

# Oracle® Private Cloud Appliance

## Installation Guide for Release 2.4.2

**ORACLE®**

F23453-04  
March 2020

---

## Oracle Legal Notices

Copyright © 2013, 2020, 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, then 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 about 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 unless otherwise set forth in an applicable agreement between you and Oracle. 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, except as set forth in an applicable agreement between you and Oracle.

---

---

# Table of Contents

Preface .....	vii
1 Introduction to Oracle Private Cloud Appliance Installation .....	1
1.1 System Installation Procedures .....	1
1.2 About Oracle Private Cloud Appliance .....	2
1.3 Features of Oracle Private Cloud Appliance .....	2
1.4 Oracle Private Cloud Appliance Rack Components .....	3
1.5 Server Components .....	4
1.6 Operational Procedures for Oracle Private Cloud Appliance .....	4
1.6.1 Emergency Power-off Considerations .....	5
1.6.2 Cautions and Warnings .....	5
2 Site Requirements .....	7
2.1 Space Requirements .....	7
2.1.1 Receiving and Unpacking Requirements .....	7
2.1.2 Maintenance Access Requirements .....	8
2.2 Flooring Requirements .....	8
2.3 Electrical Power Requirements .....	9
2.3.1 Facility Power Requirements .....	10
2.3.2 Circuit Breaker Requirements .....	10
2.3.3 Grounding Guidelines .....	10
2.4 Temperature and Humidity Requirements .....	10
2.5 Ventilation and Cooling Requirements .....	12
3 Network Requirements .....	15
3.1 Network Connection and IP Address Requirements for Ethernet-based Systems .....	15
3.1.1 Network Connection Requirements for Ethernet-based Systems .....	15
3.1.2 IP Address Requirements for Ethernet-based Systems .....	20
3.2 Network Connection and IP Address Requirements for InfiniBand-based Systems .....	21
3.2.1 Network Connection Requirements for InfiniBand-based Systems .....	21
3.2.2 IP Address Requirements for InfiniBand-based Systems .....	23
3.3 General Network Configuration .....	24
4 Installing Oracle Private Cloud Appliance .....	27
4.1 Installation Task Map .....	27
4.2 Prepare to Install Oracle Private Cloud Appliance .....	27
4.2.1 Before You Begin .....	27
4.2.2 Review Safety Guidelines .....	28
4.2.3 Prepare the Installation Site .....	28
4.2.4 Tools Required for Unpacking and Installation .....	29
4.2.5 Unpack the System .....	30
4.3 Install Oracle Private Cloud Appliance in Its Allocated Space .....	30
4.3.1 Move the System .....	30
4.3.2 Stabilize the Appliance Rack .....	31
4.3.3 Attach a Ground Cable (Optional) .....	33
4.4 What's Next? .....	34
5 Powering On Oracle Private Cloud Appliance .....	35
5.1 Power On Task Map .....	35
5.2 Before Initial Power On .....	36
5.3 Power On for the First Time .....	36
5.3.1 Review Safety Guidelines .....	36
5.3.2 Inspect the Appliance .....	37
5.3.3 Connect the Power Cords .....	37
5.3.4 Power On the Appliance .....	38
5.3.5 Connect the Appliance to the Network .....	41

5.4 What's Next? .....	45
6 Extending Oracle Private Cloud Appliance - Additional Compute Nodes .....	47
6.1 Extending Compute Capacity Task Map .....	47
6.2 Review Safety Guidelines .....	48
6.3 Rack-mount Safety Precautions .....	48
6.4 Prepare to Add a Server .....	49
6.4.1 ESD Precautions .....	49
6.4.2 Unpack the Server Upgrade Kit .....	50
6.4.3 Prepare the Appliance Rack .....	50
6.5 Install the Server Into the Rack .....	52
6.5.1 Install Mounting Brackets .....	52
6.5.2 Attach Slide-Rail Assemblies .....	53
6.5.3 Install the Server Into the Slide-Rail Assemblies .....	55
6.5.4 Install the Cable Management Arm .....	57
6.5.5 Cable the Server .....	62
6.5.6 Replace the Rack Doors .....	65
7 Extending Oracle Private Cloud Appliance - Server Memory Kits .....	67
7.1 Extending Server Memory Task Map .....	67
7.2 Selecting a Server Memory Kit .....	67
7.3 Installing the Memory Upgrade Kit .....	68
8 Extending Oracle Private Cloud Appliance - Additional Storage .....	71
8.1 Extending Storage Capacity of Ethernet-based Systems .....	71
8.1.1 Adding Disk Space to the Internal Storage Appliance .....	71
8.1.2 Adding External Ethernet Storage .....	72
8.1.3 Adding External Fibre Channel Storage .....	72
8.2 Extending Storage Capacity of InfiniBand-based Systems .....	72
8.2.1 Extending Storage Capacity Task Map .....	72
8.2.2 Adding External Fibre Channel Storage .....	73
8.2.3 Adding External InfiniBand Storage .....	85
A Default IP Addresses .....	103
A.1 Default Management IP Addresses .....	103
A.2 Default Storage Network IP Addresses .....	105
A.3 Default InfiniBand IP Addresses .....	105
B Cabling Reference .....	107
B.1 Cisco Nexus 9348GC-FXP Switch Management Ethernet Connections .....	107
B.2 Cisco Nexus 9336C-FX2 Switch Data Network Connections .....	108
B.3 Data and Spine Switch Interconnects .....	110
B.3.1 Spine Switch to Data Switch Connections .....	111
B.3.2 Data Switch to Data Switch Connections .....	111
B.3.3 Spine Switch to Spine Switch Connections .....	111
B.4 Oracle Switch ES1-24 Ethernet Connections .....	111
B.5 Storage Head Cluster Interconnects for InfiniBand-based Systems .....	113
B.6 Connections from NM2-36P Sun Datacenter InfiniBand Expansion Switch to Oracle Fabric Interconnect F1-15 .....	114
B.7 Connections from NM2-36P Sun Datacenter InfiniBand Expansion Switches to Compute Nodes .....	114
B.7.1 NM2-36P (U19) InfiniBand Switch to Compute Nodes Connections .....	114
B.7.2 NM2-36P (U20) InfiniBand Switch to Compute Nodes Connections .....	115
C Connectivity Diagrams .....	117
C.1 External Ethernet Connections .....	117
C.2 Virtual Networking .....	119
C.3 External Storage Connections .....	120
C.4 Custom and Host Networks .....	122
D Site Checklists .....	127

D.1 System Components Checklist .....	127
D.2 Data Center Room Checklist .....	127
D.3 Data Center Environmental Checklist .....	128
D.4 Access Route Checklist .....	129
D.5 Facility Power Checklist .....	131
D.6 Safety Checklist .....	132
D.7 Logistics Checklist .....	132
D.8 Network Specification Checklist .....	134
Index .....	135



---

# Preface

This document is part of the documentation set for Oracle Private Cloud Appliance (PCA) Release 2.4. All Oracle Private Cloud Appliance product documentation is available at:

<https://docs.oracle.com/en/engineered-systems/private-cloud-appliance/index.html>.

The documentation set consists of the following items:

## **Oracle Private Cloud Appliance Release Notes**

The release notes provide a summary of the new features, changes, fixed bugs and known issues in Oracle Private Cloud Appliance.

## **Oracle Private Cloud Appliance Licensing Information User Manual**

The licensing information user manual provides information about the various product licenses applicable to the use of Oracle Private Cloud Appliance.

## **Oracle Private Cloud Appliance Installation Guide**

The installation guide provides detailed instructions to prepare the installation site and install Oracle Private Cloud Appliance. It also includes the procedures to install additional compute nodes, and to connect and configure external storage components.

## **Oracle Private Cloud Appliance Safety and Compliance Guide**

The safety and compliance guide is a supplemental guide to the safety aspects of Oracle Private Cloud Appliance. It conforms to Compliance Model No. ESY27.

## **Oracle Private Cloud Appliance Administrator's Guide**

The administrator's guide provides instructions for using the management software. It is a comprehensive guide to how to configure, monitor and administer Oracle Private Cloud Appliance.

## **Oracle Private Cloud Appliance Quick Start Poster**

The quick start poster provides a step-by-step description of the hardware installation and initial software configuration of Oracle Private Cloud Appliance. A printed quick start poster is shipped with each Oracle Private Cloud Appliance base rack, and is intended for data center operators and administrators who are new to the product.

The quick start poster is also available in the documentation set as an HTML guide, which contains alternate text for ADA 508 compliance.

## **Oracle Private Cloud Appliance Expansion Node Setup Poster**

The expansion node setup poster provides a step-by-step description of the installation procedure for an Oracle Private Cloud Appliance expansion node. A printed expansion node setup poster is shipped with each Oracle Private Cloud Appliance expansion node.

The expansion node setup poster is also available in the documentation set as an HTML guide, which contains alternate text for ADA 508 compliance.

# Audience

The Oracle Private Cloud Appliance documentation is written for technicians, authorized service providers, data center operators and system administrators who want to install, configure and maintain a private cloud

environment in order to deploy virtual machines for users. It is assumed that readers have experience installing and troubleshooting hardware, are familiar with web and virtualization technologies and have a general understanding of operating systems such as UNIX (including Linux) and Windows.

The Oracle Private Cloud Appliance makes use of Oracle Linux and Oracle Solaris operating systems within its component configuration. It is advisable that administrators have experience of these operating systems at the very least. Oracle Private Cloud Appliance is capable of running virtual machines with a variety of operating systems including Oracle Solaris and other UNIXes, Linux and Microsoft Windows. The selection of operating systems deployed in guests on Oracle Private Cloud Appliance determines the requirements of your administrative knowledge.

## Related Documentation

Additional Oracle components may be included with Oracle Private Cloud Appliance depending on configuration. The documentation for such additional components is available as follows:



### Note

If your appliance contains components that are not mentioned below, please consult the related documentation list for [Oracle Private Cloud Appliance Release 2.3](#).

- Oracle Rack Cabinet 1242  
[https://docs.oracle.com/cd/E85660\\_01/index.html](https://docs.oracle.com/cd/E85660_01/index.html)
- Oracle Server X8-2  
[https://docs.oracle.com/cd/E93359\\_01/index.html](https://docs.oracle.com/cd/E93359_01/index.html)
- Oracle Server X7-2  
[https://docs.oracle.com/cd/E72435\\_01/index.html](https://docs.oracle.com/cd/E72435_01/index.html)
- Oracle Server X6-2  
[https://docs.oracle.com/cd/E62159\\_01/index.html](https://docs.oracle.com/cd/E62159_01/index.html)
- Oracle Server X5-2  
[https://docs.oracle.com/cd/E41059\\_01/index.html](https://docs.oracle.com/cd/E41059_01/index.html)
- Oracle ZFS Storage Appliance ZS7-2  
[https://docs.oracle.com/cd/F13758\\_01/index.html](https://docs.oracle.com/cd/F13758_01/index.html)
- Oracle ZFS Storage Appliance ZS5-ES  
[https://docs.oracle.com/cd/E59597\\_01/index.html](https://docs.oracle.com/cd/E59597_01/index.html)
- Oracle Integrated Lights Out Manager (ILOM)  
[https://docs.oracle.com/cd/E81115\\_01/index.html](https://docs.oracle.com/cd/E81115_01/index.html)
- Oracle Switch ES1-24  
[https://docs.oracle.com/cd/E39109\\_01/index.html](https://docs.oracle.com/cd/E39109_01/index.html)



- NM2-36P Sun Datacenter InfiniBand Expansion Switch  
[https://docs.oracle.com/cd/E76424\\_01/index.html](https://docs.oracle.com/cd/E76424_01/index.html)
- Oracle Fabric Interconnect F1-15  
[https://docs.oracle.com/cd/E38500\\_01/index.html](https://docs.oracle.com/cd/E38500_01/index.html)
- Oracle VM  
[https://docs.oracle.com/cd/E64076\\_01/index.html](https://docs.oracle.com/cd/E64076_01/index.html)
- Oracle Enterprise Manager Plug-in  
<https://docs.oracle.com/en/enterprise-manager/cloud-control/enterprise-manager-cloud-control/13.3.1/empca/index.html>

## Feedback

Provide feedback about this documentation at:

<http://www.oracle.com/goto/docfeedback>

## Conventions

The following text conventions are used in this document:

Convention	Meaning
<b>boldface</b>	Boldface type indicates graphical user interface elements associated with an action, or terms defined in text or the glossary.
<i>italic</i>	Italic type indicates book titles, emphasis, or placeholder variables for which you supply particular values.
<code>monospace</code>	Monospace type indicates commands within a paragraph, URLs, code in examples, text that appears on the screen, or text that you enter.

## Document Revision

Document generated on: 2020-03-13 (revision: 1957)

## Documentation Accessibility

For information about Oracle's commitment to accessibility, visit the Oracle Accessibility Program website at

<https://www.oracle.com/corporate/accessibility/>.

## Access to Oracle Support

Oracle customers that have purchased support have access to electronic support through My Oracle Support. For information, visit

<https://www.oracle.com/corporate/accessibility/learning-support.html#support-tab>.

---

---

# Chapter 1 Introduction to Oracle Private Cloud Appliance Installation

## Table of Contents

1.1 System Installation Procedures .....	1
1.2 About Oracle Private Cloud Appliance .....	2
1.3 Features of Oracle Private Cloud Appliance .....	2
1.4 Oracle Private Cloud Appliance Rack Components .....	3
1.5 Server Components .....	4
1.6 Operational Procedures for Oracle Private Cloud Appliance .....	4
1.6.1 Emergency Power-off Considerations .....	5
1.6.2 Cautions and Warnings .....	5

This document provides instructions for installing Oracle Private Cloud Appliance at your site. Instructions are also provided for extending an existing Oracle Private Cloud Appliance at your site. This chapter provides an overview of these system installation procedures, and describes the features and hardware components of Oracle Private Cloud Appliance.

## 1.1 System Installation Procedures

The following table lists the procedures you need to complete to install Oracle Private Cloud Appliance at your site.

Step	Description	Links
1	Review the product notes for any late-breaking information about Oracle Private Cloud Appliance.	<i>Oracle Private Cloud Appliance Product Notes</i> at: <a href="https://www.oracle.com/engineered-systems/private-cloud-appliance/on-premises/">https://www.oracle.com/engineered-systems/private-cloud-appliance/on-premises/</a>
2	Review important safety information.	Before performing the installation, refer to the following documents: <ul style="list-style-type: none"><li>• <i>Oracle Private Cloud Appliance Safety and Compliance Guide</i></li><li>• <i>Important Safety Information for Oracle's Hardware Systems</i></li></ul>
3	Review information about Oracle Private Cloud Appliance features and hardware components.	<a href="#">Chapter 1, Introduction to Oracle Private Cloud Appliance Installation</a>
4	Complete the site preparation procedures.	<a href="#">Chapter 2, Site Requirements</a>
5	Complete the site network preparation procedures.	<a href="#">Chapter 3, Network Requirements</a>
6	Install Oracle Private Cloud Appliance at your site.	<a href="#">Chapter 4, Installing Oracle Private Cloud Appliance</a>
7	Power on Oracle Private Cloud Appliance and perform initial system configuration.	<a href="#">Chapter 5, Powering On Oracle Private Cloud Appliance</a>
8	( <i>Optional</i> ) Extend Oracle Private Cloud Appliance compute capacity with additional compute nodes.	<a href="#">Chapter 6, Extending Oracle Private Cloud Appliance - Additional Compute Nodes</a>

Step	Description	Links
9	(Optional) Extend Oracle Private Cloud Appliance server memory.	<a href="#">Chapter 7, Extending Oracle Private Cloud Appliance - Server Memory Kits</a>
10	(Optional) Extend Oracle Private Cloud Appliance storage capacity.	<a href="#">Chapter 8, Extending Oracle Private Cloud Appliance - Additional Storage</a>

## 1.2 About Oracle Private Cloud Appliance

Oracle Private Cloud Appliance is an integrated hardware and software system engineered to enable rapid deployment of a converged compute, network, and storage infrastructure for hosting most applications and workloads, whether running on Linux, Oracle Solaris, or Microsoft Windows.

In combination with customer-provided storage from Oracle or other storage vendors, Oracle Private Cloud Appliance incorporates Oracle server and network hardware with Oracle operating system, virtualization, and orchestration software to automate the discovery, configuration, deployment, and management of converged infrastructure for hosting virtual machines (VMs).

Oracle Private Cloud Appliance incorporates high-speed Ethernet or InfiniBand connectivity to provide a converged, wire-once, software-defined networking and storage fabric for all servers and storage in the appliance. Leveraging software-defined networking (SDN) technology, users can rapidly and dynamically create or modify private or public networks without having to manually re-cable connections, saving time and reducing the risk of human error. The consolidation of network connections results in up to 70 percent fewer cables and cards.

In addition to rapid infrastructure provisioning, Oracle Private Cloud Appliance also accelerates complete application stack deployment through support for Oracle VM Templates and Assemblies. Oracle VM Templates and Assemblies are preconfigured applications, middleware, and databases packaged as ready-to-run virtual machines (VMs) that can be dynamically configured at deployment time. The result is an unparalleled ability to go from “bare-metal” infrastructure power on to logging in to a newly deployed, running application within days or even hours, instead of weeks or months.

## 1.3 Features of Oracle Private Cloud Appliance

Oracle Private Cloud Appliance offers the following features:

- **Converged infrastructure orchestration software:** Automated discovery, configuration, and management of included server, network, and storage resources for rapid infrastructure provisioning.
- **Unified, software-defined network fabric:** Incorporated high-speed Ethernet or InfiniBand hardware provides a single, consolidated fabric allowing new networks to be dynamically added and existing networks to be modified without the need to manually modify cabling.
- **Rapid application stack deployment, not just infrastructure:** Support for ready-to-run Oracle VM Templates and Assemblies enables users to rapidly deploy not only the infrastructure, but also the applications that are hosted on the infrastructure.
- **Support for a wide variety of guest operating systems:** Support for Linux, Oracle Solaris, and Windows operating systems and applications.
- **Support for customer-provided storage:** Oracle Private Cloud Appliance can be configured with external storage from Oracle or existing storage from other storage vendors.
- **Unified management:** Browser-based orchestration software provides at-a-glance status for all hardware components. A plugin is available to integrate appliance management and monitoring into Oracle Enterprise Manager. Role-based access control permits secure administration of the appliance across multiple administrators.

## 1.4 Oracle Private Cloud Appliance Rack Components

Oracle Private Cloud Appliance systems consist of virtualization compute nodes, storage appliances, as well as required networking components. [Table 1.1](#) lists the quantity and description of each component in an Oracle Private Cloud Appliance system, with factory-installed Controller Software Release 2.4.x. Previous generations of hardware components continue to be supported by the latest Controller Software, as described below.



### Note

Oracle Private Cloud Appliance base racks are factory installed with the latest supported generation of server nodes. When you order expansion compute nodes, you receive the latest available generation. The Oracle Private Cloud Appliance Controller Software version must match the hardware versions installed in your environment.

**Table 1.1 Oracle Private Cloud Appliance Rack Components**

Quantity	Description
2	Oracle Server X8-2, forming the HA management node cluster
2 to 25	Oracle Server X8-2, used as virtualization compute nodes
1	Oracle Rack Cabinet 1242 base
2	Cisco Nexus 9336C-FX2 Switches, providing connectivity for the data network
2	Cisco Nexus 9336C-FX2 Switches, providing connectivity for the spine network
1	Cisco Nexus 9348GC-FXP Switch, providing connectivity for the internal administration network
1	Oracle ZFS Storage Appliance ZS7-2, providing internal system storage

## Support for Previous Generations of Hardware Components

The latest version of the Oracle Private Cloud Appliance Controller Software continues to support all earlier configurations of the hardware platform. These may include the following components:

**Table 1.2 Supported Hardware**

Component Type	Component Name and Minimum Software Version
Management Nodes	<ul style="list-style-type: none"> <li>Oracle Server X5-2 (release 2.0.3 or newer)</li> <li>Sun Server X4-2 (release 1.1.3 or newer)</li> <li>Sun Server X3-2 (since initial release)</li> </ul>
Compute Nodes	<ul style="list-style-type: none"> <li>Oracle Server X7-2 (release 2.3.2 or newer)</li> <li>Oracle Server X6-2 (release 2.2.1 or newer)</li> <li>Oracle Server X5-2 (release 2.0.3 or newer)</li> <li>Sun Server X4-2 (release 1.1.3 or newer)</li> <li>Sun Server X3-2 (since initial release)</li> </ul>
Storage Appliance	<ul style="list-style-type: none"> <li>Oracle ZFS Storage Appliance ZS5-ES (release 2.3.3 or newer)</li> </ul>

Component Type	Component Name and Minimum Software Version
	<ul style="list-style-type: none"> <li>Oracle ZFS Storage Appliance ZS3-ES (release 1.1.3 or newer)</li> <li>Sun ZFS Storage Appliance 7320 (since initial release)</li> </ul>
InfiniBand Network Hardware	<ul style="list-style-type: none"> <li>Oracle Fabric Interconnect F1-15 (since initial release)</li> <li>NM2-36P Sun Datacenter InfiniBand Expansion Switch (since initial release)</li> </ul>
Internal Management Switch	<ul style="list-style-type: none"> <li>Oracle Switch ES1-24 (since initial release)</li> </ul>

## 1.5 Server Components

Table 1.3 lists the components of the server models that may be installed in an Oracle Private Cloud Appliance rack and which are supported by the current software release.



### Note

Oracle Private Cloud Appliance systems with factory-installed Controller Software release 2.4.x are delivered with Oracle Server X8-2, in the role of either management node or compute node. Previous generations of hardware components continue to be supported by the latest Controller Software. For details of previous server models, refer to the [Oracle Private Cloud Appliance Installation Guide for Release 2.3](#).

Table 1.3 Server Components

Quantity	Oracle Server X8-2 Management Node	Oracle Server X8-2 Compute Node
1	Oracle Server X8-2 base chassis	Oracle Server X8-2 base chassis
2	Intel Xeon 5218, 16 core, 2.3 GHz, 125W	Intel Xeon 8260, 24 core, 2.4 GHz, 165W
12-24	12x 32 GB DDR4 DIMMs (384 GB total)	12 or 24 DDR4 DIMMs  Available configurations: <ul style="list-style-type: none"> <li>12x 32 GB</li> <li>12x 64 GB</li> <li>24x 64 GB</li> </ul>
2	1.2 TB, 2.5-inch, 10,000-RPM drives with LSI RAID controller	1.2 TB, 2.5-inch, 10,000-RPM drives with LSI RAID controller
1	Ethernet port for remote management	Ethernet port for remote management
1	Dual-port 100Gbit Ethernet Host Channel Adapter Network Interface Card type CX5	Dual-port 100Gbit Ethernet Host Channel Adapter Network Interface Card type CX5
2	Redundant power supplies and fans	Redundant power supplies and fans
1 + 2	Embedded Ethernet: 1x Gigabit + 2x 10/25 Gigabit ports	Embedded Ethernet: 1x Gigabit + 2x 10/25 Gigabit ports

## 1.6 Operational Procedures for Oracle Private Cloud Appliance

This section provides important operational instructions to help you minimize the risk of injury or damage.

## 1.6.1 Emergency Power-off Considerations

If there is an emergency, then power to Oracle Private Cloud Appliance should be halted immediately. The following emergencies might require powering off the system:

- Natural disasters such as an earthquake, flood, hurricane, tornado or cyclone
- Abnormal noise, smell or smoke coming from the system
- Threat to human safety

### 1.6.1.1 Emergency Power-off Procedure

To perform an emergency power-off procedure for Oracle Private Cloud Appliance, turn off power at the circuit breaker or pull the emergency power-off switch in the computer room.

### 1.6.1.2 Emergency Power-off Switch

Emergency power-off (EPO) switches are required in a computer room when computer equipment contains batteries capable of supplying more than 750 volt-amperes for more than five minutes. Systems that have these batteries include internal EPO hardware for connection to a site EPO switch or relay. Use of the EPO switch will remove power from the Oracle Private Cloud Appliance.

## 1.6.2 Cautions and Warnings

The following cautions and warnings apply to Oracle Private Cloud Appliance:

- Do not touch the parts of this product that use high-voltage power. Touching them might result in serious personal injury.
- Do not power off Oracle Private Cloud Appliance unless there is an emergency. In that case, follow the [Section 1.6.1.1, “Emergency Power-off Procedure”](#).
- Keep the front and rear cabinet doors closed. Failure to do so might cause system failure or result in damage to hardware components.
- Keep the top, front, and back of cabinets clear to allow proper airflow and prevent overheating of components.
- Use only the supplied hardware.





---

# Chapter 2 Site Requirements

## Table of Contents

2.1 Space Requirements .....	7
2.1.1 Receiving and Unpacking Requirements .....	7
2.1.2 Maintenance Access Requirements .....	8
2.2 Flooring Requirements .....	8
2.3 Electrical Power Requirements .....	9
2.3.1 Facility Power Requirements .....	10
2.3.2 Circuit Breaker Requirements .....	10
2.3.3 Grounding Guidelines .....	10
2.4 Temperature and Humidity Requirements .....	10
2.5 Ventilation and Cooling Requirements .....	12

This chapter describes the site requirements for the Oracle Private Cloud Appliance.



### Note

For site checklists, refer to [Appendix D, Site Checklists](#).

## 2.1 Space Requirements

Oracle Private Cloud Appliance racks have the following space requirements:

- Height: 42U - 2000 mm (78.74 inches)
- Width: 600 mm with side panels (23.62 inches)
- Depth (front door handle to rear door handle): 1197 mm (47.12 inches)
- Depth (doors removed): 1112 mm (43.78 inches)
- Weight (base rack, fully populated): 1000 kg (2204 lbs)

The minimum ceiling height for the cabinet is 2914 mm (114.72 inches), measured from the true floor or raised floor, whichever is higher. This includes an additional 914 mm (36 inches) of space required above the rack height for maintenance access. The space surrounding the cabinet must not restrict the movement of cool air between the air conditioner and the front of the systems within the cabinet, or the movement of hot air coming out of the rear of the cabinet.

### 2.1.1 Receiving and Unpacking Requirements

Before your Oracle Private Cloud Appliance arrives, ensure that the receiving area is large enough for the package. The following are the package dimensions and weights for an Oracle Private Cloud Appliance system:

- Shipping height: 2159 mm (85 inches)
- Shipping width: 1219 mm (48 inches)
- Shipping depth: 1575 mm (62 inches)
- Shipping weight (base rack, fully populated): 1118 kg (2465 lbs)

If your loading dock meets the height and ramp requirements for a standard freight carrier truck, then you can use a pallet jack to unload the rack. If the loading dock does not meet the requirements, then you must provide a standard forklift or other means to unload the rack. You can also request that the rack be shipped in a truck with a lift gate.

When your Oracle Private Cloud Appliance arrives, leave the rack in its shipping packaging until it arrives at its installation site. Use a conditioned space to remove the packaging material to reduce particles before entering the data center. The entire access route to the installation site should be free of raised-pattern flooring that can cause vibration.

Allow enough space for unpacking the system from its shipping cartons. Ensure that there is enough clearance and clear pathways for moving the Oracle Private Cloud Appliance from the unpacking location to the installation location. [Table 2.1](#) lists the access route requirements for the Oracle Private Cloud Appliance.

**Table 2.1 Access Route Requirements**

Access Route Item	With Shipping Pallet	Without Shipping Pallet
Minimum door height	2184 mm (86 inches)	2000 mm (78.74 inches)
Minimum door width	1220 mm (48 inches)	600 mm (23.62 inches)
Minimum elevator depth	1575 mm (62 inches)	1058.2 mm (41.66 inches)
Maximum incline	6 degrees	6 degrees
Maximum elevator, pallet jack, and floor loading capacity	1145 kg (2520 lbs)	1145 kg (2520 lbs)

## 2.1.2 Maintenance Access Requirements

The maintenance area must be large enough for Oracle Private Cloud Appliance, and have the required access space. For example, the required space to remove the side panels is 675.64 mm (26.6 inches). [Table 2.2](#) lists the maintenance access requirements for the Oracle Private Cloud Appliance.

**Table 2.2 Maintenance Access Requirements**

Location	Maintenance Access Requirement
Rear maintenance	914 mm (36 inches)
Front maintenance	1232 mm (48.5 inches)
Top maintenance	914 mm (36 inches)

## 2.2 Flooring Requirements

Oracle recommends that the Oracle Private Cloud Appliance be installed on raised flooring. The site floor and the raised flooring must be able to support the total weight of the system as specified in [Section 2.1, "Space Requirements"](#).

[Table 2.3](#) lists the floor load requirements.

**Table 2.3 Floor Load Requirements**

Description	Requirement
Maximum allowable weight of installed rack equipment	952.54 kg (2100 lbs)
Maximum allowable weight of installed power distribution units	52.16 kg (115 lbs)

Description	Requirement
Maximum dynamic load (maximum allowable weight of installed equipment including power distribution units)	1004.71 kg (2215 lbs)

## 2.3 Electrical Power Requirements

The Oracle Private Cloud Appliance can operate effectively over a wide range of voltages and frequencies. However, it must have a reliable power source. Damage might occur if the ranges are exceeded. Electrical disturbances such as the following might damage the system:

- Fluctuations caused by brownouts
- Wide and rapid variations in input voltage levels or in input power frequency
- Electrical storms
- Faults in the distribution system, such as defective wiring

To protect your system from such disturbances, you should have a dedicated power distribution system, power-conditioning equipment, as well as lightning arresters or power cables to protect from electrical storms.

Each rack has two pre-installed power distribution units (PDUs). The PDUs accept different power sources. You must specify the type of PDU that is correct for your data center.

[Table 2.4](#) lists the PDU low-voltage requirements.

**Table 2.4 PDU Types and Specifications**

Item	15kVA - type 6440A	15kVA - type 6441A	15kVA - type 6442A	22kVA - type 7100873	22kVA - type 7100874	24kVA - type 6444A	24kVA - type 6445A
Region	North and South America, Japan and Taiwan	EMEA and APAC (except Japan and Taiwan)	North and South America, Japan and Taiwan	North and South America, Japan and Taiwan	EMEA and APAC (except Japan and Taiwan)	North and South America, Japan and Taiwan	EMEA and APAC (except Japan and Taiwan)
Number of inputs	1	1	3	3	3	2	2
Source voltage	200-220V	220/380 - 240/415VAC	200-240V	200-240V	200-240V	200-220V	220/380 - 240/415VAC
Frequency	50/60 Hz	50/60 Hz	50/60 Hz	50/60 Hz	50/60 Hz	50/60 Hz	50/60 Hz
Max. line current	40A	21A	24A	36.8A	32A	34.6A	16A
Phase current	23A	21A	N/A	N/A	32A	20A	16A
Recommended circuit breaker	50A	25A	30A	50A	30A	50A	25A
Outlets	42 x C13 6 x C19	42 x C13 6 x C19	42 x C13 6 x C19	42 x C13 6 x C19	42 x C13 6 x C19	42 x C13 6 x C19	42 x C13 6 x C19
Outlet groups	6	6	6	6	6	6	6
Max. group current	20A	20A	20A	20A	20A	20A	18A

Item	15kVA - type 6440A	15kVA - type 6441A	15kVA - type 6442A	22kVA - type 7100873	22kVA - type 7100874	24kVA - type 6444A	24kVA - type 6445A
Output voltage	200-240V	220-240V	200-240V	200-240V	200-240V	220-240V	200-240V
Data center receptacle	3ph 3W + ground 60A IEC60309	3ph 4W + ground 32A IEC60309 5P 6h 400V	1ph 2W + ground 30A NEMA L6-30	1ph 2W + ground 50A Hubbell CS8265C	1ph 2W + ground 32A IEC60309 3P 6H 230V	3ph 3W + ground 60A IEC60309	3ph 4W + ground 32A IEC60309 5P 6h 400V

### 2.3.1 Facility Power Requirements

Electrical work and installations must comply with applicable local, state, or national electrical codes. Contact your facilities manager or qualified electrician to determine what type of power is supplied to the building.

To prevent catastrophic failures, design the input power sources to ensure adequate power is provided to the power distribution units (PDUs). Use dedicated AC breaker panels for all power circuits that supply power to the PDU. When planning for power distribution requirements, balance the power load between available AC supply branch circuits. In the United States and Canada, ensure that the overall system AC input current load does not exceed 80 percent of the branch circuit AC current rating.

PDU power cords are 4 meters (13.12 feet) long, and 1 to 1.5 meters (3.3 to 4.9 feet) of the cord will be routed within the rack cabinet. The installation site AC power receptacle must be within 2 meters (6.6 feet) of the rack.

### 2.3.2 Circuit Breaker Requirements

To prevent catastrophic failures, the design of your power system must ensure that adequate power is provided to all of the compute nodes. Use dedicated AC breaker panels for all power circuits that supply power to the compute nodes. Electrical work and installations must comply with applicable local, state, or national electrical codes. Compute nodes require electrical circuits to be grounded to the Earth.

In addition to circuit breakers, provide a stable power source, such as an uninterruptible power supply (UPS) to reduce the possibility of component failures. If computer equipment is subjected to repeated power interruptions and fluctuations, then it is susceptible to a higher rate of component failure.



**Note**

Circuit breakers are supplied by the customer. One circuit breaker is required for each power cord.

### 2.3.3 Grounding Guidelines

The cabinets for the Oracle Private Cloud Appliance are shipped with grounding-type power cords (three-wire). Always connect the cords to grounded power outlets. Because different grounding methods are used, depending on location, check the grounding type, and refer to documentation, such as IEC documents, for the correct grounding method. Ensure that the facility administrator or qualified electrical engineer verifies the grounding method for the building, and performs the grounding work.

## 2.4 Temperature and Humidity Requirements

Airflow through the Oracle Private Cloud Appliance is from front to back. For information, see [Section 2.5, "Ventilation and Cooling Requirements"](#).

Studies have shown that temperature increases of 10 degrees Celsius (18 degrees Fahrenheit) above 20 degrees Celsius (68 degrees Fahrenheit) reduce long-term electronics reliability by 50 percent. Excessive internal temperatures might result in full or partial shutdown of the Oracle Private Cloud Appliance.

Table 2.5 lists the temperature, humidity and altitude requirements for operating and non-operating systems.

**Table 2.5 Temperature, Humidity, and Altitude Requirements**

Condition	Operating Requirement	Non-operating Requirement	Optimum
Temperature	5 ° to 32 ° Celsius (41 ° to 89.6 ° Fahrenheit)	-40 ° to 68 ° Celsius (-40 ° to 154 ° Fahrenheit)	For optimal rack cooling, data center temperatures from 21 ° to 23 ° Celsius (69.8 ° to 73.4 ° Fahrenheit)
Relative humidity	10 to 90 percent relative humidity, non-condensing	Up to 93 percent relative humidity	For optimal data center rack cooling, 45 to 50 percent non-condensing
Altitude	3,000 meters (9,840 feet) maximum	12,000 meters (39,370 feet)	Ambient temperature is reduced by 1 ° Celsius per 300 meters above 900 meters altitude above sea level

Set conditions to the optimal temperature and humidity ranges to minimize the chance of downtime due to component failure. Operating an Oracle Private Cloud Appliance for extended periods at or near the operating range limits, or installing it in an environment when it remains at or near non-operating range limits could significantly increase hardware component failure.

The ambient temperature range of 21 ° to 23 ° Celsius (69.8 ° to 73.4 ° Fahrenheit) is optimal for server reliability and operator comfort. Most computer equipment can operate in a wide temperature range, but near 22 ° Celsius (71.6 ° Fahrenheit) is desirable because it is easier to maintain safe humidity levels. Operating in this temperature range provides a safety buffer in the event that the air conditioning system goes down for a period of time.

The ambient relative humidity range of 45 to 50 percent is suitable for safe data processing operations. Most computer equipment can operate in a wide range (20 to 80 percent), but the range of 45 to 50 percent is recommended for the following reasons:

- Optimal range helps protect computer systems from corrosion problems associated with high humidity levels.
- Optimal range provides the greatest operating time buffer in the event of air conditioner control failure.
- This range helps to avoid failures or temporary malfunctions caused by intermittent interference from static discharges that may occur when relative humidity is too low.



**Note**

Electrostatic discharge (ESD) is easily generated, and hard to dissipate in areas of low relative humidity, such as below 35 percent. ESD becomes critical when humidity drops below 30 percent. It is not difficult to maintain humidity in a data center because of the high-efficiency vapor barrier and low rate of air changes normally present.

## 2.5 Ventilation and Cooling Requirements

Always provide adequate space in front of and behind the rack to allow for proper ventilation. Do not obstruct the front or rear of the rack with equipment or objects that might prevent air from flowing through the rack. Rack-mountable servers and equipment typically draw cool air in through the front of the rack and let warm air out the rear of the rack. There is no airflow requirement for the left and right sides due to front-to-back cooling.

If the rack is not completely filled with components, then cover the empty sections with filler panels. Gaps between components can adversely affect airflow and cooling within the rack.

Relative humidity is the percentage of the total water vapor that can exist in the air without condensing, and is inversely proportional to air temperature. Humidity goes down when the temperature rises, and goes up when the temperature drops. For example, air with a relative humidity of 45 percent at a temperature of 24 ° Celsius (75.2 ° Fahrenheit) has a relative humidity of 65 percent at a temperature of 18 ° Celsius (64.4 ° Fahrenheit). As the temperature drops, the relative humidity rises to more than 65 percent, and water droplets are formed.

Air conditioning facilities usually do not precisely monitor or control temperature and humidity throughout an entire computer room. Generally, monitoring is done at individual points corresponding to multiple exhaust vents in the main unit, and other units in the room. Special consideration should be paid to humidity when using underfloor ventilation. When underfloor ventilation is used, monitoring is done at each point close to an exhaust vent. Distribution of the temperature and humidity across the entire room is uneven.

The Oracle Private Cloud Appliance has been designed to function while installed in a natural convection air flow. The following requirements must be followed to meet the environmental specification:


- Ensure that there is adequate airflow through the system.
- Ensure that the system has front-to-back cooling. The air intake is at the front of the system, and the air outlet is at the rear of the system.
- Allow a minimum clearance of 1219.2 mm (48 inches) at the front of the system, and 914 mm (36 inches) at the rear of the system for ventilation.

Use perforated tiles, approximately 400 CFM/tile, in front of the rack for cold air intake. The tiles can be arranged in any order in front of the rack, as long as cold air from the tiles can flow into the rack. Inadequate cold airflow could result in a higher intake temperature in the system due to exhaust air recirculation. The following is the recommended number of floor tiles:

- Four floor tiles for an Oracle Private Cloud Appliance with up to 25 compute nodes (fully loaded)
- Three floor tiles for an Oracle Private Cloud Appliance with up to 16 compute nodes (half loaded)
- One floor tile for an Oracle Private Cloud Appliance with 8 compute nodes (quarter loaded)

[Figure 2.1](#) shows a typical installation of the floor tiles in a data center for Oracle Private Cloud Appliance with more than 16 compute nodes.

**Figure 2.1 Typical Data Center Configuration for Perforated Floor Tiles**







---

# Chapter 3 Network Requirements

## Table of Contents

3.1 Network Connection and IP Address Requirements for Ethernet-based Systems .....	15
3.1.1 Network Connection Requirements for Ethernet-based Systems .....	15
3.1.2 IP Address Requirements for Ethernet-based Systems .....	20
3.2 Network Connection and IP Address Requirements for InfiniBand-based Systems .....	21
3.2.1 Network Connection Requirements for InfiniBand-based Systems .....	21
3.2.2 IP Address Requirements for InfiniBand-based Systems .....	23
3.3 General Network Configuration .....	24

Oracle Private Cloud Appliance exists in two different types of network architecture. One is built around a physical InfiniBand fabric; the other relies on physical high speed Ethernet connectivity. While the two implementations offer practically the same functionality, there are visible hardware and configuration differences. Consequently, each network architecture has its own requirements.

New systems with factory-installed Controller Software Release 2.4.x have an Ethernet-based network architecture. For customers upgrading a system with InfiniBand-based network architecture to Release 2.4.2 or newer, those network requirements are included as well. Refer to the sections that apply to your appliance.

## 3.1 Network Connection and IP Address Requirements for Ethernet-based Systems

This section describes the network connection requirements and IP address requirements to connect the Oracle Private Cloud Appliance with Ethernet-based network architecture to your existing network infrastructure.


### 3.1.1 Network Connection Requirements for Ethernet-based Systems

Before installation, you must run network cables from your existing network infrastructure to the Oracle Private Cloud Appliance installation site.

You must connect 2 high-speed Ethernet ports on each *spine* Cisco Nexus 9336C-FX2 Switch to your data center public Ethernet network. These are ports 5/1 and 5/2, which are configured at 10Gbit and intended as the default connection for appliance management and setup. This configuration is locked and must remain at 10Gbit; it requires breakout cables with a QSFP+ transceiver on the spine switch end and four SFP+ transceivers on the other end.

[Figure 3.1](#) shows the location of the two 10 GbE breakout ports on the spine switches, which must be connected to the data center network.

Figure 3.1 Default 10GbE Breakout Ports on Spine Cisco Nexus 9336C-FX2 Switch

**Caution**

It is critical that **both** spine switches have **two** 10GbE connections each to a pair of next-level data center switches. This configuration with four cable connections provides redundancy and load splitting at the level of the spine switches, the 10GbE ports and the data center switches. This outbound cabling from the pair of spine Cisco Nexus 9336C-FX2 Switches should be crossed: each spine switch must be connected to two next-level data center switches. The cabling pattern plays a key role in the continuation of service during failover scenarios.

### 3.1.1.1 Data Center Switch Configuration Notes

When configuring the data center switches to accept incoming Oracle Private Cloud Appliance uplinks – the default uplinks as well as any custom uplinks you define – take these notes into account.

- All uplinks, default and customer, are configured to use link aggregation (LACP). All four switch ports included in an uplink configuration – two per spine switch – must belong to the same link aggregation group (LAG), also known as port channel. The switch ports on the data center side of the uplinks must be configured accordingly.
- The data center switch ports to which the Oracle Private Cloud Appliance uplinks are connected, must be set to trunk mode.
- The spine switches operate with the Virtual Port Channel (vPC) feature enabled. Therefore, Spanning Tree Protocol (STP) is not supported in the data center network topology involving Oracle Private Cloud Appliance.
- Auto-negotiation is not available for uplink ports. Transfer speed must be specified on the customer switches' end. For the default uplink ports, this is 10Gbit/s.

## Example

Below is an example of a custom uplink configuration for your reference.

In this example, port 3 on both spine switches is split into 4x 25Gbit. Ports 3:1 and 3:2 are cross-cabled to a pair of ToR switches. Like the internal spine switches of the Oracle Private Cloud Appliance, the two ToR switches have the virtual port channel (vPC) feature enabled, allowing ports from both switches to operate within the same port channel.

These are the key configuration parameters:

- The vPC feature must be enabled on both ToR switches: vpc domain 1.
- Ports 49 and 50 on **both** ToR switches must belong to port channel 200.
- Port channel 200 on **both** ToR switches must belong to virtual port channel 200.
- Port channel parameters, such as MTU, speed and spanning tree filter, must be identical on both ToR switches.
- All port channel and switch port interfaces are set to trunk mode.

ToR switch A configuration:

```
feature lacp
feature vpc
vpc domain 1
  peer-switch
  role priority 1
  peer-keepalive destination <ToR-B_IP> source <ToR-A_IP>
  peer-gateway
  layer3 peer-router
  ip arp synchronize

interface port-channel 200
  description "Uplink PCA 3:1 3:2 port channel"
  switchport mode trunk
  switchport trunk allowed vlan 1-4094      ### optional, typical default setting
  switchport trunk native vlan 500         ### optional, for data center vlan
  speed 25000
  mtu 9216
  vpc 200

interface Ethernet 1/49
  description "PCA spine 1 Port 3:1"
  switchport mode trunk
  switchport trunk allowed vlan 1-4094      ### optional, typical default setting
  switchport trunk native vlan 500         ### optional, for data center vlan
  speed 25000
  mtu 9216
  channel-group 200 mode active

interface Ethernet 1/50
  description "PCA spine 2 Port 3:1"
  switchport mode trunk
  switchport trunk allowed vlan 1-4094      ### optional, typical default setting
  switchport trunk native vlan 500         ### optional, for data center vlan
  speed 25000
  mtu 9216
  channel-group 200 mode active
```

ToR switch B configuration:

```
feature lacp
```

```

feature vpc
vpc domain 1
  peer-switch
  role priority 2
  peer-keepalive destination <ToR-A_IP> source <ToR-B_IP>
  peer-gateway
  layer3 peer-router
  ip arp synchronize

interface port-channel 200
  description "Uplink PCA 3:1 3:2 port channel"
  switchport mode trunk
  switchport trunk allowed vlan 1-4094      ### optional, typical default setting
  switchport trunk native vlan 500         ### optional, for data center vlan
  speed 25000
  mtu 9216
  vpc 200

interface Ethernet 1/49
  description "PCA spine 1 Port 3:2"
  switchport mode trunk
  switchport trunk allowed vlan 1-4094     ### optional, typical default setting
  switchport trunk native vlan 500        ### optional, for data center vlan
  speed 25000
  mtu 9216
  channel-group 200 mode active

interface Ethernet 1/50
  description "PCA spine 2 Port 3:2"
  switchport mode trunk
  switchport trunk allowed vlan 1-4094     ### optional, typical default setting
  switchport trunk native vlan 500        ### optional, for data center vlan
  speed 25000
  mtu 9216
  channel-group 200 mode active

```

### 3.1.1.2 Appliance Uplink Configuration

The uplinks are the system's physical connections between the spine switches and the data center network or top-of-rack (ToR) switches. Ports 1-4 on the spine switches are reserved for customer use. The uplink configuration must comply with requirements and restrictions, which are explained in this section.

For the purpose of high availability, uplink ports are always configured in pairs on each spine switch. Each port pair must operate in one of these modes, using the appropriate transceivers and cables:

- 100Gbit full port bandwidth
- 40Gbit full port bandwidth
- 4x 25Gbit with 4-way breakout cable
- 4x 10Gbit with 4-way breakout cable



#### Caution


The Cisco Nexus 9336C-FX2 Switch ports are of the QSFP28 100Gbit type. Depending on your data center network infrastructure, each of the customer-reserved ports can be used at maximum bandwidth or split into 25Gbit (SFP28) or 10Gbit (SFP+). External Ethernet connections from the Oracle Private Cloud Appliance require a minimum speed of 10Gbit/s. Verify the cable and transceiver requirements for your intended configuration. Only Oracle-supplied transceivers should be installed in the Oracle Private Cloud Appliance spine switches.

For custom networks through spine switch ports 1-4, the port pairs and their mode of operation are configured through the Oracle Private Cloud Appliance CLI. The configuration of the default uplink ports on spine switch port 5 cannot be modified: it operates in 4x 10Gbit mode with ports 5/1 and 5/2 cross-cabled to the ToR switches, while ports 5/3 and 5/4 remain unused.

Spine switch ports 1 and 2, and ports 3 and 4, must be configured at the same time. If one port uses a breakout cable, then the other port must also be set up that way. All four ports on the same breakout cable must be set to the same speed, but different breakout cables can be set up for different speeds.

This results in the following supported configurations:

**Figure 3.2 Supported Uplink Port Configurations**



**Caution**

Each port pair on each individual spine switch must be cross-connected to two ToR switches. Each port pair that is configured for external connectivity – identically on both spine switches by definition – corresponds with a separate VLAN. The deduplication of ports and switches is primarily for redundancy; the VLANs already separate traffic.

The spine switches allow a maximum supported number of 8 custom uplinks, corresponding with 16 port pairs. This number does not include the default uplink.

**3.1.1.3 Optional Data Center Administration Network Uplink**

In addition to the public Ethernet connection, you may connect the Cisco Nexus 9348GC-FXP Switch to a management or machine administration network at your installation site. If you choose to use such an uplink, consider it as a long-term alternative to the temporary workstation connection described in [Section 3.1.2, “IP Address Requirements for Ethernet-based Systems”](#). Configure the administration uplink *after* the initialization of the appliance, when the appliance network settings have been applied.



**Caution**

Connect **port 2** on the Cisco Nexus 9348GC-FXP Switch.

Make sure that the data center Ethernet switch used in this connection is configured to prevent DHCP leakage to the 192.168.4.0/24 subnet used by Oracle Private Cloud Appliance. Do not connect to any network with any kind of broadcast services in addition to DHCP.

For the *bastion host*, which is the name used to describe the machine that is permanently connected to the data center administration network, use the IP address 192.168.4.199 and assign it statically to its network interface. Make sure there is no other machine on the same subnet using the same IP address and causing IP conflicts.

A connection to the appliance internal management network, either directly into the Cisco Nexus 9348GC-FXP Switch or through an additional Ethernet switch in the data center, is **not required** to access the appliance management functionality of the Oracle Private Cloud Appliance Dashboard. The primary role of the appliance internal management network is to allow the controller software on the management nodes to interact with the compute nodes and other rack components. Connecting to this network from outside the appliance allows you to gain direct administrator access to each component, for example to control the ILOMs.

**Caution**

Do not make any changes to anything on this network unless directed to do so by Oracle Support.

The Oracle Private Cloud Appliance Dashboard and Oracle VM Manager UI, in contrast, are **not** accessed over the internal management network, but through the management node cluster's virtual IP address. The virtual IP address is public, in the sense that it should be accessible from within your data center network. You reserve this IP address in your data center network in advance, as described in [Section 3.3, “General Network Configuration”](#), and then configure it in the appliance network settings during software initialization. Instructions for first access to the Oracle Private Cloud Appliance Dashboard are provided in [Section 5.3.5, “Connect the Appliance to the Network”](#).

### 3.1.2 IP Address Requirements for Ethernet-based Systems

The Oracle Private Cloud Appliance requires a large number of private IP addresses. These are reserved for the normal operation and future expansion of the system and cannot be reconfigured. While system-level subnets and IPs should be avoided, the customer is allowed to choose the network configuration for external and virtual machine traffic that suits their requirements.

For lists of default IP addresses that are preassigned to your Oracle Private Cloud Appliance, see [Appendix A, Default IP Addresses](#).

To gain initial access to the Oracle Private Cloud Appliance Dashboard, you must connect a terminal or workstation with a web browser directly to the internal management network. Connect an Ethernet cable to the available port 48 in the Cisco Nexus 9348GC-FXP Switch, and statically configure the wired network connection of the workstation to use the IP address [192.168.4.254](#).

**Caution**

The IP address [192.168.4.254](#) and switch port 48 have been reserved specifically for the purpose of connecting a workstation during the first boot, initialization and provisioning of the appliance.

During the initial software configuration of the Oracle Private Cloud Appliance, you configure the network settings of the management nodes. For this purpose, you should reserve three IP addresses in the public (data center) network: one for each management node, and one to be used as virtual IP address shared by both management nodes. If the data center network traffic is tagged, make sure that the VLAN ID is also provided as part of the configuration. The virtual IP address provides access to the Dashboard once the software initialization is complete.

To avoid network interference and conflicts, you must ensure that the data center network does not overlap with any of the infrastructure networks of the Oracle Private Cloud Appliance internal configuration. These are the subnets and VLANs you should keep clear:

**Subnets:**

- 192.168.4.0/24 – internal machine administration network: connects ILOMs and physical hosts
- 192.168.32.0/21 – internal management network: traffic between management and compute nodes
- 192.168.64.0/21 – underlay network for east/west traffic within the appliance environment
- 192.168.72.0/21 – underlay network for north/south traffic, enabling external connectivity
- 192.168.40.0/21 – storage network: traffic between the servers and the ZFS storage appliance



**Note**

Each /21 subnet comprises the IP ranges of eight /24 subnets or over 2000 IP addresses. For example: 192.168.32.0/21 corresponds with all IP addresses from 192.168.32.1 to 192.168.39.255.

**VLANs:**

- 1 – the Cisco default VLAN
- 3040 – the default service VLAN
- 3041-3072 – a range of 31 VLANs reserved for customer VM and host networks
- 3073-3099 – a range reserved for system-level connectivity



**Note**

VLANs 3090-3093 are already in use for tagged traffic over the /21 subnets listed above.

- 3968-4095 – a range reserved for Cisco internal device allocation

## 3.2 Network Connection and IP Address Requirements for InfiniBand-based Systems

This section describes the network connection requirements and IP address requirements to connect the Oracle Private Cloud Appliance with InfiniBand-based network architecture to your existing network infrastructure.

### 3.2.1 Network Connection Requirements for InfiniBand-based Systems

Before installation, you must run network cables from your existing network infrastructure to the Oracle Private Cloud Appliance installation site.

You must connect two 10 Gigabit Ethernet (GbE) IO module ports labeled “Public” on each Fabric Interconnect to your public Ethernet network.




**Caution**

The IO modules only support 10 GbE transport and cannot be connected to gigabit Ethernet switches. The Oracle Private Cloud Appliance must be connected externally to 10GbE optical switch ports.

Figure 3.3 shows the location of the 10 GbE Public IO module ports on the Fabric Interconnect.



**Figure 3.3 Oracle Fabric Interconnect F1-15 10 GbE Public IO Module Ports****Caution**

It is critical that **both** Fabric Interconnects have **two** 10GbE connections each to a pair of next-level data center switches. This configuration with four cable connections provides redundancy and load splitting at the level of the Fabric Interconnects, the 10GbE ports and the data center switches. This outbound cabling should not be crossed or meshed, because the internal connections to the pair of Fabric Interconnects are already configured that way. The cabling pattern plays a key role in the continuation of service during failover scenarios involving Fabric Interconnect outages and other components.

**Caution**

Do not configure any type of link aggregation group (LAG) across the 10GbE ports: LACP, network/interface bonding or similar methods to combine multiple network connections are not supported.

To provide additional bandwidth to the environment hosted by the Oracle Private Cloud Appliance, create custom networks. For detailed information about network customization, refer to the *Oracle Private Cloud Appliance Administrator's Guide*.

**Optional Data Center Administration Network Uplink**

In addition to the public Ethernet connection, you may connect one of the two Oracle Switch ES1-24 switches to a management or machine administration network at your installation site. The system contains two Oracle Switch ES1-24 switches, making up one logical internal management network. The daisy-chained Oracle Switch ES1-24 arrangement has a special high availability (HA) configuration. Link tracking is done in software on the Oracle Switch ES1-24.

If you choose to use an uplink to the data center administration network, consider it as a long-term alternative to the temporary workstation connection described in [Section 3.2.2, "IP Address Requirements for InfiniBand-based Systems"](#). Configure the administration uplink *after* the initialization of the appliance, when the appliance default network settings have been reconfigured.

**Caution**

Connect port 24 on one Oracle Switch ES1-24 – never both.



Make sure that the data center Ethernet switch used in this connection is configured to prevent DHCP leakage to the 192.168.4.0/24 subnet used by Oracle Private Cloud Appliance.

Use a workstation connected to the data center administration network and statically assign the IP address 192.168.4.254. Make sure there is no other machine directly connected to an internal Oracle Switch ES1-24, using the same IP address and causing IP conflicts.

A connection to the appliance internal management network, either directly into an Oracle Switch ES1-24 or through an additional Ethernet switch in the data center, is not required to access the appliance management functionality of the Oracle Private Cloud Appliance Dashboard. The primary role of the appliance internal management network is to allow the controller software on the management nodes to interact with the compute nodes and other rack components. Connecting to this network from outside the appliance allows you to gain direct administrator access to each component, for example to control the ILOMs.

The Oracle Private Cloud Appliance Dashboard, in contrast, is not accessed over the internal management network, but through the management node cluster's virtual IP address. The virtual IP address is public, in the sense that it should be accessible from within your data center network. You reserve this IP address in your data center network in advance, as described in [Section 3.3, "General Network Configuration"](#), and then configure it in the appliance network settings during software initialization. Instructions for first access to the Oracle Private Cloud Appliance Dashboard are provided in [Section 5.3.5, "Connect the Appliance to the Network"](#).

## 3.2.2 IP Address Requirements for InfiniBand-based Systems

The Oracle Private Cloud Appliance requires a large number of IP addresses. For lists of default IP addresses that are preassigned to your Oracle Private Cloud Appliance, see [Appendix A, Default IP Addresses](#).

To gain initial access to the Oracle Private Cloud Appliance Dashboard, you must connect a terminal or workstation with a web browser directly to the management network. Connect an Ethernet cable to the available port 19 in one of the Oracle Switch ES1-24 switches, and statically configure the wired network connection of the workstation to use the IP address `192.168.4.254`.



### Caution

The IP address `192.168.4.254` is the only one available for customer use in the appliance management network.

During the initial software configuration of the Oracle Private Cloud Appliance, you reconfigure the network settings of the management nodes. For this purpose, you should reserve three IP addresses in the public (data center) network: one for each management node, and one to be used as virtual IP address by both management nodes. The virtual IP address provides access to the Dashboard once the software initialization is complete.

To avoid network interference and conflicts, you must ensure that the data center network does not overlap with any of the infrastructure subnets of the Oracle Private Cloud Appliance default configuration. These are the subnets you should keep clear:

- 192.168.140.0/24
- 192.168.40.0/24

- 192.168.4.0/24

### 3.3 General Network Configuration

**Table 3.1** is a general network configuration worksheet to help you configure your Oracle Private Cloud Appliance. Oracle requests that you collect the information in preparation of the installation in your data center. The parameters entered in this worksheet are used as input during the initialization of the appliance.

**Table 3.1 General Network Configuration Worksheet**

Item	Your Configuration	Description and Example
Domain Name		Company network domain name.  Example: <code>abc.example.com</code>
Region		Name of the country in which Oracle Private Cloud Appliance is installed.
Time Zone		Valid time zone.
IP address of the Domain Name Server		IP address of one or more network name servers (up to three servers).  Example: <code>10.25.45.123</code> , <code>10.25.45.125</code>
NTP Server		IP address of a Network Time Protocol server.  Example: <code>10.12.13.14</code>
Search Domains		A list of search domains for name lookup.  Example: <code>example.com</code> , <code>example.org</code>
Default Gateway		IP address of the default gateway in your organization.  Example: <code>10.203.72.2</code>
IP address of Management node 1		IP address of the first management node in your company network.  Example: <code>10.203.72.101</code>
Host name of Management node 1		Host name of the first management node in your company network.  Example: <code>manager1</code>
IP address of Management node 2		IP address of the second management node in your company network.  Example: <code>10.203.72.102</code>
Host name of Management node 2		Host name of the second management node in your company network.  Example: <code>manager2</code>
Management virtual IP address		Virtual IP address of the management node cluster in your company network.

Item	Your Configuration	Description and Example
Management VLAN ID		Example: 10.203.72.100 Data center VLAN tag for the management network, if applicable. Example: VLAN tag 20



---

# Chapter 4 Installing Oracle Private Cloud Appliance

## Table of Contents

4.1 Installation Task Map .....	27
4.2 Prepare to Install Oracle Private Cloud Appliance .....	27
4.2.1 Before You Begin .....	27
4.2.2 Review Safety Guidelines .....	28
4.2.3 Prepare the Installation Site .....	28
4.2.4 Tools Required for Unpacking and Installation .....	29
4.2.5 Unpack the System .....	30
4.3 Install Oracle Private Cloud Appliance in Its Allocated Space .....	30
4.3.1 Move the System .....	30
4.3.2 Stabilize the Appliance Rack .....	31
4.3.3 Attach a Ground Cable (Optional) .....	33
4.4 What's Next? .....	34

This chapter explains how to prepare for the installation of your Oracle Private Cloud Appliance and how to install the system at your site.

## 4.1 Installation Task Map

[Table 4.1](#) describes the steps for installing an Oracle Private Cloud Appliance at your site.

**Table 4.1 Steps for Installing Oracle Private Cloud Appliance**

Step	Description	Links
1	Review the safety precautions, guidelines, site checklists, and site requirements.	<a href="#">Section 4.2.2, “Review Safety Guidelines”</a>
2	Ensure that the site is prepared for the installation of the system.	<a href="#">Section 4.2.3, “Prepare the Installation Site”</a>
3	Remove all packaging from the system.	<a href="#">Section 4.2.5, “Unpack the System”</a>
4	Install the system in its allocated place.	<a href="#">Section 4.3, “Install Oracle Private Cloud Appliance in Its Allocated Space”</a>

## 4.2 Prepare to Install Oracle Private Cloud Appliance

This section describes the preparation steps to perform before installing your Oracle Private Cloud Appliance.

### 4.2.1 Before You Begin

Before installing your Oracle Private Cloud Appliance, or installing any server equipment into the rack, read the *Important Safety Information for Oracle’s Hardware Systems (7063567)* document that is shipped with the system.

Observe all safety notices printed on the packaging and listed in the *Oracle Rack Cabinet 1242 Safety and Compliance Guide (Part No: E87279)* and the *Oracle Rack Cabinet 1242 Power Distribution Units User’s*

*Guide* (Part No: E87281). Go to [https://docs.oracle.com/cd/E85660\\_01/index.html](https://docs.oracle.com/cd/E85660_01/index.html) to download these guides.

## 4.2.2 Review Safety Guidelines

Before your Oracle Private Cloud Appliance arrives, review the following safety precautions to ensure that the site is safe, as well as ready for delivery. Failure to observe these precautions can result in personal injury, equipment damage, or malfunction.

- Do not block ventilation openings.
- Do not install the system in a location that is exposed to direct sunlight or near a device that may become hot.
- Do not install the system in a location that is exposed to excessive dust, corrosive gases, or air with high salt concentrations.
- Do not install the system in a location that is exposed to frequent vibrations. Install the system on a flat, level surface.
- Use a power outlet that uses proper grounding. When using shared grounding, the grounding resistance must not be greater than 10 ohms. Ensure that your facility administrator or a qualified electrical engineer verifies the grounding method for the building, and performs the grounding work.
- Ensure that each grounding wire used for the system is used exclusively for the Oracle Private Cloud Appliance. Also observe the precautions, warnings, and notes about handling that appear on labels on the equipment.
- Do not place cables under the equipment or stretch the cables too tightly.
- Do not disconnect power cords from the equipment while its power is on.
- Do not place anything on top of the system or perform any work directly above it.
- If you cannot reach the connector lock when disconnecting LAN cables, then press the connector lock with a flathead screwdriver to disconnect the cable. You could damage the system board if you force your fingers into the gap rather than using a flathead screwdriver.
- Do not let the room temperature rise sharply, especially in winter. Sudden temperature changes can cause condensation to form inside the system. Allow for a sufficient warm-up period prior to server operation.
- Do not install the system near a photocopy machine, air conditioner, welding machine, or any other equipment that generates loud, electronic noises.
- Avoid static electricity at the installation location. Static electricity transferred to the system can cause malfunctions. Static electricity is often generated on carpets.
- Confirm that the supply voltage and frequency match the electrical ratings indicated on your Oracle Private Cloud Appliance.
- Do not insert anything into any Oracle Private Cloud Appliance opening, unless doing so is part of a documented procedure. The system contains high-voltage parts. If a metal object or other electrically conductive object enters an opening in the system, then it could cause a short circuit. This could result in personal injury, fire, electric shock, and equipment damage.

## 4.2.3 Prepare the Installation Site

The following procedure describes how to prepare the site prior to unpacking and situating your Oracle Private Cloud Appliance.

1. Thoroughly clean and vacuum the area in preparation for the installation.
2. Note problems or peculiarities at the site that require special equipment.
3. Verify that the installation site flooring has a strength rating to withstand the combined weight of the Oracle Private Cloud Appliance and any other installed equipment.

For more information, see [Section 2.1, "Space Requirements"](#).

4. Install all necessary electrical equipment and ensure that sufficient power is provided.

See also the *Oracle Rack Cabinet 1242 Power Distribution Units User's Guide* for the system Power Distribution Unit (PDU) power requirements.

5. Ensure that the installation site provides adequate air conditioning.

For details, see [Chapter 2, Site Requirements](#).

6. Operate the air conditioning system for 48 hours to bring the room temperature to the appropriate level.

### 4.2.4 Tools Required for Unpacking and Installation

The following tools and equipment are required for installation and service:

- Unpacking tools (provided in rack ship kit)
  - 17mm and 13mm open-ended double-sided wrench
  - Allen L-key 6mm hex wrench, 5-3/4 inches in length
  - Allen L-key 8mm hex wrench, 6-3/8 inches in length
- Rack setup tools and accessories (provided in rack ship kit)
  - 32 M6 cage nuts and washers
  - 32 M6 pan head screws
  - 8 M5 screws
  - 8 self-tapping Torx screws
  - 4 PDU brackets
  - 4 M5 cage nuts
  - 4 grounded power cables
  - 18mm and 16mm open-ended double-sided wrench
  - T-40 Torx screwdriver
  - T-25 Torx screwdriver
  - T-15 Torx screwdriver

- 8mm and 6mm Allen wrench
- Cage nut tool
- Keys to the front door, rear door, and side panel locks
- 17mm and 13mm open-ended double-sided wrench
- Customer-supplied tools (not provided in rack ship kit)
  - No. 2 Phillips screwdriver
  - Diagonal cutter
  - Antistatic wrist strap

## 4.2.5 Unpack the System

Before unpacking the rack from the shipping carton, refer to the labels on the carton and to the instructions that they provide. After unpacking the rack, follow local laws and guidelines to recycle the packaging properly.



### Caution

Carefully unpack the rack from the packaging and shipping pallet. Rocking or tilting the rack can cause it to fall over and cause serious injury or death. You should always use professional movers when unpacking and installing this rack.



### Note

After unpacking the rack from the packaging, save the shipping brackets used to secure the rack to the shipping pallet. You can use these shipping brackets to secure the rack permanently to the installation site floor. Do not dispose of these brackets, because you cannot order replacement brackets.



### Caution

Shipping brackets are not for use for bracing or anchoring the rack during seismic events.

## 4.3 Install Oracle Private Cloud Appliance in Its Allocated Space

This section describes each phase of the hardware installation procedure for the Oracle Private Cloud Appliance.

### 4.3.1 Move the System

The following procedure describes how to move the Oracle Private Cloud Appliance:

1. Ensure that the rack doors are closed and secured.
2. Ensure that the leveling and stabilizing feet on the rack are raised and out of the way.
3. Push the system from the back of the rack to the installation site.





**Warning**

Never attempt to move an Oracle Private Cloud Appliance by pushing on the rack side panels. Pushing on the rack side panels can tip over the rack. This action can cause serious personal injury or death, as well as damage to the equipment.

The front casters of the rack are fixed; they do not pivot. When moving your Oracle Private Cloud Appliance to the installation site, you must steer the unit using the rear casters. You can safely maneuver the system by carefully pushing it from behind. See [Figure 4.1](#).

It is preferred that at least three people push and guide the rack: one person in front and two persons in back to help guide the rack and keep people out of the path of the moving rack. When transporting configured racks from one location to another, take care to move them slowly, 0.65 meters per second (2.13 feet per second) or slower.


Carefully examine the transportation path. Avoid obstacles such as doorways or elevator thresholds that can cause abrupt stops or shocks. Go around obstacles by using ramps or lifts to enable smooth transport.



**Caution**

Never tip or rock the Oracle Private Cloud Appliance because the rack can fall over.

**Figure 4.1 Carefully Push the Oracle Private Cloud Appliance From the Back of the Rack**



4. When the rack is at the installation site, verify that no components or connections have become dislodged or disconnected during transport. If necessary, re-attach components and cables properly.

### 4.3.2 Stabilize the Appliance Rack

After moving the Oracle Private Cloud Appliance to the installation site, stabilize the rack to ensure that it does not move or tip over. You can stabilize the rack permanently by extending the rack leveling feet, and optionally using the shipping brackets to secure the rack permanently to the floor.



**Caution**

Shipping brackets are not for use for bracing or anchoring the rack during seismic events.

To secure the rack to the installation floor using the shipping brackets, you must drill the appropriate holes in the floor, re-attach the shipping brackets to the rack, position the rack over the mounting holes, and attach the shipping brackets to the floor firmly with bolts and washers that suit the specific environment. Oracle does not provide mounting bolts and washers for the shipping brackets, because different floors require different bolt types and strengths.

(Optional) If you plan to route data or power distribution unit (PDU) power cords down through the bottom of the rack, you will need to cut a hole in the installation site floor. Cut a rectangular hole below the rear portion of the rack, between the two rear casters and behind the RETMA (Radio Electronics Television Manufacturers Association) rails.




### Caution

Do not create a hole where the rack casters or leveling feet brackets will be placed.

When the rack is in position, the leveling feet must be deployed. The rack contains four leveling feet that can be lowered to share the load with the casters. This increases the footprint of the rack, which improves stability and helps prevent rack movement. The leveling feet must be used even when the rack is permanently secured to the floor. To adjust the leveling feet, do the following:


1. Locate the four leveling feet at the bottom four corners of the rack. See [Figure 4.2](#).

**Figure 4.2 Location of Leveling Feet on Rack**




2. Using a 6mm hex wrench, lower the leveling feet to the floor. See [Figure 4.4](#).

**Figure 4.3 Lowering the Leveling Feet**



3. Lock the four leveling feet using an 18mm open wrench. See [Figure 4.4](#).

**Figure 4.4 Locking the Leveling Feet**



When lowered correctly, the four leveling feet share the load with the casters to increase footprint, improve stability, and help support the full weight of the Oracle Private Cloud Appliance.



**Caution**

When the rack needs to be moved to a different location, including repacking, verify that the leveling feet have been retracted before moving the rack. Otherwise the leveling feet may become bent, or the rack could tip over.

### 4.3.3 Attach a Ground Cable (Optional)

Oracle Private Cloud Appliance power distribution units (PDUs) achieve earth ground through their power cords. Final chassis ground is achieved by way of the ground prong when you connect the power cord to a socket. For additional grounding, attach a chassis earth ground cable to the system. The additional ground point enables electrical current leakage to dissipate more efficiently.



**Caution**

The PDU power input lead cords and the ground cable must reference a common earth ground. If they do not, then a difference in ground potential can be introduced. If you are unsure of your facility's PDU receptacle grounding, then do not install a ground cable until you confirm that there is a proper PDU receptacle grounding. If a difference in ground potential is apparent, then you must take corrective action.




**Note**

A grounding cable is not shipped with the Oracle Private Cloud Appliance.

1. Ensure that the installation site has a properly grounded power source in the data center. The facility PDU must have earth ground.
2. Ensure that all grounding points, such as raised floors and power receptacles, reference the facility ground.
3. Ensure that direct, metal-to-metal contact is made for this installation. During manufacturing, the ground cable attachment area might have been painted or coated.
4. Attach the ground cable to one of the attachment points located at the bottom rear of the system frame. See [Figure 4.5](#).

The attachment point is an adjustable bolt that is inside the rear of the system cabinet on the right side.

**Figure 4.5 Earth Ground Attachment Bolt Location**



## 4.4 What's Next?

After installing the Oracle Private Cloud Appliance rack at your site, continue the installation process by powering on the system and performing the initial configuration. For instructions, see [Chapter 5, \*Powering On Oracle Private Cloud Appliance\*](#).

---

# Chapter 5 Powering On Oracle Private Cloud Appliance

## Table of Contents

5.1 Power On Task Map .....	35
5.2 Before Initial Power On .....	36
5.3 Power On for the First Time .....	36
5.3.1 Review Safety Guidelines .....	36
5.3.2 Inspect the Appliance .....	37
5.3.3 Connect the Power Cords .....	37
5.3.4 Power On the Appliance .....	38
5.3.5 Connect the Appliance to the Network .....	41
5.4 What's Next? .....	45

This chapter explains how to power on the Oracle Private Cloud Appliance for the first time at your site.



### Warning

Before starting the system and applying the initial configuration, read and understand the [Oracle Private Cloud Appliance Release Notes](#). The section [Known Limitations and Workarounds](#) provides information that is critical for correctly executing the procedures in this document. Ignoring the release notes may cause you to configure the system incorrectly. Bringing the system back to normal operation may require a complete factory reset.

## 5.1 Power On Task Map

[Table 5.1](#) describes the steps for powering on the Oracle Private Cloud Appliance at your site.

**Table 5.1 Steps for Powering On Oracle Private Cloud Appliance**

Step	Description	Links
1	Review the safety precautions, guidelines, site checklists, and site requirements.	<a href="#">Section 5.3.1, "Review Safety Guidelines"</a>
2	Perform preliminary checks before connecting the power cords.	<a href="#">Section 5.3.2, "Inspect the Appliance"</a>
3	Verify the hardware by visually inspecting the system.	
4	Connect the appliance public interfaces to the data center network. Also connect a workstation with monitor to the appliance internal management network switch.	<a href="#">Section 5.2, "Before Initial Power On"</a>
5	Supply rack power, and perform the power-on self-test device checks.	<a href="#">Section 5.3.3, "Connect the Power Cords"</a>
6	Switch on the two power distribution unit (PDU) circuit breakers located on the rear of the main PDU (PDU A).	<a href="#">Section 5.3.4, "Power On the Appliance"</a>
7	Power on the management nodes.	

Step	Description	Links
8	Wait 5 to 10 minutes and then connect to the Oracle Private Cloud Appliance Dashboard.	

## 5.2 Before Initial Power On

Before you power on the Oracle Private Cloud Appliance for the first time, ensure that you have made the necessary external network connections. These are different depending on the system network architecture. Refer to the section with instructions for your appliance type:

- [Section 3.1, “Network Connection and IP Address Requirements for Ethernet-based Systems”](#)
- [Section 3.2, “Network Connection and IP Address Requirements for InfiniBand-based Systems”](#)

In addition, you need to connect a workstation with a monitor to the appliance internal management network switch. The switch and port to connect to, are different depending on the system network architecture.

- **Ethernet-based appliance:** connect the workstation to port 48 in the Cisco Nexus 9348GC-FXP Switch
- **InfiniBand-based appliance:** connect the workstation to port 19 in one of the Oracle Switch ES1-24 switches

Statically configure the workstation IP address to [192.168.4.254](#). Start a web browser session in order to initially administer your Oracle Private Cloud Appliance using the pre-defined IP address ([192.168.4.216](#)). For more information, see [Section 5.3.5, “Connect the Appliance to the Network”](#).

## 5.3 Power On for the First Time

This section provides instructions for each phase of the initial power-on sequence.

### 5.3.1 Review Safety Guidelines

Before powering on the system or installing any server equipment into the rack, read the *Important Safety Information for Oracle's Hardware Systems (7063567)* document included with the system.

Before powering on the system, review the following safety precautions. Failure to observe these precautions can result in personal injury, equipment damage, or malfunction.

- Do not block ventilation openings.
- Do not place cables under the equipment or stretch the cables too tightly.
- Do not disconnect power cords from the equipment while its power is on.
- If you cannot reach the connector lock when disconnecting LAN cables, then press the connector lock with a flathead screwdriver to disconnect the cable. You could damage the system board if you force your fingers into the gap rather than using a flathead screwdriver.
- Do not place anything on top of the system or perform any work directly above it.
- Do not let the room temperature rise sharply, especially during winter. Sudden temperature changes can cause condensation to form inside the system. Allow for a sufficient warm-up period prior to system operation.

- Avoid static electricity at the installation location. Static electricity transferred to the system can cause malfunctions. Static electricity is often generated on carpets.
- Confirm that the supply voltage and frequency match the electrical ratings indicated on your Oracle Private Cloud Appliance.
- Do not insert anything into any system opening, unless doing so is part of a documented procedure. The system contains high-voltage parts. If a metal object or other electrically-conductive object enters an opening in the system, then it could cause a short circuit. This could result in personal injury, fire, electric shock, and equipment damage.

## 5.3.2 Inspect the Appliance

The following procedure describes how to visually inspect your Oracle Private Cloud Appliance after it is physically installed at your site, and prior to power being applied to the system.

1. Check for rack damage.
2. Check the rack for loose or missing screws.
3. Check your Oracle Private Cloud Appliance for the ordered configuration. Refer to the Customer Information Sheet (CIS) on the side of the packaging.



### Note

Oracle Private Cloud Appliance is preconfigured by Oracle as a self-contained system. You should not move any equipment or add any unsupported hardware to the system.

4. Check that all cable connections are secure and firmly in place as follows:
  - a. Check the power cables. Ensure that the correct connectors have been supplied for the data center facility power source.
  - b. Check the network data cables.
5. Check the site location tile arrangement for cable access and airflow.
6. Check the data center airflow that leads in to the front of the system.

For more information, see [Section 2.5, “Ventilation and Cooling Requirements”](#).

## 5.3.3 Connect the Power Cords


The following procedure describes how to connect power cords to your Oracle Private Cloud Appliance.

1. Open the rear cabinet door.
2. Ensure that the correct power connectors have been supplied.
3. Unfasten the power cord cable ties.


The ties are for shipping only and are no longer needed.

4. Route the power cords to the facility receptacles either above the rack or below the flooring. See [Figure 5.1](#) and [Figure 5.2](#).
5. Secure the power cords in bundles. See [Figure 5.1](#).

**Figure 5.1 Power Cord Routing From the Bottom of the Rack**



**Figure 5.2 Power Cord Routing From the Top of the Rack**



6. Plug the power distribution unit (PDU) power cord connectors into the facility receptacles. Ensure the breaker switches are in the OFF position before connecting the power cables.

### 5.3.4 Power On the Appliance

The following procedure describes how to power on your Oracle Private Cloud Appliance.



**Note**

You can connect to your Oracle Private Cloud Appliance using a network connection to monitor the system power-on procedure. For instructions, see [Section 5.3.5, “Connect the Appliance to the Network”](#).

1. Ensure that the three main power cords are connected to a reliable power source.

For more information, see [Section 2.3, “Electrical Power Requirements”](#).



2. Switch on the power distribution unit (PDU) circuit breakers located on the rear of PDU A and B inside the Oracle Private Cloud Appliance.


The circuit breakers are on the rear of the system cabinet as shown in [Figure 5.3](#). Press the ON (I) side of the toggle switch.



**Note**

Ensure that the PDU switches have had power applied for approximately two minutes to complete power-on configuration before starting the management nodes.

**Figure 5.3 PDU Switch Locations**



3. Make sure that the circuit breakers located on the rear left and right side power supplies of the Oracle ZFS Storage Appliance ZS7-2 are in the ON (I) position.

**Figure 5.4 Oracle ZFS Storage Appliance ZS7-2 Circuit Breakers**




4. Press the Power button located on the front of the first (bottom) management node. See [Figure 5.5](#).

The first management node is located in rack unit 5 (U5). The second management node is located in rack unit 6 (U6). It is switched on and configured automatically during the initial power-on sequence.

A management nodes take approximately five to ten minutes to power on completely. Once complete, the Power/OK LED illuminates and remains a steady green.

**Figure 5.5 Oracle Private Cloud Appliance Management Node**



**Table 5.2 Figure Legend**

Item	Description
1	Power button
2	Power/OK LED

5. After powering on the first management node, wait while the second management node and all compute nodes are powered on and discovered.



**Note**

After power is applied, the LEDs on all compute nodes and storage server heads will start to blink after approximately two minutes. From the rear of the rack, you can see the green LEDs on the power supply units (PSUs) on the compute nodes turn on instantly after power is applied. In addition, from the rear of the rack, you can see the display on the power distribution units (PDUs) illuminate once power is available.

The management nodes will verify all components within the system. The management nodes ensure that the correct networking switches and storage devices are installed in the system, and search for compute nodes to add to the compute fabric. Once compute nodes are located, the management software automatically installs the proper drivers and the Oracle VM Server version on each compute node.

Depending on your system configuration, the Oracle Private Cloud Appliance power-on process should take approximately five to ten minutes before the system becomes available for initialization, and up to several hours for the entire configuration to be ready. When initialization is completed, use the Oracle Private Cloud Appliance Dashboard web interface to configure initial system settings (for example: network properties and passwords). For instructions on connecting to the Oracle Private

Cloud Appliance and accessing the Dashboard UI, see [Section 5.3.5, “Connect the Appliance to the Network”](#).

## 5.3.5 Connect the Appliance to the Network

For initial access to Oracle Private Cloud Appliance, you must connect a workstation with a monitor to the reserved Ethernet port in the appliance internal management network switch. The Oracle Private Cloud Appliance Dashboard can be reached by means of a web browser.



### Note

You can connect the workstation to the Ethernet port at any time during the system power-on process. However, you should wait approximately two to three minutes after starting the power-on process to allow the IP address to become active. If you connect too soon to the port, you will see an HTTP error code 404 or similar code. If this occurs, wait a few minutes and try to connect again.

You can connect your Oracle Private Cloud Appliance to the network as follows:

1. Ensure that the appliance public interfaces are connected to your data center (public) network.

For details, refer to the section with instructions for your appliance type:

- [Section 3.1.1, “Network Connection Requirements for Ethernet-based Systems”](#)
- [Section 3.2.1, “Network Connection Requirements for InfiniBand-based Systems”](#)

2. Connect a workstation with a web browser directly to the management network using an Ethernet cable connected to the appropriate port:

- **Ethernet-based appliance:** port 48 in the Cisco Nexus 9348GC-FXP Switch
- **InfiniBand-based appliance:** port 19 in one of the Oracle Switch ES1-24 switches

3. Configure the wired network connection of the workstation to use the static IP address [192.168.4.254](#).

4. Using the web browser on the workstation, connect to the Oracle Private Cloud Appliance Dashboard on the master management node at <http://192.168.4.216>.

[192.168.4.216](#) is the predefined virtual IP address of the management node cluster for configuring Oracle Private Cloud Appliance. If there is another device on your network that is using this IP address, you should change the conflicting device's IP address prior to connecting to the Oracle Private Cloud Appliance.

5. Log into the Dashboard using the user name [admin](#) and default password [welcome1](#).



### Note

For security reasons, you must set a new password at your earliest convenience.

When you have logged in to the Dashboard successfully, the home page is displayed. The central part of the page contains Quick Launch buttons that provide direct access to the key functional areas. Click to open the Hardware View.

The **Hardware View** indicates the status of the appliance as it is being powered on. All the networking and storage should be green. As the compute nodes are discovered, they will change from yellow to

green. Wait until all compute nodes have been discovered and provisioned before making configuration changes in the Dashboard interface.

6. Using the Dashboard, configure the system environment parameters, including: domain name, NTP and DNS servers.



### Warning

Before reconfiguring the system network settings, make sure that no provisioning or upgrade processes are running. For details, refer to the section [“Do Not Reconfigure Network During Compute Node Provisioning or Upgrade”](#) in the Oracle Private Cloud Appliance Release Notes.

7. Go to the Network Environment window and enter the configuration details. Insert the reserved management node IP addresses and environment parameters of your data center network.
  - a. In the Management Nodes tab, enter the IP addresses and host names for both management nodes.



### Note

During the initial configuration of the Oracle Private Cloud Appliance, you change the predefined configuration IP address to the reserved virtual IP

address in your data center network. For details, see [Section 3.3, “General Network Configuration”](#).

**Figure 5.6 Management Nodes Tab**

The screenshot shows the Oracle PCA Admin Console interface. The top navigation bar includes a menu icon, the text 'ORACLE PCA Admin Console', and a button labeled 'Open OVM Manager...'. The main content area is titled 'Network Environment' and contains a sub-tabbed interface with 'Management Nodes', 'Data Center Network', and 'DNS'. The 'Management Nodes' tab is active and displays three configuration sections: 'Management Node 1' with IP Address '10.19.1.101' and Host Name 'manager1'; 'Management Node 2' with IP Address '10.19.1.102' and Host Name 'manager2'; and 'Management Virtual IP Address' with IP Address '10.19.1.100'. At the bottom, there is a password field labeled '\* PCA Admin Password(Required for changes)', a 'Reset' button, and an 'Apply Changes' button. A note above the tabs states: 'Changes to the Network Environment configuration data requires entering the PCA Admin Password before selecting the 'Apply Changes' button.'

- b. In the Data Center Network tab, enter the domain name, netmask, default gateway and NTP server of the network to which the appliance must be connected. If this network is a VLAN, set the slider to *active* and enter the VLAN ID in the field next to it.

Figure 5.7 Data Center Network Tab

The screenshot shows the Oracle PCA Admin Console interface. At the top, there is a blue header with the Oracle logo and 'PCA Admin Console' text. A 'MENU' icon is on the left, and an 'Open OVM Manager...' button with a user icon is on the right. Below the header, the main content area is titled 'Network Environment'. A note states: 'Changes to the Network Environment configuration data requires entering the PCA Admin Password before selecting the 'Apply Changes' button.' There are three tabs: 'Management Nodes', 'Data Center Network' (which is selected), and 'DNS'. The 'Data Center Network' tab contains the following configuration options:

- Management Network VLAN:** A toggle switch is turned on, and a dropdown menu shows the value '20'.
- \* Domain Name:** A text input field containing 'example.com'.
- \* Netmask:** A text input field containing '255.255.0.0'.
- \* Default Gateway:** A text input field containing '10.19.1.1'.
- \* NTP:** A text input field containing '10.19.1.1'.

At the bottom of the form, there is a field for the 'PCA Admin Password (Required for changes)', a 'Reset' button, and an 'Apply Changes' button.

- c. In the DNS tab, enter at least one and up to three name servers for the network to which the appliance must be connected.

Figure 5.8 DNS Tab

The screenshot displays the Oracle PCA Admin Console interface for configuring DNS. The top navigation bar includes a menu icon, the text 'ORACLE PCA Admin Console', and a link to 'Open OVM Manager...'. The main heading is 'Network Environment', with a sub-note: 'Changes to the Network Environment configuration data requires entering the PCA Admin Password before selecting the 'Apply Changes' button.' Below this, there are three tabs: 'Management Nodes', 'Data Center Network', and 'DNS'. The 'DNS' tab is active, showing three input fields for DNS servers. The first field is labeled '\* DNS Server 1' and contains '217.6.34.47'. The second field is labeled 'DNS Server 2' and contains '144.20.190.70'. The third field is labeled 'DNS Server 3' and is empty. At the bottom, there is a password field labeled '\* PCA Admin Password(Required for changes)', a 'Reset' button, and an 'Apply Changes' button.

- d. Enter the Administrator password and click Apply Changes to confirm and apply the configuration parameters you provided.
8. Reconnect to the Oracle Private Cloud Appliance Dashboard at the new virtual IP address of the management node cluster: [https://<virtual\\_ip>:7002/dashboard](https://<virtual_ip>:7002/dashboard).

For details about the software configuration process, and for advanced configuration and update options, refer to the *Oracle Private Cloud Appliance Administrator's Guide*.

## 5.4 What's Next?

After powering on your Oracle Private Cloud Appliance and performing the initial system configuration, the environment is ready for virtual machine deployment. Refer to the *Oracle Private Cloud Appliance Administrator's Guide* for information about working with an operational environment:

- See [Monitoring and Managing Oracle Private Cloud Appliance](#) for instructions to work with the Oracle Private Cloud Appliance Dashboard.
- See [Managing the Oracle VM Virtual Infrastructure](#) for instructions to work with virtual machines within Oracle VM.
- See [Updating Oracle Private Cloud Appliance](#) for instructions to update the Oracle Private Cloud Appliance controller software.
- See [Servicing Oracle Private Cloud Appliance Components](#) for instructions to service customer-replaceable components in the Oracle Private Cloud Appliance.



**Note**

Optionally, you can extend Oracle Private Cloud Appliance by adding more compute nodes to the system, by installing extra memory in the compute nodes, or by connecting the appliance to additional external storage. For more information and instructions, refer to the following chapters:

- [Chapter 6, \*Extending Oracle Private Cloud Appliance - Additional Compute Nodes\*](#)
- [Chapter 7, \*Extending Oracle Private Cloud Appliance - Server Memory Kits\*](#)
- [Chapter 8, \*Extending Oracle Private Cloud Appliance - Additional Storage\*](#)



---

# Chapter 6 Extending Oracle Private Cloud Appliance - Additional Compute Nodes

## Table of Contents

6.1 Extending Compute Capacity Task Map .....	47
6.2 Review Safety Guidelines .....	48
6.3 Rack-mount Safety Precautions .....	48
6.4 Prepare to Add a Server .....	49
6.4.1 ESD Precautions .....	49
6.4.2 Unpack the Server Upgrade Kit .....	50
6.4.3 Prepare the Appliance Rack .....	50
6.5 Install the Server Into the Rack .....	52
6.5.1 Install Mounting Brackets .....	52
6.5.2 Attach Slide-Rail Assemblies .....	53
6.5.3 Install the Server Into the Slide-Rail Assemblies .....	55
6.5.4 Install the Cable Management Arm .....	57
6.5.5 Cable the Server .....	62
6.5.6 Replace the Rack Doors .....	65

Extension of Oracle Private Cloud Appliance involves increasing the number of compute nodes within a rack to maximize the number of virtual machines that are capable of running on the appliance and to increase the processing power and memory available to existing virtual machines. The process of extending the Oracle Private Cloud Appliance in this way requires that you install additional servers into the rack. This operation does not require that the system is powered down, but physical changes are made to the system. Therefore you should review the task map provided in [Section 6.1, “Extending Compute Capacity Task Map”](#) before continuing. Each step to perform the operation is described in detail in this chapter.

This chapter explains how to extend your Oracle Private Cloud Appliance by adding compute nodes to the system. After hardware installation, when compute node provisioning has completed, the servers are automatically added to the default tenant group *Rack1\_ServerPool*. Depending on the tenant group configuration in the environment, you should add each newly provisioned server to the appropriate tenant group.

## 6.1 Extending Compute Capacity Task Map

[Table 6.1](#) describes the steps for extending the Oracle Private Cloud Appliance at your site by installing additional compute nodes within the rack.

**Table 6.1 Steps for Extending Oracle Private Cloud Appliance By Adding Compute Nodes**

Step	Description	Links
1	Review safety guidelines and site checklists.	<a href="#">Section 6.2, “Review Safety Guidelines”</a>
2	Review safety precautions for rack-mounting a server.	<a href="#">Section 6.3, “Rack-mount Safety Precautions”</a>
3	Prepare to add the server to the rack, including: <ul style="list-style-type: none"><li>• Unpacking the upgrade kit</li><li>• Removing the rear doors from the rack</li></ul>	<a href="#">Section 6.4, “Prepare to Add a Server”</a>

Step	Description	Links
4	Install the server into the rack, including: <ul style="list-style-type: none"> <li>• Installing mounting brackets to the server</li> <li>• Attaching slide-rail assemblies to the rack</li> <li>• Installing the server into the slide-rail assemblies</li> <li>• Installing the cable management arm</li> <li>• Cabling the server</li> <li>• Replacing the doors on the rack</li> </ul>	<a href="#">Section 6.5, "Install the Server Into the Rack"</a>

## 6.2 Review Safety Guidelines

Before expanding Oracle Private Cloud Appliance, read the following important safety information:

- *Important Safety Information for Oracle's Hardware Systems (7063567)*
- [Section 2.3, "Electrical Power Requirements"](#)
- [Section D.5, "Facility Power Checklist"](#)

## 6.3 Rack-mount Safety Precautions

This section describes rack-mount safety precautions you must follow when installing a compute node into your Oracle Private Cloud Appliance.



### Caution

**Leveling feet position:** When unpacking at the installation site, or when repackaging and moving the rack to a new location, verify that the leveling feet are up before moving the rack.



### Caution

**Stabilize rack:** Deploy the anti-tilt bar on the equipment rack before starting an installation.



### Caution

**Shipping brackets:** The rack shipping brackets are not for use for bracing or anchoring the rack during seismic events.



### Caution

**Equipment loading:** Always load equipment into a rack from the bottom up so that the rack does not become top-heavy and tip over. Deploy your rack anti-tilt bar to prevent the rack from tipping during equipment installation.



### Caution

**Lifting equipment:** Oracle Server X8-2 weighs approximately 45 lbs (20.41 kg). Two people are needed to mount the server into the rack enclosure.



**Caution**

**Communicate instructions:** When performing a two-person procedure, communicate your intentions clearly to the other person before, during, and after each step to minimize confusion.



**Caution**

**Elevated operating ambient temperature:** If you install the server in a closed or multi-unit rack assembly, the operating ambient temperature of the rack environment might be higher than the room ambient temperature. Install the equipment in an environment compatible with the maximum ambient temperature (T<sub>ma</sub>) specified for the server. For information, refer to [Section 2.5, “Ventilation and Cooling Requirements”](#).



**Caution**

**Reduced airflow:** Install the equipment in a rack so that it does not compromise the amount of airflow required for safe operation of the equipment.



**Caution**

**Mechanical loading:** Mount the equipment in the rack so that it does not cause a hazardous condition due to uneven mechanical loading.



**Caution**

**Circuit overloading:** Consider the connection of the equipment to the supply circuit and the effect that overloading the circuits might have on over-current protection and supply wiring. Also consider the equipment nameplate power ratings used when you address this concern.



**Caution**

**Reliable earthing:** Maintain reliable earthing of rackmounted equipment. Pay attention to supply connections other than direct connections to the branch circuit (for example, use of power strips).



**Caution**

**Mounted equipment:** Do not use slide-rail-mounted equipment as a shelf or a workspace.

## 6.4 Prepare to Add a Server

This section describes the preparation steps to perform before installing an additional compute node as part of your Oracle Private Cloud Appliance.



**Caution**

When adding expansion nodes, always populate the compute rack units starting from the bottom most open slot, and work your way up. If you do not populate the rack units in the right order, the controller software may no longer report the correct status of the rack components.

### 6.4.1 ESD Precautions

Electronic equipment is susceptible to damage by static electricity. To prevent electrostatic discharge (ESD) when you install or service the server:

- Use a grounded antistatic wrist strap, foot strap, or equivalent safety equipment.
- Place components on an antistatic surface, such as an antistatic discharge mat or an antistatic bag.
- Wear an antistatic grounding wrist strap connected to a metal surface on the chassis when you work on system components.



### Caution

**Equipment damage:** Electrostatic damage can permanently disable the system or require repair by authorized service technicians.

Before installing the server, read the safety information in the [Oracle Server X8-2 Safety and Compliance Guide](#), as well as the *Important Safety Information for Oracle's Hardware Systems*.

## 6.4.2 Unpack the Server Upgrade Kit

Perform the following procedure to unpack the server upgrade kit.

1. Inspect the shipping cartons for evidence of physical damage. If a shipping carton appears damaged, request that the carrier agent be present when the carton is opened. Keep all contents and packing material for the agent inspection.
2. Remove the packaging from the Oracle Private Cloud Appliance Expansion Node upgrade kit.
3. Verify that all items are included. The upgrade kit should contain the following:
  - Oracle Server
  - Cable management arm assembly with installation instructions
  - Rackmount kit including rack rails, mounting brackets, screws and Rackmounting Template
  - Getting Started Guide
  - Legal and safety documents



### Note

Cables are not included in the server upgrade kit. Oracle Private Cloud Appliance comes with pre-installed and labeled cables for all rack units. These should not be unplugged or removed from the rack.

## 6.4.3 Prepare the Appliance Rack

Perform the following procedure to prepare the rack.



### Caution

To reduce the risk of personal injury, stabilize the rack cabinet and extend all anti-tilt devices before performing this procedure.

1. Remove the Oracle Private Cloud Appliance front and rear rack doors, as follows:

- a. Unlock and open the front and rear rack doors.
- b. Detach the top and bottom grounding straps by separating the banana jacks (1).
- c. Unlatch the bottom hinge, then the top hinge.
  - Pinch the two studs on the hinge towards the center of the hinge (2).
  - Rotate the studs away from the door (3) so that they stay in place and the hinge remains unlatched.
- d. Lift the door away from the cabinet (4).
- e. Repeat these steps for the two rear doors.


These steps are illustrated in [Figure 6.1](#).



**Caution**

The doors are heavy. Have one person hold the door in place while the other operates the hinge latching mechanism. Execute the steps carefully and in the correct order, to avoid damage to the hinges.

**Figure 6.1 Removing the Rack Doors**



**Table 6.2 Figure Legend**


Item	Description
1	Detaching the grounding cables.
2	Pinching the hinge studs.
3	Rotating the studs away from the door.
4	Lifting the door away from the cabinet.

2. Remove the filler panel(s) where the upgrade server will be installed.

The filler panels have snap-in attachments. No screwdriver is required to remove them.

3. Remove the cable tray that holds the pre-installed cables in place at the rear side of the rack. See [Figure 6.2](#)
  - a. Detach the pre-installed cables from the cable tray on the rack unit for the new expansion node.
  - b. Use a No. 2 Phillips screwdriver to remove the M6 screws from the cable tray.
  - c. Remove the cable tray from the rack. Save cable tray and screws for future use.

**Figure 6.2 Removing the Cable Tray**



## 6.5 Install the Server Into the Rack


This section describes the steps to install an additional compute node as part of your Oracle Private Cloud Appliance.

### 6.5.1 Install Mounting Brackets

Perform the following procedure to install mounting brackets onto the sides of the server.

1. Position a mounting bracket against the chassis so that the slide-rail lock is at the server front, and the five keyhole openings on the mounting bracket are aligned with the five locating pins on the side of the chassis. See [Figure 6.3](#).

**Figure 6.3 Aligning the Mounting Bracket With the Server Chassis**



**Table 6.3 Figure Legend**

Item	Description
1	Chassis front
2	Slide-rail lock
3	Mounting bracket
4	Mounting bracket clip

2. With the heads of the five chassis locating pins protruding through the five keyhole openings in the mounting bracket, pull the mounting bracket toward the front of the chassis until the mounting bracket clip locks into place with an audible click.
3. Verify that the rear locating pin has engaged the mounting bracket clip.
4. Repeat these steps to install the remaining mounting bracket on the other side of the server.

## 6.5.2 Attach Slide-Rail Assemblies

Perform the following procedure to attach slide-rail assemblies to the rack.

1. Orient the slide-rail assembly so that the ball-bearing track is forward and locked in place.

Figure 6.4 Orienting the Slide-Rail With the Ball-Bearing Track





Table 6.4 Figure Legend

Item	Description
1	Slide-rail
2	Ball-bearing track
3	Locking mechanism



2. Starting with either the left or right side of the rack, align the rear of the slide-rail assembly against the inside of the rear rack rail, and push until the assembly locks into place with an audible click. See [Figure 6.5](#).

**Figure 6.5 Aligning the Slide-Rail Assembly With the Rack**



3. Align the front of the slide-rail assembly against the outside of the front rack rail, and push until the assembly locks into place with an audible click. See [Figure 6.5](#).
4. Repeat these steps to attach the slide-rail assembly to the other side of the rack.

### 6.5.3 Install the Server Into the Slide-Rail Assemblies

Perform the following procedure to install the server, with mounting brackets, into the slide-rail assemblies that are mounted to the rack.



#### Caution

This procedure requires a minimum of two people because of the weight of the server. Attempting this procedure alone could result in equipment damage or personal injury.




**Caution**

Always load equipment into a rack from the bottom up so that the rack will not become top-heavy and tip over. Extend your rack's anti-tilt bar to prevent the rack from tipping during equipment installation.

1. Push the slide-rails as far as possible into the slide-rail assemblies in the rack.
2. Position the server so that the rear ends of the mounting brackets are aligned with the slide-rail assemblies that are mounted in the rack. See [Figure 6.6](#).
3. Insert the mounting brackets into the slide-rails, and then push the server into the rack until the mounting brackets encounter the slide-rail stops (approximately 30 cm, or 12 inches).

**Figure 6.6 Inserting the Server With Mounting Brackets Into the Slide-Rails**




**Table 6.5 Figure Legend**

Item	Description
1	Inserting mounting bracket into slide-rail
2	Slide-rail release button
3	Slide-rail lock

4. Simultaneously push and hold the green slide-rail release buttons on each mounting bracket while you push the server into the rack. Continue pushing the server into the rack until the slide-rail locks (on the front of the mounting brackets) engage the slide-rail assemblies. See [Figure 6.7](#).

You will hear an audible click.

**Figure 6.7 Sliding the Server Into the Rack**



## 6.5.4 Install the Cable Management Arm

Perform the following procedure to install the cable management arm (CMA) to rear of the server.

1. Unpack the CMA.
2. To make it easier to install the CMA, extend the server approximately 13 cm (5 inches) out of the front of the rack.
3. Take the CMA to the back of the equipment rack, and ensure that you have adequate room to work at the back of the server.



**Note**

References to “left” or “right” in this procedure assume that you are facing the back of the equipment rack.



**Note**


Throughout this installation procedure, support the CMA and do not allow it to hang under its own weight until it is secured at all four attachment points.

4. Insert the CMA's connector A into the front slot on the left slide-rail until it locks into place with an audible click. See [Figure 6.8](#), frames 1 and 2.

The connector A tab (see call-out 1) goes into the slide-rail's front slot (call-out 2).

Gently tug on the left side of the front slide bar to verify that connector A is properly seated.

**Figure 6.8 Installing Connector A Into the Left Slide-Rail**



**Table 6.6 Figure Legend**


Item	Description
1	Connector A tab
2	Left slide-rail front slot

5. Insert the CMA's connector B into the front slot on the right slide-rail until it locks into place with an audible click. See [Figure 6.9](#), frames 1 and 2.

The connector B tab (call-out 1) goes into the slide-rail front slot (call-out 2).

Gently tug on the right side of the front slide bar to verify that connector B is properly seated.

**Figure 6.9 Installing Connector B Into the Right Slide-Rail**



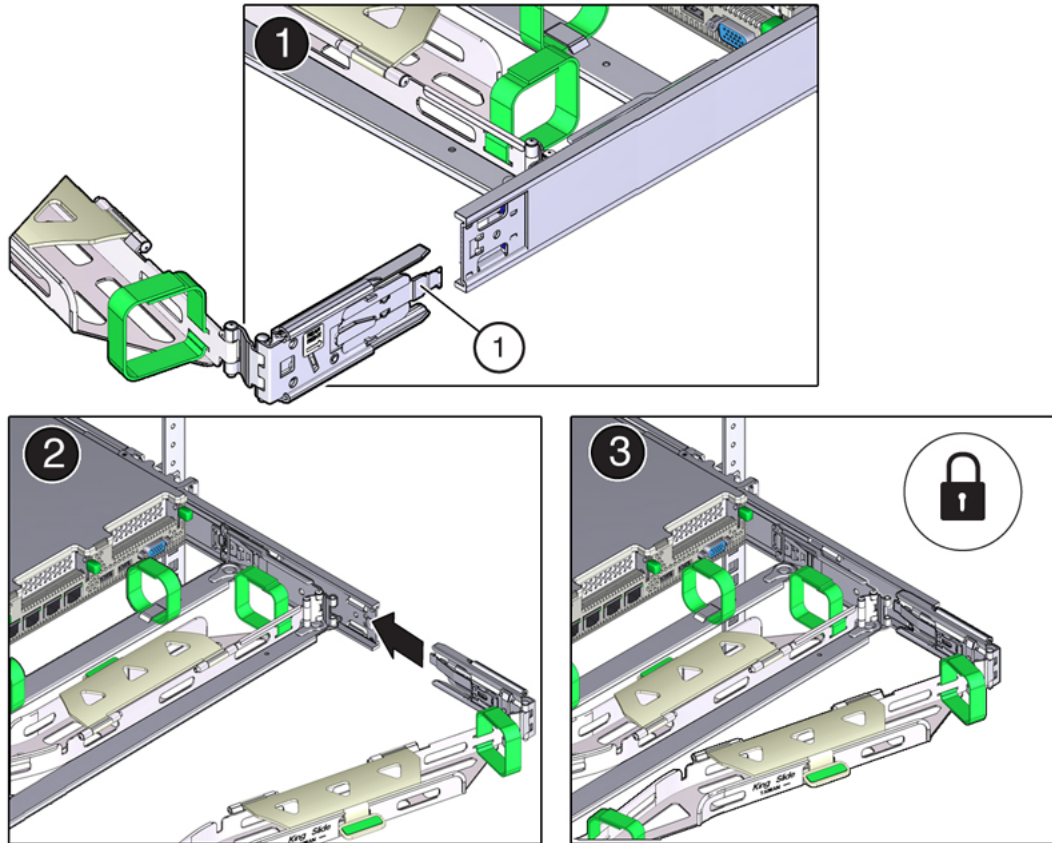
**Table 6.7 Figure Legend**

Item	Description
1	Connector B tab

Item	Description
2	Right slide-rail front slot

6. To install the CMA's connector C into the right slide-rail, perform the following steps:
  - a. Align connector C with the slide-rail so that the locking spring (call-out 1) is positioned inside (server side) of the right slide-rail. See [Figure 6.10](#), frame 1.

**Figure 6.10 Installing Connector C Into the Right Slide-Rail**



**Table 6.8 Figure Legend**

Item	Description
1	Connector C locking spring

- b. Insert connector C into the right side-rail until it locks into place with an audible click (frames 2 and 3).
- c. Gently tug on the right side of the CMA's rear slide bar to verify that connector C is properly seated.


7. To prepare the CMA's connector D for installation, remove the tape that secures the slide-rail latching bracket to connector D and ensure that the latching bracket is properly aligned with connector D. See [Figure 6.11](#), frames 1 and 2.



**Note**

The CMA is shipped with the slide-rail latching bracket taped to connector D. You must remove the tape before you install this connector.

**Figure 6.11 Aligning the CMA Slide-Rail Latching Bracket With Connector D**





8. While holding the slide-rail latching bracket in place, insert connector D and its associated slide-rail latching bracket into the left slide-rail until connector D locks into place with an audible click. See [Figure 6.12](#), frames 1 and 2.




**Note**

When inserting connector D into the slide-rail, the preferred and easier method is to install connector D and the latching bracket as one assembly into the slide-rail.

Gently tug on the left side of the CMA's rear slide bar to verify that connector D is properly seated.

**Figure 6.12 Installing Connector D Into the Left Slide-Rail**



**Note**

The slide-rail latching bracket has a green release tab. This tab is used to release and remove the latching bracket so that you can remove connector D.

9. Gently tug on the four CMA connection points to ensure that the CMA connectors are fully seated before you allow the CMA to hang by its own weight.

10. To verify that the slide-rails and the CMA are operating properly before routing cables through the CMA, perform the following steps:

- a. Extend the server from the front of the rack until the CMA is fully extended.



**Caution**

To reduce the risk of personal injury, stabilize the rack cabinet and extend all anti-tilt devices before extending the server from the rack.


- b. Simultaneously pull and hold the two green release tabs (one on each side of the server) toward the front of the server (see the following figure) while you push the server into the rack. As you push the server into the rack, verify that the CMA retracts without binding. See [Figure 6.13](#).



**Note**

To pull the green release tabs, place your finger in the center of each tab, not on the end, and apply pressure as you pull the tab toward the front of the server.

**Figure 6.13 Location of the Slide-Rail Release Tabs**



- c. Continue pushing the server into the rack until the slide-rail locks (on the front of the server) engage the slide-rail assemblies.

You will hear an audible click when the server is in the normal rack position.

11. Connect the cables to the server, as required.

Instructions for connecting server cables are provided in [Section 6.5.5, "Cable the Server"](#).

## 6.5.5 Cable the Server

After the server is installed in the rack and proper cable management arm (CMA) operation is verified, perform the procedures in this section to cable the server.





**Note**

If you are installing more than one server into the rack, perform the procedures in the previous sections for each server before performing the cabling steps in this section.

You should observe the following guidelines when attaching cables to a newly rack-mounted server:

- Do not remove existing cable connections in the rack.
- Attach and route cables in the CMA and rear panel one server at a time. Do not slide out more than one server at a time.
- Start from the bottom of the rack, and work upward. Route cables through the CMA with the dongle on the top and power cables on the bottom.




**Note**

The Oracle Private Cloud Appliance rack is entirely pre-cabled for all components, including future compute nodes. Cable bundles are cable-tied to the correct rack unit for each future server. You should not modify any of the cabling on the already installed components. For detailed cabling information, see [Appendix B, Cabling Reference](#).

1. Connect the pre-installed power and data cables of the appropriate rack unit to the back of the expansion node. The cables are labeled for your convenience, and should be connected according to [Figure 6.14](#).

**Figure 6.14 Server Cabling Requirements**



**Table 6.9 Figure Legend**


Item	Description
A	Power supplies
B	<a href="#">NET0</a> Ethernet port
C	Optional Fibre Channel HBA



**Caution**

Physical Fibre Channel expansion cards are available as an option for Oracle Server X8-2 or newer compute nodes. The appliance rack does not contain pre-installed Fibre Channel cables.

Item	Description
D	Internal appliance data network connectors


 **Caution**

Depending on the physical network architecture of your appliance, the connectors are either InfiniBand or Ethernet ports. A specific server model is available for each architecture type.

2. Open the CMA cable covers, route server cables through the CMA's cable troughs, close the cable covers, and secure the cables with the six Velcro straps. See [Figure 6.15](#).
  - a. First through the front-most cable trough.
  - b. Then through the small cable trough.
  - c. Then through the rear-most cable trough.

**Note**

When securing the cables with the Velcro straps located on the front slide bar, ensure that the Velcro straps do not wrap around the bottom of the slide bar; otherwise, expansion and contraction of the slide bar might be hindered when the server is extended from the rack and returned to the rack.

**Figure 6.15 CMA With Cables Installed**

3. Ensure that the secured cables do not extend above the top or below the bottom of the server to which they are attached; otherwise, the cables might snag on other equipment installed in the rack when the server is extended from the rack or returned to the rack.



**Note**

If necessary, bundle the cables with additional Velcro straps to ensure that they stay clear of other equipment. If you need to install additional Velcro straps, wrap the straps around the cables only, not around any of the CMA components; otherwise, expansion and contraction of the CMA slide bars might be hindered when the server is extended from the rack and returned to the rack.

4. Verify the operation of the CMA with all cables installed.



**Note**

Oracle recommends that two people execute this step. One person to move the server in and out of the rack, and another person to observe the cable and CMA operation.

- a. Slowly pull the server out of the rack until the slide rails reach their stop.
  - b. Inspect the attached cables for any binding or kinks.
  - c. Verify that the CMA extends fully from the slide rails.
5. Push the server back into the rack.
    - a. Simultaneously pull and hold the two green release tabs (one on each side of the server) toward the front of the server while you push the server into the rack. See [Figure 6.13](#).



**Note**

To pull the green release tabs, place your finger in the center of each tab, not on the end, and apply pressure as you pull the tab toward the front of the server.

- b. Verify that the CMA retracts without binding.
6. Repeat the previous steps for each additional server that you add to the rack.

The Oracle Private Cloud Appliance controller software detects that one or more new compute nodes are present in the rack. It automatically powers on the new hardware and executes the provisioning procedure. The expansion compute nodes are incorporated seamlessly into the active system. This requires no further interventions by an administrator.

## 6.5.6 Replace the Rack Doors

Perform the following procedure to replace the rack front and rear doors:

1. Lift the door into place, and align the pins on the door with the sleeve on the cabinet (1).
2. On the top hinge, rotate the two studs on the hinge away from the center of the hinge (2). The pins snap into place (3). Do the same with the bottom hinge.
3. Repeat these steps to reinstall the other rack doors.
4. Reconnect all grounding straps on all doors (4).

5. Close the doors. Optionally, lock the doors.


These steps are illustrated in [Figure 6.1](#).



**Caution**

The doors are heavy. Have one person hold the door in place while the other operates the hinge latching mechanism. Execute the steps carefully and in the correct order, to avoid damage to the hinges.

**Figure 6.16 Installing the Rack Doors**



**Table 6.10 Figure Legend**

Item	Description
1	Lifting the door into place.
2	Rotating the hinge studs away from the center.
3	Pins on the latching mechanism snap into place.
4	Attaching the grounding cables.

---

# Chapter 7 Extending Oracle Private Cloud Appliance - Server Memory Kits

## Table of Contents

7.1 Extending Server Memory Task Map .....	67
7.2 Selecting a Server Memory Kit .....	67
7.3 Installing the Memory Upgrade Kit .....	68

Extending the Oracle Private Cloud Appliance by installing additional physical RAM can enhance the product by providing more resources for virtual machines performing memory-intensive tasks beyond the capacity already available in the default compute node configuration. This process requires the server to be powered down, but virtual machines can be migrated to other compute nodes temporarily to avoid unnecessary downtime.

There are different memory kits available for the different generations of server hardware that may be installed in an Oracle Private Cloud Appliance environment. Be sure to select the RAM that matches your server model.

## 7.1 Extending Server Memory Task Map

[Table 7.1](#) describes the steps for extending the Oracle Private Cloud Appliance at your site by installing additional RAM to upgrade the compute nodes.

**Table 7.1 Steps for Extending Oracle Private Cloud Appliance By Installing Additional Memory**

Step	Description	Links
1	Review safety guidelines and site checklists.	<a href="#">Section 6.2, "Review Safety Guidelines"</a>
2	Identify the servers to upgrade. Depending on the base memory configuration and the desired amount of RAM, select the memory kit to install.	<a href="#">Section 7.2, "Selecting a Server Memory Kit"</a>
3	Install the memory upgrade kit.  This involves migrating running virtual machines to avoid service interruptions, and powering down the compute node.	<a href="#">Section 7.3, "Installing the Memory Upgrade Kit"</a>
4	When the upgraded compute nodes are back online, verify in Oracle VM Manager that the amount of RAM is correct.  When the compute node reboots after a memory upgrade, <code>dom0</code> memory settings are adjusted automatically to the new total amount of RAM. The settings are persistent across future (re-)provisioning and controller software updates.	No manual reconfiguration required.

---

## 7.2 Selecting a Server Memory Kit

Any Oracle Private Cloud Appliance may contain more than one generation of server hardware. This depends on when the system was purchased and installed, and whether it has since been extended with additional compute nodes. These server hardware generations may be part of your installation:

**Table 7.2 Oracle Private Cloud Appliance Server Hardware Generations**

Server Hardware	Memory Configuration
Oracle Server X8-2	<b>Default configuration</b>
Oracle Server X7-2	<p>orderable in three memory configurations:</p> <ul style="list-style-type: none"> <li>• 384GB – 12x 32GB DDR4-2666 DIMMs</li> <li>• 768GB – 12x 64GB DDR4-2666 DIMMs</li> <li>• 1.5TB – 24x 64GB DDR4-2666 DIMMs</li> </ul> <p><b>Expansion options</b></p> <p>The default memory configuration can be extended up to the maximum of 1.5TB. Consult your Oracle representative for the latest information about available memory kits.</p>
Oracle Server X6-2	<p><b>Default configuration</b></p> <p>256GB – 8x 32GB DDR4-2400 DIMMs, 16 free slots</p> <p><b>Expansion options</b></p> <p>add 8x or 16x 32GB DDR4-2400 DIMMs in the free slots, for a total of 512GB or 768GB</p>
Oracle Server X5-2	<p><b>Default configuration</b></p> <p>256GB – 8x 32GB DDR4-2133 DIMMs, 16 free slots</p> <p><b>Expansion options</b></p> <p>add 8x or 16x 32GB DDR4-2133 DIMMs in the free slots, for a total of 512GB or 768GB</p>
Sun Server X4-2	<b>Default configuration</b>
Sun Server X3-2	<p>256GB – 16x 16GB DDR3-1600 DIMMs, no free slots</p> <p><b>Expansion options</b></p> <p>replace with 16x 32GB DDR3-1600 DIMMs, for a total of 512GB</p>

**Caution**

Be sure to select the RAM modules with the correct speed rating that matches your server model.

## 7.3 Installing the Memory Upgrade Kit

In general, installing or replacing memory modules are tasks performed by an Oracle-qualified service technician. DIMMs are not hot-swappable parts, so the server must be powered down for this procedure. This section provides a high-level overview of the steps required for a server memory extension.

Detailed information is available in the chapter [Servicing Oracle Private Cloud Appliance Components](#) *Oracle Private Cloud Appliance Administrator's Guide* and the Service Manual for the server model in question.

**Table 7.3 Summary of the Compute Node Memory Extension Procedure**

Step	Description	Detailed Instructions
1	Read and follow the safety considerations and prerequisites for component replacement procedures.	Refer to the section “ <a href="#">Preparing Oracle Private Cloud Appliance for Service</a> ” in the Oracle Private Cloud Appliance Administrator's Guide.
2	Migrate virtual machines to another compute node and place the server in maintenance mode.	In the Service chapter of the Oracle Private Cloud Appliance Administrator's Guide, refer to the power-down instructions section for the server model in question.
3	Power down the server, disconnect all power cords and data cables, and then extend the server to the maintenance position.	Consult the Service Manual for your server model. Instructions are provided in the chapter "Preparing for Service".
4	Open the server fan door, and then remove the server top cover.	Consult the Service Manual for your server model. These specific instructions are part of the section “Preparing the Server for Component Replacement” in the chapter referenced in step 3.
5	Install the new DIMMs.	<p>For these detailed steps, you must navigate to a particular section in the Service Manual of your server model.</p> <ol style="list-style-type: none"> <li>a. In the Service chapter of the Oracle Private Cloud Appliance Administrator's Guide, refer to the Service Procedures section for the server model in question.</li> <li>b. Click the URL for the service procedures for DIMMs. This opens the Service Manual.</li> <li>c. Follow the steps in the section entitled “Install a DIMM”.</li> </ol>
6	Install the server top cover, and then close the server fan door.	Consult the Service Manual for your server model. Follow the steps in the section “Install the Server Top Cover”, which is near the start of the chapter referenced in step 7.
7	Return the server to its normal rack position, and then reconnect all data cables and power cords.	Consult the Service Manual for your server model. Refer to the chapter “Returning the Server to Operation”.
8	Power on the server. When it comes back online, take the compute node out of maintenance mode and verify that it resumes all normal operations in the server pool.	<p>In the Service chapter of the Oracle Private Cloud Appliance Administrator's Guide, refer to the power-down instructions section for the server model in question.</p> <p>The steps in this procedure are identical to placing the compute node in maintenance mode, except that you must clear the check box you selected previously.</p>





---

# Chapter 8 Extending Oracle Private Cloud Appliance - Additional Storage

## Table of Contents

8.1 Extending Storage Capacity of Ethernet-based Systems .....	71
8.1.1 Adding Disk Space to the Internal Storage Appliance .....	71
8.1.2 Adding External Ethernet Storage .....	72
8.1.3 Adding External Fibre Channel Storage .....	72
8.2 Extending Storage Capacity of InfiniBand-based Systems .....	72
8.2.1 Extending Storage Capacity Task Map .....	72
8.2.2 Adding External Fibre Channel Storage .....	73
8.2.3 Adding External InfiniBand Storage .....	85

Extending the Oracle Private Cloud Appliance by connecting and configuring additional storage can enhance the product by providing the disk space required for large repositories, for backup and recovery storage, and to provision virtual machines with virtual disks beyond the capacity already available on the appliance. This process does not require that the system is powered down, but cabling and configuration must match the Oracle Private Cloud Appliance requirements.

Different generations of the ZFS Storage Appliance have been delivered as part of Oracle Private Cloud Appliance, and the Controller Software continues to provide support for all previous generations. However, there are functional differences between the Oracle ZFS Storage Appliance ZS7-2, which is part of systems with an Ethernet-based network architecture, and the previous models of the ZFS Storage Appliance, which are part of systems with an InfiniBand-based network architecture.

There are options to extend the storage capacity of appliances with either type of network architecture, but the hardware involved and the required connectivity are different. For clarity, this chapter describes the addition of storage capacity separately for each network architecture.

## 8.1 Extending Storage Capacity of Ethernet-based Systems

This section describes the options to extend storage capacity for systems with an Ethernet-based network architecture.

### 8.1.1 Adding Disk Space to the Internal Storage Appliance

Systems with an Ethernet-based network architecture are built with an internal Oracle ZFS Storage Appliance ZS7-2. A major advantage is that the standard disk shelf already provides a significant amount of disk space at approximately 100TB. This is sufficient for the appliance internal 'system disk', as well as an Oracle VM storage repository for a virtualized production environment. Another significant feature is that, for environments with high storage demands, more disk shelves can be attached to the internal storage appliance.

A second disk shelf can be installed inside the Oracle Private Cloud Appliance base rack. It takes up the space of five compute nodes, thereby reducing the maximum capacity of the rack to a total of 20 compute nodes. If more capacity is required, up to 14 extra disk shelves can be installed in an additional rack. All of the extra disk shelves are connected to, and managed by, the two storage appliance controllers located in the bottom rack units of the Oracle Private Cloud Appliance base rack.

The addition of disk shelves to the Oracle ZFS Storage Appliance ZS7-2 is handled entirely by Oracle Advanced Customer Services. The process includes cabling, removal of compute nodes where required,

and updating server pool configurations. Please contact Oracle for more details about extending the capacity of the ZFS storage appliance.

## 8.1.2 Adding External Ethernet Storage

Ethernet-based external storage hardware must be made accessible over the data center network. Oracle Private Cloud Appliance, and the virtualized environment it hosts, access the storage through a custom network with external connectivity. Each external network corresponds with an isolated tunnel.

For instructions to set up a custom network, refer to the section [Network Customization](#) in the *Oracle Private Cloud Appliance Administrator's Guide*. For information about discovering and using storage resources within Oracle VM, refer to the section [Viewing and Managing Storage Resources](#), and the Oracle VM documentation.

## 8.1.3 Adding External Fibre Channel Storage

Oracle Server X8-2 expansion compute nodes can be ordered with *optional* 32Gbit Fibre Channel cards pre-installed. Field installation at a later time is possible as well. However, it is not possible to order new base racks with FC cards already installed in the compute nodes. Therefore, if Fibre Channel storage is part of your intended design, you should order a minimal base rack configuration and expansion compute nodes with the FC option.

Fibre Channel cables, switches and patch panels must be supplied by the customer. The installation of compute nodes with FC cards is performed by Oracle Advanced Customer Services. Once these compute nodes are integrated in the Oracle Private Cloud Appliance environment, the fibre channel HBAs can connect to standard FC switches and storage hardware in your data center. External FC storage configuration is managed through Oracle VM Manager. For more information, refer to the [Fibre Channel Storage Attached Network](#) section of the *Oracle VM Concepts Guide*.



### Caution

Oracle Server X8-2 expansion compute nodes with Fibre Channel cards can also be added to a system with InfiniBand-based network architecture. In that case, the vHBAs **must** be disabled for those compute nodes.

## 8.2 Extending Storage Capacity of InfiniBand-based Systems

This section describes the options to extend storage capacity for systems with an InfiniBand-based network architecture.

The Oracle Private Cloud Appliance can support additional Fibre Channel or InfiniBand storage devices connected to the Fabric Interconnects, depending on the technology required. This chapter describes how to extend the appliance for each of these different technologies.

### 8.2.1 Extending Storage Capacity Task Map

[Table 8.1](#) describes the steps for extending the Oracle Private Cloud Appliance storage capacity at your site by connecting additional Fibre Channel or InfiniBand storage hardware to the Fabric Interconnects.

**Table 8.1 Steps for Extending Oracle Private Cloud Appliance Storage Capacity**

Step	Description	Links
1	Review safety guidelines and site checklists.	<a href="#">Section 6.2, "Review Safety Guidelines"</a>

Step	Description	Links
2	<p>Identify the type of external storage you are installing:</p> <ul style="list-style-type: none"> <li>• Fibre Channel storage hardware</li> <li>• InfiniBand storage hardware</li> </ul>	<p><a href="#">Section 8.2.2, “Adding External Fibre Channel Storage”</a></p> <p><a href="#">Section 8.2.3, “Adding External InfiniBand Storage”</a></p>
3	<p>Connect the external storage to the appliance.</p> <p>This involves physical cabling and automated configuration resulting from it.</p>	<p>Fibre Channel:</p> <ul style="list-style-type: none"> <li>• <a href="#">Section 8.2.2.1, “Creation of Virtual Host Bus Adaptors (vHBAs)”</a></li> <li>• <a href="#">Section 8.2.2.3, “Connecting Fibre Channel Hardware”</a></li> </ul> <p>InfiniBand:</p> <ul style="list-style-type: none"> <li>• <a href="#">Section 8.2.3.1, “Connecting InfiniBand Storage Hardware”</a></li> <li>• <a href="#">Section 8.2.3.2, “IP Address Allocation”</a></li> </ul>
4	<p>Configure the external storage for use with Oracle Private Cloud Appliance.</p> <p>This refers to software configuration on the storage hardware connected to the Oracle Private Cloud Appliance.</p>	<p>Fibre Channel:</p> <ul style="list-style-type: none"> <li>• <a href="#">Section 8.2.2.4, “Zone Configuration”</a></li> <li>• <a href="#">Section 8.2.2.5, “Configuring Your Storage Appliance”</a></li> </ul> <p>InfiniBand:</p> <ul style="list-style-type: none"> <li>• <a href="#">Section 8.2.3.3, “Configuring the ZFS Storage Appliance”</a></li> </ul>
5	<p>Configure Oracle Private Cloud Appliance to enable the external storage.</p> <p>This refers to configuration in Oracle VM Manager to enable connectivity with the storage hardware connected to the Oracle Private Cloud Appliance.</p>	<p>Fibre Channel:</p> <ul style="list-style-type: none"> <li>• <a href="#">Section 8.2.2.6, “Updating Physical Disks in Oracle VM Manager”</a></li> </ul> <p>InfiniBand:</p> <ul style="list-style-type: none"> <li>• <a href="#">Section 8.2.3.4, “Enabling External IPoIB Storage in Oracle VM Manager”</a></li> </ul>

## 8.2.2 Adding External Fibre Channel Storage



### Oracle Server X8-2 compute nodes with optional Fibre Channel expansion cards

For any Oracle Private Cloud Appliance, with Ethernet-based as well as InfiniBand-based network architecture, Oracle Server X8-2 expansion compute nodes can be ordered with optional 32Gbit Fibre Channel cards pre-installed. Through these *physical* HBAs, the expansion compute nodes can be connected to standard FC switches and storage hardware in your data center. However, the *virtual* HBAs, which are normally used to connect to FC storage attached to the Fabric Interconnects, **must** be disabled.

FC storage connected through the physical HBAs in the compute nodes behaves like any standard FC storage, on both InfiniBand-based and Ethernet-based Oracle Private Cloud Appliance systems. See [Section 8.1.3, “Adding External Fibre Channel Storage”](#) for Ethernet-based systems.

The configuration described in this section refers to FC storage attached to the FC ports in the I/O modules of the Fabric Interconnects. The information in this section is specific to systems with an InfiniBand-based network architecture, and compute nodes without FC expansion cards.

Fibre Channel (FC) can be used to connect additional FC-capable storage devices to the Oracle Private Cloud Appliance (PCA) using the FC ports available on the Fabric Interconnects. The Oracle Private Cloud Appliance automatically creates virtual Host Bus Adaptors (vHBAs) on each server in the Oracle Private Cloud Appliance rack. The World Wide Port Names (WWPNs) defined for these vHBAs can be used to define Initiator Groups on your storage appliance to facilitate access to the LUNs that you wish to make available to your compute and management nodes.



### Note

Oracle Private Cloud Appliance can be ordered without Fibre Channel I/O modules in the Fabric Interconnects. However, these modules can be ordered separately and installed at a later time. For installation instructions, see [Section 8.2.2.2, “Installing Optional Fibre Channel I/O Modules”](#).

Oracle Private Cloud Appliance introduces the concept of "storage clouds" that group together the FC ports used to connect the Fabric Interconnects to your own FC switch or switches. Two FC ports on each of the Fabric Interconnects are assigned to a cloud. A vHBA is created on each server for each storage cloud. A total of four storage clouds are defined when the Oracle Private Cloud Appliance is provisioned, resulting in four vHBAs on each of the compute and management nodes.

Storage clouds allow you to cable and configure your external storage in such a way as to improve overall throughput or to build a fully HA enabled infrastructure. The storage clouds are created and configured automatically, so that all you need to do is choose the type of HA configuration that you wish to use and then cable accordingly. The design of the storage cloud configuration reduces complexity in configuration and allows the greatest possible flexibility in terms of your HA requirements.

At bare minimum, a single cloud can be used to access storage on each of your compute nodes or on your management nodes. However a simple HA configuration is recommended. To achieve this, you should use at least two clouds cross-cabled from each of the Fabric Interconnects. For a fully HA-enabled environment, you can use all four clouds cross-cabled between the Fabric Interconnects and two of your own FC switches.

If you do not require HA, you can use multiple clouds to increase overall throughput. In this configuration, you can cable three clouds into separate switches and connect them to the same storage. If one cloud should fail, you are able to bring up the fourth cloud to take over its function and to maintain the same level of throughput.

Each of these configurations is achieved entirely through the different ways in which your FC switch or switches are patched into the Fabric Interconnects. These approaches are described in more detail in [Section 8.2.2.3, “Connecting Fibre Channel Hardware”](#).

Since storage clouds map directly to the vHBAs on each server, it is possible to configure zones on your FC switch or switches to securely separate traffic for each storage cloud. This gives you the opportunity to use separate clouds for different purposes using secured channels to facilitate communication between servers and the LUNs on your storage. This is described in [Section 8.2.2.4, “Zone Configuration”](#).

To use Fibre Channel with the Oracle Private Cloud Appliance, you must supply your own NPIV-capable FC switch or switches. It is not possible to simply patch your FC-capable storage directly into the FC ports on the Fabric Interconnects. This is because the Fabric Interconnects use NPIV – Fibre Channel Node Port ID Virtualization – to map the port nodes to the World Wide Node Names (WWNNs) of the vHBAs that are created on each server. Software required to translate WWPNS to WWNNs does not exist on the storage heads of most FC storage devices, so directly attaching the storage device would prevent registration of the WWPNS for the vHBAs available on each server.

### 8.2.2.1 Creation of Virtual Host Bus Adaptors (vHBAs)

Each server in the Oracle Private Cloud Appliance is connected to the Fabric Interconnects via an InfiniBand (IB) connection. The Fabric Interconnects are capable of translating connections on their Fibre Channel ports to reroute them over these IB connections. To facilitate this, vHBAs must be defined on each server to map to a storage cloud defined on the Fabric Interconnects. The storage cloud that these vHBAs map to, determine which FC ports they relate to on the Fabric Interconnects.

During the initial configuration of the management nodes, each of the storage clouds is configured automatically on the Fabric Interconnects and vHBAs for each cloud are created on each of the management nodes. WWNNs and WWPNS are generated for the vHBAs on the management nodes. When you cable for a storage cloud and connect your storage appliance, you can add these WWPNS to a storage initiator group and export the LUNs that you wish to make available to your management nodes. These disks can then be mounted on the management nodes as needed. Management nodes are configured with vHBAs to allow them to connect to the different storage clouds so that it is possible to use externally connected storage to store backup and log data, if required.



#### Tip

It is important to distinguish between WWNNs and WWPNS. A WWNN is used to identify a device or node such as an HBA, while a WWPNS is used to identify a port that is accessible via that same device. Since some devices can have multiple ports, a device may have a single WWNN and multiple WWPNS.

In the case of the vHBAs that are generated on each compute node, there is a single WWNN and a single WWPNS for each vHBA. While these may look almost identical, the fourth hexadecimal octet that makes up the WWNN differs. This is illustrated as follows:

WWPN	WWNN
50:01:39:70:00:4F:91:00	50:01:39:71:00:4F:91:00

When configuring storage initiators and initiator groups, you should ensure that you are configuring these for the WWPNS for each vHBA. If you use the WWNN for any vHBA within an initiator group, that initiator group may not function as expected.

Compute nodes are configured similarly, although the process takes place during compute node provisioning, so that compute nodes are configured as they are provisioned. This means that if you add a compute node to the rack, it is automatically configured so that it is ready to access any externally attached storage. Creation and configuration of vHBAs has also been built into the upgrade process, so that external FC storage can be used with existing environments. Once you have cabled for your storage and have configured the initiator groups on your storage, you can view the LUNs that you make available to your compute nodes directly in Oracle VM Manager.

If you purchased an Oracle Private Cloud Appliance with factory-installed Fibre Channel I/O modules, then these processes are entirely automated and do not require any intervention.

If you purchased an Oracle Private Cloud Appliance without Fibre Channel I/O modules, and you need to install the optional modules into the Fabric Interconnects first, then carefully follow the instructions in [Section 8.2.2.2, “Installing Optional Fibre Channel I/O Modules”](#).

## 8.2.2.2 Installing Optional Fibre Channel I/O Modules

The X5-2 base rack with Oracle Private Cloud Appliance Release 2.0.3 software marks the first configuration that can be ordered with or without fibre channel I/O modules. In an Oracle Private Cloud Appliance without fibre channel connectivity, slots 3 and 12 of the two Fabric Interconnects are empty. As a consequence, the Software Clouds and vHBAs for Fibre Channel connectivity are not created during provisioning.

If you decide to add Fibre Channel connectivity at a later time, you must order and install the optional Fibre Channel I/O modules, upgrade the controller software and configure the necessary storage clouds and vHBAs.

### Adding Fibre Channel Connectivity to a Provisioned Appliance

1. If your Oracle Private Cloud Appliance has no factory-installed Fibre Channel I/O modules you must first make sure that the appliance controller software is Release 2.1.1 or newer. The update process typically takes at least 2.5 hours.

For detailed software update instructions, see [Updating Oracle Private Cloud Appliance](#) in the *Oracle Private Cloud Appliance Administrator's Guide*.

2. When the controller software update has completed successfully, install a total of four Fibre Channel I/O modules in slots 3 and 12 of each Oracle Fabric Interconnect F1-15. You can install the new modules without powering down the Fabric Interconnect.



#### Caution

You must use only slots 3 and 12. Fibre Channel I/O modules installed in other slots are likely to cause future upgrade issues and are therefore not supported within Oracle Private Cloud Appliance.

- a. Ensure that the module's component side is facing towards the right edge of the chassis. Also, ensure that the module's top and bottom edges are correctly aligned with the upper and lower tracks at the left edge of each module slot.
- b. Secure the module in the slot by tightening the securing screw located at the bottom edge of the module's rear panel.



#### Caution

Tighten the screws so that they are lightly snug. If you overtighten the screws they can break.

3. When all four Fibre Channel I/O modules have been installed, configure the storage clouds and vHBAs by running the appropriate commands from the Oracle Private Cloud Appliance Command Line Interface (CLI).
  - a. Configure the vHBAs on both management nodes.

```
PCA> configure vbas ovcamn05r1 ovcamn06r1

Compute_Node      Status
-----
ovcamn05r1       Succeeded
```

```

ovcamn06r1          Succeeded
-----
2 rows displayed
Status: Success
    
```

b. Verify that the clouds have been configured.

```

PCA> list storage-network

Network_Name          Description
-----
Cloud_A               Default Storage Cloud ru22 port1 - Do not delete or modify
Cloud_B               Default Storage Cloud ru22 port2 - Do not delete or modify
Cloud_C               Default Storage Cloud ru15 port1 - Do not delete or modify
Cloud_D               Default Storage Cloud ru15 port2 - Do not delete or modify
-----
4 rows displayed
Status: Success
    
```

c. If the 4 storage clouds have been configured correctly, configure the vHBAs on all compute nodes.

```

PCA> configure vhbas ALL

Compute_Node          Status
-----
ovcacn07r1            Succeeded
ovcacn08r1            Succeeded
[...]
ovcacn36r1            Succeeded
ovcacn37r1            Succeeded
-----
20 rows displayed
Status: Success
    
```

These steps provide only an overview of the procedure. For detailed instructions, see [Enabling Fibre Channel Connectivity on a Provisioned Appliance](#) in the *Oracle Private Cloud Appliance Administrator's Guide*.

### 8.2.2.3 Connecting Fibre Channel Hardware

Cabling requirements to attach external FC-capable storage to the Oracle Private Cloud Appliance are very specific due to the grouping of FC ports for storage clouds on the Fabric Interconnects. You must ensure that when cabling for a storage cloud, both ports defined for the storage cloud are cabled into your FC switch. Ideally, you should cable all four FC ports on each of the Fabric Interconnects within the rack for a total of eight cables, at the same time. These should either connect to a single Fibre Channel switch outside of the rack, or should be cross-connected to two Fibre Channel switches outside of the rack to improve redundancy. The following table describes how ports are grouped to create each cloud:

**Table 8.2 Fibre Channel Port and Cloud Mappings**

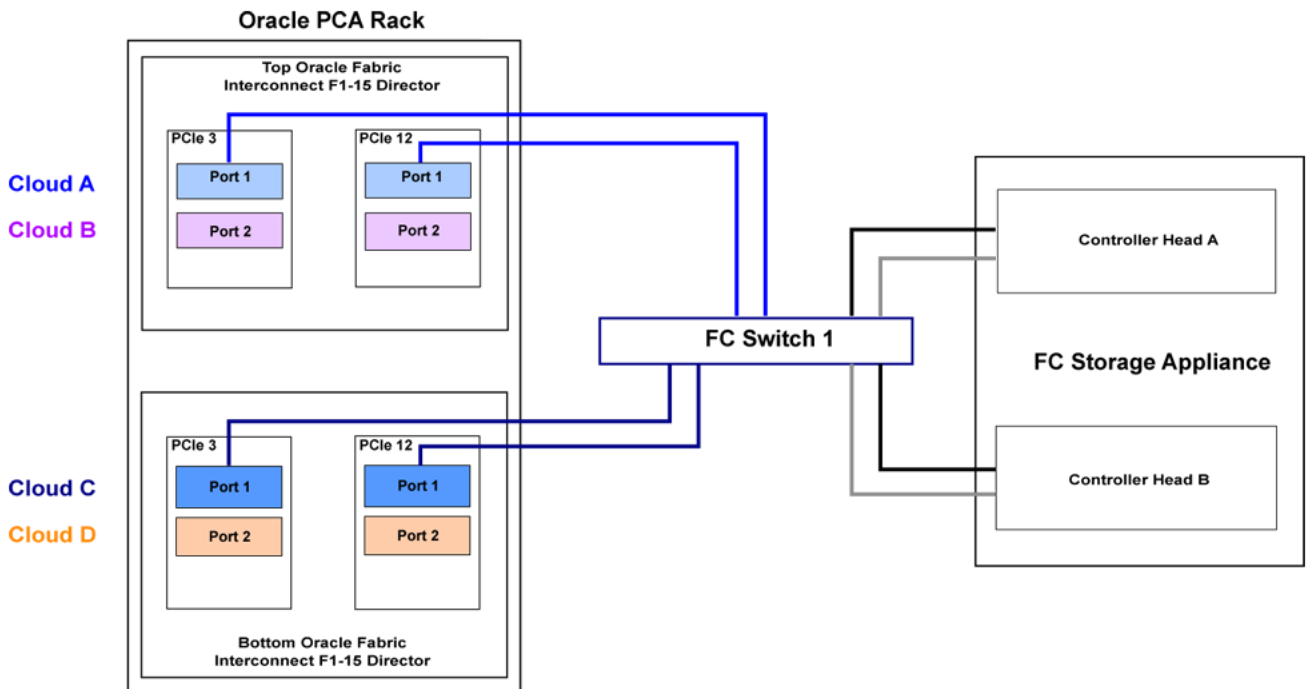
Mapping	Top Oracle Fabric Interconnect F1-15 (RU 22-25):		Bottom Oracle Fabric Interconnect F1-15 (RU 15-18):	
	I/O Module 3	I/O Module 12	I/O Module 3	I/O Module 12
A	Port 1	Port 1	none	none
B	Port 2	Port 2	none	none
C	none	none	Port 1	Port 1

Mapping	Top Oracle Fabric Interconnect F1-15 (RU 22-25):		Bottom Oracle Fabric Interconnect F1-15 (RU 15-18):	
Cloud Name	I/O Module 3	I/O Module 12	I/O Module 3	I/O Module 12
D	none	none	Port 2	Port 2

For a bare minimum configuration, you can use two cables to connect your FC switch to cloud A. This would require that you connect cables to the top Fabric Interconnect in the rack, using I/O Module 3 Port 1 and I/O Module 12 Port 1. Connect each controller head on your FC storage appliance to your FC switch. Typically you would use two cables for each controller head.

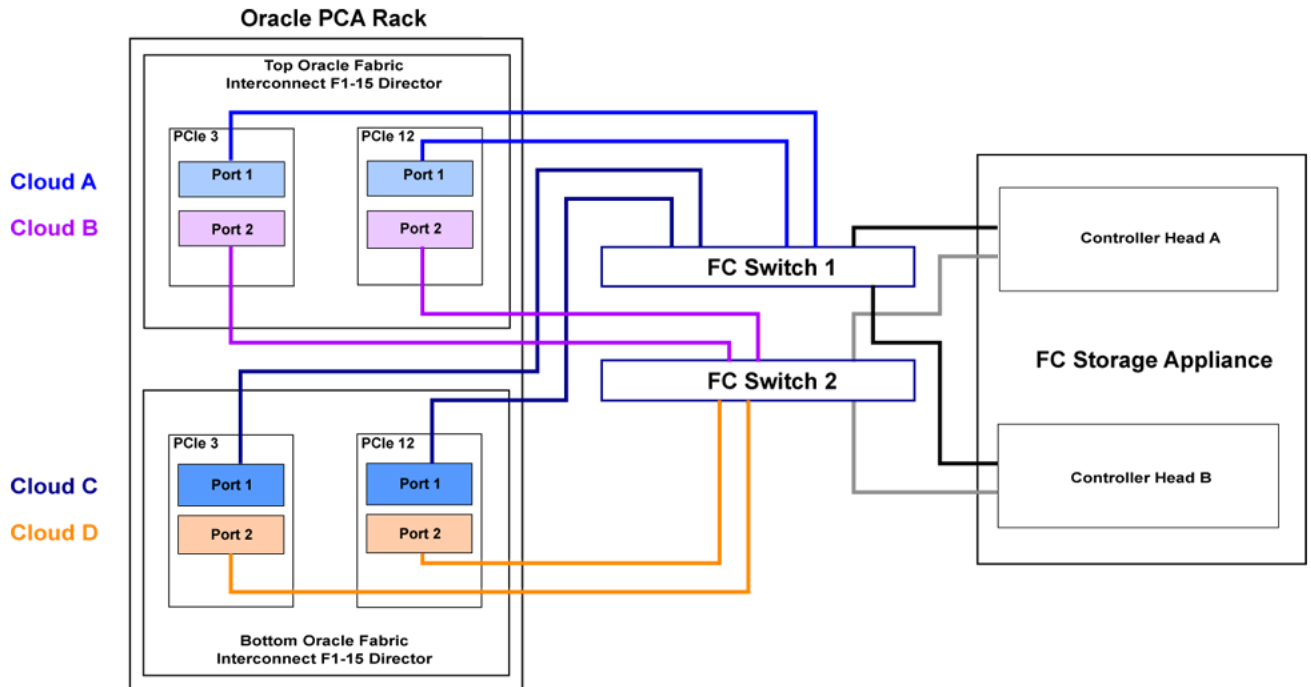
For a basic HA configuration, you should cable for cloud A and cloud C. This allows you to use both of the Fabric Interconnects in the rack to improve redundancy. This is the recommended default minimum configuration. It is important that cloud A and cloud C are cabled to the same FC switch.

**Figure 8.1 Minimum HA FC Storage Cabling Configuration**



To maximize an HA configuration, by using an additional FC switch, you can cable cloud B and cloud D into the second FC switch. In this situation you should cross-cable the controller heads on your FC storage appliance, so that each controller head is connected to each of your FC switches.



**Figure 8.2 Full HA FC Storage Cabling Configuration**

Other cabling possibilities exist, such as the aforementioned configuration where clouds are each connected to independent FC switches to maximize throughput or to achieve physical traffic separation. Usually, in this configuration, one cloud is left disconnected in standby so that in the event that one of the other clouds fails for some reason, it is possible to connect the standby cloud to take over the function of the failed cloud. In this situation, you would need to configure so that the WWPNs for the standby cloud are substituted for the failed cloud in the initiator groups on your storage appliance.

The most important points to remember when cabling are as follows:

- Clouds are comprised of two equivalent ports on alternate I/O Modules on the same Fabric Interconnect. For instance, cloud A on the top Fabric Interconnect in the rack is comprised of Port 1 on I/O Module 3 and Port 1 on I/O Module 12; while cloud D on the bottom Fabric Interconnect in the rack is comprised of Port 2 on I/O Module 3 and Port 2 on I/O Module 12.
- Although not required, you should cable both ports that belong to a cloud to provide redundancy and to allow better throughput.
- Unless you have defined zones that span more than one FC switch, the ports that belong to a cloud must be cabled to connect to the same FC switch.
- For HA configurations, clouds on alternate Fabric Interconnects should be connected to the same FC switch. For instance, cloud A and cloud C should be connected to the same FC switch; while cloud B and cloud D can be connected to an alternate switch.

You can connect FC cables to the appliance at any time, and no system shutdown is required, allowing you to dynamically add storage or improve redundancy as required.

## 8.2.2.4 Zone Configuration

The Oracle Private Cloud Appliance requires that Fibre Channel Zoning is configured on the external FC switches to control and limit the amount of traffic on the storage network for efficient LUN discovery and to maximize stability within the environment.

FC Zoning is also used to enhance security by providing an extra layer of traffic separation on the storage network. Even if you are using storage initiator groups to perform LUN masking, it is generally considered good practice to also configure FC zones to limit the exposure of LUNs and unrestricted use of this network medium. Zone configuration is very useful in the situation where the FC switch or switches are shared with other devices apart from the Oracle Private Cloud Appliance.

The Oracle Private Cloud Appliance supports single initiator pWWN zoning, in line with industry best-practice. It is highly recommended that you configure single initiator pWWN zoning for all Fibre Channel connections to the rack. This requires a Fibre Channel switch that supports NPIV.

For all storage clouds that are cable-connected, at least one zone should be configured per WWPN on each compute node. However, multiple zones may be created for each WWPN depending on your cabling. In a setup using all four storage clouds, four zones should exist on the Fibre Channel switch for each compute node. You can obtain a listing of the WWPNs for the compute nodes by running the `pca-admin list wwpn-info` command.

Using a Fibre Channel storage device with two controller heads in an active/active cluster, and two targets configured for each head, every LUN has 4 storage paths. A storage path is a connection between a target and a storage cloud. Each LUN has two active paths between a storage cloud and the controller head that owns the LUN, and an additional two standby paths between the same storage cloud and the other controller head. If the storage head that owns the LUN should fail, the other storage head takes over and the standby paths become active. This way the storage connectivity for a given cloud remains uninterrupted. To better support failover/failback operations, consider employing a zoning strategy using two zones per storage cloud. This approach involves one zone connecting a given cloud to a target of one controller head, and a second zone connecting that same cloud to a target of the other controller head.

The configuration of Fibre Channel single initiator pWWN zones is *not* optional. If you had previously attempted to extend your storage using Fibre Channel and did not configure any zones on your FC switch, you should do so now. Furthermore, if you previously configured D,P zones, it is important that you rezone your switches to use single initiator pWWN zones.

Please refer to the documentation of your switch vendor for more information on the configuration steps that you must perform to configure single initiator pWWN zoning.

## 8.2.2.5 Configuring Your Storage Appliance

Some initial configuration steps may be required on your appliance before you are able to access storage directly within Oracle VM Manager. Typically these steps involve configuring some form of 'LUN masking' achieved by mapping LUNs to particular initiator groups that ultimately define the servers and clouds that have access to each LUN. These steps are outlined as follows:

1. Create the initiators that you intend to use to identify the different servers in each cloud within your storage appliance. This is achieved by registering the World Wide Port Names (WWPNs) that were created for the vHBAs on each server in the Oracle Private Cloud Appliance rack with your appliance and assigning them aliases so that they can be easily identified as belonging to a particular server and cloud. Most appliances should be able to see the WWPNs presented by the Fabric Interconnects and you can use the Oracle Private Cloud Appliance command line interface to identify matching WWPNs and match them with the recommended aliases that you should use when creating each initiator. This is achieved by running the `pca-admin list wwpn-info` command:

```
PCA> list wwpn-info
```

```

WWPN                               vHBA      Cloud_Name  Server      Type  Alias
----                               -
50:01:39:70:00:69:F1:06           vhba01    Cloud_A     ovcacn08r1  CN    ovcacn08r1-Cloud_A
50:01:39:70:00:69:F1:08           vhba01    Cloud_A     ovcacn32r1  CN    ovcacn32r1-Cloud_A
50:01:39:70:00:69:F1:0C           vhba01    Cloud_A     ovcacn30r1  CN    ovcacn30r1-Cloud_A
50:01:39:70:00:69:F1:0A           vhba01    Cloud_A     ovcacn29r1  CN    ovcacn29r1-Cloud_A
50:01:39:70:00:69:F1:07           vhba02    Cloud_B     ovcacn08r1  CN    ovcacn08r1-Cloud_B
50:01:39:70:00:69:F1:09           vhba02    Cloud_B     ovcacn32r1  CN    ovcacn32r1-Cloud_B
50:01:39:70:00:69:F1:0B           vhba02    Cloud_B     ovcacn29r1  CN    ovcacn29r1-Cloud_B
50:01:39:70:00:69:F1:0D           vhba02    Cloud_B     ovcacn30r1  CN    ovcacn30r1-Cloud_B
50:01:39:70:00:6A:11:0A           vhba03    Cloud_C     ovcacn29r1  CN    ovcacn29r1-Cloud_C
50:01:39:70:00:6A:11:0C           vhba03    Cloud_C     ovcacn30r1  CN    ovcacn30r1-Cloud_C
50:01:39:70:00:6A:11:08           vhba03    Cloud_C     ovcacn32r1  CN    ovcacn32r1-Cloud_C
50:01:39:70:00:6A:11:06           vhba03    Cloud_C     ovcacn08r1  CN    ovcacn08r1-Cloud_C
50:01:39:70:00:6A:11:0D           vhba04    Cloud_D     ovcacn30r1  CN    ovcacn30r1-Cloud_D
50:01:39:70:00:6A:11:0B           vhba04    Cloud_D     ovcacn29r1  CN    ovcacn29r1-Cloud_D
50:01:39:70:00:6A:11:09           vhba04    Cloud_D     ovcacn32r1  CN    ovcacn32r1-Cloud_D
50:01:39:70:00:6A:11:07           vhba04    Cloud_D     ovcacn08r1  CN    ovcacn08r1-Cloud_D
-----
16 rows displayed

Status: Success

```

Note the Alias column in the example output. Use the values presented in this column for each matching WWPN when you configure the initiators on your appliance. The Oracle Private Cloud Appliance CLI is discussed in more detail in the section entitled “[The Oracle Private Cloud Appliance Command Line Interface \(CLI\)](#)” in the *Oracle Private Cloud Appliance Administrator's Guide*.

2. Create initiator groups that define how your storage should be presented to the different compute nodes and virtual machines within your environment. This step is very dependent on your own storage requirements and should be carefully planned for before you start using the storage appliance in conjunction with the Oracle Private Cloud Appliance.

Remember that there is a WWPN for each vHBA on each compute node, representing each storage cloud defined on the Fabric Interconnects. In an HA configuration, you attach the WWPNs for the different storage clouds to the same initiator group. Using this approach, if a storage cloud fails, the LUNs that are exposed for the initiator group are still available to each compute node via an alternate vHBA. Alternatively, you can sacrifice high availability for a more flexible configuration where the WWPN for each vHBA on each compute node is used independently across different initiator groups. This type of configuration may make more sense where you intend to separate out storage for different functions, such as storage that might be provided to virtual machines and storage that might be used for repositories and server related functions.

3. Map LUNs that you have created on your storage appliance to the initiator groups you have just created. You should only map LUNs to the initiator groups that you intend to use for those LUNs. Although remapping LUNs is possible, you should remember that remapping a LUN once it is in use by Oracle VM is not recommended and may cause unexpected behavior within Oracle VM. Therefore, it is important to carefully plan this step around your requirements before proceeding with this configuration.



**Caution**

A maximum of 1024 LUN paths per host applies.

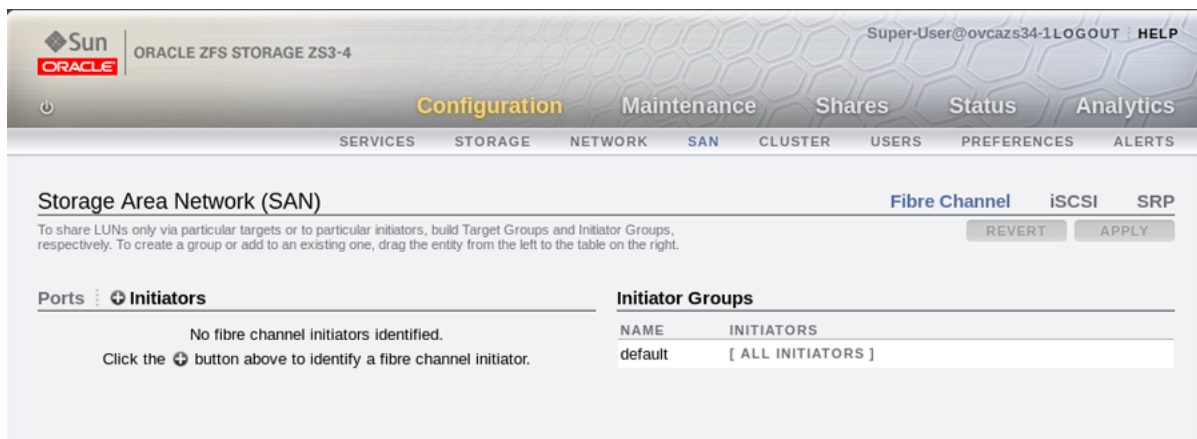
To perform these steps, you should refer to the appropriate documentation for your appliance or contact your vendor for assistance with these steps. An example is provided in [Example 8.1, “Configuring an Oracle Storage Appliance ZS3-4”](#) to show you how you can configure an Oracle Storage Appliance ZS3-4 correctly for this purpose.

### Example 8.1 Configuring an Oracle Storage Appliance ZS3-4

It is possible to configure an Oracle Storage Appliance ZS3-4 either using the web-based user interface or the command line interface. These instructions presume that you are using the web-based user interface, since the interface shows all of the WWPNs that it can actually detect.

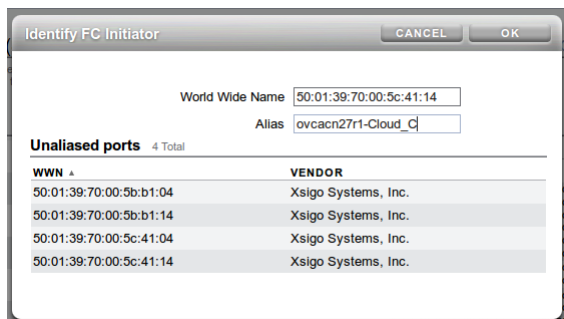
1. Log into the web-based user interface as either root or a user with adequate permissions to create initiator groups and LUNs.
2. Click on the [Configuration](#) link, then click on the [SAN](#) link. If [Fibre Channel](#) is not already selected, click on the [Fibre Channel](#) link. The SAN Summary page is displayed.
3. Click on the [+ Initiators](#) link as indicated in the image below:

**Figure 8.3 Click on the + Initiators link.**



4. The [Identify FC Initiator](#) dialog opens and lists the WWPNs that the ZS3-4 is able to detect. You should create a user-friendly alias for each WWPN listed so that it is easy to identify how these map onto the nodes and clouds configured within the Oracle Private Cloud Appliance. Use the Oracle Private Cloud Appliance command line interface on the master management node to obtain a listing of all of these WWPNs and the recommended aliases that you should use. You can do this by running the `pca-admin list wwpn-info` command. Find the matching WWPN in the [Identify FC Initiator](#) dialog and click on it to enter the recommended alias in the Alias field. Repeat this action for every WWPN. This step is illustrated in the figure below:

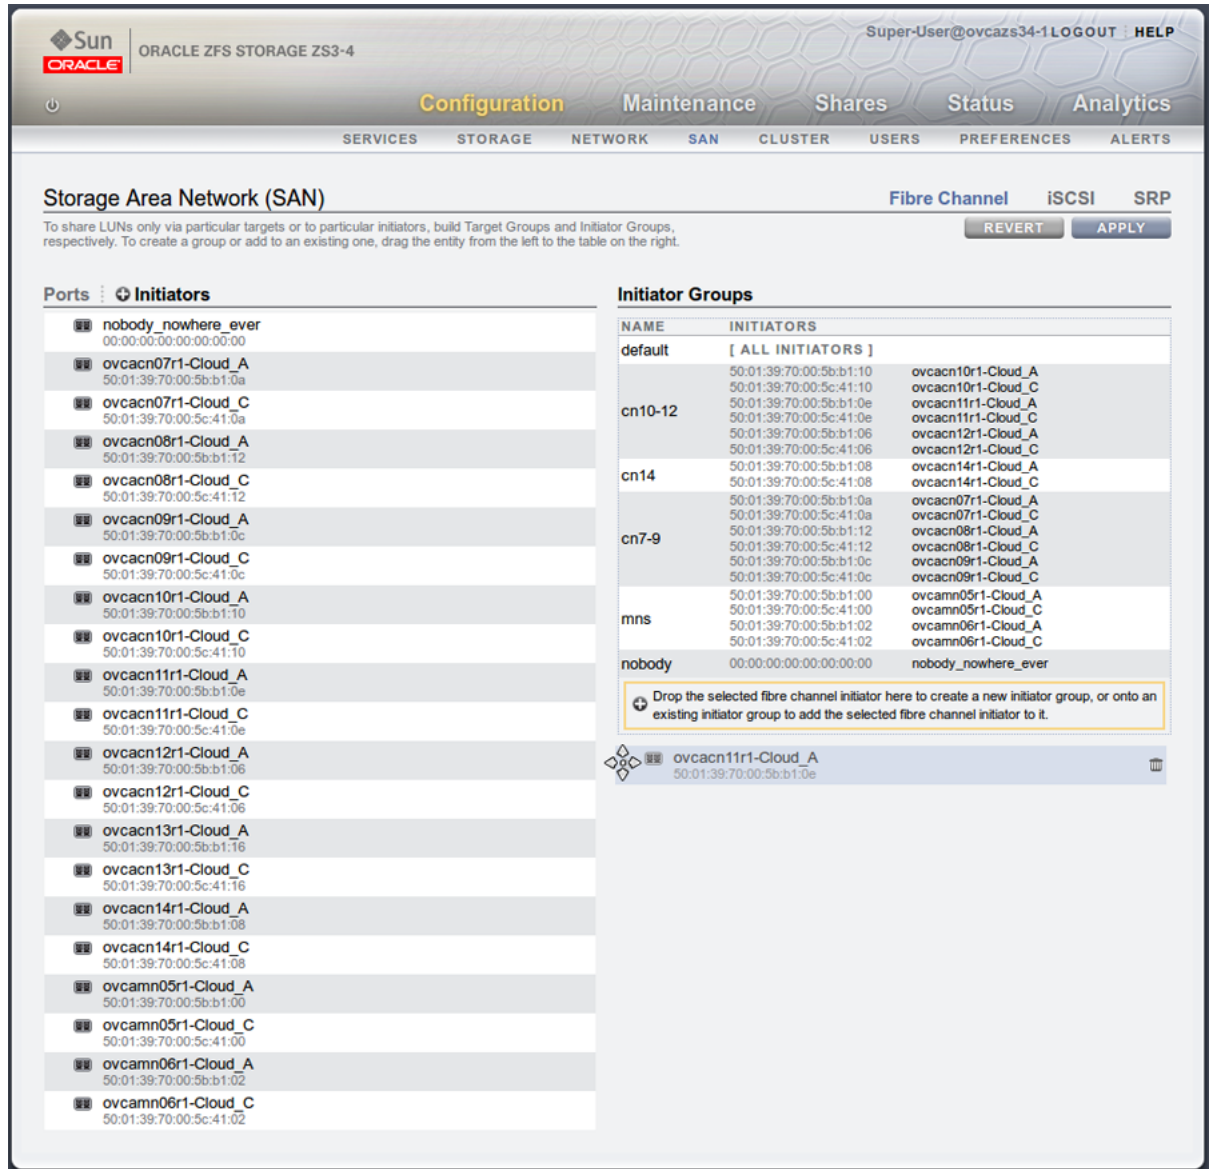
**Figure 8.4 Use the [Identify FC Initiator](#) dialog to define user-friendly aliases for the WWPNs.**



5. When you have finished creating aliases for each WWPN, you can start to create the storage initiator groups where LUNs are presented later in the configuration. To create a new storage initiator group in the SAN Summary view, move your mouse over the left hand side of the WWN name until you see

the [move](#) icon appear, then drag it to the right hand side of the screen and drop it in the [Create New Initiator Group](#) box that appears. This step is illustrated in the figure below:

**Figure 8.5** Use the [SAN Summary](#) page to create a new storage initiator group by dragging an initiator across the page.



You can create as many initiator groups as you like. In most cases, if your storage is not used outside of the Oracle Private Cloud Appliance and you do not need to segregate storage within the Oracle Private Cloud Appliance, you can simply use the default initiator group. However, if your storage is shared with appliances outside of the Oracle Private Cloud Appliance, you must, at least create an initiator group for all of the initiators that belong to nodes within the Oracle Private Cloud Appliance. If you need to segregate storage further, you can create initiator groups that only include the initiators for particular compute nodes, as is illustrated in the screenshot.

- Once you have created a new initiator group, you can roll over it using your mouse and then click on the edit icon that appears in the form of a pencil. A dialog appears where you are able to edit the initiator name to change it to something more appropriate. You can also select the rest of the initiators

that you wish to include in the initiator group. Click **OK** when you have finished editing the initiator group.

7. When you have finished creating and editing initiator groups, it is important that you click on the **Apply** button on the SAN Summary page to save the changes that you have made.
8. You can now define how LUNs map to your initiator groups depending on your requirements. This is achieved by clicking on the [Shares](#) link in the navigation bar at the top of the page. On the [Project Summary](#) page, you can view any storage pools that you have defined and the LUNs associated with those pools. You can also create new LUNs for each storage pool as required. On this page, you can click to edit a LUN and edit the storage initiator groups that it is exposed to. In this view, you can make a LUN accessible to any of the storage initiators that you defined in the previous steps.

### 8.2.2.6 Updating Physical Disks in Oracle VM Manager

Since vHBAs are created on each compute node during the provisioning or upgrade process, the storage should be automatically refreshed within Oracle VM Manager as part of this process and the physical disks should be visible within Oracle VM Manager immediately. However, in the case that a disk is not displayed within Oracle VM Manager after the initial server discovery is performed, it is possible that you may need to rediscover the compute nodes so that Oracle VM Manager sees the vHBAs and the LUNs that are accessible to them. The following steps describe actions that should be performed in Oracle VM Manager to start making use of your newly attached FC storage in the case where these disks do not appear automatically within Oracle VM Manager.



#### Caution

Reprovisioning restores a compute node to a clean state. If a compute node with active connections to external storage repositories is reprovisioned, the external storage connections need to be configured again after reprovisioning.

1. Log into the Oracle VM Manager web-interface on the Oracle Private Cloud Appliance.
2. Click on the [Servers and VMs](#) tab and select all of your compute node servers.
3. Click on the [Rediscover Servers](#) icon to rediscover all of your servers.
4. After you have finished rediscovery, check that there are four new vHBAs for each server by clicking on the server in the navigation pane and then changing the perspective to [Storage Initiators](#).
5. Click on the [Storage](#) tab.
6. Expand the [SAN Servers](#) item in the navigation pane and select the [Unmanaged Fibre Channel Storage Array](#).
7. Click on the [Edit](#) icon and navigate to the [Admin Servers](#) tab in the dialog that appears.
8. Add all of the compute node servers as admin servers for the storage array and click **OK** to close the dialog.
9. Click on the [Refresh SAN Server](#) icon and click **OK** on the confirmation dialog.
10. Wait for all of your servers to be updated. This operation can take several minutes to complete.
11. The LUNs display in the [Physical Disks](#) perspective for the [Unmanaged Fibre Channel Storage Array](#).
12. You can check that the LUNs are also available to your servers on the [Servers and VMs](#) tab by clicking on a server and changing the perspective to the [Physical Disks](#) view.



## 8.2.3 Adding External InfiniBand Storage

InfiniBand can be used to connect additional ZFS Storage Appliances to the Oracle Private Cloud Appliance (PCA) using the available ports on the Fabric Interconnects. Since ZFS Storage Appliances are dual-headed, each ZFS Controller head has its own InfiniBand connection to each of the Fabric Interconnects, providing redundancy both on the side of the ZFS Storage Appliance and on the side of the Oracle Private Cloud Appliance.

The Oracle Fabric Interconnect F1-15s have only 4 available InfiniBand ports – 2 per Fabric Interconnect – to attach external storage. If higher performance is required, or if multiple storage appliances need to be connected, you may add two external NM2-36P Sun Datacenter InfiniBand Expansion Switches between the external storage and the Fabric Interconnects inside the appliance rack. A ZFS Storage Appliance may be connected to the InfiniBand Switches in active/passive or active/active configuration. If an active/active configuration is selected, then both storage controllers must be expanded with additional HCAs for a total of 4 InfiniBand ports per ZFS controller head.



### Tip

For the latest, detailed information about InfiniBand storage with Oracle Private Cloud Appliance, please refer to the white paper entitled *Expanding Oracle Private Cloud Appliance Using Oracle ZFS Storage Appliance*.

### 8.2.3.1 Connecting InfiniBand Storage Hardware

The recommended cabling configuration that should be implemented to connect the ZFS appliance to the Oracle Private Cloud Appliance cross-connects each controller head to each of the Fabric Interconnects for a total of four connections. This configuration maximizes redundancy and throughput. The following table describes how the cabling should be connected between the Fabric Interconnects on the Oracle Private Cloud Appliance and the ZFS Storage Appliance.

**Table 8.3 InfiniBand Cabling to Attach a ZFS Storage Appliance to the Oracle Private Cloud Appliance Fabric Interconnects**

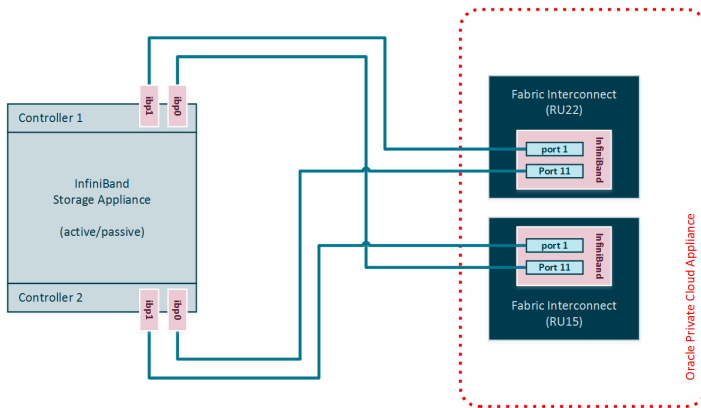
Top Oracle Fabric Interconnect F1-15 (RU 22-25):	Bottom Oracle Fabric Interconnect F1-15 (RU 15-18):	ZFS Storage Appliance Controller Head 1	ZFS Storage Appliance Controller Head 2
IB Port 1	none	IB Port 1 (ibp0)	none
IB Port 11	none	none	IB Port 1 (ibp0)
none	IB Port 1	IB Port 2 (ibp1)	none
none	IB Port 11	none	IB Port 2 (ibp1)

If you install two InfiniBand switches between the external storage and the appliance rack, then those switches connect to the Fabric Interconnect ports using the same cabling pattern as the two controllers of a directly attached ZFS Storage Appliance. On the side of the ZFS Storage Appliance, the cabling pattern in active/passive configuration remains identical whether it is connected to the InfiniBand switches or directly to the Fabric Interconnects.

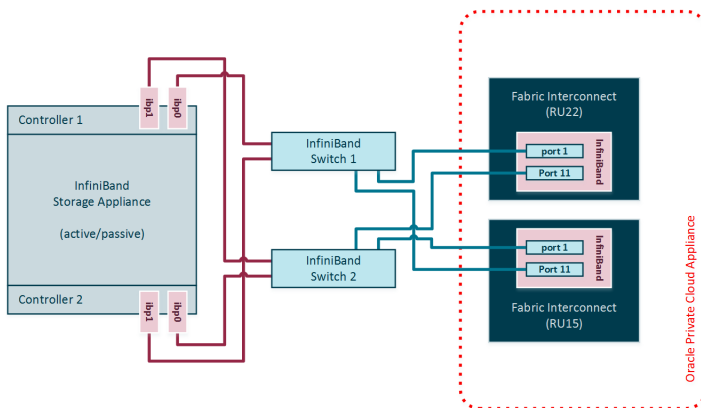
If the ZFS Storage Appliance is in active/active configuration, the cable connections must be doubled to provide redundancy between both ZFS controllers and the pair of InfiniBand switches. Because both controller heads are in active mode, they both require redundant connections to the two switches, meaning 4 cables per controller.

The different supported cabling layouts are illustrated in the figures below.

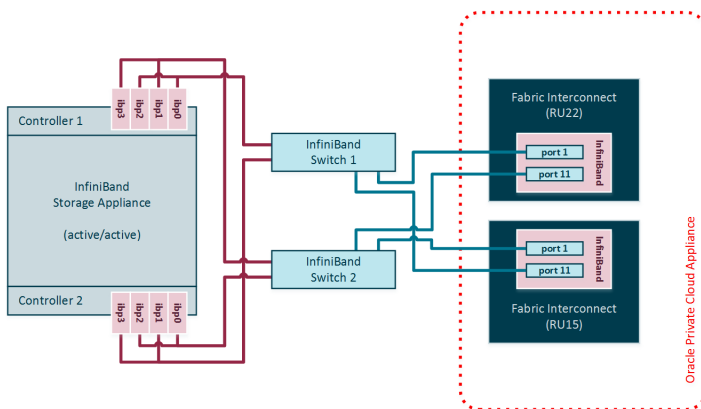
**Figure 8.6 IPoIB Storage Directly Attached**



**Figure 8.7 IPoIB Storage Connected to Switches in Active/Passive Configuration**



**Figure 8.8 IPoIB Storage Connected to Switches in Active/Active Configuration**



### 8.2.3.2 IP Address Allocation

The following IP address blocks have been reserved for use by a ZFS Storage Appliance external to the Oracle Private Cloud Appliance rack:

- 192.168.40.242
- 192.168.40.243
- 192.168.40.244



- 192.168.40.245
- 192.168.40.246
- 192.168.40.247
- 192.168.40.248
- 192.168.40.249

### 8.2.3.3 Configuring the ZFS Storage Appliance

This section describes the configuration steps that you must perform on the ZFS Storage Appliance to use IPoIB in conjunction with the Oracle Private Cloud Appliance. The description provided here assumes a typical configuration with an active/passive management cluster, using iSCSI to serve LUNs to compute nodes or virtual machines as physical disks. The ZFS Storage Appliance supports a standard NFS configuration as well.



**Note**

If you install extra expansion cards in the ZFS Storage Appliance controllers, and add a pair of InfiniBand switches between the external storage and the Fabric Interconnects inside the appliance rack, you can set up the ZFS Storage Appliance management cluster in active/active configuration. Notes in the procedure below highlight the required additional configuration.

#### Creating a Basic IPoIB Configuration on a ZFS Storage Appliance

##### 1. Create Datalinks

Log on to management user interface on the controller head you intend to use as the master. Go to **Configuration** and then to **Network**. If you have cabled correctly, two active devices are listed that map onto the cabled IB ports.

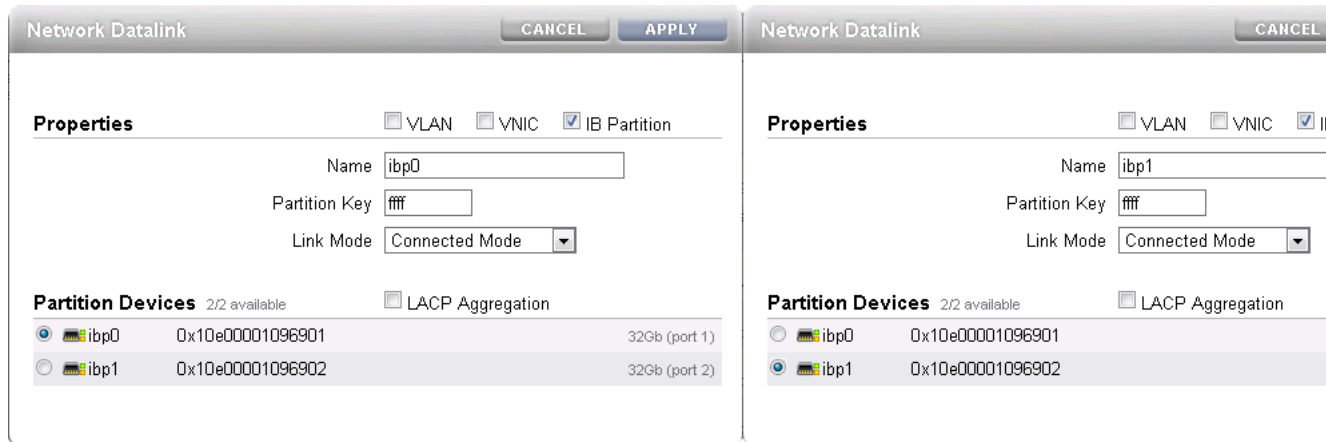


**Note**

In **active/active configuration** both controller heads have 4 active devices: typically `ibp0`, `ibp1`, `ibp2` and `ibp3`.

Drag each of these across to the Datalink menu to create a new datalink for each device. Edit each of these datalinks to provide a datalink name and partition key.


**Figure 8.9 Datalink Configuration**



## 2. Create Interfaces

Drag each configured datalink across into the Interface menu to create an interface for each datalink that you have defined. Edit each interface to provide a value for the Name field that makes it easy to identify the interface. Add the netmask 0.0.0.0/8 to each interface to prevent the system from probing and testing the routing details behind the network connection. Leave the IP MultiPathing Group unchecked. Do not configure an IP address at this stage.

**Figure 8.10 Interface Configuration**



## 3. Configure Multipathing

Click the Plus icon in the Interface menu to create an additional interface. This additional interface enables IPMP across the two interfaces you created in the previous step. Enter a value in the Name field that makes it easy to identify the IPMP interface. Assign an IP address from the list of reserved IPs in the Oracle Private Cloud Appliance storage network.

Near the bottom of the window, select the IP MultiPathing Group check box. Select the two previously configured interfaces, mark the first one as Active and the second one as Standby. The selected IP address will be configured on the active interface until a failover to the standby interface occurs.



### Note

In **active/active configuration** you create 2 IPMP interfaces: one across interfaces `ibp0` and `ibp2`, and the other across `ibp1` and `ibp3`. This ensures redundancy across the two PCIe InfiniBand expansion cards. Assign an IP

address to each IPMP interface. The configuration is replicated to the other storage controller.

**Figure 8.11 Multipathing (IPMP) Configuration**

Network Interface
CANCEL APPLY

---

**Properties**

Name

Enable Interface

Allow Administration

**Use IPv4 Protocol**

---

Configure with

Address/Mask   + -

Directly Reachable Network(s) +

**Use IPv6 Protocol**

---

**Interface Status**

**Interfaces** 3/9 available  IP MultiPathing Group


<input checked="" type="checkbox"/>	ibp0	IPv4 static, 0.0.0.0/8, via pffff_ibp0	Active
<input checked="" type="checkbox"/>	ibp1	IPv4 static, 0.0.0.0/8, via pffff_ibp1	Standby

#### 4. Apply Connectivity Configuration

Click the Apply button to commit the connectivity configuration you performed up to this point. When you select any device, datalink or interface associated with your configuration, all related items are

highlighted in the user interface, as shown in the screenshot below. All other items unrelated to the configuration in this procedure have been blurred out in the image.

Figure 8.12 Final Connectivity Configuration



## 5. Verify Management Cluster Configuration

The external storage setup proposed in this chapter has an active/passive cluster configuration. Consequently all storage resources for Oracle Private Cloud Appliance have to be available at one particular IP address at any time. The IP address and the storage resources are owned by the active head in the management cluster, and the standby head is ready to take over the IP address and expose the same storage resources if the active head should fail.

On the active storage controller, go to **Configuration** and then to **Cluster**. The active/passive cluster configuration is illustrated in the screenshot below. All resources involved in the Oracle Private Cloud Appliance external storage setup are owned by the active head: the storage appliance management

interface, the IPMP interface configured specifically for Oracle Private Cloud Appliance external storage connectivity, and the storage pool where the LUNs will be created.

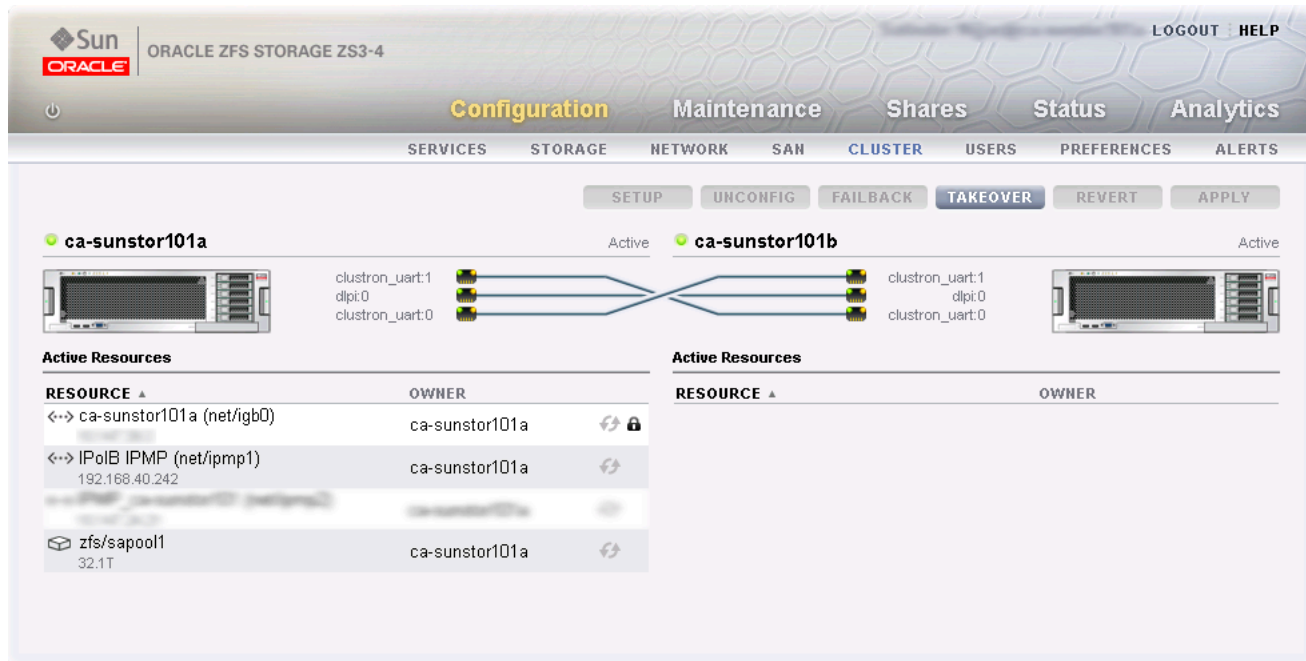
In this setup the entire configuration is applied on the active controller and automatically replicated on the standby controller. Since all resources are owned by the active head, there is no need to log on to the standby head and make any configuration changes.



**Note**

In **active/active configuration** both storage controllers take ownership of the storage resources in their storage pool. Each storage controller presents its storage resources to Oracle Private Cloud Appliance using separate redundant network paths. The active path at any given moment is determined through multipathing.

**Figure 8.13 Management Cluster Configuration**

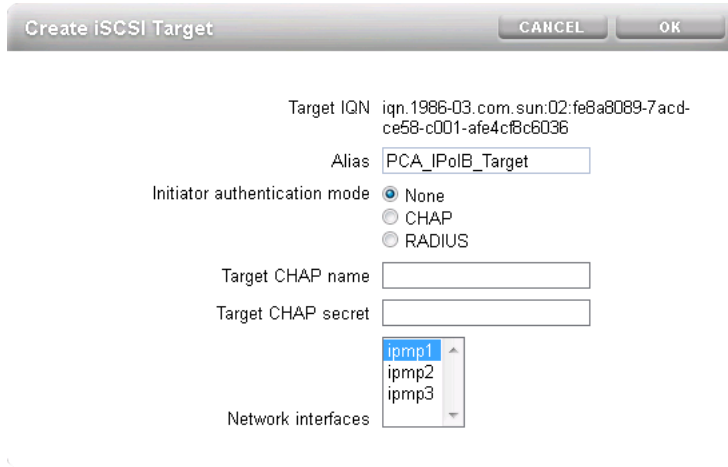


**6. Configure iSCSI Targets**

On the active controller, go to **Configuration, SAN** and then to **iSCSI**. Click the Plus icon in the Targets menu to create a new iSCSI target. Enter a value in the Alias field that makes it easy to


identify the iSCSI target. Select the IPMP interface you configured specifically for Oracle Private Cloud Appliance external storage connectivity. Click OK to create the new target.

**Figure 8.14 Create iSCSI Target**



Drag the iSCSI target that you have created into the Target Group area to create a target group. Edit the target group, give it an appropriate name, and make sure the IQN of the correct iSCSI target is selected. Click OK to save your changes.

**Figure 8.15 Create iSCSI Target Group**



**Note**

In **active/active configuration** you must create an iSCSI target and target group on each storage controller. In the target configuration select the two IPMP interfaces you created.

The iSCSI initiators and initiator group, which are configured in the steps below, are automatically replicated to the second storage controller.

**7. Configure iSCSI Initiators**

First, you need to retrieve the IQN of each compute node you wish to grant access to the external storage. Log on to the master management node of your Oracle Private Cloud Appliance and proceed as follows:

- a. Using the CLI, list all compute nodes.

```
[root@ovcamn05r1 ~]# pca-admin list compute-node
```

Compute_Node	IP_Address	Provisioning_Status	ILOM_MAC	Provisioning_State
ovcacn09r1	192.168.4.7	RUNNING	00:10:e0:3f:82:75	running
ovcacn13r1	192.168.4.11	RUNNING	00:10:e0:3f:87:73	running
ovcacn26r1	192.168.4.13	RUNNING	00:10:e0:3e:46:db	running
ovcacn12r1	192.168.4.10	RUNNING	00:10:e0:3f:8a:c7	running
ovcacn08r1	192.168.4.6	RUNNING	00:10:e0:3f:84:df	running
ovcacn27r1	192.168.4.14	RUNNING	00:10:e0:3f:9f:13	running
ovcacn07r1	192.168.4.5	RUNNING	00:10:e0:3f:75:73	running
ovcacn11r1	192.168.4.9	RUNNING	00:10:e0:3f:83:23	running
ovcacn10r1	192.168.4.8	RUNNING	00:10:e0:3f:89:83	running
ovcacn14r1	192.168.4.12	RUNNING	00:10:e0:3f:8b:5d	running

```
-----
10 rows displayed
Status: Success
```

- b. SSH into each compute node and display the contents of the file `initiatorname.iscsi`.

```
[root@ovcamn05r1 ~]# ssh ovcacn07r1
root@ovcacn07r1's password:
[root@ovcacn07r1 ~]# cat /etc/iscsi/initiatorname.iscsi
InitiatorName=iqn.1988-12.com.oracle:a72be49151

[root@ovcacn07r1 ~]# exit
logout
Connection to ovcacn07r1 closed.

[root@ovcamn05r1 ~]#
```

**Tip**

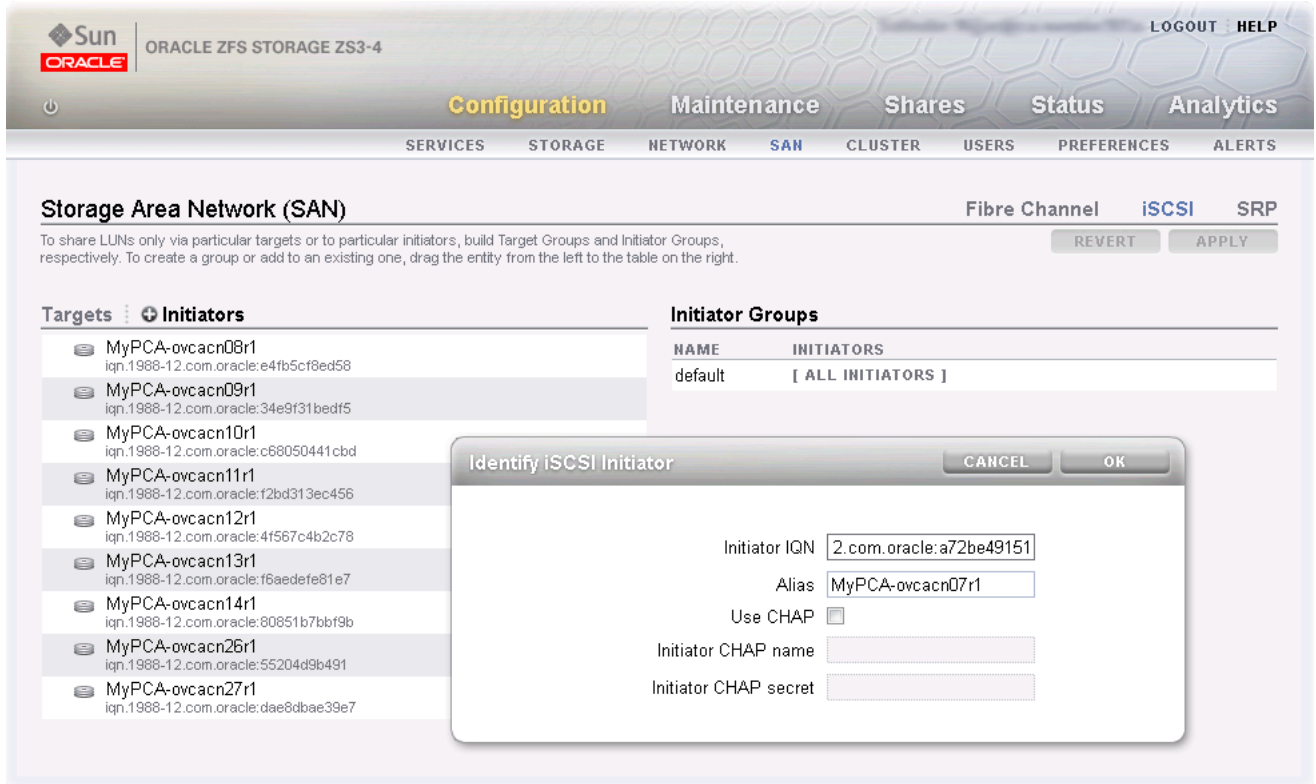
Using SSH to connect to each of the listed compute nodes is the fastest way to obtain the IQNs. However, they can also be copied from the Oracle VM Manager user interface, in the Storage Initiator perspective for each server.

- c. Copy the IQN of the compute node and use it to define a corresponding iSCSI initiator in the ZFS Storage Appliance user interface.

Click on the Initiators link to define the iSCSI initiators for the compute nodes that you wish to expose LUNs to. Click the Plus icon to identify a new iSCSI initiator. Enter the compute node Initiator IQN and

an Alias that makes it easy to identify the iSCSI initiator. Repeat this for each compute node so that all initiators appear in the list.

**Figure 8.16 Identify iSCSI Initiators**

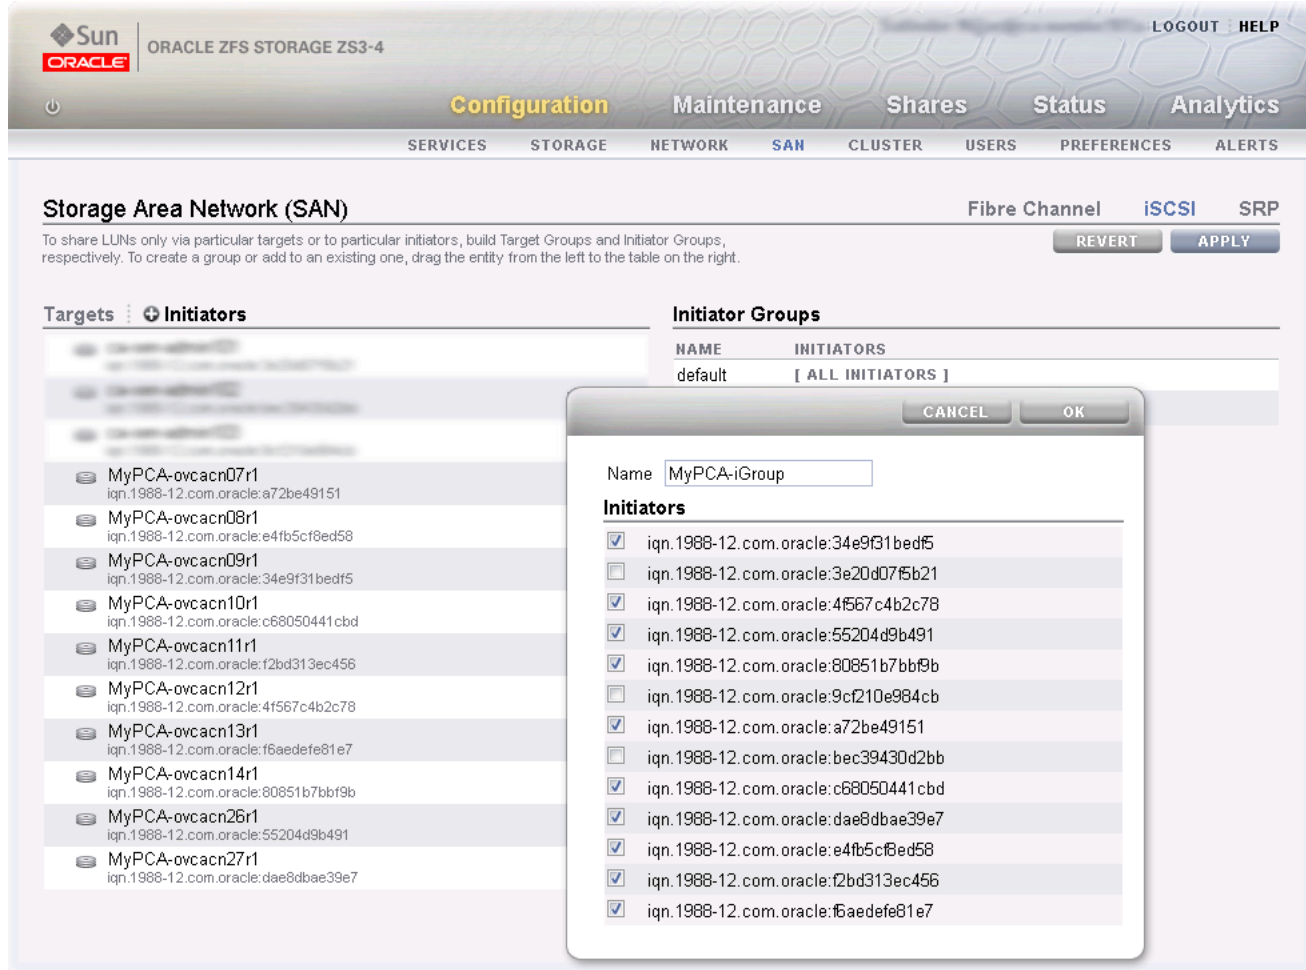


Drag the one of the iSCSI initiators that you have created into the Initiator Group area to create a new initiator group. Edit the initiator group, give it an appropriate name, and make sure that all IQNs of the



compute nodes that you want to make a member of this initiator group are selected. Click OK to save your changes.

**Figure 8.17 Create iSCSI Initiator Group**



**8. Apply iSCSI Target and Initiator Configuration**

Click the Apply button to commit the iSCSI target and initiator configuration you performed up to this point. When you select any initiator or initiator group associated with your configuration, all related

items are highlighted in the user interface, as shown in the screenshot below. Other items unrelated to the configuration in this procedure have been blurred out in the image.

**Figure 8.18 Final iSCSI Configuration**

The screenshot shows the Oracle ZFS Storage ZS3-4 web interface. The top navigation bar includes 'Configuration', 'Maintenance', 'Shares', 'Status', and 'Analytics'. Below this, there are tabs for 'SERVICES', 'STORAGE', 'NETWORK', 'SAN', 'CLUSTER', 'USERS', 'PREFERENCES', and 'ALERTS'. The 'SAN' tab is selected, and the 'Storage Area Network (SAN)' section is active. The 'iSCSI' sub-tab is selected, and the 'MyPCA-iGr...' initiator group is highlighted. The 'Targets' and 'Initiators' sections are visible on the left, and the 'Initiator Groups' table is on the right.

**Storage Area Network (SAN)** Fibre Channel **iSCSI** SRP

To share LUNs only via particular targets or to particular initiators, build Target Groups and Initiator Groups, respectively. To create a group or add to an existing one, drag the entity from the left to the table on the right. REVERT APPLY


NAME	INITIATORS
default	[ ALL INITIATORS ]
MyPCA-iGr...	iqn.1988-12.com.oracle:34e9f31bedf5 iqn.1988-12.com.oracle:4f567c4b2c78 iqn.1988-12.com.oracle:55204d9b491 iqn.1988-12.com.oracle:80851b7bbf9b iqn.1988-12.com.oracle:a72be49151 iqn.1988-12.com.oracle:c68050441cbd iqn.1988-12.com.oracle:dae8dbae39e7 iqn.1988-12.com.oracle:e4fb5cf8ed58 iqn.1988-12.com.oracle:f2bd313ec456 iqn.1988-12.com.oracle:f6aedefe81e7

## 9. Create a LUN

For easy management it is good practice to organize storage resources in separate projects. Create a project first for your Oracle Private Cloud Appliance external storage, and then add LUNs as you need them.

On the active storage controller, go to **Shares**. On the left hand side, click the Plus icon to create a new Project. In the Create Project dialog box, enter a Name to make it easy to identify the project. Click Apply to save the new project.

**Figure 8.19 Create ZFS Storage Project**



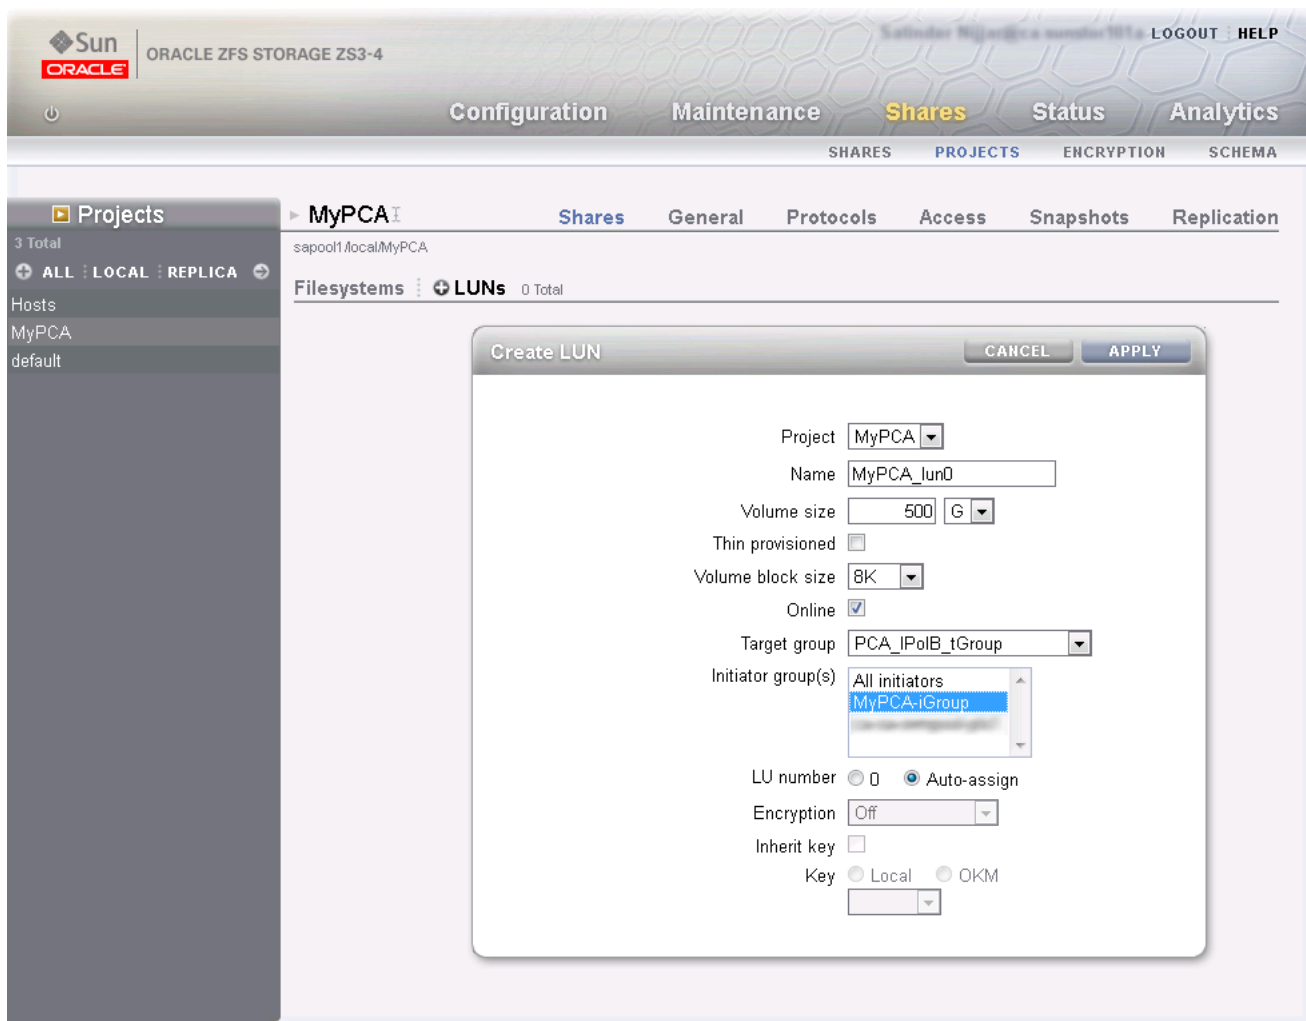
The screenshot shows a 'Create Project' dialog box. At the top, there is a title bar with the text 'Create Project' and two buttons: 'CANCEL' and 'APPLY'. Below the title bar, the form contains the following fields:

- 'Name' text box containing the text 'MyPCA'.
- 'Encryption' dropdown menu with 'Off' selected.
- 'Key' section with two radio buttons: 'Local' (selected) and 'OKM'.
- An empty dropdown menu below the radio buttons.

In the navigation pane on the left hand side, select the project you just created. In the Shares window, first select LUNs, and then click the Plus icon to create a new LUN as part of this project.

Fill out the LUN properties, select the iSCSI target group and initiator group you created earlier, and enter a Name that makes it easy to identify the LUN. Click Apply to add the LUN to the selected project.

Figure 8.20 Create LUN



**Note**

In **active/active configuration** each storage controller owns a storage pool. You must create a storage project for use with Oracle Private Cloud Appliance on each storage controller. Associate the LUNs you create with the iSCSI initiator group containing the Oracle Private Cloud Appliance compute nodes, and with the iSCSI target group of the selected storage controller. As a result, the LUNs in question are owned by that storage controller. For optimal performance the ownership of storage resources should be balanced between both active storage controllers.

If you wish to access these LUNs as physical disks within Oracle VM Manager, you must configure Oracle VM Manager first. Refer to [iSCSI Configuration](#) for more information.

**8.2.3.4 Enabling External IPoIB Storage in Oracle VM Manager**

If you intend to use your ZFS appliance to provide storage for use directly by Oracle VM, to host repositories and virtual machines, you must configure the storage within Oracle VM Manager before you are able to use it. The configuration steps that you must perform depend on whether you have configured

iSCSI or NFS on your ZFS Storage Appliance. This section provides a brief outline of the steps that you must perform to configure Oracle VM Manager for each of these technologies. For more detailed information, you should refer to the Oracle VM documentation.

If you only intend to make this storage available to individual virtual machines and do not intend to use the storage for underlying Oracle VM infrastructure, you do not need to perform any of the steps documented in this section, but you will need to configure each virtual machine directly to access the storage either over NFS or iSCSI.



**Caution**

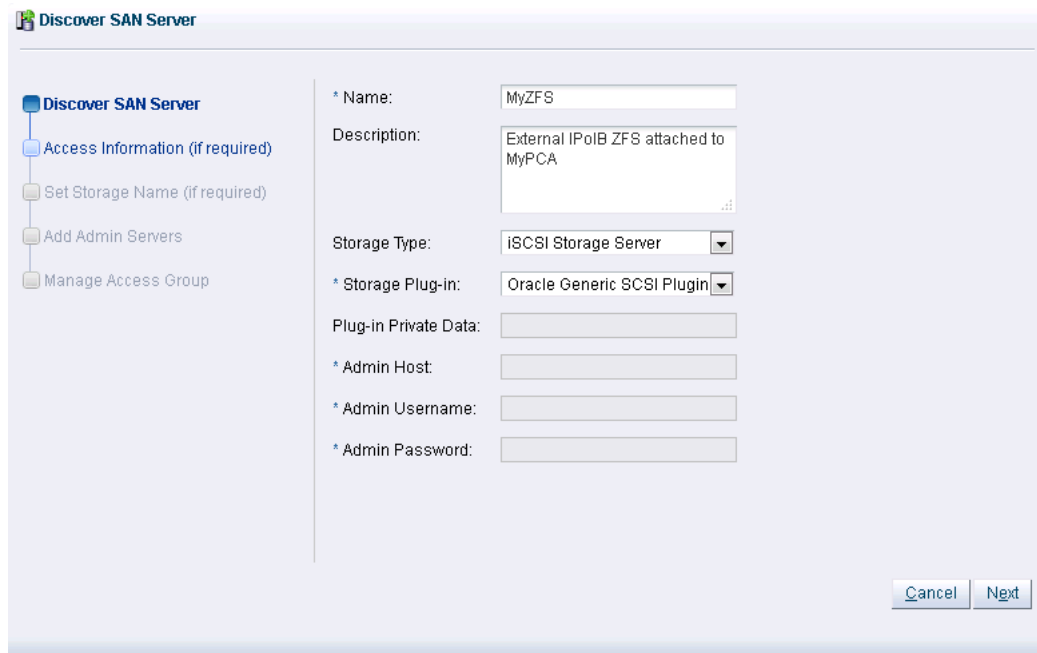
Reprovisioning restores a compute node to a clean state. If a compute node with active connections to external storage repositories is reprovisioned, the external storage connections need to be configured again after reprovisioning.

**ISCSI Configuration**

The following configuration steps should be performed in Oracle VM Manager if you have configured your storage appliance for iSCSI. The process to add a SAN Server in Oracle VM Manager is clearly documented in the *Oracle VM Manager User's Guide*. Refer to the section entitled [Discover SAN Server](#).


1. Log into Oracle VM Manager on the Oracle Private Cloud Appliance
2. Select the [Storage](#) tab to configure your storage
3. Click on the [Discover SAN Server](#) icon to load the wizard
4. Enter the DNS name of the storage appliance in the Name field. In the Storage Type field, use the drop-down selector to select [iSCSI Storage Server](#). In the Storage Plug-in field, you must select the [Oracle Generic SCSI plugin](#). Note that alternate storage plugins are not supported in this configuration. Click **Next**.

**Figure 8.21 Discover the SAN Server**



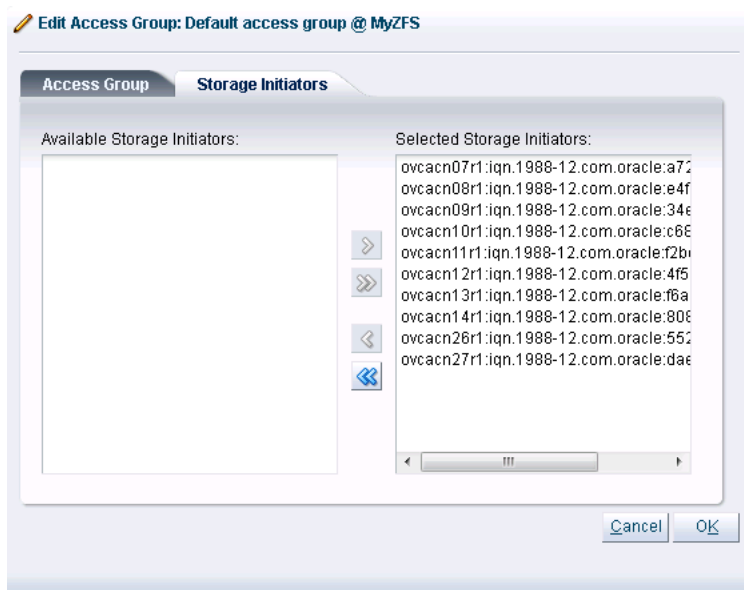
5. In the Access Information dialog, click on the icon that allows you to add a new Access Host. This opens the Create Access Host dialog. Enter the IP address that you configured for the IPMP interface of the storage appliance, for example 192.168.40.242. If you have configured your storage appliance with CHAP access, you must also enter the CHAP username and password here. Click **OK** to close the Create Access Host dialog. Click **Next**.

**Figure 8.22 Create Access Host**



6. In the Add Admin Servers dialog, select all of the servers in the Available Servers frame and move them to the Selected Servers frame. Click **Next**.
7. In the Manage Access Group dialog, click on **Default Access Group** and then click on the Edit icon. The Edit Access Group dialog is opened. Click on the **Storage Initiators** tab. Select the IQN name from all of the initiators that you have configured on the ZFS Storage Appliance and move them into the Selected Storage Initiators pane. Usually it is acceptable to move all of the IQNs across. Click **OK** to save the changes.

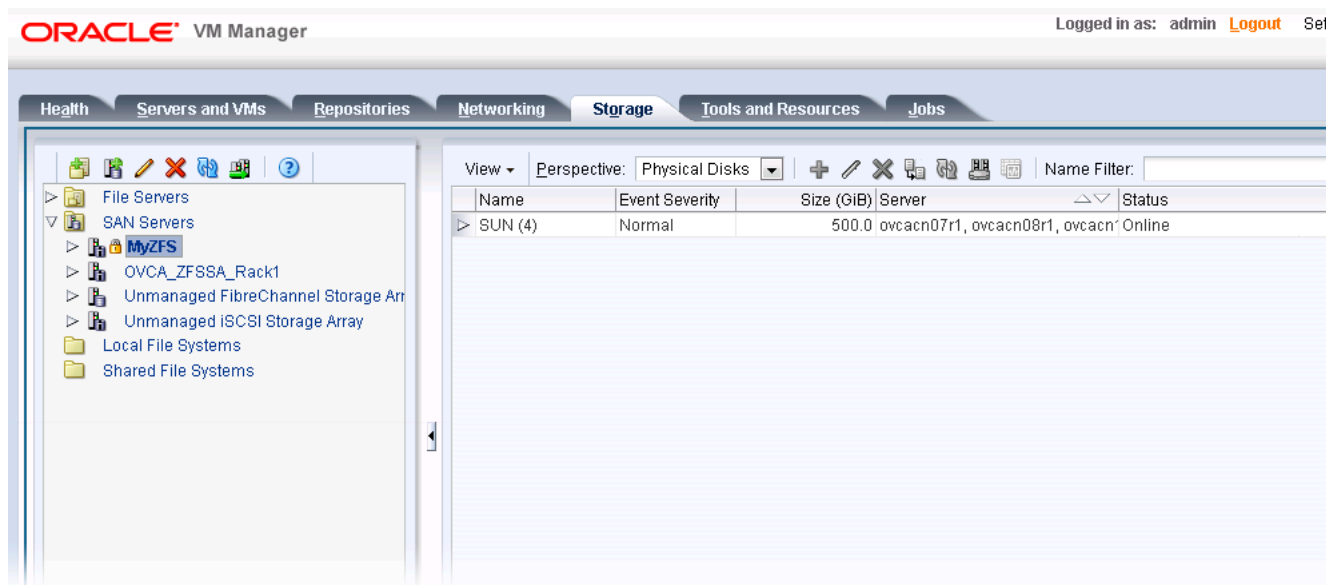
**Figure 8.23 Edit Access Group**



- Click the **Finish** button to exit the wizard and to save your changes.

The iSCSI server appears in the SAN Servers tree in the navigation pane. After storage discovery the LUN appears in the Physical Disks perspective.

**Figure 8.24 New Physical Disk Available**



- If the ZFS Storage Appliance is running in active/active configuration, the two exposed iSCSI targets are detected as separate SAN servers by Oracle VM Manager. You must execute the Discover SAN Server procedure separately for each access host IP address, so that the storage resources owned by each of the two respective controllers are available for use with Oracle Private Cloud Appliance.

## NFS Configuration

The following configuration steps should be performed in Oracle VM Manager if you have configured your storage appliance for NFS. The process to add a File Server in Oracle VM Manager is clearly documented in the *Oracle VM Manager User's Guide*. Refer to the section entitled [Discover File Server](#).

- Log into Oracle VM Manager on the Oracle Private Cloud Appliance
- Select the [Storage](#) tab to configure your storage
- Click on the [Discover File Server](#) icon to load the wizard
- Enter all required information into the wizard. Use the IP addresses that you configured for the device as the Access Host IP, for example 192.168.40.242.
- Select all of the compute nodes that should be designated as Admin Servers.
- Select two or three compute nodes that should be used as Refresh Servers.
- Select the file systems that you would like to use.
- Click the [Finish](#) button.
- The file server appears in the File Servers tree in the navigation pane.

10. If the ZFS Storage Appliance is running in active/active configuration, you must execute the Discover File Server procedure separately for each access host IP address, so that the storage resources owned by each of the two respective controllers are available for use with Oracle Private Cloud Appliance.



# Appendix A Default IP Addresses

## Table of Contents

A.1 Default Management IP Addresses .....	103
A.2 Default Storage Network IP Addresses .....	105
A.3 Default InfiniBand IP Addresses .....	105

This appendix lists the default IP addresses assigned to the Oracle Private Cloud Appliance components during manufacturing.

## A.1 Default Management IP Addresses

The management IP address represents a component's connection to the internal administration network.



### Caution

For hardware management, Oracle Private Cloud Appliance uses a network internal to the system. It is not recommended to connect the management ports or the internal administration network switches to the data center network infrastructure.

The tables in this section list the default management IP addresses assigned to servers and other hardware components in an Oracle Private Cloud Appliance base configuration rack. [Table A.1](#) applies to systems with an Ethernet-based network architecture; [Table A.2](#) applies to systems with an InfiniBand-based network architecture.

**Table A.1 Default Management IP Addresses for Oracle Private Cloud Appliance with Ethernet-based Network Architecture**

Rack Unit	Rack Component	Management IP Address Assigned During Manufacturing
23	Spine Cisco Nexus 9336C-FX2 Switch	192.168.4.205
22	Spine Cisco Nexus 9336C-FX2 Switch	192.168.4.204
21	Cisco Nexus 9348GC-FXP Switch	192.168.4.230
16	Leaf/Data Cisco Nexus 9336C-FX2 Switch	192.168.4.203
15	Leaf/Data Cisco Nexus 9336C-FX2 Switch	192.168.4.202
6	Management Node	192.168.4.4  management VIP: 192.168.4.216  ILOM: 192.168.4.104
5	Management Node	192.168.4.3  management VIP: 192.168.4.216  ILOM: 192.168.4.103
3-4	Oracle ZFS Storage Appliance ZS7-2 Controller Server (2 rack units)	192.168.4.2

Default Management IP Addresses

Rack Unit	Rack Component	Management IP Address Assigned During Manufacturing
		storage VIP: 192.168.4.100
		ILOM: 192.168.4.102
1-2	Oracle ZFS Storage Appliance ZS7-2 Controller Server (2 rack units)	192.168.4.1
		storage VIP: 192.168.4.100
		ILOM: 192.168.4.101
(none)	PDU-B	192.168.4.211
(none)	PDU-A	192.168.4.210

**Table A.2 Default Management IP Addresses for Oracle Private Cloud Appliance with InfiniBand-based Network Architecture**

Rack Unit	Rack Component	Management IP Address Assigned During Manufacturing
22-25	Oracle Fabric Interconnect F1-15 (4 rack units)	192.168.4.205
21	Oracle Switch ES1-24 (A and B)	192.168.4.200
		192.168.4.201
20	NM2-36P Sun Datacenter InfiniBand Expansion Switch	192.168.4.203
19	NM2-36P Sun Datacenter InfiniBand Expansion Switch	192.168.4.202
15-18	Oracle Fabric Interconnect F1-15 (4 rack units)	192.168.4.204
6	Management Node	192.168.4.4
		ILOM: 192.168.4.104
5	Management Node	192.168.4.3
		ILOM: 192.168.4.103
2	Oracle ZFS Storage Appliance Controller Server	192.168.4.2
		storage VIP: 192.168.4.100
		ILOM: 192.168.4.102
1	Oracle ZFS Storage Appliance Controller Server	192.168.4.1
		storage VIP: 192.168.4.100
		ILOM: 192.168.4.101
(none)	PDU-B	192.168.4.211
(none)	PDU-A	192.168.4.210

Compute nodes are assigned an IP address in the internal administration network during the provisioning process. The system IP address is DHCP-based; the ILOM is assigned the system IP plus 100. For

example: if a compute node receives IP `192.168.4.17`, then its ILOM has IP `192.168.4.117`. Once assigned to a host, these IP addresses are stored and persisted in the DHCP database.

## A.2 Default Storage Network IP Addresses



### Caution

Oracle Private Cloud Appliance configurations with an InfiniBand-based network architecture use an InfiniBand storage network. Refer to [Section A.3, “Default InfiniBand IP Addresses”](#).



### Caution

The storage network is internal to the Oracle Private Cloud Appliance. You should not create additional connections to this network.

Storage traffic between appliance components uses a subnet that is separate from the administration network and virtual machine connectivity. [Table A.3](#) lists the default storage network IP addresses assigned to hardware components in your Oracle Private Cloud Appliance.

**Table A.3 Default Storage Network IP Addresses for Oracle Private Cloud Appliance**

Rack Unit	Rack Component	Storage IP Address Assigned During Manufacturing
6	Management Node	<code>192.168.41.4</code>
5	Management Node	<code>192.168.41.3</code>
3-4	Oracle ZFS Storage Appliance ZS7-2 Controller Server (2 rack units)	<code>192.168.40.1</code> (shared)
1-2	Oracle ZFS Storage Appliance ZS7-2 Controller Server (2 rack units)	<code>192.168.40.1</code> (shared)

## A.3 Default InfiniBand IP Addresses



### Caution

The InfiniBand network is specific to Oracle Private Cloud Appliance configurations with an InfiniBand-based network architecture.



### Caution

The InfiniBand network is internal to the Oracle Private Cloud Appliance. You should not create additional connections to this network.

In systems built around the physical connectivity of the NM2-36P Sun Datacenter InfiniBand Expansion Switch and the Oracle Fabric Interconnect F1-15, InfiniBand interfaces are the default channel of communication among compute nodes and storage controller servers. If you are connecting the Oracle Private Cloud Appliance to existing or additional storage hardware on the same InfiniBand fabric, the interface enables communication between the Oracle Private Cloud Appliance and the storage devices.

[Table A.4](#) lists the default InfiniBand IP addresses assigned to hardware components in your Oracle Private Cloud Appliance.

Note that some IP addresses within this subnetwork are reserved for the addition of external InfiniBand storage appliances. These IP addresses are described in [Section 8.2.3.1, “Connecting InfiniBand Storage Hardware”](#) in the section [Section 8.2.3, “Adding External InfiniBand Storage”](#).

**Table A.4 Default Storage Network IP Addresses for Oracle Private Cloud Appliance**

<b>Rack Unit</b>	<b>Rack Component</b>	<b>InfiniBand IP Address Assigned During Manufacturing</b>
6	Management Node	192.168.41.4
5	Management Node	192.168.41.3
2	Oracle ZFS Storage Appliance Controller Server	192.168.40.1 (shared)
1	Oracle ZFS Storage Appliance Controller Server	192.168.40.1 (shared)

---

# Appendix B Cabling Reference

## Table of Contents

B.1 Cisco Nexus 9348GC-FXP Switch Management Ethernet Connections .....	107
B.2 Cisco Nexus 9336C-FX2 Switch Data Network Connections .....	108
B.3 Data and Spine Switch Interconnects .....	110
B.3.1 Spine Switch to Data Switch Connections .....	111
B.3.2 Data Switch to Data Switch Connections .....	111
B.3.3 Spine Switch to Spine Switch Connections .....	111
B.4 Oracle Switch ES1-24 Ethernet Connections .....	111
B.5 Storage Head Cluster Interconnects for InfiniBand-based Systems .....	113
B.6 Connections from NM2-36P Sun Datacenter InfiniBand Expansion Switch to Oracle Fabric Interconnect F1-15 .....	114
B.7 Connections from NM2-36P Sun Datacenter InfiniBand Expansion Switches to Compute Nodes ....	114
B.7.1 NM2-36P (U19) InfiniBand Switch to Compute Nodes Connections .....	114
B.7.2 NM2-36P (U20) InfiniBand Switch to Compute Nodes Connections .....	115

This appendix provides a reference for the cabling between components of Oracle Private Cloud Appliance.

## B.1 Cisco Nexus 9348GC-FXP Switch Management Ethernet Connections

This section lists the cabling connections for the Cisco Nexus 9348GC-FXP Switch. They are Ethernet connections intended for the internal administration network.

**Table B.1 Cabling Connections for Cisco Nexus 9348GC-FXP Switch**

From Rack Unit	From Component and Port	To Cisco Nexus 9348GC-FXP Switch Rack Unit and Port
U42	compute node: NET0	U21: port 23
U41	compute node: NET0	U21: port 22
U40	compute node: NET0	U21: port 21
U39	compute node: NET0	U21: port 20
U38	compute node: NET0	U21: port 19
U37	compute node: NET0	U21: port 18
U36	compute node: NET0	U21: port 17
U35	compute node: NET0	U21: port 16
U34	compute node: NET0	U21: port 15
U33	compute node: NET0	U21: port 42
U32	compute node: NET0	U21: port 41
U31	compute node: NET0	U21: port 40
U30	compute node: NET0	U21: port 39
U29	compute node: NET0	U21: port 38

From Rack Unit	From Component and Port	To Cisco Nexus 9348GC-FXP Switch Rack Unit and Port
U28	compute node: NET0	U21: port 37
U27	compute node: NET0	U21: port 36
U26	compute node: NET0	U21: port 35
U23	Cisco Nexus 9336C-FX2 Switch: NET MGMT	U21: port 14
U22	Cisco Nexus 9336C-FX2 Switch: NET MGMT	U21: port 13
U16	Cisco Nexus 9336C-FX2 Switch: NET MGMT	U21: port 34
U15	Cisco Nexus 9336C-FX2 Switch: NET MGMT	U21: port 33
U14	compute node: NET0	U21: port 12
U13	compute node: NET0	U21: port 11
U12	compute node: NET0	U21: port 10
U11	compute node: NET0	U21: port 9
U10	compute node: NET0	U21: port 8
U09	compute node: NET0	U21: port 7
U08	compute node: NET0	U21: port 6
U07	compute node: NET0	U21: port 5
U06	management node: NET0	U21: port 32
U05	management node: NET0	U21: port 31
U03	ZFS storage appliance controller: NET0	U21: port 28
U03	ZFS storage appliance controller: PCIe 5/0	U21: port 29
U03	ZFS storage appliance controller: PCIe 5/1	U21: port 30
U01	ZFS storage appliance controller: NET0	U21: port 25
U01	ZFS storage appliance controller: PCIe 5/0	U21: port 26
U01	ZFS storage appliance controller: PCIe 5/1	U21: port 27
N/A	PDU-A	U21: port 24
N/A	PDU-B	U21: port 43
N/A	Oracle Support	U21: port 1

## B.2 Cisco Nexus 9336C-FX2 Switch Data Network Connections

This section lists the cabling connections for the Cisco Nexus 9336C-FX2 Switch. They are high speed Ethernet connections up to 100Gbit, intended for the internal data network.

**Table B.2 Cabling Connections for Cisco Nexus 9336C-FX2 Switch**

<b>From Rack Unit</b>	<b>From Component and Port</b>	<b>To Cisco Nexus 9336C-FX2 Switch Rack Unit and Port</b>
U42	compute node: CX-5 HCA port 1	U15: port 31
U42	compute node: CX-5 HCA port 2	U16: port 31
U41	compute node: CX-5 HCA port 1	U15: port 30
U41	compute node: CX-5 HCA port 2	U16: port 30
U40	compute node: CX-5 HCA port 1	U15: port 29
U40	compute node: CX-5 HCA port 2	U16: port 29
U39	compute node: CX-5 HCA port 1	U15: port 28
U39	compute node: CX-5 HCA port 2	U16: port 28
U38	compute node: CX-5 HCA port 1	U15: port 27
U38	compute node: CX-5 HCA port 2	U16: port 27
U37	compute node: CX-5 HCA port 1	U15: port 26
U37	compute node: CX-5 HCA port 2	U16: port 26
U36	compute node: CX-5 HCA port 1	U15: port 25
U36	compute node: CX-5 HCA port 2	U16: port 25
U35	compute node: CX-5 HCA port 1	U15: port 24
U35	compute node: CX-5 HCA port 2	U16: port 24
U34	compute node: CX-5 HCA port 1	U15: port 23
U34	compute node: CX-5 HCA port 2	U16: port 23
U33	compute node: CX-5 HCA port 1	U15: port 22
U33	compute node: CX-5 HCA port 2	U16: port 22
U32	compute node: CX-5 HCA port 1	U15: port 21
U32	compute node: CX-5 HCA port 2	U16: port 21
U31	compute node: CX-5 HCA port 1	U15: port 20
U31	compute node: CX-5 HCA port 2	U16: port 20
U30	compute node: CX-5 HCA port 1	U15: port 19
U30	compute node: CX-5 HCA port 2	U16: port 19
U29	compute node: CX-5 HCA port 1	U15: port 18
U29	compute node: CX-5 HCA port 2	U16: port 18
U28	compute node: CX-5 HCA port 1	U15: port 17
U28	compute node: CX-5 HCA port 2	U16: port 17
U27	compute node: CX-5 HCA port 1	U15: port 16
U27	compute node: CX-5 HCA port 2	U16: port 16
U26	compute node: CX-5 HCA port 1	U15: port 15
U26	compute node: CX-5 HCA port 2	U16: port 15
U14	compute node: CX-5 HCA port 1	U15: port 14
U14	compute node: CX-5 HCA port 2	U16: port 14

From Rack Unit	From Component and Port	To Cisco Nexus 9336C-FX2 Switch Rack Unit and Port
U13	compute node: CX-5 HCA port 1	U15: port 13
U13	compute node: CX-5 HCA port 2	U16: port 13
U12	compute node: CX-5 HCA port 1	U15: port 12
U12	compute node: CX-5 HCA port 2	U16: port 12
U11	compute node: CX-5 HCA port 1	U15: port 11
U11	compute node: CX-5 HCA port 2	U16: port 11
U10	compute node: CX-5 HCA port 1	U15: port 10
U10	compute node: CX-5 HCA port 2	U16: port 10
U09	compute node: CX-5 HCA port 1	U15: port 9
U09	compute node: CX-5 HCA port 2	U16: port 9
U08	compute node: CX-5 HCA port 1	U15: port 8
U08	compute node: CX-5 HCA port 2	U16: port 8
U07	compute node: CX-5 HCA port 1	U15: port 7
U07	compute node: CX-5 HCA port 2	U16: port 7
U06	management node: CX-5 HCA port 1	U15: port 6
U06	management node: CX-5 HCA port 2	U16: port 6
U05	management node: CX-5 HCA port 1	U15: port 5
U05	management node: CX-5 HCA port 2	U16: port 5
U04	Oracle ZFS Storage Appliance ZS7-2 controller: slot 4 port 1	U22: port 35
U04	Oracle ZFS Storage Appliance ZS7-2 controller: slot 4 port 2	U23: port 35
U03	Oracle ZFS Storage Appliance ZS7-2 controller: slot 8 port 1	U22: port 36
U03	Oracle ZFS Storage Appliance ZS7-2 controller: slot 8 port 2	U23: port 36
U02	Oracle ZFS Storage Appliance ZS7-2 controller: slot 4 port 1	U22: port 33
U02	Oracle ZFS Storage Appliance ZS7-2 controller: slot 4 port 2	U23: port 33
U01	Oracle ZFS Storage Appliance ZS7-2 controller: slot 8 port 1	U22: port 34
U01	Oracle ZFS Storage Appliance ZS7-2 controller: slot 8 port 2	U23: port 34

## B.3 Data and Spine Switch Interconnects

This section lists the cabling connections between the Cisco Nexus 9348GC-FXP Switches performing the role of either data switch or spine switch. They are 100GbE Ethernet connections intended for high bandwidth traffic between the data and spine switches.



### B.3.1 Spine Switch to Data Switch Connections

Table B.3 lists the cabling connections between the data switches and spine switches.

**Table B.3 Cabling Connections Between Spine and Data Switches**

From Rack Unit	From Component and Port	To Rack Unit	To Component and Port
U22	spine switch: port 23	U15	data switch: port 33
U22	spine switch: port 24	U15	data switch: port 34
U22	spine switch: port 25	U16	data switch: port 35
U22	spine switch: port 26	U16	data switch: port 36
U23	spine switch: port 23	U16	data switch: port 33
U23	spine switch: port 24	U16	data switch: port 34
U23	spine switch: port 25	U15	data switch: port 35
U23	spine switch: port 26	U15	data switch: port 36

### B.3.2 Data Switch to Data Switch Connections

Table B.4 lists the cabling connections between the data switches.

**Table B.4 Cabling Connections Between Data Switches**

From Rack Unit	From Component and Port	To Rack Unit	To Component and Port
U15	data switch: port 3	U16	data switch: port 3
U15	data switch: port 4	U16	data switch: port 4

### B.3.3 Spine Switch to Spine Switch Connections

Table B.5 lists the cabling connections between the spine switches.

**Table B.5 Cabling Connections Between Spine Switches**

From Rack Unit	From Component and Port	To Rack Unit	To Component and Port
U22	spine switch: port 28	U23	spine switch: port 28
U22	spine switch: port 29	U23	spine switch: port 29
U22	spine switch: port 30	U23	spine switch: port 30
U22	spine switch: port 31	U23	spine switch: port 31
U22	spine switch: port 32	U23	spine switch: port 32

## B.4 Oracle Switch ES1-24 Ethernet Connections

This section lists the cabling connections for Oracle Switch ES1-24.

**Table B.6 Cabling Connections for Oracle Switch ES1-24**

From	Port	To	Port	Color	Notes
N/A	SER MGT	U21-A	Serial Port	1.8m console cable	7061953 cable accessible for Service only
N/A	SER MGT	U21-B	Serial Port	1.8m console cable	7061953 cable accessible for Service only
U42	NET0	U21-A	19	Blue	
U41	NET0	U21-A	18	Blue	
U40	NET0	U21-A	17	Blue	
U39	NET0	U21-A	16	Blue	
U38	NET0	U21-A	15	Blue	
U37	NET0	U21-A	14	Blue	
U36	NET0	U21-A	13	Blue	
U35	NET0	U21-A	12	Blue	
U34	NET0	U21-A	11	Blue	
U33	NET0	U21-B	18	Blue	
U32	NET0	U21-B	17	Blue	
U31	NET0	U21-B	16	Blue	
U30	NET0	U21-B	15	Blue	
U29	NET0	U21-B	14	Blue	
U28	NET0	U21-B	13	Blue	
U27	NET0	U21-B	12	Blue	
U26	NET0	U21-B	11	Blue	
U22		U21-A	9	Black	Oracle Fabric Interconnect F1-15
U20		U21-B	10	Black	NM2-36P Sun Datacenter InfiniBand Expansion Switch
U21-A	24	Customer network		Transceiver/ Black Cable	
U21-B	24	Customer network		Transceiver/ Black Cable	
U21-A	22	U21-B	22	SFP+	Switch-to-switch connection within rack
U19		U21-A	10	Black	NM2-36P Sun Datacenter InfiniBand Expansion Switch
U15		U21-B	9	Black	Oracle Fabric Interconnect F1-15
U14	NET0	U21-A	8	Blue	

From	Port	To	Port	Color	Notes
U13	NET0	U21-A	7	Blue	
U12	NET0	U21-A	6	Blue	
U11	NET0	U21-A	5	Blue	
U10	NET0	U21-A	4	Blue	
U9	NET0	U21-A	3	Blue	
U8	NET0	U21-A	2	Blue	
U7	NET0	U21-A	1	Blue	
U6	NET0	U21-B	8	Blue	
U5	NET0	U21-B	7	Blue	
U2	NET2	U21-B	6	Blue	
U2	NET1	U21-B	5	Blue	
U2	NET0	U21-B	4	Blue	
U1	NET2	U21-B	3	Blue	
U1	NET1	U21-B	2	Blue	
U1	NET0	U21-B (Left switch)	1	Blue	
PDU-A	NETMGT	U21-A	20	White	
PDU-B	NETMGT	U21-B	20	White	

## B.5 Storage Head Cluster Interconnects for InfiniBand-based Systems



### Note

This cabling information applies to the models of the Oracle ZFS Storage Appliance installed in Oracle Private Cloud Appliance systems with an InfiniBand-based network architecture.

This section contains cabling connections for the storage head cluster in Oracle ZFS Storage Appliance ZS5-ES and earlier models.

**Table B.7 Cabling Connections for Storage Head Cluster**

From	Port	To	Port	Length	Color	Part Number and Notes
U1	PCle-1 P0	U2	PCle-1 P2	2.13m	Yellow	530-4527-01
U1	PCle-1 P1	U2	PCle-1 P1	2.13m	Black	530-4436-01
U1	PCle-1 P2	U2	PCle-1 P0	2.13m	Green	530-4526-01
U2	PCle-2 P1	U3	IOM1 (top) SAS Port 1	2m	Black	2m SAS
U1	PCle-2 P1	U3	IOM1 SAS Port 0	2m	Black	2m SAS
U2	PCle-2 P0	U3	IOM0 (bottom) SAS Port 1	2m	Black	2m SAS
U1	PCle-2 P0	U3	IOM0 SAS Port 0	2m	Black	2m SAS

## B.6 Connections from NM2-36P Sun Datacenter InfiniBand Expansion Switch to Oracle Fabric Interconnect F1-15

This section contains in-rack cabling connections for the NM2-36P Sun Datacenter InfiniBand Expansion Switch to the Oracle Fabric Interconnect F1-15.

**Table B.8 Cabling Connections for NM2-36P IB Switch to Oracle Fabric Interconnect F1-15**

From	Name	Port	To	Name	Port	NM2 to F1 Interconnect
U19	NM2-36P IB	9B	U22	F1-15	8	Bottom NM2-36P IB to top F1-15 (Blue Velcro)
U19	NM2-36P IB	10B	U22	F1-15	9	Bottom NM2-36P IB to top F1-15 (Blue Velcro)
U19	NM2-36P IB	11B	U15	F1-15	10	Bottom NM2-36P IB to bottom F1-15 (Blue Velcro)
U19	NM2-36P IB	9A	U15	F1-15	8	Bottom NM2-36P IB to bottom F1-15 (Blue Velcro)
U19	NM2-36P IB	10A	U15	F1-15	9	Bottom NM2-36P IB to bottom F1-15 (Blue Velcro)
U19	NM2-36P IB	11A	U22	F1-15	10	Bottom NM2-36P IB to top F1-15 (Blue Velcro)
U20	NM2-36P IB	9B	U15	F1-15	18	Top NM2-36P IB to bottom F1-15 (Blue Velcro)
U20	NM2-36P IB	10B	U15	F1-15	19	Top NM2-36P IB to bottom F1-15 (Blue Velcro)
U20	NM2-36P IB	11B	U22	F1-15	20	Top NM2-36P IB to top F1-15 (Blue Velcro)
U20	NM2-36P IB	9A	U22	F1-15	18	Top NM2-36P IB to top F1-15 (Blue Velcro)
U20	NM2-36P IB	10A	U22	F1-15	19	Top NM2-36P IB to top F1-15 (Blue Velcro)
U20	NM2-36P IB	11A	U15	F1-15	20	Top NM2-36P IB to bottom F1-15 (Blue Velcro)

## B.7 Connections from NM2-36P Sun Datacenter InfiniBand Expansion Switches to Compute Nodes

This section contains in-rack cabling connections for the NM2-36P Sun Datacenter InfiniBand Expansion Switches to the compute nodes.

### B.7.1 NM2-36P (U19) InfiniBand Switch to Compute Nodes Connections

The following table contains cabling connections for the NM2-36P (U19) InfiniBand switch to the compute nodes.

**Table B.9 Cabling Connections for NM2-36P (U19) IB Switch to Compute Nodes**

From	Name	Port	To	Name	Port
U19	NM2-36P IB	0A	U42	Compute Node	PCIe-2 P2

From	Name	Port	To	Name	Port
U19	NM2-36P IB	0B	U41	Compute Node	PCIe-2 P2
U19	NM2-36P IB	1A	U40	Compute Node	PCIe-2 P2
U19	NM2-36P IB	1B	U39	Compute Node	PCIe-2 P2
U19	NM2-36P IB	2A	U38	Compute Node	PCIe-2 P2
U19	NM2-36P IB	2B	U37	Compute Node	PCIe-2 P2
U19	NM2-36P IB	3A	U36	Compute Node	PCIe-2 P2
U19	NM2-36P IB	3B	U35	Compute Node	PCIe-2 P2
U19	NM2-36P IB	4A	U34	Compute Node	PCIe-2 P2
U19	NM2-36P IB	4B	U33	Compute Node	PCIe-2 P2
U19	NM2-36P IB	5A	U32	Compute Node	PCIe-2 P2
U19	NM2-36P IB	5B	U31	Compute Node	PCIe-2 P2
U19	NM2-36P IB	6A	U30	Compute Node	PCIe-2 P2
U19	NM2-36P IB	6B	U29	Compute Node	PCIe-2 P2
U19	NM2-36P IB	7A	U28	Compute Node	PCIe-2 P1
U19	NM2-36P IB	7B	U27	Compute Node	PCIe-2 P1
U19	NM2-36P IB	8A	U26	Compute Node	PCIe-2 P1
U19	NM2-36P IB	12A	U14	Compute Node	PCIe-2 P1
U19	NM2-36P IB	12B	U13	Compute Node	PCIe-2 P1
U19	NM2-36P IB	13A	U12	Compute Node	PCIe-2 P1
U19	NM2-36P IB	13B	U11	Compute Node	PCIe-2 P1
U19	NM2-36P IB	14A	U10	Compute Node	PCIe-2 P1
U19	NM2-36P IB	14B	U9	Compute Node	PCIe-2 P1
U19	NM2-36P IB	15A	U8	Compute Node	PCIe-2 P1
U19	NM2-36P IB	15B	U7	Compute Node	PCIe-2 P1
U19	NM2-36P IB	16A	U6	Management Node	PCIe-2 P1
U19	NM2-36P IB	16B	U5	Management Node	PCIe-2 P1
U19	NM2-36P IB	17A	U2	Storage Appliance Controller	PCIe-0 P1
U19	NM2-36P IB	17B	U1	Storage Appliance Controller	PCIe-0 P1

## B.7.2 NM2-36P (U20) InfiniBand Switch to Compute Nodes Connections

The following table contains cabling connections for the NM2-36P (U20) InfiniBand switch to the compute nodes.

**Table B.10 Cabling Connections for NM2-36P (U20) IB Switch to Compute Nodes**

From	Name	Port	To	Name	Port
U20	NM2-36P IB	0A	U42	Compute Node	PCIe-2 P1
U20	NM2-36P IB	0B	U41	Compute Node	PCIe-2 P1

NM2-36P (U20) InfiniBand Switch to Compute Nodes Connections

<b>From</b>	<b>Name</b>	<b>Port</b>	<b>To</b>	<b>Name</b>	<b>Port</b>
U20	NM2-36P IB	1A	U40	Compute Node	PCIe-2 P1
U20	NM2-36P IB	1B	U39	Compute Node	PCIe-2 P1
U20	NM2-36P IB	2A	U38	Compute Node	PCIe-2 P1
U20	NM2-36P IB	2B	U37	Compute Node	PCIe-2 P1
U20	NM2-36P IB	3A	U36	Compute Node	PCIe-2 P1
U20	NM2-36P IB	3B	U35	Compute Node	PCIe-2 P1
U20	NM2-36P IB	4A	U34	Compute Node	PCIe-2 P1
U20	NM2-36P IB	4B	U33	Compute Node	PCIe-2 P1
U20	NM2-36P IB	5A	U32	Compute Node	PCIe-2 P1
U20	NM2-36P IB	5B	U31	Compute Node	PCIe-2 P1
U20	NM2-36P IB	6A	U30	Compute Node	PCIe-2 P1
U20	NM2-36P IB	6B	U29	Compute Node	PCIe-2 P1
U20	NM2-36P IB	7A	U28	Compute Node	PCIe-2 P2
U20	NM2-36P IB	7B	U27	Compute Node	PCIe-2 P2
U20	NM2-36P IB	8A	U26	Compute Node	PCIe-2 P2
U20	NM2-36P IB	12A	U14	Compute Node	PCIe-2 P2
U20	NM2-36P IB	12B	U13	Compute Node	PCIe-2 P2
U20	NM2-36P IB	13A	U12	Compute Node	PCIe-2 P2
U20	NM2-36P IB	13B	U11	Compute Node	PCIe-2 P2
U20	NM2-36P IB	14A	U10	Compute Node	PCIe-2 P2
U20	NM2-36P IB	14B	U9	Compute Node	PCIe-2 P2
U20	NM2-36P IB	15A	U8	Compute Node	PCIe-2 P2
U20	NM2-36P IB	15B	U7	Compute Node	PCIe-2 P2
U20	NM2-36P IB	16A	U6	Management Node	PCIe-2 P2
U20	NM2-36P IB	16B	U5	Management Node	PCIe-2 P2
U20	NM2-36P IB	17A	U2	Storage Appliance Controller	PCIe-0 P2
U20	NM2-36P IB	17B	U1	Storage Appliance Controller	PCIe-0 P2

---

## Appendix C Connectivity Diagrams

### Table of Contents

C.1 External Ethernet Connections .....	117
C.2 Virtual Networking .....	119
C.3 External Storage Connections .....	120
C.4 Custom and Host Networks .....	122

Oracle Private Cloud Appliance exists in two different types of network architecture. One is built around a physical InfiniBand fabric; the other relies on physical high speed Ethernet connectivity. While the two implementations offer practically the same functionality, there are visible hardware and configuration differences.

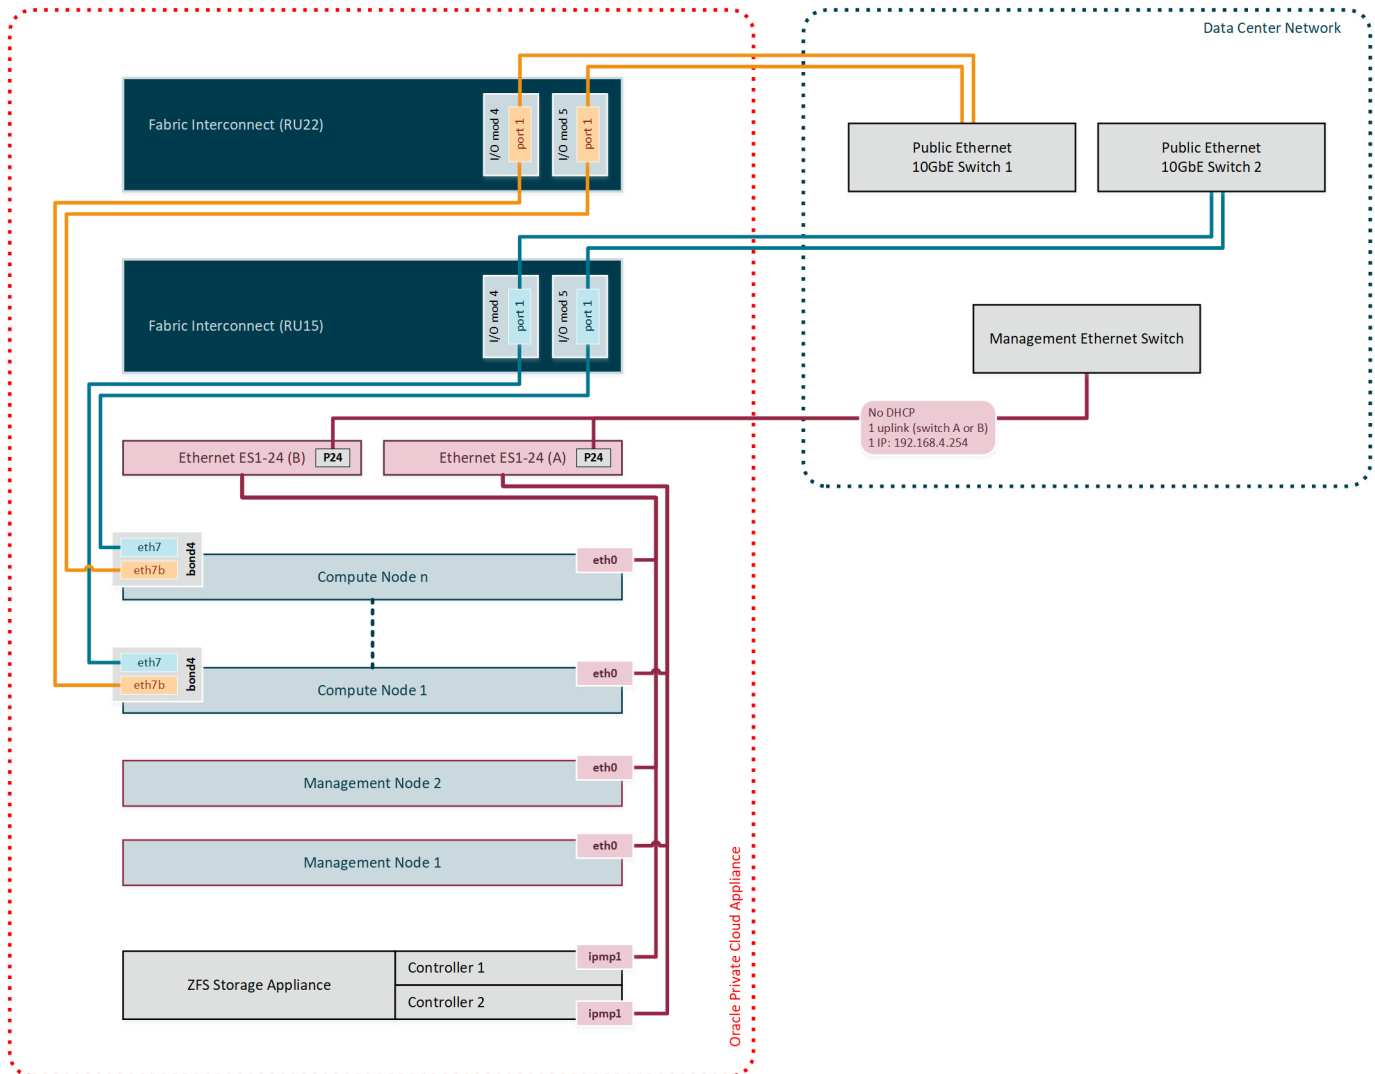
The network architecture differences also affect the options to connect external storage hardware. The Oracle Fabric Interconnect F1-15, the heart of the network configuration of InfiniBand-based systems, offers physical connectors to attach external InfiniBand and Fibre Channel storage devices.

The diagrams in this appendix apply to systems with an InfiniBand-based network architecture, controlled by means of a pair of Fabric Interconnects.

### C.1 External Ethernet Connections

This section shows the external Ethernet connectivity of an Oracle Private Cloud Appliance using the I/O modules of the Oracle Fabric Interconnect F1-15.

Figure C.1 External Ethernet Connectivity



## Diagram Notes

- The connections in red in the left half of the diagram correspond with the out-of-band appliance management network [192.168.4.0/24](#). These are physical Ethernet ports connected with Ethernet cables to the two internally inter-connected Oracle Switch ES1-24 switches. All appliance components and their ILOMs have a reserved IP address in this network. Only one address is available for customer use: [192.168.4.254](#).
- You may connect to this management Ethernet network by choosing one of these methods:
  - Connect a workstation with an Ethernet cable plugged directly into the available port 19 in an Oracle Switch ES1-24. Statically assign IP address [192.168.4.254](#).
  - Connect port 24 in **one** Oracle Switch ES1-24 – **never both** – to an Ethernet switch in your data center network. Use a workstation connected to the data center network and statically assign the IP address [192.168.4.254](#). Make sure that the management Ethernet switch used in this connection is configured to prevent DHCP leakage to the [192.168.4.0/24](#) subnet used by Oracle Private Cloud Appliance.

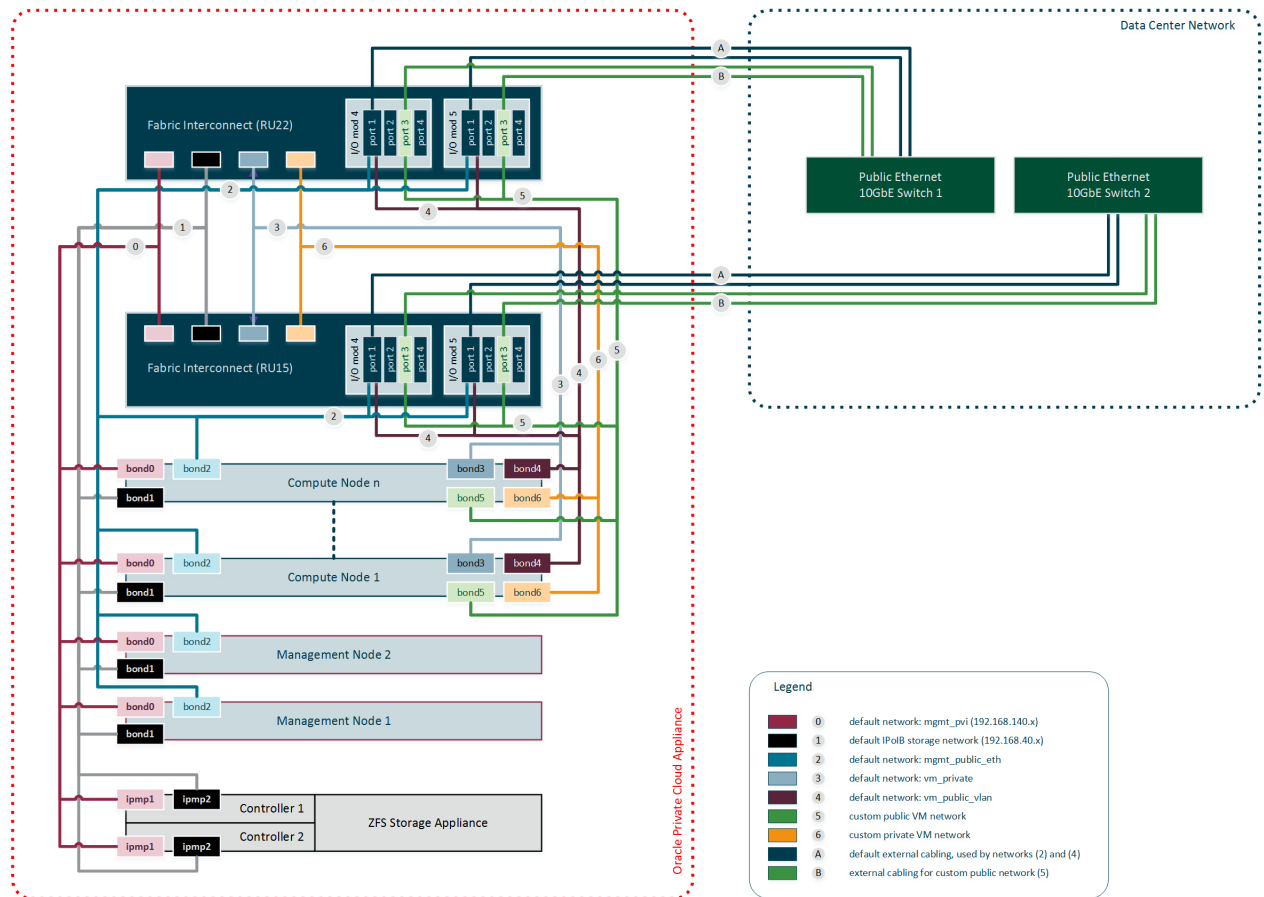


- The connections in gold and blue on the left hand side of the diagram correspond with the default network `vm_public_vlan`. This is the standard VLAN-enabled network for virtual machine traffic with public access. The compute node connections are bonds of two Ethernet vNICs that terminate on port 1 of the I/O modules 4 and 5 in both Fabric Interconnects. The underlying physical connection consists of redundant InfiniBand cabling. The physical 10GbE ports of the Fabric Interconnects are connected for redundancy to two next-level data center switches.
- The vNIC connections are distributed across the Fabric Interconnects and the I/O modules for maximum availability. If one of the redundant components fails at any level, be it an I/O module, an entire Fabric Interconnect or a next-level switch in the data center, the VMs always maintain an active external connection using either a blue or a red path. The uplinks to the data center switches are deliberately **not** cross-cabled to prevent that half of the compute nodes lose external connectivity if one of those switches fails.
- All custom public networks follow the same principle as the default `vm_public_vlan` shown in this diagram. However, for the custom networks additional bonded vNICs are created on the compute nodes, and additional I/O ports must be configured and cabled.

## C.2 Virtual Networking

This section shows how Oracle Private Cloud Appliance networking is virtualized and configured through the Fabric Interconnects.

Figure C.2 Virtual InfiniBand and Ethernet Networks With or Without Public Connectivity



## Diagram Notes

- All virtual networks in the diagram connect bonds, consisting of two virtual interfaces on the nodes, to the Fabric Interconnects. Private networks, which provide internal connectivity only, terminate on the Fabric Interconnects and are not routed externally. Public networks terminate on the I/O modules and are routed externally across the 10GbE ports, which are cabled in a redundant configuration to dedicated Ethernet switches in the data center network. Only I/O modules 4 and 5 are shown in this diagram, but each Oracle Fabric Interconnect F1-15 installed in an Oracle Private Cloud Appliance rack contains two more I/O modules in slots 10 and 11, with four 10GbE ports each, which can be used in the same way.

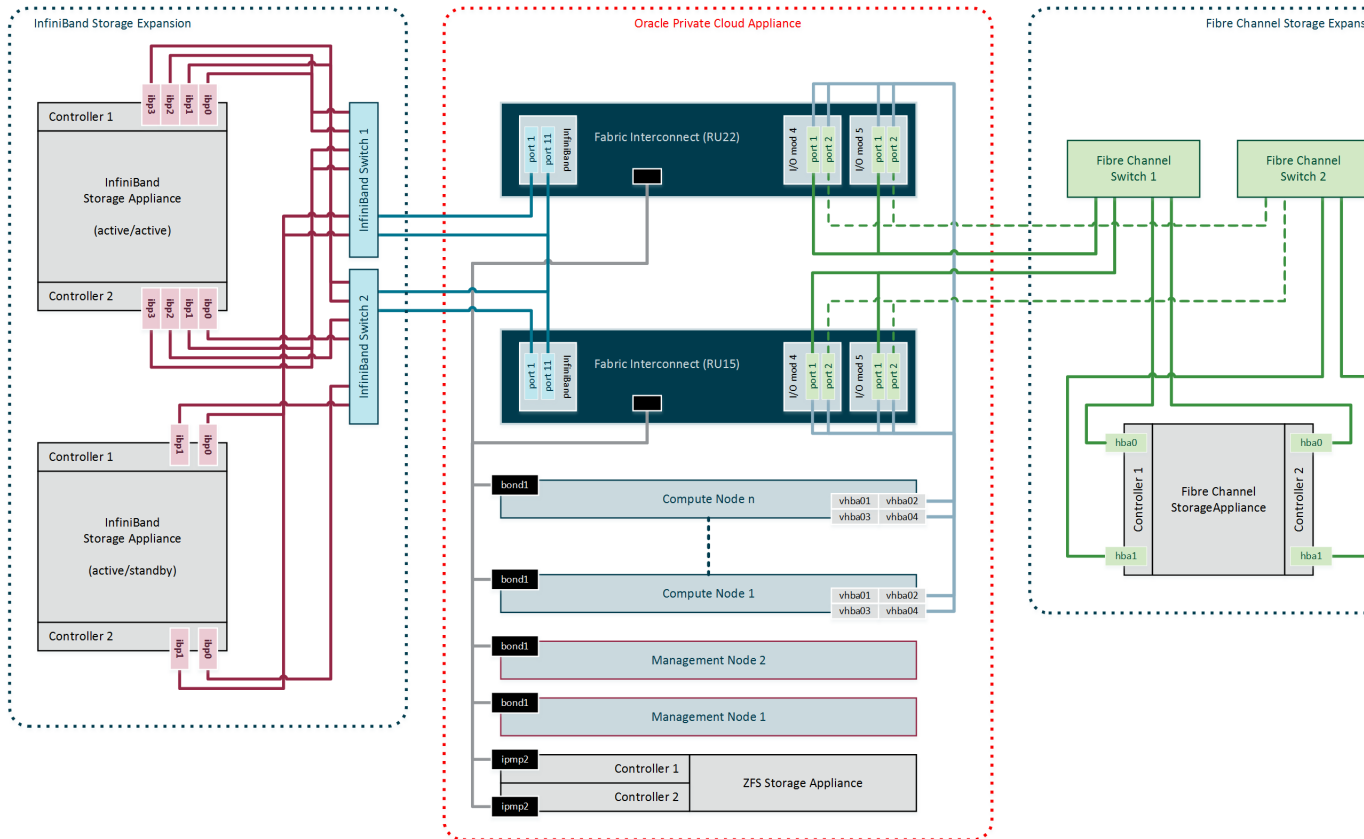
The underlying physical connection of all virtual networking consists of redundant InfiniBand cabling. All internal and external connectivity is managed by the Fabric Interconnects.

- The storage network is an IPoIB network that provides connectivity between the compute nodes, the management nodes and the internal ZFS storage appliance. All connected components use a network bond consisting of two virtual InfiniBand interfaces. The bonds have an IP address in the [192.168.40.0/24](#) subnet.
- Apart from the storage network, all virtual networks in the diagram are Ethernet networks. All connected components use a network bond consisting of two vNICs. The Ethernet networks intended for both private and public VM connectivity only include the compute nodes. All VM networks are VLAN-enabled.
- The virtual networks connecting the network bonds numbered 0-4 are configured by default on every Oracle Private Cloud Appliance. The network connecting all [bond6](#) interfaces is an example of a custom private VM network. The network connecting all [bond5](#) interfaces is an example of a custom public VM network. This type of network requires the configuration of additional I/O ports, which must also be cabled to the next-level data center switches for public access.

## C.3 External Storage Connections

This section shows how external storage is connected to the Fabric Interconnects in the Oracle Private Cloud Appliance.

Figure C.3 Fibre Channel and InfiniBand Connectivity for External Storage



## Diagram Notes

- All nodes in the Oracle Private Cloud Appliance rack are connected to the IPoB storage network (192.168.40.0/24) through their `bond1` interface, which consists of two virtual InfiniBand interfaces. For Fibre Channel storage connectivity all compute nodes have 4 virtual HBAs that terminate on the Fibre Channel I/O modules of the Fabric Interconnects. The underlying physical connection of all storage traffic consists of redundant InfiniBand cabling. All external storage connectivity is managed by the Fabric Interconnects.
- External InfiniBand storage, shown on the left-hand side of the diagram, can be connected to InfiniBand ports 1 and 11 of the Fabric Interconnects using one of these 3 supported cabling schemes:
  - In active/passive configuration, with both storage controllers cross-cabled to a pair of HA InfiniBand switches. Those switches are, in turn, cross-cabled to the available ports on the Fabric Interconnects.  
In the diagram, this cabling scheme is represented by the example ZFS Storage Appliance placed in the bottom left corner.
  - In active/passive configuration, attached directly to the available ports on the Fabric Interconnects.  
This cabling scheme is not shown in the diagram. However, it is identical to the active/passive example, but with cables running directly to the available ports on the Fabric Interconnects instead of the InfiniBand switches.
  - In active/active configuration, with both storage controllers using 4 cable connections to a pair of HA InfiniBand switches. Those switches are cross-cabled to the available ports on the Fabric

Interconnects. The additional InfiniBand ports and 4 cables per storage controller are required to provide redundancy at the level of the HCAs as well as the network paths of both active controllers.

In the diagram, this cabling scheme is represented by the example ZFS Storage Appliance placed in the top left corner.

- External Fibre Channel storage, shown on the right-hand side of the diagram, is connected to the I/O modules 3 and 12 of the Fabric Interconnects, but through one or two external Fibre Channel switches.
- The minimum configuration uses 1 port per FC I/O module for a total of 4 physical connections to one external Fibre Channel switch. In this configuration both HBAs of both storage controllers are connected to the same Fibre Channel switch.
- The full HA configuration uses both ports in each FC I/O module for a total of 8 physical connections: 4 per external Fibre Channel switch. In this configuration the HBAs of both storage controllers are cross-cabled to the Fibre Channel switches: one HBA is connected to switch 1 and the other HBA is connected to switch 2. The additional connections for the HA configuration are marked with dashed lines in the diagram.
- The vHBAs in the servers and the Fibre Channel ports in the I/O modules of the Fabric Interconnects are connected according to a round-robin scheme for high availability. In a field installation the compute nodes are provisioned in random order. Therefore it is impossible to predict which nodes are connected to which I/O ports. The software does ensure that half the compute nodes follow the connection scheme of management node 1, while the other half follows the connection scheme of management node 2.

The connection scheme presented in [Table C.1](#) shows that the four vHBAs in each server node are virtually cross-cabled to the Fabric Interconnects and their Fibre Channel I/O modules. To find out exactly how a particular compute node connects to the external FC storage you have to run the CLI command `list wwpn-info`.

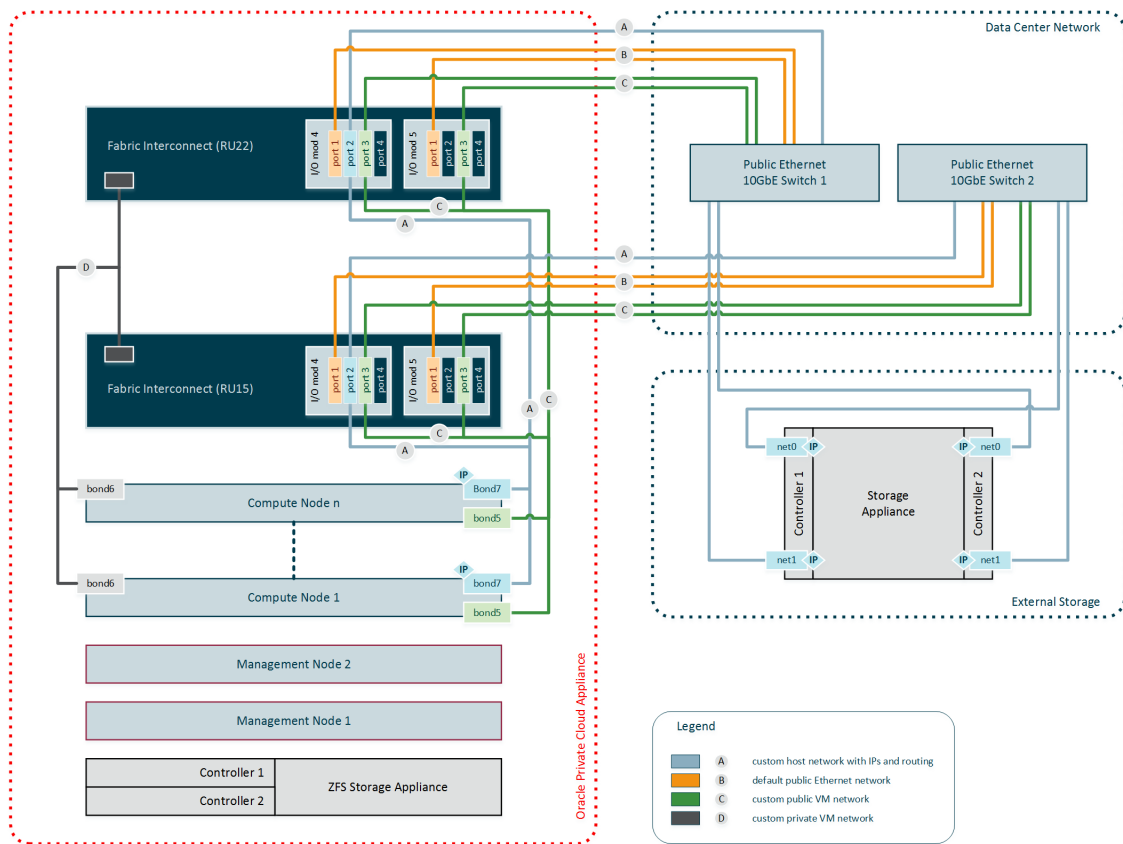
**Table C.1 Fibre Channel vHBA, I/O Port and Storage Cloud Mapping**

vHBA	Fabric Interconnect	I/O Port	Storage Cloud
vhba01	ovcasw22r1 RU22	3/1 or 12/1	A
vhba02	ovcasw22r1 RU22	3/2 or 12/2	B
vhba03	ovcasw15r1 RU15	3/1 or 12/1	C
vhba04	ovcasw15r1 RU15	3/2 or 12/2	D

## C.4 Custom and Host Networks

This section shows the various types of custom networks that are supported on an Oracle Private Cloud Appliance with Fabric Interconnects.

Figure C.4 Host Network and Other Custom Networking With or Without Public Connectivity



## Diagram Notes

- All custom networks are virtual Ethernet connections. They connect compute node bonds, consisting of two virtual interfaces on the nodes, to the Fabric Interconnects. The underlying physical connection is the InfiniBand fabric. Private custom networks, which provide internal connectivity only, terminate on the Fabric Interconnects and are not routed externally. Public custom networks terminate on the I/O modules and are routed externally across the 10GbE ports, which are cabled to Ethernet switches in the data center network. A public custom network can use one or more I/O ports.
- The host network is a particular type of public custom network: its configuration contains additional parameters for subnet and routing. The servers connected to it also receive an IP address in that subnet, and consequently can connect to an external network device.
- Only I/O modules 4 and 5 are shown in this diagram, but each Oracle Fabric Interconnect F1-15 installed in an Oracle Private Cloud Appliance rack contains two more I/O modules in slots 10 and 11, with four 10GbE ports each. Ports 4:1 and 5:1 of both Fabric Interconnects must be cabled for the default external connectivity of the appliance. All other 10GbE ports can be used to set up custom networks, provided the system limitations are respected.
- Custom networking for an Oracle Private Cloud Appliance environment is limited to a maximum of 7 custom external networks and 3 custom internal networks.
- To configure custom networks, you must use the Oracle Private Cloud Appliance Command Line Interface (CLI). The steps for each type of network are described in the section entitled [Network Customization](#) in the *Oracle Private Cloud Appliance Administrator's Guide*.

**Warning**

For the additional routing capabilities with a host network, compute nodes must be upgraded with the Oracle VM Server version included in the Oracle Private Cloud Appliance Controller Software Release 2.2.1 or later.

## Connecting to External Storage over a Host Network

The host network is particularly useful for direct access to storage devices, as illustrated by the example in blue and labeled 'A' in the diagram. This section further explains how you configure the host network for external storage connectivity.

### Subnet

Your host network can be any size subnet you prefer, as long as approximately 100 IP addresses in the range are reserved for use by the Oracle Private Cloud Appliance components. To be exact, the appliance software requires control over a contiguous set of IPs at the beginning of a static 3-octet range, where the 4th octet is 3–89. In other words, if the host network prefix is  $x.y.z$  then the required IP range is  $x.y.z.3-x.y.z.89$ . This reserved range is extended with a safety margin up to address  $x.y.z.100$ . Although technically not entirely accurate, this reserved address range is conveniently referred to as *the first 100 IPs in a subnet*. The following examples illustrate this particular requirement:

- The simplest scenario is a /24 subnet like  $10.1.1.0/24$ . In this case you must ensure that  $10.1.1.3-100$  are reserved for Oracle Private Cloud Appliance. The addresses  $10.1.1.101-255$  may be used by customer resources.
- A larger subnet configuration may seem more complicated, but the same principle applies. If the subnet is defined as  $10.100.16.0/20$  the address range is  $10.100.16.1-10.100.31.255$ . In this case you must select a 3-octet prefix for the host network, similar to a /24 subnet, and ensure that the IPs 3–100 with that prefix are reserved for Oracle Private Cloud Appliance. In this example, if you select  $10.100.20$  as the host network prefix, then the reserved address range is  $10.100.20.3-10.100.20.100$ . All other addresses in the subnet may be used by customer resources.

If you configure the host network so that it coincides or overlaps with an existing subnet in your data center, you must make sure that no address from the reserved range can be acquired – either statically or dynamically through DHCP – by a host external to the appliance environment.

Assuming this can be achieved in your data center network, you can assign non-reserved IP addresses in that same subnet to your external storage resources. As a result, the compute nodes in the appliance can reach those storage resources directly. No routing is required.

### Static Routing

Depending on the complexity of your data center network configuration, it may be convenient to define a separate subnet specifically for the host network containing the Oracle Private Cloud Appliance compute nodes. For your convenience, a /24 subnet is recommended. You can then configure a gateway inside the host network that redirects traffic to the subnet where your storage resources reside. This configuration requires additional routing arguments as part of the host network definition.

**Caution**

Static routing is possible, but not preferred. Introducing a router increases hop count and latency, which may result in decreased performance.

For storage traffic Oracle Private Cloud Appliance uses a default MTU of 9000 bytes. Make sure that all devices participating in the host network, meaning the

storage arrays as well as switches and routers, support this packet size and are configured accordingly.

Below is an example of a host network configuration with static routing:

- host network subnet: `10.1.1.0/24`
- storage resources subnet: `10.125.88.0/23`
- gateway configuration:
  - IP address 1: `10.1.1.199`
  - IP address 2: `10.125.88.199`

The gateway has an IP address in the host network and the subnet where the storage resources reside. It must be configured to route network traffic from the host network to the subnet containing the storage resources.

- When configuring the host network, you must add the following routing arguments:
  - route destination: the subnet containing the IP addresses of the external storage resources (`10.125.88.0/23`)
  - route gateway: the gateway's IP address in the host network (`10.1.1.199`)

These routing arguments result in a new static route entry in the routing table of each compute node, so that network traffic intended for subnet `10.125.88.0/23` is sent via the network interface connected to the host network to the gateway at `10.1.1.199`. The gateway routes this traffic to the destination subnet.





---

## Appendix D Site Checklists

### Table of Contents

D.1 System Components Checklist .....	127
D.2 Data Center Room Checklist .....	127
D.3 Data Center Environmental Checklist .....	128
D.4 Access Route Checklist .....	129
D.5 Facility Power Checklist .....	131
D.6 Safety Checklist .....	132
D.7 Logistics Checklist .....	132
D.8 Network Specification Checklist .....	134

This appendix contains site checklists to help you ensure that your site is prepared for installing the Oracle Private Cloud Appliance.

### D.1 System Components Checklist

Complete the following checklist to ensure that the systems component considerations have been addressed.

**Table D.1 System Component Checklist**

System Components Considerations	Yes	No	N/A	Comment
Will more than one base rack be installed?				
Is the rack installation a new system or an addition to an existing system?				
Will the base rack be connected to an expansion rack?				
Are all machine racks adjacent to each other?				
If the connecting racks are not within the specified proximity, then have the following been done:				
				<ul style="list-style-type: none"><li>• Purchased approved longer InfiniBand cables from an approved third-party vendor?</li><li>• Asked Oracle Support Services to provide and schedule the custom multi-rack cabling service with the installation?</li></ul>

### D.2 Data Center Room Checklist

Complete the following checklist to ensure that the data center room requirements are met. For information about the data center requirements, see [Section 2.2, "Flooring Requirements"](#).

**Table D.2 Data Center Room Checklist**

<b>Data Center Room Considerations</b>	<b>Yes</b>	<b>No</b>	<b>N/A</b>	<b>Comment</b>
Has the Oracle Private Cloud Appliance location been allocated?				
Is there a vacant location for the new equipment?				
Does the floor layout meet the equipment maintenance access requirements?				
Is there adequate space available for maintenance?				
Will the equipment be positioned so that the exhaust air of one rack does not enter the air inlet of another rack?				
Have cabinet stabilization measures been considered?				
Does the raised floor satisfy the weight requirements for the new hardware?				
Can floor tiles be removed without permission to accommodate service?				
Are there cable routing channels or cutouts?				
Are you providing any additional hardware?				
Is the hardware you are providing fully compatible with the Oracle Private Cloud Appliance?				
Will the new hardware location require any non-standard cable lengths?				
Is the floor to ceiling height a minimum of 3214 mm or 3.2 m (10.5 feet)?				
Is the depth of the raised floor a minimum of 46 cm (18 inches)?				

## D.3 Data Center Environmental Checklist

Complete the following checklist to ensure that the data center environment requirements are met. For information, see [Section 2.5, "Ventilation and Cooling Requirements"](#).

**Table D.3 Data Center Environmental Checklist**

<b>Data Center Environmental Considerations</b>	<b>Yes</b>	<b>No</b>	<b>N/A</b>	<b>Comment</b>
Does the computer room air handling meet temperature and humidity requirements?				
Does the installation floor layout satisfy the ventilation requirements?				
Will the equipment be positioned so the exhaust air of one rack does not enter the air intake of another rack?				
Are the perforated floor tiles each rated at 400 CFM or greater?				
Do the data center air conditioners provide sufficient front to back airflow?				
Is airflow adequate to prevent hot spots?				
Can the data center continuously satisfy environmental requirements?				
Can more vented floor tiles be obtained if required?				

## D.4 Access Route Checklist

Complete the following checklist to ensure that the access route requirements are met. For information about preparing your data center access route, see [Section 4.2, “Prepare to Install Oracle Private Cloud Appliance”](#).

**Table D.4 Access Route Checklist**

<b>Access Route Considerations</b>	<b>Yes</b>	<b>No</b>	<b>N/A</b>	<b>Comment</b>
Has the access route been checked for clearances of the packaged equipment?				
Do all the doors and entry ways conform to the width and height requirements for transportation, including the width of the unpacked unit?				
Do all the doors meet the height requirement of minimum 218 cm (86 inches) for packaged delivery?				
Does the access route provide sufficient space for transport of the packed devices?				
Are there any ramps or thresholds that are of concern? If yes, then provide details.				

Access Route Checklist

---

Access Route Considerations	Yes	No	N/A	Comment
Are there any stairs or ramps in the path when moving the new hardware?				
Have you confirmed that all route incline angles are within the permitted range?				
Have you confirmed that the access route is free of any obstacles that would expose the device to shock?				
Are all the surfaces acceptable for rolling the new unpacked and packed equipment?				
<p>If a pallet jack is to be used, then have you confirmed the following:</p> <ul style="list-style-type: none"> <li>• The pallet jack supports the device weight?</li> <li>• The pallet jack tines are compatible with the shipping pallet?</li> </ul>				
If there are stairs, then is a loading elevator accessible for the equipment?				
<p>If an elevator is to be used, then have you confirmed the following:</p> <ul style="list-style-type: none"> <li>• The elevator car is wide enough for the device to be carried into it?</li> <li>• The elevator car is high enough for the device to be carried into it?</li> <li>• The load limit of the elevator is greater than the device weight?</li> <li>• Are elevators available to handle up to 996 kg (2197 lbs) fully-loaded rack capacity?</li> <li>• The elevator door meets the minimum height requirement of 218 cm (86 inches) for packaged rack delivery?</li> </ul>				
Does the path from the receiving location to the designated data center area support the weight of the unpacked equipment?				
Is the path onto the raised floor rated for dynamic loading of the				

Access Route Considerations	Yes	No	N/A	Comment
server? Refer to <a href="#">Section 2.2</a> , "Flooring Requirements".				

## D.5 Facility Power Checklist

Complete the following checklist to ensure that the facility power requirements are met. For information about power requirements, see [Section 2.3](#), "Electrical Power Requirements".

**Table D.5 Facility Power Checklist**

Facility Power Considerations	Yes	No	N/A	Comment
Do you know the required operating voltage and electric current level of the device and peripherals?				
Will you be using single-phase (low-voltage or high-voltage) or 3-phase (low-voltage or high-voltage) power?				
Are enough power outlets provided within 2 meters (6.5 feet) for each rack?				
Do the power outlets have appropriate socket receptacles for the PDU option ordered? Options are low voltage or high voltage, single-phase or 3-phase.				
Will optional ground cables be attached to the rack?				
Are the circuit breakers for the equipment suitable in terms of voltage and current-carrying capacities?				
Does the power frequency meet the equipment specifications?				
Are power outlets available for the new equipment at the designated location?				
Will system power be delivered from two separate grids?				
Is there a UPS to power the equipment?				
Do you have the minimum required power sources to support the power load for the new hardware? Use kilowatt (kW) /kilovolt (kVA) to express power load.				

## D.6 Safety Checklist

Complete the following checklist to ensure that the safety requirements are met. For information about safety, see [Section 1.6, “Operational Procedures for Oracle Private Cloud Appliance”](#) and [Section 2.5, “Ventilation and Cooling Requirements”](#).

**Table D.6 Safety Checklist**

Safety Checklist Considerations	Yes	No	N/A	Comment
Is there an emergency power shut off?				
Is there a fire protection system in the data center room?				
Is the computer room adequately equipped to extinguish a fire?				
Is antistatic flooring installed?				
Is the floor below the raised floor free of obstacles and blockages?				

## D.7 Logistics Checklist

Complete the following checklist to ensure that the logistics requirements are met. For information about unpacking and space requirements, see [Section 2.1, “Space Requirements”](#).

**Table D.7 Logistics Checklist**

Logistics Checklist Considerations	Yes	No	N/A	Comment
Do you have contact information for the data center personnel?				
Is there security or access control for the data center?				
Are there any security background checks or security clearances required for vendor personnel to access the data center? If yes, then do you have a recommended agency?				
How many days in advance must background checks be completed?				
Are there any additional security access issues?				
Is computer room access available for installation personnel?				
Are laptops, cell phones, and cameras allowed in the data center?				
Does the building have a delivery dock?				
Is there a delivery/unpacking/staging area?				

Logistics Checklist

Logistics Checklist Considerations	Yes	No	N/A	Comment
Is the delivery inside?				
If the delivery is not inside, then is the site prepared for uncrating?				
Is the unpacking/staging area protected from the elements?				
Does the building have adequate receiving space?				
Is the unpacking area air-conditioned to avoid thermal shock for various hardware components?				
Will sufficient moving personnel be available to install the hardware?				
Is union labor required for any part of the delivery or installation?				
Are you prepared for uncrating and trash removal?				
Is uncrating of cabinet and cabinet trash removal required?				
Are there any restrictions on delivery truck length, width, or height?				
Does the customer allow cardboard boxes and other packing material in the computer room? If no, then do ground level deliveries require a truck with a side rail lift?				
Is there a time constraint on dock access? If yes, then provide time constraints.				
Is tail lift required on delivery carrier to unload the equipment at the delivery dock?				
Will any of the following be required to place equipment in computer room?				
<ul style="list-style-type: none"> <li>• Stair walkers</li> <li>• Lifters</li> <li>• Ramps</li> <li>• Steel plates</li> <li>• Floor covers</li> </ul>				
Does the delivery carrier require any special equipment, such as non-				

Logistics Checklist Considerations	Yes	No	N/A	Comment
floor damaging rollers, transport dollies, pallet jacks, or fork lifts?				

## D.8 Network Specification Checklist

Complete the following checklist to ensure that the network specification requirements are met. For information about IP addresses, see [Chapter 3, Network Requirements](#) and [Appendix A, Default IP Addresses](#).

**Table D.8 Network Specification Checklist**

Network Specification Considerations	Yes	No	N/A	Comment
Did you complete the networking worksheet, which is included in <a href="#">Section 3.3, "General Network Configuration"</a> ?				
Have you received the site-specific installation template from your Oracle technical representative?				
Did you review the installation template and consult with your Oracle technical representative regarding any site-specific changes, if any?				
Did you verify that IP addresses in the installation template are currently not in use?				
Have you performed the required configuration within your network infrastructure to allow the Oracle Private Cloud Appliance to use the IP addresses specified in the template?				
Have you registered IP addresses in the installation template with DNS?				
Did you run the required network cables from your network equipment to the location where the Oracle Private Cloud Appliance will be installed?				
Did you label the network cables that will connect to the Oracle Private Cloud Appliance?				



---

# Index

## A

- access route checklist, 129
- access route requirements, 8
- additional RAM, 67
- appliance
  - system components, 3
  - system features, 2
  - system overview, 2

## C

- cabling
  - management to data center network, 19
  - spine to data center network, 15
- cabling compute nodes, 62
- cabling reference
  - Cisco Nexus 9336C-FX2 Switch, 108
  - Cisco Nexus 9348GC-FXP Switch, 107
  - interconnects data/spine switches, 110
- cabling reference table
  - NM2-36P InfiniBand (IB) switch to F1-15 Fabric Interconnects, 114
  - NM2-36P InfiniBand switches to compute nodes, 114
  - Oracle Switch ES1-24 Ethernet connections, 111
  - storage head cluster connections, 113
- cautions and warnings, 5
- circuit breaker requirements, 10
- Cisco Nexus 9348GC-FXP Switch
  - connecting to data center network, 19
  - management port, 19
- compute node
  - attach slide-rail assemblies, 53
  - cabling, 62
  - ESD precautions for installation, 49
  - expansion, 47
  - install cable management arm, 57
  - install mounting brackets, 52
  - installation safety guidelines, 48
  - installing into rack, 52
  - preparing rack for expansion, 50
  - rack-mount safety precautions, 48
  - slide server into rails, 55
  - unpack upgrade kit, 50
- connecting to the network, 21
- connectivity diagram, 117
  - custom network, 122
  - external Ethernet, 117
  - external storage, 119, 120
- cooling requirements, 12
- custom networking
  - connectivity diagram, 122

host network, 122

## D

- data center connectivity
  - configure uplinks, 18
- data center environment checklist, 128
- data center room checklist, 127
- default IP addresses, 103
  - InfiniBand, 105
  - management IP addresses, 103
  - storage network, 105

## E

- electrical disturbances, 9
- electrical grounding guidelines, 10
- electrical power requirements, 9
- Ethernet
  - external connectivity diagram, 117
- extending the system, 67, 71
  - adding compute nodes, 47
  - ESD precautions, 49
  - installing compute nodes, 52
  - preparing the rack, 50
  - rack-mount safety precautions, 48
  - removing rack doors, 50
  - replacing rack doors, 65
  - safety guidelines, 48
  - server cabling, 62
  - server upgrade kit, 50
- external networking
  - configure uplinks, 18
- external storage
  - connectivity diagram, 119, 120

## F

- Fabric Interconnect
  - connecting to the network, 21
  - IO module ports, 21
- facility power checklist, 131
- facility power requirements, 10
- features
  - overview, 2
- floor load requirements, 8

## G

- grounding cable, 33

## H

- humidity requirements, 10

## I

- InfiniBand IP addresses, 105

---

## installation

- connecting to the network, 41
- first power-on, 35
- grounding cable, 33
- moving the system, 30
- overview of procedures, 1, 27
- power on procedures, 38
- safety documents, 27
- safety guidelines, 28
- site preparation, 28
- stabilizing the rack, 31
- tool requirements, 29
- unpacking, 30

## IP address requirements

- Ethernet-based system, 20
- InfiniBand-based system, 21, 23

## L

- logistics checklist, 132

## M

- maintenance access requirements, 8
- management IP addresses, 103
- moving the system, 30

## N

### network

- cabling requirements, 15
- configuration worksheet, 24
- internal management, 19
- IP requirements, 20
- network connection, 41
- network requirements, 15
- network specification checklist, 134
- networking
  - configure uplinks, 18
  - custom networks, 122
  - host network, 122

## O

- Oracle Private Cloud Appliance
  - package dimensions, 7
  - receiving and unpacking, 7
  - site requirements, 7
  - space requirements, 7
- Oracle Server X8-2
  - components, 4
- Oracle Switch ES1-24
  - connecting to the network, 22
  - management port, 22

## P

- Power Distribution Unit (PDU)

- high voltage requirements, 9
- low voltage requirements, 9
- power off procedures
  - emergency situations, 5
- power requirements
  - electrical disturbances, 9
  - power distribution unit, 9
  - protecting your system, 9
- powering on system, 35
  - apply power to main components, 38
  - before you start, 36
  - connecting power cords, 37
  - connecting to the network, 41
  - Dashboard, 41
  - safety guidelines, 36
  - visual inspection, 37
- preparing installation site, 28

## R

- rack components, 3
- rack doors
  - reinstall, 65
  - removing, 50
- rack-mount safety precautions, 48

## S

- safety checklist, 132
- safety guidelines, 28, 36, 48
  - ESD precautions, 49
  - rack-mount precautions, 48
- safety information
  - where to find, 1
- server components, 4
- server upgrade kit, 50
- site checklists
  - access route, 129
  - data center environment, 128
  - data center room, 127
  - facility power, 131
  - logistics, 132
  - network specification, 134
  - safety, 132
  - system components, 127
- stabilizing the rack, 31
- storage
  - external storage diagram, 119, 120
- storage extension, 71
  - Ethernet-based architecture, 71
  - InfiniBand-based architecture, 72
- storage network IP addresses, 105
- system component checklist, 127
- system components, 3
- system features, 2

---

system overview, 2

## **T**

temperature requirements, 10

tools

    unpacking and installation, 29

## **U**

unpacking the system, 7, 30

uplink

    accepted configurations, 18

## **V**

ventilation requirements, 12

