hosts already in the pool, though the specific type, (family, model and stepping numbers) need not be.

XenServer 7.0 simplifies the support of heterogeneous pools. In previous versions of XenServer, a new pool member with a feature set different to that of the pool had to be CPU masked by users before it was allowed to join the pool. Hosts can now be added to existing resource pools, irrespective of the underlying CPU type (as long as the CPU is from the same vendor family). The pool feature set is dynamically calculated every time:

- a new host joins the pool
- a pool member leaves the pool
- a pool member reconnects following a reboot

Any change in the pool feature set does not affect VMs that are currently running in the pool. A Running VM will continue to use the feature set which was applied when it was started. This feature set is fixed at boot and persists across migrate, suspend, and resume operations. In scenarios where a pool level drops when a less-capable host joins the pool, a running VM can be migrated to any host in the pool, except the newly added host. When you attempt to move or migrate a VM to a different host within or across pools, XenServer performs migration checks to compare the VM's feature set against the feature set of the destination host. If the feature sets are found to be compatible, the VM will be allowed to migrate. This enables the VM to move freely within and across pools, regardless of the CPU features the VM is using. If you use Workload Balancing (WLB) to choose an optimal destination host to migrate your VM, a host with incompatible feature set will not be recommended as the destination host.

3.5. Adding Shared Storage

For a complete list of supported shared storage types, see the [Storage chapter](#). This section demonstrates how shared storage (represented as a storage repository) can be created on an existing NFS server.

Adding NFS shared storage to a resource pool using the CLI

1. Open a console on any XenServer host in the pool.
2. Create the storage repository on `<server:/path>` by issuing the command

```
xe sr-create content-type=user type=nfs name-label=<"Example SR"> shared=true \
  device-config:server=<server> \
  device-config:serverpath=<path>
```

The `device-config:server` refers to the hostname of the NFS server and `device-config:serverpath` refers to the path on the NFS server. Since `shared` is set to true, the shared storage will be automatically connected to every XenServer host in the pool and any XenServer hosts that subsequently join will also be connected to the storage. The Universally Unique Identifier (UUID) of the created storage repository will be printed on the screen.

3. Find the UUID of the pool by the command


```
xe pool-list
```
4. Set the shared storage as the pool-wide default with the command

```
xe pool-param-set uuid=<pool_uuid> default-SR=<sr_uuid>
```

Since the shared storage has been set as the pool-wide default, all future VMs will have their disks created on shared storage by default. See [Chapter 5, Storage](#) for information about creating other types of shared storage.

3.6. Removing a XenServer Host from a Resource Pool

Note:

Before removing a XenServer host from a pool, ensure that you shut down all the VMs running on that host. Otherwise, you may see a warning stating that the host cannot be removed.

When a XenServer host is removed (*ejected*) from a pool, the machine is rebooted, reinitialized, and left in a state equivalent to that after a fresh installation. It is important not to eject a XenServer host from a pool if there is important data on the local disks.

To remove a host from a resource pool using the CLI

1. Open a console on any host in the pool.
2. Find the UUID of the host by running the command

```
xe host-list
```

3. Eject the required host from the pool:

```
xe pool-eject host-uuid=<host_uuid>
```

The XenServer host will be ejected and left in a freshly-installed state.

Warning:

Do *not* eject a host from a resource pool if it contains important data stored on its local disks. All of the data will be erased upon ejection from the pool. If you wish to preserve this data, copy the VM to shared storage on the pool first using XenCenter, or the **xe vm-copy** CLI command.

When a XenServer host containing locally stored VMs is ejected from a pool, those VMs will still be present in the pool database and visible to the other XenServer hosts. They will not start until the virtual disks associated with them have been changed to point at shared storage which can be seen by other XenServer hosts in the pool, or simply removed. It is for this reason that you are strongly advised to move any local storage to shared storage upon joining a pool, so that individual XenServer hosts can be ejected (or physically fail) without loss of data.

3.7. Preparing a Pool of XenServer Hosts for Maintenance

Before performing maintenance operations on a XenServer host that is part of a resource pool, you should disable it (which prevents any VMs from being started on it), then migrate its VMs to another XenServer host in the pool. This can most readily be accomplished by placing the XenServer host into Maintenance mode using XenCenter. See the XenCenter Help for details.

Note:

Placing the master host into maintenance mode will result in the loss of the last 24hrs of RRD updates for offline VMs. This is because the backup synchronization occurs every 24hrs.

Warning:

Citrix highly recommends rebooting all XenServers prior to installing an update and then verifying their configuration. This is because some configuration changes only take effect when a XenServer is rebooted, so the reboot may uncover configuration problems that would cause the update to fail.



To prepare a XenServer host in a pool for maintenance operations using the CLI

1. Run the command

```
xe host-disable uuid=<xenserver_host_uuid>  
xe host-evacuate uuid=<xenserver_host_uuid>
```

This will disable the XenServer host and then migrate any running VMs to other XenServer hosts in the pool.

2. Perform the desired maintenance operation.
3. Once the maintenance operation is completed, enable the XenServer host:

```
xe host-enable
```

Restart any halted VMs and/or resume any suspended VMs.

3.8. Export Resource Pool Data

Note:

Export Resource Pool Data is available for XenServer Enterprise Edition customers, or those who have access to XenServer through their XenApp/XenDesktop entitlement. To learn more about XenServer editions, and to find out how to upgrade, visit the Citrix website [here](#). For detailed information on Licensing - refer to [XenServer 7.0 Licensing FAQ](#).

The Export Resource Data option allows you to generate a resource data report for your pool and export the report into a .xls or .csv file. This report provides detailed information about various resources in the pool such as, servers, networks, storage, virtual machines, VDIs, and GPUs. This feature enables administrators to track, plan, and assign resources based on various workloads such as CPU, storage, and Network.

The following table provides a list of resources and various types of resource data included in the report:

Resource	Resource Data
Server	• Name • Pool Master • UUID • Address • CPU Usage • Network (avg/max KBs) • Used Memory • Storage • Uptime • Description
Networks	• Name • Link Status • MAC • MTU • VLAN • Type • Location

Resource	Resource Data
VDI	• Name • Type • UUID • Size • Storage • Description
Storage	• Name • Type • UUID • Size • Location • Description
VMs	• Name • Power State • Running on • Address • MAC • NIC • Operating System • Storage • Used Memory • CPU Usage • UUID • Uptime • Template • Description
GPU	Note: Information about GPUs will be available only if there are GPUs attached to your XenServer host. • Name • Servers • PCI Bus Path • UUID • Power Usage • Temperature • Used Memory • Computer Utilization

3.8.1. To Export Resource Data

1. In the XenCenter Navigation pane, click **Infrastructure** and then click on the pool

2. Click the **Pool** menu and then select **Export Resource Data**.
3. Browse to a location where you would like to save report and then click **Save**.

3.9. High Availability

3.9.1. HA Overview

HA is a set of automatic features designed to plan for, and safely recover from, issues which take down XenServer hosts or make them unreachable. For example physically disrupted networking or host hardware failures.

Note:

HA is designed to be used in conjunction with storage multipathing and network bonding to create a system which is resilient to, and can recover from hardware faults. HA should **always** be used with multipathed storage and bonded networking.

Firstly, HA ensures that in the event of a host becoming unreachable or unstable, VMs which were known to be running on that host are shut down and able to be restarted elsewhere. This avoids the scenario where VMs are started (manually or automatically) on a new host and at some point later, the original host is able to recover. This scenario could lead two instances of the same VM running on different hosts, and a corresponding high probability of VM disk corruption and data loss.

Secondly, HA recovers administrative control of a pool in the event that the pool master becomes unreachable or unstable. HA ensures that administrative control is restored automatically without any manual intervention.

Optionally, HA can also automate the process of restarting VMs on hosts which are known to be in a good state without manual intervention. These VMs can be scheduled for restart in groups to allow time to start services. This allows infrastructure VMs to started before their dependent VMs (For example, a DHCP server before its dependent SQL server).

Warning:

HA is designed to be used in conjunction with multipathed storage and bonded networking, and this should be configured before attempting to set up HA. Customers who do not set up multipathed networking and storage, may see unexpected host reboot behaviour (Self Fencing) in the event of infrastructure instability. For more information see [CTX134880 - Designing XenServer Network Configurations](#) and [CTX134881 - Configuring iSCSI Multipathing Support for XenServer](#)

3.9.1.1. Overcommitting

A pool is overcommitted if the VMs that are currently running could not be restarted elsewhere following a user-defined number of host failures.

This would happen if there was not enough free memory across the pool to run those VMs following failure. However there are also more subtle changes which can make HA guarantees unsustainable: changes to Virtual Block Devices (VBDs) and networks can affect which VMs may be restarted on which hosts. Currently it is not possible for XenServer to check all actions before they occur and determine if they will cause violation of HA demands. However an asynchronous notification is sent if HA becomes unsustainable.

XenServer dynamically maintains a *failover plan* which details what to do if a set of hosts in a pool fail at any given time. An important concept to understand is the *host failures to tolerate* value, which is defined as part of HA configuration. This determines the number of failures that is allowed without any loss of service. For example, if a resource pool consisted of 16 hosts, and the tolerated failures is set to 3, the pool calculates a failover plan that allows for any 3 hosts to fail and still be able to restart VMs on other hosts. If a plan cannot be found, then the pool is considered to be *overcommitted*. The plan is dynamically recalculated based on VM lifecycle operations and movement. Alerts are sent (either through XenCenter or e-mail) if changes (for example the addition on new VMs to the pool) cause your pool to become overcommitted.

3.9.1.2. Overcommitment Warning

If you attempt to start or resume a VM and that action causes the pool to be overcommitted, a warning alert is raised. This warning is displayed in XenCenter and is also available as a message instance through the Xen API. The message may also be sent to an email address if configured. You will then be allowed to cancel the operation, or proceed anyway. Proceeding causes the pool to become overcommitted. The amount of memory used by VMs of different priorities is displayed at the pool and host levels.

3.9.1.3. Host Fencing

If a server failure occurs such as the loss of network connectivity or a problem with the control stack is encountered, the XenServer host self-fences to ensure that the VMs are not running on two servers simultaneously. When a fence action is taken, the server immediately and abruptly restarts, causing all VMs running on it to be stopped. The other servers will detect that the VMs are no longer running and the VMs will be restarted according to the restart priorities assign to them. The fenced server will enter a reboot sequence, and when it has restarted it will try to re-join the resource pool.

3.9.2. Configuration Requirements

Note:

Citrix recommends that you enable HA only in pools that contain at least 3 XenServer hosts. For details on how the HA feature behaves when the heartbeat is lost between two hosts in a pool, see the Citrix Knowledge Base article [CTX129721](#).

To use the HA feature, you need:

- Shared storage, including at least one iSCSI, NFS or Fibre Channel LUN of size 356MB or greater- the *heartbeat* SR. The HA mechanism creates two volumes on the heartbeat SR:



4MB heartbeat volume

Used for heartbeating.

256MB metadata volume

Stores pool master metadata to be used in the case of master failover.

Note:

For maximum reliability, Citrix strongly recommends that you use a dedicated NFS or iSCSI storage repository as your HA heartbeat disk, which must not be used for any other purpose.

If you are using a NetApp or EqualLogic SR, manually provision an NFS or iSCSI LUN on the array to use as the heartbeat SR.

- A XenServer pool (this feature provides high availability at the server level within a single resource pool).
- Static IP addresses for all hosts.

Warning:

Should the IP address of a server change while HA is enabled, HA will assume that the host's network has failed, and will probably fence the host and leave it in an unbootable state. To remedy this situation, disable HA using the **host-emergency-ha-disable** command, reset the pool master using **pool-emergency-reset-master**, and then re-enable HA.

For a VM to be protected by the HA feature, it must be agile. This means that:

- it must have its virtual disks on shared storage (any type of shared storage may be used; the iSCSI, NFS or Fibre Channel LUN is only required for the storage heartbeat and can be used for virtual disk storage if you prefer, but this is not necessary)
- it must not have a connection to a local DVD drive configured
- it should have its virtual network interfaces on pool-wide networks.

Citrix strongly recommends the use of a bonded management interface on the servers in the pool if HA is enabled, and multipathed storage for the heartbeat SR.

If you create VLANs and bonded interfaces from the CLI, then they may not be plugged in and active despite being created. In this situation, a VM can appear to be not agile, and cannot be protected by HA. If this occurs, use the CLI **pif-plug** command to bring the VLAN and bond PIFs up so that the VM can become agile. You can also determine precisely why a VM is not agile by using the **xe diagnostic-vm-status** CLI command to analyze its placement constraints, and take remedial action if required.

3.9.3. Restart configuration settings

Virtual machines can be considered protected, best-effort or unprotected by HA. The value of **ha-restart-priority** defines whether a VM is treated as protected, best-effort, or unprotected. The restart behavior for VMs in each of these categories is different.

Restart behavior		ha-restart-priority value
Protected	HA guarantees to restart a protected VM that goes offline or whose host goes offline, provided the pool is not overcommitted and the VM is agile. If a protected VM cannot be restarted at the time of a server failure (for example, if the pool was overcommitted when the failure occurred), further attempts to start this VM are made when extra capacity becomes available in a pool, which might now succeed.	restart

Restart behavior		ha-restart-priority value
Best-effort	If the host where a best-effort VM is running goes offline, HA attempts to restart the best-effort VM on another host only after all protected VMs have been successfully restarted. HA makes only one attempt to restart a best-effort VM and if this attempt fails no further restart is attempted.	best-effort
Unprotected	If an unprotected VM or the host it runs on is stopped, HA does not attempt to restart the VM.	Value is an empty string

Note:

HA never stops or migrates a running VM to free resources for a protected or best-effort VM to be restarted.

If the pool experiences server failures and enters a state where the number of tolerable failures drops to zero, the protected VMs are no longer guaranteed to be restarted. If this condition is reached, a system alert is generated. In this case, if an additional failure occurs, all VMs that have a restart priority set behave according to the best-effort behavior.

3.9.3.1. Start order

The start order is the order in which XenServer HA attempts to restart protected VMs if a failure occurs. The start order is determined by the values of the `order` property for each of the protected VMs.

The `order` property of a VM is used by HA and also by other features that start and shutdown VMs. Any VM can have the `order` property set, not just those marked as protected for HA. However, HA uses the `order` property for protected VMs only.

The value of the `order` property is an integer. The default value is 0, which is the highest priority. Protected VMs with an `order` value of 0 are restarted first by HA. The higher the value of the `order` property, the later in the sequence the VM is restarted.

You can set the value of the `order` property of a VM by using the command-line interface:

```
xe vm-param-set uuid=<VM_UUID> order=<int>
```

Or in XenCenter, in the **Start Options** panel for a VM, set **Start order** to the required value.

3.10. Enabling HA on your XenServer Pool

HA can be enabled on a pool using either XenCenter or the command-line interface. In either case, you will specify a set of priorities that determine which VMs should be given highest restart priority when a pool is overcommitted.

Warning:

When HA is enabled, some operations that would compromise the plan for restarting VMs may be disabled, such as removing a server from a pool. To perform these operations, HA can be temporarily disabled, or alternately, VMs protected by HA made unprotected.

3.10.1. Enabling HA Using the CLI

1. Verify that you have a compatible Storage Repository (SR) attached to your pool. iSCSI, NFS or Fibre Channel are compatible SR types. Please refer to [Section 5.3, "Storage Configuration"](#) for details on how to configure such a storage repository using the CLI.
2. For each VM you wish to protect, set a restart priority and start order. You can do this as follows:


```
xe vm-param-set uuid=<vm_uuid> ha-restart-priority=restart order=<1>
```

3. Enable HA on the pool, and optionally, specify a timeout:

```
xe pool-ha-enable heartbeat-sr-uuids=<sr_uuid> ha-config:timeout=<timeout in seconds>
```

Timeout is the period during which networking or storage is not accessible by the hosts in your pool. If you do not specify a timeout when you enable HA, XenServer will use the default 30 seconds timeout. If any XenServer host is unable to access networking or storage within the timeout period, it will self-fence and restart.

4. Run the **pool-ha-compute-max-host-failures-to-tolerate** command. This command returns the maximum number of hosts that can fail before there are insufficient resources to run all the protected VMs in the pool.

```
xe pool-ha-compute-max-host-failures-to-tolerate
```

The number of failures to tolerate determines when an alert is sent: the system will recompute a failover plan as the state of the pool changes and with this computation the system identifies the capacity of the pool and how many more failures are possible without loss of the liveness guarantee for protected VMs. A system alert is generated when this computed value falls below the specified value for *ha-host-failures-to-tolerate*.

5. Specify the number of failures to tolerate parameter. This should be less than or equal to the computed value:

```
xe pool-param-set ha-host-failures-to-tolerate=<2> uuid=<pool-uuid>
```

3.10.2. Removing HA Protection from a VM using the CLI

To disable HA features for a VM, use the **xe vm-param-set** command to set the *ha-restart-priority* parameter to be an empty string. This does not clear the start order settings. You can enable HA for a VM again by setting the *ha-restart-priority* parameter to *restart* or *best-effort* as appropriate.

3.10.3. Recovering an Unreachable Host

If for some reason a host cannot access the HA statefile, it is possible that a host may become unreachable. To recover your XenServer installation it may be necessary to disable HA using the **host-emergency-ha-disable** command:

```
xe host-emergency-ha-disable --force
```

If the host was the pool master, then it should start up as normal with HA disabled. Slaves should reconnect and automatically disable HA. If the host was a Pool slave and cannot contact the master, then it may be necessary to force the host to reboot as a pool master (**xe pool-emergency-transition-to-master**) or to tell it where the new master is (**xe pool-emergency-reset-master**):

```
xe pool-emergency-transition-to-master uuid=<host_uuid>
xe pool-emergency-reset-master master-address=<new_master_hostname>
```

When all hosts have successfully restarted, re-enable HA:

```
xe pool-ha-enable heartbeat-sr-uuid=<sr_uuid>
```

3.10.4. Shutting Down a host When HA is Enabled

When HA is enabled special care needs to be taken when shutting down or rebooting a host to prevent the HA mechanism from assuming that the host has failed. To shutdown a host cleanly in an HA-enabled environment, first disable the host, then evacuate the host and finally shut down the host using either XenCenter or the CLI. To shutdown a host in an HA-enabled environment on the command line:

```
xe host-disable host=<host_name>
xe host-evacuate uuid=<host_uuid>
xe host-shutdown host=<host_name>
```

3.10.5. Shutting Down a VM When it is Protected by HA

When a VM is protected under a HA plan and set to restart automatically, it cannot be shut down while this protection is active. To shut down a VM, first disable its HA protection and then execute the CLI command. XenCenter offers you a dialog box to automate disabling the protection if you click on the **Shutdown** button of a protected VM.

Note:

If you shut down a VM from within the guest, and the VM is protected, it is automatically restarted under the HA failure conditions. This helps ensure that operator error (or an errant program that mistakenly shuts down the VM) does not result in a protected VM being left shut down accidentally. If you want to shut this VM down, disable its HA protection first.

3.11. Host Power On

3.11.1. Powering on Hosts Remotely

You can use the XenServer Host Power On feature to turn a server on and off remotely, either from XenCenter or by using the CLI.

To enable host power, the host must have one of the following power-control solutions:

- **Wake On LAN enabled network card.**
- **Dell Remote Access Cards (DRAC).** To use XenServer with DRAC, you must install the Dell supplemental pack to get DRAC support. DRAC support requires installing RACADM command-line utility on the server with the remote access controller and enable DRAC and its interface. RACADM is often included in the DRAC management software. For more information, see Dell's DRAC documentation.
- **Hewlett-Packard Integrated Lights-Out (iLO).** To use XenServer with iLO, you must enable iLO on the host and connect interface to the network. For more information, see HP's iLO documentation.
- A custom script based on the XenAPI that enables you to turn the power on and off through XenServer. For more information, see [Section 3.11.3, "Configuring a Custom Script for XenServer's Host Power On Feature"](#).

Using the Host Power On feature requires two tasks:

1. Ensuring the hosts in the pool support controlling the power remotely (that is, they have Wake-on-LAN functionality, a DRAC or iLO card, or you created custom script).
2. Enabling the Host Power On functionality using the CLI or XenCenter.

3.11.2. Using the CLI to Manage Host Power On

You can manage the Host Power On feature using either the CLI or XenCenter. This topic provides information about managing it with the CLI.

Host Power On is enabled at the host level (that is, on each XenServer).

After you enable Host Power On, you can turn hosts on using either the CLI or XenCenter.

3.11.2.1. To Enable Host Power On Using the CLI

1. Run the command:

```
xe host-set-power-on host=<host uuid>\
power-on-mode=("" , "wake-on-lan",
"iLO", "DRAC", "custom")
power-on-config:key=value
```

For iLO and DRAC the keys are *power_on_ip*, *power_on_user*, *power_on_password_secret*. Use *power_on_password_secret* to specify the password if you are using the secret feature.

3.11.2.2. To Turn on Hosts Remotely Using the CLI

1. Run the command:

```
xe host-power-on host=<host uuid>
```

3.11.3. Configuring a Custom Script for XenServer's Host Power On Feature

If your servers' remote-power solution uses a protocol that is not supported by default (such as Wake-On-Ring or Intel Active Management Technology), you can create a custom Linux Python script to turn on your XenServer computers remotely. However, you can also create custom scripts for iLO, DRAC, and Wake-On-LAN remote-power solutions.

This topic provides information about configuring a custom script for Host Power On using the key/value pairs associated with the XenServer API call *host.power_on*.

When you create a custom script, run it from the command line each time you want to control power remotely on XenServer. Alternatively, you can specify it in XenCenter and use the XenCenter UI features to interact with it.

The XenServer API is documented in the document, the [Citrix XenServer Management API], which is available from the Citrix Web site.

Warning:

Do not modify the scripts provided by default in the `/etc/xapi.d/plugins/` directory. You can include new scripts in this directory, but you should never modify the scripts contained in that directory after installation.

3.11.3.1. Key/Value Pairs

To use Host Power On, you must configure the `host.power_on_mode` and `host.power_on_config` keys. Their values are provided below.

There is also an API call that lets you set these fields all at once:

```
void host.set_host_power_on_mode(string mode, Dictionary<string,string> config)
```

3.11.3.1.1. host.power_on_mode

- **Definition:** This contains key/value pairs to specify the type of remote-power solution (for example, Dell DRAC).
- **Possible values:**
 - An empty string, representing power-control disabled
 - "iLO". Lets you specify HP iLO.
 - "DRAC". Lets you specify Dell DRAC. To use DRAC, you must have already installed the Dell supplemental pack.
 - "wake-on-lan". Lets you specify Wake on LAN.
 - Any other name (used to specify a custom power-on script). This option is used to specify a custom script for power management.
- **Type:** string

3.11.3.1.2. host.power_on_config

- **Definition:** This contains key/value pairs for mode configuration. Provides additional information for iLO and DRAC.
- **Possible values:**

- If you configured iLO or DRAC as the type of remote-power solution, you must also specify one of the following keys:
 - "power_on_ip". This is the IP address you specified configured to communicate with the power-control card. Alternatively, you can enter the domain name for the network interface where iLO or DRAC is configured.
 - "power_on_user". This is the iLO or DRAC user name that is associated with the management processor, which you may or may not have changed from its factory default settings.
 - "power_on_password_secret". Specifies using the secrets feature to secure your password.
- To use the secrets feature to store your password, specify the key "power_on_password_secret".
- **Type:** Map (string,string)

3.11.3.2. Sample Script

This sample script imports the XenServer API, defines itself as a custom script, and then passes parameters specific to the host you want to control remotely. You must define the parameters *session*, *remote_host*, and *power_on_config* in all custom scripts.

The result is only displayed when the script is unsuccessful.

```
import XenAPI
def custom(session,remote_host,
power_on_config):
result="Power On Not Successful"
for key in power_on_config.keys():
result=result+"
key="+key+"
value="+power_on_config[key]
return result
```

Note:

After creation, save the script in the /etc/xapi.d/plugins with a .py extension.

3.12. Communicating with XenServer Hosts and Resource Pools

XenServer 7.0 uses TLS protocols to encrypt XenAPI traffic. This means any communication between XenServer and XenAPI clients (or appliances) will now use TLS 1.2 protocol by default. However, if the XenAPI client or the appliance does not communicate using TLS 1.2, earlier protocols may be used for communication.

XenServer uses the following ciphersuites:

- TLS_RSA_WITH_AES_128_CBC_SHA256
- TLS_RSA_WITH_AES_256_CBC_SHA
- TLS_RSA_WITH_AES_128_CBC_SHA
- TLS_RSA_WITH_RC4_128_SHA
- TLS_RSA_WITH_RC4_128_MD5
- TLS_RSA_WITH_3DES_EDE_CBC_SHA

XenServer also enables you to configure your host or resource pool to allow communication through **TLS 1.2 only**. This option allows communication between XenServer and XenAPI clients (or appliances) using the TLS 1.2 protocol. The TLS 1.2 only option uses ciphersuite TLS_RSA_WITH_AES_128_CBC_SHA256.

Warning:

Do not choose the **TLS 1.2 only** option before you ensure that all XenAPI clients and appliances that communicate with the XenServer pool are compatible with TLS 1.2.

Chapter 4. Networking

This chapter provides an overview of XenServer networking, including networks, VLANs, and NIC bonds. It also discusses how to manage your networking configuration and troubleshoot it.

Important:

The XenServer default network stack is the vSwitch; however, you can use the Linux network stack if desired by using the instructions in [Section 4.2, “vSwitch Networks”](#).

If you are already familiar with XenServer networking concepts, you may want to skip ahead to one of the following sections:

- To create networks for standalone XenServer hosts, see [Section 4.4.2, “Creating Networks in a Standalone Server”](#).
- To create private networks across XenServer hosts, see [Section 4.4.1, “Cross-Server Private Networks”](#).
- To create networks for XenServer hosts that are configured in a resource pool, see [Section 4.4.3, “Creating Networks in Resource Pools”](#).
- To create VLANs for XenServer hosts, either standalone or part of a resource pool, see [Section 4.4.4, “Creating VLANs”](#).
- To create bonds for standalone XenServer hosts, see [Section 4.4.5, “Creating NIC Bonds on a Standalone Host”](#).
- To create bonds for XenServer hosts that are configured in a resource pool, see [Section 4.4.6, “Creating NIC Bonds in Resource Pools”](#).

For additional information about networking and network design, see *Designing XenServer Network Configurations* in the Citrix Knowledge Center.

For consistency with XenCenter terminology changes in this release, this chapter now uses the term *management interface* to refer to the IP-enabled NIC that carries the management traffic. In the last release, this chapter used the term *primary management interface* for this purpose. Likewise, this chapter now uses the term *secondary interface* to refer to an IP-enabled NIC configured for storage traffic.

4.1. Networking Support

XenServer supports up to 16 physical network interfaces (or up to 8 bonded network interfaces) per XenServer host and up to 7 virtual network interfaces per VM.

Note:

XenServer provides automated configuration and management of NICs using the xe command line interface (CLI). Unlike previous XenServer versions, the host networking configuration files should not be edited directly in most cases; where a CLI command is available, do not edit the underlying files.

4.2. vSwitch Networks

When used with a controller appliance, vSwitch networks support open flow and provide extra functionality, such as Access Control Lists (ACL). The controller appliance for the XenServer vSwitch is known as the vSwitch Controller: it lets you monitor your networks through a graphical user interface. The vSwitch Controller:

- Supports fine-grained security policies to control the flow of traffic sent to and from a VM.
- Provides detailed visibility into the behavior and performance of all traffic sent in the virtual network environment.



A vSwitch greatly simplifies IT administration in virtualized networking environments—all VM configuration and statistics remain bound to the VM even if it migrates from one physical host in the resource pool to another. See the *XenServer vSwitch Controller User Guide* for more details.

To determine what networking stack is currently configured, run the following command:

```
xe host-list params=software-version
```

In the command output, look for *network_backend*. If the vSwitch is configured as the network stack, the output displays:

```
network_backend: openvswitch
```

If the Linux bridge is configured as the network stack, the output displays:

```
network_backend: bridge
```

Note:

To revert to the Linux network stack, run the following command:

```
xe-switch-network-backend bridge
```

Reboot your host after running this command.

Warning:

The Linux network stack is not open-flow enabled, does not support Cross Server Private Networks, and cannot be managed by the XenServer vSwitch Controller.

4.3. XenServer Networking Overview

This section describes the general concepts of networking in the XenServer environment.

During XenServer installation, one network is created for each physical network interface card (NIC). When you add a server to a resource pool, these default networks are merged so that all physical NICs with the same device name are attached to the same network.

Typically, you would only add a new network if you wanted to create an internal network, set up a new VLAN using an existing NIC, or create a NIC bond.

You can configure four different types of networks in XenServer:

- **External networks** have an association with a physical network interface and provide a bridge between a virtual machine and the physical network interface connected to the network, enabling a virtual machine to connect to resources available through the server's physical network interface card.
- **Bonded networks** create a bond between two NICs to create a single, high-performing channel between the virtual machine and the network.
- **Single-Server Private networks** have no association to a physical network interface and can be used to provide connectivity between the virtual machines on a given host, with no connection to the outside world.
- **Cross-Server Private networks** extend the single server private network concept to allow VMs on different hosts to communicate with each other by using the vSwitch.

Note:

Some networking options have different behaviors when used with standalone XenServer hosts compared to resource pools. This chapter contains sections on general information that applies to both standalone hosts and pools, followed by specific information and procedures for each.

4.3.1. Network Objects

This chapter uses three types of server-side software objects to represent networking entities. These objects are:

- A *PIF*, which represents a physical NIC on a XenServer host. PIF objects have a name and description, a globally unique UUID, the parameters of the NIC that they represent, and the network and server they are connected to.
- A *VIF*, which represents a virtual NIC on a virtual machine. VIF objects have a name and description, a globally unique UUID, and the network and VM they are connected to.
- A *network*, which is a virtual Ethernet switch on a XenServer host. Network objects have a name and description, a globally unique UUID, and the collection of VIFs and PIFs connected to them.

Both XenCenter and the xe CLI allow configuration of networking options, control over which NIC is used for management operations, and creation of advanced networking features such as virtual local area networks (VLANs) and NIC bonds.

4.3.2. Networks

Each XenServer host has one or more networks, which are virtual Ethernet switches. Networks that are not associated with a PIF are considered *internal* and can be used to provide connectivity only between VMs on a given XenServer host, with no connection to the outside world. Networks associated with a PIF are considered *external* and provide a bridge between VIFs and the PIF connected to the network, enabling connectivity to resources available through the PIF's NIC.

4.3.3. VLANs

Virtual Local Area Networks (VLANs), as defined by the IEEE 802.1Q standard, allow a single physical network to support multiple logical networks. XenServer hosts can work with VLANs in multiple ways.

Note:

All supported VLAN configurations are equally applicable to pools and standalone hosts, and bonded and non-bonded configurations.

4.3.3.1. Using VLANs with Management Interfaces

Switch ports configured to perform 802.1Q VLAN tagging/untagging, commonly referred to as ports with a *native VLAN* or as *access mode* ports, can be used with management interfaces to place management traffic on a desired VLAN. In this case the XenServer host is unaware of any VLAN configuration.

Management interfaces cannot be assigned to a XenServer VLAN via a trunk port.

4.3.3.2. Using VLANs with Virtual Machines

Switch ports configured as 802.1Q VLAN trunk ports can be used in combination with the XenServer VLAN features to connect guest virtual network interfaces (VIFs) to specific VLANs. In this case, the XenServer host performs the VLAN tagging/untagging functions for the guest, which is unaware of any VLAN configuration.

XenServer VLANs are represented by additional PIF objects representing VLAN interfaces corresponding to a specified VLAN tag. XenServer networks can then be connected to the PIF representing the physical NIC to see all traffic on the NIC, or to a PIF representing a VLAN to see only the traffic with the specified VLAN tag.

For procedures on how to create VLANs for XenServer hosts, either standalone or part of a resource pool, see [Section 4.4.4, "Creating VLANs"](#).

4.3.3.3. Using VLANs with Dedicated Storage NICs

Dedicated storage NICs (also known as IP-enabled NICs or simply management interfaces) can be configured to use native VLAN / access mode ports as described above for management interfaces, or with trunk ports and XenServer VLANs as described above for virtual machines. To configure dedicated storage NICs, see [Section 4.4.7, "Configuring a Dedicated Storage NIC"](#).

4.3.3.4. Combining Management Interfaces and Guest VLANs on a Single Host NIC

A single switch port can be configured with both trunk and native VLANs, allowing one host NIC to be used for a management interface (on the native VLAN) and for connecting guest VIFs to specific VLAN IDs.

4.3.4. Jumbo frames

Jumbo frames can be used to optimize the performance of storage traffic. Jumbo frames are Ethernet frames containing more than 1500 bytes of payload. Jumbo frames are typically used to achieve better throughput, reducing the load on system bus memory, and reducing the CPU overhead.

Note:

XenServer supports jumbo frames only when using vSwitch as the network stack on all hosts in the pool.

Requirements for using jumbo frames

Customers should note the following when using jumbo frames:

- Jumbo frames are configured at a pool level
- vSwitch must be configured as the network backend on all hosts in the pool
- Every device on the subnet must be configured to use jumbo frames
- It is recommended that customers only enable jumbo frames on a dedicated storage network
- Enabling jumbo frames on the Management network is not a supported configuration
- Jumbo frames are not supported for use on VMs

In order to use jumbo frames, customers should set the Maximum Transmission Unit (MTU) to a value between 1500 and 9216. This can be done either by using XenCenter or the xe CLI. For more information about configuring networks with jumbo frames, see *Designing XenServer Network Configurations* in the Citrix Knowledge Center.

4.3.5. NIC Bonds


NIC bonds, sometimes also known as *NIC teaming*, improve XenServer host resiliency and/or bandwidth by enabling administrators to configure two or more NICs together so they logically function as one network card. All bonded NICs share the same MAC address.

If one NIC in the bond fails, the host's network traffic is automatically redirected through the second NIC. XenServer supports up to eight bonded networks.

XenServer provides support for active-active, active-passive, and LACP bonding modes. The number of NICs supported and the bonding mode supported varies according to network stack:

- LACP bonding is only available for the vSwitch whereas active-active and active-passive are available for both the vSwitch and Linux bridge.
- When the vSwitch is the network stack, you can bond either two, three, or four NICs.
- When the Linux bridge is the network stack, you can only bond two NICs.

In the illustration that follows, the management interface is on a bonded pair of NICs. XenServer will use this bond for management traffic.



This illustration shows a host with a management interface on a bond and two pairs of NICs bonded for guest traffic. Excluding the management interface bond, XenServer uses the other two NIC bonds and the two non-bonded NICs for VM traffic.

All bonding modes support failover; however, not all modes allow all links to be active for all traffic types. XenServer supports bonding the following types of NICs together:

- **NICs (non-management).** You can bond NICs that XenServer is using solely for VM traffic. Bonding these NICs not only provides resiliency, but doing so also balances the traffic from multiple VMs between the NICs.
- **Management interfaces.** You can bond a management interface to another NIC so that the second NIC provides failover for management traffic. Although configuring a LACP link aggregation bond provides load balancing for management traffic, active-active NIC bonding does not.
- **Secondary interfaces.** You can bond NICs that you have configured as secondary interfaces (for example, for storage). However, for most iSCSI software initiator storage, Citrix recommends configuring multipathing instead of NIC bonding as described in the *Designing XenServer Network Configurations*.

Through out this section, the term *IP-based storage traffic* is used to refer to iSCSI and NFS traffic collectively.

You can create a bond if a VIF is already using one of the interfaces that will be bonded: the VM traffic will be automatically migrated to the new bonded interface.

In XenServer, the NIC bond is represented by an additional PIF. XenServer NIC bonds completely subsume the underlying physical devices (PIFs).

Note:

Creating a bond that contains only one NIC is not supported.

Key points about IP addressing

Bonded NICs will either have one IP address or no IP addresses, as follows:

- **Management and storage networks.**
 - If you bond a management interface or secondary interface, a single IP address is assigned to the bond. That is, each NIC does not have its own IP address; XenServer treats the two NICs as one logical connection.
 - When bonds are used for non-VM traffic (to connect to shared network storage or XenCenter for management), you must configure an IP address for the bond. However, if you have already assigned an IP address to one of the NICs (that is, created a management interface or secondary interface), that IP address is assigned to the entire bond automatically.
 - If you bond a management interface or secondary interface to a NIC without an IP address, as of XenServer 6.0, the bond assumes the IP address of the respective interface automatically.
- **VM networks.** When bonded NICs are used for VM (guest) traffic, you do not need to configure an IP address for the bond. This is because the bond operates at Layer 2 of the OSI model, the data link layer, and no IP addressing is used at this layer. IP addresses for virtual machines are associated with VIFs.

Bonding types

XenServer provides three different types of bonds, all of which can be configured using either the CLI or XenCenter:

- Active/Active mode, with VM traffic balanced between the bonded NICs. See [Section 4.3.5.1, “Active-Active Bonding”](#).
- Active/Passive mode, where only one NIC actively carries traffic. See [Section 4.3.5.2, “Active-Passive Bonding”](#).
- LACP Link Aggregation, in which active and stand-by NICs are negotiated between the switch and the server. See [Section 4.3.5.3, “LACP Link Aggregation Control Protocol Bonding”](#).

Note:

Bonding is set up with an Up Delay of 31000ms and a Down Delay of 200ms. The seemingly long Up Delay is deliberate because of the time some switches take to actually enable the port. Without a delay, when a link comes back after failing, the bond could rebalance traffic onto it before the switch is ready to pass traffic. If you want to move both connections to a different switch, move one, then wait 31 seconds for it to be used again before moving the other. For information about changing the delay, see [Section 4.3.8, "Changing the Up Delay for Bonds"](#).

Bond Status

XenServer provides status for bonds in the event logs for each host. If one or more links in a bond fails or is restored, it is noted in the event log. Likewise, you can query the status of a bond's links by using the `links-up` parameter as shown in the following example:

```
xe bond-param-get uuid=<bond_uuid> param-name=links-up
```


XenServer checks the status of links in bonds approximately every 5 seconds. Consequently, if additional links in the bond fail in the five-second window, the failure is not logged until the next status check.

Bonding event logs appear in the XenCenter **Logs** tab. For users not running XenCenter, event logs also appear in `/var/log/xensource.log` on each host.

4.3.5.1. Active-Active Bonding

Active-active is an active/active configuration for guest traffic: both NICs can route VM traffic simultaneously. When bonds are used for management traffic, only one NIC in the bond can route traffic: the other NIC remains unused and provides fail-over support. Active-active mode is the default bonding mode when either the Linux bridge or vSwitch network stack is enabled.

When active-active bonding is used with the Linux bridge, you can only bond two NICs. When using the vSwitch as the network stack, you can bond either two, three, or four NICs in active-active mode. However, in active-active mode, bonding three or four NICs is generally only beneficial for VM traffic, as shown in the illustration that follows.



This illustration shows how bonding four NICs may only benefit guest traffic. In the top picture of a management network, NIC 2 is active but NICs 1, 3, and 4 are passive. For the VM traffic, all four NICs in the bond are active; however, this assumes a minimum of four VMs. For the storage traffic, only NIC 11 is active.

XenServer can only send traffic over two or more NICs when there is more than one MAC address associated with the bond. XenServer can use the virtual MAC addresses in the VIF to send traffic across multiple links. Specifically:

- **VM traffic.** Provided you enable bonding on NICs carrying only VM (guest) traffic, all links are active and NIC bonding can balance spread VM traffic across NICs. An individual VIF's traffic is never split between NICs.
- **Management or storage traffic.** Only one of the links (NICs) in the bond is active and the other NICs remain unused unless traffic fails over to them. Configuring a management interface or secondary interface on a bonded network provides resilience.
- **Mixed traffic.** If the bonded NIC carries a mixture of IP-based storage traffic and guest traffic, only the guest and control domain traffic are load balanced. The control domain is essentially a virtual machine so it uses a NIC like the other guests. XenServer balances the control domain's traffic the same way as it balances VM traffic.

Traffic Balancing

XenServer balances the traffic between NICs by using the source MAC address of the packet. Because, for management traffic, only one source MAC address is present, active-active mode can only use one NIC and traffic is not balanced. Traffic balancing is based on two factors:

- The virtual machine and its associated VIF sending or receiving the traffic
- The quantity of data (in kilobytes) being sent.

XenServer evaluates the quantity of data (in kilobytes) each NIC is sending and receiving. If the quantity of data sent across one NIC exceeds the quantity of data sent across the other NIC, XenServer rebalances which VIFs use which NICs. The VIF's entire load is transferred; one VIF's load is never split between two NICs.

While active-active NIC bonding can provide load balancing for traffic from multiple VMs, it cannot provide a single VM with the throughput of two NICs. Any given VIF only uses one of the links in a bond at a time. As XenServer periodically rebalances traffic, VIFs are not permanently assigned to a specific NIC in the bond.

Active-active mode is sometimes referred to as Source Load Balancing (SLB) bonding as XenServer uses SLB to share load across bonded network interfaces. SLB is derived from the open-source Adaptive Load Balancing (ALB) mode and reuses the ALB capability to dynamically re-balance load across NICs.

When rebalancing, the number of bytes going over each slave (interface) is tracked over a given period. When a packet to be sent contains a new source MAC address, it is assigned to the slave interface with the lowest utilisation. Traffic is re-balanced at regular intervals.

Each MAC address has a corresponding load and XenServer can shift entire loads between NICs depending on the quantity of data a VM sends and receives. For active-active traffic, all the traffic from one VM can be sent on only one NIC.

Note:


Active-active bonding does not require switch support for EtherChannel or 802.3ad (LACP).

4.3.5.2. Active-Passive Bonding

An active-passive bond routes traffic over only one of the NICs, so traffic fails over to the other NIC in the bond if the active NIC loses network connectivity. Active-passive bonds route traffic over the active NIC: only if the active NIC fails is traffic shifted to the passive NIC.

Active-passive bonding is available in the Linux bridge and the vSwitch network stack. When used with the Linux bridge, you can bond two NICs together. When used with the vSwitch, you can only bond two, three, or four NICs together. However, regardless of the traffic type, when you bond NICs together in active-passive mode, only one link is active and there is no load balancing between links.

The illustration that follows shows two bonded NICs configured in active-passive mode.



This illustration shows two NICs bonded in active-passive mode. NIC 1 is active. The bond includes a NIC for failover that is connected to a second switch. This NIC will be used only if NIC 1 fails.

Since active-active mode is the default bonding configuration in XenServer, if you are configuring bonds using the CLI, you must specify a parameter for the active-passive mode or the bond is created as active-active. However, you do not need to configure active-passive mode just because a network is carrying management traffic or storage traffic.

Active-passive can be a good choice for resiliency since it offers several benefits. With active-passive bonds, traffic does not move around between NICs. Likewise, active-passive bonding lets you configure two switches for redundancy but does not require stacking. (If the management switch dies, stacked switches can be a single point of failure.)

Active-passive mode does not require switch support for EtherChannel or 802.3ad (LACP).

Consider configuring active-passive mode in situations when you do not need load balancing or when you only intend to send traffic on one NIC.

Important:

After you have created VIFs or your pool is in production, be extremely careful about making changes to bonds or creating new bonds.

4.3.5.3. LACP Link Aggregation Control Protocol Bonding

LACP Link Aggregation Control Protocol is a type of bonding that bundles a group of ports together and treats it like a single logical channel. LACP bonding provides failover and can increase the total amount of bandwidth available.

Unlike other bonding modes, LACP bonding requires configuring both sides of the links: creating a bond on the host and, on the switch, creating a Link Aggregation Group (LAG) for each bond, as described in [Section 4.3.5.4.1, “Switch Configuration for LACP Bonds”](#). To use LACP bonding, you must configure the vSwitch as the network stack. Also, your switches must support the IEEE 802.3ad standard.

The following table compares active-active SLB bonding and LACP bonding:

	Benefits	Considerations
Active-active SLB Bonding	- Can be used with any switch on the XenServer Hardware Compatibility List. - Does not require switches that support stacking. - Supports four NICs.	- Optimal load balancing requires at least one NIC per VIF. - Storage or management traffic cannot be split on multiple NICs. - Load balancing occurs only if multiple MAC addresses are present.

	Benefits	Considerations
LACP bonding	- All links can be active regardless of traffic type. - Traffic balancing does not depend on source MAC addresses, so all traffic types can be balanced.	- Switches must support the IEEE 802.3ad standard. - Requires switch-side configuration. - Supported only for the vSwitch. - Requires a single switch or stacked switch.

Traffic Balancing

XenServer supports two LACP bonding *hashing types* —the term *hashing* refers to the way in the NICs and the switch distribute the traffic— (1) load balancing based on IP and port of source and destination addresses and (2) load balancing based on source MAC address.

Depending on the hashing type and traffic pattern, LACP bonding can potentially distribute traffic more evenly than active-active NIC bonding.

Note:


You configure settings for outgoing and incoming traffic separately on the host and the switch: the configuration does not have to match on both sides.

Load balancing based on IP and port of source and destination addresses.

This hashing type is the default LACP bonding hashing algorithm. Traffic coming from one guest can be distributed over two links provided that there is a variation in the source or destination IP or port numbers.


If one virtual machine is running several applications, which use different IP or port numbers, this hashing type distributes traffic over several links giving the guest the possibility of using the aggregate throughput. This hashing type lets one guest use the whole throughput of multiple NICs.

Likewise, as shown in the illustration that follows, this hashing type can distribute the traffic of two different applications on a virtual machine to two different NICs.



This illustration shows how, if you use LACP bonding and enable LACP with load balancing based on IP and port of source and destination as the hashing type, the traffic coming from two different applications on VM1 can be distributed to two NICs.

Configuring LACP bonding based on IP and port of source and destination address is beneficial when you want to balance the traffic of two different applications on the same VM (for example, when only one virtual machine is configured to use a bond of three NICs).




This illustration shows how, if you use LACP bonding and enable LACP with load balancing based on IP and port of source and destination as the hashing type, XenServer can send the traffic of each application in the virtual machine through one of the three NICs in the bond even though the number of NICs exceeds the number of VIFs.

The balancing algorithm for this hashing type uses five factors to spread traffic across the NICs: the source IP address, source port number, destination IP address, destination port number, and source MAC address.

Load balancing based on source MAC address.

This type of load balancing works well when there are multiple virtual machines on the same host. Traffic is balanced based on the virtual MAC address of the VM from which the traffic originated. XenServer sends outgoing traffic using the same algorithm as it does in the case of active-active bonding. Traffic coming from the same guest is not split over multiple NICs. As a result, this hashing type is not suitable if there are fewer VIFs than NICs: load balancing is not optimal because the traffic cannot be split across NICs.



This illustration shows how, if you use LACP bonding and enable LACP based on source MAC address as the hashing type, if the number of NICs exceeds the number of VIFs, not all NICs will be used. Because there are three NICs and only two VMs, only two NICs can be used at the same time and the maximum bond throughput cannot be achieved. The packets from one VM cannot be split across multiple VMs.

4.3.5.4. Switch Configuration

Depending on your redundancy requirements, you can connect the NICs in the bond to either the same or separate stacked switches. If you connect one of the NICs to a second, redundant switch and a NIC or switch fails, traffic fails over to the other NIC. Adding a second switch prevents a single point-of-failure in your configuration in the following ways:

- When you connect one of the links in a bonded management interface to a second switch, if the switch fails, the management network still remains online and the hosts can still communicate with each other.
- If you connect a link (for any traffic type) to a second switch and the NIC or switch fails, the virtual machines remain on the network since their traffic fails over to the other NIC/switch.


When you want to connect bonded NICs to multiple switches and you configured the LACP bonding mode, you must use stacked switches. The term *stacked switches* refers to configuring multiple physical switches to function as a single logical switch. You must join the switches together physically and through the switch-management software so the switches function as a single logical switching unit, as per the switch manufacturer's guidelines.

Typically, switch stacking is only available through proprietary extensions and switch vendors may market this functionality under different terms.

Note:

If you experience issues with active-active bonds, the use of stacked switches might be necessary. Active-passive bonds do not require stacked switches.

The illustration that follows shows how the cables and network configuration for the bonded NICs have to match.



This illustration shows how two NICs in a bonded pair use the same network settings, as represented by the networks in each host. The NICs in the bonds connect to different switches for redundancy.

4.3.5.4.1. Switch Configuration for LACP Bonds

While the specific details of switch configuration varies by manufacturer, there are a few key points to remember when configuring switches for use with LACP bonds:

- The switch must support LACP and the IEEE 802.3ad standard.
- When you create the LAG group on the switch, you must create one LAG group for each LACP bond on the host. This means if you have a five-host pool and you created a LACP bond on NICs 4 and 5 on each host, you must create five LAG groups on the switch. One group for each set of ports corresponding with the NICs on the host.

You may also need to add your VLAN ID to your LAG group.

- XenServer LACP bonds require setting the Static Mode setting in the LAG group to be set to **Disabled**.

As previously mentioned in [Section 4.3.5.4, "Switch Configuration"](#), stacking switches are required to connect LACP bonds to multiple switches.

4.3.6. Initial Networking Configuration after Setup

The XenServer host networking configuration is specified during initial host installation. Options such as IP address configuration (DHCP/static), the NIC used as the management interface, and hostname are set based on the values provided during installation.

When a host has multiple NICs the configuration present after installation depends on which NIC is selected for management operations during installation:

- PIFs are created for each NIC in the host
- the PIF of the NIC selected for use as the management interface is configured with the IP addressing options specified during installation
- a network is created for each PIF ("network 0", "network 1", etc.)
- each network is connected to one PIF
- the IP addressing options of all other PIFs are left unconfigured

When a XenServer host has a single NIC, the follow configuration is present after installation:

- a single PIF is created corresponding to the host's single NIC
- the PIF is configured with the IP addressing options specified during installation and to enable management of the host
- the PIF is set for use in host management operations
- a single network, network 0, is created
- network 0 is connected to the PIF to enable external connectivity to VMs

In both cases the resulting networking configuration allows connection to the XenServer host by XenCenter, the xe CLI, and any other management software running on separate machines via the IP address of the management interface. The configuration also provides external networking for VMs created on the host.

The PIF used for management operations is the only PIF ever configured with an IP address during XenServer installation. External networking for VMs is achieved by bridging PIFs to VIFs using the network object which acts as a virtual Ethernet switch.

The steps required for networking features such as VLANs, NIC bonds, and dedicating a NIC to storage traffic are covered in the sections that follow.

4.3.7. Changing Networking Configuration

You can change your networking configuration by modifying the network object. To do so, you run a command that affects either the network object or the VIF.

4.3.7.1. Modifying the Network Object

You can modify aspects of an network, such as the frame size (MTU), name-label, name-description, and other values by using the xe **network-param-set** command and its associated parameters.

When you run the xe **network-param-set** command, the only required parameter is *uuid*.

Optional parameters include:

- default_locking_mode. See [Section 4.4.10.8.10, “Simplifying VIF Locking Mode Configuration in the Cloud”](#).
- name-label
- name-description
- MTU
- other-config:

If a value for a parameter is not given, the parameter is set to a null value. To set a (key,value) pair in a map parameter, use the syntax 'map-param:key=value'.

4.3.8. Changing the Up Delay for Bonds

As described in [Section 4.3.5, “NIC Bonds”](#), by default, bonding is set up with an Up Delay of 31000ms to prevent traffic from being rebalanced onto a NIC after it fails. While seemingly long, the up delay is important for all bonding modes and not just active-active.

However, if you understand the appropriate settings to select for your environment, you can change the up delay for bonds by using the procedure that follows.

To change the up delay for bonds

1. Set the up delay in milliseconds:


```
xe pif-param-set uuid=<<uuid of bond master PIF>> other-config:bond-updelay=<<delay in ms>>
```

2. To make the change take effect, you must unplug and then replug the physical interface:

```
xe pif-unplug uuid=<<uuid of bond master PIF>>
```

```
xe pif-plug uuid=<<uuid of bond master PIF>>
```

4.4. Managing Networking Configuration

Some of the network configuration procedures in this section differ depending on whether you are configuring a stand-alone server or a server that is part of a resource pool.

4.4.1. Cross-Server Private Networks

Previous versions of XenServer allowed you to create single-server private networks that allowed VMs running on the same host to communicate with each other. The *cross-server private network* feature, which extends the single-server private network concept to allow VMs on different hosts to communicate with each other. Cross-server private networks combine the same isolation properties of a single-server private network but with the additional ability to span hosts across a resource pool. This combination enables use of VM agility features such as XenMotion live migration for VMs with connections to cross-server private networks.

Cross-server private networks are completely isolated. VMs that are not connected to the private network cannot sniff or inject traffic into the network, even when they are located on the same physical host with VIFs connected to a network on the same underlying physical network device (PIF). VLANs provide similar functionality, though unlike VLANs, cross-server private networks provide isolation without requiring configuration of a physical switch fabric, through the use of the Generic Routing Encapsulation (GRE) IP tunnelling protocol.

Private networks provide the following benefits without requiring a physical switch:

- the isolation properties of single-server private networks
- the ability to span a resource pool, enabling VMs connected to a private network to live on multiple hosts within the same pool
- compatibility with features such as XenMotion

Cross-Server Private Networks must be created on a management interface or a secondary interface, as they require an IP addressable NIC. Any IP-enabled NIC can be used as the underlying network transport. If you choose to put cross-server private network traffic on a secondary interface, this secondary interface **must** be on a separate subnet.

If any management or secondary interfaces are on the same subnet, traffic will be routed incorrectly.

Note:

To create a cross-server private network, the following conditions must be met:

- All of the hosts in the pool must be using XenServer 6.0 or greater.
- All of the hosts in the pool must be using the vSwitch for the networking stack.
- The vSwitch Controller must be running and you must have added the pool to it. (The pool must have a vSwitch Controller configured that handles the initialization and configuration tasks required for the vSwitch connection.)
- The cross-server private network must be created on a NIC configured as a management interface. This can be the management interface or a secondary interface (IP-enabled PIF) you configure specifically for this purpose, provided it is on a separate subnet.

For more information on configuring the vSwitch, see the *XenServer vSwitch Controller User Guide*. For UI-based procedures for configuring private networks, see the XenCenter Help.

4.4.2. Creating Networks in a Standalone Server

Because external networks are created for each PIF during host installation, creating additional networks is typically only required to:

- use a private network
- support advanced operations such as VLANs or NIC bonding

To add or remove networks using XenCenter, refer to the XenCenter online Help.

To add a new network using the CLI

1. Open the XenServer host text console.
2. Create the network with the `network-create` command, which returns the UUID of the newly created network:

```
xe network-create name-label=<mynetwork>
```

At this point the network is not connected to a PIF and therefore is internal.

4.4.3. Creating Networks in Resource Pools

All XenServer hosts in a resource pool should have the same number of physical network interface cards (NICs), although this requirement is not strictly enforced when a XenServer host is joined to a pool.

Having the same physical networking configuration for XenServer hosts within a pool is important because all hosts in a pool share a common set of XenServer networks. PIFs on the individual hosts are connected to pool-wide networks based on device name. For example, all XenServer hosts in a pool with an `eth0` NIC will have a corresponding PIF plugged into the pool-wide `Network 0` network. The same will be true for hosts with `eth1` NICs and `Network 1`, as well as other NICs present in at least one XenServer host in the pool.

If one XenServer host has a different number of NICs than other hosts in the pool, complications can arise because not all pool networks will be valid for all pool hosts. For example, if hosts *host1* and *host2* are in the same pool and *host1* has four NICs while *host2* only has two, only the networks connected to PIFs corresponding to `eth0` and `eth1` will be valid on *host2*. VMs on *host1* with VIFs connected to networks corresponding to `eth2` and `eth3` will not be able to migrate to host *host2*.

4.4.4. Creating VLANs

For servers in a resource pool, you can use the **pool-vlan-create** command. This command creates the VLAN and automatically creates and plugs in the required PIFs on the hosts in the pool. See [Section A.4.25.2, “pool-vlan-create”](#) for more information.

To connect a network to an external VLAN using the CLI

1. Open the XenServer host console.
2. Create a new network for use with the VLAN. The UUID of the new network is returned:

```
xe network-create name-label=network5
```

3. Use the **pif-list** command to find the UUID of the PIF corresponding to the physical NIC supporting the desired VLAN tag. The UUIDs and device names of all PIFs are returned, including any existing VLANs:

```
xe pif-list
```

4. Create a VLAN object specifying the desired physical PIF and VLAN tag on all VMs to be connected to the new VLAN. A new PIF will be created and plugged into the specified network. The UUID of the new PIF object is returned.

```
xe vlan-create network-uuid=<network_uuid> pif-uuid=<pif_uuid> vlan=5
```

5. Attach VM VIFs to the new network. See [Section 4.4.2, “Creating Networks in a Standalone Server”](#) for more details.

4.4.5. Creating NIC Bonds on a Standalone Host

Citrix recommends using XenCenter to create NIC bonds. For instructions, see the XenCenter help.

This section describes how to use the `xe` CLI to bond NIC interfaces on a XenServer host that is not in a pool. See [Section 4.4.6, “Creating NIC Bonds in Resource Pools”](#) for details on using the `xe` CLI to create NIC bonds on XenServer hosts that comprise a resource pool.

4.4.5.1. Creating a NIC Bond

When you bond a NIC, the bond absorbs the PIF/NIC currently in use as the management interface. From XenServer 6.0 onwards, the management interface is automatically moved to the bond PIF.

To bond two or four NICs

1. Use the **network-create** command to create a new network for use with the bonded NIC. The UUID of the new network is returned:

```
xe network-create name-label=<bond0>
```

2. Use the **pif-list** command to determine the UUIDs of the PIFs to use in the bond:

```
xe pif-list
```

3. Do one of the following:

- To configure the bond in active-active mode (default), use the **bond-create** command to create the bond. Using commas to separate the parameters, specify the newly created network UUID and the UUIDs of the PIFs to be bonded:

```
xe bond-create network-uuid=<network_uuid> pif-uuids=<pif_uuid_1>,<pif_uuid_2>,<pif_uuid_3>
```

Enter two UUIDs if you are bonding two NICs and four UUIDs if you are bonding four NICs. The UUID for the bond is returned after running the command.

- To configure the bond in active-passive or LACP bond mode, use the same syntax but add the optional *mode* parameter and specify *lACP* or *active-backup*:

```
xe bond-create network-uuid=<network_uuid> pif-uuids=<pif_uuid_1>,<pif_uuid_2>,<pif_uuid_3>  
mode=<balance-slb | active-backup | lACP>
```

Note:

In previous releases, you specified the *other-config:bond-mode* to change the bond mode. While this command still works, it may not be supported in future releases and it is not as efficient as the *mode* parameter. *other-config:bond-mode* requires running **pif-unplug** and **pif-plug** to get the mode change to take effect.

4.4.5.2. Controlling the MAC Address of the Bond

When you bond the management interface, the PIF/NIC currently in use as the management interface is subsumed by the bond. If the host uses DHCP, in most cases the bond's MAC address is the same as the PIF/NIC currently in use, and the management interface's IP address can remain unchanged.

You can change the bond's MAC address so that it is different from the MAC address for the (current) management-interface NIC. However, as the bond is enabled and the MAC/IP address in use changes, existing network sessions to the host will be dropped.

You can control the MAC address for a bond in two ways:

- An optional *mac* parameter can be specified in the **bond-create** command. You can use this parameter to set the bond MAC address to any arbitrary address.
- If the *mac* parameter is not specified, from XenServer 7.0 onwards, XenServer uses the MAC address of the management interface if this is one of the interfaces in the bond. If the management interface is not part of the bond, but another management interface is, the bond uses the MAC address (and also the IP address) that management interface. If none of the NICs in the bond are management interfaces, the bond uses the MAC of the first named NIC.

4.4.5.3. Reverting NIC Bonds

If reverting a XenServer host to a non-bonded configuration, be aware that the **bond-destroy** command automatically configures the *primary-slave* as the interface to be used for the management interface. Consequently, all VIFs will be moved to the management interface.

The term primary-slave refers to the PIF that the MAC and IP configuration was copied from when creating the bond. When bonding two NICs, the primary slave is:

1. The management interface NIC (if the management interface is one of the bonded NICs).
2. Any other NIC with an IP address (if the management interface was not part of the bond).
3. The first named NIC. You can find out which one it is by running the following:

```
xe bond-list params=all
```

4.4.6. Creating NIC Bonds in Resource Pools

Whenever possible, create NIC bonds as part of initial resource pool creation prior to joining additional hosts to the pool or creating VMs. Doing so allows the bond configuration to be automatically replicated to hosts as they are joined to the pool and reduces the number of steps required. Adding a NIC bond to an existing pool requires one of the following:

- Using the CLI to configure the bonds on the master and then each member of the pool.
- Using the CLI to configure the bonds on the master and then restarting each member of the pool so that it inherits its settings from the pool master.
- Using XenCenter to configure the bonds on the master. XenCenter automatically synchronizes the networking settings on the member servers with the master, so you do not need to reboot the member servers.

For simplicity and to prevent misconfiguration, Citrix recommends using XenCenter to create NIC bonds. For details, refer to the XenCenter Help.

This section describes using the *xe* CLI to create bonded NIC interfaces on XenServer hosts that comprise a resource pool. See [Section 4.4.5.1, “Creating a NIC Bond”](#) for details on using the *xe* CLI to create NIC bonds on a standalone XenServer host.

Warning:

Do not attempt to create network bonds while HA is enabled. The process of bond creation will disturb the in-progress HA heartbeating and cause hosts to self-fence (shut themselves down); subsequently they will likely fail to reboot properly and will need the **host-emergency-ha-disable** command to recover.

4.4.6.1. Adding NIC Bonds to New Resource Pools

1. Select the host you want to be the master. The master host belongs to an unnamed pool by default. To create a resource pool with the CLI, rename the existing nameless pool:

```
xe pool-param-set name-label=<"New Pool"> uuid=<pool_uuid>
```

2. Create the NIC bond as described in [Section 4.4.5.1, “Creating a NIC Bond”](#).

3. Open a console on a host that you want to join to the pool and run the command:

```
xe pool-join master-address=<host1> master-username=root master-password=<password>
```

The network and bond information is automatically replicated to the new host. The management interface is automatically moved from the host NIC where it was originally configured to the bonded PIF (that is, the management interface is now absorbed into the bond so that the entire bond functions as the management interface).

- Use the **host-list** command to find the UUID of the host being configured:

```
xe host-list
```

4.4.6.2. Adding NIC Bonds to an Existing Pool

Warning:

Do not attempt to create network bonds while HA is enabled. The process of bond creation disturbs the in-progress HA heartbeating and causes hosts to self-fence (shut themselves down); subsequently they will likely fail to reboot properly and you will need to run the **host-emergency-ha-disable** command to recover them.

Note:

If you are not using XenCenter for NIC bonding, the quickest way to create pool-wide NIC bonds is to create the bond on the master, and then restart the other pool members. Alternatively, you can use the **service xapi restart** command. This causes the bond and VLAN settings on the master to be inherited by each host. The management interface of each host must, however, be manually reconfigured.

Follow the procedure in previous sections to create a NIC Bond, see [Section 4.4.6.1, “Adding NIC Bonds to New Resource Pools”](#).

4.4.7. Configuring a Dedicated Storage NIC

You can use either XenCenter or the xe CLI to assign a NIC an IP address and dedicate it to a specific function, such as storage traffic. When you configure a NIC with an IP address, you do so by creating a secondary interface. (The IP-enabled NIC XenServer used for management is known as the management interface.)

When you want to dedicate a secondary interface for a specific purpose, you must ensure the appropriate network configuration is in place to ensure the NIC is used only for the desired traffic. For example, to dedicate a NIC to storage traffic, the NIC, storage target, switch, and/or VLAN must be configured so that the target is only accessible over the assigned NIC. If your physical and IP configuration do not limit the traffic that can be sent across the storage NIC, it is possible to send other traffic, such as management traffic, across the secondary interface.

When you create a new secondary interface for storage traffic, you must assign it an IP address that (a) is on the same subnet as the storage controller, if applicable, and (b) is not on the same subnet as any other secondary interfaces or the management interface.

When you are configuring secondary interfaces, each secondary interface must be on a separate subnet. For example, if you want to configure two additional secondary interfaces for storage, you will require IP addresses on three different subnets – one subnet for the management interface, one subnet for Secondary Interface 1, and one subnet for Secondary Interface 2.

If you are using bonding for resiliency for your storage traffic, you may want to consider using LACP instead of the Linux bridge bonding. To use LACP bonding, you must configure the vSwitch as your networking stack. For more information, see [Section 4.2, “vSwitch Networks”](#).

Note:

When selecting a NIC to configure as a secondary interface for use with iSCSI or NFS SRs, ensure that the dedicated NIC uses a separate IP subnet that is not routable from the management interface. If this is not enforced, then storage traffic may be directed over the main management interface after a host reboot, due to the order in which network interfaces are initialized.

To assign NIC functions using the xe CLI

1. Ensure that the PIF is on a separate subnet, or routing is configured to suit your network topology in order to force the desired traffic over the selected PIF.
2. Setup an IP configuration for the PIF, adding appropriate values for the mode parameter and if using static IP addressing the IP, netmask, gateway, and DNS parameters:

```
xe pif-reconfigure-ip mode=<DHCP | Static> uuid=<pif-uuid>
```

3. Set the PIF's disallow-unplug parameter to true:

```
xe pif-param-set disallow-unplug=true uuid=<pif-uuid>
```

```
xe pif-param-set other-config:management_purpose="Storage" uuid=<pif-uuid>
```

If you want to use a secondary interface for storage that can be routed from the management interface also (bearing in mind that this configuration is not the best practice), you have two options:

- After a host reboot, ensure that the secondary interface is correctly configured, and use the **xe pbd-unplug** and **xe pbd-plug** commands to reinitialize the storage connections on the host. This restarts the storage connection and routes it over the correct interface.
- Alternatively, you can use **xe pif-forget** to remove the interface from the XenServer database and manually configure it in the control domain. This is an advanced option and requires you to be familiar with how to manually configure Linux networking.

4.4.8. Using SR-IOV Enabled NICs

Single Root I/O Virtualization (SR-IOV) is a PCI device virtualization technology that allows a single PCI device to appear as multiple PCI devices on the physical PCI bus. The actual physical device is known as a **Physical Function** (PF) while the others are known as **Virtual Functions** (VF). The purpose of this is for the hypervisor to directly assign one or more of these VFs to a Virtual Machine (VM) using SR-IOV technology: the guest can then use the VF as any other directly assigned PCI device.

Assigning one or more VFs to a VM allows the VM to directly exploit the hardware. When configured, each VM behaves as though it is using the NIC directly, reducing processing overhead and improving performance.

Warning:

If your VM has an SR-IOV VF, functions that require VM mobility, for example, Live Migration, Rolling Pool Upgrade, High Availability and Disaster Recovery, are not possible. This is because the VM is directly tied to the physical SR-IOV enabled NIC VF. In addition, VM network traffic sent via an SR-IOV VF bypasses the vSwitch, so it is not possible to create ACLs or view QoS.

Assigning a SR-IOV NIC VF to a VM

Note:

SR-IOV is supported only with SR-IOV enabled NICs listed on the XenServer [Hardware Compatibility List](#) and only when used in conjunction with a Windows Server 2008 guest operating system.

1. Open a local command shell on your XenServer host.
2. Run the command `lspci` to display a list of the Virtual Functions (VF). For example:

07:10.0 Ethernet controller: Intel Corporation 82559 \\
Ethernet Controller Virtual Function (rev 01)

In the example above, 07:10.0 is the **bus:device.function** address of the VF.

- Assign the required VF to the target VM by running the following commands:

```
xe vm-param-set other-config:pci=0/0000:<bus:device.function> uuid=<vm-uuid>
```

- Start the VM, and install the appropriate VF driver for your specific hardware.

Note:

You can assign multiple VFs to a single VM, however the same VF cannot be shared across multiple VMs.

4.4.9. Controlling the Rate of Outgoing Data (QoS)

To limit the amount of *outgoing* data a VM can send per second, you can set an optional Quality of Service (QoS) value on VM virtual interfaces (VIFs). The setting lets you specify a maximum transmit rate for outgoing packets in kilobytes per second.

The QoS value limits the rate of transmission *from* the VM. The QoS setting does not limit the amount of data the VM can receive. If such a limit is desired, Citrix recommends limiting the rate of incoming packets higher up in the network (for example, at the switch level).

Depending on networking stack configured in the pool, you can set the Quality of Service (QoS) value on VM virtual interfaces (VIFs) in one of two places—either a) on the vSwitch Controller or b) in XenServer (using the CLI or XenCenter)—as described in the following table:

Networking Stack	Configuration Methods Available
vSwitch	vSwitch Controller. This is the preferred method of setting the maximum transmission rate on a VIF when the vSwitch is the networking stack. When using the vSwitch stack, the XenCenter QoS option is not available. xe commands. It is possible to set the QoS transmit rate using the commands in the example that follows. However, the preferred method is through the vSwitch Controller UI, which provides more finely grained control.
Linux bridge	XenCenter. You can set the QoS transmit rate limit value in the properties dialog for the virtual interface. xe commands. You can set the QoS transmit rate using the CLI using the commands in the section that follow.

Important:

When the vSwitch is configured as the networking stack, it is possible to inadvertently configure a QoS value on the vSwitch Controller *and* inside of the XenServer host. In this case, XenServer limits the outgoing traffic using the lowest rate that you set.

Example of CLI command for QoS:

To limit a VIF to a maximum transmit rate of 100 kilobytes per second using the CLI, use the **vif-param-set** command:



```
xe vif-param-set uuid=<vif_uuid> qos_algorithm_type=ratelimit
xe vif-param-set uuid=<vif_uuid> qos_algorithm_params:kbps=100
```

Note:

If you are using the vSwitch Controller, Citrix recommends setting the transmission rate limit in the vSwitch Controller instead of this CLI command. For directions on setting the QoS rate limit in the vSwitch Controller, see the *vSwitch Controller User Guide*.

4.4.10. Changing Networking Configuration Options

This section discusses how to change the networking configuration of a XenServer host. This includes:

- changing the hostname (that is, the Domain Name System (DNS) name)
- adding or removing DNS servers
- changing IP addresses
- changing which NIC is used as the management interface
- adding a new physical NIC to the server
- enabling ARP filtering (switch-port locking)

4.4.10.1. Hostname

The system hostname, also known as the domain or DNS name, is defined in the pool-wide database and modified using the **xe host-set-hostname-live** CLI command as follows:

```
xe host-set-hostname-live host-uuid=<host_uuid> host-name=<host-name>
```

The underlying control domain hostname changes dynamically to reflect the new hostname.

4.4.10.2. DNS Servers

To add or remove DNS servers in the IP addressing configuration of a XenServer host, use the **pif-reconfigure-ip** command. For example, for a PIF with a static IP:

```
pif-reconfigure-ip uuid=<pif_uuid> mode=static DNS=<new_dns_ip>
```

4.4.10.3. Changing IP Address Configuration for a Standalone Host

Network interface configuration can be changed using the xe CLI. The underlying network configuration scripts should not be modified directly.

To modify the IP address configuration of a PIF, use the **pif-reconfigure-ip** CLI command. See [Section A.4.14.5, “pif-reconfigure-ip”](#) for details on the parameters of the **pif-reconfigure-ip** command.

Note:

See [Section 4.4.10.4, “Changing IP Address Configuration in Resource Pools”](#) for details on changing host IP addresses in resource pools.

4.4.10.4. Changing IP Address Configuration in Resource Pools

XenServer hosts in resource pools have a single management IP address used for management and communication to and from other hosts in the pool. The steps required to change the IP address of a host's management interface are different for master and other hosts.

Note:

Caution should be used when changing the IP address of a server, and other networking parameters. Depending upon the network topology and the change being made, connections to network storage may be lost. If this happens the storage must be replugged using the

Repair Storage function in XenCenter, or the **pbd-plug** command using the CLI. For this reason, it may be advisable to migrate VMs away from the server before changing its IP configuration.

To change the IP address of a member host (not pool master)

1. Use the **pif-reconfigure-ip** CLI command to set the IP address as desired. See [Appendix A, Command Line Interface](#) for details on the parameters of the **pif-reconfigure-ip** command:

```
xe pif-reconfigure-ip uuid=<pif_uuid> mode=DHCP
```

2. Use the **host-list** CLI command to confirm that the member host has successfully reconnected to the master host by checking that all the other XenServer hosts in the pool are visible:

```
xe host-list
```

Changing the IP address of the master XenServer host requires additional steps because each of the member hosts uses the advertised IP address of the pool master for communication and will not know how to contact the master when its IP address changes.

Whenever possible, use a dedicated IP address that is not likely to change for the lifetime of the pool for pool masters.

To change the IP address of the pool master

1. Use the **pif-reconfigure-ip** CLI command to set the IP address as desired. See [Appendix A, Command Line Interface](#) for details on the parameters of the **pif-reconfigure-ip** command:

```
xe pif-reconfigure-ip uuid=<pif_uuid> mode=DHCP
```

2. When the IP address of the pool master host is changed, all member hosts will enter into an emergency mode when they fail to contact the master host.
3. On the master XenServer host, use the **pool-recover-slaves** command to force the master to contact each of the member hosts and inform them of the new master IP address:

```
xe pool-recover-slaves
```

4.4.10.5. Management Interface

When XenServer is installed on a host with multiple NICs, one NIC is selected for use as the management interface. The management interface is used for XenCenter connections to the host and for host-to-host communication.

To change the NIC used for the management interface

1. Use the **pif-list** command to determine which PIF corresponds to the NIC to be used as the management interface. The UUID of each PIF is returned.

```
xe pif-list
```

2. Use the **pif-param-list** command to verify the IP addressing configuration for the PIF that will be used for the management interface. If necessary, use the **pif-reconfigure-ip** command to configure IP addressing for the PIF to be used. See [Appendix A, Command Line Interface](#) for more detail on the options available for the **pif-reconfigure-ip** command.

```
xe pif-param-list uuid=<pif_uuid>
```

3. Use the **host-management-reconfigure** CLI command to change the PIF used for the management interface. If this host is part of a resource pool, *this command must be issued on the member host console*:

```
xe host-management-reconfigure pif-uuid=<pif_uuid>
```

Warning:

Putting the management interface on a VLAN network is not supported.

4.4.10.6. Disabling Management Access

To disable remote access to the management console entirely, use the **host-management-disable** CLI command.

Warning:

Once the management interface is disabled, you will have to log in on the physical host console to perform management tasks and external interfaces such as XenCenter will no longer work.

4.4.10.7. Adding a New Physical NIC

Install a new physical NIC on a XenServer host in the usual manner. Then, after restarting the server, run the xe CLI command **pif-scan** to cause a new PIF object to be created for the new NIC.

4.4.10.8. Using Switch Port Locking

The XenServer switch-port locking feature lets you control traffic being sent from unknown, untrusted, or potentially hostile VMs by limiting their ability to pretend they have a MAC or IP address that was not assigned to them. You can use the port-locking commands in this feature to block all traffic on a network by default or define specific IP addresses from which an individual VM is allowed to send traffic.

Switch-port locking is a feature designed for public cloud-service providers in environments concerned about internal threats. This functionality may help public cloud-service providers who have a network architecture in which each VM has a public, Internet-connected IP address. Because cloud tenants are always untrusted, it may be desirable to use security measures, such as spoofing protection, to ensure tenants cannot attack other virtual machines in the cloud.

Using switch-port locking lets you simplify your network configuration by enabling all of your tenants or guests to use the same Layer 2 network.

One of the most important functions of the port-locking commands is they can restrict the traffic that an untrusted guest can send, which, in turn, restricts the guest's ability to pretend it has a MAC or IP address it does not actually possess. Specifically, you can use these commands to prevent a guest from:

- Claiming an IP or MAC address other than the ones the XenServer administrator has specified it can use
- Intercepting, spoofing, or disrupting the traffic of other VMs

4.4.10.8.1. Requirements

- The XenServer switch-port locking feature is supported on the Linux bridge and vSwitch networking stacks.
- When Role Based Access Control (RBAC) is enabled in your environment, the user configuring switch-port locking must be logged in with an account that has at least a Pool Operator or Pool Admin role. When RBAC is not enabled in your environment, the user must be logged in with the root account for the pool master.
- When you run the switch-port locking commands, networks can be online or offline.
- In Windows guests, the disconnected Network icon only appears when XenServer Tools are installed in the guest.

4.4.10.8.2. Notes

Without any switch-port locking configurations, VIFs are set to "network_default" and Networks are set to "unlocked."

Configuring switch-port locking is not supported when the vSwitch controller and other third-party controllers are in use in the environment.

Switch port locking does not prevent cloud tenants from:

- Performing an IP-level attack on another tenant/user. However, switch-port locking prevents them performing the IP-level attack if they attempt to use the following means to do so and switch-port locking

is configured: a) impersonating another tenant in the cloud or user or b) initiating an intercept of traffic intended for another user.

- Exhausting network resources.
- Receiving some traffic intended for other virtual machines through normal switch flooding behaviors (for broadcast MAC addresses or unknown destination MAC addresses).

Likewise, switch-port locking does not restrict where a VM can send traffic to.

4.4.10.8.3. Implementation Notes

You can implement the switch-port locking functionality either by using the command line or the XenServer API. However, in large environments, where automation is a primary concern, the most typical implementation method might be by using the API.

4.4.10.8.4. Examples

This section provides examples of how switch-port locking can prevent certain types of attacks. In these examples, VM-c is a virtual machine that a hostile tenant (Tenant C) is leasing and using for attacks. VM-a and VM-b are virtual machines leased by non-attacking tenants.

Example 1: How Switch Port Locking Can Prevent ARP Spoofing Prevention

ARP spoofing refers to an attacker's attempts to associate his or her MAC address with the IP address for another node, which could potentially result in the node's traffic being sent to the attacker instead. To achieve this goal the attacker sends fake (*spoofed*) ARP messages to an Ethernet LAN.

Scenario:

Virtual Machine A (VM-a) wants to send IP traffic from VM-a to Virtual Machine B (VM-b) by addressing it to VM-b's IP address. The owner of Virtual Machine C wants to use ARP spoofing to pretend his VM, VM-c, is actually VM-b.

1. VM-c sends a speculative stream of ARP replies to VM-a. These ARP replies claim that the MAC address in the reply (c_MAC) is associated with the IP address, b_IP

Result: Because the administrator enabled switch-port locking, these packets are all dropped because enabling switch-port locking prevents impersonation.

2. VM-b sends an ARP reply to VM-a, claiming that the MAC address in the reply (b_MAC) is associated with the IP address, b_IP.

Result: VM-a receives VM-b's ARP response.

Example 2: IP Spoofing Prevention

IP address spoofing is a process that conceals the identity of packets by creating Internet Protocol (IP) packets with a forged source IP address.

Scenario:

Tenant C is attempting to perform a Denial of Service attack using his host, Host-C, on a remote system to disguise his identity.

Attempt 1

Tenant C sets Host-C's IP address and MAC address to VM-a's IP and MAC addresses (a_IP and a_MAC). Tenant C instructs Host-C to send IP traffic to a remote system.

Result: The Host-C packets are dropped. This is because the administrator enabled switch-port locking; the Host-C packets are dropped because enabling switch-port locking prevents impersonation.

Attempt 2

Tenant C sets Host-C's IP address to VM-a's IP address (a_IP) and keeps his original c_MAC.

Tenant C instructs Host-C to send IP traffic to a remote system.

Result: The Host-C packets are dropped. This is because the administrator enabled switch-port locking, which prevents impersonation.

Example 3: Web Hosting

Scenario:

Alice is an infrastructure administrator.

One of her tenants, Tenant B, is hosting multiple websites from his VM, VM-b. Each website needs a distinct IP address hosted on the same virtual network interface (VIF).

Alice reconfigures Host-B's VIF to be locked to a single MAC but many IP addresses.

4.4.10.8.5. How Switch-port Locking Works


The switch-port locking feature lets you control packet filtering at one or more of two levels:

- **VIF level.** Settings you configure on the VIF determine how packets are filtered. You can set the VIF to prevent the VM from sending any traffic, restrict the VIF so it can only send traffic using its assigned IP address, or allow the VM to send traffic to any IP address on the network connected to the VIF.
- **Network level.** The XenServer network determines how packets are filtered. When a VIF's locking mode is set to *network_default*, it refers to the network-level locking setting to determine what traffic to allow.

Regardless of which networking stack you use, the feature operates the same way. However, as described in more detail in the sections that follow, the Linux bridge does not fully support switch-port locking in IPv6.

4.4.10.8.6. VIF Locking-Mode States

The XenServer switch-port locking feature provides a locking mode that lets you configure VIFs in four different states. These states only apply when the VIF is plugged into a running virtual machine.



This illustration shows how three different VIF locking mode states behave when the network locking mode is set to unlocked and the VIF state is configured. In the first image, the VIF state is set to default so no traffic from the VM is filtered. In the second image, the VIF does not send or receive any packets because the locking mode is set to disabled. In the third image, the VIF state is set to locked, so the VIF can only send packets if those packets contain the correct MAC and IP address.

- **Network_default.** When the VIF's state is set to *network_default*, XenServer uses the network's *default-locking-mode* parameter to determine if and how to filter packets travelling through the VIF. The behavior varies according to if the associated network has the network default locking mode parameter set to disabled or unlocked:
 - *default-locking-mode=disabled*, XenServer applies a filtering rule so that the VIF drops all traffic.
 - *default-locking-mode=unlocked*, XenServer removes all the filtering rules associated with the VIF. By default, the default locking mode parameter is set to *unlocked*.

For information about the *default-locking-mode* parameter, see [Section A.4.11, “Network Commands”](#).

The default locking mode of the network has no effect on attached VIFs whose locking state is anything other than `network_default`.

Note:

You cannot change the *default-locking-mode* of a network that has active VIFs attached to it.

- **Locked.** XenServer applies filtering rules so that only traffic sent to or sent from the specified MAC and IP addresses is allowed to be sent out through the VIF. In this mode, if no IP addresses are specified, then the VM cannot send any traffic through that VIF (on that network).

To specify the IP addresses from which the VIF will accept traffic, use the IPv4 and/or IPv6 IP addresses by using the *ipv4_allowed* or *ipv6_allowed* parameters. However, if you have the Linux bridge configured, do not enter IPv6 addresses.

XenServer lets you enter IPv6 addresses when the Linux bridge is active; however, XenServer cannot filter based on the IPv6 addresses entered. (The reason is the Linux bridge does not have modules to filter Neighbor Discovery Protocol (NDP) packets, so complete protection cannot be implemented and guests would be able to impersonate another guest by forging NDP packets.) As result, if you specify even one IPv6 address, XenServer lets all IPv6 traffic pass through the VIF. If you do not specify any IPv6 addresses, XenServer will not let any IPv6 traffic pass through to the VIF.

- **Unlocked.** All network traffic can pass through the VIF. That is, no filters are applied to any traffic going to or from the VIF.
- **Disabled.** No traffic is allowed to pass through the VIF. (That is, XenServer applies a filtering rule so that the VIF drops all traffic.)

4.4.10.8.7. Configuring Switch Port Locking

This section provides three different procedures:

- Restrict VIFs to use a specific IP address
- Add an IP address to an existing restricted list (for example, if you need to add an IP address to a VIF while the VM is still running and connected to the network (for example, if you are taking a network offline temporarily))
- Remove an IP address from an existing restricted list

If a VIF's locking-mode is set to `locked`, it can only use the addresses specified in the *ipv4_allowed* or *ipv6_allowed* parameters.

Because, in some relatively rare cases, VIFs may have more than one IP address, it is possible to specify multiple IP addresses for a VIF.

You can perform these procedures before or after the VIF is plugged in (or the VM is started).

To restrict VIFs to a specific IP address

1. Change the default-locking mode to `locked`, if it is not using that mode already, by running the following command:

```
xe vif-param-set uuid=<vif-uuid> locking-mode=locked
```

The `vif-uuid` represents the UUID of the VIF you want to allow to send traffic. To obtain the UUID, run the `xe vif-list` command on the host. The `vm-uuid` indicates virtual machine for which the information is displayed. The device ID indicates the device number of the VIF.

2. Run the **vif-param-set** command to specify the IP addresses from which the virtual machine can send traffic. Do one or more of the following:

- Specify one or more IPv4 IP addresses destinations. For example:

```
xe vif-param-set uuid=<vif-uuid> ipv4-allowed=<comma separated list of ipv4-addresses>
```

- Specify one or more IPv6 IP addresses destinations. For example:

```
xe vif-param-set uuid=<vif-uuid> ipv6-allowed=<comma separated list of ipv6-addresses>
```

You can specify multiple IP addresses by separating them with a comma, as shown in the preceding example.

To restrict a VIF to an additional IP address

After performing the previous procedure to restrict a VIF to using a specific IP address, you can add one or more IP addresses the VIF can use.

- Run the **vif-param-add** command to add the IP addresses to the existing list. Do one or more of the following:

- Specify the IPv4 IP address. For example:

```
xe vif-param-add uuid=<vif-uuid> ipv4-allowed=<comma separated list of ipv4-addresses>
```

- Specify the IPv6 IP address. For example:

```
xe vif-param-add uuid=<vif-uuid> ipv6-allowed=<comma separated list of ipv6-addresses>
```

To remove an IP address from a VIF's list of IP addresses

If you restrict a VIF to use two or more IP addresses, you can remove one of those IP addresses from the list.

- Run the **vif-param-remove** command to remove the IP addresses from the existing list. Do one or more of the following:

- Specify the IPv4 IP address to remove. For example:

```
xe vif-param-remove uuid=<vif-uuid> ipv4-allowed=<comma separated list of ipv4-addresses>
```

- Specify the IPv6 IP address to remove. For example:

```
xe vif-param-remove uuid=<vif-uuid> ipv6-allowed=<comma separated list of ipv6-addresses>
```

4.4.10.8.8. Preventing a Virtual Machine from Sending or Receiving Traffic from a Specific Network

The following procedure prevents a virtual machine from communicating through a specific VIF. Since a VIF connects to a specific XenServer network, you can use this procedure to prevent a virtual machine from sending or receiving any traffic from a specific network. This provides a more granular level of control than disabling an entire network.

If you use the CLI command, you do not need to unplug the VIF to set the VIF's locking mode; the command changes the filtering rules while the VIF is running (live). In this case the network connection still appears to be present; however, the VIF drops any packets the VM attempts to send.

Tip:

To find the UUID of a VIF, run the **xe vif-list** command on the host. The device ID indicates the device number of the VIF.

To prevent a VIF from receiving traffic

- Disable the VIF connected to the network from which you want to stop the VM from receiving traffic:

```
xe vif-param-set uuid=<vif-uuid> locking-mode=disabled
```

You can also disable the VIF in XenCenter by selecting the virtual network interface in the VM's **Networking** tab and clicking **Deactivate**.

4.4.10.8.9. Removing a VIF's Restriction to an IP Address

To revert back to the default (original) locking mode state, use the following procedure. By default, when you create a VIF, XenServer configures it so that it is not restricted to using a specific IP address.

To revert a VIF to an unlocked state

- Change the VIF default-locking mode to unlocked, if it is not using that mode already, by running the following command:

```
xe vif-param-set uuid=<vif-uuid> locking-mode=unlocked
```

4.4.10.8.10. Simplifying VIF Locking Mode Configuration in the Cloud

Rather than running the VIF locking mode commands for each VIF, you can ensure all VIFs are disabled by default. To do so, you must modify the packet filtering at the network level, which causes the XenServer network to determine how packets are filtered, as described in [Section 4.4.10.8.5, "How Switch-port Locking Works"](#).


Specifically, a network's *default-locking-mode* setting determines how new VIFs with default settings behave. Whenever a VIF's *locking-mode* is set to default, the VIF refers to the network-locking mode (*default-locking-mode*) to determine if and how to filter packets travelling through the VIF:

- **Unlocked.** When the network *default-locking-mode* parameter is set to unlocked, XenServer lets the VM send traffic to any IP address on the network the VIF connects to.
- **Disabled.** When the default-locking-mode parameter is set to disabled, XenServer applies a filtering rule so that the VIF drops all traffic.

By default, the *default-locking-mode* for all networks created in XenCenter and using the CLI are set to unlocked.

By setting the VIF's locking mode to its default (*network_default*), you can use this setting to create a basic default configuration (at the network level) for all newly created VIFs that connect to a specific network.

This illustration shows how, when a VIF's *locking-mode* is set to its default setting (*network_default*), the VIF uses the network *default-locking-mode* to determine its behavior.



This illustration shows how a VIF, when configured at its default setting (locking-mode=network_default), checks to see the setting associated with the default-locking-mode. In this illustration, the network is set to default-locking-mode=disabled so no traffic can pass through the VIF.

For example, since, by default, VIFs are created with their *locking-mode* set to *network_default*, if you set a network's *default-locking-mode*=disabled any new VIFs for which you have not configured the locking mode are disabled until you either (a) change the individual VIF's *locking-mode* parameter or (b) explicitly set the VIF's *locking-mode* to unlocked (for example, if you trust a specific VM enough to not want to filter its traffic at all).

To change a network's default locking mode setting

- After creating the network, change the default-locking mode by running the following command:

```
xe network-param-set uuid=<network-uuid> default-locking-mode=[unlocked|disabled]
```

Note:

To get the UUID for a network, run the `xe network-list` command. This command displays the UUIDs for all the networks on the host on which you ran the command.

To check a network's default locking mode setting

- Run one of the following commands:

```
xe network-param-get uuid=<network-uuid> param-name=default-locking-mode
```

OR

```
xe network-list uuid=<network-uuid> params=default-locking-mode
```

4.4.10.8.11. Using Network Settings for VIF Traffic Filtering

The following procedure instructs a VIF on a virtual machine to use the XenServer network *default-locking-mode* settings on the network itself to determine how to filter traffic.

To use network settings for VIF traffic filtering

1. Change the VIF locking state to *network_default*, if it is not using that mode already, by running the following command:

```
xe vif-param-set uuid=<vif-uuid> locking-mode=network_default
```

2. Change the default-locking mode to *unlocked*, if it is not using that mode already, by running the following command:

```
xe network-param-set uuid=<network-uuid> default-locking-mode=unlocked
```

4.5. Networking Troubleshooting

If you are having problems with configuring networking, first ensure that you have not directly modified any of the control domain `ifcfg-*` files directly. These files are directly managed by the control domain host agent, and changes will be overwritten.

4.5.1. Diagnosing Network Corruption

Some network card models require firmware upgrades from the vendor to work reliably under load, or when certain optimizations are turned on. If you are seeing corrupted traffic to VMs, then you should first try to obtain the latest recommended firmware from your vendor and apply a BIOS update.

If the problem still persists, then you can use the CLI to disable receive / transmit offload optimizations on the physical interface.

Warning:

Disabling receive / transmit offload optimizations can result in a performance loss and / or increased CPU usage.

First, determine the UUID of the physical interface. You can filter on the *device* field as follows:

```
xe pif-list device=eth0
```

Next, set the following parameter on the PIF to disable TX offload:


```
xe pif-param-set uuid=<piif_uuid> other-config:ethtool-tx=off
```

Finally, re-plug the PIF or reboot the host for the change to take effect.

4.5.2. Emergency Network Reset

Incorrect networking settings can cause loss of network connectivity, and a XenServer host may become inaccessible via XenCenter or remote SSH. Emergency Network Reset provides a simple mechanism to recover and reset a host's networking.

This feature is available from the Command Line Interface (CLI) using the `xe - reset - networking` command and within the **Network and Management Interface** section of **xsconsole**.

Incorrect settings which could cause a loss of network connectivity could include renaming network interfaces, creating bonds or VLANs, or mistakes when changing the management interface (for example, entering the wrong IP address). In addition, you may want to run this utility if a rolling pool upgrade, manual upgrade, hotfix installation or driver installation causes a lack of network connectivity, or if a Pool master or host in a resource pool is unable to contact with other hosts.

This utility **should only be used in an emergency** as it will remove the configuration for all PIFs, Bonds, VLANs and tunnels associated with the host. Guest Networks and VIFs are preserved. As part of this utility, VMs will be shutdown forcefully, where possible before running this command, VMs should be cleanly shutdown. Before applying a reset, users can make changes to the management interface and specify which IP configuration, DHCP or Static, should be used.

If the Pool Master requires a network reset, it must be carried out **before** a network reset of any other pool members. It should then be followed a network reset on all remaining hosts in the pool to ensure that the pool's networking configuration is homogeneous. This is a particularly important factor for XenMotion.

Note:

If the Pool Master's IP address (the management interface) changes, as a result of a network reset or `xe host-management-reconfigure`, you must also apply the network reset command to other hosts in the pool, so that they can reconnect to the Pool Master on its new IP address. In this situation, the IP address of the Pool Master must be specified.

Network reset is **NOT** supported if High Availability (HA) is enabled. To reset network configuration in this scenario, you must first **manually disable HA**, and then run the network reset command.

4.5.2.1. Verifying the Network Reset

After specifying the configuration mode to be used after the network reset, **xsconsole** and the CLI will display the settings which will be applied after host reboot. This offers a final chance to make any modifications before applying the emergency network reset command. After reboot, the new network configuration can be verified in XenCenter and **xsconsole**. In XenCenter, with the host selected, click the **Networking** tab, this displays the new network configuration. In **xsconsole**, this information is displayed in the **Network and Management Interface** section.

Note:

Emergency Network Reset should also be applied on other pool members to replicate bonds, VLANs or tunnels from the Pool Master's new configuration.

4.5.2.2. Using the CLI for Network Reset

The following table shows the available optional parameters which can be used with the `xe-reset-networking` command.

Warning:



Users are responsible for ensuring the validity of parameters for the **xe-reset-networking** command, check the parameters carefully. If invalid parameters are specified, network connectivity and configuration will be lost. In this situation, Citrix advises customers to re-run the command **xe-reset-networking** without using **any** parameters.

Resetting the networking configuration of a whole pool **must** begin on the Pool Master, and should then be followed by network reset on all remaining hosts in the pool.

Parameter	Required/Optional	Description
-m, --master	Optional	IP address of the Pool Master's management interface. Defaults to the last known Pool Master's IP address.
--device	Optional	Device name of the management interface. Defaults to the device name specified during installation.
--mode=static	Optional	Enables the following four networking parameters for static IP configuration for the management interface. If not specified, networking will be configured using DHCP.
--ip	Required if mode=static	IP address for the host's management interface. Only valid if mode=static.
--netmask	Required if mode=static	Netmask for the management interface. Only valid if mode=static.
--gateway	Optional	Gateway for the management interface. Only valid if mode=static.
--dns	Optional	DNS Server for the management interface. Only valid if mode=static.

4.5.2.2.1. Pool Master Command Line Examples

Examples of commands that could be applied on a Pool Master:

To reset networking for DHCP configuration:

```
xe-reset-networking
```

To reset networking for Static IP configuration:

```
xe-reset-networking --mode=static --ip=<ip-address> \
--netmask=<netmask> --gateway=<gateway> \
--dns=<dns>
```

To reset networking for DHCP configuration if another interface became the management interface after initial setup:

```
xe-reset-networking --device=<device-name>
```

To reset networking for Static IP configuration if another interface became the management interface after initial setup:

```
xe-reset-networking --device=<device-name> --mode=static \
--ip=<ip-address> --netmask=<netmask> \
--gateway=<gateway> --dns=<dns>
```

4.5.2.2.2. Pool Member Command Line Examples

All previous examples also apply to pool members. Additionally the Pool Master's IP address can be specified (which will be necessary if it has changed.)

To reset networking for DHCP configuration:

```
xe-reset-networking
```

To reset networking for DHCP if the Pool Master's IP address was modified:



```
xe-reset-networking --master=<master-ip-address>
```

To reset networking for Static IP configuration, assuming the Pool Master's IP address didn't change:

```
xe-reset-networking --mode=static --ip=<ip-address> --netmask-<netmask> \  
--gateway=<gateway> --dns=<dns>
```

To reset networking for DHCP configuration if the management interface and the Pool Master's IP address was modified after initial setup:

```
xe-reset-networking --device=<device-name> --master<master-ip-address>
```

Chapter 5. Storage

This chapter describes the way in which physical storage hardware is mapped to Virtual Machines (VMs), and the software objects used by the XenServer host API to perform storage-related tasks. Detailed sections on each of the supported storage types include procedures for creating storage for VMs using the CLI, with type-specific device configuration options, generating snapshots for backup purposes and some best practices for managing storage in XenServer host environments. Finally, the virtual disk QoS (quality of service) settings are described.

5.1. Storage Overview

This section explains what the XenServer storage objects are and how they are related to each other.

5.1.1. Storage Repositories (SRs)

A Storage Repository (SR) is a particular storage target, in which Virtual Machine (VM) Virtual Disk Images (VDIs) are stored. A VDI is a storage abstraction which represent a virtual hard disk drive (HDD).

SRs are flexible, with built-in support for IDE, SATA, SCSI and SAS drives that are locally connected, and iSCSI, NFS, SAS and Fibre Channel remotely connected. The SR and VDI abstractions allow for advanced storage features such as *Thin Provisioning*, VDI snapshots, and fast cloning to be exposed on storage targets that support them. For storage subsystems that do not inherently support advanced operations directly, a software stack is provided based on Microsoft's Virtual Hard Disk (VHD) specification which implements these features.

SR commands provide operations for creating, destroying, resizing, cloning, connecting and discovering the individual VDIs that they contain.

A storage repository is a persistent, on-disk data structure. For SR types that use an underlying block device, the process of creating a new SR involves erasing any existing data on the specified storage target. Other storage types such as NFS, create a new container on the storage array in parallel to existing SRs.

Each XenServer host can use multiple SRs and different SR types simultaneously. These SRs can be shared between hosts or dedicated to particular hosts. Shared storage is pooled between multiple hosts within a defined resource pool. A shared SR must be network accessible to each host. All hosts in a single resource pool must have at least one shared SR in common.

CLI operations to manage storage repositories are described in [Section A.4.17, "SR Commands"](#).

5.1.2. Virtual Disk Image (VDI)

Virtual Disk Images (VDI) are a storage abstraction which represent a virtual hard disk drive (HDD). VDIs are the fundamental unit of virtualized storage in XenServer. VDIs are persistent, on-disk objects that exist independently of XenServer hosts. CLI operations to manage VDIs are described in [Section A.4.23, "VDI Commands"](#). The actual on-disk representation of the data differs by SR type and is managed by a separate storage plug-in interface for each SR, called the SM API.

5.1.3. Physical Block Devices (PBDs)


Physical Block Devices represent the interface between a physical server and an attached SR. PBDs are connector objects that allow a given SR to be mapped to a XenServer host. PBDs store the device configuration fields that are used to connect to and interact with a given storage target. For example, NFS device configuration includes the IP address of the NFS server and the associated path that the XenServer host mounts. PBD objects manage the run-time attachment of a given SR to a given XenServer host. CLI operations relating to PBDs are described in [Section A.4.13, "PBD Commands"](#).

5.1.4. Virtual Block Devices (VBDs)

Virtual Block Devices are connector objects (similar to the PBD described above) that allows mappings between VDIs and VMs. In addition to providing a mechanism for attaching (also called plugging) a VDI into a VM, VBDs allow for the fine-tuning of parameters regarding QoS (quality of service), statistics, and the bootability of a given VDI. CLI operations relating to VBDs are described in [Section A.4.22, “VBD Commands”](#).

5.1.5. Summary of Storage objects

The following image is a summary of how the storage objects presented so far are related:



Graphical overview of storage repositories and related objects

5.1.6. Virtual Disk Data Formats

In general, there are two types of mapping of physical storage to a VDI:

1. *Logical Volume-based VHD on a LUN*; The default XenServer blockdevice-based storage inserts a Logical Volume manager on a disk, either a locally attached device (LVM type SR) or a SAN attached LUN over either Fibre Channel (LVMoHBA type SR), iSCSI (LVMoISCSI type SR) or SAS (LVMoHBA type Sr). VDIs are represented as volumes within the Volume manager and stored in VHD format to allow thin provisioning of reference nodes on snapshot and clone.
2. *File-based VHD on a filesystem*; VM images are stored as thin-provisioned VHD format files on either a local non-shared filesystem (EXT type SR) or a shared NFS target (NFS type SR)

5.1.6.1. VDI Types

In general, VHD format VDIs will be created. You can opt to use raw at the time you create the VDI; this can only be done using the xe CLI. .

To check if a VDI was created with `type=raw`, check its `sm-config` map. The `sr-param-list` and `vdv-param-list` xe commands can be used respectively for this purpose.

5.1.6.2. Creating a Raw Virtual Disk Using the xe CLI

1. Run the following command to create a VDI given the UUID of the SR you want to place the virtual disk in:

```
xe vdi-create sr-uuid=<sr-uuid> type=user virtual-size=<virtual-size> \
  name-label=<VDI name> sm-config:type=raw
```

2. Attach the new virtual disk to a VM and use your normal disk tools within the VM to partition and format, or otherwise make use of the new disk. You can use the **vbd-create** command to create a new VBD to map the virtual disk into your VM.

5.1.6.3. Converting Between VDI Formats

It is not possible to do a direct conversion between the raw and VHD formats. Instead, you can create a new VDI (either raw, as described above, or VHD) and then copy data into it from an existing volume. Citrix recommends that you use the xe CLI to ensure that the new VDI has a virtual size at least as big as the VDI you are copying from (by checking its virtual-size field, for example by using the **vd-param-list** command). You can then attach this new VDI to a VM and use your preferred tool within the VM (standard disk management tools in Windows, or the **dd** command in Linux) to do a direct block-copy of the data. If the new volume is a VHD volume, it is important to use a tool that can avoid writing empty sectors to the disk so that space is used optimally in the underlying storage repository — in this case a file-based copy approach may be more suitable.

5.1.6.4. VHD-based VDIs

VHD files may be *chained*, allowing two VDIs to share common data. In cases where a VHD-backed VM is cloned, the resulting VMs share the common on-disk data at the time of cloning. Each proceeds to make its own changes in an isolated copy-on-write (CoW) version of the VDI. This feature allows VHD-based VMs to be quickly cloned from templates, facilitating very fast provisioning and deployment of new VMs.

This leads to a situation where trees of chained VDIs are created over time as VMs and their associated VDIs get cloned. When one of the VDIs in a chain is deleted, XenServer rationalizes the other VDIs in the chain to remove unnecessary VDIs. This *coalescing* process runs asynchronously. The amount of disk space reclaimed and the time taken to perform the process depends on the size of the VDI and the amount of shared data. Only one coalescing process will ever be active for an SR. This process thread runs on the SR master host.

If you have critical VMs running on the master server of the pool and experience occasional slow IO due to this process, you can take steps to mitigate against this:

- Migrate the VM to a host other than the SR master
- Set the disk IO priority to a higher level, and adjust the scheduler. See [Section 5.7.10, “Virtual Disk QoS Settings”](#) for more information.

The VHD format used by LVM-based and File-based SR types in XenServer uses *Thin Provisioning*. The image file is automatically extended in 2MB chunks as the VM writes data into the disk. For File-based VHD, this has the considerable benefit that VM image files take up only as much space on the physical storage as required. With LVM-based VHD the underlying logical volume container must be sized to the virtual size of the VDI, however unused space on the underlying CoW instance disk is reclaimed when a snapshot or clone occurs. The difference between the two behaviors can be characterized in the following way:

- For *LVM-based VHDs*, the difference disk nodes within the chain consume only as much data as has been written to disk but the leaf nodes (VDI clones) remain fully inflated to the virtual size of the disk. Snapshot leaf nodes (VDI snapshots) remain deflated when not in use and can be attached Read-only to preserve the deflated allocation. Snapshot nodes that are attached Read-Write will be fully inflated on attach, and deflated on detach.
- For *file-based VHDs*, all nodes consume only as much data as has been written, and the leaf node files grow to accommodate data as it is actively written. If a 100GB VDI is allocated for a new VM and an OS is installed, the VDI file will physically be only the size of the OS data that has been written to the disk, plus some minor metadata overhead.

When cloning VMs based on a single VHD template, each child VM forms a chain where new changes are written to the new VM, and old blocks are directly read from the parent template. If the new VM was converted into a

further template and more VMs cloned, then the resulting chain will result in degraded performance. XenServer supports a maximum chain length of 30, but it is generally not recommended that you approach this limit without good reason. If in doubt, "copy" the VM using XenCenter or use the **vm-copy** command, which resets the chain length back to 0.

5.2. Storage Repository Formats

New storage repositories can be created in XenCenter using the **New Storage Repository** wizard. This wizard guides you through the configuration steps. Alternatively use the CLI, and the **sr-create** command. This command creates a new SR on the storage substrate (potentially destroying any existing data), and creates the SR API object and a corresponding PBD record, enabling VMs to use the storage. On successful creation of the SR, the PBD is automatically plugged. If the SR `shared=true` flag is set, a PBD record is created and plugged for every XenServer in the resource pool.

If you are creating an SR for IP-based storage (iSCSI or NFS), you can configure one of the following as the storage network: the NIC that handles the management traffic or a new NIC for the storage traffic. To assign an IP address to a NIC, see [Section 4.4.7, "Configuring a Dedicated Storage NIC"](#).

All XenServer SR types support VDI resize, fast cloning and snapshot. SRs based on the LVM SR type (local, iSCSI, or HBA) provide thin provisioning for snapshot and hidden parent nodes. The other SR types support full thin provisioning, including for virtual disks that are active.

Warning:

When VHD VDIs are not attached to a VM, for example in the case of a VDI snapshot, they are stored as thinly-provisioned by default. When attempting to re-attach the VDI, customers should ensure that there is sufficient disk-space available for the VDI to become thickly provisioned. VDI clones are thickly-provisioned.

The maximum supported VDI sizes are:

Storage Repository Format	Maximum VDI size
EXT3	2TB
LVM	2TB
NFS	2TB
iSCSI	2TB
HBA	2TB

5.2.1. Local LVM

The Local LVM type presents disks within a locally-attached Volume Group.

By default, XenServer uses the local disk on the physical host on which it is installed. The Linux Logical Volume Manager (LVM) is used to manage VM storage. A VDI is implemented in VHD format in an LVM logical volume of the specified size.

5.2.1.1. LVM Performance Considerations

The snapshot and fast clone functionality provided in XenServer 5.5 and later for LVM-based SRs comes with an inherent performance overhead. In cases where optimal performance is desired, XenServer supports creation of VDIs in the *raw* format in addition to the default VHD format. The XenServer snapshot functionality is not supported on raw VDIs.

Note:

Non-transportable snapshots using the default Windows VSS provider will work on any type of VDI.

Warning:

Do not try to snapshot a VM that has type=raw disks attached. This could result in a partial snapshot being created. In this situation, you can identify the orphan snapshot VDIs by checking the snapshot -of field and then deleting them.

5.2.1.2. Creating a Local LVM SR (lvm)

An LVM SR is created by default on host install.

Device-config parameters for lvm SRs are:

Parameter Name	Description	Required?
Device	device name on the local host to use for the SR	Yes

To create a local lvm SR on /dev/sdb use the following command.

```
xe sr-create host-uuid=<valid_uuid> content-type=user \
name-label=<"Example Local LVM SR"> shared=false \
device-config:device=/dev/sdb type=lvm
```

5.2.2. Local EXT3

Using EXT3 enables Thin Provisioning on local storage. However, the default Storage Repository type is LVM as it gives a consistent write performance and, prevents storage over-commit. Customers using EXT3 may see reduced performance when carrying out VM lifecycle operations such as VM Create, suspend/resume and when initially creating large files from within the VM.

Local disk EXT SRs must be configured using the XenServer CLI.

5.2.2.1. Creating a Local EXT3 SR (ext)

Device-config parameters for ext SRs:

Parameter Name	Description	Required?
Device	device name on the local host to use for the SR	Yes

To create a local ext SR on /dev/sdb use the following command:

```
xe sr-create host-uuid=<valid_uuid> content-type=user \
name-label=<"Example Local EXT3 SR"> shared=false \
device-config:device=/dev/sdb type=ext
```

5.2.3. udev

The udev type represents devices plugged in using the udev device manager as VDIs.

XenServer has two SRs of type *udev* that represent removable storage. One is for the CD or DVD disk in the physical CD or DVD-ROM drive of the XenServer host. The other is for a USB device plugged into a USB port of the XenServer host. VDIs that represent the media come and go as disks or USB sticks are inserted and removed.

5.2.4. ISO

The ISO type handles CD images stored as files in ISO format. This SR type is useful for creating shared ISO libraries. For storage repositories that store a library of ISOs, the `content - type` parameter must be set to `iso`.

For example:

```
xe sr-create host-uuid=<valid_uuid> content-type=iso \
  type=iso name-label=<"Example ISO SR"> \
  device-config:location=<nfs server:path>
```

5.2.5. Software iSCSI Support

XenServer provides support for shared SRs on iSCSI LUNs. iSCSI is supported using the Open-iSCSI software iSCSI initiator or by using a supported iSCSI Host Bus Adapter (HBA). The steps for using iSCSI HBAs are identical to those for Fibre Channel HBAs, both of which are described in [Section 5.2.8.2, "Creating a Shared LVM over Fibre Channel / Fibre Channel over Ethernet / iSCSI HBA or SAS SR \(lvmohba\)"](#).

Shared iSCSI support using the software iSCSI initiator is implemented based on the Linux Volume Manager (LVM) and provides the same performance benefits provided by LVM VDIs in the local disk case. Shared iSCSI SRs using the software-based host initiator are capable of supporting VM agility using XenMotion: VMs can be started on any XenServer host in a resource pool and migrated between them with no noticeable downtime.

iSCSI SRs use the entire LUN specified at creation time and may not span more than one LUN. CHAP support is provided for client authentication, during both the data path initialization and the LUN discovery phases.

Note:

The block size of an iSCSI LUN **must** be 512 bytes.

5.2.5.1. XenServer Host iSCSI configuration

All iSCSI initiators and targets must have a unique name to ensure they can be uniquely identified on the network. An initiator has an iSCSI initiator address, and a target has an iSCSI target address. Collectively these are called iSCSI Qualified Names, or IQNs.

XenServer hosts support a single iSCSI initiator which is automatically created and configured with a random IQN during host installation. The single initiator can be used to connect to multiple iSCSI targets concurrently.

iSCSI targets commonly provide access control using iSCSI initiator IQN lists, so all iSCSI targets/LUNs to be accessed by a XenServer host must be configured to allow access by the host's initiator IQN. Similarly, targets/LUNs to be used as shared iSCSI SRs must be configured to allow access by all host IQNs in the resource pool.

Note:

iSCSI targets that do not provide access control will typically default to restricting LUN access to a single initiator to ensure data integrity. If an iSCSI LUN is intended for use as a shared SR across multiple XenServer hosts in a resource pool, ensure that multi-initiator access is enabled for the specified LUN.

The XenServer host IQN value can be adjusted using XenCenter, or using the CLI with the following command when using the iSCSI software initiator:

```
xe host-param-set uuid=<valid_host_id> other-config:iscsi_iqn=<new_initiator_iqn>
```

Warning:

It is imperative that every iSCSI target and initiator have a unique IQN. If a non-unique IQN identifier is used, data corruption and/or denial of LUN access can occur.

Warning:

Do not change the XenServer host IQN with iSCSI SRs attached. Doing so can result in failures connecting to new targets or existing SRs.

5.2.6. Software FCoE Storage

Software FCoE provides a standard framework to which hardware vendors can plug in their FCoE-capable NIC and get the same benefits of a hardware-based FCoE. This eliminates the need for using expensive HBAs.

Before creating a new software FCoE storage, configuration required to expose a LUN to the host must be completed manually, including configuring the FCoE fabric and allocating LUNs to your SAN's public world wide name (PWWN). After completing this configuration, available LUN should be mounted to the host's CNA as a SCSI device. The SCSI device can then be used to access the LUN as if it were a locally attached SCSI device. For information about configuring the physical switch and the array to support FCoE, refer to the documentation provided by the vendor.

Note:

Software FCoE can be used with Open vSwitch and Linux bridge as the network backend.

5.2.6.1. Creating a Software FCoE SR

Before creating a Software FCoE SR, customers must ensure there are FCoE-capable NICs attached to the host.

Device-config parameters for FCoE SRs are:

Parameter Name	Description	Required?
SCSIid	the SCSI bus ID of the destination LUN	Yes

Run the following command to create a shared FCoE SR:

```
xe sr-create type=lvmofcoe \
name-label=<"FCoE SR"> shared=true device-config:SCSIid=<SCSI_id>
```

5.2.7. Hardware Host Bus Adapters (HBAs)

This section covers various operations required to manage SAS, Fibre Channel and iSCSI HBAs.

5.2.7.1. Sample QLogic iSCSI HBA setup

For full details on configuring QLogic Fibre Channel and iSCSI HBAs please refer to the [QLogic](#) website.

Once the HBA is physically installed into the XenServer host, use the following steps to configure the HBA:

1. Set the IP networking configuration for the HBA. This example assumes DHCP and HBA port 0. Specify the appropriate values if using static IP addressing or a multi-port HBA.

```
/opt/QLogic_Corporation/SANsurferiCLI/iscli -ipdhcp 0
```

2. Add a persistent iSCSI target to port 0 of the HBA.

```
/opt/QLogic_Corporation/SANsurferiCLI/iscli -pa 0 <iscsi_target_ip_address>
```

3. Use the **xe sr-probe** command to force a rescan of the HBA controller and display available LUNs. See [Section 5.3.2, “Probing an SR”](#) and [Section 5.2.8.2, “Creating a Shared LVM over Fibre Channel / Fibre Channel over Ethernet / iSCSI HBA or SAS SR \(lvmohba\)”](#) for more details.

5.2.7.2. Removing HBA-based SAS, FC or iSCSI Device Entries

Note:

This step is not required. Citrix recommends that only power users perform this process if it is necessary.

Each HBA-based LUN has a corresponding global device path entry under `/dev/disk/by-scsibus` in the format `<SCSIid>:<adapter>:<bus>:<target>:<lun>` and a standard device path under `/dev`. To remove the device entries for LUNs no longer in use as SRs use the following steps:

1. Use **sr-forget** or **sr-destroy** as appropriate to remove the SR from the XenServer host database. See [Section 5.7.1, “Removing SRs”](#) for details.
2. Remove the zoning configuration within the SAN for the desired LUN to the desired host.
3. Use the **sr-probe** command to determine the ADAPTER, BUS, TARGET, and LUN values corresponding to the LUN to be removed. See [Section 5.3.2, “Probing an SR”](#) for details.
4. Remove the device entries with the following command:

```
echo "1" > /sys/class/scsi_device/<adapter>:<bus>:<target>:<lun>/device/delete
```

Warning:

Make absolutely sure you are certain which LUN you are removing. Accidentally removing a LUN required for host operation, such as the boot or root device, will render the host unusable.

5.2.8. Shared LVM Storage

The Shared LVM type represents disks as Logical Volumes within a Volume Group created on an iSCSI (FC or SAS) LUN.

Note:

The block size of an iSCSI LUN **must** be 512 bytes.

5.2.8.1. Creating a Shared LVM Over iSCSI SR Using the Software iSCSI Initiator (lvmoiscsi)

Device-config parameters for lvmoiscsi SRs:

Parameter Name	Description	Required?
target	the IP address or hostname of the iSCSI filer that hosts the SR	yes
targetIQN	the IQN target address of iSCSI filer that hosts the SR	yes
SCSIid	the SCSI bus ID of the destination LUN	yes

Parameter Name	Description	Required?
chapuser	the username to be used for CHAP authentication	no
chappassword	the password to be used for CHAP authentication	no
port	the network port number on which to query the target	no
usediscoverynumber	the specific iSCSI record index to use	no
incoming_chapuser	the username that the iSCSI filter will use to authenticate against the host	no
incoming_chappassword	the password that the iSCSI filter will use to authenticate against the host	no

To create a shared lvmoiscsi SR on a specific LUN of an iSCSI target use the following command.

```
xe sr-create host-uuid=<valid_uuid> content-type=user \
name-label=<"Example shared LVM over iSCSI SR"> shared=true \
device-config:target=<target_ip=> device-config:targetIQN=<target_iqn=> \
device-config:SCSIid=<scsci_id> \
type=lvmoiscsi
```

5.2.8.2. Creating a Shared LVM over Fibre Channel / Fibre Channel over Ethernet / iSCSI HBA or SAS SR (lvmohba)

SRs of type lvmohba can be created and managed using the xe CLI or XenCenter.

Device-config parameters for lvmohba SRs:

Parameter name	Description	Required?
SCSIid	Device SCSI ID	Yes

To create a shared lvmohba SR, perform the following steps on each host in the pool:

1. Zone in one or more LUNs to each XenServer host in the pool. This process is highly specific to the SAN equipment in use. Please refer to your SAN documentation for details.
 2. If necessary, use the HBA CLI included in the XenServer host to configure the HBA:
 - Emulex: /bin/sbin/ocmanager
 - QLogic FC: /opt/QLogic_Corporation/SANsurferCLI
 - QLogic iSCSI: /opt/QLogic_Corporation/SANsurferiCLI
- See [Section 5.2.7, "Hardware Host Bus Adapters \(HBAs\)"](#) for an example of QLogic iSCSI HBA configuration. For more information on Fibre Channel and iSCSI HBAs please refer to the [Emulex](#) and [QLogic](#) websites.
3. Use the **sr-probe** command to determine the global device path of the HBA LUN. **sr-probe** forces a re-scan of HBAs installed in the system to detect any new LUNs that have been zoned to the host and returns a list of properties for each LUN found. Specify the host-uuid parameter to ensure the probe occurs on the desired host.

The global device path returned as the <path> property will be common across all hosts in the pool and therefore must be used as the value for the device-config:device parameter when creating the SR.

If multiple LUNs are present use the vendor, LUN size, LUN serial number, or the SCSI ID as included in the <path> property to identify the desired LUN.

```
xe sr-probe type=lvmotha \
host-uuid=1212c7b3-f333-4a8d-a6fb-80c5b79b5b31
Error code: SR_BACKEND_FAILURE_90
Error parameters: , The request is missing the device parameter, \
<?xml version="1.0" ?>
<Devlist>
  <BlockDevice>
    <path>
      /dev/disk/by-id/scsi-360a9800068666949673446387665336f
    </path>
    <vendor>
      HITACHI
    </vendor>
    <serial>
      730157980002
    </serial>
    <size>
      80530636800
    </size>
    <adapter>
      4
    </adapter>
    <channel>
      0
    </channel>
    <id>
      4
    </id>
    <lun>
      2
    </lun>
    <hba>
      qla2xxx
    </hba>
  </BlockDevice>
  <Adapter>
    <host>
      Host4
    </host>
    <name>
      qla2xxx
    </name>
    <manufacturer>
      QLogic HBA Driver
    </manufacturer>
    <id>
      4
    </id>
  </Adapter>
</Devlist>
```

4. On the master host of the pool create the SR, specifying the global device path returned in the <path> property from **sr-probe**. PBDs will be created and plugged for each host in the pool automatically.

```
xe sr-create host-uuid=<valid_uuid> \
content-type=user \
name-label=<"Example shared LVM over HBA SR"> shared=true \
device-config:SCSIid=<device_scsi_id> type=lvmotha
```

Note:

You can use the XenCenter **Repair Storage Repository** function to retry the PBD creation and plugging portions of the **sr-create** operation. This can be valuable in cases where the LUN zoning was incorrect for one or more hosts in a pool when the SR was created. Correct

the zoning for the affected hosts and use the **Repair Storage Repository** function instead of removing and re-creating the SR.

5.2.9. NFS and SMB

NFS and SMB SR types store disks as VHD files on a remote filesystem.

XenServer allows shares on NFS servers (that support NFSv4 or NFSv3), or shares on SMB servers (that support SMB 3.0), over TCP/IP to be used immediately as a storage repository for virtual disks (VDIs). VDIs are stored in the Microsoft VHD format only. Additionally, as these SRs can be shared, VDIs stored on shared SRs allow:

- VMs to be started on any XenServer hosts in a resource pool
- VM migrate between XenServer hosts in a resource pool using XenMotion (without noticeable downtime)

Important:

- Support for SMB 3.0 is limited to the ability to connect to a share using the 3.0 protocol. Additional features like Transparent Failover are dependent on feature availability in the upstream Linux kernel and are not supported in XenServer 7.0.
- For NFSv4, only the authentication type AUTH_SYS is supported in XenServer 7.0.

Note:

SMB storage is available for XenServer Enterprise edition customers, or those who have access to XenServer through their XenApp/XenDesktop entitlement. To learn more about XenServer editions, and to find out how to upgrade, visit the Citrix website [here](#). For detailed information on Licensing, refer to the [XenServer 7.0 Licensing FAQ](#).

VDIs stored on file-based SRs are *Thinly Provisioned*. The image file is allocated as the VM writes data into the disk. This has the considerable benefit that the VM image files take up only as much space on the storage as is required. For example, if a 100GB VDI is allocated for a new VM and an OS is installed, the VDI file will only reflect the size of the OS data that has been written to the disk rather than the entire 100GB.

VHD files may also be chained, allowing two VDIs to share common data. In cases where a file-based VM is cloned, the resulting VMs will share the common on-disk data at the time of cloning. Each VM will proceed to make its own changes in an isolated copy-on-write version of the VDI. This feature allows file-based VMs to be quickly cloned from templates, facilitating very fast provisioning and deployment of new VMs.

Note:

The maximum supported length of VHD chains is 30.

XenServer's file-based SRs and VHD implementations assume that they have full control over the SR directory on the file server. Administrators should not modify the contents of the SR directory, as this can risk corrupting the contents of VDIs.

XenServer has been tuned for enterprise-class storage that use non-volatile RAM to provide fast acknowledgments of write requests while maintaining a high degree of data protection from failure. XenServer has been tested extensively against Network Appliance FAS2020 and FAS3210 storage, using Data OnTap 7.3 and 8.1

Warning:

As VDIs on file-based SRs are created as Thin Provisioned, administrators must ensure that there is enough disk space on the file-based SRs for all required VDIs. XenServer hosts do not enforce that the space required for VDIs on file-based SRs is actually present.

5.2.9.1. Creating a Shared NFS SR (NFS)

To create an NFS SR, you will need to provide the hostname or IP address of the NFS server. You can create the SR on any valid destination path; use the `sr - probe` command to display a list of valid destination paths exported by the server.

In scenarios where XenServer is used with lower-end storage, it will cautiously wait for all writes to be acknowledged before passing acknowledgments on to guest VMs. This will incur a noticeable performance cost, and might be solved by setting the storage to present the SR mount point as an asynchronous mode export. Asynchronous exports acknowledge writes that are not actually on disk, and so administrators should consider the risks of failure carefully in these situations.

Note:

The NFS server must be configured to export the specified path to all XenServer hosts in the pool, or the creation of the SR and the plugging of the PBD record will fail.

The XenServer NFS implementation uses TCP by default. If your situation allows, you can configure the implementation to use UDP in scenarios where there may be a performance benefit. To do this, when creating an SR, specify the device-config parameter `useUDP=true`.

Device-config parameters for NFS SRs:

Parameter Name	Description	Required?
server	IP address or hostname of the NFS server	Yes
serverpath	path, including the NFS mount point, to the NFS server that hosts the SR	Yes

For example, to create a shared NFS SR on `192.168.1.10:/export1`, use the following command:

```
xe sr-create content-type=user \
name-label=<"shared NFS SR"> shared=true \
device-config:server=192.168.1.10 device-config:serverpath=/export1 type=nfs \
nfsversion=<"3", "4">
```

To create a non-shared NFS SR, run the following command:

```
xe sr-create host-uuid=<host_uuid> content-type=user \
name-label=<"Non-shared NFS SR"> \
device-config:server=192.168.1.10 device-config:serverpath=/export1 type=nfs \
nfsversion=<"3", "4">
```

5.2.9.2. Creating a Shared SMB SR (SMB)

To create an SMB SR, you will need to provide the hostname or IP address of the SMB server, the full path of the exported share and appropriate credentials.

Note:

SMB SR has been tested against Network Appliance storage running OnTap 8.3 and Windows Server 2012 R2.

Device-config parameters for SMB SRs:

Parameter Name	Description	Required?
server	Full path to share on server	Yes
username	User account with RW access to share	Optional
password	Password for the user account	Optional

For example, to create a shared SMB SR on `192.168.1.10:/share1`, use the following command:


```
xe sr-create content-type=user \
name-label=<"Example shared SMB SR"> shared=true \
device-config:server=//192.168.1.10/share1 \
device-config:username=<valid_username> device-config:password=<valid_password> type=smb
```

To create a non-shared SMB SR, run the following command:

```
xe sr-create host-uuid=<host_uuid> content-type=user \
name-label=<"Non-shared SMB SR"> \
device-config:server=//192.168.1.10/share1 \
device-config:username=<valid_username> device-config:password=<valid_password> type=smb
```

5.2.10. LVM over Hardware HBA

The LVM over hardware HBA type represents disks as VHDs on Logical Volumes within a Volume Group created on an HBA LUN providing, for example, hardware-based iSCSI or FC support.

XenServer hosts support Fibre Channel (FC) storage area networks (SANs) through Emulex or QLogic host bus adapters (HBAs). All FC configuration required to expose a FC LUN to the host must be completed manually, including storage devices, network devices, and the HBA within the XenServer host. Once all FC configuration is complete the HBA will expose a SCSI device backed by the FC LUN to the host. The SCSI device can then be used to access the FC LUN as if it were a locally attached SCSI device.

Use the **sr-probe** command to list the LUN-backed SCSI devices present on the host. This command forces a scan for new LUN-backed SCSI devices. The path value returned by **sr-probe** for a LUN-backed SCSI device is consistent across all hosts with access to the LUN, and therefore must be used when creating shared SRs accessible by all hosts in a resource pool.

The same features apply to QLogic iSCSI HBAs.

See [Section 5.3.1, "Creating Storage Repositories"](#) for details on creating shared HBA-based FC and iSCSI SRs.

Note:

XenServer support for Fibre Channel does not support direct mapping of a LUN to a VM. HBA-based LUNs must be mapped to the host and specified for use in an SR. VDIs within the SR are exposed to VMs as standard block devices.

5.3. Storage Configuration

This section covers creating storage repository types and making them available to a XenServer host. The examples provided show how to do this using the CLI. See the XenCenter Help for details on using the **New Storage Repository** wizard.

5.3.1. Creating Storage Repositories

This section explains how to create Storage Repositories (SRs) of different types and make them available to a XenServer host. The examples provided cover creating SRs using the xe CLI. See the XenCenter help for details on using the **New Storage Repository** wizard to add SRs using XenCenter.

Note:

Local SRs of type `lvm` and `ext3` can only be created using the xe CLI. After creation all SR types can be managed by either XenCenter or the xe CLI.

There are two basic steps involved in creating a new storage repository for use on a XenServer host using the CLI:

1. Probe the SR type to determine values for any required parameters.
2. Create the SR to initialize the SR object and associated PBD objects, plug the PBDs, and activate the SR.

These steps differ in detail depending on the type of SR being created. In all examples the **sr-create** command returns the UUID of the created SR if successful.



SRs can also be *destroyed* when no longer in use to free up the physical device, or *forgotten* to detach the SR from one XenServer host and attach it to another. See [Section 5.7.1, “Removing SRs”](#) for details.

5.3.2. Probing an SR

The **sr-probe** command can be used in two ways:

1. To identify unknown parameters for use in creating a SR.
2. To return a list of existing SRs.

In both cases **sr-probe** works by specifying an SR type and one or more `device-config` parameters for that SR type. When an incomplete set of parameters is supplied the **sr-probe** command returns an error message indicating parameters are missing and the possible options for the missing parameters. When a complete set of parameters is supplied a list of existing SRs is returned. All **sr-probe** output is returned as XML.

For example, a known iSCSI target can be probed by specifying its name or IP address, and the set of IQNs available on the target will be returned:

```
xe sr-probe type=lvmoiscsi device-config:target=<192.168.1.10>
```

Error code: SR_BACKEND_FAILURE_96

Error parameters: , The request is missing or has an incorrect target IQN parameter, \

```
<?xml version="1.0" ?>
```

```
<iscsi-target-iqns>
```

```
  <TGT>
```

```
    <Index>
```

```
      0
```

```
    </Index>
```

```
    <IPAddress>
```

```
      192.168.1.10
```

```
    </IPAddress>
```

```
    <TargetIQN>
```

```
      iqn.192.168.1.10:filer1
```

```
    </TargetIQN>
```

```
  </TGT>
```

```
</iscsi-target-iqns>
```

Probing the same target again and specifying both the name/IP address and desired IQN returns the set of SCSIids (LUNs) available on the target/IQN.

```
xe sr-probe type=lvmoiscsi device-config:target=192.168.1.10 \
device-config:targetIQN=iqn.192.168.1.10:filer1
```

Error code: SR_BACKEND_FAILURE_107

Error parameters: , The SCSIid parameter is missing or incorrect, \

```
<?xml version="1.0" ?>
```

```
<iscsi-target>
```

```
  <LUN>
```

```
    <vendor>
```

```
    IET
```

```
    </vendor>
```

```
    <LUNid>
```

```
      0
```

```
    </LUNid>
```

```
    <size>
```

```
      42949672960
```

```
    </size>
```

```
    <SCSIid>
```

```
      149455400000000000000000002000000b70200000f000000
```

```
    </SCSIid>
```

```
  </LUN>
```

```
</iscsi-target>
```

Probing the same target and supplying all three parameters will return a list of SRs that exist on the LUN, if any.

```
xe sr-probe type=lvmoiscsi device-config:target=192.168.1.10 \
device-config:targetIQN=192.168.1.10:filer1 \
device-config:SCSIid=149455400000000000000000002000000b70200000f000000
```

```
<?xml version="1.0" ?>
<SRlist>
  <SR>
    <UUID>
      3f6e1ebd-8687-0315-f9d3-b02ab3adc4a6
    </UUID>
    <Devlist>
      /dev/disk/by-id/scsi-149455400000000000000000002000000b70200000f000000
    </Devlist>
  </SR>
</SRlist>
```

The following parameters can be probed for each SR type:

SR type	device-config parameter, in order of dependency	Can be probed?	Required for sr-create?
lvmoiscsi	target	No	Yes
	chapuser	No	No
	chappassword	No	No
	targetIQN	Yes	Yes
	SCSIid	Yes	Yes
lvmoaha	SCSIid	Yes	Yes
NetApp	target	No	Yes
	username	No	Yes
	password	No	Yes
	chapuser	No	No
	chappassword	No	No
	aggregate	No*	Yes
	FlexVols	No	No
	allocation	No	No
	asis	No	No
nfs	server	No	Yes
	serverpath	Yes	Yes
lvm	device	No	Yes
ext	device	No	Yes
EqualLogic	target	No	Yes

SR type	device-config parameter, in order of dependency	Can be probed?	Required for sr-create?
	username	No	Yes
	password	No	Yes
	chapuser	No	No
	chappassword	No	No
	storagepool	No [†]	Yes

*Aggregate probing is only possible at **sr-create** time. It needs to be done there so that the aggregate can be specified at the point that the SR is created.

†Storage pool probing is only possible at **sr-create** time. It needs to be done there so that the aggregate can be specified at the point that the SR is created.

5.4. Storage Multipathing

Dynamic multipathing support is available for Fibre Channel and iSCSI storage backends. By default, XenServer uses round-robin mode load balancing, so both routes have active traffic on them during normal operation. You can enable multipathing in XenCenter or on the xe CLI. For additional information about multipathing, see CTX134881—[Configuring Multipathing for XenServer](#).

Before attempting to enable multipathing, verify that multiple targets are available on your storage server. For example, an iSCSI storage backend queried for `sendtargets` on a given portal should return multiple targets, as in the following example:

```
iscsiadm -m discovery --type sendtargets --portal 192.168.0.161
192.168.0.161:3260,1 iqn.strawberry:litchie
192.168.0.204:3260,2 iqn.strawberry:litchie
```

To enable storage multipathing using the xe CLI

1. Unplug all PBDs on the host:

```
xe pbd-unplug uuid=<pbid_uuid>
```

2. Set the host's `other-config:multipathing` parameter:

```
xe host-param-set other-config:multipathing=true uuid=host_uuid
```

3. Set the host's `other-config:multipathhandle` parameter to `dmp`:

```
xe host-param-set other-config:multipathhandle=dmp uuid=host_uuid
```

4. If there are existing SRs on the host running in single path mode but that have multiple paths:

- Migrate or suspend any running guests with virtual disks in affected the SRs
- Unplug and re-plug the PBD of any affected SRs to reconnect them using multipathing:

```
xe pbd-plug uuid=<pbid_uuid>
```

To disable multipathing, first unplug your VBDs, set the host `other-config:multipathing` parameter to `false` and then replug your PBDs as described above. Do not modify the `other-config:multipathhandle` parameter as this will be done automatically.

Multipath support in XenServer is based on the device-mapper `multipathd` components. Activation and deactivation of multipath nodes is handled automatically by the Storage Manager API. Unlike the standard `dm-multipath` tools in Linux, device mapper nodes are not automatically created for all LUNs on the system, and it is only when LUNs are actively used by the storage management layer that new device mapper nodes are

provisioned. Therefore, it is unnecessary to use any of the `dm-multipath` CLI tools to query or refresh DM table nodes in XenServer. Should it be necessary to query the status of device-mapper tables manually, or list active device mapper multipath nodes on the system, use the `mpathutil` utility:

- `mpathutil list`
- `mpathutil status`

Note:

Due to incompatibilities with the integrated multipath management architecture, the standard `dm-multipath` CLI utility *should not be used* with XenServer. Please use the `mpathutil` CLI tool for querying the status of nodes on the host.

Note:

Multipath support in EqualLogic arrays does not encompass Storage IO multipathing in the traditional sense of the term. Multipathing must be handled at the network/NIC bond level. Refer to the EqualLogic documentation for information about configuring network failover for EqualLogic SRs/LVMoISCSI SRs.

5.5. XenServer and IntelliCache

Note:

This feature is only supported when using XenServer with XenDesktop.

Using XenServer with *IntelliCache* makes hosted Virtual Desktop Infrastructure deployments more cost-effective by enabling you to use a combination of shared storage and local storage. It is of particular benefit when many Virtual Machines (VMs) all share a common OS image. The load on the storage array is reduced and performance is enhanced. In addition, network traffic to and from shared storage is reduced as the local storage caches the master image from shared storage.

IntelliCache works by caching data from a VMs parent VDI in local storage on the VM host. This local cache is then populated as data is read from the parent VDI. When many VMs share a common parent VDI (for example by all being based on a particular master image), the data pulled into the cache by a read from one VM can be used by another VM. This means that further access to the master image on shared storage is not required.

A thin provisioned, local SR is an IntelliCache prerequisite. Thin Provisioning is a way of optimizing the utilization of available storage. This approach allows you to make more use of local storage instead of shared storage. It relies on on-demand allocation of blocks of data instead of the traditional method of pre-allocating all of the blocks.

Important:

Thin Provisioning changes the default local storage type of the host from LVM to EXT3. Thin Provisioning **must be** enabled in order for XenDesktop local caching to work properly.

Thin Provisioning allows the administrator to present more storage space to the VMs connecting to the Storage Repository (SR) than is actually available on the SR. There are no space guarantees, and allocation of a LUN does not claim any data blocks until the VM writes data.

Warning:

Thin provisioned SRs may run out of physical space, as the VMs within can grow to consume disk capacity on demand. IntelliCache VMs handle this condition by automatically falling back to shared storage if the local SR cache is full. It is not recommended to mix traditional virtual machines and IntelliCache VMs on the same SR, as intellicache VMs may grow quickly in size.

5.5.1. IntelliCache Deployment

IntelliCache must be enabled either during host installation or be enabled manually on a running host using the CLI.

Citrix recommends that you use a high performance local storage device to ensure the fastest possible data transfer such as a Solid State Disk or a high performance RAID array. Both data throughput and storage capacity should be considered when sizing local disks. The shared storage type, used to host the source Virtual Disk Image (VDI), must be NFS or EXT based.

5.5.1.1. Enabling on Host Installation

To enable IntelliCache during host installation, on the Virtual Machine Storage screen, select **Enable thin provisioning (Optimized storage for XenDesktop)**. This selects the host's local SR to be the one to be used for the local caching of VM VDIs.



5.5.1.2. Converting an Existing Host to Use Thin Provisioning

To destroy an existing LVM based local SR, and replace it with a thin provisioned EXT3 based SR, enter the following commands.

Warning:

These commands will destroy your existing local SR, and VMs on the SR will be permanently deleted.

```
localsr=`xe sr-list type=lvm host=<hostname> params=uuid --minimal`
echo localsr=$localsr
pbd=`xe pbd-list sr-uuid=$localsr params=uuid --minimal`
echo pbd=$pbd
xe pbd-unplug uuid=$pbd
xe pbd-destroy uuid=$pbd
xe sr-forget uuid=$localsr
sed -i "s/'lvm'/'ext'/" /etc/firstboot.d/data/default-storage.conf
rm -f /etc/firstboot.d/state/10-prepare-storage
rm -f /etc/firstboot.d/state/15-set-default-storage
service firstboot start
xe sr-list type=ext
```

To enable local caching, enter the following commands:

```
xe host-disable host=<hostname>
localsr=`xe sr-list type=ext host=<hostname> params=uuid --minimal`
xe host-enable-local-storage-caching host=<hostname> sr-uuid=$localsr
xe host-enable host=<hostname>
```

5.5.1.3. VM Boot Behavior

There are two options for the behavior of a VM VDI when the VM is booted:

1. Shared Desktop Mode

On VM boot, the VDI is reverted to the state it was in at the previous boot. All changes while the VM is running will be lost when the VM is next booted.

Select this option if you plan to deliver standardized desktops to which users cannot make permanent changes.

2. Private Desktop Mode

On VM boot, the VDI is in the state it was left in at the last shutdown.

Select this option if you plan to allow users to make permanent changes to their desktops.

5.5.1.3.1. VM Caching Behavior Settings

The VDI flag **allow-caching** dictates the caching behavior:

5.5.1.3.1.1. Shared Desktop Mode

For shared desktops, the **on-boot** option is set to `reset` and the **allow-caching** flag is set to `true`, new VM data is written only to local storage – there will be no writes to shared storage. This means that the load on shared storage is significantly reduced. However the VM cannot be migrated between hosts.

5.5.1.3.1.2. Private Desktop Mode

For private desktops, the **on-boot** option is set to `persist` and the **allow-caching** flag is set to `true`, new VM data is written to both local and shared storage. Reads of cached data do not require I/O traffic to shared storage so the load on shared storage is somewhat reduced. VM Migration to another host is permitted and the local cache on the new host is populated as data is read.

5.5.1.4. Implementation Details and Troubleshooting

Q: Is IntelliCache compatible with XenMotion and High Availability?

A: You can use XenMotion and High Availability with IntelliCache when virtual desktops are in Private mode, that is when `on-boot=persist`

Warning:

A VM cannot be migrated if any of its VDIs have caching behavior flags set to `on-boot=reset` and `allow-caching=true`. Migration attempts for VMs with these properties will fail.

Q: Where does the local cache live on the local disk?

A: The cache lives in a Storage Repository (SR). Each host has a configuration parameter (called **local-cache-sr**) indicating which (local) SR is to be used for the cache files. Typically this will be a EXT type SR. When you run VMs with IntelliCache, you will see files inside the SR with names `<uuid>.vhdcache`. This is the cache file for the VDI with the given UUID. These files are not displayed in XenCenter – the only way of seeing them is by logging into dom0 and listing the contents of `/var/run/sr-mount/<sr-uuid>`

Q: How do I specify a particular SR for use as the cache?

A: The host object field **local-cache-sr** refers to a local SR. You can view its value by running the following command:

```
xe sr-list params=local-cache-sr,uuid,name-label
```

This field is set either:

- after host installation, if the "Enable thin provisioning" option was selected in the host installer
- by executing `xe host-enable-local-storage-caching host=<host> sr-uuid=<sr>`. This command requires the specified host to be disabled, VMs must be shut down if this command is used.

The first option uses the EXT type local SR and is created during host installation. The second option, uses the SR that is specified on the command-line.

Warning:

These steps are only necessary for users who have configured more than one local SR.

Q: When is the local cache deleted?

A: A VDI cache file is only deleted when the VDI itself is deleted. The cache is reset when a VDI is attached to a VM (for example on VM start). If the host is offline when the VDI is deleted, the SR synchronisation that runs on startup will garbage collect the cache file.

Note:

The cache file is not deleted from the host when a VM is migrated to a different host or shut down.

5.6. Storage Read Caching

Note:

Storage Read Caching is available for XenServer Enterprise Edition customers or those who have access to XenServer through their XenDesktop/XenApp entitlement. To learn more about XenServer editions, and to find out how to upgrade, visit the Citrix website [here](#). For detailed information on Licensing - refer to [XenServer 7.0 Licensing FAQ](#).

Read caching improves a VMs disk performance as, after the initial read from external disk, data is cached within the host's free memory. It greatly improves performance in situations where many VMs are cloned off a single base VM for example, in XenDesktop environment Machine Creation Service (MCS) environments, as it will drastically reduce the number of blocks read from disk.

This performance improvement can be seen whenever data needs to be read from disk more than once, since it gets cached in memory. This is most noticeable in the degradation of service that occurs during heavy I/O situations. For example, when a significant number of end users boot up within a very narrow time frame (boot storm), or when a significant number of VMs are scheduled to run malware scans at the same time (anti-virus storms). Read caching is enabled by default if you have the appropriate licence type.

5.6.1. Enabling and Disabling

For file-based SRs, such as NFS and EXT3 SR types, read-caching is enabled by default. It is disabled for all other SRs.

To disable read caching for a specific SR, enter:

```
xe sr-param-set uuid=<sr-uuid> other-config:o_direct=true
```


5.6.2. Limitations

- Read caching is available only for NFS and EXT3 SRs. It is not available for other SR Types.
- Read caching only applies to read-only VDI and VDI parents; these exist where VMs are created from 'Fast Clone' or snap-shotted disks. The greatest performance improvements will be seen when many VMs are cloned from a single 'golden' image.
- Performance improvements are dependent on the amount of free memory available in the host's Control Domain (dom0). Increasing the amount of memory to dom0, will allow more memory to be allocated to the read-cache. Refer to [CTX134951 - How to Configure dom0 Memory in XenServer 6.1 and Later](#) for information on how to configure dom0 memory.

5.6.3. Comparison with IntelliCache

IntelliCache and memory based read caching are to some regards complementary. IntelliCache not only caches on a different tier, but it also caches writes in addition to reads. The major difference is that IntelliCache works by caching reads from the network onto a local disk, while in-memory read caching caches reads from either network or disk into host memory. The advantage of in-memory read caching, is that memory is still an order of magnitude faster than a solid-state disk (SSD), so performance in boot storms and other heavy I/O situations should be improved.

It is possible for both read-caching and IntelliCache to be enabled simultaneously; in this case reads from the network are cached by IntelliCache to a local disk, and reads from that local disk are cached in memory with read caching.

5.6.4. To Set the Read Cache Size

The read cache performance can be optimized, by giving more memory to XenServer's control domain (dom0).

Important:

For optimization, the read cache size should be set on **ALL** hosts in the pool individually. Any subsequent changes to the size of the read cache must also be set on all hosts in the pool.

On the XenServer host, open a local shell and log on as root.

To set the size of the read cache, run the following command:

```
/opt/xensource/libexec/xen-cmdline --set-xen dom0_mem=<nn>M, max:<nn>M
```

Both the initial and maximum values should be set to the same value, for example, to set dom0 memory to 2048MB:

```
/opt/xensource/libexec/xen-cmdline --set-xen dom0_mem=20480M, max:20480M
```

Important:

You must reboot all hosts after making any changes to the read cache size.

5.6.4.1. How to View the Current dom0 Memory Allocation

To view the current dom0 memory settings enter:

```
free -m
```

The output of `free -m` shows the current dom0 memory settings. The value may be less than expected due to various overheads. The example table below shows the output from a host with dom0 set to 752MB

	Total	Used	Free	Shared	Buffer	Cached
Mem	594	486	108	0	149	78

-/+ buffers/ cache		258	336			
Swap	511	49	462			

What Range of Values Can be Used?

As the XenServer 7.0 Control Domain (dom0) is 64-bit, large values can be used, for example 32768MB. Citrix recommends that you do **NOT** set the initial value to be lower than 752MB, as doing so may cause the host to fail to boot.

5.6.4.2. XenCenter Display Notes

Administrators should note that the entire host's memory can be considered to comprised of the Xen hypervisor, dom0, VMs and free memory. While dom0 and VM memory is usually of a fixed size, the Xen hypervisor uses a variable amount of memory. This is dependent on a number of factors including the number of VMs running on the host at any time and how those VMs are configured. It is not possible to limit the amount of memory that Xen uses, as it would then be possible for Xen to run out of memory, and could prevent new VMs from starting even when the host had free memory.

To view the memory allocated to a host, in XenCenter select the host, and then click the **Memory** tab.

The **XenServer** field displays the *sum* of the memory allocated to dom0 *and* Xen memory. Therefore the amount of memory displayed may be higher than specified by the administrator and its size will vary when starting and stopping VMs, even when the administrator has set a fixed size for dom0.

5.7. Managing Storage Repositories

This section covers various operations required in the ongoing management of Storage Repositories (SRs), including Live VDI Migration.

5.7.1. Removing SRs

A Storage Repository (SR) can be removed either temporarily or permanently.

- **Detach:** breaks the association between the storage device and the pool or host, (PBD Unplug) and the SR (and its VDIs) becomes inaccessible. The contents of the VDIs and the meta-information used by VMs to access the VDIs are preserved. Detach can be used when you need to temporarily take a SR offline, for example, for maintenance. A detached SR can subsequently be reattached.
- **Forget:** preserves the contents of the SR on the physical disk, but the information used to connect VMs to the VDIs it contains is permanently deleted (PBD and VBD Unplug). For example, allows you to re-attach the SR, to another XenServer host, without removing any of the SR contents.
- **Destroy:** deletes the contents of the SR from the physical disk.

In the case of Destroy or Forget, the PBD connected to the SR must be unplugged from the host.

1. Unplug the PBD to detach the SR from the corresponding XenServer host:

```
xe pbd-unplug uuid=<pbid_uuid>
```

2. To destroy the SR - deletes the SR and corresponding PBD from the XenServer host database and deletes the SR contents from the physical disk:

```
xe sr-destroy uuid=<sr_uuid>
```

3. To forget the SR - removes the SR and corresponding PBD from the XenServer host database but leaves the actual SR contents intact on the physical media:

```
xe sr-forget uuid=<sr_uuid>
```

Note:

It might take some time for the software object corresponding to the SR to be garbage collected.

5.7.2. Introducing an SR

To reintroduce a previously *forgotten* SR, you will need to create a PBD, and manually plug the PBD to the appropriate XenServer hosts to activate the SR.

The following example introduces a SR of type `lvmoiscsi`.

1. Probe the existing SR to determine its UUID:

```
xe sr-probe type=lvmoiscsi device-config:target=<192.168.1.10> \
device-config:targetIQN=<192.168.1.10:filer1> \
device-config:SCSIid=<14945540000000000000000002000000b70200000f000000>
```

2. Introduce the existing SR UUID returned from the **sr-probe** command. The UUID of the new SR is returned:

```
xe sr-introduce content-type=user name-label=<"Example Shared LVM over iSCSI SR">
shared=true uuid=<valid_sr_uuid> type=lvmoiscsi
```

3. Create a PBD to accompany the SR. The UUID of the new PBD is returned:

```
xe pbd-create type=lvmoiscsi host-uuid=<valid_uuid> sr-uuid=<valid_sr_uuid> \
device-config:target=<192.168.0.1> \
device-config:targetIQN=<192.168.1.10:filer1> \
device-config:SCSIid=<14945540000000000000000002000000b70200000f000000>
```

4. Plug the PBD to attach the SR:

```
xe pbd-plug uuid=<pbd_uuid>
```

5. Verify the status of the PBD plug. If successful the `currently-attached` property will be true:

```
xe pbd-list sr-uuid=<sr_uuid>
```

Note:

Steps 3 through 5 must be performed for each host in the resource pool, and can also be performed using the **Repair Storage Repository** function in XenCenter.

5.7.3. Live LUN Expansion

In order to fulfill capacity requirements, you may need to add capacity to the storage array to increase the size of the LUN provisioned to the XenServer host. Live LUN Expansion allows you to increase the size of the LUN without any VM downtime.

After you have added more capacity to your storage array, enter,

```
xe sr-scan sr-uuid=<sr_uuid>
```

This command rescans the SR, and any extra capacity is added and made available.

This operation is also available in XenCenter; select the SR to resize, and then click **Rescan**. For more information, press F1 to display the XenCenter Online Help.

Note:

In previous versions of XenServer explicit commands were required to resize the physical volume group of iSCSI and HBA SRs. These commands are no longer required.

Warning:

It is not possible to shrink or truncate LUNs. Reducing the LUN size on the storage array may/ can lead to data loss.

5.7.4. Live VDI Migration

Live VDI migration allows the administrator to relocate the VMs Virtual Disk Image (VDI) without shutting down the VM. This enables administrative operations such as:

- Moving a VM from cheap local storage to fast, resilient, array-backed storage.
- Moving a VM from a development to production environment.
- Moving between tiers of storage when a VM is limited by storage capacity.
- Performing storage array upgrades.

5.7.4.1. Limitations and Caveats

Live VDI Migration is subject to the following limitations and caveats

- There must be sufficient disk space available on the target repository.
- VDIs with more than one snapshot cannot be migrated.

5.7.4.2. To Move Virtual Disks using XenCenter

1. In the **Resources** pane, select the SR where the Virtual Disk is currently stored and then click the **Storage** tab.
2. In the **Virtual Disks** list, select the Virtual Disk that you would like to move, and then click **Move**.
3. In the **Move Virtual Disk** dialog box, select the target SR that you would like to move the VDI to.

Note:

Make sure that the SR has sufficient space for another virtual disk: the available space is shown in the list of available SRs.

4. Click **Move** to move the virtual disk.

For xe cli reference, refer to [Section A.4.23.9, “vdi-pool-migrate”](#).

5.7.5. Cold VDI Migration between SRs (Offline Migration)

VDIs associated with a VM can be copied from one SR to another to accommodate maintenance requirements or tiered storage configurations. XenCenter provides the ability to copy a VM and all of its VDIs to the same or a different SR, and a combination of XenCenter and the xe CLI can be used to copy individual VDIs.

For xe cli reference, refer to [Section A.4.26.25, “vm-migrate”](#).

5.7.5.1. Copying All of a VMs VDIs to a Different SR

The XenCenter **Copy** VM function creates copies of all VDIs for a selected VM on the same or a different SR. The source VM and VDIs are not affected by default. To move the VM to the selected SR rather than creating a copy, select the **Remove original VM** option in the **Copy Virtual Machine** dialog box.

1. Shutdown the VM.
2. Within XenCenter select the VM and then select the **VM > Copy VM** menu option.
3. Select the desired target SR.

5.7.5.2. Copying Individual VDIs to a Different SR

A combination of the xe CLI and XenCenter can be used to copy individual VDIs between SRs.

1. Shutdown the VM.
2. Use the xe CLI to identify the UUIDs of the VDIs to be moved. If the VM has a DVD drive its `vdi-uuid` will be listed as `<not in database>` and can be ignored.

```
xe vbd-list vm-uuid=<valid_vm_uuid>
```

Note:

The **vbd-list** command displays both the VBD and VDI UUIDs. Be sure to record the VDI UUIDs rather than the VBD UUIDs.

3. In XenCenter select the VM **Storage** tab. For each VDI to be moved, select the VDI and click the **Detach** button. This step can also be done using the **vbd-destroy** command.

Note:

If you use the **vbd-destroy** command to detach the VDI UUIDs, be sure to first check if the VBD has the parameter `other-config:owner` set to `true`. If so, set it to `false`. Issuing the **vbd-destroy** command with `other-config:owner=true` will also destroy the associated VDI.

4. Use the **vdi-copy** command to copy each of the VM VDIs to be moved to the desired SR.

```
xe vdi-copy uuid=<valid_vdi_uuid> sr-uuid=<valid_sr_uuid>
```

5. Within XenCenter select the VM **Storage** tab. Click the **Attach** button and select the VDIs from the new SR. This step can also be done use the `vbd-create` command.
6. To delete the original VDIs, within XenCenter select the **Storage** tab of the original SR. The original VDIs will be listed with an empty value for the VM field and can be deleted with the Delete button.

5.7.6. Converting Local Fibre Channel SRs to Shared SRs

Use the xe CLI and the XenCenter **Repair Storage Repository** feature to convert a local FC SR to a shared FC SR:

1. Upgrade all hosts in the resource pool to XenServer 7.0.
2. Ensure all hosts in the pool have the SR's LUN zoned appropriately. See [Section 5.3.2, "Probing an SR"](#) for details on using the **sr-probe** command to verify the LUN is present on each host.
3. Convert the SR to shared:

```
xe sr-param-set shared=true uuid=<local_fc_sr>
```

4. Within XenCenter the SR is moved from the host level to the pool level, indicating that it is now shared. The SR will be marked with a red exclamation mark to show that it is not currently plugged on all hosts in the pool.
5. Select the SR and then select the **Storage > Repair Storage Repository** menu option.
6. Click **Repair** to create and plug a PBD for each host in the pool.

5.7.7. Reclaiming Space for Block-Based Storage on the Backing Array Using Discard

Space reclamation allows you to free up unused blocks (for example, deleted VDIs in an SR) on a LUN that has been thinly-provisioned by the storage array. Once released, the reclaimed space can then be reused by the array.

Note:

This functionality is only available on a subset of storage arrays. Please refer to the [XenServer Hardware Compatibility List](#) and your storage vendor specific documentation to determine if your array supports this functionality, and whether it needs a specific configuration for operation.

To reclaim the space using XenCenter:

1. Select the **Infrastructure** view, and then click on the host or pool connected to the SR.
2. Click the **Storage** tab.
3. Select the SR from the list, and click **Reclaim freed space**.
4. To confirm the operation, click **Yes**.
5. Click **Notifications** and then **Events** to view the status of the operation.

For more information, in XenCenter press **F1** to access the Online Help.

Note:

- This is a XenCenter only operation.
- The operation is only available for LVM-based SRs that are based on thinly-provisioned LUNs on the array.
- Local SSDs will also benefit from space reclamation.
- It is not required for file-based SRs (such as NFS/Ext3); the **Reclaim Freed Space** button is not available in XenCenter for these SR types.
- Space Reclaim is an intensive operation and can lead to a degradation in storage array performance. Therefore it should only be initiated when space reclamation is required on the array. Citrix recommends that you schedule this work outside of peak array demand hours.

5.7.8. Automatically Reclaiming Space When Deleting Snapshots

When deleting snapshots with XenServer 7.0, all allocated space allocated on LVM-based SRs is reclaimed automatically and a VM reboot is not required; this is referred to as Online Coalescing.

Note:

Online Coalescing only applies to LVM-based SRs (LVM, LVMoISCSI, and LVMoHBA), it does not apply to EXT or NFS SRs, whose behavior remains unchanged.

In certain cases, automated space reclamation may be unable to proceed, in these cases it is advisable to use the Off-Line Coalesce tool:

- Under conditions where a VM I/O throughput is considerable
- In conditions where space is not being reclaimed after a period of time

Note:

Running the Off Line Coalesce tool will incur some downtime for the VM, due to the suspend/resume operations performed.

Before running the tool, delete any snapshots and clones you no longer want; the script will reclaim as much space as possible given the remaining snapshots/clones. If you want to reclaim all space, delete all snapshots and clones.

All VM disks must be either on shared or local storage for a single host. VMs with disks in both types of storage cannot be coalesced.

5.7.8.1. Reclaiming Space Using the Off Line Coalesce Tool

Note:

Online Coalescing only applies to LVM-based SRs (LVM, LVMoISCSI, and LVMoHBA), it does not apply to EXT or NFS SRs, whose behavior remains unchanged.

Using XenCenter, enable hidden objects (View menu -> Hidden objects). In the Resource pane, select the VM for which you want to obtain the UUID. The UUID will be displayed in the General tab.

In the Resource pane, select the resource pool master host (the first host in the list). The UUID will be displayed in the General tab. If you are not using a resource pool, select the VM host.

1. Open a console on the host and run the following command:

```
xe host-call-plugin host-uuid=<host-UUID> \
  plugin=coalesce-leaf fn=leaf-coalesce args:vm_uuid=<VM-UUID>
```

For example, if the VM UUID is 9bad4022-2c2d-dee6-abf5-1b6195b1dad5 and the host UUID is b8722062-de95-4d95-9baa-a5fe343898ea you would run this command:

```
xe host-call-plugin host-uuid=b8722062-de95-4d95-9baa-a5fe343898ea \
  plugin=coalesce-leaf fn=leaf-coalesce args:vm_uuid=9bad4022-2c2d-dee6-abf5-1b6195b1dad5
```

2. This command suspends the VM (unless it is already powered down), initiates the space reclamation process, and then resumes the VM.

Note:

Citrix recommends that, before executing the off-line coalesce tool, you shutdown or suspend the VM manually (using either XenCenter or the XenServer CLI). If you execute the coalesce tool on a VM that is running, the tool automatically suspends the VM, performs the required VDI coalesce operation(s), and resumes the VM.

If the Virtual Disk Images (VDIs) to be coalesced are on shared storage, you must execute the off-line coalesce tool on the pool master.

If the VDIs to be coalesced are on local storage, you must execute the off-line coalesce tool on the server to which the local storage is attached.

5.7.9. Adjusting the Disk IO Scheduler

For general performance, the default disk scheduler `noop` is applied on all new SR types. The `noop` scheduler provides the fairest performance for competing VMs accessing the same device. To apply disk QoS (see [Section 5.7.10, “Virtual Disk QoS Settings”](#)) it is necessary to override the default setting and assign the `cfq` disk scheduler to the SR. The corresponding PBD must be unplugged and re-plugged for the scheduler parameter to take effect. The disk scheduler can be adjusted using the following command:

```
xe sr-param-set other-config:scheduler=noop|cfq|anticipatory|deadline \
  uuid=<valid_sr_uuid>
```

Note:

This will not effect EqualLogic, NetApp or NFS storage.

5.7.10. Virtual Disk QoS Settings

Virtual disks have an optional I/O priority Quality of Service (QoS) setting. This setting can be applied to existing virtual disks using the `xe` CLI as described in this section.

In the shared SR case, where multiple hosts are accessing the same LUN, the QoS setting is applied to VBDs accessing the LUN from the same host. QoS is not applied across hosts in the pool.

Before configuring any QoS parameters for a VBD, ensure that the disk scheduler for the SR has been set appropriately. See [Section 5.7.9, “Adjusting the Disk IO Scheduler”](#) for details on how to adjust the scheduler. The scheduler parameter must be set to `cfq` on the SR for which the QoS is desired.

Note:

Remember to set the scheduler to `cfq` on the SR, and to ensure that the PBD has been re-plugged in order for the scheduler change to take effect.

The first parameter is `qos_algorithm_type`. This parameter needs to be set to the value **ionice**, which is the only type of QoS algorithm supported for virtual disks in this release.

The QoS parameters themselves are set with key/value pairs assigned to the `qos_algorithm_param` parameter. For virtual disks, `qos_algorithm_param` takes a `sched` key, and depending on the value, also requires a class key.

Possible values of `qos_algorithm_param:sched` are:

- `sched=rt` or `sched=real-time` sets the QoS scheduling parameter to real time priority, which requires a class parameter to set a value
- `sched=idle` sets the QoS scheduling parameter to idle priority, which requires no class parameter to set any value
- `sched=<anything>` sets the QoS scheduling parameter to best effort priority, which requires a class parameter to set a value

The possible values for `class` are:

- One of the following keywords: `highest`, `high`, `normal`, `low`, `lowest`
- an integer between 0 and 7, where 7 is the highest priority and 0 is the lowest, so that, for example, I/O requests with a priority of 5, will be given priority over I/O requests with a priority of 2.

To enable the disk QoS settings, you also need to set the `other-config:scheduler` to `cfq` and replug PBDs for the storage in question.

For example, the following CLI commands set the virtual disk's VBD to use real time priority 5:

```
xe vbd-param-set uuid=<vbd_uuid> qos_algorithm_type=ionice
xe vbd-param-set uuid=<vbd_uuid> qos_algorithm_params:sched=rt
xe vbd-param-set uuid=<vbd_uuid> qos_algorithm_params:class=5
xe sr-param-set uuid=<sr_uuid> other-config:scheduler=cfq
xe pbd-plug uuid=<pbd_uuid>
```


Chapter 6. Configuring VM Memory

When a VM is first created, it is allocated a fixed amount of memory. To improve the utilisation of physical memory in your XenServer environment, you can use Dynamic Memory Control (DMC), a memory management feature that enables dynamic reallocation of memory between VMs.

XenCenter provides a graphical display of memory usage in its **Memory** tab. This is described in the *XenCenter Help*.

Dynamic Memory Control (DMC) provides the following benefits:

- Memory can be added or removed without restart thus providing a more seamless experience to the user.
- When servers are full, DMC allows you to start more VMs on these servers, reducing the amount of memory allocated to the running VMs proportionally.

6.1. What is Dynamic Memory Control (DMC)?

XenServer DMC works by automatically adjusting the memory of running VMs, keeping the amount of memory allocated to each VM between specified minimum and maximum memory values, guaranteeing performance and permitting greater density of VMs per server.

Without DMC, when a server is full, starting further VMs will fail with "out of memory" errors: to reduce the existing VM memory allocation and make room for more VMs you must edit each VM's memory allocation and then reboot the VM. With DMC enabled, even when the server is full, XenServer will attempt to reclaim memory by automatically reducing the current memory allocation of running VMs within their defined memory ranges.

6.1.1. The Concept of Dynamic Range

For each VM the administrator can set a dynamic memory range – this is the range within which memory can be added/removed from the VM without requiring a reboot. When a VM is running the administrator can adjust the dynamic range. XenServer always guarantees to keep the amount of memory allocated to the VM within the dynamic range; therefore adjusting it while the VM is running may cause XenServer to adjust the amount of memory allocated to the VM. (The most extreme case is where the administrator sets the dynamic min/max to the same value, thus forcing XenServer to ensure that this amount of memory is allocated to the VM.) If new VMs are required to start on "full" servers, running VMs have their memory 'squeezed' to start new ones. The required extra memory is obtained by squeezing the existing running VMs proportionally within their pre-defined dynamic ranges

DMC allows you to configure dynamic minimum and maximum memory levels – creating a Dynamic Memory Range (DMR) that the VM will operate in.

- Dynamic Minimum Memory: A lower memory limit that you assign to the VM.
- Dynamic Higher Limit: An upper memory limit that you assign to the VM.

For example, if the Dynamic Minimum Memory was set at 512 MB and the Dynamic Maximum Memory was set at 1024 MB this would give the VM a Dynamic Memory Range (DMR) of 512 - 1024 MB, within which, it would operate. With DMC, XenServer *guarantees* at all times to assign each VM memory within its specified DMR.

6.1.2. The Concept of Static Range

Many Operating Systems that XenServer supports do not fully 'understand' the notion of dynamically adding or removing memory. As a result, XenServer must declare the maximum amount of memory that a VM will ever be asked to consume at the time that it boots. (This allows the guest operating system to size its page tables and other memory management structures accordingly.) This introduces the concept of a static memory range within XenServer. The static memory range cannot be adjusted while the VM is running. For a particular boot, the dynamic range is constrained such as to be always contained within this static range. Note that the static

minimum (the lower bound of the static range) is there to protect the administrator and is set to the lowest amount of memory that the OS can run with on XenServer.

Note:

Citrix advises not to change the static minimum level as this is set at the supported level per operating system – refer to the memory constraints table for more details.

By setting a static maximum level, higher than a dynamic max, means that in the future, if you need to allocate more memory to a VM, you can do so without requiring a reboot.

6.1.3. DMC Behavior

Automatic VM squeezing

- If DMC is not enabled, when hosts are full, new VM starts fail with ‘out of memory’ errors.
- If DMC is enabled, even when hosts are full, XenServer will attempt to reclaim memory (by reducing the memory allocation of running VMs within their defined dynamic ranges). In this way running VMs are squeezed proportionally at the same distance between the dynamic minimum and dynamic maximum for all VMs on the host

When DMC is enabled

- When the host's memory is plentiful - All running VMs will receive their Dynamic Maximum Memory level
- When the host's memory is scarce - All running VMs will receive their Dynamic Minimum Memory level.

When you are configuring DMC, remember that allocating only a small amount of memory to a VM can negatively impact it. For example, allocating too little memory:

- Using Dynamic Memory Control to reduce the amount of physical memory available to a VM may cause it to boot slowly. Likewise, if you allocate too little memory to a VM, it may start extremely slowly.
- Setting the dynamic memory minimum for a VM too low may result in poor performance or stability problems when the VM is starting.

6.1.4. How Does DMC Work?

Using DMC, it is possible to operate a guest virtual machine in one of two modes:

1. **Target Mode:** The administrator specifies a memory target for the guest. XenServer adjusts the guest's memory allocation to meet the target. Specifying a target is particularly useful in virtual server environments, and in any situation where you know exactly how much memory you want a guest to use. XenServer will adjust the guest's memory allocation to meet the target you specify.
2. **Dynamic Range Mode:** The administrator specifies a dynamic memory range for the guest; XenServer chooses a target from within the range and adjusts the guest's memory allocation to meet the target. Specifying a dynamic range is particularly useful in virtual desktop environments, and in any situation where you want XenServer to repartition host memory dynamically in response to changing numbers of guests, or changing host memory pressure. XenServer chooses a target from within the range and adjusts the guest's memory allocation to meet the target.

Note:

It is possible to change between target mode and dynamic range mode at any time for any running guest. Simply specify a new target, or a new dynamic range, and XenServer takes care of the rest.

6.1.5. Memory Constraints

XenServer allows administrators to use all memory control operations with any guest operating system. However, XenServer enforces the following memory property ordering constraint for all guests:

$$0 \leq \text{memory-static-min} \leq \text{memory-dynamic-min} \leq \text{memory-dynamic-max} \leq \text{memory-static-max}$$

XenServer allows administrators to change guest memory properties to any values that satisfy this constraint, subject to validation checks. However, in addition to the above constraint, Citrix supports only certain guest memory configurations for each supported operating system. The range of supported configurations depends on the guest operating system in use. XenServer does not prevent administrators from configuring guests to exceed the supported limit. However, customers are strongly advised to keep memory properties within the supported limits to avoid performance or stability problems. Refer to the *XenServer Virtual Machine User's Guide* for detailed guidelines on the minimum and maximum memory limits for each supported operating system.

Warning:

When configuring guest memory, Citrix advises NOT to exceed the maximum amount of physical memory addressable by your operating system. Setting a memory maximum that is greater than the operating system supported limit, may lead to stability problems within your guest.

In addition, customers should also note that the dynamic minimum should be greater than or equal to a quarter of the static maximum for all supported operating systems. Reducing the lower limit below the dynamic minimum could also lead to stability problems. Administrators are encouraged to calibrate the sizes of their VMs carefully, and make sure that their working set of applications functions reliably at dynamic-minimum.

6.2. xe CLI Commands

6.2.1. Display the Static Memory Properties of a VM

1. Find the uuid of the required VM:

```
xe vm-list
```

2. Note the uuid, and then run the command **param-name=memory-static**

```
xe vm-param-get uuid=<uuid> param-name=memory-static-{min,max}
```

For example, the following displays the static maximum memory properties for the VM with the uuid beginning ec77:

```
xe vm-param-get uuid= \
ec77a893-bff2-aa5c-7ef2-9c3acf0f83c0 \
param-name=memory-static-max;
268435456
```

This shows that the static maximum memory for this VM is 268435456 bytes (256MB).

6.2.2. Display the Dynamic Memory Properties of a VM

To display the dynamic memory properties, follow the procedure as above but use the command **param-name=memory-dynamic**:

1. Find the uuid of the required VM:

```
xe vm-list
```

2. Note the uuid, and then run the command **param-name=memory-dynamic**:

```
xe vm-param-get uuid=<uuid> param-name=memory-dynamic-{min,max}
```

For example, the following displays the dynamic maximum memory properties for the VM with uuid beginning ec77

```
xe vm-param-get uuid= \
ec77a893-bff2-aa5c-7ef2-9c3acf0f83c0 \
param-name=memory-dynamic-max;
134217728
```

This shows that the dynamic maximum memory for this VM is 134217728 bytes (128MB).

6.2.3. Updating Memory Properties

Warning:

It is essential that you use the correct ordering when setting the static/dynamic minimum/maximum parameters. In addition you must not invalidate the following constraint:

$$0 \leq \text{memory-static-min} \leq \text{memory-dynamic-min} \leq \text{memory-dynamic-max} \leq \text{memory-static-max}$$

Update the static memory range of a virtual machine:

```
xe vm-memory-static-range-set uuid=<uuid> min=<value> max=<value>
```

Update the dynamic memory range of a virtual machine:

```
xe vm-memory-dynamic-range-set \
  uuid=<uuid> min=<value> \
  max=<value>
```

Specifying a target is particularly useful in virtual server environments, and in any situation where you know exactly how much memory you want a guest to use. XenServer will adjust the guest's memory allocation to meet the target you specify. For example:

```
xe vm-target-set target=<value> vm=<vm-name>
```

Update all memory limits (static and dynamic) of a virtual machine:

```
xe vm-memory-limits-set \
  uuid=<uuid> \
  static-min=<value> \
  dynamic-min=<value> \
  dynamic-max=<value> static-max=<value>
```

Note:

- To allocate a specific amount memory to a VM that won't change, set the Dynamic Maximum and Dynamic Minimum to the same value.
- You cannot increase the dynamic memory of a VM beyond the static maximum.
- To alter the static maximum of a VM – you must shut down the VM.

6.2.4. Update Individual Memory Properties

Warning:

Citrix advises not to change the static minimum level as this is set at the supported level per operating system – refer to the memory constraints table for more details.

Update the dynamic memory properties of a VM.

1. Find the uuid of the required VM:

```
xe vm-list
```



2. Note the uuid, and then use the command **memory-dynamic-{min,max}=<value>**

```
xe vm-param-set uuid=<uuid>memory-dynamic-{min,max}=<value>
```

The following example changes the dynamic maximum to 128MB:

```
xe vm-param-set uuid=ec77a893-bff2-aa5c-7ef2-9c3acf0f83c0 memory-dynamic-max=128MiB
```

6.3. Upgrade Issues

After upgrading from Citrix XenServer 5.5, XenServer sets all VMs memory so that the dynamic minimum is equal to the dynamic maximum.

Chapter 7. XenServer Memory Usage

Two components contribute to the memory footprint of a XenServer host. Firstly, there is the memory consumed by the Xen hypervisor itself; secondly there is the memory consumed by the *Control Domain* of the host. Also known as 'Domain0', or 'dom0', the Control Domain is a secure, privileged Linux VM that runs the XenServer management toolstack (xapi). Besides providing XenServer management functions, the Control Domain also runs the driver stack that provides user created VM access to physical devices.

7.1. Control Domain Memory

The amount of memory allocated to the Control Domain is adjusted automatically and is based on the amount of physical memory on the physical host:

Host Memory (GB)	Control Domain Memory Allocated (MB)
20 - 24	752
24 - 48	2048
48 - 64	3072
64 - 1024	4096

Note:

The amount of memory reported in the **XenServer** field in XenCenter may exceed these values. This is because the amount reported includes the memory used by the Control Domain (dom0), the Xen hypervisor itself, and the crash kernel. The amount of memory used by the hypervisor will be larger for hosts with more memory.

7.1.1. Changing the Amount of Memory Allocated to the Control Domain

On hosts with smaller amounts of memory (for example, less than 16GB), you may wish to reduce the memory allocated to the Control Domain to below the default of 752MB. However, Citrix recommends that you **do not reduce the dom0 memory below 400MB**.

1. On the XenServer host, open a local shell and login as root.
2. Enter the following:

```
/opt/xensource/libexec/xen-cmdline --set-xen dom0_mem=<nn>M,max:<nn>M
```

Where <nn> represents the amount of memory, in megabytes, to be allocated to dom0.

3. Restart the XenServer host using XenCenter or the `reboot` command on the xs console.

When the host has rebooted, on the xs console, run the `free` command to verify the new memory settings.

Warning:

Increasing the amount of Control Domain memory will result in less memory being available to VMs.

The amount of memory allocated to the Control Domain can be increased beyond the amount shown above. However, this should only ever be carried out under the guidance of Citrix Support.

7.1.2. How Much Memory is Available to VMs

To find out how much host memory is actually available to be assigned to VMs, find the value of the *memory-free* field of the host, and then enter the **vm-compute-maximum-memory** command to get the actual amount of free memory that can be allocated to the VM, for example:

```
xe host-list uuid=<host_uuid> params=memory-free  
xe vm-compute-maximum-memory vm=<vm_name> total=<host_memory_free_value>
```

Chapter 8. Disaster Recovery and Backup

The XenServer Disaster Recovery (DR) feature is designed to allow you to recover virtual machines (VMs) and vApps from a catastrophic failure of hardware which disables or destroys a whole pool or site. For protection against single server failures, see [Section 3.9, “High Availability”](#)

Note:

You must be logged in as *root* or have the role of *Pool Operator* or higher to use this feature.

8.1. Understanding XenServer DR

XenServer DR works by storing all the information needed to recover your business-critical VMs and vApps on storage repositories (SRs) that are then replicated from your primary (production) environment to a backup environment. When a protected pool at your primary site goes down, the VMs and vApps in that pool can be recovered from the replicated storage and recreated on a secondary (DR) site, with minimal application or user downtime.

In the event of a disaster, the **Disaster Recovery** wizard in XenCenter can be used to interrogate this storage and import chosen VMs and vApps into a recovery pool. Once the VMs are running in the recovery pool, the recovery pool metadata is also replicated to allow any changes to VM settings to be populated back to the primary pool, should the primary pool be recovered. If the XenCenter wizard finds information for the same VM present in two or more places (for example, storage from the primary site, storage from the disaster recovery site and also in the pool that the data is to be imported into) then the wizard will ensure that only the most recent information per Virtual Machine is used.

The Disaster Recovery feature can be used both with XenCenter and the xe CLI. See [Section A.4.6, “Disaster Recovery \(DR\) Commands”](#) for details on these commands.

Tip:

You can also use the Disaster Recovery wizard to run test failovers for non-disruptive testing of your disaster recovery system. In a test failover, all the steps are the same as for failover, but the VMs and vApps are not started up after they have been recovered to the DR site, and cleanup is performed when the test is finished to remove all VMs, vApps and storage recreated on the DR site.

XenServer VMs consist of two components:

- Virtual disks that are being used by the VM, stored on configured storage repositories (SRs) in the pool where the VMs are located.
- Metadata describing the VM environment. This is all the information needed to recreate the VM if the original VM is unavailable or corrupted. Most metadata configuration data is written when the VM is created and is updated only when you make changes to the VM configuration. For VMs in a pool, a copy of this metadata is stored on every server in the pool.

In a DR environment, VMs are recreated on a secondary (DR) site from the pool metadata – configuration information about all the VMs and vApps in the pool. The metadata for each VM includes its name, description and Universal Unique Identifier (UUID), and its memory, virtual CPU and networking and storage configuration. It also includes the VM startup options – start order, delay interval and HA restart priority – which are used when restarting the VM in an HA or DR environment. For example, when recovering VMs during disaster recovery, the VMs within a vApp will be restarted in the DR pool in the order specified in the VM metadata, and with the specified delay intervals.

8.2. DR Infrastructure Requirements

To use XenServer DR, the appropriate DR infrastructure needs to be set up at both the primary and secondary sites:

- The storage used for both the pool metadata *and* the virtual disks used by the VMs must be replicated from your primary (production) environment to a backup environment. Storage replication, for example using mirroring, is best handled by your storage solution, and will vary from device to device.
- Once VMs and vApps have been recovered to a pool on your DR site and are up and running, the SRs containing the DR pool metadata and virtual disks must also be replicated to allow the recovered VMs and vApps to be restored back to the primary site (*failed back*) once the primary site is back online.
- The hardware infrastructure at your DR site does not have to match the primary site, but the XenServer environment must be at the same release and patch level, and sufficient resources should be configured in the target pool to allow all the failed over VMs to be recreated and started.

Warning:

The Disaster Recovery wizard does **not** control any Storage Array functionality.

Users of the Disaster Recovery feature **must** ensure that the metadata storage is, in some way replicated between the two sites. Some Storage Arrays contain “Mirroring” features to achieve the replication automatically, if these features are used then it is essential that the mirror functionality is disabled (“mirror is broken”) before Virtual Machines are restarted on the recovery site.

8.3. Deployment Considerations

Please review the following steps before enabling Disaster Recovery.

8.3.1. Steps to Take Before a Disaster

The following section describes the steps to take before disaster.

- Configure your VMs and vApps.
- Note how your VMs and vApps are mapped to SRs, and the SRs to LUNs. Take particular care with the naming of the `name_label` and `name_description` fields. Recovering VMs and vApps from replicated storage will be easier if your SRs are named in a way that captures how your VMs and vApps are mapped to SRs, and the SRs to LUNs.
- Arrange replication of the LUNs.
- Enable pool metadata replication to one or more SRs on these LUNs.

8.3.2. Steps to Take After a Disaster

The following section describes the steps to take after a disaster has occurred.

- Break any existing storage mirrors so that the recovery site has read/write access to the shared storage.
- Ensure that the LUNs you wish to recover VM data from are not attached to any other pool, or corruption may occur.
- If you want to protect the *recovery* site from a disaster, you must enable pool metadata replication to one or more SRs on the recovery site.

8.3.3. Steps to Take After a Recovery

The following section describes the steps to take after a successful recovery of data.

- Re-synchronise any storage mirrors.
- On the recovery site, shutdown down cleanly the VMs or vApps that you want to move back to the primary site.
- On the primary site, follow the same procedure as for the failover above, to failback selected VMs or vApps to the primary
- To protect the primary site against future disaster - you must re-enable pool metadata replication to one or more SRs on the replicated LUNs.

8.4. Enabling Disaster Recovery in XenCenter

This section describes how to enable Disaster Recovery in XenCenter. Use the **Configure DR** dialog box to identify storage repositories (SRs) where the metadata for a pool – configuration information about all the VMs and vApps in the pool – will be stored. This metadata will be updated whenever you make changes to VM or vApp configuration within the pool.

Note:

Disaster Recovery can only be enabled when using LVM over HBA or LVM over iSCSI. A small amount of space will be required on this storage for a new LUN which will contain the pool recovery information.

To do this:

1. On the primary site, select the pool that you want to protect. From the **Pool** menu, point to **Disaster Recovery**, and then click **Configure**.
2. Select up to 8 SRs where the pool metadata will be stored. A small amount of space will be required on this storage for a new LUN which will contain the pool recovery information.

Note:

Information for all VMs in the pool is stored, VMs do not need to be independently selected for protection.

3. Click **OK**. Your pool is now protected.

8.5. Recovering VMs and vApps in the Event of Disaster (Failover)

This section explains how to recover your VMs and vApps on the secondary (recovery) site.

1. In XenCenter select the secondary pool, and on the **Pool** menu, click **Disaster Recovery** and then **Disaster Recovery wizard**.

The **Disaster Recovery wizard** displays three recovery options: **Failover**, **Failback**, and **Test Failover**. To recover on to your secondary site, select **Failover** and then click **Next**.

Warning:

If you use Fibre Channel shared storage with LUN mirroring to replicate the data to the secondary site, before you attempt to recover VMs, **mirroring must be broken** so that the secondary site has Read/Write access.

2. Select the storage repositories (SRs) containing the pool metadata for the VMs and vApps that you want to recover.

By default, the list on this wizard page shows all SRs that are currently attached within the pool. To scan for more SRs, choose **Find Storage Repositories** and then select the storage type to scan for:

- To scan for all the available Hardware HBA SRs, select **Find Hardware HBA SRs**.
- To scan for software iSCSI SRs, select **Find Software iSCSI SRs** and then enter the target host, IQN and LUN details in the dialog box.

When you have selected the required SRs in the wizard, click **Next** to continue.

3. Select the VMs and vApps that you wish to recover and choose the appropriate **Power state after recovery** option to specify whether you want the wizard to start them up automatically as soon as they have been recovered, or whether you prefer to wait and start them up manually yourself after failover is complete.

Click **Next** to progress to the next wizard page and begin failover prechecks.

4. The wizard performs a number of prechecks before starting failover, for example, to ensure that all the storage required by the selected VMs and vApps is available. If any storage is missing at this point, you can click **Attach SR** on this page to find and attach the relevant SR.

Resolve any issues on the prechecks page, and then click **Failover** to begin the recovery process.

5. A progress page is displayed showing whether recovery was successful for each VM and vApp. Failover may take some time depending on the number of VMs and vApps you are recovering, as the metadata for the VMs and vApps are exported from the replicated storage. The VMs and vApps are recreated in the primary pool, the SRs containing the virtual disks are attached to the recreated VMs, and VMs are started, if specified.
6. When the failover is complete, click **Next** to see the summary report. Click **Finish** on the summary report page to close the wizard.

Once the primary site is available again, and you want to return to running your VMs on that site, work through the Disaster Recovery wizard again, but this time select the **Failback** option.

8.6. Restoring VMs and vApps to the Primary Site After Disaster (Failback)

This section explains how to restore VMs and vApps from replicated storage back to a pool on your primary (production) site when the primary site comes back up after a disaster event. To failback VMs and vApps to your primary site, use the Disaster Recovery wizard.

1. In XenCenter select the secondary pool, and on the **Pool** menu, click **Disaster Recovery** and then **Disaster Recovery wizard**.

The **Disaster Recovery wizard** displays three recovery options: **Failover**, **Failback**, and **Test Failover**. To restore VMs and vApps to your primary site, select **Failback** and then click **Next**.

Warning:

If you use Fibre Channel shared storage with LUN mirroring to replicate the data to the primary site, before you attempt to restore VMs, **mirroring must be broken** so that the primary site has Read/Write access.

2. Select the storage repositories (SRs) containing the pool metadata for the VMs and vApps that you want to recover.

By default, the list on this wizard page shows all SRs that are currently attached within the pool. To scan for more SRs, choose **Find Storage Repositories** and then select the storage type to scan for:

- To scan for all the available Hardware HBA SRs, select **Find Hardware HBA SRs**.
- To scan for software iSCSI SRs, select **Find Software iSCSI SRs** and then enter the target host, IQN and LUN details in the dialog box.

When you have selected the required SRs in the wizard, click **Next** to continue.

3. Select the VMs and vApps that you wish to restore and choose the appropriate **Power state after recovery** option to specify whether you want the wizard to start them up automatically as soon as they have been recovered, or whether you prefer to wait and start them up manually yourself after failback is complete.

Click **Next** to progress to the next wizard page and begin failback prechecks.

4. The wizard performs a number of pre-checks before starting failback, for example, to ensure that all the storage required by the selected VMs and vApps is available. If any storage is missing at this point, you can click **Attach SR** on this page to find and attach the relevant SR.

Resolve any issues on the prechecks page, and then click **Failback** to begin the recovery process.

5. A progress page is displayed showing whether recovery was successful for each VM and vApp. Failback may take some time depending on the number of VMs and vApps you are restoring, as the metadata for the VMs and vApps are exported from the replicated storage. The VMs and vApps are recreated in the primary pool, the SRs containing the virtual disks are attached to the recreated VMs, and VMs are started, if specified.
6. When the failback is complete, click **Next** to see the summary report. Click **Finish** on the summary report page to close the wizard.

8.7. Test Failover

Failover testing is an essential component in disaster recovery planning. You can use the Disaster Recovery wizard to perform non-disruptive testing of your disaster recovery system. During a test failover operation, all the steps are the same as for failover, but instead of being started after they have been recovered to the DR site, the VMs and vApps are placed in a paused state. At the end of a test failover operation, all VMs, vApps and storage recreated on the DR site are automatically removed. After initial DR configuration, and after you make significant configuration changes in a DR-enabled pool, we recommend that you verify that failover still works correctly by performing a test failover.

To perform a test failover of VMs and vApps to a secondary site

1. In XenCenter select the secondary pool, and on the **Pool** menu, click **Disaster Recovery** to open the **Disaster Recovery Wizard**.
2. Select **Test Failover** and then click **Next**.

Note:

If you use Fibre Channel shared storage with LUN mirroring to replicate the data to the secondary site, before you attempt to recover data, mirroring must be broken so that the secondary site has Read/Write access.

3. Select the storage repositories (SRs) containing the pool metadata for the VMs and vApps that you want to recover.

By default, the list on this wizard page shows all SRs that are currently attached within the pool. To scan for more SRs, choose **Find Storage Repositories** and then select the storage type to scan for:

- To scan for all the available Hardware HBA SRs, select **Find Hardware HBA SRs**.
- To scan for software iSCSI SRs, select **Find Software iSCSI SRs** and then enter the target host, IQN and LUN details in the dialog box.

When you have selected the required SRs in the wizard, click **Next** to continue.

4. Select the VMs and vApps that you wish to recover then click Next to progress to the next wizard page and begin failover prechecks.
5. Before beginning the test failover process, the wizard performs a number of pre-checks, for example, to ensure that all the storage required by the selected VMs and vApps is available.
 - **Check that storage is available** If any storage is missing, you can click **Attach SR** on this page to find and attach the relevant SR.
 - **Check that HA is not enabled on the target DR pool.** To avoid having the same VMs running on both the primary and DR pools, HA must be disabled on the secondary pool to ensure that the recovered VMs and vApps are not started up automatically by HA after recovery. To disable HA on the secondary pool, you can simply click **Disable HA** on the this page. (If HA is disabled at this point, it will be enabled again automatically at the end of the test failover process.)

Resolve any issues on the pre-checks page, and then click **Failover** to begin the test failover.

6. A progress page is displayed showing whether recovery was successful for each VM and vApp. Failover may take some time depending on the number of VMs and vApps you are recovering, as the metadata for the VMs and vApps are recovered from the replicated storage. The VMs and vApps are recreated in the DR pool, the SRs containing the virtual disks are attached to the recreated VMs.

The recovered VMs are placed in a paused state: they will not be started up on the secondary site during a test failover.

7. After you are satisfied that the test failover was performed successfully, click **Next** in the wizard to have the wizard clean up on the DR site:
 - VMs and vApps that were recovered during the test failover will be removed.
 - Storage that was recovered during the test failover will be detached.
 - If HA on the DR pool was disabled at the prechecks stage to allow the test failover to take place, it will be enabled again automatically.

The progress of the cleanup process is displayed in the wizard.

8. Click **Finish** to close the wizard.

8.8. vApps

A **vApp** is logical group of one or more related Virtual Machines (VMs) which can be started up as a single entity in the event of a disaster. When a vApp is started, the VMs contained within the vApp will start in a user predefined order, to allow VMs which depend upon one another to be automatically sequenced. This means that an administrator no longer has to manually sequence the startup of dependent VMs should a whole service require restarting (for instance in the case of a software update). The VMs within the vApp do not have to reside on one host and will be distributed within a pool using the normal rules. The vApp functionality is particularly useful in the Disaster Recovery situation where an Administrator may choose to group all VMs which reside on the same Storage Repository, or which relate to the same Service Level Agreement (SLA).

Creating vApps

To group VMs together in a vApp follow the procedure:

1. Select the pool and, on the **Pool** menu, click **Manage vApps**. This displays the **Manage vApps** window.
2. Enter a name for the vApp, and optionally a description, and then click **Next**.

You can choose any name you like, but a descriptive name is usually best. Although it is advisable to avoid having multiple vApps with the same name, it is not a requirement, and XenCenter does not enforce any uniqueness constraints on vApp names. It is not necessary to use quotation marks for names that include spaces.

3. Choose which VMs to include in the new vApp, and then click **Next**.

You can use the search box to list only VMs with names that include the specified string.

4. Specify the startup sequence for the VMs in the vApp, and then click **Next**.

Value	Description
Start Order	Specifies the order in which individual VMs will be started up within the vApp, allowing certain VMs to be restarted before others. VMs with a start order value of 0 (zero) will be started first, then VMs with a start order value of 1, then VMs with a start order value of 2, and so on.
Attempt to start next VM after	This is a delay interval that specifies how long to wait after starting the VM before attempting to start the next group of VMs in the startup sequence, that is, VMs with a lower start order.

5. On the final page of the wizard, you can review the vApp configuration. Click **Previous** to go back and modify any settings, or **Finish** to create the new vApp and close the wizard.

Note:

A vApp can span multiple servers in a single pool, but cannot span across several pools.

8.8.1. Using the Manage vApps dialog box in XenCenter

XenCenter's **Manage vApps** dialog box allows you to create, delete and modify vApps, start and shutdown vApps, and import and export vApps within the selected pool. When you select a vApp in the list, the VMs it contains are listed in the details pane on the right. See the XenCenter online help for further details. Press **F1** or click **Help** to display the Help.

8.9. Backing Up and Restoring XenServer Hosts and VMs

Citrix recommends that, whenever possible, you leave the installed state of XenServer hosts unaltered. That is, do not install any additional packages or start additional services on XenServer hosts, and treat them as if they



are appliances. The best way to restore, then, is to reinstall XenServer host software from the installation media. If you have multiple XenServer hosts, the best approach is to configure a TFTP server and appropriate answerfiles for this purpose (see the *XenServer Installation Guide*).

For VMs, the best approach is to install backup agents on them, just as if they were standard physical servers. For Windows VMs, as of this release we have tested CA [BrightStor ARCserve Backup](#), and Symantec [NetBackup](#) and [Backup Exec](#).

For more information about backup tools tested, best practices, and backups in general, see the [Citrix Knowledge Base](#).

Citrix recommends that you frequently perform as many of the following backup procedures as possible to recover from possible server and/or software failure.

To backup pool metadata

1. Run the command:

```
xe pool-dump-database file-name=<backup>
```

2. Run the command:

```
xe pool-restore-database file-name=<backup> dry-run=true
```

This command checks that the target machine has an appropriate number of appropriately named NICs, which is required for the backup to succeed.

To backup host configuration and software

- Run the command:

```
xe host-backup host=<host> file-name=<hostbackup>
```

Note:

- Do not create the backup in the control domain.
- This procedure may create a large backup file.
- To complete a restore you have to reboot to the original install CD.
- This data can only be restored to the original machine.

To backup a VM

1. Ensure that the VM to be backed up is offline.
2. Run the command:

```
xe vm-export vm=<vm_uuid> filename=<backup>
```

Note:

This backup also backs up all of the VM data. When importing a VM, you can specify the storage mechanism to use for the backed up data.

Warning:

Because this process backs up all of the VM data, it can take some time to complete.

To backup VM metadata only

- Run the command:

```
xe vm-export vm=<vm_uuid> filename=<backup> metadata=true
```

8.9.1. Backing up Virtual Machine metadata

XenServer hosts use a database on each host to store metadata about VMs and associated resources such as storage and networking. When combined with storage repositories, this database forms the complete view of all VMs available across the pool. Therefore it is important to understand how to backup this database in order to recover from physical hardware failure and other disaster scenarios.

This section first describes how to backup metadata for single-host installations, and then for more complex pool setups.

8.9.1.1. Backing up single host installations

Use the CLI to backup the pool database. To obtain a consistent pool metadata backup file, run **pool-dump-database** on the XenServer host and archive the resulting file. The backup file will contain sensitive authentication information about the pool, so ensure it is securely stored.

To restore the pool database, use the **xe pool-restore-database** command from a previous dump file. If your XenServer host has died completely, then you must first do a fresh install, and then run the **pool-restore-database** command against the freshly installed XenServer host.

After a restoration of the pool database, some VMs may still be registered as being Suspended, but if the storage repository with their suspended memory state (defined in the *suspend-VDI-uuid* field) was a local SR, it will no longer be available since the host has been reinstalled. To reset these VMs back to the Halted state so that they can be started up again, use the **xe vm-shutdown vm=vm_name -force** command, or use the **xe vm-reset-powerstate vm=<vm_name> -force** command.

Warning:

XenServer hosts restored using this method will have their UUIDs preserved. If you restore to a different physical machine while the original XenServer host is still running, there will be duplicate UUIDs. The main observable effect of this will be that XenCenter will refuse to connect to the second XenServer host. Pool database backup is not the recommended mechanism for cloning physical hosts; use the automated installation support for that (see the *XenServer Installation Guide*).

8.9.1.2. Backing up pooled installations

In a pool scenario, the master host provides an authoritative database that is synchronously mirrored to all the member hosts in the pool. This provides a degree of built-in redundancy to a pool; the master can be replaced by any member since each of them has an accurate version of the pool database. Please refer to the *XenServer Administrator's Guide* for more information on how to transition a member into becoming a master host.

This level of protection may not be sufficient; for example, if your shared storage containing the VM data is backed up in multiple sites, but your local server storage (containing the pool metadata) is not. To fully recreate a pool given just a set of shared storage, you must first backup the **pool-dump-database** file on the master host, and archive this file.

Subsequently restoring this backup on a brand new set of hosts

1. Install a fresh set of XenServer hosts from the installation media, or if applicable, network boot from your TFTP server.
2. Use the **xe pool-restore-database** on the host designated to be the new master.
3. Run the **xe host-forget** command on the new master to remove the old member machines.
4. Use the **xe pool-join** command on the member hosts to connect them to the new pool.

8.9.2. Backing up XenServer hosts

This section describes the XenServer host control domain backup and restore procedures. These procedures do *not* back up the storage repositories that house the VMs, but only the privileged control domain that runs Xen and the XenServer agent.

Note:

Because the privileged control domain is best left as installed, without customizing it with other packages, Citrix recommends that you set up a network boot environment to cleanly perform a fresh installation from the XenServer media as a recovery strategy. In many cases you will not need to backup the control domain at all, but just save the pool metadata (see [Section 8.9.1, “Backing up Virtual Machine metadata”](#)). This backup method should always be considered complementary to backing up the pool metadata.

Using the `xe` commands **host-backup** and **host-restore** is another approach that you can take. The `xe host-backup` command archives the active partition to a file you specify, and the `xe host-restore` command extracts an archive created by `xe host-backup` over the currently inactive disk partition of the host. This partition can then be made active by booting off the installation CD and choosing to restore the appropriate backup.

After completing the above steps and rebooting the host, you must ensure that the VM meta-data is restored to a consistent state. This can be achieved by running `xe pool-restore-database` on `/var/backup/pool-database-DATE`. This file is created by `xe host-backup` using `xe pool-dump-database` command before archiving the running filesystem, to snapshot a consistent state of the VM metadata.

To back up a XenServer host

- On a remote host with enough disk space, run the command:

```
xe host-backup file-name=<filename> -h <hostname> -u root -pw <password>
```

This creates a compressed image of the control domain file system in the location specified by the *file-name* argument.

To restore a running XenServer host

- If you want to restore a XenServer host from a specific backup, run the following command while the XenServer host is up and reachable:

```
xe host-restore file-name=<filename> -h <hostname> -u root -pw <password>;
```

This command restores the compressed image back to the hard disk of the XenServer host on which the command is run (not the host on which *filename* resides). In this context “restore” is something of a misnomer, as the word usually suggests that the backed-up state has been put fully in place. The restore command here only unpacks the compressed backup file and restores it to its normal form, but it is written to another partition (`/dev/sda2`) and does *not* overwrite the current version of the filesystem.

- To use the restored version of the root filesystem, reboot the XenServer host using the XenServer installation CD and select the **Restore from backup** option.

After the restore from backup is completed, reboot the XenServer host and it will start up from the restored image.

Finally, restore the VM meta-data using

```
xe pool-restore-database file-name=/var/backup/pool-database-*
```

Note:

Restoring from a backup as described here does *not* destroy the backup partition.

Restarting a crashed XenServer host

1. If your XenServer host has crashed and is not reachable anymore, you need to use the XenServer installation CD to do an upgrade install. When that is completed, reboot the machine and make sure your host is reachable with XenCenter or remote CLI.
2. Then proceed with [Section 8.9.2, “Backing up XenServer hosts”](#) above.

8.9.3. Backing up VMs

VMs are best backed up using standard backup tools running on them individually. For Windows VMs, we have tested CA BrightStor ARCserve Backup.

8.10. VM Snapshots

Important:

Virtual Machine Protection and Recovery (VMPR) and its functionality has been removed from XenServer 7.0 and later versions. Applications, code, or usage that depend on VMPR will not function in XenServer 7.0 and later versions. The VM snapshot functionality and other functionalities that depend on VM snapshot (other than VMPR) remain unaffected. Refer to [CTX137335](#) for details.

XenServer provides a convenient snapshotting mechanism that can take a snapshot of a VM storage and metadata at a given time. Where necessary, IO is temporarily halted while the snapshot is being taken to ensure that a self-consistent disk image can be captured.

Snapshot operations result in a snapshot VM that is similar to a template. The VM snapshot contains all the storage information and VM configuration, including attached VIFs, allowing them to be exported and restored for backup purposes. Snapshots are supported on all storage types, though for the LVM-based storage types the storage repository must have been upgraded if it was created on a previous version of XenServer and the volume must be in the default format (type=raw volumes cannot be snapshotted).

The snapshotting operation is a 2 step process:

- Capturing metadata as a template.
- Creating a VDI snapshot of the disk(s).

Three types of VM snapshots are supported: regular, quiesced, and snapshot with memory

8.10.1. Regular Snapshots

Regular snapshots are crash consistent and can be performed on all VM types, including Linux VMs.

8.10.2. Quiesced Snapshots

Quiesced snapshots take advantage of the Windows Volume Shadow Copy Service (VSS) to generate application consistent point-in-time snapshots. The VSS framework helps VSS-aware applications (for example Microsoft Exchange or Microsoft SQL Server) flush data to disk and prepare for the snapshot before it is taken.

Quiesced snapshots are therefore safer to restore, but can have a greater performance impact on a system while they are being taken. They may also fail under load so more than one attempt to take the snapshot may be required.

XenServer supports quiesced snapshots on:

- Windows Server 2012 R2 Server Core
- Windows Server 2012 R2



- Windows Server 2012
- Windows Server 2008 R2
- Windows Server 2008 (32/64-bit)
- Windows Server 2003 (32/64-bit)

Windows 8.1, Windows 8, Windows 7, Windows 2000, and Windows Vista are not supported. For further detail on quiesced snapshots, see [Section 8.10.9.3, “Advanced Notes for Quiesced Snapshots”](#).

8.10.3. Snapshots with memory

In addition to saving the VMs memory (storage) and metadata, snapshots with memory also save the VMs state (RAM). This can be useful if you are upgrading or patching software, or want to test a new application, but also want the option to be able to get back to the current, pre-change state (RAM) of the VM. Reverting back to a snapshot with memory, does not require a reboot of the VM.

You can take a snapshot with memory of a running or suspended VM via the XenAPI, the xe CLI, or by using XenCenter.

8.10.4. Creating a VM Snapshot

Before taking a snapshot, see the section called “Preparing for Cloning a Windows VM Using VSS” in *XenServer Virtual Machine User's Guide*, and for information about any special operating system-specific configuration and considerations to take into account, refer to the section called “Preparing to clone a Linux VM”.

Firstly, ensure that the VM is running or suspended so that the memory status can be captured. The simplest way to select the VM on which the operation is to be performed is by supplying the argument `vm=<name>` or `vm=<vm uuid>`.

Run the **vm-snapshot** and **vm-snapshot-with-quiesce** commands to take a snapshot of a VM.

```
xe vm-snapshot vm=<vm uuid> new-name-label=<vm_snapshot_name>
xe vm-snapshot-with-quiesce vm=<vm uuid> new-name-label=<vm_snapshot_name>
```

8.10.5. Creating a snapshot with memory

Run the **vm-checkpoint** command, giving a descriptive name for the snapshot with memory, so that you can identify it later:

```
xe vm-checkpoint vm=<vm uuid> new-name-label=<name of the checkpoint>
```

When XenServer has completed creating the snapshot with memory - its uuid will be displayed.

For example:

```
xe vm-checkpoint vm=2d1d9a08-e479-2f0a-69e7-24a0e062dd35 \
  new-name-label=example_checkpoint_1
b3c0f369-59a1-dd16-ecd4-a1211df29886
```

A snapshot with memory requires at least 4MB of disk space per disk, plus the size of the RAM, plus around 20% overhead. So a checkpoint with 256MB RAM would require approximately 300MB of storage.

Note:



During the checkpoint creation process, the VM is paused for a brief period of time, and cannot be used during this period.

8.10.6. To list all of the snapshots on a XenServer pool

Run the **snapshot-list** command:

```
xe snapshot-list
```

This lists all of the snapshots in the XenServer pool.

8.10.7. To list the snapshots on a particular VM

You will need to know the uuid of the particular VM; to do this run the **vm-list** command.

```
xe vm-list
```

This displays a list of all VMs and their UUIDs. For example:

```
xe vm-list
uuid ( R0): 116dd310-a0ef-a830-37c8-df41521ff72d
name-label ( RW): Windows Server 2003 (1)
power-state ( R0): halted

uuid ( R0): 96fde888-2a18-c042-491a-014e22b07839
name-label ( RW): Windows XP SP3 (1)
power-state ( R0): running

uuid ( R0): dff45c56-426a-4450-a094-d3bba0a2ba3f
name-label ( RW): Control domain on host
power-state ( R0): running
```

VMs can also be specified by filtering the full list of VMs on the values of fields.

For example, specifying **power-state=halted** will select all VMs whose power-state field is equal to 'halted'. Where multiple VMs are matching, the option **--multiple** must be specified to perform the operation. The full list of fields that can be matched can be obtained by the command **xe vm-list params=all**.

Locate the required VM and then enter the following:

```
xe snapshot-list snapshot-of=<vm uuid>
```

For example:

```
xe snapshot-list snapshot-of=2d1d9a08-e479-2f0a-69e7-24a0e062dd35
```

This lists the snapshots currently on that VM:

```
uuid ( R0): d7eefb03-39bc-80f8-8d73-2ca1bab7dcff
name-label ( RW): Regular
name-description ( RW):
snapshot-of ( R0): 2d1d9a08-e479-2f0a-69e7-24a0e062dd35
snapshot-time ( R0): 20090914T15:37:00Z

uuid ( R0): 1760561d-a5d1-5d5e-2be5-d0dd99a3b1ef
name-label ( RW): Snapshot with memory
name-description ( RW):
snapshot-of ( R0): 2d1d9a08-e479-2f0a-69e7-24a0e062dd35
snapshot-time ( R0): 20090914T15:39:45Z
```

8.10.8. Restoring a VM to its previous state

Ensure that you have the uuid of the snapshot that you wish to revert to, and then run the **snapshot-revert** command:

To do this:

1. Run the **snapshot-list** command to find the UUID of the snapshot or checkpoint that you want to revert to:

```
xe snapshot-list
```

2. Note the uuid of the snapshot, and then run the following command to revert:

```
xe snapshot-revert snapshot-uuid=<snapshot uuid>
```

For example:

```
xe snapshot-revert snapshot-uuid=b3c0f369-59a1-dd16-ecd4-a1211df29886
```

After reverting to a checkpoint, the VM will be suspended.

Note:

If there is insufficient disk space available to thickly provision the snapshot, it will not be possible to restore to the snapshot until the current disk's state has been freed. If this occurs, retry the operation.

Note:

It is possible to revert to any snapshot. Existing snapshots and checkpoints are not deleted during the revert operation.

8.10.8.1. Deleting a snapshot

Ensure that you have the UUID of the checkpoint or snapshot that you wish to remove, and then run the following command:

1. Run the **snapshot-list** command to find the UUID of the snapshot or checkpoint that you want to revert to:

```
xe snapshot-list
```

2. Note the UUID of the snapshot, and then run the **snapshot-uninstall** command to remove it:

```
xe snapshot-uninstall snapshot-uuid=<snapshot-uuid>
```

3. This command alerts you to the VM and VDI's that will be deleted. Type **yes** to confirm.

For example:

```
xe snapshot-uninstall snapshot-uuid=1760561d-a5d1-5d5e-2be5-d0dd99a3b1ef
The following items are about to be destroyed
VM : 1760561d-a5d1-5d5e-2be5-d0dd99a3b1ef (Snapshot with memory)
VDI: 11a4aa81-3c6b-4f7d-805a-b6ea02947582 (0)
VDI: 43c33fe7-a768-4612-bf8c-c385e2c657ed (1)
VDI: 4c33c84a-a874-42db-85b5-5e29174fa9b2 (Suspend image)
Type 'yes' to continue
yes
All objects destroyed
```

If you only want to remove the metadata of a checkpoint or snapshot, run the following command:



```
xe snapshot-destroy snapshot-uuid=<snapshot-uuid>
```

For example:

```
xe snapshot-destroy snapshot-uuid=d7eefb03-39bc-80f8-8d73-2ca1bab7dcff
```

8.10.9. Snapshot Templates

8.10.9.1. Creating a template from a snapshot

You can create a VM template from a snapshot, however its memory state will be removed.

To do this:

1. Use the command **snapshot-copy** and specify a **new-name-label** for the template:

```
xe snapshot-copy new-name-label=<vm-template-name> \
  snapshot-uuid=<uuid of the snapshot>
```

For example:

```
xe snapshot-copy new-name-label=example_template_1
  snapshot-uuid=b3c0f369-59a1-dd16-ecd4-a1211df29886
```

Note:

This creates a template object in the SAME pool. This template exists in the XenServer database for the current pool only.

2. To verify that the template has been created, run the command **template-list**:

```
xe template-list
```

This will list all of the templates on the XenServer host.

8.10.9.2. Exporting a snapshot to a template

When you export a VM snapshot, a complete copy of the VM (including disk images) is stored as a single file on your local machine, with a **.xva** file extension.

To do this:

1. Use the command **snapshot-export-to-template** to create a new template file:

```
xe snapshot-export-to-template snapshot-uuid=<snapshot-uuid> \
  filename=<template- filename>
```

For example:

```
xe snapshot-export-to-template snapshot-uuid=b3c0f369-59a1-dd16-ecd4-a1211df29886 \
  filename=example_template_export
```

The VM export/import feature can be used in a number of different ways:

- As a convenient backup facility for your VMs. An exported VM file can be used to recover an entire VM in the event of disaster.
- As a way of quickly copying a VM, for example, a special-purpose server configuration that you use many times. You simply configure the VM the way you want it, export it, and then import it to create copies of your original VM.
- As a simple method for moving a VM to another server.

For further information on the use of templates refer to the Creating VMs chapter in the *XenServer Virtual Machine User's Guide* and also the **Managing VMs** section in the XenCenter Help.

8.10.9.3. Advanced Notes for Quiesced Snapshots

Note:

Do not forget to install the Xen VSS provider in the Windows guest in order to support VSS. This is done using the `install-XenProvider.cmd` script provided with the XenServer Tools. For more information, see the *XenServer Virtual Machine User's Guide*.

In general, a VM can only access VDI snapshots (not VDI clones) of itself using the VSS interface. There is a flag that can be set by the XenServer administrator whereby adding an attribute of `snapmanager=true` to the VM `other-config` allows that VM to import snapshots of VDIs from other VMs.

Warning:

This opens a security vulnerability and should be used with care. This feature allows an administrator to attach VSS snapshots using an in-guest transportable snapshot ID as generated by the VSS layer to another VM for the purposes of backup.

VSS quiesce timeout: the Microsoft VSS quiesce period is set to a non-configurable value of 10 seconds, and it is quite probable that a snapshot may not be able to complete in time. If, for example the XAPI daemon has queued additional blocking tasks such as an SR scan, the VSS snapshot may timeout and fail. The operation should be retried if this happens.

Note:

The more VBDs attached to a VM, the more likely it is that this timeout may be reached. Citrix recommends attaching no more than 2 VBDs to a VM to avoid reaching the timeout. However, there is a workaround to this problem. The probability of taking a successful VSS based snapshot of a VM with more than 2 VBDs can be increased manifold, if all the VDIs for the VM are hosted on different SRs.

VSS snapshot all the disks attached to a VM: in order to store all data available at the time of a VSS snapshot, the XAPI manager will snapshot all disks and the VM metadata associated with a VM that can be snapshotted using the XenServer storage manager API. If the VSS layer requests a snapshot of only a subset of the disks, a full VM snapshot will not be taken.

vm-snapshot-with-quiesce produces bootable snapshot VM images: To achieve this, the XenServer VSS hardware provider makes snapshot volumes writable, including the snapshot of the boot volume.

VSS snap of volumes hosted on dynamic disks in the Windows Guest : The **vm-snapshot-with-quiesce CLI** and the XenServer VSS hardware provider do not support snapshots of volumes hosted on dynamic disks on the Windows VM.

Note:

Do not forget to install the Xen VSS provider in the Windows guest in order to support VSS. This is done using the `install-XenProvider.cmd` script provided with the XenServer Tools. For more information, see *XenServerVirtual Machine User's Guide*.

8.11. Coping with machine failures

This section provides details of how to recover from various failure scenarios. All failure recovery scenarios require the use of one or more of the backup types listed in [Section 8.9, "Backing Up and Restoring XenServer Hosts and VMs"](#).

8.11.1. Member failures

In the absence of HA, master nodes detect the failures of members by receiving regular heartbeat messages. If no heartbeat has been received for 600 seconds, the master assumes the member is dead. There are two ways to recover from this problem:

- Repair the dead host (e.g. by physically rebooting it). When the connection to the member is restored, the master will mark the member as alive again.
- Shutdown the host and instruct the master to forget about the member node using the **xe host-forget** CLI command. Once the member has been forgotten, all the VMs which were running there will be marked as offline and can be restarted on other XenServer hosts. Note it is very important to ensure that the XenServer host is actually offline, otherwise VM data corruption might occur. Be careful not to split your pool into multiple pools of a single host by using **xe host-forget**, since this could result in them all mapping the same shared storage and corrupting VM data.

Warning:

- If you are going to use the forgotten host as a XenServer host again, perform a fresh installation of the XenServer software.
- Do not use **xe host-forget** command if HA is enabled on the pool. Disable HA first, then forget the host, and then re-enable HA.

When a member XenServer host fails, there may be VMs still registered in the *running* state. If you are sure that the member XenServer host is definitely down, use the **xe vm-reset-powerstate** CLI command to set the power state of the VMs to `halted`. See [Section A.4.26.28, “vm-reset-powerstate”](#) for more details.

Warning:

Incorrect use of this command can lead to data corruption. Only use this command if absolutely necessary.

Before you can start VMs on another XenServer host, you are also required to release the locks on VM storage. Each disk in an SR can only be used by one host at a time, so it is key to make the disk accessible to other XenServer hosts once a host has failed. To do so, run the following script on the pool master for each SR that contains disks of any affected VMs:

```
/opt/xensource/sm/resetvdis.py <host_UUID> <SR_UUID> [master]
```

You need only supply the third string ("master") if the failed host was the SR master — pool master or XenServer host using local storage — at the time of the crash.

Warning:

Be absolutely sure that the host is down before executing this command. Incorrect use of this command can lead to data corruption.

If you attempt to start a VM on another XenServer host before running the script above, then you will receive the following error message: `VDI <UUID> already attached RW`.

8.11.2. Master failures

Every member of a resource pool contains all the information necessary to take over the role of master if required. When a master node fails, the following sequence of events occurs:

1. If HA is enabled, another master is elected automatically.
2. If HA is not enabled, each member will wait for the master to return.

If the master comes back up at this point, it re-establishes communication with its members, and operation returns to normal.

If the master is really dead, choose one of the members and run the command **xe pool-emergency-transition-to-master** on it. Once it has become the master, run the command **xe pool-recover-slaves** and the members will now point to the new master.

If you repair or replace the server that was the original master, you can simply bring it up, install the XenServer host software, and add it to the pool. Since the XenServer hosts in the pool are enforced to be homogeneous, there is no real need to make the replaced server the master.

When a member XenServer host is transitioned to being a master, you should also check that the default pool storage repository is set to an appropriate value. This can be done using the **xe pool-param-list** command and verifying that the *default-SR* parameter is pointing to a valid storage repository.

8.11.3. Pool failures

In the unfortunate event that your entire resource pool fails, you will need to recreate the pool database from scratch. Be sure to regularly back up your pool-metadata using the **xe pool-dump-database** CLI command (see [Section A.4.15.3, “pool-dump-database”](#)).

To restore a completely failed pool

1. Install a fresh set of hosts. Do not pool them up at this stage.
2. For the host nominated as the master, restore the pool database from your backup using the **xe pool-restore-database** (see [Section A.4.15.11, “pool-restore-database”](#)) command.
3. Connect to the master host using XenCenter and ensure that all your shared storage and VMs are available again.
4. Perform a pool join operation on the remaining freshly installed member hosts, and start up your VMs on the appropriate hosts.

8.11.4. Coping with Failure due to Configuration Errors

If the physical host machine is operational but the software or host configuration is corrupted:

To restore host software and configuration

1. Run the command:

```
xe host-restore host=<host> file-name=<hostbackup>
```

2. Reboot to the host installation CD and select **Restore from backup**.

8.11.5. Physical Machine failure

If the physical host machine has failed, use the appropriate procedure listed below to recover.

Warning:

Any VMs which were running on a previous member (or the previous host) which has failed will still be marked as **Running** in the database. This is for safety-- simultaneously starting a VM on two different hosts would lead to severe disk corruption. If you are sure that the machines (and VMs) are offline you can reset the VM power state to **Halting**:

```
xe vm-reset-powerstate vm=<vm_uuid> --force
```

VMs can then be restarted using XenCenter or the CLI.

Replacing a failed master with a still running member

1. Run the commands:

```
xe pool-emergency-transition-to-master
xe pool-recover-slaves
```

2. If the commands succeed, restart the VMs.

To restore a pool with all hosts failed

1. Run the command:

```
xe pool-restore-database file-name=<backup>
```

Warning:

This command will only succeed if the target machine has an appropriate number of appropriately named NICs.

2. If the target machine has a different view of the storage (for example, a block-mirror with a different IP address) than the original machine, modify the storage configuration using the **pbd-destroy** command and then the **pbd-create** command to recreate storage configurations. See [Section A.4.13, "PBD Commands"](#) for documentation of these commands.
3. If you have created a new storage configuration, use **pbd-plug** or **Storage > Repair Storage Repository** menu item in XenCenter to use the new configuration.
4. Restart all VMs.

To restore a VM when VM storage is not available

1. Run the command:

```
xe vm-import filename=<backup> metadata=true
```

2. If the metadata import fails, run the command:

```
xe vm-import filename=<backup> metadata=true --force
```

This command will attempt to restore the VM metadata on a 'best effort' basis.

3. Restart all VMs.

Chapter 9. Monitoring and Managing XenServer

XenServer provides detailed monitoring of performance metrics, including CPU, memory, disk, network, C-state/ P-state information and storage. Where appropriate, these metrics are available on a per host and a per VM basis. These metrics are available directly, or can be accessed and viewed graphically in XenCenter or other third party applications.

XenServer also provides system and performance alerts. Alerts are notifications that occur in response to selected system events, or when CPU, memory usage, network, storage throughput, or VM disk activity go over a specified threshold on a managed host, VM, or storage repository. You can configure these using the xe CLI or by using XenCenter, to create notifications based on any of the available Host or VM performance metrics see [Section 9.2, “Alerts”](#).

9.1. Monitoring XenServer Performance

Customers can monitor the performance of their XenServer hosts and Virtual Machines (VMs) using the metrics exposed through Round Robin Databases (RRDs). These metrics can be queried over HTTP or through the RRD2CSV tool; in addition, XenCenter uses this data to produce system performance graphs. See [Section 9.1.5, “Using RRDs”](#) and [Section 9.1.3, “Analyzing and Visualizing Metrics in XenCenter”](#).

The following tables list all of the available Host and VM metrics.

Note:

Latency over a period of time is defined as the average latency of operations during that period.

The availability and utility of certain metrics are SR and CPU dependent.

9.1.1. Available Host Metrics

Metric Name	Description	Condition	XenCenter Name
avgqu_sz_<sr-uuid-short>	Average I/O queue size (requests).	At least one plugged VBD in SR <sr> on the host	<sr> Queue Size
cpu<cpu>-C<cstate>	Time CPU <cpu> spent in C-state <cstate> in milliseconds.	C-state exists on CPU	CPU <cpu> C-state <cstate>
cpu<cpu>-P<pstate>	Time CPU <cpu> spent in P-state <pstate> in milliseconds.	P-state exists on CPU	CPU <cpu> P-state <pstate>
cpu<cpu>	Utilisation of physical CPU <cpu> (fraction). Enabled by default.	pCPU <cpu> exists	CPU <cpu>
cpu_avg	Mean utilisation of physical CPUs (fraction). Enabled by default.	None	Average CPU

Metric Name	Description	Condition	XenCenter Name
inflight_<sr-uuid-short>	Number of I/O requests currently in flight. Enabled by default.	At least one plugged VBD in SR <sr> on the host	<sr> Inflight Requests
io_throughput_read_<sr-uuid-short>	Data read from SR (MiB/s).	At least one plugged VBD in SR <sr> on the host	<sr> Read Throughput
io_throughput_write_<sr-uuid-short>	Data written to the SR (MiB/s).	At least one plugged VBD in SR <sr> on the host	<sr> Write Throughput
io_throughput_total_<sr-uuid-short>	All SR I/O (MiB/s).	At least one plugged VBD in SR <sr> on the host	<sr> Total Throughput
iops_read_<sr-uuid-short>	Read requests per second.	At least one plugged VBD in SR <sr> on the host	<sr> Read IOPS
iops_write_<sr-uuid-short>	Write requests per second.	At least one plugged VBD in SR <sr> on the host	<sr> Write IOPS
iops_total_<sr-uuid-short>	I/O requests per second.	At least one plugged VBD in SR <sr> on the host	<sr> Total IOPS
iowait_<sr-uuid-short>	Percentage of the time waiting for I/O.	At least one plugged VBD in SR <sr> on the host	<sr> IO Wait
latency_<sr-uuid-short>	Average I/O latency (milliseconds).	At least one plugged VBD in SR <sr> on the host	<sr> Latency

Metric Name	Description	Condition	XenCenter Name
loadavg	Domain0 loadavg. Enabled by default	None	Control Domain Load
memory_free_kib	Total amount of free memory (KiB). Enabled by default.	None	Free Memory
memory_reclaimed	Host memory reclaimed by squeeze (B).	None	Reclaimed Memory
memory_reclaimed_max	Host memory that could be reclaimed by squeezed (B).	None	Potential Reclaimed Memory
memory_total_kib	Total amount of memory (KiB) in the host. Enabled by default.	None	Total Memory
network/latency	Interval in seconds between the last two heartbeats transmitted from the local host to all Online hosts. Disabled by default.	HA Enabled	Network Latency
statefile/<t>/latency	Turn-around time in seconds of the latest State-File access from the local host. Disabled by default.	HA Enabled	HA Statefile Latency
pif_<pif>_rx	Bytes per second received on physical interface <pif>. Enabled by default.	PIF<pif>exists	<xencenter-pif-name> Receive (see note)
pif_<pif>_tx	Bytes per second sent on physical interface <pif>. Enabled by default.	PIF<pif>exists	<xencenter-pif-name> Send (see note)
pif_<pif>_rx_errors	Receive errors per second on physical interface <pif>. Disabled by default.	PIF<pif>exists	<xencenter-pif-name> Receive Errors (see note)
pif_<pif>_tx_errors	Transmit errors per second on physical interface <pif> . Disabled by default	PIF<pif>exists	<xencenter-pif-name>Send Errors (see note)
pif_aggr_rx	Bytes per second received on all physical interfaces. Enabled by default.	None	Total NIC Receive
pif_aggr_tx	Bytes per second sent on all physical interfaces. Enabled by default.	None	Total NIC Send
sr_<sr>_cache_size	Size in bytes of the IntelliCache SR. Enabled by default.	IntelliCache Enabled	IntelliCache Cache Size

Metric Name	Description	Condition	XenCenter Name
sr_<sr>_cache_hits	Cache hits per second. Enabled by default.	IntelliCache Enabled	IntelliCache Cache Hits
sr_<sr>_cache_misses	Cache misses per second. Enabled by default.	IntelliCache Enabled	IntelliCache Cache Misses
xapi_allocation_kib	Memory (KiB) allocation done by the xapi daemon. Enabled by default.	None	Agent Memory Allocation
xapi_free_memory_kib	Free memory (KiB) available to the xapi daemon. Enabled by default.	None	Agent Memory Free
xapi_healthcheck/latency_healt	Turn-around time in seconds of the latest xapi status monitoring call on the local host. Disabled by default	HA Enabled	XenServer Healthcheck Latency
xapi_live_memory_kib	Live memory (KiB) used by xapi daemon. Enabled by default.	None	Agent Memory Live
xapi_memory_usage_kib	Total memory (KiB) allocated used by xapi daemon. Enabled by default.	None	Agent Memory Usage

9.1.2. Available VM Metrics

Metric Name	Description	Condition	XenCenter Name
cpu<cpu>	Utilisation of vCPU <cpu> (fraction). Enabled by default	vCPU <cpu> exists	CPU <cpu>
memory	Memory currently allocated to VM (Bytes).Enabled by default	None	Total Memory
memory_target	Target of VM balloon driver (Bytes). Enabled by default	None	Memory target
memory_internal_free	Memory used as reported by the guest agent (KiB). Enabled by default	None	Free Memory
runstate_fullrun	Fraction of time that all VCPUs are running.	None	VCPUs full run
runstate_full_contention	Fraction of time that all VCPUs are runnable (i.e., waiting for CPU)	None	VCPUs full contention
runstate_concurrency_hazard	Fraction of time that some VCPUs are running and some are runnable	None	VCPUs concurrency hazard
runstate_blocked	Fraction of time that all VCPUs are blocked or offline	None	VCPUs idle

Metric Name	Description	Condition	XenCenter Name
runstate_partial_run	Fraction of time that some VCPUs are running, and some are blocked	None	VCPUs partial run
runstate_partial_contention	Fraction of time that some VCPUs are runnable and some are blocked	None	VCPUs partial contention
vbd_<vbd>_write	Writes to device <vbd> in bytes per second. Enabled by default	VBD <vbd> exists	Disk <vbd> Write
vbd_<vbd>_read	Reads from device <vbd> in bytes per second. Enabled by default.	VBD <vbd> exists	Disk <vbd> Read
vbd_<vbd>_write_latency	Writes to device <vbd> in microseconds.	VBD <vbd> exists	Disk <vbd> Write Latency
vbd_<vbd>_read_latency	Reads from device <vbd> in microseconds.	VBD <vbd> exists	Disk <vbd> Read Latency
vbd <vbd>_iops_read	Read requests per second.	A least one plugged VBD for non-ISO VDI on the host	Disk <vbd> Read IOPS
vbd <vbd>_iops_write	Write requests per second.	A least one plugged VBD for non-ISO VDI on the host	Disk <vbd> Write IOPS
vbd <vbd>_iops_total	I/O requests per second.	A least one plugged VBD for non-ISO VDI on the host	Disk <vbd> Total IOPS
vbd <vbd>_iowait	Percentage of time waiting for I/O.	At least one plugged VBD for non-ISO VDI on the host	Disk <vbd> IO Wait
vbd <vbd>_inflight	Number of I/O requests currently in flight.	At least one plugged VBD for non-ISO VDI on the host	Disk <vbd> Inflight Requests

Metric Name	Description	Condition	XenCenter Name
vbd <vbd>_avgqu_sz	Average I/O queue size.	At least one plugged VBD for non-ISO VDI on the host	Disk <vbd> Queue Size
vif_<vif>_rx	Bytes per second received on virtual interface number <vif>. Enabled by default.	VIF <vif> exists	<vif> Receive
vif_<vif>_tx	Bytes per second transmitted on virtual interface <vif>. Enabled by default.	VIF <vif> exists	<vif> Send
vif_<vif>_rx_errors	Receive errors per second on virtual interface <vif>. Enabled by default.	VIF <vif> exists	<vif> Receive Errors
vif_<vif>_tx_errors	Transmit errors per second on virtual interface <vif>. Enabled by default.	VIF <vif> exists	<vif> Send Errors

Note:

<xencenter -pif -name> can be any of the following:

NIC <pif>	if <pif> contains pif_eth#, where # is 0-9
<pif>	if <pif> contains pif_eth#.# or pif_xenbr# or pif_bond#
<Internal> Network <pif>	if <pif> contains pif_xapi# (note that <Internal> appears as is)
TAP <tap>	if <pif> contains pif_tap#
xapi Loopback	if <pif> contains pif_lo

9.1.3. Analyzing and Visualizing Metrics in XenCenter

The **Performance** tab in XenCenter provides real time monitoring of performance statistics across resource pools as well as graphical trending of virtual and physical machine performance. By default, graphs showing CPU, memory, network and disk I/O are included on the Performance tab, but you can add additional metrics, change the appearance of the existing graphs or create additional ones. See [Section 9.1.3.1, “Configuring Performance Graphs”](#).

- You can view up to 12 months of performance data and zoom in to take a closer look at activity spikes.
- XenCenter can generate performance alerts when CPU, memory usage, network I/O, storage I/O or disk I/O usage go over a specified threshold on a managed server, VM or storage repository. Refer to [Section 9.2.1, “Using XenCenter to View Alerts”](#)

Note:

You must install the XenServer Tools (paravirtualized drivers) in order to see full VM performance data .

9.1.3.1. Configuring Performance Graphs

To Add A New Graph

1. On the **Performance** tab, click **Actions** and then **New Graph**. The New Graph dialog box will be displayed.
2. In the **Name** field, enter a name for the graph.
3. From the list of **Datasources**, select the check boxes for the datasources you want to include in the graph.
4. Click **Save**.


To Edit An Existing Graph

1. Navigate to the **Performance** tab, and select the graph that you would like to modify.
2. Right-click on the graph and select **Actions**, or click the **Actions** button. Then select **Edit Graph**.
3. On the graph details window, make the necessary changes, and click **OK**.


9.1.3.1.1. Configuring the Graph Type

Data on the performance graphs can be displayed as lines or as areas:

Line Graph:



Area Graph:



To change the Graph Type

1. On the **Tools** menu, click **Options** and select **Graphs**.
2. To view performance data as a line graph, click on the **Line graph** radio button
3. To view performance data as an area graph, click on the **Area graph** radio button
4. Click **OK** to save your changes.

Comprehensive details for configuring and viewing XenCenter Performance Graphs can be found in the *XenCenter Help* in the section *Monitoring System Performance*.

9.1.4. Configuring Metrics

Note:

C-states and P-states are power management features of some processors. The range of states available is dependent on the physical capabilities of the host, as well power management configuration.

Both host and VM commands return the following:

- A full description of the data source
- The units applied to the metric
- The range of possible values that may be used

For example:



```
name_label: cpu0-C1
name_description: Proportion of time CPU 0 spent in C-state 1
enabled: true
standard: true
  min: 0.000
  max: 1.000
units: Percent
```

Enabling a Specific Metric

Most metrics are enabled and collected by default, to enable those that are not, enter the following:

```
xe host-data-source-record data-source=<metric name> host=<hostname>
```

Disabling a Specific Metric

You may not wish to collect certain metrics on a regular basis. To disable a previously enabled metric, enter the following:

```
xe host-data-source-forget data-source=<metric name> host=<hostname>
```

Displaying a List of Currently Enabled Host Metrics

To list the host metrics currently being collected, enter the following:

```
xe host-data-source-list host=<hostname>
```

Displaying a List of Currently Enabled VM Metrics

To host the VM metrics currently being collected, enter the following:

```
xe vm-data-source-list vm=<vm_name>
```

9.1.5. Using RRDs

In order to store performance metrics, XenServer uses (RRDs). These consist of multiple Round Robin Archives (RRAs) in a fixed size database.

Each archive in the database samples its particular metric on a specified granularity:

- Every 5 seconds for the duration of 10 minute
- Every minute for the past two hours
- Every hour for the past week
- Every day for the past year

The sampling that takes place every 5 seconds records actual data points, however the following RRAs use Consolidation Functions (CF) instead. The CFs supported by XenServer are:

- AVERAGE
- MIN
- MAX

RRDs exist for individual VMs (including dom0) and the XenServer host. VM RRDs are stored on the host on which they run, or the pool master when not running. Therefore the location of a VM must be known in order to retrieve the associated performance data.

For detailed information on how to use XenServer RRDs, refer to the Citrix Developer Network article: [Using XenServer RRDs](#).

9.1.5.1. Analysing RRDs using HTTP

RRDs can be downloaded over HTTP from the XenServer host specified using the HTTP handler registered at /host_rrd or /vm_rrd. Both addresses require authentication either by HTTP auth, or by providing a valid XenAPI session references as a query argument. For example:

Downloading a Host RRD

```
wget http://<server>/host_rrd?session_id=OpaqueRef:<SESSION HANDLE>>
```

Downloading a VM RRD

```
wget http://<server>/vm_rrd?session_id=OpaqueRef:<SESSION HANDLE>>&uuid=<VM UUID>>
```

Both of these calls will download XML in a format that can be imported into the rrdtool for analysis, or parsed directly.

For detailed information on how to use XenServer RRDs with HTTP, refer to the Citrix Developer Network article: [Using XenServer RRDs](#).

9.1.5.2. Analysing RRDs using rrd2csv

In addition to viewing performance metrics in XenCenter the **rrd2csv** tool logs RRDs to Comma Separated Value (CSV) format. Man and help pages are provided. To display the rrd2csv tool man or help pages run the following:

```
man rrd2csv
```

or

```
rrd2csv --help
```

Note:

Where multiple options are used, they should be supplied individually. For example: to return both the UUID and the name-label associated with a VM or a host, rrd2csv should be called as shown below:

```
rrd2csv -u -n
```

The UUID returned is unique and suitable as a primary key, however the name-label of an entity may not necessarily be unique.

The man page (`rrd2csv --help`) is the definitive help text of the tool.

9.2. Alerts

You can configure XenServer to generate alerts based on any of the available Host or VM Metrics. In addition, XenServer provides preconfigured alarms that trigger when hosts undergo certain conditions and states. You can view these alerts using XenCenter or the xe CLI.

9.2.1. Using XenCenter to View Alerts

You can view different types of alerts in XenCenter by clicking **Notifications** and then **Alerts**. The Alerts view displays various types of alerts, including:

- [Section 9.2.1.1, “XenCenter Performance Alerts”](#)
- [Section 9.2.1.2, “XenCenter Alerts”](#)
- [Section 9.2.1.3, “XenCenter Software Update Alerts”](#)

9.2.1.1. XenCenter Performance Alerts

Performance alerts can be generated when CPU, memory usage, network, storage throughput, or VM disk activity exceeds a specified threshold on a managed host, VM, or storage repository (SR).

By default, the alert repeat interval is set to 60 minutes, it can be modified if required. Alerts will be displayed on the **Alerts** page in the **Notifications** area in XenCenter. In addition, you can configure XenCenter to send an email for any specified performance alerts along with other serious system alerts. In addition, any customised alert configured using the xe CLI, will also display on the **Alerts** page in XenCenter.

Each alert has a corresponding priority/severity level. You can modify these and optionally choose to receive an email when the alert is triggered. The default alert priority/severity is set at 3.

Priority	Name	Description	Default Email Alert
1	Critical	Take action now or data may be permanently lost/corrupted.	Yes
2	Major	Take action now or some services may fail.	Yes
3	Warning	Take action now or a service may suffer.	Yes
4	Minor	Notice that something just improved.	No
5	Information	Day-to-day information (VM Start, Stop, Resume etc.)	No
?	Unknown	Unknown error	No

9.2.1.1.1. To Configure Performance Alerts

To configure performance alerts:

1. In the **Resources** pane, select the relevant host, VM, or SR, then click the **General** tab and then **Properties**.
2. Click the **Alerts** tab.
 - **CPU** alerts for a host or VM: select the **Generate CPU** usage alerts check box, then set the CPU usage and time threshold that will trigger the alert
 - **Network** alerts for a host or VM: select the **Generate network usage alerts** check box, then set the network usage and time threshold that will trigger the alert.
 - **Memory** alerts for a host: select the **Generate memory usage alerts** check box, and then set the memory usage and time threshold that will trigger the alert.
 - **Disk usage** alerts for a VM: select the **Generate disk usage alerts** check box, then set the disk usage and time threshold that will trigger the alert.
 - **Storage throughput** alerts for a SR: select the **Generate storage throughput alerts** check box, then set the storage throughput and time threshold that will trigger the alert.

Note:

Physical Block Devices (PBD) represent the interface between a specific XenServer host and an attached SR. When the total read/write SR throughput activity on a PBD exceeds the threshold you have specified, alerts will be generated on the host connected to the PBD. Unlike other XenServer host alerts, this must be configured on the SR.

- To change the alert repeat interval, enter the number of minutes in the **Alert repeat interval** box. Once an alert threshold has been reached and an alert generated, another alert will not be generated until after the alert repeat interval has elapsed.
3. Click **OK** to save your changes.

For comprehensive details on how to view, filter and configure severities for performance alerts, see the XenCenter online help.

9.2.1.2. XenCenter Alerts

The following table displays the system events/conditions that will trigger an alert to be displayed on the **Alerts** page in XenCenter.

Name	Priority/ Severity	Description
license_expires_soon	2	XenServer License agreement expires soon.
ha-statefile_lost	2	Lost contact with the HA Storage Repository, take action soon.
ha-heartbeat_approaching_timeout	5	HA approaching timeout, host may reboot unless action is taken.
ha_statefile_approaching_timeout	5	HA approaching timeout, host may reboot unless action is taken.
ha-xapi_healthcheck_approaching_timeout	5	HA approaching timeout, host may reboot unless action is taken.
ha_network_bonding_error	3	Potential service loss. Loss of network that sends HA heartbeat
ha_pool_overcommitted	3	Potential service loss. HA is unable to guarantee protection for configured VMs
ha_poor_drop_in_plan_exists_for	3	HA coverage has dropped, more likely to fail, no loss present yet
ha_protected_vm_restart_failed	2	Service Loss. HA was unable to restart a protected VM
ha_host_failed	3	HA detected that a host failed
ha_host_was_fenced	4	HA rebooted a host to protect against VM corruption
redo_log_healthy	4	The xapi redo log has recovered from a previous error
redo_log_broken	3	The xapi redo log has encountered an error
ip_configured_pif_can_unplug	3	An IP configured NIC can be unplugged by xapi when using HA, possibly leading to HA failure.
host_sync_data_failed	3	Failed to synchronize XenServer performance statistics
host_clock_skew_detected	3	The host clock is not synchronized with other hosts in the pool
host_clock_went_backwards	1	The host clock is corrupted
pool_master_transition	4	A new host has been specified as Pool Master
pbd_plug_failed_on_server_start	3	The host failed to connect to Storage at boot time.

Name	Priority/ Severity	Description
auth_external_init_failed	2	The host failed to enable external AD authentication
auth_external_pool_non-homogeneous	2	Hosts in a pool have different AD authentication configuration.
multipath_period_alert	3	A path to a SR has failed or recovered.
bond-status-changed	3	A link in a bond has disconnected or reconnected

9.2.1.3. XenCenter Software Update Alerts

Alert	Description
XenCenter old	XenServer expects a newer version but can still connect to the current version
XenCenter out of date	XenCenter is too old to connect to XenServer
XenServer out of date	XenServer is an old version that the current XenCenter cannot connect to
License expired alert	XenServer license has expired
Missing IQN alert	XenServer uses iSCSI storage but the host IQN is blank
Duplicate IQN alert	XenServer uses iSCSI storage, and there are duplicate host IQNs

9.2.2. Configuring Performance Alerts Using the xe CLI

Note:

Triggers for alerts are checked at a minimum interval of five minutes (this avoids placing excessive load on the system to check for these conditions and reporting of false positives); setting an alert repeat interval smaller than this will result in the alerts still being generated at the five minute minimum interval.

The performance monitoring **perfmon** tool runs once every 5 minutes and requests updates from XenServer which are averages over 1 minute. These defaults can be changed in `/etc/sysconfig/perfmon`.

The **perfmon** tool reads updates every 5 minutes of performance variables running on the same host. These variables are separated into one group relating to the host itself, and a group for each VM running on that host. For each VM and XenServer host, **perfmon** reads the parameter `other-config:perfmon` and uses this string to determine which variables to monitor, and under which circumstances to generate a message.

For example, the following shows an example of configuring a VM "CPU usage" alert by writing an XML string into the parameter `other-config:perfmon`:

```
xe vm-param-set uuid=<vm_uuid> other-config:perfmon=\
'<config>
  <variable>
    <name value="cpu_usage"/>
    <alarm_trigger_level value="0.5"/>
  </variable>
</config>'
```

Note:

Multiple *<variable>* nodes are allowed

After setting the new configuration, use the following command to refresh **perfmon** for each host:

```
xe host-call-plugin host=<host_uuid> plugin=perfmon fn=refresh
```

If this is not done, there will be a delay before the new configuration takes effect, since by default, **perfmon** checks for new configuration once every thirty minutes. This default can be changed in `/etc/sysconfig/perfmon`.

Valid VM Elements

name

The name of the variable (no default). If the name value is either `cpu_usage`, `network_usage`, or `disk_usage` the `rrd_regex` and `alarm_trigger_sense` parameters are not required as defaults for these values will be used.

alarm_priority

The priority of the alerts generated (default 3).

alarm_trigger_level

The level of value that triggers an alert (no default).

alarm_trigger_sense

`high` if *alarm_trigger_level* is a maximum value otherwise `low` if the *alarm_trigger_level* is a minimum value (the default `high`).

alarm_trigger_period

The number of seconds that values (above or below the alert threshold) can be received before an alert is sent (the default is 60).

alarm_auto_inhibit_period

The number of seconds this alarm will be disabled after an alarm is sent (the default is 3600).

consolidation_fn

Combines variables from `rrd_updates` into one value. For **cpu-usage** the default is `average`, for **fs_usage** the default is `get_percent_fs_usage` and for all others - `sum`.

rrd_regex

Matches the names of variables from **xe vm-data-sources-list uuid=<vm_uuid>**, to compute performance values. This parameter has defaults for the named variables:

- `cpu_usage`
- `network_usage`
- `disk_usage`

If specified, the values of all items returned by **xe vm-data-source-list** whose names match the specified regular expression will be consolidated using the method specified as the *consolidation_fn*.

Valid Host Elements

name

The name of the variable (no default).

alarm_priority

The priority of the alerts generated (default 3).

alarm_trigger_level

The level of value that triggers an alarm (no default).

alarm_trigger_sense

`high` if *alarm_trigger_level* is a maximum value otherwise `low` if the *alarm_trigger_level* is a minimum value. (default `high`)

alarm_trigger_period

The number of seconds that values (above or below the alert threshold) can be received before an alarm is sent (default 60).

alarm_auto_inhibit_period

The number of seconds that the alert is disabled for after an alert is sent. (default 3600).

consolidation_fn

Combines variables from **rrd_updates** into one value (default sum - or average)

rrd_regex

A regular expression to match the names of variables returned by the **xe vm-data-source-list uuid=<vm_uuid>** command that should be used to compute the statistical value. This parameter has defaults for the following named variables:

- cpu_usage
- network_usage
- memory_free_kib
- sr_io_throughput_total_xxxxxxx (where xxxxxxxx is the first eight characters of the SR-UUID).

Note:

SR Throughput: Storage throughput alerts must be configured on the SR rather than the host.

For example:

```
xe sr-param-set uuid=<sr_uuid> other-config:perfmon=\
'<config>
  <variable>
    <name value="sr_io_throughput_total_per_host"/>
    <alarm_trigger_level value="0.01"/>
  </variable>
</config>'
```

9.2.2.1. Generic Example Configuration

The following example shows a generic configuration:

```
<config>
  <variable>
    <name value="NAME_CHOSEN_BY_USER"/>
    <alarm_trigger_level value="THRESHOLD_LEVEL_FOR_ALARM"/>
    <alarm_trigger_period value="RAISE_ALARM_AFTER_THIS_MANY_SECONDS_OF_BAD_VALUES"/>
    <alarm_priority value="PRIORITY_LEVEL"/>
    <alarm_trigger_sense value="HIGH_OR_LOW"/>
    <alarm_auto_inhibit_period value="MINIMUM_TIME_BETWEEN_ALARMS_FROM_THIS_MONITOR"/>
    <consolidation_fn value="FUNCTION_FOR_COMBINING_VALUES"/>
    <rrd_regex value="REGULAR_EXPRESSION_TO_CHOOSE_DATASOURCE_METRIC"/>
  </variable>

  <variable>
    ...
  </variable>

  ...
</config>
```

9.3. Configuring Email Alerts

Customers can configure XenServer to send email notifications when alerts are generated by XenServer hosts. This can be done either by using XenCenter, or by using the xe Command Line Interface (CLI).

9.3.1. Enabling Email Alerts Using XenCenter

1. In the Resources pane, right-click on a pool and select **Properties**.
2. In the **Properties** window, select **Email Options**.
3. Select the **Send email alert notifications** check box and enter the email address and SMTP server details.

Note:

Enter the details of an SMTP server which does **not** require authentication

9.3.2. Enabling Email Alerts using the xe CLI

Important:

When using XenCenter or the xe CLI to enable email notifications, customers should enter the details of an SMTP server, which does not require authentication. Emails sent through SMTP servers which require authentication will not be delivered.

To configure this, specify the email address and SMTP server:

```
xe pool-param-set uuid=<pool_uuid> other-config:mail-destination=<joe.bloggs@domain.tld>
xe pool-param-set uuid=<pool_uuid> other-config:ssmtp-mailhub=<smtp.domain.tld[:port]>
```

You can also specify the minimum value of the priority (known as 'severity' in XenCenter) field in the message before the email will be sent:

```
xe pool-param-set uuid=<pool_uuid> other-config:mail-max-priority=<level>
```

The default priority level is 4.

Note:

Some SMTP servers only forward mails with addresses that use FQDNs. If you find that emails are not being forwarded it may be for this reason, in which case you can set the server hostname to the FQDN so this is used when connecting to your mail server.

9.3.2.1. Sending Email Alerts Through Authenticated SMTP Servers

The **mail-alarm** utility in XenServer uses sSMTP to send email notifications. Before sending email notifications, the mail-alarm utility looks for the configuration file, `mail-alarm.conf`. If the configuration file exists, the contents of the file will be used to configure sSMTP, otherwise the details available in the XAPI database (as configured using XenCenter or the xe CLI) will be used to send email alerts. To send email notifications through authenticated SMTP servers, customers should create a `mail-alarm.conf` file in `/etc/` with the following contents:

```
root=postmaster
authUser=<username>
authPass=<password>
mailhub=<server address>:<port>
```

Note:

This configuration file will be used for all alerts generated by XenServer hosts.

9.3.2.1.1. Additional Configuration Options

Each SMTP server can differ slightly in its setup and may require additional configuration. Refer to the following extract from the `ssmtp.conf` man page for the correct syntax and available options:

**NAME**

ssmtp.conf – ssmtp configuration file

DESCRIPTION

ssmtp reads configuration data from /etc/ssmtp/ssmtp.conf. The file contains keyword-argument pairs, one per line. Lines starting with '#' and empty lines are interpreted as comments.

The possible keywords and their meanings are as follows (both are case-insensitive):

Root

The user that gets all mail for userids less than 1000. If blank, address rewriting is disabled.

Mailhub

The host to send mail to, in the form host | IP_addr port [:port]. The default port is 25.

RewriteDomain

The domain from which mail seems to come. For user authentication.

Hostname

The full qualified name of the host. If not specified, the host is queried for its hostname.

FromLineOverride

Specifies whether the From header of an email, if any, may override the default domain. The default is "no".

UseTLS

Specifies whether ssmtp uses TLS to talk to the SMTP server. The default is "no".

UseSTARTTLS

Specifies whether ssmtp does a EHL0/STARTTLS before starting SSL negotiation. See RFC 2487.

TLSCert

The file name of an RSA certificate to use for TLS, if required.

AuthUser

The user name to use for SMTP AUTH. The default is blank, in which case SMTP AUTH is not used.

AuthPass

The password to use for SMTP AUTH.

AuthMethod

The authorization method to use. If unset, plain text is used. May also be set to "cram-md5".

9.4. Custom Fields and Tags

XenCenter supports the creation of tags and custom fields, which allows for organization and quick searching of VMs, storage and so on. See the *XenCenter Help* for more information.

9.5. Custom Searches

XenCenter supports the creation of customized searches. Searches can be exported and imported, and the results of a search can be displayed in the navigation pane. See the *XenCenter Help* for more information.

9.6. Determining throughput of physical bus adapters

For FC, SAS and iSCSI HBAs you can determine the network throughput of your PBDs using the following procedure.

To determine PBD throughput

1. List the PBDs on a host.
2. Determine which LUNs are routed over which PBDs.
3. For each PBD and SR, list the VBDs that reference VDIs on the SR.
4. For all active VBDs that are attached to VMs on the host, calculate the combined throughput.

For iSCSI and NFS storage, check your network statistics to determine if there is a throughput bottleneck at the array, or whether the PBD is saturated.

Chapter 10. Troubleshooting

10.1. Support

Citrix provides two forms of support: free, self-help support on the [Citrix Support](#) website and paid-for Support Services, which you can purchase from the Support site. With Citrix Technical Support, you can open a Support Case online or contact the support center by phone if you experience technical difficulties.

The [Citrix Knowledge Center](#) hosts a number of resources that may be helpful to you if you experience odd behavior, crashes, or other problems. Resources include: Forums, Knowledge Base articles, White Papers, product documentation, hotfixes and other updates .

If you experience technical difficulties with a XenServer host, this chapter is meant to help you solve the problem if possible and, failing that, describes where the application logs are located and other information that can help your Citrix Solution Provider and Citrix track and resolve the issue.

Troubleshooting of installation issues is covered in the *XenServer Installation Guide*. Troubleshooting of Virtual Machine issues is covered in the *XenServer Virtual Machine User's Guide*.

Important:

We recommend that you follow the troubleshooting information in this chapter solely under the guidance of your Citrix Solution Provider or Citrix Support.

Note:

In some support cases, serial console access is required for debug purposes. Therefore, when setting up a XenServer configuration, it is recommended that serial console access is configured. For hosts that do not have physical serial port (such as a Blade server) or where suitable physical infrastructure is not available, customers should investigate if an embedded management device, such as Dell DRAC or HP iLO can be configured. For more information on setting up serial console access, see CTX121442 , [How to Set Up a Serial Cable for Troubleshooting on XenServer](#).

10.2. Health Check

The Health Check feature enables customers to automate the process of generating and uploading the server status report to Citrix Insight Services (CIS) and receive CIS analysis reports in XenCenter.

When you connect any eligible pool to XenCenter, you will be prompted to enable Health Check for the pool. During the enrollment process, you can specify the schedule at which you would like to automatically upload the server status report to CIS, enter XenServer credentials that will be used to establish a connection with the pool, and authenticate your uploads with CIS. After the pool is successfully enrolled to Health Check, you will receive notifications in XenCenter regarding the health of the pool. This enables you to proactively monitor the health of the XenServer systems based on the report that CIS generates.

Requirements

In order to use the Health Check feature:

- All hosts in the pool should be running XenServer 7.0
- You should connect to the XenServer pool using XenCenter shipped with XenServer 7.0
- XenCenter should have access to the internet
- The Health Check Service should be installed and running on the XenCenter machine.
- If using Active Directory (AD), you should have a Pool Operator or a higher role



For detailed information about Health Check and for step-by-step instructions on enrolling a pool to Health Check, see the *XenCenter Help*.

10.3. XenServer host logs

XenCenter can be used to gather XenServer host information. Click on **Server Status Report** in the **Tools** menu to open the **Server Status Report** wizard. You can select from a list of different types of information (various logs, crash dumps, etc.). The information is compiled and downloaded to the machine that XenCenter is running on. For details, see the XenCenter Help.

Additionally, the XenServer host has several CLI commands to make it simple to collate the output of logs and various other bits of system information using the utility **xen-bugtool**. Use the **xe** command **host-bugreport-upload** to collect the appropriate log files and system information and upload them to the Citrix Support ftp site. Please refer to [Section A.4.9.4, “host-bugreport-upload”](#) for a full description of this command and its optional parameters. If you are requested to send a crashdump to Citrix Support, use the **xe** command **host-crashdump-upload**. Please refer to [Section A.4.9.6, “host-crashdump-upload”](#) for a full description of this command and its optional parameters.

It is possible that sensitive information might be written into the XenServer host logs.

10.3.1. Sending host log messages to a central server

Rather than have logs written to the control domain filesystem, you can configure a XenServer host to write them to a remote server. The remote server must have the syslogd daemon running on it to receive the logs and aggregate them correctly. The syslogd daemon is a standard part of all flavors of Linux and Unix, and third-party versions are available for Windows and other operating systems.

To write logs to a remote server

1. Set the `syslog_destination` parameter to the hostname or IP address of the remote server where you want the logs to be written:

```
xe host-param-set uuid=<xenserver_host_uuid> logging:syslog_destination=<hostname>
```

2. Issue the command:

```
xe host-syslog-reconfigure uuid=<xenserver_host_uuid>
```

to enforce the change. (You can also execute this command remotely by specifying the *host* parameter.)

10.4. XenCenter logs

XenCenter also has client-side log. This file includes a complete description of all operations and errors that occur when using XenCenter. It also contains informational logging of events that provide you with an audit trail of various actions that have occurred. The XenCenter log file is stored in your profile folder. If XenCenter is installed on Windows XP, the path is

```
%userprofile%\AppData\Citrix\XenCenter\logs\XenCenter.log
```

If XenCenter is installed on Windows Vista, the path is

```
%userprofile%\AppData\Citrix\Roaming\XenCenter\logs\XenCenter.log
```

To quickly locate the XenCenter log files, for example, when you want to open or email the log file, click on **View Application Log Files** in the XenCenter **Help** menu.

10.5. Troubleshooting connections between XenCenter and the XenServer host

If you have trouble connecting to the XenServer host with XenCenter, check the following:



- Is your XenCenter an older version than the XenServer host you are attempting to connect to?

The XenCenter application is backward-compatible and can communicate properly with older XenServer hosts, but an older XenCenter cannot communicate properly with newer XenServer hosts.

To correct this issue, install a XenCenter version that is the same, or newer, than the XenServer host version.

- Is your license current?

You can see the expiration date for your License Key in the XenServer host **General** tab under the **License Details** section in XenCenter.

For details on licensing a host, see the chapter "Licensing XenServer" in the *XenServer Installation Guide*.

- The XenServer host talks to XenCenter using HTTPS over port 443 (a two-way connection for commands and responses using the XenAPI), and 5900 for graphical VNC connections with paravirtual Linux VMs. If you have a firewall enabled between the XenServer host and the machine running the client software, make sure that it allows traffic from these ports.

Appendix A. Command Line Interface

This chapter describes the XenServer command line interface (CLI). The *xe* CLI enables the writing of scripts for automating system administration tasks and allows integration of XenServer into an existing IT infrastructure.

The *xe* command line interface is installed by default on XenServer hosts and is included with XenCenter. A stand-alone remote CLI is also available for Linux.

On Windows, the *xe.exe* CLI executable is installed along with XenCenter.

To use it, open a Windows Command Prompt and change directories to the directory where the file resides (typically `C:\Program Files\Citrix\XenCenter`), or add its installation location to your system path.

On RPM-based distributions (such as, Red Hat and CentOS), you can install the stand-alone *xe* CLI executable from the RPM named `xe-cli-6.00-@BUILD_NUMBER@.i386.rpm` on the main XenServer installation ISO, as follows:

```
rpm -ivh xe-cli-6.00-@BUILD_NUMBER@.i386.rpm
```

Basic help is available for CLI commands on-host by typing:

```
xe help command
```

A list of the most commonly-used *xe* commands is displayed if you type:

```
xe help
```

or a list of all *xe* commands is displayed if you type:

```
xe help --all
```

A.1. Basic *xe* Syntax

The basic syntax of all XenServer *xe* CLI commands is:

```
xe <command-name> <argument=value> <argument=value> ...
```

Each specific command contains its own set of arguments that are of the form *argument=value*. Some commands have required arguments, and most have some set of optional arguments. Typically a command will assume default values for some of the optional arguments when invoked without them.

If the *xe* command is executed remotely, additional connection and authentication arguments are used. These arguments also take the form *argument=argument_value*.

The *server* argument is used to specify the hostname or IP address. The *username* and *password* arguments are used to specify credentials. A *password-file* argument can be specified instead of the password directly. In this case an attempt is made to read the password from the specified file (stripping CRs and LFs off the end of the file if necessary), and use that to connect. This is more secure than specifying the password directly at the command line.

The optional *port* argument can be used to specify the agent port on the remote XenServer host (defaults to 443).

Example: On the local XenServer host:

```
xe vm-list
```

Example: On the remote XenServer host:

```
xe vm-list -user <username> -password <password> -server <hostname>
```

Shorthand syntax is also available for remote connection arguments:

-u	username
-pw	password
-pwf	password file
-p	port
-s	server

Example: On a remote XenServer host:

```
xe vm-list -u <myuser> -pw <mypassword> -s <hostname>
```

Arguments are also taken from the environment variable `XE_EXTRA_ARGS`, in the form of comma-separated key/value pairs. For example, in order to enter commands on one XenServer host that are run on a remote XenServer host, you could do the following:

```
export XE_EXTRA_ARGS="server=jeffbeck,port=443,username=root,password=pass"
```

This command means that you will not need to specify the remote XenServer host parameters anymore, in each `xe` command you execute.

Using the `XE_EXTRA_ARGS` environment variable also enables tab completion of `xe` commands when issued against a remote XenServer host, which is disabled by default.

A.2. Special Characters and Syntax

To specify argument/value pairs on the **xe** command line, write:

```
argument=value
```

Unless the value includes spaces, do not use quotes. There should be no whitespace in between the argument name, the equals sign (=), and the value. Any argument not conforming to this format will be ignored.

For values containing spaces, write:

```
argument="value with spaces"
```

If you use the CLI while logged into a XenServer host, commands have a tab completion feature similar to that in the standard Linux bash shell. If you type, for example **xe vm-l** and then press the **TAB** key, the rest of the command will be displayed when it is unambiguous. If more than one command begins with `vm-l`, pressing **TAB** a second time will list the possibilities. This is particularly useful when specifying object UUIDs in commands.

Note:

When executing commands on a remote XenServer host, tab completion does not normally work. However if you put the server, username, and password in an environment variable called `XE_EXTRA_ARGS` on the machine from which you are entering the commands, tab completion is enabled. See [Section A.1, “Basic xe Syntax”](#) for details.

A.3. Command Types

Broadly speaking, the CLI commands can be split in two halves: Low-level commands concerned with listing and parameter manipulation of API objects, and higher level commands for interacting with VMs or hosts in a more abstract level. The low-level commands are:

- `<class>-list`



- `<class>-param-get`
- `<class>-param-set`
- `<class>-param-list`
- `<class>-param-add`
- `<class>-param-remove`
- `<class>-param-clear`

where `<class>` is one of:

- bond
- console
- host
- host-crashdump
- host-cpu
- network
- patch
- pbd
- pif
- pool
- sm
- sr
- task
- template
- vbd
- vdi
- vif
- vlan
- vm

Note that not every value of `<class>` has the full set of `<class>-param-` commands; some have just a subset.

A.3.1. Parameter Types

The objects that are addressed with the `xe` commands have sets of parameters that identify them and define their states.

Most parameters take a single value. For example, the `name-label` parameter of a VM contains a single string value. In the output from parameter list commands such as `xe vm-param-list`, such parameters have an indication in parentheses that defines whether they can be read and written to, or are read-only. For example, the output of `xe vm-param-list` on a specified VM might have the lines

```
user-version ( RW): 1
is-control-domain ( RO): false
```

The first parameter, `user-version`, is writable and has the value 1. The second, `is-control-domain`, is read-only and has a value of false.

The two other types of parameters are multi-valued. A `set` parameter contains a list of values. A `map` parameter is a set of key/value pairs. As an example, look at the following excerpt of some sample output of the `xe vm-param-list` on a specified VM:



```
platform (MRW): acpi: true; apic: true; pae: true; nx: false
allowed-operations (SR0): pause; clean_shutdown; clean_reboot; \
hard_shutdown; hard_reboot; suspend
```

The *platform* parameter has a list of items that represent key/value pairs. The key names are followed by a colon character (:). Each key/value pair is separated from the next by a semicolon character (;). The M preceding the RW indicates that this is a map parameter and is readable and writable. The *allowed-operations* parameter has a list that makes up a set of items. The S preceding the RO indicates that this is a set parameter and is readable but not writable.

In xe commands where you want to filter on a map parameter, or set a map parameter, use the separator : (colon) between the map parameter name and the key/value pair. For example, to set the value of the *foo* key of the *other-config* parameter of a VM to *baa*, the command would be

```
xe vm-param-set uuid=<VM uuid> other-config:foo=baa
```

Note:

In previous releases the separator - (dash) was used in specifying map parameters. This syntax still works but is deprecated.

A.3.2. Low-level Parameter Commands

There are several commands for operating on parameters of objects: *<class>-param-get*, *<class>-param-set*, *<class>-param-add*, *<class>-param-remove*, *<class>-param-clear*, and *<class>-param-list*. Each of these takes a *uuid* parameter to specify the particular object. Since these are considered low-level commands, they must be addressed by UUID and not by the VM name label.

<class>-param-list uuid=<uuid>

Lists all of the parameters and their associated values. Unlike the *class-list* command, this will list the values of "expensive" fields.

<class>-param-get uuid=<uuid> param-name=<parameter> [param-key=<key>]

Returns the value of a particular parameter. If the parameter is a map, specifying the param-key will get the value associated with that key in the map. If param-key is not specified, or if the parameter is a set, it will return a string representation of the set or map.

<class>-param-set uuid=<uuid> param=<value>...

Sets the value of one or more parameters.

<class>-param-add uuid=<uuid> param-name=<parameter> [<key>=<value>...] [param-key=<key>]

Adds to either a map or a set parameter. If the parameter is a map, add key/value pairs using the *<key>=<value>* syntax. If the parameter is a set, add keys with the *<param-key>=<key>* syntax.

<class>-param-remove uuid=<uuid> param-name=<parameter> param-key=<key>

Removes either a key/value pair from a map, or a key from a set.

<class>-param-clear uuid=<uuid> param-name=<parameter>

Completely clears a set or a map.

A.3.3. Low-level List Commands

The *<class>-list* command lists the objects of type *<class>*. By default it will list all objects, printing a subset of the parameters. This behavior can be modified in two ways: it can filter the objects so that it only outputs a subset, and the parameters that are printed can be modified.

To change the parameters that are printed, the argument *params* should be specified as a comma-separated list of the required parameters. For example:

```
xe vm-list params=name-label,other-config
```

Alternatively, to list all of the parameters, use the syntax:



```
xe vm-list params=all
```

Note that some parameters that are expensive to calculate will not be shown by the list command. These parameters will be shown as, for example:

```
allowed-VBD-devices (SR0): <expensive field>
```

To obtain these fields, use either the command `<class>-param-list` or `<class>-param-get`

To filter the list, the CLI will match parameter values with those specified on the command-line, only printing objects that match all of the specified constraints. For example:

```
xe vm-list HVM-boot-policy="BIOS order" power-state=halted
```

This command will only list those VMs for which *both* the field `power-state` has the value `halted`, and for which the field `HVM-boot-policy` has the value `BIOS order`.

It is also possible to filter the list based on the value of keys in maps, or on the existence of values in a set. The syntax for the first of these is `map-name:key=value`, and the second is `set-name:contains=value`

For scripting, a useful technique is passing `--minimal` on the command line, causing `xe` to print only the first field in a comma-separated list. For example, the command `xe vm-list --minimal` on a XenServer host with three VMs installed gives the three UUIDs of the VMs, for example:

```
a85d6717-7264-d00e-069b-3b1d19d56ad9,aaa3eec5-9499-bcf3-4c03-af10baea96b7, \
42c044de-df69-4b30-89d9-2c199564581d
```

A.4. xe Command Reference

This section provides a reference to the `xe` commands. They are grouped by objects that the commands address, and listed alphabetically.

A.4.1. Appliance Commands

Commands for creating and modifying VM appliances (also known as vApps). For more information on vApps, see the *XenServer Virtual Machine User's Guide*.

A.4.1.1. Appliance Parameters

Appliance commands have the following parameters

Parameter Name	Description	Type
uuid	appliance uuid	required
name-description	appliance description	optional
paused		optional
force	force shutdown	optional

A.4.1.2. appliance-assert-can-be-recovered

```
appliance-assert-can-be-recovered uuid=<appliance-uuid> database:vdi-uuid=<vdi-uuid>
```

Tests whether storage is available to recover this VM appliance/vApp.

A.4.1.3. appliance-create

```
appliance-create name-label=<name-label> [name-description=<name-description>]
```

Creates an appliance/vApp. For example:



```
xe appliance-create name-label=my_appliance
```

Add VMs to the appliance:

```
xe vm-param-set uuid=<VM-UUID> appliance=<appliance-uuid> \  
xe vm-param-set uuid=<VM-UUID> appliance=<appliance-uuid>
```

A.4.1.4. appliance-destroy

```
appliance-destroy uuid=<appliance-uuid>
```

Destroys an appliance/vApp. For example:

```
xe appliance-destroy uuid=<appliance-uuid>
```

A.4.1.5. appliance-recover

```
appliance-recover uuid=<appliance-uuid> database:vdi-uuid=<vdi-uuid> [paused=<true/false>]
```

Recover a VM appliance/vAPP from the database contained in the supplied VDI.

A.4.1.6. appliance-shutdown

```
appliance-shutdown uuid=<appliance-uuid> [force=<true/false>]
```

Shuts down all VMs in an appliance/vApp. For example:

```
xe appliance-shutdown uuid=<appliance-uuid>
```

A.4.1.7. appliance-start

```
appliance-start uuid=<appliance-uuid> [paused=<true/false>]
```

Starts an appliance/vApp. For example:

```
xe appliance-start uuid=<appliance-uuid>
```

A.4.2. Audit Commands

Audit commands download all of the available records of the RBAC audit file in the pool. If the optional parameter `since` is present, it downloads only the records from that specific point in time.

A.4.2.1. audit-log-get parameters

`audit-log-get` has the following parameters

Parameter Name	Description	Type
filename	Write the audit log of the pool to <i><filename></i>	required
since	specific date/time point	optional

A.4.2.2. audit-log-get

```
audit-log-get [since=<timestamp>] filename=<filename>
```

For example, to obtain audit records of the pool since a precise millisecond timestamp, run the following command:

Run the following command:

```
xe audit-log-get since=2009-09-24T17:56:20.530Z \  
filename=/tmp/auditlog-pool-actions.out
```

A.4.3. Bonding Commands

Commands for working with network bonds, for resilience with physical interface failover. See [Section 4.4.5, “Creating NIC Bonds on a Standalone Host”](#) for details.

The bond object is a reference object which glues together *master* and *member* PIFs. The master PIF is the bonding interface which must be used as the overall PIF to refer to the bond. The member PIFs are a set of two or more physical interfaces that have been combined into the high-level bonded interface.

A.4.3.1. Bond Parameters

Bonds have the following parameters:

Parameter Name	Description	Type
uuid	unique identifier/object reference for the bond	read only
master	UUID for the master bond PIF	read only
members	set of UUIDs for the underlying bonded PIFs	read only set parameter

A.4.3.2. bond-create

```
bond-create network-uuid=<network_uuid> pif-uuids=<pif_uuid_1,pif_uuid_2,...>
```

Create a bonded network interface on the network specified from a list of existing PIF objects. The command will fail if PIFs are in another bond already, if any member has a VLAN tag set, if the referenced PIFs are not on the same XenServer host, or if fewer than 2 PIFs are supplied.

A.4.3.3. bond-destroy

```
bond-destroy uuid=<bond_uuid>
```

Deletes a bonded interface specified by its UUID from a XenServer host.

A.4.4. CD Commands

Commands for working with physical CD/DVD drives on XenServer hosts.

A.4.4.1. CD Parameters

CDs have the following parameters:

Parameter Name	Description	Type
uuid	unique identifier/object reference for the CD	read only
name-label	Name for the CD	read/write
name-description	Description text for the CD	read/write
allowed-operations	A list of the operations that can be performed on this CD	read only set parameter
current-operations	A list of the operations that are currently in progress on this CD	read only set parameter

Parameter Name	Description	Type
sr-uuid	The unique identifier/object reference for the SR this CD is part of	read only
sr-name-label	The name for the SR this CD is part of	read only
vbd-uuids	A list of the unique identifiers for the VBDs on VMs that connect to this CD	read only set parameter
crashdump-uuids	Not used on CDs since crashdumps cannot be written to them	read only set parameter
virtual-size	Size of the CD as it appears to VMs (in bytes)	read only
physical-utilisation	amount of physical space that the CD image is currently taking up on the SR (in bytes)	read only
type	Set to User for CDs	read only
sharable	Whether or not the CD drive is sharable. Default is false.	read only
read-only	Whether the CD is read-only, if false, the device is writable. Always true for CDs.	read only
storage-lock	true if this disk is locked at the storage level	read only
parent	Reference to the parent disk, if this CD is part of a chain	read only
missing	true if SR scan operation reported this CD as not present on disk	read only
other-config	A list of key/value pairs that specify additional configuration parameters for the CD	read/write map parameter
location	The path on which the device is mounted	read only
managed	true if the device is managed	read only
xenstore-data	Data to be inserted into the xenstore tree	read only map parameter
sm-config	names and descriptions of storage manager device config keys	read only map parameter
is-a-snapshot	True if this template is a CD snapshot	read only
snapshot_of	The UUID of the CD that this template is a snapshot of	read only
snapshots	The UUID(s) of any snapshots that have been taken of this CD	read only
snapshot_time	The timestamp of the snapshot operation	read only

A.4.4.2. cd-list

cd - list [params=<param1,param2,...>] [parameter=<parameter_value>...]

List the CDs and ISOs (CD image files) on the XenServer host or pool, filtering on the optional argument *params*.

If the optional argument *params* is used, the value of *params* is a string containing a list of parameters of this object that you want to display. Alternatively, you can use the keyword *all* to show all parameters. If *params* is not used, the returned list shows a default subset of all available parameters.

Optional arguments can be any number of the [CD parameters](#) listed at the beginning of this section.

A.4.5. Console Commands

Commands for working with consoles.

The console objects can be listed with the standard object listing command (**xe console-list**), and the parameters manipulated with the standard parameter commands. See [Section A.3.2, “Low-level Parameter Commands”](#) for details.

A.4.5.1. Console Parameters

Consoles have the following parameters:

Parameter Name	Description	Type
uuid	The unique identifier/object reference for the console	read only
vm-uuid	The unique identifier/object reference of the VM this console is open on	read only
vm-name-label	The name of the VM this console is open on	read only
protocol	Protocol this console uses. Possible values are <i>vt100</i> : VT100 terminal, <i>rfb</i> : Remote FrameBuffer protocol (as used in VNC), or <i>rdp</i> : Remote Desktop Protocol	read only
location	URI for the console service	read only
other-config	A list of key/value pairs that specify additional configuration parameters for the console.	read/write map parameter

A.4.6. Disaster Recovery (DR) Commands

Commands for recovering VMs in the event of disaster

A.4.6.1. drtask-create

drtask-create type=<*type*> sr-whitelist=<*sr-white-list*> device-config=<*device-config*>

Creates a disaster recovery task. For example, to connect to an iSCSI SR in preparation for Disaster Recovery:

```
xe dr-task-create type=lvmoiscsi device-config:target=<target-ip-address> \
  device-config:targetIQN=<targetIQN> device-config:SCSIid=<SCSIid> \
  sr-whitelist=<sr-uuid-list>
```

Note:

sr-whitelist lists SR UUIDs, **drtask-create** will only introduce and connect to an SR which has one of the whitelisted UUIDs

A.4.6.2. drtask-destroy

`drtask-destroy uuid=<dr-task-uuid>`

Destroys a disaster recovery task and forgets the introduced SR.

A.4.6.3. vm-assert-can-be-recovered

`vm-assert-can-be-recovered uuid=<vm-uuid> database:vdi-uuid=<vdi-uuid>`

Tests whether storage is available to recover this VM.

A.4.6.4. appliance-assert-can-be-recovered

`appliance-assert-can-be-recovered uuid=<appliance-uuid> database:vdi-uuid=<vdi-uuid>`

Checks whether the storage (containing the appliance's/vAPP disk) is visible.

A.4.6.5. appliance-recover

`appliance-recover uuid=<appliance-uuid> database:vdi-uuid=<vdi-uuid> [force=<true|false>]`

Recover an appliance/vAPP from the database contained in the supplied VDI.

A.4.6.6. vm-recover

`vm-recover uuid=<vm-uuid> database:vdi-uuid=<vdi-uuid> [force=<true|false>]`

Recovers a VM from the database contained in the supplied VDI.

A.4.6.7. sr-enable-database-replication

`sr-enable-database-replication uuid=<sr-uuid>`

Enables xapi database replication to the specified (shared) SR. For example:

`xe sr-enable-database-replication uuid=<sr-uuid>`

A.4.6.8. sr-disable-database-replication

`sr-disable-database-replication uuid=<sr-uuid>`

Disables xapi database replication to the specified SR. For example:

`xe sr-enable-database-replication uuid=<sr-uuid>`

A.4.6.9. Example Usage

The example below shows the DR CLI commands in context:

On the primary site, enable database replication:

`xe sr-database-replication uuid=<sr-uuid>`

In the event of disaster, on the secondary site, connect to the SR:(note **device-config** has the same fields as **sr-probe**).

```
xe drtask-create type=lvmoiscsi \
  device-config:target=<target ip address> \
  device-config:targetIQN=<target-iqn> \
  device-config:SCSIid=<scsi-id> \
  sr-whitelist=<sr-uuid>
```


Look for database VDIs on the SR:

```
xe vdi-list sr-uuid=<sr-uuid> type=Metadata
```

Query a database VDI for VMs present:

```
xe vm-list database:vdi-uuid=<vdi-uuid>
```

Recover a VM:

```
xe vm-recover uuid=<vm-uuid> database:vdi-uuid=<vdi-uuid>
```

Destroy the DR task; any SRs introduced by the DR task and not required by VMs are destroyed:

```
xe drtask-destroy uuid=<drtask-uuid>
```

A.4.7. Event Commands

Commands for working with events.

A.4.7.1. Event Classes

Event classes are listed in the following table:

Class name	Description
pool	A pool of physical hosts
vm	A Virtual Machine
host	A physical host
network	A virtual network
vif	A virtual network interface
pif	A physical network interface (separate VLANs are represented as several PIFs)
sr	A storage repository
vdi	A virtual disk image
vbd	A virtual block device
pbd	The physical block devices through which hosts access SRs

A.4.7.2. event-wait

```
event-wait class=<class_name> [<param-name>=<param_value>] [<param-name>!=<param_value>]
```

Blocks other commands from executing until an object exists that satisfies the conditions given on the command line. $x=y$ means "wait for field x to take value y ", and $x\neq y$ means "wait for field x to take any value other than y ".

Example: wait for a specific VM to be running.

```
xe event-wait class=vm name-label=myvm power-state=running
```

Blocks other commands until a VM called myvm is in the *power-state* "running."

Example: wait for a specific VM to reboot:

`xe event-wait class=vm uuid=$VM start-time=/=$(xe vm-list uuid=$VM params=start-time --minimal)`

Blocks other commands until a VM with UUID `$VM` reboots (i.e. has a different `start-time` value).

The class name can be any of the [event classes](#) listed at the beginning of this section, and the parameters can be any of those listed in the CLI command `class-param-list`.

A.4.8. GPU Commands

Commands for working with physical GPUs, GPU groups and virtual GPUs.

The GPU objects can be listed with the standard object listing commands (**`xe pgpu-list`**, **`xe gpu-group-list`**, and **`xe vgpu-list`**), and the parameters manipulated with the standard parameter commands. For details, see [Section A.3.2, “Low-level Parameter Commands”](#).

A.4.8.1. Physical GPU (pGPU) Parameters

pGPUs have the following parameters:

Parameter Name	Description	Type
uuid	The unique identifier/object reference for the pGPU	Read only
vendor-name	The vendor name of the pGPU	Read only
device-name	The name assigned by the vendor to this pGPU model	Read only
gpu-group-uuid	The unique identifier/object reference for the GPU group that this pGPU has been automatically assigned to by XenServer; identical pGPUs across hosts in a pool are grouped together	Read only
gpu-group-name-label	The name of the GPU group to which the pGPU is assigned	Read only
host-uuid	The unique identifier/object reference for the XenServer host to which the pGPU is connected	Read only
host-name-label	The name of the XenServer host to which the pGPU is connected	Read only
pci-id	PCI identifier	Read only
dependencies	Lists the dependent PCI devices passed-through to the same VM	Read/write map parameter
other-config	A list of key/value pairs that specify additional configuration parameters for the pGPU	Read/write map parameter
supported-VGPU-types	List of virtual GPU types supported by the underlying hardware	Read only

Parameter Name	Description	Type
enabled-VGPU-types	List of virtual GPU types which have been enabled for this pGPU	Read/Write
resident-VGPUs	List of vGPUs running on this pGPU	Read only

A.4.8.2. pGPU Operations

Commands for working with pGPU Groups

A.4.8.2.1. pgpu-param-set

```
pgpu-param-set uuid=<uuid_of_pgpu> gpu-group-uuid=<uuid_of_destination_group>
```

Moves a pGPU to a different GPU group. Command is not allowed if there are running VMs using virtual GPUs resident on the destination group.

A.4.8.2.2. pgpu-param-get-uuid

```
pgpu-param-get-uuid uuid=<uuid_of_pgpu> param-name=<supported-vGPU-types/enabled-vGPU-types>
```

Displays the supported or enabled virtual GPU types for this pGPU.

A.4.8.2.3. pgpu-param-set-uuid

```
pgpu-param-set-uuid uuid=<uuid_of_pgpu> enabled_VGPU-types=<comand_separated_list_of_vgpu_type_uuids>
```

Changes the set of enabled virtual GPU types for this type of pGPU.

A.4.8.2.4. pgpu-param-add-uuid

```
pgpu-param-add-uuid uuid=<uuid_of_pgpu> param-name=<enabled_vgpu_types> param-key=<uuid_of_vgpu>
```

Enables a virtual GPU on this pGPU.

A.4.8.3. GPU Group Parameters

GPU groups have the following parameters:

Parameter Name	Description	Type
uuid	The unique identifier/object reference for the GPU group	Read only
name-label	The name of the GPU group	Read/write
name-description	The descriptive text of the GPU group	Read/write
VGPU-uuids	Lists the unique identifier/object references for the virtual GPUs in the GPU group	Read only set parameter

Parameter Name	Description	Type
PGPU-uuids	Lists the unique identifier/object references for the pGPUs in the GPU group	Read only set parameter
other-config	A list of key/value pairs that specify additional configuration parameters for the GPU group	Read/write map parameter
supported-VGPU-types	Union of all virtual GPU types supported by the underlying hardware	Read only
enabled-VGPU-types	Union of all virtual GPU types which have been enabled on the underlying pGPUs	Read only
allocation-algorithm	Depth-first/Breadth-first setting for allocation virtual GPUs on pGPUs within the group	Read/write enum parameter

A.4.8.3.1. GPU Group Operations

Commands for working with GPU Groups

A.4.8.3.1.1. gpu-group-create

```
gpu-group-create name-label=<name_for_group> [name-description=<description>]
```

Creates a new (empty) GPU Group into which pGPUs can be moved.

A.4.8.3.1.2. gpu-group-destroy

```
gpu-group-destroy uuid=<uuid_of_group>
```

Destroys the GPU Group; only permitted for empty groups.

A.4.8.3.1.3. gpu-group-get-remaining-capacity

```
gpu-group-get-remaining-capacity uuid=<uuid_of_group> vgpu-type-uuid=<uuid_of_vgpu_type>
```

Returns how many more virtual GPUs of the specified type can be instantiated in this GPU Group.

A.4.8.3.1.4. gpu-group-param-set

```
gpu-group-param-set uuid=<uuid_of_group> allocation-algorithm=<breadth-first|depth-first>
```

Changes the GPU Group's allocation algorithm for allocating virtual GPUs to pGPUS.

A.4.8.3.1.5. gpu-group-param-get-uuid

```
gpu-group-param-get-uuid uuid=<uuid_of_group> param-name=<supported-vGPU-types/enabled-vGPU-types>
```

Returns the supported or enabled types for this GPU Group.

A.4.8.4. Virtual GPU Parameters

Virtual GPUs have the following parameters:

Parameter Name	Description	Type
uuid	The unique identifier/object reference for the virtual GPU	Read only
vm-uuid	The unique identifier/object reference for the VM to which the virtual GPU is assigned	Read only
vm-name-label	The name of the VM to which the virtual GPU is assigned	Read only
gpu-group-uuid	The unique identifier/object reference for the GPU group in which the virtual GPU is contained	Read only
gpu-group-name-label	The name of the GPU group in which the virtual GPU is contained	Read only
currently-attached	True if a VM with GPU Pass-Through is running, false otherwise	Read only
other-config	A list of key/value pairs that specify additional configuration parameters for the virtual GPU	Read/write map parameter
type-uuid	The unique identifier/object reference for the virtual GPU type of this virtual GPU	Read/write map parameter
type-model-name	Model name associated with the virtual GPU type	Read only

A.4.8.5. Virtual GPU Type Parameters

Note:

Virtual GPU and GPU Pass-through are not compatible with XenMotion, Storage XenMotion or VM Suspend. However, VMs using GPU Pass-through or virtual GPUs can still be started any host that has the appropriate resources

Virtual GPU Types have the following parameters:

Parameter Name	Description	Type
uuid	The unique identifier/object reference for the virtual GPU type	Read only
vendor-name	Name of virtual GPU vendor	Read only
model-name	Model name associated with the virtual GPU type	Read only
freeze-frame	Framebuffer size of the virtual GPU type, in bytes	Read only

Parameter Name	Description	Type
max-heads	Maximum number of displays supported by the virtual GPU type	Read only
supported-on-PGPUs	List of pGPUs that support this virtual GPU type	Read only
enabled-on-PGPUs	List of pGPUs that have this virtual GPU type enabled	Read only
VGPU-uuids	List of virtual GPUs of this type	Read only

A.4.8.6. Virtual GPU Operations

A.4.8.6.1. vgpu-create

```
vgpu-create      vm-uuid=<uuid_of_vm>      gpu_group_uuid=<uuid_of_gpu_group>      [vgpu-type-uuid=<uuid_of_vgpu-type>]
```

Creates a virtual GPU. This command attaches the VM to the specified GPU group and optionally specifies the virtual GPU type. If no virtual GPU type is specified, the 'pass-through' type is assumed.

A.4.8.6.2. vgpu-destroy

```
vgpu-destroy uuid=<uuid_of_vgpu>
```

Destroy the specified virtual GPU.

A.4.8.6.3. Disabling VNC for VMs with virtual GPU

```
xe vm-param-add uuid=<uuid_of_vm> param-name=platform vgpu_vnc_enabled=<true | false.>
```

Using `false` disables the VNC console for a VM as it passes `disablevnc=1` through to the display emulator. By default, VNC is enabled.

A.4.9. Host Commands

Commands for interacting with XenServer host.

XenServer hosts are the physical servers running XenServer software. They have VMs running on them under the control of a special privileged Virtual Machine, known as the control domain or domain 0.

The XenServer host objects can be listed with the standard object listing command (**`xe host-list`**, **`xe host-cpu-list`**, and **`xe host-crashdump-list`**), and the parameters manipulated with the standard parameter commands. See [Section A.3.2, “Low-level Parameter Commands”](#) for details.

A.4.9.1. Host Selectors

Several of the commands listed here have a common mechanism for selecting one or more XenServer hosts on which to perform the operation. The simplest is by supplying the argument `host=<uuid_or_name_label>`. XenServer hosts can also be specified by filtering the full list of hosts on the values of fields. For example, specifying `enabled=true` will select all XenServer hosts whose `enabled` field is equal to `true`. Where multiple XenServer hosts are matching, and the operation can be performed on multiple XenServer hosts, the option `--multiple` must be specified to perform the operation. The full list of parameters that can be matched is described at the beginning of this section, and can be obtained by running the command **`xe host-list params=all`**. If no parameters to select XenServer hosts are given, the operation will be performed on all XenServer hosts.

A.4.9.2. Host Parameters

XenServer hosts have the following parameters:

Parameter Name	Description	Type
uuid	The unique identifier/object reference for the XenServer host	read only
name-label	The name of the XenServer host	read/write
name-description	The description string of the XenServer host	read only
enabled	false if disabled which prevents any new VMs from starting on them, which prepares the XenServer hosts to be shut down or rebooted; true if the host is currently enabled	read only
API-version-major	major version number	read only
API-version-minor	minor version number	read only
API-version-vendor	identification of API vendor	read only
API-version-vendor-implementation	details of vendor implementation	read only map parameter
logging	logging configuration	read/write map parameter
suspend-image-sr-uuid	the unique identifier/object reference for the SR where suspended images are put	read/write
crash-dump-sr-uuid	the unique identifier/object reference for the SR where crash dumps are put	read/write
software-version	list of versioning parameters and their values	read only map parameter
capabilities	list of Xen versions that the XenServer host can run	read only set parameter
other-config	A list of key/value pairs that specify additional configuration parameters for the XenServer host	read/write map parameter
chipset-info	A list of key/value pairs that specify information about the chipset	Read only map parameter
hostname	XenServer host hostname	read only
address	XenServer host IP address	read only

Parameter Name	Description	Type
license-server	- A list of key/value pairs that specify information about the license server - The default port for communications with Citrix products is 27000. For information on changing port numbers due to conflicts, see Change port numbers topic on Citrix Product Documentation website.	Read only map parameter
supported-bootloaders	list of bootloaders that the XenServer host supports, for example, pygrub, eliloader	read only set parameter
memory-total	total amount of physical RAM on the XenServer host, in bytes	read only
memory-free	total amount of physical RAM remaining that can be allocated to VMs, in bytes	read only
host-metrics-live	true if the host is operational	read only
logging	The <i>syslog_destination</i> key can be set to the hostname of a remote listening syslog service.	read/write map parameter
allowed-operations	lists the operations allowed in this state. This list is advisory only and the server state may have changed by the time this field is read by a client.	read only set parameter
current-operations	lists the operations currently in process. This list is advisory only and the server state may have changed by the time this field is read by a client	read only set parameter
patches	Set of host patches	read only set parameter
blobs	Binary data store	read only
memory-free-computed	A conservative estimate of the maximum amount of memory free on a host	read only
ha-statefiles	The UUID(s) of all HA statefiles	read only
ha-network-peers	The UUIDs of all hosts that could host the VMs on this host in case of failure	read only

Parameter Name	Description	Type
external-auth-type	Type of external authentication, for example, Active Directory.	read only
external-auth-service-name	The name of the external authentication service	read only
external-auth-configuration	Configuration information for the external authentication service.	read only map parameter

XenServer hosts contain some other objects that also have parameter lists.

CPUs on XenServer hosts have the following parameters:

Parameter Name	Description	Type
uuid	The unique identifier/object reference for the CPU	read only
number	the number of the physical CPU core within the XenServer host	read only
vendor	the vendor string for the CPU name, for example, "GenuineIntel"	read only
speed	The CPU clock speed, in Hz	read only
modelname	the vendor string for the CPU model, for example, "Intel(R) Xeon(TM) CPU 3.00GHz"	read only
stepping	the CPU revision number	read only
flags	the flags of the physical CPU (a decoded version of the features field)	read only
Utilisation	the current CPU utilisation	read only
host-uuid	the UUID of the host the CPU is in	read only
model	the model number of the physical CPU	read only
family	the physical CPU family number	read only

Crash dumps on XenServer hosts have the following parameters:

Parameter Name	Description	Type
uuid	The unique identifier/object reference for the crashdump	read only
host	XenServer host the crashdump corresponds to	read only
timestamp	Timestamp of the date and time that the crashdump occurred, in the form <i>yyyymmdd-hhmmss-ABC</i> , where <i>ABC</i> is the timezone indicator, for example, GMT	read only

Parameter Name	Description	Type
size	size of the crashdump, in bytes	read only

A.4.9.3. host-backup

host-backup file-name=<backup_filename> host=<host_name>

Download a backup of the control domain of the specified XenServer host to the machine that the command is invoked from, and save it there as a file with the name *file-name*.

While the **xe host-backup** command will work if executed on the local host (that is, without a specific hostname specified), do *not* use it this way. Doing so would fill up the control domain partition with the backup file. The command should *only* be used from a remote off-host machine where you have space to hold the backup file.

A.4.9.4. host-bugreport-upload

host-bugreport-upload [<host-selector>=<host_selector_value>...] [url=<destination_url>]
[http-proxy=<http_proxy_name>]

Generate a fresh bug report (using xen-bugtool, with all optional files included) and upload to the Citrix Support ftp site or some other location.

The host(s) on which this operation should be performed are selected using the standard selection mechanism (see [host selectors](#) above). Optional arguments can be any number of the [host parameters](#) listed at the beginning of this section.

Optional parameters are *http-proxy*: use specified http proxy, and *url*: upload to this destination URL. If optional parameters are not used, no proxy server is identified and the destination will be the default Citrix Support ftp site.

A.4.9.5. host-crashdump-destroy

host-crashdump-destroy uuid=<crashdump_uuid>

Delete a host crashdump specified by its UUID from the XenServer host.

A.4.9.6. host-crashdump-upload

host-crashdump-upload uuid=<crashdump_uuid>
[url=<destination_url>]
[http-proxy=<http_proxy_name>]

Upload a crashdump to the Citrix Support ftp site or other location. If optional parameters are not used, no proxy server is identified and the destination will be the default Citrix Support ftp site. Optional parameters are *http-proxy*: use specified http proxy, and *url*: upload to this destination URL.

A.4.9.7. host-disable

host-disable [<host-selector>=<host_selector_value>...]

Disables the specified XenServer hosts, which prevents any new VMs from starting on them. This prepares the XenServer hosts to be shut down or rebooted.

The host(s) on which this operation should be performed are selected using the standard selection mechanism (see [host selectors](#) above). Optional arguments can be any number of the [host parameters](#) listed at the beginning of this section.

A.4.9.8. host-dmesg

host-dmesg [<host-selector>=<host_selector_value>...]

Get a Xen dmesg (the output of the kernel ring buffer) from specified XenServer hosts.

The host(s) on which this operation should be performed are selected using the standard selection mechanism (see [host selectors](#) above). Optional arguments can be any number of the [host parameters](#) listed at the beginning of this section.

A.4.9.9. host-emergency-management-reconfigure

host-emergency-management-reconfigure interface=<uuid_of_management_interface_pif>

Reconfigure the management interface of this XenServer host. Use this command only if the XenServer host is in emergency mode, meaning that it is a member in a resource pool whose master has disappeared from the network and could not be contacted for some number of retries.

A.4.9.10. host-enable

host-enable [<host-selector>=<host_selector_value>...]

Enables the specified XenServer hosts, which allows new VMs to be started on them.

The host(s) on which this operation should be performed are selected using the standard selection mechanism (see [host selectors](#) above). Optional arguments can be any number of the [host parameters](#) listed at the beginning of this section.

A.4.9.11. host-evacuate

host-evacuate [<host-selector>=<host_selector_value>...]

Live migrates all running VMs to other suitable hosts on a pool. The host must first be disabled using the **host-disable** command.

If the evacuated host is the pool master, then another host must be selected to be the pool master. To change the pool master with HA disabled, you need to use the **pool-designate-new-master** command. See [Section A.4.15.2, "pool-designate-new-master"](#) for details. With HA enabled, your only option is to shut down the server, which will cause HA to elect a new master at random. See [Section A.4.9.29, "host-shutdown"](#).

The host(s) on which this operation should be performed are selected using the standard selection mechanism (see [host selectors](#) above). Optional arguments can be any number of the [host parameters](#) listed at the beginning of this section.

A.4.9.12. host-forget

host-forget uuid=<XenServer_host_UUID>

The xapi agent forgets about the specified XenServer host without contacting it explicitly.

Use the --force parameter to avoid being prompted to confirm that you really want to perform this operation.

Warning:

Do not use this command if HA is enabled on the pool. Disable HA first, then enable it again after you've forgotten the host.

Tip:

This command is useful if the XenServer host to "forget" is dead; however, if the XenServer host is live and part of the pool, you should use **xe pool-eject** instead.

A.4.9.13. host-get-system-status

host-get-system-status filename=<name_for_status_file>
[entries=<comma_separated_list>] [output=<tar.bz2 | zip>] [<host-selector>=<host_selector_value>...]

Download system status information into the specified file. The optional parameter *entries* is a comma-separated list of system status entries, taken from the capabilities XML fragment returned by the **host-get-system-status-capabilities** command. See [Section A.4.9.14, “host-get-system-status-capabilities”](#) for details. If not specified, all system status information is saved in the file. The parameter *output* may be *tar.bz2* (the default) or *zip*; if this parameter is not specified, the file is saved in *tar . bz2* form.

The host(s) on which this operation should be performed are selected using the standard selection mechanism (see [host selectors](#) above).

A.4.9.14. host-get-system-status-capabilities

host-get-system-status-capabilities [*<host-selector>=<host_selector_value>...*]

Get system status capabilities for the specified host(s). The capabilities are returned as an XML fragment that looks something like this:

```
<?xml version="1.0" ?> <system-status-capabilities>
  <capability content-type="text/plain" default-checked="yes" key="xenserver-logs" \
    max-size="150425200" max-time="-1" min-size="150425200" min-time="-1" \
    pii="maybe"/>
  <capability content-type="text/plain" default-checked="yes" \
    key="xenserver-install" max-size="51200" max-time="-1" min-size="10240" \
    min-time="-1" pii="maybe"/>
  ...
</system-status-capabilities>
```

Each capability entity has a number of attributes.

Attribute	Description
key	A unique identifier for the capability.
content-type	Can be either <i>text/plain</i> or <i>application/data</i> . Indicates whether a UI can render the entries for human consumption.
default-checked	Can be either <i>yes</i> or <i>no</i> . Indicates whether a UI should select this entry by default.
min-size, max-size	Indicates an approximate range for the size, in bytes, of this entry. -1 indicates that the size is unimportant.
min-time, max-time	Indicate an approximate range for the time, in seconds, taken to collect this entry. -1 indicates the time is unimportant.
pii	Personally identifiable information. Indicates whether the entry would have information that would identify the system owner, or details of their network topology. This is one of: <i>no</i>: no PII will be in these entries <i>yes</i>: PII will likely or certainly be in these entries <i>maybe</i>: you might wish to audit these entries for PII <i>if_customized</i> if the files are unmodified, then they will contain no PII, but since we encourage editing of these files, PII may have been introduced by such customization. This is used in particular for the networking scripts in the control domain. Passwords are never to be included in any bug report, regardless of any PII declaration.



The host(s) on which this operation should be performed are selected using the standard selection mechanism (see [host selectors](#) above).

A.4.9.15. host-is-in-emergency-mode

`host-is-in-emergency-mode`

Returns *true* if the host the CLI is talking to is currently in emergency mode, *false* otherwise. This CLI command works directly on slave hosts even with no master host present.

A.4.9.16. host-apply-edition

`host-apply-edition [host-uuid=<XenServer_host_UUID>] [edition=xenserver_edition=<"free"><"per-socket"><"xendesktop">]`

Assigns a XenServer license to a host server. When you assign a license, XenServer contacts the Citrix License Server and requests the specified type of license. If a license is available, it is then checked out from the license server.

For Citrix XenServer for XenDesktop editions, use *<"xendesktop">*.

For initial licensing configuration, see also **license-server-address** and **license-server-port**.

A.4.9.17. host-license-add

`host-license-add [license-file=<path/license_filename>] [host-uuid=<XenServer_host_UUID>]`

For XenServer (free edition), use to parse a local license file and add it to the specified XenServer host.

A.4.9.18. host-license-view

`host-license-view [host-uuid=<XenServer_host_UUID>]`

Displays the contents of the XenServer host license.

A.4.9.19. host-logs-download

`host-logs-download [file-name=<logfile_name>] [<host-selector>=<host_selector_value>...]`

Download a copy of the logs of the specified XenServer hosts. The copy is saved by default in a time-stamped file named `hostname-yyyy-mm-dd T hh:mm:ssZ.tar.gz`. You can specify a different filename using the optional parameter *file-name*.

The host(s) on which this operation should be performed are selected using the standard selection mechanism (see [host selectors](#) above). Optional arguments can be any number of the [host parameters](#) listed at the beginning of this section.

While the **xe host-logs-download** command will work if executed on the local host (that is, without a specific hostname specified), do *not* use it this way. Doing so will clutter the control domain partition with the copy of the logs. The command should *only* be used from a remote off-host machine where you have space to hold the copy of the logs.

A.4.9.20. host-management-disable

`host-management-disable`



Disables the host agent listening on an external management network interface and disconnects all connected API clients (such as the XenCenter). Operates directly on the XenServer host the CLI is connected to, and is not forwarded to the pool master if applied to a member XenServer host.

Warning:

Be extremely careful when using this CLI command off-host, since once it is run it will not be possible to connect to the control domain remotely over the network to re-enable it.

A.4.9.21. host-management-reconfigure

```
host-management-reconfigure [interface=<device> ] | [pif-uuid=<uuid> ]
```

Reconfigures the XenServer host to use the specified network interface as its management interface, which is the interface that is used to connect to the XenCenter. The command rewrites the MANAGEMENT_INTERFACE key in /etc/xensource-inventory.

If the device name of an interface (which must have an IP address) is specified, the XenServer host will immediately rebind. This works both in normal and emergency mode.

If the UUID of a PIF object is specified, the XenServer host determines which IP address to rebind to itself. It must not be in emergency mode when this command is executed.

Warning:

Be careful when using this CLI command off-host and ensure that you have network connectivity on the new interface. Use **xe pif-reconfigure** to set one up first. Otherwise, subsequent CLI commands will reach the XenServer host.

A.4.9.22. host-power-on

```
host-power-on [host=<host_uuid> ]
```

Turns on power on XenServer hosts with Host Power On functionality enabled. Before using this command, **host-set-power-on** must be enabled on the host.

A.4.9.23. host-get-cpu-features

```
host-get-cpu-features {features=<pool_master_cpu_features>} [uuid=<host_uuid>]
```

Prints a hexadecimal representation of the host's physical-CPU features.

A.4.9.24. host-set-cpu-features

```
host-set-cpu-features {features=<pool_master_cpu_features>} [uuid=<host_uuid>]
```

Attempts to mask the host's physical-CPU features to match the given features. The given string must be a 32-digit hexadecimal number (optionally containing spaces), as given by the host-get-cpu-features command.

A.4.9.25. host-set-power-on

```
host-set-power-on {host=<host_uuid> {power-on-mode=<""> <"wake-on-lan"> <"iLO"> <"DRAC"> <"custom"> } } | [power-on-config=<"power_on_ip"><"power_on_user"><"power_on_password_secret"> ] }
```

Use to enable Host Power On functionality on XenServer hosts that are compatible with remote power solutions. When using the **host-set-power-on** command, you must specify the type of power management solution on the host (that is, the <power-on-mode>). Then specify configuration options using the <power-on-config>

argument and its associated key-value pairs. To use the secrets feature to store your password, specify the key "power_on_password_secret".

A.4.9.26. host-reboot

```
host -reboot [<host-selector>=<host_selector_value>...]
```

Reboot the specified XenServer hosts. The specified hosts must be disabled first using the **xe host-disable** command, otherwise a HOST_IN_USE error message is displayed.

The host(s) on which this operation should be performed are selected using the standard selection mechanism (see [host selectors](#) above). Optional arguments can be any number of the [host parameters](#) listed at the beginning of this section.

If the specified XenServer hosts are members of a pool, the loss of connectivity on shutdown will be handled and the pool will recover when the XenServer hosts returns. If you shut down a pool member, other members and the master will continue to function. If you shut down the master, the pool will be out of action until the master is rebooted and back on line (at which point the members will reconnect and synchronize with the master) or until you make one of the members into the master.

A.4.9.27. host-restore

```
host -restore [file-name=<backup_filename>] [<host-selector>=<host_selector_value>...]
```

Restore a backup named *file-name* of the XenServer host control software. Note that the use of the word "restore" here does not mean a full restore in the usual sense, it merely means that the compressed backup file has been uncompressed and unpacked onto the secondary partition. After you've done a **xe host-restore**, you have to boot the Install CD and use its **Restore from Backup** option.

The host(s) on which this operation should be performed are selected using the standard selection mechanism (see [host selectors](#) above). Optional arguments can be any number of the [host parameters](#) listed at the beginning of this section.

A.4.9.28. host-set-hostname-live

```
host -set -hostname host-uuid=<uuid_of_host> hostname=<new_hostname>
```

Change the hostname of the XenServer host specified by *host-uuid*. This command persistently sets both the hostname in the control domain database and the actual Linux hostname of the XenServer host. Note that *hostname* is *not* the same as the value of the *name_label* field.

A.4.9.29. host-shutdown

```
host -shutdown [<host-selector>=<host_selector_value>...]
```

Shut down the specified XenServer hosts. The specified XenServer hosts must be disabled first using the **xe host-disable** command, otherwise a HOST_IN_USE error message is displayed.

The host(s) on which this operation should be performed are selected using the standard selection mechanism (see [host selectors](#) above). Optional arguments can be any number of the [host parameters](#) listed at the beginning of this section.

If the specified XenServer hosts are members of a pool, the loss of connectivity on shutdown will be handled and the pool will recover when the XenServer hosts returns. If you shut down a pool member, other members and the master will continue to function. If you shut down the master, the pool will be out of action until the master is rebooted and back on line, at which point the members will reconnect and synchronize with the master, or until one of the members is made into the master. If HA is enabled for the pool, one of the members will be made

into a master automatically. If HA is disabled, you must manually designate the desired server as master with the **pool-designate-new-master** command. See [Section A.4.15.2, “pool-designate-new-master”](#).

A.4.9.30. host-syslog-reconfigure

`host-syslog-reconfigure [<host-selector>=<host_selector_value>...]`

Reconfigure the syslog daemon on the specified XenServer hosts. This command applies the configuration information defined in the host *logging* parameter.

The host(s) on which this operation should be performed are selected using the standard selection mechanism (see [host selectors](#) above). Optional arguments can be any number of the [host parameters](#) listed at the beginning of this section.

A.4.9.31. host-data-source-list

`host-data-source-list [<host-selectors>=<host selector value>...]`

List the data sources that can be recorded for a host.

Select the host(s) on which to perform this operation by using the standard selection mechanism (see [host selectors](#)). Optional arguments can be any number of the [host parameters](#) listed at the beginning of this section. If no parameters to select hosts are given, the operation will be performed on all hosts.

Data sources have two parameters — *standard* and *enabled* — which can be seen by the output of this command. If a data source has *enabled* set to *true*, then the metrics are currently being recorded to the performance database. If a data source has *standard* set to *true*, then the metrics are recorded to the performance database *by default* (and so, *enabled* will also be set to *true* for this data source). If a data source has *standard* set to *false*, then the metrics are *not* recorded to the performance database by default (and so, *enabled* will also be set to *false* for this data source).

To start recording data source metrics to the performance database, run the **host-data-source-record** command. This will set *enabled* to *true*. To stop, run the **host-data-source-forget**. This will set *enabled* to *false*.

A.4.9.32. host-data-source-record

`host-data-source-record data-source=<name_description_of_data-source> [<host-selectors>=<host selector value>...]`

Record the specified data source for a host.

This operation writes the information from the data source to the persistent performance metrics database of the specified host(s). For performance reasons, this database is distinct from the normal agent database.

Select the host(s) on which to perform this operation by using the standard selection mechanism (see [host selectors](#)). Optional arguments can be any number of the [host parameters](#) listed at the beginning of this section. If no parameters to select hosts are given, the operation will be performed on all hosts.

A.4.9.33. host-data-source-forget

`host-data-source-forget data-source=<name_description_of_data-source> [<host-selectors>=<host selector value>...]`

Stop recording the specified data source for a host and forget all of the recorded data.

Select the host(s) on which to perform this operation by using the standard selection mechanism (see [host selectors](#)). Optional arguments can be any number of the [host parameters](#) listed at the beginning of this section. If no parameters to select hosts are given, the operation will be performed on all hosts.

A.4.9.34. host-data-source-query

`host-data-source-query data-source=<name_description_of_data-source> [<host-selectors>=<host selector value>...]`

Display the specified data source for a host.

Select the host(s) on which to perform this operation by using the standard selection mechanism (see [host selectors](#)). Optional arguments can be any number of the [host parameters](#) listed at the beginning of this section. If no parameters to select hosts are given, the operation will be performed on all hosts.

A.4.10. Message Commands

Commands for working with messages. Messages are created to notify users of significant events, and are displayed in XenCenter as alerts.

A.4.10.1. Message Parameters

Parameter Name	Description	Type
uuid	The unique identifier/object reference for the message	read only
name	The unique name of the message	read only
priority	The message priority. Higher numbers indicate greater priority	read only
class	The message class, for example VM.	read only
obj-uuid	The uuid of the affected object.	read only
timestamp	The time that the message was generated.	read only
body	The message content.	read only

A.4.10.2. message-create

`message-create name=<message_name> body=<message_text> [[host-uuid=<uuid_of_host>] | [sr-uuid=<uuid_of_sr>] | [vm-uuid=<uuid_of_vm>] | [pool-uuid=<uuid_of_pool>]]`

Creates a new message.

A.4.10.3. message-destroy

`message-destroy {uuid=<message_uuid>}`

Destroys an existing message. You can build a script to destroy all messages. For example:

```
# Dismiss all alerts \
IFS=","; for m in $(xe message-list params=uuid --minimal); do \
  xe message-destroy uuid=$m \
done
```

A.4.10.4. message-list

`message-list`

Lists all messages, or messages that match the specified standard selectable parameters.

A.4.11. Network Commands

Commands for working with networks.

The network objects can be listed with the standard object listing command (**xe network-list**), and the parameters manipulated with the standard parameter commands. See [Section A.3.2, “Low-level Parameter Commands”](#) for details.

A.4.11.1. Network Parameters

Networks have the following parameters:

Parameter Name	Description	Type
uuid	The unique identifier/object reference for the network	read only
name-label	The name of the network	read write
name-description	The description text of the network	read write
VIF-uuids	A list of unique identifiers of the VIFs (virtual network interfaces) that are attached from VMs to this network	read only set parameter
PIF-uuids	A list of unique identifiers of the PIFs (physical network interfaces) that are attached from XenServer hosts to this network	read only set parameter
bridge	name of the bridge corresponding to this network on the local XenServer host	read only
default-locking-mode	A network object used with VIF objects for ARP filtering. Set to <i><unlocked></i> to remove all the filtering rules associated with the VIF. Set to <i><disabled></i> so the VIF drops all traffic. See Section 4.4.10.8, “Using Switch Port Locking” and VM parameters .	read write
other-config:static-routes	comma-separated list of <i><subnet>/<netmask>/<gateway></i> formatted entries specifying the gateway address through which to route subnets. For example, setting other-config:static-routes to 172.16.0.0/15/192.168.0.3,172.18.0.0/16/192.168.0.4 causes traffic on 172.16.0.0/15 to be routed over 192.168.0.3 and traffic on 172.18.0.0/16 to be routed over 192.168.0.4.	read write
other-config:ethtool-autoneg	set to no to disable autonegotiation of the physical interface or bridge. Default is yes.	read write
other-config:ethtool-rx	set to on to enable receive checksum, off to disable	read write

Parameter Name	Description	Type
other-config:ethtool-tx	set to ON to enable transmit checksum, off to disable	read write
other-config:ethtool-sg	set to ON to enable scatter gather, off to disable	read write
other-config:ethtool-tso	set to ON to enable tcp segmentation offload, off to disable	read write
other-config:ethtool-ufo	set to ON to enable UDP fragment offload, off to disable	read write
other-config:ethtool-gso	set to ON to enable generic segmentation offload, off to disable	read write
blobs	Binary data store	read only

A.4.11.2. network-create

network-create name-label=<name_for_network> [name-description=<descriptive_text>]

Creates a new network.

A.4.11.3. network-destroy

network-destroy uuid=<network_uuid>

Destroys an existing network.

A.4.12. Patch (Update) Commands

Commands for working with XenServer host patches (updates). These are for the standard non-OEM editions of XenServer for commands relating to updating the OEM edition of XenServer, see [Section A.4.20, “Update Commands”](#) for details.

The patch objects can be listed with the standard object listing command (**xe patch-list**), and the parameters manipulated with the standard parameter commands. See [Section A.3.2, “Low-level Parameter Commands”](#) for details.

A.4.12.1. Patch Parameters

Patches have the following parameters:

Parameter Name	Description	Type
uuid	The unique identifier/object reference for the patch	read only
host-uuid	The unique identifier for the XenServer host to query	read only
name-label	The name of the patch	read only
name-description	The description string of the patch	read only
applied	Whether or not the patch has been applied; true or false	read only

Parameter Name	Description	Type
size	Whether or not the patch has been applied; true or false	read only

A.4.12.2. patch-apply

```
patch-apply host-uuid=<host_uuid> uuid=<patch_file_uuid>
```

Apply the specified patch file.

A.4.12.3. patch-pool-clean

```
patch-pool-clean uuid=<patch_file_uuid>
```

Deletes the specified patch (hotfix) file from all XenServer hosts in the pool. Can be used after applying hotfix packages through XenCenter or the CLI, or when dom0 is running out of disk space.

A.4.12.4. patch-pool-apply

```
patch-pool-apply uuid=<patch_uuid>
```

Apply the specified patch to all XenServer hosts in the pool.

A.4.12.5. patch-precheck

```
patch-precheck uuid=<patch_uuid> host-uuid=<host_uuid>
```

Run the prechecks contained within the specified patch on the specified XenServer host.

A.4.12.6. patch-upload

```
patch-upload file-name=<patch_filename>
```

Upload a specified patch file to the XenServer host. This prepares a patch to be applied. On success, the UUID of the uploaded patch is printed out. If the patch has previously been uploaded, a PATCH_ALREADY_EXISTS error is returned instead and the patch is not uploaded again.

A.4.13. PBD Commands

Commands for working with PBDs (Physical Block Devices). These are the software objects through which the XenServer host accesses storage repositories (SRs).

The PBD objects can be listed with the standard object listing command (**xe pbd-list**), and the parameters manipulated with the standard parameter commands. See [Section A.3.2, “Low-level Parameter Commands”](#) for details.

A.4.13.1. PBD Parameters

PBDs have the following parameters:

Parameter Name	Description	Type
uuid	The unique identifier/object reference for the PBD.	read only

Parameter Name	Description	Type
sr-uuid	the storage repository that the PBD points to	read only
device-config	additional configuration information that is provided to the SR-backend-driver of a host	read only map parameter
currently-attached	True if the SR is currently attached on this host, False otherwise	read only
host-uuid	UUID of the physical machine on which the PBD is available	read only
host	The host field is deprecated. Use host_uuid instead.	read only
other-config	Additional configuration information.	read/write map parameter

A.4.13.2. pbd-create

```
pbd-create host-uuid=<uuid_of_host>
sr-uuid=<uuid_of_sr>
[device-config:key=<corresponding_value>...]
```

Create a new PBD on a XenServer host. The read-only *device-config* parameter can only be set on creation.

To add a mapping of 'path' -> '/tmp', the command line should contain the argument *device-config:path=/tmp*

For a full list of supported device-config key/value pairs on each SR type see [Chapter 5, Storage](#).

A.4.13.3. pbd-destroy

```
pbd-destroy uuid=<uuid_of_pbd>
```

Destroy the specified PBD.

A.4.13.4. pbd-plug

```
pbd-plug uuid=<uuid_of_pbd>
```

Attempts to plug in the PBD to the XenServer host. If this succeeds, the referenced SR (and the VDIs contained within) should then become visible to the XenServer host.

A.4.13.5. pbd-unplug

```
pbd-unplug uuid=<uuid_of_pbd>
```

Attempt to unplug the PBD from the XenServer host.

A.4.14. PIF Commands

Commands for working with PIFs (objects representing the physical network interfaces).

The PIF objects can be listed with the standard object listing command (**xe pif-list**), and the parameters manipulated with the standard parameter commands. See [Section A.3.2, “Low-level Parameter Commands”](#) for details.

A.4.14.1. PIF Parameters

PIFs have the following parameters:

Parameter Name	Description	Type
uuid	the unique identifier/object reference for the PIF	read only
device	machine-readable name of the interface (for example, eth0)	read only
MAC	the MAC address of the PIF	read only
other-config	Additional PIF configuration name:value pairs.	read/write map parameter
physical	if true, the PIF points to an actual physical network interface	read only
currently-attached	is the PIF currently attached on this host? true or false	read only
MTU	Maximum Transmission Unit of the PIF in bytes.	read only
VLAN	VLAN tag for all traffic passing through this interface; -1 indicates no VLAN tag is assigned	read only
bond-master-of	the UUID of the bond this PIF is the master of (if any)	read only
bond-slave-of	the UUID of the bond this PIF is the slave of (if any)	read only
management	is this PIF designated to be a management interface for the control domain	read only
network-uuid	the unique identifier/object reference of the virtual network to which this PIF is connected	read only
network-name-label	the name of the of the virtual network to which this PIF is connected	read only
host-uuid	the unique identifier/object reference of the XenServer host to which this PIF is connected	read only
host-name-label	the name of the XenServer host to which this PIF is connected	read only
IP-configuration-mode	type of network address configuration used; DHCP or static	read only
IP	IP address of the PIF, defined here if IP-configuration-mode is static; undefined if DHCP	read only

Parameter Name	Description	Type
netmask	Netmask of the PIF, defined here if IP-configuration-mode is static; undefined if supplied by DHCP	read only
gateway	Gateway address of the PIF, defined here if IP-configuration-mode is static; undefined if supplied by DHCP	read only
DNS	DNS address of the PIF, defined here if IP-configuration-mode is static; undefined if supplied by DHCP	read only
io_read_kbs	average read rate in kB/s for the device	read only
io_write_kbs	average write rate in kB/s for the device	read only
carrier	link state for this device	read only
vendor-id	the ID assigned to NIC's vendor	read only
vendor-name	the NIC vendor's name	read only
device-id	the ID assigned by the vendor to this NIC model	read only
device-name	the name assigned by the vendor to this NIC model	read only
speed	data transfer rate of the NIC	read only
duplex	duplexing mode of the NIC; full or half	read only
pci-bus-path	PCI bus path address	read only
other-config:ethtool-speed	sets the speed of connection in Mbps	read write
other-config:ethtool-autoneg	set to NO to disable autonegotiation of the physical interface or bridge. Default is yes.	read write
other-config:ethtool-duplex	Sets duplexing capability of the PIF, either full or half.	read write
other-config:ethtool-rx	set to ON to enable receive checksum, off to disable	read write
other-config:ethtool-tx	set to ON to enable transmit checksum, off to disable	read write
other-config:ethtool-sg	set to ON to enable scatter gather, off to disable	read write
other-config:ethtool-tso	set to ON to enable tcp segmentation offload, off to disable	read write
other-config:ethtool-ufo	set to ON to enable udp fragment offload, off to disable	read write

Parameter Name	Description	Type
other-config:ethtool-gso	set to ON to enable generic segmentation offload, OFF to disable	read write
other-config:domain	comma-separated list used to set the DNS search path	read write
other-config:bond-miimon	interval between link liveness checks, in milliseconds	read write
other-config:bond-downdelay	number of milliseconds to wait after link is lost before really considering the link to have gone. This allows for transient link loss	read write
other-config:bond-updelay	number of milliseconds to wait after the link comes up before really considering it up. Allows for links flapping up. Default is 31s to allow for time for switches to begin forwarding traffic.	read write
disallow-unplug	True if this PIF is a dedicated storage NIC, false otherwise	read/write

Note:

Changes made to the other-config fields of a PIF will only take effect after a reboot. Alternately, use the **xe pif-unplug** and **xe pif-plug** commands to cause the PIF configuration to be rewritten.

A.4.14.2. pif-forget

```
pif-forget uuid=<uuid_of_pif>
```

Destroy the specified PIF object on a particular host.

A.4.14.3. pif-introduce

```
pif-introduce host-uuid=<UUID of XenServer host> mac=<mac_address_for_pif> device=<machine-readable name of the interface (for example, eth0)>
```

Create a new PIF object representing a physical interface on the specified XenServer host.

A.4.14.4. pif-plug

```
pif-plug uuid=<uuid_of_pif>
```

Attempt to bring up the specified physical interface.

A.4.14.5. pif-reconfigure-ip

```
pif-reconfigure-ip uuid=<uuid_of_pif> [ mode=<dhcp> | mode=<static> ]  
gateway=<network_gateway_address> IP=<static_ip_for_this_pif>  
netmask=<netmask_for_this_pif> [DNS=<dns_address>]
```

Modify the IP address of the PIF. For static IP configuration, set the *mode* parameter to *static*, with the *gateway*, *IP*, and *netmask* parameters set to the appropriate values. To use DHCP, set the *mode* parameter to *DHCP* and leave the static parameters undefined.

Note:

Using static IP addresses on physical network interfaces connected to a port on a switch using Spanning Tree Protocol with STP Fast Link turned off (or unsupported) results in a period during which there is no traffic.

A.4.14.6. pif-scan

`pif-scan host-uuid=<UUID of XenServer host>`

Scan for new physical interfaces on a XenServer host.

A.4.14.7. pif-unplug

`pif-unplug uuid=<uuid_of_pif>`

Attempt to bring down the specified physical interface.

A.4.15. Pool Commands

Commands for working with pools. A *pool* is an aggregate of one or more XenServer hosts. A pool uses one or more shared storage repositories so that the VMs running on one XenServer host in the pool can be migrated in near-real time (while still running, without needing to be shut down and brought back up) to another XenServer host in the pool. Each XenServer host is really a pool consisting of a single member by default. When a XenServer host is joined to a pool, it is designated as a member, and the pool it has joined becomes the master for the pool.

The singleton pool object can be listed with the standard object listing command (**xe pool-list**), and its parameters manipulated with the standard parameter commands. See [Section A.3.2, “Low-level Parameter Commands”](#) for details.

A.4.15.1. Pool Parameters

Pools have the following parameters:

Parameter Name	Description	Type
uuid	the unique identifier/object reference for the pool	read only
name-label	the name of the pool	read/write
name-description	the description string of the pool	read/write
master	the unique identifier/object reference of XenServer host designated as the pool's master	read only
default-SR	the unique identifier/object reference of the default SR for the pool	read/write
crash-dump-SR	the unique identifier/object reference of the SR where any crash dumps for pool members are saved	read/write

Parameter Name	Description	Type
metadata-vdis	all known metadata VDIs for the pool	read only
suspend-image-SR	the unique identifier/object reference of the SR where suspended VMs on pool members are saved	read/write
other-config	a list of key/value pairs that specify additional configuration parameters for the pool	read/write map parameter
supported-sr-types	SR types that can be used by this pool	read only
ha-enabled	True if HA is enabled for the pool, false otherwise	read only
ha-configuration	reserved for future use.	read only
ha-statefiles	lists the UUIDs of the VDIs being used by HA to determine storage health	read only
ha-host-failures-to-tolerate	the number of host failures to tolerate before sending a system alert	read/write
ha-plan-exists-for	the number of hosts failures that can actually be handled, according to the calculations of the HA algorithm	read only
ha-allow-overcommit	True if the pool is allowed to be overcommitted, False otherwise	read/write
ha-overcommitted	True if the pool is currently overcommitted	read only
blobs	binary data store	read only
pool-designate-new-master		

A.4.15.2. pool-designate-new-master

`pool-designate-new-master host-uuid=<UUID of member XenServer host to become new master>`

Instruct the specified member XenServer host to become the master of an existing pool. This performs an orderly hand over of the role of master host to another host in the resource pool. This command only works when the current master is online, and is not a replacement for the emergency mode commands listed below.

A.4.15.3. pool-dump-database

`pool-dump-database file-name=<filename_to_dump_database_into_(on_client)>`

Download a copy of the entire pool database and dump it into a file on the client.

A.4.15.4. pool-eject

`pool-eject host-uuid=<UUID of XenServer host to eject>`

Instruct the specified XenServer host to leave an existing pool.

A.4.15.5. pool-emergency-reset-master

`pool-emergency-reset-master master-address=<address of the pool's master XenServer host>`

Instruct a slave member XenServer host to reset its master address to the new value and attempt to connect to it. This command should not be run on master hosts.

A.4.15.6. pool-emergency-transition-to-master

`pool-emergency-transition-to-master`

Instruct a member XenServer host to become the pool master. This command is only accepted by the XenServer host if it has transitioned to emergency mode, meaning it is a member of a pool whose master has disappeared from the network and could not be contacted for some number of retries.

Note that this command may cause the password of the host to reset if it has been modified since joining the pool (see [Section A.4.21, "User Commands"](#)).

A.4.15.7. pool-ha-enable

`pool-ha-enable heartbeat-sr-uuids=<SR_UUID_of_the_Heartbeat_SR>`

Enable High Availability on the resource pool, using the specified SR UUID as the central storage heartbeat repository.

A.4.15.8. pool-ha-disable

`pool-ha-disable`

Disables the High Availability functionality on the resource pool.

A.4.15.9. pool-join

`pool-join master-address=<address> master-username=<username> master-password=<password>`

Instruct a XenServer host to join an existing pool.

A.4.15.10. pool-recover-slaves

`pool-recover-slaves`

Instruct the pool master to try and reset the master address of all members currently running in emergency mode. This is typically used after **pool-emergency-transition-to-master** has been used to set one of the members as the new master.

A.4.15.11. pool-restore-database

`pool-restore-database file-name=<filename_to_restore_from_(on_client)> [dry-run=<true | false>]`



Upload a database backup (created with **pool-dump-database**) to a pool. On receiving the upload, the master will restart itself with the new database.

There is also a *dry run* option, which allows you to check that the pool database can be restored without actually perform the operation. By default, *dry - run* is set to false.

A.4.15.12. pool-sync-database

`pool-sync-database`

Force the pool database to be synchronized across all hosts in the resource pool. This is not necessary in normal operation since the database is regularly automatically replicated, but can be useful for ensuring changes are rapidly replicated after performing a significant set of CLI operations.

A.4.16. Storage Manager Commands

Commands for controlling Storage Manager plugins.

The storage manager objects can be listed with the standard object listing command (**xe sm-list**), and the parameters manipulated with the standard parameter commands. See [Section A.3.2, “Low-level Parameter Commands”](#) for details.

A.4.16.1. SM Parameters

SMs have the following parameters:

Parameter Name	Description	Type
uuid	the unique identifier/object reference for the SM plugin	read only
name-label	the name of the SM plugin	read only
name-description	the description string of the SM plugin	read only
type	the SR type that this plugin connects to	read only
vendor	name of the vendor who created this plugin	read only
copyright	copyright statement for this SM plugin	read only
required-api-version	minimum SM API version required on the XenServer host	read only
configuration	names and descriptions of device configuration keys	read only
capabilities	capabilities of the SM plugin	read only
driver-filename	the filename of the SR driver.	read only

A.4.17. SR Commands

Commands for controlling SRs (storage repositories).

The SR objects can be listed with the standard object listing command (**xe sr-list**), and the parameters manipulated with the standard parameter commands. See [Section A.3.2, “Low-level Parameter Commands”](#) for details.

A.4.17.1. SR Parameters

SRs have the following parameters:

Parameter Name	Description	Type
uuid	the unique identifier/object reference for the SR	read only
name-label	the name of the SR	read/write
name-description	the description string of the SR	read/write
allowed-operations	list of the operations allowed on the SR in this state	read only set parameter
current-operations	list of the operations that are currently in progress on this SR	read only set parameter
VDIs	unique identifier/object reference for the virtual disks in this SR	read only set parameter
PBDs	unique identifier/object reference for the PBDs attached to this SR	read only set parameter
physical-utilisation	physical space currently utilized on this SR, in bytes. Note that for thin provisioned disk formats, physical utilisation may be less than virtual allocation	read only
physical-size	total physical size of the SR, in bytes	read only
type	type of the SR, used to specify the SR backend driver to use	read only
introduced-by	the drtask (if any) which introduced the SR	read only
content-type	the type of the SR's content. Used to distinguish ISO libraries from other SRs. For storage repositories that store a library of ISOs, the <i>content-type</i> must be set to <i>iso</i> . In other cases, Citrix recommends that this be set either to empty, or the string user.	read only
shared	True if this SR is capable of being shared between multiple XenServer hosts; False otherwise	read/write
other-config	list of key/value pairs that specify additional configuration parameters for the SR	read/write map parameter
host	The storage repository host name	read only
virtual-allocation	sum of virtual-size values of all VDIs in this storage repository (in bytes)	read only

Parameter Name	Description	Type
sm-config	SM dependent data	read only map parameter
blobs	binary data store	read only

A.4.17.2. sr-create

```
sr - create name-label=<name> physical-size=<size> type=<type>
content-type=<content_type> device-config:<config_name>=<value>
[host-uuid=<XenServer host UUID>] [shared=<true | false>]
```

Creates an SR on the disk, introduces it into the database, and creates a PBD attaching the SR to a XenServer host. If *shared* is set to *true*, a PBD is created for each XenServer host in the pool; if *shared* is not specified or set to *false*, a PBD is created only for the XenServer host specified with *host-uuid*.

The exact *device-config* parameters differ depending on the device *type*. See [Chapter 5, Storage](#) for details of these parameters across the different storage backends.

A.4.17.3. sr-destroy

```
sr - destroy uuid=<sr_uuid>
```

Destroys the specified SR on the XenServer host.

A.4.17.4. sr-enable-database-replication

```
sr - enable-database-replication uuid=<sr_uuid>
```

Enables xapi database replication to the specified (shared) SR. For example:

```
xe sr-enable-database-replication uuid=<sr-uuid>
```

A.4.17.5. sr-disable-database-replication

```
sr - disable-database-replication uuid=<sr_uuid>
```

Disables xapi database replication to the specified SR. For example:

```
xe sr-enable-database-replication uuid=<sr-uuid>
```

A.4.17.6. sr-forget

```
sr - forget uuid=<sr_uuid>
```

The xapi agent forgets about a specified SR on the XenServer host, meaning that the SR is detached and you cannot access VDIs on it, but it remains intact on the source media (the data is not lost).

A.4.17.7. sr-introduce

```
sr - introduce name-label=<name>
physical-size=<physical_size>
type=<type>
content-type=<content_type>
```



uuid=<sr_uuid>

Just places an SR record into the database. The device-config parameters are specified by device-config:<parameter_key>=<parameter_value>, for example:

```
xe sr-introduce device-config:<device>=</dev/sdb1>
```

Note:

This command is never used in normal operation. It is an advanced operation which might be useful if an SR needs to be reconfigured as shared after it was created, or to help recover from various failure scenarios.

A.4.17.8. sr-probe

```
sr - probe type=<type> [host-uuid=<uuid_of_host>] [device-config:<config_name>=<value>]
```

Performs a backend-specific scan, using the provided *device-config* keys. If the *device-config* is complete for the SR backend, then this will return a list of the SRs present on the device, if any. If the *device-config* parameters are only partial, then a backend-specific scan will be performed, returning results that will guide you in improving the remaining *device-config* parameters. The scan results are returned as backend-specific XML, printed out on the CLI.

The exact *device-config* parameters differ depending on the device *type*. See [Chapter 5, Storage](#) for details of these parameters across the different storage backends.

A.4.17.9. sr-scan

```
sr - scan uuid=<sr_uuid>
```

Force an SR scan, syncing the xapi database with VDIs present in the underlying storage substrate.

A.4.18. Task Commands

Commands for working with long-running asynchronous tasks. These are tasks such as starting, stopping, and suspending a Virtual Machine, which are typically made up of a set of other atomic subtasks that together accomplish the requested operation.

The task objects can be listed with the standard object listing command (**xe task-list**), and the parameters manipulated with the standard parameter commands. See [Section A.3.2, “Low-level Parameter Commands”](#) for details.

A.4.18.1. Task Parameters

Tasks have the following parameters:

Parameter Name	Description	Type
uuid	the unique identifier/object reference for the Task	read only
name-label	the name of the Task	read only
name-description	the description string of the Task	read only

Parameter Name	Description	Type
resident-on	the unique identifier/object reference of the host on which the task is running	read only
status	current status of the Task	read only
progress	if the Task is still pending, this field contains the estimated percentage complete, from 0. to 1. If the Task has completed, successfully or unsuccessfully, this should be 1.	read only
type	if the Task has successfully completed, this parameter contains the type of the encoded result, that is, the name of the class whose reference is in the result field; otherwise, this parameter's value is undefined	read only
result	if the Task has completed successfully, this field contains the result value, either Void or an object reference; otherwise, this parameter's value is undefined	read only
error_info	if the Task has failed, this parameter contains the set of associated error strings; otherwise, this parameter's value is undefined	read only
allowed_operations	list of the operations allowed in this state	read only
created	time the task has been created	read only
finished	time task finished (i.e. succeeded or failed). If task-status is pending, then the value of this field has no meaning	read only
subtask_of	contains the UUID of the tasks this task is a sub-task of	read only
subtasks	contains the UUID(s) of all the subtasks of this task	read only

A.4.18.2. task-cancel

`task-cancel [uuid=<task_uuid>]`

Direct the specified Task to cancel and return.

A.4.19. Template Commands

Commands for working with VM templates.

Templates are essentially VMs with the `is-a-template` parameter set to `true`. A template is a "gold image" that contains all the various configuration settings to instantiate a specific VM. XenServer ships with a base set of templates, which are generic "raw" VMs that can boot an OS vendor installation CD (e.g. RHEL, CentOS, SLES, Windows). With XenServer you can create VMs, configure them in standard forms for your particular needs, and save a copy of them as templates for future use in VM deployment.

The template objects can be listed with the standard object listing command (**xe template-list**), and the parameters manipulated with the standard parameter commands. See [Section A.3.2, “Low-level Parameter Commands”](#) for details.

Note:

Templates cannot be directly converted into VMs by setting the *is-a-template* parameter to *false*. Setting *is-a-template* parameter to *false* is not supported and results in a VM that cannot be started.

A.4.19.1. VM Template Parameters

Templates have the following parameters:

Parameter Name	Description	Type
uuid	the unique identifier/object reference for the template	read only
name-label	the name of the template	read/write
name-description	the description string of the template	read/write
user-version	string for creators of VMs and templates to put version information	read/write
is-a-template	true if this is a template. Template VMs can never be started, they are used only for cloning other VMs. After this value has been set to true it cannot be reset to false. Template VMs cannot be converted into VMs using this parameter.	read/write
is-control-domain	true if this is a control domain (domain 0 or a driver domain)	read only
power-state	current power state; always <i>halted</i> for a template	read only
memory-dynamic-max	dynamic maximum memory in bytes. Currently unused, but if changed the following constraint must be obeyed: <i>memory_static_max</i> >= <i>memory_dynamic_max</i> >= <i>memory_dynamic_min</i> >= <i>memory_static_min</i> .	read/write
memory-dynamic-min	dynamic minimum memory in bytes. Currently unused, but if changed the same constraints for <i>memory-dynamic-max</i> must be obeyed.	read/write
memory-static-max	statically-set (absolute) maximum memory in bytes. This is the main value used to determine the amount of memory assigned to a VM.	read/write

Parameter Name	Description	Type
memory-static-min	statically-set (absolute) minimum memory in bytes. This represents the absolute minimum memory, and <i>memory-static-min</i> must be less than <i>memory-static-max</i> . This value is currently unused in normal operation, but the previous constraint must be obeyed.	read/write
suspend-VDI-uuid	the VDI that a suspend image is stored on (has no meaning for a template)	read only
VCPUs-params	configuration parameters for the selected VCPU policy. You can tune a VCPU's pinning with <pre>xe template-param-set \ uuid=<template_uuid> \ VCPUs-params:mask=1,2,3</pre> A VM created from this template will then run on physical CPUs 1, 2, and 3 only. You can also tune the VCPU priority (xen scheduling) with the <i>cap</i> and <i>weight</i> parameters; for example <pre>xe template-param-set \ uuid=<template_uuid> \ VCPUs-params:weight=512 xe template-param-set \ uuid=<template_uuid> \ VCPUs-params:cap=100</pre> A VM based on this template with a weight of 512 will get twice as much CPU as a domain with a weight of 256 on a contended XenServer host. Legal weights range from 1 to 65535 and the default is 256. The cap optionally fixes the maximum amount of CPU a VM based on this template will be able to consume, even if the XenServer host has idle CPU cycles. The cap is expressed in percentage of one physical CPU: 100 is 1 physical CPU, 50 is half a CPU, 400 is 4 CPUs, etc. The default, 0, means there is no upper cap.	read/write map parameter
VCPUs-max	maximum number of VCPUs	read/write
VCPUs-at-startup	boot number of VCPUs	read/write

Parameter Name	Description	Type
actions-after-crash	action to take if a VM based on this template crashes	read/write
console-uuids	virtual console devices	read only set parameter
platform	platform specific configuration To disable the emulation of a parallel port for HVM guests (for example, Windows guests): <pre>xe vm-param-set \ uuid=<vm_uuid> \ platform:parallel=none</pre> To disable the emulation of a serial port for HVM guests: <pre>xe vm-param-set \ uuid=<vm_uuid> \ platform:hvm_serial=none</pre> To disable the emulation of a USB controller and a USB tablet device for HVM guests: <pre>xe vm-param-set \ uuid=<vm_uuid> \ platform:usb=false</pre> <pre>xe vm-param-set \ uuid=<vm_uuid> \ platform:usb_tablet=false</pre>	read/write map parameter
allowed-operations	list of the operations allowed in this state	read only set parameter
current-operations	list of the operations that are currently in progress on this template	read only set parameter
allowed-VBD-devices	list of VBD identifiers available for use, represented by integers of the range 0-15. This list is informational only, and other devices may be used (but may not work).	read only set parameter
allowed-VIF-devices	list of VIF identifiers available for use, represented by integers of the range 0-15. This list is informational only, and other devices may be used (but may not work).	read only set parameter
HVM-boot-policy	the boot policy for HVM guests. Either BIOS Order or an empty string.	read/write

Parameter Name	Description	Type
HVM-boot-params	the <i>order</i> key controls the HVM guest boot order, represented as a string where each character is a boot method: d for the CD/DVD, c for the root disk, and n for network PXE boot. The default is dc.	read/write map parameter
PV-kernel	path to the kernel	read/write
PV-ramdisk	path to the initrd	read/write
PV-args	string of kernel command line arguments	read/write
PV-legacy-args	string of arguments to make legacy VMs based on this template boot	read/write
PV-bootloader	name of or path to bootloader	read/write
PV-bootloader-args	string of miscellaneous arguments for the bootloader	read/write
last-boot-CPU-flags	describes the CPU flags on which a VM based on this template was last booted; not populated for a template	read only
resident-on	the XenServer host on which a VM based on this template is currently resident; appears as <not in database> for a template	read only
affinity	a XenServer host which a VM based on this template has preference for running on; used by the xe vm-start command to decide where to run the VM	read/write
other-config	list of key/value pairs that specify additional configuration parameters for the template	read/write map parameter
start-time	timestamp of the date and time that the metrics for a VM based on this template were read, in the form <code>yyyymmddThh:mm:ss z</code> , where <code>z</code> is the single-letter military timezone indicator, for example, <code>Z</code> for UTC (GMT); set to 1 Jan 1970 Z (beginning of Unix/POSIX epoch) for a template	read only

Parameter Name	Description	Type
install-time	timestamp of the date and time that the metrics for a VM based on this template were read, in the form <code>yyyymmddThh:mm:ss Z</code> , where Z is the single-letter military timezone indicator, for example, Z for UTC (GMT); set to 1 Jan 1970 Z (beginning of Unix/POSIX epoch) for a template	read only
memory-actual	the actual memory being used by a VM based on this template; 0 for a template	read only
VCPUs-number	the number of virtual CPUs assigned to a VM based on this template; 0 for a template	read only
VCPUs-Utilization	list of virtual CPUs and their weight	read only map parameter
os-version	the version of the operating system for a VM based on this template; appears as <code><not in database></code> for a template	read only map parameter
PV-drivers-version	the versions of the paravirtualized drivers for a VM based on this template; appears as <code><not in database></code> for a template	read only map parameter
PV-drivers-detected	flag for latest version of the paravirtualized drivers for a VM based on this template; appears as <code><not in database></code> for a template	read only
memory	memory metrics reported by the agent on a VM based on this template; appears as <code><not in database></code> for a template	read only map parameter
disks	disk metrics reported by the agent on a VM based on this template; appears as <code><not in database></code> for a template	read only map parameter
networks	network metrics reported by the agent on a VM based on this template; appears as <code><not in database></code> for a template	read only map parameter
other	other metrics reported by the agent on a VM based on this template; appears as <code><not in database></code> for a template	read only map parameter

Parameter Name	Description	Type
guest-metrics-last-updated	timestamp when the last write to these fields was performed by the in-guest agent, in the form <code>yyyymmddThh:mm:ss Z</code> , where Z is the single-letter military timezone indicator, for example, Z for UTC (GMT)	read only
actions-after-shutdown	action to take after the VM has shutdown	read/write
actions-after-reboot	action to take after the VM has rebooted	read/write
possible-hosts	list of hosts that could potentially host the VM	read only
HVM-shadow-multiplier	multiplier applied to the amount of shadow that will be made available to the guest	read/write
dom-id	domain ID (if available, -1 otherwise)	read only
recommendations	XML specification of recommended values and ranges for properties of this VM	read only
xenstore-data	data to be inserted into the xenstore tree (<code>/local/domain/<domid>/vm-data</code>) after the VM is created.	read/write map parameter
is-a-snapshot	True if this template is a VM snapshot	read only
snapshot_of	the UUID of the VM that this template is a snapshot of	read only
snapshots	the UUID(s) of any snapshots that have been taken of this template	read only
snapshot_time	the timestamp of the most recent VM snapshot taken	read only
memory-target	the target amount of memory set for this template	read only
blocked-operations	lists the operations that cannot be performed on this template	read/write map parameter
last-boot-record	record of the last boot parameters for this template, in XML format	read only
ha-always-run	True if an instance of this template will always restarted on another host in case of the failure of the host it is resident on	read/write
ha-restart-priority	restart or best effort	read/write

Parameter Name	Description	Type
blobs	binary data store	read only
live	only relevant to a running VM.	read only

A.4.19.2. template-export

`template-export template-uuid=<uuid_of_existing_template> filename=<filename_for_new_template>`

Exports a copy of a specified template to a file with the specified new filename.

A.4.20. Update Commands

Commands for working with updates to the OEM edition of XenServer. For commands relating to updating the standard non-OEM editions of XenServer, see [Section A.4.12, “Patch \(Update\) Commands”](#) for details.

A.4.20.1. update-upload

`update-upload file-name=<name_of_upload_file>`

Streams a new software image to a OEM edition XenServer host. You must then restart the host for this to take effect.

A.4.21. User Commands

A.4.21.1. user-password-change

`user-password-change old=<old_password> new=<new_password>`

Changes the password of the logged-in user. The old password field is not checked because you require supervisor privilege to make this call.

A.4.22. VBD Commands

Commands for working with VBDs (Virtual Block Devices).

A VBD is a software object that connects a VM to the VDI, which represents the contents of the virtual disk. The VBD has the attributes which tie the VDI to the VM (is it bootable, its read/write metrics, and so on), while the VDI has the information on the physical attributes of the virtual disk (which type of SR, whether the disk is shareable, whether the media is read/write or read only, and so on).

The VBD objects can be listed with the standard object listing command (**xe vbd-list**), and the parameters manipulated with the standard parameter commands. See [Section A.3.2, “Low-level Parameter Commands”](#) for details.

A.4.22.1. VBD Parameters

VBDs have the following parameters:

Parameter Name	Description	Type
uuid	the unique identifier/object reference for the VBD	read only
vm-uuid	the unique identifier/object reference for the VM this VBD is attached to	read only
vm-name-label	the name of the VM this VBD is attached to	read only

Parameter Name	Description	Type
vdi-uuid	the unique identifier/object reference for the VDI this VBD is mapped to	read only
vdi-name-label	the name of the VDI this VBD is mapped to	read only
empty	if true, this represents an empty drive	read only
device	the device seen by the guest, for example hda	read only
userdevice	device number specified by the device parameter during vbd-create, for example, 0 for hda, 1 for hdb, etc	read/write
bootable	true if this VBD is bootable	read/write
mode	the mode the VBD should be mounted with	read/write
type	how the VBD appears to the VM, for example disk or CD	read/write
currently-attached	True if the VBD is currently attached on this host, false otherwise	read only
storage-lock	True if a storage-level lock was acquired	read only
status-code	error/success code associated with the last attach operation	read only
status-detail	error/success information associated with the last attach operation status	read only
qos_algorithm_type	the QoS algorithm to use	read/write
qos_algorithm_params	parameters for the chosen QoS algorithm	read/write map parameter
qos_supported_algorithms	supported QoS algorithms for this VBD	read only set parameter
io_read_kbs	average read rate in kB per second for this VBD	read only
io_write_kbs	average write rate in kB per second for this VBD	read only

Parameter Name	Description	Type
allowed-operations	list of the operations allowed in this state. This list is advisory only and the server state may have changed by the time this field is read by a client.	read only set parameter
current-operations	links each of the running tasks using this object (by reference) to a <code>current_operation</code> enum which describes the nature of the task.	read only set parameter
unpluggable	true if this VBD will support hot-unplug	read/write
attachable	True if the device can be attached	read only
other-config	additional configuration	read/write map parameter

A.4.22.2. vbd-create

```
vbd-create vm-uuid=<uuid_of_the_vm> device=<device_value>
vdi-uuid=<uuid_of_the_vdi_the_vbd_will_connect_to> [bootable=true] [type=<Disk | CD>] [mode=<RW | RO>]
```

Create a new VBD on a VM.

The allowable values for the *device* field are integers from 0-15, and the number must be unique for each VM. The current allowable values can be seen in the *allowed-VBD-devices* parameter on the specified VM. This will be seen as *userdevice* in the vbd parameters.

If the *type* is *Disk*, *vdi-uuid* is required. Mode can be *RO* or *RW* for a Disk.

If the *type* is *CD*, *vdi-uuid* is optional; if no VDI is specified, an empty VBD will be created for the CD. Mode must be *RO* for a CD.

A.4.22.3. vbd-destroy

```
vbd-destroy uuid=<uuid_of_vbd>
```

Destroy the specified VBD.

If the VBD has its *other-config:owner* parameter set to *true*, the associated VDI will also be destroyed.

A.4.22.4. vbd-eject

```
vbd-eject uuid=<uuid_of_vbd>
```

Remove the media from the drive represented by a VBD. This command only works if the media is of a removable type (a physical CD or an ISO); otherwise an error message `VBD_NOT_REMOVABLE_MEDIA` is returned.

A.4.22.5. vbd-insert

```
vbd-insert uuid=<uuid_of_vbd> vdi-uuid=<uuid_of_vdi_containing_media>
```

Insert new media into the drive represented by a VBD. This command only works if the media is of a removable type (a physical CD or an ISO); otherwise an error message `VBD_NOT_REMOVABLE_MEDIA` is returned.

A.4.22.6. vbd-plug

```
vbd-plug uuid=<uuid_of_vbd>
```

Attempt to attach the VBD while the VM is in the running state.

A.4.22.7. vbd-unplug

```
vbd-unplug uuid=<uuid_of_vbd>
```

Attempts to detach the VBD from the VM while it is in the running state.

A.4.23. VDI Commands

Commands for working with VDIs (Virtual Disk Images).

A VDI is a software object that represents the contents of the virtual disk seen by a VM, as opposed to the VBD, which is a connector object that ties a VM to the VDI. The VDI has the information on the physical attributes of the virtual disk (which type of SR, whether the disk is shareable, whether the media is read/write or read only, and so on), while the VBD has the attributes which tie the VDI to the VM (is it bootable, its read/write metrics, and so on).

The VDI objects can be listed with the standard object listing command (**xe vdi-list**), and the parameters manipulated with the standard parameter commands. See [Section A.3.2, “Low-level Parameter Commands”](#) for details.

A.4.23.1. VDI Parameters

VDIs have the following parameters:

Parameter Name	Description	Type
uuid	the unique identifier/object reference for the VDI	read only
name-label	the name of the VDI	read/write
name-description	the description string of the VDI	read/write
allowed-operations	a list of the operations allowed in this state	read only set parameter
current-operations	a list of the operations that are currently in progress on this VDI	read only set parameter
sr-uuid	SR in which the VDI resides	read only
vbd-uuids	a list of VBDs that refer to this VDI	read only set parameter
crashdump-uuids	list of crash dumps that refer to this VDI	read only set parameter
virtual-size	size of disk as presented to the VM, in bytes. Note that, depending on the storage backend type, the size may not be respected exactly	read only

Parameter Name	Description	Type
physical-utilisation	amount of physical space that the VDI is currently taking up on the SR, in bytes	read only
type	type of VDI, for example, System or User	read only
sharable	true if this VDI may be shared	read only
read-only	true if this VDI can only be mounted read-only	read only
storage-lock	true if this VDI is locked at the storage level	read only
parent	references the parent VDI, if this VDI is part of a chain	read only
missing	true if SR scan operation reported this VDI as not present	read only
other-config	additional configuration information for this VDI	read/write map parameter
sr-name-label	name of the containing storage repository	read only
location	location information	read only
managed	true if the VDI is managed	read only
xenstore-data	data to be inserted into the xenstore tree (/local/domain/0/backend/vbd/<domid>/<device-id>/sm-data) after the VDI is attached. This is generally set by the SM backends on vd_i_attach .	read only map parameter
sm-config	SM dependent data	read only map parameter
is-a-snapshot	true if this VDI is a VM storage snapshot	read only
snapshot_of	the UUID of the storage this VDI is a snapshot of	read only
snapshots	the UUID(s) of all snapshots of this VDI	read only
snapshot_time	the timestamp of the snapshot operation that created this VDI	read only
metadata-of-pool	the uuid of the pool which created this metadata VDI	read only
metadata-latest	flag indicating whether the VDI contains the latest known metadata for this pool	read only

A.4.23.2. vdi-clone

`vd_i_clone uuid=<uuid_of_the_vdi> [driver-params:<key=value>]`

Create a new, writable copy of the specified VDI that can be used directly. It is a variant of **vd_i-copy** that is capable of exposing high-speed image clone facilities where they exist.

The optional *driver-params* map parameter can be used for passing extra vendor-specific configuration information to the back end storage driver that the VDI is based on. See the storage vendor driver documentation for details.

A.4.23.3. vdi-copy

```
vdi-copy uuid=<uuid_of_the_vdi> sr-uuid=<uuid_of_the_destination_sr>
```

Copy a VDI to a specified SR.

A.4.23.4. vdi-create

```
vdi-create sr-uuid=<uuid_of_the_sr_where_you_want_to_create_the_vdi>
name-label=<name_for_the_vdi>
type=<system | user | suspend | crashdump>
virtual-size=<size_of_virtual_disk>
sm-config-*=<storage_specific_configuration_data>
```

Create a VDI.

The *virtual-size* parameter can be specified in bytes or using the IEC standard suffixes KiB (2^{10} bytes), MiB (2^{20} bytes), GiB (2^{30} bytes), and TiB (2^{40} bytes).

Note:

SR types that support thin provisioning of disks (such as Local VHD and NFS) do not enforce virtual allocation of disks. Users should therefore take great care when over-allocating virtual disk space on an SR. If an over-allocated SR does become full, disk space must be made available either on the SR target substrate or by deleting unused VDIs in the SR.

Note:

Some SR types might round up the *virtual-size* value to make it divisible by a configured block size.

A.4.23.5. vdi-destroy

```
vdi-destroy uuid=<uuid_of_vdi>
```

Destroy the specified VDI.

Note:

In the case of Local VHD and NFS SR types, disk space is not immediately released on **vdi-destroy**, but periodically during a storage repository scan operation. Users that need to force deleted disk space to be made available should call **sr-scan** manually.

A.4.23.6. vdi-forget

```
vdi-forget uuid=<uuid_of_vdi>
```

Unconditionally removes a VDI record from the database without touching the storage backend. In normal operation, you should be using **vdi-destroy** instead.

A.4.23.7. vdi-import

```
vdi-import uuid=<uuid_of_vdi> filename=<filename_of_raw_vdi>
```

Import a raw VDI.

A.4.23.8. vdi-introduce

```
vdi-introduce uuid=<uuid_of_vdi>
sr-uuid=<uuid_of_sr_to_import_into>
name-label=<name_of_the_new_vdi>
type=<system | user | suspend | crashdump>
location=<device_location_(varies_by_storage_type)>
[name-description=<description_of_vdi>]
[sharable=<yes | no>]
[read-only=<yes | no>]
[other-config=<map_to_store_misc_user_specific_data>]
[xenstore-data=<map_to_of_additional_xenstore_keys>]
[sm-config<storage_specific_configuration_data>]
```

Create a VDI object representing an existing storage device, without actually modifying or creating any storage. This command is primarily used internally to automatically introduce hot-plugged storage devices.

A.4.23.9. vdi-pool-migrate

```
vdi-pool-migrate <uuid>=<VDI_uuid> <sr-uuid>=<destination-sr-uuid>
```

Migrate a VDI to a specified SR, while the VDI is attached to a running guest. (Storage XenMotion)

See the chapter "VM Migration" in the *XenServer Virtual Machine User's Guide*.

A.4.23.10. vdi-resize

```
vdi-resize uuid=<vdi_uuid> disk-size=<new_size_for_disk>
```

Resize the VDI specified by UUID.

A.4.23.11. vdi-snapshot

```
vdi-snapshot uuid=<uuid_of_the_vdi> [driver-params=<params>]
```

Produces a read-write version of a VDI that can be used as a reference for backup and/or template creation purposes. You can perform a backup from a snapshot rather than installing and running backup software inside the VM. The VM can continue running while external backup software streams the contents of the snapshot to the backup media. Similarly, a snapshot can be used as a "gold image" on which to base a template. A template can be made using any VDIs.

The optional *driver-params* map parameter can be used for passing extra vendor-specific configuration information to the back end storage driver that the VDI is based on. See the storage vendor driver documentation for details.

A clone of a snapshot should always produce a writable VDI.

A.4.23.12. vdi-unlock

```
vdi-unlock uuid=<uuid_of_vdi_to_unlock> [force=true]
```

Attempts to unlock the specified VDIs. If *force=true* is passed to the command, it will force the unlocking operation.

A.4.24. VIF Commands

Commands for working with VIFs (Virtual network interfaces).

The VIF objects can be listed with the standard object listing command (**xe vif-list**), and the parameters manipulated with the standard parameter commands. See [Section A.3.2, “Low-level Parameter Commands”](#) for details.

A.4.24.1. VIF Parameters

VIFs have the following parameters:

Parameter Name	Description	Type
uuid	the unique identifier/object reference for the VIF	read only
vm-uuid	the unique identifier/object reference for the VM that this VIF resides on	read only
vm-name-label	the name of the VM that this VIF resides on	read only
allowed-operations	a list of the operations allowed in this state	read only set parameter
current-operations	a list of the operations that are currently in progress on this VIF	read only set parameter
device	integer label of this VIF, indicating the order in which VIF backends were created	read only
MAC	MAC address of VIF, as exposed to the VM	read only
MTU	Maximum Transmission Unit of the VIF in bytes. This parameter is read-only, but you can override the MTU setting with the <i>mtu</i> key using the other-config <i>map</i> parameter. For example, to reset the MTU on a virtual NIC to use jumbo frames: <pre>xe vif-param-set \ uuid=<vif_uuid> \ other-config:mtu=9000</pre>	read only
currently-attached	true if the device is currently attached	read only
qos_algorithm_type	QoS algorithm to use	read/write
qos_algorithm_params	parameters for the chosen QoS algorithm	read/write map parameter
qos_supported_algorithms	supported QoS algorithms for this VIF	read only set parameter
MAC-autogenerated	True if the MAC address of the VIF was automatically generated	read only
other-config	Additional configuration key:value pairs	read/write map parameter
other-config:ethtool-rx	set to on to enable receive checksum, off to disable	read write

Parameter Name	Description	Type
other-config:ethtool-tx	set to ON to enable transmit checksum, off to disable	read write
other-config:ethtool-sg	set to ON to enable scatter gather, off to disable	read write
other-config:ethtool-tso	set to ON to enable tcp segmentation offload, off to disable	read write
other-config:ethtool-ufo	set to ON to enable udp fragment offload, off to disable	read write
other-config:ethtool-gso	set to ON to enable generic segmentation offload, off to disable	read write
other-config:promiscuous	true to a VIF to be promiscuous on the bridge, so that it sees all traffic over the bridge. Useful for running an Intrusion Detection System (IDS) or similar in a VM.	read write
network-uuid	the unique identifier/object reference of the virtual network to which this VIF is connected	read only
network-name-label	the descriptive name of the virtual network to which this VIF is connected	read only
io_read_kbs	average read rate in kB/s for this VIF	read only
io_write_kbs	average write rate in kB/s for this VIF	read only
locking_mode	Affects the VIFs ability to filter traffic to/from a list of MAC and IP addresses. Requires additional parameters.	read/write
locking_mode:default	Varies according to the default locking mode for the VIF network. If the <i>default-locking-mode</i> is set to <i><disabled></i> , XenServer applies a filtering rule so that the VIF cannot send or receive traffic. If the <i>default-locking-mode</i> is set to <i><unlocked></i> , XenServer removes all the filtering rules associated with the VIF. See also Section A.4.11, "Network Commands" .	read/write
locking_mode:locked	Only traffic sent to or sent from the specified MAC and IP addresses is allowed on the VIF. If no IP addresses are specified, no traffic is allowed.	read/write
locking_mode:unlocked	No filters are applied to any traffic going to or from the VIF.	read/write
locking_mode:disabled	XenServer applies a filtering rule is applied so that the VIF drops all traffic.	read/write

A.4.24.2. vif-create

```
vif-create vm-uuid=<uuid_of_the_vm> device=<see below>
network-uuid=<uuid_of_the_network_the_vif_will_connect_to> [mac=<mac_address>]
```

Create a new VIF on a VM.

Appropriate values for the *device* field are listed in the parameter *allowed-VIF-devices* on the specified VM. Before any VIFs exist there, the allowable values are integers from 0-15.

The *mac* parameter is the standard MAC address in the form aa : bb : cc : dd : ee : ff. If you leave it unspecified, an appropriate random MAC address will be created. You can also explicitly set a random MAC address by specifying *mac=random*.

A.4.24.3. vif-destroy

```
vif-destroy uuid=<uuid_of_vif>
```

Destroy a VIF.

A.4.24.4. vif-plug

```
vif-plug uuid=<uuid_of_vif>
```

Attempt to attach the VIF while the VM is in the running state.

A.4.24.5. vif-unplug

```
vif-unplug uuid=<uuid_of_vif>
```

Attempts to detach the VIF from the VM while it is running.

A.4.24.6. vif-configure-ipv4

Configure IPv4 settings for this virtual interface. Set IPv4 settings as below:

```
vif-configure-ipv4      uuid=<uuid_of_vif>      mode=<static>      address=<CIDR_address>
gateway=<gateway_address>
```

For example:

```
VIF.configure_ipv4(vifObject,"static", " 192.168.1.10/24", " 192.168.1.1")
```

Clean IPv4 settings as below:

```
vif-configure-ipv4 uuid=<uuid_of_vif> mode=<none>
```

A.4.24.7. vif-configure-ipv6

Configure IPv6 settings for this virtual interface. Set IPv6 settings as below:

```
vif-configure-ipv6      uuid=<uuid_of_vif>      mode=<static>      address=<IP_address>
gateway=<gateway_address>
```

For example:

```
VIF.configure_ipv6(vifObject,"static", "fd06:7768:b9e5:8b00::5001/64", "fd06:7768:b9e5:8b00::1")
```

Clean IPv6 settings as below:

```
vif-configure-ipv6 uuid=<uuid_of_vif> mode=<none>
```

A.4.25. VLAN Commands

Commands for working with VLANs (virtual networks). To list and edit virtual interfaces, refer to the PIF commands, which have a VLAN parameter to signal that they have an associated virtual network (see [Section A.4.14, “PIF Commands”](#)). For example, to list VLANs you need to use **xe pif-list**.

A.4.25.1. vlan-create

```
vlan-create pif-uuid=<uuid_of_pif> vlan=<vlan_number> network-uuid=<uuid_of_network>
```

Create a new VLAN on a XenServer host.

A.4.25.2. pool-vlan-create

```
pool-vlan-create pif-uuid=<uuid_of_pif> vlan=<vlan_number> network-uuid=<uuid_of_network>
```


Create a new VLAN on all hosts on a pool, by determining which interface (for example, eth0) the specified network is on (on each host) and creating and plugging a new PIF object one each host accordingly.

A.4.25.3. vlan-destroy

`vlan-destroy uuid=<uuid_of_pif_mapped_to_vlan>`

Destroy a VLAN. Requires the UUID of the PIF that represents the VLAN.

A.4.26. VM Commands

Commands for controlling VMs and their attributes.

A.4.26.1. VM Selectors

Several of the commands listed here have a common mechanism for selecting one or more VMs on which to perform the operation. The simplest way is by supplying the argument `vm=<name_or_uuid>`. An easy way to get the uuid of an actual VM is to, for example, execute `xe vm-list power-state=running`. (The full list of fields that can be matched can be obtained by the command `xe vm-list params=all`.) For example, specifying `power-state=halted` will select all VMs whose `power-state` parameter is equal to `halted`. Where multiple VMs are matching, the option `--multiple` must be specified to perform the operation. The full list of parameters that can be matched is described at the beginning of this section, and can be obtained by the command `xe vm-list params=all`.

The VM objects can be listed with the standard object listing command (`xe vm-list`), and the parameters manipulated with the standard parameter commands. See [Section A.3.2, “Low-level Parameter Commands”](#) for details.

A.4.26.2. VM Parameters

VMs have the following parameters:

Note:

All writable VM parameter values can be changed while the VM is running, but the new parameters are *not* applied dynamically and will not be applied until the VM is rebooted.

Parameter Name	Description	Type
appliance	the appliance/vApp to which the VM belongs	read/write
uuid	the unique identifier/object reference for the VM	read only
name-label	the name of the VM	read/write
name-description	the description string of the VM	read/write
order	start order for vApp startup/shutdown and for startup after HA failover	read/write
version	the number of times this VM has been recovered - if a user wishes to over write a new VM with an older version, then they must call vm-recover	read only
user-version	string for creators of VMs and templates to put version information	read/write

Parameter Name	Description	Type
is-a-template	False unless this is a template; template VMs can never be started, they are used only for cloning other VMs. After this value has been set to true it cannot be reset to false. Template VMs cannot be converted into VMs using this parameter.	read/write
is-control-domain	True if this is a control domain (domain 0 or a driver domain)	read only
power-state	current power state	read only
start delay	the delay to wait before a call to start up the VM returns	read/write
shutdown-delay	the delay to wait before a call to shutdown the VM returns	read/write
memory-dynamic-max	dynamic maximum in bytes	read/write
memory-dynamic-min	dynamic minimum in bytes	read/write
memory-static-max	statically-set (absolute) maximum in bytes. If you want to change this value, the VM must be shut down.	read/write
memory-static-min	statically-set (absolute) minimum in bytes. If you want to change this value, the VM must be shut down.	read/write
suspend-VDI-uuid	the VDI that a suspend image is stored on	read only

Parameter Name	Description	Type
VCPUs-params	configuration parameters for the selected VCPU policy. You can tune a VCPU's pinning with <pre>xe vm-param-set \ uuid=<vm_uuid> \ VCPUs-params:mask=1,2,3</pre> The selected VM will then run on physical CPUs 1, 2, and 3 only. You can also tune the VCPU priority (xen scheduling) with the <i>cap</i> and <i>weight</i> parameters; for example <pre>xe vm-param-set \ uuid=<vm_uuid> \ VCPUs-params:weight=512 xe vm-param-set \ uuid=<vm_uuid> \ VCPUs-params:cap=100</pre> A VM with a weight of 512 will get twice as much CPU as a domain with a weight of 256 on a contended XenServer host. Legal weights range from 1 to 65535 and the default is 256. The <i>cap</i> optionally fixes the maximum amount of CPU a VM will be able to consume, even if the XenServer host has idle CPU cycles. The <i>cap</i> is expressed in percentage of one physical CPU: 100 is 1 physical CPU, 50 is half a CPU, 400 is 4 CPUs, etc. The default, 0, means there is no upper cap.	read/write map parameter
VCPUs-max	maximum number of virtual CPUs.	read/write
VCPUs-at-startup	boot number of virtual CPUs	read/write

Parameter Name	Description	Type
actions-after-crash	action to take if the VM crashes. For PV guests, valid parameters are: preserve (for analysis only), coredump_and_restart (record a coredump and reboot VM), coredump_and_destroy (record a coredump and leave VM halted), restart (no coredump and restart VM), and destroy (no coredump and leave VM halted).	read/write
console-uuids	virtual console devices	read only set parameter
platform	platform-specific configuration	read/write map parameter
allowed-operations	list of the operations allowed in this state	read only set parameter
current-operations	a list of the operations that are currently in progress on the VM	read only set parameter
allowed-VBD-devices	list of VBD identifiers available for use, represented by integers of the range 0-15. This list is informational only, and other devices may be used (but may not work).	read only set parameter
allowed-VIF-devices	list of VIF identifiers available for use, represented by integers of the range 0-15. This list is informational only, and other devices may be used (but may not work).	read only set parameter
HVM-boot-policy	the boot policy for HVM guests. Either BIOS Order or an empty string.	read/write
HVM-boot-params	the <i>order</i> key controls the HVM guest boot order, represented as a string where each character is a boot method: d for the CD/DVD, c for the root disk, and n for network PXE boot. The default is dc.	read/write map parameter
HVM-shadow-multiplier	Floating point value which controls the amount of shadow memory overhead to grant the VM. Defaults to 1.0 (the minimum value), and should only be changed by advanced users.	read/write

Parameter Name	Description	Type
PV-kernel	path to the kernel	read/write
PV-ramdisk	path to the initrd	read/write
PV-args	string of kernel command line arguments	read/write
PV-legacy-args	string of arguments to make legacy VMs boot	read/write
PV-bootloader	name of or path to bootloader	read/write
PV-bootloader-args	string of miscellaneous arguments for the bootloader	read/write
last-boot-CPU-flags	describes the CPU flags on which the VM was last booted	read only
resident-on	the XenServer host on which a VM is currently resident	read only
affinity	a XenServer host which the VM has preference for running on; used by the xe vm-start command to decide where to run the VM	read/write
other-config	A list of key/value pairs that specify additional configuration parameters for the VM For example, a VM will be started automatically after host boot if the other-config parameter includes the key/value pair <i>auto_poweron: true</i>	read/write map parameter
start-time	timestamp of the date and time that the metrics for the VM were read, in the form <code>yyyymmddThh:mm:ss Z</code> , where Z is the single-letter military timezone indicator, for example, Z for UTC (GMT)	read only
install-time	timestamp of the date and time that the metrics for the VM were read, in the form <code>yyyymmddThh:mm:ss Z</code> , where Z is the single-letter military timezone indicator, for example, Z for UTC (GMT)	read only
memory-actual	the actual memory being used by a VM	read only

Parameter Name	Description	Type
VCPUs-number	the number of virtual CPUs assigned to the VM For a paravirtualized Linux VM, this number can differ from <i>VCPUS-max</i> and can be changed without rebooting the VM using the vm-vcpu-hotplug command. See Section A.4.26.34, “vm-vcpu-hotplug”. Windows VMs always run with the number of vCPUs set to <i>VCPUS-max</i> and must be rebooted to change this value. Note that performance will drop sharply if you set <i>VCPUS-number</i> to a value greater than the number of physical CPUs on the XenServer host.	read only
VCPUs-Utilization	a list of virtual CPUs and their weight	read only map parameter
os-version	the version of the operating system for the VM	read only map parameter
PV-drivers-version	the versions of the paravirtualized drivers for the VM	read only map parameter
PV-drivers-detected	flag for latest version of the paravirtualized drivers for the VM	read only
memory	memory metrics reported by the agent on the VM	read only map parameter
disks	disk metrics reported by the agent on the VM	read only map parameter
networks	network metrics reported by the agent on the VM	read only map parameter
other	other metrics reported by the agent on the VM	read only map parameter
guest-metrics-last-updated	timestamp when the last write to these fields was performed by the in-guest agent, in the form <i>yyyymmddThh:mm:ssZ</i> , where Z is the single-letter military timezone indicator, for example, Z for UTC (GMT)	read only

Parameter Name	Description	Type
actions-after-shutdown	action to take after the VM has shutdown	read/write
actions-after-reboot	action to take after the VM has rebooted	read/write
possible-hosts	potential hosts of this VM	read only
dom-id	domain ID (if available, -1 otherwise)	read only
recommendations	XML specification of recommended values and ranges for properties of this VM	read only
xenstore-data	data to be inserted into the xenstore tree (/local/domain/<domid>/vm-data) after the VM is created	read/write map parameter
is-a-snapshot	True if this VM is a snapshot	read only
snapshot_of	the UUID of the VM this is a snapshot of	read only
snapshots	the UUID(s) of all snapshots of this VM	read only
snapshot_time	the timestamp of the snapshot operation that created this VM snapshot	read only
memory-target	the target amount of memory set for this VM	read only
blocked-operations	lists the operations that cannot be performed on this VM	read/write map parameter
last-boot-record	record of the last boot parameters for this template, in XML format	read only
ha-always-run	True if this VM will always restarted on another host in case of the failure of the host it is resident on	read/write
ha-restart-priority	restart or best effort	read/write
blobs	binary data store	read only
live	True if the VM is running, false if HA suspects that the VM may not be running.	read only

A.4.26.3. vm-assert-can-be-recovered

vm-assert-can-be-recovered <uuid> [<database>] <vdi-uuid>

Tests whether storage is available to recover this VM.

A.4.26.4. vm-cd-add

```
vm-cd-add cd-name=<name_of_new_cd> device=<integer_value_of_an_available_vbd>
[<vm-selector>=<vm_selector_value>...]
```

Add a new virtual CD to the selected VM. The *device* parameter should be selected from the value of the *allowed-VBD-devices* parameter of the VM.

The VM or VMs on which this operation should be performed are selected using the standard selection mechanism (see [VM selectors](#)). Optional arguments can be any number of the [VM parameters](#) listed at the beginning of this section.

A.4.26.5. vm-cd-eject

```
vm-cd-eject [<vm-selector>=<vm_selector_value>...]
```

Eject a CD from the virtual CD drive. This command only works if there is one and only one CD attached to the VM. When there are two or more CDs, use the command **xe vbd-eject** and specify the UUID of the VBD.

The VM or VMs on which this operation should be performed are selected using the standard selection mechanism (see [VM selectors](#)). Optional arguments can be any number of the [VM parameters](#) listed at the beginning of this section.

A.4.26.6. vm-cd-insert

```
vm-cd-insert cd-name=<name_of_cd> [<vm-selector>=<vm_selector_value>...]
```

Insert a CD into the virtual CD drive. This command will only work if there is one and only one empty CD device attached to the VM. When there are two or more empty CD devices, use the **xe vbd-insert** command and specify the UUIDs of the VBD and of the VDI to insert.

The VM or VMs on which this operation should be performed are selected using the standard selection mechanism (see [VM selectors](#)). Optional arguments can be any number of the [VM parameters](#) listed at the beginning of this section.

A.4.26.7. vm-cd-list

```
vm-cd-list [vbd-params] [vdi-params] [<vm-selector>=<vm_selector_value>...]
```

Lists CDs attached to the specified VMs.

The VM or VMs on which this operation should be performed are selected using the standard selection mechanism (see [VM selectors](#)). Optional arguments can be any number of the [VM parameters](#) listed at the beginning of this section.

You can also select which VBD and VDI parameters to list.

A.4.26.8. vm-cd-remove

```
vm-cd-remove cd-name=<name_of_cd> [<vm-selector>=<vm_selector_value>...]
```

Remove a virtual CD from the specified VMs.

The VM or VMs on which this operation should be performed are selected using the standard selection mechanism (see [VM selectors](#)). Optional arguments can be any number of the [VM parameters](#) listed at the beginning of this section.

A.4.26.9. vm-clone

```
vm-clone new-name-label=<name_for_clone>
[new-name-description=<description_for_clone>] [<vm-selector>=<vm_selector_value>...]
```

Clone an existing VM, using storage-level fast disk clone operation where available. Specify the name and the optional description for the resulting cloned VM using the *new-name-label* and *new-name-description* arguments.

The VM or VMs on which this operation should be performed are selected using the standard selection mechanism (see [VM selectors](#)). Optional arguments can be any number of the [VM parameters](#) listed at the beginning of this section.

A.4.26.10. vm-compute-maximum-memory

```
vm-compute-maximum-memory total=<amount_of_available_physical_ram_in_bytes>
[approximate=<add overhead memory for additional vCPUS? true | false>]
[<vm-selector>=<vm_selector_value>...]
```

Calculate the maximum amount of static memory which can be allocated to an existing VM, using the total amount of physical RAM as an upper bound. The optional parameter *approximate* reserves sufficient extra memory in the calculation to account for adding extra vCPUs into the VM at a later date.

For example:

```
xe vm-compute-maximum-memory vm=testvm total=`xe host-list params=memory-free --minimal`
```

This command uses the value of the *memory-free* parameter returned by the **xe host-list** command to set the maximum memory of the VM named *testvm*.

The VM or VMs on which this operation will be performed are selected using the standard selection mechanism (see [VM selectors](#)). Optional arguments can be any number of the [VM parameters](#) listed at the beginning of this section.

A.4.26.11. vm-copy

```
vm-copy new-name-label=<name_for_copy> [new-name-description=<description_for_copy>]
[sr-uuid=<uuid_of_sr>] [<vm-selector>=<vm_selector_value>...]
```

Copy an existing VM, but without using storage-level fast disk clone operation (even if this is available). The disk images of the copied VM are guaranteed to be "full images", that is, not part of a copy-on-write (CoW) chain.

Specify the name and the optional description for the resulting copied VM using the *new-name-label* and *new-name-description* arguments.

Specify the destination SR for the resulting copied VM using the *sr-uuid*. If this parameter is not specified, the destination is the same SR that the original VM is in.

The VM or VMs on which this operation should be performed are selected using the standard selection mechanism (see [VM selectors](#)). Optional arguments can be any number of the [VM parameters](#) listed at the beginning of this section.

A.4.26.12. vm-crashdump-list

```
vm-crashdump-list [<vm-selector>=<vm_selector_value>...]
```

List crashdumps associated with the specified VMs.

If the optional argument *params* is used, the value of *params* is a string containing a list of parameters of this object that you want to display. Alternatively, you can use the keyword *all* to show all parameters. If *params* is not used, the returned list shows a default subset of all available parameters.

The VM or VMs on which this operation should be performed are selected using the standard selection mechanism (see [VM selectors](#)). Optional arguments can be any number of the [VM parameters](#) listed at the beginning of this section.

A.4.26.13. vm-data-source-list

```
vm-data-source-list [<vm-selector>=<vm selector value>...]
```

List the data sources that can be recorded for a VM.

Select the VM(s) on which to perform this operation by using the standard selection mechanism (see [VM selectors](#)). Optional arguments can be any number of the [VM parameters](#) listed at the beginning of this section. If no parameters to select hosts are given, the operation will be performed on all VMs.

Data sources have two parameters — `standard` and `enabled` — which can be seen by the output of this command. If a data source has `enabled` set to `true`, then the metrics are currently being recorded to the performance database. If a data source has `standard` set to `true`, then the metrics are recorded to the performance database *by default* (and so, `enabled` will also be set to `true` for this data source). If a data source has `standard` set to `false`, then the metrics are *not* recorded to the performance database by default (and so, `enabled` will also be set to `false` for this data source).

To start recording data source metrics to the performance database, run the **vm-data-source-record** command. This will set `enabled` to `true`. To stop, run the **vm-data-source-forget**. This will set `enabled` to `false`.

A.4.26.14. vm-data-source-record

```
vm-data-source-record data-source=<name_description_of_data-source> [<vm-selector>=<vm selector value>...]
```

Record the specified data source for a VM.

This operation writes the information from the data source to the persistent performance metrics database of the specified VM(s). For performance reasons, this database is distinct from the normal agent database.

Select the VM(s) on which to perform this operation by using the standard selection mechanism (see [VM selectors](#)). Optional arguments can be any number of the [VM parameters](#) listed at the beginning of this section. If no parameters to select hosts are given, the operation will be performed on all VMs.

A.4.26.15. vm-data-source-forget

```
vm-data-source-forget data-source=<name_description_of_data-source> [<vm-selector>=<vm selector value>...]
```

Stop recording the specified data source for a VM and forget all of the recorded data.

Select the VM(s) on which to perform this operation by using the standard selection mechanism (see [VM selectors](#)). Optional arguments can be any number of the [VM parameters](#) listed at the beginning of this section. If no parameters to select hosts are given, the operation will be performed on all VMs.

A.4.26.16. vm-data-source-query

```
vm-data-source-query data-source=<name_description_of_data-source> [<vm-selector>=<vm selector value>...]
```

Display the specified data source for a VM.

Select the VM(s) on which to perform this operation by using the standard selection mechanism (see [VM selectors](#)). Optional arguments can be any number of the [VM parameters](#) listed at the beginning of this section. If no parameters to select hosts are given, the operation will be performed on all VMs.

A.4.26.17. vm-destroy

```
vm-destroy uuid=<uuid_of_vm>
```

Destroy the specified VM. This leaves the storage associated with the VM intact. To delete storage as well, use **xe vm-uninstall**.

A.4.26.18. vm-disk-add

```
vm-disk-add disk-size=<size_of_disk_to_add> device=<uuid_of_device>
[<vm-selector>=<vm_selector_value>...]
```

Add a new disk to the specified VMs. Select the *device* parameter from the value of the *allowed-VBD-devices* parameter of the VMs.

The *disk-size* parameter can be specified in bytes or using the IEC standard suffixes KiB (2^{10} bytes), MiB (2^{20} bytes), GiB (2^{30} bytes), and TiB (2^{40} bytes).

The VM or VMs on which this operation should be performed are selected using the standard selection mechanism (see [VM selectors](#)). Optional arguments can be any number of the [VM parameters](#) listed at the beginning of this section.

A.4.26.19. vm-disk-list

```
vm-disk-list [vbd-params] [vdi-params] [<vm-selector>=<vm_selector_value>...]
```

Lists disks attached to the specified VMs. The *vbd-params* and *vdi-params* parameters control the fields of the respective objects to output and should be given as a comma-separated list, or the special key *all* for the complete list.

The VM or VMs on which this operation should be performed are selected using the standard selection mechanism (see [VM selectors](#)). Optional arguments can be any number of the [VM parameters](#) listed at the beginning of this section.

A.4.26.20. vm-disk-remove

```
vm-disk-remove device=<integer_label_of_disk> [<vm-selector>=<vm_selector_value>...]
```

Remove a disk from the specified VMs and destroy it.

The VM or VMs on which this operation should be performed are selected using the standard selection mechanism (see [VM selectors](#)). Optional arguments can be any number of the [VM parameters](#) listed at the beginning of this section.

A.4.26.21. vm-export

```
vm-export filename=<export_filename>
[metadata=<true | false>]
[<vm-selector>=<vm_selector_value>...]
```

Export the specified VMs (including disk images) to a file on the local machine. Specify the filename to export the VM into using the *filename* parameter. By convention, the filename should have a *.xva* extension.

If the *metadata* parameter is *true*, then the disks are not exported, and only the VM metadata is written to the output file. This is intended to be used when the underlying storage is transferred through other mechanisms, and permits the VM information to be recreated (see [Section A.4.26.22, “vm-import”](#)).

The VM or VMs on which this operation should be performed are selected using the standard selection mechanism (see [VM selectors](#)). Optional arguments can be any number of the [VM parameters](#) listed at the beginning of this section.

A.4.26.22. vm-import

```
vm-import filename=<export_filename>
[metadata=<true | false>]
[preserve=<true | false>]
[sr-uuid=<destination_sr_uuid>]
```

Import a VM from a previously-exported file. If *preserve* is set to *true*, the MAC address of the original VM will be preserved. The *sr-uuid* determines the destination SR to import the VM into, and is the default SR if not specified.

The *filename* parameter can also point to an XVA-format VM, which is the legacy export format from XenServer 3.2 and is used by some third-party vendors to provide virtual appliances. This format uses a directory to store the VM data, so set *filename* to the root directory of the XVA export and not an actual file. Subsequent exports of the imported legacy guest will automatically be upgraded to the new filename-based format, which stores much more data about the configuration of the VM.

Note:

The older directory-based XVA format does not fully preserve all the VM attributes. In particular, imported VMs will not have any virtual network interfaces attached by default. If networking is required, create one using **vif-create** and **vif-plug**.

If the *metadata* is *true*, then a previously exported set of metadata can be imported without their associated disk blocks. Metadata-only import will fail if any VDIs cannot be found (named by SR and *VDI.location*) unless the *--force* option is specified, in which case the import will proceed regardless. If disks can be mirrored or moved out-of-band then metadata import/export represents a fast way of moving VMs between disjoint pools (e.g. as part of a disaster recovery plan).

Note:

Multiple VM imports will be performed faster in serial than in parallel.

A.4.26.23. vm-install

```
vm-install new-name-label=<name>
[template-uuid=<uuid_of_desired_template> | [template=<uuid_or_name_of_desired_template>]]
[sr-uuid=<sr_uuid> | sr-name-label=<name_of_sr> ]
[copy-bios-strings-from=<uuid of host> ]
```

Install or clone a VM from a template. Specify the template name using either the *template-uuid* or *template* argument. Specify an SR using either the *sr-uuid* or *sr-name-label* argument. Specify to install BIOS-locked media using the *copy-bios-strings-from* argument.

Note:

When installing from a template that has existing disks, by default, new disks will be created in the same SR as these existing disks. Where the SR supports it, these will be fast copies. If a different SR is specified on the command line, the new disks will be created there. In this case a fast copy is not possible and the disks will be full copies.

When installing from a template that does not have existing disks, any new disks will be created in the SR specified, or the pool default SR if not specified.

A.4.26.24. vm-memory-shadow-multiplier-set

```
vm-memory-shadow-multiplier-set [<vm-selector>=<vm_selector_value>...]
[multiplier=<float_memory_multiplier>]
```

Set the shadow memory multiplier for the specified VM.

This is an advanced option which modifies the amount of *shadow memory* assigned to a hardware-assisted VM. In some specialized application workloads, such as Citrix XenApp, extra shadow memory is required to achieve full performance.

This memory is considered to be an overhead. It is separated from the normal memory calculations for accounting memory to a VM. When this command is invoked, the amount of free XenServer host memory will decrease according to the multiplier, and the *HVM_shadow_multiplier* field will be updated with the actual value which Xen has assigned to the VM. If there is not enough XenServer host memory free, then an error will be returned.

The VMs on which this operation should be performed are selected using the standard selection mechanism (see [VM selectors](#) for more information).

A.4.26.25. vm-migrate

```
vm-migrate [<copy>=<true/false>] [host-uuid=<destination_host_UUID>] [host=<name or UUID of
destination host>] [<force>=<true/false>] [<live>=<true/false>] [<vm-selector>=<vm_selector_value>...]
[<remote-master>=<destination_pool_master_uuid>] [<remote-username>=<destination_pool_username>]
[<remote-password>=<destination_pool_password>] [<remote-network>=<destination_pool_network_uuid>]
[<vif:>=<vif_uuid>] [<vdi>=<vdi_uuid>]
```

This command migrates the specified VMs between physical hosts. The *host* parameter can be either the name or the UUID of the XenServer host. For example, to migrate the VM to another host in the pool, where the VM disks are located on storage shared by both hosts (known as XenMotion):

```
xe vm-migrate uuid=<vm_uuid>
host-uuid=<host_uuid>
```

To move VMs between hosts in the same pool, which do **not** share storage (Storage XenMotion):

```
xe vm-migrate uuid=<vm_uuid> remote-master=12.34.56.78 \
remote-username=<username> remote-password=<password> \
host-uuid=<desination_host_uuid> vdi=<vdi_uuid>
```

You can choose the SR where each VDI gets stored:

```
xe vm-migrate uuid=<vm_uuid> host-uuid=<destination_host_uuid> \
vdi1:<vdi_1_uuid>=<destination_sr_uuid> \
vdi2:<vdi_2_uuid>=<destination_sr2_uuid> \
vdi3:<vdi_3_uuid>=<destination_sr3_uuid>
```

Additionally, you can choose which network to attach the VM after migration:

```
xe vm-migrate uuid=<vm_uuid> \
vdi1:<vdi_1_uuid>=<destination_sr_uuid> \
vdi2:<vdi_2_uuid>=<destination_sr2_uuid> \
vdi3:<vdi_3_uuid>=<destination_sr3_uuid> \
vif:<vif_uuid>=<network_uuid>
```

For cross-pool migration:

```
xe vm-migrate uuid=<vm_uuid> remote-master=12.34.56.78
remote-username=<username> remote-password=<password> \
host-uuid=<desination_host_uuid> vdi=<vdi_uuid>
```

For more information on Storage XenMotion, XenMotion and Live VDI Migration, refer to the *VM User's Guide*.

By default, the VM will be suspended, migrated, and resumed on the other host. The *live* parameter activates XenMotion and keeps the VM running while performing the migration, thus minimizing VM downtime to less than a second. In some circumstances such as extremely memory-heavy workloads in the VM, XenMotion automatically falls back into the default mode and suspends the VM for a brief period of time before completing the memory transfer.

The VM or VMs on which this operation should be performed are selected using the standard selection mechanism (see [VM selectors](#)). Optional arguments can be any number of the [VM parameters](#) listed at the beginning of this section.

A.4.26.26. vm-reboot

```
vm-reboot [<vm-selector>=<vm_selector_value>...] [force=<true>]
```

Reboot the specified VMs.

The VM or VMs on which this operation should be performed are selected using the standard selection mechanism (see [VM selectors](#)). Optional arguments can be any number of the [VM parameters](#) listed at the beginning of this section.

Use the *force* argument to cause an ungraceful shutdown, akin to pulling the plug on a physical server.

A.4.26.27. vm-recover

```
vm-recover <vm-uuid> [<database>] [<vdi-uuid>] [<force>]
```

Recovers a VM from the database contained in the supplied VDI.

A.4.26.28. vm-reset-powerstate

```
vm-reset-powerstate [<vm-selector>=<vm_selector_value>...] {force=true}
```

The VM or VMs on which this operation should be performed are selected using the standard selection mechanism (see [VM selectors](#)). Optional arguments can be any number of the [VM parameters](#) listed at the beginning of this section.

This is an *advanced* command only to be used when a member host in a pool goes down. You can use this command to force the pool master to reset the power-state of the VMs to be `halted`. Essentially this forces the lock on the VM and its disks so it can be subsequently started on another pool host. This call *requires* the force flag to be specified, and fails if it is not on the command-line.

A.4.26.29. vm-resume

```
vm-resume [<vm-selector>=<vm_selector_value>...] [force=<true | false>] [on=<XenServer host UUID>]
```

Resume the specified VMs.

The VM or VMs on which this operation should be performed are selected using the standard selection mechanism (see [VM selectors](#)). Optional arguments can be any number of the [VM parameters](#) listed at the beginning of this section.

If the VM is on a shared SR in a pool of hosts, use the *on* argument to specify which host in the pool on which to start it. By default the system will determine an appropriate host, which might be any of the members of the pool.

A.4.26.30. vm-shutdown

```
vm-shutdown [<vm-selector>=<vm_selector_value>...] [force=<true | false>]
```

Shut down the specified VM.

The VM or VMs on which this operation should be performed are selected using the standard selection mechanism (see [VM selectors](#)). Optional arguments can be any number of the [VM parameters](#) listed at the beginning of this section.

Use the *force* argument to cause an ungraceful shutdown, similar to pulling the plug on a physical server.

A.4.26.31. vm-start

```
vm-start [<vm-selector>=<vm_selector_value>...] [force=<true | false>] [on=<XenServer host UUID>] [--multiple]
```

Start the specified VMs.

The VM or VMs on which this operation should be performed are selected using the standard selection mechanism (see [VM selectors](#)). Optional arguments can be any number of the [VM parameters](#) listed at the beginning of this section.

If the VMs are on a shared SR in a pool of hosts, use the *on* argument to specify which host in the pool on which to start the VMs. By default the system will determine an appropriate host, which might be any of the members of the pool.

A.4.26.32. vm-suspend

```
vm-suspend [<vm-selector>=<vm_selector_value>...]
```

Suspend the specified VM.

The VM or VMs on which this operation should be performed are selected using the standard selection mechanism (see [VM selectors](#)). Optional arguments can be any number of the [VM parameters](#) listed at the beginning of this section.

A.4.26.33. vm-uninstall

```
vm-uninstall [<vm-selector>=<vm_selector_value>...] [force=<true | false>]
```

Uninstall a VM, destroying its disks (those VDI's that are marked RW and connected to this VM only) as well as its metadata record. To simply destroy the VM metadata, use **xe vm-destroy**.

The VM or VMs on which this operation should be performed are selected using the standard selection mechanism (see [VM selectors](#)). Optional arguments can be any number of the [VM parameters](#) listed at the beginning of this section.

A.4.26.34. vm-vcpu-hotplug

```
vm-vcpu-hotplug new-vcpus=<new_vcpu_count> [<vm-selector>=<vm_selector_value>...]
```

Dynamically adjust the number of VCPUs available to a running paravirtual Linux VM within the number bounded by the parameter *VCPUS-max*. Windows VMs always run with the number of VCPUs set to *VCPUS-max* and must be rebooted to change this value.

The paravirtualized Linux VM or VMs on which this operation should be performed are selected using the standard selection mechanism (see [VM selectors](#)). Optional arguments can be any number of the [VM parameters](#) listed at the beginning of this section.

Note:

When running certain Linux VMs (such as Debian) without XenServer Tools installed, customers should run the following command on the VM as *root* to ensure the newly hotplugged vCPUs are used:

```
# for i in /sys/devices/system/cpu/cpu[1-9]*/online; do if
[ "$(cat $i)" = 0 ]; then echo 1 > $i; fi; done
```

A.4.26.35. vm-vif-list

```
vm-vif-list [<vm-selector>=<vm_selector_value>...]
```



Lists the VIFs from the specified VMs.

The VM or VMs on which this operation should be performed are selected using the standard selection mechanism (see [VM selectors](#)). Note that the selectors operate on the VM records when filtering, and *not* on the VIF values. Optional arguments can be any number of the [VM parameters](#) listed at the beginning of this section.