

Oracle® Fabric Manager 4.3.0 User's Guide

ORACLE®

Part No: E49557-03
May 2015

Part No: E49557-03

Copyright © 2014, 2015, Oracle and/or its affiliates. All rights reserved.

This software and related documentation are provided under a license agreement containing restrictions on use and disclosure and are protected by intellectual property laws. Except as expressly permitted in your license agreement or allowed by law, you may not use, copy, reproduce, translate, broadcast, modify, license, transmit, distribute, exhibit, perform, publish, or display any part, in any form, or by any means. Reverse engineering, disassembly, or decompilation of this software, unless required by law for interoperability, is prohibited.

The information contained herein is subject to change without notice and is not warranted to be error-free. If you find any errors, please report them to us in writing.

If this is software or related documentation that is delivered to the U.S. Government or anyone licensing it on behalf of the U.S. Government, the following notice is applicable:

U.S. GOVERNMENT END USERS. Oracle programs, including any operating system, integrated software, any programs installed on the hardware, and/or documentation, delivered to U.S. Government end users are "commercial computer software" pursuant to the applicable Federal Acquisition Regulation and agency-specific supplemental regulations. As such, use, duplication, disclosure, modification, and adaptation of the programs, including any operating system, integrated software, any programs installed on the hardware, and/or documentation, shall be subject to license terms and license restrictions applicable to the programs. No other rights are granted to the U.S. Government.

This software or hardware is developed for general use in a variety of information management applications. It is not developed or intended for use in any inherently dangerous applications, including applications that may create a risk of personal injury. If you use this software or hardware in dangerous applications, then you shall be responsible to take all appropriate fail-safe, backup, redundancy, and other measures to ensure its safe use. Oracle Corporation and its affiliates disclaim any liability for any damages caused by use of this software or hardware in dangerous applications.

Oracle and Java are registered trademarks of Oracle and/or its affiliates. Other names may be trademarks of their respective owners.

Intel and Intel Xeon are trademarks or registered trademarks of Intel Corporation. All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. AMD, Opteron, the AMD logo, and the AMD Opteron logo are trademarks or registered trademarks of Advanced Micro Devices. UNIX is a registered trademark of The Open Group.

This software or hardware and documentation may provide access to or information on content, products, and services from third parties. Oracle Corporation and its affiliates are not responsible for and expressly disclaim all warranties of any kind with respect to third-party content, products, and services. Oracle Corporation and its affiliates will not be responsible for any loss, costs, or damages incurred due to your access to or use of third-party content, products, or services.

Référence: E49557-03

Copyright © 2014, 2015, Oracle et/ou ses affiliés. Tous droits réservés.

Ce logiciel et la documentation qui l'accompagne sont protégés par les lois sur la propriété intellectuelle. Ils sont concédés sous licence et soumis à des restrictions d'utilisation et de divulgation. Sauf disposition de votre contrat de licence ou de la loi, vous ne pouvez pas copier, reproduire, traduire, diffuser, modifier, breveter, transmettre, distribuer, exposer, exécuter, publier ou afficher le logiciel, même partiellement, sous quelque forme et par quelque procédé que ce soit. Par ailleurs, il est interdit de procéder à toute ingénierie inverse du logiciel, de le désassembler ou de le décompiler, excepté à des fins d'interopérabilité avec des logiciels tiers ou tel que prescrit par la loi.

Les informations fournies dans ce document sont susceptibles de modification sans préavis. Par ailleurs, Oracle Corporation ne garantit pas qu'elles soient exemptes d'erreurs et vous invite, le cas échéant, à lui en faire part par écrit.

Si ce logiciel, ou la documentation qui l'accompagne, est concédé sous licence au Gouvernement des Etats-Unis, ou à toute entité qui délivre la licence de ce logiciel ou l'utilise pour le compte du Gouvernement des Etats-Unis, la notice suivante s'applique:

U.S. GOVERNMENT END USERS. Oracle programs, including any operating system, integrated software, any programs installed on the hardware, and/or documentation, delivered to U.S. Government end users are "commercial computer software" pursuant to the applicable Federal Acquisition Regulation and agency-specific supplemental regulations. As such, use, duplication, disclosure, modification, and adaptation of the programs, including any operating system, integrated software, any programs installed on the hardware, and/or documentation, shall be subject to license terms and license restrictions applicable to the programs. No other rights are granted to the U.S. Government.

Ce logiciel ou matériel a été développé pour un usage général dans le cadre d'applications de gestion des informations. Ce logiciel ou matériel n'est pas conçu ni n'est destiné à être utilisé dans des applications à risque, notamment dans des applications pouvant causer des dommages corporels. Si vous utilisez ce logiciel ou matériel dans le cadre d'applications dangereuses, il est de votre responsabilité de prendre toutes les mesures de secours, de sauvegarde, de redondance et autres mesures nécessaires à son utilisation dans des conditions optimales de sécurité. Oracle Corporation et ses affiliés déclinent toute responsabilité quant aux dommages causés par l'utilisation de ce logiciel ou matériel pour ce type d'applications.

Oracle et Java sont des marques déposées d'Oracle Corporation et/ou de ses affiliés. Tout autre nom mentionné peut correspondre à des marques appartenant à d'autres propriétaires qu'Oracle.

Intel et Intel Xeon sont des marques ou des marques déposées d'Intel Corporation. Toutes les marques SPARC sont utilisées sous licence et sont des marques ou des marques déposées de SPARC International, Inc. AMD, Opteron, le logo AMD et le logo AMD Opteron sont des marques ou des marques déposées d'Advanced Micro Devices. UNIX est une marque déposée d'The Open Group.

Ce logiciel ou matériel et la documentation qui l'accompagne peuvent fournir des informations ou des liens donnant accès à des contenus, des produits et des services émanant de tiers. Oracle Corporation et ses affiliés déclinent toute responsabilité ou garantie expresse quant aux contenus, produits ou services émanant de tiers. En aucun cas, Oracle Corporation et ses affiliés ne sauraient être tenus pour responsables des pertes subies, des coûts occasionnés ou des dommages causés par l'accès à des contenus, produits ou services tiers, ou à leur utilisation.

Contents

Using This Documentation	19
Understanding Oracle Fabric Manager	21
Oracle Fabric Manager Overview	21
Support for the Fabric Interconnect F1-15	21
Support for the Fabric Interconnect F1-4	22
Support for the Oracle SDN Controller	22
Oracle Fabric Manager as a Stand-Alone Application	22
Installing or Updating Oracle Fabric Manager	25
Minimum Requirements	25
Oracle Fabric Manager Server Requirements	26
Client Requirements	26
▼ Preparing for Installation	27
Installing Oracle Fabric Manager	27
▼ Install Oracle Fabric Manager on an Oracle Solaris Server	27
Linux Server Installation Considerations	31
▼ Install Oracle Fabric Manager on a Linux Server (RPM File)	31
Windows Server Installation Considerations	32
▼ Install Oracle Fabric Manager on a Windows Server	33
Validating and Troubleshooting the Installation	36
▼ Check the Status of Oracle Fabric Manager Services	36
▼ Stop Oracle Fabric Manager Services	37
▼ Start Oracle Fabric Manager Services	39
Service Logs	40
Upgrading the Oracle Fabric Manager Software	40
▼ Upgrade Oracle Fabric Manager on an Oracle Solaris Server	40
▼ Upgrade Oracle Fabric Manager on a Windows or Linux Server	42
Removing the Oracle Fabric Manager Software	43
▼ Remove the Oracle Fabric Manager Software From an Oracle Solaris Server	43

▼ Remove the Oracle Fabric Manager Software From a Linux Server	45
▼ Remove the Oracle Fabric Manager Software From a Windows Server	46
Configuring a Certificate for Oracle Fabric Manager	46
Certificate Overview	46
Certificate Installation Verification	48
Creating and Installing a CA-Signed Certificate	49
Starting Oracle Fabric Manager and Logging In	55
Credentials for Intital Login	56
Oracle Fabric Manager and Authentication	57
▼ Log In to Oracle Fabric Manager	57
Working With the Interface	61
Banner Features	61
Understanding the Navigation Panel	62
Navigation Panel	62
General Option	65
Server Resource Manager	66
Network Cloud Manager	67
Storage Cloud Manager	68
Service Manager	69
Security Manager	70
Managed Devices	70
Plugins	71
Understanding the Work Panel	72
Work Panel	72
Filtering and Sorting Table Displays	74
Using the Maintenance Menu	77
Maintenance Menu	77
▼ Configure Backup Locations	77
▼ Back Up the Oracle Fabric Manager Configuration	79
▼ Back Up and Restore the Oracle Fabric Manager Configuration (Fresh Install)	81
▼ Restore the Oracle Fabric Manager Server Configuration	82
▼ Allow Unlisted User Access	83
▼ Download Oracle Fabric Manager Log Files	84
▼ Clean Up the Oracle Fabric Manager Database	88
Working With Oracle Fabric Manager Jobs	91
Understanding Jobs	91

Jobs Summary	91
Recent Jobs Summary	93
Displaying the Jobs Summary	94
▼ Display Job Summary	94
▼ Show All Jobs	94
▼ Show All Active Jobs	95
▼ Clear All Jobs in the Jobs Summary	96
▼ Cancel Selected Jobs in the Jobs Summary	97
Working With Network Clouds	99
Understanding Network Clouds	99
Network Clouds Overview	99
Network QoS Profiles and Network Clouds	100
QoS Assignment at Cloud Level and vNIC Level	100
Creating a Network Cloud	101
Network Cloud Creation Overview	101
▼ Create a Network Cloud	101
Fabric Device HA Designation	104
▼ Set the Fabric Device HA Designation for a Network Cloud	104
Configuring Advanced Features for a Network Cloud	107
Network Cloud Advanced Features	107
▼ Set Network QoS on a Network Cloud	107
▼ Set a VLAN for a Network Cloud	109
Private vNICs Overview	111
▼ Create Private vNICs on a Network Cloud	112
▼ Set the Port Allocation Policy	113
▼ Set Link Aggregation for a Network Cloud	115
Managing Network Cloud Properties	117
Summary Information for Network Clouds	118
Detailed Information for Network Clouds	118
▼ Edit Network Cloud Properties	120
▼ Rename a Network Cloud	121
▼ Push Cloud Changes to an I/O Template's vNICs	123
▼ Display the Ethernet Ports or LAGs in a Network Cloud	125
▼ Display the vNICs Associated With a Network Cloud	125
▼ Apply Network Cloud Changes to vNICs in the Cloud	126
▼ Terminate vNICs in a Network Cloud to Another vNIC Cloud	127
▼ Display the vNIC Templates Associated With a Network Cloud	128

Working With Storage Clouds	131
Understanding Storage Clouds	131
SAN QoS Profiles and Storage Clouds	131
QoS Assignment at Cloud Level and vHBA Level	132
▼ Create a Storage Cloud	132
Configuring Features in a Storage Cloud	135
▼ Set SAN QoS on a Storage Cloud	135
▼ Set an Allocation Policy for a Storage Cloud	138
▼ Change the Port Priority for a Storage Cloud	142
Displaying Storage Clouds	143
Display Summary Information for Storage Clouds	143
Displaying Detailed Information for Storage Clouds	144
Working With Default Gateways	153
Default Gateways Overview	153
Displaying Default Gateways	153
Default Gateways Summary	153
▼ Create a Default Gateway	154
▼ Delete a Default Gateway	156
Working With the Fabric Device	157
Fabric Interconnects Summary	157
▼ Scan for Fabric Interconnects	158
▼ Discover Fabric Interconnects With Oracle Fabric Manager	160
Displaying Fabric Interconnect Details	161
Fabric Interconnect Summary	162
Fabric Interconnect General Information	163
Ethernet Card Information	164
▼ Set the Allowed VLAN Ranges for an Ethernet Port	165
Fibre Channel Card Information	167
Fabric Interconnect User Information	168
▼ Configuring Local Users	169
Phone Home Information	169
▼ Display SNMP Properties	173
▼ Display SNMP Secure Users	174
▼ Display SNMP Trap Destinations	177
▼ Configure SNMP Trap Destinations	178
Fan Information	179
Power Supply Information	180

IMS Information	181
Active Directory Information	183
RADIUS Servers Information	188
RADIUS Users Information	192
▼ Unmanage a Fabric Interconnect	194
▼ Back Up and Restore a Fabric Interconnect Configuration	195
Performing Tech Support Functions on the Fabric Interconnect	197
Troubleshooting Functions	197
▼ Collect Tech Support Information	198
▼ Send an On-Demand Phone Home Message to Oracle Support	199
▼ Snooze Phone Home for a Window of Time	200
▼ Collect Fabric Interconnect Log Files	202
Working With I/O Templates	205
I/O Templates	205
Creating an I/O Template	206
I/O Template Creation Overview	207
▼ Create a New I/O Template	208
▼ Add a Network Cloud to the I/O Template	211
▼ Specify Advanced Properties for the Network Cloud	213
▼ Add a Storage Cloud to the Template	215
▼ Specify Advanced Properties for the Storage Cloud	217
▼ Add a PVI Cloud to the Template	219
▼ Add vNICs to the I/O Template	221
▼ Configure Advanced vNIC Properties	223
▼ Add PVI vNICs to the I/O Template	225
▼ Configure Advanced vNIC Properties for a PVI vNIC	228
▼ Add vHBAs to the I/O Template	229
▼ Configure Advanced vHBA Properties	232
▼ Save the I/O Template	233
Creating an HA vNIC Template	234
Creating an HA vHBA Overview	239
I/O Template Summary	245
▼ Rename an I/O Template	246
Editing an I/O Template	247
▼ Change I/O Template General Properties	247
▼ Add or Delete a Template's vNICs	248
▼ Add or Delete a Template's vHBA	248
▼ Apply an I/O Template to Servers	249

I/O Template Re-Apply Overview	251
▼ Re-Apply an I/O Template	252
Allowed VLANs Feature	254
▼ Configure Allowed VLANs in an I/O Template	254
▼ Delete an I/O Template	256
Working With I/O Profiles	259
I/O Profiles Overview	259
I/O Profiles Summary	260
▼ Create an I/O Profile	262
▼ Connect Specific I/O Profiles to Specific Fabric Devices	264
▼ Save an I/O Profile as an I/O Template	266
▼ Connecting an I/O Profile to a Server	266
Process for Merging Multiple I/O Profiles Into One I/O Profile	267
▼ Merge I/O Profiles	268
▼ Disconnect an I/O Profile From a Server	273
▼ Link an I/O Profile to an I/O Template	273
▼ Delete an I/O Profile	275
Displaying I/O Profile Details	276
I/O Profile General Properties	276
vNICs in an I/O Profile	277
vHBAs in an I/O Profile	277
Server Profiles in an I/O Profile	278
Working With Link Aggregation	281
Link Aggregation Groups	281
▼ Configure Link Aggregation Groups	281
Displaying LAG Properties	285
LAG Summary	285
LAG Details	285
LAG Ethernet Properties	286
Port Properties in the LAG	287
▼ Add More Ports to an Existing LAG	288
▼ Delete Ports From an Existing LAG	289
Working With Network QoS	291
Understanding Network QoS	291
Network QoS Profiles	291
Predefined Network QoS Profiles	292

▼ Display the Network QoS Summary	293
Network QoS Details	294
vNICs Using a Network QoS Profile	295
Understanding MAC-Based QoS	296
MAC-Based QoS	297
MAC-Based QoS Summary	297
MAC-Based QoS Profile Configuration	298
▼ Create a MAC-Based QoS Profile	299
MAC-Based QoS Profiles Deletion	301
Working With Physical Servers	303
Physical Servers	303
▼ Display Physical Servers	304
▼ Display the Physical Servers Details	307
Managing vNICs on a Physical Server	307
vNICs on a Physical Server	308
▼ Display Details About a vNIC	309
▼ Edit vNIC General Properties	309
▼ Display vNIC Ethernet Properties	310
▼ Edit vNIC Ethernet Properties	311
▼ Configure VLAN Ranges for vNICs on a Server	311
Managing vHBAs on a Physical Server	312
vHBAs on a Physical Server	313
▼ Display Details About a vHBA	314
▼ Display vHBA Fibre Channel Properties	315
▼ Edit vHBA Fibre Channel Properties	316
▼ Display vHBA Targets	316
▼ Scan for New Servers	317
▼ Apply an I/O Template to Servers	318
▼ Save a Server Configuration as an I/O Template	319
▼ Connect an I/O Profile to a Server	320
▼ Disconnect an I/O Profile From a Server	321
▼ Migrate Resources to Different Servers	321
▼ Create a Server Group From Selected Servers	323
▼ Remove Offline or Disconnected Servers	323
Working With Server Groups	325
Server Groups	325
▼ Display Server Groups	325

▼ Rename Server Groups	326
▼ Configure Server Groups	328
▼ Add Servers to Server Groups	329
▼ Remove Servers From Server Groups	330
▼ Delete a Server Group	331
Working With Boot Profiles	333
SAN Boot Profiles	333
SAN Boot Sequence	334
SAN Boot Setup Overview	336
Creating a SAN Boot Profile	336
SAN Boot Profile	336
▼ Create a SAN Boot Profile	337
▼ Create a Static SAN Boot Profile	338
▼ Create a Direct-Attach SAN Boot Profile	339
▼ Create a Logical Volume Manager SAN Boot Profile	340
Creating the Bootable I/O Template	342
Bootable I/O Template	342
▼ Create a Single-Path Bootable I/O Template	342
▼ Create a Dual-Path Bootable I/O Template	346
SAN Boot Profile Summary	349
SAN Boot Profile Details	350
Create an I/O Profile From the I/O Template	352
Connect the I/O Profile to a Physical Server	352
▼ Delete a SAN Boot Profile	352
iSCSI Boot Overview	353
iSCSI Boot Configuration Overview	354
iSCSI Boot Setup Overview	354
Creating an iSCSI Boot Profile	355
iSCSI Boot Profiles	355
▼ Create an iSCSI Boot Profile	356
▼ Create a Static iSCSI Boot Profile	358
▼ Create a Direct Attach iSCSI Boot Profile	360
▼ Create a Logical Volume Manager iSCSI Boot Profile	361
Creating the Bootable I/O Template	363
▼ Create a Single Path Bootable I/O Template	363
▼ Create a Dual Path Bootable I/O Template	365
iSCSI Boot Profiles Summary	369
iSCSI Boot Profile Details	370

Create an I/O Profile From the I/O Template	372
Connect the I/O Profile to a Physical Server	372
▼ Delete an iSCSI Boot Profile	373
Working With the Virtual Topology	375
Virtual Topology	375
Displaying the Topology Overview	375
Topology Overview	376
▼ Scroll Columns Vertically	378
Display the Physical Server Details	378
Display the Server Group Details	378
Display Physical Connectivity for the Topology Overview	378
Displaying the Server Cloud View	379
Server Cloud View	379
Server Group Topology	380
Individual Server Topology	381
Server Cloud Physical Connections Topology	382
Displaying the Fabric Device Cloud View	382
Fabric Device Cloud View	383
Fabric Device Details	383
Fabric Device Cloud Physical Connections Topology	384
Displaying the Server Fabric Device View	385
Server Fabric Device View	385
Detailed Connectivity Topology	386
Server Fabric Device Virtual Connections Topology	387
Displaying the Target Topology View	388
Target Topology View	388
vHBA-to-LUN Connections	389
Performance Meters	390
Physical and Virtual View Options	392
Working With Discovery Subnets	395
▼ Display the Discovery IP Subnet	395
▼ Display Discovery Subnet Details	396
▼ Add a Discovery Subnet	397
Working With Alarms	399
Alarm Summary Overview	399
▼ Display the Alarm Summary	399

▼ Display the Alarm History Summary	401
▼ Filter the Alarm History Summary	402
▼ Displaying Detailed Alarm Information	402
▼ View an Alarm's General Properties	403
▼ View an Alarm's Historical Information	404
▼ Clear Alarms From the Alarm Summary	406
▼ Clear Alarms From the Alarm History	406
▼ Filter the Alarm History Summary	407
Working With High Availability Oracle Fabric Manager	409
High Availability Oracle Fabric Manager Overview	409
Consideration for HA Oracle Fabric Manager and Plug-Ins	410
Fabric Performance Monitoring in HA Oracle Fabric Manager Environment	411
▼ Edit the PostgreSQL File to Include the HA Partners	412
Failover and Failback	413
Local Host and Remote Host	414
HA States	414
Configuring HA Oracle Fabric Manager Servers	416
▼ Configure the Active Server	416
▼ Configure the Passive Server	417
▼ Set an HA Partner's Mode	418
▼ Set a Sync Interval	420
▼ Force Sync Up	421
Performing Failover and Failback	421
▼ Perform a Failover	421
▼ Perform a Failback	422
▼ Install Plug-ins on the Passive Server	423
Displaying HA Oracle Fabric Manager Information	424
Displaying HA Partner Detailed Information	425
▼ Display HA Partner Detailed Information	426
HA Partner General Properties	427
HA Oracle Fabric Manager Statistics	427
▼ Clear Partner Statistics	429
Display Unsynced Commands for HA Oracle Fabric Manager	429
▼ Delete an HA Partner	430
Working With Live Monitoring	433
Live Monitoring	433
Displaying Host Server Throughput	435

Real-Time Grapher	435
▼ Display Real-Time Grapher Statistics	436
Historical and Calculated Statistics	437
vNIC Throughput	438
▼ Display vNIC Throughput	438
vHBA Throughput	440
▼ Display vHBA Throughput	440
Working With the Task Scheduler	443
Understanding Task Scheduler	443
Task Scheduler	443
Oracle Fabric Manager Backups Considerations	444
Backup and Restore of Fabric Device Through Oracle Fabric Manager	445
Considerations for Fabric Device Backups	446
Fabric Device Backup With HA Fabric Devices	447
Scheduled Backups	447
On-Demand Backups	448
Displaying Oracle Fabric Manager Schedules	449
Oracle Fabric Manager Backup Schedule Summary	449
▼ Edit Oracle Fabric Manager Backup Schedules	450
▼ Rename Oracle Fabric Manager Backup Schedules	451
▼ Create Oracle Fabric Manager Backup Schedules	452
▼ Run an On-Demand Oracle Fabric Manager Backup	454
▼ Enable or Disable Oracle Fabric Manager Backup Schedules	455
▼ Create Fabric Device Backup Schedules	456
Displaying Fabric Device Backup Schedules	458
Fabric Device Backups	459
▼ Edit Fabric Device Backup Schedules	460
▼ Rename Fabric Device Backup Schedules	461
▼ Run an On-Demand Fabric Device Backup	462
▼ Enable or Disable Fabric Device Backup Schedules	463
Working With LUN Masks	465
LUN Masks Overview	465
▼ Create LUN Masks	465
Displaying LUN Masks	467
LUN Mask Profile Summary	467
LUN Mask Profile Details	468
▼ Add LUNs to an Existing LUN Mask	469

▼ Remove LUNs From the LUN Mask	471
Working With SAN QoS	473
SAN QoS Profile Overview	473
Default SAN QoS Profiles	473
Displaying SAN QoS Profiles	474
▼ Display the SAN QoS Summary	474
SAN QoS Details	475
vHBAs Associated With a SAN QoS Profile	476
Working With Domains	477
Resource Domains	477
Domain Resources and Unmanaging a Fabric Device	477
Default Domain	478
Default Domain Users	478
Non-Default Domains	479
Non-Default Domain Users	480
Users With Administrator Role	480
Users With Network Role	480
Users With Storage Role	481
Displaying Domain Information	481
▼ Display Domain Information	482
▼ Display Detailed Information for Domains	482
Servers in a Domain	483
Fabric Devices in a Domain	484
I/O Cards in a Domain	485
Network or PVI Clouds in a Domain	486
Storage Clouds in a Domain	487
Creating a Domain	488
▼ Create a Domain	488
▼ Add Fabric Devices to a Domain	489
▼ Add Physical Servers to a Domain	491
▼ Add I/O Modules to a Domain	492
▼ Add Network Connectivity to a Domain	494
▼ Add Storage Connectivity to a Domain	495
Working With User Roles	497
Oracle Identity Management System	497
Using the Internal IMS	497

Local Users	498
Roles	498
Oracle Fabric Manager User	499
▼ Define a User in Oracle Fabric Manager	500
▼ Delete a User From Oracle Fabric Manager	502
Define a Local User on the Fabric Device	503
▼ Rename a Oracle Fabric Manager User	504
User Account Roles	506
▼ Change the Role Assigned to a User Account	506
Mapping Users in External Groups Into Oracle Fabric Manager	508
Group Mapping Overview	508
▼ Map an External IMS Group to a Oracle Fabric Manager Role	509
▼ Map a Group to a Domain	511
Configuring Oracle Fabric Manager to Authenticate Users Against AD	513
▼ Set Up the AD Server	513
▼ Set Up the Group in Oracle Fabric Manager and Assign Roles	515
Index	519

Using This Documentation

- **Overview** – Provides overview, installation, configuration and administration information for Oracle Fabric Manager 4.3.0 software.
- **Audience** – Technicians, system administrators, and authorized service providers.
- **Required knowledge** – Advanced experience troubleshooting and replacing hardware.

Product Documentation Library

Late-breaking information and known issues for this product are included in the documentation library at http://docs.oracle.com/cd/E38500_01.

Access to Oracle Support

Oracle customers have access to electronic support through My Oracle Support. For information, visit <http://www.oracle.com/pls/topic/lookup?ctx=acc&id=info> or visit <http://www.oracle.com/pls/topic/lookup?ctx=acc&id=trs> if you are hearing impaired.

Documentation Accessibility

For information about Oracle's commitment to accessibility, visit the Oracle Accessibility Program web site at <http://www.oracle.com/pls/topic/lookup?ctx=acc&id=docacc>.

Feedback

Provide feedback about this documentation at <http://www.oracle.com/goto/docfeedback>.

Understanding Oracle Fabric Manager

This chapter contains the following topics:

- [“Oracle Fabric Manager Overview” on page 21](#)
- [“Oracle Fabric Manager as a Stand-Alone Application” on page 22](#)

Oracle Fabric Manager Overview

Oracle Fabric Manager is a multi-device management system created by Oracle to inventory and manage Oracle Fabric Devices (one or more Oracle Fabric Interconnects and/or Oracle SDN Controllers) and Oracle virtual I/O. Throughout this document, the term “Fabric Device” will be used to refer to either Oracle's Fabric Interconnect or Oracle's SDN Controller. Oracle Fabric Manager product is a server-client model that supports configuring and managing virtual resources at Fabric Device, module, and server or virtual machine level.

Oracle Fabric Manager is a browser-based management system that runs on a remote server. The remote Oracle Fabric Manager server translates configuration and management tasks from the Oracle Fabric Manager interface, and relays that information to Fabric Devices that are managed by Oracle Fabric Manager.

Oracle Fabric Manager can run in one of the following ways:

- Stand-alone application
- As a plug-in to other GUI management systems

Oracle Fabric Manager configuration and management capabilities are the same regardless of whether Oracle Fabric Manager is running in stand-alone mode or as a plug-in to other GUI management systems.

Support for the Fabric Interconnect F1-15

Oracle Fabric Manager supports management of the Oracle Fabric Interconnect F1-15, which is a 4 rack-unit virtual I/O platform supporting connectivity to network and storage resources for host servers through 24 InfiniBand ports at up to 40 Gbps per server connection. The Fabric

Interconnect F1-15 has 15 slots which support a mix of network and storage modules that is appropriate for your network.

Support for the Fabric Interconnect F1-4

Oracle Fabric Manager supports management of Oracle's Fabric Interconnect F1-4, which is a 2 rack-unit virtual I/O platform supporting connectivity to network and storage resources for host servers through 24 InfiniBand ports at up to 40 Gbps per server connection. The Fabric Interconnect F1-4 supports has 4 slots which support a mix of network and storage modules that is appropriate for your network.

Support for the Oracle SDN Controller

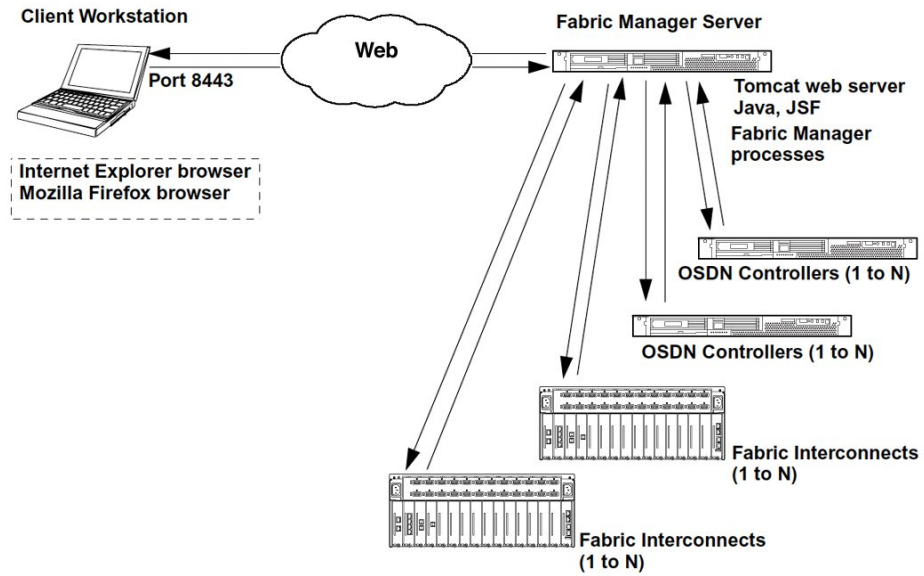
Oracle Fabric Manager supports management of Oracle's SDN Controller, which is a virtual I/O platform supporting connectivity to network resources for host servers. The Oracle SDN Controller is a software solution that resides on off-the-shelf server hardware and allows configuration and management of network I/O connections.

Oracle Fabric Manager as a Stand-Alone Application

In stand-alone mode, Oracle Fabric Manager runs on a remote server that controls Fabric Devices. The Oracle Fabric Manager system uses additional services to provide manageability:

- Tomcat web server
- Java, JSF
- Oracle proprietary processes

The following figure shows a sample topology in which Oracle Fabric Manager is running in stand-alone mode.



After Oracle Fabric Manager and its components are installed on the physical server that hosts the Oracle Fabric Manager software, any client workstation can use Oracle Fabric Manager by pointing one of the supported browsers at the Oracle Fabric Manager server. Configuration and management requests are sent through the Oracle Fabric Manager interface to the Oracle Fabric Manager server, which handles the back-end processing and ensures that the client's task is communicated to the Fabric Device and written to the Fabric Device's configuration. When the Fabric Device has acknowledged the client's task, Oracle Fabric Manager relays the change back to the Oracle Fabric Manager interface where the updated configuration is displayed by the client workstation.

Installing or Updating Oracle Fabric Manager

This chapter documents the following topics:

- [“Minimum Requirements” on page 25](#)
- [“Oracle Fabric Manager Server Requirements” on page 26](#)
- [“Client Requirements” on page 26](#)
- [“Preparing for Installation” on page 27](#)
- [“Installing Oracle Fabric Manager” on page 27](#)
- [“Validating and Troubleshooting the Installation” on page 36](#)
- [“Upgrading the Oracle Fabric Manager Software” on page 40](#)
- [“Removing the Oracle Fabric Manager Software” on page 43](#)
- [“Configuring a Certificate for Oracle Fabric Manager” on page 46](#)
- [“Starting Oracle Fabric Manager and Logging In” on page 55](#)

Minimum Requirements

Oracle Fabric Manager runs on a separate Oracle Solaris, Linux, or Windows host server, or in a virtual machine. Oracle Fabric Manager does not actually run on the Oracle Fabric Device (Oracle Fabric Interconnect or Oracle SDN Controller) itself. The following platforms support Oracle Fabric Manager:

- Windows:
 - Microsoft Windows Server 2003 (various), Windows Server 2008 (various), and Windows Server 2012 R2. To install Oracle Fabric Manager on this platform, use the install executable named `install-xms-win32-version_OFM.exe` for 32-bit servers, or use `install-xms-win64-version_OFM.exe` for 64-bit servers. Double-click the executable to begin extracting the file and running the install wizard. For information about installing Oracle Fabric Manager on a Windows Server, see [“Install Oracle Fabric Manager on a Windows Server” on page 33](#) after reading [“Oracle Fabric Manager Server Requirements” on page 26](#).
- Linux:
 - Oracle Enterprise Linux 6 Update 3 and higher.
- Oracle Solaris:

- Oracle Solaris 10 Update 11 or Oracle Solaris 11 Update 1 on x86 architecture. To install Oracle Fabric Manager on this platform, use the install package named `xsgo-xms-version-solaris-x86.pkg`.

To run Oracle Fabric Manager, host servers have been tested to operate when configured to meet the following server and client requirements.

Oracle Fabric Manager Server Requirements

For interoperability with Oracle Fabric Manager, the physical server that hosts Oracle Fabric Manager has been tested to operate when configured with the following:

- Server OS
 - Windows Server 2003 SP2 and R2, Windows Server 2003 32-bit and 64-bit architectures, Windows Server 2008 32-bit and 64-bit architectures, Windows Server 2008, Windows Server 2008 R2, and Windows Server 2012 R2.
 - Oracle Enterprise Linux 6 Update 3 and higher
 - VMware ESX Server 5.5 supported. Oracle Fabric Manager does not actually run on the ESX Server. It runs within a VM, but there are some requirements for running Oracle Fabric Manager server has some requirements for the VMware server.
 - Oracle Solaris 10 Update 11 and Oracle Solaris 11 Update 1 (x86 architecture only)
- Memory
 - 2 GB minimum, 4 GB or more recommended. (Four GB or more is required if Oracle Fabric Manager will be running while integrated into VMware VirtualCenter Manager when both applications co-located on the same physical server.) The same requirements exist if Oracle Fabric Manager will be integrated into vSphere 5.
 - vSphere Client—dual core and 4 GB or more of memory (recommended).
- CPU—any server-based, current X86 architecture CPU.
- Java 1.7 (JRE), which is the only version of Java that supports Oracle Fabric Manager. This software is not included with the Oracle Fabric Manager software.

Client Requirements

For interoperability with Oracle Fabric Manager, the clients that will use Oracle Fabric Manager have been tested to operate with the following:

- Browsers
 - Mozilla® Firefox 6.0 and higher. This release of Oracle Fabric Manager software was tested with FireFox 17.

- Microsoft® Internet Explorer 7.0 and later, with all cumulative security updates. Any version of Internet Explorer less than 7.0 is not supported.
- Display—1280 x 1024 resolution, 16-bit Medium color mode.
- JavaScript and cookies enabled.

▼ Preparing for Installation

Oracle Fabric Manager relies on certain ports for communication and connection. Perform the following steps to open the relevant ports.

1. Ensure that ports 80 and 443 (HTTPs) are open.

Oracle Fabric Manager requires certain ports. Make sure that ports 80 and 443 (HTTPs) are open for both incoming and outgoing directions for Oracle Fabric Manager to communicate between the Oracle Fabric Manager Server and the Fabric Device. By default, port 443 is used.

2. Ensure that ports 8880 or 8443 (HTTPs) are open.

When Oracle Fabric Manager is installed, IP Forwarding tables need to be updated in order to allow connection to the Oracle Fabric Manager server. As a result, it is a best practice that before you install the Oracle Fabric Manager software, you should make sure that ports 8880 or 8433 in any firewalls are open.

As an alternative, you can leave those port closed, but open an RDP session and use the local browser for added security.

Installing Oracle Fabric Manager

This section contains the following topics:

- [“Install Oracle Fabric Manager on an Oracle Solaris Server” on page 27](#)
- [“Linux Server Installation Considerations” on page 31](#)
- [“Install Oracle Fabric Manager on a Linux Server \(RPM File\)” on page 31](#)
- [“Windows Server Installation Considerations” on page 32](#)
- [“Install Oracle Fabric Manager on a Windows Server” on page 33](#)

▼ Install Oracle Fabric Manager on an Oracle Solaris Server

Oracle Fabric Manager is supported on x86 architecture servers running Oracle Solaris 10 Update 11 and Oracle Solaris 11 Update 1. At this time, SPARC servers are not supported.

Oracle Fabric Manager is installed through a PKG file which can be downloaded to the Oracle Fabric Manager server. To install Oracle Fabric Manager on a server running one of the supported Oracle Solaris OSes, follow this procedure:

1. **Log in to the server where Oracle Fabric Manager will be installed.**
2. **Copy the PKG file to a directory on the server. Make sure that you have at least write and execute privileges on the directory where Oracle Fabric Manager will be installed.**
3. **Run the pkgadd command and specify the Oracle Fabric Manager PKG filename.**

Note - During this part of the installation process, you will be prompted for input. Bold text indicates required input.

For example:

```
pkgadd -d xsigo-xms-version-solaris-x86.pkg
```

The following packages are available:

```
1 xms xsigo-xms-version-solaris-x86.pkg
```

(x86) Oracle Fabric Manager

Select package(s) you wish to process (or 'all' to process all packages). (default: all) [?,??,q]: all

Processing package instance <xms> from </xsigo-xms-private_build-solaris-x86.pkg>

```
xsigo-xms(x86) xms-version
```

Oracle

The selected base directory </opt/xsigo/xms> must exist before installation is attempted

Do you want this directory created now [y,n,?,q] y

Using </opt/xsigo/xms> as the package base directory.

```
## Processing package information.
```

```
## Processing system information.
```

```
## Verifying disk space requirements.
```

```
## Checking for conflicts with packages already installed.
```

```
## Checking for setuid/setgid programs.
```

This package contains scripts which will be executed with super-user permission during the process of installing this package.

Do you want to continue with the installation of <xms> [y,n,?] y

```
Installing xsigo-xms as <xms>
```

```
## Executing preinstall script.
## Installing part 1 of 1.
/opt/xsigo/xms/.openlogic.yaml
/opt/xsigo/xms/LICENSE
/opt/xsigo/xms/NOTICE
/opt/xsigo/xms/RELEASE-NOTES
/opt/xsigo/xms/RUNNING.txt
/opt/xsigo/xms/bin/bootstrap.jar
/opt/xsigo/xms/bin/catalina-tasks.xml
/opt/xsigo/xms/bin/catalina.bat
/opt/xsigo/xms/bin/catalina.sh
/opt/xsigo/xms/bin/commons-daemon-native.tar.gz
/opt/xsigo/xms/bin/commons-daemon.jar
/opt/xsigo/xms/bin/cpappend.bat
/opt/xsigo/xms/bin/digest.bat
/opt/xsigo/xms/bin/digest.sh
/opt/xsigo/xms/bin/directortest
/opt/xsigo/xms/bin/directortestxmscli.bat
/opt/xsigo/xms/bin/msvcr71.dll
/opt/xsigo/xms/bin/service.bat
/opt/xsigo/xms/bin/setclasspath.bat
/opt/xsigo/xms/bin/setclasspath.sh
/opt/xsigo/xms/bin/setenv.bat
/opt/xsigo/xms/bin/setenv.sh
/opt/xsigo/xms/bin/shutdown.bat
/opt/xsigo/xms/bin/shutdown.sh
/opt/xsigo/xms/bin/sleep.vbs
/opt/xsigo/xms/bin/start-xms.vbs
/opt/xsigo/xms/bin/startup.bat
/opt/xsigo/xms/bin/startup.sh
/opt/xsigo/xms/bin/stop-xms.vbs
```

```
/opt/xsigo/xms/bin/tcnative-1.dll
/opt/xsigo/xms/bin/tomcat-juli.jar
/opt/xsigo/xms/bin/tomcat-native.tar.gz
/opt/xsigo/xms/bin/tomcat6.exe
/opt/xsigo/xms/bin/tomcat6w.exe
/opt/xsigo/xms/bin/tool-wrapper.bat
/opt/xsigo/xms/bin/tool-wrapper.sh
<<display of installed files continues...>>
/opt/xsigo/xms/xmstools/xmscli/sandbox/bin/pwd
/opt/xsigo/xms/xmstools/xmscli/sandbox/bin/rm
/opt/xsigo/xms/xmstools/xmscli/sandbox/bin/sedit
[ verifying class <none> ]
## Executing postinstall script.
Installation of <xms> was successful.
#
```

4. **When the software is installed, issue the `svcs` command to check the status of the Oracle Fabric Manager services. For example:**

```
# svcs xms
STATE STIME FMRI
online 16:18:38 svc:/site/xms:default
#
```

Check the beginning of the line, which is where the services status is listed. Valid statuses are:

- **Online, which indicates that the XMS service is online and operating correctly.**
- **Offline, which indicates that the XMS service is not online, or is not operating correctly.**
- *** (an asterisk), which indicates that service with the asterisk is running, but some part of the service or its status is in the process of changing. This status should be temporary because when the service's change is complete, the asterisk disappears.**

Note - On Oracle Solaris servers, Oracle Fabric Manager services automatically attempt to recover. As a result, if an error occurs, the Oracle Fabric Manager service will restart on the Oracle Fabric Manager server. The only way to permanently stop or halt the Oracle Fabric Manager service is to explicitly stop it by issuing the `svcadm disable xms` command. If you explicitly stop the service, you must explicitly re-enable it (when needed) by issuing the `svcadm enable xms` or `svcadm restart xms` command.

Linux Server Installation Considerations

Oracle Fabric Manager is supported on Linux servers running Oracle Linux 6 Update 3 and higher. Installing the Oracle Fabric Manager software uses the standard Linux commands, similar to installing or updating the Linux host OS software.

Oracle Fabric Manager requires certain ports. Make sure that ports 80 and 443 (HTTPs) are open for both incoming and outgoing directions for Oracle Fabric Manager to communicate between the Oracle Fabric Manager Server and the Fabric Device. By default, port 443 is used.

Note - The installation procedure documented in this section assumes a fresh install (no Oracle Fabric Manager software exists on the Oracle Fabric Manager server). If you are upgrading from an older version of Oracle Fabric Manager (not doing a fresh install), you will need to clear any completed jobs from the Jobs summary in the older version of Oracle Fabric Manager before running the upgrade. For information about clearing completed jobs from the Jobs Summary, see [“Clear All Jobs in the Jobs Summary” on page 96](#).

▼ Install Oracle Fabric Manager on a Linux Server (RPM File)

You can install Oracle Fabric Manager directly from an RPM file. A README file is included with the RPM file. Oracle recommends that you read the README file.

To install Oracle Fabric Manager directly from the RPM file, follow this procedure:

1. **Log in to the server where Oracle Fabric Manager will be installed.**
2. **Download the RPM file to a directory on the server.**
Make sure that you have at least write and execute privileges on the directory where Oracle Fabric Manager will be installed.
3. **Extract the file. While extracting, the `/xms_install` directory is created by default. For example:**

4. **Change directory to the `/xms_install` directory. For example:**

```
root@terminus tmp]# cd ./xms_install/
```

5. **From the `/xms_install` directory, install the `xsigo-xms-version_OFM.noarch RPM` file. For example:**

```
root@terminus tmp/xms_install]# rpm -ivh xsigo-xms-version_OFM.noarch.rpm
```

```
Preparing...      ##### [100%]
```

```
1:xsigo-xms      ##### [100%]
```

When the Oracle Fabric Manager RPM is installed, the Oracle Fabric Manager services are automatically brought up and should be running. You do not need to restart the Linux server on which you have just installed Oracle Fabric Manager.

6. **As an option, you can check the status of the Oracle Fabric Manager services:**

```
root@terminus tmp# service xms status
```

```
Checking Oracle Fabric Manager: Running
```

Note - As an option, some service log files exist that you can check for messages in case the installation does not appear to have completed correctly. For information, see [“Validating and Troubleshooting the Installation”](#) on page 36.

7. **Proceed to [“Starting Oracle Fabric Manager and Logging In”](#) on page 55.**

Windows Server Installation Considerations

Oracle Fabric Manager is supported on various versions of Microsoft Windows Server 2003 and Windows Server 2008 hosts as well as on Microsoft Windows Server 2012. For a complete list of Microsoft OSes supported for the Oracle Fabric Manager Server, see [“Oracle Fabric Manager Server Requirements”](#) on page 26.

Note - If you will be using LDAP/AD for authentication in a Windows domain, install the Oracle Fabric Manager software on a Windows server that is a member of that domain.

Oracle Fabric Manager is installed through a Windows-based installer called IZpack, which is similar to an InstallShield installer. Oracle Fabric Manager is provided as an executable file which starts the installer when you double-click the EXE file.

When the installer runs for the first time, you will be prompted to create a default installation directory. You must accept the prompt (click Yes) to successfully install Oracle Fabric Manager.

For subsequent installations, you will not see this prompt as long as the Oracle Fabric Manager install directory is present on the Windows Oracle Fabric Manager Server.

You also have the option of installing Oracle Fabric Manager in a non-default install directory. To install to a non-default directory, you can click **Browse** at the appropriate step in the installation procedure, then select the directory where you want Oracle Fabric Manager installed.

Note - The installation procedure documented in this section assumes a fresh install (no Oracle Fabric Manager software exists on the Oracle Fabric Manager server). If you are upgrading from an older version of Oracle Fabric Manager (not doing a fresh install), you will need to clear any completed jobs from the Jobs summary in the older version of Oracle Fabric Manager before running the upgrade. For information about clearing completed jobs from the Jobs Summary, see [“Clear All Jobs in the Jobs Summary” on page 96](#).

Note - It is possible to install Oracle Fabric Manager to a non-default device (for example, a USB drive). However, some functions of Oracle Fabric Manager expect the default device on the server. For example, the xms-backups, techsupport, and director-backups directories are installed to the local device. Be aware that if you install Oracle Fabric Manager to a non-default device (for example, a network drive, a USB drive, and so on) functions that rely on the xms-backup, techsupport, and director-backups directories might not work predictably.

Note - Currently, the Oracle Fabric Manager installer does not check the OS version installed on the Windows server.

Oracle Fabric Manager requires certain ports. Make sure that ports 80 or 443 (HTTPs) are open for both incoming and outgoing directions for communication between the Oracle Fabric Manager Server and the Fabric Device. By default, port 443 is used.

▼ Install Oracle Fabric Manager on a Windows Server

To install Oracle Fabric Manager, follow this procedure:

1. **Log in to the Windows server that will be running Oracle Fabric Manager.**
2. **Get the Oracle Fabric Manager installer for Windows (for example, `install-xms-win32-version_OFM.exe` or `install-xms-win64-version_OFM.exe`) onto the Windows server that will be running Oracle Fabric Manager.**

The installer can be downloaded from the Oracle Customer Support site, or the Oracle Customer Support secure FTP site. If you need assistance with getting Oracle Fabric Manager Installer, contact Oracle Customer Support as documented in [“Access to Oracle Support” on page 19](#).

3. **Double-click the ZIP file to extract the Oracle Fabric Manager executable file and its related content.**
4. **Double-click the Oracle Fabric Manager executable to run the installer.**
The Oracle Fabric Manager splash screen is displayed momentarily, and then the welcome screen is displayed.



5. **Read the Welcome screen and note the support contact information and Oracle home page, in case you require additional correspondence with Oracle, then click Next to display the License Agreement.**
6. **When you accept the license agreement and click Next, the Installation Path dialog is displayed.**
7. **On the Installation Path dialog, you can either accept the default installation directory, or select a non-default installation directory:**
 - **If you accept the installation directory, proceed to [Step 9](#).**
 - **If you select a non-default installation directory, click Browse to display the Select Path dialog, and proceed to [Step 8](#).**
8. **Select the directory in which you want to install Oracle Fabric Manager, click Save, then proceed to [Step 11](#).**
9. **Click Next to display the confirmation dialog that alerts you that the installation directory will be created.**
This confirmation dialog occurs on a first-time install. If the installation directory already exists on the Windows Oracle Fabric Manager server, this dialog is not displayed.
10. **Click OK to create the installation directory.**

The installer resumes guiding the installation of Oracle Fabric Manager software.

11. **When the Oracle Fabric Manager package installation is complete, click Next to display the Setup Shortcuts dialog.**
12. **Use the checkboxes to determine how Oracle Fabric Manager shortcuts are installed:**
 - **Use the Create shortcuts in the Start-Menu checkbox to allow the installer to put a Oracle Fabric Manager shortcut on the Windows Start menu. This option is selected by default, so a shortcut to Oracle Fabric Manager will be installed on the Windows Start menu. This checkbox is a toggle, so clicking alternates between a checkmark (allow the shortcut) and no checkmark (do not allow a shortcut). If you are allowing Oracle Fabric Manager shortcuts on the Start menu, you also have the option of installing shortcuts at additional locations (for example, on the desktop).**
 - **Use the Create additional shortcuts on the desktop checkbox to allow the installer to put a Oracle Fabric Manager shortcut on the Windows desktop. This option is selected by default, so a shortcut to Oracle Fabric Manager will be installed on the Windows desktop. This checkbox is a toggle, so clicking alternates between a checkmark (allow the shortcut) and no checkmark (do not allow a shortcut).**
13. **If you are allowing shortcuts on the Start menu, click an item in the Select a Program Groups for the Shortcuts list to determine where the shortcut will be displayed.**
14. **If you are allowing shortcuts on the desktop, in the text box below the Select a Program Groups for the Shortcuts list, you can enter a name for the Oracle Fabric Manager shortcut or select the default location.**
15. **Use the create shortcut for: controls to determine which users on the Windows server will readily see the installed Oracle Fabric Manager shortcuts on either the Start menu or the desktop:**
 - **Click Current User to allow Oracle Fabric Manager shortcuts to be created for the user currently logged in to the Windows server.**
 - **Click All Users to allow Oracle Fabric Manager shortcuts to be created for anyone who can log in to the Windows server.**
16. **As an option, you can click the Default button if you want to reset the name and path of the Oracle Fabric Manager program to its default value. Clicking Default resets the program name in the text-entry field next to the Default button. No other controls on the dialog are reset.**

17. **When the shortcut installation options have been setup, click Next to display the Installation Finished dialog.**
18. **Click Done to close the Oracle Fabric Manager installer.**

At this point, Oracle Fabric Manager is installed on the Oracle Fabric Manager Windows Server. You can open a supported browser and log in to Oracle Fabric Manager as documented in [“Starting Oracle Fabric Manager and Logging In” on page 55](#).

Validating and Troubleshooting the Installation

After Oracle Fabric Manager is installed, there are a few ways to validate that Oracle Fabric Manager is up and running, or discover some error messages if you need to begin troubleshooting. The following sections document how to check Oracle Fabric Manager status and gather some troubleshooting information when Oracle Fabric Manager is operating in run-time.

When Oracle Fabric Manager is installed and running, you can run some easy-to-use commands to verify that all required services are operating. The commands are different based on the type of server where Oracle Fabric Manager is installed. For all server types, you can perform the following tasks:

- [“Check the Status of Oracle Fabric Manager Services” on page 36](#)
- [“Stop Oracle Fabric Manager Services” on page 37](#)
- [“Start Oracle Fabric Manager Services” on page 39](#)
- [“Service Logs” on page 40](#)

▼ Check the Status of Oracle Fabric Manager Services

- **Determine the status of Oracle Fabric Manager services on an Oracle Solaris or Linux server.**
 - **On Oracle Solaris:**
 - a. **Use the `svcs xms` command to display the current status for the services that support Oracle Fabric Manager. Issue this command on the Linux server where you installed Oracle Fabric Manager. For example**

```
[root@terminus xms_install]# svcs xms  
STATE      STIME     FMRI
```

```
online 16:18:38 svc:/site/xms:default
```

```
#
```

- b. In this example, all Oracle Fabric Manager services are operating. However, in the following example, Oracle Fabric Manager is not running.

```
[root@terminus xms_install]# svcs xms
```

```
STATE STIME FMRI
```

```
offline 16:18:38 svc:/site/xms:default
```

```
#
```

- c. In this situation, you will need to stop all Oracle Fabric Manager services and start them again to get Oracle Fabric Manager running.

■ On Linux:

- a. Use the service xms status command to display the current status for the services that support Oracle Fabric Manager. Issue this command on the Linux server where you installed Oracle Fabric Manager. For example:

```
[root@terminus xms_install]# service xms status
```

```
Checking Oracle Fabric Manager: Running
```

- b. In this example, all Oracle Fabric Manager services are operating. However, in the following example, Tomcat is not running.

```
[root@terminus xms_install]# service xms status
```

```
Checking Oracle Fabric Manager: Stopped
```

- c. In this situation, you will need to stop all Oracle Fabric Manager services and start them again to get Tomcat running.

▼ Stop Oracle Fabric Manager Services

- If a service is not running, stop, then start, all services to allow them to gracefully terminate and start up in the correct order.

■ On Oracle Solaris:

- a. You can use the `svcadm disable xms` command to stop services gracefully and in the proper order. This command stops all services that support Oracle Fabric Manager. You cannot use this command to stop individual services.
- b. All services are stopped when the Oracle Fabric Manager status shows OFFLINE.
- c. Stopping services might take a short time.
- d. When you issue `svcadm disable xms`, wait for all services to stop before starting them up again.
- e. Issue the command on the Oracle Solaris server where you installed Oracle Fabric Manager. After the services stop, issue the `svcs xms` to query Oracle Fabric Manager's status:

```
[root@terminus xms_install]# svcs xms
STATE STIME FMRI
offline 16:18:38 svc:/site/xms:default
#
```

■ On Linux:

- a. You can use the `service xms stop` command to stop all services gracefully and in the proper order. This command stops all services that support Oracle Fabric Manager. You cannot use this command to stop individual services.
 - All services are stopped when the status for each service shows OK
 - Stopping services might take a short time
 - When you issue `service xms stop`, wait for all services to stop before starting them up again.
- b. Issue the command on the Linux server where you installed Oracle Fabric Manager. The following example is for a Linux server:

```
[root@terminus xms_install]# service xms stop
Stopping Oracle Fabric Manager:          [ OK ]
```

▼ Start Oracle Fabric Manager Services

- When all services are stopped, you can restart all Oracle Fabric Manager services in the proper order with a single command.
 - On Oracle Solaris:
 - a. Use the `svcadm enable xms` command to simultaneously start all Oracle Fabric Manager services. You cannot start individual services because of a dependency that some services come up before others.

- b. Issue the command on the Oracle Solaris server where you installed Oracle Fabric Manager. After the services start, issue the `svcs xms` to query Oracle Fabric Manager's status:

```
[root@terminus xms_install]# svcs xms

STATE STIME FMRI

online 16:18:38 svc:/site/xms:default

#
```

- c. This example shows Oracle Fabric Manager shows Oracle Fabric Manager successfully started and running (online). However, as all services are coming up you might see the STATE field with an asterisk (*) which indicates that bring-up of services is progressing, but not yet complete. After a short time, all services are brought online and started. At this point, the STATE field displays online.

- On Linux:

- a. Use the `service xms start` command to simultaneously start all Oracle Fabric Manager services. You cannot start individual services because of a dependency that some services come up before others. Issue the command on the Linux server where you installed Oracle Fabric Manager. For example:

```
[root@terminus xms_install]# service xms start

Starting Oracle Fabric Manager:                [ OK ]
```

- b. All services are started when the status for each service shows OK.

Service Logs

This section provides basic information about logs that receive events if errors occur when Oracle Fabric Manager is online, as well as where to look on the Linux, Windows, or Oracle Solaris Oracle Fabric Manager Server for the logs. The information contained in these logs can be useful for Oracle Technical Support.

Oracle Fabric Manager runs as the sum of multiple different services on the server. Most services have a log file associated with them. Information about runtime Oracle Fabric Manager errors are written to different logs on the Oracle Fabric Manager Server.

The Oracle Solaris logs are located in `/opt/xsigo/xms/logs`, and you can display the contents the logs' contents by using any standard UNIX or LINUX editor—for example, `vi`, `cat`, `emacs`, and so on.

Window Server	Linux Server	Oracle Solaris Server
director-<name>.log	director-<name>.log	director-<name>.log
commons-daemon.log	catalina.out	catalina.out
xmscli	catalina.pid	tomcat.log
xms-stderr.log	tomcat.log	xmsaudit.log.1
xms-stdout.log	xmsaudit.log.1	xms-ha.log.1
xms-ha.log.1	xms-ha.log.1	xmsjobs.log.1
xmsjobs.log.1	xmsjobs.log.1	xms.log.1
xms.log.1	xms.log.1	xms-schedule.log.1
xms-schedule.log.1	xms-schedule.log.1	
xmsaudit.log.1		

Upgrading the Oracle Fabric Manager Software

- [“Upgrade Oracle Fabric Manager on an Oracle Solaris Server” on page 40](#)
- [“Upgrade Oracle Fabric Manager on a Windows or Linux Server” on page 42](#)

▼ Upgrade Oracle Fabric Manager on an Oracle Solaris Server

Oracle Fabric Manager does not support an update operation on Solaris servers. You must remove the previous version of Oracle Fabric Manager from your server and install the new version.

Because Oracle Fabric Manager store some data (ie the database, backups, logs, etc) in the install path (`/opt/xsigo/xms`), backup any files you want to keep before removing the previous version of Oracle Fabric Manager.

1. **Perform a Oracle Fabric Manager backup from the user interface.**
 - a. See [“Back Up the Oracle Fabric Manager Configuration” on page 79](#).
 - b. Copy the XML file from the `/xms-backups` directory to somewhere else—for example, to a network drive.
2. **Shut down Oracle Fabric Manager, then confirm that the Oracle Fabric Manager services are not running.**

```
# svcadm disable xms
```

```
# svcs xms
```

3. **Make a database backup and copy it outside of your `/xsigo` directory, for example:**

```
# cp -r /opt/xsigo/xms/db /tmp/db
```

If you want to copy other files such as `xms-backups`, `director-backups`, `techsupport`, `logs`, do so at this time. The next step will uninstall your previous version of Oracle Fabric Manager, including those files.

4. **Remove the previous version of Oracle Fabric Manager, then confirm that the Oracle Fabric Manager installation was removed.**

```
# pkgrm xms
```

```
# svcs xms
```

5. **Install the new version of Oracle Fabric Manager.**

```
# pkgadd -d xsigo-xms-version-solaris-x86.pkg
```

Once installed, Oracle Fabric Manager will automatically start up. To restore the files you backed up earlier, you need to shut it down.

6. **Shut down Oracle Fabric Manager.**

```
# svcadm disable xms
```

7. **Restore the files you backed up in [Step 3](#), for example:**

```
# rm -r /opt/xsigo/xms/db
```

```
# cp -r /tmp/db /opt/xsigo/xms/db  
# chown -R xsigo:xsigo /opt/xsigo/xms/db
```

8. **Start Oracle Fabric Manager, then confirm the Oracle Fabric Manager services are running.**

```
# svcadm enable xms  
# svcs xms
```

▼ Upgrade Oracle Fabric Manager on a Windows or Linux Server

Upgrading the Oracle Fabric Manager software to version 4.3.0 also requires an update to JRE 1.7.

1. **Perform a Oracle Fabric Manager backup from the user interface.**
See “[Back Up the Oracle Fabric Manager Configuration](#)” on page 79.
2. **Copy the XML file from the `xms-backups` directory to somewhere else—for example, to a network drive.**
3. **Shut down Oracle Fabric Manager.**
See “[Stop Oracle Fabric Manager Services](#)” on page 37.
4. **Upgrade to JRE 1.7.**
 - a. **Perform a query to determine what programs are using JRE, for example:**

```
rpm -q --whatrequires jre  
xsigo-xms-4.2.0_OFM-1.noarch
```

If packages other than Oracle Fabric Manager (and associated plugins) use JRE, consider what impact a JRE upgrade will have on those programs before you continue.

- b. **Download the JRE 1.7 package for your OS from here:**
<http://www.oracle.com/technetwork/java/javase/downloads/java-se-jre-7-download-432155.html>
Note that the JRE package might be inside another package, such as a JDK package.
- c. **Uninstall the previous JRE package, for example:**

```
rpm -e jre-1.6.0_45-fcs.i586 --nodeps
```

d. Install the new JRE package, for example:

```
rpm -i jre-7u51-linux-i586.rpm
```

e. Verify that JRE 1.7 has been installed.

```
java -version
```

```
java version "1.7.0_51"
```

5. Upgrade Oracle Fabric Manager to 4.3.0.

- On Linux: [“Install Oracle Fabric Manager on a Linux Server \(RPM File\)”](#) on page 31.
- On Windows: [“Install Oracle Fabric Manager on a Windows Server”](#) on page 33.

6. Start Oracle Fabric Manager.

See [“Start Oracle Fabric Manager Services”](#) on page 39.

Removing the Oracle Fabric Manager Software

The Oracle Fabric Manager software can be removed from the Oracle Fabric Manager Server like you would remove any Oracle Solaris, Linux, or Windows software. When Oracle Fabric Manager is removed, you cannot use the Oracle Fabric Manager GUI to configure or manage servers or virtual I/O. However, you can still use Oracle's XgOS CLI to manage the fabric in your data center.

This section contains the following topics:

- [“Remove the Oracle Fabric Manager Software From an Oracle Solaris Server”](#) on page 43
- [“Remove the Oracle Fabric Manager Software From a Linux Server”](#) on page 45
- [“Remove the Oracle Fabric Manager Software From a Windows Server”](#) on page 46

▼ Remove the Oracle Fabric Manager Software From an Oracle Solaris Server

To remove the Oracle Fabric Manager software from a Oracle Solaris server, follow this procedure:

1. **Use `pkginfo | grep xms` to query the name of the installed Oracle Fabric Manager package.**

```
# pkginfo | grep xms
system    xms                xsigo-xms
```

2. **When the full package name is known, issue the `pkgrm` command and specify the package name:**

Note - During the removal process, you will be prompted for input. Bold text indicates required input.

```
# rpm -e xsigo-xms-version_OFM.noarch
```

The following package is currently installed:

```
xms xsigo-xms
(x86) xsigo-xms-version
```

Do you want to remove this package? [y,n,?,q] **y**

Removing installed package instance <xms>

This package contains scripts which will be executed with super-user permission during the process of removing this package.

Do you want to continue with the removal of this package [y,n,?,q] **y**

Verifying package <xms> dependencies in global zone

Processing package information.

Executing preremove script.

Removing pathnames in class <none>

```
/opt/xsigo/xms/xmstools/xmscli/sandbox/bin/sedit
```

```
/opt/xsigo/xms/xmstools/xmscli/sandbox/bin/rm
```

```
/opt/xsigo/xms/xmstools/xmscli/sandbox/bin/pwd
```

```
/opt/xsigo/xms/xmstools/xmscli/sandbox/bin/mv
```

```
/opt/xsigo/xms/xmstools/xmscli/sandbox/bin/mkdir
```

```
/opt/xsigo/xms/xmstools/xmscli/sandbox/bin/ls
```

```
/opt/xsigo/xms/xmstools/xmscli/sandbox/bin/grep
```

```
/opt/xsigo/xms/xmstools/xmscli/sandbox/bin/echo
```

```
/opt/xsigo/xms/xmstools/xmscli/sandbox/bin/cp
/opt/xsigo/xms/xmstools/xmscli/sandbox/bin/cligen
/opt/xsigo/xms/xmstools/xmscli/sandbox/bin/chmod
<<display of installed files continues...>>
/opt/xsigo/xms/NOTICE
/opt/xsigo/xms/LICENSE
/opt/xsigo/xms/.openlogic.yaml
## Executing postremove script.
## Updating system information.
Removal of <xms> was successful.
```

3. **As an option, you can issue the `svcs xms` command to determine if any records still remain after removing the software. If no records are displayed, then the Oracle Fabric Manager software is completely removed. For example:**

```
root@terminus:~# svcs xms
svcs: Pattern 'xms' doesn't match any instances
STATE      STIME     FMRI
```

▼ Remove the Oracle Fabric Manager Software From a Linux Server

To remove the Oracle Fabric Manager software from a Linux server, you can use the `rpm -e` command. Follow this procedure:

1. **Use `rpm -qa | grep xms` to query the name of the installed Oracle Fabric Manager package.**
2. **From the directory where you installed Oracle Fabric Manager:**

```
[root@terminus xms_install]# rpm -e xsigo-xms-version_0FM.noarch.rpm
```

▼ Remove the Oracle Fabric Manager Software From a Windows Server

The Oracle Fabric Manager Windows installer does not currently have an uninstall option.

- **In the unlikely event that you will want to remove Oracle Fabric Manager, you can remove it from a Windows server by using either of the following methods:**
 - **Selecting the Oracle Fabric Manager uninstall option of the Oracle Fabric Manager install program which is available through the server's Start menu (for example, Start->Programs->Oracle Fabric Manager->Uninstall).**
 - **For Windows Server 2003, standard Windows Add or Remove Programs option (Start->Settings->Control Panel->Add or Remove Programs).**
 - **For Windows Server 2008, use the standard Uninstall a Program option (Start->Control Panel->Uninstall a Program).**

Configuring a Certificate for Oracle Fabric Manager


This section contains the following topics:

- [“Certificate Overview” on page 46](#)
- [“Certificate Installation Verification” on page 48](#)
- [“Creating and Installing a CA-Signed Certificate” on page 49](#)

Certificate Overview

A signed certificate should be installed on the Oracle Fabric Manager Server. When installed, Oracle Fabric Manager can be validated through the certificate that the Oracle Fabric Manager Server presents. The certificate identifies the Oracle Fabric Manager user interface as a trusted application. Even though the signed certificate is not a strict requirement, it is recommended. Certificates can be signed through any of the common Certificate Authorities (CAs), such as Verisign[®], Digi-Sign[®], Thawte[®], and so on.

When you configure a security certificate on Oracle Fabric Manager, you are allowing the client browser to validate the Oracle Fabric Manager Server and create a trust association.



In this example, when you attempt to log in to the Oracle Fabric Manager Server (the black line) the application (Oracle Fabric Manager) starts. At this point, the browser searches for a certificate, and the server presents its certificate to the client browser (the blue line). When the client browser recognizes the certificate as signed by a valid Certificate Authority (CA), the client browser then establishes that the application running in the browser (Oracle Fabric Manager) is trusted.

If the certificate is not present, not valid, or cannot be verified, the browser cannot validate the integrity of the Oracle Fabric Manager application, and the application running in the browser (Oracle Fabric Manager) is not trusted. As a result, the browser displays a warning page whenever you attempt to start Oracle Fabric Manager by logging into the Oracle Fabric Manager Server.

Configuring a certificate for Oracle Fabric Manager requires sending a certificate signing request and installing signed certificates.

Certificate Signing Request (CSR)

The Certificate Signing Request (CSR) is a request that is sent to a Certificate Authority (CA). The request consists of an public key and other information that is generated from the Oracle Fabric Manager Server. The public key is unique and identifies each Oracle Fabric Manager Server.

Besides the public key, you will provide some additional information such as contact emails, country of origin and so on. This information will be included in the CSR so that the CA can assign the signed certificate. Before generating the CSR, you will find it helpful to gather the names of the Oracle Fabric Manager Servers on which you will be installing the certificate(s) as well as the host name(s).

Signed Certificate

The signed certificate is a series of alphanumeric characters that is installed on the Oracle Fabric Manager Server. The signed certificate is generated by the Certificate Authority (CA) and uniquely guarantees the authenticity of Oracle Fabric Manager when it attempts to establish

HTTPS connectivity to the Oracle Fabric Manager Server. When the signed certificate is imported into the Oracle Fabric Manager Server, the certificate is used to verify the authenticity of the Oracle Fabric Manager Server. Through the HTTPS connection to the Oracle Fabric Manager Server, the client attempts to validate Oracle Fabric Manager and establish trust. If trust is established, the HTTPS connection is successful, and subsequent HTTPS transactions are secure.

Because a CSR is generated by an individual Oracle Fabric Manager Server, you cannot move the resulting signed certificate to a different Oracle Fabric Manager Server. Also, when you are importing signed certificates from the CA to the Oracle Fabric Manager Server, make sure that you import the correct certificate onto each Oracle Fabric Manager Server.

keytool Utility

The `keytool` utility is a powerful tool for generating and managing public and private keys for server certificates. With `keytool`, you can also manage the keystore on individual servers as well as the required license pairs. `keytool` is included in Java Runtime Environment 1.7 package, and the utility runs on any standard server operating system.

Extensive documentation about `keytool` is available on the web, for example <http://docs.oracle.com/javase/7/docs/technotes/tools/solaris/keytool.html>. Within the utility itself, help is available by using the `-help` option after Java is installed. For the purpose of creating and installing the certificate on the Oracle Fabric Manager server, you will use the following command options for `keytool`:

- `-genkeypair`
- `-alias`
- `-keyalg`
- `-keystore`
- `-keysize`
- `-certreq`
- `-file`
- `-import`
- `-trustcacerts`

Certificate Installation Verification

When you start Oracle Fabric Manager, you might see an error page that indicates a certificate error. This error occurs because Oracle Fabric Manager cannot be verified as a trusted application because it does not have a certificate that the client browser recognizes as signed by a trusted CA. By default, Oracle Fabric Manager has a self-signed certificate which allows the application to be usable without a certificate signed by a CA. As a result, you can log in to Oracle Fabric Manager and use the application even if the certificate error occurs.

You can skip past the certificate error by simply clicking Continue to this website, and if you do so, you can use Oracle Fabric Manager for all supported configuration and management workflows. The functionality of Oracle Fabric Manager is not affected by the presence or absence of a CA-signed certificate. If you choose not to install the certificate, you will see an error displayed in the browser's Address bar. Additional information is available through the Certificate Error notification in the browser's Security Status bar.

Configuring the Oracle Fabric Manager Security Certificate is not mandatory. However, by not installing a certificate, the certificate error page will be displayed every time you log in to Oracle Fabric Manager.

Creating and Installing a CA-Signed Certificate

This section contains the following topics:

- [“Oracle Fabric Manager Certificate Process” on page 49](#)
- [“Create and Install a Certificate on a Linux or Oracle Solaris Server” on page 49](#)
- [“Create and Install a Certificate on a Windows Server” on page 52](#)

Oracle Fabric Manager Certificate Process

You can generate a CSR from the Oracle Fabric Manager Server. When the Certificate Authority (CA) responds, the certificate is installed and the certificate error is cleared on the next log in to Oracle Fabric Manager. Configuring the Oracle Fabric Manager security certificate occurs through commands on the management console running on the host Oracle Fabric Manager Server. Therefore, make sure that you have a management console available to run the commands. Configuring the Oracle Fabric Manager Certificate takes the following steps:

1. Generate a public key for use in the CSR
2. Transmit (export) the CSR to the CA
3. Receive the signed certificate from the CA
4. Install (import) the Certificate on the Oracle Fabric Manager Server

Proceed to the appropriate section depending on which type of Oracle Fabric Manager server you have in use:

- [“Create and Install a Certificate on a Windows Server” on page 52](#)

▼ Create and Install a Certificate on a Linux or Oracle Solaris Server

The CSR is generated on the Oracle Fabric Manager server by creating a public key for the Oracle Fabric Manager Server then embedding that key into the CSR. The public key is

uniquely associated with the Oracle Fabric Manager Server, and cannot be moved or used on a different server. If you provision a new Oracle Fabric Manager server and decommission the existing one, you will need to generate a new CSR from the new Oracle Fabric Manager server, and install the signed certificate on the new server.

The CSR is generated from the Oracle Fabric Manager Server through the `keytool` utility. By default, Oracle Fabric Manager uses a self-signed certificate to allow for installation and use. However, the Oracle Fabric Manager signed certificate is not a CA-signed certificate, therefore the browser will not be able to validate Oracle Fabric Manager, and an error page will be displayed until you install a CA-signed certificate.

As part of generating a CSR, you will be prompted to enter a password for the keystore, which is a database on the Oracle Fabric Manager server that contains all the public keys and signed certificates. The keystore is password protected to keep it secure from anyone who might want to use the keys and certificates for malicious intent. Make sure to remember the password because you will need to enter it when you install the signed certificate.

Note - Certificates have a life span and can expire if not renewed. If you installed a certificate and Oracle Fabric Manager was running without the Certificate Error, and suddenly starts displaying the Certificate Error page again, check the expiration date on the certificate.

The following procedure shows how to create and import a certificate on a Linux server. The same procedure applies to Oracle Solaris servers as well, with the exception of how the Oracle Fabric Manager service is started. To create and import a certificate on a Linux or Oracle Solaris server, follow this procedure:

1. Remove any existing certificates:

```
rm /opt/xsigo/xms/conf/xms_cacerts
```

2. Use the `keytool` utility with the `-genkeypair` argument to generate a certificate for the Oracle Fabric Manager server:

```
./keytool -genkeypair -alias xms -keyalg RSA -keystore /opt/xsigo/xms/conf/xms_cacerts -keysize 2048
```

3. Create a certificate signature request:

```
./keytool -certreq -keyalg RSA -alias xms -file /opt/xsigo/xms/conf/certreq.csr -keystore /opt/xsigo/xms/conf/xms_cacerts
```

4. Edit the CSR to verify that it contains the correct information. If it doesn't, enter the correct information:

```
vi /opt/xsigo/xms/conf/certreq.csr
```

5. Submit the CSR to a certificate authority (CA), such as Verisign. After you submit the CSR to the CA, you receive an email from CA.

6. **Complete the directions in the email to import the intermediate certificates on to the Oracle Fabric Manager server.**
7. **On the Oracle Fabric Manager server, use keytool, to create a primary and secondary intermediate certificate files.**

For the primary intermediate file:

```
./keytool -import -trustcacerts -alias primaryintermediate -keystore
/opt/xsigo/xms/conf/xms_cacerts -file /opt/xsigo/xms/conf/primary_inter.cer
```

For the secondary intermediate file:

```
./keytool -import -trustcacerts -alias secondaryintermediate -keystore
/opt/xsigo/xms/conf/xms_cacerts -file /opt/xsigo/xms/conf/secondary_inter.cer
```

8. **Create an SSL certificates file for the intermediate licenses you just created:**
9. **Verify that the intermediate key files you created (for example "primaryintermediate" and "secondaryintermediate" are present in /opt/xsigo/xms/conf/.**
10. **Send the contents of the xms_cacerts directory as a .txt file. To do so, you will need to provide the password to access the keystore on the Oracle Fabric Manager server.**

```
./keytool -list -v -keystore /opt/xsigo/xms/conf/xms_cacerts >keystorelist.txt
```

Enter keystore password: <enter your keystore password>

11. **Go to the xms/conf directory, and find the server.xml file.**
12. **Using vi, emacs, or another standard UNIX editor, add the following lines to server.xml:**

```
<!-- Define a SSL HTTP/1.1 Connector on port 8443
```

```
    This connector uses the JSSE configuration, when using APR, the
```

```
    connector should be using the OpenSSL style configuration
```

```
    described in the APR documentation -->
```

```
<Connector port="8443" protocol="org.apache.coyote.http11.Http11NioProtocol" SSLEnabled="true"
```

```
    maxThreads="150" scheme="https" secure="true"
```

```
    keystoreFile="/opt/xsigo/xms/conf/xms_cacerts"
```

```
keystorePass=<enter the keystore password>
```

```
keyAlias="xms"
```

```
enableLookups="true"
```

13. Restart the Oracle Fabric Manager service:

- **On a Linux server, use service xms stop. After all services are stopped, use service xms start**
- **On an Oracle Solaris server, use svcadm disable xms. After all services are stopped, user svcadm enable xms.**

14. Consult the email you received from the CA and follow the instruction for importing the certificate into your browser.

15. When the certificate is correctly imported, start your browser and check the security option to verify that the certificate was installed.

16. Examine the certificate to make sure that the information you specified is correct, for example:

- **in a FireFox browser, you can check the certificate by following this menu path:**

Tools → Options... → Advanced tab → View Certificates

- **in an Internet Explorer browser, you can check the certificate by following this menu path:**

ToolsInternet Options → Content tab → Certificates

17. Log in to Oracle Fabric Manager. If the certificate was correctly installed, the certificate error page should no longer be displayed. If you had a previous browser session running Oracle Fabric Manager, close the browser completely before logging back in to Oracle Fabric Manager after installing the certificate. By doing so, you will flush the browser cache.

▼ **Create and Install a Certificate on a Windows Server**

The CSR is generated on the Oracle Fabric Manager server by creating a public key for the Oracle Fabric Manager Server then embedding that key into the CSR. The public key is uniquely associated with the Oracle Fabric Manager Server, and cannot be moved or used on a different server. If you provision a new Oracle Fabric Manager server and decommission the

existing one, you will need to generate a new CSR from the new Oracle Fabric Manager server, and install the signed certificate on the new server.

The CSR is generated from the Oracle Fabric Manager Server through the keytool utility. By default, Oracle Fabric Manager uses a self-signed certificate to allow for installation and use. However, the Oracle Fabric Manager signed certificate is not a CA-signed certificate, therefore the browser will not be able to validate Oracle Fabric Manager, and an error page will be displayed until you install a CA-signed certificate.

As part of generating a CSR, you will be prompted to enter a password for the keystore, which is a database on the Oracle Fabric Manager server that contains all the public keys and signed certificates. The keystore is password protected to keep it secure from anyone who might want to use the keys and certificates for malicious intent. Make sure to remember the password because you will need to enter it when you install the signed certificate.

Note - Certificates have a life span and can expire if not renewed. If you installed a certificate and Oracle Fabric Manager was running without the Certificate Error, and suddenly starts displaying the Certificate Error page again, check the expiration date on the certificate.

To create and import a certificate on a Windows server, follow this procedure:

1. Remove any existing certificates by deleting:

C:\Program Files\xms\conf\xms_cacerts file

2. Use the keytool utility to generate a certificate for the Oracle Fabric Manager server:

```
C:\Program Files\Java\jre7\bin>keytool.exe -genkeypair -alias xms -keyalg RSA -keystore "C:\Program Files\xms\conf\xms_cacerts" -keysize 2048
```

3. Create a certificate signature request:

```
C:\Program Files\Java\jre7\bin>keytool.exe -certreq -keyalg RSA -alias xms -file c:\Program Files\xms\conf\certreq.csr -keystore xms_cacerts
```

4. Using Microsoft Notepad or another standard Windows file editor, view the CSR to verify that it contains the correct information. If it doesn't, enter the correct information:

C:\Program Files\xms\conf\certreq.csr

5. Submit the CSR to a certificate authority (CA), such as Verisign as shown in this example:

After you submit the CSR to the CA, you receive an email from CA.

6. Complete the directions in the email to import the intermediate certificates on to the Oracle Fabric Manager server.

- 7. On the Oracle Fabric Manager server, use keytool to create a primary and secondary intermediate certificate files.**

For the primary intermediate file:

```
C:\Program Files\Java\jre7\bin>keytool.exe -import -trustcacerts -alias primaryintermediate -keystore "c:\Program Files\xms\conf\xms_cacerts" -file c:\Program Files\xms\conf\primary_inter.cer
```

For the secondary intermediate file:

```
C:\Program Files\Java\jre7\bin>keytool.exe -import -trustcacerts -alias secondaryintermediate -keystore "c:\Program Files\xms\conf\xms_cacerts" -file c:\Program Files\xms\conf\secondary_inter.cer
```

- 8. Create an SSL certificates file for the intermediate licenses you just created:**

```
C:\Program Files\Java\jre7\bin>keytool.exe -import -trustcacerts -alias xms -keystore "c:\Program Files\xms\conf\xms_cacerts" -file c:\ssl\ssl_cert.cer
```

- 9. Verify that the intermediate key files you created (for example "primaryintermediate" and "secondaryintermediate" are present in /opt/xsigo/xms/conf/.**
- 10. Send the contents of the xms_cacerts directory as a .txt file. To do so, you will need to provide the password to access the keystore on the Oracle Fabric Manager server.**

```
C:\Program Files\Java\jre7\bin>keytool.exe -list -v -keystore xms_cacerts >keystorelist.txt
```

Enter keystore password: <enter your keystore password>

- 11. Go to the xms/conf directory, and find the server.xml file.**
- 12. Using Windows Notepad of another standard Windows file editor, add the following lines to server.xml:**

```
<!-- Define a SSL HTTP/1.1 Connector on port 8443  
    This connector uses the JSSE configuration, when using APR, the  
    connector should be using the OpenSSL style configuration  
    described in the APR documentation -->  
<Connector port="8443" protocol="org.apache.coyote.http11.Http11NioProtocol" SSLEnabled="true"  
    maxThreads="150" scheme="https" secure="true"  
    keystoreFile=""c:\Program Files\xms\conf\xms_cacerts"  
    keystorePass=<enter the keystore password>  
    keyAlias="xms"
```

```
enableLookups="true"
```

13. **On the Windows server's Services tab, select the Oracle Fabric Manager service and restart it.**

You can use any of the available methods to restart the service.

14. **Consult the email you received from the CA and follow the instruction for importing the certificate into your browser.**
15. **When the certificate is correctly imported, start your browser and check the security option to verify that the certificate was installed.**
16. **Examine the certificate to make sure that the information you specified is correct, for example:**

- **in a FireFox browser, you can check the certificate by following this menu path:**

Tools → Options... → Advanced tab → View Certificates

- **in an Internet Explorer browser, you can check the certificate by following this menu path:**

Tools → Internet Options → Content tab → Certificates

17. **Log in to Oracle Fabric Manager.**

If the certificate was correctly installed, the certificate error page should no longer be displayed. If you had a previous browser session running Oracle Fabric Manager, close the browser completely before logging back in to Oracle Fabric Manager after installing the certificate. By doing so, you will flush the browser cache.

Starting Oracle Fabric Manager and Logging In

This section contains the following topics:

- [“Credentials for Intital Login” on page 56](#)
- [“Oracle Fabric Manager and Authentication” on page 57](#)
- [“Log In to Oracle Fabric Manager” on page 57](#)

Credentials for Initial Login

To log in to the Oracle Fabric Manager interface, you must provide a valid user name and password. The default user name and password are different for Oracle Fabric Manager on an Oracle Solaris, Linux, or Windows Oracle Fabric Manager server:

- For Oracle Solaris, the default user name and password are the same as the default user account on the Oracle Solaris server (root). Additional user accounts can be configured for Oracle Fabric Manager, but they must be configured as non-root user accounts on the Oracle Solaris server.
- For Linux, the default user name and password are the same as the default user account on the Linux server (root). Additional user accounts can be configured for Oracle Fabric Manager, but they must be configured as non-root user accounts on the Linux server.
- For Windows, the default user name and password are the same as the default administrator account on the Windows server (administrator). Additional, non-admin accounts can be configured for Oracle Fabric Manager, but they must be configured as non-administrator accounts on the Windows server.

Note - Because some installations do not allow administrator or root access to servers, Oracle also provides the ofmadmin account which can be used to log in to the Oracle Fabric Manager Server.

You will use this user name and password the first time you access Oracle Fabric Manager. One of these accounts (root, administrator, or ofmadmin) is required to install additionally purchaseable features such as plug ins. Oracle strongly recommends that after logging into Oracle Fabric Manager, you create specific user accounts (instead of continuing to use the root, administrator, or ofmadmin accounts). When you create these specific user accounts, make sure to assign them with the appropriate roles.

Note - You should have the root default user name and password kept as securely as possible. After initially logging in to Oracle Fabric Manager, you should create additional user accounts with specific privileges, as documented in [“Working With User Roles” on page 497](#). When you create the additional non-root user accounts in Oracle Fabric Manager, be aware that these user accounts must also exist in the underlying OS (either an Oracle Solaris, Linux, or Windows user account). You cannot log in to Oracle Fabric Manager if no underlying user account exists at the server's OS level, or if there is any difference between the user names—for example, user name joey at the OS level and joe at Oracle Fabric Manager will not pass authentication and allow login to Oracle Fabric Manager.

Oracle Fabric Manager and Authentication

Oracle Fabric Manager does no authentication of its own. Instead, it presents a log in screen, then passes the user name and password string to the underlying Oracle Solaris, Linux, or Windows OS (pass-through), which allows or disallows the authentication. Because Oracle Fabric Manager login is a pass through, additional authentication methods are supported, such as LDAP AD. Be aware that the additional user accounts on the Active Directory server(s), the Windows server OS, and the Oracle Fabric Manager server must all be congruent. Additional information is available for specific additional authentication methods such as LDAP AD. For more information, see [“Working With User Roles” on page 497](#). When the OS and any other authentication methods complete, the appropriate blob of information is sent to Oracle Fabric Manager, which then sets the correct Oracle Fabric Manager-level privileges for the authenticated user. When those privileges are applied, they determine which objects in the Oracle Fabric Manager GUI are available to the authenticated user. Be aware that the user name can have different rights or privileges at each level, but for Oracle Fabric Manager, only the Oracle Fabric Manager privileges granted to the user are enforced. For example, user joey at the OS level might be a super user with all rights, but user joey at the Oracle Fabric Manager level might be a network administrator user and therefore, only the privileges for a network-admin role are applied, not the super user privileges. In this case, the user joey has different privileges, but only the Oracle Fabric Manager privileges are enforced for managing Fabric Devices and virtual I/O through Oracle Fabric Manager.

▼ Log In to Oracle Fabric Manager

When you are logged in, your management session remains active as long as you are actively configuring or managing through Oracle Fabric Manager. However, Oracle Fabric Manager has a management session inactivity timer of 30 minutes. As a result, if you do not perform a click action (for example, selecting an object or starting a wizard) in Oracle Fabric Manager for 30 minutes, you are logged out and must log back in.

You can start Oracle Fabric Manager in either of the following ways:

- Through a client browser window. Proceed to [Step 1](#).
- Through the VMware vCenter if you have installed the correct plug-in. For information, see the *Oracle Fabric Manager VMware Integrator Plug-In User Guide*.

To start Oracle Fabric Manager as stand-alone application, follow this procedure:

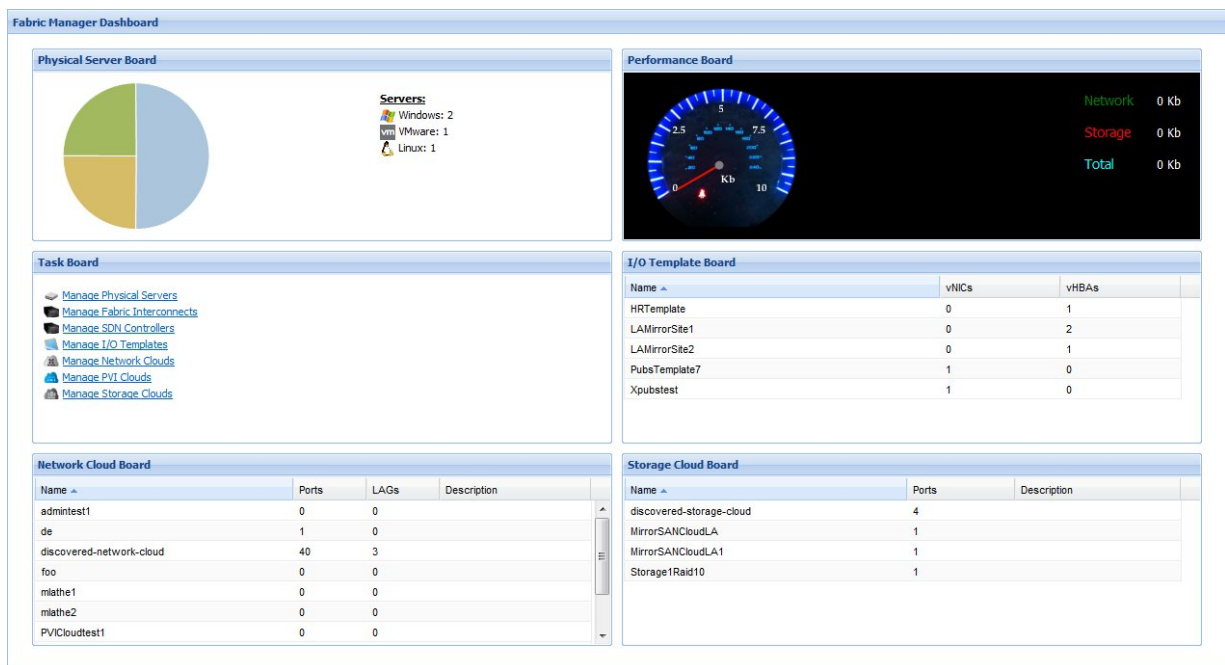
1. **Start one of the supported browsers. For information about supported browsers, see [“Client Requirements” on page 26](#).**
2. **Point your browser to `https://<server-name>:8443/xms` where `<server-name>` is the IP address or host name of the Oracle Fabric Manager server, plus port 8443. For example:**

https://gorilla:8443/xms

would start a Oracle Fabric Manager session on the Oracle Fabric Manager server named “gorilla”.

Note - You might see a pop-up dialog notifying you of the suggested screen resolution.

3. In the User Name field, enter a valid user name.
4. In the Password field, enter the password for the user name you are specifying.
5. Click Login to display the Welcome page.



6. If you will be managing Oracle Fabric Manager through VMware Virtual Infrastructure, you will need to register Oracle Fabric Manager with the Virtual Center Server (Service Manager → VMware Integration).

Otherwise, you can manage Oracle Fabric Manager through one of the supported browsers. Proceed to the next step.

- 7. On the navigation panel, click Managed Devices → Fabric Interconnects or Managed Devices → SDN Controller to discover one or more Oracle Fabric Interconnects or Oracle SDN Controllers, respectively.**

For information, see [“Working With the Fabric Device” on page 157](#) or the *Oracle SDN Controller User Guide, 1.0.0*.

Working With the Interface

The Oracle Fabric Manager interface is a graphical user interface that enables you to configure and manage Oracle Fabric Devices and the virtual network and storage resources associated with the Fabric Devices. Additional management and configuration is supported for additional functionality, such as host servers, storage targets, access control, and so on.

Oracle Fabric Manager runs on a stand-alone server called the Oracle Fabric Manager Server, which can be either Oracle Solaris, Linux, or Windows based, or as an add-on service that can be integrated into a VMware vSphere or Virtual Infrastructure Server. When Oracle Fabric Manager is installed, the Oracle Fabric Manager interface provides an intuitive and robust management suite for virtualized I/O that is extended seamlessly to various types of Oracle Solaris, Windows, Linux, or ESX host servers.

The Oracle Fabric Manager interface contains the following components:

- [“Banner Features” on page 61](#)
- [“Understanding the Navigation Panel” on page 62](#)
- [“Understanding the Work Panel” on page 72](#)
- [“Using the Maintenance Menu” on page 77](#)

Banner Features



The banner contains buttons and icons that enable you to get general information about Oracle Fabric Manager, and the Fabric Devices that Oracle Fabric Manager has discovered. Through the Banner, you can:

- See a quick, at-a-glance tally of the current alarms through icons that summarize the number of critical, major, warning, and minor alarms reported by the Oracle Fabric Manager Server or Fabric Device. Additional alarm information is provided in more detail through the Alarm Log and Alarm History options on the Navigation panel.
- Perform Oracle Fabric Manager maintenance tasks, such as backing up and restoring the Oracle Fabric Manager configuration.

- Set the UI to projector mode, which is a conditional setting that is available on pages that contain dark backgrounds. This option supports setting the dark backgrounds to a light background that is easier to see when Oracle Fabric Manager is displayed over a video projector.
- See a quick, at-a-glance status of the HA Oracle Fabric Manager feature as well as click the HA state to configure HA Partners (two Oracle Fabric Manager servers that provide high availability) and configure the current Oracle Fabric Manager servers with HA properties if needed.
- See which user account is logged in to the Oracle Fabric Manager Server as well as what role is assigned to that user account
- Determine which domain the Oracle Fabric Manager Server and its managed Fabric Device (s) are part of
- Logout of the Oracle Fabric Manager Server. In addition to explicitly logging out, Oracle Fabric Manager has an inactivity timer of 30 minutes. As long as you are actively using Oracle Fabric Manager, the inactivity timer is not triggered. However, if no click or keyboard action is detected for 30 minutes, the Oracle Fabric Manager automatically closes any currently active session. You can log back in as usual, and the inactivity timer is reset.
- Get version information for the Oracle Fabric Manager software installed on the Oracle Fabric Manager Server

Understanding the Navigation Panel

This section describes the navigation panel interface.

- [“Navigation Panel” on page 62](#)
- [“General Option” on page 65](#)
- [“Server Resource Manager” on page 66](#)
- [“Network Cloud Manager” on page 67](#)
- [“Storage Cloud Manager” on page 68](#)
- [“Service Manager” on page 69](#)
- [“Security Manager” on page 70](#)
- [“Managed Devices” on page 70](#)
- [“Plugins” on page 71](#)

Navigation Panel

The navigation panel appears as a list on the left side of the Oracle Fabric Manager. The navigation panel is divided into logical groupings of related functionality. These groupings are known as “managers” with the exception of the first grouping, which is the General option.

Each manager in the list expands and collapses one level to display links that you can use for configuration and management of one or more entities within it. For example, the Server Resource Manager contains functionality related to physical servers. The one exception is the General option, which contains system-wide information, such as Topology and Alarm information. The navigation panel is a common point for starting most work flows, so you can use it as a first step in most configuration and management tasks.

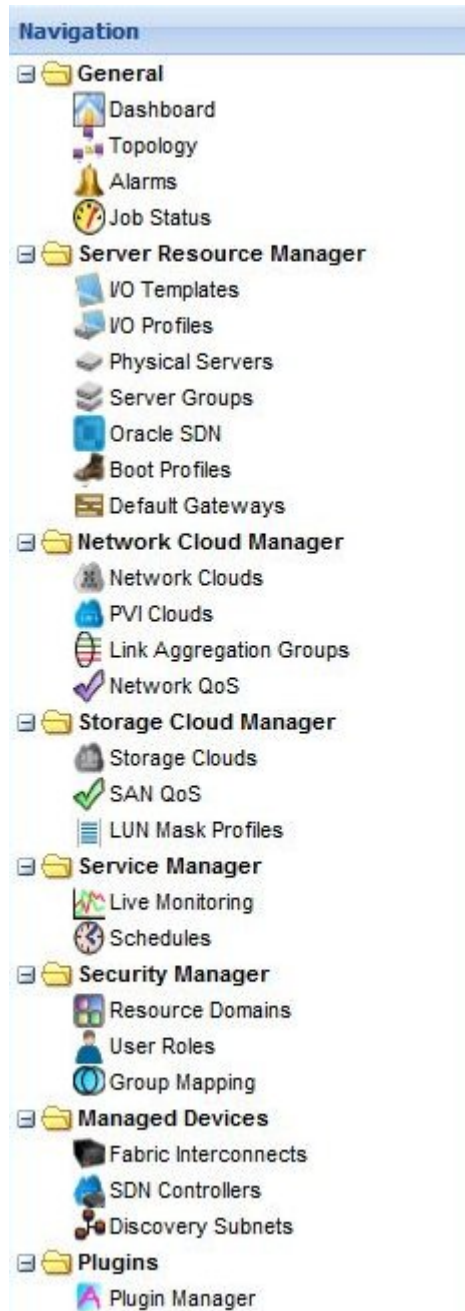
Note - The navigation panel is a starting point, but you can also use the “Task Board” and other of the sub-boards on the Dashboard. Through the Dashboard and its contents you can access some of the most common tasks in configuring your virtual I/O. The Dashboard is useful for accessing common tasks, but the Navigation panel provides a more complete starting point for configuration and management tasks.

The following example shows the Navigation panel with all options closed.



You can use the navigation panel to display appropriate icons that allow configuration and management in the main work area of the work panel. For example, if you wanted to create QoS on a vNIC, you would open the Network Cloud Manager on the navigation panel, then click Network QoS to display the Network QoS Summary in the work panel of the Oracle Fabric Manager interface.

You can toggle the navigation panel between hide and display by clicking the button in the upper right corner of the navigation panel. Hiding the navigation panel gives you more room in the Oracle Fabric Manager work area for configuration and management, and in wide tables, allows more columns to be displayed. The following example shows the hide and display toggle switch for the navigation panel.



The navigation panel contains the following items:

General Option

The General option in the navigation panel contains the Dashboard, Topology, Alarms, and Job Status options.



Through the Dashboard link you can display the Dashboard. The Dashboard contains multiple sub-boards that give you quick, at-a-glance information about various elements of the Fabric Device(s) and servers under management by the Oracle Fabric Manager. The Dashboard contains the following sections:

- The Physical Server Board, which shows a pie graph of the distribution of different server types discovered in the network, based on host OS type
- The Performance Board, which shows a speedometer with the following information:
 - Aggregate network throughput of all Fabric Devices managed by Oracle Fabric Manager
 - Aggregate storage throughput of all Fabric Devices managed by Oracle Fabric Manager
 - Total traffic throughput, which is the sum of network and storage traffic on all Fabric Devices that are managed by Oracle Fabric Manager
- The Task Board, which contains the following links and controls:
 - Manage Physical Servers
 - Manage Fabric Devices, including discovering specific Fabric Interconnects or Oracle SDN Controllers by either IP address or system name, or scanning the network for all connected Fabric Interconnects or OSDN Controllers that are online and running.
 - Manage Private Virtual Interconnect (PVI) Clouds
 - Manage I/O Templates
 - Manage Network Clouds
 - Manage Storage Clouds
- The I/O Template Board, which shows a table of all the configured I/O Templates in Oracle Fabric Manager. I/O Templates are displayed regardless of whether they are deployed on a server or not
- The Network Cloud Board, which shows a table of all the configured Network Clouds in Oracle Fabric Manager. Network Clouds are displayed regardless of whether they are deployed on a server or not

- The Storage Cloud Board, which shows a table of all the configured Storage Clouds in Oracle Fabric Manager. Storage Clouds are displayed regardless of whether they are deployed on a server or not

Through the Topology link you can display the Topology Overview. In the Topology Overview, you can see how individual servers and server groups (if any) are connected to their respective Network and Storage Clouds through the Fabric Device(s) in the network. Additional views are available through the Topology page. These additional views allow for different ways to see the virtual interconnectivity of servers, Fabric Devices, Network and Storage clouds, targets, ports, vNICs, and vHBAs. For more information about the Topology, see [“Working With Discovery Subnets” on page 395](#).

Through the Alarms link you can display the Alarms page which displays the current active alarms on the Alarms Summary tab and the historical information for alarms on the Alarm History Summary tab. For more information about the Alarms Summary, see [“Working With Alarms” on page 399](#).

Through the Job Status link you can display information about jobs that are occurring on the Fabric Device and/or the Oracle Fabric Manager Server. Some jobs are self contained (called primary jobs), but others contain sub-jobs. Any primary job is listed as an entry on the Jobs Status page, but jobs that contain sub-jobs are listed as an expandable or collapsible job on the Job Status page. For such jobs, you can expand the sub-jobs to see more granular information about the status of each sub-job. For more information about the Job Status, see [“Working With Oracle Fabric Manager Jobs” on page 91](#).

Server Resource Manager

The Server Resource Manager in the navigation panel displays information about the physical servers connected to the Fabric Devices that Oracle Fabric Manager has discovered.



The Server Resource Manager contains a list of the following links:

- I/O Templates, which links to information about all the I/O Templates configured in Oracle Fabric Manager. If you click the I/O Templates link, a table is displayed that contains all

configured I/O Templates and information about them. For more information about creating and managing I/O Templates, see [“Working With I/O Templates” on page 205](#).

- I/O Profiles, which links to information about all the I/O Profiles configured in Oracle Fabric Manager. If you click the link, a table is displayed that contains all of the created I/O Profiles as well as information about them, and other functions that are related to configuring and managing I/O Profiles. For more information about I/O Profiles, see [“Working With I/O Profiles” on page 259](#).
- Physical Servers, which links to information about the physical servers discovered by Oracle Fabric Manager. If you click the link, a table is displayed that contains all of the physical servers as well as information about each of them. For more information about physical servers, see [“Working With Physical Servers” on page 303](#).
- Server Groups, which links to information about all the server groups configured in Oracle Fabric Manager. If you click the Server Groups link, a table is displayed that contains all configured server groups and information about them. For more information about server groups, see [“Working With Server Groups” on page 325](#).
- Oracle SDN, which links to information about all the individual InfiniBand fabrics under management in Oracle Fabric Manager. Fabrics are used by Oracle Fabric Interconnects and Oracle Software Defined Networking (SDN) Controllers. When you click this link, a table of all the individual Oracle Fabrics under management is displayed. Oracle Fabrics allow the configuration of Private Virtual Interconnects (PVI), which support “east-west” traffic—for example, vMotion traffic.
- Boot Profiles, which links to information about optional configurations for supporting SAN Boot and iSCSI Boot of physical servers. If you click the Boot Profiles link, a table is displayed that contains all the Boot Profiles configured in Oracle Fabric Manager for physical host servers. For more information about Boot Profiles, see [“Working With Boot Profiles” on page 333](#).
- Default Gateways, which links to information about default gateways configured in Oracle Fabric Manager and information about the gateways and how they are associated with host servers. If you click the Default Gateways link, a table is displayed that contains all configured default gateways and information about them. For information about default gateways, see [“Working With Default Gateways” on page 153](#).

Network Cloud Manager

The Network Cloud Manager in the navigation panel displays information about features configurable in Oracle Fabric Manager's IP network clouds and their child elements.




The Network Cloud Manager contains the following options:

- Network Clouds, which links to information about Network Clouds that have been configured in Oracle Fabric Manager. If you click the Network Clouds link, a table is displayed that contains all Network Clouds, ports, and network properties associated with the Network Clouds. For more information about Network Clouds, see [“Working With Network Clouds” on page 99](#).
- PVI Clouds, which links to information about Private Virtual Interconnect Clouds which are available to either Oracle Fabric Interconnects or Oracle Software Defined Network (SDN) Controllers. When you click this link, a table of all the configured PVI Clouds is displayed as well as information about each of them. For more information, see the *Oracle Fabric Accelerator Quick Start Guide* or the *Oracle Software Defined Network Controller User Guide*.
- Link Aggregation Groups, which links to information about link aggregation groups (LAGs) configured in Oracle Fabric Manager and the ports assigned to them. If you click the Link Aggregation Groups link, a table is displayed that contains the LAG configured in Oracle Fabric Manager as well as the ports that constitute each LAG. For more information about LAGs, see [“Working With Link Aggregation” on page 281](#).
- Network QoS, which links to information about network Quality of Service (QoS) Profiles that have been configured in Oracle Fabric Manager. If you click the Network QoS link, a table is displayed that contains all Network QoS Profiles and information about them. For more information about creating and managing Network QoS, see [“Working With Network QoS” on page 291](#).

Storage Cloud Manager

The Storage Cloud Manager in the navigation panel displays information about all the Fibre Channel-connected storage targets discovered by Oracle Fabric Manager, as well as controls for managing them. Through this option, you can create Storage Clouds, LUN Masks, set SAN QoS parameters for vHBAs, and so on.



The Storage Cloud Manager option contains the following links:

- Storage Clouds, which links to information about Storage Clouds that have been configured in Oracle Fabric Manager. If you click the Storage Clouds link, a table is displayed that contains all Storage Clouds, ports, and storage properties associated with the Storage Clouds. For more information about creating and managing Storage Clouds, see [“Working With Storage Clouds” on page 131](#).
- SAN QoS, which links to information about storage Quality of Service (QoS) profiles that have been configured on vHBAs on discovered Fabric Devices. If you click the SAN QoS link, a table is displayed that contains all SAN QoS Profiles and information about them. For more information about creating and managing SAN QoS, see [“Working With SAN QoS” on page 473](#).
- LUN Mask Profiles, which links to information about LUN Masks to filter the storage targets that are available to host servers. If you click this link, a table is displayed that contains all the LUN Masks, as well as properties associated with each LUN Mask. For more information about creating and managing LUN Masks, see [“Working With LUN Masks” on page 465](#).

Service Manager

The Service Manager supports features that allow for live, real-time performance monitoring and setting schedules for recurring tasks on the Oracle Fabric Manager server—for example, backing up the Oracle Fabric Manager server's config.



The Service Manager contains the following topics:

- Live Monitoring, which links to a statistics and performance grapher that shows live, real-time throughput on a per-server, per-vNIC, or per-vHBA basis. For more information about Live Monitoring, see [“Working With Live Monitoring” on page 433](#).
- Schedules, which links to a table of configurable schedules. The schedules enable you to configure a specific date and time at which scheduled tasks (such as Oracle Fabric Manager Server or Fabric Device backup) will occur. If you click this link, a table

of all the configured scheduled tasks will occur on a daily, weekly, or monthly basis. Through the table, you can also perform an on-demand backup of Oracle Fabric Manager Servers or Fabric Devices as needed. For more information, see [“Working With the Task Scheduler” on page 443](#).

Security Manager

The Security Manager contains options that allow the configuration of Fabric Devices in specific network domains as well as the configuration of Oracle Fabric Manager and Fabric Device users and roles.



The Security Manager contains the following topics:

- Resource Domains, which links to a wizard and summary tables that allow for flexibility in where the Fabric Device(s) and Oracle Fabric Manager Server are configured in the overall network, or specific sub-domains within the overall network. For information about resource domains, see [“Working With Domains” on page 477](#).
- User Roles, which links to dialogs that enable configuring specific user accounts for network administrators and assigning Oracle Fabric Manager roles to each account. The Oracle Fabric Manager role for a user account supersedes any roles or privileges assigned through the Oracle Solaris, Linux, or Windows OS. For information about User Roles, see [“Working With User Roles” on page 497](#).
- Group Mapping, which links to dialogs that enable specifying a mapping between a user's group configured in an external identity management system (IMS), such as Active Directory (AD) or Lightweight Directory Access Protocol (LDAP) and a Oracle Fabric Manager role or domain. For more information, see [“Working With User Roles” on page 497](#).

Managed Devices

The Managed Devices option on the navigation panel contains links related to the Fabric Device and the IP subnet in which they are installed and where the Fabric Devices can be discovered. Through this option you can view detailed information about each managed Fabric Device, all hardware inventoried in Oracle Fabric Manager, and software and hardware features that available on the Fabric Device.



The Managed Devices option contains the following links:

- Fabric Interconnects, which links to information about the Fabric Interconnects that Oracle Fabric Manager has discovered. If you click the link, a table is displayed that contains all of the Fabric Interconnects as well as information about each of them. For information about managing Oracle Fabric Interconnects, see [“Working With the Fabric Device” on page 157](#).
- SDN Controllers, which links to information about the Oracle SDN Controllers that Oracle Fabric Manager has discovered. If you click the link, a table is displayed that contains all of the Oracle SDN Controllers as well as information about each of them. For information about managing Oracle SDN Controllers, see the *Oracle SDN Controller User Guide, 1.0.0*.
- Discovery Subnets, which allows Oracle Fabric Manager to discover Fabric Devices that are located in a different IP subnet by establishing a connection between the Oracle Fabric Manager server and a Fabric Device which acts as a proxy within that remote subnet. Discovery subnets must be manually configured to allow contact between the Oracle Fabric Manager Server and the proxy off of its local subnet. For more information about Discovery Subnets, see [“Working With Discovery Subnets” on page 395](#).

Plugins

The Plugins option on the navigation panel contains links related to various Oracle Fabric Manager plug-ins that can be added to the Oracle Fabric Manager GUI to provide additional functionality. Through this option you can install, configure, and manage additional applications.



The contents of the Plugins option will vary based on which additional plugins you have purchased and installed. For example, one of Oracle Fabric Manager's application called Oracle Fabric Performance Monitor (which is separately purchaseable) might be installed. If so, that application would be listed by name in the Plugins option. The name would function as a link to a separate set of pages where the application can be managed.

The Plugins option contains Plugin Manager, option, which links to a information about the individual applications that have been installed into Oracle Fabric Manager. This option is always present because it is the utility in Oracle Fabric Manager that supports additional applications to seamlessly plug in to the core Oracle Fabric Manager product.

As additional Oracle applications are made available, the Plugin Manager will be used to install them into Oracle Fabric Manager.

Understanding the Work Panel

This section describes the work panel interface.

- [“Work Panel” on page 72](#)
- [“Filtering and Sorting Table Displays” on page 74](#)

Work Panel

The work panel is the main work area for Oracle Fabric Manager. Most configuration and management tasks occur through the work panel. The work panel has consistent layout for most features, but there are some exceptions.

The following example shows a typical Oracle Fabric Manager work panel.

The screenshot displays the Oracle Fabric Manager interface, divided into three main sections:

Physical Server Summary

Host Name	Host OS	vNI...	vHB...	Bound	Busy	State	I/O Profile Name	Fabric Device Ports	Groups
bering	Linux/2.6.18-238.el5.xg-3.7.1.LX1/x8...	2	2	✓		up	LAMirrorSite1_607	delaware:ServerPort22 oregon:ServerPort15	
ovn87-36.us.oracle.com	Linux/2.6.32-220.el6.x86_64/x86_64	0	0					ExtISw-2c90200443198-Port8	
ovn87-34.us.oracle.com	Linux/2.6.32-220.el6.x86_64/x86_64	0	0					ExtISw-2c90200443198-Port10	
ovn87-35.us.oracle.com	Linux/2.6.32-220.el6.x86_64/x86_64	0	0					ExtISw-2c90200443198-Port9	
ovn87-22	Linux/2.6.32-220.el6.x86_64.xg-5.0...	1	0	✓		up	iotemplate	ExtISw-2c90200443198-Port7	HRServerGroup9 HRServerGroup
ovn86-45	VMware/ESX-4.1.0.xg-3.2.0-r5773.E...	0	0			offline			

10 items

Server: bering

General | vNICs | vHBAs | Server Groups

Name: bering
Host OS: Linux/2.6.18-238.el5.xg-3.7.1.LX1/x86_64
I/O Profile Name: LAMirrorSite1_607
State: up
Fabric Devices Ports: delaware:ServerPort22,oregon:ServerPort15
Adapter FW Version: 5.3.0/3.0.0

Recent Jobs Summary

Time Updated	Job ID	State	Username	Job Detail
2013-10-01 18:25:23.011	RescanServers	completed	root	Rescan for new servers
2013-10-01 05:00:00.001	ScheduledJob	completed	scheduled	Execute scheduled job: DailyDelaware
2013-10-01 00:45:00.051	ScheduledJob	completed	scheduled	Execute scheduled job: lorenzo

5 items

The work panel is typically divided into the following parts:

- The Summary, which is the top frame in the work panel. The Summary is a high-level listing of all instances of specific object—for example, the Physical Server Summary contains a list of all physical servers that Oracle Fabric Manager is managing. The Summary contains some basic information about general properties for the objects.
- The Details frame, which is the middle frame in the work panel. The Details frame is a single instance of an object selected in a Summary. For example, the Physical Server Details frame contains information about a single physical server that you selected in the Physical Server Summary. The Details frame contains additional information that is specific to the selected object.
- The Details frame typically contains tabs that organize chunks of similar information in an intuitive and easy-to-use way. Some details frames are nested. For example, the Physical Server Details frame contains a vNICs tab, which lists all vNICs in the selected physical server. Each vNIC name on the vNICs tab is a link to another level of details frame—in this case, the vNIC Details. By using a combination of tabs and links in the Details frame, it is possible to drill-down to low-level details about a selected object.
- Due to the amount of nesting in the Details frame, a breadcrumb trail (“breadcrumbs” in this documentation) are available on the detail frame’s banner. Breadcrumbs are a user interface navigation tool of incremental hyperlinks that allow you to see the current level of depth for the current dialog or frame. You can also use the breadcrumbs to retrace your clicks.
- Recent Jobs, which is the bottom frame in the work panel. The Recent Jobs frame is a list of the status of recently attempted configuration or management tasks. For example, if you attempt to discover a Fabric Device, the Recent Jobs frame will show an interpretable result

based on the interaction between the Oracle Fabric Manager client, Oracle Fabric Manager Server, and Fabric Device. By default, the Recent Jobs frame shows results for the three most recently attempted tasks, but scroll bars on the frame allow you to view older results.

Note - Some exceptions to this layout are: * The Overview page, which has a number of sub-boards * The I/O Template Editor, which has one top frame for general properties and one large workspace for assembling the I/O Template's building blocks

Filtering and Sorting Table Displays

The Oracle Fabric Manager Summary pages, and also some of the details pages, display their data in tables. The following example shows the Physical Server Summary.

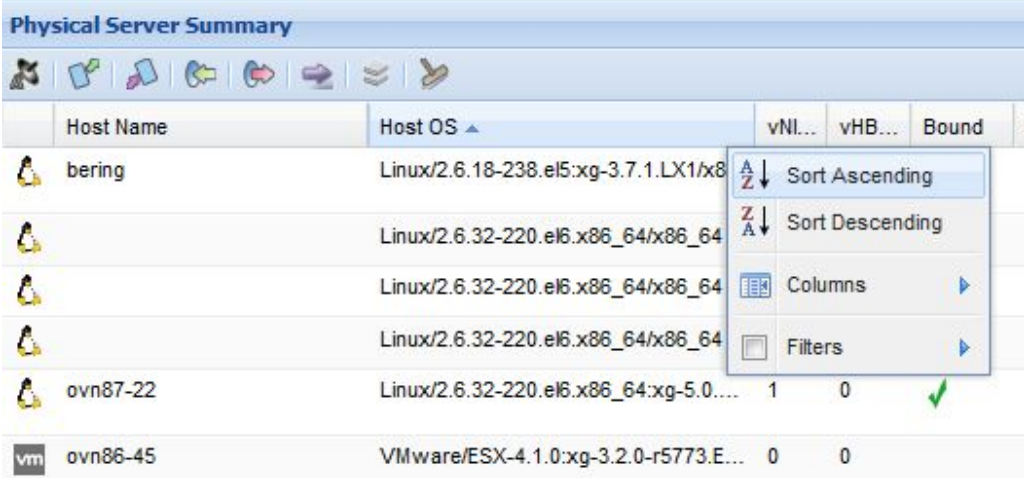
Host Name	Host OS	vNL...	vHB...	Bound	Busy	State	I/O Profile Name	Fabric Device Ports	Groups
bering	Linux/2.6.18-238.el5.xg-3.7.1.LX1/x8...	2	2	✓		up	LAMirrorSite1_607	delaware:ServerPort22 oregon:ServerPort15	
ovn87-36.us.oracle.com	Linux/2.6.32-220.el6.x86_64/x86_64	0	0					ExtSw-2c90200443198-Port8	
ovn87-34.us.oracle.com	Linux/2.6.32-220.el6.x86_64/x86_64	0	0					ExtSw-2c90200443198-Port10	
ovn87-35.us.oracle.com	Linux/2.6.32-220.el6.x86_64/x86_64	0	0					ExtSw-2c90200443198-Port9	
ovn87-22	Linux/2.6.32-220.el6.x86_64.xg-5.0...	1	0	✓		up	iotemplate	ExtSw-2c90200443198-Port7	HRServerGroup9 HRServerGroup
ovn86-45	VMware/ESX-4.1.0.xg-3.2.0-r5773.E...	0	0			offline			

On a table, you can sort and filter information to make displaying pertinent information much easier. Sorting and filtering occur on a per-column basis through a dropdown menu at the right edge of each column heading. To activate the sort and filter controls, click the heading of the column that you want to sort, and a downward arrow appears. The downward arrow is the dropdown menu where the sort and filter options are available.

Sorting

Sorting involves selecting a column and rearranging each entry on the table in either ascending or descending order. All table entries are arranged based on the column(s) that you select. For example, you can arrange the host OS versions installed on servers in ascending order to bring a specific host OS type to the top of the table, as shown in the following example.

Physical Server Summary



Host Name	Host OS	vNI...	vHB...	Bound
bering	Linux/2.6.18-238.el5.xg-3.7.1.LX1/x8			
	Linux/2.6.32-220.el6.x86_64/x86_64			
	Linux/2.6.32-220.el6.x86_64/x86_64			
ovn87-22	Linux/2.6.32-220.el6.x86_64/x86_64	1	0	✓
vm ovn86-45	VMware/ESX-4.1.0.xg-3.2.0-r5773.E...	0	0	

Filtering

Filtering involves either masking out columns in the table display, or entering criteria into a text-entry box and filtering out any table entries that do not use the criteria.

- To completely filter out a column in the table, click the column heading and select Columns. Then, click the resulting checkbox as shown in the following example. The presence of a checkmark indicates that the field is displayed. The checkbox is a toggle, so clicking the checkbox to remove the checkmark causes the selected field to be filtered. You can filter or display the column as needed by adding or removing the checkbox. In the following example, the Adapter FW Version field is filtered out of the table display.

Physical Server Summary

Host Name	Host OS	vNI...	vHB...	Bound	Busy	State	I/O F
bering	Linux/2.6.18-238.el5.xg-3.7.1.LX1/x8					up	LAM
ovn87-36.us.oracle.com	Linux/2.6.32-220.el6.x86_64/x86_64						
ovn87-34.us.oracle.com	Linux/2.6.32-220.el6.x86_64/x86_64						
ovn87-35.us.oracle.com	Linux/2.6.32-220.el6.x86_64/x86_64						
ovn87-22	Linux/2.6.32-220.el6.x86_64.xg-5.0....	1	0	✓			
vm ovn86-45	VMware/ESX-4.1.0.xg-3.2.0-r5773.E...	0	0				
vm ibsrv3.lab.xsigo.com	VMware/ESXi-5.1.1.ESX1M.1/x86_64	0	0				
CHARCOT-W2K8	Windows/6.1.7100/x64-2.2.0.36	1	0	✓			
COLEMAN	Windows/6.1.7600/x64-3.0.0.26	15	2	✓			
IBSRV1	Windows/6.1.7601/x64-5.0.0.46	0	0				

10 items

- To filter table contents by specific criteria, click the column heading and select Filters to display a text-entry box where you specify the filter criteria. Then, enter the filter criteria in the text-entry box as shown in the following example. Either press Enter or click the binoculars icon to start filtering. The following example shows an example of filtering the Adapter FW Version field based on version 2.6 as the criteria.

Physical Server Summary

Host Name	Host OS	vNI...	vHB...	Bound	Busy	State	I/O P
bering	Linux/2.6.18-238.el5.xg-3.7.1.LX1/x8					up	LAMi
ovn87-36.us.oracle.com	Linux/2.6.32-220.el6.x86_64/x86_64						
ovn87-34.us.oracle.com	Linux/2.6.32-220.el6.x86_64/x86_64						
ovn87-35.us.oracle.com	Linux/2.6.32-220.el6.x86_64/x86_64						
ovn87-22	Linux/2.6.32-220.el6.x86_64.xg-5.0...	1	0	✓		un	intern

2.6

- When the filter operation completes, the filtered table contents are displayed. Also, the checkbox to the left of the Filter menu option contains a checkmark. The checkmark indicates that the table contents displayed are the filtered contents.

Using the Maintenance Menu

This section describes the maintenance menu options.

- “Maintenance Menu” on page 77
- “Configure Backup Locations” on page 77
- “Back Up the Oracle Fabric Manager Configuration” on page 79
- “Back Up and Restore the Oracle Fabric Manager Configuration (Fresh Install)” on page 81
- “Restore the Oracle Fabric Manager Server Configuration” on page 82
- “Allow Unlisted User Access” on page 83
- “Download Oracle Fabric Manager Log Files” on page 84
- “Clean Up the Oracle Fabric Manager Database” on page 88

Maintenance Menu

The Maintenance Menu contains common tasks for managing the Oracle Fabric Manager Server and Oracle Fabric Manager GUI. These features tend to be outside the realm of actually configuring and managing host servers and their vNIC and vHBA connections. The Maintenance menu is a dropdown menu available on the Oracle Fabric Manager banner, and appears as a screwdriver icon.




▼ Configure Backup Locations

Through Oracle Fabric Manager, you can configure specific locations for Oracle Fabric Manager backup configs and Fabric Device backup configs. By default, Oracle Fabric Manager server configs are backed up to `/opt/xsigo/xms/xms-backups` and Fabric Device configs are backed up to `/opt/xsigo/xms/director-backups`. However, you can specify a non-default backup location for either config or both.

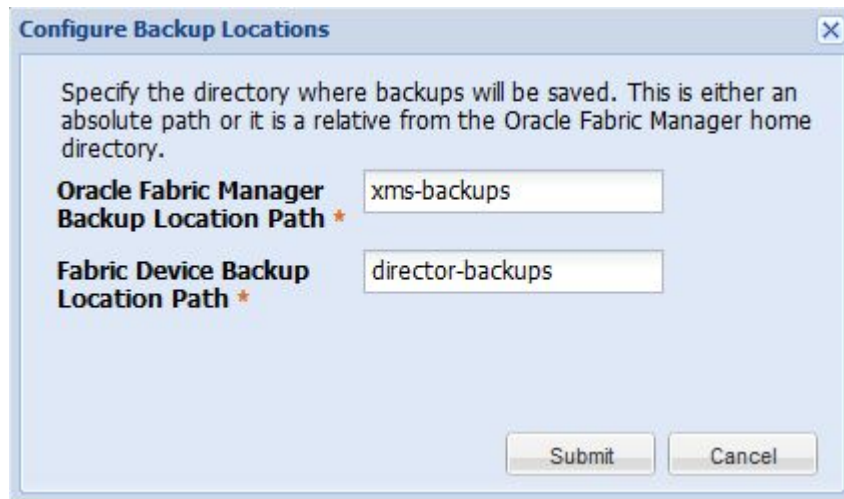
To back up the Oracle Fabric Manager configuration, follow this procedure:

1. **Click the Maintenance icon, which looks like a screwdriver on the banner.**

2. From the Maintenance Tasks dropdown menu, select the Configure Backup Locations options.



When you select this option, the Configure Backup Locations dialog is displayed.



3. In the Oracle Fabric Manager Backup Location Path field, enter the directory location as either an absolute or relative file path from the Tomcat home directory.
4. In the Fabric Interconnect Location Path field, enter the directory location as either an absolute or relative path from the Tomcat home directory.

5. **Click Submit.**

▼ **Back Up the Oracle Fabric Manager Configuration**

The current Oracle Fabric Manager configuration can be backed up on the Oracle Fabric Manager Server. When the configuration is backed up, the existing settings, virtual resources, and all other configured functionality in the Oracle Fabric Manager GUI are saved to a config file that you can name.

Be aware that there is a running a config and one or more backed up configs (if a Oracle Fabric Manager config has been saved). These two configurations can be different.

When the configuration is backed up, it is saved into `/opt/xsigo/xms/xms-backups`. It is not required that you provide a unique name for the configuration when you back it up. When a configuration is backed up, a time stamp is appended to the name that you give to the backed up configuration. For example, when you save the configuration as `xms-Config-Indy` for the Oracle Fabric Manager Server named “Indy”, the backup processes add a time stamp so that the file name might then become `Indy__2010-11-30_13_59_22_865`.

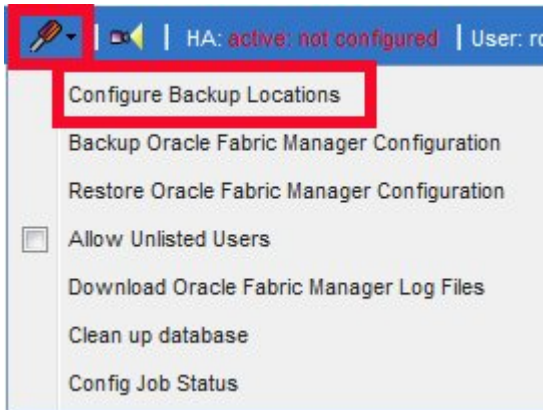
Backing up the Oracle Fabric Manager configuration occurs locally on the Oracle Fabric Manager Server, and no option exists to save the configuration to a central repository (for example, a LUN in the SAN). However, if you want to place the Oracle Fabric Manager configuration off of the Oracle Fabric Manager Server, you can simply back up the config on the Oracle Fabric Manager Server, then copy it off of the server to whichever location you desire.

Note - Backup the Oracle Fabric Manager Server configuration and the Fabric Device configuration at the same time. Backup the Oracle Fabric Manager Server, and when that operation is done, then back up the Fabric Device. Also, when restoring the backed up configuration, restore the Fabric Device configuration first, and when that operation is complete, then restore the Oracle Fabric Manager Server configuration.

To back up the Oracle Fabric Manager configuration, follow this procedure:

1. **Click the Maintenance icon, which looks like a screwdriver on the banner.**

2. **From the Maintenance Tasks dropdown menu, select the Configure Backup Locations options.**



When you select this option, the Configure Backup Locations dialog is displayed.

3. **In the Oracle Fabric Manager Backup Location Path field, enter the directory location as either an absolute or relative file path from the Tomcat home directory.**
4. **In the Fabric Interconnect Location Path field, enter the directory location as either an absolute or relative path from the Tomcat home directory.**
5. **Click Submit.**
6. **From Maintenance Tasks dropdown menu, select the Backup Oracle Fabric Manager Configuration option.**
7. **In the File Name field, enter an alphanumeric character string for the name that you want to give to the Oracle Fabric Manager configuration.**
Because multiple configs can be saved, Oracle Fabric Manager appends a timestamp to the file name that you enter.
8. **As an option, in the Description field, you can enter an alphanumeric character string that describes the config file that you are backing up.**
9. **When the Oracle Fabric Manager config file has been named, and any description provided, click Submit.**

At this point, the configuration is saved. You do not need to reset the server, or start a new Oracle Fabric Manager session.

▼ Back Up and Restore the Oracle Fabric Manager Configuration (Fresh Install)

Some cases might require the Oracle Fabric Manager software to be completely removed. Some examples are:

- A disk crashes and users want to keep the configuration
- A VM is rebuilt and users want to keep the configuration
- You completely uninstall a version of Oracle Fabric Manager and install a newer version, instead of upgrading

To successfully back up and restore the Oracle Fabric Manager configuration when you will be removing Oracle Fabric Manager, you will need to use a specific procedure. This restriction does not exist if you upgrade the Oracle Fabric Manager server. When the Oracle Fabric Manager Server is upgraded, the configuration database is kept, and so is the relevant record of a saved configuration. Therefore, the config database has knowledge of the saved config file, and can restore the config.

To perform a backup of an existing config prior to doing a fresh install of the Oracle Fabric Manager software, follow this procedure:

1. Perform a Oracle Fabric Manager backup.

This step places the Oracle Fabric Manager config as an XML file into the xms-backups directory.

2. Copy the XML file from the xms-backups directory to somewhere else—for example, to a network drive.

3. Remove Oracle Fabric Manager:

- On Oracle Solaris servers, use `pkgrm xms`
- On Linux servers, use `rpm -e xsigo-xms`
- On Windows servers, use **Start-Programs->Oracle Fabric Manager->Uninstall**

At this point, Oracle Fabric Manager is completely removed, including the xms-backups directory. However, since the saved config is in a safe location, it can be restored when Oracle Fabric Manager is re-installed.

4. Re-install Oracle Fabric Manager:

- On Oracle Solaris servers, use `pkgadd -d <Fabric-Manager-package-name>.pkg`

- **On Linux servers, use `rpm -ivh <Fabric-Manager-package-name>.rpm`**
- **On Windows servers, use the installer, either `install-xms-win32-version_OFM.exe` or `install-xms-win64-version_OFM.exe`.**

At this point, Oracle Fabric Manager is installed, including the `xms-backups` directory which is where the Oracle Fabric Manager restore feature looks for the backed up config file.

Note - For Windows servers, copy the Oracle Fabric Manager software to the server's hard drive, then running the installer.

5. **Copy the saved file from its current location (for example, from the network drive) to the new `xms-backups` directory.**
6. **Restore the Oracle Fabric Manager config from the backup by selecting the appropriate backup config file.**
See [“Restore the Oracle Fabric Manager Server Configuration”](#) on page 82.

▼ Restore the Oracle Fabric Manager Server Configuration

Through Oracle Fabric Manager you can restore any saved configuration. By restoring the configuration, you are reloading the saved configuration into Oracle Fabric Manager. Be aware that the restored configuration can be different than the current running configuration, and Oracle Fabric Manager does not automatically take a snapshot or save the current configuration before restoring the save configuration. As a result, any pending configuration (for example, I/O Template changes in progress on the I/O Template Editor) or any differences between the current Oracle Fabric Manager configuration and the config that will be restored are not saved.


When you restore the Oracle Fabric Manager configuration, Oracle Fabric Manager reads the contents of `/opt/xsigo/xms/xms-backups`, and displays a list of available configurations to load. When selected, that configuration is loaded instantly, without needing to reboot the Oracle Fabric Manager server. If multiple Oracle Fabric Manager configurations exist with unique names, you can select whichever Oracle Fabric Manager configuration you want to load.

Restoring the Oracle Fabric Manager configuration occurs locally from the `/opt/xsigo/xms/xms-backups` directory. No option exists to restore a Oracle Fabric Manager configuration that is stored in a central repository (for example, a LUN in the SAN). However, if the Oracle Fabric Manager configuration you want to restore is off of the Oracle Fabric Manager server itself, you can simply copy the Oracle Fabric Manager configuration you want into `/opt/xsigo/xms/xms-backups`, then restore the configuration from that directory.

To restore a Oracle Fabric Manager config, follow this procedure:

1. **Click the Maintenance icon, which looks like a screwdriver on the banner.**

From Maintenance Tasks dropdown menu, select the Restore Oracle Fabric Manager Configuration option.



When you select the restore option, a confirmation dialog is displayed that prompts you for verification that you actually intend to restore the Oracle Fabric Manager configuration.

2. **Read the confirmation dialog carefully, then select Yes to restore the Oracle Fabric Manager configuration.**

When you select Yes, the Oracle Fabric Manager Restore Configuration dialog is displayed.

3. **On the Oracle Fabric Manager Restore dialog, select the backup file that you want to restore.**
4. **Click Submit to load that Oracle Fabric Manager configuration.**

▼ Allow Unlisted User Access


Oracle Fabric Manager user accounts are created in the Oracle Solaris, Linux, or Windows OS of the Oracle Fabric Manager Server. When the user passes authentication, that user then receives a role that is granted and enforced by Oracle Fabric Manager. If no specific Oracle Fabric Manager role is assigned, the user account defaults to the operator role (read-only mode).

However, for tighter data center security, Oracle provides the ability to completely block an unassigned user account instead of allowing it to default to the operator role. When a user account is blocked, the Oracle Fabric Manager Login dialog is displayed, but login will fail when the user name and password are specified.

The Unlisted Users feature is accessible through the Allow Unlisted Users checkbox on the Oracle Fabric Manager Maintenance Tasks menu. By default, unlisted users are allowed to access Oracle Fabric Manager in read-only mode, but you can change this behavior by toggling the checkbox.

To prevent unlisted users from accessing Oracle Fabric Manager, follow this procedure:

1. **Click the Oracle Fabric Manager Maintenance icon, which looks like a screwdriver on the banner.**



2. **Click the check box to allow or prevent unlisted users access to Oracle Fabric Manager as needed:**
 - **Place a check mark in the checkbox to allow unlisted users to access Oracle Fabric Manager with read-only rights. This setting is the default.**
 - **Remove the check mark from the checkbox to prevent unlisted users from accessing Oracle Fabric Manager at all. When this setting is applied, unlisted users are presented with the log in dialog, but will not be able to successfully log in to Oracle Fabric Manager.**

▼ Download Oracle Fabric Manager Log Files

Oracle Fabric Manager Log files are kept in the `/opt/xsigo/xms/logs` directory. Different log files are kept for different purposes including tomcat, catalina, installation, and error logs. Logs are either current or historical, which is indicated by the file name:

- A current file, which is the most recent log file for the Oracle Fabric Manager server. A current file is named with the format `<name>.log.<highest-number>`—for example, assume there are 10 log files. In this example, `xms.log.10` is the current Oracle Fabric Manager Server log file, `xms-schedule.log.10` is the current Oracle Fabric Manager schedule log file, and `director-infiltrator.log.10` is the current log file for a Fabric Device named "infiltrator".
- A historical file, which is a log file that has rotated from a current file to an older file that contains data that is no longer current, but has been kept. A historical file is named with the format `<name>.log.<lowest-number>`—for example, `xms.log.1` is the first historical file for the Oracle Fabric Manager Server, and `director-infiltrator.log.1` is the first historical log file for the Fabric Device named "infiltrator".

Through Oracle Fabric Manager, you can download the logs from the Oracle Fabric Manager Server to a local client (for example, a laptop or desktop workstation) so that the logs can be examined if needed. Logs are packaged as a ZIP archive before downloading to minimize the size of the download and the amount of disk space required on the client when the files complete their download. Because the files are zipped, you should ensure that you have enough disk space to fully unzip them, or depending on your compression program, you can selectively extract and unzip only certain files that you need. The choice is completely up to you.


Note - Depending on the number of files in the logs directory, and the number of Oracle Fabric Devices under management, gathering log files, zipping, them and downloading them to a client will take some time. Please be patient and allow the process to complete.

You have some choice in the log files that you want to download:

- All, which downloads the entire contents of `/opt/xsigo/xms/logs` as one large ZIP file. This option downloads current files and historical files.
- Recent Logs only, which downloads only the current files in `/opt/xsigo/xms/logs` as one large ZIP file.

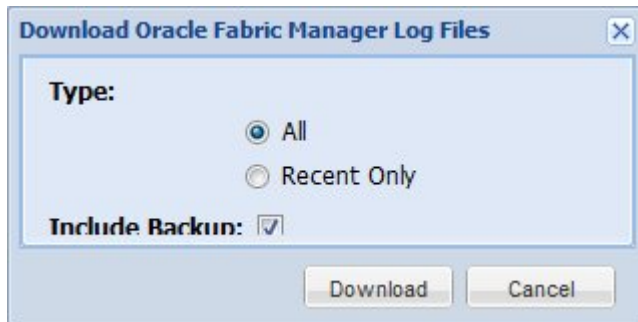
Downloading the Oracle Fabric Manager Log files is supported through a menu item on the Oracle Fabric Manager Maintenance dropdown menu. To download the Oracle Fabric Manager logs, follow this procedure:

1. On the Oracle Fabric Manager Maintenance menu, click **Download Oracle Fabric Manager Log Files**.



When you select Download Oracle Fabric Manager Files, a popup dialog that will enable you to choose the type of download you want, either Recent Files or All files.

2. On the **Download Oracle Fabric Manager Log Files** dialog, select either **Recent Only** or **All**.

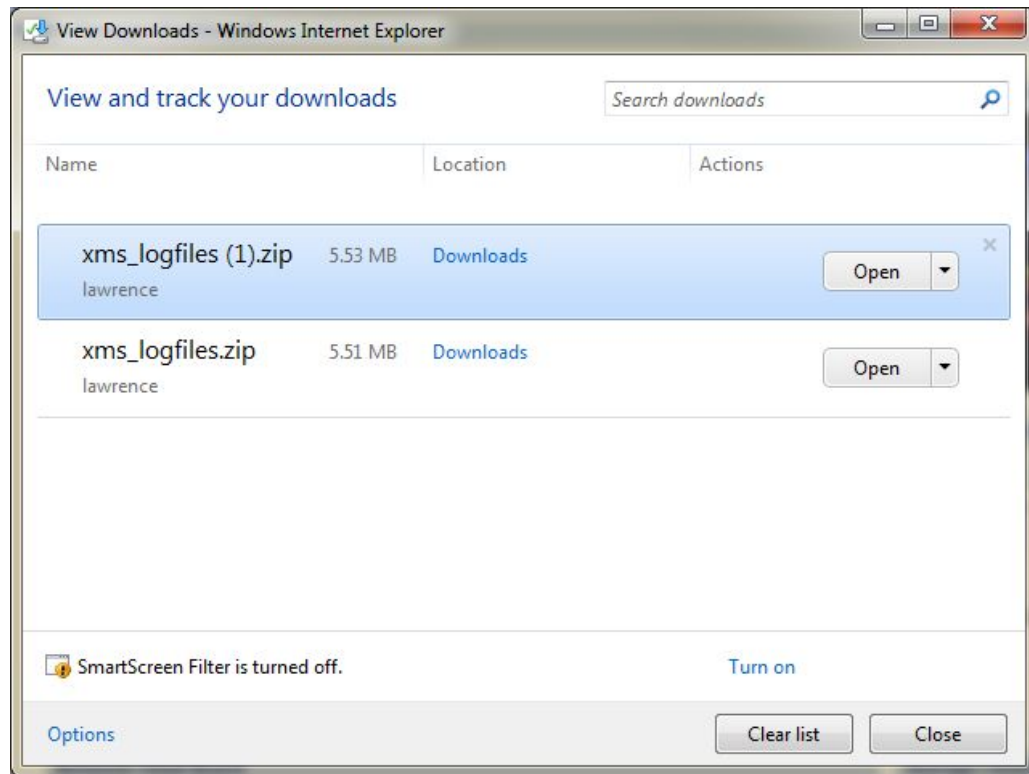


3. On the **Download Oracle Fabric Manager Log Files** dialog, select the appropriate button:
 - **Recent Only**
 - **All**

- Click **Include Backup** if you want to include the backup file.
4. **Click Download.**

A confirmation dialog is displayed by your browser to either save the files to a specific location or open them. Choose one of the following:

 - Select **Save** to download the files to a location that you select
 - Select **Open** to save the files to a default download location
 5. **When the files complete their download, you can unzip them and open the file(s) you want to see. The files are named xms_logs.zip.**



6. **When the files are on your Oracle Fabric Manager client, you can unzip them and use them as needed.**

▼ Clean Up the Oracle Fabric Manager Database

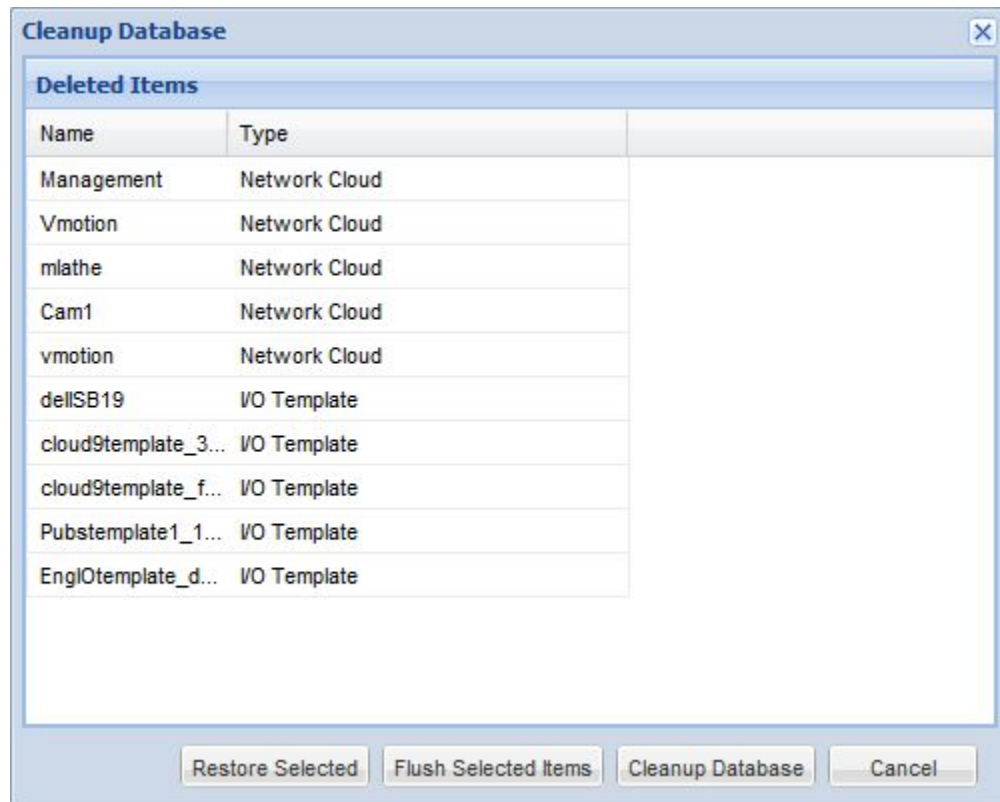
On each Oracle Fabric Manager Server, there is a database of objects that are under management. When upgrading between versions of Oracle Fabric Manager, new objects might be added, and older objects might not be used anymore. Because of upgrades, you are not completely removing the existing database and installing a fresh database (like you do when you perform a fresh install of Oracle Fabric Manager). As an option, Oracle Fabric Manager offers a way to clean up the database. When you clean up the Oracle Fabric Manager database, you have options to:

- Restore selected objects in the Oracle Fabric Manager database
- Flush selected objects from the Oracle Fabric Manager database
- Clean up the entire database

The first two options allow you to selectively administer specific items in the database, and the third option does a widespread cleanup of old items.

To clean up the Oracle Fabric Manager database, follow this procedure:

1. On the Oracle Fabric Manager Maintenance menu (the screwdriver icon on the Oracle Fabric Manager toolbar), click the Clean Up Database option to display the Cleanup Database dialog.



2. From the dialog, select one or more objects that you want to clean up in the database.
3. Click either of the following options:
 - **Restore Selected**, to restore an earlier version of the selected item in the current Oracle Fabric Manager database.
 - **Flush Selected Items**, to clear the selected items out of the database. This option is useful for uncluttering the database of old objects that are left over.
4. As an option, you can click **Cleanup Database** to have Oracle Fabric Manager programmatically cleanup everything that it can.

Working With Oracle Fabric Manager Jobs

This chapter contains the following topics:

- [“Understanding Jobs” on page 91](#)
- [“Displaying the Jobs Summary” on page 94](#)

Understanding Jobs

Oracle Fabric Manager is a client-server software bundle that enables you to configure server and virtual I/O needs based on an application's or network's needs.

- The Oracle Fabric Manager Server accepts configuration and management inputs from the Oracle Fabric Manager Client, then relays that information to Oracle Fabric Devices.
- Oracle Fabric Manager also retrieves information from Oracle Fabric Device(s), host servers, and other devices in the network and relays that information to the Oracle Fabric Manager Client.

Each time you click to submit a configuration or management task that has multiple steps, you create a job, which is the requested configuration action. Jobs are always asynchronous, but jobs can be either quick or long-running. All jobs in Oracle Fabric Manager are displayed in one of two locations:

- [“Jobs Summary” on page 91](#)
- [“Recent Jobs Summary” on page 93](#)

Jobs Summary

The Jobs summary shows a list of all jobs that have occurred recently. The Jobs Summary contains all the most recent jobs, to a maximum of the last 30 jobs. Jobs are displayed in the Jobs Summary regardless of whether they completed successfully or not.

Note - You can use the Jobs Summary to retrieve information about whether a feature is in the process of being configured, or if the job has stalled. Depending on the message displayed in the Jobs Summary, you might find it helpful to use this information if you have to contact Oracle Technical Support.

The Jobs Summary also contains sub-jobs which are an expandable option for each completed job in the Jobs Summary. By clicking the plus sign (+) to expand the primary job, you display the sub-jobs for that primary job. Through the sub-jobs you can see more granular information about step-by-step tasks that comprise the primary job.

The following is an example Jobs Summary. Notice that a few of the primary jobs have been expanded to display their sub-jobs.

Time Updated	Sub Jobs	Job Steps	Last Update	Username	Job Detail	State
2013-09-26 19:11:58.392	0	0		root	Backup Oracle Fabric Manager configuration to /opt/xsigo/xms/xms-backups/FabMan-A...	completed
2013-09-26 16:12:43.096	0	0	2013-09-26 16:12:43.153	root	Manage Backup OSDN Controller: ovn87-33	completed
2013-09-26 16:08:32.429	0	0	2013-09-26 16:08:32.504	root	Manage OSDN Controller: ovn87-31	completed
2013-09-26 16:07:24.098	1	0	2013-09-26 16:07:26.614	root	Unmanage Devices	completed
2013-09-26 16:06:51.082	0	0	2013-09-26 16:06:51.098	root	Manage Backup OSDN Controller: ovn87-33	failed
2013-09-26 16:06:17.029	0	0	2013-09-26 16:06:17.037	root	Manage OSDN Controller: aaaa	failed
2013-09-26 16:05:55.545	0	0	2013-09-26 16:05:55.553	root	Manage OSDN Controller: aaaaa	failed

Field	Indicates
Time Updated	The timestamp at which the job was completed. For primary jobs with sub-jobs, this field also indicates when each of the sub-jobs was attempted.
Last Update	The last timestamp that any activity occurred for the primary job.
Username	The name (if available) of the logged in user or administrator that initiated the job.
Job Detail	A string that provides additional details (if any) for the requested job and the object associated with that job. For example, the string "ApplyIOProfile 'pubstest' indicates that Oracle Fabric Manager is attempting to apply (bind) the I/O Profile named "pubstest" to one or more physical servers. The server name is also displayed if known.
State	The current state of the job, either complete, pending, or failed.
Detail Status	A string that summarizes the requested job and the object associated with that job. For example, the string "Successfully Applied IOProfile" indicates that Oracle Fabric Manager has successfully applied (bound) an I/O Profile to one or more physical servers.

The Jobs Summary lists jobs by the timestamp at which the job was started. Jobs are listed in sequential order with the most recent jobs at the top of the Job Summary. The Jobs Summary contains a maximum of 30 jobs.

The Jobs Summary contains a calendar tool which allows you to select a date or date range as a filter for the Jobs Summary's contents. By selecting a date or date range, you can display any jobs that occurred at that time. All alarms that do not match the date criteria are not displayed. When using the calendar tool, the filtered content should be displayed after you select the date or date range. However, you can always click the Refresh button to display the filtered content in the Jobs Summary.

Also, the Jobs Summary contains a keyword search function. The Keywords field enables additional filtering of the Jobs Summary by allowing you to enter specific words or phrases

to filter the contents of the Jobs Summary. For example, you could enter “cre” to filter the Jobs Summary to display all “create” jobs. The keywords you enter are matched against the characters in the Job Description field.


Recent Jobs Summary

On the page for each feature, you can see information about the 5 most recent jobs that occurred from any page in Oracle Fabric Manager. The Recent Jobs Summary is at the bottom of a feature or device's page.

Note - You can use the Recent Jobs Summary to retrieve information about whether a feature is in the process of being configured, or if the job has stalled. Depending on the message displayed in the Recent Jobs Summary, you might find it helpful to use this information if you have to contact Oracle Technical Support.

The Recent Jobs Summary is related to the Jobs Summary in that clearing the Jobs Summary also clears the Recent Jobs frame.

Recent Jobs Summary				
Time Updated	Job ID	State	Username	Job Detail
2012-05-24 21:47:43.304	remove objects	failed	root	remove objects
2012-05-23 18:49:26.499	Remove vnics	completed	root	Remove Vnics :pubsnic4 from IOProfile: coke
2012-05-23 18:47:22.31	AddVnic	failed	root	AddVnic (pubsnic4) to server coke
2012-05-21 13:30:32.831	RescanServers	completed	root	Rescan for new servers
2012-05-04 13:26:00.753	RescanServers	completed	root	Rescan for new servers

5 items 

Field	Indicates
Time Updated	The timestamp at which the job was either started, restarted, completed, or failed.
Job ID	A string that summarizes the requested job and the object associated with that job. For example, the string “ApplyIOProfile” indicates that Oracle Fabric Manager is attempting to apply (bind) an I/O Profile to one or more physical servers.
State	The current state of the job, either complete, pending, or failed.
User Name	The name (if available) of the logged in user or administrator that initiated the job.
Job Detail	A string that provides additional details (if any) for the requested job and the object associated with that job. For example, the string "ApplyIOProfile 'pubstest' indicates that Oracle Fabric Manager is attempting to apply (bind) the I/O Profile named "pubstest" to one or more physical servers. The server name(s) is also displayed if known.

Displaying the Jobs Summary

This section contains the following topics:

- [“Display Job Summary” on page 94](#)
- [“Show All Jobs” on page 94](#)
- [“Show All Active Jobs” on page 95](#)
- [“Clear All Jobs in the Jobs Summary” on page 96](#)
- [“Cancel Selected Jobs in the Jobs Summary” on page 97](#)

▼ Display Job Summary

The Jobs Summary contains a table of all the jobs that Oracle Fabric Manager has attempted regardless of their state, success, or failure. The Jobs Summary contains all jobs and lists summary information. Oracle Fabric Manager does not have a separate Details frame for each of the jobs in the Jobs Summary.

To display the Jobs Summary, follow this procedure:

- **On the Navigation Frame, select [General](#) → [Jobs Status](#) to display the Jobs Summary. See [“Jobs Summary” on page 91](#).**

The Jobs Summary supports the following options:

- [“Show All Jobs” on page 94](#)
- [“Show All Active Jobs” on page 95](#)
- [“Clear All Jobs in the Jobs Summary” on page 96](#)
- [“Cancel Selected Jobs in the Jobs Summary” on page 97](#)

▼ Show All Jobs

Jobs in the Jobs Summary are all jobs that have been attempted regardless of state. As a result, jobs that are in any state are displayed.

To show all jobs, follow this procedure:

- On the Navigation Frame, select **General → Jobs Status** to display the **Jobs Summary**.

Time Updated	Sub Jobs	Job Steps	Last Update	Username	Job Detail	State
2013-09-26 19:11:58.392	0	0		root	Backup Oracle Fabric Manager configuration to /opt/xsigo/xms/xms-backups/FabMan-A...	completed
2013-09-26 16:12:43.096	0	0	2013-09-26 16:12:43.153	root	Manage Backup OSDN Controller: ovn87-33	completed
2013-09-26 16:08:32.429	0	0	2013-09-26 16:08:32.504	root	Manage OSDN Controller: ovn87-31	completed
2013-09-26 16:07:24.098	1	0	2013-09-26 16:07:26.614	root	Unmanage Devices	completed
2013-09-26 16:06:51.082	0	0	2013-09-26 16:06:51.098	root	Manage Backup OSDN Controller: ovn87-33	failed
2013-09-26 16:06:17.029	0	0	2013-09-26 16:06:17.037	root	Manage OSDN Controller: aaaa	failed
2013-09-26 16:05:55.545	0	0	2013-09-26 16:05:55.553	root	Manage OSDN Controller: aaaaa	failed

▼ Show All Active Jobs

The Jobs Summary shows all jobs that have been attempted regardless of state. As a result, jobs that are in any state are displayed. You can filter the Jobs Summary to display all active jobs, which are jobs that are currently in progress. Be aware that jobs can take a while to complete based on the number of subjobs that are contained in the parent job. For example, if you attempt to unbind a Server Profile from a server, that job might contain individual subjobs for removing the vNICs and vHBAs that are connected to the host. While those subjobs are completing, the overall unbind job will still be active.

To filter the active jobs in the Jobs Summary, follow this procedure:

1. On the Navigation Frame, select **General → Jobs Status** to display the **Jobs Summary**.
2. On the Jobs Status summary, click the **Show All Active Jobs** button.

Time Updated	Sub Jobs	Job S
2013-09-26 19:11:58.392	0	0
2013-09-26 16:12:43.096	0	0
2013-09-26 16:08:32.429	0	0
2013-09-26 16:07:24.098	1	0

When you click the Show All Active Jobs button, the Jobs Status Summary is filtered to display only the jobs with “active” state. This button is a toggle, so clicking repeatedly switches between displaying all active jobs and all jobs (including non-active ones). You will notice that the button changes when it toggles between Show All Jobs and Show All Active Jobs.

▼ Clear All Jobs in the Jobs Summary

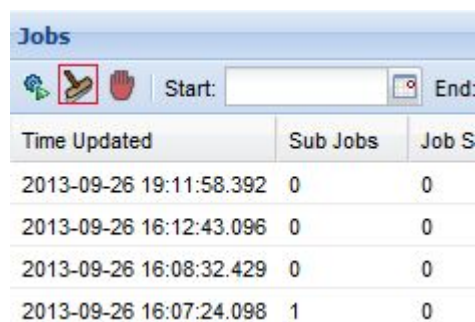
Jobs are persistent across reboot, and Oracle Fabric Manager service restarts:

- Jobs are kept in the Jobs Summary until they are explicitly batch removed.
- Jobs are not aged out, or removed when a pending job completes.
- Jobs can be cleared only by removing them as a batch by an administrator-level account.
- Individual jobs cannot be removed. Jobs must be explicitly removed as a batch by an administrator.

When you clear the jobs in the Jobs Summary, you are removing all completed or failed jobs only. Jobs that are in any other state are not cleared.

To clear all Jobs from the Jobs Summary, follow this procedure:

1. **On the Navigation Frame, select General → Jobs Status to display the Jobs Summary.**
2. **Click the eraser icon to remove all jobs from the Job Summary.**



Time Updated	Sub Jobs	Job S
2013-09-26 19:11:58.392	0	0
2013-09-26 16:12:43.096	0	0
2013-09-26 16:08:32.429	0	0
2013-09-26 16:07:24.098	1	0

The Job Summary will be repopulated when the next administration or configuration tasks are submitted.

▼ Cancel Selected Jobs in the Jobs Summary

Some jobs in Oracle Fabric Manager can be long running. For example, applying an I/O Profile to many servers can take a significant amount of time. While a job is running, you can cancel it before completion, regardless of what stage the job is in.

Through the Jobs Summary, you can cancel any of the jobs that are still in process. By cancelling a job, you cause it to completely stop. For example, if you apply an I/O Profile to a server, and the job shows as running or pending in the Job Summary, you can cancel the job to stop the I/O Profile from being applied to the server. Cancelling a job simply stops it. The job is not cleaned up and removed from the Jobs Summary.

Unlike the Undo Jobs function, there is no limit on the type of job that can be cancelled. The only requirement for cancelling a job is that the job state is either running or pending.

To cancel a job, follow this procedure:

1. **On the Navigation Frame, select General → Jobs Status to display the Jobs Summary.**

2. **Select a running or pending job that you want to cancel.**

You can check the State column to verify that the job you want to cancel is either running or pending. When you select a valid, cancellable job, the Cancel Jobs button becomes active.

3. **When a valid job is selected, click the Cancel Jobs button.**

4. **On the confirmation dialog, click Yes to proceed with the Undo function.**

Allow the undo job to run to completion, which might take some time.

While the job is being undone, the text of the selected job will turn green to indicate the success of the undo.

5. **When the undo is complete, check the State column in the Jobs Summary again for a message about the undo function.**

Working With Network Clouds

This chapter contains the following topics:

- [“Understanding Network Clouds” on page 99](#)
- [“Creating a Network Cloud” on page 101](#)
- [“Configuring Advanced Features for a Network Cloud” on page 107](#)
- [“Set Link Aggregation for a Network Cloud” on page 115](#)
- [“Managing Network Cloud Properties” on page 117](#)

Understanding Network Clouds

This section contains the following topics:

- [“Network Clouds Overview” on page 99](#)
- [“Network QoS Profiles and Network Clouds” on page 100](#)
- [“QoS Assignment at Cloud Level and vNIC Level” on page 100](#)

Network Clouds Overview

Oracle Fabric Devices contain a set of I/O ports that are connected to your data center network. The network and SAN administrator connect these ports to their Ethernet switches in order to provide I/O resources to the servers.

Each I/O port on the Oracle system can be thought of as providing access to a set of “clouds.” For example, you might have a Network Cloud that provides access to the HR network. As far as the network administrator is concerned, you are connecting wires to the Oracle I/O ports for the sole purpose of providing the server administrator access to the resources.

The server administrator then needs to connect servers to a set of network resources provided by the network administrators. The server administrator doesn't care what physical Oracle ports are being connected to as long as the servers have access to the required resources.

Network Clouds can be configured and managed by Network and Administrator roles in Oracle Fabric Manager.

Network QoS Profiles and Network Clouds

If you want a specific Network QoS Profile assigned to the Network Cloud, that specific profile must already exist. If it does not exist, it will not be assignable to the Network Cloud at the time you create the cloud.

As an alternative, Oracle Fabric Manager has some commonly used network QoS Profiles that are pre-configured for various network link speeds. These pre-configured Network QoS Profiles use typical CIR and PIR values calculated from the link's total throughput. If you are using pre-configured Network QoS Profiles, make sure that you use a reasonable profile. For example, for a 1 Gigabit Ethernet link, do not select a network QoS profile for a 10 Gigabit Ethernet link (for example, the Network QoS profile named 7g_10g which has throughput values that are in excess of what the 1 Gigabit link can support).

Network QoS can be configured for a Network Cloud at any of the following times:

- During Network Cloud creation, through the Create Network Cloud wizard.
- After the Network Cloud has already been created, through the Network Cloud Details frame.
- Network QoS can also be changed for a Network Cloud after the Network Cloud has already been created. Changing an assigned Network QoS Profile occurs through the Network Cloud Details frame.

QoS Assignment at Cloud Level and vNIC Level

It is possible to assign Network QoS to a Network Cloud or a vNIC. The following text explains what QoS is applied where:

- If a Network Cloud has a QoS Profile associated with it, a vNIC terminated in the Network Cloud will inherit the Network QoS profile from the Network Cloud.
- In this case, the vNICs inherit the QoS parameters from their respective clouds. This condition is true even if the vNIC has its own QoS Profile associated with it.
- If a Network Cloud has no QoS Profile associated with it, and the vNICs have their own QoS profiles, a vNIC terminated in the Network Cloud will have QoS Profile associated with it.
- In this case, the Network Cloud has no QoS Profile to apply. This situation does not automatically apply no QoS to the vNICs. Instead, it causes the cloud to defer to the QoS profiles assigned to the individual vNICs (if any).
- If you intend for no QoS to be applied anywhere for the virtual I/O in a cloud, make sure that no QoS Profile is associated with the Network Cloud, and also that no QoS is associated with the vNIC.

Creating a Network Cloud

This section contains the following topics:

- [“Network Cloud Creation Overview” on page 101](#)
- [“Create a Network Cloud” on page 101](#)
- [“Fabric Device HA Designation” on page 104](#)
- [“Set the Fabric Device HA Designation for a Network Cloud” on page 104](#)

Network Cloud Creation Overview

When you create a Network Cloud, you assign one or more Ethernet ports or LAGs, plus other network characteristics. When the Network Cloud is attached to an I/O Template, the Network Cloud provides a way for vNICs to connect hosts to the data network when an I/O Profile is deployed to the host. The Network Cloud provides needed network access points with specified characteristics—for example, Network QoS profiles, VLANs, and so on.

A default Network Cloud exists, which contains all the discovered Ethernet ports available to Oracle Fabric Manager through the discovered Fabric Device(s). This default Network Cloud is a way to show you the total number of termination points that are under management by Oracle Fabric Manager. The total number shown does not decrement when ports are assigned to a specific Network Cloud, since those ports have not been removed from Oracle Fabric Manager—they are just reassigned to a different object.

Network Clouds typically are created around some similar theme for the ports or LAGs that are associated with it. For example, a Network Cloud might have ports or LAGs that will provide vNIC connections to a specific domain or specific upstream Ethernet switch, for a specific set of host servers.






Creating a Network Cloud occurs through the Network Cloud option in the Network Resources Manager on the navigation panel.


[“Create a Network Cloud” on page 101](#) shows creating a Network Cloud for a single Fabric Device. However, in most cases, Fabric Devices are deployed as redundant pairs, which has a minor difference in the properties you will set for the Network Cloud. For more information about deploying Fabric Device pairs, see [“Set the Fabric Device HA Designation for a Network Cloud” on page 104](#).

▼ Create a Network Cloud

To create a Network Cloud, follow this procedure:

1. On the Navigation panel, select Network Cloud Manager → Network Clouds to display the Network Cloud Summary.

Network Cloud Summary							
 							
	Name ▲	Number of Ports	Number of Lags	QoS	Number of vNL...	Number of vNL...	Description
	de	1	0	none	0	0	
	discovered-network-cloud	40	3	none	20	0	
	techpubs9	2	0	none	0	0	

3 items 

Notice the default Network Cloud (discovered-network-cloud) in this example. This example also shows that specific Network Clouds (non-default clouds) are configured.

Adding and deleting Network Clouds occurs through the plus sign and garbage can icons, respectively.

- Click the plus sign (+) to display the Add a New Network Cloud dialog.

New Network cloud

Name: * pubstest10

Description:

Ethernet Ports /LAGs: *

Name	Type	State	vNICs	Capa...	Trunk...	vLAN...	vLAN Range	Description
delaware/10/1	nwEthernet...	up/down	0	10G	access	1		
delaware/14/1	nwEthernet...	up/up	13	10G	access	1		
delaware/7/1	nwEthernet...	up/down	0	10G	access	1		
delaware/8.1	LAG	up/down	0		access	1		
delaware/8/2	nwEthernet...	up/down	1	1G	access	1		
delaware/8/3	nwEthernet...	up/down	0	1G	access	1		
delaware/8/4	nwEthernet...	up/down	0	1G	access	1		
delaware/8/5	nwEthernet...	up/down	0	1G	access	1		

43 items

HA Designation: Check To set primary and secondary Fabric Interconnects

Advanced Configuration

Submit Cancel

- In the Name field, enter a name for the Network Cloud that you are creating.
- As an option, in the Description field, enter a description for the Network Cloud that you are creating.
- From the Ethernet Ports/LAGs table, select the port(s) or LAGs that will be used in this Network Cloud. Ports and LAGs are selected when they are highlighted (as shown). Multiple ports and LAGs can be assigned to the same Network Cloud, and multiple Network Clouds can be assigned to the same port(s).
- If you are deploying two Fabric Devices as a redundant pair, you can use the HA Designation checkbox to specify which Fabric Device is the primary and which Fabric Device is the secondary.
For more information about deploying Fabric Device pairs, see [“Set the Fabric Device HA Designation for a Network Cloud”](#) on page 104.
- When the correct properties have been assigned for the Network Cloud, click Submit to create the Network Cloud.

8. **Check the Network Cloud Summary to verify that the Network Cloud was correctly configured.**

Fabric Device HA Designation

In a typical deployment, redundant Fabric Devices are present, and the Network Cloud will be associated with both Fabric Devices. When you create a Network Cloud, Oracle Fabric Manager supports a simple and easy way to associate the Network Cloud with two Fabric Devices simultaneously, and also set the connection priority so that one Fabric Device is considered the primary and the other is considered the secondary. When you initially create the cloud, any HA vNICs must be manually connected to the cloud and their primary and secondary status must be manually set. However, after the Network Cloud is initially created, the Fabric Device connection priority is set. Any additional HA pairs that attach to the Network Cloud after the HA designation is made will first attach to the primary Fabric Device and then to the Secondary Fabric Device.

The HA Fabric Device designation is configured when you create the Network Cloud, and is available on the Create Network Cloud dialog.

When a host is discovered or connected to a Fabric Device that is discovered, Oracle Fabric Manager will query the Oracle host drivers on the host for the capability to support HA vNICs. This reporting facilitates the creation of HA vNICs on hosts because some hosts do not support true HA vNICs. For example, ESX servers support NIC Teaming, which requires two individual vNICs instead of one HA vNIC.

- If the host driver reports that HA vNICs are available on the host, a Oracle Fabric Manager HA vNIC will create HA vNICs on the server.
- If the host driver reports that HA vNICs are not available on the host, a Oracle Fabric Manager HA vNIC will create a pair of single vNICs on the server.

▼ Set the Fabric Device HA Designation for a Network Cloud

To create a Network Cloud with HA, follow this procedure:

1. **Display the Network Cloud Summary (Network Cloud Manager – Network Clouds).**

2. On the Network Cloud Summary, click the plus sign (+) to display the New Network Cloud dialog.

New Network cloud

Name: * pubstest9

Description:

Ethernet Ports / LAGs: *

Name	Type	State	vNICs	Capa...	Trunk...	vLAN...	vLAN Range	Description
delaware/9/9	nwEthernet...	up/down	0	1G	access	1		
oregon/10/1	nwEthernet...	up/down	0	10G	access	1		
oregon/14/1	nwEthernet...	up/down	0	10G	access	1		
oregon/7/1	nwEthernet...	up/down	0		access	1		
oregon/8.1	LAG	up/down	0		access	1		
oregon/8/2	nwEthernet...	up/down	5	1G	access	1		
oregon/8/3	nwEthernet...	up/down	2		access	1		
oregon/8/4	nwEthernet...	up/down	0		access	1		

43 items

HA Designation: Check To set primary and secondary Fabric Interconnects

Primary Fabric Interconnect: delaware

Secondary Fabric Interconnect: oregon

Advanced Configuration

Submit Cancel

3. In the Name field, enter a name for the Network Cloud that you are creating.
4. As an option, in the Description field, enter a description for the Network Cloud that you are creating.
5. From the Ethernet Ports/LAGs table, select at least one port or LAG from each Fabric Device that will be used in the redundant Fabric Device Network Cloud.

When one port from two different Fabric Devices are selected, the HA Designation checkbox at the bottom of the dialog becomes active.

6. In the HA Designation checkbox, click to specify that this is an HA Fabric Device configuration, and also to display the Primary Fabric Interconnect and Secondary Fabric Interconnect dropdown menus.

Name: * pubstest9

Description:

Ethernet Ports / LAGs: *

Name	Type	State	vNICs	Capa...	Trunk...	vLAN...	vLAN Range	Description
delaware/9/9	nwEthernet...	up/down	0	1G	access	1		
oregon/10/1	nwEthernet...	up/down	0	10G	access	1		
oregon/14/1	nwEthernet...	up/down	0	10G	access	1		
oregon/7/1	nwEthernet...	up/down	0		access	1		
oregon/8.1	LAG	up/down	0		access	1		
oregon/8/2	nwEthernet...	up/down	5	1G	access	1		
oregon/8/3	nwEthernet...	up/down	2		access	1		
oregon/8/4	nwEthernet...	up/down	0		access	1		

43 items

HA Designation: Check To set primary and secondary Fabric Interconnects

Primary Fabric Interconnect: delaware

Secondary Fabric Interconnect: oregon

Advanced Configuration

Submit Cancel

7. From the Primary Fabric Interconnect dropdown menu, select which of the two Fabric Devices will be the primary.
The primary is the Fabric Device to which the primary vNIC in an HA pair is connected first when new HA vNICs are connected to the Network Cloud.
8. From the Secondary Fabric Interconnect dropdown menu, select which of the two Fabric Devices will be the secondary.
The secondary is the Fabric Device to which the secondary vNIC in an HA Pair is connected when new HA vNICs are connected to the Network Cloud.
9. If you want to configure additional properties for the cloud (for example, a VLAN ID or network QoS Profile) continue to [“Configuring Advanced Features for a Network Cloud” on page 107.](#)

Otherwise, you can complete the configuration of a basic network Cloud by proceeding to [Step 10](#).

10. **When the correct properties have been assigned for the Network Cloud, click Submit to create the Network Cloud.**
11. **Check the Network Cloud Summary to verify that the Network Cloud was correctly configured.**

Configuring Advanced Features for a Network Cloud

This section contains the following topics:

- [“Network Cloud Advanced Features” on page 107](#)
- [“Set Network QoS on a Network Cloud” on page 107](#)
- [“Set a VLAN for a Network Cloud” on page 109](#)
- [“Private vNICs Overview” on page 111](#)
- [“Create Private vNICs on a Network Cloud” on page 112](#)
- [“Set the Port Allocation Policy” on page 113](#)

Network Cloud Advanced Features

Because Network Clouds contain Ethernet ports or LAGs, many standard features that are applicable to network ports and LAGs (for example, Network QoS) are configurable on Network Clouds. Network Clouds support configuration of the following features:

- Network QoS CIR and PIR values. (CBS and PBS values are calculated internally from the CIR and PIR values). The QoS parameters are associated with vNICs in the cloud by either having the cloud assign the Network QoS parameters, or by allowing the vNIC to have its own QoS parameters if the cloud itself has no Network QoS associated with it.
- VLAN properties. Any vNIC VLAN properties are inherited from the Network Cloud with which the vNIC is associated.
- Trunk Mode vNICs
- Private vNICs
- Port Allocation Policy in the Network Cloud

▼ Set Network QoS on a Network Cloud

Setting Network QoS on an Oracle Network Cloud can occur at the cloud level or at the vNIC level. When Network QoS is applied to a Network Cloud, all vNICs associated with the cloud

will get the specified CIR and PIR values. If a vNIC also has Network QoS and is added to the Network Cloud, the cloud's Network QoS takes precedence. For more information about Network QoS and Network Clouds, see [“Working With Network QoS” on page 291](#).

The following procedure assumes that the Network QoS Profile already exists. If it does not, you will need to assign one of the pre-defined Network QoS Profiles that Oracle has provided.

To configure Network QoS for a Network Cloud, follow this procedure:

1. **If the Create a Network Cloud dialog is not already displayed, display it now.**
2. **Click the Advanced Configuration dropdown menu to display the additional properties for the Network Cloud.**

New Network cloud

Name: * pubstest10

Description:

Ethernet Ports /LAGs: *

Name	Type	State	vNICs	Capa...	Trunk...	vLAN...	vLAN Range	Description
delaware/10/1	nwEthernet...	up/down	0	10G	access	1		
delaware/14/1	nwEthernet...	up/up	13	10G	access	1		
delaware/7/1	nwEthernet...	up/down	0	10G	access	1		
delaware/8.1	LAG	up/down	0		access	1		
delaware/8/2	nwEthernet...	up/down	1	1G	access	1		
delaware/8/3	nwEthernet...	up/down	0	1G	access	1		
delaware/8/4	nwEthernet...	up/down	0	1G	access	1		
delaware/8/5	nwEthernet...	up/down	0	1G	access	1		

43 items

HA Designation: Check To set primary and secondary Fabric Interconnects

Advanced Configuration

QoS: Select a QoS Profile...

Access VLAN ID:

Trunk Mode:

Private:

Allocation Policy: Round Robin

3. **From the QoS dropdown menu, select the new Network QoS Profile for the selected Network Cloud.**

4. **Click Submit to assign the new Network QoS Profile to the selected Network Cloud.**

You will be prompted with a confirmation message.

5. **Click Yes to assign the Network QoS profile or No to abort assigning the Network QoS Profile to the Network Cloud.**

▼ Set a VLAN for a Network Cloud

VLANs are used for traffic isolation and security to prevent some hosts on a network from seeing traffic that is intended for other hosts. When traffic is tagged with a VLAN ID, only the hosts that can transmit or receive packets for that VLAN are able to see and use that traffic.

VLANs can be set for a Network Cloud. When a VLAN is set on the Network Cloud, vNICs added to that cloud inherit the VLAN ID unless additional configuration exists at the port level to enforce different tagging rules.

Note - In addition to VLANs on a Network Cloud, you can control the VLAN IDs that are supported on a specific Fabric Device through the Allowed VLAN Range feature. For more information, see [“Set the Allowed VLAN Ranges for an Ethernet Port” on page 165](#).

VLANs can be set for a Network Cloud through the Advanced Configuration dropdown menu on the New Network Cloud dialog. When you set a VLAN on the Network Cloud, it is applied to all ports in the Network Cloud. As a result, when a vNIC is terminated on that Network Cloud it will be assigned to the specified VLAN.

To set a VLAN for a Network Cloud, follow this procedure:

1. **If the Create a Network Cloud dialog is not already displayed, display it now.**

2. Click the Advanced Configuration dropdown menu to display the additional properties for the Network Cloud.

Name: * pubstest10

Description:

Ethernet Ports / LAGs: *

Name	Type	State	vNICs	Capa...	Trunk...	vLAN...	vLAN Range	Description
delaware/10/1	nwEthernet...	up/down	0	10G	access	1		
delaware/14/1	nwEthernet...	up/up	13	10G	access	1		
delaware/7/1	nwEthernet...	up/down	0	10G	access	1		
delaware/8.1	LAG	up/down	0		access	1		
delaware/8/2	nwEthernet...	up/down	1	1G	access	1		
delaware/8/3	nwEthernet...	up/down	0	1G	access	1		
delaware/8/4	nwEthernet...	up/down	0	1G	access	1		
delaware/8/5	nwEthernet...	up/down	0	1G	access	1		

43 items

HA Designation: Check To set primary and secondary Fabric Interconnects

Advanced Configuration

QoS: Select a QoS Profile...

Access VLAN ID:

Trunk Mode:

Private:

Allocation Policy: Round Robin

3. In the VLAN ID field, set the specific VLAN number which will be applied to the vNICs that are associated with the cloud.

Note - The VLAN ID you set for the Network Cloud must be within the Allowed VLAN Range set for the Fabric Device on which the vNIC(s) are terminated.

4. If the Network Cloud will be supporting vNICs that trunk the VLAN tags, click the Trunk Mode checkbox.
5. When the VLAN properties are configured, and any other properties for the Network Cloud, click Submit to complete the configuration.

Private vNICs Overview

By default, when two vNICs are terminated on the same I/O module, Oracle Fabric Device uses a feature called vNIC-to-vNIC Switching which forwards the traffic between the two vNICs across the I/O module, instead of sending them over the midplane. The vNIC-to-vNIC Switching feature provides better performance for network traffic. However, some situations exist where uncontrolled vNIC-to-vNIC Switching might not be desirable due to security reasons.

Oracle Fabric Manager provides a way to add security to vNIC-to-vNIC switching, through a feature called Private vNICs. With Private vNICs, you can add isolation for a set of vNICs, and also provide enhanced compatibility for existing and new methods of external switching. If a vNIC-to-vNIC Switching configuration does not explicitly have one or more vNICs set as a Private vNIC, then those vNICs are assumed to be Public vNICs, which do not support isolation and external switching benefits that Private vNICs do. Specifically, Private vNICs provide the following isolation and switching enhancements:

- A private vNIC can communicate with the Ethernet network and Public vNICs, but it cannot communicate with other Private vNICs.
- If a packet's destination is assigned to a Private vNIC, the packet will be dropped. Packets originating on a private vNIC and destined for another private vNIC will be dropped regardless of whether both vNICs are in the same VLAN.
- Public vNICs can communicate without restriction with the exception of traffic isolation enforced by standard VLANs.
- Packets arriving from private vNICs are dropped or forwarded based on the destination. If the destination is a Public vNIC or the Ethernet network, the packet will be forwarded. Otherwise, it will be dropped. VLAN restrictions will still apply to packets that are forwarded.
- Broadcast or multicast packets are not a special case. They are also not forwarded between Private vNICs.

Note - Private vNICs are not supported with HA vNICs. Configure your Private vNICs as standalone vNICs only. Public vNICs, however, can be configured as HA vNICs.

Private vNICs can be specified through the Private checkbox, which toggles the Private property on (Private vNIC) or off (Public vNIC). The Private checkbox is available through the Network Cloud so that all vNICs that attach to the cloud are considered private.

This feature is also settable at the I/O Template level and the vNIC level:

- To set one or more Private vNICs at the I/O Template level, double-click the vNIC on the I/O Template Editor and click the Private checkbox as needed. However, any setting configured on an individual vNIC (either Public or Private) will be overridden by the setting

at the Network Cloud level. For example, if the vNIC is configured as a Public vNIC, then you move the vNIC to a Network Cloud that is set for Private vNICs, the vNIC will become a Private vNIC after the move completes and the vNIC is terminated on a port in the new Network Cloud.

- To set one or more Private vNICs at the vNIC-level, when the vNIC is deployed to the server, select the server on the Physical Servers page to populate the details frame. Then, click the vNICs tab to display the server's vNICs. Then, click the name of the individual vNICs that you want to make into Private vNICs. When the vNIC is displayed, click the General tab, then click the Edit button and mark the Private checkbox.

▼ Create Private vNICs on a Network Cloud

To configure a Network Cloud with Private vNICs, follow this procedure:

1. **If the Create a Network Cloud dialog is not already displayed, display it now.**
2. **Click the plus sign (+) to display the New Network Cloud dialog.**

3. Click the **Advanced Configuration** dropdown menu to display the additional properties for the Network Cloud.

The screenshot shows the 'New Network cloud' configuration window. The 'Name' field is set to 'pubstest10'. Below it is a 'Description' text area. A table lists 'Ethernet Ports / LAGs' with columns for Name, Type, State, vNICs, Capacity, Trunk Mode, vLAN, vLAN Range, and Description. The table contains 9 rows of data. Below the table is an 'HA Designation' checkbox labeled 'Check To set primary and secondary Fabric Interconnects'. The 'Advanced Configuration' section is highlighted with a red box and contains the following fields:

- QoS:** A dropdown menu with the text 'Select a QoS Profile...'.
- Access VLAN ID:** A text input field.
- Trunk Mode:** A checkbox.
- Private:** A checkbox.
- Allocation Policy:** A dropdown menu with 'Round Robin' selected.

At the bottom of the window are 'Submit' and 'Cancel' buttons.

4. In the **Private** checkbox, click to specify that vNICs terminated on this network cloud are private vNICs.
When configured, any vNIC that is terminated on this cloud will not be able to communicate with any public vNICs in any other Network Clouds.
5. When the cloud's properties are specified, click **Submit** to complete the configuration.

▼ Set the Port Allocation Policy

For Network and Storage Clouds, you can set the manner in which ports are allocated to vNICs or vHBAs that connect to the clouds.

To set the port allocation policy for a Network Cloud, follow this procedure:

1. **If the Create a Network Cloud dialog is not already displayed, display it now.**
2. **Click the Advanced Configuration dropdown menu to display the additional properties for the Network Cloud.**

New Network cloud

Name: * pubstest10

Description:

Ethernet Ports / LAGs: *

Name	Type	State	vNICs	Capa...	Trunk...	vLAN...	vLAN Range	Description
delaware/10/1	nwEthernet...	up/down	0	10G	access	1		
delaware/14/1	nwEthernet...	up/up	13	10G	access	1		
delaware/7/1	nwEthernet...	up/down	0	10G	access	1		
delaware/8.1	LAG	up/down	0		access	1		
delaware/8/2	nwEthernet...	up/down	1	1G	access	1		
delaware/8/3	nwEthernet...	up/down	0	1G	access	1		
delaware/8/4	nwEthernet...	up/down	0	1G	access	1		
delaware/8/5	nwEthernet...	up/down	0	1G	access	1		

43 items

HA Designation: Check To set primary and secondary Fabric Interconnects

Advanced Configuration

QoS: Select a QoS Profile... **Access VLAN ID:**

Trunk Mode: **Private:**

Allocation Policy: Round Robin

Submit Cancel

3. **From the Allocation Policy field, select the following policy for how network ports are assigned to vNICs that terminate on the cloud.**
 - **Round Robin, which is a systematic way for ports to be assigned. With round robin allocation, you specify multiple ports in a list of available ports for a Network Cloud. When the list is constructed, you then assign a port as the next available port. When you assign the next available port, it gets a rank that is different than any other port in the available ports list. The port you assign is the next port to be assigned, and after that, any additional**

vHBAs that are connected to the Network Cloud will receive an Ethernet port based on its rank. For example, assume you create a Network Cloud with Ethernet port 2/1 as the port used to create the cloud. Then, assume you add ports 2/2, 4/1, and 5/1 to the available ports list, and set port 4/1 as the next available port. When new vNICs are connected to the Network Cloud, the first vNIC is connected to port 4/1. After that, additional vNICs are assigned to whichever port has the next lowest rank.

4. When the allocation policy is configured, and any other properties for the Network Cloud, click **Submit** to complete the configuration.

▼ Set Link Aggregation for a Network Cloud

Link Aggregation Groups (LAGs) function the same as physical Ethernet port in the Network Cloud. Virtual NICs can be terminated on a LAG, just as they would be on a physical port. If you are configuring a LAG, it will typically have all ports in the Network Cloud, but the same Network Cloud can have a LAG and one or more physical ports in it. Multiple LAGs can be assigned to the same Network Cloud as long as each LAG complies with the standard LAG guidelines for minimum and maximum number of ports.

To configure LAG for a Network Cloud, follow this procedure. This procedure assumes that the Link Aggregation Groups already exist. If they do not, you will need to create them before assigning them to the Network Cloud. For information about Link Aggregation Groups, see [“Working With Link Aggregation” on page 281](#).

1. **Display the Network Cloud Summary (Network Resources Manager → Network Clouds).**
2. **Click the name of the Network Cloud for which you want to assign a Link Aggregation Group (LAG).**

This step displays the Network Cloud in the details frame.
3. **Click the Ethernet Ports/LAGs tab to display the LAG currently associated with the Network Cloud.**
 - If no LAG is currently assigned, you can add a LAG to the Network Cloud.
 - If a LAG is currently assigned, you can change the LAG assigned to the Network Cloud, or delete the current LAG from the Network Cloud.

The following example shows a Network Cloud with no LAG assigned.

Next	Name	Type	State	vNICs	Capacity	HA Prefer...	Rank	Trunk Mode	vLAN ID	vLAN Range	MTU	Description
	delaware/10/1	10G Port	up/down	0	10G	Primary	1	access	1		1500	
	oregon/8/2	1G Port	up/down	5	1G	Secondary	2	access	1		1500	

4. Click the plus sign (+) to display the Change Network Cloud LAGs dialog.

Note - To add a LAG, it must already be configured. If it is not yet configured, it is not a selectable item that can be associated with the LAG.

If you are changing the LAG associated with a Network Cloud, select the LAG, click the delete icon (the red dash), then click the plus sign (+) to associate a different LAG with the Network Cloud.

Name	Type	State	vNICs	Capacity	Trunk Mo...	vLAN ID	vLAN Range	Description
arkansas/14/1	nwEthernet1GbPort	up/up	8	1G	access	1		Private Network
arkansas/14/10	nwEthernet1GbPort	up/down	3	1G	trunk	1		tagged-mtu1500
arkansas/14/2	nwEthernet1GbPort	up/up	0	1G	access	1		Private Network
arkansas/14/3	nwEthernet1GbPort	up/up	8	1G	access	1		Private Network
arkansas/14/4	nwEthernet1GbPort	up/down	1	1G	access	5		
arkansas/14/5	nwEthernet1GbPort	up/up	2	1G	trunk	1		Compellent_iSCSI
arkansas/14/6	nwEthernet1GbPort	up/up	3	1G	access	1		LAB Network
arkansas/14/7	nwEthernet1GbPort	up/up	4	1G	access	1		LAB Network
arkansas/14/8	nwEthernet1GbPort	up/up	6	1G	access	1		LAB Network
arkansas/14/9	nwEthernet1GbPort	up/down	0	1G	trunk	1		tagged-mtu9000
arkansas/5/1	nwEthernet10Gb...	up/up	11	10G	access	1		vMotion Network
montana/5.5	LAG	up/down	0		access	1		
montana/5/1	nwEthernet1GbPort	up/down	0	1G	access	1		

5. **Click the LAG you want to assign to the Network Cloud.**
The LAG is selected when it is highlighted, as shown.
6. **Click Submit to assign the LAG to the Network Cloud in Oracle Fabric Manager.**
7. **Check the Ethernet Ports/LAG tab to verify that the LAG was assigned.**
The following example shows a LAG assigned to a Network Cloud.

Next	Name ▲	Type	State	vNICs	Capa
➡	delaware/10/1	10G Port	up/down	0	10G
	delaware/8/5	1G Port	up/down	0	1G
➡	oregon/8/2	1G Port	up/down	5	1G

Managing Network Cloud Properties

Network Clouds are displayed in the summary and details frames of the work panel. Each frame provides different levels of information about the configured Network Clouds.

This section contains the following topics:

- [“Summary Information for Network Clouds” on page 118](#)
- [“Detailed Information for Network Clouds” on page 118](#)
- [“Edit Network Cloud Properties” on page 120](#)
- [“Rename a Network Cloud” on page 121](#)
- [“Push Cloud Changes to an I/O Template's vNICs” on page 123](#)
- [“Display the Ethernet Ports or LAGs in a Network Cloud” on page 125](#)
- [“Display the vNICs Associated With a Network Cloud” on page 125](#)
- [“Apply Network Cloud Changes to vNICs in the Cloud” on page 126](#)
- [“Terminate vNICs in a Network Cloud to Another vNIC Cloud” on page 127](#)
- [“Display the vNIC Templates Associated With a Network Cloud” on page 128](#)

Summary Information for Network Clouds

The Network Cloud Summary is a table of all the configured Network Clouds, including the default Network Cloud (discovered-network-cloud). Additional, detailed information about each Network Cloud is contained in the Network Cloud Details frame, which is located directly below the Network Cloud Summary.

Network Cloud Summary							
Name	Number of Ports	Number of Lags	QoS	Number of vNL...	Number of vNL...	Description	
de	1	0	none	0	0		
discovered-network-cloud	40	3	none	20	0		
techpubs9	2	0	none	0	0		

3 items

Field	Indicates
Name	The name of each configured Network Cloud.
Number of Ports	The number of Ethernet ports in the Network Cloud. This number is the total of all ports, regardless of whether they are up or down.
Number of LAGs	The number of link aggregation groups (LAGs) in the Network Cloud. This number is the total of all LAGs configured, not the total number of ports in LAGs.
Number of vNICs	The total number of vNICs that are connected to the Network Cloud.
Number of vNIC Templates	The total number of I/O Templates that are associated with each Network Cloud.
Description	The description string (if any) that was applied to the Network Cloud. If this field is blank, either no description string was specified when the Network Cloud was created, or the Network Cloud was originally created with a description string, but the Network Cloud was later edited and the description string was removed.

Detailed Information for Network Clouds

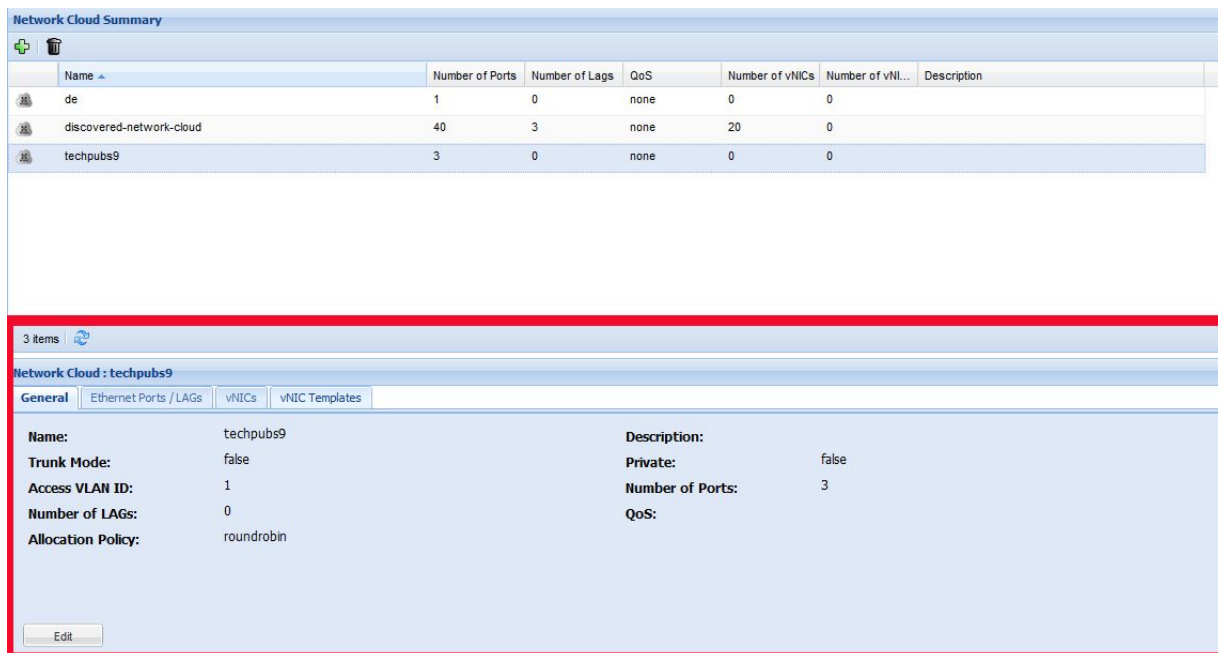
The Network Cloud Details frame is a section of the work panel that is located below the Network Cloud Summary. The Network Cloud Details frame is typically a list of fields for a selected Network Cloud, a list of Ethernet ports or LAGs contained in the Network Cloud, or a list of fields for a selected Ethernet port in the Network Cloud.

The Network Cloud Details frame enables you to display additional, detailed information for a single Network Cloud including the default Network Cloud (discovered-network-ports). For non-default Network Clouds, you can use the Network Cloud Details frame to edit various aspects of the cloud itself, or the components of it. The Network Cloud Details frame has the following controls, which are common on most details frames for features throughout the entire Oracle Fabric Manager interface:

- Add and delete buttons, a plus sign and a red dash, respectively
- Tabs, which logically organized additional information about the contents of the details frame
- An Edit button, which unlocks editable parts of the details frame so that you can set or change information elements for the contents of the details frame

To use the Network Cloud Details frame, you must first select a configured cloud from the Network Cloud Summary. By selecting a cloud from the summary, you set the focus of the Network Cloud Details frame. When the cloud is selected in the summary, you will see its details displayed in the Details frame.

The following example shows the Network Cloud Details frame. Notice that the details frame is contextual, so that it displays detailed information for the item selected in the Network Cloud Summary.



Through the Network Cloud Details frame, you can also see information about the Ethernet Ports or LAGs that are associated with the specific cloud.

▼ Edit Network Cloud Properties

Through the Network Cloud Details frame, you can edit parameters associated with the Network Cloud without having to completely redefine the cloud. When you edit the Network Cloud's properties, the new properties take effect as soon as you save them, and the properties are applied to the vNICs terminated in the cloud.

Note - Anytime you edit a Network Cloud to make changes, it is a best practice to manually push those changes to the vNICs associated with that cloud. For more information, see [“Push Cloud Changes to an I/O Template's vNICs” on page 123](#).

To edit the Network Cloud properties, follow this procedure:

1. **Display the Network Cloud Summary.**
2. **Click a configured Network Cloud to highlight it.**
When the cloud is highlighted, the Network Cloud Details frame is populated with the details for the selected Network Cloud.
3. **Click the Edit button to unlock the editable fields.**

The screenshot shows the 'Network Cloud : techpubs1' dialog box with the 'General' tab selected. The fields are as follows:

Name: *	techpub4	Description:	
Trunk Mode:	<input type="checkbox"/>	Private:	<input type="checkbox"/>
VLAN ID:	1	Number of Ports:	0
Number of LAGs:	0	QoS:	none
Allocation Policy:	Round Robin		

At the bottom of the dialog are 'Submit' and 'Cancel' buttons.

4. **Make the appropriate changes. Editable fields are:**
 - **Name, for renaming the Network Cloud after it has been created.**

- **Description**, for specifying an optional description string for the cloud.
 - **Trunk Mode**, for vNICs in the cloud that are, or will be, participating in a VLAN.
 - **VLAN ID**, for vNICs in the cloud that are, or will be, participating in a VLAN.
 - **Private**, which is a checkbox that toggles the state of Private vNIC on or off for the Network Cloud.
 - **Allocation Policy**, for specifying how Ethernet ports get attached to vNICs that are configured in the cloud.
 - **QoS**, for setting or changing the Network QoS Policer Profile associated with the vNICs in the Network Cloud. You can also use this dropdown menu to remove cloud-level QoS by depopulating any assigned Network QoS Profile from the field.
5. **When the Network Cloud changes are complete, click Save.**

▼ **Rename a Network Cloud**

Oracle Fabric Manager supports renaming a Network Cloud without having to completely delete and recreate the entire Network Cloud. When the Network Cloud is renamed, all other properties for the Network Cloud are retained, including the ports and vNICs associated with the cloud. As an option, you can also set or change the description for the Network Cloud.

You can rename a Network Cloud through the Network Cloud Details frame. To rename the Network Cloud, follow this procedure:

1. **On the Navigation Frame, select Network Cloud Manager → Network Clouds.**

2. Select a Network Cloud to populate the details frame with its properties.

Name	Number of Ports	Number of Lags	Number of vNICs	Number of vNIC T...	Description
backup	3	0	0	0	
Cloud9	3	0	0	3	
devtest	2	0	0	0	
discovered-network-cloud	31	1	103	0	
EngNWcloudA	1	0	0	1	Engineering Network Activity
management	2	0	0	0	
netapp1	2	0	0	0	
production	2	0	0	0	
techpubs1	0	0	0	3	

3. Click the Edit button to edit the properties of the selected Network Cloud.

Network Cloud Summary

Name	Number of Ports	Number of Lags	Number of vNICs	Number of vNIC T...	Description
backup	3	0	0	0	
Cloud9	3	0	0	3	
devtest	2	0	0	0	
discovered-network-cloud	31	1	103	0	
EngNWcloudA	1	0	0	1	Engineering Network Activity
management	2	0	0	0	
netapp1	2	0	0	0	
production	2	0	0	0	
techpubs1	0	0	0	3	

Network Cloud : techpubs1

General | Ethernet Ports / LAGs | vNICs | vNIC Templates

Name: **Description:**

Trunk Mode: **Private:**

VLAN ID: **Number of Ports:**

Number of LAGs: **QoS:**

Allocation Policy:

4. In the Name field, enter the new name for the Network Cloud.
5. As an option, you also can set or change the description for the selected Network Cloud.
6. As an option, you also can set or change any of the other editable fields on the details frame.

7. **When the new name has been specified for the Network Cloud, click Submit.**

▼ **Push Cloud Changes to an I/O Template's vNICs**

When you configure an I/O Template, you reference Network or Storage Clouds for the termination of the vNICs and vHBAs that connect the host server. These clouds are external to the template, and they are controlled through a different workflow. Because the cloud(s) and the templates are separate entities, any updates on a cloud must be explicitly forced to the I/O Template. Conceptually, this is intuitive because you are changing a cloud, you are not actually changing anything within the template itself.

If you need to change a cloud property, and that cloud is associated with an I/O Template, be aware that the change you make is not automatically updated in the I/O Template that is using that cloud. However, you can push changes in a Network or Storage Cloud to the I/O Template that is using that cloud by using the Re-Apply Cloud Changes to vNIC button for Network Clouds or the Re-Apply Cloud Changes to vHBA button for Storage Clouds.

To reapply Network Cloud changes to vNICs in a template, follow this procedure:

1. **Select Network Cloud Manager → Network Clouds to display the Network Cloud Summary.**
2. **Select the Network Cloud for which you want to make changes.**

This step populates the details frame with information about the selected Network Cloud.

The screenshot shows the 'Network Cloud Summary' table with three entries. Below it, the 'Network Cloud : techpubs9' details are shown in the 'General' tab.

Name	Number of Ports	Number of Lags	QoS	Number of vNICs	Number of vNI...	Description
de	1	0	none	0	0	
discovered-network-cloud	40	3	none	20	0	
techpubs9	3	0	none	0	0	

Network Cloud : techpubs9	
General	
Name:	techpubs9
Trunk Mode:	false
Access VLAN ID:	1
Number of LAGs:	0
Allocation Policy:	roundrobin
Description:	
Private:	false
Number of Ports:	3
QoS:	roundrobin

- In the Cloud Details frame, click the vNICs tab.**
- On the vNICs tab, select the vNIC(s) you want to receive the Network Cloud Changes.**

This step activates the Re-Apply Cloud Changes to vNIC button.

The screenshot shows the 'vNICs' tab for the 'pubscloud' network cloud. A table lists the vNICs with columns for Name, Network Cloud, Server Name, Termination, State, IP Address, Netmask, IP Type, MAC Address, HA, QoS, Private, TSO, and Batching.

Name	Network Cloud	Server Name	Termination	State	IP Address	Netmask	IP Type	MAC Address	HA	QoS	Private	TSO	Batching
vnic	pubscloud	bacchus		up/null			hostManaged		false	Disabled	false	false	false

- Click the button to apply the changes to the selected vNIC(s).**

6. Click Yes to confirm the changes.

When you do, the Storage Cloud changes are pushed to the vHBA.

▼ Display the Ethernet Ports or LAGs in a Network Cloud

Each Network Cloud contains one or more Ethernet ports (or LAGs) that provide the termination point for vNICs connected to the cloud. For each Network Cloud, you can display a list of all the Ethernet ports or LAGs in a particular cloud. In addition to displaying the ports or LAGs in the cloud, you can:

- Remove Ethernet ports (or LAGs) from the Network Cloud
- Control the port allocation policy for each Ethernet port (or LAG) in the Network Cloud.

Note - Ethernet ports are listed in a modified ascending order. So, for example, port iowa1/1 is listed before port iowa2/1 as expected. However, port iowa1/10 and iowa11/1 would also be listed before port iowa2/1.

Due to the listing order in the Ethernet Ports/LAGs tab, you will need to pay attention to which port you select as the next available by moving the arrows around. For example, if you want the round robin policy to start with port iowa1/1 and the next port to be assigned as port iowa1/2, you cannot just move the arrow to the next port down in the list because this might be port iowa11/1 (not iowa1/2). Without paying close attention to the listing order, you might accidentally assign the wrong port(s). As a result, the round robin policy will operate correctly, but do so based on the incorrect ports you inadvertently assigned.

To display the Ethernet ports, follow this procedure:

1. **On the Navigation Frame, select Network Cloud Manager → Network Clouds.**
2. **Select a Network Cloud to populate the details frame with its properties.**
3. **Click the Ethernet Ports/LAGs tab to display the Fibre Channel ports associated with the Storage Cloud.**

▼ Display the vNICs Associated With a Network Cloud

When vNICs are either associated with an I/O Profile or deployed to a host server, those vNICs are displayed through the vNICs tab on the Network Cloud details frame. The vNICs

tab displays only vNICs that are actually deployed on a host server, not the vNICs that are connected to a Network Cloud. As a result, it is possible that the vNICs tab is empty even though a Network Cloud actually has vNICs connected to it.

To display the vNICs associated with a Network Cloud, follow this procedure:

1. **Display the Network Cloud Details frame.**
2. **Click the vNICs tab.**

Name	Network Cloud	Server Name	Termination	State	IP Address	Netmask	IP Type	MAC Address	HA	QoS	Private
vnic7	NTCL8_1m_10m	WHQL3	iowa/7/1	up/initializing	169.254.198.2...	255.0.0.0	hostMana...	00:13:97:01:89:54	fa...	1000_100...	false
vnic7	NTCL8_1m_10m	coke	south-carolina...	up/up	0.0.0.0	255.255.255.2...	hostMana...	00:13:97:03:7D:69	fa...	1000_100...	false
vnic7	NTCL8_1m_10m	police.lab.xsigo.com	iowa/7/1	up/up	0.0.0.0	255.255.255.2...	hostMana...	00:13:97:01:89:2C	fa...	1000_100...	false
vnic7	NTCL8_1m_10m	DTM2	south-carolina...	up/up	169.254.155.1...	255.0.0.0	hostMana...	00:13:97:03:79:0C	fa...	1000_100...	false
vnic7	NTCL8_1m_10m	LOGO5	iowa/7/1	up/up	169.254.100.2...	255.0.0.0	hostMana...	00:13:97:01:89:40	fa...	1000_100...	false

▼ Apply Network Cloud Changes to vNICs in the Cloud

Properties from the Network Cloud can be inherited by the vNICs that attach to that cloud. For example, VLAN IDs and Network QoS profiles are properties for the Network Cloud that can be inherited by the vNICs that attach to it. For more information, see [“Understanding Network Clouds” on page 99](#).

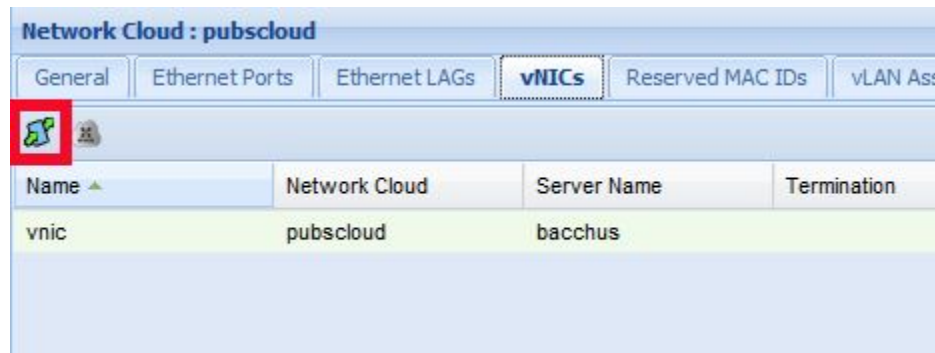
If a Network Cloud's properties have been edited after vNICs are attached. The new properties for the Network Cloud are not automatically inherited by the currently attached vNICs. Instead, you must manually push those changes to the vNICs that are currently attached to the cloud. New vNICs that are attached after the cloud's properties were changed will receive the new properties, but currently connected vNICs do not receive the new cloud properties until you apply those changes. You can apply network cloud changes to currently attached vNICs by using the Apply Network Cloud changes button on the Network Cloud details frame.

To apply changed cloud properties to currently attached vNICs, follow this procedure:

1. **Display the Network Cloud details frame.**
2. **Click the vNICs tab.**

For an example see [“Display the vNICs Associated With a Network Cloud”](#) on page 125.

3. **Select the vNIC to which you want to push the Network Cloud changes.**
This step activates the Apply Cloud Changes to vNICs button.
4. **Click the Apply Cloud Changes to vNICs button to push all the new Network Cloud properties to the currently selected vNICs.**



After a brief period of time, the vNICs have the new cloud properties.

▼ Terminate vNICs in a Network Cloud to Another vNIC Cloud

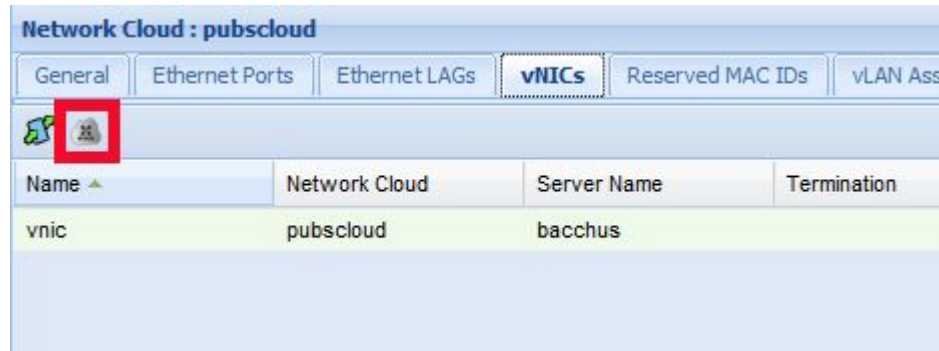
Some management tasks might require you to move a vNIC to a different Network Cloud. For example, if you need to run diagnostics on an Ethernet port, or perhaps service an I/O Module. In such cases, if other Network Clouds are available and have available Ethernet ports, you can change the termination of a vNIC from one cloud to another through Oracle Fabric Manager. If you change the termination of a vNIC with live traffic, a service interruption will occur.

To terminate a vNIC on another Network Cloud, follow this procedure:

1. **Display the Network Cloud details frame.**
2. **Click the vNICs tab.**
For an example see [“Display the vNICs Associated With a Network Cloud”](#) on page 125.
3. **Select the vNIC in the Network Cloud you want to terminate on a different Network Cloud.**

This step activates the Terminate vNICs on another Cloud button.

4. **Click the Terminate vNICs on another Cloud button to push all the new Network Cloud properties to the currently selected vNICs.**



5. **From the resulting dialog, select the Network Cloud on which you want to terminate the selected vNIC.**

This step activates the Submit button.

6. **Click Submit to reterminate the selected vNIC.**

After a brief period of time, the vNIC are terminated on available Gigabit Ethernet ports in the new Network Cloud.

▼ Display the vNIC Templates Associated With a Network Cloud

Network Clouds are part of I/O Templates, which in turn, are used in creating the I/O Profiles that are deployed to servers. Through the vNIC Templates tab, Oracle Fabric Manager supports a way of seeing all the vNIC Templates that are associated with a specific Network Cloud. This tab shows the association between vNICs in a template and the I/O Cloud regardless of the state of the vNIC being bound or not.

The vNIC Templates tab shows information only. You cannot manage the vNIC Templates associated with a specific Network Cloud.

To display the vNIC Templates for a specific Network Cloud, follow this procedure:


1. **Display the Network Cloud details frame.**

2. Click the vNIC Templates tab.

Network Cloud : techpubs1

General | Ethernet Ports / LAGs | vNICs | **vNIC Templates**

Name ▲	VO Cloud	QoS	IP Type	Bootable	HA	VLAN Id	Trunk Mode	Private
vnic.PoSprocessing_1	techpubs1	1g_10g	Host Managed	false	false	1	false	false
vnic.pubscsiboot	techpubs1		Host Managed	true	true	1	false	false
vnic.Pubstemplate1	techpubs1	1g_10g	Host Managed	false	false	1	false	false
vnic.scsi4	techpubs1		Host Managed	true	true	1	false	false

4 items 

Working With Storage Clouds

This chapter contains the following topics:

- [“Understanding Storage Clouds” on page 131](#)
- [“Create a Storage Cloud” on page 132](#)
- [“Configuring Features in a Storage Cloud” on page 135](#)
- [“Displaying Storage Clouds” on page 143](#)

Understanding Storage Clouds

Oracle Fabric Devices contain a set of I/O ports that are connected to your data center SAN. The SAN administrator connects these ports to their Fibre Channel switches in order to provide I/O resources to the servers.

Each I/O port on the Oracle system can be thought of as providing access to a set of “clouds”. For example, you might have a Storage Cloud that is zoned to give access to a set of LUNs used by HR. As far as the storage administrator is concerned, you are connecting wires to the Oracle I/O ports for the sole purpose of providing the server administrator access to the resources.

The server administrator then needs to connect the servers to a set of storage resources provided by the storage administrators. The server administrator doesn't care what physical Oracle ports are being connected to as long as the servers have access to the required resources.

Storage Clouds can be configured and managed by Storage and Administrator roles.

SAN QoS Profiles and Storage Clouds

If you want a specific SAN QoS Profile assigned to a Storage Cloud, that specific profile must already exist. If it does not exist, it will not be assignable to the Storage Cloud at the time you create the cloud.

Oracle Fabric Manager has some commonly used SAN QoS Profiles that are pre-configured for various Fibre Channel link speeds. These pre-configured SAN QoS Profiles use typical CIR

and PIR values calculated from the link's total throughput. If you are using pre-configured SAN QoS Profiles, make sure that you use a reasonable profile.

SAN QoS can be configured for a Storage Cloud at any of the following times:

- During Storage Cloud creation, through the Create Storage Cloud wizard.
- After the Storage Cloud has already been created, through the Storage Cloud Details frame.
- SAN QoS can also be changed for a Storage Cloud after the Storage Cloud has already been created. Changing an assigned SAN QoS Profile occurs through the Storage Cloud Details frame.

For information about SAN QoS Profiles, see [“Working With SAN QoS” on page 473](#).

QoS Assignment at Cloud Level and vHBA Level

It is possible to assign SAN QoS to a Storage Cloud or a vHBA. The following text explains where and how QoS get applied:

- If a Storage Cloud has a QoS Profile associated with it, a vHBA terminated in the Storage Cloud will inherit the Storage QoS Profile from the Storage Cloud.
- In this case, the vHBAs inherit the QoS parameters from their respective clouds. This condition is true even if the vHBA has its own QoS Profile associated with it.
- If a Storage Cloud has no QoS Profile associated with it, and the vHBAs have their own QoS profiles, a vHBA terminated in the Storage Cloud will have the QoS profile associated with it.
- In this case, the Storage Cloud has no QoS Profile to apply. This situation does not automatically apply no QoS to the vHBAs. Instead, it causes the cloud to defer to the QoS profiles assigned to the individual vHBAs (if any).
- If you intend for no QoS to be applied anywhere for the virtual I/O in a cloud, make sure that no QoS Profile is associated with the Storage Cloud, and also that no QoS is associated with the vHBA.

▼ Create a Storage Cloud

When you create a Storage Cloud, you associate one or more Fibre Channel ports that have been assigned to a Fibre Channel switch with the cloud. Other SAN characteristics (such as SAN QoS) are also associated with the Storage Cloud. When the Storage Cloud is attached to an I/O Template, the Storage Cloud eventually provides a way for hosts to attach to the storage network after an I/O Profile containing the Storage Cloud is pushed to the hosts. The Storage Cloud provides needed storage access points with specified characteristics—for example, SAN QoS.

A default Storage Cloud exists, which contains all the discovered Fibre Channel ports available to Oracle Fabric Manager through the discovered Fabric Device(s). This default Storage Cloud is a way to show you the total number of termination points that are under management by Oracle Fabric Manager. The total number shown does not decrement when ports are assigned to a specific Storage Cloud, since those ports have not been removed from Oracle Fabric Manager—they are just reassigned to a different object.

Storage Clouds typically are created around some similar theme for the Fibre Channel ports that are associated with it. For example, a Storage Cloud might have ports that provide vHBA connections to a specific zone or controller, for a specific set of host servers.

Creating a Storage Cloud occurs through the Storage Clouds option in the Storage Resources Manager on the navigation panel. To create a Storage Cloud, follow this procedure:

1. **On the Navigation panel, select Storage Resource Manager → Storage Clouds to display the Storage Cloud Summary.**

	Name	Number of Ports	QoS	Number of vHBAs	Number of vHBA temp...	Description
	discovered-storage-cloud	22	none	102	0	
	EngSTcloudA	1	none	0	1	Engineering Storage Activity
	S_Cloud0	1	none	0	3	
	techpubs1	0	none	0	2	

4 items

Notice the default Storage Cloud (discovered-storage-cloud) in this example. This example also shows that specific Storage Clouds (non-default clouds) are configured.

Adding and deleting Storage Clouds occurs through the plus sign and garbage can icons, respectively.

2. Click the plus sign (+) to display the Create a New Storage Cloud dialog.

New Storage Cloud

Name: * techpubs1

Description:

FC Ports: *

Port Name	Type	State	vHBAs	Topology	Rate	Description
delaware/11/1	sanFc...	up/resourceMis...	0	F	auto	
delaware/11/2	sanFc...	up/resourceMis...	1	F	auto	
oregon/11/1	sanFc...	up/down	7	F	auto	
oregon/11/2	sanFc...	up/down	4	F	auto	

4 items

HA Designation: Check to set Primary and Secondary Fabric Interconnects

Advanced Configuration

Submit Cancel

3. In the Name field, enter a name for the Storage Cloud that you are creating.
4. As an option, in the Description field, enter a description for the Storage Cloud that you are creating.
5. From the FC Ports table, select the port(s) that will be used in this Storage Cloud. Ports are selected when they are highlighted.
Multiple ports can be assigned to the same Storage Cloud, and multiple Storage Clouds can be assigned to the same port(s).
6. When the correct properties have been assigned for the Storage Cloud, click **Submit** to create the Storage Cloud.
When you click Submit, a dialog is briefly displayed that informs you that Oracle Fabric Manager client and server are exchanging data.
7. Check the Storage Cloud Summary to verify that the Storage Cloud was successfully created.

Configuring Features in a Storage Cloud

Because Storage Clouds contain Fibre Channel ports, many standard features that are applicable to physical FC ports (for example, SAN QoS) are configurable on Storage Clouds. Storage Clouds support configuration of the following features:

- SAN QoS
- Port Allocation Policy

This section contains the following topics:

- [“Set SAN QoS on a Storage Cloud” on page 135](#)
- [“Set an Allocation Policy for a Storage Cloud” on page 138](#)
- [“Change the Port Priority for a Storage Cloud” on page 142](#)

▼ Set SAN QoS on a Storage Cloud

Setting Storage QoS on a Storage Cloud can occur at the cloud level or on the vHBA level. When SAN QoS is applied to a Storage Cloud, all vHBAs associated with the cloud will get the specified CIR value. If a vHBA also has SAN QoS and is added to the Storage Cloud, the cloud's SAN QoS takes precedence. For more information about SAN QoS and Storage Clouds, see [“Working With SAN QoS” on page 473](#).

To configure SAN QoS for a Storage Cloud, follow this procedure. This procedure assumes that the SAN QoS Profile already exists. If it does not, you will need to create the SAN QoS Profile.

1. **Display the Storage Cloud Summary (Storage Resources Manager → Storage Clouds).**

See [“Create a Storage Cloud” on page 132](#) for an example.

2. Click the plus sign to display the New Storage Cloud dialog.

New Storage Cloud

Name: * techpubs1

Description:

FC Ports: *

Port Name ▲	Type	State	vHBAs	Topology	Rate	Description
delaware/11/1	sanFc...	up/resourceMis...	0	F	auto	
delaware/11/2	sanFc...	up/resourceMis...	1	F	auto	
oregon/11/1	sanFc...	up/down	7	F	auto	
oregon/11/2	sanFc...	up/down	4	F	auto	

4 items | ↻

HA Designation: Check to set Primary and Secondary Fabric Interconnects

▼ **Advanced Configuration**

Submit Cancel

3. Select at least one fibre channel port for creating the Storage Cloud.

4. On the New Storage Cloud dialog, click the Advanced Configuration button to display the QoS dropdown menu.

New Storage Cloud

Name: * techpubs1

Description:

FC Ports: *

Port Name ▲	Type	State	vHBAs	Topology	Rate	Description
delaware/11/1	sanFc...	up/resourceMis...	0	F	auto	
delaware/11/2	sanFc...	up/resourceMis...	1	F	auto	
oregon/11/1	sanFc...	up/down	7	F	auto	
oregon/11/2	sanFc...	up/down	4	F	auto	

4 items

HA Designation: Check to set Primary and Secondary Fabric Interconnects

▲ **Advanced Configuration**

QoS: 2g_4g ▼

Allocation Policy: Round Robin ▼

Submit Cancel

In this example, a new SAN QoS Profile configured for 2 Gbps of CIR and 4 Gbps Mbps of PIR (2g_4g) is being applied to the Storage Cloud.

5. Click **Submit** to assign the new SAN QoS Profile to the selected Storage Cloud. You will be prompted with a confirmation message.
Click Yes to assign the SAN QoS profile or No to abort assigning the SAN QoS Profile to the Storage Cloud.
6. Check the QoS field of the Storage Cloud Details Frame to verify that the correct SAN QoS Policy is assigned.

▼ Set an Allocation Policy for a Storage Cloud

An allocation policy determines how ports are assigned from the Storage Cloud to any vHBA that is connected to the cloud. In order for the port to be allocated, it must already be configured in the Storage Cloud. Multiple ports can be configured in a Storage Cloud, and ports from different Fibre Channel modules, and even different Oracle Fabric Devices, can be configured in the same Storage Cloud.

The following allocation policy exists:

- Round Robin, which is a systematic way for ports to be assigned. With round robin allocation, you specify multiple ports in a list of available ports for a Storage Cloud. When the list is constructed, you then assign a port as the next available port. When you assign the next available port, it gets a rank that is different than any other port in the available ports list. The port you assign is the next port to be assigned, and after that, any additional vHBAs that are connected to the Storage Cloud will receive a fibre channel port based on its rank. For example, assume you create a Storage Cloud with fibre channel port 2/1 as the port used to create the cloud. Then, assume you add ports 2/2, 4/1, and 5/1 to the available ports list, and set port 4/1 as the next available port. When new vHBAs are connected to the Storage Cloud, the first vHBA is connected to port 4/1. After that, additional vHBAs are assigned to whichever port has the next lowest rank.

To set an allocation policy for a storage cloud, follow this procedure:

New Storage Cloud

Name: * techpubs1

Description:

FC Ports: *

Port Name	Type	State	vHBAs	Topology	Rate	Description
delaware/11/1	sanFc...	up/resourceMis...	0	F	auto	
delaware/11/2	sanFc...	up/resourceMis...	1	F	auto	
oregon/11/1	sanFc...	up/down	7	F	auto	
oregon/11/2	sanFc...	up/down	4	F	auto	

4 items

HA Designation: Check to set Primary and Secondary Fabric Interconnects

Advanced Configuration

QoS: 2g_4g

Allocation Policy: Round Robin

Submit Cancel

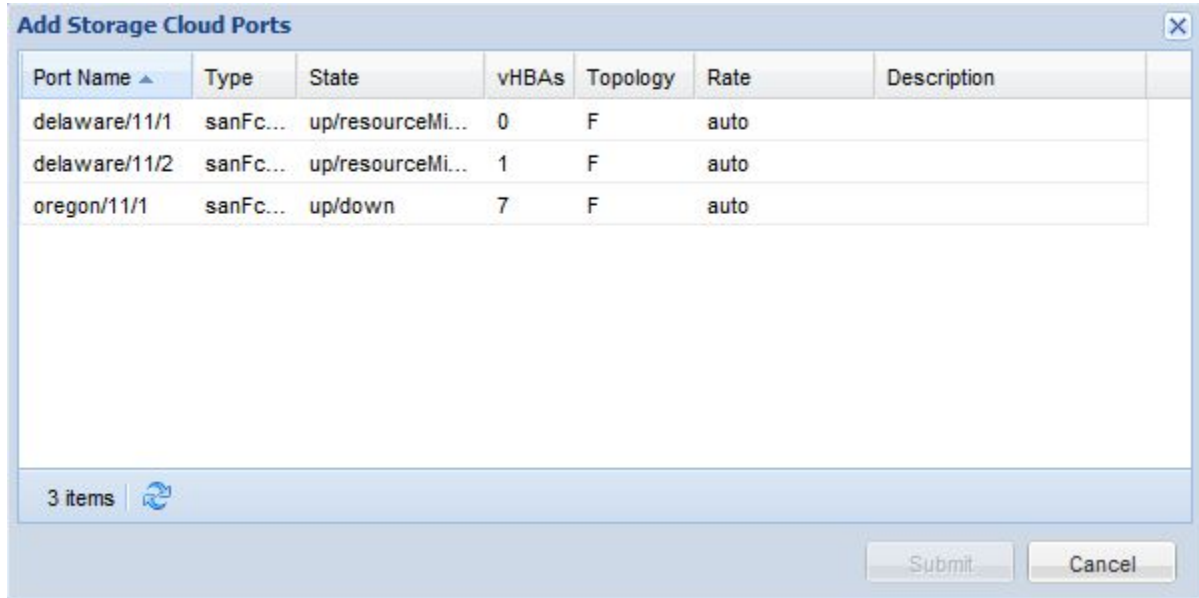
1. If you have not already done so, create a Storage Cloud with at least one fibre channel port in it, as documented in [“Create a Storage Cloud” on page 132](#).
2. When the Storage Cloud is created, click the Advanced Configuration button to display the Allocation Policy dropdown menu.
3. From the Allocation Policy dropdown menu, select the method you want to use to allocate fibre channel ports from the Storage Cloud:
 - Round Robin allocates port in the cloud by rotating through the unallocated fibre channel ports in a systematic way starting from the port you select when you create the Storage Cloud. When using round robin port allocation, you will need to select a list of the ports available in the round robin queue. When the list is specified, anytime a vHBA is connected to the Storage Cloud, Oracle Fabric Manager will step through the list and allocate the next available port based on numerical value of the port's slot. For example, For

example, assume the following ports are available in the Storage Cloud: port 2 in slot 2, port 1 in slot 4, and port 2 in slot 9. When new vHBAs are connected to the Storage Cloud, the first vHBA will be terminated on port 2 slot 2, the next vHBA will be connected to port 1 in slot 4, and the last vHBA will be terminated on port 2 in slot 9.

4. **When the Allocation Policy is selected, click Submit to create the Storage Cloud.**
5. **On the Storage Cloud Summary, select the Storage Cloud you just created.**
When you do so, the details frame will contain additional information about the selected Storage Cloud.
6. **Click the FC Ports tab to display the list of fibre channel ports in the Storage Cloud.**



- Click the plus sign (+) to display the Add Storage Cloud Ports dialog.



- Select the port(s) that you want to add to the list of ports.
When you do, the Submit button becomes active.
- Click Submit to add the selected ports to the list.
- After the correct storage ports have been added to the Storage Cloud, click the blue arrow on the toolbar to set the next available port in the list.



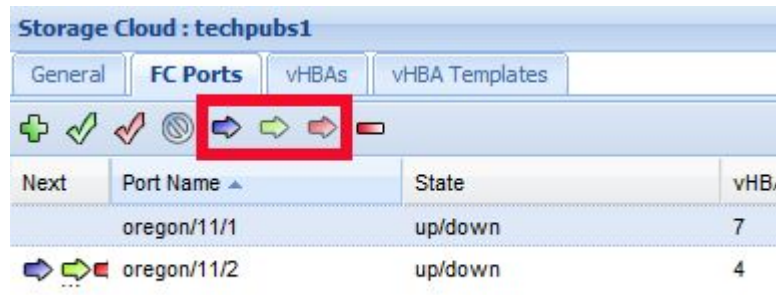
After the next available port is selected, Oracle Fabric Manager will begin the round robin queue from that port, and assign fibre channel ports in ascending numerical order based on the Rank field. You can set the rank by specifying the next port for each port in the list. When each “next port” is specified, it gets a unique rank number, which is used in selecting the next available port in the round robin algorithm.

11. When all ports in the list are selected, click Submit to finalize the list.

It is a good idea to double check the rank associated with each port in the list because this is the criteria used by the round robin algorithm. Ports with lower ranks are assigned before ports with higher ranks.

▼ Change the Port Priority for a Storage Cloud

The port priority for a storage cloud can be changed by selecting a different option from the toolbar on the FC Ports tab.



This toolbar contains the following buttons for setting port priority:

- Mark the next port to be assigned (the blue arrow). This option is available for standalone vHBAs (non-HA/non-multipath).
- Mark the next primary port to be assigned (the green arrow). This option is available only for HA/multipath vHBAs.
- Mark the next secondary port to be assigned (the red arrow). This option is available only for HA/multipath vHBAs.

If the available ports list has only one port, and the port priority is set incorrectly, you might need to delete the port list, then recreate it. You can delete ports from the port availability list by selecting the port you want to delete, then clicking the red dash.

Note - Fibre Channel ports are listed in a modified ascending order. So, for example, port iowa1/1 is listed before port iowa 2/1 as expected. However, port iowa1/10 and iowa11/1 would also be listed before port iowa2/1.

Due to the listing order in the FC Ports tab, you will need to pay attention to which port you select as the next available by moving the arrows around. For example, if you want the round robin policy to start with port iowa1/1 and the next port to be assigned as port iowa2/1, you cannot just move the arrow to the next port down in the list because this might be port iowa11/1 (not iowa1/2). Without paying close attention to the listing order, you might accidentally assign the wrong port(s). As a result, the round robin policy will operate correctly, but do so based on the incorrect ports you inadvertently assigned.

To change the port priority in the list, follow this procedure:

- 1. Select the port for which you want to set or change the port priority.**
The port is selected when it is highlighted in the FC Ports tab. You will also notice that the port priority toolbar buttons become active.
- 2. Click the appropriate button to set the port priority for the selected port:**
 - For non-HA/non-multipath vHBAs, only the blue arrow is valid.
 - For HA/multipath vHBAs, use the green arrow to set the port to be assigned as the primary port for the vHBA. Then, use the red arrow to set the port that will be next assigned as the secondary vHBA.

Displaying Storage Clouds

Storage Clouds are displayed in the summary and details frame of the work panel, which provide information about the configured Storage Clouds.

This section contains the following topics:

- [“Display Summary Information for Storage Clouds” on page 143](#)
- [“Displaying Detailed Information for Storage Clouds” on page 144](#)

Display Summary Information for Storage Clouds

The Storage Cloud Summary is a table of all the configured Storage Clouds, including the default Storage Cloud (discovered-storage-cloud). Additional, detailed information about each Storage Cloud is contained in the Storage Cloud Details frame, which is located directly below the Storage Cloud Summary.

Storage Cloud Summary						
Name	Number of Ports	QoS	Number of vHBAs	Number of vHBA temp...	Description	
discovered-storage-cloud	22	none	102	0		
EngSTcloudA	1	none	0	1	Engineering Storage Activity	
S_Cloud0	1	none	0	3		
techpubs1	0	none	0	2		

Field	Indicates
Name	The name of each configured Storage Cloud.
Number of Ports	The number of Fibre Channel Ports in the Storage Cloud. This number is the total of all ports, regardless of whether they are up or down.
QoS	The name of any SAN QoS Profile associated with the Storage Cloud. If no SAN QoS is configured, this field indicates “none”. If a SAN QoS Profile is configured for the Storage Cloud, the name indicates the CIR and PIR values. For example, a QoS Profile named 125M_250M indicates 125 Mbps of CIR and 250 Mbps of PIR.
Number of vHBAs	The total number of vHBAs connected to the Storage Cloud.
Number of vHBA Templates	The total number of vHBAs in I/O Templates that are associated with the Storage Cloud.
Description	The description string (if any) that was applied to the Storage Cloud. If this field is blank, either no description string was specified when the Storage Cloud was created, or the Storage Cloud was originally created with a description string, but the Storage Cloud was later edited and the description string was removed.

Displaying Detailed Information for Storage Clouds

This section contains the following topics:

- [“Storage Cloud Details Overview” on page 145](#)
- [“Rename a Storage Cloud” on page 146](#)
- [“Push Cloud Changes to vHBAs in an I/O Template” on page 148](#)
- [“Display the Fibre Channel Ports in a Storage Cloud” on page 149](#)
- [“Display the vHBAs Associated With a Storage Cloud” on page 150](#)
- [“Apply Storage Cloud Changes to vHBAs in the Cloud” on page 150](#)
- [“Terminate vHBAs in a Storage Cloud to Another Storage Cloud” on page 151](#)
- [“Display the vHBA Templates Associated With a Storage Cloud” on page 152](#)

Storage Cloud Details Overview

The Storage Cloud Details frame is a section of the work panel that is located below the Storage Cloud Summary. The Storage Cloud Details frame is typically a list of fields for a selected Storage Cloud, a list of FC ports contained in the Storage Cloud, or a list of fields for a selected FC port in the Storage Cloud.

The Storage Cloud Details frame enables you to display additional, detailed information for a single Storage Cloud including the default Storage Cloud (discovered-storage-ports). For non-default Storage Clouds, you can use the Storage Cloud Details frame to edit various aspects of the cloud itself, or the components of it. The Storage Cloud Details has the following controls, which are common on most details frames for features throughout the entire Oracle Fabric Manager interface:

- Add and delete buttons, a plus sign and a red dash, respectively
- Tabs, which logically organize additional information about the contents of the details frame
- An Edit button, which unlocks editable parts of the details frame so that you can set or change information elements for the contents of the details frame

To use the Storage Cloud Details frame, you must first select a configured cloud from the Storage Cloud Summary. By selecting a cloud from the summary, you set the focus of the Storage Cloud Details frame. When the cloud is selected in the summary, you will see its details displayed in the Details frame.

Through the Storage Cloud details frame, you can edit properties for the Storage Cloud.

Note - Anytime you edit a Storage Cloud to make changes, it is a best practice to manually push those changes to the vHBAs that are associated with that cloud. For more information, see [“Push Cloud Changes to vHBAs in an I/O Template” on page 148](#).

The following example shows the Storage Cloud Details frame. Notice that the details frame is contextual, so that it displays detailed information for the item selected in the Storage Cloud Summary.

Rename a Storage Cloud

The screenshot shows the Oracle Fabric Manager interface. At the top is a 'Storage Cloud Summary' table with columns: Name, Number of Ports, QoS, Number of vHBAs, Number of vHBA te..., and Description. Below the table is a '5 items' indicator and a refresh icon. The main area is titled 'Storage Cloud : techpubs1' and has tabs for 'General', 'FC Ports', 'vHBAs', and 'vHBA Templates'. The 'General' tab is active, showing fields for Name (techpubs1), QoS (2g_4g), Number of Ports (2), and Allocation Policy (roundrobin). A description field is empty. An 'Edit' button is located at the bottom left of the details frame.

Name	Number of Ports	QoS	Number of vHBAs	Number of vHBA te...	Description
discovered-storage-cloud	4	none	1	0	
MirrorSANCloudLA	1	none	5	2	
MirrorSANCloudLA1	2	none	3	1	
Storage1Raid10	1	none	0	1	
techpubs1	2	2g_4g	0	0	

5 items

Storage Cloud : techpubs1

General | FC Ports | vHBAs | vHBA Templates

Name: techpubs1 **Number of Ports:** 2
QoS: 2g_4g **Allocation Policy:** roundrobin
Description:

Edit

▼ Rename a Storage Cloud

Oracle Fabric Manager supports renaming a Storage Cloud, which enables you to change the name without having to completely delete and recreate the entire Storage Cloud. When the Storage Cloud is renamed, all other properties for the Storage Cloud are retained, including the ports and vHBAs associated with the cloud. As an option, you can also set or change the description for the Storage Cloud.

You can rename the Storage Cloud through the Storage Cloud Details frame. To rename the Storage Cloud, follow this procedure:

1. On the Navigation Frame, select Storage Cloud Manager → Storage Clouds.

Storage Cloud Summary

Name	Number of Ports	QoS	Number of vHBAs	Number of vHBA temp...	Description
discovered-storage-cloud	22	none	102	0	
EngSTcloudA	1	none	0	1	Engineering Storage Activity
S_Cloud0	1	none	0	3	
techpubs1	2	none	0	2	

4 items

2. Select a Storage Cloud to populate the details frame with its properties.
3. Click the Edit button to edit the properties of the selected Storage Cloud.

Storage Cloud : techpubs1

General | FC Ports | vHBAs | vHBA Templates

Name: * Number of Ports: 2

QoS: Allocation Policy:

Description:

4. In the Name field, enter the new name for the Storage Cloud.
5. As an option, you also can set or change the description for the selected Storage Cloud.
6. As an option, you also can set or change any of the other editable fields on the details frame.
7. When the new name has been specified for the Storage Cloud, click Submit.

▼ Push Cloud Changes to vHBAs in an I/O Template

When you configure an I/O Template, you reference Storage Clouds for the termination of the vHBAs that connect the host server. These clouds are external to the template, and they are controlled through a different workflow. Because the cloud(s) and the templates are separate entities, any updates on a cloud must be explicitly forced to the I/O Template. Conceptually, this is intuitive because you are changing a cloud, you are not actually changing anything within the template itself.

If you need to change a cloud property, and that cloud is associated with an I/O Template, be aware that the change you make is not automatically updated in the I/O Template that is using that cloud. However, you can push changes in a Network or Storage Cloud to the vHBAs that are using that cloud by using the Re-Apply Cloud Changes to vHBA button for Storage Clouds.

To reapply Storage Cloud changes to vHBAs in the cloud, follow this procedure:

1. **Select Storage Cloud Manager → Storage Clouds to display the Storage Cloud Summary.**
2. **Select the Storage Cloud for which you want to make changes.**

This step populates the details frame with information about the selected Storage Cloud.

Storage Cloud Summary

Name	Number of Ports	QoS	Number of vHBAs	Number of vHBA te...	Description
discovered-storage-cloud	4	none	1	0	
MirrorSANCloudLA	1	none	5	2	
MirrorSANCloudLA1	2	none	3	1	
Storage1Raid10	1	none	0	1	
techpubs1	2	2g_4g	0	0	

5 items

Storage Cloud : techpubs1

General | FC Ports | vHBAs | vHBA Templates

Name: techpubs1 **Number of Ports:** 2
QoS: 2g_4g **Allocation Policy:** roundrobin
Description:

Edit

3. **In the Cloud Details frame, click the vHBAs tab.**

4. On the vHBAs tab, select the vHBA you want to receive the Storage Cloud changes. This step activates the Re-Apply Cloud Changes to vHBA button.



5. Click the button to apply the changes to the vHBA in the template.
6. Click Yes to confirm the changes.
The Storage Cloud changes are pushed to the vHBA.

▼ Display the Fibre Channel Ports in a Storage Cloud

Each Storage Cloud contains one or more Fibre Channel ports that provide the termination point for vHBAs connected to the cloud. For each Storage Cloud, you can display a list of all the Fibre Channel ports in a particular cloud. In addition to displaying the ports in the cloud, you can:

- Remove Fibre Channel ports from the Storage Cloud
- Control the port allocation policy for each Fibre Channel port in the Storage Cloud

Note - Fibre Channel ports are listed in a modified ascending order. So, for example, port iowa1/1 is listed before port iowa 2/1 as expected. However, port iowa1/10 and iowa11/1 would also be listed before port iowa2/1.

Due to the listing order in the FC Ports tab, you will need to pay attention to which port you select as the next available by moving the arrows around. For example, if you want the round robin policy to start with port iowa1/1 and the next port to be assigned as port iowa2/1, you cannot just move the arrow to the next port down in the list because this might be port iowa11/1 (not iowa1/2). Without paying close attention to the listing order, you might accidentally assign the wrong port(s). As a result, the round robin policy will operate correctly, but do so based on the incorrect ports you inadvertently assigned.

To display the Fibre Channel ports, follow this procedure:

1. **On the Navigation Frame, select Storage Cloud Manager → Storage Clouds. See “[Display Summary Information for Storage Clouds](#)” on page 143.**
2. **Select a Storage Cloud to populate the details frame with its properties.**
3. **Click the FC Ports tab to display the Fibre Channel ports associated with the Storage Cloud.**

▼ Display the vHBAs Associated With a Storage Cloud

When vHBAs are either associated with an I/O Profile or deployed to a host server, those vHBAs are displayed through the vHBAs tab on the Storage Cloud details frame. The vHBAs tab displays only vHBAs that are actually used in an I/O Profile regardless of whether or not the I/O Profile is connected to a host server, not the vHBAs that are connected to a Storage Cloud. As a result, it is possible that the vHBAs tab is empty even though a Storage Cloud actually has vHBAs connected to it.

To display the vHBAs associated with a Storage Cloud, follow this procedure:

1. **Display the Storage Cloud Details frame.**
2. **Click the vHBAs tab.**

▼ Apply Storage Cloud Changes to vHBAs in the Cloud

Properties from the Storage Cloud can be inherited by the vHBAs that attach to that cloud. For example, LUN Masks and SAN QoS profiles are properties for the Storage Cloud that can be

inherited by the vHBAs that attach to it. For more information, see [“Understanding Storage Clouds” on page 131](#).

If a Storage Cloud's properties have been edited after vHBAs are attached. The new properties for the Storage Cloud are not automatically inherited by the currently attached vHBAs. Instead, you must manually push those changes to the vHBAs that are currently attached to the cloud. New vHBAs that are attached after the cloud's properties were changed will receive the new properties, but currently connected vHBAs do not receive the new cloud properties until you apply those changes. You can apply Storage Cloud changes to currently attached vHBAs by using the Apply Storage Cloud changes button on the Storage Cloud details frame.

To apply changed cloud properties to currently attached vHBAs, follow this procedure:

1. **Display the Storage Cloud details frame.**
2. **Click the vHBAs tab.**
3. **Select the vHBA to which you want the Storage Cloud's changes pushed.**
This step activates the Apply Cloud Changes to vHBAs button.
4. **Click the Apply Cloud Changes to vHBAs button to push all the new Storage Cloud properties to the currently selected vHBA.**
After a brief period of time, the vHBA has the new cloud properties.

▼ Terminate vHBAs in a Storage Cloud to Another Storage Cloud

Some management tasks might require you to move a vHBA to a different Storage Cloud. For example, if you need to run diagnostics on a Fibre Channel port, or perhaps service an I/O Module. In such cases, if other Storage Clouds are available and have available Fibre Channel ports, you can change the termination of a vHBA from one cloud to another through Oracle Fabric Manager. If you change the termination of a vHBA with live traffic, a service interruption will occur.

To terminate a vHBA on another Storage Cloud, follow this procedure:

1. **Display the Storage Cloud details frame.**
2. **Click the vHBAs tab.**
3. **Select the vHBA that you want to terminate on a different Storage Cloud.**
This step activates the Terminate vHBAs on another Cloud button.
4. **Click the Terminate vHBAs on another Cloud button to change the vHBA to a new Storage Cloud.**

5. **From the resulting dialog, select the Storage Cloud on which you want to terminate the selected vHBA.**

This step activates the Submit button.

6. **Click Submit to reterminate the selected vHBAs.**

After a brief period of time, the vHBAs are terminated on available Fibre Channel ports in the new Storage Cloud.

▼ **Display the vHBA Templates Associated With a Storage Cloud**

Storage Clouds are part of I/O Templates, which in turn, are used in creating the I/O Profiles that are deployed to servers. Through the vHBAs Templates tab, Oracle Fabric Manager supports a way of seeing all the vHBA Templates that are associated with a specific Storage Cloud. This tab shows the association between vHBAs in a template and the Storage Cloud regardless of the state of the vHBAs being bound or not.

The vNIC Templates tab shows information only. You cannot manage the vNIC Templates associated with a specific Storage Cloud.

To display the vHBA Templates for a specific Storage Cloud, follow this procedure:

1. **Display the Storage Cloud details frame.**
2. **Click the vHBA Templates tab.**

Working With Default Gateways

This chapter documents the following topics:

- [“Default Gateways Overview” on page 153](#)
- [“Displaying Default Gateways” on page 153](#)

Default Gateways Overview

A default gateway allows host servers to forward packets with unknown destination addresses off of the server's local network. When this feature is configured, the Oracle Fabric Device gets a pool of default gateways for all the servers attached to the Fabric Device, and with this pool of default gateways the Fabric Device can support forwarding packets to and from the various networks that hosts use. Each Default Gateway must be manually specified for it to be added to the Default Gateways Summary.

Displaying Default Gateways





This section contains the following topics:

- [“Default Gateways Summary” on page 153](#)
- [“Create a Default Gateway” on page 154](#)
- [“Delete a Default Gateway” on page 156](#)

Default Gateways Summary

Through the Default Gateways Summary, you can display or configure one or more default gateways, which enables host servers to forward unknown packets to vNICs. Each Default Gateway configured in Oracle Fabric Manager is displayed through the Default Gateways Summary.

Create a Default Gateway

Default Gateways Summary							
Name	Status	Address	DNS Address	Domain	I/O Templates	I/O Profiles	Description
miathe1		12.12.12.12	127.0.0.1		0	0	
miathe2		0.0.0.0	0.0.0.0		0	0	
miathe3		12.12.12.12	12.12.12.13	domainName	0	0	comment
miathe4		20.20.20.20	44.44.44.44	foobar	0	0	
renoGate		192.168.1.100	192.168.1.1	de.lab.xsigo.com	0	0	miatheXMS

5 items 

Notice that default gateways can be configured or deleted through the Default Gateways Summary.

Field...	Indicates...
Name	The default gateway's name.
Status	An icon that determines whether a problem exists (the warning triangle) or the default gateway is configured correctly (the green circle).
Discovered From	The name or IP address of the Fabric Device from which the default gateway was discovered.
Address	The gateway router's IP address. The default address is 0.0.0.0.
DNS Address	The DNS server's IP address. The default address is 0.0.0.0.
Domain	The domain in which the default gateway is configured.
I/O Templates	The number of I/O Templates that are assigned to hosts that are using each default gateway.
Description	The optional description assigned to the default gateway.

▼ Create a Default Gateway

A default gateway is configured for each host server network attached to the Fabric Device. By doing, you specify the default gateway address for all host servers on that network, and packets will be forwarded off of the server networks as needed.

Default gateways are configured manually, and if you want a specific default gateway assigned to a specific server(s), you can associate the default gateway with an I/O Template, create an I/O Profile from the I/O Template, then deploy that I/O Profile to the server.

Default gateway creation is supported through the Create Default Gateway dialog. To create a default gateway, follow this procedure:

- 1. On the navigation frame, select Default Gateways to display the Default Gateways summary.**

See [“Displaying Default Gateways” on page 153](#) for an example.

Notice that you can add or delete default gateways by using the plus sign and garbage can icons, respectively.

2. **Click the plus sign to display the Create Default Gateway dialog.**



The screenshot shows a dialog box titled "New Default Gateway" with a close button in the top right corner. The dialog contains the following fields and values:

Field	Value
Name: *	pubstest
IP Address: *	192.168.113.101
DNS Server:	0.0.0.0
Domain Name:	all
Description:	

At the bottom of the dialog are two buttons: "Submit" and "Cancel".

3. **In the Name field, enter the name you are assigning to the default gateway.**
4. **In the IP Address field, enter the IP Address of the default gateway.**
5. **In the DNS Address field, enter either the IP address or the fully qualified name of the DNS server for the default gateway.**
6. **In the Domain Name field, enter the name of the domain in which the default gateway is configured.**
7. **As an option, in the Description field, enter a short description for the default gateway.**
8. **Click Submit to configure the default gateway information. A confirmation dialog is displayed to verify that you want to configure the default gateway.**
9. **On the confirmation dialog, click Yes.**

▼ Delete a Default Gateway

Any default gateway that is configured on an Oracle Fabric Device can be deleted through the Default Gateways Summary. When a default gateway is deleted, the vNICs can no longer forward traffic to destinations on a subnet that is different than the server where the vNIC is deployed.

To delete a default gateway, follow this procedure:

1. **On the navigation frame, select Default Gateways to display the Default Gateway summary.**

See [“Displaying Default Gateways” on page 153](#) for an example.

2. **Select the default gateway(s) that you want to delete, then click the garbage can icon.**



Default Gateways Summary	
Name ▲	Status
mlathe1	
mlathe2	
mlathe3	
mlathe4	

3. **On the confirmation dialog, click Yes to delete the selected default gateway.**

Working With the Fabric Device

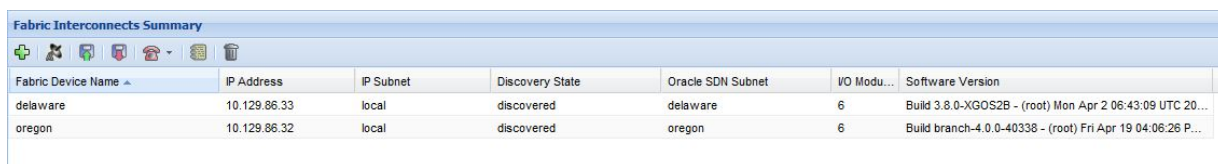
This chapter documents the following topics:

- [“Fabric Interconnects Summary” on page 157](#)
- [“Scan for Fabric Interconnects” on page 158](#)
- [“Discover Fabric Interconnects With Oracle Fabric Manager” on page 160](#)
- [“Displaying Fabric Interconnect Details” on page 161](#)
- [“Unmanage a Fabric Interconnect” on page 194](#)
- [“Back Up and Restore a Fabric Interconnect Configuration” on page 195](#)
- [“Performing Tech Support Functions on the Fabric Interconnect” on page 197](#)
- [“Collect Fabric Interconnect Log Files” on page 202](#)

Fabric Interconnects Summary

Information about the Oracle Fabric Interconnects that are managed through Oracle Fabric Manager is contained on the Fabric Interconnects Summary. Through the Fabric Interconnects Summary, you can select a specific Oracle Fabric Interconnect and display more details for it by using the Fabric Interconnect Details frame. For more information about Fabric Interconnect Details, see [“Displaying Fabric Interconnect Details” on page 161](#).

The Fabric Interconnects Summary is available through the navigation panel by selecting Managed Devices->Fabric Interconnects.



Fabric Device Name	IP Address	IP Subnet	Discovery State	Oracle SDN Subnet	I/O Modu...	Software Version
delaware	10.129.86.33	local	discovered	delaware	6	Build 3.8.0-XGOS2B - (root) Mon Apr 2 06:43:09 UTC 20...
oregon	10.129.86.32	local	discovered	oregon	6	Build branch-4.0.0-40338 - (root) Fri Apr 19 04:06:26 P...

Column	Contains
Fabric Interconnect Name	The name of each managed Fabric Interconnect in the network. The name is a link to the Details frame for the Oracle Fabric Interconnect. For more information about the Fabric Interconnect Details frame, see “Displaying Fabric Interconnect Details” on page 161 .

Column	Contains
IP Address	The Fabric Device IP address of each managed Fabric Interconnect.
IP Subnet	The IP subnet on which each Fabric Interconnect is discovered. The subnet is either a name (for example, "local") or displayed in dotted decimal notation.
Discovery State	<p>The Fabric Interconnect's current state in Oracle Fabric Manager. The state takes in the format state1/state2.</p> <p>Valid states include:</p> <ul style="list-style-type: none"> ■ Discovered ■ Unmanaged ■ Upgrading ■ Up/Initializing ■ Up/Indeterminate ■ Up/VersionMismatch ■ Up/HostControlPathDown ■ Up/Down ■ Up/Up
Oracle SDN Subnet	<p>Indicates the name or InfiniBand ID of the Fabric Interconnect that is the master node, which provides the Infiniband subnet.</p> <p>If this value is the same as the Fabric Interconnect's name, that Fabric Interconnect is providing the IB subnet manager.</p> <p>If this value is different than the Fabric Interconnect's name, the IB subnet manager in use is provided by the device named in this field. It is possible that a Fabric Interconnect named in the Fabric Subnet field that is not managed by Oracle Fabric Manager is providing subnet manager functionality for a Fabric Interconnect that is managed by Oracle Fabric Manager.</p>
I/O Modules	The total number of I/O Modules installed in each Fabric Interconnect.
Software Version	The version of Oracle's XgOS currently installed on each managed Fabric Interconnect.

▼ Scan for Fabric Interconnects

Before discovering one or more Fabric Interconnects, you can perform a scan process that queries for Fabric Interconnects in the network. The scan broadcasts to the Fabric Interconnect's management IP address, and listens for a response. The query and acknowledgement occurs over the discovery subnet, and can occur for Fabric Interconnects on the local network, or through a proxy Fabric Interconnect on a remote network. (For more information about discovery subnets, see [“Working With Discovery Subnets” on page 395.](#)) When a Fabric Interconnect responds, the management IP address is learned and added to the Fabric Interconnects Summary Page.

When a scan occurs, the listed Fabric Interconnects are not actually added to Oracle Fabric Manager for management. Instead, they are simply listed as available. The responding Fabric Interconnects are put into a list of all the Fabric Interconnects in the network that can be managed through Oracle Fabric Manager.

You can also scan through the Scan button on the Oracle Fabric Manager on the Overview page's Common Tasks section.

To scan for Fabric Interconnects, follow this procedure:

1. **Display the Fabric Interconnects Summary.**
See [“Fabric Interconnects Summary” on page 157](#) for an example.
2. **On the Task Board, click the Scan... button to display the results of the Fabric Interconnect scan.**

The scan found the following Fabric Interconnects

Fabric Interconnect Name	Subnet	IP Address
delaware.lab.xsigo.com	local	10.129.86.33
iplat5.lab.xsigo.com	local	10.129.86.97
maine-hw2.us.oracle.com	local	10.129.86.88
ontario.lab.xsigo.com	local	10.129.86.31
oregon.lab.xsigo.com	local	10.129.86.32
ovn86-86.us.oracle.com	local	10.129.86.86
puertorico.lab.xsigo.com	local	10.129.86.93
skylab.lab.xsigo.com	local	10.129.86.95

8 items

Submit Cancel

3. **You can discover one or more of the scanned Fabric Interconnects by selecting the Fabric Interconnect(s) in the table, and clicking Submit.**
The Login dialog is displayed.
4. **In the User Name field, enter a valid user that can log in to the Fabric Interconnect as an administrator.**
The user account must be able to login to the Fabric Interconnect to be able to add the Fabric Interconnect to Oracle Fabric Manager.
5. **In the Password field, enter the password for the user account entered in the User Name field.**

6. **Click Submit to discover the scanned Fabric Interconnect and add it to Oracle Fabric Manager.**

When the Fabric Interconnect is successfully added to Oracle Fabric Manager, you can begin configuring and managing Oracle virtual I/O and its features.

▼ Discover Fabric Interconnects With Oracle Fabric Manager

When Fabric Interconnects are added to the network, they are discovered through either their Fabric Interconnect IP address or host name, and then added to Oracle Fabric Manager. The process of discovery actually adds the Fabric Interconnect into Oracle Fabric Manager so that a discovered Fabric Interconnect is one that is actively being managed by Oracle Fabric Manager. This is different than a scanned Fabric Interconnect, which simply provides a list of known Fabric Interconnects that can be managed.

When discovery occurs, the Oracle Fabric Manager server communicates with the Fabric Interconnect's system control processor (SCP) to learn about the Fabric Interconnect, take an inventory of each Fabric Interconnect's hardware, learn the software version installed and the virtual resources configured on the Fabric Interconnect (if any). After a Fabric Interconnect is discovered, Oracle Fabric Manager adds it to the Fabric Interconnects Summary. After the Fabric Interconnect is displayed in the Fabric Interconnects Summary, it can be managed through Oracle Fabric Manager.

Oracle Fabric Manager requires the management IP address to open a communication path for learning the Fabric Interconnect contents. Do not use a vNIC IP address for discovery.

To discover a Fabric Interconnect, follow this procedure:

1. **Display the Fabric Interconnects Summary.**
See [“Fabric Interconnects Summary” on page 157](#) for an example.
2. **On the toolbar, click the Manage a Fabric Interconnect... button.**

The Manage Fabric Interconnect dialog is displayed.



The screenshot shows a dialog box titled "Manage a Fabric Interconnect". It contains the following fields and values:

- IP Address/DNS Name: * iowa
- Subnet name: local
- User Name: * admin
- Password: *

At the bottom right, there are two buttons: "Submit" and "Cancel".

3. In the IP Address/DNS Name field, enter the Fabric Interconnect's IP address or name.
4. In the Subnet Name field, enter name of the subnet on which the Fabric Interconnect you want to discover is currently located. By default, the subnet is the local subnet.
5. In the User Name, specify the admin account that you will use to log in to the Fabric Interconnect for the purpose of discovering it. The user name can be anything, but by default it is administrator, root, or ofmadmin.
6. In the Password field, enter the password for the user name that you specified in the preceding field.
7. When the Fabric Interconnect discovery information is complete, click Submit to discover the Fabric Interconnect.

Displaying Fabric Interconnect Details

This section contains the following topics:

- [“Fabric Interconnect Summary” on page 162](#)
- [“Fabric Interconnect General Information” on page 163](#)

- [“Ethernet Card Information” on page 164](#)
- [“Set the Allowed VLAN Ranges for an Ethernet Port” on page 165](#)
- [“Fibre Channel Card Information” on page 167](#)
- [“Fabric Interconnect User Information” on page 168](#)
- [“Configuring Local Users” on page 169](#)
- [“Phone Home Information” on page 169](#)
- [“Display SNMP Properties” on page 173](#)
- [“Display SNMP Secure Users” on page 174](#)
- [“Display SNMP Trap Destinations” on page 177](#)
- [“Configure SNMP Trap Destinations” on page 178](#)
- [“Fan Information” on page 179](#)
- [“Power Supply Information” on page 180](#)
- [“IMS Information” on page 181](#)
- [“Active Directory Information” on page 183](#)
- [“RADIUS Servers Information” on page 188](#)
- [“RADIUS Users Information” on page 192](#)

Fabric Interconnect Summary

In addition to the Fabric Interconnect Summary, Oracle Fabric Manager contains a Fabric Interconnect details frame where additional information is contained. The Fabric Interconnects Details frame is available through the Fabric Interconnects Summary by clicking a Fabric Device name. When you click the Fabric Device name, the Fabric Interconnect details frame is populated with information about the selected Fabric Interconnect.

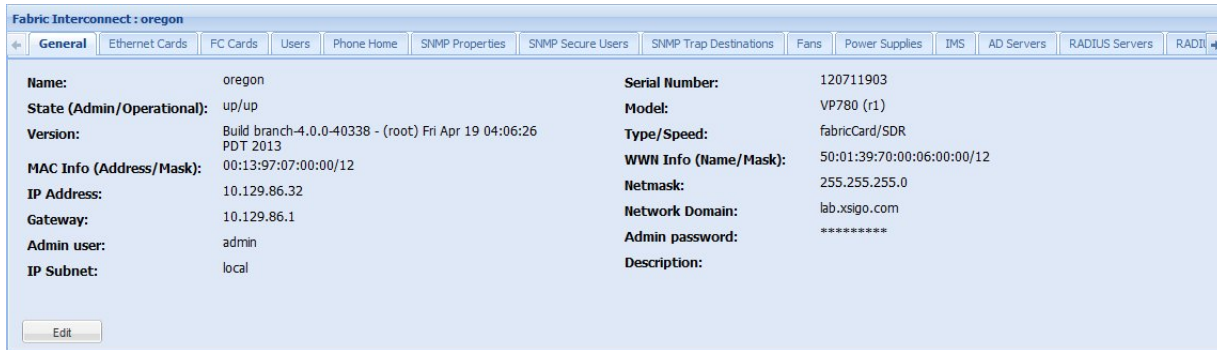
Oracle Fabric Manager supports gathering and inventory and managing both Oracle's Fabric Interconnect F1-15 and the Fabric Interconnect F1-4., and Fabric Interconnect details are available for both models. There are some minor hardware differences between the Fabric Interconnect F1-15 and Fabric Interconnect F1-4, such as a smaller form factor and fewer I/O modules in the Fabric Interconnect F1-4, which is a two rack-unit Fabric Interconnect.

The Fabric Interconnects details frame shows information about a single Fabric Interconnect. The Fabric Interconnects details frame is separated into tabs, which arrange different types of Fabric Interconnect information into sections. Some tabbed pages have additional tables and controls that allow more in-depth configuration and management of a single Fabric Interconnect.

It is important to understand that Oracle Fabric Manager has many of the same types of information as the Fabric Interconnect. However, when you display this information through the Fabric Interconnect, the information is relevant only to the context of that particular Fabric Interconnect.

Fabric Interconnect General Information

The Fabric Interconnect's General tab contains some basic, top-level information about the Fabric Interconnect, its Fabric type and speed, non-I/O Module hardware inventory, management interface, location, and so on.



Column	Contains
Name	The name of the Fabric Interconnect displayed in the details frame.
Serial Number	The unique serial number assigned to the Fabric Interconnect.
State (Admin/Operational)	The current administrative and operational state of the Fabric Interconnect. The administrative state is the state that the Fabric Interconnect should be in, and the operational state is what state the Fabric Interconnect actually is in.
Model	The model number of the Fabric Interconnect.
Version	The current version of XgOS operating system software running on the Fabric Interconnect.
MAC Info (Name/Mask)	The beginning MAC address available in the Fabric Interconnect's embedded MAC address pool. The mask bits determine how many MAC addresses can be assigned from the Fabric Interconnect's MAC address pool. For example, /12 indicates that 12 bits worth of addresses can be assigned incrementally starting with the MAC address listed.
WWN Info (Name/Mask)	The beginning WWN number available in the Fabric Interconnect's embedded WWN pool. The mask bits determine how many WWN numbers can be assigned from the Fabric Interconnect's WWN pool. For example, /12 indicates that 12 bits worth of addresses can be assigned incrementally starting with the WWN number listed.
Netmask	The network mask assigned to the Fabric Interconnect's management address.
IP Address	The Fabric Interconnect's management address.
Network Domain	The network domain in which the Fabric Interconnect is currently deployed.
Gateway	The IP address of the Fabric Interconnect's gateway switch or router.
Admin Password	The currently assigned administrator password for accessing the Fabric Interconnect. For security purposes, the password is displayed as a series of asterisks (*****).
Admin User	The user name of the currently assigned admin user.
Description	An optional string that helps to describe the Fabric Interconnect.

Column	Contains
IP Subnet	The IP subnet for the Fabric Interconnect.

Ethernet Card Information

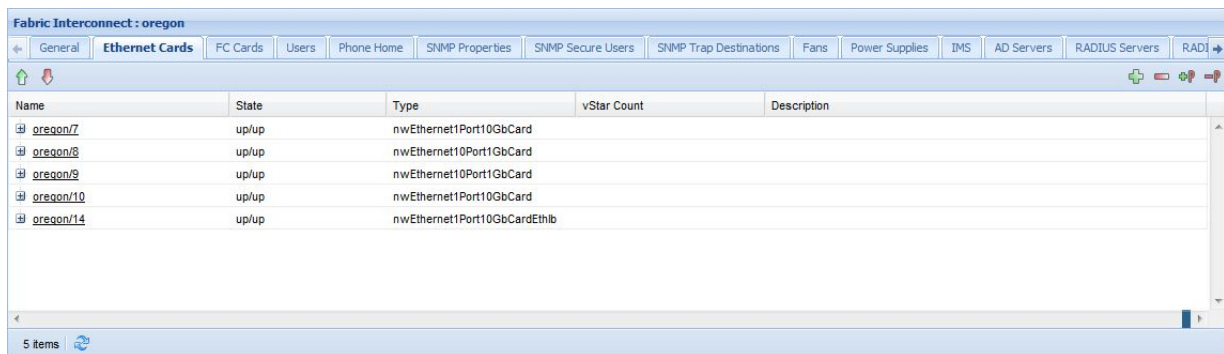
Each Fabric Interconnect supports Ethernet I/O modules that are the underlying physical connectivity to the IP network. These hardware modules then support the creation of vNICs for virtual IP connectivity. The following I/O modules are supported in each Fabric Interconnect:

- 10 Gig Ethernet module (10 GE), which supports 1 physical link at 10 Gbps.
- 4-Port 10 Gig Ethernet module (4-Port 10 GE), which support 4 physical links at 10 Gbps per link.
- 10-Port Gig Ethernet module (10-Port GE), which supports 10 individual physical links at 1 Gbps.

For more information about the Fabric Interconnect's I/O modules, see the XgOS Hardware and Drivers Installation Guide.

In addition, some hardware-specific features, such as Link Aggregation Groups (LAGs) are supported through the Ethernet Cards page. For information, see [“Working With Link Aggregation” on page 281](#).

When a Fabric Interconnect is discovered by Oracle Fabric Manager, the Fabric Interconnect's I/O Modules are inventoried. The I/O Module information for the Ethernet card is displayed through the Ethernet Cards tab on the Fabric Interconnect Details frame.



Column	Contains
Name	The name of each Ethernet card that is inventoried in the Fabric Interconnect. Each I/O card is named with the format Fabric Interconnect/slot, so the card texas/1 is the Ethernet card in slot 1 of the Fabric Interconnect named “Texas.”

Column	Contains
	If a module contains a termination port supporting a vNIC, the module is expandable and collapsible to display its configured Ethernet ports.
	Ethernet port information is also shown for any port that is terminating a vNIC. The port is named with the format Fabric Interconnect/slot/port, so the port texas/1/1 is the Ethernet port 1, on slot 1, of the Fabric Interconnect named "Texas."
State	The current state of the Ethernet card. The state is displayed as administrative state/operational state. Ethernet cards in up/up state are operating correctly.
Type	The type of Ethernet card installed in the Fabric Interconnect: The following modules are supported: <ul style="list-style-type: none"> ■ 10 Gig Ethernet module (10 GE), which supports 1 physical link at 10 Gbps. In the Type field, Oracle Fabric Manager displays this module appears as nwEthernet1Port10GBCardEthIB ■ 4-Port 10 Gig Ethernet module (4-Port 10 GE), which support 4 physical links at 10 Gbps per link. In the Type field, this module appears as nwEthernet4Port10GbCardEthIB ■ 10-Port Gig Ethernet module (10-Port GE), which supports 10 individual physical links at 1 Gbps. In the Type, field, this module appears as nwEthernet10Port1GbCardEthIB
vStar Count	The total number of virtual resources configured on each port on the Ethernet card.
Description	An optional description field for the Ethernet Card.

Notice that the Ethernet Card tab has a list of individual modules, which is expandable to the port level. On the port level, additional features are available, such as general properties, Ethernet properties for the port (for example, MTU size), and Allowed VLAN Range.

▼ Set the Allowed VLAN Ranges for an Ethernet Port

By default, the allowable VLAN range is from 1 to 4095. However, you can set a custom range of VLANs on the port so that only the VLAN-tagged packets within the specified range are allowed to ingress or egress the port. Traffic that has a VLAN tag not in the specified range is blocked from ingress or egress on the port.

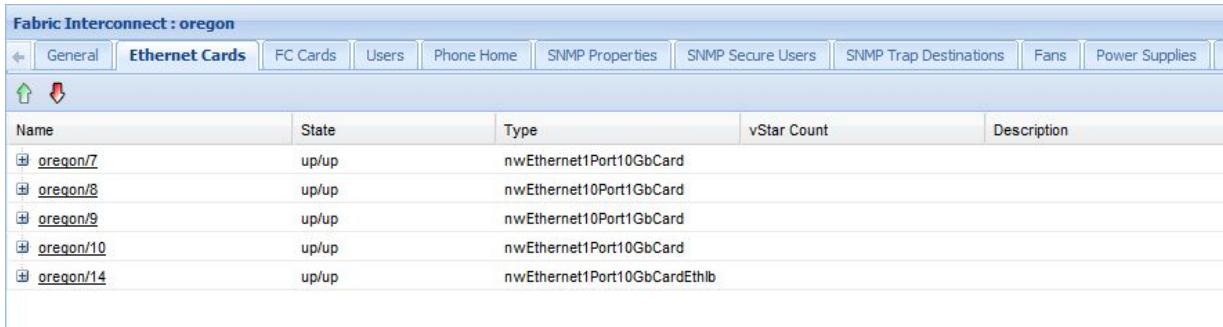
You can set or change the allowed VLAN range for the port through the Allowed VLANs tab at the port level of the Ethernet Card Details Frame.

To set the allowed VLAN range for a specific Fabric Interconnect, follow this procedure:

1. **Select Managed Devices → Fabric Interconnects to display the Fabric Interconnect Summary page.**
2. **Select the Fabric Interconnect that contains the Ethernet Port or LAG for which you want to set the Allowed VLAN range by clicking the Fabric Interconnect in the summary.**

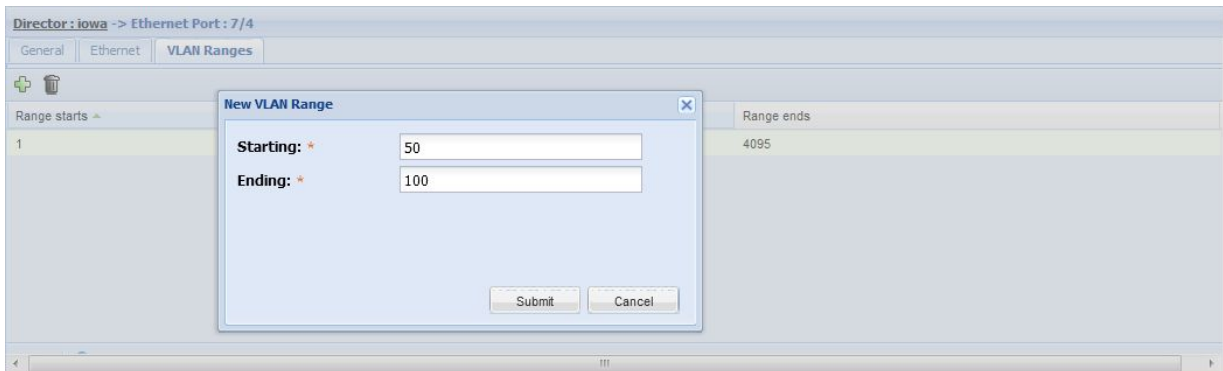
This step populates the details frame with information for that Fabric Interconnect.

3. On the details frame, click the Ethernet Cards tab to display each Ethernet module in the selected Fabric Interconnect.



Name	State	Type	vStar Count	Description
oregon/7	up/up	nwEthernet1Port10GbCard		
oregon/8	up/up	nwEthernet10Port1GbCard		
oregon/9	up/up	nwEthernet10Port1GbCard		
oregon/10	up/up	nwEthernet1Port10GbCard		
oregon/14	up/up	nwEthernet1Port10GbCardEth1b		

4. Click the plus sign for the card on which you want to set the VLAN range.
This step expands the card to display individual ports.
5. Click the port (which is indented under the card) to display the port properties in the details frame.
6. Click the VLAN Ranges tab.
7. Click the plus sign (+) to display the New VLAN Range dialog.



8. In the Starting field, enter the first VLAN ID that you want to be available.
9. In the Ending field, enter the last VLAN ID you want to be available.
10. When the Allowed VLAN Range is configured, click Submit.

Fibre Channel Card Information

Each Fabric Interconnect supports Fibre Channel modules that are the underlying physical connectivity to an FC SAN. These hardware modules then support the creation of vHBAs for host servers' SAN connectivity. The Fabric Interconnect supports the following Fibre Channel cards:

- The Line Rate Fibre Channel module (2x4 FC Line Rate card). This module supports 2 Fibre Channel port pairs (one transmit port of 4 Gbps and one receive port at 4 Gbps for each port pair) for a total of 8 Gbps per module.
- The 8 Gbps Line Rate Channel module (2x8 FC Line Rate card). This module supports 2 Fibre Channel port pairs (one transmit port of 8 Gbps and one receive port at 8 Gbps for each port pair) for a total of 8 Gbps (full line rate) for each port pair.

For more information about the Fabric Interconnect I/O modules, see the Fabric Interconnect Hardware and Drivers Installation Guide.

When a Fabric Interconnect is discovered by Oracle Fabric Manager, the Fabric Interconnect's I/O Modules are inventoried. The I/O Module information for the Fibre Channel card is displayed through the FC Cards tab on the Fabric Interconnect Details frame.

Fabric Interconnect : oregon									
General	Ethernet Cards	FC Cards	Users	Phone Home	SNMP Properties	SNMP Secure Users	SNMP Trap Destinations	Fans	Power Supplies
Name	State	Type	vStar Count	Description					
oregon/11	up/resourceMissing	sanFc2Port4GbCard							

Column	Contains
Name	<p>The name of each Fibre Channel card that is inventoried in the Fabric Interconnect. Each I/O card is named with the format Fabric Interconnect/slot, so the card <code>texas/1</code> is the Fibre Channel card in slot 1 of the Fabric Interconnect named "Texas."</p> <p>If a module contains a termination port supporting a vHBA, the module is expandable and collapsible to display its configured FC ports.</p> <p>Fibre Channel port information is also shown for any port that is terminating a vHBA. The port is named with the format Fabric Interconnect/slot/port, so the port <code>texas/1/1</code> is the FC port 1, on slot 1, of the Fabric Interconnect named "Texas."</p>
State	The current state of the Fibre Channel card. The state is displayed as administrative state/operational state. Fibre Channel cards in up/up state are operating correctly.
Type	The type of Fibre Channel card installed in the Fabric Interconnect. The following modules are supported:

Column	Contains
	<ul style="list-style-type: none"> ■ the Line Rate Fibre Channel module, which supports 2 Fibre Channel port pairs (one transmit port of 4 Gbps and one receive port at 4 Gbps for each port pair) for a total of 8 Gbps per module. In the Type field, this module appears as sanFc2Port4GbLrCard ■ the 8 Gbps Line Rate Fibre Channel module, which supports 2 Fibre Channel port pairs (one transmit port of 8 Gbps and one receive port at 8 Gbps for each port pair) for a total of 8 Gbps (full line rate) for each port pair. In the Type field, this module appears as sanFc2Port8GbLrCardEthIB
vStar Count	The total number of virtual resources configured on each port on the FC card.
Description	An optional description field for the FC Card.

Fabric Interconnect User Information


Oracle Fabric Manager supports local user accounts. Local user accounts are local because they are independent accounts that exist on individual Fabric Interconnects.

Typically, a local user account is mapped into a role group. That role group provides the permissions for what objects the user can and cannot write to. When a user account is assigned to a role, the conditions are created for what actions the user can perform, and what objects the user can configure, change, or manage.

It is possible to create a local user without assigning the user to a specific role group. In this case, the user is assigned to the “operator” role group by default. The operator role group is the most restrictive. It provides read-only access.

All users have a role group—either a role group that is explicitly assigned to the user, or the operator role group that is assigned by default.

At the Fabric Interconnect level, local user accounts are used. You can manage local user accounts through the Users tab on the Fabric Interconnects Details frame.



You can add local users to the Fabric Interconnect by clicking the plus sign, and you can delete a configured user by selecting it in the Users tab, then clicking the garbage can icon.

▼ Configuring Local Users

Local user accounts define what users can log in to the Fabric Interconnect and what privileges that user has while logged in.

To configure local user accounts, follow this procedure:

1. **On the Users tab, click the plus sign to display the Create a New User dialog.**
2. **In the User Name field, enter the name for the local user account that you are creating.**
3. **As an option, in the Description field, enter an alphanumeric character string that describes the user account.**
4. **From the Roles dropdown menu, select the appropriate role for the user account.**

The user's privileges on the Fabric Interconnect are a result of the role that is assigned to the user account. The following table shows the roles and summarizes what privileges belong to each role.

Role	Privileges for...
operator	read access to Fabric Interconnect features
network	vNIC configuration and management, Network QoS
storage	vHBA configuration and management, SAN QoS
server	compute resource configuration and management
administrator	full admin responsibilities

5. **When the user account is configured, click Submit.**
6. **Check the Users tab to verify the user account is configured correctly.**

Phone Home Information

The Oracle Phone Home feature supports periodically transmitting the following information to Oracle Technical Support:

- The contents of selected log files
- The output of the show tech-support command from the XgOS CLI

By transmitting this information, Oracle Technical Support can proactively look for and diagnose potential problems without requiring you to collect data, package it, and transmit it to Oracle.

No sensitive customer data is gathered and transmitted to Oracle. To ensure that private information is kept safe, the Oracle Phone Home feature provides ways to:

- Send a copy of the information to an internal website for auditing purposes
- Remove private data, such as IP addresses, from the data sent to Oracle Customer Support

Also, the data is transmitted in an encrypted form so that it cannot easily be read.

The following table shows defaults for Phone Home. Most of these parameters are determined through the Oracle First Boot script which automatically runs when you install the Fabric Interconnect for the first time.

However, you can customize the Phone Home properties and transmission schedules through the Phone Home tab of the Fabric Interconnects Details frame.



Phone Home Parameter	Default
proxy	Oracle Phone Home communication occurs to port 3128.
notify	No notification when Oracle Phone Home communication occurs
strip-private	IP addresses are removed from the Oracle Phone Home information and replaced with a <privip:x> string—for example, <privip:1> <privip:2>
mode	Periodic, transmissions occur as scheduled by you

▼ Edit Phone Home Properties

To edit the Phone Home properties, follow this procedure:

1. **Display the Fabric Interconnects Details frame by selecting Managed Devices → Fabric Interconnects.**

2. On the Fabric Interconnects Details frame, click the Phone Home tab to display the Phone Home properties configured on the Fabric Interconnect.
3. On the Phone Home tab, click the Edit button to unlock the editable phone home parameters.

The screenshot shows the 'Phone Home' configuration page for a Fabric Interconnect named 'oregon'. The page has a navigation bar with tabs for General, Ethernet Cards, FC Cards, Users, Phone Home (selected), SNMP Properties, SNMP Secure Users, SNMP Trap Destinations, Fans, Power Supplies, IMS, AD Servers, RADIUS Servers, and RADI. The main content area is divided into several sections:

- Phone Home Enabled:** Radio buttons for True and False. 'False' is selected.
- Strip Private:** Radio buttons for True and False. 'True' is selected.
- Frequency:** A dropdown menu currently set to 'Weekly'.
- Notification:** Radio buttons for True and False. 'False' is selected.
- Send Alarms:** Radio buttons for True and False. 'True' is selected.
- HTTP Proxy Host:** An empty text input field.

At the bottom left, there are 'Submit' and 'Cancel' buttons.

4. As needed, enable or disable scheduled transmission of phone home data:
 - True to enable scheduled transmission.
 - False to disable scheduled transmission.
5. From the Notification section, determine whether you want Oracle Technical Support to contact you whenever Phone Home information is received:
 - True causes Oracle to contact the person in the Contact Name field by any method listed in the Contact Phone Number or Contact Email fields.
 - False causes Oracle to not contact you, even if contact information exists in the Phone Home configuration.
6. From the Strip Private section, select whether you want private IP address information removed from the Phone Home information:
 - True causes the IP address information to be represented with <privip:X> labels.
 - False causes your IP addresses to be visible in the Phone Home information.
7. From the Send Alarms section, select whether you want Phone Home information sent whenever a major alarm is raised in the network:

- **True causes the Phone Home information to be sent whenever a major alarm occurs.**
 - **False causes Phone Home information to be transmitted as configured, and major alarms will not trigger a Phone Home message.**
8. **From the Frequency dropdown menu, select the periodicity at which the Fabric Interconnect will transmit Phone Home data.**
 9. **If you are configuring HTTP Proxy, in the HTTP Proxy Host field, enter the name of the HTTP Proxy Host.**
 10. **If you are configuring HTTP Proxy, in the HTTP Proxy Port field, enter the port number on which the HTTP Proxy will communicate Phone Home information.**
By default, port 3128 is used, but you can set any other port that is not supporting a service.
 11. **If you are configuring HTTP Proxy, in the HTTP Proxy User field, enter the user name for the HTTP Proxy.**
 12. **In the URL field, enter the URL of where the Phone Home information will be sent.**
By default, this field contains the default site for Oracle Technical Support.
 13. **In the Copy To URL field, enter the URL of an audit server in your network to which you want a copy of any Phone Home message sent.**
 14. **In the Contact Customer Name field, enter the name of someone on site that manages the Fabric Interconnect.**
This person is who Oracle will contact if the “Notifications” section is set to “True” and a Phone Home message is received.
 15. **In the Contact Customer Phone field, enter the phone number (if any) for the person named in the Contact Customer Name field.**
 16. **In the Contact Customer Email field, enter the email address (if any) for the person named in the Contact Customer Name field.**
 17. **Click Save to configure Phone Home on the Fabric Interconnect.**

▼ Display SNMP Properties

Oracle Fabric Manager tracks general properties for SNMP on each Fabric Interconnect. These properties are available through the SNMP Properties tab on the Fabric Interconnects Details frame.



The screenshot shows a web interface for a Fabric Interconnect named 'oregon'. The 'SNMP Properties' tab is selected, displaying the following configuration:

Description:	test
Read Community:	public
System Name:	VP780 (r1)
System Location:	
System Contact:	

An 'Edit' button is located at the bottom left of the configuration area.

You can also edit the SNMP Properties to set or change current values. To edit the SNMP properties, follow this procedure:

1. **Display the Fabric Interconnects Details frame by selecting Managed Devices → Fabric Interconnects.**
2. **On the Fabric Interconnects Details frame, click the SNMP Properties tab to display the configured SNMP properties.**

3. On the **SNMP Properties** tab, click **Edit** button to unlock the editable **SNMP** properties.

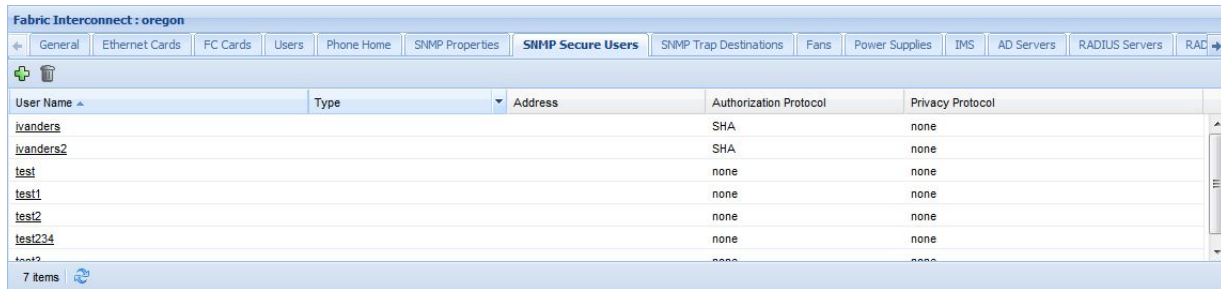


4. Edit the **SNMP** fields as needed, then click **Submit**.
5. Check the **SNMP Properties** tab to verify that the correct properties are configured.

▼ Display SNMP Secure Users

Oracle Fabric Manager supports Secure SNMP. If you are configuring Secure SNMP, you must configure the SNMP user and password which the security algorithms will use as inputs. When the user name and password are specified, you can select the type of security to use. As an option, you can use simple password authentication by not specifying any security parameters.

SNMP secure users for each Fabric Interconnect are available through the **SNMP Secure Users** tab of the Fabric Interconnects Details frame.



The screenshot shows the configuration page for a Fabric Interconnect named 'oregon'. The 'SNMP Secure Users' tab is active, displaying a table of configured users. The table has five columns: 'User Name', 'Type', 'Address', 'Authorization Protocol', and 'Privacy Protocol'. There are seven rows of data, including users named 'ivanders', 'ivanders2', 'test', 'test1', 'test2', 'test234', and 'test2'. The 'Authorization Protocol' for 'ivanders' and 'ivanders2' is 'SHA', while for the others it is 'none'. The 'Privacy Protocol' is 'none' for all users. A plus sign icon is visible in the top left of the table area, and a trash can icon is in the top right. A status bar at the bottom left indicates '7 items'.

User Name	Type	Address	Authorization Protocol	Privacy Protocol
ivanders			SHA	none
ivanders2			SHA	none
test			none	none
test1			none	none
test2			none	none
test234			none	none
test2			none	none

You can also add SNMP secure users by clicking the plus sign, or delete secure users by selecting one or more of them in the SNMP Secure Users tab, then clicking the garbage can icon.

To add SNMP Secure Users to a Fabric Interconnect, follow this procedure:

1. **Display the Fabric Interconnects Details frame by selecting Managed Devices → Fabric Interconnects.**
2. **On the Fabric Interconnects Details frame, click the SNMP Secure Users tab to display the trap destinations configured on the Fabric Interconnect.**

3. On the SNMP Secure Users tab, click the plus sign to display the Create a New SNMP Secure User dialog.

The screenshot shows a dialog box titled "Create a new SNMP Secure User". It contains the following fields and controls:

- User Name:** * (required), text input field containing "paulw".
- Authorization Protocol:** dropdown menu with "SHA (Preferred)" selected.
- Authorization Password:** * (required), empty text input field.
- Privacy Protocol:** dropdown menu with "AES128 (Preferred)" selected.
- Privacy Password:** * (required), empty text input field.
- Description:** empty text area.
- Buttons:** "Submit" and "Cancel" buttons at the bottom right.

4. In the User Name field, enter the username that will be used to log in to the trap destination.
5. From the Authorization Protocol dropdown menu, select the type of authorization that will be used for logging in to the trap destination:
 - None for no authorization
 - MD5 to use MD5 hashing algorithm
 - SHA to use Secure Hashing Algorithm

6. In the **Authorization Password** field, enter the password that will be used by the **SNMP secure user**.
7. From the **Privacy Protocol** dropdown menu, select the protocol to be used:
 - **None** for no authorization
 - **DES** to use **Data Encryption Standard**
8. In the **Privacy Password** field, enter the password that will be used.
9. As an option, in the **Description** field, enter an optional description string for this secure user.

▼ Display SNMP Trap Destinations

Oracle Fabric Manager supports simple network management protocol (SNMP). The Oracle implementation of SNMP supports SNMPv1, v2, and v3. Get, getnext, and getbulk operations are all supported. Set operations are not supported. Community strings are read only. Some standard Enterprise MIBs as well as Oracle proprietary MIBs are supported.

Through SNMP you can configure trap hosts (trap destinations) that will receive events and errors if they occur. Oracle Fabric Manager supports configuring SNMP variables also, such as system ID strings.

To configure SNMP trap destinations, follow this procedure:

1. **Display the Fabric Interconnects Details frame by selecting *Managed Devices* → *Fabric Interconnects*.**
2. **On the Fabric Interconnects Details frame, click the *SNMP Trap Destinations* tab to display the trap destinations configured on the Fabric Interconnect.**

Fabric Interconnect : oregon		
General Ethernet Cards FC Cards Users Phone Home SNMP Properties SNMP Secure Users SNMP Trap Destinations Fans Power Supplies IMS AD Servers RADIUS Servers RADIUS		
Address	SNMP Community	SNMP Version
10.10.10.12:162	public	SNMPv2

1 item

▼ Configure SNMP Trap Destinations

You can use the SNMP Trap Destinations tab to add or delete trap destinations. Add trap destinations by clicking the plus sign, or delete configured trap destinations by selecting one or more of them, then clicking the garbage can icon.

To add an SNMP trap destination, follow this procedure:

1. **Display the Fabric Interconnects Details frame by selecting Managed Devices → Fabric Interconnects.**
2. **On the Fabric Interconnects Details frame, click the SNMP Trap Destinations tab to display the trap destinations configured on the Fabric Interconnect.**
3. **On the SNMP Trap Destinations tab, click the plus sign to display the Create a New SNMP Trap Destination dialog.**



The screenshot shows a dialog box titled "Create a new SNMP Trap Destination". It contains the following fields and values:

Field	Value
IP Address: *	192.168.129.119
Port: *	162
SNMP Community:	public
SNMP Version:	SNMPv2
User Name: *	
Authorization Protocol:	none
Authorization Password:	
Privacy Protocol:	none
Privacy Password:	

Buttons: Submit, Cancel

4. **In the IP Address field, enter the network address of the trap destination in dotted decimal notation.**

5. **In the Port field, specify the port number on which traps will be sent to the trap destination.**

By default, port 162 is used, but you can set another port as long as it is not supporting other traffic. Valid ports are in the range of 1 to 65535.
6. **In the SNMP Community field, enter the read community string for the trap destination.**
7. **From the SNMP Version dropdown menu, select the version of SNMP is in use in your network:**
 - **SNMPv2**
 - **SNMPv3**
8. **In the User Name field, enter the username that will be used to log in to the trap destination.**
9. **In the Authorization Protocol dropdown menu, select the type authorization that will be used for logging in to the trap destination:**
 - **None for no authorization**
 - **MD5 to use MD5 hashing algorithm**
 - **SHA to use Secure Hashing Algorithm**
10. **In the Authorization Password field, enter the password that will be used by the user specified in User Name.**
11. **From the Privacy Protocol dropdown menu, select the protocol to be used:**
 - **None for no authorization**
 - **DES to use Data Encryption Standard**
 - **AES128 to use Advanced Encryption Standard 128-bit encryption**
12. **In the Privacy Password field, enter the password that will be used.**

Fan Information

When a Fabric Interconnect is managed by Oracle Fabric Manager, each Fabric Interconnect registers its hardware information with the Oracle Fabric Manager server. Part of the hardware

information that Oracle Fabric Manager monitors is the fan state. Each Fabric Interconnect has built-in fans for cooling.

Oracle Fabric Manager tracks the operational state of the fans and if an error is detected, an alarm is posted. Fans are monitored for operational states, such as:

- Operational, up/up
- Non-operational, up/down
- Unknown, indeterminate, which can occur if the fan is operating but not supplying proper cooling to the Fabric Interconnect.

The fan information is available through the Fans tab on the Fabric Interconnect Details frame.

Name	State	Actual Speed	Expected Speed	Deviation	Description
Fan-1/1	up	4560	4200	360	
Fan-1/2	up	4800	4200	600	
Fan-2/1	up	4320	4200	120	
Fan-2/2	up	4560	4200	360	

Column	Contains
Name	The name of each fan unit in the Fabric Interconnect. Each fan unit contains more than one fan, so the fan named 1/1 is for fan unit 1, fan 1 and the fan named 1/2 is for fan unit 1, fan 2 and so on.
State	The current state of the fan. The state is displayed as administrative state/operational state. Fans in up/up state are operating correctly.
Actual Speed	The actual speed in RPMs of each fan.
Expected Speed	The expected speed of the fan based on internal calculations.
Deviation	The difference between the actual speed and the expected speed. This number can be a positive number if a fan is spinning faster than expected, or a negative number if the fan is spinning slower than expected.
Description	An optional description for each fan.

Power Supply Information

When a Fabric Interconnect is managed by Oracle Fabric Manager, each Fabric Interconnect registers its hardware information with the Oracle Fabric Manager server. Part of the hardware information that Oracle Fabric Manager monitors is the power supply state. The number of power supplies (PSUs) is different depending on the model of Fabric Interconnect:

- Oracle's Fabric Interconnect F1-15 two built-in PSUs that condition facility power, numbered one and two.

- Oracle's Fabric Interconnect F1-4 has two built-in PSUs, numbered one and two

Oracle Fabric Manager tracks the operational state of the PSUs and if an error is detected, an alarm is posted. Fans are monitored for operational states, such as:

- Operational, up/up
- Non-operational, up/down
- unknown, indeterminate, which can occur if the PSU is operating, but not at proper speed.

The PSU information is available on the Power Supplies tab of the Fabric Interconnect Details frame.

Name	State	Model	Serial	Vendor Model	Description
PowerSupply-1	up/down	VP780-FRU-PS	TE2675600	CAR1212FPBCX-Y04A	
PowerSupply-2	up/up	VP780-FRU-PS	SH3066200	CAR1212FPCXX-Y04A	

Column	Contains
Name	The name of each power supply unit (PSU) in the Fabric Interconnect. Each Fabric Interconnect contains two power supplies that are redundant and load share when both are operating. If one fails, the Fabric Interconnect can operate on one PSU.
State	The current state of the PSU. The state is displayed as administrative state/operational state. PSUs in up/up state are operating correctly.
Model	The Oracle model number for each PSU.
Serial	The Oracle serial number associated with each PSU.
Vendor Model	The PSUs model number assigned by the fan manufacturer.
Description	An optional description field for the PSU.

IMS Information

The Oracle Identity Management System (IMS) uses some general parameters for synchronizing information between the Oracle Fabric Manager Server and an external authentication system, such as an Active Directory or RADIUS server. Although the IMS properties have sensible default values, you might need to set more appropriate values for your network. You can set the IMS properties for Oracle Fabric Manager through the IMS tab on the Fabric Interconnect Details frame.

Default parameters are displayed, but you can set or change individual parameters by clicking the Edit button to unlock the editable options.

Configure SNMP Trap Destinations

The following table shows the contents of the IMS tab, and explains the tab's contents.

Field...	Means...
Cache Timeout	<p>Specifies the periodicity of flushing the IMS Cache on the Oracle Fabric Manager Server and resynchronizing with the external authentication server. The IMS Cache is encrypted and contains user name, password, and role(s) for all configured users.</p> <p>Can be set to a value between 1 and 1440 minutes. The default is 240 minutes, and setting the Cache Timeout to zero (0) disables cache flushing and resynchronization.</p>
Maps To Root	<p>Specifies where the user account information is located on the authentication server. The location is typically where you configure your users and groups. For Active Directory, you will typically enter users. The default is root.</p>
Search Order	<p>Specifies which IMS entity is checked first for user account information:</p> <ul style="list-style-type: none"> ■ InternalFirst sets the IMS to look at the IMS Server's local users and groups first. If user account information is not found in the internal IMS, then check an external IMS (for example, the AD Server). ■ ExternalFirst sets the IMS to look at the external IMS server first (for example, the AD Server) for user and groups information. If user account information is not found in the external IMS, then check the IMS Server's local users database. <p>These two options are not mutually exclusive, and their use is determined by where user accounts are—either on the Fabric Interconnect or on the AD server. In cases where a user account is configured on both the Fabric Interconnect's internal IMS and an external IMS (the AD Server), the roles and user privileges on the Fabric Interconnect are used.</p>
Server Type	<p>Specifies the type of external authentication is currently in use.</p> <ul style="list-style-type: none"> ■ For AD configuration, this setting is not optional. It must be set to ldap_ad. ■ For RADIUS authentication, this value must be RADIUS.
Token Timeout	<p>Specifies the amount of time that the IMS will wait for authentication to occur before timing out.</p> <p>When a log in attempt occurs, the authentication token is sent to the AD Server or domain controller. This field allows you specify how long the AD Server (or domain controller) will hold the token before closing the login attempt. Valid values are from 1 to 1440 seconds. The default is 5 seconds. Setting the value to zero (0) disables the time out and allows the login attempt to remain in progress indefinitely.</p>

Active Directory Information

This section documents how to configure a Fabric Interconnect in an AD authentication model by using Oracle Fabric Manager so that members of an AD group can log into the Fabric Interconnect after they are authenticated. When you perform the procedure in this section, you will use Oracle Fabric Manager to specify the Fabric Interconnect(s) and other properties associated with AD, such as the AD server IP (or host name), the authentication method and so on. When the Fabric Interconnect is integrated into the AD environment, users will be able to log in to the Fabric Interconnect through ssh or other methods and undergo AD authentication.

Oracle Fabric Manager supports an internal identity management system (IMS) for users and roles, but also supports external IMS functionality through LDAP, Active Directory, Simple Password, or Kerberos authentication.

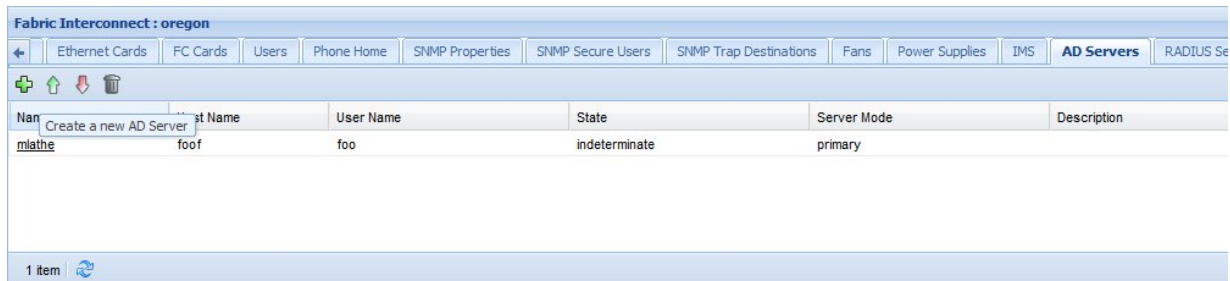
- Simple password authentication provides a basic level of security. An encrypted database of user names and passwords exists on the AD server. When a user logs in through Oracle Fabric Manager, that user is challenged by the IMS (AD) server to provide valid user name and password. When the user enters valid information, the user name and password are sent “in the clear” to the server which looks up the user name and password in the list. If the user name and password are present, the user's credentials are sent. With simple password authentication, the user name and password as well as other content are sent in clear text.
- Kerberos authentication relies on some of the generic Active Directory settings, but also requires additional parameters that are greyed out if you are doing simple password authentication through AD. As a result, you will need to have the generic AD parameters configured, just like with simple password, but the authentication type will be changed to Kerberos. When Kerberos is selected, the AD server configuration wizard unlocks the Kerberos parameters.

On the AD server, users and roles must be configured. Roles can be defined on the AD server in either of the following ways:

- Legacy, the Oracle roles (administrators, operators, network, storage, server, and no-access) and can be prefaced with “xg-” for example, xg-administrators). Using the “xg-” prefix was required, but no longer is. Using “xg-” is still supported so that you do not need to delete and recreate the user accounts and roles on the AD server. However, if you do have existing groups that use the “xg-” prefix, you will need to create a group mapping to map them to the role(s) on the Oracle Fabric Manager.
- Group Mapping, the users can be mapped to Oracle Fabric Manager RBAC roles through a group mapping. The group mapping allows groups and roles to be created on the AD servers without the “xg-” prefix. Then, on Oracle Fabric Manager, the group can be mapped into Oracle Fabric Manager's RBAC roles. Through the group mapping, the groups can be mapped into different roles, and users in a mapped group can then have multiple roles. Group mappings can be set up at either the Oracle Fabric Manager and Fabric Interconnect level or at the individual domain level. For more information, see [“Mapping Users in External Groups Into Oracle Fabric Manager” on page 508](#).

For additional information about how to configure users on an Active Directory server, see documentation that accompanied your AD server.

In Oracle Fabric Manager, additional properties must be configured that reference, or point to, the AD server in your network. These properties are documented in the following section.



Through the AD Servers tab, you can add a new AD Server for a Fabric Interconnect, or you can delete a configured AD server by selecting it in the AD Servers tab, then clicking the garbage can icon. You can also control the state of a selected AD server by bringing it up by clicking the up arrow, or bringing the AD server down by clicking the down arrow.

▼ Configure AD Properties for Oracle Fabric Manager

In addition to configuring the Active Directory server, you must configure a set of parameters in Oracle Fabric Manager that specify the AD server in use in your network. These AD properties are configurable for up to two AD servers, a primary and a secondary. You can configure the AD properties for Oracle Fabric Manager through the AD Servers tab.

To configure AD properties for Oracle Fabric Manager, follow this procedure:

1. **Display the Fabric Interconnects Summary by selecting Managed Devices → Fabric Interconnects.**
2. **Select the Fabric Interconnect for which you want to configure the AD Server. This step displays the Fabric Interconnect in the Fabric Interconnects Summary.**
3. **In the Fabric Interconnects Summary, click the AD Servers tab.**

- On the AD Servers tab, click the plus sign to display the Create a new AD Server dialog.

Create a new AD Server

Name of AD Server: *	tabby
Description:	
Name of Host Server: *	192.168.1.112
Port:	3268
UserDn: *	users@tabby.oracle.com
BaseDn: *	"CN=users,DC=tabby,DC=oracle,DC=com"
Password: *	••••••
Server Mode:	primary
Authentication Type:	simple
Formal User Dn: *	
Kerberos Default Realm: *	
Kerberos Default Domain: *	
Kerberos Host Name: *	
Kerberos Host Port:	

Submit Cancel

- In the Name of AD Server field, enter the name of the AD server in your network. The AD server name can be a nickname, alias, or other name that is not a fully qualified domain name (non-FQDN).
- As an option, in the Description field, you can provide an alphanumeric string that describes the AD server that you are configuring.
- In the Name of Host Server field, enter the name of the host through which one or more users will be authenticating. The host server name must be a fully qualified domain name (FQDN).

8. **As an option, in the Port field, enter the number of a particular port you want the Oracle Fabric Manager Server and Active Directory to use for communication.**

By default, port 3268 is used.

If you specify a non-default port, the port you specify must be dedicated specifically to the Oracle Fabric Manager Server and AD server. The port you specify cannot be used for any other traffic or service.

9. **In the User DN field, enter the user domain name for the AD server. The user DN is a string that consists of:**

- **User name**
- **The “at” sign (@)**
- **The host server name as an FQDN**

For example, users@fatman.oracle.com (as shown) is a validly formed user DN.

10. **In the Base DN field, enter the base domain name that the server will be using.**

The base DN is a string of all the individual components of the domain name plus the DC= prefix to indicate each domain component of the domain name. Each user DN string has the following syntax requirements:

- **The entire user DN string must be enclosed in quotation marks**
- **The user DN string cannot contain any blank spaces**
- **Each domain component in the Base DN must separated by a comma**

For example, “DC=pubstest,DC=oracle,DC=com” (as shown) is a properly formed Base DN.

11. **In the Password field enter the password for the server that you are configuring.**

This password is used to log in from the host server to the AD server.

Note - This chapter focuses on Oracle Fabric Interconnects. The other Fabric Device that Oracle Fabric Manager can manage is called an Oracle SDN Controller. The procedures displayed in this chapter is also applicable to the Oracle SDN Controllers. For information about Oracle SDN Controller, see the Oracle SDN Controller User Guide, 1.0.0.

12. **As an option, from the Server Mode dropdown menu, select whether the server instance you are configuring is the primary or secondary AD server.**

The Oracle IMS implementation supports setting a primary and secondary server. Select the appropriate value for your network:

- **Primary**—The primary server is where the user login is attempted first. If the primary AD server is available, it will always be used for user authentication and authorization. Primary is the default value. Only one primary server is allowed in each AD configuration.
- **Secondary**—For redundancy, you can also specify a secondary AD server, which will be used if the primary AD server cannot respond. The secondary server will perform authentication and authorization as long as the primary server is offline.

Note - For redundancy, Oracle suggests configuring a primary AD server, and a secondary AD server that is a separate physical server.

13. **As an option, from the Authentication Type dropdown menu, select the type of authentication that will be used for this user:**

- **Simple, for simple password authentication, which is the default.**
- **Kerberos, if Kerberos authentication will be used as the IMS.**

At this point, simple password authentication parameters are complete. Proceed to [Step 19](#).

Additional parameters exist on this step of the wizard for Kerberos authentication. If you are configuring Kerberos authentication on the AD server, continue to [Step 14](#).

14. **In the Formal User DN field, enter the user domain name that the server will be processing. The Formal User DN is a string consisting of:**

- **User name**
- **The “at” sign (@)**
- **The host server name as an FQDN**

For example, pubs@pubstest.oracle.com (as shown) is a valid Formal User DN.

15. **In the Kerberos Default Realm field, enter the default realm name that the server will be using. The Kerberos Default Realm is a string of all the individual components of the default realm plus the DC= prefix to indicate each individual component of the default realm string. Each default realm string has the following syntax requirements:**

- **The entire default realm string must be enclosed in quotation marks**
- **The entire default realm string cannot contain any blank spaces**

- **Each component of the default realm string must be separated by a comma**
For example, “DC=pubstest,DC=oracle,DC=com” (as shown) indicates a valid Kerberos Default Realm.
- 16. **In the Kerberos Default Domain field, enter the base domain name that the server will be using.**
This string consists of the minimum domain components of the realm for which the AD server will be authenticating through Kerberos. For example, oracle.com is a valid default domain.
- 17. **In the Kerberos Host Name field, enter the name of the host server through which one or more users will be authenticating.**
The host server name must be a fully qualified domain name (FQDN).
- 18. **As an option, in the Kerberos Host Port field, enter the number of a particular port you want the Oracle Fabric Manager Server and Active Directory to use for communication. By default, port 88 is used.**
If you specify a non-default port, the port you specify must be open and available to the Oracle Fabric Manager Server and AD Server. The port you specify cannot be used for any other traffic or service.
- 19. **When the AD server properties are configured, click Submit to configure the AD server.**
- 20. **Check the AD Servers tab to verify that the AD Server is configured for the Fabric Interconnect.**

RADIUS Servers Information

RADIUS authentication is supported in Oracle Fabric Manager. When you configure RADIUS, you are specifying parameters that do the following:

- Point to the specific RADIUS server that Oracle Fabric Manager can use
- Allow the Oracle Fabric Manager Server to login to the RADIUS database so that users in the RADIUS database are authenticated and authorized through the RADIUS server.

The Oracle implementation of RADIUS supports authentication and authorization and is based on RFC 2138. No Oracle proprietary attributes currently exist. Oracle's RADIUS support does not support accounting. For information about installing and configuring the RADIUS server, see the documentation that accompanied your RADIUS server.

Be aware that you will need to configure specific RADIUS users in the Oracle Fabric Manager GUI. These users are different than the local users on the Fabric Interconnect, and options specific to each user must be set. Additional information about configuring RADIUS users exists in the following section.

RADIUS users are required as part of configuring the RADIUS server for Oracle Fabric Manager. As a result, you must configure the RADIUS users before configuring the RADIUS server.

Note - When configuring RADIUS, be aware that you must edit the RADIUS Clients database to include the Fabric Interconnect, and Oracle Fabric Manager Server that will be using RADIUS for authentication and authorization.

RADIUS server information is contained in the RADIUS Servers tab of the Fabric Interconnects Details frame.

Name	Host Name	User Name	State	Authentication Type	Description
No Record found					

You can configure a RADIUS server for Oracle Fabric Manager by clicking the plus sign, or delete a configured RADIUS server by selecting it on the RADIUS Servers tab, then clicking the garbage can icon.

▼ Configure RADIUS Servers for Oracle Fabric Manager

When configured, RADIUS servers handle authorization and authentication of users logging in to the Oracle Fabric Manager or Fabric Interconnects. To provide RADIUS functionality, you must configure some parameters in Oracle Fabric Manager that reference the particular RADIUS server that will be used. You can configure one or more RADIUS servers for your network, through the RADIUS Servers tab in the Fabric Interconnects Details frame.

When configuring a RADIUS Server for use by Oracle Fabric Manager, be aware that you must also configure RADIUS user entries which have some prerequisites. For more information, see [“RADIUS Users Information” on page 192](#). RADIUS users must be created before configuring the RADIUS server for Oracle Fabric Manager, so make sure to create all RADIUS user accounts first. See [“Configure RADIUS Users” on page 192](#).

To configure RADIUS Servers for Oracle Fabric Manager, follow this procedure:

1. **If RADIUS users have not been configured yet, create them now as documented in [“Configure RADIUS Users” on page 192](#).**

When the correct RADIUS user accounts are created, continue with this procedure.

2. **Display the Fabric Interconnects Details frame by selecting Managed Devices → Fabric Interconnects, then selecting the Fabric Interconnect for which you want to configure a RADIUS Server.**
3. **In the Fabric Interconnects Details frame, click the RADIUS Servers tab.**
See [“RADIUS Users Information” on page 192.](#)
4. **Click the plus sign to display the Create a New RADIUS Server dialog.**

The screenshot shows a dialog box titled "Create a new Radius Server". It contains the following fields and values:

- Name of Radius Server: * doolittle
- Name of Host Server: * raider
- Port: (empty)
- Name of User: * mccarth
- Password: * (masked with dots)
- Secret: * (masked with dots)
- Authentication Type: Select Authentication type (dropdown menu)
- Retries: (empty)
- Timeout: (empty)
- Description: (empty)

Buttons: Submit, Cancel

5. **In the Name of Radius Server field, enter the name of the RADIUS server in your network.**
The RADIUS server name can be a nickname, alias, or other name that is not a fully qualified domain name (non-FQDN).
6. **In the Name of Host Server field, enter the name of the host through which one or more users will be authenticating.**
The host server name must be a fully qualified domain name (FQDN).

7. **As an option, in the Port field, enter the number of a particular port you want the Oracle Fabric Manager Server and RADIUS Server to use for communication. By default, port 3268 is used.**

If you specify a non-default port, the port you specify must be dedicated specifically to the Oracle Fabric Manager Server and RADIUS server. The port you specify cannot be used for any other traffic or service.

8. **In the Name of User field, enter the name of the user(s) that will be authenticated by RADIUS when they log in to Oracle Fabric Manager.**

This user name is the RADIUS user name that you configured in the preceding section.

Note - The named users must exist in the RADIUS users database to be allowed to authenticate. If a RADIUS user has not been created for each person who will log in to Oracle Fabric Manager, you must cancel the Create a New RADIUS Server wizard and add their RADIUS user accounts now.

9. **In the Password field, enter that password that the user will be required to enter in order to log in.**

This password will be checked as part of the RADIUS authentication.

10. **In the Secret field, enter the RADIUS secret password, which is used between the Oracle Fabric Manager Server and the RADIUS server to allow the Oracle Fabric Manager Server to log in to the RADIUS for the purpose of handing off user names for authentication and authorization.**

11. **From the Authentication type dropdown menu, select the type of authentication that the RADIUS user will be enforcing for the user:**

- **PAP (password authentication protocol), which is a simple password authentication method. PAP is the default authentication protocol.**
- **CHAP (challenge handshake authentication protocol), which is a method of combining the user's password with a computation, and comparing that to information that the user enters when the RADIUS challenges the user. CHAP is the more secure authentication protocol of the two.**


12. **As an option, in the Retries field, you can enter the number of retries that can occur between the Oracle Fabric Manager Server and the RADIUS server. Enter a number between 0 and 100. The default is 3 retries. Zero (0) sets no retry, so any failed connection attempt between the Oracle Fabric Manager Server and the RADIUS server halts the authentication attempt.**

Each unsuccessful log in attempt between the Oracle Fabric Manager Server and the RADIUS server causes a retry (if the retry value is greater than 0), and the number specified in the Retries field is the maximum number of retries before authentication is aborted. The default value is 3 retries.

13. **As an option, you can enter the Timeout value (in seconds) for log in attempts between the Oracle Fabric Manager Server and the RADIUS server. Enter a number between 0 and 120. The default is 3 seconds. Zero (0) sets no timeout value, and causes a failed log in attempt to abort the log in attempt.**
 If a login attempt between the Oracle Fabric Manager Server and the RADIUS server is hung or slow, the attempt will remain alive for the length of the timeout. If the timeout value is met without a successful log in, the log in attempt is aborted.
14. **As an option, in the Description field, you can provide an alphanumeric string that describes the AD server that you are configuring.**
15. **When the RADIUS configuration has been specified, click Submit to create the instance of the RADIUS Server in Oracle Fabric Manager.**
16. **Check the RADIUS Servers tab to verify that the RADIUS server was configured correctly.**

RADIUS Users Information

RADIUS users must exist in the RADIUS users database in order for authentication to occur for each user. The RADIUS user name must match the name of the user who will log in to Oracle Fabric Manager for RADIUS authentication to successfully occur. If there is a mismatch between the user name used at Oracle Fabric Manager login and the user name in the RADIUS database, authentication will not complete and the user will not be allowed to log in to Oracle Fabric Manager.



You can add a new RADIUS user by clicking the plus sign, and you can delete a configured RADIUS user by selecting it in the RADIUS Users tab, then clicking the garbage can icon.

▼ Configure RADIUS Users

RADIUS users must exist in the RADIUS users database in order for authentication to occur for each user. The RADIUS user name must match the name of the user who will log in to

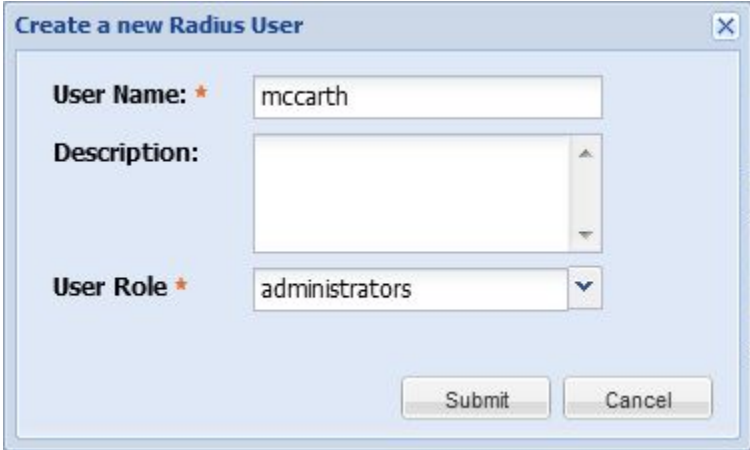
Oracle Fabric Manager for RADIUS authentication to successfully occur. If there is a mismatch between the user name used at Oracle Fabric Manager login and the user name in the RADIUS database, authentication will not complete and the user will not be allowed to log in to Oracle Fabric Manager.

RADIUS users must be created before configuring the RADIUS server for Oracle Fabric Manager, so make sure to create all RADIUS user accounts first.

RADIUS users can be created through Oracle Fabric Manager by using the RADIUS Users list. As an alternative, you can edit the `raddb/users` file and add the users through a command-line session. Be aware that if a user name is not configured in the RADIUS users database, or if there is a mismatch between the RADIUS users user name and the name used to log in to Oracle Fabric Manager, authentication will fail. If RADIUS users are not yet configured, configure them now.

To configure RADIUS Users, follow this procedure:

1. **Display the Fabric Interconnects Details frame by selecting Managed Devices → Fabric Interconnects, then selecting the Fabric Interconnect for which you want to configure a RADIUS user.**
2. **In the Fabric Interconnects Details frame, click the RADIUS Users tab.**
See [“RADIUS Users Information”](#) on page 192.
3. **Click the plus sign to display the Create a New RADIUS User dialog.**



The screenshot shows a dialog box titled "Create a new Radius User". It has a close button in the top right corner. The dialog contains three input fields: "User Name" with a red asterisk and the value "mccarth", "Description" with an empty text area, and "User Role" with a red asterisk and a dropdown menu showing "administrators". At the bottom are "Submit" and "Cancel" buttons.

4. **In the User Name field, enter the name of the user who will log in to Oracle Fabric Manager.**

The user name must exactly match the name you enter. This name will be passed to the RADIUS server, written to its users database, and will be checked against the user name that is entered when someone attempts to log in to Oracle Fabric Manager.

5. **As an option, in the Description field, you can enter an alphanumeric sting that describes the user you are creating.**
6. **From the User Role dropdown menu, select the role that the user will be granted when authentication occurs. Select from the following roles:**
 - administrators
 - network
 - operators
 - server
 - storage
7. **When the RADIUS user properties have been configured, click Submit.**
8. **Check the RADIUS Users tab to verify that the RADIUS user was correctly configured.**

▼ Unmanage a Fabric Interconnect

At any time, you can unmanage a Fabric Interconnect in Oracle Fabric Manager. Unmanaging a Fabric Interconnect is intended to be used only when the Fabric Interconnect is no longer going to be managed by Oracle Fabric Manager. Unmanaging a Fabric Interconnect permanently removes a Fabric Interconnect from Oracle Fabric Manager; use it only for this purpose.

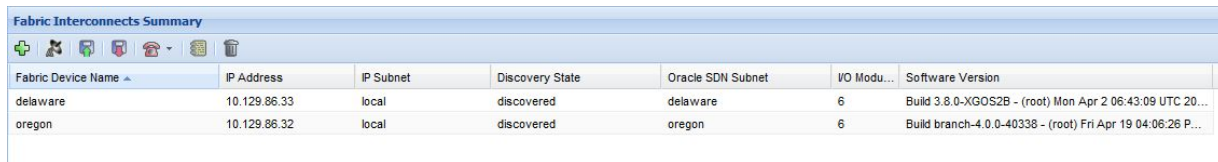
Note - Unmanaging Fabric Interconnects is not to be used for clearing errors or transient states (for example, if a large job is running for a long time, and you want to “force quit” the job), unless you are otherwise directed to do so by Oracle personnel.

When a Fabric Interconnect is unmanaged, then remanaged within a reasonable amount of time, the Oracle Fabric Manager database retains selected configuration information (for example, port associations in clouds, vNIC and vHBA terminations, and so on). As a result, when the Fabric Interconnect is remanaged, the general rule is that this information is put back into Oracle Fabric Manager after the Fabric Interconnect has returned to up/up state when it is remanaged.

However, there are some exceptions to this rule. Domains contain their own secured objects —Fabric Interconnects, users, I/O modules and so on. If a domain has secured objects assigned to it, and the Oracle Fabric Manager that controlling those secured items is unmanaged, when the Oracle Fabric Manager is managed again, the secured objects are not automatically put back into their original domain(s). For more information, see [“Working With Domains” on page 477](#).

To unmanage a Fabric Interconnect, follow this procedure:

1. **Display the Fabric Interconnect summary by selecting Managed Devices → Fabric Interconnects.**



Fabric Device Name	IP Address	IP Subnet	Discovery State	Oracle SDN Subnet	I/O Modu...	Software Version
delaware	10.129.86.33	local	discovered	delaware	6	Build 3.8.0-XGOS2B - (root) Mon Apr 2 06:43:09 UTC 20...
oregon	10.129.86.32	local	discovered	oregon	6	Build branch-4.0.0-40338 - (root) Fri Apr 19 04:06:26 P...

2. **Select the Fabric Interconnect that you want to unmanage.**
This step activates the Unmanage button (garbage can) on the toolbar.
3. **Click Unmanage to remove the selected Fabric Interconnect from Oracle Fabric Manager.**
4. **Click Yes on the confirmation dialog to complete unmanaging the Fabric Interconnect.**

▼ Back Up and Restore a Fabric Interconnect Configuration

As a best practice, you should be running scheduled backups of the Fabric Interconnect. By doing so, you have the ability to replace parts of the Fabric Interconnect (or in more serious situations, the entire Fabric Interconnect itself) while still keeping the virtual I/O configuration.

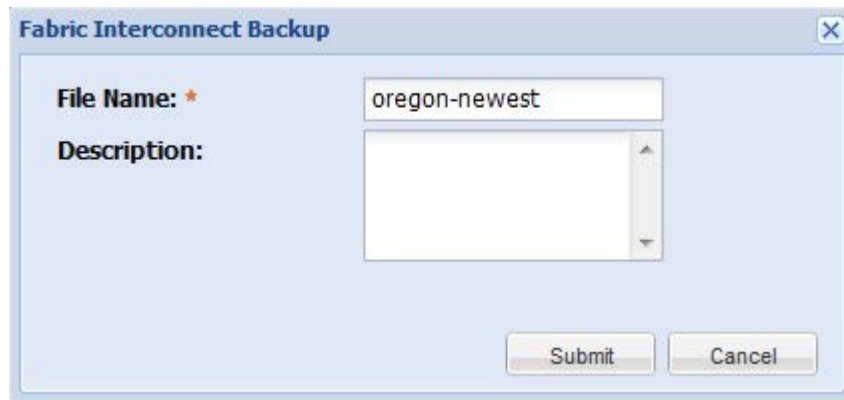
As an alternative to scheduled backups, you can use on-demand backups which you must manually start. While scheduled backups are preferred because they operate automatically, on-demand backups are also useful for capturing the Fabric Interconnect configuration as long as you are consistent with manually taking the on-demand backups.

In situations where you need to replace the Fabric Interconnect and want to save the current configuration through Oracle Fabric Manager, use this procedure, which requires that you have a Fabric Interconnect configuration saved prior to the need to replace the Fabric Interconnect:

1. Display the Fabric Interconnects Summary page by selecting Managed Devices → Fabric Interconnects.
2. On the Fabric Interconnect Summary, select the Fabric Interconnect that you want to backup, and click the Backup Configuration from Selected Fabric Interconnects button.



3. When you click the button, the Fabric Interconnect Backup dialog is displayed.



4. In the File Name field, enter a name for the configuration that you are saving.
5. Click Submit to save the selected Fabric Interconnect's configuration to the file you just named.
6. Make the necessary repairs or replacements to the Fabric Interconnect.
7. When the Fabric Interconnect has recovered from the errored state, log in to Oracle Fabric Manager again.

8. **Display the Fabric Interconnects Summary page by selecting Managed Devices → Fabric Interconnects.**
9. **On the Fabric Interconnect Summary, select the Fabric Interconnect that you want to backup, and click the Restore a Configuration Back to the Selected Fabric Interconnects button.**



Performing Tech Support Functions on the Fabric Interconnect

This section contains the following topics:

- [“Troubleshooting Functions” on page 197](#)
- [“Collect Tech Support Information” on page 198](#)
- [“Send an On-Demand Phone Home Message to Oracle Support” on page 199](#)
- [“Snooze Phone Home for a Window of Time” on page 200](#)

Troubleshooting Functions

For each Fabric Interconnect managed by Oracle Fabric Manager, you can perform a selected number of troubleshooting functions, including:

- Collecting Tech Support Information
- Sending an On-Demand Phone Home Message to Oracle Support
- Snoozing Phone Home for a Window of Time

Because these functions are available for each Fabric Interconnect, if you have two Fabric Interconnects in your environment, you might need to perform the troubleshooting function on both Fabric Interconnects to get the information you need.

Each of these features is supported through the Technical Support Actions toolbar button (the telephone) on the Fabric Interconnects Summary.



▼ Collect Tech Support Information

The Fabric Interconnect can prepare specific information for Oracle Technical Support services to diagnose. Through Oracle Fabric Manager, you can trigger the Fabric Interconnect to prepare and gather the information contained in tech support logs, and transmit that information to Oracle Support. Collecting tech support information occurs from one Fabric Interconnect at a time, so if you have HA Fabric Interconnects on a shared fabric, you might need to collect the tech support information twice—once for each Fabric Interconnect. This function in Oracle Fabric Manager operates the same as the set system tech-support command in Oracle's XgOS CLI.

Typically, you will collect the tech support information when directed to do so by Oracle Support, but you can collect this information at any time it is needed. When you collect tech support information, the information is put into log files and stored at the following locations on the Oracle Fabric Manager server:

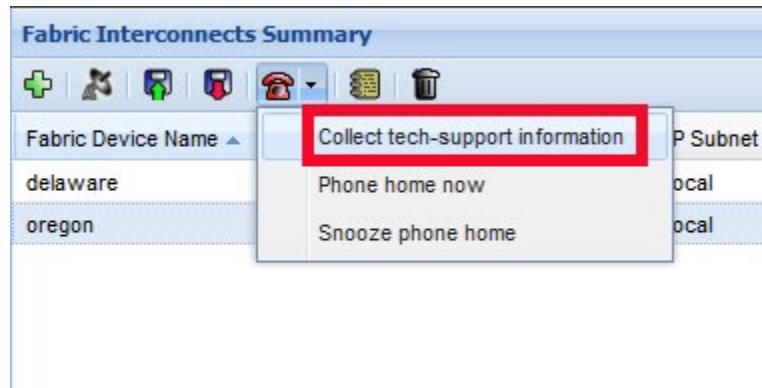
- On Oracle Solaris and Linux Oracle Fabric Manager servers, tech support logs are at `/opt/xsigo/xms/techsupport`
- On Windows Oracle Fabric Manager servers, tech support logs are at `\Program Files\XMS\techsupport`

Note - Collecting the tech support information can take a considerable amount of time based on the amount of data that will be collected. Please be patient and allow this function to complete.

To collect the Technical Support information from a Fabric Interconnect, follow this procedure:

1. **On the navigation panel, select Managed Devices → Fabric Interconnects to display the Fabric Interconnects Summary.**

2. Click the Tech Support Actions toolbar button to display the dropdown menu.
3. Select Collect tech-support information.



4. Read the confirmation message, and click Yes to begin gathering the information.
5. When directed to do so by Oracle Support, send the information to Oracle Support.

▼ Send an On-Demand Phone Home Message to Oracle Support

Phone Home information can be set up to occur on a regular schedule, but you can also send a Phone Home message to Oracle Support whenever you need to by sending an on-demand phone home message. This type of phone home message provides the same information to Oracle Support as a scheduled phone home, and an on-demand phone home message also uses the same parameters. For example, if Strip Private is set for scheduled phone home messages, no private IP addresses will be sent in the on-demand phone home message either. For information about setting up Phone Home parameters, see [“Phone Home Information” on page 169](#).

Typically, a Phone Home Now message is sent when Oracle Support directs you to do so, but you can send a Phone Home message without Oracle explicitly telling you to.

Note - Sending a Phone Home message to Oracle Support creates a secure HTTP message. This message varies in size, but based on the information gathered from the Fabric Interconnect, can be very large. Please be patient and allow the Phone Home message to complete.

To send an on-demand phone home message to Oracle Support, follow this procedure:

1. **On the navigation panel, select Managed Devices → Fabric Interconnects to display the Fabric Interconnects Summary.**
2. **Click the Tech Support Actions toolbar button to display the dropdown menu.**
3. **Select Phone Home Now.**



4. **Read the confirmation message, then click Yes to send the Phone Home message to Oracle Support.**
5. **If needed, follow up with Oracle Support to verify that they received the Phone Home message.**

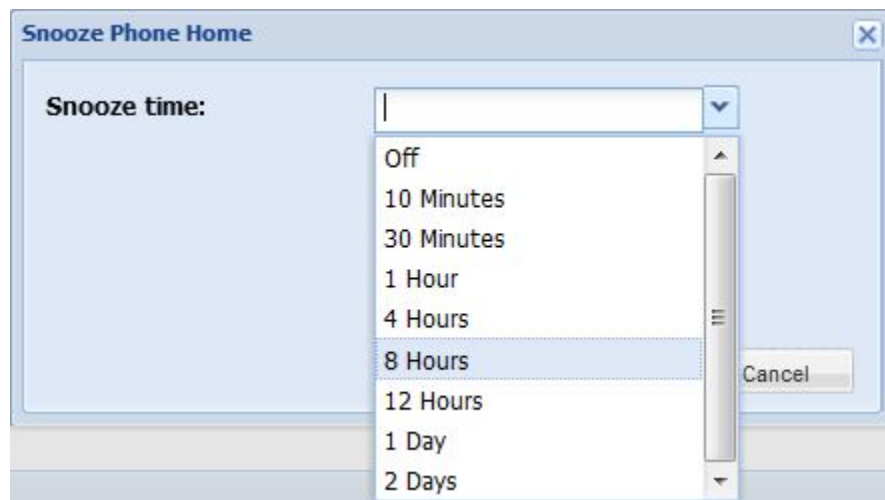
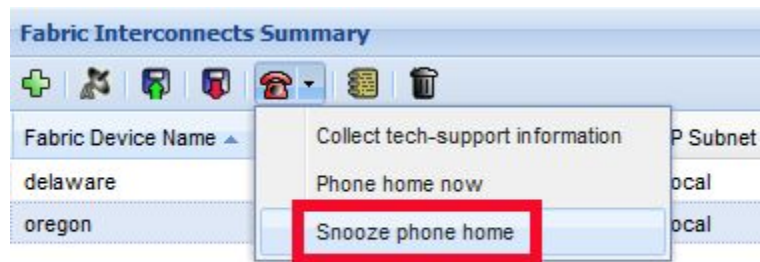
▼ Snooze Phone Home for a Window of Time

For scheduled Phone Home, you can temporarily halt the scheduled transmission of Phone Home information through the use of the snooze feature. The snooze feature lasts for a customizable amount of time from 10 minutes to two days. During the snooze interval, no scheduled phone messages occur. After the snooze interval expires, the scheduled Phone Home messages resume transmission at their regularly scheduled time. During the snooze interval, you can still send on-demand Phone Home messages. (See [“Send an On-Demand Phone Home Message to Oracle Support”](#) on page 199 if needed).

Snoozing Phone Home occurs on one Fabric Interconnect at a time, so you might need to use this feature twice if you have an HA Fabric Interconnect environment—once for each of the Fabric Interconnects. Typically, Phone Home is snoozed by Oracle Support's request, for example, while they are diagnosing Phone Home messages that have already been sent. However, you can snooze phone home at any time for any reason.

To snooze Phone Home, follow this procedure:

1. **On the navigation panel, select Managed Devices – Fabric Interconnects to display the Fabric Interconnects Summary.**
2. **Click the Tech Support Actions toolbar button to display the dropdown menu.**
3. **Select Snooze Phone Home.**



4. **From the Snooze Time dropdown menu, select the amount of time from 10 minutes to 2 days that you want to snooze the scheduled phone home messages.**

5. When the snooze time is selected, click **Submit** to set Phone Home into snooze mode.

▼ Collect Fabric Interconnect Log Files

Oracle Fabric Manager supports collecting core files (if any) and a subset of the log files from a Fabric Interconnect. Typically, collecting log files is done at Oracle Support's request, but it can be done at any time. Depending on the amount of data in a Fabric Interconnect's log files, collecting the log files can take a considerable amount of time. Please be patient and allow the process to run to completion.

Collecting log files from a Fabric Interconnect is supported on one Fabric Interconnect at a time. In an HA Fabric Interconnect environment, you might need to run this operation twice—once for each Fabric Interconnect. This procedure is supported through the Collect Log Files from Selected Fabric Interconnect toolbar button on the Fabric Interconnects Summary.

Note - The other Fabric Device supported through Oracle Fabric Manager is the Oracle SDN Controller. Collecting log files from an Oracle SDN Controller is a similar process to what is documented here. For more information, see the Oracle SDN Controller User's Guide.

To collect all log files from the Fabric Interconnect, follow this procedure:

1. On the navigation panel, select **Managed Devices** → **Fabric Interconnects** to display the **Fabric Interconnects Summary**.
2. Click the **Collect Log Files from Selected Fabric Interconnect** toolbar button.



Fabric Device Name ▲	IP Address	IP S
delaware	10.129.86.33	loc
oregon	10.129.86.32	loc

Note - Depending on the amount of vNICs, vHBAs, and hosts connected to the Fabric Interconnect, collecting log files might take a noticeable amount of time. Please be patient and allow the system to complete collecting all files.

After the logs are downloaded from the Fabric Interconnect to the Oracle Fabric Manager server, they will be in either of the following locations:

- **On Oracle Solaris and Linux Oracle Fabric Manager servers, the logs will be in /opt/xsigo/xms/techsupport**
 - **On Windows Oracle Fabric Manager servers, the logs will be in \Program Files\Xsigo\xms\techsupport**
3. **When directed to do so by Oracle Support, send the files to Oracle Support for diagnostics.**

Working With I/O Templates

This chapter contains the following topics:

- [“I/O Templates” on page 205](#)
- [“Creating an I/O Template” on page 206](#)
- [“I/O Template Summary” on page 245](#)
- [“Rename an I/O Template” on page 246](#)
- [“Editing an I/O Template” on page 247](#)
- [“Apply an I/O Template to Servers” on page 249](#)
- [“I/O Template Re-Apply Overview” on page 251](#)
- [“Re-Apply an I/O Template” on page 252](#)
- [“Allowed VLANs Feature” on page 254](#)
- [“Configure Allowed VLANs in an I/O Template” on page 254](#)
- [“Delete an I/O Template” on page 256](#)

I/O Templates

An I/O Template is a Oracle Fabric Manager feature that allows the server administrator to create the shape of the I/O for a set of servers. Different servers have different I/O requirements. For example, a server that will be used by the engineering department may require access to a couple of Network Clouds and a Storage Cloud, but a server that is used by Finance may only require access to the internet.

By setting up an I/O Template, the server administrator is building a way to deploy a set of servers with the desired I/O requirements. An I/O Template specifies a blueprint (or general configuration) for a server, which you can use for servers of a similar type or that need the same connectivity. For example, for servers that require HA network connections and multipathed storage connections, you could create one I/O Template that provides these connections.

An I/O Template contains the plan for a server's connectivity, but not the actual connectivity itself. The I/O Template provides a connection plan that is then used to create an I/O Profile. The I/O Profile derived from the I/O Template is what pushes the vNICs and vHBAs onto the server, in turn creating the actual connections. So, when an I/O Template is deployed to a server, an I/O Profile is created “behind the scenes” to provide the live connections on

the server. connection is established. The actual connectivity occurs when an I/O Profile is connected to a server.

I/O Templates don't always need to be applied directly to servers for “live” connections. Instead, you can create an instance of an I/O Profile and provision it with the resource required for a server, then connect the I/O Profile (and, in turn, the live connections) at a later time. Consider the following example.

You are building out a data center. You will receive a truckload of servers later today, but you need more time for the SAN configuration. You know what network and storage connections are required for each server. However, configuring your targets, zoning, and LUN masking will take a significant amount of time. In this situation, you have enough network and storage connectivity information, so you can create some I/O Profiles for the servers. These I/O Profiles can remain configured with vNICs and vHBAs, but not yet connected, until you have completed the SAN configuration. For more information about I/O Profiles, see [“Working With I/O Profiles” on page 259](#).

As part of creating an I/O Template, you will be using a Network Cloud, a Storage Cloud, or both. You will find it helpful if your Network and Storage Clouds are already created before building I/O Templates.

Creating an I/O Template

This section contains the following topics:

- [“I/O Template Creation Overview” on page 207](#)
- [“Create a New I/O Template” on page 208](#)
- [“Add a Network Cloud to the I/O Template” on page 211](#)
- [“Specify Advanced Properties for the Network Cloud” on page 213](#)
- [“Add a Storage Cloud to the Template” on page 215](#)
- [“Specify Advanced Properties for the Storage Cloud” on page 217](#)
- [“Add a PVI Cloud to the Template” on page 219](#)
- [“Add vNICs to the I/O Template” on page 221](#)
- [“Configure Advanced vNIC Properties” on page 223](#)
- [“Add PVI vNICs to the I/O Template” on page 225](#)
- [“Configure Advanced vNIC Properties for a PVI vNIC” on page 228](#)
- [“Add vHBAs to the I/O Template” on page 229](#)
- [“Configure Advanced vHBA Properties” on page 232](#)
- [“Save the I/O Template” on page 233](#)
- [“Creating an HA vNIC Template” on page 234](#)
- [“Creating an HA vHBA Overview” on page 239](#)

I/O Template Creation Overview

An I/O Template must be created by you with information that is pertinent to your network. Oracle does not provide pre-configured I/O Templates.

Creating an I/O Template can occur by either adding a new I/O Template and building it from scratch, or when at least one I/O Template exists, copying that I/O Template and editing it as needed.

This section documents creating a new I/O Template from scratch. For information about editing a configured I/O Template, see [“Editing an I/O Template” on page 247](#).

When you create a new I/O Template, you are essentially creating an empty container that you add building blocks to until the I/O Template has the connectivity and virtual resources required for your network. For example, creating a new I/O Template might consist of the following phases:

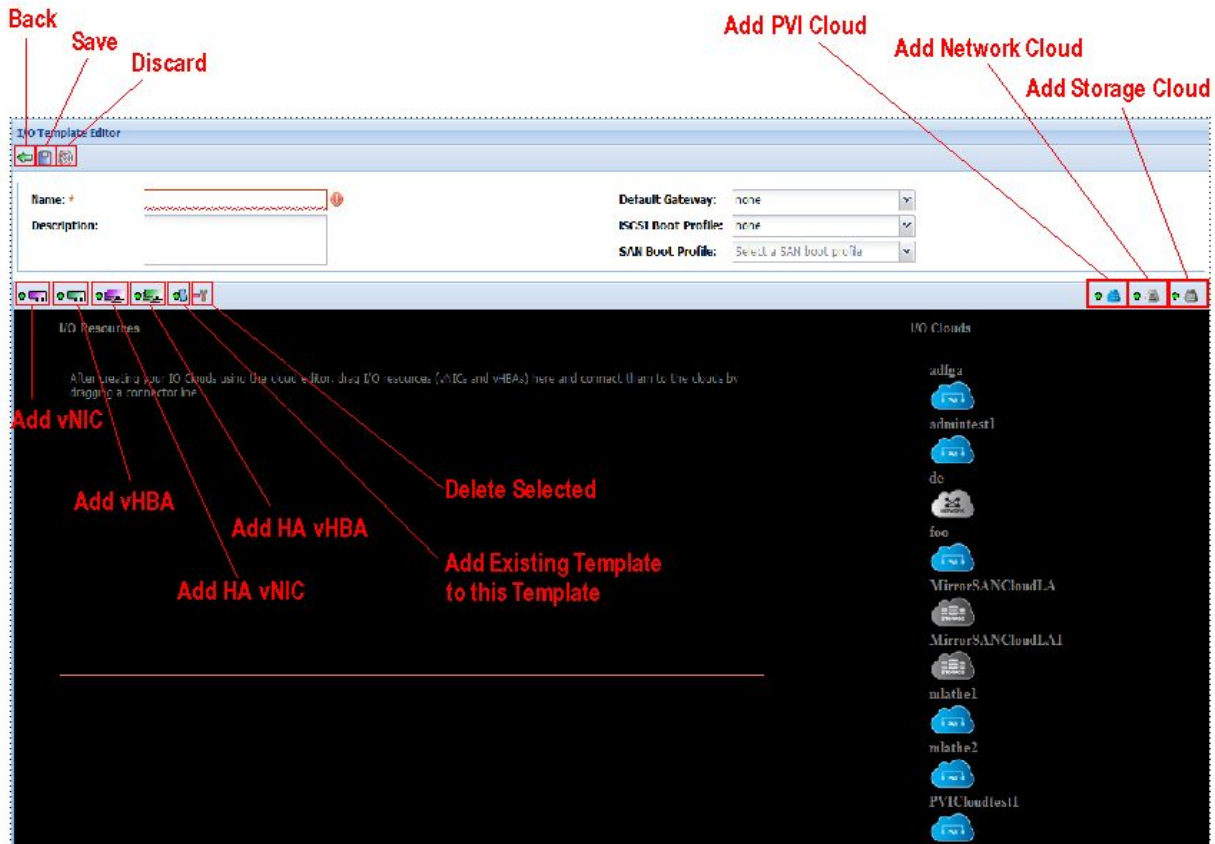
- Naming the I/O Template
- Adding a Network Cloud and/or a Storage Cloud (building blocks). Typical deployments for end-to-end server connectivity for both network and storage traffic requires each type of cloud. However, different deployments might require only one type of cloud. For example, if your deployment has numerous switches and virtual switches for network traffic, you would need a minimum of one Network Cloud, but no Storage Clouds since no storage traffic exists.
- Editing the network and Storage Clouds to support specific features (for example, Network and SAN QoS Profiles)
- Adding one or more vNICs and vHBAs (more building blocks)
- Editing the vNICs and vHBAs with specific Ethernet and Fibre Channel properties
- Connecting the vNIC to the Network Cloud, and connecting the vHBA to the Storage Cloud
- Saving the named I/O Template, and deploying it to one or more host servers

Creating an I/O Template is supported through the I/O Template Editor, which has two main frames:

- The top frame is for general properties of the I/O Template—for example, name and description.
- The bottom frame is the work area where the I/O Template's building blocks are assembled and the bulk of the work occurs. The bottom frame also has two specific parts:
 - I/O Resources, which is where vNICs and vHBAs are added as building blocks to the I/O Template.
 - I/O Clouds, which is where Network Clouds, PVI Clouds, and Storage Clouds are added as building blocks to the overall I/O Template.

Each area (the top frame, I/O Resources in the bottom frame, and I/O Clouds in the bottom frame) has a group of buttons that control those parts of the I/O Template editor.

The following figure shows the I/O Template Editor and indicates the controls available in the I/O Template.



You can add a new I/O Template by clicking the plus sign, and you can delete the currently displayed I/O Template by clicking the Discard icon. If the template has not been saved once, the entire template is discarded. If the template has been saved at least once, then any changes since the last save are discarded.

▼ Create a New I/O Template

Creating an I/O Template is supported through the I/O Template Editor, which is a work space that allows adding building blocks to the empty I/O Template. Because the I/O Template Editor is a work space, you can add, delete, and make changes as needed without impacting the Oracle Fabric Manager or Oracle Fabric Device. Oracle Fabric Manager does not save the in-progress

I/O Template configuration at intervals during the configuration, and Oracle Fabric Manager also does not automatically save the finished configuration. To save the configuration when it is complete, you must explicitly click the Save icon.

The I/O Template Editor is accessible from the navigation panel by selecting Server Resource Manager->I/O Templates.

Note - You will find it helpful if the following elements are already created before building I/O Templates:

- * server boot profiles (SAN Boot Profiles or iSCSI Boot Profiles)
- * default gateways

To create a new I/O Template, follow this procedure:

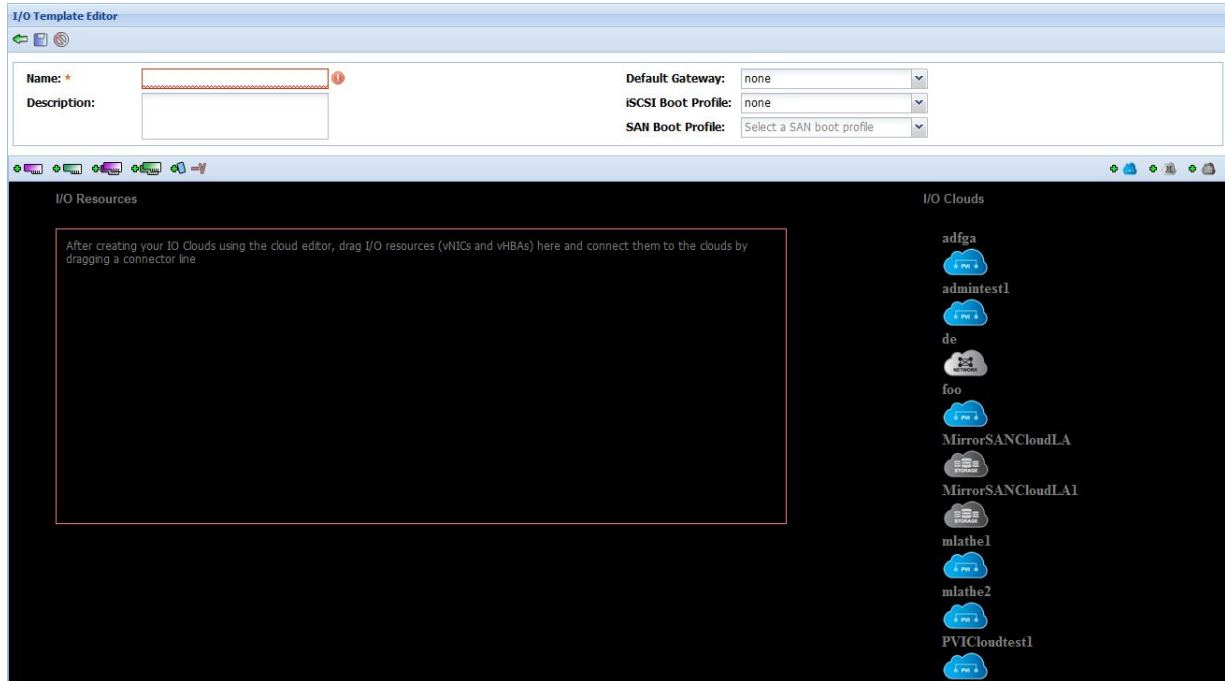
1. **Display the I/O Template Summary by selecting Server Resource Manager → I/O Templates.**

I/O Template Summary							
Name	iSCSI Boot Profile	SAN Boot Profile	Status	vNICs	vHBAs	Default Gateway	Description
HRTemplate			✓	0	1		
LAMirrorSite1			✓	0	1		
LAMirrorSite2			✓	0	1		
Xpubstest			✓	1	0		

4 items

You can add a new I/O Template by clicking the plus sign, and you can delete a configured I/O Template by selecting it in the I/O Template Summary, then clicking the garbage can icon.

2. Click the plus sign to display the I/O Template Editor.



3. In the Name field, enter an alphanumeric string that names the I/O Template that you are creating.
4. As an option, in the Description field, enter an alphanumeric string that describes the I/O Template that you are creating.
5. From the Default Gateway dropdown menu, select the default gateway for the I/O Template.
The default gateway must already be created in Oracle Fabric Manager for it to be a selectable item on the dropdown menu.
6. If the hosts using this I/O Template will be booting up through iSCSI boot, from the iSCSI Boot Profile dropdown menu, select the iSCSI Boot Profile for the I/O Template.
 - The iSCSI Boot Profile must already be created in Oracle Fabric Manager for it to be a selectable item on the dropdown menu.

- At least one bootable vNIC must be configured in the iSCSI Boot Profile for the server to bootable through iSCSI.
7. If the hosts using this I/O Template will be booting up through SAN boot, from the SAN Boot Profile dropdown menu, select the SAN Boot Profile for the I/O Template.
- The SAN Boot Profile must already be created in Oracle Fabric Manager for it to be a selectable item on the dropdown menu.
 - At least one bootable vHBA must be configured in the SAN Boot Profile for the server to bootable through SAN Boot.

▼ Add a Network Cloud to the I/O Template

1. Click the Add Network Cloud button to display the New Network Cloud dialog.

For this procedure, only one Network Cloud is being assigned, but you can assign multiple Network Clouds to the same I/O Template.

New Network cloud

Name: *

Description:

Ethernet Ports / LAGs:

Name	Type	State	vNICs	Capac...	Trunk ...	vLAN ID	vLAN Range	Description
arkansas/14/1	nwEthernet1...	up/up	7	1G	access	1		Private Network
arkansas/14/10	nwEthernet1...	up/down	3	1G	trunk	1		tagged-mtu1500
arkansas/14/2	nwEthernet1...	up/up	5	1G	access	1		Private Network
arkansas/14/3	nwEthernet1...	up/down	6	1G	access	1		Private Network
arkansas/14/4	nwEthernet1...	up/down	0	1G	access	1		
arkansas/14/5	nwEthernet1...	up/down	0	1G	access	1		
arkansas/14/6	nwEthernet1...	up/up	4	1G	access	1		LAB Network
arkansas/14/7	nwEthernet1...	up/up	4	1G	access	1		LAB Network

52 items

HA Designation: Check To set primary and secondary directors

Advanced Configuration

Submit Cancel

2. In the Name field, enter an alphanumeric string that names the Network Cloud that you are creating.
3. As an option, in the Description field, enter an alphanumeric string that describes the Network Cloud that you are creating.
4. From the Ethernet Ports table, select the port(s) that will be used in this Network Cloud.

Ports are selected when they are highlighted. Multiple ports can be assigned to the same Network Cloud, and multiple Network Clouds can be assigned to the same port(s).

Note - If you are creating a Network Cloud that will be terminating an HA vNIC, you must have at least two separate Ethernet ports in the Network Cloud. You can use standard key combinations (for example, CTRL + click) to select multiple ports. For more information about creating an I/O Template for an HA vNIC, see [“Create an HA vNIC Template” on page 235](#).

5. **As an option, if two Fabric Devices will be deployed as a high availability pair, you can set the HA designation for each Fabric Device by using the HA Designation check box.**

The HA designation determines the first and second Fabric Device to which any HA objects in Oracle Fabric Manager will connect. The primary Fabric Device is where the primary HA object connects first, and the secondary Fabric Device is where the secondary HA object connects second. For additional information, see [“Set the Fabric Device HA Designation for a Network Cloud” on page 104](#).

At this point, you can create a basic Network Cloud, or you can specify advanced configuration parameters for the Network Cloud.

- **To create the basic cloud, proceed to [Step 6](#).**
 - **To specify advanced parameters for the Network Cloud, proceed to [“Specify Advanced Properties for the Network Cloud” on page 213](#).**
6. **To complete the configuration of a basic Network Cloud, click Submit.**

▼ **Specify Advanced Properties for the Network Cloud**

Advanced properties are available for the Network Cloud. Advanced properties give you additional control and customization of the vNICs that are available to servers. With advanced properties, you can enable powerful features like VLANs, or traffic shaping features like Network QOS. While these properties are not required for the basic configuration of a vNIC to carry traffic, they are useful for more complex deployments.

To specify advanced properties, follow this procedure:

1. Click the Advanced Configuration button to display the advanced properties for the Network Cloud.

The screenshot shows the 'New Network cloud' configuration window. The 'Advanced Configuration' section is highlighted with a red box. It contains the following fields:

- QoS:** A dropdown menu with the text 'Select a QoS Profile...'.
- VLAN ID:** A text input field.
- Trunk Mode:** A checkbox.
- Private:** A checkbox.
- Allocation Policy:** A dropdown menu with 'Round Robin' selected.

Below the 'Advanced Configuration' section are 'Submit' and 'Cancel' buttons.

2. From the QoS dropdown menu, if the vNICs in the Network Cloud will require QoS, select the Network QoS profile that will be applied to the Network Cloud. The Network Cloud must already exist for it to be a selectable object. If you do not see the Network QoS Profile you want to assign, verify that it exists.
3. If the vNICs that connect to the Network Cloud will all be in the same VLAN, in the VLAN ID field, enter the VLAN number.
4. In the Trunk Mode checkbox, if the Network Cloud will be supporting trunk-mode VLANs, click the checkbox.
5. As an option, if you want the vNICs that connect to the Network Cloud to be private vNICs, click the Private checkbox. Private vNICs are used for Oracle's

vNIC-to-vNIC switching to provide add isolation for a set of vNICs, and also provide enhanced compatibility for existing and new methods of external switching.

6. From the Allocation Policy dropdown menu, select how the Network Cloud will be applied to the host server when the cloud gets deployed. Select:
 - Round Robin, which is a systematic way for ports to be assigned. With round robin allocation, you specify multiple ports in a list of available ports for a Network Cloud. When the list is constructed, you then assign an Ethernet port as the next available port. When you assign the next available port, it gets a rank that is different than any other port in the available ports list. The port you assign is the next port to be assigned, and after that, any additional vNICs that are connected to the Network Cloud will receive an Ethernet port based on its rank. For example, assume you create a Network Cloud with Ethernet port 2/1 as the port used to create the cloud. Then, assume you add ports 2/2, 4/1, and 5/1 to the available ports list, and set port 4/1 as the next available port. When new vNICs are connected to the Network Cloud, the first vNIC is connected to port 4/1. After that, additional vNICs are assigned to whichever port has the next lowest rank.
7. When the correct properties have been assigned for the Network Cloud, click Submit to create the Network Cloud.
8. Check the I/O Template Editor to verify that the Network Cloud was added correctly.

▼ Add a Storage Cloud to the Template

1. Click the Add Storage Cloud button to display the New Storage Cloud dialog.

For this procedure, only one Storage Cloud is being assigned, but you can assign multiple Storage Clouds to the same I/O Template.

New Storage Cloud

Name: *

Description:

FC Ports: *

Port Name	Type	vHBAs	State	Description
arkansas/10/1	sanFcPort	0	up/up	
arkansas/10/2	sanFcPort	0	up/down	
arkansas/12/1	sanFcPort	0	up/up	
arkansas/12/2	sanFcPort	0	up/up	
arkansas/7/1	sanFcPort	0	up/up	
arkansas/7/2	sanFcPort	0	up/down	
arkansas/9/1	sanFc8GbPort	25	up/up	NetAPP
arkansas/9/2	sanFc8GbPort	0	up/up	Compellant
iowa/8/1	sanFcPort	45	up/up	

24 items

HA Designation: Check To set primary and secondary directors

Advanced Configuration

Submit Cancel

2. In the Name field, enter a name for the Storage Cloud that you are creating.
3. As an option, in the Description field, enter a description for the Storage Cloud that you are creating.
4. From the FC Ports table, select the port(s) that will be used in this Storage Cloud.

Ports are selected when they are highlighted. Multiple ports can be assigned to the same Storage Cloud, and multiple Storage Clouds can be assigned to the same port(s).

Note - If you are creating a Storage Cloud that will be terminating an HA vHBA, you must have at least two separate Fibre Channel ports in the Storage Cloud. You can use standard key combinations (for example, CTRL + click) to select multiple ports. For more information about creating an I/O Template for an HA vHBA, see [“Create an HA vHBA Template” on page 240](#).

5. **As an option, if two Fabric Devices will be deployed as a high availability pair, you can set the HA designation for each Fabric Device by using the HA Designation check box.**

The HA designation determines the first and second Fabric Device to which any HA objects in Oracle Fabric Manager will connect. The primary Fabric Device is where the primary HA object connects first, and the secondary Fabric Device is where the secondary HA object connects second. For additional information, see [“Set the Fabric Device HA Designation for a Network Cloud” on page 104](#).

At this point, you can create a basic Storage Cloud, or you can specify advanced configuration parameters for the Network Cloud.

- **To create the basic cloud, proceed to [Step 6](#).**
 - **To specify advanced parameters for the Storage Cloud, proceed to [“Specify Advanced Properties for the Storage Cloud” on page 217](#).**
6. **To complete the configuration of a basic Storage Cloud, click Submit.**

▼ **Specify Advanced Properties for the Storage Cloud**

Advanced properties are available for the Storage Cloud. Advanced properties give you additional control and customization of the vHBAs that are available to servers. With advanced properties, you can enable traffic shaping features like SAN QOS. While these properties are not required for the basic configuration of a vNIC to carry traffic, they are useful for more complex deployments.

To specify advanced properties for a Storage Cloud, follow this procedure:

1. Click the Advanced Configuration button to display the advanced properties for the Storage Cloud.

The screenshot shows the 'New Storage Cloud' configuration window. The 'Advanced Configuration' section is highlighted with a red box. It contains the following fields:

- Name:** * (empty text box with a red dashed border and an information icon)
- Description:** (empty text area)
- FC Ports:** * (table with 5 columns: Port Name, Type, vHBAs, State, Description)
- HA Designation:** Check To set primary and secondary directors
- Advanced Configuration:** (highlighted section)
 - QoS:** Select a QoS Profile... (dropdown menu)
 - Allocation Policy:** Round Robin (dropdown menu)
- Buttons:** Submit, Cancel

Port Name	Type	vHBAs	State	Description
arkansas/10/1	sanFcPort	0	up/up	
arkansas/10/2	sanFcPort	0	up/down	
arkansas/12/1	sanFcPort	0	up/up	
arkansas/12/2	sanFcPort	0	up/up	
arkansas/7/1	sanFcPort	0	up/up	
arkansas/7/2	sanFcPort	0	up/down	
arkansas/9/1	sanFc8GbPort	25	up/up	NetAPP
arkansas/9/2	sanFc8GbPort	0	up/up	Compellant
iowa/8/1	sanFcPort	45	up/up	

2. From the QoS dropdown menu, if the vHBAs in the Storage Cloud will require QoS, select the SAN QoS profile that will be applied to the Storage Cloud (if any).
3. From the Allocation Policy dropdown menu, select how the Storage Cloud will be applied to the host server when the cloud gets deployed. Select:
 - Round Robin, which is a systematic way for ports to be assigned. With round robin allocation, you specify multiple ports in a list of available ports for a Storage Cloud. When the list is constructed, you then assign a port as the next available port. When you assign the next available port, it gets a rank that is different than any other port in the available ports list. The port you assign is the next port to be assigned, and after that, any additional vHBAs that are connected to the Storage Cloud will receive a fibre channel port based on its rank. For example, assume you create a Storage Cloud

with fibre channel port 2/1 as the port used to create the cloud. Then, assume you add ports 2/2, 4/1, and 5/1 to the available ports list, and set port 4/1 as the next available port. When new vHBAs are connected to the Storage Cloud, the first vHBA is connected to port 4/1. After that, additional vHBAs are assigned to whichever port has the next lowest rank.

4. **When the correct properties have been assigned for the Storage Cloud, click Submit to create the Storage Cloud.**

When you click Submit, a dialog is briefly displayed that informs you that Oracle Fabric Manager client and server are exchanging data.

5. **Check the I/O Template Editor to verify that the Storage Cloud was added correctly.**

▼ Add a PVI Cloud to the Template

Private Virtual Interconnect (PVI) Clouds are used with the Fabric Device to provide purpose-specific connections over the IB fabric that are used to move non-data, non-IO traffic around the network quickly. For example, you can use PVI Clouds (and the PVI vNIC within them) for “east-west” traffic such as evacuating one virtual machine from one server to another. For more information about PVI Clouds, see the Fabric Accelerator User Guide, 1.0.0 or the Oracle SDN Controller User Guide.

Assuming you have created the basic I/O Template as documented in [“Create a New I/O Template” on page 208](#), you can continue the workflow for creating an I/O Template by adding PVI Clouds to the template:

1. **Click the Add PVI Cloud button to display the New PVI Cloud dialog.**

For this procedure, only one PVI Cloud is being assigned, but you can assign multiple PVI Clouds to the same I/O Template.

Oracle SDN Name	Oracle SDN Subnet	Fabric Device List	Number Of F
fabric_5514059420008689	delaware	delaware	7
fabric_5514059420009009	oregon	oregon	2

2. In the Name field, enter an alphanumeric string that names the PVI Cloud that you are creating.
3. As an option, in the Description field, enter an alphanumeric string that describes the PVI Cloud that you are creating.
4. From the MTU dropdown menu, select the largest size frame that can be transmitted on the PVI Cloud without fragmentation.
5. From the Oracle SDN table, select the SDN Fabric on which you want to create the PVI Cloud.
The fabric name will be either a Fabric Interconnect name or an Oracle SDN Controller name.
6. Click Submit to add the PVI Cloud to the I/O Template.
7. Proceed to [“Add PVI vNICs to the I/O Template” on page 225.](#)

▼ Add vNICs to the I/O Template

Continue the main workflow for creating an I/O Template by adding one or more vNICs to the I/O Template:

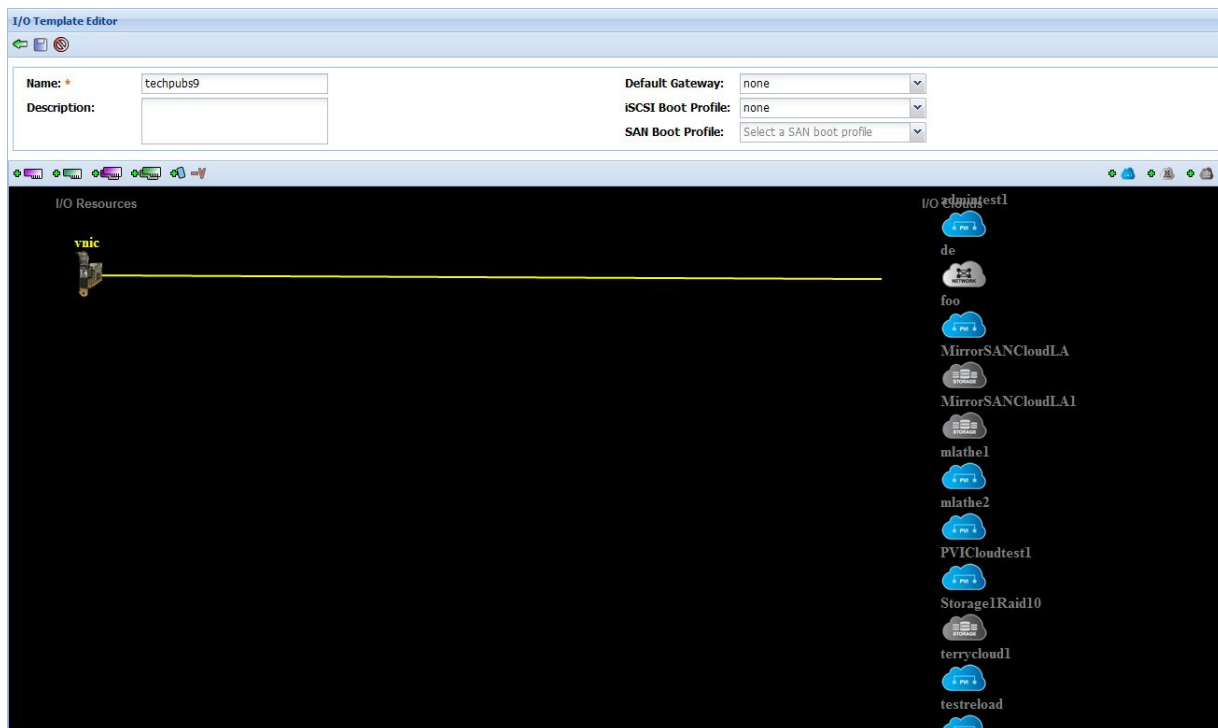
1. **Click the Add vNIC button to add an instance of a vNIC to the I/O Template Editor.**

This procedure only adds one, but you can add as many vNICs to the I/O Template as needed.

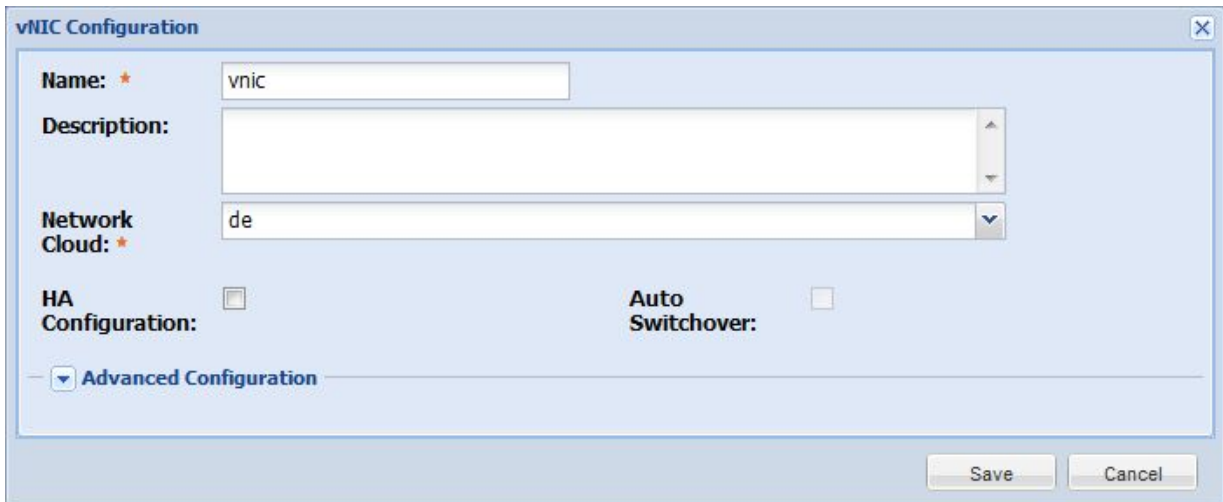
2. **When the vNIC is added to the I/O Template Editor's work space, click and drag a connection from the vNIC to the correct Network Cloud.**

This step creates a vNIC with default parameters and terminates it on a Network Cloud. You will want to edit the vNIC parameters.

Note - This step does not provide the virtual I/O connection to the host server. It simply creates a vNIC and terminates it on a port in the Network Cloud. The vNIC will not be pushed to the host server until the I/O Template is saved, and after that an I/O Profile is created from the I/O Template, and then that I/O Profile is connected to a server.



3. When the vNIC is connected to the Network Cloud, double click the vNIC icon in the work space to display the Edit vNIC Resource dialog and set or change vNIC properties as needed.



The screenshot shows the 'vNIC Configuration' dialog box. It has a title bar with the text 'vNIC Configuration' and a close button. The main area contains the following fields and controls:

- Name: ***: A text input field containing the value 'vnic'.
- Description:**: A large, empty text area for entering a description.
- Network Cloud: ***: A dropdown menu currently showing 'de'.
- HA Configuration:**: A checkbox that is currently unchecked.
- Auto Switchover:**: A checkbox that is currently unchecked.
- Advanced Configuration:**: A section header with a small downward-pointing arrow to its left.
- Buttons:** 'Save' and 'Cancel' buttons are located at the bottom right of the dialog.

4. In the Name field, set or change the vNIC name.
5. As an option, in the Description field, enter an alphanumeric string that describes the vNIC.
6. As needed, from the Network Cloud dropdown menu, select the Network Cloud to which the vNIC will belong.

Note - Because you have already connected the vNIC to the Network Cloud by clicking and dragging a connection from the vNIC icon to the cloud icon, you should not need to do this step. However, you can use the Network Cloud dropdown menu to change the vNIC's termination to a different cloud if needed.

7. **If the vNIC will be one of an HA vNIC pair, click the HA Configuration checkbox.**
This checkbox toggles, so clicking multiple times alternatively sets and unsets the HA configuration flag for the vNIC.

Note - In this example, the HA Configuration checkbox is empty because the vNIC you are creating is not an HA vNIC. For information about configuring an I/O Template for an HA vNIC, see [“Create an HA vNIC Template” on page 235](#).

8. If the vNIC will be one of an HA vNIC pair, click the Auto Switchover checkbox if you want the secondary vNIC to give traffic back to the primary vNICs when the primary comes back online.

At this point you have completed configuring the basic vNIC properties. However, you can configure additional properties if needed by using the procedure in the following section. If you do not want to configure advanced properties, proceed to [“Add vHBAs to the I/O Template” on page 229](#).

▼ Configure Advanced vNIC Properties

If you want to configure advanced features for the vNIC, display the Advanced Configuration options.

The screenshot shows the 'vNIC Configuration' dialog box. The 'Advanced Configuration' section is highlighted with a red border. It contains the following fields and options:

- Name:** vnic
- Description:** (empty text box)
- Network Cloud:** de
- HA Configuration:**
- Auto Switchover:**
- Advanced Configuration (highlighted):**
 - QoS Configuration:** Select a QoS Profile (dropdown menu)
 - IP Type:**
 - DHCP
 - Host Managed
 - Trunk Mode:**
 - PXE Boot:**
 - User defined MAC Address:**
 - TSO:**
 - VLAN ID:** 1
 - Checksum Offload:**
 - Community Name:** (empty text box)
 - Private:**
 - iSCSI Boot:**
 - MAC address range:** (empty text box)

Buttons: Save, Cancel

1. As needed, from the QoS Configuration dropdown menu, select the Network QoS Profile required for the vNIC.

2. **If the vNIC will be required to participate in a VLAN, enter the VLAN number in the VLAN ID field.**
3. **In the IP Type field, select whether the vNIC's IP address will be assigned by the host server, or by DHCP.**
4. **If the vNIC will be supporting checksum offload, which allows the module to send checksumming tasks to the I/O module instead of the port terminating the vNIC, click the Checksum Offload checkbox.**
5. **If the vNIC needs to be part of a specific community, in the Community Name field, enter the name of the community.**
6. **If the vNIC will be configured in a VLAN, click the Trunk Mode checkbox to determine whether the vNIC will be trunk or access.**

If the Trunk Mode checkbox contains a checkmark, the vNIC will operate in trunk mode. If the Trunk Mode checkbox is empty, the vNIC will operate in access mode.
7. **If the vNIC needs to be accessible only by private vNICs (non-public vNICs), click the Private checkbox.**
8. **If the vNIC will be booting a server from a PXE boot server that contains boot information for the server where the vNIC is hosted, click the PXE Boot checkbox.**
9. **If the vNIC will be booting a server from an iSCSI target that contains boot information for the server where the vNIC is hosted, click the iSCSI Boot checkbox.**
10. **User Defined MAC Address, for releasing the port's MAC address from the Fabric Device's MAC address pool and allowing you to specify a particular MAC address for the port supporting the vNIC.**

This option combines with the MAC Address Range to determine the appropriate MAC address.
11. **MAC Address Range, to specify the MAC address range for the port supporting the vNIC.**
12. **If the vNIC will be supporting TCP Segmentation Offload, click the TSO checkbox.**
13. **When the vNIC properties are specified, click Save.**

▼ Add PVI vNICs to the I/O Template

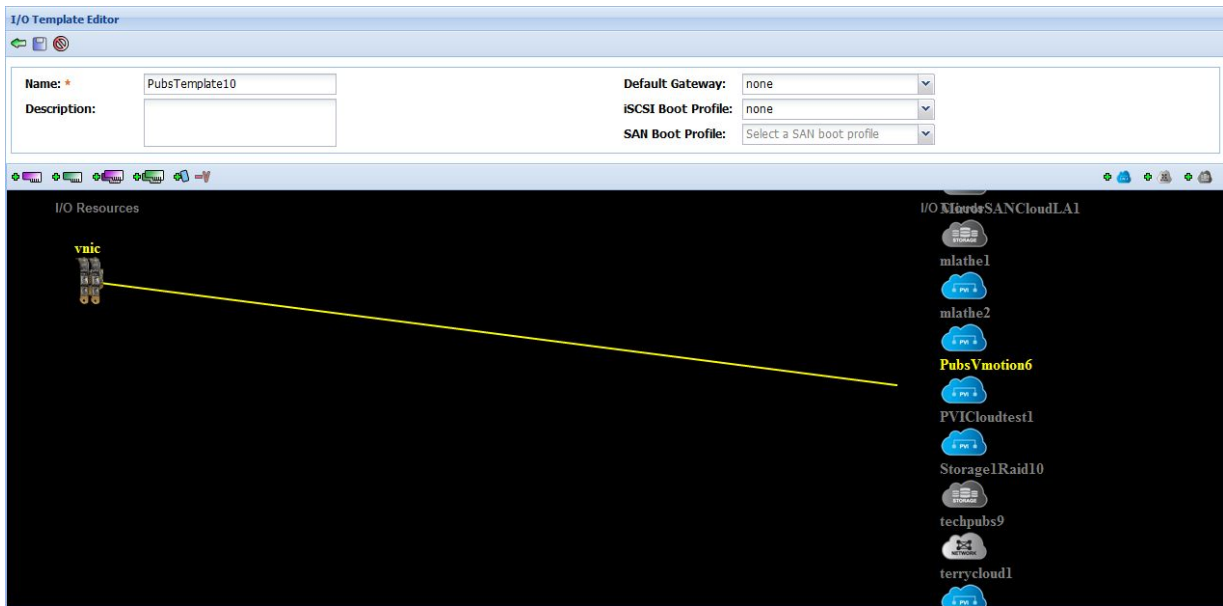
PVI vNICs are standard network vNICs that are terminated on PVI Clouds. PVI vNICs are named for the type of cloud that they connect to (a PVI Cloud). Although PVI vNICs are standard network vNICs, because they are terminated on a PVI Cloud, the vNICs are used for the purposes of PVI traffic. Namely, PVI vNICs are used for specific non-data traffic in an “east-west” direction through the network, such as evacuating a virtual machine from one server to another. PVI vNICs are used by either Oracle Fabric Interconnects or Oracle SDN Controllers. For information about Oracle SDN Controllers, see the Oracle SDN Controller User Guide, 1.0.0.

Despite the difference in usage between PVI vNICs and network vNICs, you can add PVI vNICs to an I/O Template in a similar way to adding network vNICs. PVI vNICs can be either single vNICs or HA vNICs, but Oracle recommends that you configure PVI vNICs as HA vNICs. The following procedure documents adding HA PVI vNICs to an I/O Template.

Continue the workflow of creating an I/O Template by adding vNICs to the template and terminating them on the PVI Cloud you created earlier:

1. **Click the Add an HA vNIC button to add an instance of an HA vNIC to the I/O Template.**
2. **When the vNIC is added to the I/O Template Editor's work space, click and drag a connection from the vNIC to the correct Network Cloud.**

This step creates a vNIC with default parameters and terminates it on a Network Cloud. You will want to edit the vNIC parameters.



Note - This step does not provide the virtual I/O connection to the host server. It simply creates a vNIC and terminates it on a port in the PVI Cloud. The PVI vNIC will not be pushed to the host server until the I/O Template is saved, and then the I/O Template is connected to a server.

3. When the vNIC is connected to the PVI Cloud, double click the vNIC icon in the work space to display the Edit vNIC Resource dialog and set or change vNIC properties as needed.

4. In the Name field, set or change the vNIC name.
5. As an option, in the Description field, enter an alphanumeric string that describes the vNIC.
6. As needed, from the Network Cloud dropdown menu, select the Network Cloud to which the vNIC will belong.

Note - Because you have already connected the vNIC to the PVI Cloud by clicking and dragging a connection from the vNIC icon to the PVI Cloud icon, you should not need to do this step. However, you can use the PVI Cloud dropdown menu to change the vNIC's termination to a different PVI Cloud if needed.

7. Because the PVI vNIC is an HA PVI vNIC, click the Auto Switchover checkbox if you want the secondary vNIC to give traffic back to the primary vNIC when the primary comes back online.

At this point you have completed configuring the basic vNIC properties. However, you can configure additional properties if needed by using the procedure in the following section. For templates that contain PVI Clouds and PVI vNICs, you will not configure vHBAs. If you do not need to configure advanced properties for the PVI vNIC, proceed to [“Save the I/O Template” on page 233](#).

▼ Configure Advanced vNIC Properties for a PVI vNIC

If you want to configure advanced features for the HA PVI vNIC, display the Advanced Configuration options.

The screenshot shows the 'vNIC Configuration' dialog box. The 'Name' field is 'vnic', 'Description' is empty, and 'Network Cloud' is 'PubsVmotion6'. 'HA Configuration' is checked, and 'Auto Switchover' is unchecked. The 'Advanced Configuration' section is expanded and highlighted with a red border. It contains the following fields: 'QoS Configuration' (dropdown menu), 'IP Type' (radio buttons for DHCP and Host Managed, with Host Managed selected), 'Trunk Mode' (checkbox), 'PXE Boot' (checkbox), 'User defined MAC Address' (checkbox), 'TSO' (checkbox), 'VLAN ID' (text field with '1'), 'Checksum Offload' (checkbox), 'Community Name' (text field), 'Private' (checkbox), 'iSCSI Boot' (checkbox), and 'MAC address range' (text field). 'Save' and 'Cancel' buttons are at the bottom right.

Note - Only some of the advanced configuration properties are available for PVI vNICs. If a property is greyed out (disabled), then it is not configurable for a PVI vNIC.

1. If the PVI vNIC will be required to participate in a VLAN, enter the VLAN number in the VLAN ID field.

2. In the IP Type field, select whether the PVI vNIC's IP address will be assigned by the host server, or by DHCP.
3. If the PVI vNIC will be supporting checksum offload, which allows the module to send checksumming tasks to the I/O module instead of the port terminating the PVI vNIC, click the Checksum Offload checkbox.
4. MAC Address Range, to specify the MAC address range for the port supporting the PVI vNIC.
5. When the PVI vNIC properties are specified, click Save.

▼ Add vHBAs to the I/O Template

Continue the main workflow for creating an I/O Template by adding one or more vHBAs to the I/O Template:

1. **Click the Add vHBA button to add an instance of a vHBA to the I/O Template Editor.**

This procedure only adds one, but you can add as many vHBAs to the I/O Template as needed. The example shows adding a single vHBA to the pending I/O Template.

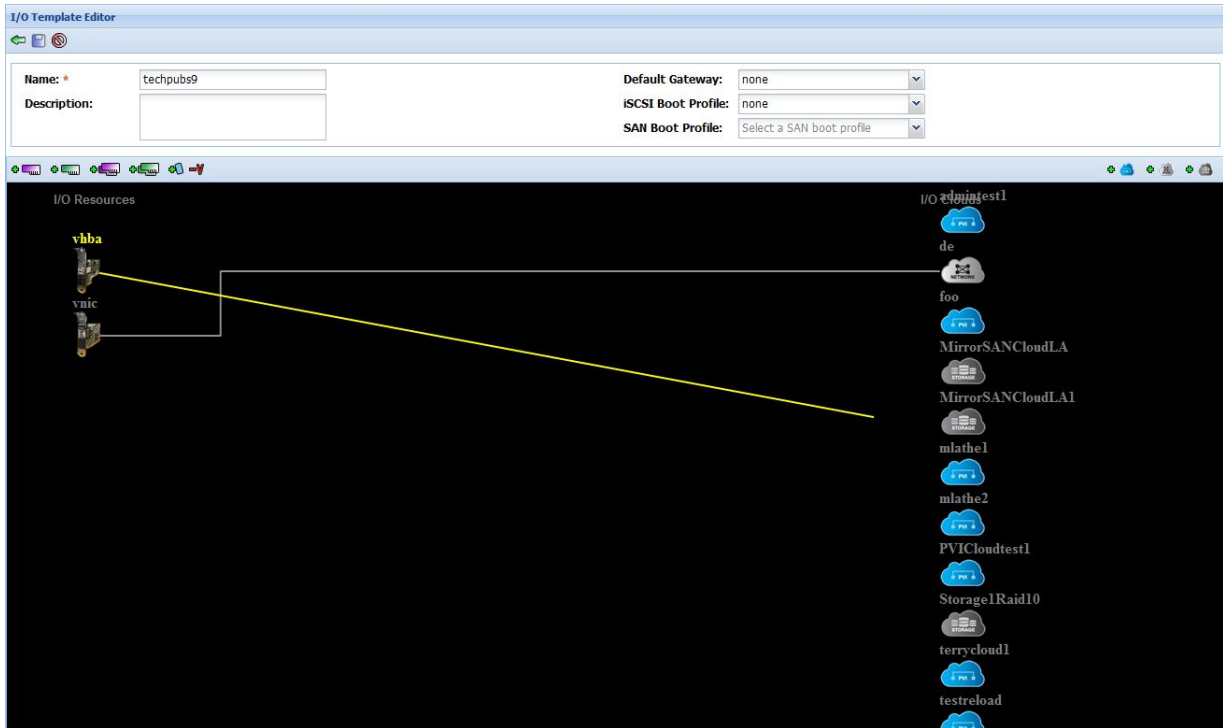
The screenshot shows a dialog box titled "vHBA Configuration". It has a close button in the top right corner. The fields are as follows:

- Name: ***: A text input field containing "vhba".
- Description:**: A large empty text area.
- Storage Cloud: ***: A dropdown menu with "MirrorSANCloudLA1" selected.
- HA Configuration:**: A checkbox that is currently unchecked.
- Advanced Configuration:**: A section with a dropdown arrow pointing down.
- Buttons:** "Save" and "Cancel" buttons are located at the bottom right of the dialog.

2. **When the vHBA is added to the I/O Template Editor's work space, click and drag a connection from the vHBA to the correct Storage Cloud.**

Add vHBAs to the I/O Template

Note - This step creates a vHBA with default parameters and terminates it on a Storage Cloud. You will want to edit the vHBA parameters.



3. When the vHBA is connected to the Storage Cloud, double click the vHBA icon in the work space to display the Edit vHBA Properties dialog and set or change vHBA parameters as needed.

The screenshot shows the 'vHBA Configuration' dialog box. It has a title bar with the text 'vHBA Configuration' and a close button. The dialog contains the following fields and controls:

- Name: ***: A text input field containing the text 'vhba'.
- Description:**: A large, empty text area for entering a description.
- Storage Cloud: ***: A dropdown menu with 'MirrorSANCloudLA1' selected.
- HA Configuration:**: A checkbox that is currently unchecked.
- Advanced Configuration:**: A section with a dropdown arrow pointing down.
- Buttons:** 'Save' and 'Cancel' buttons are located at the bottom right of the dialog.

4. In the Name field, set or change the vHBA name.
5. As an option, in the Description field, enter an alphanumeric string that describes the vHBA.
6. As needed, from the Storage Cloud dropdown menu, select the Storage Cloud to which the vHBA will belong.

Note - Because you have already connected the vHBA to the Storage Cloud by clicking and dragging a connection from the vHBA icon to the cloud icon, you should not need to do this step. However, you can use the Storage Cloud dropdown menu to change the vHBA's termination to a different cloud if needed.

7. **If the vHBA will be part of an HA vHBA, click the HA Configuration checkbox.**
At this point you have completed configuring the basic vHBA properties. However, you can configure additional properties if needed by using the procedure in the following section. If you do not want to configure advanced properties, proceed to [“Add vHBAs to the I/O Template”](#) on page 229.

▼ Configure Advanced vHBA Properties

If you want to configure advanced features for the vHBA, display the Advanced Configuration options.

1. **As needed, from the QoS Configuration dropdown menu, select the SAN QoS Profile required for the vHBA.**
The SAN QoS Profile must already exist to be a selectable item in the dropdown menu.
2. **As needed, from the LUN Mask Profile dropdown menu, select the LUN Mask Profile required for the vHBA.**
The LUN Mask Profile must already exist to be a selectable item in the dropdown menu.
3. **If the vHBA's WWN ID must be assigned from a specific range of WWNs, click the User Defined WWN ID checkbox to have the Fabric Device bypass automatically assigning the vHBA's WWN from the Fabric Device's pool if WWN IDs.**
Use this option if you, or a SAN device, will be assigning WWNs.

4. **If the vHBA's WWN ID must be assigned from a specific range of WWNs, enter the WWN range in the WWN Range field.**

Enter the WWN range as a starting WWN ID, then either a dash (-) or colon (:), then the ending WWN ID.

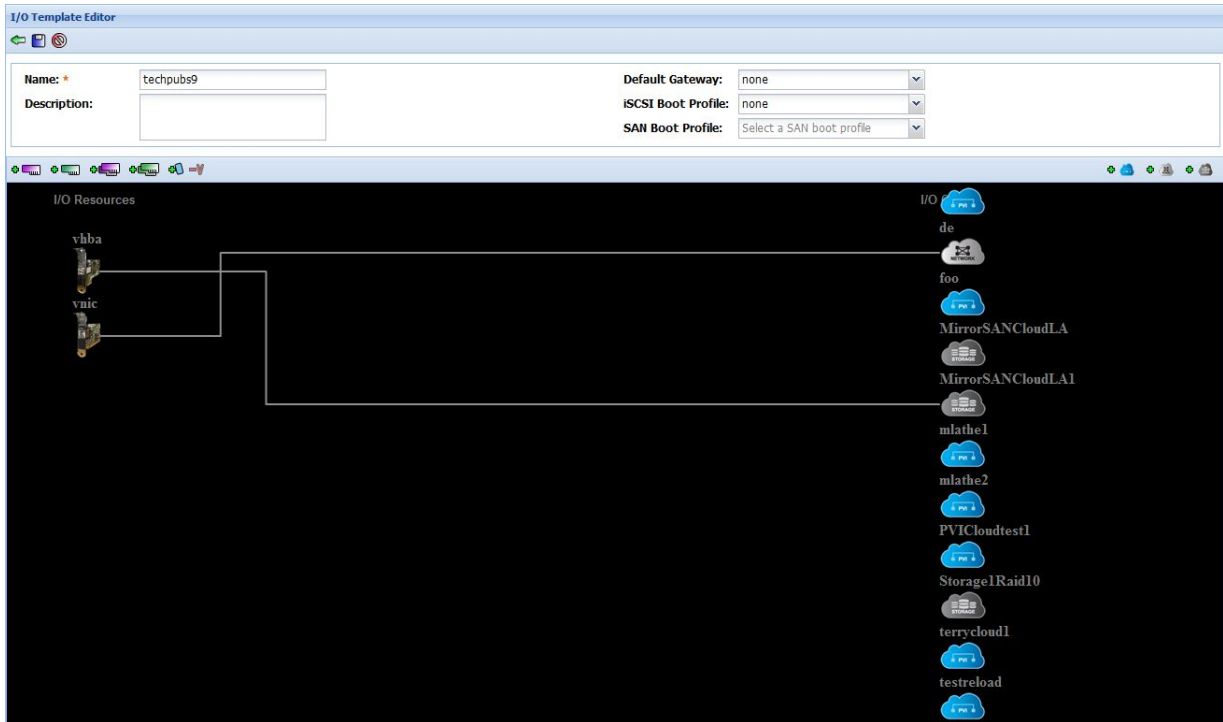
5. **If the vHBA must support SAN Booting the host server on which the vHBA is deployed, click the SAN Boot checkbox.**
6. **When the vHBA properties are specified, click Save.**

▼ Save the I/O Template

Complete the workflow of configuring an I/O Template by following these steps:

1. **Check the work area to make sure that all vNICs, vHBAs, Network Clouds, and Storage Clouds are in the I/O Template, and have been connected appropriately.**
 - **If vNICs and vHBAs need to be deleted, you can click to select them, then click Delete to remove them from the I/O Template before saving the I/O Template.**
 - **If vNICs and vHBAs are not connected, you can either double click the vNIC or vHBA and set the Network and Storage Cloud as needed, or you can click and drag to draw a line between the virtual resource and the I/O Clouds.**

2. When the I/O Template contains the proper virtual connections, click the Save icon to complete creating the I/O Template and display the I/O Template Summary.



Note - At this point, the I/O Template is created, but not yet deployed. You will need to create an I/O Profile from the I/O Template, and connect that I/O Profile to a physical server in order to push the vNICs and vHBAs in the I/O Template to the server.

Creating an HA vNIC Template

This section contains the following topics:

- [“HA vNIC Overview” on page 235](#)
- [“Create an HA vNIC Template” on page 235](#)

HA vNIC Overview

HA vNICs allow for redundancy and fault tolerance for network-connected hosts. With each HA vNIC, a pair of vNICs is created and assigned to two ports in the cloud, or two separate Fabric Devices.

When you create an HA vNIC, the OS where the vNIC is deployed determines how the vNIC operates:

- For Oracle Solaris, Linux, and Windows, the HA vNIC is a pair of vNICs in which one is the primary (which is the online, active vNIC) and a secondary vNIC (which is the online standby). The primary vNIC in the HA pair carries traffic to and from the host server. The secondary vNIC does not carry traffic unless the primary encounters an error that downs the vNIC, and after a very brief failover period, the secondary takes over and the traffic resumes on the secondary vNIC. As an option, you can set the secondary to failback, which allows the secondary to return the traffic to the primary vNIC when the primary comes back online.
- For ESX, Window Server 2012 R2 and OVM, the HA vNIC is simply two separate vNICs. There is no Primary or Secondary vNIC. Instead the ESX OS settings on the host determine how the HA vNIC is used. Typically, on an ESX Server, VMware's native NIC Teaming is used to provide the high availability mechanism, and the Oracle vNICs are just two vNICs that provide the connectivity.

HA vNICs can be terminated on different physical hardware depending on the level of fault tolerance you require:

- To avoid a single point of failure at the module level, terminate the primary and secondary vNICs on different modules in the same Fabric Interconnect.
- To avoid a single point of failure at the chassis level, terminate the primary and secondary vNICs on different Fabric Devices (if you have a multi-chassis configuration).

When you are configuring an HA vNIC pair through Oracle Fabric Manager, two vNICs are added in the I/O Template to represent the primary and secondary vNICs in the HA pair. When the HA vNIC pair is assigned to a Network Cloud, the termination ports are assigned to the primary and secondary vNIC based on which ports were configured in the Network Cloud.

▼ Create an HA vNIC Template

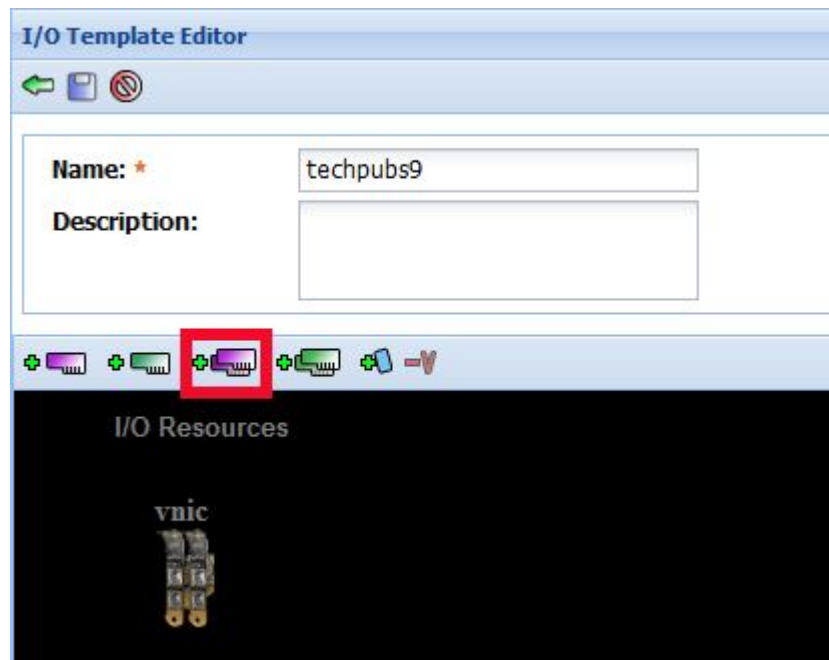
To create an HA vNIC Template, follow this procedure:

1. **Display the I/O Template Summary by selecting Server Resource Manager → I/O Templates.**
2. **Click the plus sign to display the I/O Template Editor.**
3. **In the Name field, enter the name for the I/O Template that you are creating.**

4. As an option, from the Default Gateway dropdown menu, select a default gateway to assign to the vNIC I/O Template.
5. As an option, you can enter a description in the Description field.
6. As an option, if the host for the I/O Template will be iSCSI Booting, from the iSCSI Boot Profile dropdown menu, select an iSCSI Boot Profile.
7. As an option, if the host for the I/O Template will SAN Booting, from the SAN Boot Profile dropdown menu, select a SAN Boot Profile.
8. If a Network Cloud is not already created, create one now by clicking the Add a Network Cloud button. For more information, see [Step 1](#) through [Step 8](#).

Note - For an HA vNIC, the Network Cloud you create must have at least 2 separate ports to terminate the HA vNIC that you will be creating.

9. Click the Add an HA vNIC button to add an instance of an HA vNIC to the I/O Template.

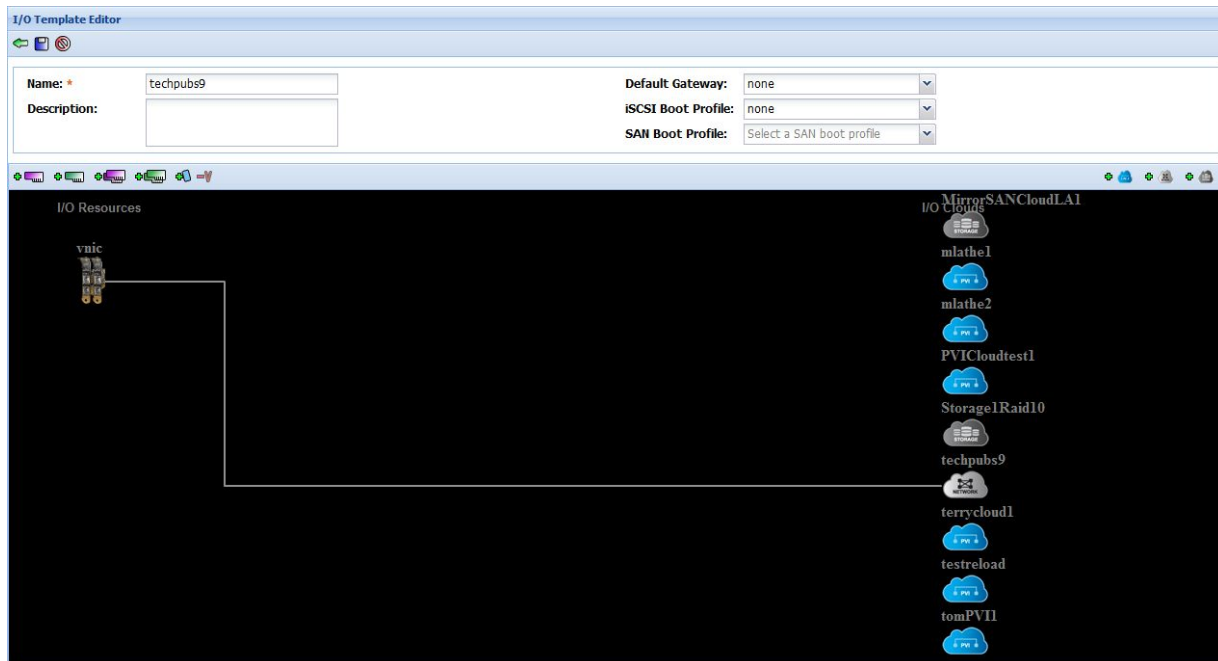


10. When the HA vNIC is added to the I/O Template Editor's work space, click and drag a connection from the HA vNIC to the correct Network Cloud, as shown

in the example. When the line is completely drawn, the HA vNIC has been terminated on Ethernet ports in the Network Cloud.

This step creates an HA vNIC with default parameters and terminates it on a Network Cloud. You will want to edit the HA vNIC parameters.

Note - This step does not provide the virtual I/O connection to the host server. It simply creates two vNICs and terminates them on a two Ethernet ports in the same Network Cloud. The HA vNIC will not be pushed to the host server until the I/O Template is saved, and after that an I/O Profile is created from the I/O Template, and then that I/O Profile is connected to a server.



11. When the HA vNIC is connected to the Network Cloud, double click the HA vNIC to display the Edit vNIC Resource dialog for the HA vNIC that you are adding.

Note - Notice the HA Configuration checkbox is filled, which indicates that the vNIC you are configuring is an HA vNIC.

12. Edit the HA vNIC parameters as needed. For information, see steps [Step 4](#) through [Step 7](#).

13. As an option, you can configure additional options for the HA vNIC by clicking the Advanced Configuration checkbox.

14. As needed, from the QoS Configuration dropdown menu, select the Network QoS Profile required for the vNIC.
15. If the vNIC will be required to participate in a VLAN, enter the VLAN number in the VLAN ID field.
16. In the IP Type field, select whether the vNIC's IP address will be assigned by the host server, or by DHCP.
17. If the vNIC will be supporting checksum offload, which allows checksumming tasks to be sent to the Gigabit Ethernet I/O module where the vNIC is terminated instead of the port terminating the vNIC, click the Checksum Offload checkbox.

18. If the vNIC needs to be part of a specific community, in the Community Name field, enter the name of the community.
19. If the vNIC will be configured in a VLAN, click the Trunk Mode checkbox to determine whether the vNIC will be trunk or access. If the Trunk Mode checkbox contains a checkmark, the vNIC will operate in trunk mode. If the Trunk Mode checkbox is empty, the vNIC will operate in access mode.
20. If the vNIC needs to be accessible only by private vNICs (non-public vNICs), click the Private checkbox.
21. If the vNIC will be booting over a PXE boot vNIC attached to a PXE boot server that contains boot information for the server where the vNIC is hosted, click the PXE Boot checkbox.
22. If the vNIC will be used as an iSCSI vNIC attached to an iSCSI target that contains boot information for the server where the vNIC is hosted, click the iSCSI Boot checkbox.
23. To allow you to specify a particular MAC address for the port supporting the vNIC (instead of allowing the Fabric Device to assign the MAC addresses for vNICs), click the User Defined MAC Address checkbox. This option combines with the MAC Address Range to determine the appropriate MAC address.
24. If you are assigning MAC addresses for ports supporting the vNICs, in the MAC Address Range, specify the MAC address range for the port supporting the vNIC.
25. If the vNIC will be supporting TCP Segmentation Offload, click the TSO checkbox.
26. When the appropriate parameters have been set for the HA vNIC, click Save.
27. On the I/O Template Editor toolbar, click Save to save the HA vNIC I/O Template.

Note - All changes will be lost if you do not save the HA vNIC I/O Template.

Creating an HA vHBA Overview

This section contains the following topics:

- [“HA vHBA Overview” on page 240](#)
- [“Create an HA vHBA Template” on page 240](#)
- [“Configure Advanced vHBA Properties” on page 244](#)

HA vHBA Overview

HA vHBAs provide some redundancy for hosts connected to a Fibre Channel SAN. With HA vHBAs, you create a pair of vHBAs on two separate ports from the same Storage Cloud, or two separate Fabric Devices.

- When configuring HA vHBAs, there is no implied primary or secondary vHBA. HA vHBAs do not have a failover mechanism like common applications of HA vNICs. Instead, the two vHBAs allow for a second logical path from the host to the storage target. How the data path is managed for reads and writes must still be configured through either the FC switch or the host (for example, through MPIO on a Linux host).
- HA vHBAs are not inherently multipathing. HA vHBAs simply create two instances of separate vHBAs on different Fibre Channel ports.

The vHBAs functionality differs depending on the OS of the server where the vHBA is deployed. However, HA vHBAs are typically used to provide two HBAs to the same host to support the connectivity for host-side multipathing.

HA vHBAs can be terminated on different physical hardware depending on the level of fault tolerance you require:

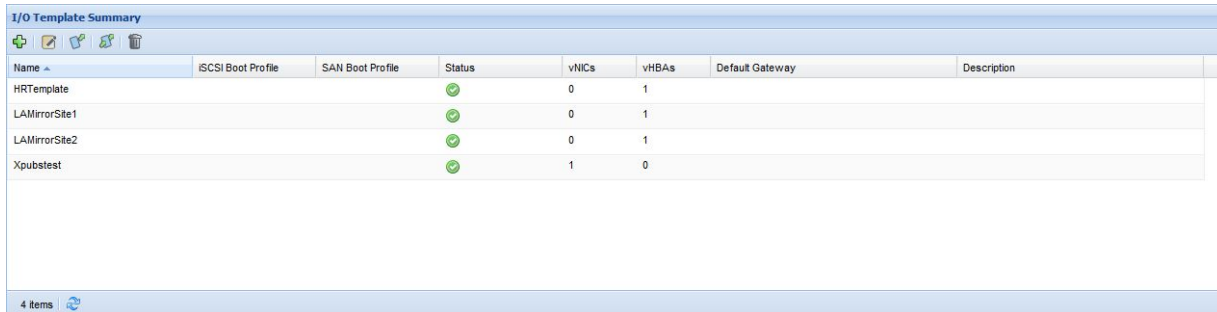
- to avoid a single point of failure at the module level, terminate the primary and secondary vHBAs on different modules in the same Fabric Interconnect.
- to avoid a single point of failure at the chassis level, terminate the primary and secondary vHBAs on different Fabric Devices (if you have a multi-chassis configuration).

When you are configuring an HA vHBA pair through Oracle Fabric Manager, two vHBAs are added in the I/O Template to represent the HA pair. When the HA vHBA pair is associated to a Storage Cloud, the two termination ports are assigned randomly to support the HA vHBA pair.

▼ Create an HA vHBA Template

To create an HA vHBA Template, follow this procedure:

1. Display the I/O Template Summary by selecting Server Resource Manager → I/O Templates.



Name ^	iSCSI Boot Profile	SAN Boot Profile	Status	vNICs	vHBAs	Default Gateway	Description
HRTemplate			✓	0	1		
LAMirrorSite1			✓	0	1		
LAMirrorSite2			✓	0	1		
Xpubstest			✓	1	0		

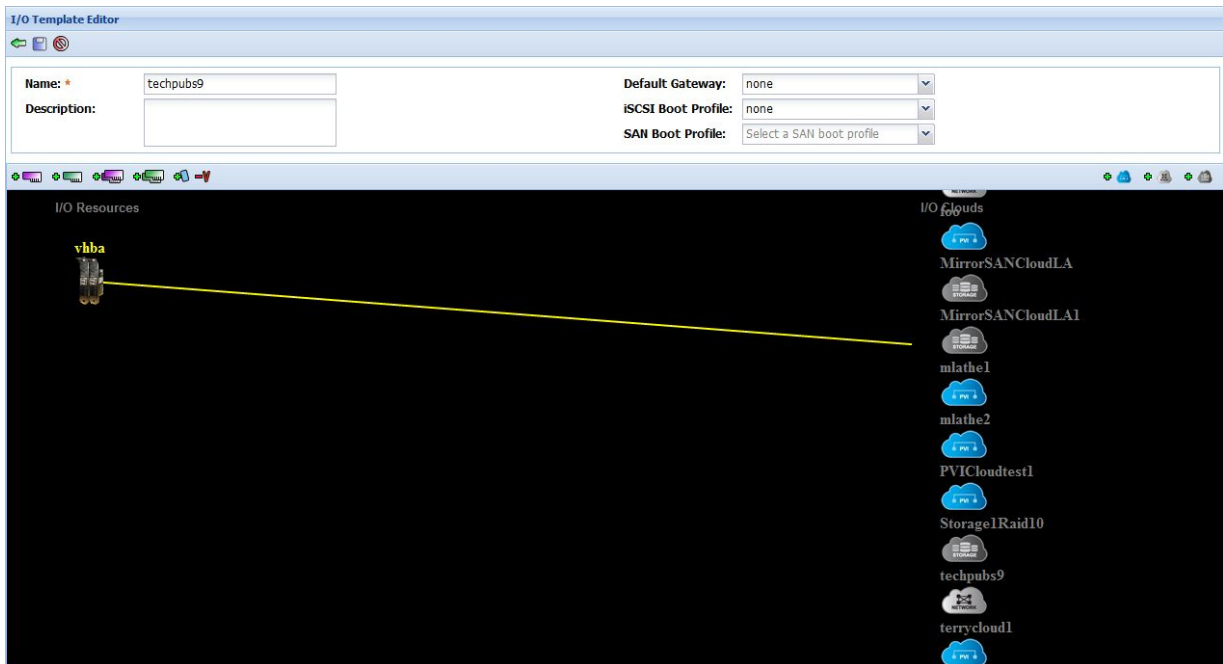
2. Click the plus sign to display the I/O Template Editor.
3. In the Name field, enter the name for the I/O Template that you are creating.
See [“Create a New I/O Template”](#) on page 208.
4. As an option, from the Default Gateway dropdown menu, select a default gateway to assign to the I/O Template.
5. As an option, if the host for the I/O Template will be iSCSI Booting, from the iSCSI Boot Profile dropdown menu, select an iSCSI Boot Profile.
6. As an option, if the host for the I/O Template will SAN Booting, from the SAN Boot Profile dropdown menu, select a SAN Boot Profile.
7. As an option, you can enter a description in the Description field.
8. If a Storage Cloud is not already created, add one now by clicking the Add a Storage Cloud button on the toolbar. For more information, see [Step 1](#) through [Step 5](#).
9. When the Storage Cloud is added to the I/O Template's work area, click the Add an HA vHBA button to add an instance of an HA vHBA to the I/O Template.
10. When the HA vHBA is added to the I/O Template Editor's work space, click and drag a connection from the HA vHBA icon to the correct Storage Cloud.

When the line is completely drawn, HA vHBA has been terminated on Fibre Channel ports in the Storage Cloud.

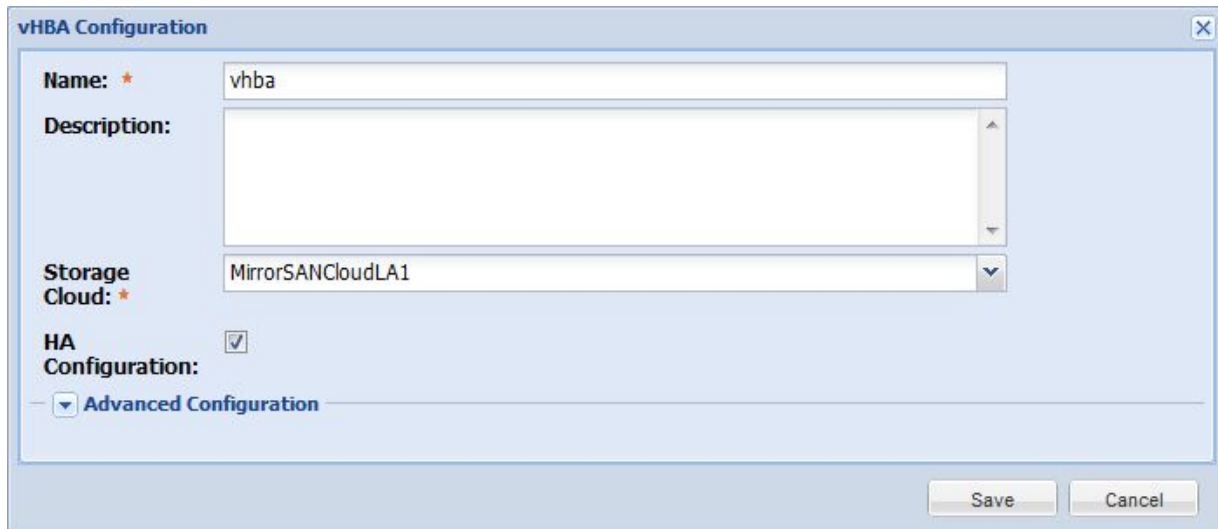
This step creates an HA vHBA with default parameters and terminates it on a Storage Cloud. You will want to edit the HA vHBA parameters.

Create an HA vHBA Template

Note - This step does not provide the virtual I/O connection to the host server. It simply creates two vHBAs and terminates them on two ports in the same Storage Cloud. The HA vHBA will not be pushed to the host server until the I/O Template is saved, and after that, an I/O Profile is created from the I/O Template, and then that I/O Profile is connected to a server.



11. When the HA vHBA is connected to the Storage Cloud, double click the HA vHBA icon in the work space to display the Edit vHBA Properties dialog and set or change HA vHBA parameters as needed.



The screenshot shows the 'vHBA Configuration' dialog box. The 'Name' field is set to 'vhba'. The 'Description' field is empty. The 'Storage Cloud' dropdown is set to 'MirrorSANCloudLA1'. The 'HA Configuration' checkbox is checked. There is an 'Advanced Configuration' section with a dropdown arrow. The 'Save' and 'Cancel' buttons are at the bottom right.

12. In the Name field, set or change the HA vHBA name.
13. As an option, in the Description field, enter an alphanumeric string that describes the HA vHBA.
14. As needed, from the Storage Cloud dropdown menu, select the Storage Cloud to which the HA vHBA will belong.

Note - Because you have already connected the HA vHBA to the Storage Cloud by clicking and dragging a connection from the HA vHBA icon to the cloud icon, you should not need to do this step. However, you can use the Storage Cloud dropdown menu to change the HA vHBA's termination to a different cloud if needed.

15. If the vHBA will be part of an HA vHBA, click the HA Configuration checkbox.

Note - Notice the HA Configuration checkbox is filled, which indicates that the vHBA you are configuring is an HA vHBA.

At this point you have completed configuring the basic vHBA properties. However, you can configure additional properties if needed by using the procedure in the following section.

▼ Configure Advanced vHBA Properties

If you want to configure advanced features for the HA vHBAs, display the Advanced Configuration options.

1. As needed, from the QoS Configuration dropdown menu, select the SAN QoS Profile required for the HA vHBA.
2. As needed, from the LUN Mask Profile dropdown menu, select the LUN Mask Profile required for the HA vHBA.
The LUN Mask Profile must already exist to be a selectable item in the dropdown menu.
3. If the HA vHBA's WWN ID must be assigned from a specific range of WWNs, click the User Defined WWN ID checkbox to have the Fabric Device bypass automatically assigning the HA vHBA's WWN from the Oracle Fabric Device's pool of WWN IDs. Use this option if you, or a SAN device, will be assigning WWNs.
4. If the HA vHBA's WWN ID must be assigned from a specific range of WWNs, enter the WWN range in the WWN Range field.

Enter the WWN range as a starting WWN ID, then either a dash (-) or colon (:), then the ending WWN ID.

5. **If the HA vHBA will need to support SAN Booting the host server on which the vHBA is deployed, click the SAN Boot checkbox.**
6. **When the HA vHBA properties are specified, click Save.**
7. **On the I/O Template Editor toolbar, click Save to save the HA vHBA configuration.**

Note - All changes will be lost if you do not save the HA vHBA I/O Template.

I/O Template Summary

When I/O Templates are configured, they are listed in the I/O Template Summary, which is a listing of all I/O Templates regardless of whether they are deployed to a host server or not. Through the I/O Templates Summary you can create and delete I/O Templates, and edit existing I/O Templates to set or change properties.

Name	iSCSI Boot Profile	SAN Boot Profile	Status	vNICs	vHBAs	Default Gateway	Description
HRTemplate			✓	0	1		
LAMirrorSite1			✓	0	1		
LAMirrorSite2			✓	0	1		
Xpubstest			✓	1	0		

Field	Indicates
Name	The name of each configured I/O Template.
iSCSI Boot Profile	The name of the iSCSI Boot Profile (if any) assigned to the I/O Template.
SAN Boot Profile	The name of the SAN Boot Profile (if any) assigned to the I/O Template.
Status	The operational status of the I/O Template displayed as an icon: <ul style="list-style-type: none"> ■ a green checkmark indicates that the I/O Template is up and available. ■ an empty field indicates that the I/O Template is not finished, or is in an indeterminate state. ■ a red "X" indicates that the I/O Template is not up or is not available for deployment.

Field	Indicates
vNICs	The total number of vNICs configured in each I/O Template.
vHBAs	The total number of vHBAs configured in each I/O Template.
Default Gateway	The IP address of the default gateway configured in the I/O Template.
Description	The description string (if any) that was applied to the I/O Template. If this field is blank, either no description string was specified when the I/O Template was created, or the I/O Template was originally created with a description string, but the I/O Template was later edited and the description string was removed.

▼ Rename an I/O Template

Oracle Fabric Manager supports renaming an I/O Template, which enables you to change the name without having to completely delete and recreate the entire virtual I/O configuration. When the I/O Template is renamed, all other properties for it are retained, including the vNIC and vHBA configuration. As an option, you can also set or change the description for the I/O Template.

You can rename an I/O Template through the I/O Template Details frame. To rename an I/O Template, follow this procedure:

- 1. On the Navigation Frame, select I/O Templates.**
See [“Create a New I/O Template” on page 208](#).
- 2. Select an I/O Template to populate the details frame with its properties.**
- 3. Click the Edit button to edit the properties of the selected I/O Template.**

- 4. In the Name field, enter the new name for the I/O Template.**
- 5. As an option, you also can set or change the description for the selected I/O Template.**
- 6. When the new name has been specified for the I/O Template, click Submit.**

Editing an I/O Template

When an I/O Template exists, you can edit it to set or change features for the I/O Template and its building blocks. For example, if you find a host server is not getting enough bandwidth, you can edit the I/O Template to change the Network QoS properties in the Network Cloud in the I/O Template. You can also edit an I/O Template to add or delete vNICs or vHBAs, or set a specific boot policy for the host (for example, set up the I/O Template for SAN Booting).

Note - Although the procedures documented in this section are valid, an easier way to edit an I/O Template is to use the I/O Template Editor, which provides a more graphical approach to modifying vNICs, vHBAs, and I/O Templates.

Editing an I/O Template occurs through the I/O Template Editor, which supports the following functionality:

- [“Change I/O Template General Properties” on page 247](#)
- [“Add or Delete a Template's vNICs” on page 248](#)
- [“Add or Delete a Template's vHBA” on page 248](#)

▼ Change I/O Template General Properties

Through the I/O Template Details frame, you can display the general properties of an individual I/O Template. The general properties appear on the General tab.

The General tab also contains an Edit button that allows you to set or change parameters as needed. For example, through the General tab you could add or change a server boot profile.

The screenshot shows the 'I/O Template: PubsTemplate7' editor window. The 'General' tab is selected, showing the following fields:

- Name:** PubsTemplate7
- Default Gateway:** none
- Total vHBAs:** 0
- SAN Boot Profile:** Select a SAN boot profile
- Apply Template Name:**
- Total vNICs:** 1
- iSCSI Boot Profile:** none
- Description:** (empty text area)

Buttons for 'Submit' and 'Cancel' are located at the bottom left.

To edit the general properties for an I/O Template, follow this procedures:

1. **Display the I/O Templates Summary (Server Resource Manager → I/O Templates).**

2. **On the I/O Template Summary, select an I/O Template to display it in the details frame.**
3. **Click the General tab, then click the Edit button to unlock the editable fields on the General tab.**
4. **Make the necessary modifications.**
5. **When the properties are configured correctly, click Submit.**

▼ **Add or Delete a Template's vNICs**

Through the I/O Template Details frame, you can display the vNICs tab, which is a list of the vNICs that are assigned to an I/O Template.

The individual vNICs on this tab are also links to vNIC properties, which you can use to edit properties of the vNICs in the I/O Template. The general properties appear on the General tab. Clicking the vNIC name displays the vNIC properties for that specific vNIC.

To edit the vNIC properties for an I/O Template's vNIC, follow this procedure:

1. **Display the I/O Templates Summary (Server Resource Manager → I/O Templates).**
2. **On the I/O Template Summary, select an I/O Template to display it in the details frame.**
3. **Click the vNICs tab.**
4. **On the vNICs tab, click a vNIC name to display the vNIC Details frame, then click the Edit button to unlock the editable fields.**
5. **Make the required modifications.**
6. **When the properties are configured correctly, click Submit.**

▼ **Add or Delete a Template's vHBA**

Through the I/O Template Details frame, you can display the vHBAs tab, which is a list of the vHBAs that are assigned to an I/O Template.

The individual vHBAs on this tab are also links to vHBA properties, which you can use to edit the general properties appear on the General tab. Clicking the vHBA name displays the vHBA properties for that specific vHBA.

To edit the vHBA properties for an I/O Template's vHBAs, follow this procedure:

1. **Display the I/O Templates Summary (Server Resource Manager → I/O Templates).**
2. **On the I/O Template Summary, select an I/O Template to display it in the details frame.**
3. **Click the vHBAs tab.**
4. **On the vHBAs tab, click a vHBA name to display the vHBA Details frame, then click the Edit button to unlock the editable fields on the vHBAs tab.**

The screenshot shows a configuration window titled "I/O Template: LAMirrorSite2 -> vHBA Template: vhma". It contains the following fields and controls:

- Name:** A text input field containing "vhba".
- Storage Cloud:** A dropdown menu with "MirrorSANCloudLA" selected.
- QoS Configuration:** A dropdown menu with "none" selected.
- User defined WWN ID:** A checkbox that is unchecked.
- SAN Boot:** A checkbox that is unchecked.
- Description:** A text input field.
- HA Configuration:** A checkbox that is unchecked.
- LUN Mask Profile:** A dropdown menu with "none" selected.
- WWN Range:** A text input field.

At the bottom left, there are "Submit" and "Cancel" buttons.

5. **Make the required modifications.**
6. **When the properties are configured correctly, click Submit.**

▼ Apply an I/O Template to Servers

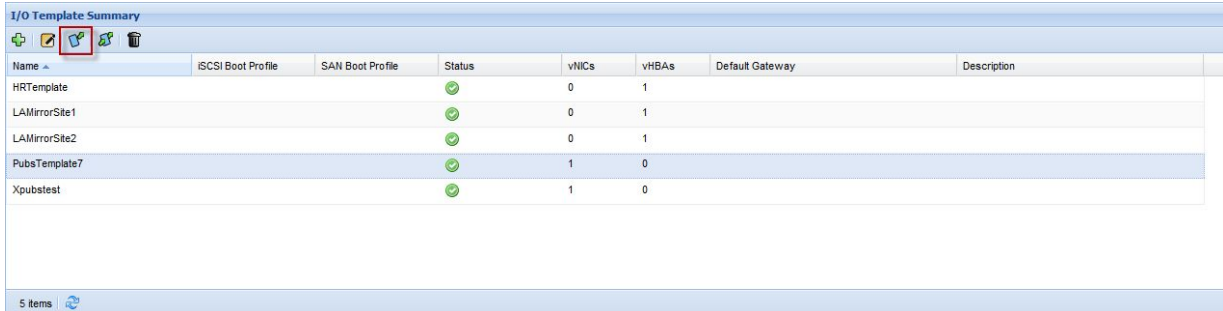
When you create and save an I/O Template, it appears in the I/O Template Summary. At this point the template is available to assign to a server, but the I/O Template is not actually deployed to a host server until you explicitly associate the I/O Template with a host. Applying an I/O Template occurs through the Assign an I/O Template to a Set of Servers button on the I/O Template Summary. The physical server must be "up" for the I/O Template to be successfully applied.

Note - You can also apply a template to a physical server through the Physical Servers page. (See ["Working With Physical Servers" on page 303.](#))

To assign an I/O Template to one or more servers, follow this procedure:

1. **On the I/O Template Summary, select the I/O Template that you want to deploy to one or more hosts.**

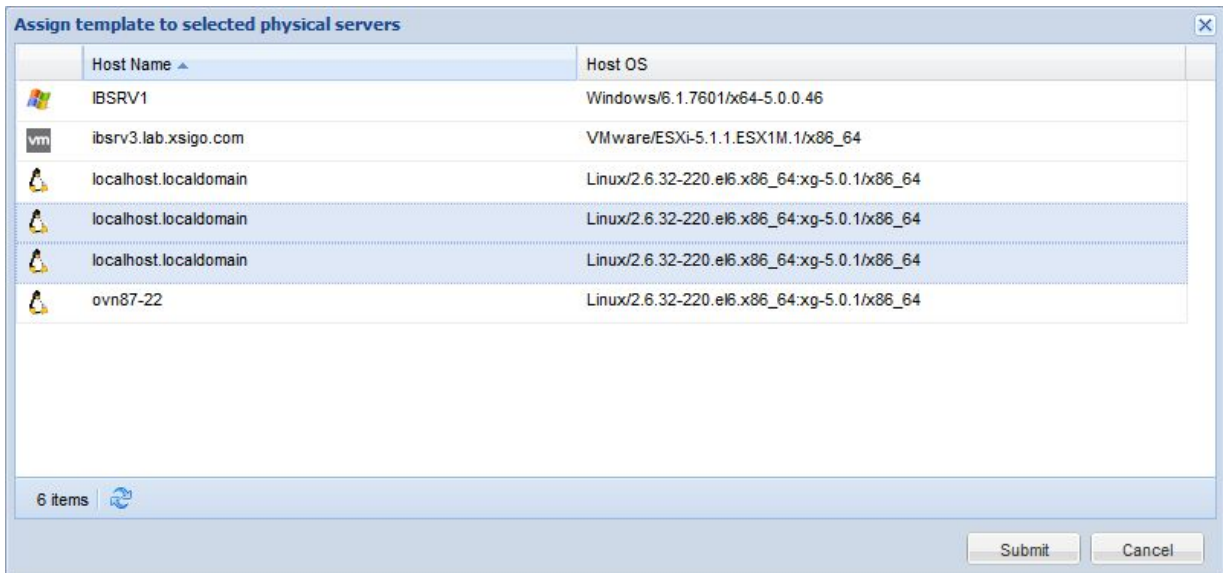
At least one I/O Template must be selected to activate the Assign an I/O Template to a Set of Servers button.



Name	iSCSI Boot Profile	SAN Boot Profile	Status	vNICs	vHBAs	Default Gateway	Description
HRTemplate			✓	0	1		
LAMirrorSite1			✓	0	1		
LAMirrorSite2			✓	0	1		
PubsTemplate7			✓	1	0		
Xpubstest			✓	1	0		

5 items

2. **Click the Assign an I/O Template to a Set of Servers button to display the Assign Template to Servers dialog.**



Host Name	Host OS
IBSRV1	Windows/6.1.7601/x64-5.0.0.46
ibsrv3.lab.xsigo.com	VMware/ESXi-5.1.1.ESX1M.1/x86_64
localhost.localdomain	Linux/2.6.32-220.el6.x86_64:yg-5.0.1/x86_64
localhost.localdomain	Linux/2.6.32-220.el6.x86_64:yg-5.0.1/x86_64
localhost.localdomain	Linux/2.6.32-220.el6.x86_64:yg-5.0.1/x86_64
ovn87-22	Linux/2.6.32-220.el6.x86_64:yg-5.0.1/x86_64

6 items

Submit Cancel

3. **In the Assign Template to Selected Physical Servers dialog, select one or more servers where you want the I/O Template deployed.**

At least one server must be selected to activate the Submit button.

4. Click Submit to deploy the I/O Template to the selected server(s).

Note - If a server is bound (considered “in use”) that server already has an I/O Template. You cannot deploy an I/O Template to an “in use” server. You can, however, migrate the existing I/O Template off of the server, and then deploy another I/O Template to the vacant server.

I/O Template Re-Apply Overview

If an I/O Template has been applied on a physical server, you can use the Re-apply the Template to the Physical Server button to update any changes made to the I/O Template. When the template is changed and re-applied, its associated I/O Profiles receive that change and update the connectivity on the server(s) that originally had the I/O Profiles. Remember that the I/O Profiles were already created and connected to a server through the original template before any changes were made to it.

Re-applying an I/O Template updates the server's I/O Profiles with user-defined configuration changes. The scope of those changes differs based on the type of change that was made, and depending on the change a service interruption might occur:

- For additions, the new vHBAs or vNICs added to the template are pushed to the corresponding I/O Profile. Existing vNICs and vHBAs are not changed, and therefore, remain connected to the server and able to support traffic. Consider the following example. Assume server A has 4 vNICs and 4 vHBAs, and you want to add an HA vNIC and an HA vHBA. When you reapply the template, the I/O Profile on the server gains the additional HA vNIC and HA vHBA. During this process, the existing 4 vNICs and vHBAs remain available and able to pass traffic while the new HA vNICs and HA vHBAs are added.
- For changes to existing virtual I/O resources, vNICs and vHBAs are removed and recreated using the updated settings in the I/O Template. Consider the following example. Assume server B has 4 vNICs and 4 vHBAs, and you want to add a Network QoS Profile to the vNICs. When you re-apply the template, the changed vNICs will be removed and recreated with the new QoS parameters. In this example, a service interruption would occur while the vNICs are removed and recreated, but the vHBAs will remain available and able to pass traffic because no changes were made to them.

Note - Be aware that if you re-apply vHBA changes to an I/O Template can sometimes cause a service interruption. With vHBAs changes, relatively minor changes to existing resources can cause a service interruption while the changed resource is deleted and recreated. For example, applying a more restrictive QoS Profile to a vHBA will cause some amount of service interruption while the vHBA is disconnected, the SAN QoS Profile is applied, then the newly changed vHBA is reconnected. If you will be changing vHBAs, consider doing so during a scheduled service window.

Re-applying is similar to assigning an I/O Template except the list of available physical servers is limited to only those servers that originally had the I/O Template. Reapplying an I/O Template is useful when you have made some changes to an I/O Template, and want to “put it back” on the servers.

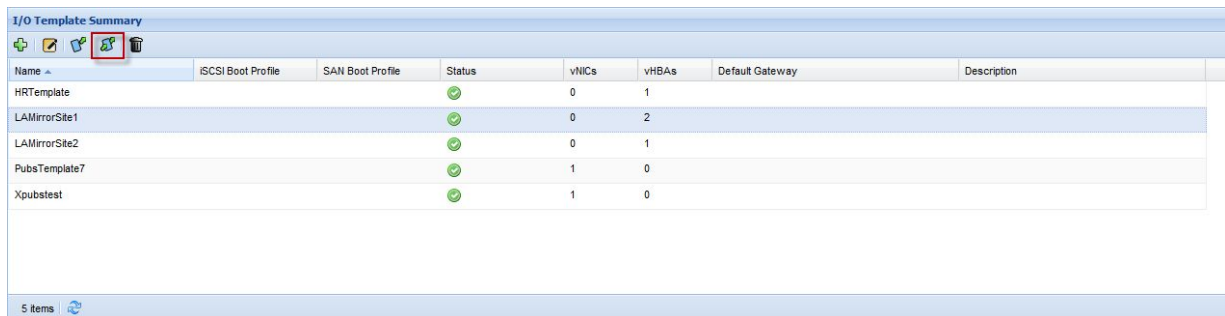
Re-applying an I/O Template to a server is supported only through the I/O Template Summary. Only one template can be applied at a time.

▼ Re-Apply an I/O Template

To re-apply an I/O Template to one or more servers, follow this procedure:

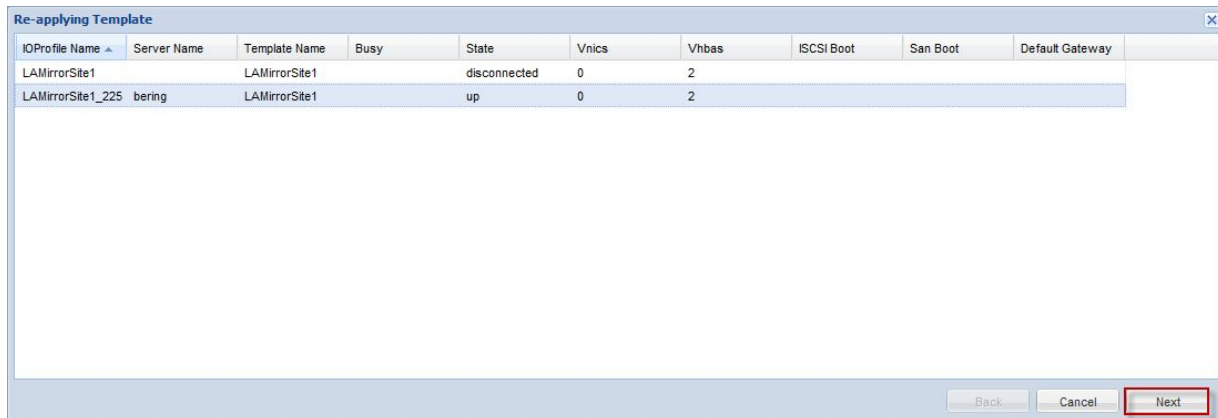
1. **Select an I/O Template that is already assigned to a server.**
2. **Edit the I/O Template as needed and save the changes.**
3. **On the I/O Template Summary, select the I/O Template that you just edited.**

The I/O Template must be selected to activate the Re-apply the Template to the Physical Servers button.



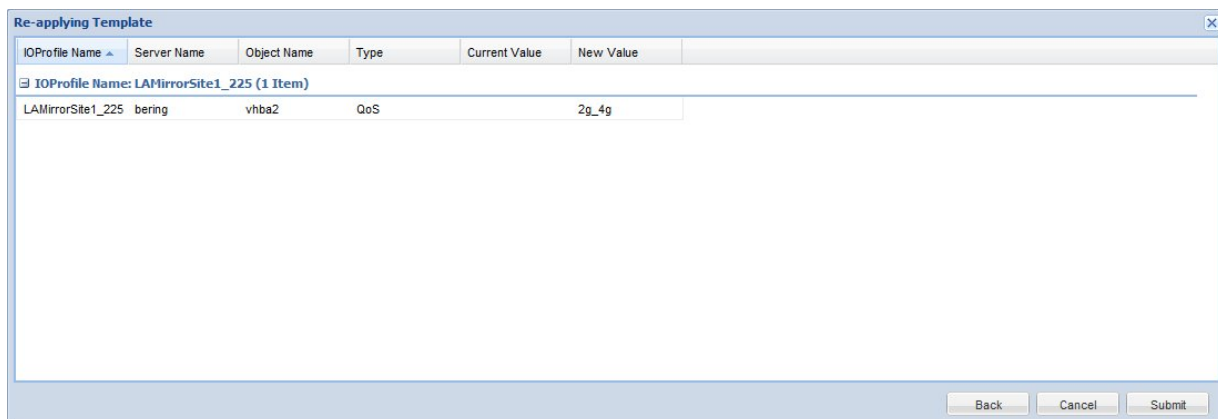
Name	iSCSI Boot Profile	SAN Boot Profile	Status	vNICs	vHBAs	Default Gateway	Description
HRTemplate			✓	0	1		
LAMirrorSite1			✓	0	2		
LAMirrorSite2			✓	0	1		
PubsTemplate7			✓	1	0		
Xpubstest			✓	1	0		

- Click the Re-Apply the I/O Template to the Physical Servers button to display the Re-applying Template dialog.



This dialog shows all the I/O Profiles derived from the I/O Template. In this example, the SAN QoS has been added to a vHBA in the I/O Template, and a new I/O Profile for the that change appears as selectable for the template.

- Select the I/O Profile that represents the updated I/O Template.
- Click Next to display the Re-Applying I/O Template dialog.



In this example, the template being reapplied has an I/O Profile that contains "vHBA2" which has had a QoS Profile of 2g_4g applied.

7. Click Submit to redeploy the I/O Template to the selected server(s).

The changes to the template are reflected in the I/O Profile that is assigned to the server.

Allowed VLANs Feature

With the Allowed VLANs feature, you can specify a list of VLANs that are allowed to pass over any trunk vNICs. (Access mode vNICs cannot receive the Allowed VLANs setting.) The Network Cloud has a set of ranges of VLAN IDs that are allowed, and only traffic that is in the specified VLAN range is allowed to pass over trunk vNICs connected to the Network Cloud. By default, this option is set to allow all VLANs (1-4095) on the Network Cloud. Allowed VLANs are configured per network, so the server will receive the associated VLAN traffic when one or more vNICs terminated in a Network Cloud are deployed to any number of servers.

Note - This feature is also configurable at the Network Cloud level so that all vNICs that connect to the Network Cloud receive the same range of allowed VLANs. For information, see [“Working With Network Clouds” on page 99](#)

As an alternative, you can also assign Allowed VLANs to individual vNICs deployed to a Physical Server. To do so, click the vNICs tab on the Physical Server Details frame, then click the vNIC name, then click the plus sign (+) on the VLAN Ranges tab.

▼ Configure Allowed VLANs in an I/O Template


By specifying the Allowed VLANs at the I/O Template level, you gain some flexibility in which vNICs have Allowed VLANs. When the I/O Template is deployed to a server, only the vNICs that are pushed to that server get the configured Allowed VLANs. By contrast, if you assign the Allowed VLANs feature at the Network Cloud level, all vNICs connected to the cloud have the same Allowed VLANs.

By using an I/O Template, you can set the allowed VLANs feature on specific vNICs in an I/O Template, regardless of which Network Cloud the I/O Template is using.

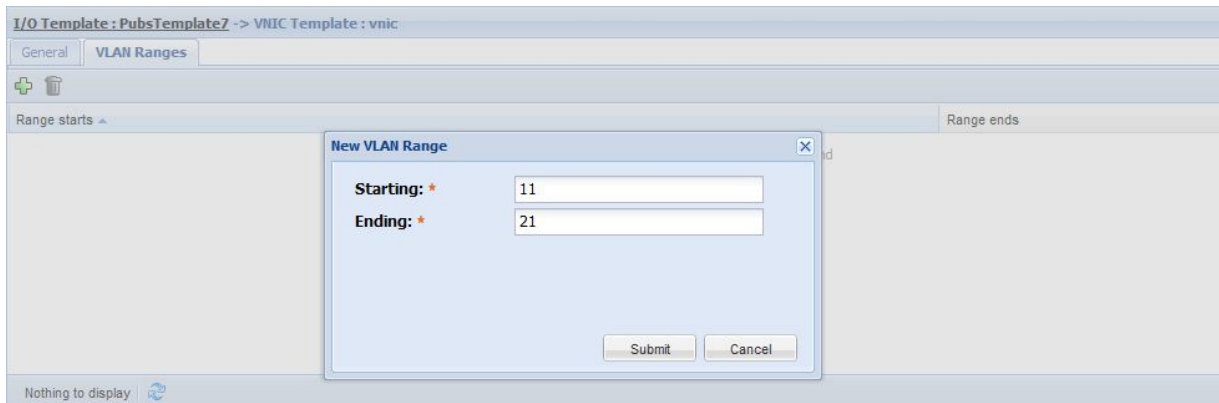
To configure allowed VLANs for an I/O Template, you first create the I/O Template with at least one vNIC in it. The vNIC on which you want to create allowed VLANs must be configured as a trunk-mode vNIC. After the trunk mode vNIC is part of a saved template, you then edit the vNIC to apply the VLAN ranges. You can check the Trunk Mode column of the vNIC Details frame. If the column displays "true", the vNIC will accept Allowed VLANs. Otherwise, you will need to edit the vNIC to click the Trunk Mode checkbox, which sets the vNIC as a trunk-mode vNIC, before continuing.

Follow this procedure:

1. Create the I/O Template as documented in [“Creating an I/O Template” on page 206](#).
2. When the I/O Template is created, select Server Resource Manager → I/O Templates to display the I/O Template Summary.
3. On the I/O Template Summary, click the name of the I/O Template where you want to set the allowed VLANs.
This step populates the I/O Template details frame.
4. In the I/O Template Details frame, click the vNICs tab to display the vNICs associated with the I/O Template.



5. Click the name of the vNIC for which you want to configure the Allowed VLANs.
6. Click the VLAN Ranges tab for the selected vNIC.
7. On the VLAN Ranges tab, click the plus sign (+) to display the New VLAN Range dialog.



8. **In the Starting field, enter the first VLAN ID that you want carried on the trunk VLAN.**
 - **If you want only one VLAN on the trunk vNIC, set the same VLAN ID for the starting and ending fields. For example, to have only VLAN ID 256 carried on the vNIC, set 256 as the starting value and ending value.**
 - **If you want multiple individual VLANs on the trunk vNIC, set the single VLAN ID (as documented in the previous bullet), and do this one time for each of the VLAN IDs. For example, if you want VLANs 256, 512, and 1024 carried on the trunk vNIC, you would need to create one range for each VLAN.**
 - **If you want multiple VLAN ranges on the trunk vNIC, set the appropriate VLAN ranges on the same vNIC. You will need to complete this procedure multiple times.**
9. **In the Ending field, enter the last VLAN ID that you want carried on the trunk VLAN.**
10. **When the VLAN Range has been specified, click Submit to complete the configuration of the Allowed VLAN Range.**

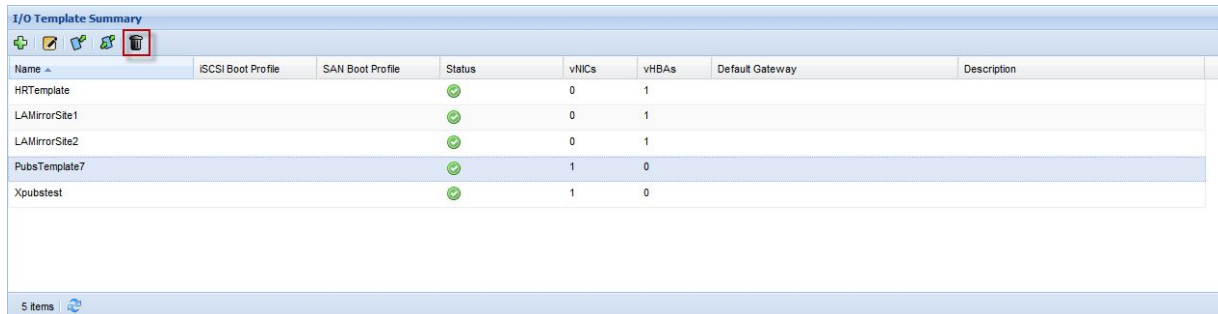
▼ Delete an I/O Template

Anytime an I/O Template is configured, it can be deleted from the I/O Template Summary. When you delete the I/O Template, it is no longer available to any servers that use it. Any servers that are running with vNICs and vHBAs provided through the I/O Template will continue to run with the vNICs and vHBAs. If the template is deleted, any existing I/O Profiles derived from that I/O Template remain, but new I/O Profile cannot be created from the non-existent I/O Template.

To delete an I/O Template, follow this procedure:

1. **Display the I/O Template Summary.**

2. Select the I/O Template(s) that you want to delete, then click the garbage can icon.



The screenshot shows a table titled "I/O Template Summary" with a toolbar at the top containing icons for add, edit, refresh, and delete. The delete icon (a trash can) is highlighted with a red box. The table has the following columns: Name, ISCSI Boot Profile, SAN Boot Profile, Status, vNICs, vHBAs, Default Gateway, and Description. The rows are: HRTemplate, LAMirrorSite1, LAMirrorSite2, PubsTemplate7, and Xpubstest. The Status column contains green checkmarks for all rows.

Name ^	ISCSI Boot Profile	SAN Boot Profile	Status	vNICs	vHBAs	Default Gateway	Description
HRTemplate			✓	0	1		
LAMirrorSite1			✓	0	1		
LAMirrorSite2			✓	0	1		
PubsTemplate7			✓	1	0		
Xpubstest			✓	1	0		

5 items

3. On the confirmation dialog, click Yes to confirm the deletion. Or, you can click No to abort the deletion and leave the I/O Template configured.

Working With I/O Profiles

This chapter contains the following topics:

- [“I/O Profiles Overview” on page 259](#)
- [“I/O Profiles Summary” on page 260](#)
- [“Create an I/O Profile” on page 262](#)
- [“Connect Specific I/O Profiles to Specific Fabric Devices” on page 264](#)
- [“Save an I/O Profile as an I/O Template” on page 266](#)
- [“Connecting an I/O Profile to a Server” on page 266](#)
- [“Process for Merging Multiple I/O Profiles Into One I/O Profile” on page 267](#)
- [“Merge I/O Profiles” on page 268](#)
- [“Disconnect an I/O Profile From a Server” on page 273](#)
- [“Link an I/O Profile to an I/O Template” on page 273](#)
- [“Delete an I/O Profile” on page 275](#)
- [“Displaying I/O Profile Details” on page 276](#)

I/O Profiles Overview

I/O Profiles provide a more simplified connection policy that specifies the number of vNICs, vHBAs, or both that are required for a specific server.

I/O Profiles are a software construction that act as a container for virtual I/O resources. I/O Profiles hold all the virtual I/O information for a particular server under management by Oracle Fabric Manager. The virtual I/O information in the I/O Profile is related to vNICs and HA vNICs, vHBAs and HA vHBAs, all server profiles for a server, and all Oracle objects on the Fabric Devices that connect to the server.

An I/O Profile is derived from an I/O Template. The I/O Template provides the basic scope and type of connections that the server has, and the I/O Template is used to create one or more I/O Profiles for a server. However, it is the I/O Profile—not the I/O Template—that is the object that actually gets connected to a server. When the I/O Profile is connected to the server, the vNICs and vHBAs in the I/O Template are pushed to the server.

Because I/O Profiles are created from I/O Templates, the workflow for deploying vNICs and vHBAs on a server has changed if you are deploying connectivity through I/O Templates. The new workflow is:

1. Create an I/O Template
2. Create one or more I/O Profiles from the I/O Template
3. Connect an I/O Profile to a server to deploy the vNICs or vHBAs to the server

I/O Profiles are also used when you create vNICs and vHBAs directly on a physical server (without an I/O Template) through Physical Server details page. In this case, an I/O Profile is created automatically and assigned the name of the server so that the server, I/O Profile, and all Server Profiles for the server are intuitively and consistently named.

For ease-of use, you can create an I/O Profile without actually connecting it, which enables you to pre-provision an I/O Profile then deploy it when it is needed. For example, if you will be deploying a new server in your datacenter, you can create the I/O Profile before the server is actually built up if you know its connection needs. The ability to create a server's I/O connectivity without actually having to connect that I/O to the server is the main advantage of I/O Profiles.

I/O Profiles Summary

I/O Profiles are displayed in the I/O Profiles Summary, which shows a list of all the I/O Profiles created in Oracle Fabric Manager. I/O Profiles are listed in the I/O Profile Summary regardless of whether or not they are connected to servers. Through the I/O Profile you can see which I/O Profile is connected to which server, as well as other information about connectivity and boot options supported through the I/O Profile.

Add a New I/O Profile

Save I/O Profile as a Template

Connect an I/O Profile to a Physical Server

Disconnect an I/O Profile to a Physical Server

Link the I/O Profile to an I/O Template

Delete an I/O Profile

Name	Server Name	Template Name	Busy	State	vNICs	vHBAs	Boot Profile	Default Gateway
brack	brack.lbb.xsigo.com			up	13	4		
brick	brick.lbb.xsigo.com			up	11	4		
probe				disconnected	6	2		
bugshaker00	bugshaker00			up	0	0		
crooby	crooby			up	5	1		
daisy				disconnected	4	3		
delhede10				disconnected	2	1		
delhede13	delhede13@csxaco.com			up	6	4		
frack				disconnected	4	3		
frick	frick.lbb.xsigo.com			up	0	4		
lodd				disconnected	1	1		

Field	Means
Name	The name of each configured I/O Profile.
Server Name	The name of the server on which the I/O Profile is deployed. If no server name is displayed, the I/O Profile is not connected to any server.
Template Name	The name of the I/O Template (if any) that is linked to the I/O Profile. It is possible to create an I/O Profile and connect it to a server without creating an I/O Template. If no I/O Template name is present, then the I/O Profile is not linked to an I/O Template.
Busy	Indicates the state of the I/O Profile as either in the process of binding to the server or disconnecting from the server.
State	Indicates the current state of the listed I/O Profiles. Valid states are: <ul style="list-style-type: none"> ■ Up, when the I/O Profile is successfully connected to a server. ■ Connected, when the I/O Profile is connected to a server. ■ Disconnected, when the I/O Profile is not connected to a server. This state is the default state of an I/O Profile after it has been created. ■ Partial, when the I/O Profile is not completely connected to a server. This state requires your attention because it usually means that an error has occurred.
vNICs	The number of vNICs that are controlled by the I/O Profile.
vHBAs	The number of vHBAs that are controlled by the I/O Profile.
Boot Profile	The name of the Boot Profile(s) that are controlled by the I/O Profile.
Default Gateway	The default gateway that is controlled by the I/O Profile.

▼ Create an I/O Profile

When you create an I/O Profile, you are specifying the individual vNIC and vHBA connections for each servers. The I/O Profile is the policy that contains the vNICs and vHBAs, and it is the object that can be connected to, or disconnected from, a server to deploy or remove the network and storage connectivity for the server.

I/O Profiles are created from I/O Templates, in any number from 1 to 20. For example, after you have created an I/O Template, if 6 servers need the same connectivity, you can use the I/O Template to create 6 different I/O Profiles. Each I/O Profile can then be connected to a server as needed. This method allows you granular control over which servers get connectivity without requiring you to take the time to create 6 different sets of vNICs and vHBAs. Instead, creation occurs once through the I/O Template, and control over which servers receive connectivity is accomplished on a server-by-server basis.

When I/O Profiles are created from an I/O Template, they are displayed in the I/O Profile Summary. If multiple I/O Profiles are created from the same I/O Template, you give them a name and Oracle Fabric Manager adds an enumerated suffix in the form of _X where X is a number. For example, if you created 6 I/O Profiles for your security auditing servers with the name "AuditServers", then the I/O Profiles would be named "AuditServer_1" through "AuditServer_6". Because of the suffix, each of the I/O Profiles is a unique object in Oracle Fabric Manager, and as a result, can be used and managed separate of the other 5 I/O Profiles.

Creating an I/O Profile starts from the I/O Profile Summary, and requires that an I/O Template has already been created and saved. If you do not have an appropriate I/O Template created, create one now. See ["Working With I/O Templates" on page 205](#).

To create an I/O Profile, follow this procedure:

1. **Select Display the I/O Profile Summary. For an example, see ["I/O Profiles Summary" on page 260](#).**

2. Click the plus sign (+) to display the New I/O Profile dialog.



The screenshot shows a dialog box titled "New I/O Profile". It contains the following fields and controls:

- I/O Profile Name ***: Text input field containing "PoSProcessing".
- Template Name ***: Dropdown menu showing "Pubstemplate1".
- Number of I/O Profiles ***: Text input field containing "5".
- I/O Profile Description**: Empty text area.
- Advanced Configuration**: A section header with a dropdown arrow.
- Submit** and **Cancel**: Buttons at the bottom right.

3. From the Template Name dropdown menu, select the I/O Template that you want to use to create I/O Profiles.

If the I/O Template you want to use is not listed, it has not been successfully created. You will need to cancel out of this procedure and create an I/O Template before returning to this procedure.

4. In the Number of I/O Profiles field, enter a number from 1 to 20 that specifies the number of I/O Profiles that will be created from the selected I/O Template.

When more than one I/O Profile is created, each I/O Profile is created with a numerical suffix appended to the I/O Profile name.

5. In the I/O Profile Name field, enter a name for the I/O Profile that you want to created.

If more than one I/O Profile will be created, the name you specify is a baseline to which a number is appended (for example, PoSProcessing_1).

6. As an option, in the I/O Profile Description field, you can enter an alphanumeric character string that describes the I/O Profile(s) that you are creating.

Note - At this point, you have completed the basic configuration of an I/O Profile. However, you can specify which Fabric Devices are connecting the I/O Profiles by using options in the Advanced Configuration section. See [“Connect Specific I/O Profiles to Specific Fabric Devices” on page 264](#).

7. **Click Submit to create the I/O Profile(s) from the selected I/O Template.**

▼ **Connect Specific I/O Profiles to Specific Fabric Devices**

As part of configuring an I/O Profile, you can select the individual Fabric Devices that are connecting the I/O Profile to a server. For example, assume you have a dual-Fabric Device environment, but you want only one of the Fabric Devices to connect the I/O Profile to the server. Through the New I/O Profile dialog, you can specify which Fabric Devices can support one or more I/O Profiles. By default, any Fabric Device that Oracle Fabric Manager is managing can support the connection, but using the Advanced Configuration section, you can control where the I/O Profile is supported.

To connect I/O Profiles to specific Fabric Devices, follow this procedure:

1. On the New I/O Profile dialog, click the Advanced Configuration button to display the advanced properties for an I/O Profile.

The screenshot shows the 'New I/O Profile' dialog box. The fields are as follows:

- I/O Profile Name ***: PoSProcessing
- Template Name ***: techpubs10
- Number of I/O Profiles ***: 1
- I/O Profile Description**: (Empty text area)
- Advanced Configuration** (Expanded):
 - Select Fabric Devices:**
 - aaaa
 - oregon
 - delaware

Buttons: Submit, Cancel

As shown in the example, all Fabric Devices that Oracle Fabric Manager is managing are listed. If you have specific connection needs for the I/O Profile, you can make sure that the correct Fabric Device(s) connect the I/O Profile to the server by selecting the Fabric Device(s) from the pick list.

2. **Select the Oracle Fabric Devices that need to support the I/O Profile.**
3. **Click Submit.**

▼ Save an I/O Profile as an I/O Template

When an I/O Profile is created, you can save it as an I/O Template. If you do, the current I/O Profile is created as a brand new I/O Template and is added to the I/O Template Summary. This functionality allows you to create a new I/O Template after customizing an I/O Profile. For example, you could create a simple I/O Template with a single vNIC and vHBA that you use to create an I/O Profile. After using the I/O Profile, you might promote a server to a more important role which requires HA vNICs and multipath vHBAs. You could then modify the I/O Profile, and save it as an I/O Template which could then be used as the source for creating additional I/O Profiles featuring HA.

To save an I/O profile, follow this procedure:

- 1. Display the I/O Profile Summary.**
For an example, see [“I/O Profiles Summary” on page 260](#).
- 2. Click the Save as I/O Template button to display the Save as I/O Template dialog.**
- 3. In the Template Name field, enter an alphanumeric character string that names the I/O Template you are creating, or accept the I/O Profile name which is supplied by default.**
- 4. Click Save.**

▼ Connecting an I/O Profile to a Server

An I/O Profile is the container for vNICs and vHBAs, but it doesn't actually provide any connectivity until you connect that I/O Profile to a server. The I/O Profile is the object that you connect to a server, and when you do so, the vNICs and vHBAs contained in the I/O Profile are pushed to the host.

An I/O Profile in the “disconnected” state can be applied to a server. If you want to connect an I/O Profile that is already connected to a server, you must explicitly disconnect that I/O Profile first. Also, an I/O Profile can be connected to a server that is in any state, but only when the server is fully up and online will the I/O Profile support traffic.

But remember that an I/O Profile is not automatically connected to a server when it is created. You must manually connect the I/O Profile to a server to deploy the vNICs or vHBAs to the server.

To connect an I/O Profile follow this procedure:

- 1. Display the I/O Profile Summary.**
For an example, see [“I/O Profiles Summary” on page 260](#).

2. **Click an I/O Profile to select it.**
This step activates the Connect I/O Profile to a Server button.
3. **Click the Connect to a Server button to display the Choose a Server to Connect dialog.**
4. **Select the server(s) to which you want to connect the I/O Profile.**
5. **Click Submit to connect the I/O Profile to the selected server.**

Process for Merging Multiple I/O Profiles Into One I/O Profile

In an HA Fabric Device deployment, multiple I/O Profiles can be connected to a server. This feature is useful in situations where a high-powered server (for example, an Oracle T5 or M5 server) has the processing power to support more than 2 Server Profiles without easily exhausting the bandwidth of the HCA port connection.

Multiple server profiles per Fabric Device are supported for each server HCA port by creating individual I/O Profiles then merging them two-at-a-time until they all become one I/O Profile. For example, if you had 3 I/O Profiles that you wanted to apply a single server, you would merge I/O Profile 1 and I/O Profile 2. When that merge completed, you would then have one I/O Profile. You would then merge I/O Profile 3 with the single I/O Profile.

Be aware of the following considerations when merging I/O Profiles:

- Two I/O Profiles are required.
- Any vNICs and vHBAs in the two I/O Profiles must be uniquely named. The same vNIC or vHBA name cannot be used in both I/O Profiles.

If you add another vNIC or vHBA from Oracle Fabric Manager after both I/O Profiles are merged, the new vNIC or vHBA will be added on either of the Server Profiles on that Fabric Interconnect where the vNIC or vHBA is terminated. As a result, vNIC and vHBAs are randomly assigned to either I/O Profile but cannot deterministically be assigned to a specific I/O Profile.

Merging multiple I/O Profiles is supported on two Fabric Devices, each of which has one port connection to the server. The merging procedure takes the following phases:

1. Create one I/O Profile (either manually or through an I/O Template).
2. Connect the I/O Profile to a server with 4 ports.
3. Create one I/O Profile (either manually or through different I/O Template). It is critical that this I/O Profile is not connected to anything. This second I/O Profile must be created in disconnected state.

Note - You do not necessarily need to create a different I/O Template but you do need a second I/O Profile. What is important is that the vNICs and vHBAs are uniquely named in each of the I/O Profiles.

If you use the same I/O Template to create a second I/O Profile, make sure that the vNICs and vHBAs have different names in each I/O Profile.

4. Connect the second I/O Profile to the same server that has the first I/O Profile.

The specific procedure for this configuration is documented later in this section.

▼ Merge I/O Profiles

Assuming you have cabled a T5- or M5-series server (or another server that has 4 or more server ports) to Oracle Fabric Manager, and that Oracle Fabric Manager sees the Fabric Device and the server, you can connect the server as follows:

Note - This procedure assume that you are creating an I/O Profile from an I/O Template. If you will be creating the I/O Profile manually, you can start the procedure at [Step 5](#).

1. **In the Physical Server Summary, locate the server with multiple ports/GUIDs.**

It will be the server that has 4 or more IB ports displayed on the Fabric Device—for example:

Host Name	Host OS	Adapter FW Vers...	vNL...	vHB...	Bound	Busy	State	I/O Profile Name	Director Ports	Groups
bering	Linux/2.6.18-238.el5.xg-3.7.1.LX1/x...	5.3.0/3.0.0	0	0					oregon:ServerPort15 delaware:ServerPort22	
CHARCOT-W2K8	Windows/6.1.7100/x64-2.2.0.36	5.3.0/2.1.1	0	0					delaware:ServerPort20 oregon:ServerPort17	
COLEMAN	Windows/6.1.7600/x64-3.0.0.26	5.3.0/2.1.1	0	0					oregon:ServerPort13 delaware:ServerPort11	
component	Linux/2.6.32-131.0.15.el6.x86_64.xg...	5.3.0/3.0.0,5.2.0/...	0	0					delaware:ServerPort19 delaware:ServerPort23 oregon:ServerPort14 oregon:ServerPort3	
ovn86-45	VMware/ESX-4.1.0.xg-3.2.0-r5773.E...	5.3.0/3.0.0	0	0					oregon:ServerPort1 delaware:ServerPort9	

2. **Under Server Resource Manager on the navigation pane, select I/O Templates to display the I/O Template Editor.**
3. **In the Name field, enter the name for the I/O Template that you are creating.**

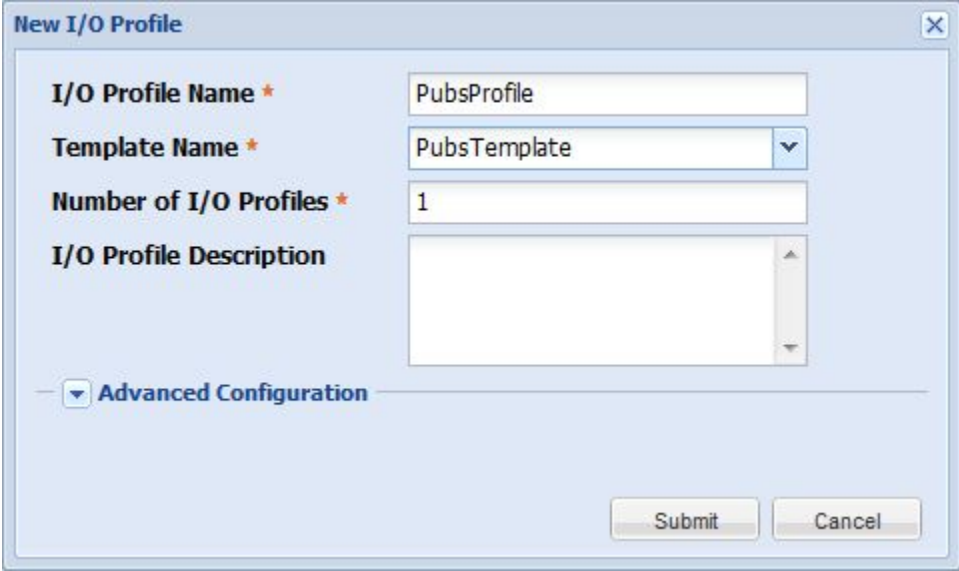
If the server will boot through PXE iSCSI or PXE SAN Boot, select the name of the iSCSI Boot or SAN Boot profile from the corresponding dropdown menu.

4. Create the I/O Template for the I/O Profile.

For information about creating an I/O Template, see [“Create a New I/O Template” on page 208](#).

Note - When the I/O Template is created, click Save to save the template to the Oracle Fabric Manager server.

- 5. Under Server Resource Manager on the navigation pane, select I/O Profiles to display the I/O Profile Summary.**
- 6. On the I/O Profile Summary, click the green plus sign (+) to display the Create I/O Profile dialog.**

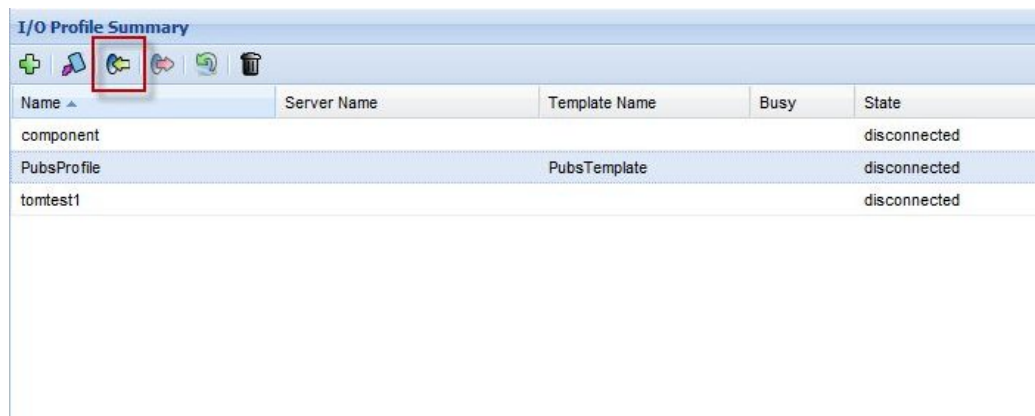


The screenshot shows a dialog box titled "New I/O Profile". It contains the following fields and controls:

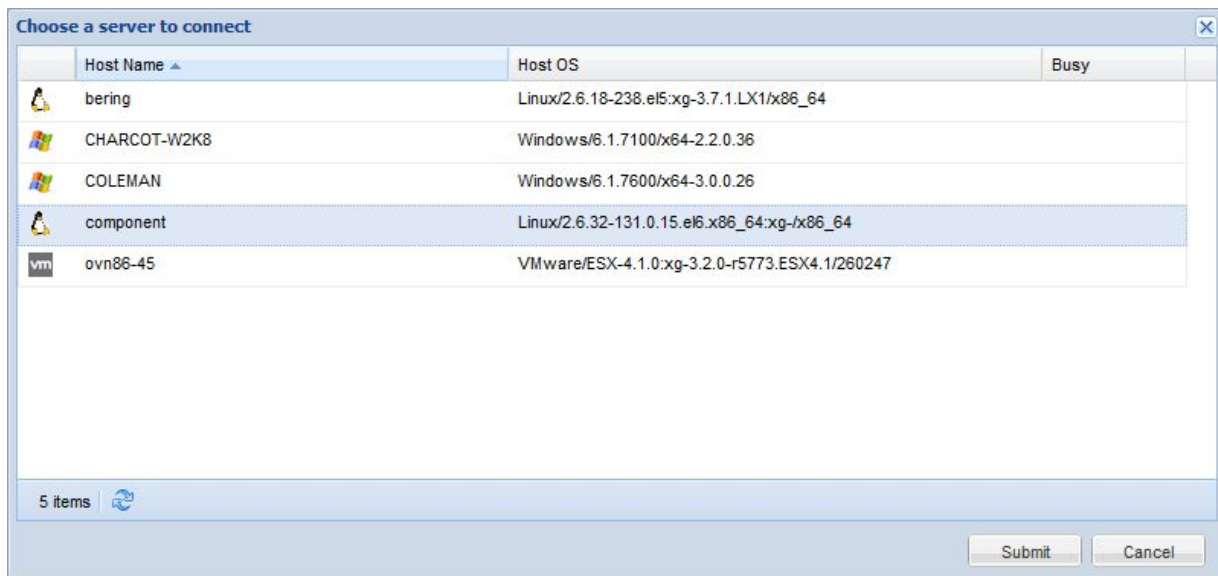
- I/O Profile Name ***: Text input field containing "PubsProfile".
- Template Name ***: Dropdown menu with "PubsTemplate" selected.
- Number of I/O Profiles ***: Text input field containing "1".
- I/O Profile Description**: Empty text area.
- Advanced Configuration**: A checkbox that is currently unchecked.
- Submit** and **Cancel** buttons at the bottom right.

- 7. In the I/O Profile Name field, enter the name for the I/O Profile that you are creating.**
- 8. From the Template Name dropdown menu, select the I/O Template you created earlier in this procedure.**
- 9. In the Number of I/O Profiles field, select the number of profiles that you want to be derived from this I/O Template.**

10. As an option, in the Description field, you can provide a description for the I/O Profile you are creating.
11. When the I/O Profile parameters are specified, click Submit to finish creating the I/O Profile(s) which will appear in the I/O Profile Summary.
12. On the I/O Profile Summary, select the I/O Profile you just created. This step activates the toolbar buttons.



13. On the toolbar, click **Connect the Selected I/O Profile to a Physical Server** to display the **Choose a Server to Connect** dialog.

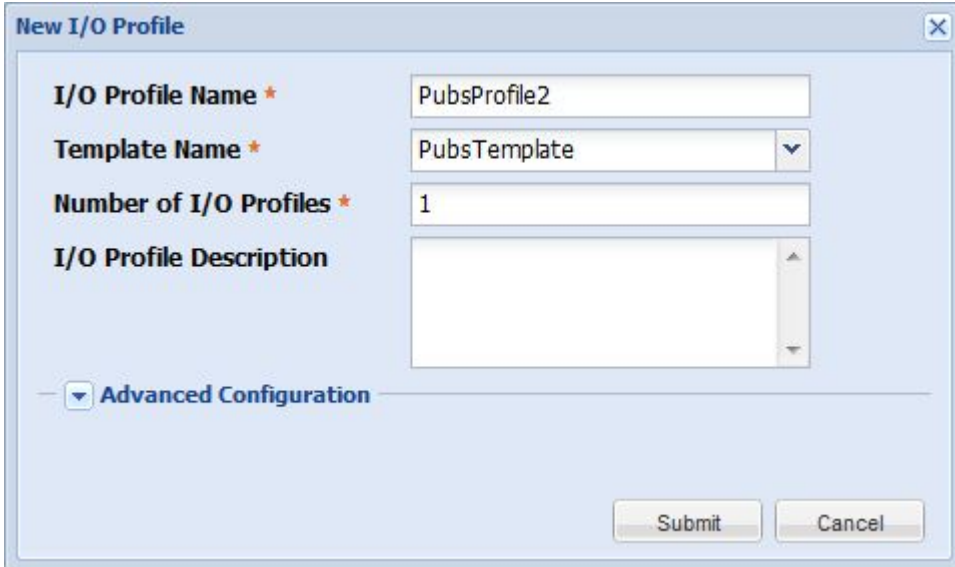


14. **Select the T5/M5 server.**
15. **Click Submit to push the I/O Profile to the selected server.**
 When a confirmation dialog is displayed, click Yes to accept deploying the I/O Profile.
 The first I/O Profile should be connected to the server.

Name	Server Name	Template Name	Busy	State	VNI
component				disconnected	1
PubsProfile	component	PubsTemplate		up	1

16. **Create the second I/O Profile by repeating Step 5 through Step 10.** The following example shows creating the second I/O Profile from the same I/O Template.

Note - When you create the second I/O Profile (either from the same I/O Template or a completely different one), make sure that the vNICs are named differently between the two I/O Profiles, and make sure that the vHBAs are named differently between the two I/O Profiles.



The screenshot shows a dialog box titled "New I/O Profile". It contains the following fields and controls:

- I/O Profile Name ***: Text input field containing "PubsProfile2".
- Template Name ***: Dropdown menu showing "PubsTemplate".
- Number of I/O Profiles ***: Text input field containing "1".
- I/O Profile Description**: Empty text area.
- Advanced Configuration**: A collapsed section indicated by a minus sign and a dropdown arrow.
- Submit** and **Cancel**: Buttons at the bottom right.

17. **When the parameters for the second I/O Profile are specified, click Submit to finish creating the I/O Profile(s).**

The second I/O Profile appears in the I/O Profile Summary in disconnected state, which is correct and required to connect the second I/O Profile to the T5/M5 server.

18. **On the I/O Profile Summary, select the disconnected I/O Profile.**
19. **On the toolbar, click Connect the Selected I/O Profile to a Physical Server to display the Choose a Server to Connect dialog.**
20. **Select the same T5/M5 server where the first I/O Profile was connected.**
21. **Click Submit to connect the second I/O Profile to the T5/M5 server.**

On the I/O Profile Summary, both I/O Profiles will be displayed for a short time. The second I/O Profile will transition through some state changes, typically from disconnected to down and down to up.

When the state changes are complete, the second I/O Profile will merge into the first so that only one I/O Profile shows connected to the server. The one I/O Profile will be in up state.

I/O Profile Summary					
Name	Server Name	Template Name	Busy	State	VNICs
component				disconnected	1
PubsProfile	component	PubsTemplate		up	2
tomtest1				disconnected	2

You can verify that the two I/O Profiles merged by checking the vNICs and vHBAs columns on the I/O Profile Summary. The total shown should be the sum of the vNICs and vHBAs in each of the I/O Profiles.

▼ Disconnect an I/O Profile From a Server

An I/O Profile in the “connected” state can be removed from a server. If you want to remove the vNICs and vHBAs from a server without deleting the I/O Profile, you can simply disconnect the I/O Profile. All the connectivity in the I/O Profile remains intact and can be used again on the same server or another. By disconnecting the I/O Profile, you will remove the network and storage connections for a server, which will cause a service interruption, but does not affect the admin or operational state of the server itself.

If you want to completely delete an I/O Profile, you must first disconnect the I/O Profile.

To disconnect an I/O Profile, follow this procedure:

1. **Display the I/O Profile Summary.**
2. **Select the I/O Profile that you want to disconnect.**
3. **Click the Disconnect I/O Profile from Server button, which causes a confirmation dialog to be displayed.**
4. **On the confirmation dialog, click Yes to disconnect the selected I/O Profile, or lick No to abort the disconnect.**

▼ Link an I/O Profile to an I/O Template

When an I/O Profile is not associated with an I/O Template, you can link it to an I/O Template, as long as the number and name of the vNIC and vHBAs are the same. If you do, the vNICs and

vHBAs in the I/O Profile are added to an existing I/O Template. With this feature, you can start with an I/O Template and add virtual connectivity to it over time to create a more robust I/O Template.

To link an I/O Profile to an I/O Template, follow this procedure:


- 1. Display the I/O Profile Summary.**
For an example, see [“I/O Profiles Summary” on page 260](#).
- 2. Select the I/O Profile that you want to link to an I/O Template.**
- 3. Click the Link I/O Profile to I/O Template button to display the Choose a Template to Link dialog.**



- 4. Select an I/O Template from the dialog, then click Sumbit.**

Only Templates with a green checkmark in the Status column can be linked to an I/O Profile.

Choose a template to link							
Name ▲	iSCSI Boot Pr...	SAN Boot Pr...	Status	vNICs	vHBAs	Default Gateway	Description
cloud9template			✓	2	2		
EnglOtemplate			✓	1	1		IO Template for Engineering R...
HRtemplate			✗	0	0		Standard Human Resources i...
PoSprocessing_1			✓	1	1		
Pubstemplate1			✓	1	1		for Tech Pubs use
scsi4	iscsiboot_4		✓	1	0		
test1			✗	0	0		
ttttt			✓	1	0		
vmcto			✓	2	1		

9 items | 

▼ Delete an I/O Profile

An I/O Profile can be deleted at any time, even if the I/O Profile is connected to a server. When you delete an I/O Profile, you are removing the I/O Profile and its vNICs and vHBAs from Oracle Fabric Manager. A deleted I/O Profile is not kept, so if you need to use it later, you will need to recreate it. Deleting an I/O Profile removes all the vNICs and vHBAs bound to that server, so if you delete an I/O Profile, you will disconnect its network and storage connections, and service will be interrupted.

To delete an I/O Profile, follow this procedure:

1. **Display the I/O Profile Summary.**
2. **Select the I/O Profile that you want to delete.**
3. **Click Delete I/O Profile (the garbage can icon) to display a confirmation dialog.**
4. **Click Yes to delete the selected I/O Profile.**

Displaying I/O Profile Details

Through the I/O Profile Summary, you can display additional details about individual I/O Profiles. By clicking an I/O Profile, the details frame is populated with additional details about the profile. The I/O Profile Details frame contains the following tabs:

- [“I/O Profile General Properties” on page 276](#)
- [“vNICs in an I/O Profile” on page 277](#)
- [“vHBAs in an I/O Profile” on page 277](#)
- [“Server Profiles in an I/O Profile” on page 278](#)
- Boot Info, which is conditional. The presence of this tab depends on whether a Boot Profile for iSCSI, SAN Boot, or PXE Boot is configured in the I/O Template that was used to create the I/O Profile.

I/O Profile General Properties

General properties are available for the I/O Profile through the General tab. General properties include the I/O Profile's name, information about the server where the I/O Profile is connected, and the I/O Template that was used to derive the I/O Profile.

Name	Server Name	Template Name	Busy	State	vNICs	VHBAs	Boot Profile	Default Gateway
ABC	ovn86-43			up	4	3		
mlathe4	ovn86-42			up	1	0		
mlathe5				disconnected	1	0		
mlathe55_2				disconnected	1	0		
mlathe55_3				disconnected	1	0		
ovn86-40	ovn86-40			up	3	0		
test		test		disconnected	0	1		
test2_1				disconnected	3	0		
test_3				disconnected	3	0		

Field...	Means...
Display Name	The name of the I/O Profile
Description	Is an optional alphanumeric character string that describes the I/O Profile.
Host Name	Is the name of the server to which the I/O Profile is connected.
Template Name	Is the name of the I/O Template that was used to derive the I/O Profile.

Field...	Means...
State	Is the state of the I/O Profile. Valid states are: <ul style="list-style-type: none"> ■ Connected ■ Partial ■ Disconnected ■ Down
Default Gateway	The name of the default gateway for the server where the I/O Profile is connected.

vNICs in an I/O Profile

An I/O Profile controls one or more vNICs that are originally assigned through the I/O Template that was used to create the I/O Profile. However, after the I/O Profile is created, you can customize the I/O Profile by using the toolbar buttons on the vNICs tab.

Add vNIC to Profile

Set vNIC to Up State

Set vNIC to Down State

Set vNIC Termination to Different Network Cloud

Change vNIC Termination to Different Port or LAG

Merge Single vNICs into HA vNIC

Delete Selected vNIC

Name	Network Cloud	Termination	State	IP Address	Network	IP Type	MAC Address	HA	QoS
m1h01	discovered-network-c...	oregon/6/2	up/hul up/resource/Java...	0.0.0.0	255.255.255.255	hostVenaqst1		true	Disabled
m1h02	discovered-network-c...	oregon/6/3	up/hul up/resource/Java...	0.0.0.0	255.255.255.255	hostVenaqst1		true	Disabled

vHBAs in an I/O Profile

An I/O Profile controls one or more vHBAs that are originally assigned through the I/O Template that was used to create the I/O Profile. However, after the I/O Profile is created, you can customize the I/O Profile by using the toolbar buttons on the vHBAs tab.



Server Profiles in an I/O Profile

I/O Profiles contain all the objects on a server, including I/O Templates and Server Profiles. You can display the server profile(s) that are controlled through an I/O Profile by selecting the Server Profiles tab on the I/O Profile details frame.

In addition to displaying detailed information about the server profiles associated with an I/O Profile, you can also perform basic management functions for the server profile, such as turn up, shut down, and reset of the server profile.

Set the Server Profile to Up State

Shutdown the Server Profile

Reset the Server Profile

Name	State (Admin/Oper)	Fabric Device	Fabric Device Port	Busy	Default Gateway	Description
m	up/unassigned	mrg37		false		

1 item

Working With Link Aggregation

This chapter contains the following topics:

- [“Link Aggregation Groups” on page 281](#)
- [“Configure Link Aggregation Groups” on page 281](#)
- [“Displaying LAG Properties” on page 285](#)

Link Aggregation Groups

A Link Aggregation Group (LAG) enables you to combine multiple individual physical Ethernet ports into one logical port group. As a result, the ports combined into a LAG can operate in parallel with the benefit of increased link speed and high availability.

Oracle Fabric Manager supports LAG at the I/O module level. When you configure LAG, you specify a group name for the LAG, then assign ports from the same Ethernet I/O module to the group. LAGs are supported on the 4-Port 10 GE module and the 10-Port GE module only (not the 10 GE module).

LAGs are associated with a Network Cloud and, just like Ethernet ports, can be the termination point for vNICs. If you will be provisioning LAGs in your network, you will need to create them before associating a Network Cloud with them. By doing so, you make the LAG available as a selectable object in the dialog that supports Network Cloud creation.

Each LAG can be supported as a static LAG or passive-mode LAG:

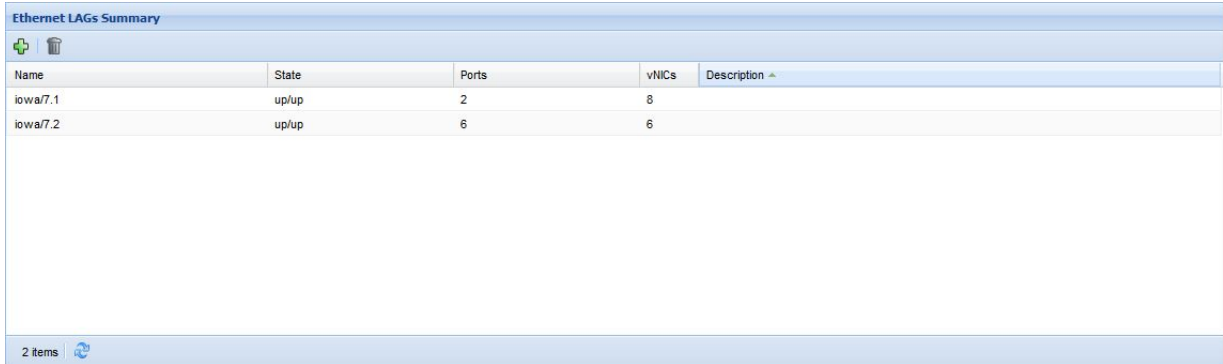
- After initially specifying the ports that are members of the LAG, static LAGs require user intervention to add ports to and delete ports from a LAG.
- After initially specifying the ports that are members of the LAG, passive-mode LAGs use Link Aggregation Control Protocol (LACP) in combination with a neighboring Gigabit Ethernet switch to dynamically control port additions and deletions within a LAG.

▼ Configure Link Aggregation Groups

To configure Link Aggregation Groups, follow this procedure:

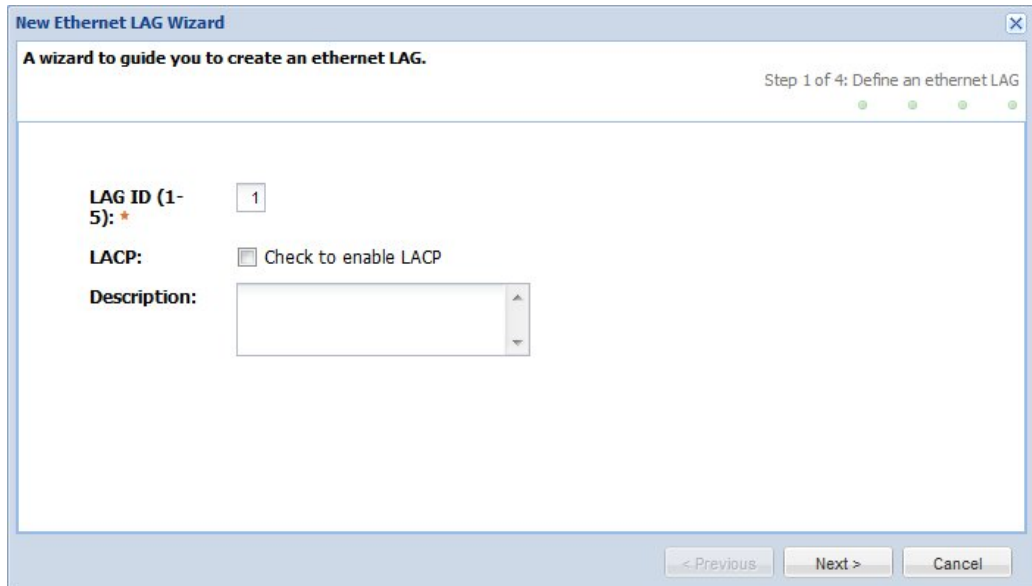
1. **Display the Ethernet LAG Summary by selecting Network Cloud Manager → Link Aggregation Groups.**

LAGs can be added by clicking the plus sign (+), and LAGs can be deleted by clicking the garbage can icon.



Name	State	Ports	vNICs	Description
iowa/7.1	up/up	2	8	
iowa/7.2	up/up	6	6	

2. **Click the plus sign (+) to start the Create Link Aggregation wizard and add a LAG to Oracle Fabric Manager.**



New Ethernet LAG Wizard

A wizard to guide you to create an ethernet LAG.

Step 1 of 4: Define an ethernet LAG

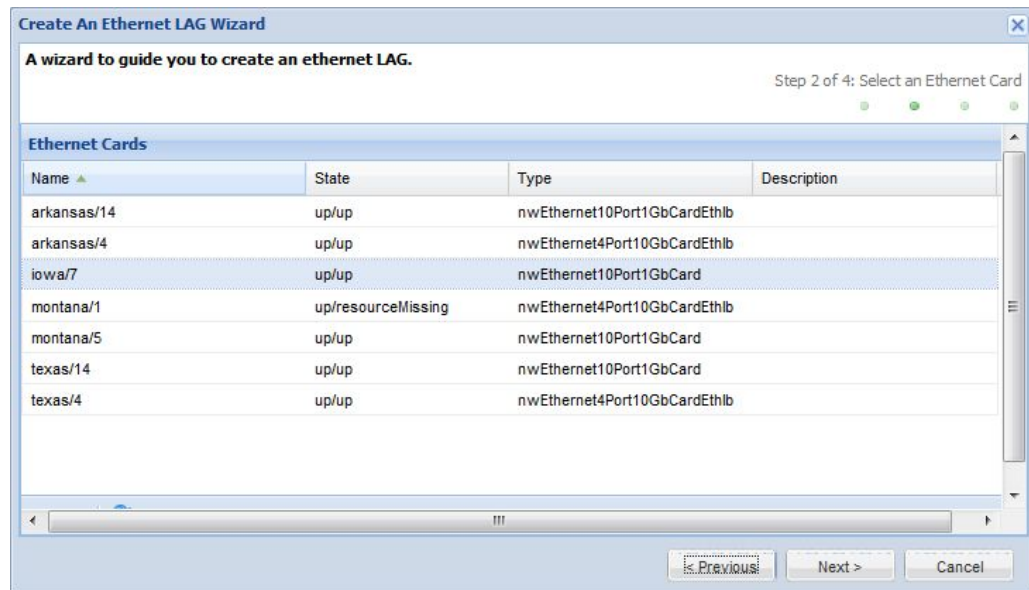
LAG ID (1-5): *

LACP: Check to enable LACP

Description:

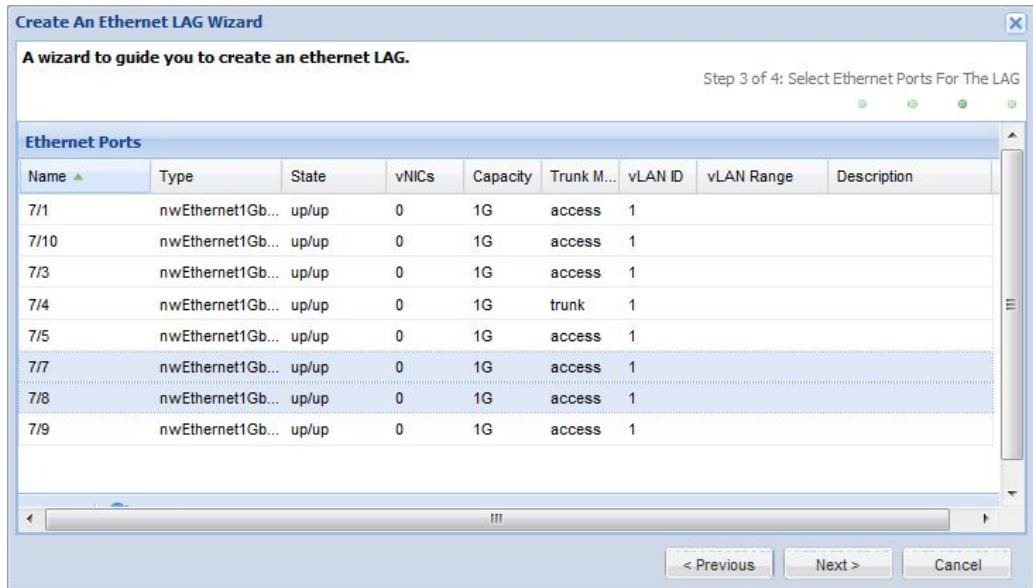
< Previous Next > Cancel

3. In the LAG ID field, enter a number from 1 to 5 to create the LAG Profile into which you will put individual LAG ports.
4. As an option, check the LACP checkbox to enable LACP. By default, LAG ports are not managed by LACP. Enable LACP if the peer network switch for the LAG ports is using LACP.
5. As an option, in the Description field, enter a description string for the link aggregation group.
6. Click Next to display the Select Ethernet Card dialog.



7. Select the Ethernet module on which you will be creating the LAG.

8. **Click Next to display the Select LAG Port dialog.**



9. **Select one or more Ethernet ports that will be used to create the LAG.**

Note - Standard keyboard shortcuts can be used to select multiple ports in the table. For example, Shift + click and Ctrl + click allow you to select all ports or individual ports in the table.

10. **When the ports you want to add to the LAG are selected, click Next to display the Summary for the LAG you are creating.**

11. **Review the information on the Summary.**

If any of the information is incorrect, click Previous to page backward through the wizard until you find the page where you can make the appropriate corrections.

12. **When the Summary contains the correct information, click Finish to create the LAG in Oracle Fabric Manager.**

13. **Check the LAG Summary table to verify that the LAG was created.**

At this point, the LAG is available for use in other wizards and can be assigned to other Oracle Fabric Manager entities—for example, Network Clouds.

Displaying LAG Properties

Detailed information about each LAG is displayed in the LAG Details frame. When you click to select a single LAG in the LAG Summary table, details about that LAG are displayed in the frame below.

- [“LAG Summary” on page 285](#)
- [“LAG Details” on page 285](#)
- [“LAG Ethernet Properties” on page 286](#)
- [“Port Properties in the LAG” on page 287](#)
- [“Add More Ports to an Existing LAG” on page 288](#)
- [“Delete Ports From an Existing LAG” on page 289](#)

LAG Summary

In the LAG Summary, you will notice that the naming convention uses a dot instead of a slash to separate the slot and port notation. For example, the LAG name “arkansas/14.4” indicates that LAG 4 exists on slot 14 of the Oracle Fabric Device “arkansas”. This notation is standard for Oracle LAGs and is also used in Oracle’s XgOS CLI.

Ethernet LAGs Summary				
Name	State	Ports	vNICs	Description ↕
iowa/7.1	up/up	2	8	
iowa/7.2	up/up	6	6	

2 items

LAG Details

Details for a specific LAG are displayed in the LAG Details frame, which exists below the LAG Summary. Through the LAG Details frame, you can see additional information for a specific LAG, and also edit some LAG parameters.

The screenshot displays the Oracle Fabric Manager interface. At the top, there is a section titled "Ethernet LAGs Summary" which contains a table with the following data:

Name	State	Ports	vNICs	Description
delaware/8.1	up/down	1,10	0	
delaware/9.4	up/down	2,3	0	
oregon/8.1	up/down	1,10	0	

Below the summary table, there is a section titled "Ethernet LAG" with three tabs: "General", "Ethernet", and "Ports". The "General" tab is currently selected and displays the following details for the LAG named "oregon/8.1":

- Name:** oregon/8.1
- State (Admin/Oper):** up/down
- Description:**
- LAG Ports:** 1,10

An "Edit" button is located at the bottom of the "General" tab.

The LAG Details frame contains the following tabs, which allow basic management and editing for a configured LAG:

- General
- Ethernet
- Ports

LAG Ethernet Properties

Oracle Fabric Manager supports Ethernet properties for the traffic that is carried on the LAG. Properties such as MTU size, VLAN ID, and others can be displayed or edited through the Ethernet tab on the LAG Details frame.

Ethernet LAGs Summary

+

Name ^	State	Ports	vNICs	Description
delaware/8.1	up/down	1,10	0	
delaware/9.4	up/down	2,3	0	
oregon/8.1	up/down	1,10	0	

3 items

Ethernet LAG

General **Ethernet** Ports

Admin Rate:	autoNegotiate	MTU:	1500 bytes
Port Mode:	access	Access VLAN ID:	1
Tag Native:	false	Flow Control:	false
IGMP Snooping:	true	LACP:	false

Property...	Means...
Admin Rate	The admin rate at which the network traffic is supported. This is typically 10 Gbps or lower, or auto negotiated to the highest value supported on the LAG.
MTU	The maximum transmission unit, which is the largest size data packet supported without fragmentation.
Port Mode	The VLAN mode for the port — either trunk or access.
VLAN ID	The VLAN ID supported on the port in the LAG.
Tag Native	Whether or not traffic on the port is retagged or retains its VLAN ID when the traffic is from the native VLAN.
Flow Control	Whether flow control is enabled or not on the port.
IGMP Snooping	Whether IGMP Snooping is enabled or not on the port.
LACP	Whether LACP is enabled or not on the port.

Port Properties in the LAG

Oracle Fabric Manager supports Ethernet properties for the traffic that is carried on the LAG. Properties such as MTU size, VLAN ID, and others can be displayed or edited through the Ethernet tab on the LAG Details frame.

Add More Ports to an Existing LAG

The screenshot shows two panels from the Oracle Fabric Manager. The top panel, titled "Ethernet LAGs Summary", contains a table with the following data:

Name	State	Ports	vNICs	Description
delaware/8.1	up/down	1,10	0	
delaware/9.4	up/down	2,3	0	
oregon/8.1	up/down	1,10	0	

The bottom panel, titled "Ethernet LAG", has tabs for "General", "Ethernet", and "Ports". The "Ports" tab is active, showing a table with the following data:

Port Number	Type	State	Capacity	Description
1	nwEthernet1GbPort	up/down	1G	
10	nwEthernet1GbPort	up/down	1G	

Property...	Means...
Port Number	The number of individual ports in the LAG
Type	The type of port in the LAG. Typically, ports in the LAG are 1 Gbps Ethernet ports as indicated by the type string nwEthernet1GbPort.
State	The admin and operational state of the port in the LAG. When the port is up/up it is online and able to pass traffic in the LAG.
Capacity	The maximum rate of traffic that each port in the LAG can support.
Description	An optional field that contains a description for each port in the LAG.

▼ Add More Ports to an Existing LAG

When a LAG is configured, you can add more ports to it as long as other ports are available that are not already part of another LAG. If ports are assigned to another LAG, they cannot be directly assigned into a different LAG. If you need to take a port from another LAG, you must delete it from its LAG before assigning it to another LAG.

To add more ports to an existing LAG, follow this procedure:

1. **On the navigation panel, select Network Cloud Manager → Link Aggregation Groups to display the Ethernet LAG Summary.**
2. **On the Ethernet LAG Summary, select the LAG to which you want to add one or more ports.**

This step populates the Ethernet LAG Details frame.

3. Click the Ports tab to display the current port contents of the LAG.

The screenshot displays two sections of the Oracle Fabric Manager interface. The top section, titled "Ethernet LAGs Summary", shows a table with three rows of LAGs. The bottom section, titled "Ethernet LAG", shows the "Ports" tab selected, displaying a table with two rows of ports.

Name	State	Ports	vNICs	Description
delaware/8.1	up/down	1,10	0	
delaware/9.4	up/down	2,3	0	
oregon/8.1	up/down	1,10	0	

Port Number	Type	State	Capacity	Description
1	nwEthernet1GbPort	up/down	1G	
10	nwEthernet1GbPort	up/down	1G	

4. Click the plus sign (+) to display the Select Ports to Add to LAG dialog.
5. Select the port(s) you want to add to the LAG, and click Submit.

▼ Delete Ports From an Existing LAG

When ports are in a LAG, they can be removed to become individual ports. When you delete one or more ports from a LAG, they are then configurable as an individual network port, or they can be assigned to another LAG. It is possible to remove all ports from a LAG without deleting the LAG. In this case, the LAG exists in Oracle Fabric Manager as an empty LAG, which can then be assigned more ports (See [“Add More Ports to an Existing LAG” on page 288.](#))

To delete one or more ports from an existing LAG, follow this procedure:

1. On the navigation panel, select **Network Cloud Manager** → **Link Aggregation Groups** to display the **Ethernet LAG Summary**.
2. On the **Ethernet LAG Summary**, select the LAG to which you want to add one or more ports. This step populates the **Ethernet LAG Details** frame.

3. Click the Ports tab to display the current port contents of the LAG.

The screenshot displays two sections of the Oracle Fabric Manager interface. The top section, titled "Ethernet LAGs Summary", contains a table with the following data:

Name	State	Ports	vNICs	Description
delaware/8.1	up/down	1,10	0	
delaware/9.4	up/down	2,3	0	
oregon/8.1	up/down	1,10	0	

Below this table is a "3 items" indicator with a refresh icon. The bottom section, titled "Ethernet LAG", has three tabs: "General", "Ethernet", and "Ports". The "Ports" tab is active, showing a table with the following data:

Port Number	Type	State	Capacity	Description
1	nwEthernet1GbPort	up/down	1G	
10	nwEthernet1GbPort	up/down	1G	

Below this table is a "2 items" indicator with a refresh icon.

4. On the Ports tab, select the port(s) that you want to delete from the LAG, and click the Delete a Port button.

Working With Network QoS

This chapter contains the following topics:

- [“Understanding Network QoS” on page 291](#)
- [“Understanding MAC-Based QoS” on page 296](#)

Understanding Network QoS

This section contains the following topics:

- [“Network QoS Profiles” on page 291](#)
- [“Predefined Network QoS Profiles” on page 292](#)
- [“Display the Network QoS Summary” on page 293](#)
- [“Network QoS Details” on page 294](#)
- [“vNICs Using a Network QoS Profile” on page 295](#)

Network QoS Profiles

A Network QoS Profile allows you to place bandwidth usage parameters on a Network Cloud or a vNIC so that specific amounts of traffic are allowed, or specific amounts of throughput are available.

Through Oracle Fabric Manager, network Quality of Service (QoS) ensures that bandwidth is available by using parameters to control how much traffic is allowed on Network Cloud at any given time. QoS parameters are configured within a Network QoS Profile, which is then bound to a Network Cloud or vNIC. For more information, see [“QoS Assignment at Cloud Level and vNIC Level” on page 100](#).

Oracle Fabric Manager supports Network QoS through policers on Network Clouds and their underlying Ethernet links. Policers guarantee bandwidth by controlling traffic through dropping packets that exceed the peak information rate (PIR).

The Network QoS Profile uses the following parameters to guarantee bandwidth:

- Committed Information Rate (CIR), which you specify
- Peak Information Rate (PIR), which is you specify

Oracle provides a number of predefined Network QoS Profiles, complete with commonly used CIR and PIR parameters for typical network usage scenarios.

Predefined Network QoS Profiles

Oracle Fabric Manager contains predefined Network QoS Profiles. Within the predefined QoS Profile, one or more entries exist. Each entry is pre-configured for efficient bandwidth availability and resource usage. You can associate a predefined QoS Profile directly with a Network Cloud.

Multiple predefined QoS Profiles exist, and they have been created for different bandwidth configurations—for example a 1 Gbps bandwidth predefined QoS Profile, a 2 Gbps predefined QoS Profile, and so on.

When you are binding a QoS Profile to a Network Cloud, you can select a predefined QoS Profile from the list of all available QoS Profiles for that particular module.

Predefined QoS Profiles are supported for Network QoS. The following table shows each of the predefined Policer Profiles and the QoS parameters that the Policer Profiles contain.

Default QoS Profile	CIR	PIR
100m_1g	100 mbps	1 gbps
100m_250m	100 mbps	250 mbps
10m_100m	10 mbps	100 mbps
10m_1g	10 mbps	1 gbps
10m_50m	10 mbps	50 mbps
1g_10g	1 gbps	9.9297 gbps
1m_10m	1 mbps	10 mbps
250m_500m	250 mbps	500 mbps
2g_10g	2 gbps	9.9297 gbps
3g_10g	3.00293 gbps	9.9297 gbps
4g_10g	4 gbps	9.9297 gbps
500m_750m	500 mbps	750 mbps
50m_100m	50 mbps	100 mbps
5g_10g	5.00122 gbps	9.9297 gbps
64k_1m	66 kbps	1 mbps
6g_10g	6.00587 gbps	9.9297 gbps
750m_1g	750 mbps	1 gbps
7g_10g	7.00171 gbps	9.9297 gbps
8g_10g	8 gbps	10 gbps
9g_10g	9.00212 gbps	9.9297 gbps

▼ Display the Network QoS Summary

When Network QoS Profiles are configured, you can display them through the Network QoS Summary, which is a table of all Network QoS Profiles configured in Oracle Fabric Manager. Additional, detailed information about individual Network QoS Profiles is available in the Network QoS Profile Details frame below the Network QoS Profile Summary.

The Network QoS Summary contains all configured Network QoS Profiles in Oracle Fabric Manager regardless of whether they are assigned to a Network Cloud.

To display the Network QoS Summary, follow this procedure:

- **Display the Network QoS Summary by selecting Network Cloud Manager → Network QoS.**

Network QoS Policing Profile		MAC-based QoS		
Name	CIR	PIR	Number of vNICs	Description
100m_1g	100M	1G	0	100Mbps CIR, 1Gbps PIR
100m_250m	100M	250M	0	100Mbps CIR, 250Mbps PIR
10g_10g	9.929696G	9.929696G	0	10Gbps CIR, 10Gbps PIR
10m_100m	10M	100M	0	10Mbps CIR, 100Mbps PIR
10m_1g	10M	1G	0	10Mbps CIR, 1Gbps PIR
10m_50m	10M	50M	0	10Mbps CIR, 50Mbps PIR
1g_10g	1G	9.929696G	0	1Gbps CIR, 10Gbps PIR
1m_10m	1M	10M	0	1Mbps CIR, 10Mbps PIR
250m_500m	250M	500M	0	250Mbps CIR, 500Mbps PIR
2g_10g	2G	9.929696G	0	2Gbps CIR, 10Gbps PIR
3g_10g	3.002932G	9.929696G	0	3Gbps CIR, 10Gbps PIR
4g_10g	4G	9.929696G	0	4Gbps CIR, 10Gbps PIR
500m_750m	500M	750M	0	500Mbps CIR, 750Mbps PIR

21 items

Field	Indicates
Name	The name of each configured Network QoS Profile. The name is in the format CIR_PIR, so the “100m_1g” profile sets 100 Mbps of CIR and 1 Gbps of PIR.
CIR	The committed information rate, which is the amount of guaranteed bandwidth for constant traffic.
PIR	The peak information rate, which the amount of peak bandwidth for constant traffic.
Number of vNICs	The total number of vNICs that are associated with each Network QoS Profile.
Description	The description string (if any) that was applied to the Network QoS Profile. If this field is blank, either no description string was specified when the Network QoS Profile was created, or the Network QoS Profile was originally created with a description string, but the Network QoS Profile was later edited and the description string was removed.

Network QoS Details

The Network QoS Profile Details frame is a section of the work panel that is located below the Network QoS Profile Summary. This frame is a list of fields for a selected Network QoS Profile that shows the CIR and PIR policing parameters that control bandwidth usage.

The Network QoS Profile Details frame enables you to display additional, detailed information for a single Network QoS Profile and contains an Edit button, which unlocks editable parts of the details frame so that you can set or change information elements of the details frame.

To use the Network QoS Profile Details frame, you must first select a configured Network QoS Profile from the Network QoS Profile Summary. By selecting a Network QoS Profile from the summary, you provide an element that will be the focus of the Network QoS Profile Details frame. When the Network QoS Profile is selected in the summary, you will see its details displayed in the Details frame.

The following example shows the Network QoS Profile Details frame. Notice that the details frame is contextual, so that it displays detailed information for the item selected in the Network Cloud Summary.

The screenshot shows a management console interface. At the top, there is a tab labeled "Network QoS Policing Profile" with a sub-tab "MAC-based QoS". Below this is a table listing various Network QoS profiles. The table has columns for Name, CIR, PIR, Number of vNICs, and Description. The profiles listed include 100m_1g, 100m_250m, 10g_10g, 10m_100m, 10m_1g, 10m_50m, 1g_10g, 1m_10m, 250m_500m, 2g_10g, 3g_10g, 4g_10g, and 500m_750m. Below the table, there is a section titled "Network QoS: 1g_10g" with a sub-tab "vNICs". This section displays the details for the 1g_10g profile, including Name, Description, CIR(kbps), and PIR(kbps).

Name	CIR	PIR	Number of vNICs	Description
100m_1g	100M	1G	0	100Mbps CIR, 1Gbps PIR
100m_250m	100M	250M	0	100Mbps CIR, 250Mbps PIR
10g_10g	9.929696G	9.929696G	0	10Gbps CIR, 10Gbps PIR
10m_100m	10M	100M	0	10Mbps CIR, 100Mbps PIR
10m_1g	10M	1G	0	10Mbps CIR, 1Gbps PIR
10m_50m	10M	50M	0	10Mbps CIR, 50Mbps PIR
1g_10g	1G	9.929696G	0	1Gbps CIR, 10Gbps PIR
1m_10m	1M	10M	0	1Mbps CIR, 10Mbps PIR
250m_500m	250M	500M	0	250Mbps CIR, 500Mbps PIR
2g_10g	2G	9.929696G	0	2Gbps CIR, 10Gbps PIR
3g_10g	3.002932G	9.929696G	0	3Gbps CIR, 10Gbps PIR
4g_10g	4G	9.929696G	0	4Gbps CIR, 10Gbps PIR
500m_750m	500M	750M	0	500Mbps CIR, 750Mbps PIR

21 items

Network QoS: 1g_10g

General | vNICs

Name: 1g_10g
Description: 1Gbps CIR, 10Gbps PIR
CIR(kbps): 1000000
PIR(kbps): 9929696

vNICs Using a Network QoS Profile

Through the vNICs tab, you can display a table of the vNICs that are currently using a specific Network QoS Profile. To display the vNICs that are using a Network QoS Profile, select a Network QoS Profile in the Network QoS Profile Summary. Then, click the vNICs tab to display the vNICs that are using that profile.

Network QoS: 1g_10g											
vNICs											
Name	Network Cloud	Server Name	Termination	State	IP Address	Netmask	IP Type	MAC Address	HA	QoS	Private
vnic	techpubs1		iowa/7/1	up/resourceU...	0.0.0.0	255.255.255.2...	hostMana...	00:13:97:01:8B:D3	f...	1000000_...	false
vnic	techpubs1		iowa/7/10	up/resourceU...	0.0.0.0	255.255.255.2...	hostMana...	00:13:97:01:8B:D2	f...	1000000_...	false
vnic	techpubs1		iowa/7/1	up/resourceU...	0.0.0.0	255.255.255.2...	hostMana...	00:13:97:01:8B:D1	f...	1000000_...	false
vnic	techpubs1		iowa/7/10	up/resourceU...	0.0.0.0	255.255.255.2...	hostMana...	00:13:97:01:8B:D0	f...	1000000_...	false
vnic	techpubs1		iowa/7/1	up/resourceU...	0.0.0.0	255.255.255.2...	hostMana...	00:13:97:01:8B:CF	f...	1000000_...	false

Field	Indicates
Name	The name of each configured vNIC that is using the Network QoS profile.
Network Cloud	The name of the Network Cloud to which the vNICs are connected.
Server Name	The name of all servers connected to the vNIC.
Termination	The termination point for the vNICs. The termination point is either a port which is displayed in slot/port notation, or a LAG which is displayed in slot.port notation.
State	The administrative and operational state of the vNICs that are using the selected Network QoS profile.
IP Address	The IP address of each vNIC that is using the Network QoS Profile.
Netmask	The network mask for each vNIC that is using the Network QoS Profile.
IP Type	The method by which the IP address for the vNIC was assigned. Valid values are: <ul style="list-style-type: none"> ■ Static, for a statically assigned IP address. ■ DHCP, for an address that was assigned through DHCP. ■ Host managed, for an address that was assigned by the host instead of through the Oracle Fabric Device.
MAC Address	The MAC address of each vNIC that is using the Network QoS Profile. MAC addresses can either be statically assigned by you, or automatically assigned from the Oracle Fabric Device's MAC address pool.
HA	The state of high availability for each vNIC that is using the Network QoS Profile: <ul style="list-style-type: none"> ■ If this field displays false, then the vNIC is not part of an HA pair. ■ If this field displays true, then the vNIC is either the primary or secondary vNIC in an HA vNIC pair.
QoS	The bandwidth usage parameters for each vNIC. The name is in the format CIR_PIR in megabits per second.
Private	Whether the individual vNICs shown are Private vNICs: <ul style="list-style-type: none"> ■ If the field displays true, the vNIC is a Private vNIC. Private vNICs are used in vNIC-to-vNIC switching to ensure enhanced security and isolation from standard "public" vNICs. ■ If the field displays false, the vNIC is not a Private vNICs. Instead, it is a standard public vNIC.

Understanding MAC-Based QoS

This section contains the following topics:

- [“MAC-Based QoS” on page 297](#)

- [“MAC-Based QoS Summary” on page 297](#)
- [“MAC-Based QoS Profile Configuration” on page 298](#)
- [“Create a MAC-Based QoS Profile” on page 299](#)
- [“MAC-Based QoS Profiles Deletion” on page 301](#)

MAC-Based QoS

In addition to network QoS which governs the amount of bandwidth used on vNICs, Oracle Fabric Manager also supports MAC-Based QoS. MAC-Based QoS is a method of controlling which devices on a vNIC can use specific amounts of bandwidth. Through MAC-Based QoS, Oracle Fabric Manager supports assigning a usage condition to traffic sent or received by a specific network device, which is identified by its MAC address. The usage condition are enforced by QoS application flows.

When you configure MAC-Based QoS, you set conditions for a specific device address on a vNIC. When MAC-Based QoS is configured on a vNIC, the actual match against traffic on the vNIC occurs at the I/O Card level, so traffic will be controlled at the chip that controls the vNIC's termination port.

A MAC-Based QoS Profile is composed of the following information elements:

- AMAC Address, which is the MAC address of a shared vNIC—for example, a vNIC deployed on a VM. If a vNIC is deployed on a physical server, there is no need to configure a MAC-Based QoS Profile because the vNIC has only one MAC address. As a result, MAC-Based QoS Profiles are not configurable in a non-VM environment.
- Network QoS Profile, which controls the amount of bandwidth and traffic that can be used on the shared vNIC.
- Direction, which determines a traffic flow for the Network QoS Profile.
- An optional description.

MAC-Based QoS Profiles are useful on a shared vNIC, which occurs when multiple devices are supported on the same vNIC, and each device requires a different QoS flow to or from the device (ingress or egress) or both.

MAC-Based QoS Summary

Oracle Fabric Manager supports MAC-Based QoS as an option to the Network QoS feature. MAC-Based QoS is available as a tab through the Network QoS Summary. This option displays the MAC-Based QoS Summary, which is a table of all the configured MAC-Based QoS profiles.

The screenshot shows a web-based configuration interface for Network QoS Policing Profiles. The 'MAC-based QoS' tab is active. A table lists two configured rules:

Name	Condition	QoS Profile	Description
rule-00:11:22:33:44:55-408	00:11:22:33:44:55 dest	10m_50m	
rule-00:11:22:33:44:55-440	00:11:22:33:44:55 src	10m_100m	

At the bottom of the interface, it indicates '2 items'.

Field	Indicates
Name	The name of the MAC-Based QoS Profile, which is generated from the specified MAC Address.
Condition	The specific MAC-Based QoS Profile conditions created. These conditions consist of the MAC Address specified plus either source (src) or destination (dest) which refers to the direction on which the QoS flow will be applied to traffic.
QoS Profile	The name of each configured Network QoS Profile associated with each MAC-Based QoS Profile. The name is in the format CIR_PIR, so the "100m_1g" profile sets 100 Mbps of CIR and 1 Gbps of PIR.
Description	The description string (if any) that was applied to the MAC-Based QoS Profile. If this field is blank, either no description string was specified when the MAC-Based QoS Profile was created, or the MAC-Based QoS Profile was originally created with a description string which was later edited and the description string was removed.

MAC-Based QoS Profile Configuration

Oracle Fabric Manager supports MAC-Based QoS through the Network QoS Profile icon in the Navigation Panel. The Network QoS Profile object leads to the Network QoS Profile Summary which has a separate tab for MAC-Based QoS. The MAC-Based QoS tab displays the summary list of all configured MAC-Based QoS profiles. Through the MAC-Based QoS tab, you can add, configure, and delete specific MAC-Based QoS Profiles.

As part of configuring a MAC-Based QoS Profile, you will associate the shared vNIC's MAC address with a Network QoS Policing Profile. Before beginning the MAC-Based QoS Profile

configuration procedure, make sure that the Network QoS Profile exists so that it is a selectable option.

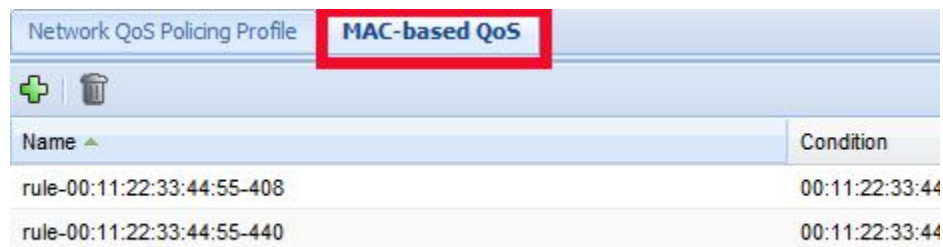
▼ Create a MAC-Based QoS Profile

When you create a MAC-Based QoS Profile, you are creating a profile that applies a predefined Network QoS Profile to a specific MAC address. You are also determining the direction on which the Network QoS flow is applied, either ingress, egress, or both. Traffic on a shared vNIC that originates from, or is destined to, the MAC address will be controlled based on the Network QoS profile.

You can create an MAC-Based QoS Profile through the plus sign (+) on the MAC-Based QoS Profile Summary.

To create a MAC-Based QoS Profile, follow this procedure:

1. On the navigation panel, click **Network Cloud Manager – Network QoS** to display the **Network QoS Profile Summary** page. See [“Display the Network QoS Summary” on page 293](#) for an example.
2. On the **Network QoS Summary**, click the **MAC-Based QoS** tab to display the **MAC-Based QoS Summary**.

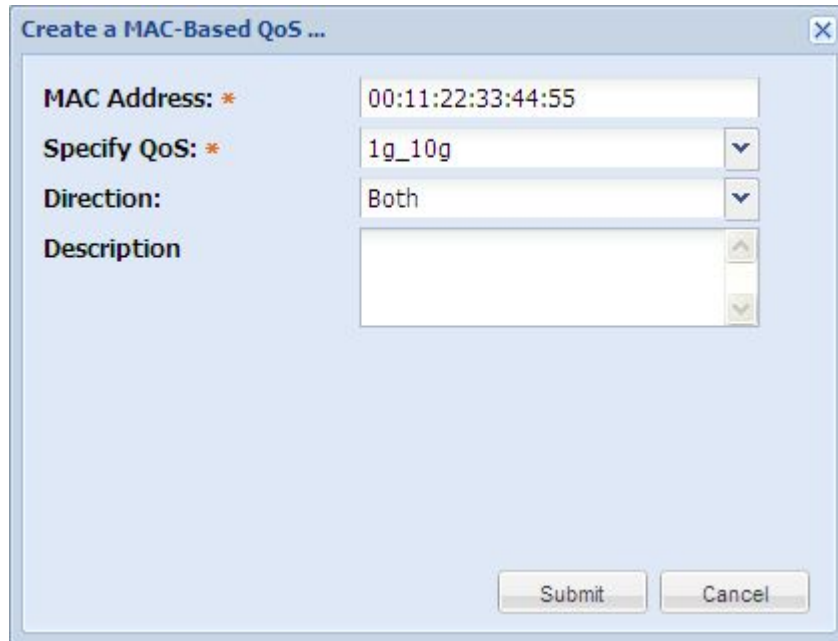


The screenshot shows the 'Network QoS Policing Profile' interface with the 'MAC-based QoS' tab highlighted in a red box. Below the tab are icons for adding (+) and deleting (trash) profiles. A table lists two profiles with their names and conditions.

Name ▲	Condition
rule-00:11:22:33:44:55-408	00:11:22:33:44
rule-00:11:22:33:44:55-440	00:11:22:33:44

Notice that you can create or delete MAC-Based QoS Profiles by clicking the plus sign (+) or garbage can, respectively. Also, notice that the MAC-Based QoS feature has a Summary, but no details frame for individual MAC-Based QoS Profiles.

3. Click the plus sign (+) to display the Create a MAC-Based QoS dialog.



The screenshot shows a dialog box titled "Create a MAC-Based QoS ...". It contains the following fields and controls:

- MAC Address: ***: A text input field containing the value "00:11:22:33:44:55".
- Specify QoS: ***: A dropdown menu with "1g_10g" selected.
- Direction:**: A dropdown menu with "Both" selected.
- Description**: A text area for entering a description.
- Submit** and **Cancel**: Buttons at the bottom right of the dialog.

4. In the MAC Address field enter the MAC Address for the shared vNIC.
5. From the Specify QoS dropdown menu, select the QoS Profile that you want to assign to the vNIC for the specific application.
6. From the Direction dropdown menu, select a direction (or both) to which the MAC-Based QoS Profile will be applied.
7. As an option, in the Description field, enter a description for the MAC-Based QoS Profile that you are creating.
8. Click Submit to create the MAC-Based QoS Profile.
9. As needed, repeat this procedure as needed to create any additional MAC-Based QoS Profiles for individual MAC addresses on the shared vNIC.

MAC-Based QoS Profiles Deletion

When a MAC-Based QoS Profile is configured, it is displayed in the MAC-Based QoS Summary page. Any configured MAC-Based QoS Profile can be deleted through the garbage can icon on the MAC-Based QoS Profile Summary.

If multiple MAC-Based QoS Profiles exist, you can delete multiple profiles by selecting multiple entries then deleting them. Selecting multiple instances in the MAC-Based QoS Profile Summary is supported through standard keyboard conventions—for example, CTRL + click.

When you attempt to delete a MAC-Based QoS Profile, Oracle Fabric Manager prompts you with a pop-up dialog to requires you to confirm the deletion. As a result, if you have selected the wrong MAC-Based QoS Profile, you have a chance to correct the error before committing the deletion.

▼ Delete a MAC-Based QoS Profile

When you delete a MAC-Based QoS Profile, the deletion is immediate after the confirmation dialog is answered. When deleted, the MAC-Based QoS Profile that matches against a specific traffic flow is no longer applied, so the traffic flow is not affected.

To delete an MAC-Based QoS Profile, follow this procedure:

1. **From the navigation panel, select Network Cloud Manager → Network QoS to display the Network QoS Profile Summary.**
2. **Click the MAC-Based QoS tab to display the MAC-Based QoS Summary.**
3. **Click the MAC-Based QoS Profile that you want to delete.**
4. **Click the garbage can icon to delete the selected MAC-Based QoS Profile.**
5. **Click OK to confirm the deletion.**

Working With Physical Servers

This chapter documents the following topics:

- [“Physical Servers” on page 303](#)
- [“Display Physical Servers” on page 304](#)
- [“Display the Physical Servers Details” on page 307](#)
- [“Managing vNICs on a Physical Server” on page 307](#)
- [“Managing vHBAs on a Physical Server” on page 312](#)
- [“Scan for New Servers” on page 317](#)
- [“Apply an I/O Template to Servers” on page 318](#)
- [“Save a Server Configuration as an I/O Template” on page 319](#)
- [“Connect an I/O Profile to a Server” on page 320](#)
- [“Disconnect an I/O Profile From a Server” on page 321](#)
- [“Migrate Resources to Different Servers” on page 321](#)
- [“Create a Server Group From Selected Servers” on page 323](#)
- [“Remove Offline or Disconnected Servers” on page 323](#)

Physical Servers

Physical servers are the host devices on which your applications run. They connect to the Oracle Fabric Device through high-speed interconnect HCAs that are installed in each server. The Fabric Device and Oracle Fabric Manager support the commonly used server types including Linux, Oracle Solaris, Oracle VM, Oracle Linux, Red Hat Enterprise Linux, Windows, and VMware. For information about specific OSes and hypervisors supported, contact Oracle Support or your account team.

Oracle Fabric Manager discovers servers that are connected through the Fabric Device and have Oracle host drivers installed. When Oracle Fabric Manager discovers physical servers, they are displayed in the Physical Servers Summary. Oracle Fabric Manager supports displaying summary and detailed information for physical servers.

▼ Display Physical Servers

To display the Physical Server Summary, follow this procedure:

- **Select Server Resource Manager → Physical Servers.**

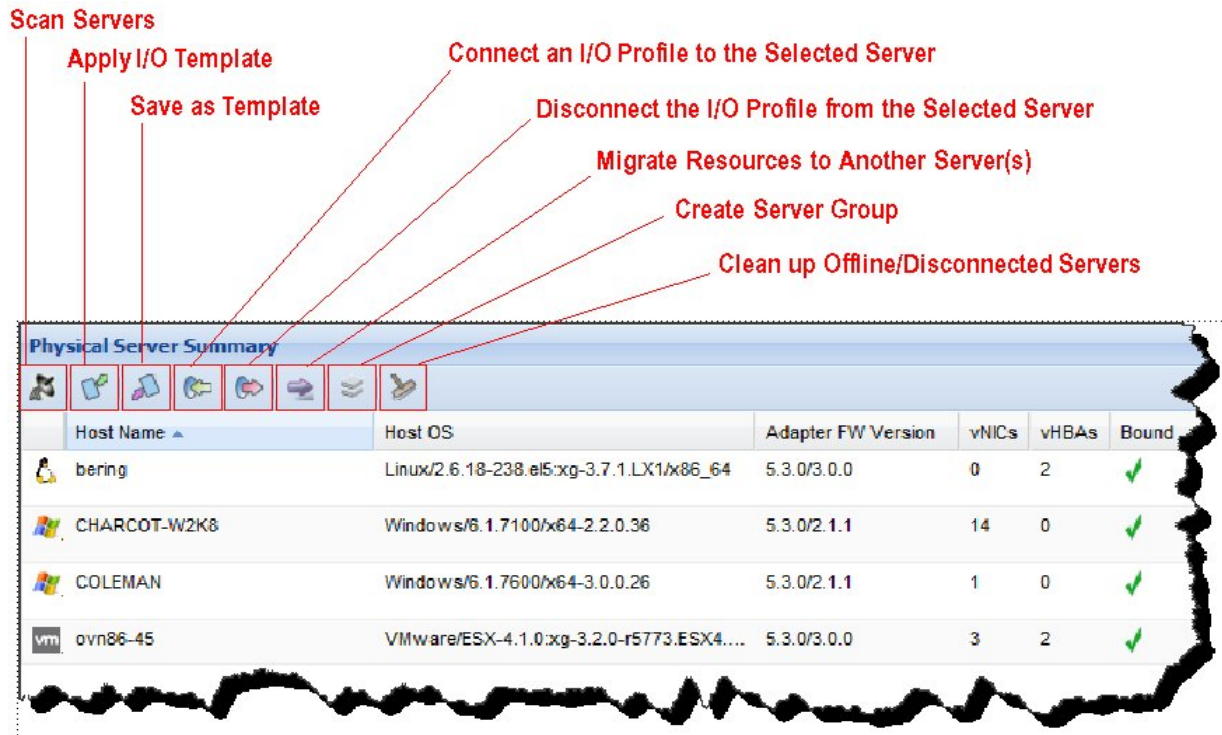
The Physical Server Summary is displayed.

Host Name	Host OS	vNICs	vHBAs	Bound	Busy	State	I/O Profile Name	Fabric Device Ports	Groups
bering	Linux/2.6.18-238.el5.xg-3.7.1.LX1/x8...	2	2	✓		up	LAMirrorSite1_607	delaware:ServerPort22 oregon:ServerPort15	
ovn87-36.us.oracle.com	Linux/2.6.32-220.el6.x86_64/x86_64	0	0					ExtSw-2c90200443198-Port8	
ovn87-34.us.oracle.com	Linux/2.6.32-220.el6.x86_64/x86_64	0	0					ExtSw-2c90200443198-Port10	
ovn87-35.us.oracle.com	Linux/2.6.32-220.el6.x86_64/x86_64	0	0					ExtSw-2c90200443198-Port9	
ovn87-22	Linux/2.6.32-220.el6.x86_64:yg-5.0....	1	0	✓		up	iotemplate	ExtSw-2c90200443198-Port7	HRServerGroup9 HRServerGroup
ovn86-45	VMware/ESX-4.1.0:xg-3.2.0-r5773.E...	0	0			offline			

Field	Indicates
Host Name	The name of each physical server that Oracle Fabric Manager has discovered. If a name has been assigned to the server, that name appears in this field. Otherwise, the GUID of the HCA that connects the server and Fabric Device is displayed. The name listed is a link to the Server Details frame. You can click this link to display additional information about a selected physical server.
Host OS	The Operating System currently in use on the host server.
Adapter FW Version	The version of Oracle Virtual Networking (OVN) host driver or HCA firmware currently in use on the host or HCA.
vNICs	The total number of vNICs that are configured on the physical server.
vHBAs	The total number of vHBAs that are configured on the physical server.
Bound	The state of whether or not the server is bound to an I/O Profile. A checkmark indicates that the server is bound. If this field shows no checkmark, the server is available for binding to an I/O Profile, but is not currently bound.
Busy	Whether Oracle Fabric Manager is busy with an operation related to the server, such as binding or unbinding an I/O Profile.
State	The current administrative state of the physical server. Valid values are: <ul style="list-style-type: none"> ■ Up ■ Down ■ Unbound ■ Initializing ■ Partial
I/O Profile Name	The name of the I/O Profile bound (if any) to each server.

Field	Indicates
Fabric Device Ports	<p>The port string for the InfiniBand port on which the physical server is connected to (the Fabric Interconnect or Oracle SDN Controller). The port string consists of the Fabric Interconnect or Oracle SDN Controller name, and the Server Port number separated by a colon (:). For example, iowa:ServerPort18 indicates that the host server is connected to the Fabric Interconnect named “iowa” through port 18 on iowa’s InfiniBand fabric board.</p> <p>If the host server is connected to multiple Fabric Interconnects or Oracle SDN Controllers, a port string is displayed for each connection, and the individual port strings are arranged in a vertical list (as shown below). As an example, see the entry</p> <pre data-bbox="841 625 1008 646">iowa:ServerPort2</pre> <pre data-bbox="786 674 1062 695">south-carolina:ServerPort21</pre> <p>This entry indicates two individual connections—one connection to the Fabric Interconnect name “iowa” and the other connection to the Fabric Interconnect named “south-carolina.”</p> <p>If the host server is connected through an intervening InfiniBand switch, the switch’s GUID is listed in this field. For example, the value ExtSw-2c902004126e8-Port20 indicates a connection to an intervening IB switch between the host server and the Fabric Interconnect.</p>
Groups	The name of the Server Group(s) (if any) to which the selected server belongs.

The Physical Server Summary also contains various controls for creating and managing listed servers.



Through the Physical Server Summary, you can perform the following tasks:

- “Scan for New Servers” on page 317
- “Apply an I/O Template to Servers” on page 318
- “Save a Server Configuration as an I/O Template” on page 319
- “Connect an I/O Profile to a Server” on page 320
- “Disconnect an I/O Profile From a Server” on page 321
- “Migrate Resources to Different Servers” on page 321
- “Create a Server Group From Selected Servers” on page 323
- “Remove Offline or Disconnected Servers” on page 323

▼ Display the Physical Servers Details

For each physical server in the Physical Servers table, you can display detailed information. Detailed server information is available through the Physical Server Details frame. Server details are available for individual servers by selecting a server in the Physical Server Summary, which provides the focus for the details frame.

The Physical Server Details frame contains tabs that organize information about the server into logical groups. The following tabs are supported:

- General
- vNICs
- Server Groups
- vHBAs

Note - Fabric Devices can contain vHBAs, but Oracle SDN Controllers do not. As a result, the vHBA tab is not displayed for Oracle Fabric Manager systems that are managing Oracle SDN Controllers.

To display the Physical Server Details frame, follow this procedure:

1. **Select Server Resource Manager → Physical Servers. See “Physical Servers” on page 303.**
2. **Select a server in the Physical Server Summary to display the server in the details frame.**

Server: bering	
General vNICs vHBAs Server Groups	
Name:	bering
Host OS:	Linux/2.6.18-238.el5.xg-3.7.1.LX1/x86_64
I/O Profile Name:	LAMirrorSite1_225
State:	up
Fabric Devices Ports:	delaware:ServerPort22,oregon:ServerPort15
Adapter FW Version:	5.3.0/3.0.0

Managing vNICs on a Physical Server

- [“vNICs on a Physical Server” on page 308](#)

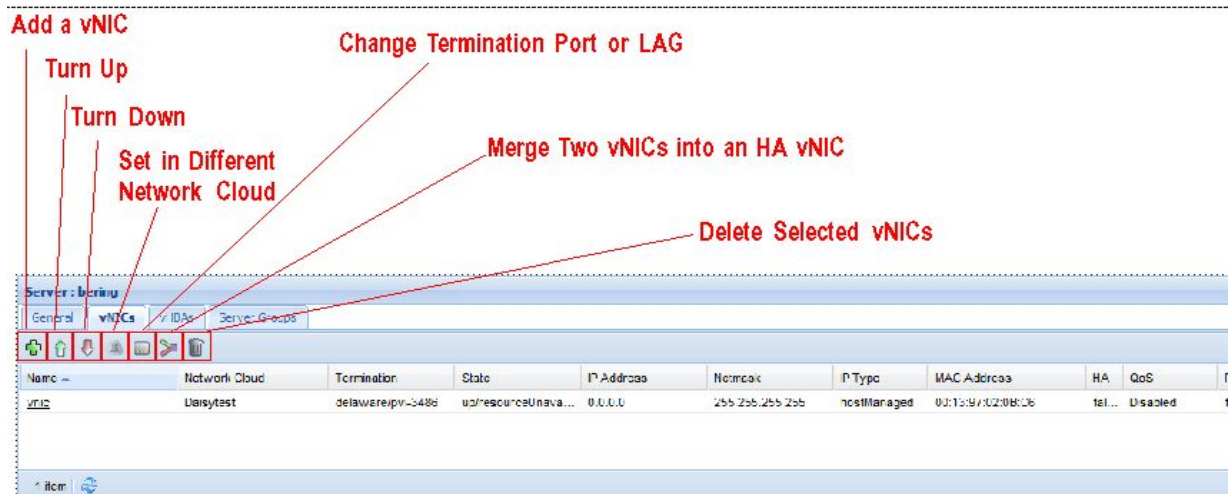
- “Display Details About a vNIC” on page 309
- “Edit vNIC General Properties” on page 309
- “Display vNIC Ethernet Properties” on page 310
- “Edit vNIC Ethernet Properties” on page 311
- “Configure VLAN Ranges for vNICs on a Server” on page 311

vNICs on a Physical Server

When a physical server is bound to an I/O Profile, it has virtual connectivity through vNICs, vHBAs, or both. When a server has vNICs deployed on it, they are listed on the vNICs tab for the selected server. The vNICs tab provides the following functionality:

- It displays information about each vNIC deployed on the server.
- It provides ways to manage the vNICs.
- It lists vNICs by name, and the name is a link to additional details about each vNIC's individual properties.

The vNICs tab offers various controls that allow you to affect the state of vNICs on the selected server.



The Physical Server Details frame has a toolbar that allows you to control the following for one or more vNICs on the selected server:

- Add a vNIC
- Turn Up the vNIC to set it to up/up state

- Turn Down the vNIC to set its state to up/down
- Set in Different Network Cloud, to associate a vNIC on the server with a different Network or PVI Cloud.
- Change Termination Port or LAG to assign the vNIC to a different Ethernet port, or different link aggregation group.
- Merge Two vNICs into an HA vNIC, which takes two selected vNICs and creates one HA pair out of them. To merge the vNICs, they must be part of the same Network Cloud.
- Delete selected vNICs, which removes any number of selected vNICs from the server and deletes them from their respective Network Clouds.

▼ Display Details About a vNIC

When you click a vNIC name, a separate details frame is displayed containing tabs for the vNIC. It is important to remember that this vNIC is in the context of the physical server you selected. The vNIC information is for only the selected vNIC on the selected server. Other vNICs on the server (or on other servers in the network) might or might not have the same vNIC properties.

General information about the vNIC properties is available through the General tab of the vNIC Details for the selected server.

To display a vNIC's general properties, follow this procedure:

1. **Select Server Resource Manager → Physical Servers to display the Physical Server Summary.**
2. **Select a server in the Physical Server Summary to display the server in the details frame.**
3. **On the Physical Server details frame, click the vNICs tab.**
4. **On the vNICs tab, click a vNIC name to display the vNIC Details frame for that vNIC.**
5. **On the vNIC details frame, click the General tab.**

The General tab contains an Edit button that allows you to unlock editable fields to set or change existing vNIC properties.

▼ Edit vNIC General Properties

To edit vNIC General properties, follow this procedure:

1. Click the Edit button.

The screenshot shows a configuration window for a vNIC named 'vnic'. The 'General' tab is active, displaying various properties. The 'Name' is 'vnic', 'State (Admin/Oper)' is 'up/resourceUnavailable', 'Port State' is 'Daisytest', and 'Network Cloud' is 'Daisytest'. The 'iSCSI Boot' checkbox is unchecked. On the right side, 'Description' is empty, 'Termination' is 'pvi-3486', 'QoS Configuration' is 'none', 'PXE Boot' is unchecked, and 'HA' is 'false'. At the bottom, there are 'Submit' and 'Cancel' buttons.

2. Edit the fields as needed, then click Submit to activate the changes.
3. Check the General tab to verify that the correct properties are configured.

Note - While you are editing the vNIC's general properties, you can also set the vNIC up or down by clicking the up or down arrow, respectively.

▼ Display vNIC Ethernet Properties

A vNIC is a software construction that virtualizes a physical NIC, and provides the same functionality as a physical NIC. As a result, a vNIC has the same Ethernet properties as a standard NIC.

vNIC Ethernet properties are displayed in the Ethernet Properties tab on the vNIC Details frame.

To display vNIC Ethernet properties, follow this procedure:

1. Select **Server Resource Manager** → **Physical Servers** to display the **Physical Server Summary**.
2. Select a server in the **Physical Server Summary** to display the server in the details frame.
3. On the **Physical Server** details frame, click the **vNICs** tab.
4. On the **vNICs** tab, click a vNIC name to display the vNIC Details frame for that vNIC.
5. On the **Physical Server** details frame, click the **Ethernet Properties** tab.

The Ethernet Properties tab also contains an Edit button that allows you to unlock editable field on the tab to set or change the existing Ethernet properties for the vNIC.

▼ Edit vNIC Ethernet Properties

To edit vNIC Ethernet Properties, follow this procedure:

1. **On the Ethernet Properties tab, click the Edit button to unlock the editable fields.**

2. **Edit the fields as needed, then click Submit to make the changes.**
3. **Check the Ethernet Properties tab to verify that the correct properties are configured.**

▼ Configure VLAN Ranges for vNICs on a Server

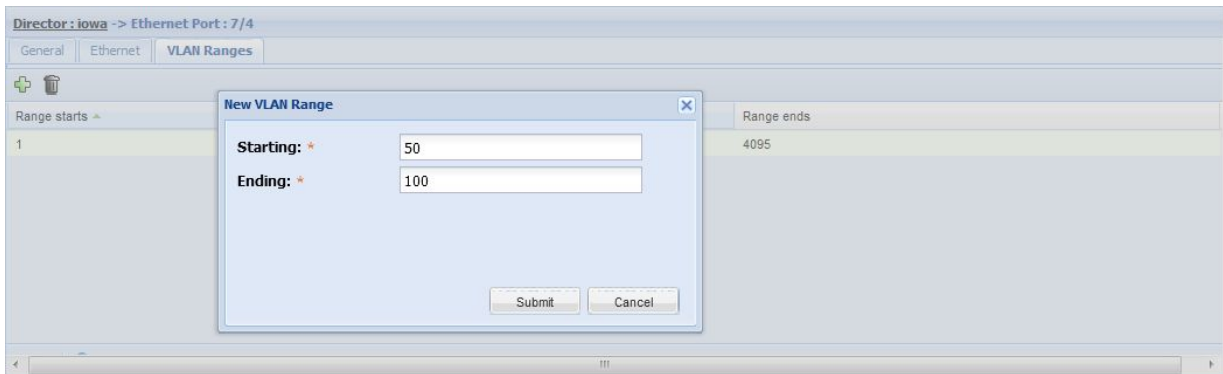
By default, the allowable VLAN range is from 1 to 4095. However, you can set a custom range of VLANs for a vNIC so that only the VLAN-tagged packets within the specified range are allowed on a vNIC. Traffic that has a VLAN tag not in the specified range is blocked from ever being transmitted or received on one or more specified vNICs.

You can set or change the allowed VLAN range for a vNIC through the Allowed VLANs tab at the port level of the Ethernet Card Details Frame.

To set the allowed VLAN range for a specific Fabric Device, follow this procedure:

1. **Select Server Resource Manager → Physical Servers to display the Physical Server Summary page.**
2. **Select the server for which you want to set the Allowed VLAN range by clicking the server in the summary. This step populates the details frame with information for that server.**

3. On the details frame, click the vNICs tab to display each vNIC configured on the selected server.
4. Double click the vNIC on which you want to set the Allowed VLAN range. This step displays the vNIC details for that vNIC.
5. Click the VLAN Ranges tab.
6. Click the plus sign (+) to display the New VLAN Range dialog.



7. In the Starting field, enter the first VLAN ID that you want to be available.
8. In the Ending field, enter the last VLAN ID you want to be available.
9. When the Allowed VLAN Range is configured, click Submit.

Managing vHBAs on a Physical Server

- [“vHBAs on a Physical Server” on page 313](#)
- [“Display Details About a vHBA” on page 314](#)
- [“Display vHBA Fibre Channel Properties” on page 315](#)
- [“Edit vHBA Fibre Channel Properties” on page 316](#)
- [“Display vHBA Targets” on page 316](#)

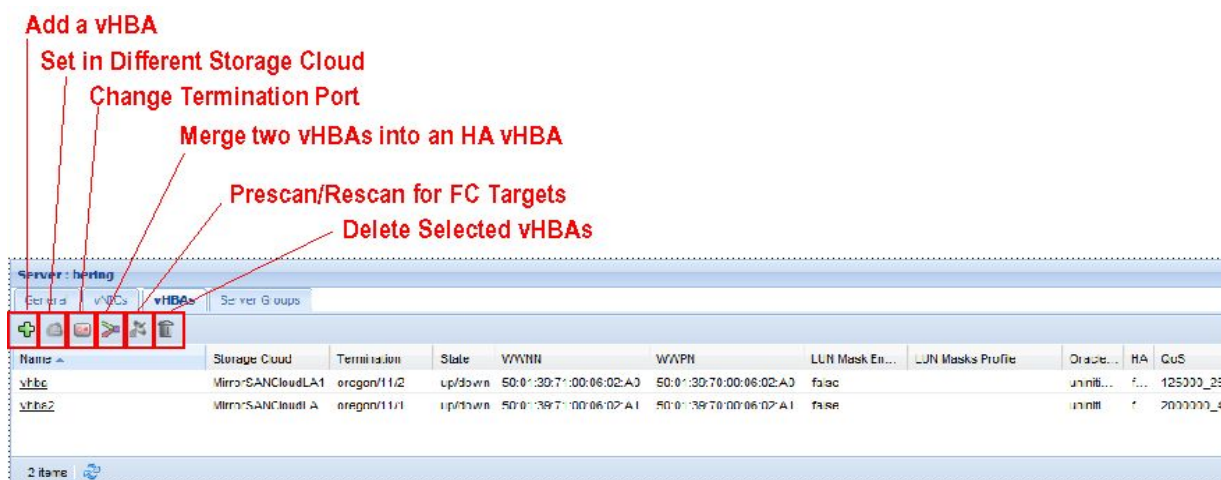
vHBAs on a Physical Server

When a physical server is bound to an I/O Profile, it has virtual connectivity through vNICs, vHBAs, or both. When a server has vHBAs deployed on it, they are listed on the vHBAs tab for the selected server. The vHBAs tab provides the following functionality:

- It displays information about each vHBA deployed on the server.
- It provides ways to manage the vHBAs.
- It lists vHBAs by name, and the name is a link to additional details about each vHBA's individual properties.

Note - Oracle SDN Controllers allow software-defined networking through vNICs connected to a virtual switch. No storage is supported. As a result, the vHBAs tab is not displayed for Oracle Fabric Manager systems that are managing Oracle SDN Controllers. However, the vHBAs tab is present for any Fabric Interconnects that are managed by Oracle Fabric Manager.

The vHBAs tab also contains various controls that allow you to affect the state of vHBAs on the selected server.



The Physical Server Details frame has a toolbar that allows you to control the following for one or more vHBAs on the selected server:

- Add a vHBA
- Set in Different Storage Cloud, to associate a vHBA on the server with a different Storage Cloud.

- Change Termination Port, to assign the vHBA to a different Fibre Channel port.
- Merge Two vHBAs into an HA vHBA, which takes two selected vHBAs and creates one HA pair out of them. To merge the vHBAs, they must be part of the same Storage Cloud.
- Prescan/Rescan for Fibre Channel targets on the vHBA.
- Delete selected vHBAs, which removes any number of selected vHBAs from the server and deletes them from their respective Storage Clouds.

▼ Display Details About a vHBA

When you click a vHBA name, a separate details frame is displayed containing tabs for the vHBA. It is important to remember that this vHBA is in the context of the physical server you selected. The vHBA information is for only the selected vHBA on the selected server. Other vHBAs on the server (or on other servers in the network) might or might not have the same vHBA properties.

General information about the vHBA properties is available through the General tab of the vHBA Details for the selected server.

To display a vHBA's general properties, follow this procedure:

1. **Select Server Resource Manager → Physical Servers to display the Physical Server Summary.**
2. **Select a server in the Physical Server Summary to display the server in the details frame.**
3. **On the Physical Server details frame, click the vHBAs tab.**
4. **On the vHBAs tab, click a vHBA name to display the vHBA Details frame for that vHBA.**
5. **On the vHBA details frame, click the General tab.**

The General tab also has an Edit button that allows you to unlock the editable fields on the tab.

To edit the vHBA's general properties, follow this procedure:

6. On the **General** tab, click the **Edit** button to unlock the editable fields.

7. Edit the fields as needed, then click **Submit** to make the changes.
8. Check the **General** tab to verify that the correct properties are configured.

▼ Display vHBA Fibre Channel Properties

A vHBA is a software construction that virtualizes a physical HBA, and provides the same functionality as a physical HBA. As a result, a vHBA has the same Fibre Channel properties as a standard HBA.

The vHBA Fibre Channel properties are displayed in the Fibre Channel Properties tab on the vHBA Details frame.

To display a vHBA's Fibre Channel properties, follow this procedure:

1. Select **Server Resource Manager – Physical Servers** to display the **Physical Server Summary**.
2. Select a server in the **Physical Server Summary** to display the server in the **details** frame.
3. On the **Physical Server details** frame, click the **vHBAs** tab.
4. On the **vHBAs** tab, click a vHBA name to display the **vHBA Details** frame for that vHBA.
5. On the **vHBA details** frame, click the **Fibre Channel Properties** tab.

The Fibre Channel Properties tab has an **Edit** button allows you to unlock the editable fields on the tab to set or change existing properties.

▼ Edit vHBA Fibre Channel Properties

To edit vHBA Fibre Channel Properties, follow this procedure:

1. **On the Fibre Channel Properties tab, click the Edit button to unlock the editable fields.**
2. **Edit the fields as needed, then click Submit to make the changes.**
3. **Check the Fibre Channel Properties tab to verify that the correct properties are configured.**

▼ Display vHBA Targets

A vHBA connects host servers to SAN resources (targets) such as arrays, JBODs, and so on. The vHBA can connect either directly to storage through the Fabric Device, or indirectly through an intervening Fibre Channel switch between the Fabric Device and the storage.

When an I/O Profile is bound to a host server, the vHBA that connects to the storage is deployed on the server. That vHBA provides the connection for read and write data between the host server and the storage.

Oracle Fabric Manager tracks the storage connected to a host server through the Targets tab on the vHBA Details frame.

To display the storage targets available to a server, follow this procedure:

1. **Select Server Resource Manager → Physical Servers to display the Physical Server Summary.**
2. **Select a server in the Physical Server Summary to display the server in the details frame.**
3. **On the Physical Server details frame, click the vHBAs tab.**
4. **On the vHBAs tab, click a vHBA name to display the vHBA Details frame for that vHBA.**

- On the vHBA details frame, click the Targets tab.

WWNN	WWPN	LUN IDs
20:00:00:50:CC:20:0E:8E	24:00:00:50:CC:20:0E:8E	0,7

▼ Scan for New Servers

When servers are connected to a Fabric Device, they become available on the fabric. Also, if servers drop offline for any reason, or are disconnected from the fabric, scanning will remove them from Oracle Fabric Manager. Adding and deleting servers can occur automatically in Oracle Fabric Manager, but you can also perform a manual scan for physical servers.

To scan the fabric, follow this procedure:

- Select **Server Resource Manager → Physical Servers** to display the **Physical Server Summary**.
- Click the **Scan for New or Stale Servers** button.
A confirmation dialog is displayed.

Host Name	Host OS	Adapter FW Version	vNICs	vHBAs	Bound	Busy	State	IO Profile Name	Fabric Device Ports	Groups
bering	Linux/2.6.18-238.el5.xg-3.7.1.LX1/x86_64	5.3.0/3.0.0	0	2	✓		up	LA MirrorSite1_225	delaware:ServerPort22 oregon:ServerPort15	
CHARCOT-W2K8	Windows/6.1.7100/x64-2.2.0.36	5.3.0/2.1.1	14	0	✓		up	tomtest3	oregon:ServerPort17 delaware:ServerPort20	
COLEMAN	Windows/6.1.7600/x64-3.0.0.26	5.3.0/2.1.1	1	0	✓		up	tomtest2	oregon:ServerPort13 delaware:ServerPort11	
ovn86-45	VMware/ESX-4.1.0.xg-3.2.0-5773.ESX4...	5.3.0/3.0.0	3	2	✓		down	ml		

- Click **Yes** to initiate the rescan.

Depending on the number of servers that have been added or the number of stale servers in Oracle Fabric Manager, the scan process can take some time. Please be patient and allow the scan to complete.

▼ Apply an I/O Template to Servers

An I/O Template can be directly applied to a server to create a real-time connection for the vNICs and vHBAs in the template. When the template is applied, the vNICs and vHBAs in the template are automatically connected to the server, and an I/O Profile is added in connected state to the server. When the template is applied, all connections are up and I/O traffic is supported.

Note - You can also create an I/O configuration in disconnected state—for example, if you need to pre-provision a server's vHBA connectivity and then allow the network admin to create the network I/O config at a later time—by creating an I/O Profile. For information, see [“Working With I/O Profiles” on page 259](#).

Assuming an I/O Template is created and saved, you can apply the template to one or more servers as long as the servers do not already have an I/O Template connected to them.

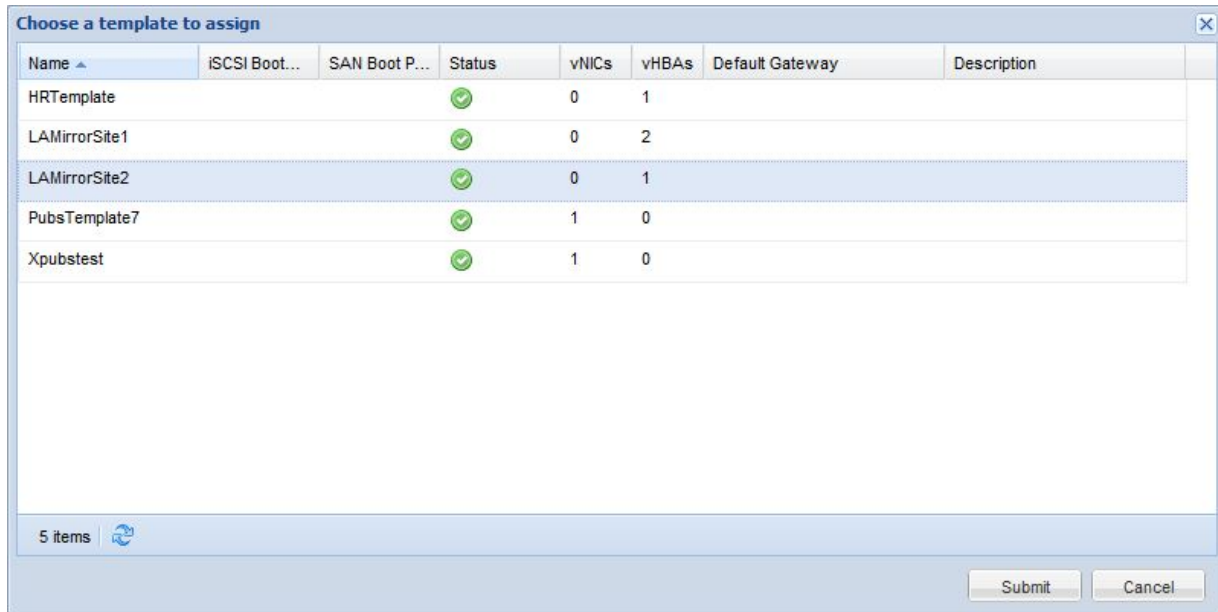
To apply a template to servers, follow this procedure:

1. **On the Physical Server Summary, select the unbound server to which you want to apply a template.**

Note - You can check the Status column on the Physical Server Summary to determine if the server is already connected. If so, select the server, then click disconnect it as documented in [“Disconnect an I/O Profile From a Server” on page 321](#).

2. **When the unbound server is selected, click the Apply Template to the Selected Servers button.**

The Choose A Template to Assign dialog is displayed.



3. Select the Template that you want to apply to the unbound server.
4. Click Submit and then click yes.

▼ Save a Server Configuration as an I/O Template

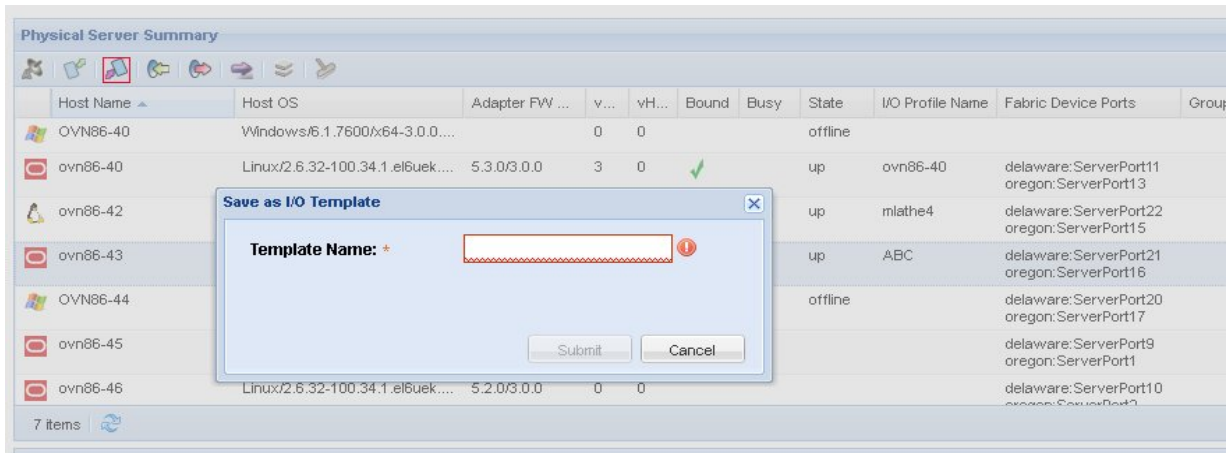
Oracle Fabric Manager supports saving a server's configuration as an I/O Template. When the configuration is saved as an I/O Template, you can then use that I/O Template for other servers that need the same configuration. By saving the server configuration as an I/O Template, you can create a “master configuration” that can be used to create I/O Profiles as needed. This process is basically cloning one server's configuration, which can then be used (by deploying I/O Profiles) on individual servers that need the same vNIC and vHBAs configuration.

Note - Remember that you can edit individual I/O Templates to change settings.

To save a server configuration as an I/O Template, follow this procedure:

1. **Select Server Resource Manager – Physical Servers to display the Physical Server Summary.**

2. **On the Physical Server Summary, select a physical server that is already bound to an I/O Profile.**
3. **Click the Save the selected server configuration as an I/O Template icon to display the Template Name dialog.**



4. **Enter the name that will be given to the I/O Template you are creating from the selected server's configuration, then click Submit.**
5. **Check the I/O Template Summary to verify that the I/O Template was saved correctly.**

▼ Connect an I/O Profile to a Server

Configured I/O Profiles do not automatically provide I/O to a server. The I/O Profile must be connected to the server for the server to have the vNICs and vHBAs pushed to the host. If a server is not already connected to an I/O Profile, you can connect it to any I/O Profile in disconnected state. When the I/O Profile is connected to a server, it will take a short time to push the network and storage connectivity to the host. Once established, the I/O Profile transitions from “disconnected” state to “up” state. While in “up” state, traffic can flow to and from the server. If the I/O Profile transitions from “disconnected” state to “partial” state, an error has prevented the I/O Profile from completely connecting to the server. This state requires your attention to fix.

To connect an I/O Profile to a server, follow this procedure:

1. **Display the Physical Server summary.**

2. **Select a physical server that is not bound to an I/O Profile.**
3. **Click the Connect I/O Profile to Server button to display the choose an I/O Profile to Select dialog.**
4. **Select an I/O Profile that you want to connect to the Physical Server.**
5. **Click Submit to connect the selected server to the selected I/O Profile.**

▼ Disconnect an I/O Profile From a Server

You can remove an I/O Profile that is already bound to a physical server. When you unbind an I/O Profile, all vNICs and vHBAs that are assigned to the server are completely removed. As a result, all traffic no longer moves between the server and the data and storage networks to which it is connected.

When the I/O Profile is disconnected from a server, that I/O Profile is in the disconnected state. The server itself remains online, but it is no longer connected to the Fabric Device at the network or storage level. The server is still physically connected to the Fabric Device through the cable between server's HCA and the Oracle Fabric Device.

Disconnecting a server is supported through the Physical Server Summary.

To disconnect a currently bound I/O Template, follow this procedure:

1. **Select Server Resource Manager → Physical Servers to display the Physical Server Summary.**
2. **Select a physical server that is already bound to an I/O Template.**
3. **Click the Unbind the Server button.**
4. **On the confirmation dialog, click Yes to confirm that you want to unbind the selected server from its I/O Template.**

▼ Migrate Resources to Different Servers

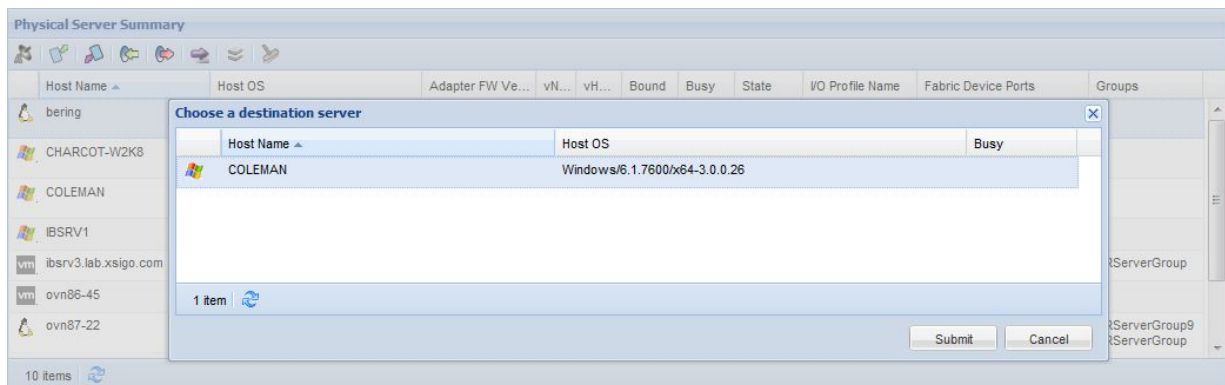
Oracle Fabric Manager supports migrating one server's I/O resources to another. When you migrate the resources, the virtual I/O is sent from a source physical server to a destination server(s) that you specify. Migrating a server's resources actually moves (not copies) the resources, so the underlying vNICs and vHBAs are disconnected from the source server, and reconnected to the destination server. As a result, a service interruption occurs while the resources are being deleted and reconnected on the new server.

Server resources can be migrated between individual servers. You cannot migrate server resources to a destination server that already has virtual I/O bound to it, so make sure that the destination server is not already bound, or if it does, unbind the current virtual I/O before migrating a server's configuration to another.

Migrating a server configuration is supported through the Physical Server Summary and presents a list of target servers to which the virtual resources can be moved. The process of actually moving the resources occurs by disconnecting the I/O profile on the source server, and reconnecting that I/O Profile to the destination server. As a result, a service interruption will occur while the I/O Profile is disconnecting and reconnecting.

To migrate a server configuration, follow this procedure:

1. **Select Server Resource Manager → Physical Servers to display the Physical Server Summary.**
2. **On the Physical Server Summary, select a bound physical server that has the resources you want to migrate.**
3. **Click Migrate the Virtual Resources to Another Server to display the Choose a Destination Server dialog.**



4. **Select the destination server that will receive the source server's configuration.**
5. **Click Submit.**
6. **On the confirmation dialog, click Yes.**
7. **Check the Physical Server Summary to verify that the source server no longer has the migrated resources, and the destination server now has the resources.**

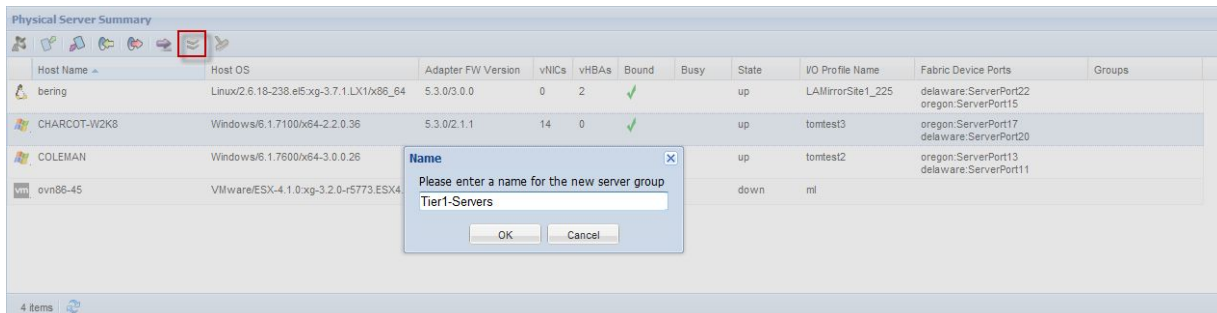
▼ Create a Server Group From Selected Servers

A Server Group is a software construction that enables you to select one or more physical servers and treat them as one logical unit for the purpose of certain management functions. Server Groups allow adding and removing individual servers.

You can create a Server Group through the Physical Server Summary by selecting one or more servers. For more information about Server Groups, see [“Working With Server Groups” on page 325](#).

To create a Server Group, follow this procedure:

1. **Select Server Resource Manager – Physical Servers to display the Physical Server Summary.**
2. **On the Physical Server Summary, select one or more physical servers.**
3. **Click the Create a Server Group from a Set of Servers button to display the Name dialog.**



4. **Enter the name for the new Server Group and click OK.**

▼ Remove Offline or Disconnected Servers

Servers that are managed by Oracle Fabric Manager are listed in the Physical Server Summary regardless of their state. As a result, the Physical Server Summary will display servers even if a server was online and has been taken offline—for example, by shutting it down, disconnecting its IB cables, experiencing different kinds of hardware failures, and so on. Such servers are effectively gone from Oracle Fabric Manager as if the server has been physically removed and will not be re-installed.

For transitory conditions (for example, a server is taken offline to replace a disk and then will be brought back online), the “offline” server(s) remain in Oracle Fabric Manager by design so that they can be remanaged after they are brought back online or connected with new resources without having to re-add them to the Oracle Fabric Manager system. For permanent conditions (for example, a server is being removed and not replaced), you will need to explicitly delete that server from Oracle Fabric Manager by using the Clean Up Offline Servers feature.

To permanently remove offline servers, follow this procedure:

1. **On the Physical Server Summary, select the offline or disconnected server that you want to remove.**

Note - You can check the Bound or State columns to determine if a server is offline or disconnected.

2. **Click Clean Up Offline Servers.**
3. **Click Yes to delete the offline server.**

When deleted, the server is no longer recognized by Oracle Fabric Manager. If you want to manage the server at a later date, you will need to add it back and rescan for it.

Working With Server Groups

This chapter documents the following topics:

- [“Server Groups” on page 325](#)
- [“Display Server Groups” on page 325](#)
- [“Rename Server Groups” on page 326](#)
- [“Configure Server Groups” on page 328](#)
- [“Add Servers to Server Groups” on page 329](#)
- [“Remove Servers From Server Groups” on page 330](#)
- [“Delete a Server Group” on page 331](#)

Server Groups

Server Groups are logical constructions that are assembled from individual servers.

When you create server groups, you are gathering individual servers in the network and putting them together in a logical construction that acts like a container.

You are free to group servers as needed for your network, but it is common to have a unifying theme for the server group. For example, grouping servers by department or business unit, by application type, or by hardware configuration is common.

▼ Display Server Groups

When server groups are configured, they are displayed in the Server Group Summary. To display the Server Groups configured in Oracle Fabric Manager, follow this procedure:

- On the navigation frame, select **Server Resource Manager** → **Server Groups**.

Server Group Summary			
Group Name	Number of Servers	Bound	Description
HRServerGroup	1		

1 item

Field	Indicates...
Group Name	The name of the server group(s) that are configured within Oracle Fabric Manager.
Number of Servers	The number of physical servers in each of the listed Server Groups.
Bound	<p>The presence or absence of an I/O Profile on any of the servers in the server group.</p> <ul style="list-style-type: none"> ■ A check mark indicates that one or more of the servers in the Server Group has an I/O Profile. ■ If no check mark is displayed, one or more of the servers has no I/O Profile. <p>Server Groups that do not contain a check mark are viable candidates for receiving virtual resources that will be migrated from another server group.</p>
Description	An optional description for each server group.

Also, when Server Groups are configured, you can see them in other physical server-based displays within Oracle Fabric Manager, for example, the Virtual Topology.

▼ Rename Server Groups

Oracle Fabric Manager supports renaming a server group, which enables you to change the name without having to completely delete and recreate the entire server group. When the server group is renamed, all other properties for the server group are retained, including the server membership for the group. As an option, you can also set or change the description for the Server Group.

You can rename the server group through the Server Group Details frame. To rename the server group, follow this procedure:

1. On the Navigation Frame, select Server Resource Manager → Server Groups.

The screenshot shows two parts of the vSphere interface. The top part is the 'Server Group Summary' for 'HRServerGroup', showing 1 server. The bottom part is the 'Server Groups: HRServerGroup' details view, showing a table of servers.

Group Name	Number of Servers	Bound	Description
HRServerGroup	1		

Host Name	Host OS	Adapter FW Vers...	vNI...	vHB...	Bound	Busy	State	I/O Profile Name	Fabric Device Ports	Groups
lbrsv3.lab.xsigo.com	VMware/ESXi-5.1.1.ESX1M.1x86_64	2.9.1000/3.0.0	0	0			offline		ExtSw-2c90200443198-Port13	HRServerGroup

2. Select a server group to populate the details frame with its properties.
3. Click the General tab to display general properties for the selected server group.
4. On the General tab, click the Edit button to edit the properties of the selected Server Group.

The screenshot shows the 'General' tab of the 'Server Groups: HRServerGroup' details view. The 'Group Name' field is highlighted with a red asterisk and contains 'HRServerGroup'. The 'Description' field contains 'Adding a description'. There are 'Submit' and 'Cancel' buttons at the bottom.

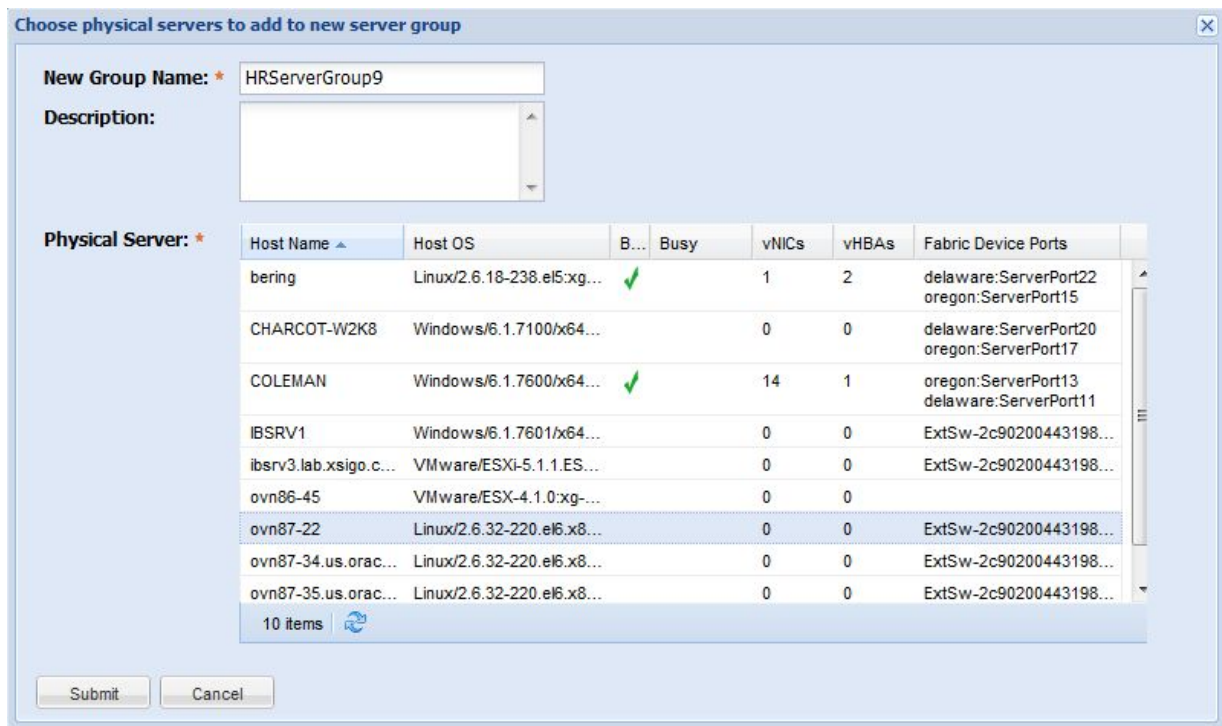
5. In the Group Name field, enter the new name for the Server Group.
6. As an option, you also can set or change the description for the selected Server Group.
7. When the new name has been specified for the Server Group, click Submit.

▼ Configure Server Groups

When you configure server groups, you create the logical “container” by selecting one or more servers and naming the server group. After servers are added to the server group, they are still displayable and manageable as individual entities, but since servers in the server group are typically similar in one way, you also gain the ability to manage the servers as one entity.

Server Groups are configurable through the Server Group Summary. To configure a Server Group, follow this procedure:

1. On the navigation frame, select **Server Resource Manager** → **Server Groups**.
2. Click the plus sign to display the **Create a Server Group** dialog.



3. In the **New Group Name** field, enter the name of the Server Group that you are creating.
4. As an option, in the **Description** field, you can enter a description for the server group that you are creating.

5. **From the Physical Servers table, select the servers that you want in the server group that you are creating.**

You can select one server to create the server group, then add more servers. Or, you can select multiple servers at once in the Physical Servers table.

- You can use bound servers to create a server group.
- You can mix bound and unbound servers to create a server group.

6. **Click Submit to create the server group.**
7. **On the confirmation dialog, click Yes to create the server group.**

When you click Yes, the Server Group Summary is displayed.

8. **Check the Server Group Summary to verify that the server group was created with the correct number of servers.**

▼ Add Servers to Server Groups

If a server group is configured, you can add more standalone servers to it. The servers you add can be either unassigned to an existing server group, or part of an existing server group. Servers can be bound or unbound, and of different types, but typically server groups are created from servers that have some kind of similarity (hardware, OS, hypervisor, or function). The following procedure assumes that a server group already exists. If it doesn't, it must be created before adding more servers to it. (The following procedure is valid even for an “empty” server group, which is a server group that is created with no servers in it.)

To add servers to an existing server group, follow this procedure:

1. **On the Server Group Summary, click to select the server group to which you want to add more servers.**

This step populates the details frame with the information for the server group you select.

2. **On the Server Group details frame, click the Servers tab to display the Servers in the Server Group.**

3. On the Servers tab, click the plus sign (+) to display the Choose Physical Server to Add to Server Group dialog.

Host Name	Host OS	Bo...	Busy	vNICs	vHBAs	Fabric Device Ports
bering	Linux/2.6.18-238.el5.xg-3.7.1.LX...	✓		1	2	delaware:ServerPort22 oregon:ServerPort15
CHARCOT-W2K8	Windows/6.1.7100/x64-2.2.0.36			0	0	delaware:ServerPort20 oregon:ServerPort17
COLEMAN	Windows/6.1.7600/x64-3.0.0.26	✓		14	1	oregon:ServerPort13 delaware:ServerPort11
IBSRV1	Windows/6.1.7601/x64-5.0.0.46			0	0	ExtSw-2c90200443198-Port15
ovn86-45	VMware/ESX-4.1.0.xg-3.2.0-r57...			0	0	
ovn87-22	Linux/2.6.32-220.el6.x86_64.xg-...			0	0	ExtSw-2c90200443198-Port7
ovn87-34.us.oracle.com	Linux/2.6.32-220.el6.x86_64/x8...			0	0	ExtSw-2c90200443198-Port10
ovn87-35.us.oracle.com	Linux/2.6.32-220.el6.x86_64/x8...			0	0	ExtSw-2c90200443198-Port9
ovn87-36.us.oracle.com	Linux/2.6.32-220.el6.x86_64/x8...			0	0	ExtSw-2c90200443198-Port8

9 items

Submit Cancel

4. Selecte the server(s) that you want to add to the existing Server Group, and then click Submit to add the selected servers.

▼ Remove Servers From Server Groups

When a Server Group is configured, you can delete individual servers from the Server Group at any time. When you delete a server, it reverts to being a single, standalone manageable object. When a server is deleted from a Server Group, it remains in the same operating state and has all the same software features (OS/hypervisor, host drivers, vNICs, vHBAs, and so on). The only difference is that it is no longer a member of a Server Group.

To remove servers from an existing server group, follow this procedure:

1. On the Server Group Summary, click to select the server group from which you want to remove one or more servers.
2. On the Server Group details frame, click the Servers tab to display the Servers in the Server Group.
3. On the Servers tab, click in the row(s) that contains the server(s) you want to delete.

Do not click the server name or you will go to a different page where you cannot delete the server.

4. **When you have selected the server(s) that you want to add to the existing Server Group, click the Delete Server button (the red minus sign) to delete the selected server(s) from the Server Group.**



It is possible to delete all servers from a Server Group. In this case, the Server Group is not deleted. It remains in Oracle Fabric Manager as an empty Server Group, which you can use as needed at any time.

▼ Delete a Server Group

You can delete a server group at any time after the group is created. Deleting a server group removes the logical container around the servers, and as a result, the servers must be configured or managed individually. The servers are otherwise not affected. For example, any virtual resources that were migrated from servers in one server group to another remain deployed on the target server(s). Any I/O Templates that were assigned to the servers in the Server Group remain assigned.

To delete a server group, follow this procedure:

1. **From the navigation frame, display the Server Group Summary (Server Resource Manager → Server Groups).**
2. **Click one or more server groups to select them, then click the garbage can icon.**
When you click the garbage can icon, a confirmation dialog is displayed to verify that you actually want to delete the selected Server Groups.
3. **On the confirmation dialog, click Yes to accept the deletion and complete removing the selected Server Group(s).**

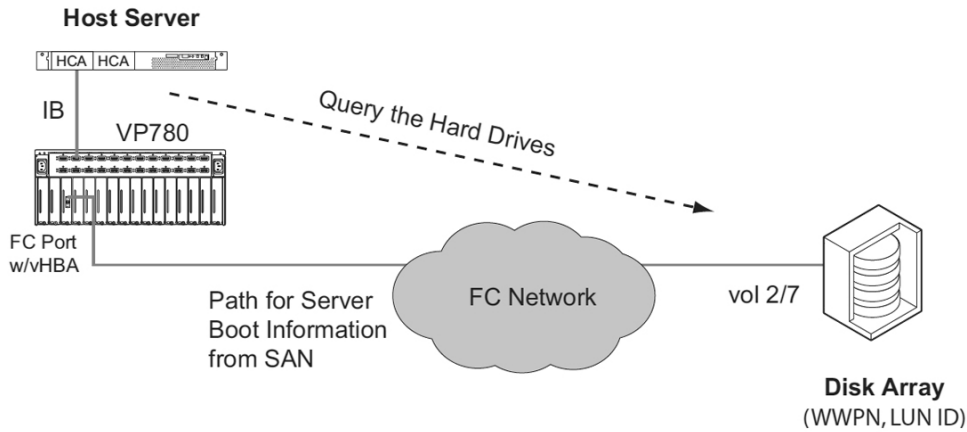
Working With Boot Profiles

This chapter documents the creation and use of SAN Boot Profiles and iSCSI Boot Profiles.

- [“SAN Boot Profiles” on page 333](#)
- [“SAN Boot Setup Overview” on page 336](#)
- [“Creating a SAN Boot Profile” on page 336](#)
- [“Creating the Bootable I/O Template” on page 342](#)
- [“SAN Boot Profile Summary” on page 349](#)
- [“SAN Boot Profile Details” on page 350](#)
- [“Create an I/O Profile From the I/O Template” on page 352](#)
- [“Connect the I/O Profile to a Physical Server” on page 352](#)
- [“Delete a SAN Boot Profile” on page 352](#)
- [“iSCSI Boot Overview” on page 353](#)
- [“iSCSI Boot Setup Overview” on page 354](#)
- [“Creating an iSCSI Boot Profile” on page 355](#)
- [“Creating the Bootable I/O Template” on page 363](#)
- [“iSCSI Boot Profiles Summary” on page 369](#)
- [“iSCSI Boot Profile Details” on page 370](#)
- [“Create an I/O Profile From the I/O Template” on page 372](#)
- [“Connect the I/O Profile to a Physical Server” on page 372](#)
- [“Delete an iSCSI Boot Profile” on page 373](#)

SAN Boot Profiles

SAN Boot allows you to boot a server or virtual machine from a SAN disk accessed through a vHBA. The disk is identified by a target World Wide Port Name (WWPN) and Logical Unit Number (LUN) ID on a storage disk array.

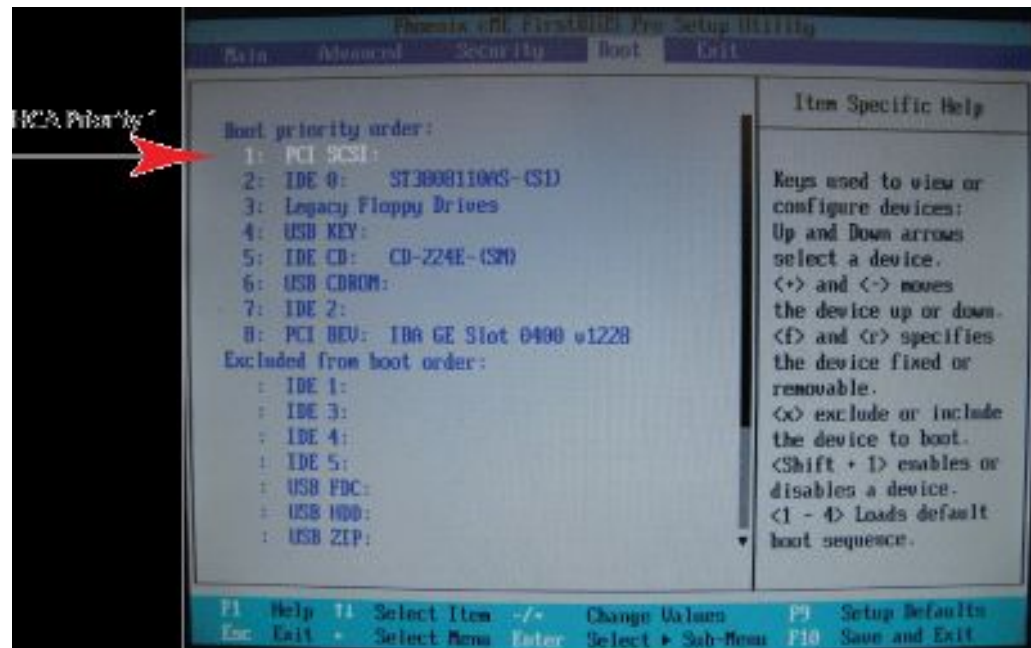


SAN Boot Sequence

All computers boot from local boot devices. However to boot from a “remote” device, you need to make the device appear to be local. Oracle implemented ROM BIOS extensions for HCA cards that follow these boot-sequence steps:

1. On power up, the server's BIOS performs basic hardware initialization.
2. The host server establishes a connection to the Oracle Fabric Device, where the system determines which vHBA to use and how to set up the communication path from the host server to the hard disk in the storage array.
3. BIOS reads the boot sequence, where the Oracle HCA is the first boot device (highest priority) in the list. Be sure the Oracle HCA is moved up from the “Excluded from boot order” list into the “Boot priority order” list.

The HCA card appears as a generic “PCI SCSI” device.



1. The OS Loader is installed, which is specific to each OS:

For Linux, the loader is GRand Unified Bootloader (GRUB). This loader resides in the boot sector of a bootable disk. The software responsible for loading the GRUB loader is held in the Option ROM of the HCA and runs in the context of the BIOS.

When the loader runs, it typically reads a configuration file from the disk and allows the user to boot the operating system in a number of configurations. For GRUB, this file is called grub.conf.

For Linux, GRUB will load the kernel and initrd into memory and begin running the kernel. The root file system (rootfs) given to the kernel will be the initrd. An initial RAM disk (initrd) is a compressed CPIO image (or compressed ext2 image for older kernels).

The kernel runs a program in the initrd in the location /init. Typically this program is a shell script that loads the kernel modules and mounts the real root file system.

2. Interrogate the hard disk to determine if it is present and if the disk is bootable (or not). If the hard disk is bootable, the Oracle HCA functions as a hard-disk controller. BIOS begins to treat the Oracle HCA as the local ID controller for the local hard disk.
3. The data (Linux kernel) is sent from the hard disk and loaded into memory. The kernel begins to work and loads all the necessary drivers into the host server.

Additional information about preparing the host server for SAN Booting over an Oracle vHBA is available through the Remote Booting Guide.

SAN Boot Setup Overview

SAN Boot has the following phases:

1. Create the SAN Boot Profile
2. Create a Bootable I/O Template for either a single or dual path vHBA to the LUN that contains the SAN Boot information.
3. Create an I/O Profile.

For this procedure, you will find it helpful to have the following information available before starting the procedure:

- The Server GUID
- The LUN from which the server will be booting. This is the LUN where the server's boot information is located in the SAN. The Oracle vHBA must be connected to this LUN to provide a path for the boot information to reach the server on which the vHBA is deployed.

Creating a SAN Boot Profile

- [“SAN Boot Profile” on page 336](#)
- [“Create a SAN Boot Profile” on page 337](#)
- [“Create a Static SAN Boot Profile” on page 338](#)
- [“Create a Direct-Attach SAN Boot Profile” on page 339](#)
- [“Create a Logical Volume Manager SAN Boot Profile” on page 340](#)

SAN Boot Profile

When the physical server has been installed with an Oracle HCA with SAN Boot Option ROM, SAN Boot requires a vHBA and a SAN Boot Profile to support server bootup from SAN disk.

As part of creating a SAN Boot Profile, you specify the root file system through a LUN on a target. The SAN Boot Profile supports different ways of mounting the root file system:

- Direct, which allows you to specify a device name for the LUN that contains the boot image and the root file system, and always use that device name. Otherwise, when devices are discovered you cannot guarantee that the SAN boot device is used consistently.
- Static, which allows you to manually specify the location of the SAN boot information and configure SAN booting on the host server.
- Logical volume manager, which allows you to specify a group and volume that contain the root file system.

Be aware that the boot image must be entirely contained within one LUN. The boot image file cannot be striped across multiple LUNs. The root file system must also be entirely contained within one LUN.

When you configure a SAN Boot vHBA, the vHBA supports both SAN Boot and vHBA functionality.

▼ Create a SAN Boot Profile

To create a SAN Boot Profile, follow this procedure:

1. **On the navigation panel, select Server Resource Manager → Boot Profiles → SAN Boot Profile Summary tab.**

The SAN Boot Profile Summary supports adding a new SAN Boot Profile or deleting a configured SAN Boot Profile through the plus sign (+) and garbage can icon, respectively.

2. **Click the plus sign (+) to start the Create SAN Boot Profile dialog.**

The screenshot shows a dialog box titled "Create SAN Boot Profile". It contains the following fields and controls:

- Name:** A text box containing "sanboot_3".
- Description:** A large empty text area.
- Mount Type:** A dropdown menu currently set to "Static".
- Device Name:** An empty text box.
- Group Name:** An empty text box.
- Volume Name:** An empty text box.
- Buttons:** "Submit" and "Cancel" buttons at the bottom right.

Mount types Static, Direct Attach, and Logical Volume Manager are supported for SAN Boot Profiles. For more information, see the following sections:

- [“Create a Static SAN Boot Profile” on page 338.](#)
- [“Create a Direct-Attach SAN Boot Profile” on page 339.](#)
- [“Create a Logical Volume Manager SAN Boot Profile” on page 340.](#)

▼ Create a Static SAN Boot Profile

To configure a static SAN Boot Profile:

1. **In the Name field, enter the name for the SAN Boot Profile that you are creating.**
The name can be an alphanumeric character string, and typically relates to the server(s) that use the SAN Boot Profile—for example, LinuxFinance for all the Linux servers in the Finance server group.
2. **As an option, in the Description field, you can enter a description for the SAN Boot Profile that you are creating.**
3. **From the Mount Type dropdown menu, select Static.**



The screenshot shows a dialog box titled "Create SAN Boot Profile" with a close button (X) in the top right corner. The dialog contains several input fields and a dropdown menu:

- Name:** A text field containing "sanboot_3".
- Description:** A large empty text area.
- Mount Type:** A dropdown menu with "Static" selected.
- Device Name:** An empty text field.
- Group Name:** An empty text field.
- Volume Name:** An empty text field.

At the bottom of the dialog are two buttons: "Submit" and "Cancel".

4. Click Submit to create the SAN Boot Profile.

At this point, the SAN Boot Profile is configured. However, for the server to receive the SAN Boot information, it must have a vHBA configured on the server, which occurs by configuring an I/O Profile from the I/O Template that contains the bootable vHBA, then connecting the I/O Profile to the server. For information about I/O Profiles, see [“Working With I/O Profiles” on page 259](#).

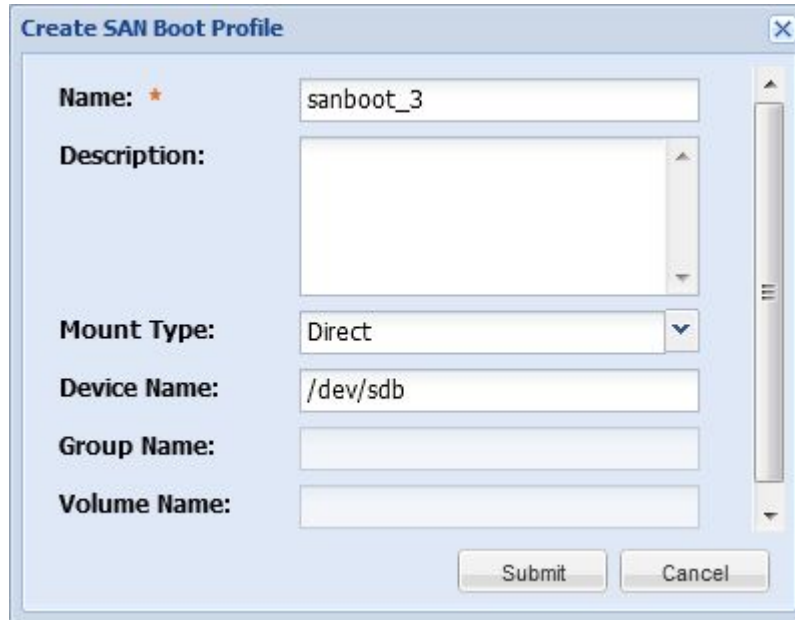
Also, the server's BIOS boot order must be edited to insert the Oracle HCA as the highest priority boot device in the server's boot priority list. See [“SAN Boot Profiles” on page 333](#).

▼ Create a Direct-Attach SAN Boot Profile

To configure the SAN Boot Profile for Direct Attached storage:

- 1. In the Name field, enter the name for the SAN Boot Profile that you are creating.**
The name can be an alphanumeric character string, and typically relates to the server(s) that use the SAN Boot Profile—for example, LinuxFinance for all the Linux servers in the Finance server group.
- 2. As an option, in the Description field, you can enter a description for the SAN Boot Profile that you are creating.**

3. From the Mount Type dropdown menu, select Direct.



The screenshot shows a dialog box titled "Create SAN Boot Profile". It contains the following fields and values:

- Name:** * sanboot_3
- Description:** (empty text area)
- Mount Type:** Direct
- Device Name:** /dev/sdb
- Group Name:** (empty text field)
- Volume Name:** (empty text field)

At the bottom of the dialog are "Submit" and "Cancel" buttons.

4. In the Device Name field, enter the name of the server's boot device that will receive the SAN boot information (kernel and initrd) from the SAN.
5. Click Submit to create the SAN Boot Profile.

At this point, the SAN Boot Profile is configured. However, for the server to receive the SAN Boot information, it must have a vHBA configured on the server, which occurs by configuring an I/O Profile from the I/O Template that contains the bootable vHBA, then connecting the I/O Profile to the server. For information about I/O Profiles, see [“Working With I/O Profiles” on page 259](#).

Also, the server's BIOS boot order must be edited to insert the Oracle HCA as the highest priority boot device in the server's boot priority list. See [“SAN Boot Profiles” on page 333](#).

▼ Create a Logical Volume Manager SAN Boot Profile

To configure the SAN Boot Profile for Logical Volume Manager (LVM):

1. In the Name field, enter the name for the SAN Boot Profile that you are creating.

The name can be an alphanumeric character string, and typically relates to the server(s) that use the SAN Boot Profile—for example, LinuxFinance for all the Linux servers in the Finance server group.

2. **As an option, in the Description field, you can enter a description for the SAN Boot Profile that you are creating.**
3. **From the Mount Type dropdown menu, select Logical Volume Manager.**



The screenshot shows a dialog box titled "Create SAN Boot Profile". It contains the following fields and values:

- Name:** * sanboot_3
- Description:** (empty text area)
- Mount Type:** Logical Volume Manager
- Device Name:** /dev/sdb
- Group Name:** VolGroup00
- Volume Name:** LogVol00

Buttons at the bottom: Submit, Cancel.

4. **In the Group Name field, enter the name of the Volume Group that contains the volume where SAN Boot information is located.**
5. **In the Volume Name field, enter the name of the volume on which the SAN Boot information is located.**
6. **Click Submit to create the SAN Boot Profile.**
7. **Proceed to the next section.**

Also, the server's BIOS boot order must be edited to insert the Oracle HCA as the highest priority boot device in the server's boot priority list. See ["SAN Boot Profiles" on page 333](#).

Creating the Bootable I/O Template

This section contains the following topics:

- [“Bootable I/O Template” on page 342](#)
- [“Create a Single-Path Bootable I/O Template” on page 342](#)
- [“Create a Dual-Path Bootable I/O Template” on page 346](#)

Bootable I/O Template

When the SAN Boot Profile is created, you must associate an I/O Template with it. The I/O Template provides the vHBA that connects the LUN in the SAN where the server's boot information (kernel, boot image, and so on) are located. This section documents how to create an I/O Template with a bootable vHBA. After creating the bootable template, you will need to also create an I/O Profile from that I/O Template. For more information about I/O Profiles, see [“Working With I/O Profiles” on page 259](#).

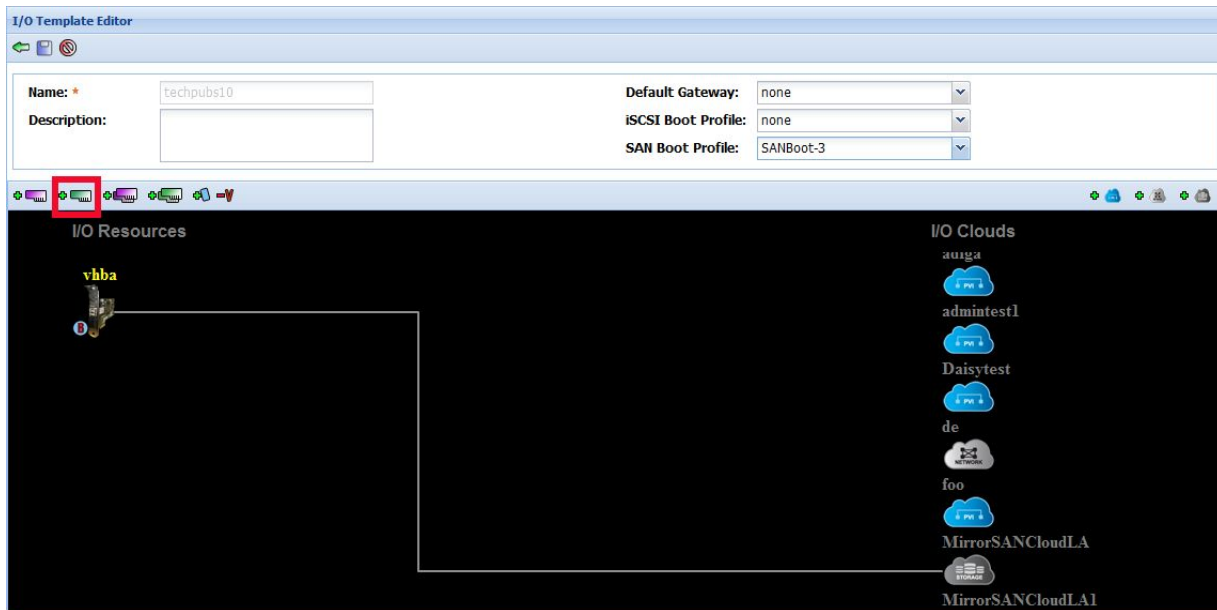
When you create the I/O Template for the SAN Boot Profile, make sure that you select the bootable field. Without this setting the vHBA in the I/O Template will not be able to support SAN Boot functionality. When the bootable option is set, it does not preclude the vHBA from carrying standard read and write data. Instead, the bootable option allows the vHBA to support SAN Booting the server in addition to read and write I/O.

Oracle Fabric Manager supports SAN Boot functionality through one or two paths, which are configured through the I/O Template.

▼ Create a Single-Path Bootable I/O Template

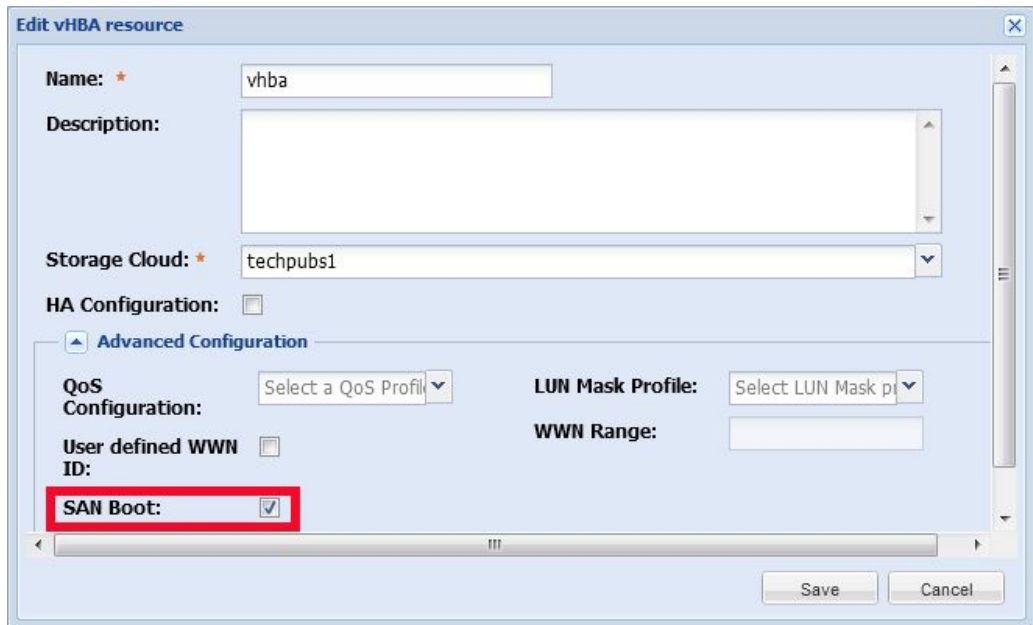
To create the bootable I/O Template, follow this procedure:

1. Select Server Resource Manager → I/O Templates to display the I/O Template summary, then click the plus sign (+) to display the I/O Template Editor. The following example shows the I/O Template Editor.



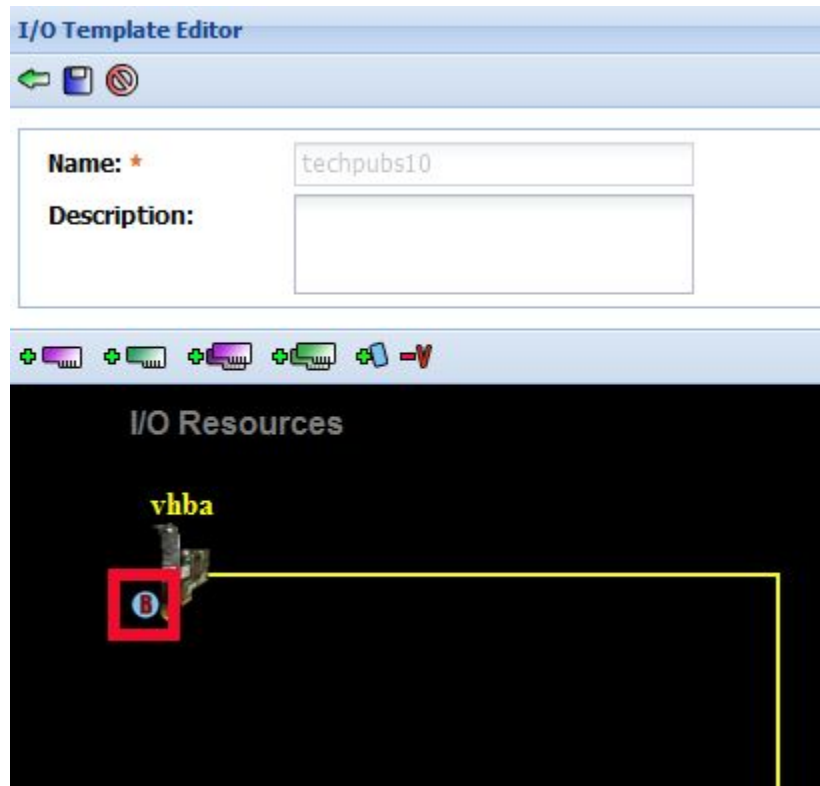
2. In the Name field, enter the name of the I/O Template that will be used for SAN Booting.
3. From the SAN Boot Profile dropdown menu, make sure to select the SAN Boot Profile you just created in the preceding section.
Without creating the I/O Template as a bootable I/O Template, the server where the vHBA is deployed cannot be SAN Booted.
4. Click and drag the vHBA to the appropriate Storage Cloud to terminate the vHBA on an FC port.
5. Double click the vHBA icon on the I/O Template work space to display the Edit vHBA Resource dialog.

6. On the Edit vHBA Resource dialog, click the Advanced Configuration button to display the advanced properties for the selected vHBA.



7. Click the SAN Boot checkbox.
The checkbox must contain a check mark for SAN Booting capabilities to be supported on the vHBA.
8. Click Save to return to the I/O Template Editor.

9. Check the I/O Template Editor for the bootable icon (a circle with a red “B” within) associated with the server’s boot vHBA.



The presence of this icon indicates that the vHBA connected to the host is bootable. If the “bootable” icon is not present, make sure the SAN Boot checkbox contains a checkmark.

10. **On the I/O Template Editor, click Save.**
Without saving the I/O Template, all in-progress configuration will be lost.
11. **After the bootable I/O Template is configured, proceed to [“Create an I/O Profile From the I/O Template”](#) on page 352.**

▼ Create a Dual-Path Bootable I/O Template

Dual paths provide two paths to the same LUN so that a single point of failure through the Fabric Device's fabric is eliminated and availability of the SAN Boot information is increased. If one path is not available, the second one usually is.

Be aware that the two paths in dual pathing are not true HA, so there is no automatic failover or failback if one of the paths is not available. Instead, dual pathing simply provides to connections for the server to SAN Boot. If one path is not available, user intervention is required to select the other path through host-side methods, such as multipathing software.

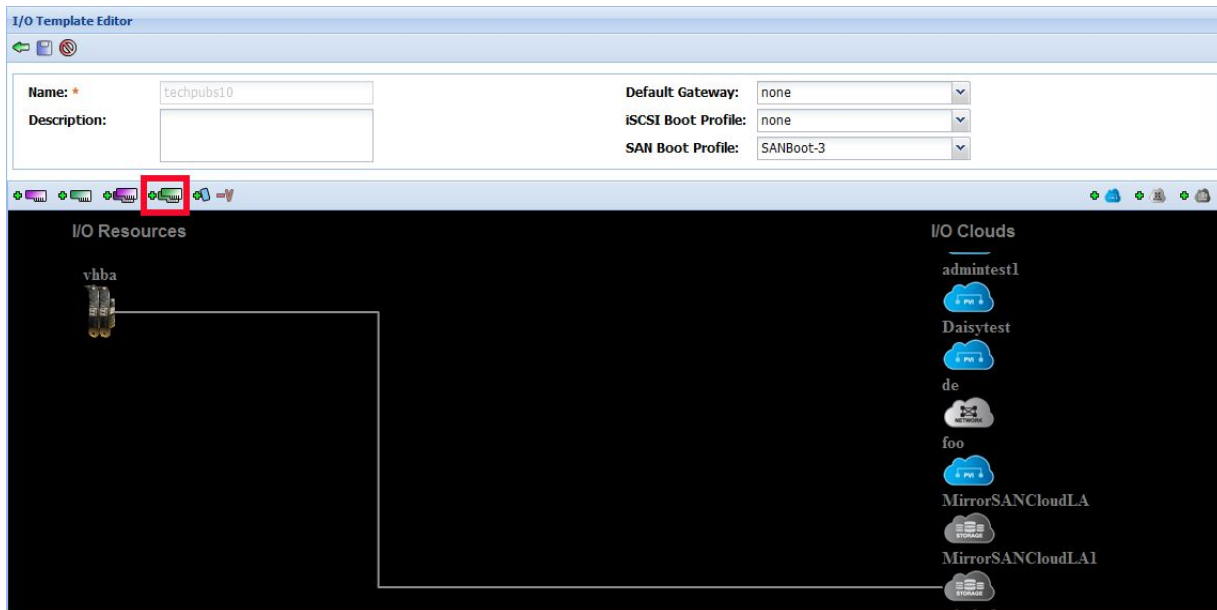
For dual path bootable I/O Templates, two vHBAs are configured with the following considerations:

- Two fibre channel ports must be configured in the Storage Cloud. The HA vHBA cannot be created unless a minimum of two ports exist in the Storage Cloud. Depending on how you want the dual pathing to work, you can have storage ports on the same module or Fabric Device as long as the minimum of two ports exist in the Storage Cloud.
- On the I/O Template Editor, you must use the HA vHBAs option instead of creating two single vHBAs.
- The two vHBAs must connect to the same LUN. You cannot put the server's SAN Boot information on multiple LUNs. Since the two vHBAs will be connected to the same LUN, you will terminate the two vHBAs in the same Storage Cloud.

To create a dual path bootable I/O Template, follow this procedure:

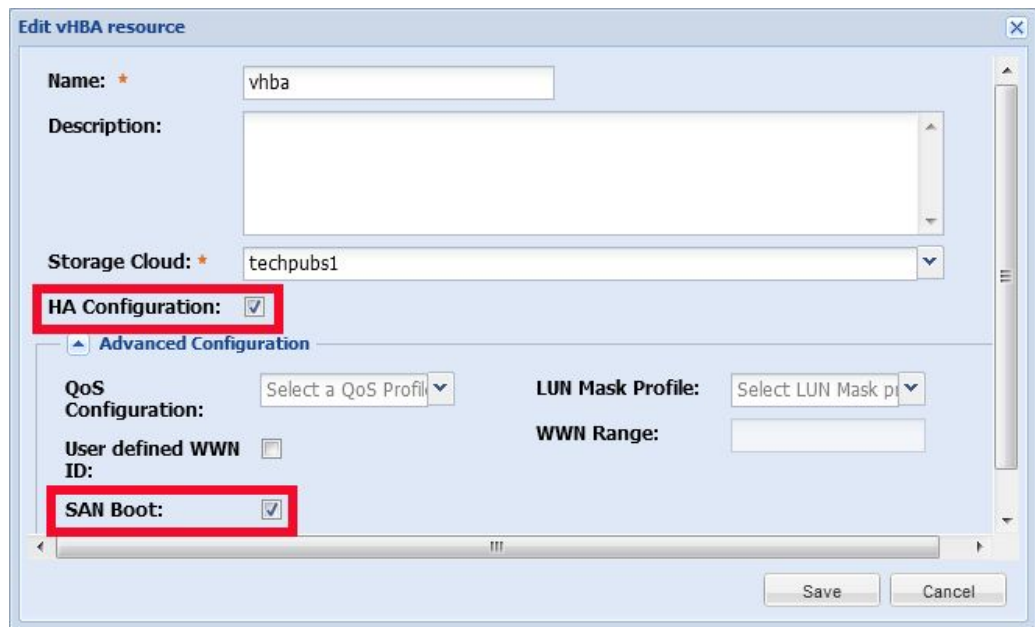
1. **Select Server Resource Manager → I/O Templates to display the I/O Template summary, then click the plus sign (+) to display the I/O Template Editor.**

- Click the Add an HA vHBA to Template button, then click and drag to connect the vHBA to the appropriate storage cloud.



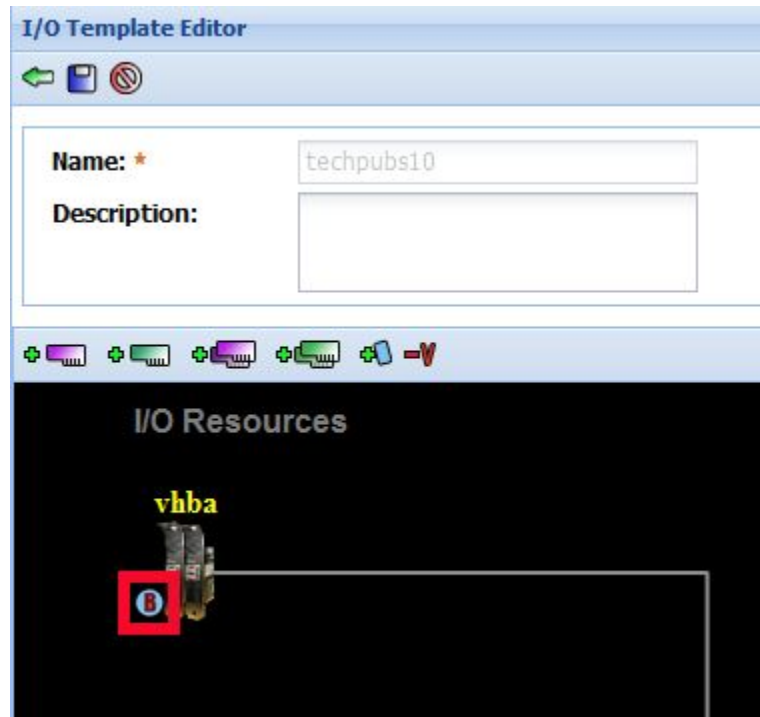
- In the Name field, enter the name of the I/O Template that will be used for SAN Booting.
- From the SAN Boot Profile dropdown menu, make sure to select the SAN Boot Profile you just created in the preceding section.
Without creating the I/O Template as a bootable I/O Template, the server where the vHBA is deployed cannot be SAN Booted.
- Click and drag the vHBA to the appropriate Storage Cloud to terminate the vHBA on an FC port.
- Double click the vHBA icon on the I/O Template work space to display the Edit vHBA Resource dialog.

7. On the Edit vHBA Resource dialog, click the Advanced Configuration button to display the advanced properties available to the HA vHBA that you are creating.



8. On the Edit vHBA Resource dialog, make sure that the HA Configuration checkbox contains a checkmark. If it does not, click it now.
Without this checkbox properly set, only one vHBA (and therefore one path) will be supported for SAN Booting the server.
9. Click the SAN Boot checkbox. The checkbox must contain a check mark for SAN Booting capabilities to be supported on the vHBA.
10. Click Save to return to the I/O Template Editor.
11. Check the I/O Template Editor for the bootable icon (a circle with a red "B" within) associated with the server's boot vHBAs.

The presence of this icon indicates that the dual vHBAs connected to the host are bootable. If the “bootable” icon is not present, make sure the SAN Boot checkbox contains a checkmark.



12. **On the I/O Template Editor, click Save.**
Without saving the I/O Template, all in-process configuration will be lost.
13. **After the bootable I/O Template is configured, proceed to [“Create an I/O Profile From the I/O Template”](#) on page 352.**

SAN Boot Profile Summary

The SAN Boot Profiles are displayed in the SAN Boot Profile Summary, which contains information about all the SAN Boot Profiles that have been created through Oracle Fabric Manager. The table shows general information about the SAN Boot Profile such as the server profile and vHBA that support the connection to the SAN disk that contains the server's kernel and initrd.

The SAN Boot Profile Summary also contains a link to the SAN Boot Details page where additional information about the configuration is available for the vHBA that is supporting the connection to the SAN disk that contains the server's kernel and initrd.

The following figure shows the SAN Boot Profile Summary, which is available through the Server Resource Manager->Boot Profiles->SAN Boot Profile Summary.

Name	Mount Type	Group Name	Volume Name	Mount Device	IOTemplates	IOProfiles
sanboot_1	lvm	VolGroup00	LogVol00		0	2
sanboot_2	static				0	3
sanboot_3	lvm	VolGrop00	LogVol00		0	0

Field	Indicates
Name	The name of the SAN Boot Profile.
Mount Type	The type of mount configured on the SAN Boot vHBA. Valid values are: <ul style="list-style-type: none"> ■ static, for a Static SAN Boot Profile. ■ lvm, for the Logical Volume Manager. ■ direct, for assigning the device name that has the root file system.
Group Name	The logical volume manager (LVM) group name. If the mount type is LVM, this field displays the group name. If the mount type is other than LVM, no value is displayed.
Volume Name	The name of the logical volume that LVM uses for the root file system mount point.
Mount Device	The name of the device that contains the root file system mount point for the SAN Boot Profile.
I/O Template	The number of I/O Templates that are assigned to each SAN Boot Profile.
I/O Profiles	The number of I/O Profiles that are associated with each SAN Boot Profile.

SAN Boot Profile Details

The SAN Boot Profile Details frame contains information about any configured SAN Boot Profiles. The SAN Boot Details frame contains only one instance of a SAN Boot, which is determined by the SAN Boot Profile that you select in the SAN Boot Profile Summary. By clicking a specific SAN Boot Profile in the summary, all of its corresponding details are available in the details frame.

Through the SAN Boot Profile Details Frame you can also edit a configured SAN Boot Profile to make changes to the profile without having to delete then re-add the profile. Be aware that:

- Changes to the profile might also require resetting the server, or making changes to the server's configuration (for example, if a new boot device is specified, you might need to change the order of boot devices in the server's boot priorities list.)
- Editing a SAN Boot Profile is allowed only when the SAN Boot Profile is not currently associated with an I/O Template.

The SAN Boot Profile Details Frame is in the Oracle Fabric Manager GUI just below the SAN Boot Summary Profile.

The screenshot shows the Oracle Fabric Manager GUI. At the top, there are two tabs: "SAN Boot Profile Summary" (selected) and "ISCSI Boot Profile Summary". Below the tabs is a table with the following data:

Name	Mount Type	Group Name	Volume Name	Mount Device	IOTemplates	IOProfiles
pubstest	lvm	VolGroup99	LogVol00		0	0
sanboot_1	lvm	VolGroup00	LogVol00		0	2
sanboot_2	static				0	3
sanboot_3	lvm	VolGrop00	LogVol00		1	0

Below the table, there is a "4 items" indicator and a refresh icon. Below that, the "SAN Boot Profile : pubstest" details frame is shown, which is highlighted with a red border. It has a "General" tab and the following fields:

Name: pubstest
Mount Type: lvm
Group Name: VolGroup99
Volume Name: LogVol00
Mount Device:

At the bottom of the details frame is an "Edit" button.

The following table shows the contents of the SAN Boot Profile Details frame, and explains what each field means. Not all fields will be populated based on the options specified during the configuration of the SAN Boot Profile. For example, if you created a direct connect SAN Boot Profile, none of the fields for logical volume manager will contain information.

Field	Means
Name	The name assigned to the selected SAN Boot Profile.

Field	Means
Mount Type	The method used to mount the SAN location of the server's SAN Boot information: <ul style="list-style-type: none"> ■ Direct ■ Static ■ Logical Volume Manager
Group Name	For Logical Volume Manager mount types, this field shows the logical group that contains the SAN location of the server's SAN boot information.
Volume Name	For Logical Volume Manager mount types, this field shows the logical volume that contains the SAN location of the server's SAN boot information.
Mount Device	For Direct mount types, this field shows the device that contains the root file system mount point for the server's SAN Boot information.
I/O Template	The number of I/O Templates that are assigned to each SAN Boot Profile.
I/O Profiles	The number of I/O Profiles that are associated with each SAN Boot Profile.

Create an I/O Profile From the I/O Template

If you are setting up SAN Boot for a server that has already booted up and connected to the Fabric Device, you can now create an I/O Profile from the I/O Template. When the I/O Profile is created, it will contain the bootable vHBA that you created. You can create one more I/O Profiles from the same I/O Template.

For information about creating an I/O Profile, follow the procedure for creating an I/O Profile as documented in [“Create an I/O Profile” on page 262](#).

Connect the I/O Profile to a Physical Server

When you assign the I/O Profile to the server, you begin pushing the bootable vHBA and the associated SAN Boot Profile to the server.

To assign the I/O Template to a server, follow this procedure for connecting an I/O Profile to a Physical server, as documented in [“Connecting an I/O Profile to a Server” on page 266](#).

▼ Delete a SAN Boot Profile

Through the SAN Boot Profile Summary, you can delete a SAN Boot Profile. To delete a SAN Boot Profile, follow this procedure:

1. **On the navigation panel, click Server Resource Manager → Boot Profiles → SAN Boot Profile Summary.**

The SAN Boot Profile Summary is displayed.

2. Select the SAN Boot Profile that you want to delete.

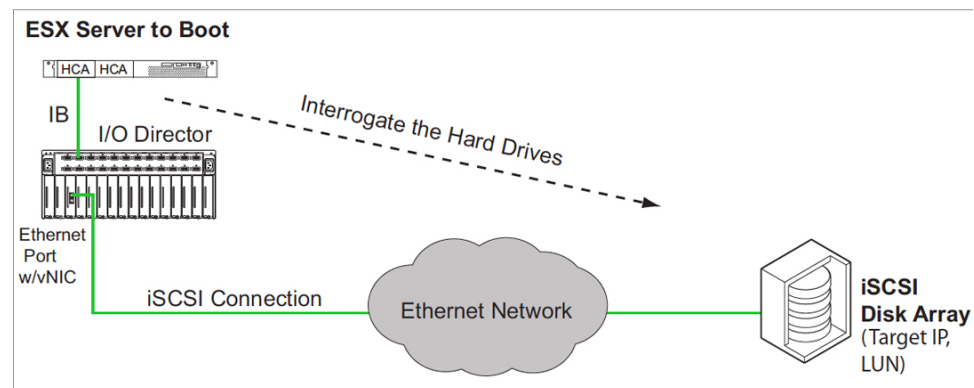
You can select multiple SAN Boot profiles through standard keyboard conventions, for example by pressing the CTRL + click or Shift + click.

3. Click the garbage can icon to delete the selected SAN Boot Profile.

4. Click OK in the confirmation window to delete the selected SAN Boot Profile(s), or Cancel to abort the deletion.

iSCSI Boot Overview

The Fabric Interconnect supports booting a server over a vNIC using an iSCSI connection. The following figure illustrates the topology used to achieve iSCSI booting.



The following table lists some terminology specific to iSCSI booting.

Term	Definition
initiator	The host server that is booting over an iSCSI connection
target	The iSCSI array
IQN	An iSCSI qualified name of an initiator or target
target IP	The IP address of the target filer or array

iSCSI Boot Configuration Overview

When configuring iSCSI boot, you perform the same general steps as for any remote booting setup:

1. Install the SAN volume with the necessary bits.
 - iSCSI boot uses the same initrd that Linux SAN Boot uses. You cannot use the supported Linux installers over iSCSI. You can find additional affirmation about patching the Anaconda installer and installing the boot disk in the Remote Booting Guide's "iSCSI Booting" chapter.
 - Alternatively, you can install over FC to the disk and then have your server boot over iSCSI. The iSCSI install options require a temporary FC or Ethernet connection to the volume.
2. Configure the Fabric Interconnect with the required virtual I/O resources.
Create a vNIC and server profile, using either DHCP or static addressing.
3. Configure the host server with drivers and firmware to enable remote booting.

iSCSI Boot Setup Overview

Setting up iSCSI Boot has the following phases:

1. Create the iSCSI Boot Profile
2. Create a Bootable I/O Template for either a single or dual path to the LUN that contains the SAN Boot information.
3. Create an I/O Profile. While the I/O Profile is being pushed to the host, you will be prompted to provide pertinent information, such as the iSCSI initiator to complete connecting the I/O Profile while the iSCSI Boot configuration is being pushed to the server.

For this procedure, you will find it helpful to have the following information available before starting the procedure:

- The Server GUID
- The LUN from which the server will be booting. This is the LUN where the server's boot information is located on the iSCSI storage array. The Oracle vNIC must be connected to this LUN to provide a path for the boot information to reach the server on which the vNIC is deployed.
- The target IQN (T-IQN), which you can get by logging into the storage.
- The initiator IQN (I-IQN), which you can get by displaying the vNIC properties. On the Fabric Device, you can get this information by issuing the show server-profile <profile-name> iscsi-boot -detail command. The -detail qualifier is required to display the I-IQN.

Creating an iSCSI Boot Profile

This section contains the following topics:

- [“iSCSI Boot Profiles” on page 355](#)
- [“Create an iSCSI Boot Profile” on page 356](#)
- [“Create a Static iSCSI Boot Profile” on page 358](#)
- [“Create a Direct Attach iSCSI Boot Profile” on page 360](#)
- [“Create a Logical Volume Manager iSCSI Boot Profile” on page 361](#)

iSCSI Boot Profiles

Creating an iSCSI Boot Profile enables the host server to access its boot information over a boot capable vNIC which connects the server to its boot information. The iSCSI boot information is supported on the server through an Oracle HCA and option ROM, which must be present in the host server that will be iSCSI booting. If the Oracle HCA and option ROM are not yet installed, install them now. This section documents how to create an iSCSI Boot Profile.

When an iSCSI Boot Profile is created, the necessary information is available for booting the server over a vNIC that connects to its boot information on iSCSI storage. However, the bootable vNIC that connects the server to the network is not configured as part of creating an iSCSI Boot Profile. The iSCSI Boot Profile must be connected to an I/O Template which contains a valid bootable vNIC for iSCSI booting. For information about creating an I/O Template, see [“Creating an I/O Template” on page 206](#).

As part of creating an iSCSI Boot Profile, you specify the root file system through a LUN on a target. The iSCSI Boot Profile supports different ways of mounting the root file system:

Through the iSCSI Profile Summary, you can add a new iSCSI Boot Profile or delete a selected iSCSI Boot Profile through the plus sign (+) and garbage can icon, respectively.

- Direct, which allows you to specify a device name for the LUN that contains the boot image and the root file system, and always use that device name. Otherwise, when devices are discovered you cannot guarantee that the iSCSI boot device is used consistently.
- Static, which allows you to manually specify the location of the iSCSI boot information and configure iSCSI booting on the host server.
- Logical volume manager, which allows you to specify a group and volume that contain the root file system.

Be aware that the boot image must be entirely contained within one LUN. The boot image file cannot be striped across multiple LUNs. The root file system must also be entirely contained within one LUN.

You also must specify the target IP Address Group, which is the IP address of the target filer or array.

When you configure an iSCSI Boot vNIC, the vNIC supports both iSCSI Boot and standard virtual I/O functionality.

▼ Create an iSCSI Boot Profile

To create an iSCSI Boot Profile, follow this procedure:

1. **On the navigation panel, select Server Resource Manager → Boot Profiles → iSCSI Boot Profile Summary tab.**

The iSCSI Boot Profile Summary is displayed.

Name	Target IP Address	Target Portal Group	Port	Proto...	Mount Ty...	Group Name	Volume Name	Mount Device	IO Templates	IO Profiles
iscsiboot_1	172.16.10.70		3260	6	direct				0	0
iscsiboot_10	192.168.8.108		3260	6	direct				0	0
iscsiboot_11	172.16.1.10		3260	6	direct				0	0
iscsiboot_12	172.16.2.10		3260	6	direct				0	0
iscsiboot_13	1.1.1.4		3260	6	direct				0	0
iscsiboot_14	1.1.2.4		3260	6	direct				0	0
iscsiboot_15	3.3.3.3		3260	6	static				0	0
iscsiboot_2	172.16.10.80		3260	6	direct				0	0
iscsiboot_3	192.168.46.3		3260	6	direct				0	0
iscsiboot_4	172.16.1.4		3260	6	direct				1	0

15 items

2. Click the plus sign to display the Create iSCSI Boot Profile dialog.

The screenshot shows a dialog box titled "Create iSCSI Boot Profile". It contains the following fields and values:

- Name:** * iscsiboot_19
- Target IP Address:** * 192.168.111.112
- Target Portal Group:** (empty)
- Protocol ID:** 6
- Port ID:** 3260
- Description:** (empty text area)
- Mount Type:** Static
- Device Name:** (empty)
- Group Name:** (empty)
- Volume Name:** (empty)

Buttons: Submit, Cancel

3. In the Name field, enter an alphanumeric character string that will name the iSCSI Boot Profile that you are creating.
4. In the Target IP Address field, enter the IP Address group for the filer or array that contains the server's iSCSI Boot information.
5. In the Target Portal Group field, enter the IP address for the iSCSI Portal Group. Depending on your iSCSI storage array, this field might not be required.
6. In the Protocol ID field, enter the number of the protocol that will be supporting iSCSI communication.

This field is populated with a default value, but you can change the communication protocol number if you need to.

- 7. In the Port ID field, enter the port number that will support the iSCSI communication between the server and the location of the server's iSCSI Boot information.**

This field is populated with a default value, but you can change the communication port number if you need to.

- 8. As an option, in the Description field, you can enter an alphanumeric character string that describes the iSCSI Boot Profile that you are creating.**
- 9. From the Mount Type dropdown menu, select one of the following methods of mounting the location of the server's SAN Boot Profile:**
 - **Static**, in which the iSCSI Boot profile connects to the statically assigned storage through the same vNIC. To configure this iSCSI Boot Profile, proceed to [“Create a Static iSCSI Boot Profile” on page 358](#).
 - **Direct**, in which the iSCSI Boot Profile connects to the same server boot device. To configure this iSCSI Boot Profile, proceed to [“Create a Direct Attach iSCSI Boot Profile” on page 360](#).
 - **Logical Volume Manager**, which contains a pointer to a LUN and Volume Group that contains the iSCSI Boot information. To configure this iSCSI Boot Profile, proceed to [“Create a Logical Volume Manager iSCSI Boot Profile” on page 361](#).

▼ Create a Static iSCSI Boot Profile

To configure a static iSCSI Boot Profile:

- 1. In the Name field, enter the name for the iSCSI Boot Profile that you are creating.**

The name can be an alphanumeric character string, and typically relates to the server(s) that use the iSCSI Boot Profile.
- 2. As an option, in the Description field, you can enter an alphanumeric character string that describes the iSCSI Boot Profile that you are creating.**
- 3. From the Mount Type dropdown menu, select Static.**

The Create iSCSI Boot Profile window is displayed.

The screenshot shows a dialog box titled "Create iSCSI Boot Profile". The fields are filled as follows:

- Name:
- Target IP Address:
- Target Portal Group:
- Protocol ID:
- Port ID:
- Description:
- Mount Type:
- Device Name:
- Group Name:
- Volume Name:

Buttons: Submit, Cancel

4. Click Submit to create the iSCSI Boot Profile.

At this point, the iSCSI Boot Profile is configured. However, for the server to receive the iSCSI Boot information, it must have a bootable vNIC configured on the server, which occurs by configuring an I/O Profile from the I/O Template that contains the bootable vNIC, then connecting the I/O Profile to the server. For information about I/O Profiles, see [“Working With I/O Profiles” on page 259](#).

Also, the server's BIOS boot order must be edited to insert the Oracle HCA as the highest priority boot device in the server's boot priority list. See [“SAN Boot Profiles” on page 333](#).

5. Proceed to [“Creating the Bootable I/O Template” on page 363](#).

▼ Create a Direct Attach iSCSI Boot Profile

To configure a direct attach iSCSI Boot Profile:

- 1. In the Name field, enter the name for the iSCSI Boot Profile that you are creating.**
The name can be an alphanumeric character string, and typically relates to the server(s) that use the iSCSI Boot Profile.
- 2. As an option, in the Description field, you can enter an alphanumeric character string that describes the iSCSI Boot Profile that you are creating.**
- 3. From the Mount Type dropdown menu, select Direct.**
This step activates the Device Name field.

Create iSCSI Boot Profile

Name: * iscsiboot_19

Target IP Address: * 192.168.111.112

Target Portal Group:

Protocol ID: 6

Port ID: 3260

Description:

Mount Type: Direct

Device Name: /dev/sdc

Group Name:

Volume Name:

Submit Cancel

4. **In the Device Name field, enter the name of the server's boot device that will receive the iSCSI boot information.**

5. **Click Submit to create the iSCSI Boot Profile.**

At this point, the iSCSI Boot Profile is configured. However, for the server to receive the iSCSI Boot information, it must have a bootable vNIC configured on the server, which occurs by configuring an I/O Profile from the I/O Template that contains the bootable vNIC, then connecting the I/O Profile to the server. For information about I/O Profiles, see [“Working With I/O Profiles” on page 259](#).

Also, the server's BIOS boot order must be edited to insert the Oracle HCA as the highest priority boot device in the server's boot priority list. See [“SAN Boot Profiles” on page 333](#).

6. **Proceed to [“Creating the Bootable I/O Template” on page 363](#).**

▼ **Create a Logical Volume Manager iSCSI Boot Profile**

To configure the iSCSI Boot Profile to use a logical volume manager:

1. **In the Name field, enter the name for the iSCSI Boot Profile that you are creating.**
The name can be an alphanumeric character string, and typically relates to the server(s) that use the iSCSI Boot Profile.
2. **As an option, in the Description field, you can enter an alphanumeric character string that describes the iSCSI Boot Profile that you are creating.**
3. **From the Mount Type dropdown menu, select Logical Volume Manager.**

This step activates the Group Name and Volume Name fields.

The screenshot shows a dialog box titled "Create iSCSI Boot Profile". It contains the following fields and values:

- Name:
- Target IP Address:
- Target Portal Group:
- Protocol ID:
- Port ID:
- Description:
- Mount Type:
- Device Name:
- Group Name:
- Volume Name:

At the bottom right, there are "Submit" and "Cancel" buttons.

4. In the Group Name field, enter the name of the Volume Group that contains the volume where the server's iSCSI Boot information is located.
5. In the Volume Name field, enter the name of the volume on which the server's iSCSI Boot information is located.
6. **Click Submit to create the iSCSI Boot Profile.**

At this point, the iSCSI Boot Profile is configured. However, for the server to receive the iSCSI Boot information, it must have a bootable vNIC configured on the server, which occurs by configuring an I/O Profile from the I/O Template that contains the bootable vNIC, then connecting the I/O Profile to the server. For information about I/O Profiles, see [“Working With I/O Profiles” on page 259](#).

Also, the server's BIOS boot order must be edited to insert the Oracle HCA as the highest priority boot device in the server's boot priority list. See [“SAN Boot Profiles” on page 333](#).

7. Proceed to [“Creating the Bootable I/O Template” on page 363](#).

Creating the Bootable I/O Template

When the iSCSI Boot Profile is created, you must associate an I/O Template with it. The I/O Template provides the vNIC that connects the server to the LUN on the iSCSI storage array where the server's boot information (kernel, boot image, and so on) are located. This section documents how to create an I/O Template with a bootable vNIC. For more information about I/O Templates, see [“Working With I/O Templates” on page 205](#).

When you create the I/O Template for the iSCSI Boot Profile, make sure that you select the bootable field. Without this setting, the vNIC in the I/O Template will not be able to support iSCSI Boot functionality. When the bootable option is set, it does not preclude the vNIC from carrying standard network traffic. Instead, the bootable option allows the vNIC to support iSCSI Booting the server in addition to standard Ethernet and Gigabit Ethernet network traffic.

Oracle Fabric Manager supports iSCSI Boot functionality through one or two paths, which are configured through the I/O Template.

- [“Create a Single-Path Bootable I/O Template” on page 342](#)
- [“Create a Dual-Path Bootable I/O Template” on page 346](#)

▼ Create a Single Path Bootable I/O Template

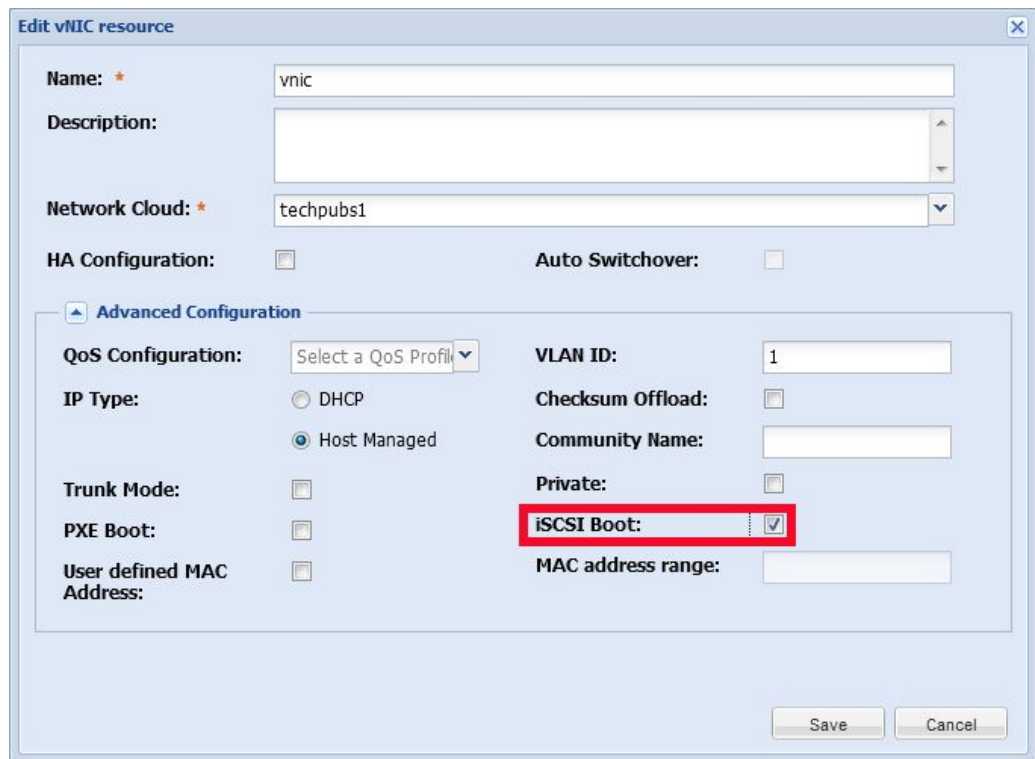
To create the bootable I/O Template, follow this procedure:

1. **Select Server Resource Manager → I/O Templates to display the I/O Template summary, then click the plus sign (+) to display the I/O Template Editor.**
2. **In the Name field, enter the name of the I/O Template that will be used for SAN Booting.**
3. **From the iSCSI Boot Profile dropdown menu, select the iSCSI Boot Profile you just created in the preceding section.**

Without creating the I/O Template as a bootable I/O Template, the server where the vNIC is deployed cannot be iSCSI Booted.

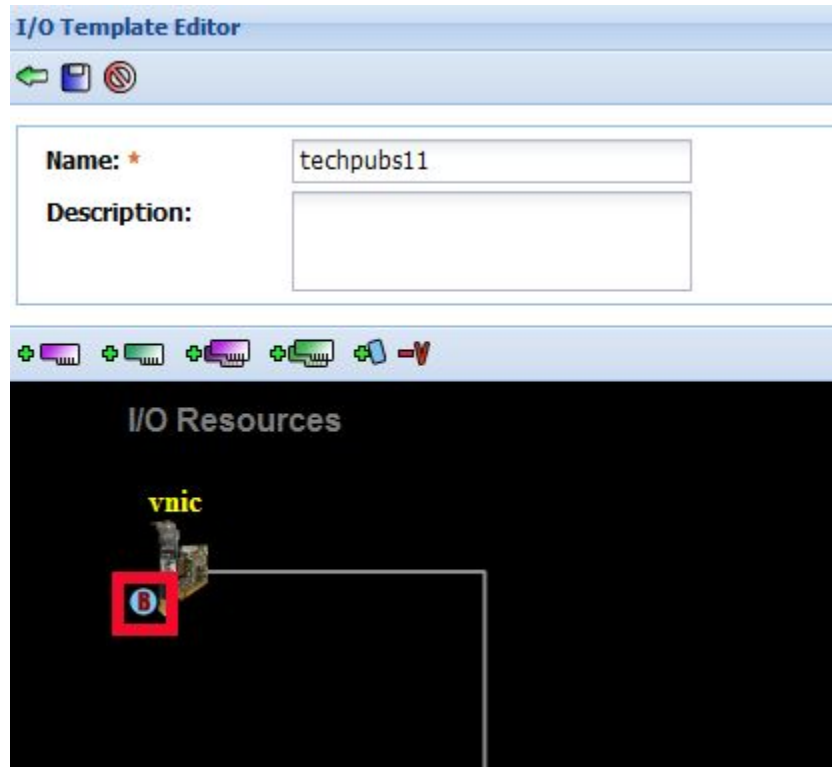
4. **Click and drag the vNIC to the appropriate Network Cloud to terminate the vNIC on an Ethernet port.**

5. Double click the vNIC icon on the I/O Template work space to display the Edit vNIC Resource dialog.
6. On the Edit vNIC Resource dialog, click the Advanced Configuration button to display additional options for the vNIC—including the iSCSI Boot checkbox.



7. Click the iSCSI Boot checkbox.
The checkbox must contain a check mark for iSCSI Booting capabilities to be supported on the vNIC.
8. Click Save to return to the I/O Template Editor.
9. Check the I/O Template Editor for the bootable icon (a circle with a red "B" within) associated with the server's boot vNIC.

The presence of this icon indicates that the vNIC connected to the host is bootable. The following figure shows a bootable vNIC.



If the “bootable” icon is not present, make sure the iSCSI Boot checkbox contains a checkmark.

10. **On the I/O Template Editor toolbar, click Save.**
Without saving the I/O Template, the pending I/O Template configuration will be lost.
11. **After the bootable I/O Template is configured, proceed to [“Create an I/O Profile From the I/O Template” on page 372.](#)**

▼ Create a Dual Path Bootable I/O Template

Dual path provides two paths to the same LUN so that a single point of failure through the Fabric Device's fabric is eliminated and availability of the iSCSI Boot information is increased. If one path is not available, the second one is.

Be aware that the two paths in dual pathing are not true HA, so there is no automatic failover or failback if one of the paths is not available. Instead, dual pathing simply provides two connections for the server to iSCSI Boot. If one path is not available, user intervention is required to select the other path through host-side methods, such as multipathing software.

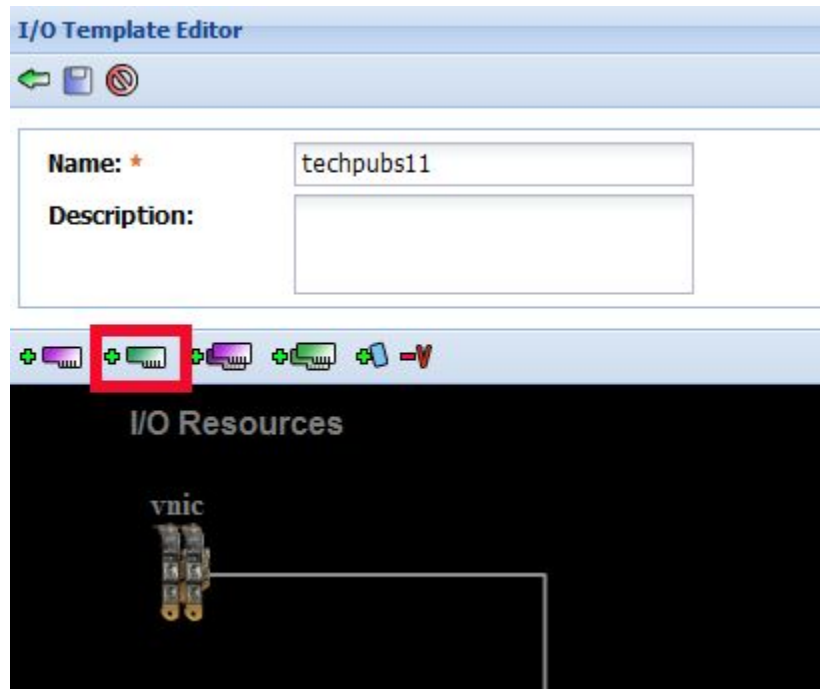
For dual path bootable I/O Templates, two vNICs are configured with the following considerations:

- On the I/O Template Editor, you must use the HA vNICs option instead of creating two single vNICs.
- The two vNICs must connect to the same LUN. You cannot put the server's iSCSI Boot information on multiple LUNs. Since the two vNICs will be connected to the same LUN, you will terminate the two vNICs in the same Network Cloud.

To create a dual path bootable I/O Template, follow this procedure:

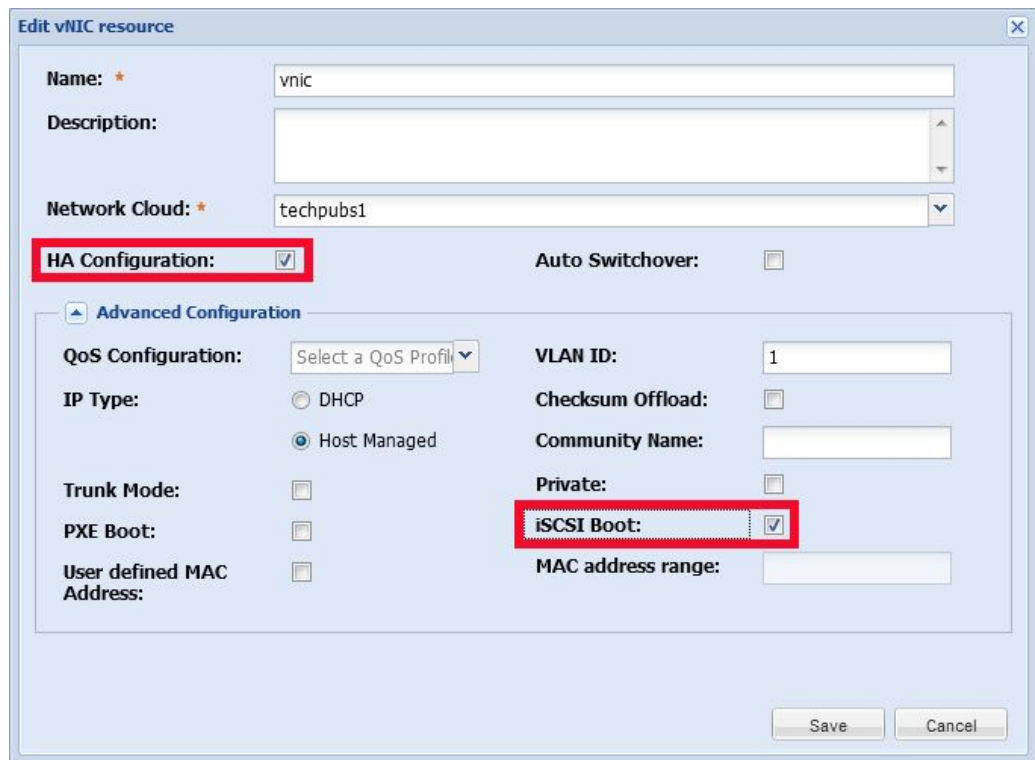
1. **Select Server Resource Manager → I/O Templates to display the I/O Template summary, then click the plus sign (+) to display the I/O Template Editor.**
2. **Click the Add an HA vNIC to Template button, then click and drag to connect the vNIC to the appropriate Network Cloud.**

Note - To terminate an HA vNIC, the Network Cloud must have more than one Gigabit Ethernet port in it.



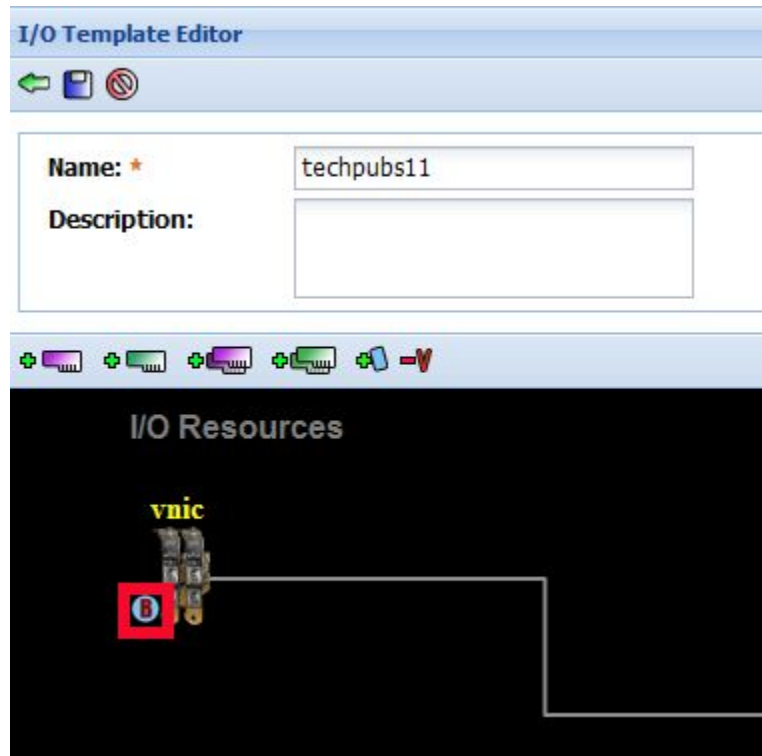
3. In the Name field, enter the name of the I/O Template that will be used for iSCSI Booting.
4. From the iSCSI Boot Profile dropdown menu, make sure to select the iSCSI Boot Profile you just created in the preceding section.
Without creating the I/O Template as a bootable I/O Template, the server where the vNICs are deployed cannot be iSCSI Booted.
5. Click and drag the vNICs to the appropriate Network Cloud to terminate the vNICs on Ethernet ports.
6. Double click the vNIC icon on the I/O Template work space to display the Edit vNIC Resource dialog.

7. Click the **Advanced Configuration** checkbox to display the additional configuration options for the HA vNIC—including the iSCSI Boot checkbox.



8. On the Edit vNIC Resource dialog, make sure that the HA Configuration checkbox contains a checkmark.
If it does not, click it now. Without this checkbox properly set, only one vNIC (and therefore one path) will be supported for iSCSI Booting the server.
9. Click the iSCSI Boot checkbox.
This checkbox must contain a checkmark for the HA vNICs to support iSCSI boot of the server where the vNICs are deployed.
10. Click Save to return to the I/O Template Editor.
11. Check the I/O Template Editor for the bootable icon (a circle with a red “B” within) associated with the server's boot vNIC.

The presence of this icon indicates that the vNIC connected to the host is bootable. The following figure shows a bootable vNIC.



12. **On the I/O Template Editor, click Save.**
Without saving the I/O Template, all pending I/O Template configuration will be lost.
13. **After the bootable I/O Template is configured, proceed to [“Create an I/O Profile From the I/O Template” on page 372.](#)**

iSCSI Boot Profiles Summary

The iSCSI Boot Profiles are displayed in the iSCSI Boot Profile Summary, which contains information about all the iSCSI Boot Profiles that have been created through Oracle Fabric Manager. The table shows general information about the iSCSI Boot Profile such as the IP address of the device containing the server's boot information and the mount type that will be used to access the boot information, kernel, and initrd.

The following figure shows the iSCSI Boot Profile Summary, which is available through the Server Resource Manager->Boot Profiles->iSCSI Boot Profile Summary.

SAN Boot Profile Summary		iSCSI Boot Profile Summary								
Name	Target IP Address	Target Portal Group	Port	Proto...	Mount Ty...	Group Name	Volume Name	Mount Device	IOTemplates	IOProfiles
iscsiboot_1	172.16.10.70		3260	6	direct				0	0
iscsiboot_10	192.168.8.108		3260	6	direct				0	0
iscsiboot_11	172.16.1.10		3260	6	direct				0	0
iscsiboot_12	172.16.2.10		3260	6	direct				0	0
iscsiboot_13	1.1.1.4		3260	6	direct				0	0
iscsiboot_14	1.1.2.4		3260	6	direct				0	0
iscsiboot_15	3.3.3.3		3260	6	static				0	0
iscsiboot_2	172.16.10.80		3260	6	direct				0	0
iscsiboot_3	192.168.46.3		3260	6	direct				0	0
iscsiboot_4	172.16.1.4		3260	6	direct				1	0

Field	Indicates
Name	The name of the iSCSI Boot Profile.
Target IP Address	Shows the IP address of the storage target filer or array where the server's iSCSI Boot information is located.
Target Portal Group	Shows the name of the target portal group (if any) in which the server's iSCSI Boot information is located.
Port	Shows the port number on which the iSCSI communication between the server and the target containing the iSCSI boot information.
Protocol	Shows the protocol number used for communication between the server and the target containing its iSCSI boot information.
Mount Type	The type of mount configured on the iSCSI Boot vNIC. Valid values are: <ul style="list-style-type: none"> ■ Static, for a Static iSCSI Boot Profile. ■ LVM, for the Logical Volume Manager. ■ Direct, for assigning the device name that has the root file system.
Group Name	The logical volume manager (LVM) group name. If the mount type is LVM, this field displays the group name. If the mount type is other than LVM, no value is displayed.
Volume Name	The name of the logical volume that LVM uses for the root file system mount point.
Mount Device	The name of the device that contains the root file system mount point for the iSCSI Boot Profile.
I/O Template	The number of I/O Templates that are assigned to each iSCSI Boot Profile.
I/O Profiles	The number of I/O Profiles that are associated with each iSCSI Boot Profile.

iSCSI Boot Profile Details

The iSCSI Boot Profile Details frame contains information about any configured iSCSI Boot Profiles. The iSCSI Boot Profile Details frame contains only one instance of an iSCSI Boot

Profile, which is determined by the iSCSI Boot Profile that you select in the iSCSI Boot Profile Summary. By clicking a specific iSCSI Boot Profile in the summary, all of its corresponding details are available in the details frame.

Through the iSCSI Details Frame you can also edit a configured iSCSI Boot Profile to make changes to the profile without having to delete then re-add the profile. Be aware that:

- Changes to the profile might also require resetting the server, or making changes to the server's configuration (for example, if a new boot device is specified, you might need to change the order of boot devices in the server's boot priorities list).
- Editing an iSCSI Boot Profile is allowed only when the iSCSI Boot Profile is not currently associated with an I/O Template.

The iSCSI Boot Profile Details Frame is in Oracle Fabric Manager GUI just below the iSCSI Boot Summary Profile.

The screenshot displays the Oracle Fabric Manager GUI. At the top, there are two tabs: 'SAN Boot Profile Summary' and 'iSCSI Boot Profile Summary'. Below the tabs is a table listing various iSCSI boot profiles. The table has columns for Name, Target IP Address, Target Portal Group, Port, Protocol, Mount Type, Group Name, Volume Name, Mount Device, IO Templates, and IO Profiles. The profile 'iscsiboot_19' is highlighted in blue. Below the table, there is a section titled 'iSCSI Boot Profile : iscsiboot_19' with a 'General' tab. This section contains the following details:

Name:	iscsiboot_19
Target IP Address:	192.168.111.112
Target Portal Group:	
Port:	3260
Protocol:	6
Mount Type:	lvm
Group Name:	VolGroup01
Volume Name:	LogVol01
Mount Device:	

At the bottom of the details section, there is an 'Edit' button.

When the iSCSI Boot Details frame is populated with an iSCSI Boot Profile, you can edit the profile to make changes to the current profile without having to delete and re-add the profile.

The following table shows the contents of the iSCSI Boot Profile Details frame, and explains what each field means. Not all fields will be populated based on the options specified during the configuration of the iSCSI Boot Profile. For example, if you created a direct connect iSCSI Boot Profile, none of the fields for logical volume manager will contain information.

Field	Means
Name	The name assigned to the selected iSCSI Boot Profile.
Target IP Address	Shows the IP address of the storage target filer or array where the server's iSCSI Boot information is located.
Target Portal Group	Shows the name of the target portal group (if any) in which the server's iSCSI Boot information is located.
Port	Shows the port number on which the iSCSI communication between the server and the target containing the iSCSI boot information.
Protocol	Shows the protocol number used for communication between the server and the target containing its iSCSI boot information.
Mount Type	The method used to mount the SAN location of the server's iSCSI Boot information: <ul style="list-style-type: none">■ Direct■ Static■ Logical Volume Manager
Group Name	For Logical Volume Manager mount types, this field shows the logical group that contains the SAN location of the server's iSCSI boot information.
Volume Name	For Logical Volume Manager mount types, this field shows the logical volume that contains the SAN location of the server's iSCSI boot information.
Mount Device	For Direct mount types, this field shows the device that contains the root file system mount point for the server's iSCSI Boot information.
I/O Template	The number of I/O Templates that are assigned to each iSCSI Boot Profile.
I/O Profiles	The number of I/O Profiles that are associated with each iSCSI Boot Profile.

Create an I/O Profile From the I/O Template

If you are setting up iSCSI Boot for a server that has already booted up and connected to the Oracle Fabric Device, you can now create an I/O Profile from the I/O Template. When the I/O Profile is created, it will contain the bootable vNIC that you created. You can create multiple I/O Profiles from the same I/O Template if needed.

For information about creating an I/O Profile, follow the procedure for creating an I/O Profile as documented in [“Working With I/O Profiles” on page 259](#).

Connect the I/O Profile to a Physical Server

When you assign the I/O Profile to the server, you begin pushing the bootable vNIC and the associated iSCSI Boot Profile to the server.

While Oracle Fabric Manager is pushing the iSCSI bootable vNIC and iSCSI Boot Profile to hosts, it prompts you for IQN information to complete the path between the host and the LUN containing the host's iSCSI boot information. When needed, Oracle Fabric Manager prompts you for the necessary information through a series of popup dialogs. IQN information is required for either a single path or dual path iSCSI configuration. Answer the prompts as needed to complete assigning the I/O Profile to the relevant servers. To assign the I/O Profile to a server, follow the procedure documented in [“Connecting an I/O Profile to a Server” on page 266](#).

▼ Delete an iSCSI Boot Profile

Through the iSCSI Boot Profile Summary, you can delete an iSCSI Boot Profile. Any iSCSI Boot Profile that is not associated with an I/O Template or a physical server can be deleted.

To delete an iSCSI Boot Profile, follow this procedure:

1. **On the navigation panel, click Server Resource Manager → Boot Profiles → iSCSI Boot Profile Summary tab.**

This iSCSI Boot Profile Summary is displayed.

2. **On the iSCSI Boot Profile Summary, select the iSCSI Boot Profile that you want to delete.**

You can select multiple iSCSI Boot profiles through standard keyboard conventions, for example by pressing the CTRL + click or Shift + click.

3. **Click the garbage can icon to delete the selected iSCSI Boot Profile.**
4. **Click OK on the confirmation window to delete the selected iSCSI Boot Profile(s), or Cancel to abort the deletion.**

Working With the Virtual Topology

This chapter contains the following topics:

- [“Virtual Topology” on page 375](#)
- [“Displaying the Topology Overview” on page 375](#)
- [“Displaying the Server Cloud View” on page 379](#)
- [“Displaying the Fabric Device Cloud View” on page 382](#)
- [“Displaying the Server Fabric Device View” on page 385](#)
- [“Displaying the Target Topology View” on page 388](#)

Virtual Topology

The Topology is a series of network illustrations of connectivity from different perspectives. Each perspective is called a “view.” When you are displaying the Topology, you are displaying how different parts of the datacenter that Oracle Fabric Manager is managing are connected together over Oracle virtual I/O and other products.

The Topology is different than the Physical Topology, which displays a diagram of how the Oracle Fabric Device and physical hosts are physically connected (cabled) together.

In addition to the top-level information available on the Topology, you can drill-down to lower, more detailed levels of information about the contents of the Topology. The detailed information is displayed based on the context of the view you are in. For example, if you are in Server Fabric Device view, details displayed will be in the context of the servers' connections to one or more Fabric Interconnects or Oracle SDN Controllers including any connections through intermediary switches. If you are in Server Cloud view, the detailed information will be displayed in the context of the servers' connections to clouds.

You can display detailed information by double-clicking individual elements (servers, clouds, Fabric Interconnects or Oracle SDN Controllers, and so on) within each view. Additional information is documented in the sections for each view.

Displaying the Topology Overview

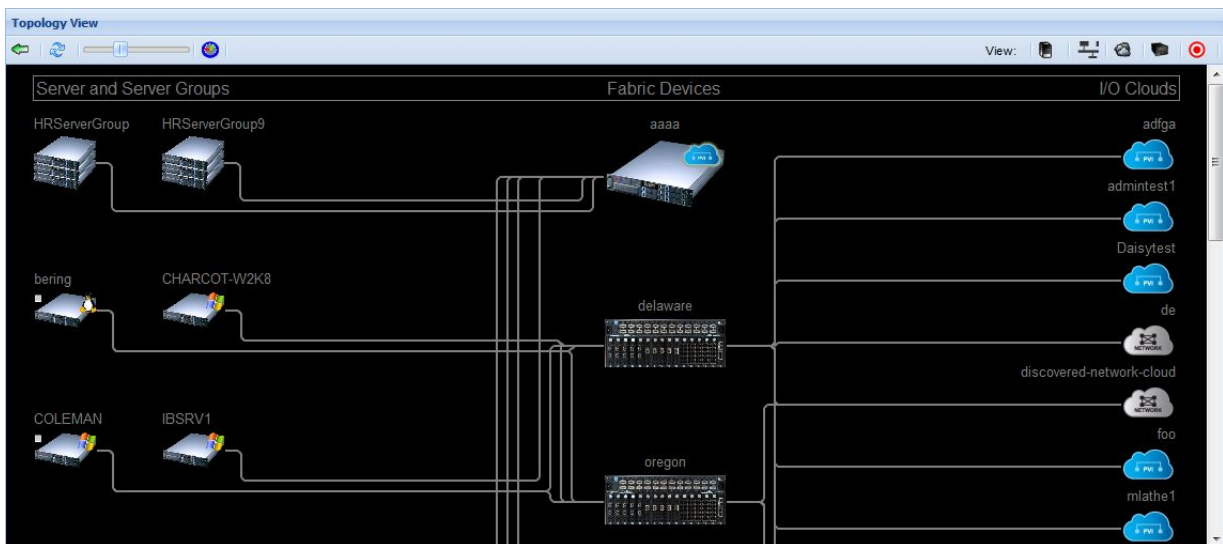
This sections contains the following topics:

- [“Topology Overview” on page 376](#)
- [“Scroll Columns Vertically” on page 378](#)
- [“Display the Physical Server Details” on page 378](#)
- [“Display the Server Group Details” on page 378](#)
- [“Display Physical Connectivity for the Topology Overview” on page 378](#)

Topology Overview

The general topology is displayed through the Topology Overview, and shows the virtual connections (vNICs and vHBAs) for each Network and Storage Cloud to all the connected hosts. In the Topology Overview, toolbar controls are available for:

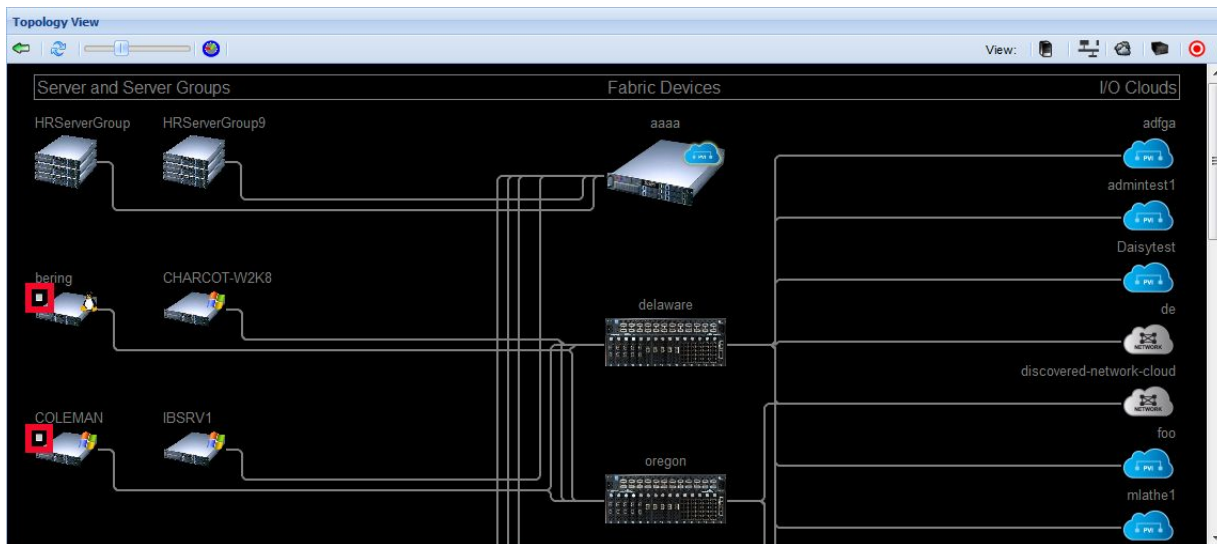
- [“Performance Meters” on page 390](#)



Through the toolbar on the Topology Overview, you can control various aspects of how the Topology Overview contents are displayed:

- A button toggles the display of performance meters for each server in the Topology Overview. (See [“Performance Meters” on page 390.](#))
- A slider controls zooming in and out on the Topology. Zooming in enables you to focus in on a smaller area of the Topology and enlarge that area. Zooming out enables you to pan out to a wider view of the Topology and reduce the size of the individual elements in the Topology.

- The right side of the toolbar contains controls for displaying different views of the Topology page, including:
 - The Topology view, which shows the overall topology from the Network and Storage Cloud through the Fabric Interconnect or Oracle SDN Controller to the discovered host servers.
 - The Server Cloud view, which shows connectivity between hosts and clouds. The Fabric Interconnects or Oracle SDN Controllers providing connectivity are not shown.
 - The Fabric Device Cloud view, which shows connectivity between the clouds and the Fabric Interconnects or Oracle SDN Controllers that are providing the connections. Host servers are not displayed.
 - The Server Fabric Device view, which shows connectivity between the physical servers and server groups and the Fabric Interconnects or Oracle SDN Controllers. The connectivity to the network and storage clouds is not displayed.
 - The Target Topology view, which shows the connectivity of individual vHBAs and their respective targets. The Fabric Interconnects or Oracle SDN Controllers and hosts are not displayed.
- The main area of the Topology also shows whether or not a Server has an I/O Profile bound to it. A small icon to the upper left of each server or server group indicates whether the server or group has Oracle I/O bound to it or not. The following figure shows which servers have I/O Profiles bound.



▼ Scroll Columns Vertically

In the main work area of each view in the Topology, elements (such as servers, Fabric Interconnect or Oracle SDN Controllers, and Clouds) are arranged in vertical columns. Column headings on the Topology indicate the individual columns. Usually, two or three columns are displayed on the work area. In a larger deployment, these elements can scroll off the Topology, and it can become difficult to see all the elements of the topology on one page even if you zoom out to the widest display. Each vertical column on the Topology supports a vertical click and drag to scroll that column up and down.

To click and drag a column, follow this procedure:

1. **Hover with the mouse near the element in that vertical column until the mouse pointer switches to a hand icon.**
2. **When the hand icon appears, you can click with the mouse and drag the column up or down.**

Each column is individually controllable, so you can arrange each column until you get all elements arranged so that they show the information you need.

Display the Physical Server Details

Through the Topology overview, you can get additional information about server connections to clouds by double-clicking a server. For information, see [“Individual Server Topology” on page 381](#).

Display the Server Group Details

Through the Topology overview, you can get additional information about the connections from a server group to clouds by double-clicking a server group. For information, see [“Server Group Topology” on page 380](#).

Display Physical Connectivity for the Topology Overview

Through the Topology Details, you can display physical connectivity information for the Fabric Interconnect or Oracle SDN Controller physical ports that connect vNICs and vHBAs to servers. The information displayed shows the same connectivity as the Server Cloud view. For more information, see [“Server Cloud Physical Connections Topology” on page 382](#).

Displaying the Server Cloud View

This sections contains the following topics:

- [“Server Cloud View” on page 379](#)
- [“Server Group Topology” on page 380](#)
- [“Individual Server Topology” on page 381](#)
- [“Server Cloud Physical Connections Topology” on page 382](#)

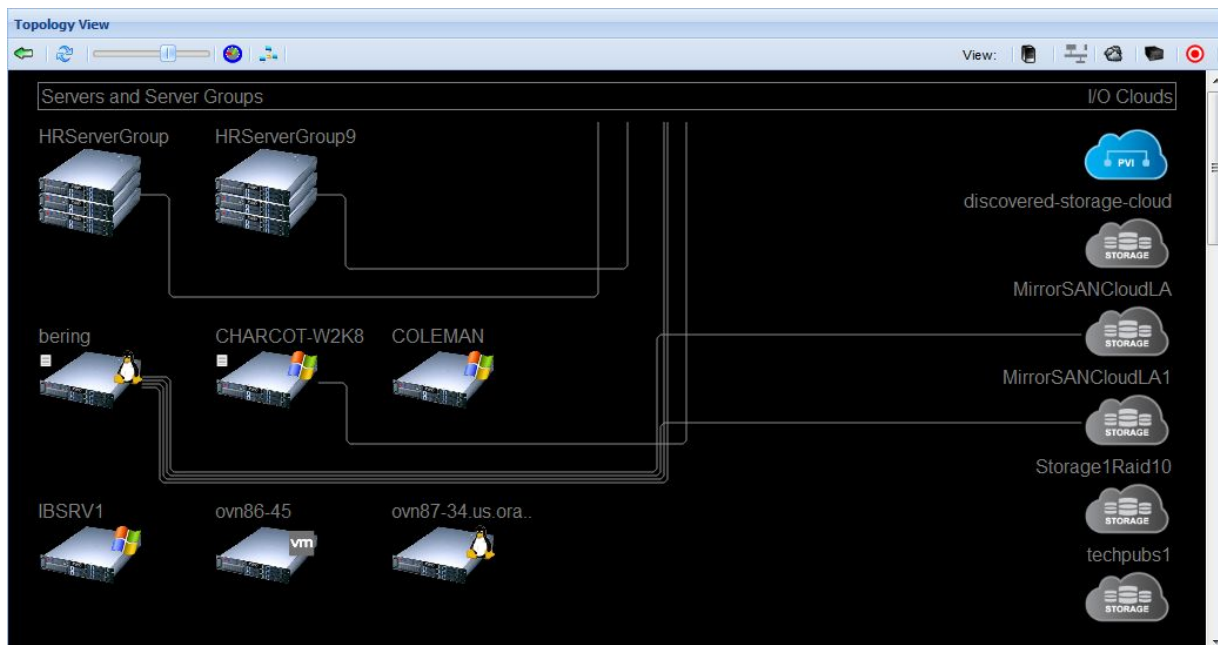
Server Cloud View

The Server Cloud view shows information about which server(s) are connected to which Network/PVI and Storage Clouds. This view shows individual servers as well as server groups.

In Server Fabric Device view, two options are available as toolbar buttons:

- [“Performance Meters” on page 390](#)
- [“Server Cloud Physical Connections Topology” on page 382](#)

The following figure shows an example of the Server Cloud view.

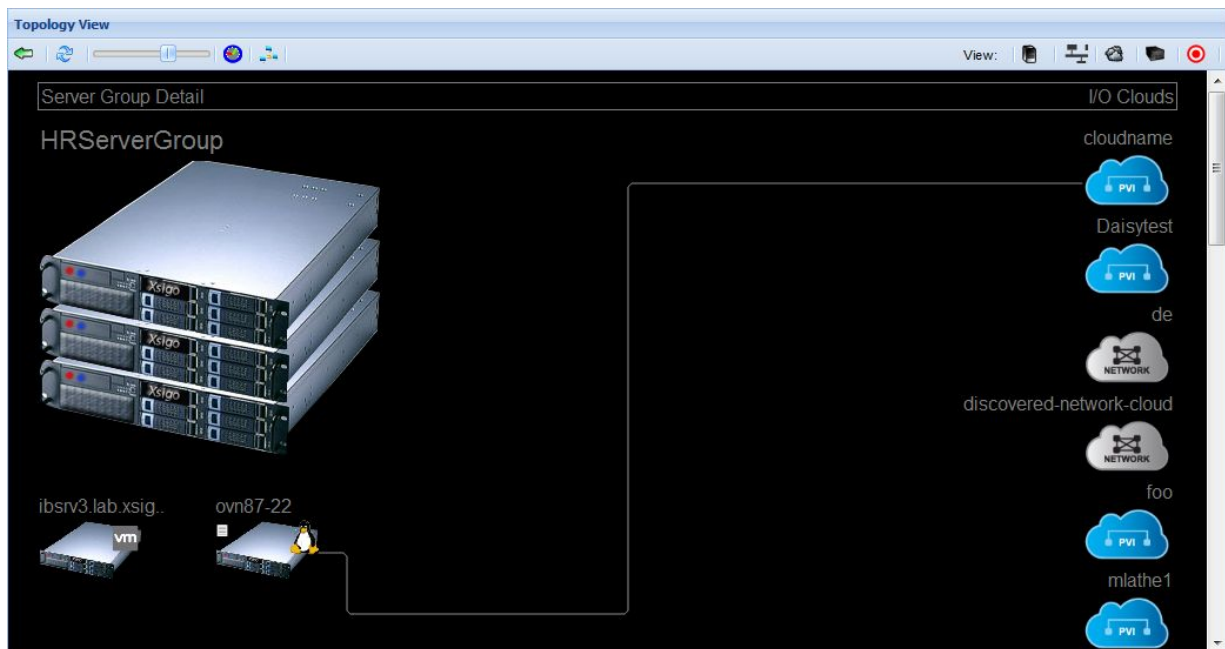


Server Group Topology

A server group is logical collection of individual servers that you have created. (For information see, [“Working With Server Groups” on page 325.](#)) The Topology supports displaying information about server groups, including:

- Which servers constitute the server group
- Whether servers in the server group are bound to an I/O Profile
- How individual servers in the server group are connected to individual clouds

To display the Server Group details, double-click a server group on the Topology from either Server Cloud view, the Topology Overview, or Server Fabric Device view.



In this example, you can see that the Server Group “HRServerGroup” consists of the servers ibsrv3 and ovn87-22. One server is bound to an I/O Profile (as indicated by the icon to the upper left of the server), and one server is connected to a PVI cloud.

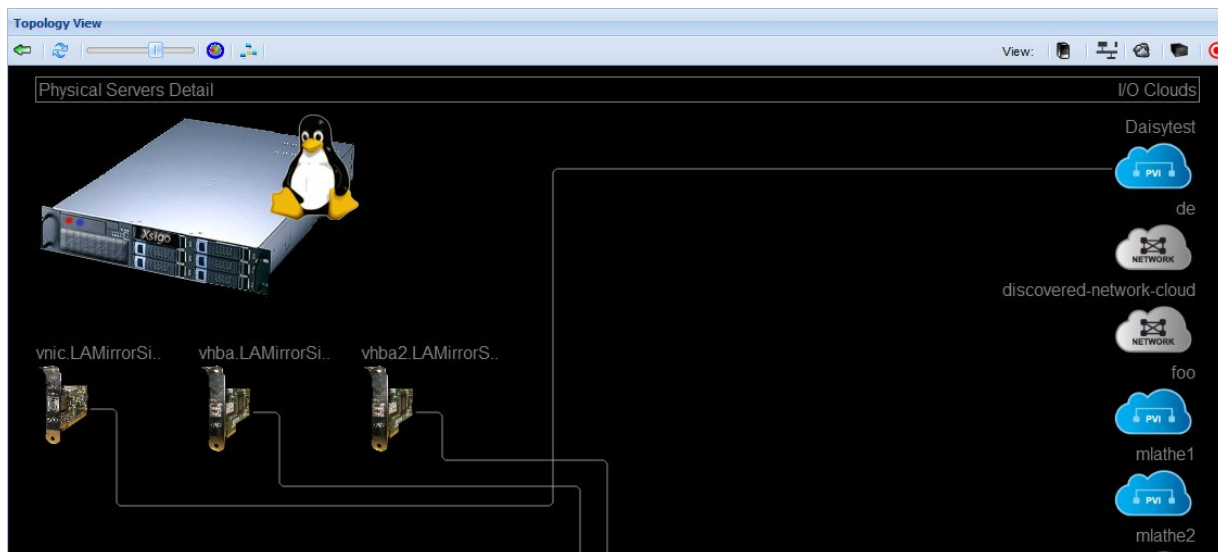
Individual Server Topology

Detailed information is available for individual servers that are managed by Oracle Fabric Manager. Through the Topology, you can see the following detailed information for individual servers:

- The individual server
- Individual vNICs and vHBA connections from the server
- How individual vNICs and vHBAs on the server are connected to individual clouds

Note - Additional details about physical servers are available based on the context of the view you are in. For example, if you are in Server Fabric Device view, details displayed will be in the context of the servers' connections to one or more Fabric Interconnects or Oracle SDN Controllers including any connections through intermediary switches.

To display the details for individual servers, double-click a server group on the Topology from either Server Cloud view, the Topology Overview, or the Server Fabric device view.

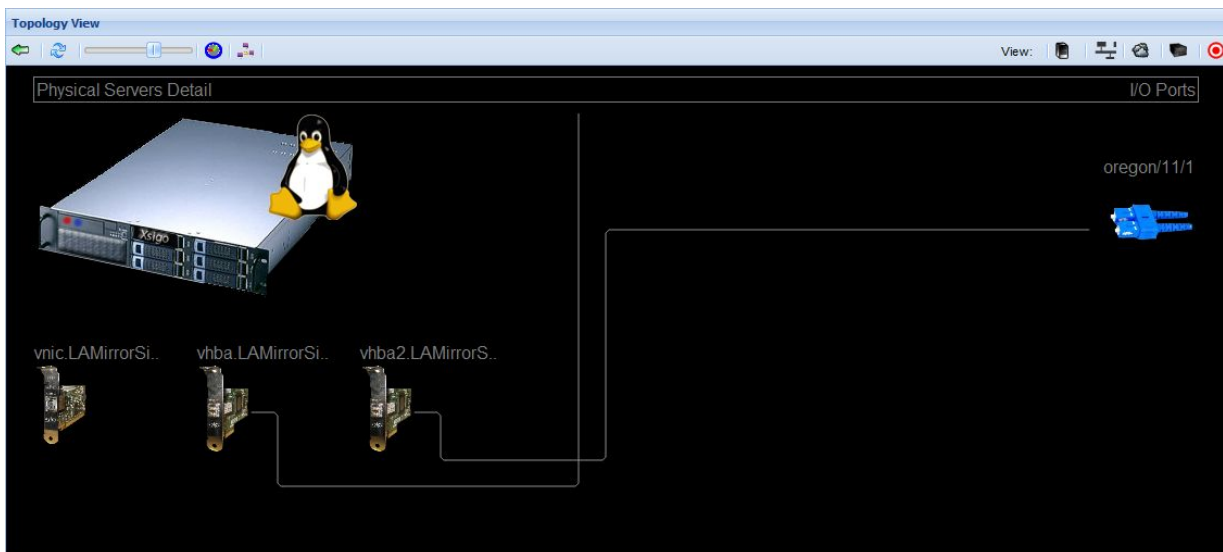


In this example, the server has a vNIC and multiple vHBAs. In this example, the vNIC is shown connected to a PVI cloud, and the vHBAs are connected to Storage Clouds (not displayed, but connections to it run off the bottom of the Physical Server Details page).

Server Cloud Physical Connections Topology

The Server Cloud Details contains an option to show physical connection information, such as what the server(s) in the server group are connected to on one or more Fabric Interconnects or Oracle SDN Controllers. The Show Physical View toolbar button is supported on the toolbar for the Server Cloud view and the Server Cloud Details. Different information is displayed based on which view is used when the Show Physical View toolbar button is clicked.

- On the Server Cloud view, the Show Physical View button shows the same information as the Server Fabric Device view. For information, see [“Displaying the Server Cloud View” on page 379](#).
- On the Server Cloud Details page, you can see which Fabric Interconnect or Oracle SDN Controller port each server in the group is connected to.



In this example, you can see the individual vHBAs in the server. One of these vHBAs are connected to physical port 11/1 Fabric Device “oregon”.

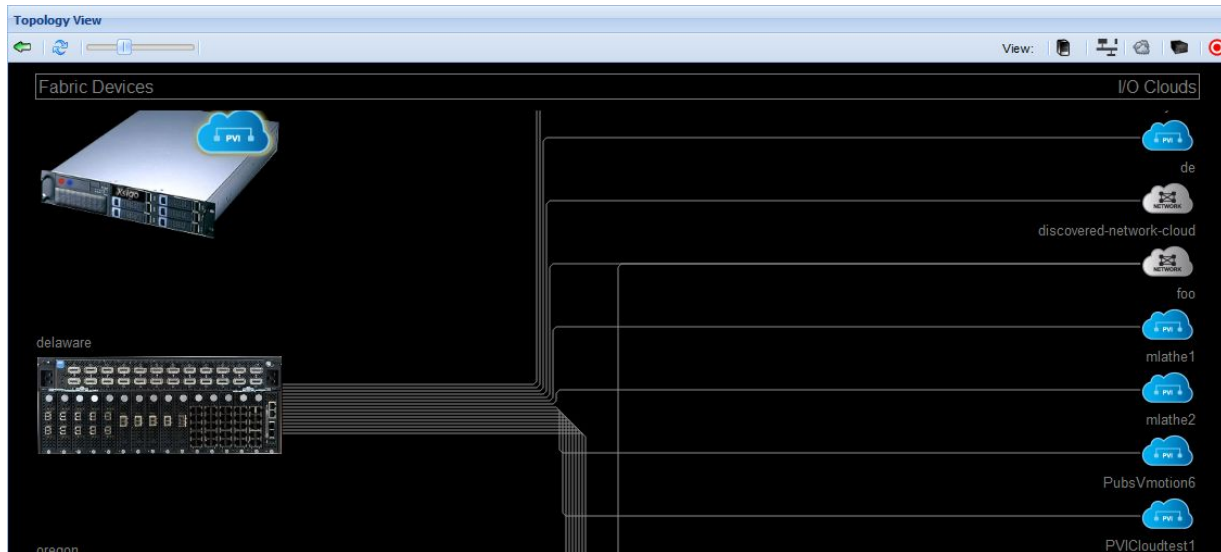
Displaying the Fabric Device Cloud View

This sections contains the following topics:

- [“Fabric Device Details” on page 383](#)
- [“Fabric Device Cloud Physical Connections Topology” on page 384](#)

Fabric Device Cloud View

The Fabric Device Cloud view shows information about which Fabric Interconnect(s) or Oracle SDN Controller(s) are connected to which Network and Storage Clouds.




The Fabric Device Cloud view shows which Fabric Interconnects or Oracle SDN Controllers are connected to which clouds. In this example, the Fabric Interconnect “delaware” is connected to multiple clouds, some of which are off the current topology page. This example also shows an Oracle SDN Controller is present in Oracle Fabric Manager.

Fabric Device Details

Detailed information about the Fabric Interconnects or Oracle SDN Controllers in the Topology are available. Through the Fabric Device Details, you can view information about a selected Fabric Interconnect or Oracle SDN Controller, including:

- The number and type of I/O Modules in the Fabric Device. All installed I/O Modules are displayed regardless of whether or not they are currently supporting a vNIC or vHBA connection.
- The connection from each I/O Module to the cloud where the module's vNIC or vHBA is terminated. Only modules that are currently supporting a vNIC or vHBA connection are displayed.
- The cloud(s) to which the I/O Module is connected through the listed I/O Module(s).

To display the Fabric Device details, double-click a Fabric Interconnect or Oracle SDN Controller on the Fabric Device Cloud view. The following figure shows the Fabric Device details.



In this example, the Fabric Interconnect “delaware” is displayed. This Fabric Interconnect has multiple network and storage I/O modules:

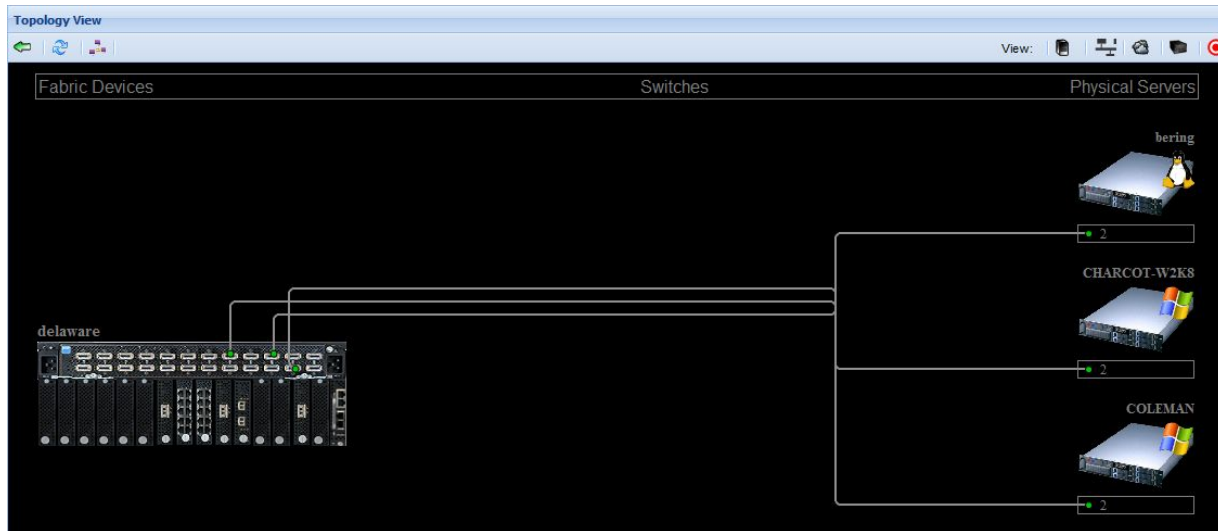
- One 8 Gbps line-rate Fibre Channel module in slot 11 is connected to Storage Clouds (not shown)
- Three 10 Gbps Ethernet I/O Modules in slots 10, 14, and 4 are connected to various Network Clouds
- Two 4-Port 10 Gbps Ethernet modules exist in slots 8 and 9, and are connected to various Network Clouds

By hovering over the module or a cloud, Oracle Fabric Manager highlights the connection between the module and a cloud.

Fabric Device Cloud Physical Connections Topology

Through the Topology, you can display a diagram of physical connections between the Fabric Interconnect or Oracle SDN Controllers and physical servers that Oracle Fabric Manager is

managing. This physical connection diagram also includes intervening InfiniBand switches if any are present. Through the Show Physical View toolbar button on the Fabric Device Cloud Details view, you can display physical connections.



In this example, all physical connections from the Fabric Interconnect server ports to individual servers are displayed. The diagram shows each physical port's connection directly to a server. If an intervening InfiniBand switch is present, it will also be displayed along with the ports on the IB switch where connections are made. By hovering over Fabric Interconnect or Oracle SDN Controllers, Fabric Interconnect or Oracle SDN Controller ports, InfiniBand ports, or servers, you can trace either parts of the physical connection, or the entire end-to-end path.

Displaying the Server Fabric Device View

This sections contains the following topics:

- [“Server Fabric Device View” on page 385](#)
- [“Detailed Connectivity Topology” on page 386](#)
- [“Server Fabric Device Virtual Connections Topology” on page 387](#)

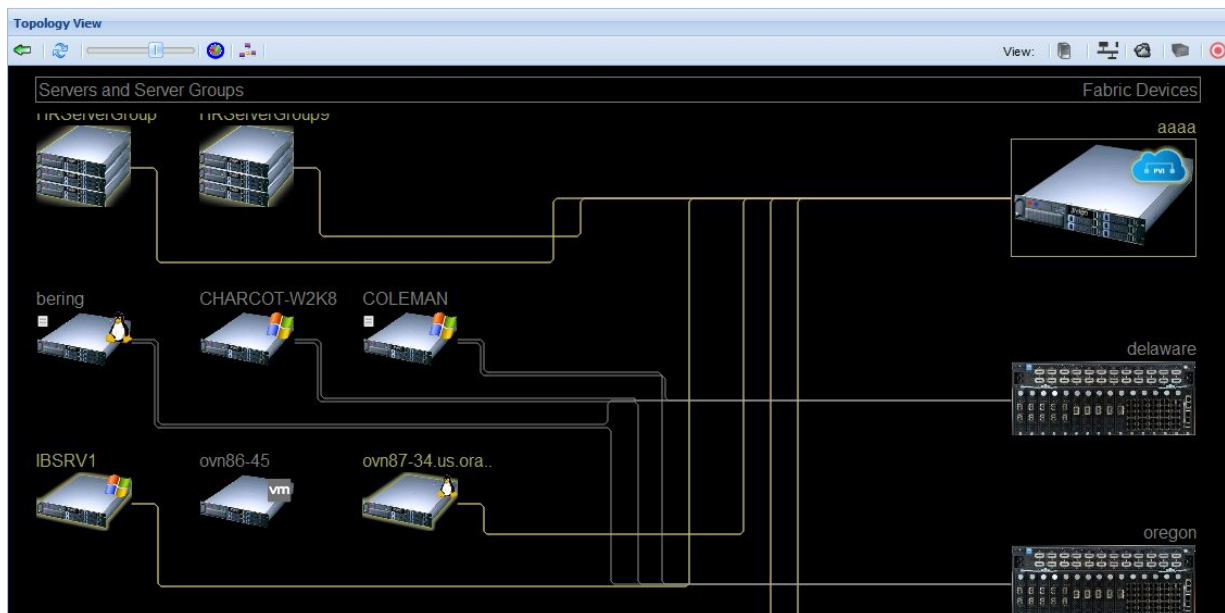
Server Fabric Device View

The Server Fabric Device view shows information about which physical servers are connected to which Fabric Interconnects or Oracle SDN Controller(s). If a Fabric Interconnect or Oracle

SDN Controller is not displayed in this view, it is not managed by Oracle Fabric Manager. Also, if a Fabric Interconnect or Oracle SDN Controller is not connected to any server, it will still be displayed in the Topology, but it will appear unconnected (no lines leading to other equipment) in the Server Fabric Device view. Since a server must contain a server profile and at least one vNIC or vHBA, you can use this display to determine if a server has been deployed with virtual I/O.

In Server Fabric Device view, two options are available as toolbar buttons:

- [“Fabric Device Cloud Physical Connections Topology”](#) on page 384
- [“Performance Meters”](#) on page 390



The Server Fabric Device view shows the Fabric Devices and the host server that are connected to them. In this example, two Fabric Interconnects “delaware” and “oregon” are displayed. Also, an Oracle SDN Controller is present in Oracle Fabric Manager. By hovering over each of the Fabric Devices, connections to host servers are highlighted.

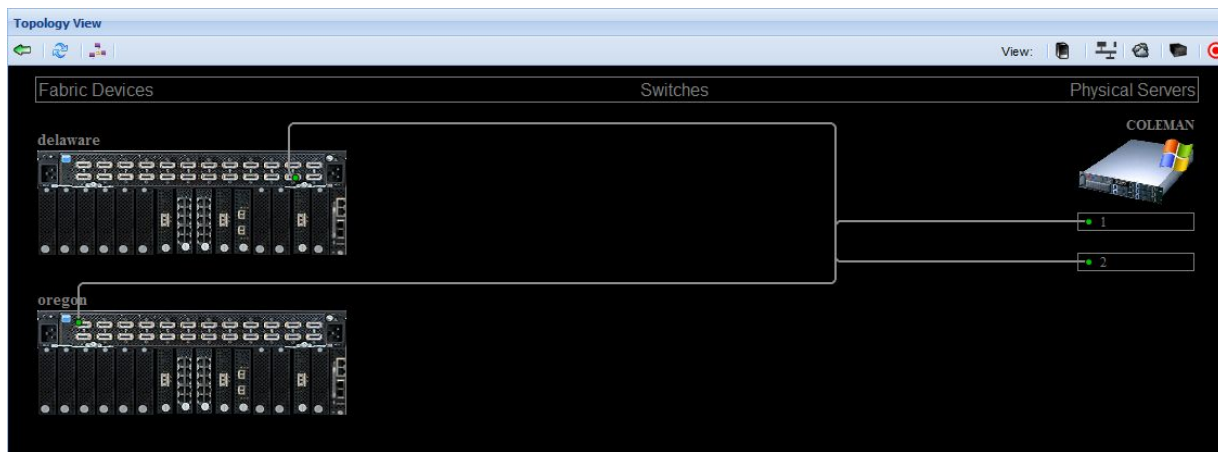
Detailed Connectivity Topology

Through the Server Fabric Device view you can see detailed connection diagrams for the Fabric Interconnect or Oracle SDN Controllers and servers in your data center. The detailed connectivity information shows a diagram of how the Fabric Interconnect or Oracle SDN Controller is cabled to servers.

The detailed information shows:

- A server or server groups that is connected to Fabric Interconnects or Oracle SDN Controllers, and the port number on which the server is connected
- Fabric Interconnects or Oracle SDN Controllers that are providing connection to the server or server group, and the IB server port on the Fabric Interconnect or Oracle SDN Controller that is connected to the server or server group
- IB switches (if any) that are connecting the servers, and the switch ports that are providing the connections.

You can display the detailed connectivity information by double-clicking a server in the Server Fabric Device view.



In this example, the server Coleman is connected to two Fabric Interconnects (“delaware” and “oregon”). The Fabric Interconnect “delaware” has its connection on server port 11 which is on HCA port 1 on the “Coleman.” The Fabric Interconnect “oregon” has its connection on server port 13, which is on HCA port 2 on “Coleman.”

Server Fabric Device Virtual Connections Topology

By default, physical connections are shown in the Server Fabric Device view. You can display the virtual connections between servers and their clouds, by clicking the Show Virtual View toolbar button. This button is available on either the top-level Server Fabric Device view or the Server Fabric Device Details:

- On the Server Fabric Device view, when you can click the Show Virtual View toolbar button, the virtual connections displayed are the same as the Fabric Device Cloud view. See [“Displaying the Fabric Device Cloud View” on page 382](#).
- On Server Fabric Device Details page, when you click the Show Virtual View toolbar button, the virtual connections displayed are the same as the Server Cloud view. See [“Displaying the Server Cloud View” on page 379](#).

Displaying the Target Topology View

This sections contains the following topics:

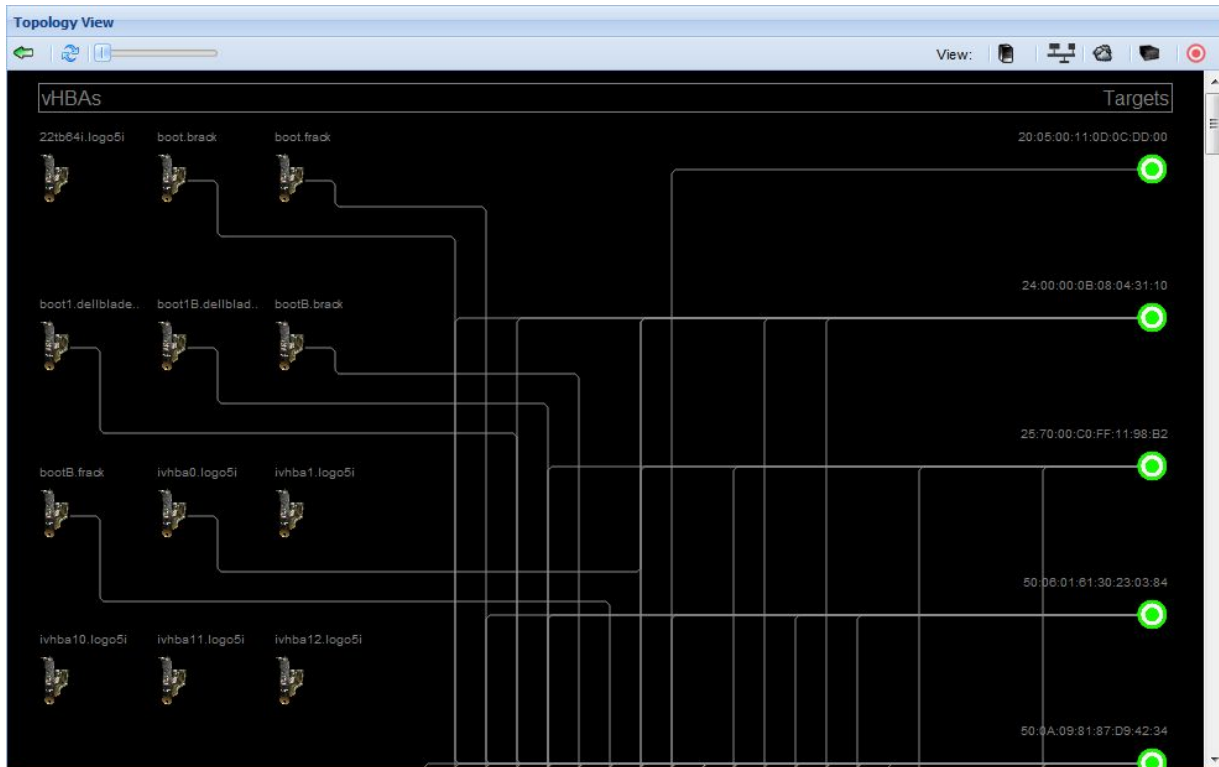
- [“Target Topology View” on page 388](#)
- [“vHBA-to-LUN Connections” on page 389](#)
- [“Performance Meters” on page 390](#)
- [“Physical and Virtual View Options” on page 392](#)

Target Topology View

The Target Topology view shows information about which storage targets (listed by WWN) are connected to which vHBAs. This view can be displayed by clicking the far right button in the View bar. The Target Topology button looks like a red target or bull's-eye.

Through the Target Topology view, you can determine if a vHBA is currently connected. The presence of a line between the target and a vHBA indicates that the vHBA is currently connected. If the line is absent, the vHBA is not currently connected to any storage target.

You can highlight the connection between a vHBA and its storage by mousing over either the storage target icon or the vHBA icon.



The Target Topology view shows the individual vHBAs managed by Oracle Fabric Manager, and the WWPN to which each vHBA is connected.

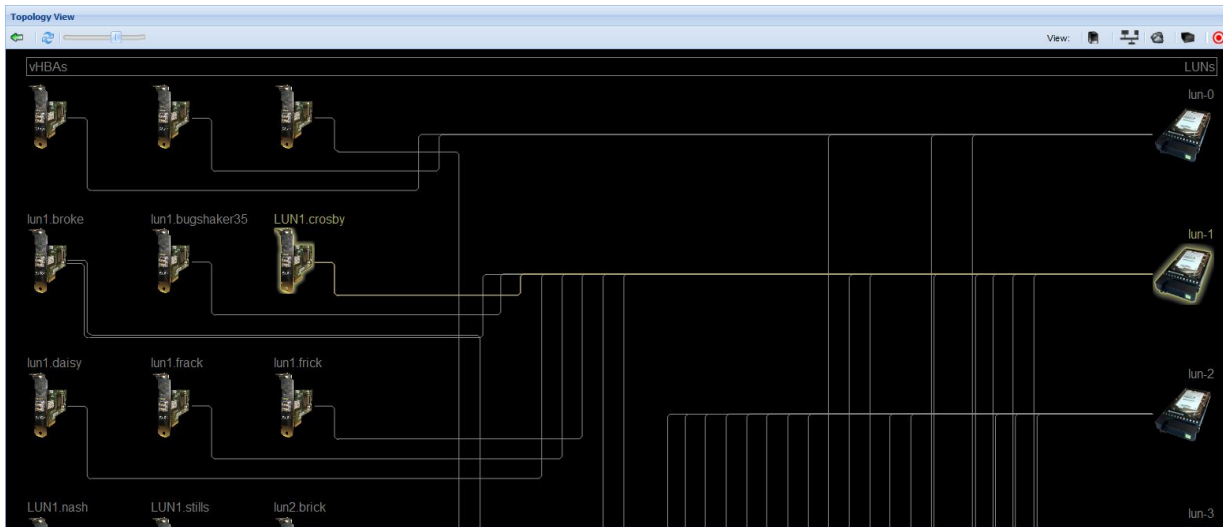
vHBA-to-LUN Connections

Through the Topology you can display detailed information about which vHBAs are connected to which LUNs on a specific target. The detailed information includes:

- A list of all vHBAs connected to a particular fibre channel storage target
- A list of all LUNs on that target, and how each vHBA is connected

By hovering over a vHBA you will highlight its connection to a particular LUN. Also, by clicking one of the listed LUNs, you will highlight all of its connections to listed vHBAs.

You can double-click a target in the Storage Topology view to display the detailed target information.



In this example, all vHBAs connections are displayed for the storage target that you click. Also, the LUNs within that target are displayed. For this example, the mouse was hovered over the vHBA named “LUN1.crosby” to highlight how that vHBA is connected. You can see that it is connected to lun-1, so the server “crosby” which is where vHBA “LUN1.crosby” exists is connected to lun-1 on the selected target.

Performance Meters

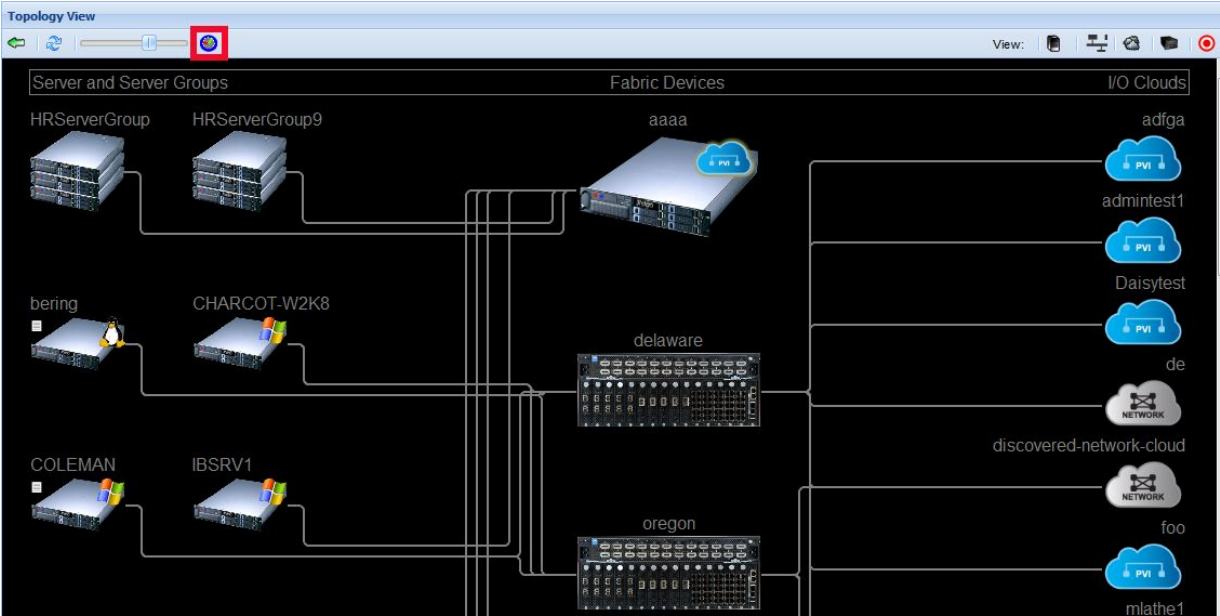
Through the Topology, you can display performance meters for individual servers through a toolbar button. Through the meters you can display aggregate network and storage throughput for servers in the topology or I/O modules connected to clouds.

Performance meters are available on most pages in the Topology where elements with measurable throughput are present (typically, servers or I/O Modules). The following table shows which views have performance meters available.

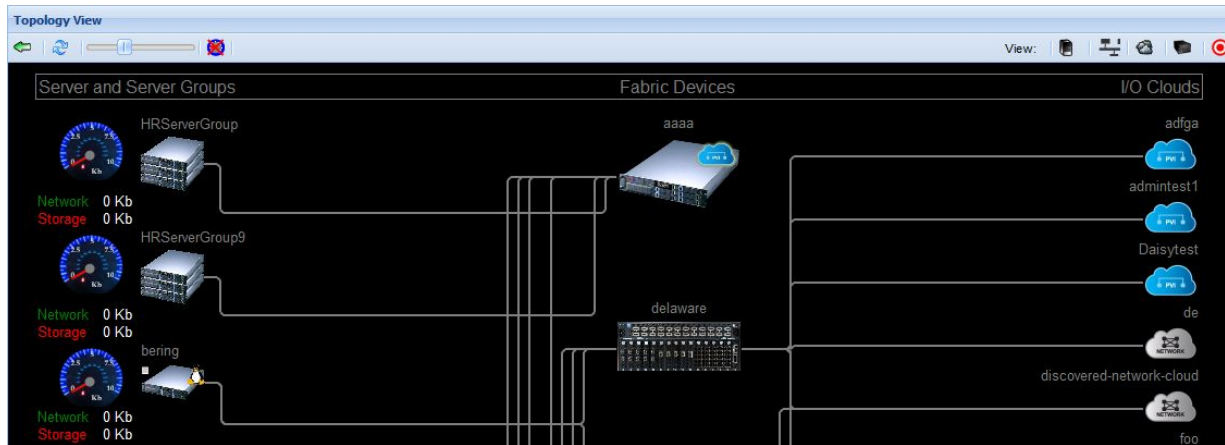
View	Shows
Topology Overview	Throughput of servers in the Topology
Topology Server Details	Double-click a server in the Topology, and performance meters are available to show throughput for that particular server.
Server Cloud view	Throughput of servers in the Server Cloud view
Server Cloud Details	Double-click a server in the Topology, and performance meters are available to show throughput for that particular server.
Fabric Device Cloud Details	Double-click a Fabric Interconnect or Oracle SDN Controller in the Fabric Device Cloud view, and performance meters are available to show throughput for individual I/O Modules in that Fabric Device.

View	Shows
Server Fabric Device view	Throughput of servers in the Server Fabric Device view

The following figure shows an example of the Topology Overview without performance meters, and the location of the Show Performance Meters toolbar button.



By default, the meters are not displayed. However, since this button is a toggle, click it once to display the performance meters as shown in the following figure, and click once again to hide the performance meters.



As shown, each performance meter consists of two needles on a speedometer. The Green needle represents the total network throughput across vNICs on the server. The red needle represents the total I/O across all vHBAs on the server.

When you no longer want to see the performance meters, click the Hide Performance Meters.

Physical and Virtual View Options

Many of the individual pages in the Topology show virtual connections from Oracle resources to other devices, but you might also want to know how a physical connection exists in your data center. For this reason, Oracle Fabric Manager offers a toggle to show physical and virtual connections through the Show Physical View or Show Virtual View toolbar button that exists on many of the Topology pages. The following table shows the pages that contain the toolbar button and explains what the button does.

View	Toggle	See...
Topology Details	Default: Server virtual connections to clouds. Use Show Physical View toolbar button to show the physical ports on the Fabric Interconnect or Oracle SDN Controller to which HCAs are connected.	“Display Physical Connectivity for the Topology Overview” on page 378
Server Cloud view	Default: Server virtual connections to clouds. Use Show Physical View toolbar button to show the Fabric Interconnect or Oracle SDN Controller to which servers are connected.	“Server Cloud Physical Connections Topology” on page 382
Server Cloud Details	Default: HCA virtual connections to clouds. Use Show Physical View toolbar button to show the physical ports on the Fabric Interconnect or Oracle SDN Controller to which HCAs are connected.	“Server Cloud Physical Connections Topology” on page 382
Fabric Device Cloud Details	Default: One Fabric Interconnect or Oracle SDN Controller's virtual connections to clouds. Use Show Physical View to show	“Fabric Device Cloud Physical Connections Topology” on page 384

View	Toggle	See...
Server Fabric Device view	the InfiniBand server ports connections to individual servers and InfiniBand switches (if present)	“Server Fabric Device Virtual Connections Topology” on page 387
Server Fabric Device Details	Default: Shows physical connections between servers and Fabric Interconnect or Oracle SDN Controllers. Use Show Virtual View to show the virtual connections for all servers to their clouds	“Server Fabric Device Virtual Connections Topology” on page 387
Server Fabric Device	Default: Shows the physical connection between servers and Fabric Interconnect or Oracle SDN Controller ports including connections to InfiniBand switches (if present). Use Show Virtual View to show the virtual connections between servers and Oracle Fabric Interconnect or Oracle SDN Controller.	“Server Fabric Device Virtual Connections Topology” on page 387

Working With Discovery Subnets

This chapter contains the following topics:

- [“Display the Discovery IP Subnet” on page 395](#)
- [“Display Discovery Subnet Details” on page 396](#)
- [“Add a Discovery Subnet” on page 397](#)

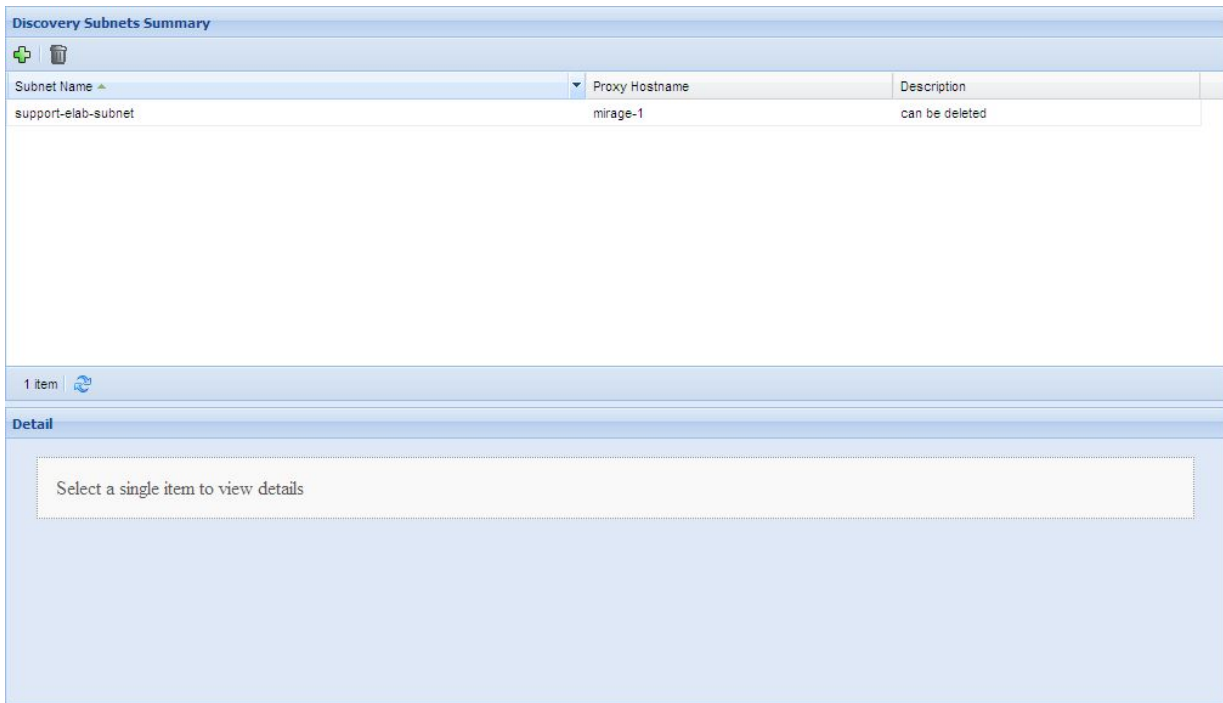
▼ Display the Discovery IP Subnet

By default, Oracle Fabric Manager discovers the IP subnet that it is currently on. However, there might be a need to discover other subnets. In this case, Oracle Fabric Manager can query the remote IP subnet through the use of a proxy Oracle Fabric Interconnect you specify in that remote network. This query occurs only for the purpose of discovering devices on other subnets. The discovery subnet shows the IP-connected devices that can be found on that subnet, but not all subnets, so if you want to discover other subnets, you must manually discover them individually.

Note - Discovery IP Subnet functionality is supported on Oracle Fabric Interconnects only (not Oracle SDN Controllers).

To display the discovery subnet, follow this procedure:

- **From the Navigation frame, select Managed Devices – Discovery Subnet to display the Discovery Subnet Summary.**



Notice that through the Discovery Subnet Summary, you can add new discovery subnets or delete existing discovery subnets.

▼ Display Discovery Subnet Details

The IP Discovery Subnet has a details frame that displays additional information for a configured discovery subnet. Through the IP Discovery Subnet Details frame, you can also edit selected parameters.

To display details for a discovery subnet, follow this procedure:

1. **From the Navigation frame, select Managed Devices – Discovery Subnet to display the Discovery Subnet Summary.**

For an example, see [“Display the Discovery IP Subnet” on page 395](#).

2. Click an IP Discovery Subnet in the Discovery Subnet Summary to display the selected item in the details frame.

The screenshot shows the 'Discovery Subnets Summary' interface. At the top, there is a table with the following data:

Subnet Name	Proxy Hostname	Description
support-elab-subnet	mirage-1	can be deleted

Below the table, there is a section titled 'Discovery Subnet: support-elab-subnet' which displays the following details:

- Name:** support-elab-subnet
- Proxy Hostname:** mirage-1
- Description:** can be deleted

An 'Edit' button is visible at the bottom left of the details section.

Notice that on the Discovery Subnet details frame, you can edit the details to change the proxy host name used for discovery, or the description.

▼ Add a Discovery Subnet

Through the Discovery Subnet Summary you can add additional discovery subnet profiles so that Oracle Fabric Manager can scan additional IP subnets. By adding a discovery subnet, you allow Oracle Fabric Manager to learn about hosts and other network devices that are connected on a different subnet than the Oracle Fabric Manager Server.

You can add multiple discovery subnets by creating multiple discovery subnet profiles. To add a discovery subnet, follow this procedure:

1. From the Navigation frame, select **Managed Devices** → **Discovery Subnet** to display the **Discovery Subnet Summary**.
2. Click the plus sign to display the **New Discovery Subnet dialog**.

3. In the **Name** field, enter a name for the discovery subnet profile that you are configuring.
4. In the **Proxy Hostname** field, enter the name of the proxy device that will be used to discover devices on the IP subnet. The Proxy host name can be either a DNS resolvable name or an IP address.
5. As an option, in the **Description** field, enter a description for the discovery subnet.
6. When the Discovery Subnet profile's parameters have been specified, click **Submit** to configure the discovery subnet profile.

Working With Alarms

This chapter contains the following topics:

- [“Alarm Summary Overview” on page 399](#)
- [“Display the Alarm Summary” on page 399](#)
- [“Display the Alarm History Summary” on page 401](#)
- [“Filter the Alarm History Summary” on page 402](#)
- [“Displaying Detailed Alarm Information” on page 402](#)
- [“View an Alarm's General Properties” on page 403](#)
- [“View an Alarm's Historical Information” on page 404](#)
- [“Clear Alarms From the Alarm Summary” on page 406](#)
- [“Clear Alarms From the Alarm History” on page 406](#)
- [“Filter the Alarm History Summary” on page 407](#)

Alarm Summary Overview

Alarms are posted to the Alarm Summary, which contains instances of existing alarms. Alarms in the Alarm Summary remain until they are explicitly cleared. When cleared, the “cleared” alarm is no longer present in the Alarm Summary. It is possible for an alarm to be cleared, then reoccur in the Alarm Summary. However, for this to happen, the condition(s) that spawned the alarm must occur again.

▼ Display the Alarm Summary

Oracle Fabric Manager tracks system events and network management alarms and displays them in a table called the Alarm Summary. The Alarm Summary contains real-time information about alarms and events. An alarm is a network management fault of one of the following severities as defined by the TMN TMF model:

- Critical
- Major
- Warning

Display the Alarm Summary

- Indeterminate
- Minor
- Info
- Conditional

Through the Filter menu, you can control the displayed contents of the Alarm Summary, by filtering based on any of these severities.

To display the Alarm Summary, follow this procedure:

- **From the navigation panel, select General → Alarms. The Alarm Summary is displayed.**

Alarm Summary						
Object Name	Director Name	Severity	Type	Cause	Time Created	Description
ethernet	texas	info	ethernetTermination...		2010.07.28.22.04.11.832	LAG ports speed mismatch
ethernet	texas	info	ethernetTermination...		2010.07.28.22.04.11.505	LAG ports speed mismatch
ethernet	texas	info	ethernetTermination...		2010.07.20.18.55.04.174	LAG ports speed mismatch
phys-con-1b78ffff...	florida	major	resource	unavailable	2010.07.29.18.33.10.754	Server connection down. Check conne...
havn63	florida	major	service	linkDown	2010.07.29.18.32.27.448	IO Link is down.
havn62	florida	major	service	linkDown	2010.07.29.18.32.25.959	IO Link is down.
havn61	florida	major	service	linkDown	2010.07.29.18.32.24.536	IO Link is down.
phys-con-2c90200...	florida	major	resource	unavailable	2010.07.29.18.23.20.473	Server connection down. Check conne...
phys-con-2c90300...	texas	major	resource	unavailable	2010.07.01.20.31.30.980	Server connection down. Check conne...
phys-con-2c90300...	arkansas	major	resource	unavailable	2010.07.01.19.03.22.604	Server connection down. Check conne...
phys-con-2c90300...	arkansas	major	resource	unavailable	2010.07.01.19.03.22.604	Server connection down. Check conne...

41 items

Field	Indicates
Object Name	The object in the Oracle model on which the alarm occurred. This field typically contains server profile names, vNIC or vHBA names, or physical server names.
Fabric Device Name	The name of the Oracle Fabric Device where the alarm condition occurred.
Severity	The level of alarm that was raised.
Type	The type of error that occurred.
Cause	The cause of the alarm.
Time Created	The date and time stamp of when the alarm occurred.
Description	A brief description of the alarm condition.

▼ Display the Alarm History Summary

Through the Alarm History you can display a history of all alarms reported to Oracle Fabric Manager. This information is displayed through the Alarm History Summary tab, which contains a table of each alarm's activity and any state changes that have occurred for each alarm. The Alarm Summary History tracks alarm historical information between Oracle Fabric Manager reboots, so if the server is rebooted (for example, as part of a software upgrade) the Alarm History Summary will be cleared. However, if the alarm condition still exists, it will be reported to Oracle Fabric Manager, and in turn, the Alarm History Summary will begin filling again.

- You can display the Alarm History Summary by selecting **Service Manager** → **Alarms**, then clicking the **Alarm History Summary** tab.

Object Name	Director Name	Severity	Type	Cause	Time Created	Time Updated	Description
blaine	texas	warning	configuration	terminationUnspecified	2010.07.28.19.3...	2010.07.28.19.3...	Physical equipment or equipment group t...
blaine	texas	cleared	configuration	terminationUnspecified	2010.07.28.19.3...	2010.07.28.19.3...	Physical equipment or equipment group t...
eth3	arkansas	warning	configuration	terminationUnspecified	2010.07.28.19.3...	2010.07.28.19.3...	Physical equipment or equipment group t...
eth3	arkansas	cleared	configuration	terminationUnspecified	2010.07.28.19.3...	2010.07.28.19.3...	Physical equipment or equipment group t...
eth4	arkansas	warning	configuration	terminationUnspecified	2010.07.28.19.3...	2010.07.28.19.3...	Physical equipment or equipment group t...
eth4	arkansas	cleared	configuration	terminationUnspecified	2010.07.28.19.3...	2010.07.28.19.3...	Physical equipment or equipment group t...
foo	texas	warning	configuration	terminationUnspecified	2010.07.28.19.3...	2010.07.28.19.3...	Physical equipment or equipment group t...
foo	texas	cleared	configuration	terminationUnspecified	2010.07.28.19.3...	2010.07.28.19.3...	Physical equipment or equipment group t...
frack-hclAR	arkansas	warning	configuration	terminationUnspecified	2010.07.28.19.1...	2010.07.28.19.1...	No physical compute resource provision...
frack-hclAR	arkansas	cleared	configuration	terminationUnspecified	2010.07.28.19.1...	2010.07.28.19.4...	No physical compute resource provision...
frack-hcITX	texas	warning	configuration	terminationUnspecified	2010.07.28.19.1...	2010.07.28.19.1...	No physical compute resource provision...
frack-hcITX	texas	cleared	configuration	terminationUnspecified	2010.07.28.19.1...	2010.07.28.19.4...	No physical compute resource provision...
frick-hclAR	arkansas	warning	configuration	terminationUnspecified	2010.07.28.19.1...	2010.07.28.19.1...	No physical compute resource provision...
frick-hclAR	arkansas	cleared	configuration	terminationUnspecified	2010.07.28.19.1...	2010.07.28.19.4...	No physical compute resource provision...
frick-hcITX	texas	warning	configuration	terminationUnspecified	2010.07.28.19.1...	2010.07.28.19.1...	No physical compute resource provision...
frick-hcITX	texas	cleared	configuration	terminationUnspecified	2010.07.28.19.1...	2010.07.28.19.4...	No physical compute resource provision...
havn61	florida	major	service	linkDown	2010.07.28.19.1...	2010.07.28.19.1...	IO Link is down.
havn61	florida	cleared	service	linkDown	2010.07.28.19.1...	2010.07.28.19.1...	IO Link is down.
havn61	florida	cleared	service	linkDown	2010.07.28.16.5...	2010.07.28.19.2...	IO Link is down.

Because the Alarm History Summary can contain a large number of entries, you can filter the entries to display pertinent information, and you can also manually clear one or more alarm entries from the table.

The Alarm History Summary has no corresponding Alarm History Details frame. However, you can display alarm history information for individual alarms. See [“View an Alarm's Historical Information” on page 404](#).

▼ Filter the Alarm History Summary

The Alarm History Summary can be filtered by clicking the magnifying glass icon, which allows you to select predefined date options for filtering. When you specify filtering criteria, you select a starting date and an ending date that creates a date range. All alarms that occurred during that date range are then displayed with a highlight.

Using these options, you can create the date range to use for filtering. For example, setting the start date as yesterday, and the ending date as today creates a 1-day date range, and all alarms that occurred from yesterday to today would be displayed.

You can filter the Alarm History Summary by using the Filter dialog.

To filter the Alarm History Summary, follow this procedure:

- 1. From the navigation panel, select General → Alarms.**
The Alarm Summary is displayed.
- 2. Click the Alarm History Summary tab to display the Alarm History Summary.**
See [“Alarm Summary Overview” on page 399](#) for an example.
- 3. On the Alarm History Summary, click the magnifying glass icon to display the Filter dialog.**
- 4. From the Start Date dropdown menu, click to start the calendar utility, and select the start date from the calendar. As an alternative, you can manually enter the start date in the format shown.**
- 5. In the End Date dropdown menu, click to start the calendar utility, and select the end date from the calendar.**
As an alternative, you can manually enter the end date in the format shown.
- 6. When the Filter dialog is completed, click Submit.**
- 7. Check the Alarm History Summary, which displays the filtered results.**

▼ Displaying Detailed Alarm Information

The Alarm Details page contains additional information about network alarms. Detailed alarm information is available through the Alarm Details page.

To display detailed alarm information, follow this procedure:

- 1. From the navigation panel, select General → Alarms to display the Alarm Summary.**

See “[Display the Alarm Summary](#)” on page 399 for an example.

2. In the Object Name field, click the name of the alarm to display detailed information for that alarm.

When you click the object name, the Alarm Details page is displayed.

The screenshot shows a table of alarm summaries. The selected row is highlighted in red. Below the table, the 'Alarm Details' frame is expanded for the selected alarm, showing its properties in a key-value format.

Object Name	Director Name	Severity	Type	Cause	Time Created	Description
ethernet	texas	info	ethernetTermination...		2010.07.28.22.04.11.832	LAG ports speed mismatch
ethernet	texas	info	ethernetTermination...		2010.07.28.22.04.11.505	LAG ports speed mismatch
ethernet	texas	info	ethernetTermination...		2010.07.20.18.55.04.174	LAG ports speed mismatch
phys-con-1b78ffff...	florida	major	resource	unavailable	2010.07.29.18.33.10.754	Server connection down. Check conne...
havn63	florida	major	service	linkDown	2010.07.29.18.32.27.448	IO Link is down.
havn62	florida	major	service	linkDown	2010.07.29.18.32.25.959	IO Link is down.
havn61	florida	major	service	linkDown	2010.07.29.18.32.24.536	IO Link is down.
phys-con-2c90200...	florida	major	resource	unavailable	2010.07.29.18.23.20.473	Server connection down. Check conne...
phys-con-2c90300...	texas	major	resource	unavailable	2010.07.01.20.31.30.980	Server connection down. Check conne...
phys-con-2c90300...	arkansas	major	resource	unavailable	2010.07.01.19.03.22.604	Server connection down. Check conne...
phys-con-2c90300...	arkansas	major	resource	unavailable	2010.07.01.19.03.22.604	Server connection down. Check conne...

Alarm : phys-con-1b78ffff34096e			
Detail		History	
Name:	resource:phys-con-1b78ffff34096e	Cause:	unavailable
Severity:	major	Description:	Server connection down. Check connection to Server.
Detailed:	oper-state is down;	Time Created:	2010.07.29.18.33.10.754 (yyyy.MM.dd.HH.mm.ss.SSS)
Time Updated:	2010.07.29.18.33.10.754 (yyyy.MM.dd.HH.mm.ss.SSS)		

Notice that the Alarms Details frame contains the Detail and History tabs:

- Detail shows the general properties for the alarm.
- History shows detailed historical information about changes that have occurred for the selected alarm.

▼ View an Alarm's General Properties

Through the Alarm Details frame, you can display general properties about a single alarm. The general properties display additional pertinent information about the selected alarm. An alarm's general properties can be displayed through the Detail tab, and historical information about the selected alarm can be displayed through the History tab.

To display an Alarm's general properties, follow this procedure:

1. **From the navigation panel, select **General** → **Alarms** to display the **Alarm Summary**.**
See [“Display the Alarm Summary” on page 399](#) for an example.
2. **In the **Object Name** field, click the name of the alarm (as shown) to display detailed information for that alarm.**

When you click the object name, the Alarm Details frame is displayed. See [“Displaying Detailed Alarm Information” on page 402](#) for an example.

Field	Indicates
Name	The object on which the alarm was raised.
Cause	The cause of the alarm.
Severity	The level of alarm that was raised.
Description	A brief description of the alarm condition.
Detailed	More information (if available) about the error condition on the object that spawned the alarm.
Time Created	The date and time stamp of when the alarm occurred.
Time Updated	The date and time stamp of when the alarm was last posted to the Alarm Summary.

▼ View an Alarm's Historical Information

Through the Alarm Details frame, you can display historical information about a single alarm. The historical information is a table that lists each time an alarm was modified, such as when an alarm has changed state, or an admin has materially changed the alarm so that a new record is added to the table due to its new status resulting from the change to the alarm.

An alarm's historical information can be displayed through the Details tab, and historical information about the selected alarm can be displayed through the History tab.

To display an alarm's historical information, follow this procedure:

1. **From the navigation panel, select **General** → **Alarms** to display the **Alarm Summary**.**
See [“Display the Alarm Summary” on page 399](#) for an example.
2. **In the **Object Name** field, click the name of the alarm (as shown) to display detailed information for that alarm.**

When you click the object name, the Alarm Details page is displayed.

The screenshot shows the 'Alarm Summary' page with a table of alarms. The selected alarm is 'phys-con-1b78ffff...' with a severity of 'major' and a description of 'Server connection down. Check connection to Server...'. Below the summary, the 'History' tab is active, showing a list of events:

Severity	Time Created	Time Updated	Description
major	2010.07.22.23.33.59.206	2010.07.22.23.33.59.206	Server connection down. Check connection to Server.
cleared	2010.07.22.23.33.59.206	2010.07.28.17.29.01.514	Server connection down. Check connection to Server.
major	2010.07.28.19.18.23.499	2010.07.28.19.18.23.499	Server connection down. Check connection to Server.
cleared	2010.07.28.19.18.23.499	2010.07.28.19.18.23.499	Server connection down. Check connection to Server.

The History tab shows the history of changes made to the selected alarm, including the action that cleared the alarm. Historical records for an alarm are arranged oldest to newest so that the most recent changes to an alarm are at the bottom of the list.

The following table shows the contents of the Alarm Details History tab, and explains what each field on the tab means.

Field	Indicates
Severity	The level of alarm that was raised.
Time Created	The date and time stamp of when the alarm occurred.
Time Updated	The date and time stamp of when the alarm was last posted to the Alarm Summary.
Description	A brief description of the alarm condition.

▼ Clear Alarms From the Alarm Summary

When the condition that created the alarm has been fixed, you can remove the alarm from the Alarm Summary by deleting the entry. An alarm that has been removed from the Alarm Summary is considered a cleared alarm. When an alarm is cleared, it is no longer displayed in the Alarm Summary, but is still displayed in the Alarm History Summary.

To clear alarms from the Alarm Summary, follow this procedure:

- 1. From the navigation panel, select General → Alarms.**
The Alarm Summary is displayed. See [“Display the Alarm Summary” on page 399](#) for an example.
- 2. On the Alarm Summary, click one or more alarms to clear them.**
- 3. Click the garbage can icon to delete the selected alarms.**
- 4. On the Confirmation dialog, select Yes to delete the selected alarm, or No to abort deleting the selected alarm.**

▼ Clear Alarms From the Alarm History

The Alarm History Summary contains a table of historical information for all alarms reported to Oracle Fabric Manager. Since all alarms are kept for historical reference, the Alarm History Summary can grow to a large size. You can manually clear the Alarm History Summary to keep it at a manageable size. The Alarm History Summary does not automatically delete entries.

To clear alarms from the Alarm History summary, follow this procedure:

- 1. From the navigation panel, select General → Alarms.**
The Alarm Summary is displayed. See [“Display the Alarm Summary” on page 399](#) for an example.
- 2. Click the Alarm History Summary tab.**
- 3. Click the eraser icon to display the Clean Up Alarm History dialog.**
- 4. In the From Date field, click the calendar utility and select the start date for the time range in which alarms will be cleaned out of the Alarm History.**
As an alternative, you can manually enter the start date in the format shown.
- 5. In the To Date field, click the calendar utility and select the start date for the time range in which alarms will be cleaned out of the Alarm History.**
As an alternative, you can manually enter the end date in the format shown.

- When the dates are entered, click **Submit** to clean alarms that fall within the date range out of the Alarm History.

▼ Filter the Alarm History Summary

When the Alarm History Summary is large, finding specific alarm entries can be difficult. To facilitate finding the appropriate information in the Alarm History Summary, you can use the filtering tool. The filtering tool allows you to specify a starting date and an ending date that create a range. When you apply the filter, the Alarm History Summary displays the alarms for only that date range. All other alarm history summary entries are not displayed.

To filter the Alarm History Summary, follow this procedure:

- From the navigation panel, select **General → Alarms**.

The Alarm Summary is displayed. See [“Display the Alarm Summary” on page 399](#) for an example.

- Click the **Alarm History Summary** tab.

Object Name	Director Name	Severity	Type	Cause	Time Created	Time Updated	Description
blaine	texas	warning	configuration	terminationUnspecified	2010.07.28.19.3...	2010.07.28.19.3...	Physical equipment or equipment group t...
blaine	texas	cleared	configuration	terminationUnspecified	2010.07.28.19.3...	2010.07.28.19.3...	Physical equipment or equipment group t...
eth3	arkansas	warning	configuration	terminationUnspecified	2010.07.28.19.3...	2010.07.28.19.3...	Physical equipment or equipment group t...
eth3	arkansas	cleared	configuration	terminationUnspecified	2010.07.28.19.3...	2010.07.28.19.3...	Physical equipment or equipment group t...
eth4	arkansas	warning	configuration	terminationUnspecified	2010.07.28.19.3...	2010.07.28.19.3...	Physical equipment or equipment group t...
eth4	arkansas	cleared	configuration	terminationUnspecified	2010.07.28.19.3...	2010.07.28.19.3...	Physical equipment or equipment group t...
foo	texas	warning	configuration	terminationUnspecified	2010.07.28.19.3...	2010.07.28.19.3...	Physical equipment or equipment group t...
foo	texas	cleared	configuration	terminationUnspecified	2010.07.28.19.3...	2010.07.28.19.3...	Physical equipment or equipment group t...
frack-hclAR	arkansas	warning	configuration	terminationUnspecified	2010.07.28.19.1...	2010.07.28.19.1...	No physical compute resource provision...
frack-hclAR	arkansas	cleared	configuration	terminationUnspecified	2010.07.28.19.1...	2010.07.28.19.4...	No physical compute resource provision...
frack-hcTX	texas	warning	configuration	terminationUnspecified	2010.07.28.19.1...	2010.07.28.19.1...	No physical compute resource provision...
frack-hcTX	texas	cleared	configuration	terminationUnspecified	2010.07.28.19.1...	2010.07.28.19.4...	No physical compute resource provision...
frick-hclAR	arkansas	warning	configuration	terminationUnspecified	2010.07.28.19.1...	2010.07.28.19.1...	No physical compute resource provision...
frick-hclAR	arkansas	cleared	configuration	terminationUnspecified	2010.07.28.19.1...	2010.07.28.19.4...	No physical compute resource provision...
frick-hcTX	texas	warning	configuration	terminationUnspecified	2010.07.28.19.1...	2010.07.28.19.1...	No physical compute resource provision...
frick-hcTX	texas	cleared	configuration	terminationUnspecified	2010.07.28.19.1...	2010.07.28.19.4...	No physical compute resource provision...
havn61	florida	major	service	linkDown	2010.07.28.19.1...	2010.07.28.19.1...	IO Link is down.
havn61	florida	cleared	service	linkDown	2010.07.28.19.1...	2010.07.28.19.1...	IO Link is down.
havn61	florida	cleared	service	linkDown	2010.07.28.16.5...	2010.07.28.19.2...	IO Link is down.
havn61	florida	cleared	service	linkDown	2010.07.28.16.5...	2010.07.28.19.2...	IO Link is down.

- Click the magnifying glass icon to display the **Filter Alarm History** dialog.

4. **From the Start Date dropdown menu, select the first date in the range you will use for filtering.**
5. **From the End Date dropdown menu, select the last date in the range you will use for filtering.**
6. **When the correct timestamp range is specified, click Submit.**

All alarm history entries that comply with the date range are displayed as highlighted rows in the Alarm History Summary.

Working With High Availability Oracle Fabric Manager

This chapter documents the following topics:

- [“High Availability Oracle Fabric Manager Overview” on page 409](#)
- [“Consideration for HA Oracle Fabric Manager and Plug-Ins” on page 410](#)
- [“Fabric Performance Monitoring in HA Oracle Fabric Manager Environment” on page 411](#)
- [“Edit the PostgreSQL File to Include the HA Partners” on page 412](#)
- [“Failover and Failback” on page 413](#)
- [“Local Host and Remote Host” on page 414](#)
- [“HA States” on page 414](#)
- [“Configuring HA Oracle Fabric Manager Servers” on page 416](#)
- [“Set an HA Partner's Mode” on page 418](#)
- [“Set a Sync Interval” on page 420](#)
- [“Force Sync Up” on page 421](#)
- [“Performing Failover and Failback” on page 421](#)
- [“Install Plug-ins on the Passive Server” on page 423](#)
- [“Displaying HA Oracle Fabric Manager Information” on page 424](#)
- [“Displaying HA Partner Detailed Information” on page 425](#)
- [“Delete an HA Partner” on page 430](#)

High Availability Oracle Fabric Manager Overview

Oracle Fabric Manager supports high availability mode, in which multiple Oracle Fabric Manager servers are associated with each other to provide a system of Oracle Fabric Manager servers that operate in active/passive roles. This high availability Oracle Fabric Manager (HA Oracle Fabric Manager) system consists of the following components:

- One active Oracle Fabric Manager server, and one passive Oracle Fabric Manager server. Together, the two servers are called HA partners
- A Oracle Fabric Manager server configuration (for the primary Oracle Fabric Manager server)

- An HA configuration (one for each of the HA partners)

When HA Oracle Fabric Manager partners are configured for a Oracle Fabric Manager server, the active and passive partner work together as a pair to retain the same configuration (or an extremely close match) to provide high availability to the Fabric Devices and virtual resources configured within the HA Oracle Fabric Manager deployment. After initial configuration, the active partner syncs up automatically with the passive partner. From that point forward, all the HA partners keep in contact with the active partner by sending HA ping packets which verify connectivity between the partners. (Note that HA pings are a separate proprietary message, not a standard ICMP ping.) All nodes use pings to verify each other's mode, and also to update their records with information about changes made since the last ping. Through the pings, HA partners can determine if one of the Oracle Fabric Manager servers in the HA system has gone offline, and will update that server's state in the Oracle Fabric Manager user interface so that you can take any corrective action.

Note - Currently, if an HA partner is determined to be not operational, there is no alarm or notification that the server. Also, there is no recovery or self-healing algorithm to get the offline server back online. If an HA partner is determined to be offline, you must take actions to get the server back online.

Also, the active partner periodically syncs a backup file to all passive nodes to ensure that the HA system has the same configuration. On the passive partner, the backup file is stored in the `xms-backups` directory. You can customize the sync interval through the Oracle Fabric Manager user interface if your network requires a quicker or slower sync between the passive and active partners.

The High Availability Oracle Fabric Manager system has the following requirements:

- The OS installed must be the same on both servers
- The version of Java Runtime Environment (JRE) must be the same on both servers
- Both servers must be pingable by host name
- For HA Oracle Fabric Managers with Fabric Performance Monitor running on both servers, you must install the PostgreSQL data on a separate server that is reachable by both partners in the HA system. For more information, see [“Configuring HA Oracle Fabric Manager Servers” on page 416](#).

Consideration for HA Oracle Fabric Manager and Plug-Ins

Oracle Fabric Manager supports numerous additional robust tools as plug-ins to the Oracle Fabric Manager core graphical user interface. However, with HA Oracle Fabric Manager, two instances of Oracle Fabric Manager are running—one on the active partner, and one on the passive partner.

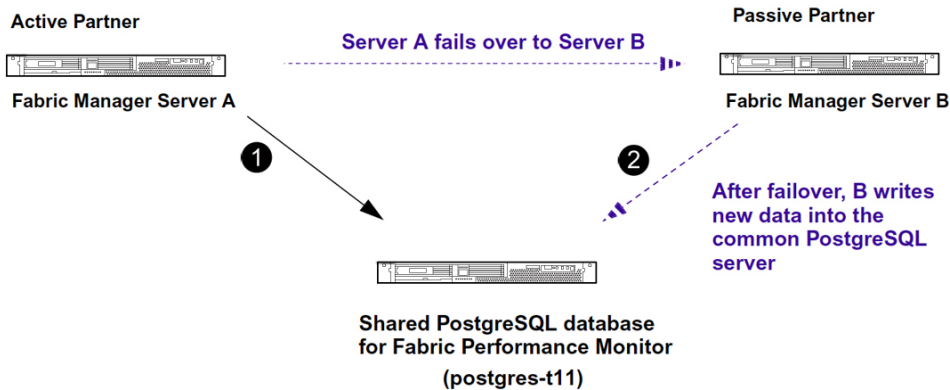
If you will be using plug-ins in your HA Oracle Fabric Manager system, it is a best practice to install the plug ins on both the partners. To do so, completely install the plug ins on both Oracle Fabric Manager servers before connecting them together in an HA system.

By doing installing the plug ins on both servers that will be in the HA system, both servers are “active” at that time. As a result, you ensure that important information, such as config files, application files, and database records, are backed up when the active partner syncs with the passive partner. Also, in the unlikely event that the active partner becomes unavailable, the same applications will be available to the passive partner after it is promoted to the new active node. By installing plug ins on both servers in the HA system before configuring them in the HA system, you retain identical functionality after a failover event, and minimize downtime. If you already have HA Oracle Fabric Manager configured and need to install one or more plug-ins on the passive partner, see [“Install Plug-ins on the Passive Server” on page 423](#).

Fabric Performance Monitoring in HA Oracle Fabric Manager Environment

For HA Oracle Fabric Manager, if you will be installing Fabric Performance Monitoring, be aware that you will need to install the PostgreSQL database on a separate server that is reachable by both HA partners. This requirement exists to allow the database to be available to both servers during a failover. As part of this requirement, you should follow the PostgreSQL best practices to install and backup the database. Documentation for such best practices is out of the scope of this documentation, but is available in the public domain such as on the web.

In an HA system, if the PostgreSQL database is installed on only the active server, a failover can cause the Fabric Performance Monitoring application with current information to move away from the PostgreSQL database. But, when the PostgreSQL database is installed on a commonly accessible server, Fabric Performance Monitoring survives a failover. Consider the following example.



However, by installing the PostgreSQL database on a shared server and pointing both servers to shared server's IP address, the database remains available after the failover regardless of which HA partner is active. During normal runtime operation (see “1” above) the active HA partner writes information to the PostgreSQL database. However, when the active HA partner “Server A” fails over, (see “2” above), the PostgreSQL database is still online and available on the shared PostgreSQL server (“postgres-t11”). As long as both HA partners point to the shared server as the location of Fabric Performance Monitoring's PostgreSQL database. The database will be available to both the HA partners. As a result, after you bring “Server B” back online, it can access the PostgreSQL database and Fabric Performance Monitoring can easily resume operation.

Note - This requirement exists only on an HA system. On a stand-alone server (non-HA), the PostgreSQL database is available on the same server where Fabric Performance Monitoring is installed and running. As long as that server remains up and running, the database is available.

▼ Edit the PostgreSQL File to Include the HA Partners

In order to allow both partners to access PostgreSQL database, you will need to configure the database with either the specific IP addresses of the Oracle Fabric Manager servers, or the subnet address on which the Oracle Fabric Manager servers are configured. Specifying this information determines what nodes are authorized to access the database.

1. **To do so, you will need to edit the PostgreSQL Client Authentication Configuration File (`pg_hba.conf.conf`) file to add the Oracle Fabric Manager server IP information.**
 - **On Windows Oracle Fabric Manager servers, the file is located in `C:\Program Files\PostgreSQL\9.1\data`**

- On Linux Oracle Fabric Manager servers, the file is located in `opt/postgres/9.1/data`
2. Using any common file editor, you will need to open the `pg_hba.conf.conf` file and edit it to include either the specific IP addresses of each of the Oracle Fabric Manager servers in the HA Oracle Fabric Manager system, or add the subnet and mask on which the servers are configured.

The following example `pg_hba.conf.conf` file is provided for reference. The blue text shows an example of the server IP address information that you will need to add.

```
##### start change oracle

# TYPE DATABASE USER ADDRESS METHOD

host all all 192.168.38.131 md5

host all all 192.168.38.132 md5

##### end change oracle
```

Failover and Failback

With HA Oracle Fabric Managers, you have two servers with the same configuration, which provides a level of high availability. If one server goes offline, the other server can continue operating as the Oracle Fabric Manager server controlling your deployment by using a manual failover. Manual failover is a corrective action that you can take when the active partner goes offline. When the original active partner goes down, you can manually promote the passive partner to become the active partner so that configuration and management experiences a minimum of interruption. In a manual failover, it might be required to manually restore a backup file to bring the latest configuration onto the new active partner.

Note - In the HA Oracle Fabric Manager system, failover does not occur automatically, so you must take action to move the configuration and resolve the condition that caused the HA partner to go offline.

A failback occurs when the offline server is brought back online and the configuration is moved back to the original server. Failback is not mandatory. You can keep the existing active partner, and bring the original active partner online as a passive node. For information about performing the failover and failback procedure, see [“Performing Failover and Failback” on page 421](#).

Local Host and Remote Host

HA Oracle Fabric Manager consists of two servers (hosts) that are either the localhost or the remote host. At initial configuration, and upgrade, both servers are the local host. As local host, each server has a record added to the database. The name of the record is the fully-qualified name for the host. At this point, both servers are independent servers that are seen as “active” role servers.

When you connect the two servers into an HA system, you will add them by name. You must use the exact name for each localhost server as it was entered into the configuration database—for example, the DNS server name. If you enter a different name, then the HA system will not be correctly configured between the two servers. For example, if you have a server with a DNS name of larry.lab.companyA, you cannot use “larry” when you add that server into the HA system because the database record associated with that server is “larry.lab.company”

After the HA system is configured the localhost is not tied a specific server. Instead, the local host is the server where you currently have the management session. The other server is the remote host. For example, if you are currently on the active server, it is the localhost and the passive server is the remote host. But, if you are on the passive server then it is localhost and the active server is the remote host.

It is important to understand local host and remote host because some functions in the HA Oracle Fabric Manager system are started from the local host but actually run on the remote host. For example, displaying statistics occurs when a command is sent from the local host to the remote host. The statistics are gathered from the remote host and sent back to the local host where they are displayed. Even though the statistics are displayed on the local host, they are from the “other” server which is the remote host.

HA States

HA Oracle Fabric Manager has different states that depend on the presence of HA partners and their state. Displaying the operational state of HA Oracle Fabric Manager is supported through the HA toolbar object on the Oracle Fabric Manager toolbar. This object also is a dropdown menu which supports configuration of the current instance of HA Oracle Fabric Manager, as well as configuring HA Partners. The HA object is on the Oracle Fabric Manager toolbar, which is also a dropdown menu that is the first step in configuring the HA Oracle Fabric Manager system.

The state is partially derived from a number of factors, including the presence or absence of the remote partner. When the state is determined, it is displayed as a concatenation to two information elements:

- The first information element indicates which partner

- The second information element indicates the actual administrative state of that partner.

So, for example, `active:up` indicates that the active partner is in “up” state. In the user interface, the states in the table are color coded to indicate any error. Red text indicates some kind of error, and green text indicates no error (correct functionality). The only correct states, which is shown in green text, are `active:up` and `passive:up`. The following table shows the HA states and explains what each means.

HA State...	Means...
<code>active: up</code>	The active node is present and HA is correctly configured for HA Oracle Fabric Manager. In this state, the number of remote partners detected is equal to the number of remote partners connected to the active partner. This state is a correct runtime state for the active partner.
<code>active:down</code>	The active node is present but HA is not correctly configured. In this state, a partner might be detected, but that partner is not connected to the current active server. This is an error state.
<code>active: not configured</code>	The active node is present and HA is not configured because no remote partners are defined in the current instance of HA. This situation occurs typically when the passive node is not present. This is an error state.
<code>active: remote node not configured</code>	The active node is present, but the HA configuration is only half complete because the local host has a configuration for the remote host, but the remote host does not have a configuration for the local host. In this state, the active partner knows detects the remote partner and can ping it, but the remote partner does not respond to the pings. This situation can occur when the both partners are present, but the intended passive partner is not yet in “passive” mode. This is an error state.
<code>active: more than 1 active partner</code>	The active node is present and HA is not correctly configured because too many active partners exist. The HA Oracle Fabric Manager system supports a 1:1, active-passive configuration. If more than one active server exists in the HA system, a conflict occurs. This situation can occur during a failback, when the previous active partner has been brought back online, and the interim active partner has not yet been demoted to its original passive role. This is an error state.
<code>passive: up</code>	The passive node is present and HA Oracle Fabric Manager is configured, and the node is connected to its active partner. In this state, the passive server is synced up with its active node. This state is the correct runtime state for the passive partner. On the passive server, the Oracle Fabric Manager navigation frame is disabled. Any configuration or management tasks are complete, it is a recommended that you perform a backup so that the passive partner's configuration is a close match to the active partner.
<code>passive:down</code>	The passive node is present but HA is not correctly configured. In this state, a partner might be detected, but that partner is not connected to the current passive server. This is an error state.
<code>passive: not configured</code>	The passive node is present and HA is not configured because no remote partners are detected. This situation occurs typically when the active node is not present. This is an error state.
<code>passive: remote node not configured</code>	The passive node is present and configured but HA is not configured correctly. In this state, the passive partner detects the remote partner and can ping it, but the remote partner does not respond to the pings. This situation can occur when the both partners are present, but the remote partner is not correctly configured. This is an error state.
<code>passive: no active partner</code>	The remote node is present and HA is not configured because no active partner is connected. In this state, no active partner exists, and as a result, the passive mode

HA State...	Means...
passive: more than 1 active partner	<p>has nothing to sync up with. This situation can occur when an active node has been changed to passive and two passive servers exist—for example, if a failover has occurred without promoting the original passive to active. This is a serious error state because configuration and management tasks are supported on the active partner only. In this state, both HA partner are effectively read-only.</p> <p>The passive node is present and HA is not correctly configured because too many active partners exist. The HA Oracle Fabric Manager system supports a 1:1, active-passive configuration. If more than one active server exists in the HA system, a conflict occurs. This situation can occur when too many servers in the HA system have been configured with the active role. This is an error state.</p>

Configuring HA Oracle Fabric Manager Servers

HA Oracle Fabric Managers work as a pair of servers that retain configuration information and allow a standby servers (the passive server) to take over if the primary server (the active server) becomes unavailable as long as the servers have been synced up at least once.

Each HA partner must be specified by a host name., and communication between the HA Partners occurs through a secure connection. When an HA pair is configured, both the Oracle Fabric Manager Server and its HA partner are online. One of the servers is active and the other is passive.

- [“Configure the Active Server” on page 416](#)
- [“Configure the Passive Server” on page 417](#)

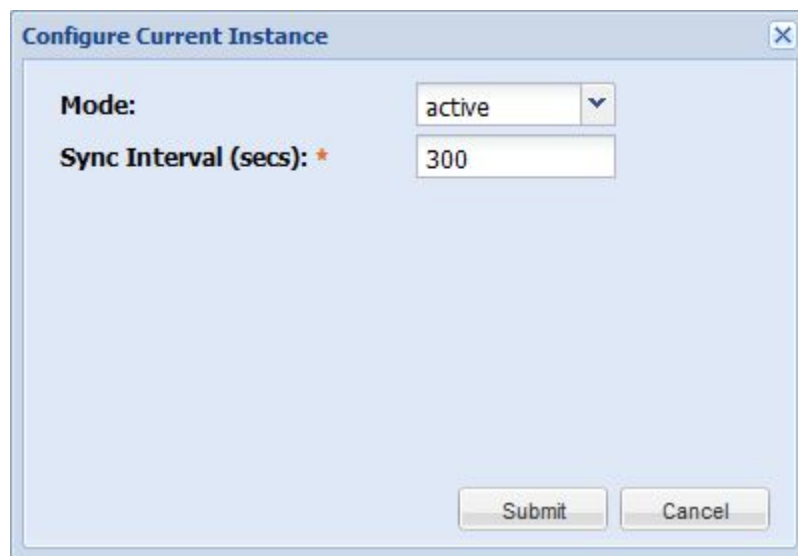
▼ Configure the Active Server

HA Partners can be in either active or passive states. The Active partner is the one on which commands are issued, and the server on which the sync up is initiated. When the active HA partner is configured, it can push data to the passive HA partner to complete a backup. Configuring the Active HA partner occurs through the HA dropdown menu on the Oracle Fabric Manager toolbar.

To configure the active HA partner, follow this procedure which requires two Oracle Fabric Manager servers and assumes the first server will be the active HA partner.

1. **Log in to the first Oracle Fabric Manager server.**
By default, this server is the active server and the HA feature is not yet configured.
2. **On the first server's toolbar, click the Add Servers button (the plus sign) to display the New HA Partner dialog.**

3. In the **HostName** field, enter the name of the server that you are configuring as the active HA partner.
4. As an option, in the **Description** field, enter an alphanumeric string that describes the HA partner you are adding.
5. Click **Submit** to add the partner the High Availability Summary.
6. The HA partner you added should be present.
This partner will be unknown.
7. On the HA toolbar icon, select **HA → Configure Current Instance** to display the **Configure HA Partner** dialog.



The screenshot shows a dialog box titled "Configure Current Instance". It has a close button in the top right corner. The dialog contains two fields: "Mode" with a dropdown menu currently showing "active", and "Sync Interval (secs): *" with a text input field containing the value "300". At the bottom right of the dialog, there are two buttons: "Submit" and "Cancel".

8. From the **Mode** dropdown menu, select **active**.
9. Click **Submit** to add the server as the active HA Partner

▼ Configure the Passive Server

The passive HA partner is the second server in the HA Oracle Fabric Manager system. After it syncs up with the active partner, the passive HA partner has a copy of the Oracle Fabric

Manager config, and the server provides high availability by being an online backup server. The passive server is easily recognizable by checking the navigation panel. If the configuration options are disabled (greyed out), then that node is the passive server.

After configuring the active HA partner, you must configure the passive HA server. To do so, follow this procedure:

1. **Log in to the second Oracle Fabric Manager server.**
2. **Repeat [Step 2](#) through [Step 5](#) to add the current server to the High Availability Summary.**
3. **On the HA toolbar icon, select HA → Configure Current Instance to display the Configure HA Partner dialog.**
4. **From the Mode dropdown menu, select passive.**
5. **Click Submit to add the server as the passive HA Partner.**

Note - At this point, one active HA partner and one passive HA partner should be present. The two HA partners should automatically sync up after both are added to the High Availability Summary. You can also initiate a manual sync up. For information, see [“Force Sync Up” on page 421](#).

6. **After the sync up occurs, log back in to both Oracle Fabric Manager servers in the system, and check the HA icon on the toolbar to verify that the state of HA is configured.**
7. **On the active HA partner, the HA toolbar icon should show active:up.**
8. **On the passive HA partner, the HA toolbar icon should show passive:up.**

▼ Set an HA Partner's Mode

The mode of an HA partner indicates how it participates in the HA Oracle Fabric Manager system:

- If the mode is active, this Oracle Fabric Manager server is the online server that is used to manage your Fabric Devices and support virtual I/O connections to host servers. Only one active HA Partner is supported in each HA Oracle Fabric Manager system.
- If the mode is passive, this Oracle Fabric Manager servers is the standby that acts as a redundant server in your deployment. In passive mode, the Oracle Fabric Manager servers is not used for configuration and management of Fabric Devices and servers. In fact, the

objects on the navigation frame are greyed out to indicate that you cannot use the passive server for network management functions.

- In this release, one passive HA partner is supported in each HA Oracle Fabric Manager system.

Changing an HA Partner's mode is useful in situations where one of the Oracle Fabric Manager servers in the HA system becomes unavailable. For example, assume partner A is the active-mode server, and partner B is the passive-mode server and both servers are online and configured. If server A becomes unavailable, you can change partner B's mode to “promote” it to the active server, then add another server (partner C) as the passive partner to regain high availability. Setting an HA partner's mode also is useful for predictable tasks such as scheduled maintenance of a Oracle Fabric Manager server when you manually trigger a failover and failback. (For more information about failover and failback, see [“Failover and Failback” on page 413](#)).

- If you change the HA mode from active to passive, Oracle Fabric Manager disconnects from the Fabric Devices that it is managing, and you are prompted to log out to allow the state change to complete. During the state change, the partner is reported to other partners as passive. When you log back in, the Oracle Fabric Manager navigation frame is disabled to prevent configuration and management from the passive server.
- If you change the HA mode from passive to active, Oracle Fabric Manager verifies that no other active partners exist since there can be only one. Then, the most recent sync up is stored through Oracle Fabric Manager's backup and restore feature. (For information, see [“Backing up the Oracle Fabric Manager Config”](#)). Oracle Fabric Manager connects to all the Fabric Devices that were managed by the original active partner, then prompts you to log out to allow the state change to complete. During the state change, the partner is reported to other partners as the new active node. When you log back in, the new active partner begins to take periodic backups and send them to the other nodes.

Setting an HA Partner's mode occurs through the HA toolbar icon. To set an HA Partner's mode, follow this procedure:

1. **From the server's toolbar, select *HA->Configure Current Instance* to display the **Configure Current Instance dialog**.**
2. **From the *Mode* dropdown menu, select the mode that you want to assign to the current server.**

Note - Remember that within an HA Oracle Fabric Manager system, only one active server is supported.

3. **In the *Sync Interval (secs)* field, specify the sync up interval (if needed).**
4. **When the mode is set, click *Submit* to complete setting the server mode.**

▼ Set a Sync Interval

The sync interval determines the periodicity of sending sync messages from the active node to the passive node. By default, the sync interval between the active and passive HA partners is set to 300 seconds (5 minutes). However, you can set any sync interval that is required in your network. The sync interval can be set to a non-default value at initial configuration time, or at any time after the HA partners have been configured.

Setting a custom sync interval can be beneficial to your deployment based on network conditions. For example:

- Setting a custom sync interval can be useful in highly volatile environments where frequent backups are required. For example, in such an environment you can set the sync interval to a lower number so that backups occur more frequently. Be aware that a smaller sync interval requires more overhead due to sync up occurring more frequently.
- Setting a custom sync interval can be useful also in stable networks or low latency environments where frequent backups are not required. For example, in such an environment you can set the sync interval to a higher number so that backups occur less frequently. Be aware that a larger sync interval creates a larger window between syncs, and as a result, more data can be lost if the active partner goes down.

Setting a sync interval is supported through the *HA->Configure Current Instance* dialog. To set a custom sync interval, follow this procedure:

1. **From the server's toolbar, select *HA->Configure Current Instance* to display the **Configure Current Instance** dialog.**
2. **In the *Sync Interval (secs)* field, you can enter the number of seconds that you want to pass between configuration synchronizations.**

In most cases, the default value will be sufficient. However, this value is customizable. To set a non-default value, enter a number of seconds that indicate the amount of time that passes before HA partners sync up. You must enter at least 10 seconds.

Note - When configuring the *Sync Interval*, make sure that you set a realistic value. It is possible to set an interval large enough that configurations are not synced up in a timely manner, and the result is that configurations might not be the same on both partners. In most cases, the default value will be acceptable.

3. **When you have set the new sync interval, click *Submit* to complete the change.**
4. **Login in to the other HA partner and repeat this procedure making sure to set the same sync interval.**

▼ Force Sync Up

By default, all partners sync up automatically based on the sync interval specified when you configured each instance. When the sync interval is complete, the active partner pushes a backup file to the passive node(s). For example, with a sync interval of 60 seconds, the backup file is pushed from the active node to all passive nodes every minute.

However, there might be instances when you want to manually sync up databases. For example, if a passive node was taken offline for maintenance, you might want to sync it up immediately after it comes back online to get redundant Oracle Fabric Manager servers as soon as possible. In such cases, you can manually synchronize the HA partners through the Force Sync Up button on the HA Summary.

You can force a manual sync up at any time as long as multiple Oracle Fabric Manager servers are online and the HA status is active:up. To manually sync up, follow this procedure:

- 1. Log in to one of the HA partners.**
- 2. On the Oracle Fabric Manager toolbar, select *HA->Configure HA Partners*.**
The High Availability Summary is displayed.
- 3. On the High Availability Summary, click the *Force Sync Up* button.**

When you click the *Force Sync* button, a short amount of time passes while information is exchanged between the Oracle Fabric Manager servers. When complete, a popup dialog indicates that the partners synced up and no errors occurred.

Performing Failover and Failback

This section describes the procedure to manually failover from an active node that has gone offline to its partner, then failback to the original active node when you have brought it back online.

- [“Perform a Failover” on page 421](#)
- [“Perform a Failback” on page 422](#)

▼ Perform a Failover

During a failover, the active node goes offline, and the other partners will see no active partner. Configuration sync up cannot occur between the offline server and the partners. However, after an active partner is detected, the most recent sync up will be restored. Failover is a manual process.

To perform a failover, follow this procedure:

1. **When you notice the active partner is down, start a browser and log in to the passive partner.**
2. **On the toolbar, select *HA-> Configure Current Instance*.**
3. **From the Mode dropdown, select active to promote this node the active node.**
At this point, you can use this node to configure and managed Fabric Devices through Oracle Fabric Manager if needed.
4. **Click *Submit*.**

Note - Remember that at this point only one server is operating, so you will want to resolve the issue on the offline Oracle Fabric Manager server as soon as possible to regain redundancy. If another Oracle Fabric Manager server is available, you can add it as an interim passive node. For information, see [“Configure the Passive Server” on page 417](#).

▼ Perform a Failback

After the original server is brought back online, the server reads its database to determine its role. Because the server was the active partner when it went offline, it comes back online as the active partner.

Note - If you have an interim Oracle Fabric Manager server acting as the passive node and you do not want that server in the HA system when the original active server comes back online, delete the server now as documented in [“Delete an HA Partner” on page 430](#).

To perform a failback to use the original active partner again, follow this procedure:

1. **Make sure that both HA instances are running. At this point, both partners are active so you'll see an error.**
2. **As an option, take a backup on both nodes. Although this is optional, it is recommended. For information, see [“Back Up the Oracle Fabric Manager Configuration” on page 79](#).**
3. **On the node that just came back online, set the mode to passive.**
4. **At this point, the HA status (on the toolbar) should be correct since the HA Oracle Fabric Manager system has one active node and one passive node.**

Note - Make sure that the sync interval is the same between the servers.

5. **From the active node, initiate a sync up. For information, see [“Force Sync Up” on page 421](#).**
6. **On the active node, set the mode to passive.**
7. **At this point, the HA status (on the toolbar) should show an error because both nodes are passive.**
8. **On the original active node, set the mode to active.**
9. **At this point, the HA status (on the toolbar) should show active:up state because the HA Oracle Fabric Manager system has one active node and one passive node.**

▼ Install Plug-ins on the Passive Server

Both servers in the HA Oracle Fabric Manager system must have the exact same plug-ins installed. It is a best practice to make sure that both Oracle Fabric Manager servers are identically configured before creating the HA pair. However, in some deployments, this is not always possible. The procedure in this section documents how to install the plug-ins on a passive HA Oracle Fabric Manager server that is already configured in an HA Oracle Fabric Manager pair.

To install the plug-ins, you will need to have them available to you. Before attempting this procedure, make sure that you have the correct plug-ins and versions ready for installation on the passive server.

The procedure for installing the plug-ins on the passive server require you to stop the HA connection, promote the passive server to active, install the necessary plug-ins, then revert the server back to passive mode. Follow this procedure:

1. **On both Oracle Fabric Manager servers, take a Oracle Fabric Manager backup as documented in [“Back Up the Oracle Fabric Manager Configuration” on page 79](#).**
2. **On the passive server, set the mode to active as documented in [“Set an HA Partner's Mode” on page 418](#). As a result of changing the mode to active, you will need to log back in.**
3. **On the server, which is now in active mode, check the navigation panel. The icons should no longer be greyed out. If the icons are still greyed out, the server is not in active mode.**

4. **Install the necessary plug-in(s). Make sure that the versions you install match the same versions in use on the other server. For information about installing the plug-ins, see the “Installation” chapter of the user guide that accompanied your plug-in(s).**
5. **At this point, the plug-in has been installed on the server, but it has not yet been added to Oracle Fabric Manager.**
6. **On the new active server, click *Plugins->Plugin Manager* to display the Installed Applications Summary.**
7. **On the Installed Applications Summary, click the plus sign (+) to add the plug-in to Oracle Fabric Manager.**
8. **When all the necessary plug-ins are installed, set the server's mode back to passive as documented in [“Set an HA Partner's Mode” on page 418](#). You will need to log in to fully reset the state to passive.**
9. **Log back in to the passive server and check the navigation panel. The icons should be greyed out since this is no longer an active server.**
10. **Force a sync up as documented in [“Force Sync Up” on page 421](#) to get the two servers synchronized.**
11. **Take a backup of the new configuration as documented in [“Back Up the Oracle Fabric Manager Configuration” on page 79](#).**

Displaying HA Oracle Fabric Manager Information

The HA configuration for Oracle Fabric Manager is displayed through the High Availability Summary, which contains a list of all the configured HA Partners for the current Oracle Fabric Manager Server. Be aware that the HA Summary shows all configured HA Partners for the current Oracle Fabric Manager Server, but does not show any specific logical group or connection mapping.



Notice that through the High Availability Summary you can add or delete new HA Partners by using the plus sign or garbage can icon, respectively. You can also clear operational and performance statistics for management traffic between the HA partners, as well as initiate a sync up between the HA partners.

Field	Means...
Host Name	The host name of the server that is the HA Partner
Mode	The mode (either active or passive) for each HA Partner
State	The operational state of the HA Partner. Valid states are: <ul style="list-style-type: none"> ■ Connected, when the HA Partners recognize each other and are correctly connected and configured. ■ Retrying (the connection is not up yet, but is trying) ■ Unknown, when the configuration is partial or not yet complete ■ Not Configured, when an HA Partner does not have a mode (either active or passive) assigned to it. ■ Not Connected ■ Down, when one of the HA Partners is not connected to the other, or that HA Partner is not operational.
Description	An optional alphanumeric character string that describes the HA Partner or the team of Oracle Fabric Manager Server and HA Partner.

Displaying HA Partner Detailed Information

This section contains the following topics:

- [“Display HA Partner Detailed Information” on page 426](#)
- [“HA Partner General Properties” on page 427](#)
- [“HA Oracle Fabric Manager Statistics” on page 427](#)
- [“Clear Partner Statistics” on page 429](#)

- [“Display Unsynced Commands for HA Oracle Fabric Manager” on page 429](#)

▼ Display HA Partner Detailed Information

In addition to the summary information available for each HA partner through High Availability Summary, you can also display additional information through the details frame of each server in the HA Oracle Fabric Manager system.

Each server contains information about the general properties configured for the HA Oracle Fabric Manager system. This information is available through the *General* tab. This tab is available on both the local host and the other HA partner (remote host). For information, see [“HA Partner General Properties” on page 427](#).

However, the following additional information is available, but is conditional based on a combination of which HA partner you are currently administering (the local host), and which partner you select on the High Availability Summary:

- Unsynced Commands, which are available through Unsynced Commands tab. This tab is available on the local host (when your browser is on a server and you select that server in the High Availability Summary). For information, see [“Working With High Availability Oracle Fabric Manager” on page 409](#).
- Statistics, which are available through the Statistics tab. This tab is available on the local host when you select the other HA partner. For information, see [“HA Oracle Fabric Manager Statistics” on page 427](#).

To display details for the HA Oracle Fabric Manager configuration, follow this procedure:

- 1. From the Oracle Fabric Manager toolbar, select the *HA* object to display the High Availability Summary.**

See [“Displaying HA Oracle Fabric Manager Information” on page 424](#).

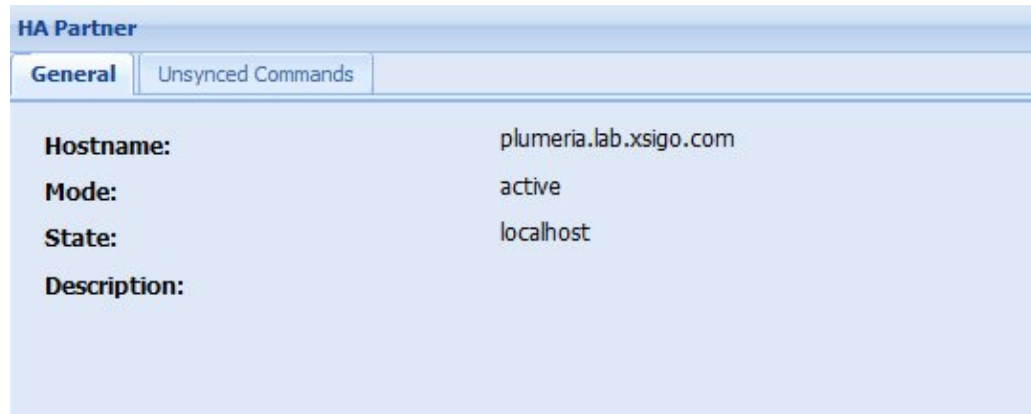
- 2. Select the HA Partner for which you want to display detailed information.**

When the HA Partner is selected, the High Availability details frame is populated with the following types of data:

- **General Properties.** For more information, see [“HA Partner General Properties” on page 427](#).
- **Statistics.** For information see, [“HA Oracle Fabric Manager Statistics” on page 427](#).
- **Unsynced Commands.** For more information, see [“Working With High Availability Oracle Fabric Manager” on page 409](#).

HA Partner General Properties

General properties for the current HA partner are displayed through the General tab. After selecting the HA partner in the High Availability Summary, click the General tab to display the general properties.



Field	Means
Host Name	The host name of the Oracle Fabric Manager server
Mode	The mode in which the host is currently operating. Valid values are: <ul style="list-style-type: none"> ■ active, for the Oracle Fabric Manager server that is maintaining real-time information and pushing that information to the Standby HA ■ passive, for the Oracle Fabric Manager server that is the passive standby partner in the HA configuration. ■ unknown, for a Oracle Fabric Manager server with a mode that is not recognizable as either active or passive
State	The current state of the Oracle Fabric Manager server in the HA environment. Valid values are: <ul style="list-style-type: none"> ■ connected ■ not connected ■ localhost
Description	An optional alphanumeric character string that describes the HA Partner(s), or combination of Oracle Fabric Manager Server and HA Partner(s).

HA Oracle Fabric Manager Statistics

Oracle Fabric Manager shows operational and performance statistics for management traffic sent between HA partners. Statistics and counters for information such as sync up functions, pings, errors, and partner changes are tracked through the non local host partner's *Statistics* tab.

After selecting the HA partner in the High Availability Summary, click the *Statistics* tab to display the HA Oracle Fabric Manager statistics.

HA Partner	
General	Statistics
Start Time:	2012-12-27 11:12:59
Connected Since:	2012-12-27 11:12:59
Last Ping Sent:	2012-12-27 11:37:05
Number of Ping Sent:	474
Number of Ping Failures:	0
Last Sync Sent:	never
Number of Sync Sent:	0
Number of Sync Sent Failures:	0
Last Sync Received:	never
Number of Sync Received:	0
Number of Sync Received Failures:	0

Field	Means...
Start Time	The time and date stamp that indicates when the local host successfully connected to the remote host for the first time.
Connected Since	The time and date stamp since the HA pings have consecutively been successful
Last Ping Sent	The time and date stamp of the last ping sent from the local host.
Number of Pings Sent	The total number of pings sent from the local host to the remote host. This counter is tracked continuously while HA Oracle Fabric Manager is running, but it resets after such actions as restoring a backup, restarting Oracle Fabric Manager, or forcing a sync up.
Number of Ping Failures	The total number of pings failures during failover/failback, and when the statistics are cleared by using the Clear Partner Statistics button on the High Availability Summary page.
Last Sync Sent	The time and date stamp of the last sync sent from the local host to the remote host during automatic sync up and forced sync up.
Number of Syncs Sent	The total number of synchronizations sent from the local host to the remote host during automatic sync up and forced sync up.

Field	Means...
Number of Sync Sent Failures	The total number of failed synchronizations sent from the local host to the remote host during automatic sync up and forced sync up. Sync failures can occur for various reasons, but one common reason is that remote host is not present or connected to the local host.
Last Sync Received	The time and date stamp of the last synchronization successfully received by the local host.
Number of Syncs Received	The total number of synchronizations received by the local host.
Number of Sync Received Failures	The total number of Synchronization received failures on the local host.

▼ Clear Partner Statistics

When statistics exist, you can clear them at any time. Clearing statistics can be helpful when you have brought an HA Oracle Fabric Manager server back online (for example, as part of a failover and failback) to see how the HA partners are communicating.

Clearing partner statistics is supported through the High Availability Summary.

To clear the partner statistics, follow this procedure:

1. **Display the High Availability Summary.**
2. **Select the HA partner (not the local host).**
3. **Click the *Clear Partner Statistics* button.**
4. **On the confirmation dialog, click *Yes* to complete clearing the statistics.**

Display Unsynced Commands for HA Oracle Fabric Manager

Commands and management traffic between the HA partners are displayed through the *Unsynced Commands* tab. By seeing unsynced commands, in the unlikely event that the active partner goes offline, the passive node will have some record of changes that occurred on the active node since the last sync up.

Due to the periodic sync up between the active and passive partners, the possibility exists that some data can be lost if the active partner goes offline right before the scheduled sync interval has completed. By using the Unsynced Commands tab, you have a way to see the unsynced content, and manually recreate it after the passive partner is promoted to the new active partner.

After selecting an HA partner on the High Availability Summary, click the Unsynced Commands tab to display the unsynced commands table as shown in the following example.

Note - This option is available on the local host when you select the local host in the High Availability Summary.

HA Partner	
General	Unsynced Commands
Date	Name
2012-12-04 17:11:52	xmsCommand(root)% execute ha.EditHA {syncInterval=300, mode=passive}
2012-12-04 17:12:01	xmsCommand(root)% execute ha.EditHA {syncInterval=300, mode=passive}
2012-12-04 17:45:51	xmsCommand(root)% logout
2012-12-05 09:30:31	xmsCommand(root)% login
2012-12-05 10:01:26	xmsCommand(root)% logout
2012-12-05 10:33:48	xmsCommand(root)% login
2012-12-05 11:04:36	xmsCommand(root)% logout
2012-12-05 11:34:49	xmsCommand(root)% login
2012-12-05 12:45:53	xmsCommand(root)% logout
2012-12-05 13:44:15	xmsCommand(root)% login
2012-12-05 14:16:33	xmsCommand(root)% logout
2012-12-05 18:15:14	xmsCommand(root)% login
2012-12-05 18:24:38	xmsCommand(root)% remove ha.PartnerData {objectDns=dns}
2012-12-05 18:29:27	xmsCommand(root)% remove ha.PartnerData {objectDns=dns}

▼ Delete an HA Partner

If you no longer want high availability for your Oracle Fabric Manager Servers, you can remove this functionality by deleting all the HA Partner(s). When the HA partner is deleted, it reverts to a single server running Oracle Fabric Manager:

- Configuration changes made to the previous Oracle Fabric Manager Server are no longer synchronized
- The deleted HA Partner, which is now a stand-alone Oracle Fabric Manager Server, has a baseline configuration which is the last config sent from the previous Oracle Fabric Manager Server when the HA Partner was still the HA Oracle Fabric Manager system.
- The deleted HA Partner, which is now a stand-alone Oracle Fabric Manager Server, can be used to manage I/O, just like any other Oracle Fabric Manager Server. No reboot is required

on either of the servers (the previous Oracle Fabric Manager Server or the previous HA Partner).

The HA Partner can be deleted at anytime through the High Availability Summary. To delete an HA Partner and return it to the role of a stand-alone Oracle Fabric Manager Server, follow this procedure:

- 1. From the Oracle Fabric Manager toolbar, select *HA->Configure HA Partners* to display the High Availability Summary.**
See [“Displaying HA Oracle Fabric Manager Information” on page 424](#).
- 2. Select the HA Partner that you want to delete from the HA Oracle Fabric Manager system.**
- 3. Click the garbage can icon to delete the HA Partner from the High Availability Summary.**
- 4. On the confirmation dialog, click *Yes* to complete deleting the HA Partner from the HA Oracle Fabric Manager system, or *No* to abort the deletion and leave the selected HA Partner in the HA Oracle Fabric Manager setup.**

Working With Live Monitoring

This chapter documents the following topics:

- [“Live Monitoring” on page 433](#)
- [“Displaying Host Server Throughput” on page 435](#)
- [“Display vNIC Throughput” on page 438](#)
- [“Display vHBA Throughput” on page 440](#)

Live Monitoring

Live Monitoring is a statistics grapher built into Oracle Fabric Manager GUI that supports monitoring of live, real-time statistics and usage. Live Monitoring is divided into two main sections:

- The Selected Server frame, which contains information about a single server that you select.
- The Virtual Resources frame, which is subdivided into individual charts for each vNIC or vHBA deployed on the physical server.



Live Monitoring monitors throughput and usage statistics for physical servers, vNICs, and vHBAs for specific intervals, and displays the statistics in real time.

For physical servers, the following information is displayed:

- Average throughput, which is displayed by interval in the real-time grapher
- Current usage, which is the real-time usage of the entire physical server
- Average, the average throughput for the physical server since the grapher was invoked through Oracle Fabric Manager.
- Maximum, the maximum throughput for the physical server since the grapher was invoked through Oracle Fabric Manager.

For vNICs and vHBAs, the following information is displayed:

- Ingress bandwidth usage is graphed for any vNIC or vHBA that is deployed on the selected host server. The ingress throughput is tracked in intervals just like the per-server throughput.
- Egress throughput is graphed for any vNIC or vHBA that is deployed on the selected host server. The egress throughput is tracked in intervals just like the per-server throughput.

Displaying Host Server Throughput

Through the Selected Server frame you can see the throughput of an entire selected server. This frame will be blank until you select a server.

The selected server frame contains a real-time grapher for the server's throughput and historical and calculated statistics.


Note - For each server, a Reload button allows you to flush the current grapher session and start tracking the statistics in a new set of intervals. The Reload button resets the grapher for the server and any vNICs and vHBAs deployed on it. Even though the Reload button is in the Server Throughput frame, be aware that it also reloads statistics for virtual I/O on the server.

- [“Real-Time Grapher” on page 435](#)
- [“Display Real-Time Grapher Statistics” on page 436](#)
- [“Historical and Calculated Statistics” on page 437](#)


Real-Time Grapher

The grapher tracks the throughput over an interval and displays the throughput as standard line graph, which consists of an X axis and a Y axis.

- The X axis shows an interval-based time line. By default, each interval is 50 seconds, but this value can be decreased to smaller amounts of time.
- The Y axis shows the throughput in Kbps.




The real-time grapher can display the statistics for a non-default interval time by clicking and dragging on the section of the graph that you want to “zoom in.” The following figure shows an example of zooming in on a specific statistic you want to display.



In this example, the darker blue box shows the statistics that will be displayed in a smaller interval (less than 50 seconds). This area of statistics was selected by clicking and dragging over the statistics in a top-down, left-to-right motion. Notice that the intervals at this point are 50 seconds, for example from 14:33:20 to 14:34:10.

When the click and drag completes, the selected statistics area is zoomed in, so that a smaller interval is displayed and a more granular focus is given to the statistics.



▼ Display Real-Time Grapher Statistics

To display the Selected Server throughput, follow this procedure:


1. **From the Navigation Frame, display Live Monitoring by selecting Service Manager → Live Monitoring.**

Live Monitoring is blank by default until you select a physical server.

2. **From the dropdown menu, select a server for Live Monitoring.**
After a brief amount of time, the Live Monitoring Summary is displayed.
See [“Live Monitoring” on page 433](#) for an example.
3. **As an option, click and drag left-to-right in a downward motion to draw a box over the statistics that you want to zoom in on.**
You can click and drag right-to-left in an upward motion to zoom out.

Historical and Calculated Statistics

Based on the monitored throughput, Oracle Fabric Manager displays additional information for the selected server in a table on the side of a graph. The additional information is tracked for historical information, unless documented otherwise. The following figure shows the table of historical and calculated statistics.



The following additional information is displayed.

- the aggregate amount of throughput for the current graphing session (Total Current).
- the maximum amount of through for the current graphing session (Total Maximum).
- the average amount of throughput for the current graphing session (Total Average). This statistic is calculated.
- the amount of network throughput for the current graphing session (Network Current).
- the maximum amount of network throughput for the current graphing session (Network Maximum).
- the average amount of network throughput for the current graphing session (Network Average). This statistic is calculated.
- the amount of storage throughput for the current graphing session (Storage Current).
- the maximum amount of storage throughput for the current graphing session (Storage Maximum).
- the average amount of storage throughput for the current graphing session (Storage Average).

vNIC Throughput

When vNICs are configured and deployed on the host server, Live Monitoring can graph the throughput of each vNIC through the I/O Resource Charts. Each chart is a separate instance of a vNIC that graphs the ingress and egress traffic. Each I/O Resource Chart is a standard line graph consisting of an X axis and a Y axis:

- The X axis shows an interval-based time line. By default, each interval is 50 seconds, but this value can be decreased to smaller amounts of time.
- The Y axis shows the throughput in Kbps.

Each Chart has the ability to track throughput for one or both directions of traffic (ingress only, egress only, or both ingress and egress) depending on what traffic exists on the virtual resource.

The following figure shows an example of the I/O Resource Charts for vNICs.



You'll notice that each vNIC has its own chart, which is identified by the vNIC name. Also, each vNIC I/O chart can be zoomed in to show the statistics for a smaller interval. To zoom in, click and drag in a downward motion to create a box over the statistics you want to zoom in on. When you complete the click and drag motion, the interval becomes smaller and the graphed section of statistics becomes larger, as documented in [“Display Real-Time Grapher Statistics”](#) on page 436.

▼ Display vNIC Throughput

The vNIC I/O Resource charts are contained in a separate frame on the Live Monitoring summary. To display the Live Monitoring Summary, follow this procedure:

1. **From the Navigation Frame, display Live Monitoring by selecting Service Manager → Live Monitoring.**

Live Monitoring is blank by default until you select a physical server.

2. From the dropdown menu, select a server for Live Monitoring.

The Live Monitoring Summary is displayed.



If enough vNIC are deployed, scroll bars appear on the I/O Resource Charts frame. You might need to scroll the frame to see all the deployed vNICs. Also, if vHBAs are also present, they will be displayed below the vNICs. If you want to see the vHBAs, you might need to scroll the frame to see all the deployed vHBAs.

3. As an option, click and drag left-to-right in a downward motion to draw a box over the vNIC statistics that you want to zoom in on.

You can click and drag right-to-left in an upward motion to zoom out.

vHBA Throughput

When vHBAs are configured and deployed on the host server, Live Monitoring can graph the throughput of each vHBA through the I/O Resource Charts. Each chart is a separate instance of a vHBA that graphs the ingress and egress traffic. Each I/O Resource Chart is a standard line graph consisting of an X axis and a Y axis:

- The X axis shows an interval-based time line. By default, each interval is 50 seconds, but this value can be decreased to smaller amounts of time.
- The Y axis shows the throughput in Kbps.

Each chart has the ability to track throughput for one or both directions of traffic (ingress only, egress only, or both ingress and egress) depending on what traffic exists on the virtual resource.

The following figure shows the vHBA I/O Resource Charts.



You'll notice that each vHBA has its own chart, which is identified by the vHBA name. Also, each vHBA I/O chart can be zoomed in to show the statistics for a smaller interval. To zoom in, click and drag in a downward motion to create a box over the statistics you want to zoom in on. When you complete the click and drag motion, the interval becomes smaller and the graphed section of statistics becomes larger, as documented in [“Display Real-Time Grapher Statistics” on page 436](#).

▼ Display vHBA Throughput

The vHBA I/O Resource charts are contained in a separate frame on the Live Monitoring summary. To display the Live Monitoring Summary, follow this procedure:

1. **From the Navigation Frame, display Live Monitoring by selecting Service Manager → Live Monitoring.**

Live Monitoring is blank by default until you select a physical server.

2. From the dropdown menu, select a server for Live Monitoring.

The Live Monitoring Summary is displayed. See [“Live Monitoring” on page 433](#) for an example.

If enough vHBAs are deployed, scroll bars appear on the I/O Resource Charts frame. You might need to scroll the frame to see all the deployed vHBAs. Also, if vNICs are also present, they will be displayed above the vHBAs. If you want to see the vNICs, you might need to scroll the frame to see all the deployed vNICs.

3. As an option, click and drag left-to-right in a downward motion to draw a box over the vHBA statistics that you want to zoom in on.

You can click and drag right-to-left in an upward motion to zoom out.

Working With the Task Scheduler

This chapter contains the following topics:

- “Understanding Task Scheduler” on page 443
- “Displaying Oracle Fabric Manager Schedules” on page 449
- “Run an On-Demand Oracle Fabric Manager Backup” on page 454
- “Enable or Disable Oracle Fabric Manager Backup Schedules” on page 455
- “Create Fabric Device Backup Schedules” on page 456
- “Displaying Fabric Device Backup Schedules” on page 458
- “Run an On-Demand Fabric Device Backup” on page 462
- “Enable or Disable Fabric Device Backup Schedules” on page 463

Understanding Task Scheduler

- “Task Scheduler” on page 443
- “Oracle Fabric Manager Backups Considerations” on page 444
- “Backup and Restore of Fabric Device Through Oracle Fabric Manager” on page 445
- “Considerations for Fabric Device Backups” on page 446
- “Fabric Device Backup With HA Fabric Devices” on page 447
- “Scheduled Backups” on page 447
- “On-Demand Backups” on page 448

Task Scheduler

Oracle Fabric Manager Task Scheduler allows you to set times and dates for backing up Oracle Fabric Manager Server configurations and Oracle Fabric Device configurations. Backing up a Oracle Fabric Manager Server or a Fabric Device can occur either on-demand or on a recurring schedule. Regardless of the schedule type, backups are stored in the following locations on the Oracle Fabric Manager Server:

- Linux or Oracle Solaris Oracle Fabric Manager Server:
 - Oracle Fabric Manager backups are stored in `/opt/xsigo/xms/xms-backups`
 - Fabric Interconnect backups are stored in `/opt/xsigo/xms/director-backups`

- Oracle SDN Controller backups are stored in /opt/xsigo/xms/director-backups
- Windows Oracle Fabric Manager Server:
 - Oracle Fabric Manager backups are stored in C:\Program Files\xms\xms-backups
 - Fabric Interconnect backups are stored in C:\Program Files\xms\director-backups
 - Oracle SDN Controller backups are stored in C:\Program Files\xms\director-backups

Each backup captures the entire configuration of the Oracle Fabric Manager Server or Fabric Device.

Oracle Fabric Manager Backups Considerations

Oracle Fabric Manager Backups can occur in multiple ways. Being aware of these considerations will facilitate your backup and restore operations for Oracle Fabric Manager Servers:

- Through the Backup Oracle Fabric Manager Configuration option on the Maintenance menu (Maintenance->Backup Oracle Fabric Manager Config), which is represented by the screwdriver icon on the Oracle Fabric Manager banner.
- Through a Schedule backup (Schedules->Oracle Fabric Manager Backups)
- Through an on-demand backup (Schedules->Oracle Fabric Manager Backups->Execute Now)

In all cases the same information is backed up, the Oracle Fabric Manager Server configuration and all objects that Oracle Fabric Manager has under management. However, there are some subtle differences between each method of Oracle Fabric Manager backup:

- The difference between Maintenance->Backup Oracle Fabric Manager Config and Schedules->Oracle Fabric Manager Backups->Execute Now, is that Maintenance->Backup Oracle Fabric Manager Config requires you to specify a file name whereas Schedules->Oracle Fabric Manager Backups->Execute Now requires a schedule to exist and executes the scheduled backup immediately, not at the scheduled time(s).
- The file names are different depending on which method of backup you choose:
 - If you perform Oracle Fabric Manager backup through Maintenance->Backup Oracle Fabric Manager Config, Oracle Fabric Manager prompts you for a specific file name, so the file will be named <name> plus a date and time stamp—for example, foobar_2011-04-05_04_01_59_325
 - If you perform the Oracle Fabric Manager Backup through Schedules->Oracle Fabric Manager Backups, Oracle Fabric Manager creates its own file name in the format scheduled_ plus a name, plus a time and date stamp—for example, scheduled_foobar_2011_04_05_04_01_59_325
 - If you perform the backup through Schedules->Oracle Fabric Manager Backups->Execute Now button, Oracle Fabric Manager creates its own file name in the format scheduled_executeNow_ plus a name, plus a time and date stamp—for example, scheduled_executeNow_foobar_2011_04_05_04_01_59_325

Regardless of the method you choose for Oracle Fabric Manager Server backups, the backup files are stored as XML files and can be restored through Maintenance->Restore Oracle Fabric Manager Server Config.

When you need to restore from a specific backup, you will find it helpful to know the different file naming methods listed above so that you will be able to easily restore the specific Oracle Fabric Manager Server backup you want.

Simultaneous Oracle Fabric Manager and Fabric Device Backups

When you need to do a backup of either Oracle Fabric Manager or the Fabric Device, it is a best practice to take a backup of both at the same time. The reason for this requirement is objects will not get out of sync, thereby restore operations will occur without problems. Take a backup of all Fabric Devices first, then take the Oracle Fabric Manager backups.

When restoring a saved backup, Oracle requires that you restore both the Fabric Device and the Oracle Fabric Manager Server at the same time. Restore the Oracle Fabric Manager backups first, then restore the Fabric Devices.

Backup and Restore of Fabric Device Through Oracle Fabric Manager

Oracle Fabric Manager supports backup and restore of either the Oracle Fabric Manager configuration or the Fabric Device configuration. The backup option for the Fabric Device captures information at a single point in time, and the restore option for the Fabric Device brings back the configuration that was present at the last backup.

Note - The backup and restore options for Oracle Fabric Manager affect only the Oracle Fabric Manager Server and Oracle Fabric Manager configuration. Restoring the Oracle Fabric Manager configuration does not bring back the configurations of any Fabric Devices that Oracle Fabric Manager is managing. Instead, only the Oracle Fabric Manager database is backed up and restored.

Oracle Fabric Manager's backup option for the Fabric Device operates the same as issuing the system export command on the Fabric Device. The backup option captures the following information at the time that you initiate the Fabric Device backup through Oracle Fabric Manager:

- Hardware inventory of which modules are present in which slots. Any hardware changes made after the backup completes are not written to the last backup, and the changes are not automatically captured. You can always take another backup after the hardware changes are

complete, if needed. If you take an additional backup, the previous backup is not deleted. The previous backup is kept along with new one.

- Administrative states of all objects in the Fabric Device. Operational states are not captured in the backup.

Oracle Fabric Manager's restore option for the Fabric Device operates the same as issuing the system import command on the Fabric Device. The restore option brings back the hardware inventory and administrative states contained in the last backup.

Some common use cases for Oracle Fabric Manager backup and restore of a Fabric Device are:

- Restoring deleted vNICs or vHBAs, but only if no changes have been made to the vNICs. For example, assume you created some vNICs through Oracle Fabric Manager and took a Fabric Device backup. If someone else then deletes the vNICs, you can restore the configuration on the Fabric Device to replace the deleted vNICs provided that no other changes have been made to the vNICs. If changes have been made, the vNICs might not be restored correctly.
- Restore a clean Fabric Device configuration. Assume a Fabric Device has been returned to its default configuration—for example, through the system clear config command. You can then restore the last backed up configuration instead of having to re-create the entire configuration.

Considerations for Fabric Device Backups

Fabric Device backups can occur in multiple ways. Being aware of these considerations will facilitate your backup and restore operations for Fabric Devices:

- Through the Backup Fabric Device configuration toolbar button on the summary page for whichever Fabric Device you have configured.
- Through a Scheduled backup
- Through an on-demand backup

Note - For information about Oracle SDN Controller backups, see the Oracle SDN User Guide.

In all cases the same information is backed up. However, there are some subtle differences between each method of Fabric Device backup:

- The difference between Managed Devices->Fabric Interconnects->Backup Fabric Interconnect Config button and Schedules-> Fabric Interconnect Backups->Execute Now button, is that Managed Devices->Fabric Interconnects->Backup Fabric Interconnect Config button requires you to specify a file name whereas Schedules->Director Backups->Execute Now button requires a schedule to exist and executes the scheduled backup immediately, not at the scheduled time.

- The file names are different depending on which method of backup you choose:
 - If you perform a Fabric Interconnect backup through Managed Devices->Fabric Interconnects->Backup Fabric Interconnect Config button, Oracle Fabric Manager prompts you for a file name, so the file will be named <director-name> plus <file-name> plus a date and time stamp— for example, destroyer_backup1_2011-04-05_04_01_59_325
 - If you perform a Fabric Interconnect backup through Schedules->Director Backups, Oracle Fabric Manager creates its own file name in the format <director-name>_scheduled_ plus <name>, plus a time and date stamp—for example, destroyer_scheduled_backup1_2011_04_05_04_01_59_325
 - If you perform a Fabric Interconnect backup through Schedules->Director Backups->Execute Now button, Oracle Fabric Manager creates its own file name in the format <director-name>_scheduled_executeNow_ plus <name>, plus a time and date stamp—for example, destroyer_scheduled_executeNow_destroyer_2011_04_05_04_01_59_325

Regardless of the method you choose for Fabric Device backups, the backup files are stored as XML files and can be restored through Managed Devices->Fabric Interconnects->Restore Fabric Interconnect Config button.

Note - You should take both types of backup simultaneously as documented in [“Simultaneous Oracle Fabric Manager and Fabric Device Backups” on page 445](#).

When you need to restore from a specific backup, you will find it helpful to know the different file naming methods listed above so that you will be able to easily restore the specific Fabric Interconnect you want.

Fabric Device Backup With HA Fabric Devices

When multiple Fabric Devices are deployed, they share a common Ethernet or InfiniBand switching fabric. In a dual-Fabric Device deployment, there is no method of synchronization for backups. As a result, each Fabric Device should be configured with its own Fabric Device backup schedule. Any on-demand backups either through Managed Devices->Fabric Interconnects->Backup Fabric Interconnect Config button or Schedules->Director Backups->Execute Now button should be manually duplicated across both Fabric Devices to keep their backups in sync.

Scheduled Backups

When you configure the task scheduler, you have the option of running backups either daily, weekly, or monthly:

- daily backups occur every day at time(s) of day that you specify. You can specify more than one scheduled backup time for the same day. For example, it would be a valid configuration

to set one daily schedule to run at 6:00 a.m., and also set another schedule that runs at midnight.

- weekly backups occur each week on the day(s) that you specify. You can specify multiple days each week. For example, it would be a valid configuration to set one weekly schedule to run once a week on Monday and another weekly scheduled to run once a week on Friday.
- monthly backups occur each month on the day(s) that you specify at the time(s) of day that you specify. You can specify multiple days each month. For example, it would be a valid configuration to set one monthly schedule to run on the 1st of each month and another monthly schedule to run on the 14th of each month.

When you configure a schedule, you will set the maximum number of backups that will be kept. The maximum number of backups is a sliding window based on the number that you specify, and only the most recent number of backups is kept. For example, if you specify 5, the first 5 backups are kept. When 6 or more backups occur, the oldest are deleted so that only the newest 5 backups are kept. This sliding windows allows the most recent data to be kept, and also prevents the backups directory from becoming full and consuming an unreasonable amount of disk space.

Scheduled backups complete with silent acknowledgement, so unless an error message is displayed, the scheduled backup has completed successfully.

For information about running scheduled backups, see the appropriate section:

- To backup Oracle Fabric Manager Servers based on a recurring schedule, see [“Enable or Disable Oracle Fabric Manager Backup Schedules” on page 455](#)
- To backup Fabric Devices based on a recurring schedule, see [“Enable or Disable Fabric Device Backup Schedules” on page 463](#)

On-Demand Backups

On-demand backups can be run at any time. When you run them, they complete immediately.

On-demand backups and scheduled backups are not mutually exclusive, so any scheduled backups will run in addition to the on-demand backup. For example, if you have a daily backup set for 4:00 a.m., you can run an on-demand backup at any time, and the scheduled backup still runs at 4:00 a.m. On-demand backups complete with silent acknowledgement, so unless an error message is displayed, the on-demand backup has completed successfully. You can, however, check the Execution History tab for a timestamp that is close to the time(s) when you started the backup. The presence of the backup in the Execution History verifies that the backup completed successfully. For information about running an on-demand backup, see the appropriate section:

- To backup up a Oracle Fabric Manager Server on-demand, see [“Run an On-Demand Oracle Fabric Manager Backup” on page 454](#)
- To backup a Fabric Device on-demand, see [“Run an On-Demand Fabric Device Backup” on page 462](#)

Displaying Oracle Fabric Manager Schedules

- [“Oracle Fabric Manager Backup Schedule Summary” on page 449](#)
- [“Edit Oracle Fabric Manager Backup Schedules” on page 450](#)
- [“Rename Oracle Fabric Manager Backup Schedules” on page 451](#)
- [“Create Oracle Fabric Manager Backup Schedules” on page 452](#)

Oracle Fabric Manager Backup Schedule Summary

Oracle Fabric Manager displays the backup schedules for the Oracle Fabric Manager Server in the Oracle Fabric Manager Backup Schedule Summary, which is a tabbed display. The Oracle Fabric Manager Backup Schedule Summary contains a list of all the configured backup schedules, regardless of their operational state (enabled or disabled).

The Oracle Fabric Manager Backup Schedule Summary is available through Schedules option on the Navigation Panel. The following example shows the Oracle Fabric Manager Backup Schedule Summary tab.

Name	Enabled	Schedule	Max Backups	Description
lorenzo	true	At 12:45 AM every day	14	Daily backup at 12:45 a.m. Max 14 for FabMan server "lor..."

The Oracle Fabric Manager Backup Schedule Summary contains information about configured backup schedules for Oracle Fabric Manager Server backup.

Field...	Means...
Name	The name assigned to the specific Oracle Fabric Manager backup schedule.
Enabled	The operational state (either enabled or disabled) for a particular Oracle Fabric Manager backup schedule.
Schedule	The specific day(s), time(s), or date(s) on which the backup will run.

Field...	Means...
Max Backups	The maximum number of backups that are allowed for each particular Oracle Fabric Manager Backup Schedule.
Description	An optional alphanumeric string that describes a Oracle Fabric Manager Backup Schedule.

▼ Edit Oracle Fabric Manager Backup Schedules

When a backup schedule is created for Oracle Fabric Manager, it can be edited at any time. When the edits are completed and the details frame is locked for editing, the new changes take effect based on the new properties. For example, if you had a daily backup schedule set to occur at 7:00 a.m., but decided to set the backup to occur at 3:00 a.m. instead, the new backup will occur at 3:00 a.m. the next morning after you complete edits. Of course, you can always use the on-demand backup feature, to complete an instantaneous backup at any time.

To edit a configured Oracle Fabric Manager backup schedule, follow this procedure:

1. **Display the Oracle Fabric Manager Backup Schedule Summary.**
2. **In the Oracle Fabric Manager Backup Schedule Summary, click the backup schedule that you want to edit.**

This step populates the Details frame with the properties of the selected backup schedule.

The screenshot shows two frames from the Oracle Fabric Manager interface. The top frame, titled "Oracle Fabric Manager Backup Schedule Summary", contains a table with the following data:

Name	Enabled	Schedule	Max Backups	Description
lorenzo	true	At 12:45 AM every day	14	Daily backup at 12:45 a.m.. Max 14 for FabMan server "lor..."

The bottom frame, titled "Fabric Manager Backup Schedule : lorenzo", shows the details for the selected schedule. The "General" tab is active, displaying the following properties:

- Name:** lorenzo
- Enabled:** true
- Schedule:** At 12:45 AM every day
- Max Backups:** 14
- Description:** Daily backup at 12:45 a.m.. Max 14 for FabMan server "lorenzo"

An "Edit" button is located at the bottom left of the details frame and is highlighted with a red rectangular box.

3. **In the Details frame, click the Edit button to unlock the properties for editing.**

This example shows editing a daily backup schedule, but similar properties (name, enable/disable, backup schedule times and dates, maximum number of backups, and description) are editable regardless of whether the backup schedule is daily, weekly, or monthly.

The screenshot shows the 'Fabric Manager Backup Schedule : lorenzo' window with the 'General' tab selected. The configuration is as follows:

- Name:** lorenzo
- Enabled:**
- Schedule:**
 - Hour: 12
 - Minute: 45
 - AM/PM: AM
 - Frequency: Daily, Weekly, Monthly
- Max Backups:** 14
- Description:** Daily backup at 12:45 a.m.. Max 14 for FabMan server "lorenzo"

Buttons for 'Submit' and 'Cancel' are located at the bottom left of the window.

4. **Edit the properties as needed.**
5. **When the properties are set correctly, click Submit.**
The new properties do not take effect unless the Submit button is clicked.

▼ Rename Oracle Fabric Manager Backup Schedules

Oracle Fabric Manager supports renaming a Oracle Fabric Manager Schedule, which enables you to change the name without having to completely delete and recreate the entire Oracle Fabric Manager Schedule. When the Oracle Fabric Manager Schedule is renamed, all other properties for the Oracle Fabric Manager Schedule are retained, including the schedule and server(s) affected. As an option, you can also set or change the description for the Oracle Fabric Manager Schedule.

You can rename the Oracle Fabric Manager Schedule through the Oracle Fabric Manager Schedule Details frame. To rename the Oracle Fabric Manager Schedule, follow this procedure:

1. **On the Navigation Frame, select Service Manager → Schedules.**
See [“Oracle Fabric Manager Backup Schedule Summary” on page 449](#) for an example.
2. **Click the backup schedule that you want to rename.**

This step populates the details frame with its properties.

- 3. Click the Edit button to edit the properties of the selected Oracle Fabric Manager Schedule.**

See [“Edit Oracle Fabric Manager Backup Schedules” on page 450](#) for an example.

- 4. In the Name field, enter the new name for the Oracle Fabric Manager Schedule.**
- 5. As an option, you also can set or change the description for the selected Oracle Fabric Manager Schedule.**
- 6. When the new name has been specified for the Oracle Fabric Manager Schedule, click Submit.**

▼ Create Oracle Fabric Manager Backup Schedules

Through Oracle Fabric Manager, you can create a set of backup schedules for the Oracle Fabric Manager Server(s) in your data center. A Oracle Fabric Manager backup schedule sets time (s) and date(s) on which the Oracle Fabric Manager Server configuration is copied to a file. Each time the backup runs, the entire configuration is saved, not just the differences from an established baseline.

To backup the Oracle Fabric Manager Server config, you set a schedule for when the config is backed up and copied to the xms-backups directory. The following backup schedules are supported, and you can configure more than one type for the same Oracle Fabric Manager Server:

- Daily
- Weekly
- Monthly

To create a Oracle Fabric Manager backup schedule, follow this procedure, which uses a weekly backup schedule for illustrative purposes:

- 1. Display the Oracle Fabric Manager Backup Schedule Summary.**

See [“Oracle Fabric Manager Backup Schedule Summary” on page 449](#) for an example.

- Click the plus sign to display the New Oracle Fabric Manager backup schedule dialog.

New Oracle Fabric Manager Backup Schedule

Name: * lorenzo

Enable:

Job Schedule:

Hour: 12 **Minute:** 45 **AM/PM:** AM

Daily Weekly Monthly

Max Backups: * 14

Description: Daily bakcup at 12:45 a.m.. Max 14 for FabMan server "lorenzo"

Submit Cancel

- In the Name field, enter an alphanumeric character string that identifies the backup schedule.
- In the Enable checkbox, make sure that the backup schedule is enabled.
- Or, as an option, if you want to create the schedule without it being active, you can click to deselect the checkbox. When the checkbox is empty, the backup schedule is disabled.
- In the Job Schedule section, select the hour, minute, and a.m. or p.m. designation for when the backup will run.
You must enter additional information if you are creating a weekly or monthly backup schedule:
- For a weekly schedule, you will also need to specify the day(s) on which the backup will run at the specified time(s) (hour, minute, a.m. or p.m.).
- For a monthly schedule, you will also need to select a date(s) of the month on which the backup will run at the specified time(s). The date(s) you select is

applied to all months, so be aware of the implications of the date(s) that you select.

For example, selecting the 30th of the month is acceptable for months that have 30 or 31 days. However, in February, the backup will not occur on the 30th since that month does not have 30 days. In this example, the backup would not run in February, but would resume in March.

- 9. In the Max Backups field, enter a number from 1 to 99999 that quantifies the maximum number of backups that will occur for the backup schedule that you are creating. When the maximum number of backups is reached, the sliding window keeps only the last X backups.**

For example, if you have daily backups, and are keeping a maximum of 5, on the sixth day, the first backup is deleted, and the last 5 are kept.

- 10. As an option, in the Description field, enter a description for the backup schedule that you are configuring.**
- 11. When the backup schedule properties are configured, click Submit. The schedule is not created until you click Submit.**

▼ Run an On-Demand Oracle Fabric Manager Backup

An on-demand Oracle Fabric Manager Backup is one that you manually start. On-demand backups can be completed at any time, and they are in addition to any scheduled backups.

An on-demand backup runs for a configured backup schedule, just not at the scheduled time(s). As a result, a backup schedule must exist for the Oracle Fabric Manager Server that you want to backup on an on-demand basis.

You can run an on-demand Oracle Fabric Manager backup through the Oracle Fabric Manager Backup Schedule Summary tab.

To run an on-demand Oracle Fabric Manager backup, follow this procedure:

- 1. Display the Oracle Fabric Manager Backup Schedule Summary.**
See [“Oracle Fabric Manager Backup Schedule Summary” on page 449](#) for an example.
- 2. Click the backup schedule that you want to run on an on-demand basis. This step activates the Execute Now button, which is the green arrow on the Oracle Fabric Manager Backup Schedule Summary toolbar.**
- 3. Click the Execute Now button to begin the backup.**
- 4. On the popup dialog, click Yes to continue.**

5. **Check the Recent Jobs Summary to verify that the backup job has completed successfully.**

▼ **Enable or Disable Oracle Fabric Manager Backup Schedules**

Oracle Fabric Manager Backup Schedules are in either of the following states:

- Enabled, in which the scheduled backup will run at its designated date(s) and time(s) until the maximum number of backups has occurred.
- Disabled, in which the scheduled backup will not run until explicitly re-enabled. Even if a scheduled backup is disabled, you still can run an on-demand backup as needed.

By default, when you create a Oracle Fabric Manager Backup Schedule, it is enabled. However, you can manually disable a Oracle Fabric Manager Backup schedule, or re-enable a disabled Backup Scheduled. If you will be enabling or disabled Oracle Fabric Manager Backup Schedules, you must do so individually. Oracle Fabric Manager does not currently support a “disable all” or “enable all” option for Oracle Fabric Manager Server Backup Schedules.

You can enable or disable individual Oracle Fabric Manager backup Schedules by editing a currently configured Oracle Fabric Manager Backup Schedule:

1. **Display the Oracle Fabric Manager Backup Schedule Summary.**
See [“Oracle Fabric Manager Backup Schedule Summary” on page 449](#) for an example.
2. **In the Oracle Fabric Manager Backup Schedule Summary, click the backup schedule that you want to edit.**
This step populates the Details frame with the properties of the selected backup schedule.
See [“Edit Oracle Fabric Manager Backup Schedules” on page 450](#) for an example.
3. **In the Details frame, click the Edit button to unlock the properties for editing.**

This example shows editing a daily backup schedule, but similar properties (name, enable/disable, backup schedule times and dates, maximum number of backups, and description) are editable regardless of whether the backup schedule is daily, weekly, or monthly.

The screenshot shows the 'Fabric Manager Backup Schedule : lorenzo' window. The 'General' tab is active. The 'Name' field contains 'lorenzo'. The 'Enabled' checkbox is checked and highlighted with a red box. The 'Schedule' section shows 'Hour' set to 12, 'Minute' set to 45, and 'AM/PM' set to AM. The 'Daily' radio button is selected. The 'Max Backups' field is set to 14. The 'Description' field contains the text: 'Daily backup at 12:45 a.m.. Max 14 for FabMan server "lorenzo"'. At the bottom, there are 'Submit' and 'Cancel' buttons.

4. Click the Enabled checkbox to set the appropriate state for the backup schedule:
5. Enabled state occurs when the Enabled checkbox contains a check mark. This checkbox is a toggle, so each click alternates between Enabled and Disabled.
6. Disabled state occurs when the Enabled checkbox does not contain a check mark.
This checkbox is a toggle, so each click alternates between Enabled and Disabled.
7. When the appropriate properties have been edited, click Submit to save the changes.

▼ Create Fabric Device Backup Schedules

Through Oracle Fabric Manager, you can create a set of backup schedules for the Fabric Device(s) in your data center. A Fabric Device backup schedule sets a time(s) and date(s) on which the Fabric Device configuration is copied to a file. Each time the backup runs, the entire configuration is saved, not just the differences from an established baseline.

Also, you can keep a specific number of backups by setting the maximum number of backups to keep. This maximum number is a sliding window, so that only the last X number of backups is kept. For example, if you are doing weekly backups, and set a maximum of 12 backups, on the 13th week that backup is deleted and only the last 12 are kept.

To backup the Fabric Device config, you set a schedule for when the config is backed up and copied to the director-backups directory. The following backup schedules are supported, and you can configure more than one type for the same Fabric Device:

- Daily
- Weekly
- Monthly

To create a Fabric Device backup schedule, follow this procedure, which uses a weekly backup schedule for illustrative purposes:

1. Display the Fabric Device Schedule Summary.

See [“Oracle Fabric Manager Backup Schedule Summary” on page 449](#) for an example.

2. Click the plus sign to displayed the New Fabric Device Schedule dialog.

New Fabric Device Backup Schedule

Name: * DailyDelaware

Fabric Device: * delaware

Enable:

Job Schedule:

Hour: 05 **Minute:** 00 **AM/PM:** AM

Daily Weekly Monthly

Max Backups: * 7

Description: Daily backup for Fabric Device "delaware" @ 5:00 a.m.

Submit Cancel

3. In the Name field, enter an alphanumeric character string that identifies the backup schedule.
4. From the Fabric Interconnect dropdown menu, select the Fabric Interconnect for which you will be creating the backup schedule. If the Fabric Interconnect is not displayed, it has not yet been discovered by Oracle Fabric Manager.
5. In the Enable checkbox, make sure that the backup schedule is enabled.
6. Or, as an option, if you want to create the schedule without it being active, you can click to deselect the checkbox. When the checkbox is empty, the backup schedule is disabled.
7. In the Job Schedule section, select the hour, minute, and a.m. or p.m. designation for when the backup will run. You must enter additional information if you are creating a weekly or monthly backup schedule:
8. For a weekly schedule, you will also need to specify the day(s) on which the backup will run at the specified time(s) (hour, minute, a.m. or p.m.).
9. For a monthly schedule, you will also need to select a date(s) of the month on which the backup will run at the specified time(s). The date(s) you select is applied to all months, so be aware of the implications of the date(s) that you select. For example, selecting the 30th of the month is acceptable in months that have 30 or 31 days. However, in February, the backup will not occur on the 30th since that month does not have 30 days. In this example, the backup would not run in February, but would resume in March.
10. In the Max Backups field, enter a number from 1 to 99999 that quantifies the maximum number of backups that will occur for the backup schedule that you are creating. The maximum backups number you enter sets the size of the sliding window, so that only the last X number of backups is kept. For example, if you are doing daily backups, and have set a maximum of 5, on the sixth day that backup is deleted so that only the most recent 5 backups are kept.
11. As an option, in the Description field, enter a description for the backup schedule that you are configuring.
12. When the backup schedule properties are configured, click Submit. The schedule is not created until you click Submit.

Displaying Fabric Device Backup Schedules

This sections contains the following topics:

- [“Fabric Device Backups” on page 459](#)
- [“Edit Fabric Device Backup Schedules” on page 460](#)
- [“Rename Fabric Device Backup Schedules” on page 461](#)

Fabric Device Backups

Oracle Fabric Manager displays the backup schedules for the Fabric Device in the Fabric Device Backup Schedule Summary, which is a tabbed display. The Fabric Device Schedule Summary contains a list of all the configured backup schedules, regardless of their operational state (enabled or disabled).

The Fabric Device Schedule Summary is available through Schedules option on the Navigation Panel.

Name	Fabric Device	Enabled	Schedule	Max Backups	Description
DailyDelaware	delaware	true	At 5:00 AM every day	7	Daily backup for Fabric Device "delaware" @ 5...

The Fabric Device Schedule Summary contains information about configured backup schedules for Fabric Interconnects and SDN controllers or fabric devices.

Field...	Means...
Name	The name assigned to the specific Fabric Device Backup schedule.
Fabric Device	The name of the Fabric Device that will be backed up.
Enabled	The operational state (either enabled or disabled) for a particular Fabric Device Backup schedule.
Schedule	The specific day(s), time(s), or date(s) on which the backup will run.
Max Backups	The maximum number of backups that are allowed for each particular Fabric Device Backup Schedule.
Description	An optional alphanumeric string that describes a Fabric Device Backup Schedule.

▼ Edit Fabric Device Backup Schedules

When a backup schedule is created for a Fabric Device, it can be edited at any time. When the edits are completed and the details frame is locked for editing, the new changes take effect based on the new properties. For example, if you had a daily backup schedule set to occur at 7:00 a.m., but decided to set the backup to occur at 3:00 a.m. instead, the new backup will occur at 3:00 a.m. the next morning after you complete edits. Of course, you can always use the on-demand backup feature, to complete an instantaneous backup at any time.

To edit a configured Fabric Device Backup schedule, follow this procedure:

1. **Display the Fabric Device Backup Schedule Summary.**
See [“Oracle Fabric Manager Backup Schedule Summary” on page 449](#) for an example.
2. **In the Fabric Device Backup Schedule Summary, click the backup schedule that you want to edit.**

This step populates the Details frame with the properties of the selected backup schedule.

Name	Fabric Device	Enabled	Schedule	Max Backups	Description
DailyDelaware	delaware	true	At 5:00 AM every day	7	Daily backup for Fabric Device "delaware" @ 5:00 a.m.

1 item

Fabric Device Backup Schedule : DailyDelaware

General

Name: DailyDelaware
Fabric Device: delaware
Enabled: true
Schedule: At 5:00 AM every day
Max Backups: 7
Description: Daily backup for Fabric Device "delaware" @ 5:00 a.m.

Edit

3. **In the Details frame, click the Edit button to unlock the properties for editing.**

This example shows editing a daily backup schedule, but similar properties (name, enable/disable, backup schedule times and dates, maximum number of backups, and description) are editable regardless of whether the backup schedule is daily, weekly, or monthly.

Fabric Device Backup Schedule : DailyDelaware

General

Name: * DailyDelaware

Fabric Device: * delaware

Enabled:

Schedule:

Hour	Minute	AM/PM
05	00	AM

Daily Weekly Monthly

Max Backups: * 7

Description: Daily backup for Fabric Device "delaware" @ 5:00 a.m.

Submit Cancel

4. **Edit the properties as needed.**
5. **When the properties are set correctly, click Submit.**

▼ Rename Fabric Device Backup Schedules

Oracle Fabric Manager supports renaming a Fabric Device Backup Schedule, which enables you to change the name without having to completely delete and recreate the entire Fabric Device Backup Schedule. When the Fabric Device Backup Schedule is renamed, all other properties for the Fabric Device Backup Schedule are retained, including the schedule and Fabric Device(s) affected. As an option, you can also set or change the description for the Fabric Device Backup Schedule.

You can rename the Fabric Device Backup Schedule through the Fabric Device Backup Schedule Details frame. To rename the Fabric Device Backup Schedule, follow this procedure:

1. **On the Navigation Frame, select Service Manager → Schedules.**

2. **Select a Fabric Device Backup Schedule to populate the details frame with its properties.**
3. **Click the Edit button to edit the properties of the selected Fabric Device Backup Schedule.**

Fabric Device Backup Schedule : DailyDelaware

General

Name: * DailyDelaware

Fabric Device: * delaware

Enabled:

Schedule:

Hour	Minute	AM/PM
05	00	AM

Daily Weekly Monthly

Max Backups: * 7

Description: Daily backup for Fabric Device "delaware" @ 5:00 a.m.

Submit Cancel

4. **In the Name field, enter the new name for the Fabric Device Backup Schedule.**
5. **As an option, you also can set or change the description for the selected Fabric Device Backup Schedule.**
6. **When the new name has been specified for the Fabric Device Backup Schedule, click Submit.**

▼ Run an On-Demand Fabric Device Backup

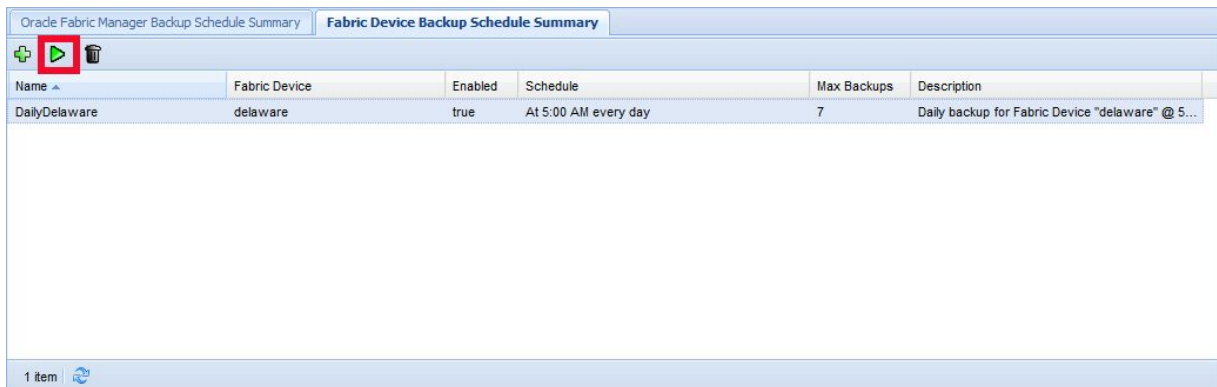
An on-demand Fabric Device Backup is one that you manually start. On-demand backups can be completed at any time, and they are in addition to any scheduled backups.

An on-demand backup runs for a configured backup schedule, just not at the scheduled time(s). As a result, a backup schedule must exist for the Fabric Device that you want to backup on an on-demand basis.

You can run an on-demand Fabric Device backup through the Fabric Device Backup Schedule Summary tab.

To run an on-demand Fabric Interconnect backup, follow this procedure:

1. **Display the Fabric Interconnect Backup Schedule Summary, available through the Fabric Device Backup Schedule Summary.**
See [“Oracle Fabric Manager Backup Schedule Summary” on page 449](#) for an example.
2. **Select the backup schedule that you want to run on an on-demand basis.**
This step activates the Execute Now button, which is the green arrow on the Fabric Interconnect Backup Schedule Summary toolbar.



3. **Click the Execute Now button to begin the backup.**
4. **On the popup dialog, click Yes to continue.**
5. **Check the Recent Jobs Summary to verify that the operation completed successfully.**

▼ Enable or Disable Fabric Device Backup Schedules

Fabric Device backup schedules are in either of the following states:

- Enabled, in which the scheduled backup will run at its designated date(s) and time(s) until the maximum number of backups has occurred.

- Disabled, in which the schedule backup will not run until explicitly re-enabled.

By default, when you create a Fabric Device Backup Schedule, it is enabled. However, you can manually disable a Fabric Device Backup schedule, or re-enable a disabled Backup Scheduled. If you will be enabling or disabled Fabric Device Backup Schedules, you must do so individually. Oracle Fabric Manager does not currently support a disable all or enable all option for Oracle Fabric Device Server Backup Schedules.

Note - Even if a backup schedule is disabled, you still can run an on-demand backup as needed.

You can enable or disable individual Fabric Device Backup Schedules by editing a currently configured Fabric Device Backup Schedule:

1. **Display the Fabric Interconnect Backup Schedule Summary, available through the Fabric Device Backup Schedule Summary.**

See [“Oracle Fabric Manager Backup Schedule Summary”](#) on page 449 for an example.

2. **In the Fabric Interconnect Backup Schedule Summary, click the backup schedule that you want to edit.**

This step populates the Details frame with the properties of the selected backup schedule.

3. **In the Details frame, click the Edit button to unlock the properties for editing.**

This example shows editing a daily backup schedule, but similar properties (name, enable/disable, backup schedule times and dates, maximum number of backups, and description) are editable regardless of whether the backup schedule is daily, weekly, or monthly.

4. **Click the Enabled checkbox to set the appropriate state for the backup schedule:**

- **Enabled state occurs when the Enabled checkbox contains a check mark. This checkbox is a toggle, so each click alternates between Enabled and Disabled.**
- **Disabled state occurs when the Enabled checkbox does not contain a check mark. This checkbox is a toggle.**

Working With LUN Masks

This chapter documents the following topics:

- [“LUN Masks Overview” on page 465](#)
- [“Create LUN Masks” on page 465](#)
- [“Displaying LUN Masks” on page 467](#)
- [“Add LUNs to an Existing LUN Mask” on page 469](#)
- [“Remove LUNs From the LUN Mask” on page 471](#)

LUN Masks Overview

LUN Masks enable you to zone out specific LUNs or storage targets from servers or initiators. With LUN Masks, you can keep security in the storage network by keeping LUNs that contain sensitive data in a private, restricted section of the Fibre Channel network.

Through Oracle Fabric Manager, LUN Masks are created and applied to vHBAs to determine which hosts can see which storage resources. A LUN Mask must be associated with a vHBA to control whether or not the LUNs are reported. If you create a LUN Mask, but do not associate it with one or more vHBAs, all hosts will be able to see all storage resources.

▼ Create LUN Masks

LUN Masks are created through the LUN Mask Summary, and must be created before they can be assigned to a vHBA. The LUN Mask will be assigned to individual vHBAs, not the Storage Cloud that provides vHBA terminations.

To create a LUN Mask, follow this procedure:

1. **Select Storage Cloud Manager → LUN Mask Profiles to display the LUN Mask Summary.**
2. **You can add a new LUN Mask by clicking the plus sign (+), and you can delete one or more LUN Masks by selecting them in the LUN Mask Profile Summary, then clicking the garbage can icon.**

3. Click the plus sign to display the Create LUN Mask Profile Dialog.
4. In the LUN Mask Profile Name field, enter an alphanumeric string that names the LUN Mask Profile that you are creating.
5. As an option, in the Description field, enter an alphanumeric string that describes the LUN Mask Profile that you are creating.
6. Click Submit to create the LUN Mask Profile.

Note - At this point, the LUN Mask Profile exists, but it is an empty profile. You will need to populate the empty LUN Mask with individual records that determine the WWN and LUN IDs to which the LUN Mask will be applied.

7. When the LUN Mask Profile has been created, select it in the LUN Mask Profile Summary.

The screenshot shows two parts of the Oracle Fabric Manager interface. The top part is the 'LUN Mask Profile Summary' window, which displays a table with one entry:

Name	Status	Number of Targets	In use	Description
mlathe1		1		

The bottom part is the 'LUN Mask : mlathe1' configuration window, showing the 'Target/LUNs' tab. It contains a table with one entry:

WWPN	LUN IDs
12:21:21:21:21:21:22	100:101,103

In the 'Target/LUNs' tab, a red box highlights the plus sign icon in the toolbar, which is used to add new target entries.

8. You can add target WWPN and LUN information for a storage resource by clicking the plus sign, and you can delete one or more WWPN and LUN entries by selecting the entry in on the Target/LUNs tab, then clicking the garbage can icon.

9. **By default, when the selected LUN Mask is displayed in the details frame, the General tab is displayed.**
10. **Click the Target/LUNs tab.**
11. **On the Target/LUNs tab, click the plus sign to display the Add LUN Mask Target dialog.**
12. **In the WWPN field, enter the world-wide port number for the storage resource you want added to the LUN Mask Profile.**
13. **In the LUN IDs field, enter the LUN number(s) for the storage resource you want added to the LUN Mask Profile.**

Multiple LUN numbers can be specified as either a comma-separated list, or a range of LUNs can be specified by using a colon (:) as a separator.

14. **Click Submit to configure the target and LUN(s) in the selected LUN Mask Profile.**
15. **Check the Target/LUNs tab to verify that the correct WWPN and LUN entry was added to the selected LUN Mask Profile.**

You can also check the LUN Mask Profile Summary to verify that the “Number of Targets” column has incremented (or decremented if you are deleting one or more entries) by the number of entries you are adding or deleting.

Displaying LUN Masks

When LUN Mask Profiles are configured, you can display them through the LUN Mask Profiles Summary, which is a table of all LUN Mask Profiles configured in Oracle Fabric Manager. Additional, detailed information about individual LUN Mask Profiles and their contents is available in the LUN Mask Profiles Details frame below the LUN Mask Profile Summary.

- [“LUN Mask Profile Summary” on page 467](#)
- [“LUN Mask Profile Details” on page 468](#)

LUN Mask Profile Summary

The LUN Mask Profile Summary contains all configured LUN Mask Profiles in Oracle Fabric Manager regardless of whether they are assigned to a Storage Cloud.

The screenshot shows a table titled "LUN Mask Profile Summary". At the top left, there are icons for adding (+) and deleting (trash) items. The table has five columns: "Name", "Status", "Number of Targets", "In use", and "Description". There is one row with the name "mathe1", a status icon of a red 'X', and the number "1" in the "Number of Targets" column. Below the table, there is a footer that says "1 item" with a refresh icon.

Name	Status	Number of Targets	In use	Description
mathe1	X	1		

Field	Indicates
Name	The name of each configured LUN Mask Profile.
Status	The name of the Oracle Fabric Device that discovered storage target in the LUN Mask Profile.
Number of Targets	The number of storage targets contained in each LUN Mask.
In Use	The state of the LUN Mask with regard to whether or not it is used by an I/O Profile that is deployed to a host server. If no value is displayed, the LUN Mask is configured but not actually in use on a server.
Description	The description string (if any) that was applied to the LUN Mask Profile. If this field is blank, either no description string was specified when the LUN Mask Profile was created, or the LUN Mask Profile was originally created with a description string, but the LUN Mask Profile was later edited and the description string was removed.

LUN Mask Profile Details

The LUN Mask Profile Details frame is a section of the work panel that is located below the LUN Mask Profile Summary. This frame is a list of fields for a selected LUN Mask Profile.

The LUN Mask Profile Details frame enables you to display additional, detailed information for a single LUN Mask Profile and contains an Edit button, which unlocks editable parts of the details frame so that you can set or change information elements of the details frame.

To use the LUN Mask Profile Details frame, you must first select a configured LUN Mask Profile from the LUN Mask Profile Summary. By selecting a LUN Mask Profile from the summary, you provide an element that will be the focus of the LUN Mask Profile Details frame.

The following example shows the LUN Mask Profile Details frame. Notice that the details frame is contextual, so that it displays detailed information for the item selected in the LUN Mask Profile Summary.

The screenshot shows the 'LUN Mask Profile Summary' window with a table containing one entry:

Name	Status	Number of Targets	In use	Description
mlathe1		1		

Below the table, there is a '1 item' indicator and a refresh icon. The 'LUN Mask : mlathe1' details window is open, showing the 'General' tab with the following information:

Name: mlathe1
Description:

An 'Edit' button is visible at the bottom of the details window.

▼ Add LUNs to an Existing LUN Mask

Through the LUN Mask Details frame, you can add one or more LUNs to the existing LUN Mask as needed. When LUNs are added to the target mask, they will be masked as needed.

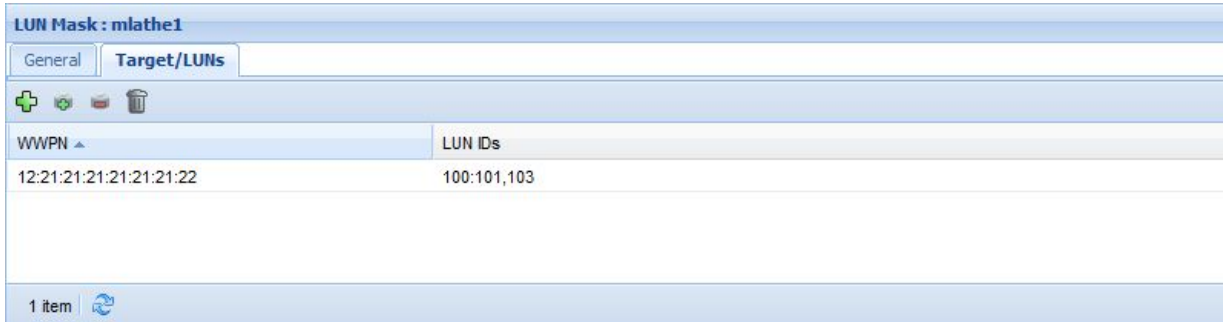
Note -Adding LUNs to a LUN Mask might require a rescan of the vHBA, which is supported through a toolbar button (the satellite dish) on the vHBAs tab of the Physical Server details page.

To add LUNs to a specified target, follow this procedure:

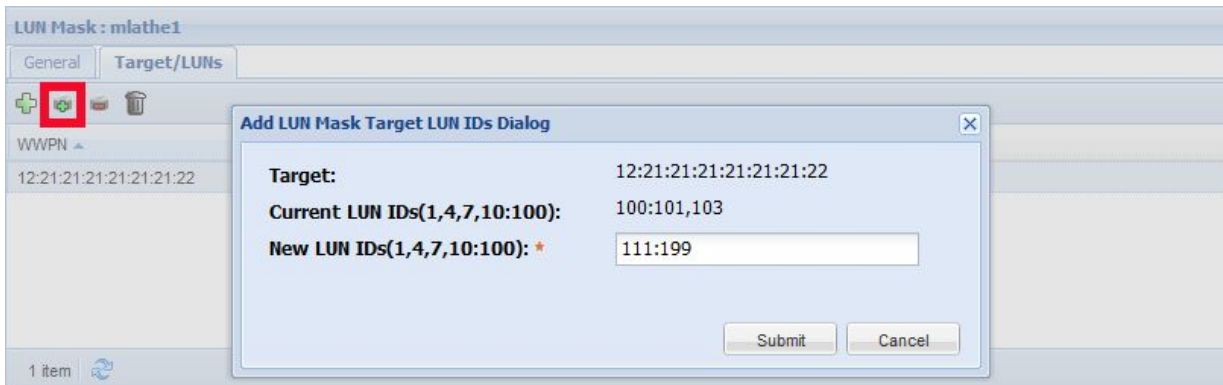
1. **Display the LUN Mask Summary.**
2. **Select a LUN Mask from the summary to highlight it.**

This step populates the LUN Mask Details frame with additional details for the selected LUN Mask.

3. Click the Target/LUNs tab.



4. On the Target/LUNs tab, click the WWPN for which you want to add more LUNs to the LUN Mask. This step activates the Add LUNs to Target button.
5. Click the Add LUN IDs to the Selected Target button to display the Add LUN Mask Target LUN IDs dialog.



6. In the New LUN IDs field, enter the LUN ID(s) that you want to add to the existing LUN Mask.
Any LUNs already configured in the LUN Mask are kept, and the LUNs you specified are added to the mask. A single LUN can be added, or multiple LUNs can be added by specifying a comma-separated list, or a hyphenated range of LUNs.
7. When the LUN(s) are specified, click Submit to update the LUNs in the LUN Mask.

▼ Remove LUNs From the LUN Mask

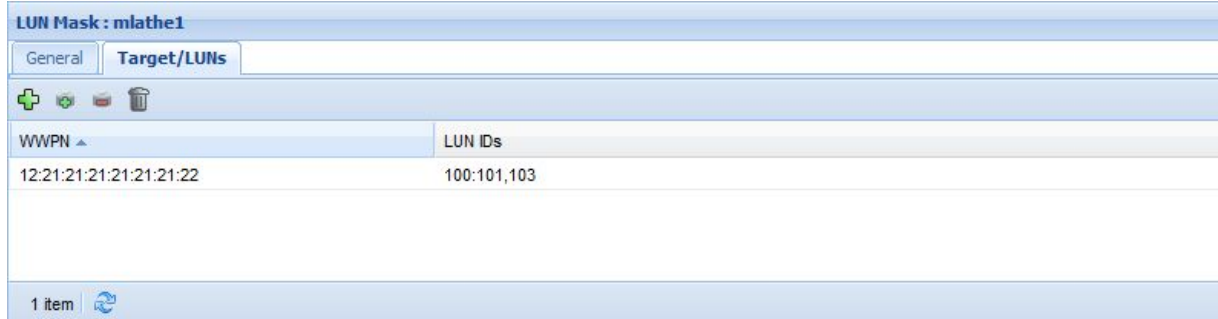
Through the LUN Mask Details frame, you can remove one or more LUNs from an existing LUN Mask as needed. When LUNs are removed from the target mask, they will no longer be masked. Removing LUNs from a LUN Mask might require a rescan of the vHBA. If the LUNs are removed, but are still being masked, you can attempt to rescan the vHBA from the Physical Server details frame (Server Resource Manager->Physical Server details frame->vHBAs tab ->Prescan/Rescan for Fibre Channel).

To remove LUNs from an existing LUN Mask, follow this procedure:

1. **Display the LUN Mask Summary.**
2. **Select a LUN Mask from the summary to highlight it.**

This step populates the LUN Mask Details frame with additional details for the selected LUN Mask.

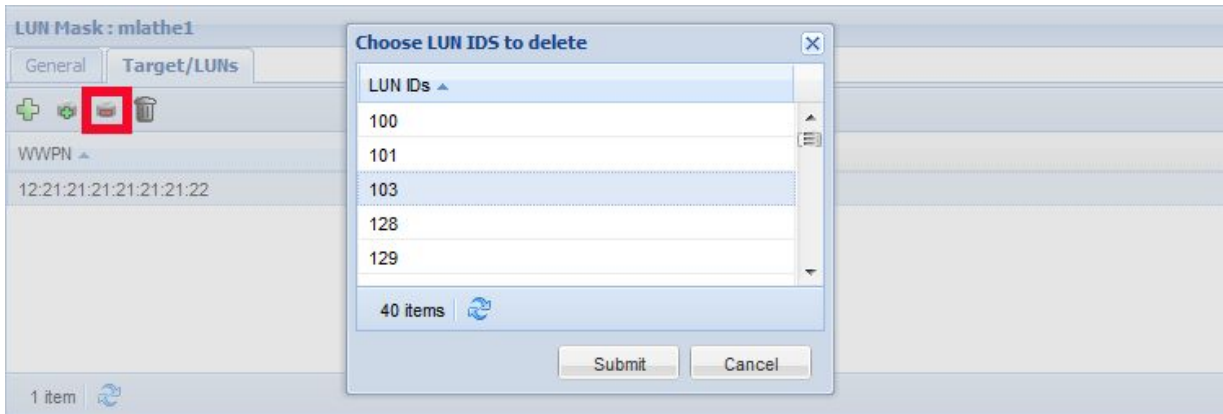
3. **Click the Target/LUNs tab.**



4. **On the Target/LUNs tab, click the WWPN for which you want to remove more LUNs from the LUN Mask.**

This step activates the Remove LUNs from the Selected Target button.

5. Click the Remove LUNs IDs from Selected Target button to display the Choose LUN IDs to delete dialog.



6. From the Choose LUN IDs to Delete dialog, click as many LUNs as you want to delete from the LUN Mask.
7. Click Submit to remove the LUNs from the LUN Mask.

Working With SAN QoS

This chapter documents the following topics:

- [“SAN QoS Profile Overview” on page 473](#)
- [“Displaying SAN QoS Profiles” on page 474](#)

SAN QoS Profile Overview

A SAN QoS Profile allows you to place bandwidth usage parameters on a Storage Cloud so that specific amounts of traffic are allowed, or specific amounts of the throughput is available. SAN QoS is enforced by a shaper profile, which attempts to guarantee bandwidth by controlling traffic through delaying and queuing frames that exceed the Committed Information Rate (CIR) value you set. When you configure a SAN QoS Profile on a Storage Cloud or vHBA, you are assigning shaping parameters to the read and write data that affect the host server that uses that Storage Cloud.

The SAN QoS Profile uses Committed Information Rate (CIR) to guarantee bandwidth.

You can assign one SAN QoS Profile to a Storage Cloud or vHBA, and you can assign the same SAN QoS Profile to multiple, different Storage Clouds or vHBAs. For more information about SAN QoS on a Storage Cloud as opposed to a vHBA, see [“vHBAs Associated With a SAN QoS Profile” on page 476](#).

Predefined SAN QoS profiles are supplied by Oracle in different bandwidth configurations to facilitate QoS use on Storage Cloud(s) or vHBAs.

Default SAN QoS Profiles

Oracle Fabric Manager contains default SAN QoS Profiles, which are pre-configured for efficient bandwidth availability and resource usage. You can apply the default SAN QoS profiles directly to a Storage Cloud or a vHBA.

Default SAN QoS Profiles are pre-configured in Oracle Fabric Manager for application to Storage Clouds or vHBAs. Oracle Fabric Manager supports the following default SAN QoS Profiles:

- 50M_125M (50 Mbps CIR and 125 Mbps PIR)
- 125M_250m (125 Mbps CIR and 250 Mbps PIR)
- 250M_500M (250 Mbps CIR and 500 Mbps PIR)
- 500M_1G (500 Mbps CIR and 1 Gbps PIR)
- 1G_2G (1 Gbps CIR and 2 Gbps PIR)
- 2G_4G (2 Gbps CIR and 4 Gbps PIR)
- 2G_8G (2 Gbps CIR and 8 Gbps PIR)
- 4G_6G (4 Gbps CIR and 6 Gbps PIR)
- 4G_8G (4 Gbps CIR and 8 Gbps PIR)
- 6G_8G (6 Gbps CIR and 6 Gbps PIR)

When you are binding a SAN QoS Profile to a Storage Cloud, you can select a default QoS Profile from the list of all available QoS Profiles for that particular module.

Displaying SAN QoS Profiles

When SAN QoS Profiles are assigned, you can display them through the SAN QoS Summary, which is a table of all SAN QoS Profiles configured in Oracle Fabric Manager. Additional, detailed information about individual SAN QoS Profiles is available in the SAN QoS Profile Details frame below the SAN QoS Profile Summary.

- [“Display the SAN QoS Summary” on page 474](#)
- [“SAN QoS Details” on page 475](#)
- [“vHBAs Associated With a SAN QoS Profile” on page 476](#)

▼ Display the SAN QoS Summary

The SAN QoS Summary contains all configured SAN QoS Profiles in Oracle Fabric Manager regardless of whether they are assigned to a Storage Cloud or a vHBA.

To display the SAN QoS Summary, follow this procedure:

- Display the SAN QoS Summary by selecting Storage Cloud Manager → SAN QoS.

SAN QoS Shaper Profile Summary				
Name ▲	CIR	PIR	Number of vHBAs	Description
125m_250m	125M	250M	0	125-250Mbps
1g_2g	1G	2G	0	1-2Gbps
250m_500m	250M	500M	0	250-500Mbps
2g_4g	2G	4G	0	2-4Gbps
2g_8g	2G	8G	0	2-8Gbps
4g_6g	4G	6G	0	4-6Gbps
4g_8g	4G	8G	0	4-8Gbps
500m_1g	500M	1G	0	500Mbps-1Gbps
50m_125m	50M	125M	0	50-125Mbps
6g_8g	6G	8G	0	6-8Gbps

10 items |

Field	Indicates
Name	The name of each configured SAN QoS Profile.
CIR	The committed information rate, which is the amount of guaranteed bandwidth for constant traffic.
PIR	The peak information rate, which is the maximum amount of total bandwidth that can be consumed by traffic.
Number of vHBAs	The total number of vHBAs that each SAN QoS Profile is attached to.
Description	The description string (if any) that was applied to the SAN QoS Profile. If this field is blank, either no description string was specified when the SAN QoS Profile was created, or the SAN QoS Profile was originally created with a description string, but the SAN QoS Profile was later edited and the description string was removed.

SAN QoS Details

The SAN QoS Profile Details frame is a section of the work panel that is located below the SAN QoS Profile Summary. This frame is a list of fields for a selected SAN QoS Profile that shows the CIR that shapes bandwidth usage on the vHBA.

The SAN QoS Profile Details frame enables you to display additional, detailed information for a single SAN QoS Profile.

To use the SAN QoS Profile Details frame, you must first select a configured SAN QoS Profile from the SAN QoS Profile Summary. By selecting a SAN QoS Profile from the summary, you set the focus of the SAN QoS Profile Details frame. When the SAN QoS Profile is selected in the summary, you will see its details displayed in the Details frame.

The following figure shows the SAN QoS Profile Details frame. Notice that the details frame is contextual, so that it displays detailed information for the item selected in the SAN QoS Shaper Profile Summary.

Display the SAN QoS Summary

SAN QoS Shaper Profile Summary				
Name	CR	PR	Number of vHBAs	Description
125m_250m	125M	250M	0	125Mbps
1g_2g	1G	2G	0	1-2Gbps
250m_500m	250M	500M	0	250-500Mbps
2g_4g	2G	4G	0	2-4Gbps
4g_8g	4G	8G	0	4-8Gbps
500m_1g	500M	1G	0	500Mbps-1Gbps
50m_125m	50M	125M	0	50Mbps

Shaper Profile: 2g_4g	
General	vHBAs
Name:	2g_4g
Description:	2-4Gbps
CR (kbps):	2000000
PIR (kbps):	4000000

vHBAs Associated With a SAN QoS Profile

The vHBAs tab displays a summary of all the vHBAs that are currently associated with the selected SAN QoS Profile. If the Number of vHBAs field on the SAN QoS Shaper Profile Summary shows a non-zero value, you can use the vHBAs tab to display information about which vHBAs are using the selected SAN QoS Profile.

The following figure shows an example of the vHBAs associated with a SAN QoS Profile.

Shaper Profile: 125m_250m										
General	vHBAs									
Name	Storage Cloud	Server Name	Termination	State	WWNN	WWPN	LUN Mask...	LUN Masks Profile	HA	QoS
vhba	MirrorSANClo...		oregon/11/2	up/r...	50:01:39:71:00:06...	50:01:39:70:00:06...	false		f...	125000_250000
vhba	MirrorSANClo...	bering	oregon/11/2	up/d...	50:01:39:71:00:06...	50:01:39:70:00:06...	false		f...	125000_250000
vhba	MirrorSANClo...		oregon/11/2	up/r...	50:01:39:71:00:06...	50:01:39:70:00:06...	false		f...	125000_250000

Working With Domains

This chapter contains the following topics:

- [“Resource Domains” on page 477](#)
- [“Default Domain” on page 478](#)
- [“Default Domain Users” on page 478](#)
- [“Non-Default Domains” on page 479](#)
- [“Non-Default Domain Users” on page 480](#)
- [“Displaying Domain Information” on page 481](#)
- [“Creating a Domain” on page 488](#)

Resource Domains

Oracle Fabric Devices, Oracle Fabric Manager Server, physical servers, and Network and Storage Clouds reside in domains, which are logical groupings of resources in the network. Typically, domains are arranged by a functional group, such as a business unit or department, but domains can be created with virtually any theme—a lab domain, a production domain, a domain of top-quality hardware, a domain of mid-quality hardware, a domain of services or applications, and so on.

Oracle Fabric Manager enables you to create the individual domains within your network by carving out the resources required and grouping them into a needed domain. Be aware that domain boundaries are strictly enforced, so the Oracle Fabric Manager Server and the Fabric Devices it is managing must be in the same domain.

Domain Resources and Unmanaging a Fabric Device

When domains are created, they can be configured with specific resources—for example, Fabric Devices, users, clouds, and I/O Modules. If a domain contains specific resources (including a Fabric Device), and that Fabric Device is unmanaged and remanaged, the individual resources (including the Fabric Device) are not automatically put back into the respective domain(s). In this situation, the resources that were assigned to a domain(s) are put into the default domain.

The individual domain(s) still exist but will not be populated with the resources present before the unmanage operation. This behavior occurs by design to add speed and flexibility in reassigning resources in real time. By clearing the resources from their previous domain, they can be quickly reassigned to new domains as needed. Be aware that if you intend the same resources to be put back into the domain(s) that contained them before the Fabric Device was unmanaged, you must manually re-assign the resources to the correct domains.

Default Domain

By default, all resources discovered and managed by Oracle Fabric Manager reside in the default domain, which exists without any need to configure it. However, when you create additional domains, you are pulling resources out of the default domain and adding them to the specific domain that you are creating.

The process of configuring a domain is available only to the administrator-level (i.e., root) Oracle Fabric Manager user. The domain admin user (which is used to manage individual domains) cannot create individual domains. The act of creating a domain must be done through the root or admin user of the default domain. When you log in to Oracle Fabric Manager as root, the Security Manager->Resource Domains option is available of the Navigation Panel. Use this option to create individual domains that will be later managed by the admin users. When the domain is created, individual resources are added to it to allow for virtual connectivity to hosts—for example, a domain might contain one or more Oracle Fabric Manager Server(s), Fabric Devices, FC and GE Switches, and hosts. When that domain and resources are created by the administrator-level (root) user, the admin user for each domain can be created to manage that domain and its subset of the overall resources.

Default Domain Users

In the default domain, users that have administrator role can create, update, or remove non-default domains. These users can also add Fabric Devices and I/O Modules to the non-default domains, as well as delete them from the non-default domains.

In a standard usage model the administrator-level user (root) is a site administrator, and the non-default administrator (domain admin) is the administrator of individual domains. For example, think of a data center that sells rack space and services to other companies. The default domain administrator (site administrator), has access and control of the Fabric Devices and creates the individual domains for each paying company. Within the domain for each company, perhaps a few I/O Modules are available for end-to-connectivity. The I/O Modules are put into the company's domain by the default domain administrator, and from that point the configuration of a company's domain occurs through the non-default domain administrator for that domain.

Non-Default Domains

Non-default domains allow you to create logical partitions in order to subdivide physical environments from a configuration and management perspective. For example, you could create a sub-domain for finance, engineering, customer support, and so on. Within each of these non-default domains you can then assign host servers, Fabric Devices or I/O modules, Network or PVI Clouds, and Storage Clouds as needed to provide the connections required for each domain.

Individual non-default domains can be created by individual domain admins (admin users, not the root user). After the individual non-default domains are created by root, the individual domain admin users can then administer their domain(s). Individual non-default domains enable the Fabric Device and Oracle Fabric Manager to create individual logical partitions of resources that can be assigned and used as needed.

The following resources are available in non-default domains:

- Host servers. When the host servers are added to a non-default domain, they are available to the domain admins for that domain. The servers are not available to other users or admins from other domains.
- Fabric Devices or I/O Modules (but not both) which provide termination points for Oracle virtual I/O. The Oracle RBAC model supports configuring the required roles for domain administrators and users. The Fabric Devices or I/O Modules that are added to a non-default domain are not available to other users or admins from other domains. Of the users for a non-default domain, only the domain admin has rights to manage the Fabric Devices or I/O Modules.
- Fabric Devices or I/O modules that are managed by Oracle Fabric Manager, but not part of a user's domain, cannot be managed.
- I/O Modules can be added into (or deleted from) a non-default domain as needed by the domain admin. Other functionality for I/O Modules are available to the appropriate users in that domain. For example, users with network role can administer the Network or PVI Cloud and Gigabit Ethernet I/O Modules and ports, but not Storage Clouds, SAN resources or features, or Fibre Channel cards.
- Fabric Devices and I/O Modules can be assigned to only one domain at a time, and they cannot be shared across domain boundaries.
- Network and Storage clouds. Clouds can be added to the domain by the administrator of the default domain or, the domain administrator of the sub-domain. Functionality for other features are available to the appropriate users in that domain. For example, users with the storage role can administer the Storage Cloud, Fibre Channel cards and ports, and SAN resources and features. However, they cannot administer the Network or PVI Cloud or any of its associated resources or features.
- Network and Storage Clouds can contain only the resources that are assigned to the domain, and these resources can exist in only one domain at a time including the default domain. The resources in a domain cannot be shared across multiple domains, however, a Network or Storage Cloud can be used by multiple domains.

Users can be configured in multiple domains, but a user can be logged in to only one domain at a time. For example, assume user Joe belong to the engineering and customer support domains. Joe can log in to engineering and make changes as needed. Joe can also log in to customer-support and make changes as needed, but he must first log out of engineering.

- [“Non-Default Domain Users” on page 480](#)

Non-Default Domain Users

Non-default domain users are those uses that can access a particular domain in the Oracle Fabric Manager management framework, but not the default domain. The most powerful user of a non-default domain is the administrator, which in a typical usage model is the domain admin. The domain administrator takes the elements that are provided to the domain by the default domain administrator (the site administrator), then uses the elements to create the required connectivity for the domain.

Users in a non-default domain have some additional rights, but also some restrictions.

Users With Administrator Role

A user with the administrator role in a non-default domain can do the following within the respective domain:

- View and make changes on all the Fabric Devices and I/O Modules in the domain.
- Create, update, or delete a Network, Storage or PVI Cloud that has I/O modules that are assigned to the domain. However, Network or Storage Clouds that are shared by multiple domains cannot be modified.
- View all user role mappings in the domain.
- Add or remove any I/O Module from a Fabric Device in the domain.

The administrator of a non-default domain cannot do the following:

- Discover a new Fabric Device
- Remove a Fabric Device from a domain
- Use any options on the Oracle Fabric Manager maintenance menu. As a result, all the Maintenance menu items will be greyed out.

Users With Network Role

A user with the network role in a non-default domain can do the following within the respective domain:

- Create, update, or delete a Network or PVI Cloud that has I/O modules that are assigned to the domain. However, Network that are shared by multiple domains cannot be modified.
- View the I/O Module summary for the domain. This page is relevant to the current domain only, so only those Gigabit Ethernet modules that are added to the domain will be viewable.
- Make changes to the Gigabit Ethernet modules in the domain.

A user with network role in a non-default domain can view only the Fabric Device Summary of all Fabric Devices in the domain. A list of all Fabric Devices in all domains is not accessible to the non-domain user with the Network role.

Users With Storage Role

A user with storage role in a non-default domain can do the following within the respective domain:

- Create, update, or delete a Storage Cloud that has I/O modules that are assigned to the domain. However, Storage Clouds that are shared by multiple domains cannot be modified.
- View the I/O Module summary for the domain. This page is relevant to the current domain only, so only those Fibre Channel modules that are added to the domain will be viewable.
- Make changes to the Fibre Channel modules in the domain.

A user with storage role in a non-default domain can view only the Fabric Device Summary of all Fabric Devices in the domain. A list of all Fabric Devices in all domains is not accessible to the non-domain user with the Storage role.

Displaying Domain Information

This section contains the following topics:

- [“Display Domain Information” on page 482](#)
- [“Display Detailed Information for Domains” on page 482](#)
- [“Servers in a Domain” on page 483](#)
- [“Fabric Devices in a Domain” on page 484](#)
- [“I/O Cards in a Domain” on page 485](#)
- [“Network or PVI Clouds in a Domain” on page 486](#)
- [“Storage Clouds in a Domain” on page 487](#)

▼ Display Domain Information

Domain information is available through the Domain Summary, which is a list of individual domains that are configured. The Domain Summary shows a different amount of domains depending on whether you are logged in as root or standard admin:

- As root, all configured domains are displayed (including the default domain)
- As admin, the domain(s) are displayed only if that admin user has rights to them. With this model, you can keep domain information completely separate from admin users who are not supposed to see other domains. For example, you might not want the admin for the Engineering domain seeing or having access to the Finance domain.

To display the Domain Summary, follow this procedure:

1. **On the Navigation panel, select Security Manager → Resource Domains to display the Domain Summary.**

This option will be available only if you are logged in with an administrator account.

Name	Description
HRresourceDomain	Performs HR functions
ITresources	
netapp2	
pubdomain	
vm3	

5 items

2. **Notice that the Domain Summary contains toolbar controls to add a new domain (+) and delete an existing domain.**

▼ Display Detailed Information for Domains

Additional information is available for the configured domains through the Domain Details frame.

1. **On the Navigation panel, select Security Manager → Resource Domains to display the Domain Summary.**

This option will be available only if you are logged in with an administrator account. For an example, see [“Display Domain Information” on page 482](#).

2. **On the Domain Summary, select a configured domain to populate the Domain Details frame with information about the selected domain.**
3. **The Domain Summary contains toolbar controls to add a new domain (+) and delete an existing domain.**

When domains are added, the resources for the new domain are removed from the default domain. When domains are deleted, any resources in them are returned to the default domain. By default, the General tab is displayed in the details frame. The general information contains the name of the domain and any optional description. You can edit the description in the details frame by clicking the Edit button.

Servers in a Domain

You can display information about the physical servers in a specific domain through the Physical Servers tab. By clicking this tab, you display a list of physical servers that are in the domain. Through the Physical Servers tab, you can add more physical servers to the domain, or delete selected servers from the domain. It is possible to have an empty domain (one with no servers) by deleting all servers in the domain. In this case, the domain remains configured as a software entity that contains no servers.

Domain : foo						
General Physical Servers Fabric Devices I/O Cards Network Clouds Storage Clouds						
+ -						
Host Name ▲	Host OS	Bound	Busy	vNICs	vHBAs	Fabric Device Ports
ovn87-22	Linux/2.6.32-220.el6.x86_64:yg-5.0.1/x86_64			0	0	ExtSw-2c90200443198-Port7
1 item						

Field	Means
Host Name	The name of the physical server(s) in the domain.
Host OS	The version (or build) string for the operating system currently running on the physical server.
Bound	The state of whether or not the server is bound to an I/O Template. If the server is bound, this field contains a check mark. If the server is not bound, the field is empty.
Busy	An indicator of whether or not the server is currently busy processing a configuration or management task (for example, busy binding to or unbinding from an I/O Profile).
vNICs	Shows the total number of vNICs configured on the server. These vNICs can be in any state, not just the operational vNICs.

Field	Means
vHBAs	Shows the total number of vHBAs configured on the server. These vHBAs can be in any state, not just the operational vHBAs.
Fabric Device Ports	Shows the individual high-speed interconnect ports on the Fabric Device on which the server is connected.

Fabric Devices in a Domain

You can display information about the Fabric Devices in a specific domain through the Fabric Devices tab. By clicking this tab, you display a list of all the Fabric Devices that are partitioned into a non-default domain. Through this tab, you can add more Fabric Devices to, or delete them from, a specific domain. It is possible to have an empty domain (one with no Fabric Devices) by deleting all Fabric Devices in the domain. In this case, the domain remains configured as a software entity. Fabric Devices can be in only one domain at a time. If you delete one or more Fabric Devices, they return to the default domain, where they can then be reassigned to another non-default domain.

Domain : foo					
General	Physical Servers	Fabric Devices	I/O Cards	Network Clouds	Storage Clouds
Fabric Device Name	IP Address	IP Subnet	State (Admin/Oper)	I/O Modules	Software Version
aaaa	10.129.87.31	local	up/up	0	Build branch-6.0.0-40985 - (xicchen) Mo...

1 item

Field	Means
Fabric Device Name	The name of the Fabric Device in the domain
IP Address	The IP address of the Fabric Device.
IP Subnet	The subnet on which the Fabric Device is currently configured.
State (Admin/Oper)	The administrative and operational states of the Fabric Device.
I/O Modules	The total number of I/O Modules available on each Fabric Device in the domain.
Software Version	The version of Oracle's XgOS software that is currently installed and running on the Fabric Device.

It is possible that this dialog is empty depending on the type of termination you use for the domain. Each domain needs some kind of termination point for the virtual I/O within it. The virtual I/O can be terminated in either of the following ways:

- Terminated on a Fabric Device. In this case, one or more entire Fabric Devices are added to the domain. All I/O Modules in the Fabric Device are then added to the domain, and virtual I/O connections can be terminated on any of the ports on any of the cards in any of the Fabric Devices. This option might be preferable depending on the domain that you are

configuring. For example, if the domain you are configuring is large or mission critical, you might need an entire Fabric Device's worth of I/O modules to connect the servers in the domain.

- Terminated on an I/O Module. In this case, one or more I/O modules are added to the domain. Only the I/O Modules that you explicitly add will be contained in the domain. This option might be preferable depending on the domain that you are configuring. For example, if the domain you are configuring is small or not mission critical, you might need only a few I/O modules in the domain (instead of an entire Fabric Device) to connect the servers in the domain.

These options are mutually exclusive, so either Fabric Devices are added to the domain, or I/O cards are added to the domain, but not both. If Fabric Devices are added to the domain, the Fabric Devices tab will contain information when you add one or more Fabric Devices to the domain. If I/O Modules are added to the domain, or if no Fabric Devices have been added to the domain yet, this tab will be empty.

I/O Cards in a Domain

You can display information about the I/O cards in a specific domain through the I/O Cards tab. By clicking this tab, you display a list of all the I/O cards that are partitioned into a non-default domain. Through this tab, you can add more I/O cards to, or delete them from, a specific domain. It is possible to have an empty domain (one with no I/O cards) by deleting all I/O cards in the domain. In this case, the domain remains configured as a software entity. I/O cards can be in only one domain at a time. If you delete one or more I/O cards, they return to the default domain, where they can then be reassigned to another non-default domain.

Domain : foo			
General	Physical Servers	Fabric Devices	I/O Cards
Name	State	Type	Description
delaware/11	up/resourceMissing	sanFc2Port4GbCard	

Field	Means
Name	The name of the I/O Card that is assigned to the current domain. The name is a slot/port notation for a standard I/O card, or slot.port notation for a LAG on the I/O card.
State	The administrative and operational state of the I/O card
Type	The type string for the I/O card. In this example, the I/O card is a 10-Port 1 Gbps card, as indicated by the string nwGbEthernet10Port1GbCard.
Description	The optional description string (if any) that is assigned to each Network or PVI Cloud in the domain.

It is possible that this dialog is empty depending on the type of termination you use for the domain. Each domain needs some kind of termination point for the virtual I/O within it. The virtual I/O can be terminated in either of the following ways:

- Terminated on a Fabric Device. In this case, one or more entire Fabric Devices are added to the domain. All I/O Modules in the Fabric Device are then added to the domain, and virtual I/O connections can be terminated on any of the ports on any of the cards in any of the Fabric Devices. This option might be preferable depending on the domain that you are configuring. For example, if the domain you are configuring is large or mission critical, you might need an entire Fabric Device's worth of I/O modules to connect the servers in the domain.
- Terminated on an I/O Module. In this case, one or more I/O Modules are added to the domain. Only the I/O Modules that you explicitly add will be contained in the domain. This option might be preferable depending on the domain that you are configuring. For example, if the domain you are configuring is small or not mission critical, you might need only a few I/O modules in the domain (instead of an entire Fabric Device) to connect the servers in the domain.

These options are mutually exclusive, so either Fabric Devices are added to the domain, or I/O cards are added to the domain, but not both. If Fabric Devices are added to the domain, the Fabric Devices tab will contain information when you add one or more Fabric Devices to the domain. If no I/O Modules are added to the domain, or if no Fabric Devices have been added to the domain yet, this tab will be empty.

Network or PVI Clouds in a Domain

You can display information about the Network or PVI Clouds in a specific domain through the Network Clouds tab. By clicking this tab, you display a list of Network or PVI Clouds that are in the domain. Through the Network Clouds tab, you can add more Network or PVI Clouds to the domain, or delete selected Network or PVI Clouds from the domain. It is possible to have an empty domain (one with no Network or PVI Clouds) by deleting all Network or PVI Clouds in the domain. In this case, the domain remains configured as a software entity that contains no Network or PVI Clouds.

Name	Number of P...	Number of L...	QoS	Number of v...	Number of v...	Description
techpubs10	1	0	none	0	0	

Field Name	Means
Name	The name of the Network or PVI Cloud(s) in the domain.
Number of Ports	The total number of Ethernet ports in each Network or PVI Cloud, not the total number of ports in the domain.
Number of LAGs	The total number of Ethernet link aggregation groups (if any) configured in the domain. This number displayed in this column shows both static and dynamic (LACP) LAGs.
QoS	The name of the SAN QoS profile (if any) that is applied to each Network or PVI Cloud in the domain.
Number of vNICs	The total number of vNICs connected to the Network or PVI Cloud.
Number of vNIC Templates	The total number of I/O Templates containing vNICs that are used by the I/O Profile connected to the Network or PVI Cloud.
Description	The optional description string (if any) that is assigned to each Network or PVI Cloud in the domain.

Storage Clouds in a Domain

You can display information about the Storage Clouds in a specific domain through the Storage Clouds tab. By clicking this tab, you display a list of Storage Clouds that are in the domain. Through the Storage Clouds tab, you can add more Storage Clouds to the domain, or delete selected Storage Clouds from the domain. It is possible to have an empty domain (one with no Storage Clouds) by deleting all Storage Clouds in the domain. In this case, the domain remains configured as a software entity that contains no Storage Clouds. The following example shows the Storage Clouds tab.

Domain : foo						
General Physical Servers Fabric Devices I/O Cards Network Clouds Storage Clouds						
Name	Number of Ports	QoS	Number of vHBAs	Number of vHBA...	Description	
Storage1Raid10	1	none	0	1		

1 item

Field Name	Means
Name	The name of the Storage Cloud(s) in the domain.
Number of Ports	The total number of Fibre Channel ports in each Storage Cloud, not the total number of ports in the domain.
QoS	The name of the SAN QoS profile (if any) that is applied to each Storage Cloud in the domain.
Number of vHBAs	The total number of vHBAs configured in the Storage Cloud.
Number of vHBA Templates	The total number of I/O Templates containing vHBAs that are used by the I/O Profile connected to the Storage Cloud.

Field Name	Means
Description	The optional description string (if any) that is assigned to each Storage Cloud in the domain.

Creating a Domain

This section contains the following topics:

- [“Create a Domain” on page 488](#)
- [“Add Fabric Devices to a Domain” on page 489](#)
- [“Add Physical Servers to a Domain” on page 491](#)
- [“Add I/O Modules to a Domain” on page 492](#)
- [“Add Network Connectivity to a Domain” on page 494](#)
- [“Add Storage Connectivity to a Domain” on page 495](#)

▼ Create a Domain

When you are creating a domain, you are carving out a set of resources from the overall default domain. The Domain will include Ethernet ports, Fibre Channel ports, and physical servers that will be required in the domain. Domains are created by using the Create Domain Wizard, which will guide you through the steps of identifying which Network Clouds, Storage Clouds, and physical servers will be used to create the domain. When you run the wizard, the following objects must be available to be selectable in the wizard:

- Network Cloud(s) for the domain. The Network Cloud consists of one or more termination GE ports. The port cannot be assigned to another cloud.
- Storage Clouds for the domain. The Storage Cloud consists of one or more termination FC ports. The port cannot be assigned to another cloud.
- Physical hosts, which must be connected to a Fabric Device managed by Oracle Fabric Manager. The host(s) cannot already be connected to a Network or Storage Cloud.

When you create a domain, you must be logged in as a root-level administrator account so that you can carve out the resources from the default domain. After the domain is configured, you can log back in with a domain admin account for that domain, and continue to configure and manage that individual domain. When you are logged in to the domain as a domain admin, the Security Manager->Resource Domains is no longer available, and no other domains are accessible.

In typical deployments, creating a domain takes the following steps:

- Adding Physical Servers to a Domain
- Adding either Fabric Devices or I/O Modules to a domain, but not both
- Adding Network Connectivity to a Domain.

- Add Storage Connectivity to a Domain

To create a domain, follow this procedure:

1. **On the Navigation panel, select Security Manager → Resource Domains to display the Domain Summary.**

This option will be available only if you are logged in with an administrator account.

For an example, see [“Display Domain Information” on page 482](#).

2. **The Domain Summary contains toolbar controls to add a new domain and delete an existing domain.**

When domains are added, the resources for the new domain are removed from the default domain. When domains are deleted, any resources in them are returned to the default domain.

3. **On the Domain Summary, click the plus sign (+) to start the Create Domain Wizard.**

4. **In the Name field, enter the name for the domain that you are creating.**

5. **As an option, you can enter an alphanumeric character string that describes the domain that you are creating.**

6. **Click Submit to create the domain.**

The domain is created as an empty container to which you will add the required components to provide virtual connectivity to the host(s).

▼ Add Fabric Devices to a Domain

When you have a domain created, you need to populate it with the network entities that the domain needs. One of these network entities is the termination port for virtual I/O. You have an option for specifying the termination points for the domain:

- **Fabric Devices.** If you add Fabric Devices to the domain, all I/O Modules in the Fabric Device are placed into the domain and available for connections.
- **I/O Modules.** If you add I/O Modules to the domain, only the I/O Modules you select are added. This option provides you with an additional level of granularity. For example, if the domain is relatively small, you might not want to use an entire Fabric Device's worth of I/O modules. In this case, you can use only a few of the modules for the domain, and reserve the other I/O Modules for other uses or other domains.

The choice of adding Fabric Devices or I/O Modules is a mutually exclusive one. You can add either, but not both. You will need to determine which option suits the domain.


- If you will add Fabric Devices to the domain, complete the procedure in this section.

- If you will be adding I/O modules to the domain, skip the procedure in this section, and [“Add I/O Modules to a Domain” on page 492.](#)

You can add Fabric Devices to the domain through the Fabric Devices tab on the Domain Details page.

To add a Fabric Device to a domain, follow this procedure:

- 1. On the Domain Summary, click the domain to which you want to add a Fabric Device.**
This step populates the Domain Details frame with the selected domain.
- 2. On the Domain Details frame, click the Fabric Devices tab.**
- 3. On the Details frame, click the plus sign (+) to add Fabric Devices to the domain.**



- 4. On the dialog, click the Fabric Device(s) you want to add to the domain.**
For an HA deployment, you will select two Fabric Devices. You can use standard key sequences to select multiple servers. For example, on a Windows Oracle Fabric Manager client, CTRL + click selects multiple Fabric Devices, Shift + click selects all Fabric Devices.

5. Click **Submit** to add the selected **Fabric Devices** to the domain.
6. Proceed to the next section.

▼ Add Physical Servers to a Domain

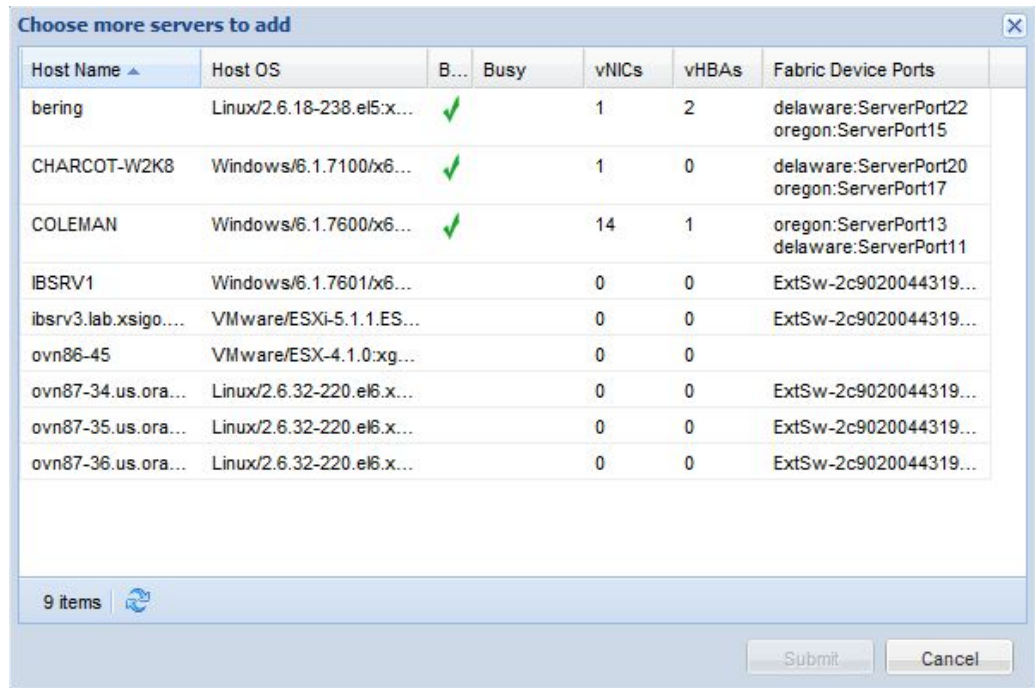
When you have a domain created, you need to populate it with the network entities that the domain needs. One of these network entities is the physical host server(s). You can add physical servers to the domain through the Physical servers tab on the Domain Details page.

To add servers to the domain, follow this procedure:

1. On the **Domain Summary**, click the domain in which you want the server configured.
This step populates the Domain Details frame with the selected domain.
2. On the **Domain Details** frame, click the **Physical Servers** tab.



3. On the Physical servers tab, click the plus sign (+) to start the add server dialog.



4. Select the server(s) that you want to add to the domain.
You can use standard key sequences to select multiple servers—for example, on a Windows Oracle Fabric Manager client, CTRL + click or Shift + click.
5. When the servers have been specified, click Submit to add them to the domain.
6. Proceed to the next section.

▼ Add I/O Modules to a Domain

When you have a domain created, you need to populate it with the network entities that the domain needs. One of these network entities is the termination port for virtual I/O. You have an option for specifying the termination points for the domain:

- Fabric Devices. If you add Fabric Devices to the domain, all I/O Modules in the Fabric Device are placed into the domain and available for connections.

- **I/O Modules.** If you add I/O Modules to the domain, only the I/O Modules you select are added. This option provides you with an additional level of granularity. For example, if the domain is relatively small, you might not want to use an entire Fabric Device's worth of I/O modules. In this case, you can use only a few of the modules for the domain, and reserve the other I/O Modules for other uses or other domains.

The choice of adding Fabric Devices or I/O Modules is a mutually exclusive one. You can add either, but not both. You will need to determine which option suits the domain.

- If you will add Oracle Fabric Devices to the domain, skip the procedure in this section, and proceed to [“Add Fabric Devices to a Domain” on page 489](#).
- If you will be adding I/O modules to the domain, complete the procedure in this section.

You can add I/O modules to the domain through the I/O Cards tab on the Domain Details page.

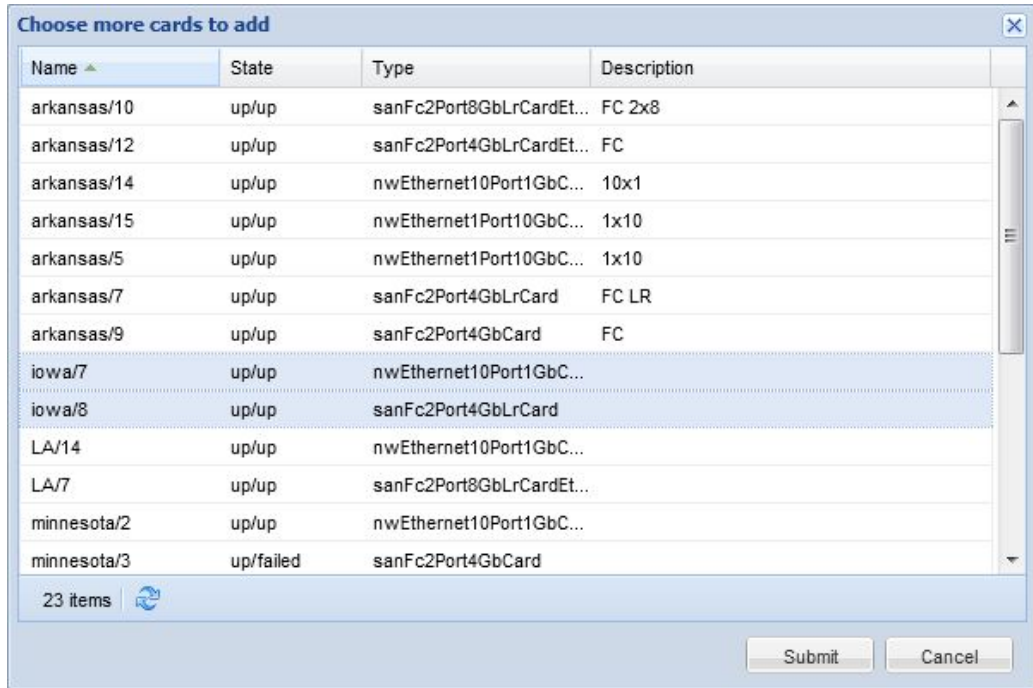
To add an I/O module to a domain, follow this procedure:

- 1. On the Domain Summary, click the domain to which you want to add an I/O module.**

This step populates the Domain Details frame with the selected domain.

- 2. Click the I/O Cards tab.**

3. Click the plus sign (+) to display a dialog through which you can add one or more I/O Modules to the domain.



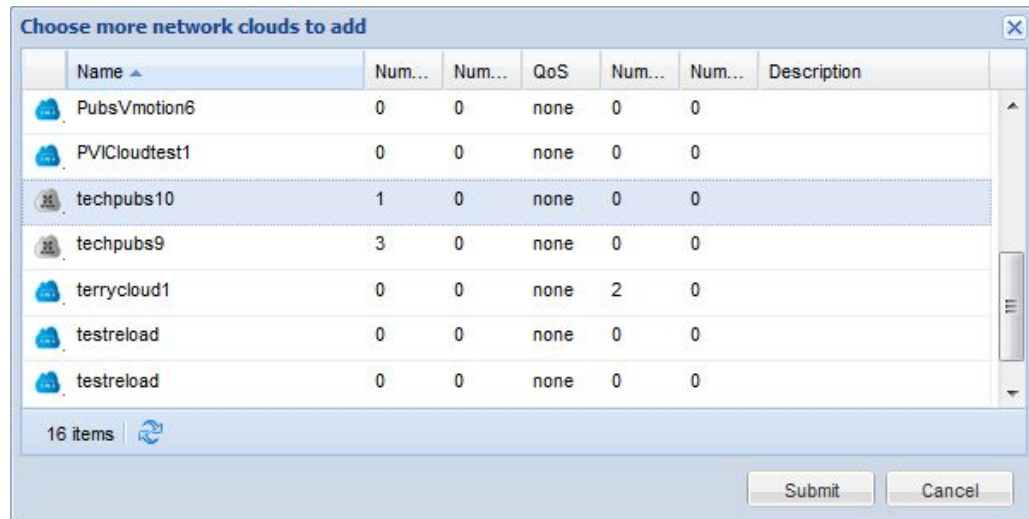
4. On the dialog select one or more I/O Modules to add them to the domain.
For an HA deployment, you will select two I/O cards of the same type. You can use standard key sequences to select multiple I/O cards. For example, on a Windows Oracle Fabric Manager client, CTRL + click selects multiple I/O cards, Shift + click selects all I/O cards.
5. Click Submit to add the selected I/O cards to the domain.
6. Proceed to the next section.

▼ Add Network Connectivity to a Domain

When you have a domain created, you need to populate it with the network entities that the domain needs. One of these network entities is a Network Cloud, which provides connectivity to network resources. You can add Network Clouds to the domain through the Network Clouds tab on the Domain Details page.

To add a Network/PVI Cloud to the domain, follow this procedure:

1. **On the Domain Summary, click the domain to which you want to add a Network/PVI Cloud.**
This step populates the Domain Details frame with the selected domain.
2. **On the details frame, click the Network Clouds tab.**
3. **Click the plus sign (+) to display a dialog that adds the Network/PVI Cloud.**



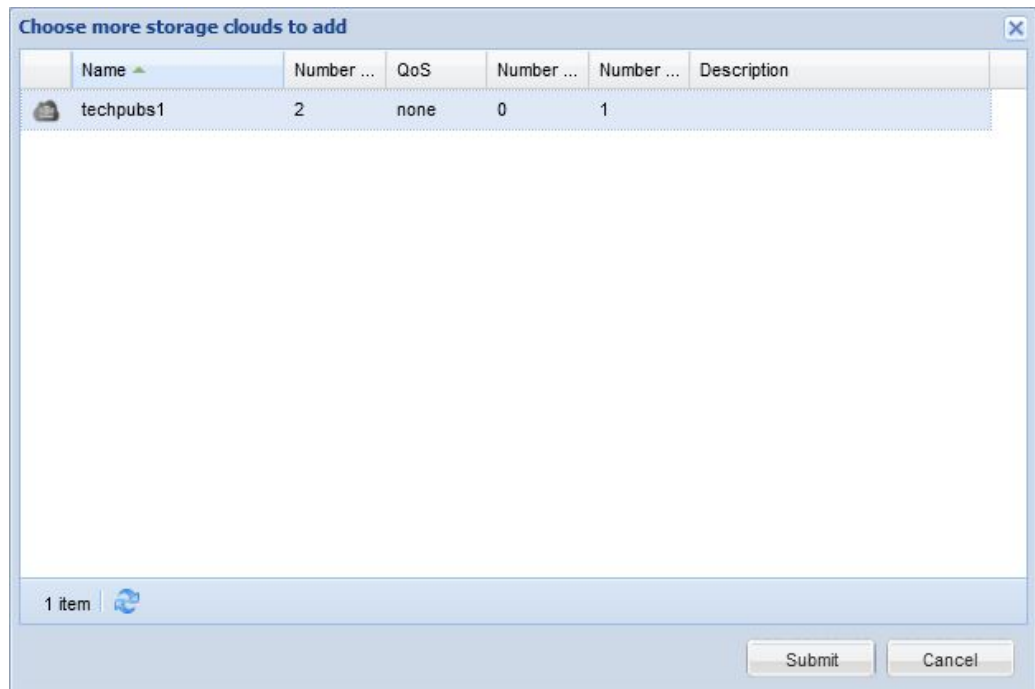
4. **On the dialog, select the Network/PVI Cloud that you want to add to the domain.**
For an HA deployment, you will select two I/O cards of the same type. You can use standard key sequences to select multiple I/O cards. For example, on a Windows Oracle Fabric Manager client, CTRL + click selects multiple I/O cards, Shift + click selects all I/O cards.
5. **Click Submit to add the Network/PVI Cloud to the domain.**
6. **Proceed to the next section.**

▼ Add Storage Connectivity to a Domain

When you have a domain created, you need to populate it with the SAN entities that the domain needs. One of these storage entities is a Storage Cloud, which provides connectivity to storage resources. You can add Storage Clouds to the domain through the Storage Clouds tab on the Domain Details page.

To add a Storage Cloud to the domain, follow this procedure:

1. **On the Domain Summary, click the domain to which you want to add a Storage Cloud.**
This step populates the Domain Details frame with the selected domain.
2. **On the details frame, click the Storage Clouds tab.**
3. **Click the plus sign (+) to display a dialog through which you can add one or more I/O Modules to the domain.**



4. **On the dialog, select the Storage Cloud that you want to add to the domain.**
For an HA deployment, you will select two I/O cards of the same type. You can use standard key sequences to select multiple I/O cards. For example, on a Windows Oracle Fabric Manager client, CTRL + click selects multiple I/O cards, Shift + click selects all I/O cards.
5. **Click Submit to add the Storage Cloud to the domain.**

Working With User Roles

This chapter contains the following topics:

- [“Oracle Identity Management System” on page 497](#)
- [“Using the Internal IMS” on page 497](#)
- [“Mapping Users in External Groups Into Oracle Fabric Manager” on page 508](#)
- [“Configuring Oracle Fabric Manager to Authenticate Users Against AD” on page 513](#)

Oracle Identity Management System

Oracle Fabric Device has a robust system of user management called the Identity Management System (IMS). The IMS has an internal system (called the internal IMS) and also has an external system called the external IMS:

- Internal IMS, is the Fabric Device's assignment of a user to a role group. The role group has permissions associated with it that determine what system object the user can use when logged in and a member of that role. This assignment of a user to a role is called internal because the mapping of user-to-role as well as the authorization and authentication of that user occurs on the Fabric Device. For information about the internal IMS, see [“Using the Internal IMS” on page 497](#).
- External IMS is the assignment of a user to a role group also. The authorization and authentication aspects of logging in through that user occur off of the Fabric Device and Oracle Fabric Manager Server. Authentication and authorization occur typically on an external device—for example, through a RADIUS server—that is connected to the Oracle Fabric Manager Server through a Fabric Device.

Using the Internal IMS

The Fabric Device's internal IMS consists of user accounts and roles. Users and roles are interrelated:

- User accounts enable login and access to the Oracle Fabric Manager system and the Fabric Device. User accounts are created on the Fabric Device and the Oracle Fabric Manager Server to grant people access to the Fabric Interconnect.

- The role that a user belongs to determines which objects the user can modify. When a user is logged in and becomes a member of a particular role, the user gets only the permissions that are granted to that role.

When a user account is created, it is assigned to a role. The role determines the permissions that the user has in Oracle Fabric Manager.

This section contains the following topics:

- [“Local Users” on page 498](#)
- [“Roles” on page 498](#)
- [“Oracle Fabric Manager User” on page 499](#)
- [“Define a User in Oracle Fabric Manager” on page 500](#)
- [“Delete a User From Oracle Fabric Manager” on page 502](#)
- [“Define a Local User on the Fabric Device” on page 503](#)
- [“Rename a Oracle Fabric Manager User” on page 504](#)
- [“User Account Roles” on page 506](#)
- [“Change the Role Assigned to a User Account” on page 506](#)

Local Users

Each Fabric Device supports local users which are user accounts on each Fabric Device. Local user accounts are local because they are independent accounts that exist on individual Fabric Devices.

Typically, a local user account is mapped into a role group. That role group provides the permissions for what objects the user can and cannot write to. When a user account is assigned to a role, the conditions are created for what actions the user can perform, and what objects the user can configure, change, or manage. For information about roles, see [“Roles” on page 498](#).

It is possible to create a local user without assigning the user to a specific role group. In this case, the user is assigned to the “no access” role group by default. The no access role group is the most restrictive. It provides read-only access.

All users have a role group—either a role group that is explicitly assigned to the user, or the operator role group that is assigned by default.

Roles

Roles are software constructions that are associated with one or more user accounts. The roles set the amount of control that a user has within the Oracle Fabric Manager and Fabric Device

system. Roles typically are created based on the division of system administration tasks in the data center. For example, some common roles are:

- Network, for managing IP network and routing connectivity.
- Storage, for managing storage capacity, configuration, and connectivity.

Oracle Fabric Manager supports a set of default roles. Each group has permissions on different hardware, software, and network components. These are the roles with which a user can be associated.

Role	Privileges for...
operator	read access to Fabric Device features
network	vNIC configuration and management, Network QoS
storage	vHBA configuration and management, SAN QoS
compute	compute resource configuration and management
administrator	full admin responsibilities
no access	Nothing. Users with this role cannot use Oracle Fabric Manager and cannot display any information in it.

Note - A special case exists for any user that is assigned to the administrator role. In such a case, when the user logs in, Oracle Fabric Manager maps that individual user account to the root account (admin/admin), which grants the user all permissions on the Fabric Device.

Oracle Fabric Manager User

Oracle Fabric Manager users are defined on each Oracle Fabric Manager server. For proper access, each Oracle Fabric Manager user account must also exist in the OS of the server where Oracle Fabric Manager is loaded.

When you create a Oracle Fabric Manager user account, you will define the user account name (which should be the same as the OS level user account) and specify a role. You do not need to specify a different password for the Oracle Fabric Manager user account because the user's network (domain) password is used.

Users will use the Oracle Fabric Manager user account and their network/domain password on the Oracle Fabric Manager login page. After successfully logging in, the current user name is displayed at the right side of the banner.

This section documents how to create a global user account at the Oracle Fabric Manager level. For information about how to create a local user on a Fabric Device, see [“Define a Local User on the Fabric Device” on page 503](#).

Be aware of the following considerations:

- The Oracle Fabric Manager user must also exist as a user at the OS level of the Oracle Fabric Manager Server. If the user does not exist in the Linux OS of a Linux Oracle Fabric Manager Server, or in the Windows OS of a Windows Oracle Fabric Manager Server, that OS user must be created now. The OS level user account must exist to allow for authentication and log in to Oracle Fabric Manager. After authentication, the OS level user account is not used. At that point, Oracle Fabric Manager user accounts are used, and they are what enforce role-based access (RBAC) privileges.
- If a specific Oracle Fabric Manager role is not configured for a Oracle Fabric Manager user, after logging in, that user is assigned to the “operator” role by default, which is the most restrictive role available.

▼ Define a User in Oracle Fabric Manager


To define a user at the Oracle Fabric Manager level, follow this procedure:

- 1. From the navigation panel, select Security Manager → User Roles.**

The Security Role Mapping Summary is displayed.

2. Click New... to display the Create Security Role Mapping dialog.

Create Security Role Mapping

User Name: * 

Domain: * default

Security Roles: *

- operator
- administrator
- network
- storage
- compute
- no-access

Description:

Apply Template Name: True

Session Timeout (min): *

Submit Cancel

3. In the User Name field, enter a string from 1 to 128 characters in length that names the user account that you are creating.
4. From the Domain dropdown menu, select a specific domain in which this user and role will be created. By default, all users and their roles are created in the default domain unless a custom domain is selected.
5. From the Security Roles list, click the checkbox for the role that you want to assign to the user that you are creating. If no user is specified, the no access role is assigned.

6. **As an option, in the Description field, you can enter a string that describes the user account that you are creating.**
7. **In the Apply Template Name checkbox, click to determine how IO Profile Names are derived.**
8. **If set to True (checkmark is present) and you apply an I/O Template to a server, the IO Profile will be named after the I/O Template.**
9. **If set to False (no checkmark), the I/O Profile will be named after the server.**
10. **When all information on the Create Security Role Mapping dialog is correct, click Finish to close the dialog and complete creating a user and its roles.**

▼ **Delete a User From Oracle Fabric Manager**


Any global user configured in Oracle Fabric Manager can be deleted through the User List. When a Oracle Fabric Manager user is deleted, that user account can no longer pass authentication at the Oracle Fabric Manager server. As a result, that user can no longer use Oracle Fabric Manager to manage the Fabric Device. However, if a local user account is still configured on the Fabric Device (not the Oracle Fabric Manager server), that user can manage the Fabric Device through the CLI.

You can delete a Oracle Fabric Manager user account through the User List summary page. To delete a Oracle Fabric Manager user, follow this procedure:

1. **From the navigation panel, select Security Manager → User Roles.**

The Security Role Mapping Summary is displayed.

Security Role Mapping Summary			
User Name ▲	Security Roles	Description	Domain
johnq	administrator		default
ofmadmin	administrator	Default administrator	default
root	administrator	Default administrator	default

3 items 

2. On the Security Role Mapping Summary, click the name of the user(s) that you want to delete.
3. Click the garbage can icon to delete the selected user(s).
4. Click Yes on the confirmation dialog to delete the specified users. Click No to abort the deletion of the selected users.

Note - Remember that this Oracle Fabric Manager user can no longer log in and manage Fabric Devices through Oracle Fabric Manager when the account is deleted. However, if there is a local account on a specific Fabric Device, that user can still log in to that Fabric Interconnect CLI through Oracle's XgOS. You will need to remove that local account on any Fabric Device if you do not want the deleted user to have any access to Fabric Devices at all.

Define a Local User on the Fabric Device

After creating the user at the Oracle Fabric Manager level, you should also create the local user account on each Fabric Device where you want the account available. You do not need to create the local user accounts on all Fabric Devices, just on the Fabric Devices that the user will need to access. The local user account is pushed to the Fabric Device. When a local user account is created, it exists on the Fabric Device and can be used for Fabric Interconnect access through Oracle's XgOS CLI.

Defining an local user occurs on each individual Fabric Device. For information about creating a local user, see [“Local Users” on page 498](#).

▼ Rename a Oracle Fabric Manager User

Oracle Fabric Manager supports renaming a Oracle Fabric Manager user, which enables you to change the name without having to completely delete and recreate the Oracle Fabric Manager user. When the Oracle Fabric Manager user is renamed, all other properties for the Oracle Fabric Manager user are retained, including the domain and security roles associated with the user. As an option, you can also set or change the description for the Oracle Fabric Manager user.

You can rename the Oracle Fabric Manager user through the Oracle Fabric Manager User Details frame. To rename the Oracle Fabric Manager user, follow this procedure:

1. **On the Navigation Frame, select Security Manager → User Roles to display the Security Role Mapping Summary.**

Security Role Mapping Summary			
User Name ▲	Security Roles	Description	Domain
johnq	administrator		default
ofmadmin	administrator	Default administrator	default
root	administrator	Default administrator	default

3 items |

2. **Select a Oracle Fabric Manager user to populate the details frame with its properties.**
3. **Click the Edit button to edit the properties of the selected Oracle Fabric Manager user.**

4. In the User Name field, enter the new name for the Oracle Fabric Manager user.

Role Mapping : staff

User Name: *

Domain: *

Security Roles: *

- operator
- administrator
- network
- storage
- compute
- no-access

Description:

Apply Template Name: True

Session Timeout (min): *

Default Topology Scale: *

5. As an option, from the Domain dropdown menu, you can select or change the role for the Oracle Fabric Manager user.
6. As an option, you also can set or change the roles for the selected Oracle Fabric Manager user.
7. As an option, you also can set or change the description for the selected Oracle Fabric Manager user.
8. As an option, you can also determine how I/O Profiles are named by using the Apply Template Name checkbox.

9. **As an option, you can also change the number of minutes that the Oracle Fabric Manager session remains open but inactive before logging out and closing the session.**
10. **As an option, you can also set the default scale of the Topology for the current user.**
11. **When the new name has been specified for the Oracle Fabric Manager user, click Submit.**

User Account Roles

After a user account has been created, you can change the role that is assigned to the account. When you change the role assigned to the account, it is possible to affect the privileges assigned to the user. For example, if the user was part of the compute role, that user had the ability to affect Server Profiles. However, if you change that user account from the compute group to the operator group, that user no longer has the ability to affect Server Profiles. In this example, the user would be able to only view information (read-only) because of being configured in the new role.

Each user account can be assigned to multiple roles. Depending on which user account is used for logging in, the appropriate amount of control can be assigned without giving super user privileges.

A role (and the privileges associated with it) are applied to the session for a user account whenever that user account logs in.

- If you change the role assigned to a user account, the new role (and its associated privileges) are applied the next time that user account logs in.
- If you change the role for an account that is currently logged in, the old role (and the associated privileges) remains in effect until that user account logs out of the current Oracle Fabric Manager session. For example, if two user accounts with administrator role are currently logged in, and you set the other user account's role to something other than "administrator" that other account will still have "administrator" role (and its privileges) as long as it remains logged in. When the other account logs out, the new non-administrator role is applied whenever the other account logs in again.

You can change the role assigned to an existing user account through the User Details page for that user account.

▼ Change the Role Assigned to a User Account

To change the role assigned to a user account, follow this procedure:

1. From the navigation panel, select **Service Manager** → **User Roles** to display the **Security Role Mapping Summary**.
See [“Rename a Oracle Fabric Manager User” on page 504](#) for an example.
2. Select the user account for which you want to change the security role.
3. In the **Role Mapping details** frame, click **Edit...** to unlock the editable fields in the **Role Mapping Details** frame.

The screenshot shows a web form titled "Role Mapping : pwilson". It contains the following fields and controls:

- User Name:** A text input field containing "pwilson".
- Domain:** A dropdown menu showing "default".
- Security Roles:** A list of roles with checkboxes:
 - operator
 - administrator
 - network
 - storage
 - compute
- Description:** A large text area, currently empty.
- Session Timeout:** A text input field containing "180".
- Default Topology Scale:** A text input field containing "3".
- At the bottom, there are two buttons: "Submit" and "Cancel".

4. As an option, in the **User Name** field, you can change the name of the Oracle Fabric Manager user for which roles will be set or changed.
5. As an option, in the **Domain** dropdown menu, you can change the domain in which the user account exists.
6. As an option, in the **Security Roles** list, you can change the security role assigned to the user account by clicking the checkbox next to the new role that you want to assign. A user can have only one role assigned to it.
7. As an option, in the **Description** field, you can enter an alphanumeric character string that describes the user account that you are editing.
8. When the user and its role information is correct, click **Submit** complete editing the user account and close the **Role Mapping Details Frame**.
When you click **Submit**, a confirmation dialog is displayed that requires to you confirm that you do actually want to make changes to the user account.
9. On the confirmation dialog, click **Yes** to make the changes to the user account, or **No** to abort editing the user account without making any changes.

Mapping Users in External Groups Into Oracle Fabric Manager

This section contains the following topics:

- [“Group Mapping Overview” on page 508](#)
- [“Map an External IMS Group to a Oracle Fabric Manager Role” on page 509](#)
- [“Map a Group to a Domain” on page 511](#)

Group Mapping Overview

Oracle Fabric Manager has enhanced group mapping functionality that enables you to specify a mapping between a user's group and either a role defined in Oracle Fabric Manager or a domain, or both. With this feature, the user's group is defined on an external Identity Management System (IMS) such as AD or LDAP, and when that user logs into the group, that group is mapped to a Oracle Fabric Manager role (if desired), or a specific domain in Oracle Fabric Manager (if desired) or both. As a result of the group mapping to role or domain (or both), the user then gets either the corresponding role, or access to only the resources in the mapped domain, or both the corresponding role and resources in the mapped domain.

Groups are identified by a regular expression that is matched against the list of groups to which a user belongs. (Regular expressions are parts of strings that represent a pattern for the overall larger string, and they can contain wildcards for the substitution of the rest of the overall string. For example, `app.*` is a regular expression for “applications”.) When a match is found, the role (or domain) is set for that user to that specified in the mapping.

Note - The regular expression must match the whole group name not part of it. For example, `abc` would not match for the group `abcd`. However, `abc.*` would match the group name.

With the Group Mapping feature, you can specify a mapping of a group name to a Oracle Fabric Manager role, or group name to domain so that a user does not have to add special groups just for Oracle Fabric Manager.

The Group Mapping feature can be used instead of creating Oracle Fabric Manager users and roles and domains. However, if the user is configured through a Group Mapping and a Oracle Fabric Manager local user, priority is given to the Oracle Fabric Manager user and its associated roles or domains.

Consider the following example. Assume a user belongs to the following groups:

`admins, administrators, adminguys, customerA`

The group mapping can be created to allow a user to log in as `admin` in a domain called `customer_a`.

Group	Role/Domain
admin.*	Role: administrator
CustomerA	Domain: customer_a

In this example, the same user would be able to login as a member of any admin.* group and receive the “administrator” role in Oracle Fabric Manager, and if the same user is able to login to the CustomerA group, then that user would also have access to the resources that belong to the customer_a domain.

A mapping can be created for both a role and a domain. In such a case, the two options are additive, so the user would log into any admin group and receive the administrator role for only the domain customer_a. If the user is part of a different group that receives a different role, or the group is mapped to a different domain, then different conditions could apply.

Group Role mapping is accessible through the Security Manager on the navigation frame. See the following sections:

- [“Map an External IMS Group to a Oracle Fabric Manager Role” on page 509](#)
- [“Map a Group to a Domain” on page 511](#)

▼ Map an External IMS Group to a Oracle Fabric Manager Role

The Group Mapping feature has a separate tab for mapping external IMS groups (for example, AD groups) to Oracle Fabric Manager roles. On the Group Mapping summary, the Group Role Mapping Summary tab supports linking one or more groups in an external IMS to one or more specific Oracle Fabric Manager roles. When multiple roles are set for a mapping, the users that are in the affected group are mapped into the corresponding roles. For example, assume the group Oracle_Server_Storage_Admins is created on an AD server. This group is then mapped to the Server and Storage roles in the Oracle Fabric Manager RBAC roles. This mapping is supported, and in this case, users in the Oracle_Server_Storage_Admins group will now have multiple roles.

To support enhanced security, the no-access role can be used. This role is a special case that allows no access to Oracle Fabric Manager, not even read-only access. With the no-access role, you can block certain groups from even seeing any of the information within Oracle Fabric Manager. For strictest security, Oracle recommends putting the domain users of Oracle Fabric Manager and Fabric Devices into their respective groups, and putting everyone else into the no-access group.

Note - When mapping external IMS groups, be aware that currently you cannot have an administrator, network, storage, or compute role in a non-default domain. If these roles are used in a non-default domain, users will be logged in as operator (read-only). In the default domain, administrator, network, storage, and compute roles can be assigned and function predictably.

To map an external IMS group to a Oracle Fabric Manager role, follow this procedure:

1. **On the navigation frame, select Security Manager → Group Mapping to display the group mapping summary.**

By default, the Group Role Mapping Summary tab is displayed.

2. **On the Group Role Summary tab, click the plus sign (+) to display the Create Group Domain Mapping dialog..**

The screenshot shows a dialog box titled "Create Group Role Mapping". It contains the following fields and options:

- Mapping Name:** * HumanResourcesMapping
- Description:** Access to all Human Resources networks and storage
- Group Name:** * HumanResourcesGroup
- Security Roles:** *
 - operator
 - network
 - compute
 - administrator
 - storage
 - no-access

Buttons: Submit, Cancel

3. **In the Mapping Name field, enter the name of the mapping that you are creating for the external IMS group.**
4. **As an option, in the Description field, enter an alphanumeric character string that describes the mapping that you are creating.**
5. **In the Group Name field, enter the name of the external IMS group that you are mapping to a Oracle Fabric Manager role.**

Enter the group name exactly as it appears on the external IMS (for example, exactly as it appears on the AD server) or in the form of a regular expression.

6. **From the Security Roles checkboxes, select the Oracle Fabric Manager role that you want the External IMS group mapped to.**

This step determines the exact privileges that will be granted to any user in the mapped group.

See [“Mapping Users in External Groups Into Oracle Fabric Manager” on page 508](#) for role descriptions.

Note - For Network, Storage, and Compute roles, you can assign some or all of these roles to the same group. For Administrator and Operator roles, these roles are mutually exclusive with each other, and all other roles. You can assign either Administrator or Operator to the same group.

7. **When the mapping name, external group name and Oracle Fabric Manager roles have been specified, click Submit to create the mapping.**

When the mapping is successfully created, any user that logs in to the specified group will receive the correct role in Oracle Fabric Manager.

▼ Map a Group to a Domain

The Group Mapping feature has a separate tab for mapping external IMS groups (for example, AD groups) to Oracle Fabric Manager domains. On the Group Mapping summary, the Group Domain Mapping Summary tab supports linking one or more groups in an external IMS to one or more specific Oracle Fabric Manager domains.

Note - Oracle Fabric Manager has a default domain, which encompasses all resources discovered and managed by Oracle Fabric Manager. The default domain always exists even if specific domains exist. If you will be mapping an external IMS group to a specific Oracle Fabric Manager domain, that specific domain must already exist to be a selectable option. If you need to create a specific domain, see [“Resource Domains” on page 477](#).

To map an external IMS group to a specific domain in Oracle Fabric Manager follow this procedure:

1. **From the navigation frame, select Security Manager → Group Mapping.**
By default, the Group Role Mapping Summary is displayed.
2. **Click the Group Domain Mapping Summary tab to display the Group Domain Mapping Summary.**

3. Click the plus sign (+) to display the Create Role Domain Mapping dialog.



The screenshot shows a dialog box titled "Create Group Domain Mapping". It has a close button in the top right corner. The dialog contains the following fields and values:

- Mapping Name:** * Publications
- Description:** Mapping for AD group "techpubs" to Fabric Manager domain "pubsdomain2"
- Group Name:** * techpubs
- Domain:** * pusdomain2

At the bottom right of the dialog are two buttons: "Submit" and "Cancel".

4. In the Mapping Name field, enter the name of the mapping that you are creating for the external IMS group.
5. As an option, in the Description field, enter an alphanumeric character string that describes the mapping that you are creating.
6. In the Group Name field, enter the name of the external IMS group that you are mapping to a Oracle Fabric Manager role.
Enter the group name exactly as it appears on the external IMS (for example, exactly as it appears on the AD server) or in the form of a regular expression.
7. From the Domain dropdown menu, select the domain to which users in the external IMS group will be mapped.
The domain must be configured for it to be a selectable option. If the domain you want is not listed, it does not exist in Oracle Fabric Manager and must be created. See [“Resource Domains” on page 477](#).
8. When the mapping name, external group name, and Oracle Fabric Manager domain have been specified, click Submit to create the mapping.
When the mapping is successfully created, any user that logs in to the specified group will be granted access to only the resources available in the specified domain.

Configuring Oracle Fabric Manager to Authenticate Users Against AD

This section documents how to configure Oracle Fabric Manager so that it authenticates users against AD. When this procedure is complete, users in the AD domain will be able to log in to Oracle Fabric Manager and use the RBAC user role you assign to them through Oracle Fabric Manager. If you will be using LDAP/AD for authentication in a Windows domain, install the Oracle Fabric Manager software on a Windows server that is a member of that domain as mentioned in [“Oracle Identity Management System” on page 497](#).

Note - A separate procedure exists if you want users to be able to log in to the Fabric Device itself. You can use Oracle Fabric Manager to configure this method of access—for example, allowing AD to authenticate users that ssh into Oracle Fabric Device. For more information, see [“Oracle Identity Management System” on page 497](#) and [“Oracle Identity Management System” on page 497](#).

To configure authentication of Oracle Fabric Manager users against AD LDAP users, the procedure occurs on both the AD server and on in Oracle Fabric Manager.

- [“Set Up the AD Server” on page 513](#)
- [“Set Up the Group in Oracle Fabric Manager and Assign Roles” on page 515](#)

▼ Set Up the AD Server

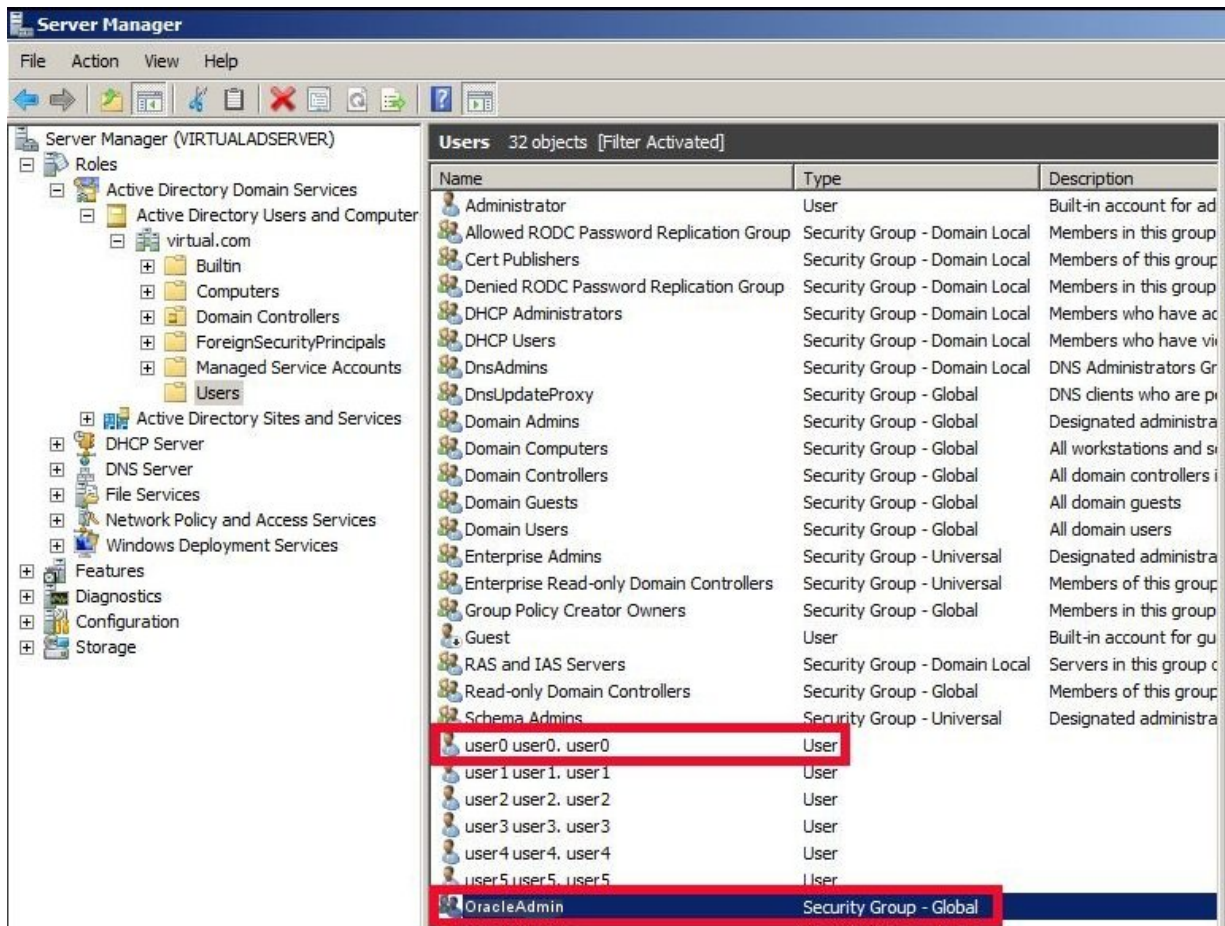
Users and groups need to exist on the AD server.

- If they do not exist, you will need to configure the users and group(s) for Oracle Fabric Manager that you want AD LDAP to authenticate. Then, you will need to add the users to their respective group(s). This section documents creating the users and groups on the AD server. For illustrative purposes, this procedure uses user0 and OracleAdmin as the user and group that will be accessing Oracle Fabric Manager and using AD authentication to do so. If your users and groups do not already exist, perform the procedure in this section.
- If they already exist, you can use a regular expression when mapping the existing group to a Oracle Fabric Manager user role. For information about regular expressions and mapping the role groups, see [“Mapping Users in External Groups Into Oracle Fabric Manager” on page 508](#). When you are familiar with using regular expressions in the group role mapping, proceed to [“Set Up the Group in Oracle Fabric Manager and Assign Roles” on page 515](#). You do not need to perform the procedure in this section if the users and groups already exist.

To configure the users and groups on the AD server, follow this procedure:

Note - The following procedure provides general guidelines. If you need additional help with creating AD groups and users, see the documentation that accompanied your AD server.

1. **Log in to the AD server as a domain administrator.**
2. **Using AD Server Manager (or whichever method you choose), create the group for Oracle Fabric Manager that will need authentication through AD.**
3. **Using AD Server Manager (or whichever method you choose), create the users in the group that will be authenticated when accessing Oracle Fabric Manager.**



4. **In this example, user0 is a member of the OracleAdmin group.**

5. Repeat this procedure as needed to create additional AD groups and users.

When the AD groups and users are completely configured, you need to map the group(s) into one or more Oracle Fabric Manager RBAC role groups. Proceed to [“Set Up the Group in Oracle Fabric Manager and Assign Roles” on page 515.](#)

▼ **Set Up the Group in Oracle Fabric Manager and Assign Roles**

When you set up Oracle Fabric Manager, you will map the existing AD group(s) into Oracle Fabric Manager and assign one or more user roles to those group(s). Any users that log in through the group (OracleAdmin in this example) will get the role assigned to the group.

To map the AD group into Oracle Fabric Manager and assign roles, follow this procedure:

1. Verify that the server where Oracle Fabric Manager software is (or will be) installed is a member of the AD domain.

For example, if the AD domain covers companyA.com, make sure that the Windows Oracle Fabric Manager Server is a member of a fully qualified domain name for companyA.com., such as ofmsserver.companyA.com.

Note - The Oracle Fabric Manager Server and the AD server(s) must all be in the same domain. If any of this equipment is in a different domain, authentication will not complete successfully.

2. Log in to Oracle Fabric Manager with a Oracle Fabric Manager local administrator account (not the AD domain account).

3. Click the Maintenance icon, which looks like a screwdriver on the banner.

4. Click the check box to Allow Unlisted Users to determine if users not in a specific group can access Oracle Fabric Manager.

This option allows users not specifically configured in Oracle Fabric Manager (for example, users configured on an AD server) to be allowed access to Oracle Fabric Manager.

Note - For more information about the Allowed Unlisted User options, see [“Oracle Identity Management System” on page 497.](#)

5. Select Security Manager → Group Mapping to display the Group Role Mapping Summary tab.

6. On the Group Role Mapping Summary, click the plus sign (+) to display the Create New Group Role Mapping dialog.

The screenshot shows a dialog box titled "Create Group Role Mapping". It contains the following fields and options:

- Mapping Name: ***: Text box containing "Virtual_AD_Admins".
- Description:**: Text area containing "Admin users from "OracleAdmin" group on AD/LDAP".
- Group Name: ***: Text box containing "OracleAdmin".
- Security Roles: ***: A grid of checkboxes:
 - operator
 - network
 - compute
 - administrator
 - storage
 - no-access

At the bottom right, there are "Submit" and "Cancel" buttons.

7. In the Mapping Name field, enter an alphanumeric character string that names or describes the group role mapping that you are configuring.
8. As an alternative, in the Description field, you can enter a string that describes the mapping from the AD group to the Oracle Fabric Manager role.
9. In the Group Name field, enter the name of AD Group you configured on the AD server (OracleAdmin in this example).

The group name must be an exact match between the AD server and this field. This field supports regular expressions. (For more information about regular expressions in a group role mapping, see ["Mapping Users in External Groups Into Oracle Fabric Manager"](#) on page 508.)
10. From the Security Roles checkboxes, click the appropriate role(s) that you are granting to the user(s) that are members of the AD group named in the Group Name field.

In this example, any user (such as user0) from the OracleAdmin group that logs in to Oracle Fabric Manager will be granted administrator privileges in Oracle Fabric Manager after AD authentication occurs successfully.
11. Click Submit and verify that the role group mapping was created correctly.

Make sure that the AD group displayed in the Group Name field correctly represents the AD group configured on the AD server.

- 12. On the Oracle Fabric Manager banner, click Logout.**
- 13. Log in to Oracle Fabric Manager as one of the users from the AD group you just mapped into Oracle Fabric Manager.**

In this example, the user user0 (who is a member of the OracleAdmin group that was just configured) is logging in to Oracle Fabric Manager.
- 14. Check the Oracle Fabric Manager banner for the currently logged in user, which should be the user from the AD group that you just used to log in to Oracle Fabric Manager.**

Also, make sure that the correct user role is assigned to the user.
- 15. Repeat this procedure as needed to create additional mappings for additional AD groups.**

Index

Numbers and Symbols

- 10 Gig Ethernet module, 164
- 10-Port GE module, Link Aggregation Groups, 281
- 10-Port Gig Ethernet module, 164
- 4-Port 10 Gig Ethernet module, 164
- 8 Gig Line-Rate Fibre Channel module, 167

A

- Active Directory, 183
- Active Directory, authenticating Fabric Manager, 513
- Active Directory, configuring, 184
- Active Directory, creating AD groups, 513
- Active Directory, creating AD users, 513
- AD groups, mapping to Fabric Manager, 515
- aggregation, link, 281
- Alarm History Summary, 401, 402
- Alarm History Summary, clearing, 406
- Alarm History Summary, displaying, 404
- Alarm History Summary, displaying details, 404
- Alarm History Summary, filtering, 407
- Alarm Summary, 399
- Alarm Summary, clearing, 406
- Alarm Summary, displaying, 399
- Alarm Summary, displaying details, 402
- Alarms, 399, 401, 402
- Alarms icons, 62
- Alarms, clearing from Alarm History Summary, 406
- Alarms, displaying general properties, 403
- Alarms, filtering historical information, 407
- Allowed VLANs, 254
- assigning iSCSI Boot Templates, 373
- authentication, AD, 183
- authentication, Fabric Manager in AD, 513, 515

B

- backup jobs, Fabric Interconnect, 443
- backup jobs, Fabric Manager Server, 443
- backup schedules
 - Fabric Interconnect, 448
 - Fabric Manager Server, 448
- backup, Fabric Manager config, 79
- backup, Fabric Manager config for fresh install, 81
- banner, 61
- Boot Profiles link, 66
- boot up sequence, iSCSI Boot, 353
- boot up sequence, SAN Boot, 334
- bootable I/O Template, iSCSI Boot, 363
- bootable I/O Template, SAN Boot, 342
- bootable template, assigning to server (iSCSI), 373
- bootable template, dual vHBA, 346
- bootable template, dual vNIC, 365
- bootable template, single vHBA, 342
- bootable template, single vNIC, 363
- booting physical servers through iSCSI Boot, 354
- booting physical servers through SAN Boot, 333
- buttons, About, 62
- buttons, Help, 62
- buttons, Login, 56
- buttons, Logout, 62
- buttons, Maintenance, 62
- buttons, Projector/Screen mode, 62

C

- CA (certificate authority), 46
- CA, sending CSR to, 50, 53
- cards, displaying (FC), 167, 167
- cards, displaying (GE), 164
- certificate error, 49
- certificate signing request (CSR), 47

- certificate signing request, generating, 50, 53
- certificate, for Fabric Manager, 46
- certificate, installing, 46
- certificate, security, 46
- certificate, sending to signing authority, 50, 53
- certificate, signed, 48
- CIR (committed information rate)
 - for Network QoS Profile, 292
 - for SAN QoS Profile, 473
- clearing alarms, 406
- client requirements, 26
- cloud, configuring network features, 107
- cloud, configuring storage features, 135
- cloud, creating Network, 101
- cloud, creating Storage, 133
- cloud, displaying network, 118
- cloud, displaying network details, 119
- cloud, displaying storage, 143
- cloud, displaying storage details, 145
- cloud, editing network properties, 120
- cloud, editing storage properties, 145
- clouds, displaying domains, 485, 486, 487
- committed information rate (CIR)
 - for Network QoS, 292
 - for SAN QoS, 473
- configuration, Fabric Manager backup, 79
- configuration, Fabric Manager backup for fresh install, 81
- configuration, Fabric Manager restoring, 82
- creating default gateway, 154
- creating direct-attached, 360
- creating I/O Profiles, 262
- creating iSCSI Boot Profile, 355
- creating iSCSI Boot Template, 363
- creating local users, 503
- creating LVM iSCSI Boot Profile, 361
- creating Network Clouds, 101
- creating SAN Boot Profile, 336
- creating SAN Boot Template, 342
- creating static iSCSI Boot Profile, 358
- creating Storage Clouds, 133
- CSR (certificate signing request), 47
- CSR, generating, 50, 53
- CSR, sending to CA, 50, 53
- current alarms, clearing from Alarm Summary, 406
- current alarms, displaying, 399

- current alarms, displaying details, 402

D

- default domain, 478
- Default Gateway, 153
- Default Gateway, creating, 154
- Default Gateway, deleting, 156
- Default Gateways link, 66
- default user roles, 499
- deleting gateway, default, 156
- deleting iSCSI Boot Profiles, 373
- deleting SAN Boot Profiles, 352
- Details frame, Fabric Interconnect, 162
- Details frame, iSCSI Boot Profile, 369
- discovering a Fabric Interconnect, 160
- Discovery Subnets, 395
- Discovery Subnets, adding, 397
- Discovery Subnets, displaying, 395
- Discovery Subnets, displaying details, 396
- displaying alarm details, 402
- displaying alarms, 399
- displaying Ethernet cards, 164
- displaying Fabric Interconnect, 157
- displaying fan information, 180
- displaying Fibre Channel cards, 167, 167
- displaying iSCSI Boot details, 369, 371
- displaying physical servers, 303
- displaying power supply information, 180
- displaying SAN Boot details, 350
- displaying SAN Boot Profiles, 349
- Domain details, 482
- Domain Summary, 477, 482
- domains, creating, 488
- domains, default, 478
- domains, displaying, 482
- domains, displaying details, 482
- domains, displaying Network Cloud, 485, 486
- domains, displaying Storage Cloud, 487
- domains, managing, 478
- domains, managing sub-domains, 479
- domains, mapping external users, 508

E

- editing Network Clouds, 109

- editing Storage Clouds, 145
 - error, security certificate, 49
 - Ethernet cards, displaying, 164
 - Ethernet LAG Summary, 285
 - external users, group mapping, 508
 - external users, mapping domain, 508
- F**
- Fabric Interconnect
 - discovering, 160
 - displaying all, 157
 - displaying Ethernet cards, 164
 - displaying fans, 180
 - displaying Fibre Channel cards, 167, 167
 - displaying power supplies, 180
 - roles, 499
 - scanning all, 158
 - users, 503
 - Fabric Interconnect backup, 443
 - Fabric Interconnect Backup schedule, 448
 - Fabric Interconnect discovery subnet, 395
 - Fabric Interconnect discovery subnet, adding, 397
 - Fabric Interconnect discovery subnet, displaying details for, 396
 - Fabric Interconnect on-demand schedule, running, 462
 - Fabric Interconnect schedules, 69
 - Fabric Interconnect schedules, creating, 456
 - Fabric Interconnect schedules, disabling, 463
 - Fabric Interconnect schedules, displaying, 459
 - Fabric Interconnect schedules, editing, 460
 - Fabric Interconnect schedules, enabling, 463
 - Fabric Interconnect schedules, renaming, 461
 - Fabric Interconnect Summary, 162
 - Fabric Interconnect users, 168, 498
 - Fabric Interconnect users, configuring, 169
 - Fabric Interconnect users, defining, 504
 - Fabric Interconnect, discovering in Fabric Manager, 160
 - Fabric Interconnect, displaying all, 157
 - Fabric Interconnect, displaying details, 162
 - Fabric Interconnect, scanning all, 158
 - Fabric Manager authentication with AD, 513
 - Fabric Manager backup schedule, editing, 450
 - Fabric Manager backup schedule, enabling, 455
 - Fabric Manager backup schedules, disabling, 455
 - Fabric Manager backup schedules, displaying, 449
 - Fabric Manager banner, 61
 - Fabric Manager checking status, 36
 - Fabric Manager client requirement, 26
 - Fabric Manager config, backing up, 79
 - Fabric Manager config, backing up for fresh install, 81
 - Fabric Manager config, restoring, 82
 - Fabric Manager high availability, 409
 - Fabric Manager high availability, configuring partner, 416
 - Fabric Manager high availability, deleting HA partner, 430
 - Fabric Manager high availability, displaying, 424
 - Fabric Manager installation, Linux server, 27, 31
 - Fabric Manager installation, Windows server, 32, 40
 - Fabric Manager interface, learning, 61
 - Fabric Manager jobs, 91
 - Fabric Manager Jobs Summary, 91
 - Fabric Manager jobs, clearing, 96
 - Fabric Manager login, authentication, 57
 - Fabric Manager mapping for AD group, 515
 - Fabric Manager navigation panel, 63
 - Fabric Manager Network Cloud, 99
 - Fabric Manager on-demand backup, 454
 - Fabric Manager roles, 169
 - Fabric Manager schedules, 69
 - Fabric Manager schedules, creating, 452
 - Fabric Manager schedules, renaming, 451
 - Fabric Manager security certificate error, 49
 - Fabric Manager security certificate, configuring, 46
 - Fabric Manager security certificate, generating, 50, 53
 - Fabric Manager Server backup, 443
 - Fabric Manager Server Backup schedule, 448
 - Fabric Manager Server requirements, 26
 - Fabric Manager services, 36
 - Fabric Manager Storage Cloud, 131
 - Fabric Manager Task Scheduler, 443
 - Fabric Manager templates, 205
 - Fabric Manager templates, creating, 207
 - HA vHBA, 240
 - HA vNIC, 235
 - Fabric Manager templates, deleting, 256
 - Fabric Manager templates, deploying, 249
 - Fabric Manager templates, displaying, 245

- Fabric Manager templates, editing, 247
- Fabric Manager templates, re-deploying, 252
- Fabric Manager Topology, 375, 379, 383, 388
- Fabric Manager Topology Overview, 375
- Fabric Manager users, changing roles, 506
- Fabric Manager users, creating, 499
- Fabric Manager users, deleting, 502
- Fabric Manager work panel, 72
- Fabric Manager, authentication, 57
- Fabric Manager, logging in, 57
- Fabric Manager, removing, 43
- Fabric Manager, starting and logging in, 56
- fan state, displaying, 180
- FC cards, 167
- Fibre Channel cards, displaying, 167, 167
- Fibre Channel links, setting QoS on, 473
- filtering and sorting, 74

G

- GE cards, 164
- generating a CSR, 50, 53
- grapher, real time, 435
- group role mapping, AD authentication, 515
- Group, Link Aggregation, 281
- groups (AD), creating, 513

H

- HA Fabric Manager Server partners, configuring, 416
- HA Fabric Manager Server partners, deleting, 430
- HA Fabric Manager Servers, 409
- HA Fabric Manager Servers, displaying, 424
- HA vHBA templates, creating, 240
- HA vNIC templates, creating, 235
- historical alarms, displaying, 399, 401
- historical alarms, filtering, 402
- hosts, SNMP traps
 - configuring, 178
 - displaying, 177

I

- I/O Profiles, 259, 260
- I/O Profiles link, 66
- I/O Profiles, connecting, 264

- I/O Profiles, connecting servers, 266
- I/O Profiles, creating, 262
- I/O Profiles, deleting, 275
- I/O Profiles, disconnecting, 273
- I/O Profiles, displaying, 260
- I/O Profiles, displaying details, 276
- I/O Profiles, linking to template, 274
- I/O Profiles, save as template, 266
- I/O Template Editor, bootable vHBA, 342
- I/O Template Editor, bootable vNICs, 363
- I/O Template, dual vHBA, 346
- I/O Template, dual vNIC, 365
- I/O Template, single vHBA, 342
- I/O Template, single vNIC, 363
- I/O Templates, 205
- I/O Templates link, 66
- I/O Templates, creating, 207
 - HA vHBA, 240
 - HA vNIC, 235
- I/O Templates, deleting, 256
- I/O Templates, deploying, 249
- I/O Templates, displaying, 245
- I/O Templates, editing, 247, 247
- I/O templates, editing, 247
- I/O Templates, migrating between servers, 321
- I/O Templates, re-deploying, 252
- I/O Templates, saving server config to, 319
- I/O Templates, unbinding from a server, 321
- identity management system, 497
- IMS, 497
- IMS users, group mapping, 508
- IMS users, mapping domain, 508
- IMS, displaying properties, 181
- initiator, 353
- installing Fabric Manager
 - Linux server, 27, 31
 - Windows server, 32, 40
- Interconnect details, displaying, 162
- interface, filtering and sorting, 74
- IP discovery subnet, adding, 397
- IP discovery subnet, displaying details for, 396
- IP discovery subnet, for Fabric Interconnect, 395
- IQN, 353
- iSCSI Boot, 353, 354
- iSCSI Boot overview, 354
- iSCSI Boot Profile, 358, 360, 360

- iSCSI Boot Profile Details, 369, 371
 - iSCSI Boot Profile logical volume manager, 361
 - iSCSI Boot Profile Summary, 355, 373
 - iSCSI Boot Profile, creating, 355
 - iSCSI Boot Profile, deleting, 373
 - iSCSI Boot Profile, direct-attach, 360
 - iSCSI Boot Profile, displaying, 369
 - iSCSI Boot Profile, displaying details, 371
 - iSCSI Boot Profile, logical volume manager, 361
 - iSCSI Boot Profile, static, 358
 - iSCSI Boot sequence, 353
 - iSCSI Boot template, deploying, 373
 - iSCSI bootable template, creating, 363
 - iSCSI bootable template, dual path, 365
 - iSCSI bootable template, single path, 363
 - iSCSI qualified name, 353
- J**
- Job Summary, displaying, 94
 - Jobs Summary, 91
 - Jobs Summary, clearing, 96
 - jobs, clearing, 96
 - jobs, Fabric Interconnect backup, 443
 - jobs, Fabric Manager Server backup, 443
 - jobs, recent, 93
- L**
- LAGs, configuring, 281
 - LAGs, displaying details, 285
 - LAGs, displaying properties, 285
 - LAGs, setting on Network Cloud, 115
 - LDAP, 183, 184
 - learning the Fabric Manager interface, 61
 - Lightweight Directory Access Protocol, 183, 184
 - Line-Rate Fibre Channel module, 167
 - Link Aggregation Groups, 10-Port GE module, 281
 - Link Aggregation Groups, 4-Port 10 GE module, 281
 - link aggregation, configuring, 281
 - link aggregation, displaying details, 285
 - link aggregation, displaying properties, 285
 - link aggregation, on Network Cloud, 115
 - links, Ethernet LAGs, 67
 - links, LUN Mask Profiles, 68
 - links, Network Clouds, 67
 - links, Network QoS, 67
 - links, Resource Domains, 70
 - links, SAN QoS, 68
 - links, Storage Clouds, 68
 - links, Storage Targets, 68
 - links, User Roles, 70
 - Linux server, installing Fabric Manager on, 27, 31
 - Linux server, removing Fabric Manager from, 43
 - live monitoring, 433
 - Live Monitoring, calculated statistics, 437
 - Live Monitoring, historical statistics, 437
 - Live Monitoring, real-time grapher, 435
 - local users, 168, 498
 - local users, configuring, 169
 - local users, creating, 503
 - local users, defining, 504
 - logging in and authenticating, 57
 - logging in to Fabric Manager, 56, 57
 - logical volume manager, 340
 - Login button, 56
 - logs, catalina, 40
 - logs, Fabric Interconnect, 40
 - logs, post-installation, 40
 - LUN Mask Profile Details, 468
 - LUN Mask Profile Summary, 467
 - LUN Mask Summary, 465
 - LUN Masks, adding LUNs to, 469
 - LUN Masks, creating, 465, 465
 - LUN Masks, displaying, 467
 - LUN Masks, displaying details, 468
 - LUN Masks, removing LUNs from, 471
 - LVM iSCSI, 361
 - LVM, SAN Boot connectivity, 340
- M**
- MAC-Based QoS Profiles, 297
 - MAC-Based QoS Profiles, configuring, 299
 - MAC-Based QoS Profiles, deleting, 301
 - MAC-Based QoS Profiles, displaying, 297
 - maintenance tasks, 62
 - managing domains, 477
 - managing domains, default domain, 478
 - mapping domains, 508
 - mapping groups, 508

- mapping, role group, 509
- masking LUNs for a host, 465
- masking LUNs, adding more LUNs, 469
- masking LUNs, creating, 465
- masking LUNs, displaying, 467
- masking LUNs, displaying details, 468
- masking LUNs, removing LUNs from, 471
- module
 - 10 Gig Ethernet, 164
 - 10-Port Gig Ethernet module, 164
 - 4-Port 10 Gig Ethernet, 164
- module, 8 Gig Line-Rate Fibre Channel, 167
- module, Line-Rate Fibre Channel, 167
- monitoring live performance, 433

N

- navigation panel, 63
- navigation panel, Network Cloud Manager, 67
- navigation panel, Schedules, 69
- navigation panel, Security Manager, 70
- navigation panel, Server Resource Manager, 66
- navigation panel, Service Manager, 69
- navigation panel, Storage Cloud Manager, 68
- Network Cloud, 99
- Network Cloud and QoS, 99
- Network Cloud Manager, 67
 - Ethernet LAGs, 67
 - Network Clouds, 67
 - Network QoS, 67
- Network Cloud, configuring features, 107
- Network Cloud, creating, 101
- Network Cloud, displaying, 118
- Network Cloud, displaying details, 119
- Network Cloud, editing, 120
- Network Cloud, setting LAG, 115
- Network Cloud, setting QoS, 108
- Network Clouds, displaying domains, 485, 486
- Network QoS
 - CIR, 292
 - PIR, 292
- Network QoS and Network Cloud, 99
- Network QoS assignment, 99
- Network QoS Profiles, displaying, 293
- Network QoS, displaying details, 294

O

- On-demand backup, 448
- Oracle Fabric Manager, 21, 21, 93
- Oracle Fabric Manager overview, 21
- overview, Fabric Manager interface, 61
- overview, iSCSI Boot, 354
- overview, Oracle Fabric Manager, 21

P

- partners, HA Fabric Manager
 - configuring, 416
- peak information rate (PIR), Network QoS, 292
- phone home (ProWatch), displaying, 170
- Physical Servers Details, 307
- Physical Servers link, 66
- physical servers, displaying, 303
- physical servers, displaying details, 307
- physical servers, displaying throughput, 435
- physical servers, managing vHBAs on, 313
- physical servers, managing vNICs on, 308
- physical servers, managing vNICs on, 310
- physical servers, masking LUNs for, 465
- physical servers, migrating configs between, 321
- physical servers, SAN Boot of, 333
- physical servers, save config as template, 319
- physical servers, unbinding templates from, 321
- PIR (peak information rate) for Network QoS Set, 292
- Policer Profile and predefined QoS, 292
- post-installation logs, Fabric Manager audit, 40
- predefined Network QoS Policer Profile, 292
- predefined Policer Profile, 292
- predefined QoS profiles, Network QoS, 473
- predefined QoS, Network, 292
- primary jobs, 91
- Private vNIC, 111
- Profile, I/O, 262
 - connecting, 264, 266
 - creating, 262
 - deleting, 275
 - disconnecting, 273
 - displaying details, 276
 - link to template, 274
 - saving as template, 266
- profile, Network QoS, 292

profile, Network QoS (predefined), 292
 profile, SAN QoS predefined Shaper Profile, 473
 profiles, I/O, 259
 properties, AD
 configuring, 184
 displaying, 183
 properties, alarms
 displaying, 403
 properties, editing
 Network Cloud, 120
 properties, I/O Template
 editing, 247
 properties, IMS
 displaying, 181
 properties, link aggregation
 displaying, 285
 ProWatch (phone home), configuring, 172
 ProWatch (phone home), displaying, 170
 PSU state, displaying, 180
 Public vNIC, 111

Q

QoS parameters, vHBA, 473
 QoS Profile, SAN QoS, 473
 QoS, configuring on vHBA, 473
 QoS, displaying
 Network QoS Profiles, 293
 SAN QoS Profiles, 474
 QoS, displaying details
 Network QoS Profiles, 294
 SAN QoS, 475
 QoS, MAC-based, 297
 QoS, MAC-based, configuring, 299
 QoS, MAC-based, deleting, 301
 QoS, MAC-based, displaying, 297
 QoS, Network, 99
 QoS, predefined Policer Profiles, 292
 QoS, predefined profiles
 SAN, 292, 473
 QoS, predefined sets, 292
 QoS, SAN, 473
 for vHBAs, 474
 QoS, SAN Profiles
 displaying, 474
 displaying details, 475

QoS, setting on Network Cloud, 108
 QoS, setting on Storage Cloud, 135
 QoS, Storage, 131
 quality of service, 473
 vHBA or Storage Cloud, 473
 vNIC or Network Cloud, 292
 quality of service, displaying
 network, 293
 storage, 474
 quality of service, displaying details
 network, 294
 storage, 475
 quality of service, MAC-based, 297
 configuring, 299
 deleting, 301
 displaying, 297
 quality of service, network, 99
 quality of service, predefined profiles, 292
 Policer Profile, 292
 quality of service, storage, 131

R

RADIUS, servers
 configuring, 189
 RADIUS, users
 configuring, 193
 displaying, 192
 RBAC roles, 169
 real-time grapher, 435
 recent jobs, 93, 93
 Recent Jobs Summary, 93
 removing Fabric Manager
 as a stand-alone application, 46
 Linux Server, 43
 removing fabric Manager
 Windows server, 43
 requirements
 Fabric Manager client, 26
 Fabric Manager server, 26
 restore, Fabric Manager config, 82
 role group mapping, 509
 roles, Fabric Manager, 169, 499
 roles. mapping external user, 508

S

- SAN Boot, 333, 334
- SAN Boot details frame, 350
- SAN Boot Profile Summary, 336, 349, 352
- SAN Boot Profile, creating, 336
- SAN Boot Profile, deleting, 352
- SAN Boot Profile, displaying, 349
- SAN Boot sequence, 334
- SAN Boot vHBA connection, 340
- SAN bootable template, creating, 342
- SAN bootable template, dual path, 346
- SAN bootable template, single path, 342
- SAN QoS and Storage Cloud, 131
- SAN QoS assignment, 131
- SAN QoS parameters, CIR, 473
- SAN QoS Profile, 473
- SAN QoS Profile Details, 475
- SAN QoS Profile Summary, 474
- SAN QoS, displaying, 474
- SAN QoS, displaying details, 475
- SAN QoS, predefined, 473
- scanning Fabric Interconnect, 158
- Scheduled backups, 448
- schedules, 448
- secure users, displaying, 174
- security certificate error, 49
- security certificate, certificate authority (CA), 46
- security certificate, Fabric Manager, 46
- Security Manager
 - Resource Domains, 70
 - User Roles, 70
- server bootup through SAN Boot, 333
- server group, creating from servers, 323
- Server Groups, 325
- Server Groups link, 66
- Server Groups, configuring, 328
- Server Groups, deleting, 331
- Server Groups, displaying, 325
- server requirements Fabric Manager server, 26
- Server Resource Manager
 - Boot Profiles, 66
 - Default Gateways, 66
 - Group Migration, 66
 - I/O Template, 66
 - Physical Servers, 66
 - Server Groups, 66
- server templates, 205
- server templates, creating, 207
 - HA vHBA, 240
 - HA vNIC, 235
- server templates, deleting, 256
- server templates, deploying, 249
- server templates, displaying, 245
- server templates, editing, 247, 247
- server templates, re-deploying, 252
- server, assigning iSCSI Boot template, 373
- server, deploying templates on, 249
- server, displaying templates on, 245
- server, editing templates, 247
- server, masking LUNs, 465
- server, re-deploying templates on, 252
- servers, Active Directory
 - configuring, 184
- servers, configuring
 - HA Fabric Manager Partners, 416
- servers, creating into server group, 323
- servers, deleting HA Fabric Manager partner, 430
- servers, displaying
 - HA Fabric Manager, 424
- servers, displaying details, 307
- servers, displaying throughput, 435
- servers, HA Fabric Manager, 409
- servers, managing vHBAs on, 313
- servers, managing vNICs on, 308, 310
- servers, migrating a config, 321
- servers, RADIUS
 - configuring, 189
- servers, saving config as an I/O Template, 319
- servers, unbinding I/O Template, 321
- Service Manager, 69
 - Live Monitoring, 69
 - VMware Integration, 69
- Service Manager link, 69
- services, Fabric Manager, 36
- shaper profiles, SAN QoS, 473
- simple network management protocol, 177
- site, untrusted, 49
- SNMP, 177
- SNMP secure, 177
- SNMP secure users, displaying, 174
- SNMP trap destinations, configuring, 178
- SNMP trap destinations, displaying, 177

- SNMP v2, 177
- SNMP v3, 177
- sorting and filtering, 74
- standalone Fabric Manager, 22
- starting Fabric Manager, 56
- statistics, historical and calculated, 437
- Storage Cloud, 131
- Storage Cloud and QoS, 131
- Storage Cloud Manager
 - LUN Mask Profiles, 68
 - SAN QoS, 68
 - Storage Clouds, 68
 - Storage Targets, 68
- Storage Cloud, configuring features, 135
- Storage Cloud, creating, 133
- Storage Cloud, displaying, 143
- Storage Cloud, displaying details, 145
- Storage Cloud, editing, 145
- Storage Cloud, setting QoS, 135
- Storage Clouds, displaying domains, 487
- sub-domains, 477
- sub-domains, managing, 479
- sub-jobs, 91
- summary, iSCSI Boot Profile, 355
- summary, SAN Boot Profile, 336
- system requirements
 - Fabric Manager server, 26
- system requirements Fabric Manager client, 26

T

- table displays, filtering and sorting, 74
- target IP, 353
- target, storage, 353
- Task Scheduler, 448
- template, assigning iSCSI Boot, 373
- templates, 207, 209
- templates, I/O, 205
 - creating, 207
 - deleting, 256
 - deploying, 249
 - displaying, 245
 - editing, 247, 247
 - re-deploying, 252
- throughput, host server
 - displaying, 435

- throughput, vHBA
 - displaying, 440
- throughput, vNIC
 - displaying, 438
- Topology Overview, displaying, 376
- Topology, Fabric Device Cloud View, 383
- Topology, Server Cloud View, 379
- Topology, Target Topology View, 388
- trap hosts, SNMP, 177
 - configuring, 178

U

- untrusted site, 49
- user accounts, 498
- user roles, default, 499
- user roles, specific, 499
- users (AD) creating, 513
- users, changing roles, 506
- users, creating, 499
- users, creating on Fabric Interconnect, 503
- users, deleting, 502
- users, Fabric Interconnect, 168, 498
 - configuring, 169
- users, Fabric Manager, 498, 498
- users, local, 168, 498
 - configuring, 169
 - creating, 503
- users, RADIUS
 - configuring, 193
 - displaying, 192

V

- vHBA, configuring QoS on, 473
- vHBA, configuring SAN Boot, 333
- vHBA, displaying SAN Boot Profile for, 349
- vHBA, SAN QoS parameters, 473
- vHBAs, displaying throughput, 440
- VLAN ranges, allowed per Template, 254
- VLANs, Allowed per Template, 254
- VMware consideration, installation, 26
- vNIC, displaying iSCSI Boot Profiles, 369
- vNIC, displaying iSCSI Boot Profiles for, 371
- vNIC, Private, 111

vNIC, Public, 111
vNICs, displaying throughput, 438

W

Windows server, installing Fabric Manager on, 32, 40
Windows server, removing Fabric Manager from, 43
work panel, 72