Oracle® Exadata Database Machine Extending and Multi-Rack Cabling Guide





Oracle Exadata Database Machine Extending and Multi-Rack Cabling Guide, 19.1.0

E93785-01

Copyright © 2008, 2018, Oracle and/or its affiliates. All rights reserved.

Primary Author: Janet Stern

Contributing Authors: Lypp-Tek Khoo-Ellis

Contributors: Doug Archambault, Leo Agranonik, Andrew Babb, Bharat Baddepudi, Michael Belef, Nilesh Choudhury, Henry Chow, Kevin Closson, Bob Cyphers, Ravindra Dani, Boris Erlikhman, Jaime Figueroa, Gurmeet Goindi, Roger Hansen, Shrikumar Hariharasubrahmanian, Kevin Jernigan, Caroline Johnston, Sarat Karkarla, Leslie Keller, Frank Kobylanski, Poojan Kumar, René Kundersma, Robert S. Lee, Sue Lee, Jennifer Leung, Yang Liu, Juan Loaiza, Barb Lundhild, Catherine Luu, Varun Malhotra, Bagish Mishra, Valarie Moore, Philip Newlan, Adrian Ng, Dan Norris, Michael Nowak, Bharat Paliwal, Umesh Panchaksharaiah, Hector Pujol, Sugam Pandey, Dmitry Potapov, Darryl Presley, Vince Pulice, Ashish Ray, Samitha Samaranayake, Richard Scales, Olly Sharwood, Jia Shi, Kesavan Srinivasan, Mahesh Subramaniam, Michelle Sun, Krishnadev Telikicherla, Cliff Thomas, Alex Tsukerman, Shreyas Udgaonkar, Kothanda Umamageswaran, Doug Utzig, James Viscusi, Vern Wagman, Zheren Zhang

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.

Contents

Preface

Audience Docume	tation Accessibility	v v
Related I	ocumentation	Vi
Conventi	ons	
Extend	ing Oracle Exadata Database Machine	
1.1 Ab	out Extending Oracle Exadata Database Machine	1-
1.2 Pre	paring to Extend Oracle Exadata Database Machine	1-
1.2.1	Reviewing the Safety Precautions	1-
1.2.2	Reviewing the InfiniBand Cable Precautions	1-
1.2.3	Estimating InfiniBand Cable Path Lengths	1-
1.2.4	Bundling InfiniBand Cables	1-
1	2.4.1 Floor and Underfloor Delivery of InfiniBand Ca	bles 1-
1.2.5	Reviewing the Cable Management Arm Guidelines	1-
1.2.6	Obtaining Current Configuration Information	1-
1.2.7	Preparing the Network Configuration	1-
1.2.8	Moving Audit and Diagnostic Files	1-
1.2.9	Reviewing Release and Patch Levels	1-
1.2.1	Performing Preliminary Checks	1-
1.2.1	Preparing to Add Servers	1-
1.3 Ex	ending the Hardware	1-1
1.3.1	Extending an Eighth Rack to a Quarter Rack in Orac Database Machine X4-2 and Later	le Exadata 1-1
1	3.1.1 Reviewing and Validating Current Configuratio Oracle Exadata Database Machine X4-2 or La	•
1	3.1.2 Activating Database Server Cores in Eighth Ra Database Machine X4-2 or Later	ack Oracle Exadata 1-1
1	3.1.3 For Oracle Exadata Database Machine X6-2: A Capacity Disks and Flash Cards	Adding High 1-1
1	3.1.4 For Oracle Exadata Database Machine X7-2: URack Systems to a Quarter Rack	Jpgrading Eighth 1-1



	1.3.1.	5 Activating Storage Server Cores and Disks in Eighth Rack Oracle Exadata Database Machine X4-2 or Later	1-14
	1.3.1.		1-14
	1.3.1.		1-14
	1.3.1.	Oracle Exadata Database Machine X4-2 or Later	1-18
	1.3.1.8	Validating New Quarter Rack Configuration for Oracle Exadata Database Machine X4-2 or Later	1-20
		xtending an Eighth Rack to a Quarter Rack in Oracle Exadata atabase Machine X3-2	1-21
	1.3.2.	Reviewing and Validating Current Configuration of Oracle Exadata Database Machine X3-2 Eighth Rack	1-21
	1.3.2.	Activating Database Server Cores in Oracle Exadata Database Machine X3-2 Eighth Rack	1-22
	1.3.2.	Activating Storage Server Cores and Disks in Oracle Exadata Database Machine X3-2 Eighth Rack	1-22
	1.3.2.	Creating Grid Disks in Oracle Exadata Database Machine X3-2 Eighth Rack	1-23
	1.3.2.	Adding Grid Disks to Oracle ASM Disk Groups in Oracle Exadata Database Machine X3-2 Eighth Rack	1-25
	1.3.2.0	Validating New Oracle Exadata Database Machine X3-2 Quarter Rack Configuration	1-26
	1.3.3 Ex	ctending a Quarter Rack or Half Rack	1-27
	1.3.3.	Removing the Doors	1-27
	1.3.3.	2 Adding a Sun Datacenter InfiniBand Switch 36 Switch (Optional)	1-29
	1.3.3.	3 Adding New Servers	1-31
	1.3.3.4	4 Cabling Database Servers	1-34
	1.3.3.	5 Cabling Exadata Storage Servers	1-39
	1.3.3.	6 Closing the Rack	1-41
	1.3.4 Ex	ctending a Rack by Adding Another Rack	1-42
	1.3.4.	Cabling Two Racks Together	1-42
	1.3.4.	2 Cabling Several Racks Together	1-44
1.4	Configu	ring the New Hardware	1-45
	1.4.1 CI	nanging the Interface Names	1-46
	1.4.2 Se	etting Up New Servers	1-47
	1.4.3 Se	etting up a New Rack	1-50
	1.4.4 Se	etting User Equivalence	1-51
	1.4.5 St	arting the Cluster	1-53
	1.4.6 Ad	dding Grid Disks to Oracle ASM Disk Groups	1-54
	1.4.7 Ad	dding Servers to a Cluster	1-57
	1.4.8 Co	onfiguring Cell Alerts for New Oracle Exadata Storage Servers	1-62
	1.4.9 Ad	dding Oracle Database Software to the New Servers	1-63
	1.4.10 A	Adding Database Instance to the New Servers	1-64



1.5	Returning	the	Rack to	Service
-----	-----------	-----	---------	---------

1-65

2 Multi-Rack Cabling Tables

2.1	Und	erstanding Multi-Rack Cabling	2-1
	2.1.1	Preparing to Cable Racks Together	2-4
	2.1.2	Cabling Oracle Exadata Quarter Racks and Oracle Exadata Eighth Racks	2-7
2.2	Two	-Rack Cabling	2-9
2.3	Thre	e-Rack Cabling	2-11
2.4	Fou	r-Rack Cabling	2-13
2.5	Five	-Rack Cabling	2-16
2.6	Six-l	Rack Cabling	2-19
2.7	Seve	en-Rack Cabling	2-23
2.8	Eigh	t-Rack Cabling	2-28
Inc	lex		



List of Tables

1-1	Commands to Create Disk Groups When Extending Oracle Exadata Database	
	Machine X4-2 Eighth Rack or Later	1-16
1-2	Commands to Add Disk Groups When Extending Eighth Rack Oracle Exadata	
	Database Machine X4-2 and Later	1-19
1-3	Commands to Create Disk Groups When Extending Oracle Exadata Database	
	Machine X3-2 Eighth Rack	1-24
1-4	Commands to Add Disk Groups When Extending an Oracle Exadata Database	
	Machine X3-2 Eighth Rack	1-26
2-1	Leaf Switch Connections for the First Rack in a Three-Rack System	2-12
2-2	Leaf Switch Connections for the Second Rack in a Three-Rack System	2-12
2-3	Leaf Switch Connections for the Third Rack in a Three-Rack System	2-13
2-4	Leaf Switch Connections for the First Rack in a Four-Rack System	2-14
2-5	Leaf Switch Connections for the Second Rack in a Four-Rack System	2-14
2-6	Leaf Switch Connections for the Third Rack in a Four-Rack System	2-15
2-7	Leaf Switch Connections for the Fourth Rack in a Four-Rack System	2-15
2-8	Leaf Switch Connections for the First Rack in a Five-Rack System	2-17
2-9	Leaf Switch Connections for the Second Rack in a Five-Rack System	2-17
2-10	Leaf Switch Connections for the Third Rack in a Five-Rack System	2-18
2-11	Leaf Switch Connections for the Fourth Rack in a Five-Rack System	2-18
2-12	Leaf Switch Connections for the Fifth Rack in a Five-Rack System	2-19
2-13	Leaf Switch Connections for the First Rack in a Six-Rack System	2-20
2-14	Leaf Switch Connections for the Second Rack in a Six-Rack System	2-20
2-15	Leaf Switch Connections for the Third Rack in a Six-Rack System	2-21
2-16	Leaf Switch Connections for the Fourth Rack in a Six-Rack System	2-21
2-17	Leaf Switch Connections for the Fifth Rack in a Six-Rack System	2-22
2-18	Leaf Switch Connections for the Sixth Rack in a Six-Rack System	2-23
2-19	Leaf Switch Connections for the First Rack in a Seven-Rack System	2-24
2-20	Leaf Switch Connections for the Second Rack in a Seven-Rack System	2-24
2-21	Leaf Switch Connections for the Third Rack in a Seven-Rack System	2-25
2-22	Leaf Switch Connections for the Fourth Rack in a Seven-Rack System	2-25
2-23	Leaf Switch Connections for the Fifth Rack in a Seven-Rack System	2-26
2-24	Leaf Switch Connections for the Sixth Rack in a Seven-Rack System	2-27
2-25	Leaf Switch Connections for the Seventh Rack in a Seven-Rack System	2-27
2-26	Leaf Switch Connections for the First Rack in a Eight-Rack System	2-28
2-27	Leaf Switch Connections for the Second Rack in a Eight-Rack System	2-29



2-28	Leaf Switch Connections for the Third Rack in a Eight-Rack System	2-29
2-29	Leaf Switch Connections for the Fourth Rack in a Eight-Rack System	2-30
2-30	Leaf Switch Connections for the Fifth Rack in a Eight-Rack System	2-31
2-31	Leaf Switch Connections for the Sixth Rack in a Eight-Rack System	2-31
2-32	Leaf Switch Connections for the Seventh Rack in a Eight-Rack System	2-32
2-33	Leaf Switch Connections for the Eighth Rack in a Eight-Rack System	2-32



Preface

This guide describes how to extend and cable Oracle Exadata Database Machine. It includes information about the cables, new server installation, and cabling tables.

Audience

This guide is intended for Oracle Exadata Database Machine customers and those responsible for data center site planning, configuration, and maintenance of Oracle Exadata Database Machine.

Documentation Accessibility

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

Access to Oracle Support

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

Related Documentation

In addition to the Oracle Exadata Database Machine documentation set, the following guides contain hardware information for Oracle Exadata Database Machine:

- Oracle Exadata Database Machine System Overview
- Oracle Exadata Database Machine Security Guide
- Oracle Engineered System Safety and Compliance Guide, Compliance Model No.: ESY27
- Oracle Exadata Database Machine Extending and Multi-Rack Cabling Guide
- Oracle Exadata Database Machine Maintenance Guide
- Oracle Exadata System Software User's Guide
- Oracle Server X7-2 Service Manual at http://docs.oracle.com/cd/E72435_01/html/E72445/index.html
- Oracle Server X7-2L Service Manual at http://docs.oracle.com/cd/E72463_01/html/E72474/index.html



- Oracle Server X7-8 Service Manual at http://docs.oracle.com/cd/E71925_01/html/E71936/index.html
- Oracle Server X6-2 Service Manual at http://docs.oracle.com/cd/E62159_01/html/E62171/index.html
- Oracle Server X6-2L Service Manual at http://docs.oracle.com/cd/E62172_01/html/E62184/index.html
- Oracle Server X5-2 Service Manual at http://docs.oracle.com/cd/E41059_01/html/E48312/ napsm.html
- Oracle Server X5-2L Service Manual at http://docs.oracle.com/cd/E41033_01/html/E48325/ cnpsm.html#scrolltoc
- Sun Server X4-8 Service Manual at http://docs.oracle.com/cd/E40591_01/html/E40317/index.html
- Sun Server X4-2 Service Manual at http://docs.oracle.com/cd/E36975_01/html/E38045/gentextid-14757.html#scrolltoc
- Sun Server X4-2L Service Manual at http://docs.oracle.com/cd/E36974_01/html/E38145/ gentextid-14728.html#scrolltoc
- Sun Server X3-2 (formerly Sun Fire X4170 M3) Service Manual at http:// docs.oracle.com/cd/E22368 01/html/E27242/gentextid-14840.html#scrolltoc
- Sun Server X3-2L (formerly Sun Fire X4270 M3) Service Manual at http://docs.oracle.com/cd/E23393_01/html/E27229/gentextid-14804.html#scrolltoc
- Sun Server X2-8 (formerly Sun Fire X4800 M2) Service Manual at http://docs.oracle.com/cd/E20815_01/html/E20819/index.html
- Sun Fire X4800 Server Service Manual at http://docs.oracle.com/cd/E19140-01/ html/821-0282/index.html
- Sun Fire X4270 M2 Server Service Manual at http://docs.oracle.com/cd/ E19245-01/E21671/index.html
- Sun Fire X4170 M2 Server Service Manual at http://docs.oracle.com/cd/ E19762-01/E22369-02/index.html
- Sun Fire X4170, X4270, and X4275 Servers Service Manual at http://docs.oracle.com/cd/E19477-01/820-5830-13/index.html
- Sun Datacenter InfiniBand Switch 36 Firmware Version 2.1 Documentation at http://docs.oracle.com/cd/E36265 01/index.html
- Sun Datacenter InfiniBand Switch 36 Firmware Version 2.2 Documentation at http://docs.oracle.com/cd/E76424_01/index.html
- Sun Flash Accelerator F20 PCIe Card User's Guide at http://docs.oracle.com/cd/ E19682-01/E21358/index.html
- Sun Flash Accelerator F40 PCIe Card User's Guide at http://docs.oracle.com/cd/ E29748_01/html/E29741/index.html
- Sun Flash Accelerator F80 PCIe Card User's Guide at http://docs.oracle.com/cd/ E41278_01/html/E41251/index.html
- Oracle Flash Accelerator F160 PCle Card User Guide at http://docs.oracle.com/cd/E54943_01/html/E54947/index.html
- Oracle Flash Accelerator F320 PCIe Card User Guide at http://docs.oracle.com/cd/E65386_01/html/E65387/index.html



- Sun Storage 6 Gb SAS PCIe RAID HBA Documentation at http://docs.oracle.com/cd/E19221-01/
- Oracle Storage 12 Gb/s SAS PCIe RAID HBA, Internal Documentation Library at http://docs.oracle.com/cd/E52363_01/index.html
- Oracle Integrated Lights Out Manager (ILOM) Documentation at http://www.oracle.com/goto/ilom/docs
- "Cisco Catalyst 4948E and 4948E-F Ethernet Switches Data Sheet" at https://www.cisco.com/c/en/us/products/collateral/switches/catalyst-4948e-ethernet-switch/data_sheet_c78-598933.html
- "Cisco Nexus 9300-EX and 9300-FX Platform Switches Data Sheet at https://www.cisco.com/c/en/us/products/collateral/switches/nexus-9000-series-switches/datasheet-c78-736651.html"

Conventions

The following text conventions are used in this document:

Convention	Meaning
boldface	Boldface type indicates graphical user interface elements associated with an action, or terms defined in text or the glossary.
italic	Italic type indicates book titles, emphasis, or placeholder variables for which you supply particular values.
monospace	Monospace type indicates commands within a paragraph, URLs, code in examples, text that appears on the screen, or text that you enter.
\$ prompt	The dollar sign (\$) prompt indicates a command run as the oracle user.
# prompt	The pound (#) prompt indicates a command that is run as the root user.



1

Extending Oracle Exadata Database Machine

This chapter describes how to extend Oracle Exadata Database Machine.



For ease of reading, the name "Oracle Exadata Rack" is used when information refers to both Oracle Exadata Database Machine and Oracle Exadata Storage Expansion Rack.

See Also:

- Multi-Rack Cabling Tables
- Oracle Exadata Database Machine System Overview for the cabling tables for each Exadata version

1.1 About Extending Oracle Exadata Database Machine

You can extend Oracle Exadata Database Machine as follows:

- You can extend Oracle Exadata Database Machine from a fixed configuration to another fixed configuration, for example, from an eighth rack to a quarter rack, a quarter rack to a half rack, or a half rack to a full rack.
- You can also extend Oracle Exadata Database Machine from a fixed or a custom configuration to another custom configuration by adding any combination of database or storage servers up to the allowed maximum. This is known as elastic configuration.
- Any combination of Oracle Exadata Database Machine half racks and full racks can be cabled together.
- A Sun Datacenter InfiniBand Switch 36 switch, and cables must be ordered before extending Oracle Exadata Database Machine X4-2 racks.

Note:

- The cable lengths shown in Multi-Rack Cabling Tables assume the racks are adjacent to each other. If the racks are not adjacent or use overhead cabling trays, then they may require longer cables lengths. Up to 100 meters is supported.
 - Only optical cables are supported for lengths greater than 5 meters.
- Earlier Oracle Exadata Racks can be extended with later Oracle Exadata Racks.
- When extending Oracle Exadata Database Machine Eighth Rack with Oracle Exadata Storage Expansion Rack make sure that there are two separate disk groups. There should be one disk group for the drives in the Oracle Exadata Database Machine Eighth Rack rack and one disk group for the drives in Oracle Exadata Storage Expansion Rack.

Multiple Oracle Exadata Database Machines can be run as separate environments, and connect through the InfiniBand network. If you are planning to use multiple Oracle Exadata Database Machines in this manner, then note the following:

All servers on the InfiniBand network must have a unique IP address. When Oracle Exadata Database Machine is deployed, the default InfiniBand network is 192.168.10.1. You must modify the IP addresses before re-configuring the InfiniBand network. Failure to do so causes duplicate IP addresses. After modifying the network, run the verify-topology and infinicheck commands to verify the network is working properly. You need to create a file that contains IP addresses for Exadata Storage Servers, such as combined_cellip.ora. The following is an example of the commands:

```
# cd /opt/oracle.SupportTools/ibdiagtools
# ./verify-toplogy -t fattree
# ./infinicheck -c /tmp/combined_cellip.ora -b
```

- When Oracle Exadata Database Machines run in separate clusters, do not modify the cellip.ora files. The cellip.ora file on a database server should only include the IP addresses for the Exadata Storage Servers used with that database server.
- Cells with disk types different from what is already installed can be added, but the
 disk types cannot be mixed in the same Oracle ASM disk group. For example, if
 the existing disk groups all use high performance disks, and cells with high
 capacity disks are being added, then it is necessary to create new disk groups for
 the high capacity disks.
 - When adding the same type of disk, ensure that the grid disk sizes are exactly the same even if the new disks are larger than the existing ones. For example, if the existing disks are 3 TB, and the additional disks are 4 TB, then it is necessary to create grid disks that match the size on the 3 TB disks. A new disk group can be created using the extra 1 TB of disk space.
- In order to access Exadata Storage Servers in one Oracle Exadata Database
 Machine by another Oracle Exadata Database Machine when they are not running
 as a single cluster, Exadata Storage Servers must have unique Oracle ASM disk
 group and failure group names on each Oracle Exadata Database Machine. For
 example, for two Oracle Exadata Database Machines cabled together but run as
 separate clusters, the following names should be unique:



- Cell name
- Cell disk name
- Grid disk name
- Oracle ASM failure group name
- All equipment receives a Customer Support Identifier (CSI). Any new equipment
 for Oracle Exadata Database Machine has a new CSI. Contact Oracle Support
 Services to reconcile the new CSI with the existing Oracle Exadata Database
 Machine CSI. Have the original instance numbers or serial numbers available, as
 well as the new numbers when contacting Oracle Support Services.

The InfiniBand network can be used for external connectivity. The external connectivity ports in the Sun Datacenter InfiniBand Switch 36 switches can connect to media servers for tape backup, data loading, and client and application access. Use the available ports on the leaf switches for external connectivity. There are 12 ports per rack. The available ports are 5B, 6A, 6B, 7A, 7B, and 12A in each leaf switch. For high availability connections, connect one port to one leaf switch and the other port to the second leaf switch. The validated InfiniBand cable lengths are as follows:

- Up to 5 meters passive copper 4X QDR QSFP cables
- Up to 100 meters fiber optic 4X D=QDR QSFP cables

See Also:

- Oracle Automatic Storage Management Administrator's Guide for information about renaming disk groups
- Oracle Exadata Database Machine System Overview for details about elastic configuration

1.2 Preparing to Extend Oracle Exadata Database Machine

Before extending any rack hardware, review the safety precautions and cabling information, and collect information about the current rack in this section.

1.2.1 Reviewing the Safety Precautions

Before upgrading Oracle Exadata Database Machines, read *Important Safety Information for Sun Hardware Systems* included with the rack.



Contact your service representative or Oracle Advanced Customer Services to confirm that Oracle has qualified your equipment for installation and use in Oracle Exadata Database Machine. Oracle is not liable for any issues when you install or use non-qualified equipment.



See Also:

- Oracle Exadata Database Machine Installation and Configuration Guide for safety guidelines
- Oracle Engineered System Safety and Compliance Guide, Compliance Model No.: ESY27 for safety notices

1.2.2 Reviewing the InfiniBand Cable Precautions

Review the following InfiniBand cable precautions before working with the InfiniBand cables:

- Fiber optic InfiniBand cables with laser transceivers must be of type Class 1.
- Do not allow any copper core InfiniBand cable to bend to a radius tighter than 127 mm (5 inches). Tight bends can damage the cables internally.
- Do not allow any optical InfiniBand cable to bend to a radius tighter than 85 mm (3.4 inches). Tight bends can damage the cables internally.
- Do not use zip ties to bundle or support InfiniBand cables. The sharp edges of the ties can damage the cables internally. Use hook-and-loop straps.
- Do not allow any InfiniBand cable to experience extreme tension. Do not pull or drag an InfiniBand cable. Pulling on an InfiniBand cable can damage the cable internally.
- Unroll an InfiniBand cable for its length.
- Do not twist an InfiniBand cable more than one revolution for its entire length.
 Twisting an InfiniBand cable can damage the cable internally.
- Do not route InfiniBand cables where they can be stepped on, or experience rolling loads. A crushing effect can damage the cable internally.

1.2.3 Estimating InfiniBand Cable Path Lengths

Cable paths should be as short as possible. When the length of a cable path has been calculated, select the shortest cable to satisfy the length requirement. When specifying a cable, consider the following:

- Bends in the cable path increase the required length of the cable. Rarely does a
 cable travel in a straight line from connector to connector. Bends in the cable path
 are necessary, and each bend increases the total length.
- Bundling increases the required length of the cables. Bundling causes one or more
 cables to follow a common path. However, the bend radius is different in different
 parts of the bundle. If the bundle is large and unorganized, and there are many
 bends, one cable might experience only the inner radius of bends, while another
 cable might experience the outer radius of bends. In this situation, the differences
 of the required lengths of the cables is guite substantial.
- If you are routing the InfiniBand cable under the floor, consider the height of the raised floor when calculating cable path length.



1.2.4 Bundling InfiniBand Cables

When bundling InfiniBand cables in groups, use hook-and-loop straps to keep cables organized. If possible, use color-coordinated straps to help identify cables and their routing. The InfiniBand splitter and 4X copper conductor cables are fairly thick and heavy for their length. Consider the retention strength of the hook-and-loop straps when supporting cables. Bundle as few cables as reasonably possible. If the InfiniBand cables break free of their straps and fall free, the cables might break internally when they strike the floor or from sudden changes in tension.

You can bundle the cables using many hook-and-loop straps. Oracle recommends that no more than eight cables be bundled together.

Place the hook-and-loop straps as close together as reasonably possible, for example, one strap every foot (0.3 m). If a cable breaks free from a strap, then the cable cannot fall far before it is retained by another strap.

1.2.4.1 Floor and Underfloor Delivery of InfiniBand Cables

Sun Datacenter InfiniBand Switch 36 switch accepts InfiniBand cables from floor or underfloor delivery. Floor and underfloor delivery limits the tension in the InfiniBand cable to the weight of the cable for the rack height of the switch.



Overhead cabling details are not included in this guide. For details on overhead cabling, contact a certified service engineer.

1.2.5 Reviewing the Cable Management Arm Guidelines

Review the following cable management arm (CMA) guidelines before routing the cables:

- Remove all required cables from the packaging, and allow cables to acclimate or reach operating temperature, if possible. The acclimation period is usually 24 hours. This improves the ability to manipulate the cables.
- Label both ends of each cable using a label stock that meets the ANSI/TIA/EIA 606-A standard, if possible.
- Begin the installation procedure in ascending order.
- Only slide out one server at a time. Sliding out more than one server can cause cables to drop cause problems when sliding the servers back.
- Separate the installation by dressing cables with the least stringent bend radius requirements first. The following bend radius requirements are based on EIA/TIA 568-x standards, and may vary from the manufacturer's requirements:
 - CAT5e UTP: 4 x diameter of the cable or 1 inch/25.4 mm minimum bend radius
 - AC power cables: 4 x diameter of the cable or 1 inch/ 25.4 mm minimum bend radius



- TwinAx: 5 x diameter of the cable or 1.175 inch/33 mm.
- Quad Small Form-factor Pluggable (QSFP) InfiniBand cable: 6 x diameter of the cable or 2 inch/55 mm.
- Fiber core cable: 10 x diameter of the cable or 1.22 inch/31.75 mm for a 0.125 cable.
- Install the cables with the best longevity rate first.

1.2.6 Obtaining Current Configuration Information

The current configuration information is used to plan patching requirements, configure new IP addresses, and so on. The following information should be collected as described before extending the rack:

- The Exachk report for the current rack.
- Image history information using the following command:

```
dcli -g ~/all_group -l root "imagehistory" > imagehistory.txt
```

 Current IP addresses defined for all Exadata Storage Servers and database servers using the following command:

```
dcli -g ~/all_group -l root "ifconfig" > ifconfig_all.txt
```

 Information about the configuration of the cells, cell disks, flash logs, and IORM plans using the following commands:

```
dcli -g ~/cell_group -l root "cellcli -e list cell detail" > cell_detail.txt

dcli -g ~/cell_group -l root "cellcli -e list physicaldisk detail" > \
    physicaldisk_detail.txt

dcli -g ~/cell_group -l root "cellcli -e list griddisk attributes \
    name,offset,size,status,asmmodestatus,asmdeactivationoutcome" > griddisk.txt

dcli -g ~/cell_group -l root "cellcli -e list flashcache detail" > \
    fc_detail.txt

dcli -g ~/cell_group -l root "cellcli -e list flashlog detail" > fl_detail.txt

dcli -g ~/cell_group -l root "cellcli -e list iormplan detail" > \
    iorm_detail.txt
```

 HugePages memory configuration on the database servers using the following command:

InfiniBand switch information using the following command:

```
ibswitches > ibswitches.txt
```

- Firmware version of the Sun Datacenter InfiniBand Switch 36 switches using the nm2version command on each switch.
- The following network files from the first database server in the rack:
 - /etc/resolv.conf
 - /etc/ntp.conf



- /etc/network
- /etc/sysconfig/network-scripts/ifcfg-*
- Any users, user identifiers, groups and group identifiers created for clustermanaged services that need to be created on the new servers, such as Oracle GoldenGate.
 - /etc/passwd
 - /etc/group
- Output of current cluster status using the following command:

```
crsctl stat res -t > crs_stat.txt
```

 Patch information from the Grid Infrastructure and Oracle homes using the following commands. The commands must be run as Grid Infrastructure home owner, and the Oracle home owner.

In the preceding commands, *GRID_HOME* is the path for the Grid Infrastructure home directory, and *ORACLE_HOME* is the path for the Oracle home directory.

Related Topics

- Oracle Autonomous Health Framework User's Guide
- Oracle Exadata Database Machine exachk or HealthCheck (Doc ID 1070954.1)

1.2.7 Preparing the Network Configuration

When adding additional servers or rack to an existing rack, the IP addresses for the new servers are obtained using Oracle Exadata Deployment Assistant. If adding additional servers to an existing rack, then the application should only include the new servers. If adding an additional rack, then the new rack should use its own Oracle Exadata Deployment Assistant. The exact Oracle ASM disk group configuration currently in use may not be reflected by the application. This is not an issue, as the grid disks and disk groups are configured manually. All other items, such as the Oracle home location and owner, should be defined exactly as the existing configuration.

When adding Oracle Exadata X4-2 Database Server or later or Oracle Exadata Storage Server X4-2L or later, the bonding configuration must match the existing servers in the rack. The Oracle Exadata Deployment Assistant InfiniBand configuration page has an option to select the type of bonding. Select the option for active-active bonding, or deselect the option for active-passive bonding.

The configuration file generated by the application is used by Oracle Exadata Deployment Assistant. After completing Oracle Exadata Deployment Assistant, use the <code>checkip.sh</code> and <code>dbm.dat</code> files to verify the network configuration. The only errors that should occur are from the <code>ping</code> command to the SCAN addresses, Cisco switch, and Sun Datacenter InfiniBand Switch 36 switches.





Configuring Servers Manually

1.2.8 Moving Audit and Diagnostic Files

The files in the \$GRID HOME/rdbms/audit directory and the \$GRID HOME/log/diagnostics directory should be moved or deleted before extending a cluster. Oracle recommends moving or deleting the files a day or two before the planned extension because it may take time.

1.2.9 Reviewing Release and Patch Levels

The new rack or servers most-likely have a later release or patch level than the current rack. In some cases, you may want to update the current rack release to the later release. In other cases, you may want to stay at your current release, and choose to reimage the new rack to match the current rack. Whatever you choose to do, ensure that the existing and new servers and Sun Datacenter InfiniBand Switch 36 switches are at the same patch level. Note the following about the hardware and releases:



Tip:

Check My Oracle Support note 888828.1 for latest information on minimum releases.

- When expanding Oracle Exadata Database Machine X4-2 with Sun Server X4-2 Oracle Database Servers and Oracle Exadata Storage Server X4-2L Servers, the minimum release for the servers is release 11.2.3.3.0.
- When expanding with Oracle Exadata Database Machine X4-8 Full Rack, the minimum release for the servers is release 11.2.3.3.1.
- When expanding Oracle Exadata Database Machine X2-2 (with X4170 and X4275 servers) or Oracle Exadata Database Machine X2-2 (with X4170 M2 and X4270 M2 servers) with Sun Server X3-2 Oracle Database Servers and Exadata Storage Server X3-2 Servers, the minimum release for servers is release 11.2.3.2.0.
- When expanding Oracle Exadata Database Machine X2-2 (with X4170 and X4275 servers) with Sun Fire X4170 M2 Oracle Database Servers and Oracle Exadata Storage Server with Sun Fire X4270 M2 Servers, the minimum release for servers is release 11.2.2.2.0.
- The earlier servers may need to be patched to a later release to meet the minimum. In addition, earlier database servers may use Oracle Linux release 5.3. Those servers need to be updated to the latest Oracle Linux release.

Additional patching considerations include the Grid Infrastructure and database home releases and bundle patch updates. If new patches will be applied, then Oracle recommends changing the existing servers so the new servers will inherit the releases as part of the extension procedure. This way, the number of servers being patched is lower. Any patching of the existing servers should be performed in advance so they



are at the desired level when the extension work is scheduled, thereby reducing the total amount of work required during the extension.

Related Topics

- Exadata Database Machine and Exadata Storage Server Supported Versions (My Oracle Support Doc ID 888828.1)
- Updating key software components on database hosts to match those on the cells (My Oracle Support Doc ID 1284070.1)

1.2.10 Performing Preliminary Checks

Perform a visual check of Oracle Exadata Database Machine physical systems before extending the hardware.

- 1. Check the rack for damage.
- 2. Check the rack for loose or missing screws.
- 3. Check Oracle Exadata Database Machine for the ordered configuration.
- 4. Check that all cable connections are secure and well seated.
- Check power cables.
- Ensure the correct connectors have been supplied for the data center facility power source.
- 7. Check network data cables.
- 8. Check the site location tile arrangement for cable access and airflow.
- 9. Check the data center airflow into the front of Oracle Exadata Database Machine.

1.2.11 Preparing to Add Servers

Perform the following tasks before adding the servers:

- 1. Unpack the Oracle Exadata Database Machine expansion kit.
- 2. Unpack all Oracle Exadata Database Machine server components from the packing cartons. The following items should be packaged with the servers:
 - Oracle Database servers or Exadata Storage Server
 - Power cord, packaged with country kit
 - Cable management arm with installation instructions
 - Rackmount kit containing rack rails and installation instructions
 - (Optional) Sun server documentation and media kit



If you are extending Oracle Exadata Database Machine X4-2, Oracle Exadata Database Machine X3-8 Full Rack, or Oracle Exadata Database Machine X2-2 (with X4170 and X4275 servers) half rack, then order the expansion kit that includes a Sun Datacenter InfiniBand Switch 36 switch.



Figure 1-1 shows the components in the server expansion kit.




Figure 1-1 Server Components for Upgrade

- 3. Lay out the cables for the servers.
- 4. Unroll the cables and stretch them to remove the bends.
- 5. Apply the cable labels. Oracle recommends labeling all cables before installation.
- 6. Install the servers.
- Cable the servers.

See Also:

- Oracle Exadata Database Machine Maintenance Guide for information about cable labels
- "Adding New Servers" for information about installing the servers
- "Cabling Database Servers" and "Cabling Exadata Storage Servers" for information about cabling the servers

1.3 Extending the Hardware

Oracle Exadata Database Machine can be extended from Oracle Exadata Database Machine Quarter Rack to Oracle Exadata Database Machine Half Rack, from Oracle Exadata Database Machine Half Rack to Oracle Exadata Database Machine Full Rack, and by cabling racks together.



Note:

All new equipment receives a Customer Support Identifier (CSI). Any new equipment for Oracle Exadata Database Machine has a new CSI. Contact Oracle Support Services to reconcile the new CSI with the existing Oracle Exadata Database Machine CSI. Have the original instance numbers or serial numbers available, as well as the new numbers when contacting Oracle Support Services.

1.3.1 Extending an Eighth Rack to a Quarter Rack in Oracle Exadata Database Machine X4-2 and Later

Extending Oracle Exadata Database Machine X4-2 or X5-2 from an eighth rack to a quarter rack is done using software. No hardware modifications are needed to extend the rack.

However, hardware modifications may be needed for other Oracle Exadata Database Machine versions. See For Oracle Exadata Database Machine X6-2: Adding High Capacity Disks and Flash Cards and For Oracle Exadata Database Machine X7-2: Upgrading Eighth Rack Systems to a Quarter Rack for details.

This procedure can be done with no downtime or outages, other than a rolling database outage.

Note:

In the following procedures, the disk group names and sizes are examples. The values should be changed in the commands to match the actual system.

The procedures assume user equivalence exists between the root user on the first database server and all other database servers, and to the celladmin user on all storage cells.

The text files <code>cell_group</code> and <code>db_group</code> should be created to contain lists of cell host names and database server host names, respectively.

1.3.1.1 Reviewing and Validating Current Configuration of Eighth Rack Oracle Exadata Database Machine X4-2 or Later

The following procedure describes how to review and validate the current configuration.

- 1. Log in as the root user on the first database server.
- 2. Review the current configuration of the database servers using the following command:

dcli -g db_group -l root 'dbmcli -e list dbserver attributes coreCount'



The following is an example of the output from the command for Oracle Exadata Database Machine X5-2 Eighth Rack:

dm01db01: 18
dm01db02: 18



The number of active cores in Oracle Exadata Database Machine X5-2 Eighth Rack database server is 18. The number of active cores in Oracle Exadata Database Machine X4-2 Eighth Rack database server is 12.

If the number of cores on a database server configured as an eighth rack differs, then contact Oracle Support Services.

3. Review the current configuration of the storage servers using the following command. The expected output is ${\tt TRUE}$.

dcli -g cell_group -l celladmin 'cellcli -e LIST CELL attributes eighthrack'

1.3.1.2 Activating Database Server Cores in Eighth Rack Oracle Exadata Database Machine X4-2 or Later

The following procedure describes how to activate the database server cores.

- 1. Log in as the root user on the first database server.
- 2. Activate all the database server cores using the following dcli utility command on the database server group:

```
# dcli -g db_group -l root 'dbmcli -e \
ALTER DBSERVER pendingCoreCount = number_of_cores'
```

In the preceding command, *number_of_cores* is the total number of cores to activate. The value includes the existing core count and the additional cores to be activated. The following command shows how to activate all the cores in Oracle Exadata Database Machine X5-2 Eighth Rack:

dcli -g db_group -l root 'dbmcli -e ALTER DBSERVER pendingCoreCount = 36'



The maximum number of total active cores for Oracle Exadata Database Machine X5-2 Eighth Rack is 36. The maximum number of total active cores for Oracle Exadata Database Machine X4-2 Eighth Rack is 24.

3. Restart each database server.



Note:

If this procedure is done in a rolling fashion with the database and Grid Infrastructure active, then ensure the following before restarting the database server:

- All Oracle ASM grid disks are online...
- No Oracle ASM rebalance operations are active. You can query the V\$ASM_OPERATION view for the status of the rebalance operation.
- Shut down the database and Grid Infrastructure in a controlled manner, failing over services as needed.
- 4. Verify the following items on the database server after the restart completes and before proceeding to the next server:
 - The database and Grid Infrastructure services are active.
 - See "Using SRVCTL to Verify That Instances are Running" in *Oracle Real Application Clusters Administration and Deployment Guide* and the crsctl status resource -w "TARGET = ONLINE" -t command.
 - The number of active cores is correct. Use the <code>dbmcli</code> -e list <code>dbserver</code> attributes <code>coreCount</code> command to verify the number of cores.

See Also:

- "Changing a Disk to Offline or Online." in *Oracle Exadata System Software User's Guide*
- "Stopping One or More Instances and Oracle RAC Databases Using SRVCTL" in Oracle Real Application Clusters Administration and Deployment Guide
- "Stopping Oracle Clusterware" in *Oracle Database 2 Day + Real Application Clusters Guide*
- Oracle Exadata Database Machine Maintenance Guide for additional information about activating a subset of cores
- Oracle Exadata Database Machine Licensing Information User's Guide for information about licensing a subset of cores

1.3.1.3 For Oracle Exadata Database Machine X6-2: Adding High Capacity Disks and Flash Cards

Upgrade of Oracle Exadata Database Machine X6-2 Eighth Rack High Capacity systems require hardware modification, but upgrade of X6-2 Extreme Flash does not require hardware modification.

Eighth Rack High Capacity storage servers have half the cores enabled, but half the disks and flash cards are removed. Eighth Rack Extreme Flash storage servers have half the cores and flash drives enabled.



Eighth Rack database servers have half the cores enabled.

On Oracle Exadata Database Machine X6-2 Eighth Rack systems with High Capacity disks, you can add high capacity disks and flash cards to extend the system to a Quarter Rack:

- 1. Install the six 8 TB disks in HDD slots 6 11.
- 2. Install the two F320 flash cards in PCIe slots 1 and 4.

1.3.1.4 For Oracle Exadata Database Machine X7-2: Upgrading Eighth Rack Systems to a Quarter Rack

Upgrade of Oracle Exadata Database Machine X7-2 Eighth Rack systems requires hardware modification. Eighth Rack database servers have one of the CPUs removed, and all of the memory for CPU1 was moved to CPU0. Storage servers have half the cores enabled, but half the disks and flash cards were removed.

On Oracle Exadata Database Machine X7-2 Eighth Rack systems, you can add CPUs, high capacity disks and flash cards to extend the system to a Quarter Rack:

- 1. On the Exadata X7 database server, install CPU1, move half of CPU0's memory to CPU1, and move the 10/25GbE PCI card to PCIe slot 1.
- On Exadata High Capacity Storage Servers, install six 10 TB high capacity SAS disks in HDD6-11 and two F640 flash cards in PCIe slots 4 and 6.
- 3. On Exadata Extreme Flash Storage Server, install four F640 flash cards in PCIe slots 2,3,8, and 9.

1.3.1.5 Activating Storage Server Cores and Disks in Eighth Rack Oracle Exadata Database Machine X4-2 or Later

The following procedure describes how to activate the storage server cores and disks.

- 1. Log in as the root user on the first database server.
- 2. Activate the cores on the storage server group using the following command. The command uses the dcli utility, and runs the command as the celladmin user.

```
# dcli -g cell_group -l celladmin cellcli -e "alter cell eighthRack=false"
```

3. Create the cell disks using the following command:

```
# dcli -g cell_group -l celladmin cellcli -e "create celldisk all"
```

4. Recreate the flash log using the following commands:

```
# dcli -g cell_group -l celladmin cellcli -e "drop flashlog all force"
# dcli -g cell_group -l celladmin cellcli -e "create flashlog all"
```

5. Expand the flash cache using the following command:

```
# dcli -q cell_group -l celladmin cellcli -e "alter flashcache all"
```

1.3.1.6 Creating Grid Disks in Eighth Rack Oracle Exadata Database Machine X4-2 or Later

Grid disk creation must follow a specific order to ensure the proper offset.



The order of grid disk creation must follow the same sequence that was used during initial grid disks creation. For a standard deployment using Oracle Exadata Deployment Assistant, the order is DATA, RECO, and DBFS_DG. Create all DATA grid disks first, followed by the RECO grid disks, and then the DBFS_DG grid disks.

The following procedure describes how to create the grid disks:



The commands shown in this procedure use the standard deployment grid disk prefix names of DATA, RECO and DBFS_DG. The sizes being checked are on cell disk 02. Cell disk 02 is used because the disk layout for cell disks 00 and 01 are different from the other cell disks in the server.

1. Check the size of the grid disks using the following commands. Each cell should return the same size for the grid disks starting with the same grid disk prefix.

```
# dcli -g cell_group -l celladmin cellcli -e
"list griddisk attributes name, size where name like \'DATA.*_02_.*\'"

# dcli -g cell_group -l celladmin cellcli -e
"list griddisk attributes name, size where name like \'RECO.*_02_.*\'"

# dcli -g cell_group -l celladmin cellcli -e
"list griddisk attributes name, size where name like \'DBFS_DG.*_02_.*\'"
```

The sizes shown are used during grid disk creation.

2. Create the grid disks for the disk groups using the sizes shown in step 1. Table 1-1 shows the commands to create the grid disks based on rack type and disk group.



Table 1-1 Commands to Create Disk Groups When Extending Oracle Exadata Database Machine X4-2 Eighth Rack or Later

Rack	Commands
Extreme Flash Oracle Exadata Database Machine	<pre>dcli -g cell_group -l celladmin "cellcli -e create griddisk \ DATA_FD_04_\'hostname -s\' celldisk=FD_04_\'hostname -s\',size=datasize"</pre>
X5-2 and later	dcli -g cell_group -l celladmin "cellcli -e create griddisk \ DATA_FD_05_\'hostname -s\' celldisk=FD_05_\'hostname -s\',size=datasize"
	dcli -g cell_group -l celladmin "cellcli -e create griddisk \ DATA_FD_06_\'hostname -s\' celldisk=FD_06_\'hostname -s\',size=datasize"
	dcli -g cell_group -l celladmin "cellcli -e create griddisk \ DATA_FD_07_\'hostname -s\' celldisk=FD_07_\'hostname -s\',size=datasize"
	<pre>dcli -g cell_group -l celladmin "cellcli -e create griddisk \ RECO_FD_04_\'hostname -s\' celldisk=FD_04_\'hostname -s\',size=recosize, \ cachingPolicy=none"</pre>
	<pre>dcli -g cell_group -l celladmin "cellcli -e create griddisk \ RECO_FD_05_\'hostname -s\' celldisk=FD_05_\'hostname -s\',size=recosize, \ cachingPolicy=none"</pre>
	<pre>dcli -g cell_group -l celladmin "cellcli -e create griddisk \ RECO_FD_06_\'hostname -s\' celldisk=FD_06_\'hostname -s\',size=recosize, \ cachingPolicy=none"</pre>
	<pre>dcli -g cell_group -l celladmin "cellcli -e create griddisk \ RECO_FD_07_\'hostname -s\' celldisk=FD_07_\'hostname -s\',size=recosize, \ cachingPolicy=none"</pre>
	<pre>dcli -g cell_group -l celladmin "cellcli -e create griddisk \ DBFS_DG_FD_04_\'hostname -s\' celldisk=FD_04_\'hostname -s\',size=dbfssize, \ cachingPolicy=none"</pre>
	<pre>dcli -g cell_group -l celladmin "cellcli -e create griddisk \ DBFS_DG_FD_05_\'hostname -s\' celldisk=FD_05_\'hostname -s\',size=dbfssize, \ cachingPolicy=none"</pre>
	<pre>dcli -g cell_group -l celladmin "cellcli -e create griddisk \ DBFS_DG_FD_06_\'hostname -s\' celldisk=FD_06_\'hostname -s\',size=dbfssize, \ cachingPolicy=none"</pre>
	<pre>dcli -g cell_group -l celladmin "cellcli -e create griddisk \ DBFS_DG_FD_07_\'hostname -s\' celldisk=FD_07_\'hostname -s\',size=dbfssize, \ cachingPolicy=none"</pre>



Table 1-1 (Cont.) Commands to Create Disk Groups When Extending Oracle Exadata Database Machine X4-2 Eighth Rack or Later

Rack	Commands
High Capacity Oracle Exadata Database Machine X5-2 or	dcli -g cell_group -l celladmin "cellcli -e create griddisk \ DATA_CD_06_\'hostname -s\' celldisk=CD_06_\'hostname -s\',size=datasize"
Oracle Exadata Database Machine X4-2 and later	<pre>dcli -g cell_group -l celladmin "cellcli -e create griddisk \ DATA_CD_07_\'hostname -s\' celldisk=CD_07_\'hostname -s\',size=datasize"</pre>
	<pre>dcli -g cell_group -l celladmin "cellcli -e create griddisk \ DATA_CD_08_\'hostname -s\' celldisk=CD_08_\'hostname -s\',size=datasize"</pre>
	<pre>dcli -g cell_group -l celladmin "cellcli -e create griddisk \ DATA_CD_09_\'hostname -s\' celldisk=CD_09_\'hostname -s\',size=datasize"</pre>
	dcli -g cell_group -l celladmin "cellcli -e create griddisk \ DATA_CD_10_\'hostname -s\' celldisk=CD_10_\'hostname -s\',size=datasize"
	dcli -g cell_group -l celladmin "cellcli -e create griddisk \ DATA_CD_11_\'hostname -s\' celldisk=CD_11_\'hostname -s\',size=datasize"
	<pre>dcli -g cell_group -l celladmin "cellcli -e create griddisk \ RECO_CD_06_\'hostname -s\' celldisk=CD_06_\'hostname -s\',size=recosize, \ cachingPolicy=none"</pre>
	<pre>dcli -g cell_group -l celladmin "cellcli -e create griddisk \ RECO_CD_07_\'hostname -s\' celldisk=CD_07_\'hostname -s\',size=recosize, \ cachingPolicy=none"</pre>
	<pre>dcli -g cell_group -l celladmin "cellcli -e create griddisk RECO_CD_08_\'hostname -s\' celldisk=CD_08_\'hostname -s\',size=recosize, \ cachingPolicy=none"</pre>
	<pre>dcli -g cell_group -l celladmin "cellcli -e create griddisk \ RECO_CD_09_\'hostname -s\' celldisk=CD_09_\'hostname -s\',size=recosize, \ cachingPolicy=none"</pre>
	<pre>dcli -g cell_group -l celladmin "cellcli -e create griddisk \ RECO_CD_10_\'hostname -s\' celldisk=CD_10_\'hostname -s\',size=recosize, \ cachingPolicy=none"</pre>
	<pre>dcli -g cell_group -l celladmin "cellcli -e create griddisk \ RECO_CD_11_\'hostname -s\' celldisk=CD_11_\'hostname -s\',size=recosize, \ cachingPolicy=none"</pre>
	<pre>dcli -g cell_group -l celladmin "cellcli -e create griddisk \ DBFS_DG_CD_06_\'hostname -s\' celldisk=CD_06_\'hostname -s\',size=dbfssize, \ cachingPolicy=none"</pre>
	<pre>dcli -g cell_group -l celladmin "cellcli -e create griddisk \ DBFS_DG_CD_07_\'hostname -s\' celldisk=CD_07_\'hostname -s\',size=dbfssize, \ cachingPolicy=none"</pre>
	<pre>dcli -g cell_group -l celladmin "cellcli -e create griddisk \ DBFS_DG_CD_08_\'hostname -s\' celldisk=CD_08_\'hostname -s\',size=dbfssize, \ cachingPolicy=none"</pre>

Table 1-1 (Cont.) Commands to Create Disk Groups When Extending Oracle Exadata Database Machine X4-2 Eighth Rack or Later

Rack	Commands
	<pre>dcli -g cell_group -l celladmin "cellcli -e create griddisk \ DBFS_DG_CD_09_\'hostname -s\' celldisk=CD_09_\'hostname -s\',size=dbfssize, \ cachingPolicy=none"</pre>
	<pre>dcli -g cell_group -l celladmin "cellcli -e create griddisk \ DBFS_DG_CD_10_\'hostname -s\' celldisk=CD_10_\'hostname -s\',size=dbfssize, \ cachingPolicy=none"</pre>
	<pre>dcli -g cell_group -l celladmin "cellcli -e create griddisk \ DBFS_DG_CD_11_\'hostname -s\' celldisk=CD_11_\'hostname -s\',size=dbfssize, \ cachingPolicy=none"</pre>

1.3.1.7 Adding Grid Disks to Oracle ASM Disk Groups in Eighth Rack Oracle Exadata Database Machine X4-2 or Later

The following procedure describes how to add the grid disks to Oracle ASM disk groups.

The grid disks created in "Creating Grid Disks in Eighth Rack Oracle Exadata Database Machine X4-2 or Later" must be added as Oracle ASM disks to their corresponding, existing Oracle ASM disk groups.

- 1. Validate the following:
 - No rebalance operation is currently running.
 - All Oracle ASM disks are active.
- 2. Log in to the first database server as the owner who runs the Grid Infrastructure software.
- 3. Set the environment to access the +ASM instance on the server.
- 4. Log in to the ASM instance as the sysasm user using the following command:

```
$ sqlplus / as sysasm
```

5. Validate the current settings, as follows:

```
SQL> set lines 100

SQL> column attribute format a20

SQL> column value format a20

SQL> column diskgroup format a20

SQL> SELECT att.name attribute, upper(att.value) value, dg.name diskgroup

FROM V$ASM_ATTRIBUTE att, V$ASM_DISKGROUP DG

WHERE DG.group_number=att.group_number AND att.name LIKE '%appliance.mode%'

ORDER BY att.group number;
```

The output should be similar to the following:

ATTRIBUTE	VALUE	DISKGROUP
appliance.mode	TRUE	DATAC1



```
appliance.mode TRUE DBFS_DG appliance.mode TRUE RECOC1
```

6. Disable the appliance.mode attribute for any disk group that shows TRUE using the following commands:

```
SQL> ALTER DISKGROUP data_diskgroup set attribute 'appliance.mode'='FALSE';
SQL> ALTER DISKGROUP reco_diskgroup set attribute 'appliance.mode'='FALSE';
SQL> ALTER DISKGROUP dbfs_dg_diskgroup set attribute 'appliance.mode'='FALSE';
```

In the preceding commands, data_diskgroup, reco_diskgroup, and dbfs_dg_diskgroup are the names of the DATA, RECO and DBFS_DG disk groups, respectively.

Add the grid disks to the Oracle ASM disk groups. Table 1-2 shows the commands
to create the grid disks based on rack type and disk group. Adding the new disks
requires a rebalance of the system.

Table 1-2 Commands to Add Disk Groups When Extending Eighth Rack Oracle Exadata Database Machine X4-2 and Later

Rack	Commands
Extreme Flash Oracle Exadata Database Machine X5-2 and later	SQL> ALTER DISKGROUP data_diskgroup ADD DISK 'o/*/DATA_FD_0[4-7]*' \ REBALANCE POWER 32;
	SQL> ALTER DISKGROUP reco_diskgroup ADD DISK 'o/*/RECO_FD_0[4-7]*' \ REBALANCE POWER 32;
	SQL> ALTER DISKGROUP dbfs_dg_diskgroup ADD DISK 'o/*/DBFS_DG_FD_0[4-7]*'\ REBALANCE POWER 32;
High Capacity Oracle Exadata Database Machine X5-2 or Oracle Exadata Database Machine X4-2 and later	SQL> ALTER DISKGROUP data_diskgroup ADD DISK 'o/*/DATA_CD_0[6-9]*',' o/*/DATA_CD_1[0-1]*' REBALANCE POWER 32;
	SQL> ALTER DISKGROUP reco_diskgroup ADD DISK 'o/*/RECO_CD_0[6-9]*',' o/*/RECO_CD_1[0-1]*' REBALANCE POWER 32;
	SQL> ALTER DISKGROUP dbfs_dg_diskgroup ADD DISK ' \ o/*/DBFS_DG_CD_0[6-9]*',' o/*/ DBFS_DG_CD_1[0-1]*' REBALANCE POWER 32;

The preceding commands return Diskgroup altered, if successful.

8. (Optional) Monitor the current rebalance operation using the following command:

```
SQL> SELECT * FROM gv$asm_operation;
```

9. Re-enable the appliance.mode attribute, if it was disabled in step 6 using the following commands:



```
SQL> ALTER DISKGROUP data_diskgroup set attribute 'appliance.mode'='TRUE';
SQL> ALTER DISKGROUP reco_diskgroup set attribute 'appliance.mode'='TRUE';
SQL> ALTER DISKGROUP dbfs_dq_diskgroup set attribute 'appliance.mode'='TRUE';
```

1.3.1.8 Validating New Quarter Rack Configuration for Oracle Exadata Database Machine X4-2 or Later

After adding the grid disks to the Oracle ASM disk groups, validate the configuration.

- 1. Log in as the root user on the first database server.
- 2. Check the core count using the following command:

```
# dcli -g db_group -l root 'dbmcli -e list dbserver attributes coreCount'
```

3. Review the storage server configuration using the following command.

```
# dcli -g cell_group -l celladmin 'cellcli -e list cell attributes eighthrack'
```

The output should show FALSE.

4. Review the appliance mode for each disk group using the following commands:

```
SQL> set lines 100
SQL> column attribute format a20
SQL> column value format a20
SQL> column diskgroup format a20
SQL> SELECT att.name attribute, upper(att.value) value, dg.name diskgroup
FROM V$ASM_ATTRIBUTE att, V$ASM_DISKGROUP DG
WHERE DG.group_number = att.group_number AND
att.name LIKE '%appliance.mode%' ORDER BY DG.group_number;
```

5. Validate the number of Oracle ASM disks using the following command:

```
SQL> SELECT g.name,d.failgroup,d.mode_status,count(*)
FROM v$asm_diskgroup g, v$asm_disk d
WHERE d.group_number=g.group_number
QROUP BY g.name,d.failgroup,d.mode_status;
```

NAME COUNT(*)	FAILGROUP		MODE_ST	
DATAC1	EXA01CELADM01	ONLINE	12	
DATAC1	EXA01CELADM02	ONLINE	12	
DATAC1	EXA01CELADM03	ONLINE	12	
RECOC1	EXA01CELADM01	ONLINE	12	
RECOC1	EXA01CELADM02	ONLINE	12	
RECOC1	EXA01CELADM03	ONLINE	12	
RECOC2	EXA01CELADM01	ONLINE	12	
RECOC2	EXA01CELADM02	ONLINE	12	
RECOC2	EXA01CELADM03	ONLINE	12	
DBFS_DG	EXA01CELADM01	ONLINE	10	
DBFS_DG	EXA01CELADM02	ONLINE	10	
DBFS_DG	EXA01CELADM03	ONLINE	10	

All two-socket systems (except eighth rack configurations) will have 12 disks per cell for any system model. Eighth rack configurations will have 6 disks per cell.

1.3.2 Extending an Eighth Rack to a Quarter Rack in Oracle Exadata Database Machine X3-2

Extending Oracle Exadata Database Machine X3-2 or earlier rack from an eighth rack to a quarter rack is done using software. No hardware modifications are needed to extend the rack. This procedure can be done with no downtime or outages, other than a rolling database outage. The following procedures in this section describe how to extend an Oracle Exadata Database Machine X3-2 eighth rack to a quarter rack:

1.3.2.1 Reviewing and Validating Current Configuration of Oracle Exadata Database Machine X3-2 Eighth Rack

The following procedure describes how to review and validate the current configuration:

- 1. Log in as the root user on the first database server.
- Review the current configuration of the database servers using the following command:

```
# dcli -g db_group -l root /opt/oracle.SupportTools/resourcecontrol -show
```

The following is an example of the output from the command:

```
dm01db01: [INFO] Validated hardware and OS. Proceed.
dm01db01:
dm01db01: system_bios_version: 25010600
dm01db01: restore_status: Ok
dm01db01: config_sync_status: Ok
dm01db01: reset_to_defaults: Off
dm01db01: [SHOW] Number of cores active per socket: 4
dm01db02: [INFO] Validated hardware and OS. Proceed.
dm01db02:
dm01db02: system_bios_version: 25010600
dm01db02: restore_status: Ok
dm01db02: config_sync_status: Ok
dm01db02: reset_to_defaults: Off
dm01db02: [SHOW] Number of cores active per socket: 4
```

Note:

The number of active cores in Oracle Exadata Database Machine X3-2 Eighth Rack database server is 4.

If the number of cores on a database server configured as an eighth rack differs, then contact Oracle Support Services.

Ensure the output for restore_status and config_sync_status are shown as Ok before continuing this procedure.

3. Review the current configuration of the storage servers using the following command. The expected output is TRUE.

```
# dcli -g cell_group -l celladmin 'cellcli -e LIST CELL attributes eighthrack'
```



4. Ensure that flash disks are not used in Oracle ASM disk groups using the following command. Flash cache is dropped and recreated during this procedure:

```
# dcli -g cell_group -l celladmin cellcli -e "list griddisk attributes \
asmDiskgroupName,asmDiskName,diskType where diskType ='FlashDisk' \
and asmDiskgroupName !=null"
```

No rows should be returned by the command.

1.3.2.2 Activating Database Server Cores in Oracle Exadata Database Machine X3-2 Eighth Rack

The following procedure describes how to activate the database server cores:

- 1. Log in as the root user on the first database server.
- 2. Activate all the database server cores using the following dcli utility command on the database server group:

```
# dcli -g db_group -l root /opt/oracle.SupportTools/resourcecontrol
-core number_of_cores
```

In the preceding command, *number_of_cores* is the total number of cores to activate. To activate all the cores, enter All for the number of cores.

See Also:

- Oracle Exadata Database Machine Maintenance Guide for additional information about activating a subset of cores
- Oracle Exadata Database Machine Licensing Information for information about licensing a subset of cores
- 3. Restart the database servers in a rolling manner using the following command:

reboot

Note:

Ensure the output for <code>restore_status</code> and <code>config_sync_status</code> are shown as <code>ok</code> before activating the storage server cores and disks. Getting the status from the BIOS after restarting may take several minutes.

1.3.2.3 Activating Storage Server Cores and Disks in Oracle Exadata Database Machine X3-2 Eighth Rack

The following procedure describes how to activate the storage server cores and disks:

- 1. Log in as the root user on the first database server.
- 2. Activate the cores on the storage server group using the following command. The command uses the dcli utility, and runs the command as the celladmin user.



```
# dcli -g cell_group -l celladmin cellcli -e "alter cell eighthRack=false"
```

3. Create the cell disks using the following command:

```
# dcli -g cell_group -l celladmin cellcli -e "create celldisk all"
```

4. Recreate the flash log using the following commands:

```
# dcli -g cell_group -l celladmin cellcli -e "drop flashlog all force"
# dcli -g cell_group -l celladmin cellcli -e "create flashlog all"
```

5. Expand the flash cache using the following command:

```
# dcli -q cell_group -l celladmin cellcli -e "alter flashcache all"
```

1.3.2.4 Creating Grid Disks in Oracle Exadata Database Machine X3-2 Eighth Rack

Grid disk creation must follow a specific order to ensure the proper offset. The order of grid disk creation must follow the same sequence that was used during initial grid disks creation. For a standard deployment using Oracle Exadata Deployment Assistant, the order is DATA, RECO, and DBFS_DG. Create all DATA grid disks first, followed by the RECO grid disks, and then the DBFS_DG grid disks.

The following procedure describes how to create the grid disks:

Note:

The commands shown in this procedure use the standard deployment grid disk prefix names of DATA, RECO and DBFS_DG. The sizes being checked are on cell disk 02. Cell disk 02 is used because the disk layout for cell disks 00 and 01 are different from the other cell disks in the server.

1. Check the size of the grid disks using the following commands. Each cell should return the same size for the grid disks starting with the same grid disk prefix.

```
# dcli -g cell_group -l celladmin cellcli -e
"list griddisk attributes name, size where name like \'DATA.*02.*\'"

# dcli -g cell_group -l celladmin cellcli -e
"list griddisk attributes name, size where name like \'RECO.*02.*\'"

# dcli -g cell_group -l celladmin cellcli -e
"list griddisk attributes name, size where name like \'DBFS_DG.*02.*\'"
```

The sizes shown are used during grid disk creation.

2. Create the grid disks for the disk groups using the sizes shown in step 1. Table 1-3 shows the commands to create the grid disks based on rack type and disk group.



Table 1-3 Commands to Create Disk Groups When Extending Oracle Exadata Database Machine X3-2 Eighth Rack

Rack	Commands
High Performance or High Capacity Oracle Exadata Database	<pre>dcli -g cell_group -l celladmin "cellcli -e create griddisk \ DATA_CD_06_\`hostname -s\` celldisk=CD_06_\`hostname -s\`,size=datasize"</pre>
Machine X3-2	<pre>dcli -g cell_group -l celladmin "cellcli -e create griddisk \ DATA_CD_07_\`hostname -s\` celldisk=CD_07_\`hostname -s\`,size=datasize"</pre>
	<pre>dcli -g cell_group -l celladmin "cellcli -e create griddisk \ DATA_CD_08_\`hostname -s\` celldisk=CD_08_\`hostname -s\`,size=datasize"</pre>
	<pre>dcli -g cell_group -l celladmin "cellcli -e create griddisk \ DATA_CD_09_\`hostname -s\` celldisk=CD_09_\`hostname -s\`,size=datasize"</pre>
	<pre>dcli -g cell_group -l celladmin "cellcli -e create griddisk \ DATA_CD_10_\`hostname -s\` celldisk=CD_10_\`hostname -s\`,size=datasize"</pre>
	<pre>dcli -g cell_group -l celladmin "cellcli -e create griddisk \ DATA_CD_11_\`hostname -s\` celldisk=CD_11_\`hostname -s\`,size=datasize"</pre>
	<pre>dcli -g cell_group -l celladmin "cellcli -e create griddisk \ RECO_CD_06_\`hostname -s\` celldisk=CD_06_\`hostname -s\`,size=recosize"</pre>
	<pre>dcli -g cell_group -l celladmin "cellcli -e create griddisk \ RECO_CD_07_\`hostname -s\` celldisk=CD_07_\`hostname -s\`,size=recosize"</pre>
	<pre>dcli -g cell_group -l celladmin "cellcli -e create griddisk \ RECO_CD_08_\`hostname -s\` celldisk=CD_08_\`hostname -s\`,size=recosize"</pre>
	<pre>dcli -g cell_group -l celladmin "cellcli -e create griddisk \ RECO_CD_09_\`hostname -s\` celldisk=CD_09_\`hostname -s\`,size=recosize"</pre>
	<pre>dcli -g cell_group -l celladmin "cellcli -e create griddisk \ RECO_CD_10_\`hostname -s\` celldisk=CD_10_\`hostname -s\`,size=recosize"</pre>
	<pre>dcli -g cell_group -l celladmin "cellcli -e create griddisk \ RECO_CD_11_\`hostname -s\` celldisk=CD_11_\`hostname -s\`,size=recosize"</pre>
	<pre>dcli -g cell_group -l celladmin "cellcli -e create griddisk \ DBFS_DG_CD_06_\`hostname -s\` celldisk=CD_06_\`hostname -s\`,size=dbfssize"</pre>
	<pre>dcli -g cell_group -l celladmin "cellcli -e create griddisk \ DBFS_DG_CD_07_\`hostname -s\` celldisk=CD_07_\`hostname -s\`,size=dbfssize"</pre>
	<pre>dcli -g cell_group -l celladmin "cellcli -e create griddisk \ DBFS_DG_CD_08_\`hostname -s\` celldisk=CD_08_\`hostname -s\`,size=dbfssize"</pre>
	<pre>dcli -g cell_group -l celladmin "cellcli -e create griddisk \ DBFS_DG_CD_09_\`hostname -s\` celldisk=CD_09_\`hostname -s\`,size=dbfssize"</pre>
	<pre>dcli -g cell_group -l celladmin "cellcli -e create griddisk \ DBFS_DG_CD_10_\`hostname -s\` celldisk=CD_10_\`hostname -s\`,size=dbfssize"</pre>
	<pre>dcli -g cell_group -l celladmin "cellcli -e create griddisk \ DBFS_DG_CD_11_\`hostname -s\` celldisk=CD_11_\`hostname -s\`,size=dbfssize"</pre>



1.3.2.5 Adding Grid Disks to Oracle ASM Disk Groups in Oracle Exadata Database Machine X3-2 Eighth Rack

The grid disks created in "Creating Grid Disks in Oracle Exadata Database Machine X3-2 Eighth Rack" must be added as Oracle ASM disks to their corresponding, existing Oracle ASM disk groups. The following procedure describes how to add the grid disks to Oracle ASM disk groups:

- 1. Validate the following:
 - No rebalance operation is currently running.
 - All Oracle ASM disks are active.
- Log in to the first database server as the owner who runs the Grid Infrastructure software.
- 3. Set the environment to access the +ASM instance on the server.
- 4. Log in to the ASM instance as the sysasm user using the following command:

```
$ sqlplus / as sysasm
```

5. Validate the current settings, as follows:

```
SQL> set lines 100
SQL> column attribute format a20
SQL> column value format a20
SQL> column diskgroup format a20
SQL> SELECT att.name attribute, upper(att.value) value, dg.name diskgroup \
FROM V$ASM_ATTRIBUTE att, V$ASM_DISKGROUP DG \
WHERE DG.group_number = att.group_number AND \
att.name LIKE '%appliance.mode%' ORDER BY att.group_number;
```

The output should be similar to the following:

ATTRIBUTE	VALUE	DISKGROUP
appliance.mode	TRUE	DATAC1
appliance.mode	TRUE	DBFS_DG
appliance.mode	TRUE	RECOC1

6. Disable the appliance.mode attribute for any disk group that shows TRUE using the following commands:

```
SQL> ALTER DISKGROUP data_diskgroup set attribute 'appliance.mode'='FALSE';
SQL> ALTER DISKGROUP reco_diskgroup set attribute 'appliance.mode'='FALSE';
SQL> ALTER DISKGROUP dbfs_dg_diskgroup set attribute 'appliance.mode'='FALSE';
```

In the preceding commands, data_diskgroup, reco_diskgroup, and dbfs_dg_diskgroup are the names of the DATA, RECO and DBFS_DG disk groups, respectively.

Add the grid disks to the Oracle ASM disk groups. Table 1-4 shows the commands
to create the grid disks based on rack type and disk group. Adding the new disks
requires a rebalance of the system.



Table 1-4 Commands to Add Disk Groups When Extending an Oracle Exadata Database Machine X3-2 Eighth Rack

Rack	Commands
High Capacity or High Performance Oracle Exadata Database Machine X3-2	SQL> ALTER DISKGROUP data_diskgroup ADD DISK 'o/*/DATA_CD_0[6-9]*',' o/*/DATA_CD_1[0-1]*' REBALANCE POWER 32;
	SQL> ALTER DISKGROUP reco_diskgroup ADD DISK 'o/*/RECO_CD_0[6-9]*',' o/*/RECO_CD_1[0-1]*' REBALANCE POWER 32;
	SQL> ALTER DISKGROUP dbfs_dg_diskgroup ADD DISK '

The preceding commands return Diskgroup altered, if successful.

8. (Optional) Monitor the current rebalance operation using the following command:

```
SQL> SELECT * FROM gv$asm_operation;
```

Re-enable the appliance.mode attribute, if it was disabled in step 6 using the following commands:

```
SQL> ALTER DISKGROUP data_diskgroup set attribute 'appliance.mode'='TRUE';
SQL> ALTER DISKGROUP recodiskgroup set attribute 'appliance.mode'='TRUE';
SQL> ALTER DISKGROUP dbfs_dg_diskgroup set attribute 'appliance.mode'='TRUE';
```

1.3.2.6 Validating New Oracle Exadata Database Machine X3-2 Quarter Rack Configuration

After adding the grid disks to the Oracle ASM disk groups, validate the configuration. The following procedure describes how to validate the configuration:

- 1. Log in as the root user on the first database server.
- 2. Check the core count using the following command:

```
# dcli -g db_group -l root 'dbmcli -e list dbserver attributes coreCount'
```

3. Review the storage server configuration using the following command.

```
# dcli -g cell_group -l celladmin 'cellcli -e list cell attributes eighthrack'
```

The output should show False.

4. Review the appliance mode for each disk group using the following commands:

```
SQL> set lines 100

SQL> column attribute format a20

SQL> column value format a20

SQL> column diskgroup format a20

SQL> SELECT att.name attribute, upper(att.value) value, dg.name diskgroup

FROM V$ASM_ATTRIBUTE att, V$ASM_DISKGROUP DG
```



```
WHERE DG.group_number =att.group_number AND att.name LIKE '%appliance.mode%' ORDER BY DG.group_number;
```

5. Validate the number of Oracle ASM disks using the following command:

```
SQL> SELECT g.name,d.failgroup,d.mode_status,count(*)
FROM v$asm_diskgroup g, v$asm_disk d
WHERE d.group_number=g.group_number
GROUP BY g.name,d.failgroup,d.mode status;
```

1.3.3 Extending a Quarter Rack or Half Rack

Extending Oracle Exadata Database Machine from quarter rack to half rack, or half rack to full rack consists of adding new hardware to the rack. The following sections describe how to extend Oracle Exadata Database Machine with new servers:



It is possible to extend the hardware while the machine is online, and with no downtime. However, extreme care should be taken. In addition, patch application to existing switches and servers should be done before extending the hardware.

1.3.3.1 Removing the Doors

The following procedure describes how to remove the doors on Oracle Exadata Database Machine.



For Exadata X7 systems, refer to "Remove the Doors" in *Oracle Rack Cabinet 1242 User's Guide* at https://docs.oracle.com/cd/E85660_01/html/E87280/gshfw.html#scrolltoc

- 1. Remove the Oracle Exadata Database Machine front and rear doors, as follows:
 - a. Unlock the front and rear doors. The key is in the shipping kit.
 - **b.** Open the doors.
 - c. Detach the grounding straps connected to the doors by pressing down on the tabs of the grounding strap's quick-release connectors, and pull the straps from the frame.
 - d. Lift the doors up and off their hinges.






Figure 1-2 Removing the Rack Doors

Description of callouts in Figure 1-2:

- 1: Detaching the grounding cable.
- 2: Top rear hinge.
- 3: Bottom rear hinge.
- 4: Top front hinge.
- 5: Bottom front hinge.
- 2. Remove the filler panels where the servers will be installed using a No. 2 screwdriver to remove the M6 screws. The number of screws depends on the type of filler panel. Save the screws for future use.






Figure 1-3 Removing the Filler Panels

Note:

If you are replacing the filler panels, then do not remove the Duosert cage-nuts from the RETMA (Radio Electronics Television Manufacturers Association) rail holes.

1.3.3.2 Adding a Sun Datacenter InfiniBand Switch 36 Switch (Optional)

This procedure is necessary as follows:

- Upgrading a rack with Sun Fire X4170 Oracle Database Servers to Oracle Exadata Database Machine Half Rack or Oracle Exadata Database Machine Full Rack.
- Extending an Oracle Exadata Database Machine Quarter Rack or Oracle Exadata Database Machine Eighth Rack to another rack.
- Extending an Oracle Exadata Database Machine X4-2 rack to another rack.

Note:

The steps in this procedure are specific to Oracle Exadata Database Machine. They are not the same as the steps in the Sun Datacenter InfiniBand Switch 36 manual.



- 1. Unpack the Sun Datacenter InfiniBand Switch 36 switch components from the packing cartons. The following items should be in the packing cartons:
 - Sun Datacenter InfiniBand Switch 36 switch
 - Cable bracket and rackmount kit
 - Cable management bracket and cover
 - Two rack rail assemblies
 - Assortment of screws and captive nuts
 - Sun Datacenter InfiniBand Switch 36 documentation

The service label procedure on top of the switch includes descriptions of the preceding items.

- 2. X5 racks only: Remove the trough from the rack in RU1 and put the cables aside while installing the IB switch. The trough can be discarded.
- 3. Install cage nuts in each rack rail in the appropriate holes.
- **4.** Attach the brackets with cutouts to the power supply side of the switch.
- 5. Attach the C-brackets to the switch on the side of the InfiniBand ports.
- 6. Slide the switch halfway into the rack from the front. You need to keep it to the left side of the rack as far as possible while pulling the two power cords through the C-bracket on the right side.
- 7. Slide the server in rack location U2 out to the locked service position. This improves access to the rear of the switch during further assembly.
- 8. Install the slide rails from the rear of the rack into the C-brackets on the switch, pushing them up to the rack rail.
- 9. Attach an assembled cable arm bracket to the slide rail and using a No. 3 Phillips screwdriver, screw these together into the rack rail:
 - **a.** Install the lower screw loosely with the cable arm bracket rotated 90 degrees downward. This allows better finger access to the screw.
 - **b.** Rotate the cable arm bracket to the correct position.
 - c. Install the upper screw.
 - d. Tighten both screws.

If available, a screwdriver with a long-shaft (16-inch / 400mm) will allow easier installation such that the handle is outside the rack and beyond the cabling.

- **10.** Push the switch completely into the rack from the front, routing the power cords through the cutout on the rail bracket.
- **11.** Secure the switch to the front rack rail with M6 16mm screws. Tighten the screws using the No. 3 Phillips screwdriver.
- 12. Install the lower part of the cable management arm across the back of the switch.
- **13.** Connect the cables to the appropriate ports.
- 14. Install the upper part of the cable management arm.
- **15.** Slide the server in rack location U2 back into the rack.
- **16.** Install power cords to the InfiniBand switch power supply slots on the front.



17. Loosen the front screws to install the vented filler panel brackets. Tighten the screws, and snap on the vented filler panel in front of the switch.

See Also:

- Oracle Exadata Database Machine System Overview to view the rack layout
- Oracle Exadata Database Machine System Overview for information about InfiniBand networking cables

1.3.3.3 Adding New Servers

Oracle Exadata Database Machine Quarter Rack can be upgraded to Oracle Exadata Database Machine Half Rack, and Oracle Exadata Database Machine Half Rack can be upgraded to Oracle Exadata Database Machine Full Rack. The upgrade process includes adding new servers, cables, and, when upgrading to Oracle Exadata Database Machine X2-2 Full Rack, Sun Datacenter InfiniBand Switch 36 switch.

A Oracle Exadata Database Machine Quarter Rack to Oracle Exadata Database Machine Half Rack upgrade consists of installing the following:

- Two Oracle Database servers
- Four Exadata Storage Servers
- One Sun Datacenter InfiniBand Switch 36 switch (for Oracle Exadata Database Machine X2-2 with Sun Fire X4170 M2 Oracle Database Server only)
- Associated cables and hardware

A Oracle Exadata Database Machine Half Rack to Oracle Exadata Database Machine Full Rack upgrade consists of installing the following:

- Four Oracle Database servers
- Seven Exadata Storage Servers
- Associated cables and hardware



Note:

- If you are extending Oracle Exadata Database Machine X5-2, Oracle Exadata Database Machine X4-2, Oracle Exadata Database Machine X3-8 Full Rack, or Oracle Exadata Database Machine X2-2 (with X4170 and X4275 servers) half rack, then order the expansion kit that includes a Sun Datacenter InfiniBand Switch 36 switch.
- The new servers need to be configured manually when extending Oracle Exadata Database Machine Quarter Rack to Oracle Exadata Database Machine Half Rack, or Oracle Exadata Database Machine Half Rack to Oracle Exadata Database Machine Full Rack. Refer to Setting Up New Servers for additional information.
- Always load equipment into the rack from the bottom up, so that the rack does not become top-heavy and tip over. Extend the rack anti-tip bar to prevent the rack from tipping during equipment installation.

See Also:

- Oracle Exadata Database Machine System Overview to view the rack layout
- Installing the Server Into a Rack in Sun Server X4-2L Installation Guide

The following tasks describes how to add the servers, and cables:

Task 1 Preparing for Installation

The following procedure describes the pre-installation steps:

- 1. Identify the rack unit where the server will be installed. Fill the first available unit, starting from the bottom of the rack.
- Remove and discard the trough, which attaches the cable harness when no server is installed in the unit.
- Remove and discard the solid filler.


Task 2 Installing the Rack Assembly

The following procedure describes how to install the rack assembly:

- 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.
- 2. Orient the slide-rail assembly so that the ball-bearing track is forward and locked in place.
- Starting on either 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.



Figure 1-4 Locking the Slide-Rail Assembly Against the Inside of the Rear Rack Rail



- 4. Align the front of the slide-rail assembly against the outside of the front rack rail, and push until the assembly locks into place and you hear the click.
- 5. Repeat steps 2 to 4 on the other side on the rack.

Task 3 Installing the Server

WARNING:

- Installing a server requires a minimum of two people or a lift because of the weight of each server. Attempting this procedure alone can result in equipment damage, personal injury, or both.
- Always load equipment into the rack from the bottom up, so that the rack does not become top-heavy and tip over. Extend the rack anti-tip bar to prevent the rack from tipping during equipment installation.

The following procedure describes how to install the sever:

- Read the service label on the top cover of the server before installing a server into the rack.
- 2. Push the server into the slide rail assembly:
 - a. Push the slide rails into the slide rail assemblies as far as possible.
 - **b.** Position the server so the rear ends of the mounting brackets are aligned with the slide rail assemblies mounted in the equipment rack.






Figure 1-5 Aligning the Rear Ends of the Mounting Brackets with the Slide Rail Assemblies in the Rack

The callouts in the preceding image highlight the following:

- 1: Mounting bracket inserted into slide rail
- 2: Slide-rail release lever
- c. Insert the mounting brackets into the slide rails, and push the server into the rack until the mounting brackets encounter the slide rail stops, approximately 30 cm (12 inches).
- **d.** Simultaneously push down and hold the slide rail release levers on each mounting bracket while pushing the server into the rack.
- e. Continue pushing until the slide rail locks on the front of the mounting brackets engage the slide rail assemblies, and you hear the click.
- 3. Cable the new server as described in "Cabling Exadata Storage Servers".



Oracle recommends that two people push the servers into the rack: one person to move the server in and out of the rack, and another person to watch the cables and CMA.

1.3.3.4 Cabling Database Servers


After the new database servers are installed, they need to be cabled to the existing equipment. The following procedure describes how to cable the new equipment in the rack. The images shown in the procedure are of a Sun Fire X4170 M2 Oracle Database Server.



Note:


- The existing cable connections in the rack do not change.
- The blue cables connect to Oracle Database servers, and the black cables connect to Exadata Storage Servers. These network cables are for the NETO Ethernet interface port.
- Attach and route the management cables on 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 the cables through the CMA with the dongle on the top and power cables on the bottom.
- Longer hook and loop straps are needed when cabling three CAT5e cables or two TwinAx cables.
- Connect the CAT5e cables, AC power cables, and USB to their respective ports on the rear of the server. Ensure the flat side of the dongle is flush against the CMA inner rail.

Figure 1-6 Cables at the Rear of the Server



2. Adjust the green cable management arm (CMA) brackets

Figure 1-7 Cable Management Arm (CMA) Brackets



Description of the CMA callouts in the preceding image"

a. Connector A



- b. Front slide bar
- c. Velcro straps (6)
- d. Connector B
- e. Connector C
- f. Connector D
- g. Slide-rail latching bracket (used with connector D)
- h. Rear slide bar
- i. Cable covers
- j. Cable covers
- 3. Attach the CMA to the server.
- **4.** Route the CAT5e and power cables through the wire clip.

Figure 1-8 Cables Routed Through the Cable Management Arm



- 5. Bend the CAT5e and power cables to enter the CMA, while adhering to the bend radius minimums.
- **6.** Secure the CAT5e and power cables under the cable clasps.


Figure 1-9 Cables Secured under the Cable Clasps





7. Route the cables through the CMA, and secure them with hook and loop straps at equal intervals.

Figure 1-10 Cables Secured with Hook and Loop Straps at Regular Intervals



8. Connect the InfiniBand or TwinAx cables with the initial bend resting on the CMA. The TwinAx cables are for client access to the database servers.

Figure 1-11 InfiniBand or TwinAx Cables Positioned on the CMA



9. Secure the InfiniBand or TwinAx cables with hook and loop straps at equal intervals.

Figure 1-12 InfiniBand or TwinAx Cables Secured with Hook and Loop Straps at Regular Intervals



- 10. Route the fiber core cables.
- 11. Rest the InfiniBand cables over the green clasp on the CMA.
- **12.** Attach the red ILOM cables to the database server.
- 13. Attach the network cables to the Oracle Database server.
- **14.** Attach the InfiniBand cables from Oracle Database server to the Sun Datacenter InfiniBand Switch 36 switches.
- 15. Connect the orange Ethernet cable to the KVM switch.
- 16. Connect the red and blue Ethernet cables to the Cisco switch.
- 17. Verify operation of the slide rails and CMA for each server, as follows:



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

- a. Slowly pull the server out of the rack until the slide rails reach their stops.
- Inspect the attached cables for any binding or kinks.
- c. Verify the CMA extends fully from the slide rails.
- 18. Push the server back into the rack, as follows:
 - a. Release the two sets of slide rail stops.
 - b. Push in both levers simultaneously, and slide the server into the rack. The first stop in the set are levers located on the inside of each slide rail, just behind the back panel of the server. The levers are labeled PUSH. The server slides approximately 46 cm (18 inches) and stop.
 - c. Verify the cables and CMA retract without binding.
 - **d.** Simultaneously push or pull both slide rail release buttons, and push the server completely into the rack until both slide rails engage. The second stop



in the set are the slide rail release buttons located near the front of each mounting bracket.

- 19. Dress the cables, and then tie off the cables with the straps. Oracle recommends the InfiniBand cables should be dressed in bundles of eight or less.
- **20.** Extend and then fully retract the server to check cable travel by sliding each server out and back fully to ensure that the cables are not binding or catching.
- **21.** Repeat the procedure for the rest of the servers.
- 22. Connect the power cables to the power distribution units (PDUs). Ensure the breaker switches are in the OFF position before connecting the power cables. Do not plug the power cables into the facility receptacles at this time.

See Also:

- Oracle Exadata Database Machine System Overview for cabling tables
- "Reviewing the Cable Management Arm Guidelines" for the bend radius minimums
- "Floor and Underfloor Delivery of InfiniBand Cables"

1.3.3.5 Cabling Exadata Storage Servers

After the new Exadata Storage Servers are installed, you need to connect them to the existing equipment.

The following procedure describes how to cable the new equipment in the rack.

Note:

- The existing cable connections in the rack do not change.
- The blue cables connect to Oracle Database servers, and the black cables connect to Exadata Storage Servers. These network cables are for the NETO Ethernet interface port.
- Attach and route the management cables on 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.
- Longer hook and loop straps are needed when cabling three CAT5e cables or two TwinAx cables.
- Attach a CMA to the server.
- 2. Insert the cables into their ports through the hook and loop straps, then route the cables into the CMA in this order:
 - a. Power
 - **b.** Ethernet



c. InfiniBand

Figure 1-13 Rear of the Server Showing Power and Network Cables



Route the cables through the CMA and secure them with hook and loop straps on both sides of each bend in the CMA.

Figure 1-14 Cables Routed Through the CMA and Secured with Hook and Loop Straps



- 4. Close the crossbar covers to secure the cables in the straightaway.
- 5. Verify operation of the slide rails and the CMA for each server:



Oracle recommends that two people do this step: one person to move the server in and out of the rack, and another person to watch the cables and the CMA.

- a. Slowly pull the server out of the rack until the slide rails reach their stops.
- Inspect the attached cables for any binding or kinks.
- c. Verify that the CMA extends fully from the slide rails.
- 6. Push the server back into the rack:
 - a. Release the two sets of slide rail stops.
 - b. Locate the levers on the inside of each slide rail, just behind the back panel of the server. They are labeled PUSH.
 - c. Simultaneously push in both levers and slide the server into the rack, until it stops in approximately 46 cm (18 inches).
 - d. Verify that the cables and CMA retract without binding.



- e. Locate the slide rail release buttons near the front of each mounting bracket.
- f. Simultaneously push in both slide rail release buttons and slide the server completely into the rack, until both slide rails engage.
- Dress the cables, and then tie off the cables with the straps. Oracle recommends that you dress the InfiniBand cables in bundles of eight or fewer.
- 8. Slide each server out and back fully to ensure that the cables are not binding or catching.
- 9. Repeat the procedure for all servers.
- 10. Connect the power cables to the power distribution units (PDUs). Ensure the breaker switches are in the OFF position before connecting the power cables. Do not plug the power cables into the facility receptacles now.



Multi-Rack Cabling Tables

Oracle Exadata Database Machine System Overview for the cabling tables for your system

1.3.3.6 Closing the Rack

The following procedure describes how to close the rack after installing new equipment.

- Replace the rack front and rear doors as follows:
 - a. Retrieve the doors, and place them carefully on the door hinges.
 - **b.** Connect the front and rear door grounding strap to the frame.
 - c. Close the doors.
 - d. (Optional) Lock the doors. The keys are in the shipping kit.
- 2. (Optional) Replace the side panels, if they were removed for the upgrade, as follows:
 - a. Lift each side panel up and onto the side of the rack. The top of the rack should support the weight of the side panel. Ensure the panel fasteners line up with the grooves in the rack frame.
 - b. Turn each side panel fastener one-quarter turn clockwise using the side panel removal tool. Turn the fasteners next to the panel lock clockwise. There are 10 fasteners per side panel.
 - c. (Optional) Lock each side panel. The key is in the shipping kit. The locks are located on the bottom, center of the side panels.
 - **d.** Connect the grounding straps to the side panels.

After closing the rack, proceed to "Configuring the New Hardware" to configure the new hardware.



1.3.4 Extending a Rack by Adding Another Rack

Extending Oracle Exadata Database Machine by adding another rack consists of cabling and configuring the racks together. Racks can be cabled together with no downtime. During the cabling procedure, the following should be noted:

- There is some performance degradation while cabling the racks together. This
 degradation results from reduced network bandwidth, and the data retransmission
 due to packet loss when a cable is unplugged.
- The environment is not a high-availability environment because one leaf switch will need to be off. All traffic goes through the remaining leaf switch.
- Only the existing rack is operational, and any new rack that is added is powered down.
- The software running on the systems cannot have problems related to InfiniBand restarts.
- It is assumed that Oracle Exadata Database Machine Half Racks have three InfiniBand switches already installed.
- The new racks have been configured with the appropriate IP addresses to be migrated into the expanded system prior to any cabling, and there are no duplicate IP addresses.
- The existing spine switch is set to priority 10 during the cabling procedure. This
 setting gives the spine switch a higher priority than any other switch in the fabric,
 and is the first to take the Subnet Manager Master role whenever a new Subnet
 Manager Master is being set during the cabling procedure.

The following sections describe how to extend Oracle Exadata Database Machine with another rack:

1.3.4.1 Cabling Two Racks Together

The following procedure describes how to cable two racks together. This procedure assumes that the racks are adjacent to each other. In the procedure, the existing rack is R1, and the new rack is R2.

- 1. Set the priority of the current, active Subnet Manager Master to 10 on the spine switch, as follows:
 - a. Log in to any InfiniBand switch on the active system.
 - b. Use the getmaster command to determine that the Subnet Manager Master is running on the spine switch. If it is not, then follow the procedure in *Oracle Exadata Database Machine Installation and Configuration Guide*.
 - c. Log in to the spine switch.
 - d. Use the disablesm command to stop Subnet Manager.
 - e. Use the setsmpriority 10 command to set the priority to 10.
 - f. Use the enablesm command to restart Subnet Manager.
 - g. Repeat step 1.b to ensure the Subnet Manager Master is running on the spine switch.



- 2. Ensure the new rack is near the existing rack. The InfiniBand cables must be able to reach the servers in each rack.
- 3. Completely shut down the new rack (R2).
- 4. Cable the two leaf switches R2 IB2 and R2 IB3 in the new rack according to Two-Rack Cabling. Note that you need to first remove the seven existing inter-switch connections between each leaf switch, as well as the two connections between the leaf switches and the spine switch in the new rack R2, not in the existing rack R1.
 - Verify both InfiniBand interfaces are up on all database nodes and storage cells. You can do this by running the ibstat command on each node and verifying both interfaces are up.
- 5. Power off leaf switch R1 IB2. This causes all the database servers and Exadata Storage Servers to fail over their InfiniBand traffic to R1 IB3.
- 6. Disconnect all seven inter-switch links between R1 IB2 and R1 IB3, as well as the one connection between R1 IB2 and the spine switch R1 IB1.
- 7. Cable leaf switch R1 IB2 according to Two-Rack Cabling.
- 8. Power on leaf switch R1 IB2.
- 9. Wait for three minutes for R1 IB2 to become completely operational.
 - To check the switch, log in to the switch and run the ibswitches command. The output should show three switches, R1 IB1, R1 IB2, and R1 IB3.
- 10. Verify both InfiniBand interfaces are up on all database nodes and storage cells. You can do this by running the ibstat command on each node and verifying both interfaces are up.
- 11. Power off leaf switch R1 IB3. This causes all the database servers and Exadata Storage Servers to fail over their InfiniBand traffic to R1 IB2.
- 12. Disconnect the one connection between R1 IB3 and the spine switch R1 IB1.
- 13. Cable leaf switch R1 IB3 according to Two-Rack Cabling.
- 14. Power on leaf switch R1 IB3.
- 15. Wait for three minutes for R1 IB3 to become completely operational.
 - To check the switch, log in to the switch and run the ibswitches command. The output should show three switches, R1 IB1, R1 IB2, and R1 IB3.
- 16. Power on all the InfiniBand switches in R2.
- 17. Wait for three minutes for the switches to become completely operational.
 - To check the switch, log in to the switch and run the <code>ibswitches</code> command. The output should show six switches, R1 IB1, R1 IB2, R1 IB3, R2 IB1, R2 IB2, and R2 IB3.
- **18.** Ensure the Subnet Manager Master is running on R1 IB1 by running the getmaster command from any switch.
- 19. Power on all servers in R2.
- 20. Log in to spine switch R1 IB1, and lower its priority to 8 as follows:
 - a. Use the disablesm command to stop Subnet Manager.
 - **b.** Use the setsmpriority 8 command to set the priority to 8.
 - c. Use the enablesm command to restart Subnet Manager.



21. Ensure Subnet Manager Master is running on one of the spine switches.

After cabling the racks together, proceed to Configuring the New Hardware to configure the racks.

See Also:

Multi-Rack Cabling Tables

Oracle Exadata Database Machine System Overview for the cabling tables for your system

1.3.4.2 Cabling Several Racks Together

The following procedure describes how to cable several racks together. This procedure assumes that the racks are adjacent to each other. In the procedure, the existing racks are R1, R2, ... Rn, the new rack is Rn+1, and the Subnet Manager Master is running on R1 IB1.

- 1. Set the priority of the current, active Subnet Manager Master to 10 on the spine switch, as follows:
 - a. Log in to any InfiniBand switch on the active system.
 - b. Use the getmaster command to determine that the Subnet Manager Master is running on the spine switch. If it is not, then follow the procedure in *Oracle Exadata Database Machine Installation and Configuraton Guide*.
 - c. Log in to the spine switch.
 - d. Use the disablesm command to stop Subnet Manager.
 - e. Use the setsmpriority 10 command to set the priority to 10.
 - f. Use the enablesm command to restart Subnet Manager.
 - g. Repeat step 1.b to ensure the Subnet Manager Master is running on the spine switch.
- 2. Ensure the new rack is near the existing rack. The InfiniBand cables must be able to reach the servers in each rack.
- 3. Completely shut down the new rack (Rn+1).
- 4. Cable the leaf switch in the new rack according to the appropriate table in Multi-Rack Cabling Tables. For example, if rack Rn+1 was R4, then use Table 2-7.
- 5. Complete the following procedure for each of the original racks:
 - a. Power off leaf switch Rx IB2. This causes all the database servers and Exadata Storage Servers to fail over their InfiniBand traffic to Rx IB3.
 - **b.** Cable leaf switch Rx IB2 according to Multi-Rack Cabling Tables.
 - c. Power on leaf switch Rx IB2.
 - **d.** Wait for three minutes for Rx IB2 to become completely operational.

To check the switch, log in to the switch and run the ibswitches command. The output should show n*3 switches for IB1, IB2, and IB3 in racks R1, R2, ... Rn.



- e. Power off leaf switch Rx IB3. This causes all the database servers and Exadata Storage Servers to fail over their InfiniBand traffic to Rx IB2.
- f. Cable leaf switch Rx IB3 according to Multi-Rack Cabling Tables.
- g. Power on leaf switch Rx IB3.
- h. Wait for three minutes for Rx IB3 to become completely operational.

To check the switch, log in to the switch and run the ibswitches command. The output should show n*3 switches for IB1, IB2, and IB3 in racks R1, R2, ... Rn.

All racks should now be rewired according to Multi-Rack Cabling Tables.

- 6. Power on all the InfiniBand switches in Rn+1.
- 7. Wait for three minutes for the switches to become completely operational.

To check the switch, log in to the switch and run the <code>ibswitches</code> command. The output should show (n+1)*3 switches for IB1, IB2, and IB3 in racks R1, R2, ... Rn +1.

- 8. Ensure the Subnet Manager Master is running on R1 IB1 by running the getmaster command from any switch.
- 9. Power on all servers in Rn+1.
- 10. Log in to spine switch R1 IB1, and lower its priority to 8 as follows:
 - a. Use the disablesm command to stop Subnet Manager.
 - **b.** Use the setsmpriority 8 command to set the priority to 8.
 - c. Use the enablesm command to restart Subnet Manager.
- 11. Ensure Subnet Manager Master is running on one of the spine switches using the getmaster command from any switch.
- **12.** Ensure Subnet Manager is running on every spine switch using the following command from any switch:

```
ibdiagnet -r
```

Each spine switch should show as running in the Summary Fabric SM-state-priority section of the output. If a spins switch is not running, then log in to the switch and enable Subnet Manager using the <code>enablesm</code> command.

13. If there are now four or more racks, then log in to the leaf switches in each rack and disable Subnet Manager using the disablesm command.

1.4 Configuring the New Hardware

This section contains the following tasks needed to configure the new hardware:

- Changing the Interface Names
- 2. Setting Up New Servers
- 3. Setting up a New Rack
- 4. Setting User Equivalence
- 5. Starting the Cluster
- 6. Adding Grid Disks to Oracle ASM Disk Groups
- 7. Adding Servers to a Cluster



- 8. Configuring Cell Alerts for New Oracle Exadata Storage Servers
- 9. Adding Oracle Database Software to the New Servers
- 10. Adding Database Instance to the New Servers



The new and existing racks must be at the same patch level for Oracle Exadata Database Servers and Oracle Exadata Storage Servers, including the operating system. Refer to "Reviewing Release and Patch Levels" for additional information.

1.4.1 Changing the Interface Names

Earlier releases of Oracle Exadata Database Machine X2-2 (with X4170 and X4275 servers) used BONDO and BONDO as the names for the bonded InfiniBand and bonded Ethernet client networks, respectively. In the current release, BONDO and BONDETHO are used for the bonded InfiniBand and bonded Ethernet client networks.

If you are adding new servers to an existing Oracle Exadata Database Machine X2-2 (with X4170 and X4275 servers), then ensure the database servers use the same names for bonded configuration. You can either change the new database servers to match the existing server interface names, or change the existing server interface names and Oracle Cluster Registry (OCR) configuration to match the new servers.

Do the following after changing the interface names:

 Edit the entries in /etc/sysctl.conf file on the database servers so that the entries for the InfiniBand network match. The following is an example of the file entries before editing. One set of entries must be changed to match the other set.

```
Found in X2 node

net.ipv4.neigh.bondib0.locktime = 0
net.ipv4.conf.bondib0.arp_ignore = 1
net.ipv4.conf.bondib0.arp_accept = 1
net.ipv4.neigh.bondib0.base_reachable_time_ms = 10000
net.ipv4.neigh.bondib0.delay_first_probe_time = 1

Found in V2 node

net.ipv4.conf.bond0.arp_accept=1
net.ipv4.neigh.bond0.base_reachable_time_ms=10000
net.ipv4.neigh.bond0.delay_first_probe_time=1
```

- 2. Save the changes to the sysctl.conf file.
- 3. Use the oifcfg utility to change the OCR configuration, if the new names differ from what is currently in OCR. The interface names for Oracle Exadata Storage Servers do not have to be changed.
- 4. Continue configuring the new hardware, as follows:
 - If the hardware is new servers, then go to Setting Up New Servers to configure the servers.
 - If the hardware is a new rack, then go to Setting up a New Rack to configure the rack.



See Also:

Oracle Exadata Database Machine Maintenance Guide for information about changing the InfiniBand network information

1.4.2 Setting Up New Servers

New servers need to be configured when extending Oracle Exadata Database Machine Quarter Rack or Oracle Exadata Database Machine Half Rack.

The new servers do not have any configuration information, and you cannot use Oracle Enterprise Manager Cloud Control to configure them. The servers are configured using the Oracle Exadata Deployment Assistant (OEDA) or manually.

Configuring Servers Using OEDA



In order to configure the servers with OEDA, the new server information must be entered in OEDA, and configuration files generated.

- 1. Download the latest release of OEDA listed in My Oracle Support note 888828.1.
- Enter the new server information in OEDA. Do not include information for the existing rack.

Note:

- When extending an existing rack that has database servers earlier than Oracle Exadata Database Machine X4-2, be sure to deselect the active bonding option for the InfiniBand network so the new database servers are configured with active-passive bonded interfaces.
- When extending an existing Oracle Exadata Database Machine X4-2 or later with active-active bonding, select the active bonding option to configure the new database servers for active-active bonding.
- 3. Generate the configuration files.
- 4. Prepare the servers as follows, starting with the first database server of the new servers:
 - a. Configure the servers as described in Preparing the Servers in *Oracle Exadata System Software User's Guide*.



Note:

OEDA checks the performance level of Oracle Exadata Storage Servers so it is not necessary to check them using the CellCLI CALIBRATE command at this time.

- b. Create the cell disks and grid disks as described in Configuring Cells, Cell Disks, and Grid Disks with CellCLI in Oracle Exadata System Software User's Guide.
- c. Create the flash cache and flash log as described in Creating Flash Cache and Flash Grid Disks in *Oracle Exadata System Software User's Guide*.

Note:

When creating the flash cache, enable write back flash cache.

- **5.** Ensure the InfiniBand and bonded client Ethernet interface names are the same on the new database servers as on the existing database servers.
- 6. When using the same, earlier style bonding names, such as BONDO, for the new database servers, then update the <code>/opt/oracle.cellos/cell.conf</code> file to reflect the correct bond names.

Note:

If the existing servers use ${\tt BONDIBO}$ as the InfiniBand bonding name, then this step can be skipped.

7. Install OEDA on the first new database server.

See Also:

My Oracle Support note 888828.1 for information about OEDA

- 8. Copy the configuration files to the first database server of the *new* servers in the /opt/oracle.SupportTools/onecommand directory. This is the information completed in step 2.
- 9. Run OEDA up to, but not including, the <code>CreateGridDisk</code> step, and then run the <code>SetupCellEmailAlerts</code> step and the Oracle Auto Service Request (ASR) configuration steps.



Note:

- The OEDA ValidateEnv step may display an error message about missing files, pxx.zip. This is expected behavior because the files are not used for this procedure. You can ignore the error message.
- When using capacity-on-demand, OEDA has the SetUpCapacityOnDemand step. This step uses the resourcecontrol command to set up the cores correctly.
- 10. Configure the storage servers, cell disks and grid disks as described in Configuring Cells, Cell Disks and Grid Disks with CellCLI in Oracle Exadata System Software User's Guide.

Note:

Use the data collected from the existing system, as described in Obtaining Current Configuration Information to determine the grid disk names and sizes.

11. Run reclaimdisks.sh on each database server.

The <code>/opt/oracle.SupportTools/reclaimdisks.sh</code> -free -reclaim command reclaims disk space reserved for the deployment type not selected. The command takes less than 5 minutes, approximately. Systems are imaged with disks configured with RAID5; a RAID rebuild is no longer part of the <code>reclaimdisks.sh</code> process.

Do not skip this step. Skipping this step results in unused space that can no longer be reclaimed by reclaimdisks.sh.

- **12.** Verify the time is the same on the new servers as on the existing servers. This check is performed for the storage servers and database servers.
- **13.** Ensure the NTP settings are the same on the new servers as on the existing servers. This check is performed for storage servers and database servers.
- **14.** Configure HugePages on the new servers to match the existing servers.
- **15.** Ensure the values in the /etc/security/limits.conf file for the new database servers match the existing database servers.
- 16. Go to Setting User Equivalence to continue the hardware configuration.

Configuring Servers Manually

- 1. Prepare the servers using the procedure described in Preparing the Servers in *Oracle Exadata System Software User's Guide*.
- 2. Ensure the InfiniBand and bonded client Ethernet interface names are the same on the new database servers as on the existing database servers.
- 3. Configure the storage servers, cell disks and grid disks as described in Configuring Cells, Cell Disks and Grid Disks with CellCLI in *Oracle Exadata System Software User's Guide*.
- 4. Configure the database servers as described in Setting Up Configuration Files for a Database Server Host in *Oracle Exadata System Software User's Guide*.



5. Run reclaimdisks.sh on each database server.

The <code>/opt/oracle.SupportTools/reclaimdisks.sh</code> -free -reclaim command reclaims disk space reserved for the deployment type not selected. The command takes less than 5 minutes, approximately. Systems are imaged with disks configured with RAID5; a RAID rebuild is no longer part of the <code>reclaimdisks.sh</code> process.

Do not skip this step. Skipping this step results in unused space that can no longer be reclaimed by reclaimdisks.sh.

- **6.** Verify the time is the same on the new servers as on the existing servers. This check is performed for the storage servers and database server.
- 7. Ensure the NTP settings are the same on the new servers as on the existing servers. This check is performed for the storage servers and database servers.
- 8. Configure HugePages on the new servers to match the existing servers.
- 9. Go to Setting User Equivalence to continue the hardware configuration.

Related Topics

- Preparing the Network Configuration
- Exadata Database Machine and Exadata Storage Server Supported Versions (My Oracle Support Doc ID 888828.1)

1.4.3 Setting up a New Rack

A new rack is configured at the factory. However, it is necessary to set up the network and configuration files for use with the existing rack.

- 1. Check the storage servers as described in Checking Exadata Storage Servers in Oracle Exadata Database Machine Installation and Configuration Guide.
- 2. Check the database servers as described in Checking Oracle Database Servers in *Oracle Exadata Database Machine Installation and Configuration Guide*.
- 3. Perform the checks as described in Performing Additional Checks and Configuration in *Oracle Exadata Database Machine Installation and Configuration Guide*.
- 4. Verify the InfiniBand network as described in Verifying the InfiniBand Network in Oracle Exadata Database Machine Installation and Configuration Guide.
- 5. Perform initial configuration as described in Performing Initial Elastic Configuration of Oracle Exadata Database Machine in *Oracle Exadata Database Machine Installation and Configuration Guide*.
- 6. Reclaim disk space as described in Configuring Oracle Database and Oracle ASM Instances for Oracle Exadata Database Machine Manually in *Oracle Exadata Database Machine Installation and Configuration Guide*.
- 7. Verify the time is the same on the new servers as on the existing servers. This check is performed for storage servers and database servers.
- **8.** Ensure the NTP settings are the same on the new servers as on the existing servers. This check is performed for storage servers and database servers.
- 9. Configure HugePages on the new servers to match the existing servers.
- **10.** Ensure the InfiniBand and bonded client Ethernet interface names on the new database servers match the existing database servers.



11. Configure the rack as described in Loading the Configuration Information and Installing the Software in Oracle Exadata Database Machine Installation and Configuration Guide. You can use either the Oracle Exadata Deployment Assistant (OEDA) or Oracle Enterprise Manager Cloud Control to configure the rack.

Note:

- Only run OEDA up to the CreateGridDisks step, then configure storage servers as described in Configuring Cells, Cell Disks, and Grid Disks with CellCLI in Oracle Exadata System Software User's Guide.
- When adding servers with 3 TB High Capacity (HC) disks to existing servers with 2TB disks, it is recommended to follow the procedure in My Oracle Support note 1476336.1 to properly define the grid disks and disk groups. At this point of setting up the rack, it is only necessary to define the grid disks. The disk groups are created after the cluster has been extended on to the new nodes.
- If the existing storage servers have High Performance (HP) disks and you are adding storage servers with High Capacity (HC) disks or the existing storage servers have HC disks and you are adding storage servers HP disks, then you must place the new disks in new disk groups. It is not permitted to mix HP and HC disks within the same disk group.
- 12. Go to Setting User Equivalence to continue the hardware configuration.

Related Topics

 How to Add Exadata Storage Servers Using 3 TB (or larger) Disks to an Existing Database Machine (My Oracle Support Doc ID 1476336.1)

1.4.4 Setting User Equivalence

User equivalence can be configured to include all servers once the servers are online.

This procedure must be done before running the post-cabling utilities.

- 1. Log in to each new server manually using SSH to verify that each server can accept log ins and that the passwords are correct.
- 2. Modify the dbs_group and cell_group files on all servers to include all servers.
 - **a.** Create the new directories on the first existing database server.

```
# mkdir /root/new_group_files
# mkdir /root/old_group_files
# mkdir /root/group_files
```

- **b.** Copy the group files for the new servers to the /root/new_group_files directory.
- c. Copy the group files for the existing servers to the /root/ old group files directory.
- d. Copy the group files for the existing servers to the /root/group_files directory.



e. Update the group files to include the existing and new servers.

```
cat /root/new_group_files/dbs_group >> /root/group_files/dbs_group
cat /root/new_group_files/cell_group >> /root/group_files/cell_group
cat /root/new_group_files/all_group >> /root/group_files/all_group
cat /root/new_group_files/dbs_ib_group >> /root/group_files/dbs_ib_group
cat /root/new_group_files/cell_ib_group >> /root/group_files/cell_ib_group
cat /root/new_group_files/all_ib_group >> /root/group_files/all_ib_group
```

f. Make the updated group files the default group files. The updated group files contain the existing and new servers.

```
cp /root/group_files/* /root
cp /root/group_files/* /opt/oracle.SupportTools/onecommand
```

- g. Put a copy of the updated group files in the root user, oracle user, and Oracle Grid Infrastructure user home directories, and ensure that the files are owned by the respective users.
- 3. Modify the /etc/hosts file on the existing and new database server to include the existing InfiniBand IP addresses for the database servers and storage servers. The existing and new priv_ib_hosts files can be used for this step.

Note:

Do not copy the /etc/hosts file from one server to the other servers. Edit the file on each server.

4. Run the setssh-Linux.sh script as the root user on one of the existing database servers to configure user equivalence for all servers using the following command. Oracle recommends using the first database server.

```
# /opt/oracle.SupportTools/onecommand/setssh-Linux.sh -s -c N -h \
   /path_to_file/all_group -n N
```

In the preceding command, *path_to_file* is the directory path for the all_group file containing the names for the existing and new servers.

Note:

For Oracle Exadata Database Machine X2-2 (with X4170 and X4275 servers) systems, use the <code>setssh.sh</code> command to configure user equivalence.

The command line options for the setssh.sh command differ from the setssh-Linux.sh command. Run setssh.sh without parameters to see the proper syntax.

5. Add the known hosts using InfiniBand. This step requires that all database servers are accessible by way of their InfiniBand interfaces.

```
# /opt/oracle.SupportTools/onecommand/setssh-Linux.sh -s -c N -h
/path_to_file/all_ib_group -n N -p password
```

6. Verify equivalence is configured.

```
# dcli -g all_group -l root date
# dcli -g all_ib_group -l root date
```



7. Run the setssh-Linux.sh script as the oracle user on one of the existing database servers to configure user equivalence for all servers using the following command. Oracle recommends using the first database server. If there are separate owners for the Oracle Grid Infrastructure software, then run a similar command for each owner.

In the preceding command, *path_to_file* is the directory path for the dbs_group file. The file contains the names for the existing and new servers.

Note:

- For Oracle Exadata Database Machine X2-2 (with X4170 and X4275 servers) systems, use the setssh.sh command to configure user equivalence.
- It may be necessary to temporarily change the permissions on the setssh-Linux.sh file to 755 for this step. Change the permissions back to the original settings after completing this step.
- 8. Add the known hosts using InfiniBand. This step requires that all database servers are accessible by way of their InfiniBand interfaces.

```
$ /opt/oracle.SupportTools/onecommand/setssh-Linux.sh -s -c N -h
/root/group_files/dbs_ib_group -n N
```

9. Verify equivalence is configured.

```
$ dcli -g dbs_group -l oracle date
$ dcli -g dbs_ib_group -l oracle date
```

If there is a separate Oracle Grid Infrastructure user, then also run the preceding commands for that user, substituting the grid user name for the oracle user.

1.4.5 Starting the Cluster

The following procedure describes how to start the cluster if it was stopped earlier for cabling an additional rack.

Note:

- Oracle recommends you start one server, and let it come up fully before starting Oracle Clusterware on the rest of the servers.
- It is not necessary to stop a cluster when extending Oracle Exadata
 Database Machine Half Rack to a Full Rack, or a Quarter Rack to a Half
 Rack or Full Rack.
- 1. Log in as the root user on the original cluster.
- 2. Start one server of the cluster.



Grid_home/grid/bin/crsctl start cluster

3. Check the status of the server.

```
Grid_home/grid/bin/crsctl stat res -t
```

Run the preceding command until it shows that the first server has started.

4. Start the other servers in the cluster.

```
# Grid_home/grid/bin/crsctl start cluster -all
```

5. Check the status of the servers.

```
Grid_home/grid/bin/crsctl stat res -t
```

It may take several minutes for all servers to start and join the cluster.

1.4.6 Adding Grid Disks to Oracle ASM Disk Groups

Grid disks can be added to Oracle ASM disk groups before or after the new servers are added to the cluster. The advantage of adding the grid disks before adding the new servers is that the rebalance operation can start earlier. The advantage of adding the grid disks after adding the new servers is that the rebalance operation can be done on the new servers so less load is placed on the existing servers.

The following procedure describes how to add grid disk to existing Oracle ASM disk groups.

Note:

 It is assumed in the following examples that the newly-installed storage servers have the same grid disk configuration as the existing storage servers, and that the additional grid disks will be added to existing disk groups.

The information gathered about the current configuration should be used when setting up the grid disks.

- If the existing storage servers have High Performance (HP) disks and you are adding storage servers with High Capacity (HC) disks or the existing storage servers have HC disks and you are adding storage servers HP disks, then you must place the new disks in new disk groups.
 It is not permitted to mix HP and HC disks within the same disk group.
- 1. Ensure the new storage servers are running the same version of software as storage servers already in use. Run the following command on the first database server:

```
dcli -q dbs_group -l root "imageinfo -ver"
```





If the Oracle Exadata System Software on the storage servers does not match, then upgrade or patch the software to be at the same level. This could be patching the existing servers or new servers. Refer to Reviewing Release and Patch Levels for additional information.

2. Modify the /etc/oracle/cell/network-config/cellip.ora file on all database servers to have a complete list of all storage servers. This can be done by modifying the file for one database server, and then copying the file to the other database servers. The cellip.ora file should be identical on all database servers.

When adding Oracle Exadata Storage Server X4-2L servers, the cellip.ora file contains two IP addresses listed for each cell. Copy each line completely to include the two IP addresses, and merge the addresses in the cellip.ora file of the existing cluster.

The following is an example of the cellip.ora file after expanding Oracle Exadata Database Machine X3-2 Half Rack to Full Rack using Oracle Exadata Storage Server X4-2L servers:

```
cell="192.168.10.9"
cell="192.168.10.10"
cell="192.168.10.11"
cell="192.168.10.12"
cell="192.168.10.13"
cell="192.168.10.14"
cell="192.168.10.15"
cell="192.168.10.15"
cell="192.168.10.17;192.168.10.18"
cell="192.168.10.21;192.168.10.20"
cell="192.168.10.21;192.168.10.22"
cell="192.168.10.23;192.168.10.24"
cell="192.168.10.25;192.168.10.26"
cell="192.168.10.27;192.168.10.28"
cell="192.168.10.29;192.168.10.30"
```

In the preceding example, lines 1 through 7 are for the original servers, and lines 8 through 14 are for the new servers. Oracle Exadata Storage Server X4-2L servers have two IP addresses each.

- 3. Ensure the updated cellip.ora file is on all database servers. The updated file must include a complete list of all storage servers.
- 4. Verify accessibility of all grid disks from one of the original database servers. The following command can be run as the root user or the oracle user.

```
$ Grid_home/grid/bin/kfod disks=all dscvgroup=true
```

The output from the command shows grid disks from the original and new storage servers.

5. Add the grid disks from the new storage servers to the existing disk groups using commands similar to the following. You cannot have both high performance disks and high capacity disks in the same disk group.

```
$ .oraenv
ORACLE_SID = [oracle] ? +ASM1
The Oracle base for ORACLE_HOME=/u01/app/11.2.0/grid is /u01/app/oracle
$ sqlplus / as sysasm
```



```
SQL> ALTER DISKGROUP data ADD DISK
2> 'o/*/DATA*dm02*'
3> rebalance power 11;
```

In the preceding commands, a Full Rack was added to an existing Oracle Exadata Rack. The prefix for the new rack is dm02, and the grid disk prefix is DATA.

The following is an example in which an Oracle Exadata Database Machine Half Rack was upgraded to a Full Rack. The cell host names in the original system were named <code>dm01cel01</code> through <code>dm01cel07</code>. The new cell host names are <code>dm01cel08</code> through <code>dm01cel14</code>.

```
$ .oraenv
ORACLE_SID = [oracle] ? +ASM1
The Oracle base for ORACLE_HOME=/u01/app/11.2.0/grid is /u01/app/oracle
$ SQLPLUS / AS sysasm
SQL> ALTER DISKGROUP data ADD DISK
2> 'o/*/DATA*dm01cel08*',
3> 'o/*/DATA*dm01cel09*',
4> 'o/*/DATA*dm01cel10*',
5> 'o/*/DATA*dm01cel11*',
6> 'o/*/DATA*dm01cel11*',
8> 'o/*/DATA*dm01cel12*',
7> 'o/*/DATA*dm01cel13*',
8> 'o/*/DATA*dm01cel14*'
9> rebalance power 11;
```

Note:

- If your system is running Oracle Database 11g release 2 (11.2.0.1), then Oracle recommends a power limit of 11 so that the rebalance completes as quickly as possible. If your system is running Oracle Database 11g release 2 (11.2.0.2), then Oracle recommends a power limit of 32. The power limit does have an impact on any applications that are running during the rebalance.
- Ensure the ALTER DISKGROUP commands are run from different Oracle ASM instances. That way, the rebalance operation for multiple disk groups can run in parallel.
- Add disks to all disk groups including SYSTEMDG or DBFS_DG.
- When adding servers with 3 TB High Capacity (HC) disks to existing servers with 2 TB disks, it is recommended to follow the procedure in My Oracle Support note 1476336.1 to properly define the grid disks and disk groups. At this point of setting up the rack, the new grid disks should be defined, but need to be placed into disk groups. Refer to the steps in My Oracle Support note 1476336.1.
- If the existing storage servers have High Performance (HP) disks and you are adding storage servers with High Capacity (HC) disks, or the existing storage servers have HC disks and you are adding storage servers with HP disks, then you must place the new disks in new disk groups. It is not permitted to mix HP and HC disks within the same disk group.



6. Monitor the status of the rebalance operation using a query similar to the following from any Oracle ASM instance:

```
SQL> SELECT * FROM GV$ASM_OPERATION WHERE STATE = 'RUN';
```

The remaining tasks can be done while the rebalance is in progress.

See Also:

- Obtaining Current Configuration Information for information about the existing grid disks.
- Setting Up New Servers for information about configuring the grid disks.
- Oracle Automatic Storage Management Administrator's Guide for information about the ASM_POWER_LIMIT parameter.

1.4.7 Adding Servers to a Cluster

This procedure describes how to add servers to a cluster.

For adding nodes to an Oracle VM cluster, refer to Expanding an Oracle VM RAC Cluster on Exadata in *Oracle Exadata Database Machine Maintenance Guide*.

Caution:

If Oracle Clusterware manages additional services that are not yet installed on the new nodes, such as Oracle GoldenGate, then note the following:

- It may be necessary to stop those services on the existing node before running the addNode.sh script.
- It is necessary to create any users and groups on the new database servers that run these additional services.
- It may be necessary to disable those services from auto-start so that
 Oracle Clusterware does not try to start the services on the new nodes.

Note:

To prevent problems with transferring files between existing and new nodes, you need to set up SSH equivalence. See Step 4 in Expanding an Oracle VM Oracle RAC Cluster on Exadata in for details.

- Ensure the /etc/oracle/cell/network-config/*.ora files are correct and consistent on all database servers. The cellip.ora file all database server should include the older and newer database servers and storage servers.
- 2. Ensure the <code>ORACLE_BASE</code> and <code>diag</code> destination directories have been created on the Oracle Grid Infrastructure destination home.



The following is an example for Oracle Grid Infrastructure 11g:

```
# dcli -g /root/new_group_files/dbs_group -l root mkdir -p \
    /u01/app/11.2.0/grid /u01/app/oraInventory /u01/app/grid/diag

# dcli -g /root/new_group_files/dbs_group -l root chown -R grid:oinstall \
    /u01/app/11.2.0 /u01/app/oraInventory /u01/app/grid

# dcli -g /root/new_group_files/dbs_group -l root chmod -R 770 \
    /u01/app/oraInventory

# dcli -g /root/new_group_files/dbs_group -l root chmod -R 755 \
    /u01/app/11.2.0 /u01/app/11.2.0/grid
```

The following is an example for Oracle Grid Infrastructure 12c:

```
# cd /
# rm -rf /u01/app/*
# mkdir -p /u01/app/12.1.0.2/grid
# mkdir -p /u01/app/oracle/product/12.1.0.2/dbhome_1
# chown -R oracle:oinstall /u01
```

3. Ensure the inventory directory and Grid home directories have been created and have the proper permissions. The directories should be owned by the Grid user and the OINSTALL group. The inventory directory should have 770 permission, and the Oracle Grid Infrastructure home directories should have 755.

If you are running Oracle Grid Infrastructure 12c or later:

- Make sure oraInventory does not exist inside /u01/app.
- Make sure /etc/oraInst.loc does not exist.
- Create users and groups on the new nodes with the same user identifiers and group identifiers as on the existing nodes.

Note:

If Oracle Exadata Deployment Assistant (OEDA) was used earlier, then these users and groups should have been created. Check that they do exist, and have the correct UID and GID values.

- 5. Log in as the Grid user on an existing host.
- 6. Verify the Oracle Cluster Registry (OCR) backup exists.

```
ocrconfig -showbackup
```

7. Verify that the additional database servers are ready to be added to the cluster using commands similar to following:

```
$ cluvfy stage -post hwos -n \
  dm02db01,dm02db02,dm02db03,dm02db04,dm02db05,dm02db06,dm02db07,dm02db08 \
  -verbose

$ cluvfy comp peer -refnode dm01db01 -n \
  dm02db01,dm02db02,dm02db03,dm02db04,dm02db05,dm02db06,dm02db07,dm02db08 \
  -orainv oinstall -osdba dba | grep -B 3 -A 2 mismatched

$ cluvfy stage -pre nodeadd -n \
  dm02db01,dm02db02,dm02db03,dm02db04,dm02db05,dm02db06,dm02db07,dm02db08 \
  -verbose -fixup -fixupdir /home/grid_owner_name/fixup.d
```



In the preceding commands, *grid_owner_name* is the name of the Oracle Grid Infrastructure software owner, *dm02db01* through *db02db08* are the new database servers, and *refnode* is an existing database server.

Note:

- The second and third commands do not display output if the commands complete correctly.
- An error about a voting disk, similar to the following, may be displayed:

If such an error occurs:

- If you are running Oracle Grid Infrastructure $11_{\rm g}$, set the environment variable as follows:

```
$ export IGNORE_PREADDNODE_CHECKS=Y
```

Setting the environment variable does not prevent the error when running the <code>cluvfy</code> command, but it does allow the <code>addNode.sh</code> script to complete successfully.

- If you are running Oracle Grid Infrastructure 12c or later, use the following addnode parameters: -ignoreSysPrereqs -ignorePrereq

In Oracle Grid Infrastructure 12c and later, addnode does not use the IGNORE PREADDNODE CHECKS environment variable.

• If a database server was installed with a certain image and subsequently patched to a later image, then some operating system libraries may be older than the version expected by the cluvfy command. This causes the cluvfy command and possibly the addNode.sh script to fail.

It is permissible to have an earlier version as long as the difference in versions is minor. For example, glibc-

common-2.5-81.el5_8.2 versus glibc-common-2.5-49. The versions are different, but both are at version 2.5, so the difference is minor, and it is permissible for them to differ.

Set the environment variable IGNORE_PREADDNODE_CHECKS=Y before running the addNode.sh Script or use the addNode parameters - ignoreSysPrereqs -ignorePrereq with the addNode.sh Script to workaround this problem.



8. Ensure that all directories inside the Oracle Grid Infrastructure home on the existing server have their executable bits set. Run the following commands as the root user.

```
find /u01/app/11.2.0/grid -type d -user root ! -perm /u+x !
-perm /g+x ! -perm o+x

find /u01/app/11.2.0/grid -type d -user grid_owner_name ! -perm /u+x ! \
-perm /g+x ! -perm o+x
```

In the preceding commands, $grid_owner_name$ is the name of the Oracle Grid Infrastructure software owner, and /u01/app/11.2.0/grid is the Oracle Grid Infrastructure home directory.

If any directories are listed, then ensure the group and others permissions are +x. The <code>Grid_home/network/admin/samples</code>, \$GI_HOME/crf/admin/run/crfmond, and <code>Grid_home/crf/admin/run/crflogd</code> directories may need the +x permissions set.

If you are running Oracle Grid Infrastructure 12c or later, run commands similar to the following:

```
# chmod -R u+x /u01/app/12.1.0.2/grid/gpnp/gpnp_bcp*
# chmod -R o+rx /u01/app/12.1.0.2/grid/gpnp/gpnp_bcp*
# chmod o+r /u01/app/12.1.0.2/grid/bin/oradaemonagent /u01/app/12.1.0.2/grid/srvm/admin/logging.properties
# chmod a+r /u01/app/oracle/product/12.1.0.2/dbhome_1/bin/*0
# chmod a+r /u01/app/oracle/product/12.1.0.2/dbhome_1/bin/*0
# chown -f gi_owner_name:dba /u01/app/12.1.0.2/grid/OPatch/ocm/bin/emocmrsp
```

The <code>Grid_home/network/admin/samples</code> directory needs the <code>+x permission</code>:

```
chmod -R a+x /u01/app/12.1.0.2/grid/network/admin/samples
```

9. Run the following command. It is assumed that the Oracle Grid Infrastructure home is owned by the Grid user.

```
$ dcli -g old_db_nodes -l root chown -f grid_owner_name:dba \
   /u01/app/11.2.0/grid/OPatch/ocm/bin/emocmrsp
```

10. This step is needed only if you are running Oracle Grid Infrastructure 11_{g} . In Oracle Grid Infrastructure 12_{c} , no response file is needed because the values are specified on the command line.

Create a response file, add-cluster-nodes.rsp, as the Grid user to add the new servers similar to the following:

In the preceding file, the host names dm02db01 through db02db08 are the new nodes being added to the cluster.





The lines listing the server names should appear on one continuous line. They are wrapped in the documentation due to page limitations.

- **11.** Ensure most of the files in the <code>Grid_home/rdbms/audit</code> and <code>Grid_home/log/diag/*</code> directories have been moved or deleted before extending a cluster.
- 12. Refer to My Oracle Support note 744213.1 if the installer runs out of memory. The note describes how to edit the <code>Grid_home/oui/ora-param.ini</code> file, and change the <code>JRE_MEMORY_OPTIONS</code> parameter to <code>-Xms512m-Xmx2048m</code>.
- **13.** Add the new servers by running the addNode.sh script from an existing server as the Grid user.
 - If you are running Oracle Grid Infrastructure 11g:

```
$ cd Grid_home/oui/bin
$ ./addNode.sh -silent -responseFile /path/to/add-cluster-nodes.rsp
```

• If you are running Oracle Grid Infrastructure 12c or later, run the addnode.sh command with the CLUSTER_NEW_NODES and CLUSTER_NEW_VIRTUAL_HOSTNAMES parameters. The syntax is:

```
$ ./addnode.sh -silent "CLUSTER_NEW_NODES={comma_delimited_new_nodes}"
"CLUSTER NEW VIRTUAL HOSTNAMES={comma_delimited_new_node_vips}"
```

For example:

```
$ cd Grid_home/addnode/
$ ./addnode.sh -silent
"CLUSTER_NEW_NODES={dm02db01,dm02db02,dm02db03,dm02db04,dm02db05,
dm02db06,dm02db07,dm02db08}" "CLUSTER_NEW_VIRTUAL_HOSTNAMES={dm02db01-
vip,dm02db02-vip,
dm02db03-vip,dm02db04-vip,dm02db05-vip,dm02db06-vip,dm02db07-vip,dm02db08-
vip}"
-ignoreSysPrereqs -ignorePrereq
```

14. Verify the grid disks are visible from each of the new database servers.

```
$ Grid_home/grid/bin/kfod disks=all dscvgroup=true
```

15. Run the orainstRoot.sh Script as the root user when prompted using the dcli utility.

```
$ dcli -g new_db_nodes -l root \
   /u01/app/oraInventory/orainstRoot.sh
```

16. Disable HAIP on the new servers.

Before running the root.sh script, on each new server, set the HAIP_UNSUPPORTED environment variable to TRUE.

```
$ export HAIP_UNSUPPORTED=true
```

17. Run the <code>Grid_home/root.sh</code> script on each server sequentially. This simplifies the process, and ensures that any issues can be clearly identified and addressed.



Note:

The node identifier is set in order of the nodes where the ${\tt root.sh}$ script is run. Typically, the script is run from the lowest numbered node name to the highest.

- **18.** Check the log file from the root.sh script and verify there are no problems on the server before proceeding to the next server. If there are problems, then resolve them before continuing.
- **19**. Check the status of the cluster after adding the servers.

```
\ cluvfy stage -post nodeadd -n \ dm02db01,dm02db02,dm02db03,dm02db04,dm02db05,dm02db06,dm02db07,dm02db08 \ -verbose
```

20. Check that all servers have been added and have basic services running.

```
crsctl stat res -t
```

Note:

It may be necessary to mount disk groups on the new servers. The following commands must be run as the oracle user.

```
$ srvctl start diskgroup -g data
$ srvctl start diskgroup -g reco
```

- 21. If you are running Oracle Grid Infrastructure releases 11.2.0.2 and later, then perform the following steps:
 - a. Manually add the CLUSTER_INTERCONNECTS parameter to the SPFILE for each Oracle ASM instance.

```
ALTER SYSTEM SET cluster_interconnects = '192.168.10.x' \sid='+ASMx' scope=spfile
```

- b. Restart the cluster on each new server.
- c. Verify the parameters were set correctly.

1.4.8 Configuring Cell Alerts for New Oracle Exadata Storage Servers

Cell alerts need to be configured for the new Oracle Exadata Storage Servers.

The configuration depends on the type of installation.

 When extending Oracle Exadata Database Machine Quarter Rack to Half Rack, or Half Rack to Full Rack:

Manually configure cell alerts on the new storage servers. Use the settings on the original storage servers as a guide. To view the settings on the original storage servers, use a command similar to the following:

```
dcli -g new_cells_nodes -l celladmin cellcli -e list cell detail
```

To view the alerts on the new storage servers, use a command similar to the following:



Note:

The backslash character (\setminus) is used as an escape character for the dcli utility, and as a line continuation character in the preceding command.

When cabling racks:

Use Oracle Exadata Deployment Assistant (OEDA) to set up e-mail alerts for storage servers as the root user from the original rack to the new rack. The utility includes the SetupCellEmailAlerts step to configure alerts.

1.4.9 Adding Oracle Database Software to the New Servers

It is necessary to add the Oracle Database software directory $\mbox{oracle_home}$ to the database servers after the cluster modifications are complete, and all the database servers are in the cluster.

1. Check the <code>Oracle_home/bin</code> directory for files ending in zero (0), such as <code>nmb0</code>, that are owned by the <code>root</code> user and do not have <code>oinstall</code> or world read privileges. Use the following command to modify the file privileges:

```
# chmod a+r $ORACLE_HOME/bin/*0
```

If you are running Oracle Database release 12c or later, you also have to change permissions for files ending in uppercase O, in addition to files ending in zero.

```
# chmod a+r $ORACLE_HOME/bin/*0
```

2. This step is required for Oracle Database 11*g* only. If you are running Oracle Database 12*c*, you can skip this step because the directory has already been created.

Create the <code>ORACLE_BASE</code> directory for the database owner, if it is different from the Oracle Grid Infrastructure software owner (Grid user) using the following commands:

```
# dcli -g root/new_group_files/dbs_group -l root mkdir -p /u01/app/oracle
# dcli -g root/new_group_files/dbs_group -l root chown oracle:oinstall \
   /u01/app/oracle
```

3. Run the following command to set ownership of the <code>emocmrsp</code> file in the Oracle Database <code>\$ORACLE_HOME</code> directory:

```
# dcli -g old_db_nodes -l root chown -f oracle:dba \
/u01/app/11.2.0/grid/OPatch/ocm/bin/emocmrsp
```

4. This step is required for Oracle Database 11*g* only. If you are running Oracle Database 12*c*, then you can skip this step because the values are entered on the command line.

Create a response file, add-db-nodes.rsp, as the oracle owner to add the new servers similar to the following:

```
RESPONSEFILE_VERSION=2.2.1.0.0

CLUSTER_NEW_NODES={dm02db01,dm02db02,dm02db03,dm02db04,dm02db05,dm02db06,dm02db07,dm02db08}
```



The lines listing the server names should appear on one continuous line. The are wrapped in the document due to page limitations.

- 5. Add the Oracle Database ORACLE_HOME directory to the new servers by running the addNode.sh script from an existing server as the database owner user.
 - If you are running Oracle Grid Infrastructure 11g:

```
$ cd $ORACLE_HOME/oui/bin
$ ./addNode.sh -silent -responseFile /path/to/add-db-nodes.rsp
```

• If you are running Oracle Grid Infrastructure 12c, then you specify the nodes on the command line. The syntax is:

```
./addnode.sh -silent "CLUSTER_NEW_NODES={comma_delimited_new_nodes}"
```

For example:

```
$ cd $Grid_home/addnode

$ ./addnode.sh -silent
"CLUSTER_NEW_NODES={dm02db01,dm02db02,dm02db03,dm02db04,dm02db05,dm02db06,dm02db07,dm02db08}" -ignoreSysPreregs -ignorePrereg
```

- **6.** Ensure the \$ORACLE_HOME/oui/oraparam.ini file has the memory settings that match the parameters set in the Oracle Grid Infrastructure home.
- 7. Run the root.sh script on each server when prompted as the root user using the dcli utility.

```
$ dcli -g new_db_nodes -l root $ORACLE_HOME/root.sh
```

In the preceding command, new_db_nodes is the file with the list of new database servers.

8. Verify the <code>ORACLE_HOME</code> directories have been added to the new servers.

```
# dcli -g /root/all_group -l root du -sm \
   /u01/app/oracle/product/11.2.0/dbhome_1
```

1.4.10 Adding Database Instance to the New Servers

Before adding the database instances to the new servers, check the following:

- Maximum file size: If any data files have reached their maximum file size, then the
 addInstance command may fail with an ORA-00740 error. Oracle recommends you
 check that none of the files listed in DBA_DATA_FILES have reached their maximum
 size. Files that have reached their maximum should be corrected.
- Online redo logs: If the online redo logs are kept in the directory specified by the DB_RECOVERY_FILE_DEST parameter, then ensure the space allocated is sufficient for

the additional redo logs for the new instances being added. If necessary, then increase the size for the <code>DB_RECOVERY_FILE_DEST_SIZE</code> parameter.

- Total number of instances in the cluster: Set the value of the initialization
 parameter cluster_database_instances in the SPFILE for each database to the total
 number of instances that will be in the cluster after adding the new servers.
- The HugePages settings are correctly configured on the new servers to match the existing servers.
- 1. Use a command similar the following from an existing database server to add instances to the new servers. In the command, the instance, dbm9, is being added for server dm02db01.

```
dbca -silent -addInstance -gdbName dbm -nodeList dm02db01 -instanceName dbm9 \backslash -sysDBAUsername sys
```

The command must be run for all servers and instances, substituting the server name and instance name, as appropriate.



If the command fails, then ensure any files that were created, such as redo log files, are cleaned up. The deleteInstance command does not clean log files or data files that were created by the addInstance command.

- 2. Add the cluster_interconnects parameter to each new instance.
 - a. Manually add the CLUSTER_INTERCONNECTS parameter to the SPFILE for each new database instance. The additions are similar to the existing entries, but are the InfiniBand addresses corresponding to the server that each instance runs on.
 - **b.** Restart the instance on each new server.
 - c. Verify the parameters were set correctly.

1.5 Returning the Rack to Service

Use the following procedure to ensure the new hardware is correctly configured and ready for use:

- 1. Run the /opt/oracle.SupportTools/ibdiagtools/verify-topology command to ensure that all InfiniBand cables are connected and secure.
- 2. Run the Oracle Exadata Database Machine HealthCheck utility using the steps described in My Oracle Support note 1070954.1.
- **3.** Verify the instance additions using the following commands:

```
srvctl config database -d dbm
srvctl status database -d dbm
```

4. Check the cluster resources using the following command:

```
crsctl stat res -t
```

5. Ensure the original configuration summary report from the original cluster deployment is updated to include all servers. This document should include the



- calibrate and network verifications for the new rack, and the InfiniBand cable checks (verify-topology and infinicheck).
- 6. Conduct a power-off test, if possible. If the new Exadata Storage Servers cannot be powered off, then verify that the new database servers with the new instances can be powered off and powered on, and that all processes start automatically.

Note:

Ensure the Oracle ASM disk rebalance process has completed for all disk groups by using the following command:

select * from gv\$asm_operation

No rows should be returned by the command.

- 7. Review the configuration settings, such as the following:
 - All parallelism settings
 - Backup configurations
 - Standby site, if any
 - Service configuration
 - Oracle Database File System (DBFS) configuration, and mount points on new servers
 - Installation of Oracle Enterprise HugePage Manager agents on new database servers
 - HugePages settings
- 8. Incorporate the new cell and database servers into Auto Service Request.
- 9. Update Oracle Enterprise Manager Cloud Control to include the new nodes.

See Also:

- Oracle Exadata Database Machine Maintenance Guide for information about verifying the InfiniBand network configuration
- Auto Service Request Quick Installation Guide for Oracle Exadata Database Machine



2

Multi-Rack Cabling Tables

This appendix contains the tables for multi-rack cabling.

2.1 Understanding Multi-Rack Cabling

Up to eight racks can be cabled together without external InfiniBand switches.

Three Sun Datacenter InfiniBand Switch 36 switches are needed when cabling racks together. If your rack does not include a third switch, then order the expansion switch kit. These switches attach to standard Quad Small Form-factor Pluggable (QSFP) connectors at the end of the InfiniBand cables. The procedures in this section assume the racks are adjacent to each other. If they are not, then longer cables may be required for the connections.

Note:

- For Eighth or Quarter Racks, which are the smallest Elastic Configurations, follow the instructions in "Cabling Oracle Exadata Quarter Racks and Oracle Exadata Eighth Racks" for direct connection without spine switch.
- For other racks (Half Rack, Full Rack, Elastic Configurations larger than Eighth or Quarter Rack) install a spine switch and follow the standard Nrack cabling using fat tree topology for two or more interconnected racks, for example "Two-Rack Cabling."

The switch at rack unit 1 (U1) is referred to as the spine switch.

The following switches are referred to as leaf switches:

- The switches at rack unit 20 (U20) and rack unit 24 (U24) in Oracle Exadata Database Machine X2-2 racks and Oracle Exadata Storage Expansion Racks
- The switches at rack unit 21 (U21) and rack unit 23 (U23) in Oracle Exadata Database Machine X2-8 and later.
- The switches at rack unit 20 (U20) and 22 (U22) in Oracle Exadata Database Machine X3-2 and later racks or Oracle Exadata Storage Expansion Rack X3-2 and later.




Note:

Oracle Exadata Database Machine X4-2 and later racks or Oracle Exadata Database Machine X3-8 Full Racks with Exadata Storage Server X4-2L Servers do not include spare cables or a third Sun Datacenter InfiniBand Switch 36 switch. To extend Oracle Exadata Database Machine X4-2 and later racks or Oracle Exadata Database Machine X3-8 Full Racks with Exadata Storage Server X4-2L Servers, you must order cables and a Sun Datacenter InfiniBand Switch 36 switch.

In a single rack, the two leaf switches are interconnected using seven connections. In addition, each leaf switch has one connection to the spine switch. The leaf switches connect to the spine switch as shown in the following graphic:

Figure 2-1 Connections Between Spine Switch and Leaf Switches in a Single Rack



The Oracle Database servers and Exadata Storage Servers connect to the leaf switches as shown in the following graphic:



Port B

Oracle Database Servers (1, 2, 3, 4)

Exadata Storage Servers (1, 2, .., 7)

Port A
Oracle Database Servers (1, 2, 3, 4)
Exadata Storage Servers (1, 2, ..., 7)

11

Leaf Switch 1

Leaf Switch 2

Figure 2-2 Connections Between Database Servers and Storage Servers and Leaf Switches

When connecting up to eight racks together, remove the seven existing inter-switch connections between each leaf switch, as well as the two connections between the leaf switches and the spine switch. From each leaf switch, distribute eight connections over the spine switches in all racks. In multi-rack environments, the leaf switches inside a rack are no longer directly interconnected, as shown in the following graphic:

Rack 1

Spine Switch

A A A A A A

Leaf Switch 1

Rack 2

Spine Switch

Spine Switch

Leaf Switch 1

Leaf Switch 2

Leaf Switch 1

Leaf Switch 2

Figure 2-3 Connections Between Spine Switches and Leaf Switches Across Two Racks

Port B

Oracle Database Servers (5, 6, 7, 8)

Exadata Storage Servers (8, 9, .., 14)

As shown in the preceding graphic, each leaf switch in rack 1 connects to the following switches:

- Four connections to its internal spine switch
- Four connections to the spine switch in rack 2


The spine switch in rack 1 connects to the following switches:



- Eight connections to both internal leaf switches
- Eight connections to both leaf switches in rack 2

As the number of racks increases from two to eight, the pattern continues as shown in the following graphic:

Figure 2-4 Connections Between Spine Switches and Leaf Switches for up to 8 Racks



As shown in the preceding graphic, each leaf switch has eight inter-switch connections distributed over all spine switches. Each spine switch has 16 inter-switch connections distributed over all leaf switches. The leaf switches are not directly interconnected with other leaf switches, and the spine switches are not directly interconnected with the other spine switches.

Note:

Cable lengths are specified for racks 1 through 8.

2.1.1 Preparing to Cable Racks Together

Racks can be added together to increase system capacity and performance. When cabling racks together, note the following:

- Oracle recommends that the names for the servers include the rack unit number.
 This helps identify the server during diagnostics.
- When completing Oracle Exadata Deployment Assistant for the additional rack, you are prompted for SCAN addresses. However, these SCAN addresses are not used because the SCAN address from the original rack are used. Manually remove the new SCAN addresses from the generated installation files.
- The software owner account names and group names, as well as their identifiers, must match the names and identifiers of the original rack.
- If the additional grid disks are used with existing disk groups, then ensure the grid disk sizes for the new rack are the same as the original rack.



- If the InfiniBand network consists of four or more racks cabled together, then disable the Subnet Manager on the leaf switches.
- Verify the Master Subnet Manager is located on the spine switch.
- Oracle Exadata Database Machine Quarter Racks can be extended as follows:
 - Connect two Oracle Exadata Database Machine Quarter Racks together. At least four of the six ports reserved for external connectivity are open on each leaf switch. The six ports are 5B, 6A, 6B, 7A, 7B, and 12A in each leaf switch. Maintain the existing seven inter-switch links between the leaf switches within each rack. Connect the leaf switches between the racks with two links each, using the ports reserved for external connectivity.
 - Connect one Oracle Exadata Database Machine Quarter Rack with one Oracle Exadata Database Machine Half Rack or one Oracle Exadata Database Machine Full Rack. At least four ports reserved for external connectivity are open on each leaf switch. The spine switch in the Oracle Exadata Database Machine Half Rack or Oracle Exadata Database Machine Full Rack remains as the spine switch. Maintain the existing seven inter-switch links between the leaf switches within each rack. Connect the leaf switches between the racks with two links each, using the ports reserved for external connectivity.
 - Connect one Oracle Exadata Database Machine Quarter Rack with two or more Oracle Exadata Database Machine Half Racks or Oracle Exadata Database Machine Full Racks. The racks are interconnected using a fat-tree topology. Connect each leaf switch in the quarter rack to the spine switch of each half rack or full rack using two links each. If there are more than four racks, then use one link instead of two. The seven inter-switch links between the leaf switches in the quarter rack are removed.



To connect more than one quarter rack to additional racks, it is necessary to purchase Sun Datacenter InfiniBand Switch 36 switches for the quarter racks.

• If you are extending Oracle Exadata Database Machine X4-2 or later, or Oracle Exadata Database Machine X3-8 Full Rack, or Oracle Exadata Database Machine X2-2 (with X4170 and X4275 servers) half rack, then order the expansion kit that includes a Sun Datacenter InfiniBand Switch 36 switch.

Perform the following tasks before cabling racks together:

- 1. Determine the number of racks that will be cabled together.
- 2. Count the spare cables from the kit, and existing inter-switch cables.



Note:

Oracle Exadata Database Machine X4-2 and later racks or Oracle Exadata Database Machine X3-8 Full Racks with Exadata Storage Server X4-2L Servers do not include spare cables or a third Sun Datacenter InfiniBand Switch 36 switch. To extend Oracle Exadata Database Machine X4-2 and later racks or Oracle Exadata Database Machine X3-8 Full Racks with Exadata Storage Server X4-2L Servers, you must order cables and a Sun Datacenter InfiniBand Switch 36 switch.

For Oracle Exadata Racks earlier than Oracle Exadata Database Machine X4-2, no additional InfiniBand cables need to be purchased when connecting up to three Oracle Exadata Database Machine Full Racks. The following table lists the spare cables for the switch:

Cable	Spare	Inter-switch	Total
InfiniBand 2 meter cable	0	7	7
InfiniBand 3 meter cable	6	2	8
InfiniBand 5 meter cable	10	0	10

When connecting four or more racks, it is necessary to purchase additional 10 meter optical QSFP-to-QSFP QDR InfiniBand cables. Two QSFP SR transceivers are needed for each optical cable. The following table lists the number of racks, and cables needed to interconnect the racks.

Number of Racks	3 Meter Cables Needed	5 Meter Cables Needed	10 Meter Cables Needed
2	16	16	0
3	18	30	0
4	16	40	8
5	20	42	18
6	24	46	26
7	28	44	40
8	16	52	60

3. Gather the following tools:

- Pen
- Note pad
- Flashlight
- Labels for cables and rack
- Cable ties (Oracle recommends hook and loop straps)



See Also:

Oracle Exadata Database Machine Maintenance Guide for information about cable labels

- 4. Determine a naming method for the rack prefixes. For example, if the original rack has the prefix dm01, then use the prefix dm02 for the second rack, the prefix dm03 for the third rack, and so on.
- 5. Verify the racks have unique host names and IP addresses. All servers interconnected in the racks must have unique names and IP addresses.

Server names and IP addresses conventions may differ in the following cases:

- Initial installation of all Oracle Exadata Database Machine Full Racks: System address assignments and host names should be complete.
- New Oracle Exadata Database Machine Full Racks are added to an existing cluster: The new rack configuration should require unique host names and IP addresses for the new Oracle Exadata Database Machines. The IP addresses on the same subnet cannot conflict with the existing systems.
- Two existing Oracle Exadata Database Machine Full Racks are clustered together: You can assign host names and IP addresses only if Oracle Exadata Database Machines are already assigned unique host names and IP addresses, or the entire cluster must be reconfigured. The machines must be on the same subnet and not have conflicting IP addresses.
- 6. Ensure the IP addresses for the new servers are in the same subnet, and do not overlap with the currently-installed servers.
- 7. Ensure the firmware on the original switches are at the same level as the new switches using the nm2version command. If the firmware is not at the same level, then apply a firmware patch.

2.1.2 Cabling Oracle Exadata Quarter Racks and Oracle Exadata Eighth Racks

Oracle Exadata Quarter Racks and Oracle Exadata Eighth Racks can be cabled as follows:

- Oracle Exadata Quarter Rack to Oracle Exadata Quarter Rack
- Oracle Exadata Quarter Rack to Oracle Exadata Half Rack, or multiple Oracle Exadata Half Racks or Oracle Exadata Full Racks
- Oracle Exadata Quarter Rack to Oracle Exadata Full Rack, or multiple Oracle Exadata Full Racks or Oracle Exadata Half Racks
- Oracle Exadata Eighth Rack to Oracle Exadata Eighth Rack
- Oracle Exadata Eighth Rack to Oracle Exadata Half Rack, or multiple Oracle Exadata Half Racks or Oracle Exadata Full Racks
- Oracle Exadata Eighth Rack to Oracle Exadata Full Rack, or multiple Oracle Exadata Full Racks or Oracle Exadata Half Racks



Note:

- To cable additional Oracle Exadata Quarter Racks or Oracle Exadata Eighth Racks, it is necessary to purchase additional Sun Datacenter InfiniBand Switch 36 switches.
- It is assumed that four of the six ports dedicated for external connectivity in Oracle Exadata Half Rack or Oracle Exadata Full Rack are unused and available for switch interconnections with Oracle Exadata Quarter Rack or Oracle Exadata Eighth Rack. The six ports are 5B, 6A, 6B, 7A, 7B, and 12A in each leaf switch.

The following graphic shows the cable connections for two Oracle Exadata Quarter Racks. The leaf switches within each rack maintain their existing seven connections. The leaf switches interconnect between the racks with two links each using the ports reserved for external connectivity.




Figure 2-5 Leaf Switch Connections Between 2 Quarter Racks

The following graphic shows the cable connections from Oracle Exadata Quarter Rack to Oracle Exadata Half Rack or Oracle Exadata Full Rack. The leaf switches within each rack maintain their existing seven connections. The leaf switches interconnect between the racks with two links each using the ports reserved for external connectivity.







Figure 2-6 Leaf and Spine Switch Connections Between a Quarter Rack and a Half or Full Rack

The following graphic shows the cable connections from Oracle Exadata Quarter Rack to two or more racks. The racks that connect to Oracle Exadata Quarter Rack must be all Oracle Exadata Half Racks or Oracle Exadata Full Racks, interconnected using a fat-tree topology. Each leaf switch in Oracle Exadata Quarter Rack connects to the spine switches in the other half racks or full racks with two links each. If there are more than four racks, then use one link instead of two. The seven inter-switch links between the leaf switches in the quarter rack are removed.

Figure 2-7 Leaf and Spine Switch Connections for a Quarter Rack Connected to One or More Half or Full Racks



2.2 Two-Rack Cabling

Note the following when cabling two racks together:

- The spine switch is in U1 (also referred to as IB1) for all racks.
- In the following lists, the leaf switches are referred to as IB2 and IB3. Use the rack unit that is appropriate for your system.
- In Oracle Exadata Database Machine X2-2 racks and Oracle Exadata Storage
 Expansion Racks with Exadata Storage Server with Sun Fire X4270 M2 Servers,
 the leaf switches are in U20 and U24, also referred to as IB2 and IB3, respectively.
- In Oracle Exadata Database Machine X5-8 Full Rack, Oracle Exadata Database Machine X4-8 Full Rack, Oracle Exadata Database Machine X3-8 Full Rack and Oracle Exadata Database Machine X2-8 Full Rack racks, the leaf switches are in U21 and U23, also referred to as IB2 and IB3, respectively.
- In Oracle Exadata Database Machine X4-2 and later racks, Oracle Exadata
 Database Machine X3-2 Full Racks and Oracle Exadata Storage Expansion Rack
 X3-2 racks, the leaf switches are in U20 and U22, also referred to as IB2 and IB3,
 respectively.
- The cable lengths shown in the following lists assume that the racks are adjacent to each other. If the racks are not adjacent or use overhead cabling trays, then they may require longer cables lengths. Cables lengths up to 100 meters are supported.

The following list shows the cable connections for the first spine switch (R1-U1) when cabling two racks together:

- From leaf switch R1 IB3 within Rack 1, using 3 meter cables, connect:
 - R1-IB3-P8A to R1-U1-P3A
 - R1-IB3-P8B to R1-U1-P4A
 - R1-IB3-P9A to R1-U1-P5A
 - R1-IB3-P9B to R1-U1-P6A
- From leaf switch R1 IB3 to Rack 2, using 5 meter cables, connect:
 - R1-IB3-P10A to R2-U1-P7A
 - R1-IB3-P10B to R2-U1-P8A
 - R1-IB3-P11A to R2-U1-P9A
 - R1-IB3-P11B to R2-U1-P10A
- From leaf switch R1 IB2 within Rack 1, using 3 meter cables, connect::
 - R1-IB2-P8A to R1-U1-P3B
 - R1-IB2-P8B to R1-U1-P4B
 - R1-IB2-P9A to R1-U1-P5B
 - R1-IB2-P9B to R1-U1-P6B
- From leaf switch R1 IB2 to Rack 2, using 5 meter cables, connect:
 - R1-IB2-P10A to R2-U1-P7B
 - R1-IB2-P10B to R2-U1-P8B
 - R1-IB2-P11A to R2-U1-P9B
 - R1-IB2-P11B to R2-U1-P10B



The following list shows the cable connections for the second spine switch (R2-U1) when cabling two full racks together.

- From leaf switch R2 IB3 within Rack 2, using 3 meter cables, connect:
 - R2-IB3-P8A to R2-U1-P3A
 - R2-IB3-P8B to R2-U1-P4A
 - R2-IB3-P9A to R2-U1-P5A
 - R2-IB3-P9B to R2-U1-P6A
- From leaf switch R2 IB3 to Rack 1, using 5 meter cables, connect:
 - R2-IB3-P10A to R1-U1-P7A
 - R2-IB3-P10B to R1-U1-P8A
 - R2-IB3-P11A to R1-U1-P9A
 - R2-IB3-P11B to R1-U1-P10A
- From leaf switch R2 IB2 within Rack 2, using 3 meter cables, connect::
 - R2-IB2-P8A to R2-U1-P3B
 - R2-IB2-P8B to R2-U1-P4B
 - R2-IB2-P9A to R2-U1-P5B
 - R2-IB2-P9B to R2-U1-P6B
- From leaf switch R2 IB2 to Rack 1, using 5 meter cables, connect:
 - R2-IB2-P10A to R1-U1-P7B
 - R2-IB2-P10B to R1-U1-P8B
 - R2-IB2-P11A to R1-U1-P9B
 - R2-IB2-P11B to R1-U1-P10B

Related Topics

· Cabling Two Racks Together

2.3 Three-Rack Cabling

Note the following when cabling three racks together:

- The spine switch is in U1 (also referred to as IB1) for all racks.
- In the following tables, the leaf switches are referred to as IB2 and IB3. Use the rack unit that is appropriate for your system.
- In Oracle Exadata Database Machine X2-2 racks and Oracle Exadata Storage Expansion Racks with Exadata Storage Server with Sun Fire X4270 M2 Servers, the leaf switches are in U20 and U24, also referred to as IB2 and IB3, respectively.
- In Oracle Exadata Database Machine X2-8 and later, the leaf switches are in U21 and U23, also referred to as IB2 and IB3, respectively.
- In Oracle Exadata Database Machine X4-2 and later racks, or Oracle Exadata
 Database Machine X3-2 Full Racks or Oracle Exadata Storage Expansion Rack
 X3-2 racks, the leaf switches are in U20 and U22, also referred to as IB2 and IB3,
 respectively.



- The cable lengths shown in the tables assume the racks are adjacent to each other. If the racks are not adjacent or use overhead cabling trays, then they may require longer cables lengths. Up to 100 meters is supported.
- Only optical cables are supported for lengths greater than 5 meters.

Table 2-1 shows the cable connections for the first spine switch (R1-U1) when cabling three racks together:

Table 2-1 Leaf Switch Connections for the First Rack in a Three-Rack System

Leaf Switch	Connection	Cable Length
R1 IB3 within Rack 1	R1-IB3-P8A to R1-U1-P3A	3 meters
	R1-IB3-P8B to R1-U1-P4A	
	R1-IB3-P9A to R1-U1-P5A	
R1 IB3 to Rack 2	R1-IB3-P9B to R2-U1-P6A	5 meters
	R1-IB3-P10A to R2-U1-P7A	
	R1-IB3-P10B to R2-U1-P8A	
R1 IB3 to Rack 3	R1-IB3-P11A to R3-U1-P9A	5 meters
	R1-IB3-P11B to R3-U1-P10A	
R1 IB2 within Rack 1	R1-IB2-P8A to R1-U1-P3B	3 meters
	R1-IB2-P8B to R1-U1-P4B	
	R1-IB2-P9A to R1-U1-P5B	
R1 IB2 to Rack 2	R1-IB2-P9B to R2-U1-P6B	5 meters
	R1-IB2-P10A to R2-U1-P7B	
	R1-IB2-P10B to R2-U1-P8B	
R1 IB2 to Rack 3	R1-IB2-P11A to R3-U1-P9B	5 meters
	R1-IB2-P11B to R3-U1-P10B	

Table 2-2 shows the cable connections for the second spine switch (R2-U1) when cabling three racks together:

Table 2-2 Leaf Switch Connections for the Second Rack in a Three-Rack System

Leaf Switch	Connection	Cable Length
R2 IB3 within Rack 2	R2-IB3-P8A to R2-U1-P3A	3 meters
	R2-IB3-P8B to R2-U1-P4A	
	R2-IB3-P9A to R2-U1-P5A	
R2 IB3 to Rack 1	R2-IB3-P11A to R1-U1-P9A	5 meters
	R2-IB3-P11B to R1-U1-P10A	
R2 IB3 to Rack 3	R2-IB3-P9B to R3-U1-P6A	5 meters
	R2-IB3-P10A to R3-U1-P7A	
	R2-IB3-P10B to R3-U1-P8A	
R2 IB2 within Rack 2	R2-IB2-P8A to R2-U1-P3B	3 meters
	R2-IB2-P8B to R2-U1-P4B	
	R2-IB2-P9A to R2-U1-P5B	



Table 2-2 (Cont.) Leaf Switch Connections for the Second Rack in a Three-Rack System

Leaf Switch	Connection	Cable Length
R2 IB2 to Rack 1	R2-IB2-P11A to R1-U1-P9B	5 meters
	R2-IB2-P11B to R1-U1-P10B	
R2 IB2 to Rack 3	R2-IB2-P9B to R3-U1-P6B	5 meters
	R2-IB2-P10A to R3-U1-P7B	
	R2-IB2-P10B to R3-U1-P8B	

Table 2-3 shows the cable connections for the third spine switch (R3-U1) when cabling three full racks together:

Table 2-3 Leaf Switch Connections for the Third Rack in a Three-Rack System

Leaf Switch	Connection	Cable Length
R3 IB3 within Rack 3	R3-IB3-P8A to R3-U1-P3A	3 meters
	R3-IB3-P8B to R3-U1-P4A	
	R3-IB3-P9A to R3-U1-P5A	
R3 IB3 to Rack 1	R3-IB3-P9B to R1-U1-P6A	5 meters
	R3-IB3-P10A to R1-U1-P7A	
	R3-IB3-P10B to R1-U1-P8A	
R3 IB3 to Rack 2	R3-IB3-P11A to R2-U1-P9A	5 meters
	R3-IB3-P11B to R2-U1-P10A	
R3 IB2 within Rack 3	R3-IB2-P8A to R3-U1-P3B	3 meters
	R3-IB2-P8B to R3-U1-P4B	
	R3-IB2-P9A to R3-U1-P5B	
R3 IB2 to Rack 1	R3-IB2-P9B to R1-U1-P6B	5 meters
	R3-IB2-P10A to R1-U1-P7B	
	R3-IB2-P10B to R1-U1-P8B	
R3 IB2 to Rack 2	R3-IB2-P11A to R2-U1-P9B	5 meters
	R3-IB2-P11B to R2-U1-P10B	

2.4 Four-Rack Cabling

Note the following when cabling four racks together:

- The spine switch is in U1 (also referred to as IB1) for all racks.
- In the following tables, the leaf switches are referred to as IB2 and IB3. Use the rack unit that is appropriate for your system.
- In Oracle Exadata Database Machine X2-2 racks and Oracle Exadata Storage Expansion Racks with Exadata Storage Server with Sun Fire X4270 M2 Servers, the leaf switches are in U20 and U24, also referred to as IB2 and IB3, respectively.
- In Oracle Exadata Database Machine X2-8 and later, the leaf switches are in U21 and U23, also referred to as IB2 and IB3, respectively.



- In Oracle Exadata Database Machine X4-2 and later racks, or Oracle Exadata
 Database Machine X3-2 Full Racks or Oracle Exadata Storage Expansion Rack
 X3-2 racks, the leaf switches are in U20 and U22, also referred to as IB2 and IB3,
 respectively.
- The cable lengths shown in the tables assume the racks are adjacent to each other. If the racks are not adjacent or use overhead cabling trays, then they may require longer cables lengths. Up to 100 meters is supported.
- Only optical cables are supported for lengths greater than 5 meters.

Table 2-4 shows the cable connections for the first spine switch (R1-U1) when cabling four racks together:

Table 2-4 Leaf Switch Connections for the First Rack in a Four-Rack System

Leaf Switch	Connection	Cable Length
R1 IB3 within Rack 1	R1-IB3-P8A to R1-U1-P3A	3 meters
	R1-IB3-P8B to R1-U1-P4A	
R1 IB3 to Rack 2	R1-IB3-P9A to R2-U1-P5A	5 meters
	R1-IB3-P9B to R2-U1-P6A	
R1 IB3 to Rack 3	R1-IB3-P10A to R3-U1-P7A	5 meters
	R1-IB3-P10B to R3-U1-P8A	
R1 IB3 to Rack 4	R1-IB3-P11A to R4-U1-P9A	10 meters
	R1-IB3-P11B to R4-U1-P10A	
R1 IB2 within Rack 1	R1-IB2-P8A to R1-U1-P3B	3 meters
	R1-IB2-P8B to R1-U1-P4B	
R1 IB2 to Rack 2	R1-IB2-P9A to R2-U1-P5B	5 meters
	R1-IB2-P9B to R2-U1-P6B	
R1 IB2 to Rack 3	R1-IB2-P10A to R3-U1-P7B	5 meters
	R1-IB2-P10B to R3-U1-P8B	
R1 IB2 to Rack 4	R1-IB2-P11A to R4-U1-P9B	10 meters
	R1-IB2-P11B to R4-U1-P10B	

Table 2-5 shows the cable connections for the second spine switch (R2-U1) when cabling four full racks together:

Table 2-5 Leaf Switch Connections for the Second Rack in a Four-Rack System

Leaf Switch	Connection	Cable Length
R2 IB3 within Rack 2	R2-IB3-P8A to R2-U1-P3A R2-IB3-P8B to R2-U1-P4A	3 meters
R2 IB3 to Rack 1	R2-IB3-P11A to R1-U1-P9A R2-IB3-P11B to R1-U1-P10A	5 meters
R2 IB3 to Rack 3	R2-IB3-P9A to R3-U1-P5A R2-IB3-P9B to R3-U1-P6A	5 meters
R2 IB3 to Rack 4	R2-IB3-P10A to R4-U1-P7A R2-IB3-P10B to R4-U1-P8A	5 meters



Table 2-5 (Cont.) Leaf Switch Connections for the Second Rack in a Four-Rack System

Leaf Switch	Connection	Cable Length
R2 IB2 within Rack 2	R2-IB2-P8A to R2-U1-P3B R2-IB2-P8B to R2-U1-P4B	3 meters
R2 IB2 to Rack 1	R2-IB2-P11A to R1-U1-P9B R2-IB2-P11B to R1-U1-P10B	5 meters
R2 IB2 to Rack 3	R2-IB2-P9A to R3-U1-P5B R2-IB2-P9B to R3-U1-P6B	5 meters
R2 IB2 to Rack 4	R2-IB2-P10A to R4-U1-P7B R2-IB2-P10B to R4-U1-P8B	5 meters

Table 2-6 shows the cable connections for the third spine switch (R3-U1) when cabling four full racks together:

Table 2-6 Leaf Switch Connections for the Third Rack in a Four-Rack System

Leaf Switch	Connection	Cable Length
Lear Switch	Connection	Cable Length
R3 IB3 within Rack 3	R3-IB3-P8A to R3-U1-P3A	3 meters
	R3-IB3-P8B to R3-U1-P4A	
R3 IB3 to Rack 1	R3-IB3-P10A to R1-U1-P7A	5 meters
	R3-IB3-P10B to R1-U1-P8A	
R3 IB3 to Rack 2	R3-IB3-P11A to R2-U1-P9A	5 meters
	R3-IB3-P11B to R2-U1-P10A	
R3 IB3 to Rack 4	R3-IB3-P9A to R4-U1-P5A	5 meters
	R3-IB3-P9B to R4-U1-P6A	
R3 IB2 within Rack 3	R3-IB2-P8A to R3-U1-P3B	3 meters
	R3-IB2-P8B to R3-U1-P4B	
R3 IB2 to Rack 1	R3-IB2-P10A to R1-U1-P7B	5 meters
	R3-IB2-P10B to R1-U1-P8B	
R3 IB2 to Rack 2	R3-IB2-P11A to R2-U1-P9B	5 meters
	R3-IB2-P11B to R2-U1-P10B	
R3 IB2 to Rack 4	R3-IB2-P9A to R4-U1-P5B	5 meters
	R3-IB2-P9B to R4-U1-P6B	

Table 2-7 shows the cable connections for the fourth spine switch (R4-U1) when cabling four full racks together:

Table 2-7 Leaf Switch Connections for the Fourth Rack in a Four-Rack System

Leaf Switch	Connection	Cable Length
R4 IB3 within Rack 4	R4-IB3-P8A to R4-U1-P3A	3 meters
	R4-IB3-P8B to R4-U1-P4A	



Table 2-7 (Cont.) Leaf Switch Connections for the Fourth Rack in a Four-Rack System

Leaf Switch	Connection	Cable Length
R4 IB3 to Rack 1	R4-IB3-P9A to R1-U1-P5A	10 meters
	R4-IB3-P9B to R1-U1-P6A	
R4 IB3 to Rack 2	R4-IB3-P10A to R2-U1-P7A	5 meters
	R4-IB3-P10B to R2-U1-P8A	
R4 IB3 to Rack 3	R4-IB3-P11A to R3-U1-P9A	5 meters
	R4-IB3-P11B to R3-U1-P10A	
R4 IB2 within Rack 4	R4-IB2-P8A to R4-U1-P3B	3 meters
	R4-IB2-P8B to R4-U1-P4B	
R4 IB2 to Rack 1	R4-IB2-P9A to R1-U1-P5B	10 meters
	R4-IB2-P9B to R1-U1-P6B	
R4 IB2 to Rack 2	R4-IB2-P10A to R2-U1-P7B	5 meters
	R4-IB2-P10B to R2-U1-P8B	
R4 IB2 to Rack 3	R4-IB2-P11A to R3-U1-P9B	5 meters
	R4-IB2-P11B to R3-U1-P10B	

2.5 Five-Rack Cabling

Note the following when cabling five racks together:

- The spine switch is in U1 (also referred to as IB1) for all racks.
- In the following tables, the leaf switches are referred to as IB2 and IB3. Use the rack unit that is appropriate for your system.
- In Oracle Exadata Database Machine X2-2 racks and Oracle Exadata Storage
 Expansion Racks with Exadata Storage Server with Sun Fire X4270 M2 Servers,
 the leaf switches are in U20 and U24, also referred to as IB2 and IB3, respectively.
- In Oracle Exadata Database Machine X2-8 and later racks, the leaf switches are in U21 and U23, also referred to as IB2 and IB3, respectively.
- In Oracle Exadata Database Machine X4-2 and later racks, or Oracle Exadata
 Database Machine X3-2 Full Racks and Oracle Exadata Storage Expansion Rack
 X3-2 racks, the leaf switches are in U20 and U22, also referred to as IB2 and IB3,
 respectively.
- The cable lengths shown in the tables assume the racks are adjacent to each other. If the racks are not adjacent or use overhead cabling trays, then they may require longer cables lengths. Up to 100 meters is supported.
- Only optical cables are supported for lengths greater than 5 meters.

Table 2-8 shows the cable connections for the first spine switch (R1-U1) when cabling five racks together:



Table 2-8 Leaf Switch Connections for the First Rack in a Five-Rack System

Leaf Switch	Connection	Cable Length
R1 IB3 within Rack 1	R1-IB3-P8A to R1-U1-P3A	3 meters
	R1-IB3-P8B to R1-U1-P4A	
R1 IB3 to Rack 2	R1-IB3-P9A to R2-U1-P5A	5 meters
	R1-IB3-P9B to R2-U1-P6A	
R1 IB3 to Rack 3	R1-IB3-P10A to R3-U1-P7A	5 meters
	R1-IB3-P10B to R3-U1-P8A	
R1 IB3 to Rack 4	R1-IB3-P11A to R4-U1-P9A	10 meters
R1 IB3 to Rack 5	R1-IB3-P11B to R5-U1-P10A	10 meters
R1 IB2 within Rack 1	R1-IB2-P8A to R1-U1-P3B	3 meters
	R1-IB2-P8B to R1-U1-P4B	
R1 IB2 to Rack 2	R1-IB2-P9A to R2-U1-P5B	5 meters
	R1-IB2-P9B to R2-U1-P6B	
R1 IB2 to Rack 3	R1-IB2-P10A to R3-U1-P7B	5 meters
	R1-IB2-P10B to R3-U1-P8B	
R1 IB2 to Rack 4	R1-IB2-P11A to R4-U1-P9B	10 meters
R1 IB2 to Rack 5	R1-IB2-P11B to R5-U1-P10B	10 meters

Table 2-9 Leaf Switch Connections for the Second Rack in a Five-Rack System

Leaf Switch	Connection	Cable Length
R2 IB3 within Rack 2	R2-IB3-P8A to R2-U1-P3A	3 meters
	R2-IB3-P8B to R2-U1-P4A	
R2 IB3 to Rack 1	R2-IB3-P11B to R1-U1-P10A	5 meters
R2 IB3 to Rack 3	R2-IB3-P9A to R3-U1-P5A	5 meters
	R2-IB3-P9B to R3-U1-P6A	
R2 IB3 to Rack 4	R2-IB3-P10A to R4-U1-P7A	5 meters
	R2-IB3-P10B to R4-U1-P8A	
R2 IB3 to Rack 5	R2-IB3-P11A to R5-U1-P9A	10 meters
R2 IB2 within Rack 2	R2-IB2-P8A to R2-U1-P3B	3 meters
	R2-IB2-P8B to R2-U1-P4B	
R2 IB2 to Rack 1	R2-IB2-P11B to R1-U1-P10B	5 meters
R2 IB2 to Rack 3	R2-IB2-P9A to R3-U1-P5B	5 meters
	R2-IB2-P9B to R3-U1-P6B	
R2 IB2 to Rack 4	R2-IB2-P10A to R4-U1-P7B	5 meters
	R2-IB2-P10B to R4-U1-P8B	
R2 IB2 to Rack 5	R2-IB2-P11A to R5-U1-P9B	10 meters



Table 2-10 shows the cable connections for the third spine switch (R3-U1) when cabling five full racks together:

Table 2-10 Leaf Switch Connections for the Third Rack in a Five-Rack System

Leaf Switch	Connection	Cable Length
R3 IB3 within Rack 3	R3-IB3-P8A to R3-U1-P3A	3 meters
	R3-IB3-P8B to R3-U1-P4A	
R3 IB3 to Rack 1	R3-IB3-P11A to R1-U1-P9A	5 meters
R3 IB3 to Rack 2	R3-IB3-P11B to R2-U1-P10A	5 meters
R3 IB3 to Rack 4	R3-IB3-P9A to R4-U1-P5A	5 meters
	R3-IB3-P9B to R4-U1-P6A	
R3 IB3 to Rack 5	R3-IB3-P10A to R5-U1-P7A	5 meters
	R3-IB3-P10B to R5-U1-P8A	
R3 IB2 within Rack 3	R3-IB2-P8A to R3-U1-P3B	3 meters
	R3-IB2-P8B to R3-U1-P4B	
R3 IB2 to Rack 1	R3-IB2-P11A to R1-U1-P9B	5 meters
R3 IB2 to Rack 2	R3-IB2-P11B to R2-U1-P10B	5 meters
R3 IB2 to Rack 4	R3-IB2-P9A to R4-U1-P5B	5 meters
	R3-IB2-P9B to R4-U1-P6B	
R3 IB2 to Rack 5	R3-IB2-P10A to R5-U1-P7B	5 meters
	R3-IB2-P10B to R5-U1-P8B	

Table 2-11 shows the cable connections for the fourth spine switch (R4-U1) when cabling five full racks together:

Table 2-11 Leaf Switch Connections for the Fourth Rack in a Five-Rack System

Leaf Switch	Connection	Cable Length
R4 IB3 within Rack 4	R4-IB3-P8A to R4-U1-P3A	3 meters
	R4-IB3-P8B to R4-U1-P4A	
R4 IB3 to Rack 1	R4-IB3-P10A to R1-U1-P7A	10 meters
	R4-IB3-P10B to R1-U1-P8A	
R4 IB3 to Rack 2	R4-IB3-P11A to R2-U1-P9A	5 meters
R4 IB3 to Rack 3	R4-IB3-P11B to R3-U1-P10A	5 meters
R4 IB3 to Rack 5	R4-IB3-P9A to R5-U1-P5A	5 meters
	R4-IB3-P9B to R5-U1-P6A	
R4 IB2 within Rack 4	R4-IB2-P8A to R4-U1-P3B	3 meters
	R4-IB2-P8B to R4-U1-P4B	
R4 IB2 to Rack 1	R4-IB2-P10A to R1-U1-P7B	10 meters
	R4-IB2-P10B to R1-U1-P8B	
R4 IB2 to Rack 2	R4-IB2-P11A to R2-U1-P9B	5 meters
R4 IB2 to Rack 3	R4-IB2-P11B to R3-U1-P10B	5 meters



Table 2-11 (Cont.) Leaf Switch Connections for the Fourth Rack in a Five-Rack System

Leaf Switch	Connection	Cable Length
R4 IB2 to Rack 5	R4-IB2-P9A to R5-U1-P5B R4-IB2-P9B to R5-U1-P6B	5 meters

Table 2-12 shows the cable connections for the fifth spine switch (R5-U1) when cabling five full racks together:

Table 2-12 Leaf Switch Connections for the Fifth Rack in a Five-Rack System

Leaf Switch	Connection	Cable Length
R5 IB3 within Rack 5	R5-IB3-P8A to R5-U1-P3A	3 meters
	R5-IB3-P8B to R5-U1-P4A	
R5 IB3 to Rack 1	R5-IB3-P9A to R1-U1-P5A	10 meters
	R5-IB3-P9B to R1-U1-P6A	
R5 IB3 to Rack 2	R5-IB3-P10A to R2-U1-P7A	10 meters
	R5-IB3-P10B to R2-U1-P8A	
R5 IB3 to Rack 3	R5-IB3-P11A to R3-U1-P9A	5 meters
R5 IB3 to Rack 4	R5-IB3-P11B to R4-U1-P10A	5 meters
R5 IB2 within Rack 5	R5-IB2-P8A to R5-U1-P3B	3 meters
	R5-IB2-P8B to R5-U1-P4B	
R5 IB2 to Rack 1	R5-IB2-P9A to R1-U1-P5B	10 meters
	R5-IB2-P9B to R1-U1-P6B	
R5 IB2 to Rack 2	R5-IB2-P10A to R2-U1-P7B	10 meters
	R5-IB2-P10B to R2-U1-P8B	
R5 IB2 to Rack 3	R5-IB2-P11A to R3-U1-P9B	5 meters
R5 IB2 to Rack 4	R5-IB2-P11B to R4-U1-P10B	5 meters

2.6 Six-Rack Cabling

Note the following when cabling six racks together:

- The spine switch is in U1 (also referred to as IB1) for all racks.
- In the following tables, the leaf switches are referred to as IB2 and IB3. Use the rack unit that is appropriate for your system.
- In Oracle Exadata Database Machine X2-2 racks and Oracle Exadata Storage Expansion Racks with Exadata Storage Server with Sun Fire X4270 M2 Servers, the leaf switches are in U20 and U24, also referred to as IB2 and IB3, respectively.
- In Oracle Exadata Database Machine X2-8 and later 8-socket racks, the leaf switches are in U21 and U23, also referred to as IB2 and IB3, respectively.
- In Oracle Exadata Database Machine X4-2 and later two-socket racks, or Oracle Exadata Database Machine X3-2 Full Racks and Oracle Exadata Storage



Expansion Rack X3-2 racks, the leaf switches are in U20 and U22, also referred to as IB2 and IB3, respectively.

- The cable lengths shown in the tables assume the racks are adjacent to each other. If the racks are not adjacent or use overhead cabling trays, then they may require longer cables lengths. Up to 100 meters is supported.
- Only optical cables are supported for lengths greater than 5 meters.

Table 2-13 shows the cable connections for the first spine switch (R1-U1) when cabling six racks together:

Table 2-13 Leaf Switch Connections for the First Rack in a Six-Rack System

Leaf Switch	Connection	Cable Length
R1 IB3 within Rack 1	R1-IB3-P8A to R1-U1-P3A	3 meters
	R1-IB3-P8B to R1-U1-P4A	
R1 IB3 to Rack 2	R1-IB3-P9A to R2-U1-P5A	5 meters
	R1-IB3-P9B to R2-U1-P6A	
R1 IB3 to Rack 3	R1-IB3-P10A to R3-U1-P7A	5 meters
R1 IB3 to Rack 4	R1-IB3-P10B to R4-U1-P8A	10 meters
R1 IB3 to Rack 5	R1-IB3-P11A to R5-U1-P9A	10 meters
R1 IB3 to Rack 6	R1-IB3-P11B to R6-U1-P10A	10 meters
R1 IB2 within Rack 1	R1-IB2-P8A to R1-U1-P3B	3 meters
	R1-IB2-P8B to R1-U1-P4B	
R1 IB2 to Rack 2	R1-IB2-P9A to R2-U1-P5B	5 meters
	R1-IB2-P9B to R2-U1-P6B	
R1 IB2 to Rack 3	R1-IB2-P10A to R3-U1-P7B	5 meters
R1 IB2 to Rack 4	R1-IB2-P10B to R4-U1-P8B	10 meters
R1 IB2 to Rack 5	R1-IB2-P11A to R5-U1-P9B	10 meters
R1 IB2 to Rack 6	R1-IB2-P11B to R6-U1-P10B	10 meters

Table 2-14 shows the cable connections for the second spine switch (R2-U1) when cabling six full racks together:

Table 2-14 Leaf Switch Connections for the Second Rack in a Six-Rack System

Leaf Switch	Connection	Cable Length
R2 IB3 within Rack 2	R2-IB3-P8A to R2-U1-P3A	3 meters
	R2-IB3-P8B to R2-U1-P4A	
R2 IB3 to Rack 1	R2-IB3-P11B to R1-U1-P10A	5 meters
R2 IB3 to Rack 3	R2-IB3-P9A to R3-U1-P5A	5 meters
	R2-IB3-P9B to R3-U1-P6A	
R2 IB3 to Rack 4	R2-IB3-P10A to R4-U1-P7A	5 meters
R2 IB3 to Rack 5	R2-IB3-P10B to R5-U1-P8A	10 meters
R2 IB3 to Rack 6	R2-IB3-P11Ato R6-U1-P9A	10 meters



Table 2-14 (Cont.) Leaf Switch Connections for the Second Rack in a Six-Rack System

Leaf Switch	Connection	Cable Length
R2 IB2 within Rack 2	R2-IB2-P8A to R2-U1-P3B	3 meters
R2 IB2 to Rack 1	R2-IB2-P8B to R2-U1-P4B R2-IB2-P11B to R1-U1-P10B	5 meters
R2 IB2 to Rack 3	R2-IB2-P9A to R3-U1-P5B R2-IB2-P9B to R3-U1-P6B	5 meters
R2 IB2 to Rack 4	R2-IB2-P10A to R4-U1-P7B	5 meters
R2 IB2 to Rack 5	R2-IB2-P10B to R5-U1-P8B	10 meters
R2 IB2 to Rack 6	R2-IB2-P11Ato R6-U1-P9B	10 meters

Table 2-15 shows the cable connections for the third spine switch (R3-U1) when cabling six full racks together:

Table 2-15 Leaf Switch Connections for the Third Rack in a Six-Rack System

Leaf Switch	Connection	Cable Length
R3 IB3 within Rack 3	R3-IB3-P8A to R3-U1-P3A	3 meters
	R3-IB3-P8B to R3-U1-P4A	
R3 IB3 to Rack 1	R3-IB3-P11A to R1-U1-P9A	5 meters
R3 IB3 to Rack 2	R3-IB3-P11B to R2-U1-P10A	5 meters
R3 IB3 to Rack 4	R3-IB3-P9A to R4-U1-P5A	5 meters
	R3-IB3-P9B to R4-U1-P6A	
R3 IB3 to Rack 5	R3-IB3-P10A to R5-U1-P7A	5 meters
R3 IB3 to Rack 6	R3-IB3-P10B to R6-U1-P8A	10 meters
R3 IB2 within Rack 3	R3-IB2-P8A to R3-U1-P3B	3 meters
	R3-IB2-P8B to R3-U1-P4B	
R3 IB2 to Rack 1	R3-IB2-P11A to R1-U1-P9B	5 meters
R3 IB2 to Rack 2	R3-IB2-P11B to R2-U1-P10B	5 meters
R3 IB2 to Rack 4	R3-IB2-P9A to R4-U1-P5B	5 meters
	R3-IB2-P9B to R4-U1-P6B	
R3 IB2 to Rack 5	R3-IB2-P10A to R5-U1-P7B	5 meters
R3 IB2 to Rack 6	R3-IB2-P10B to R6-U1-P8B	10 meters

Table 2-16 shows the cable connections for the fourth spine switch (R4-U1) when cabling six full racks together:

Table 2-16 Leaf Switch Connections for the Fourth Rack in a Six-Rack System

Leaf Switch	Connection	Cable Length
R4 IB3 within Rack 4	R4-IB3-P8A to R4-U1-P3A	3 meters
	R4-IB3-P8B to R4-U1-P4A	



Table 2-16 (Cont.) Leaf Switch Connections for the Fourth Rack in a Six-Rack System

Leaf Switch	Connection	Cable Length
R4 IB3 to Rack 1	R4-IB3-P10B to R1-U1-P8A	10 meters
R4 IB3 to Rack 2	R4-IB3-P11A to R2-U1-P9A	5 meters
R4 IB3 to Rack 3	R4-IB3-P11B to R3-U1-P10A	5 meters
R4 IB3 to Rack 5	R4-IB3-P9A to R5-U1-P5A R4-IB3-P9B to R5-U1-P6A	5 meters
R4 IB3 to Rack 6	R4-IB3-P10A to R6-U1-P7A	5 meters
R4 IB2 within Rack 4	R4-IB2-P8A to R4-U1-P3B R4-IB2-P8B to R4-U1-P4B	3 meters
R4 IB2 to Rack 1	R4-IB2-P10B to R1-U1-P8B	10 meters
R4 IB2 to Rack 2	R4-IB2-P11A to R2-U1-P9B	5 meters
R4 IB2 to Rack 3	R4-IB2-P11B to R3-U1-P10B	5 meters
R4 IB2 to Rack 5	R4-IB2-P9A to R5-U1-P5B R4-IB2-P9B to R5-U1-P6B	5 meters
R4 IB2 to Rack 6	R4-IB2-P10A to R6-U1-P7B	5 meters

Table 2-17 shows the cable connections for the fifth spine switch (R5-U1) when cabling six full racks together:

Table 2-17 Leaf Switch Connections for the Fifth Rack in a Six-Rack System

Leaf Switch	Connection	Cable Length
R5 IB3 within Rack 5	R5-IB3-P8A to R5-U1-P3A	3 meters
	R5-IB3-P8B to R5-U1-P4A	
R5 IB3 to Rack 1	R5-IB3-P10A to R1-U1-P7A	10 meters
R5 IB3 to Rack 2	R5-IB3-P10B to R2-U1-P8A	10 meters
R5 IB3 to Rack 3	R5-IB3-P11A to R3-U1-P9A	5 meters
R5 IB3 to Rack 4	R5-IB3-P11B to R4-U1-P10A	5 meters
R5 IB3 to Rack 6	R5-IB3-P9A to R6-U1-P5A	5 meters
	R5-IB3-P9B to R6-U1-P6A	
R5 IB2 within Rack 5	R5-IB2-P8A to R5-U1-P3B	3 meters
	R5-IB2-P8B to R5-U1-P4B	
R5 IB2 to Rack 1	R5-IB2-P10A to R1-U1-P7B	10 meters
R5 IB2 to Rack 2	R5-IB2-P10B to R2-U1-P8B	10 meters
R5 IB2 to Rack 3	R5-IB2-P11A to R3-U1-P9B	5 meters
R5 IB2 to Rack 4	R5-IB2-P11B to R4-U1-P10B	5 meters
R5 IB2 to Rack 6	R5-IB2-P9A to R6-U1-P5B	5 meters
	R5-IB2-P9B to R6-U1-P6B	



Table 2-18 shows the cable connections for the sixth spine switch (R6-U1) when cabling six full racks together:

Table 2-18 Leaf Switch Connections for the Sixth Rack in a Six-Rack System

Leaf Switch	Connection	Cable Length
R6 IB3 within Rack 6	R6-IB3-P8A to R6-U1-P3A	3 meters
	R6-IB3-P8B to R6-U1-P4A	
R6 IB3 to Rack 1	R6-IB3-P9A to R1-U1-P5A	10 meters
	R6-IB3-P9B to R1-U1-P6A	
R6 IB3 to Rack 2	R6-IB3-P10A to R2-U1-P7A	10 meters
R6 IB3 to Rack 3	R6-IB3-P10B to R3-U1-P8A	10 meters
R6 IB3 to Rack 4	R6-IB3-P11A to R4-U1-P9A	5 meters
R6 IB3 to Rack 5	R6-IB3-P11B to R5-U1-P10A	5 meters
R6 IB2 within Rack 6	R6-IB2-P8A to R6-U1-P3B	3 meters
	R6-IB2-P8B to R6-U1-P4B	
R6 IB2 to Rack 1	R6-IB2-P9A to R1-U1-P5B	10 meters
	R6-IB2-P9B to R1-U1-P6B	
R6 IB2 to Rack 2	R6-IB2-P10A to R2-U1-P7B	10 meters
R6 IB2 to Rack 3	R6-IB2-P10B to R3-U1-P8B	10 meters
R6 IB2 to Rack 4	R6-IB2-P11A to R4-U1-P9B	5 meters
R6 IB2 to Rack 5	R6-IB2-P11B to R5-U1-P10B	5 meters

2.7 Seven-Rack Cabling

Note the following when cabling five racks together:

- The spine switch is in U1 (also referred to as IB1) for all racks.
- In the following tables, the leaf switches are referred to as IB2 and IB3. Use the rack unit that is appropriate for your system.
- In Oracle Exadata Database Machine X2-2 racks and Oracle Exadata Storage Expansion Racks with Exadata Storage Server with Sun Fire X4270 M2 Servers, the leaf switches are in U20 and U24, also referred to as IB2 and IB3, respectively.
- In Oracle Exadata Database Machine X2-8 and later racks, the leaf switches are in U21 and U23, also referred to as IB2 and IB3, respectively.
- In Oracle Exadata Database Machine X4-2 and later racks, or Oracle Exadata
 Database Machine X3-2 Full Racks and Oracle Exadata Storage Expansion Rack
 X3-2 racks, the leaf switches are in U20 and U22, also referred to as IB2 and IB3,
 respectively.
- The cable lengths shown in the tables assume the racks are adjacent to each other. If the racks are not adjacent or use overhead cabling trays, then they may require longer cables lengths. Up to 100 meters is supported.
- Only optical cables are supported for lengths greater than 5 meters.

Table 2-19 shows the cable connections for the first spine switch (R1-U1) when cabling seven racks together:



Table 2-19 Leaf Switch Connections for the First Rack in a Seven-Rack System

Leaf Switch	Connection	Cable Length
R1 IB3 within Rack 1	R1-IB3-P8A to R1-U1-P3A	3 meters
	R1-IB3-P8B to R1-U1-P4A	
R1 IB3 to Rack 2	R1-IB3-P9A to R2-U1-P5A	5 meters
R1 IB3 to Rack 3	R1-IB3-P9B to R3-U1-P6A	5 meters
R1 IB3 to Rack 4	R1-IB3-P10A to R4-U1-P7A	10 meters
R1 IB3 to Rack 5	R1-IB3-P10B to R5-U1-P8A	10 meters
R1 IB3 to Rack 6	R1-IB3-P11A to R6-U1-P9A	10 meters
R1 IB3 to Rack 7	R1-IB3-P11B to R7-U1-P10A	10 meters
R1 IB2 within Rack 1	R1-IB2-P8A to R1-U1-P3B	3 meters
	R1-IB2-P8B to R1-U1-P4B	
R1 IB2 to Rack 2	R1-IB2-P9A to R2-U1-P5B	5 meters
R1 IB2 to Rack 3	R1-IB2-P9B to R3-U1-P6B	5 meters
R1 IB2 to Rack 4	R1-IB2-P10A to R4-U1-P7B	10 meters
R1 IB2 to Rack 5	R1-IB2-P10B to R5-U1-P8B	10 meters
R1 IB2 to Rack 6	R1-IB2-P11A to R6-U1-P9B	10 meters
R1 IB2 to Rack 7	R1-IB2-P11B to R7-U1-P10B	10 meters

Table 2-20 shows the cable connections for the second spine switch (R2-U1) when cabling seven full racks together:

Table 2-20 Leaf Switch Connections for the Second Rack in a Seven-Rack System

Leaf Switch	Connection	Cable Length
R2 IB3 within Rack 2	R2-IB3-P8A to R2-U1-P3A	3 meters
	R2-IB3-P8B to R2-U1-P4A	
R2 IB3 to Rack 1	R2-IB3-P11B to R1-U1-P10A	5 meters
R2 IB3 to Rack 3	R2-IB3-P9A to R3-U1-P5A	5 meters
R2 IB3 to Rack 4	R2-IB3-P9B to R4-U1-P6A	5 meters
R2 IB3 to Rack 5	R2-IB3-P10A to R5-U1-P7A	10 meters
R2 IB3 to Rack 6	R2-IB3-P10B to R6-U1-P8A	10 meters
R2 IB3 to Rack 7	R2-IB3-P11A to R7-U1-P9A	10 meters
R2 IB2 within Rack 2	R2-IB2-P8A to R2-U1-P3B	3 meters
	R2-IB2-P8B to R2-U1-P4B	
R2 IB2 to Rack 1	R2-IB2-P11B to R1-U1-P10B	5 meters
R2 IB2 to Rack 3	R2-IB2-P9A to R3-U1-P5B	5 meters
R2 IB2 to Rack 4	R2-IB2-P9B to R4-U1-P6B	5 meters
R2 IB2 to Rack 5	R2-IB2-P10A to R5-U1-P7B	10 meters
R2 IB2 to Rack 6	R2-IB2-P10Bto R6-U1-P8B	10 meters



Table 2-20 (Cont.) Leaf Switch Connections for the Second Rack in a Seven-Rack System

Leaf Switch	Connection	Cable Length
R2 IB2 to Rack 7	R2-IB2-P11A to R7-U1-P9B	10 meters

Table 2-21 shows the cable connections for the third spine switch (R3-U1) when cabling seven full racks together:

Table 2-21 Leaf Switch Connections for the Third Rack in a Seven-Rack System

Leaf Switch	Connection	Cable Length
R3 IB3 within Rack 3	R3-IB3-P8A to R3-U1-P3A	3 meters
	R3-IB3-P8B to R3-U1-P4A	
R3 IB3 to Rack 1	R3-IB3-P11A to R1-U1-P9A	5 meters
R3 IB3 to Rack 2	R3-IB3-P11B to R2-U1-P10A	5 meters
R3 IB3 to Rack 4	R3-IB3-P9A to R4-U1-P5A	5 meters
R3 IB3 to Rack 5	R3-IB3-P9B to R5-U1-P6A	5 meters
R3 IB3 to Rack 6	R3-IB3-P10A to R6-U1-P7A	10 meters
R3 IB3 to Rack 7	R3-IB3-P10B to R7-U1-P8A	10 meters
R3 IB2 within Rack 3	R3-IB2-P8A to R3-U1-P3B	3 meters
	R3-IB2-P8B to R3-U1-P4B	
R3 IB2 to Rack 1	R3-IB2-P11A to R1-U1-P9B	5 meters
R3 IB2 to Rack 2	R3-IB2-P11B to R2-U1-P10B	5 meters
R3 IB2 to Rack 4	R3-IB2-P9A to R4-U1-P5B	5 meters
R3 IB2 to Rack 5	R3-IB2-P9B to R5-U1-P6B	5 meters
R3 IB2 to Rack 6	R3-IB2-P10A to R6-U1-P7B	10 meters
R3 IB2 to Rack 7	R3-IB2-P10B to R7-U1-P8B	10 meters

Table 2-22 shows the cable connections for the fourth spine switch (R4-U1) when cabling seven full racks together:

Table 2-22 Leaf Switch Connections for the Fourth Rack in a Seven-Rack System

Leaf Switch	Connection	Cable Length
R4 IB3 within Rack 4	R4-IB3-P8A to R4-U1-P3A	3 meters
	R4-IB3-P8B to R4-U1-P4A	
R4 IB3 to Rack 1	R4-IB3-P10B to R1-U1-P8A	10 meters
R4 IB3 to Rack 2	R4-IB3-P11A to R2-U1-P9A	5 meters
R4 IB3 to Rack 3	R4-IB3-P11B to R3-U1-P10A	5 meters
R4 IB3 to Rack 5	R4-IB3-P9A to R5-U1-P5A	5 meters
R4 IB3 to Rack 6	R4-IB3-P9B to R6-U1-P6A	5 meters



Table 2-22 (Cont.) Leaf Switch Connections for the Fourth Rack in a Seven-Rack System

Leaf Switch	Connection	Cable Length
R4 IB3 to Rack 7	R4-IB3-P10A to R7-U1-P7A	10 meters
R4 IB2 within Rack 4	R4-IB2-P8A to R4-U1-P3B R4-IB2-P8B to R4-U1-P4B	3 meters
R4 IB2 to Rack 1	R4-IB2-P10B to R1-U1-P8B	10 meters
R4 IB2 to Rack 2	R4-IB2-P11A to R2-U1-P9B	5 meters
R4 IB2 to Rack 3	R4-IB2-P11B to R3-U1-P10B	5 meters
R4 IB2 to Rack 5	R4-IB2-P9A to R5-U1-P5B	5 meters
R4 IB2 to Rack 6	R4-IB2-P9B to R6-U1-P6B	5 meters
R4 IB2 to Rack 7	R4-IB2-P10A to R7-U1-P7B	10 meters

Table 2-23 shows the cable connections for the fifth spine switch (R5-U1) when cabling seven full racks together:

Table 2-23 Leaf Switch Connections for the Fifth Rack in a Seven-Rack System

Leaf Switch	Connection	Cable Length
R5 IB3 within Rack 5	R5-IB3-P8A to R5-U1-P3A	3 meters
	R5-IB3-P8B to R5-U1-P4A	
R5 IB3 to Rack 1	R5-IB3-P10A to R1-U1-P7A	10 meters
R5 IB3 to Rack 2	R5-IB3-P10B to R2-U1-P8A	10 meters
R5 IB3 to Rack 3	R5-IB3-P11A to R3-U1-P9A	5 meters
R5 IB3 to Rack 4	R5-IB3-P11B to R4-U1-P10A	5 meters
R5 IB3 to Rack 6	R5-IB3-P9A to R6-U1-P5A	5 meters
R5 IB3 to Rack 7	R5-IB3-P9B to R7-U1-P6A	5 meters
R5 IB2 within Rack 5	R5-IB2-P8A to R5-U1-P3B	3 meters
	R5-IB2-P8B to R5-U1-P4B	
R5 IB2 to Rack 1	R5-IB2-P10A to R1-U1-P7B	10 meters
R5 IB2 to Rack 2	R5-IB2-P10B to R2-U1-P8B	10 meters
R5 IB2 to Rack 3	R5-IB2-P11A to R3-U1-P9B	5 meters
R5 IB2 to Rack 4	R5-IB2-P11B to R4-U1-P10B	5 meters
R5 IB2 to Rack 6	R5-IB2-P9A to R6-U1-P5B	5 meters
R5 IB2 to Rack 7	R5-IB2-P9B to R7-U1-P6B	5 meters

Table 2-24 shows the cable connections for the sixth spine switch (R6-U1) when cabling seven full racks together:

Table 2-24 Leaf Switch Connections for the Sixth Rack in a Seven-Rack System

Leaf Switch	Connection	Cable Length
R6 IB3 within Rack 6	R6-IB3-P8A to R6-U1-P3A	3 meters
	R6-IB3-P8B to R6-U1-P4A	
R6 IB3 to Rack 1	R6-IB3-P9B to R1-U1-P6A	10 meters
R6 IB3 to Rack 2	R6-IB3-P10A to R2-U1-P7A	10 meters
R6 IB3 to Rack 3	R6-IB3-P10B to R3-U1-P8A	10 meters
R6 IB3 to Rack 4	R6-IB3-P11A to R4-U1-P9A	5 meters
R6 IB3 to Rack 5	R6-IB3-P11B to R5-U1-P10A	5 meters
R6 IB3 to Rack 7	R6-IB3-P9A to R7-U1-P5A	5 meters
R6 IB2 within Rack 6	R6-IB2-P8A to R6-U1-P3B	3 meters
	R6-IB2-P8B to R6-U1-P4B	
R6 IB2 to Rack 1	R6-IB2-P9B to R1-U1-P6B	10 meters
R6 IB2 to Rack 2	R6-IB2-P10A to R2-U1-P7B	10 meters
R6 IB2 to Rack 3	R6-IB2-P10B to R3-U1-P8B	10 meters
R6 IB2 to Rack 4	R6-IB2-P11A to R4-U1-P9B	5 meters
R6 IB2 to Rack 5	R6-IB2-P11B to R5-U1-P10B	5 meters
R6 IB2 to Rack 7	R6-IB2-P9A to R7-U1-P5B	5 meters

Table 2-25 shows the cable connections for the seventh spine switch (R7-U1) when cabling seven full racks together:

Table 2-25 Leaf Switch Connections for the Seventh Rack in a Seven-Rack System

Leaf Switch	Connection	Cable Length
R7 IB3 within Rack 7	R7-IB3-P8A to R7-U1-P3A	3 meters
	R7-IB3-P8B to R7-U1-P4A	
R7 IB3 to Rack 1	R7-IB3-P9A to R1-U1-P5A	10 meters
R7 IB3 to Rack 2	R7-IB3-P9B to R2-U1-P6A	10 meters
R7 IB3 to Rack 3	R7-IB3-P10A to R3-U1-P7A	10 meters
R7 IB3 to Rack 4	R7-IB3-P10B to R4-U1-P8A	10 meters
R7 IB3 to Rack 5	R7-IB3-P11A to R5-U1-P9A	5 meters
R7 IB3 to Rack 6	R7-IB3-P11B to R6-U1-P10A	5 meters
R7 IB2 within Rack 7	R7-IB2-P8A to R7-U1-P3B	3 meters
	R7-IB2-P8B to R7-U1-P4B	
R7 IB2 to Rack 1	R7-IB2-P9A to R1-U1-P5B	10 meters
R7 IB2 to Rack 2	R7-IB2-P9B to R2-U1-P6B	10 meters
R7 IB2 to Rack 3	R7-IB2-P10A to R3-U1-P7B	10 meters
R7 IB2 to Rack 4	R7-IB2-P10B to R4-U1-P8B	10 meters



Table 2-25 (Cont.) Leaf Switch Connections for the Seventh Rack in a Seven-Rack System

Leaf Switch	Connection	Cable Length
R7 IB2 to Rack 5	R7-IB2-P11A to R5-U1-P9B	5 meters
R7 IB2 to Rack 6	R7-IB2-P11B to R6-U1-P10B	5 meters

2.8 Eight-Rack Cabling

Note the following when cabling five racks together:

- The spine switch is in U1 (also referred to as IB1) for all racks.
- In the following tables, the leaf switches are referred to as IB2 and IB3. Use the rack unit that is appropriate for your system.
- In Oracle Exadata Database Machine X2-2 racks and Oracle Exadata Storage Expansion Racks with Exadata Storage Server with Sun Fire X4270 M2 Servers, the leaf switches are in U20 and U24, also referred to as IB2 and IB3, respectively.
- In Oracle Exadata Database Machine X2-8 and later racks, the leaf switches are in U21 and U23, also referred to as IB2 and IB3, respectively.
- In Oracle Exadata Database Machine X4-2 and later racks, or Oracle Exadata
 Database Machine X3-2 Full Racks and Oracle Exadata Storage Expansion Rack
 X3-2 racks, the leaf switches are in U20 and U22, also referred to as IB2 and IB3,
 respectively.
- The cable lengths shown in the tables assume the racks are adjacent to each other. If the racks are not adjacent or use overhead cabling trays, then they may require longer cables lengths. Up to 100 meters is supported.
- Only optical cables are supported for lengths greater than 5 meters.

Table 2-26 shows the cable connections for the first spine switch (R1-U1) when cabling eight racks together:

Table 2-26 Leaf Switch Connections for the First Rack in a Eight-Rack System

Leaf Switch	Connection	Cable Length
R1 IB3 within Rack 1	R1-IB3-P8A to R1-U1-P3A	3 meters
R1 IB3 to Rack 2	R1-IB3-P8B to R2-U1-P4A	5 meters
R1 IB3 to Rack 3	R1-IB3-P9A to R3-U1-P5A	5 meters
R1 IB3 to Rack 4	R1-IB3-P9B to R4-U1-P6A	10 meters
R1 IB3 to Rack 5	R1-IB3-P10A to R5-U1-P7A	10 meters
R1 IB3 to Rack 6	R1-IB3-P10B to R6-U1-P8A	10 meters
R1 IB3 to Rack 7	R1-IB3-P11A to R7-U1-P9A	10 meters
R1 IB3 to Rack 8	R1-IB3-P11B to R8-U1-P10A	10 meters
R1 IB2 within Rack 1	R1-IB2-P8A to R1-U1-P3B	3 meters
R1 IB2 to Rack 2	R1-IB2-P8B to R2-U1-P4B	5 meters
R1 IB2 to Rack 3	R1-IB2-P9A to R3-U1-P5B	5 meters



Table 2-26 (Cont.) Leaf Switch Connections for the First Rack in a Eight-Rack System

Leaf Switch	Connection	Cable Length
R1 IB2 to Rack 4	R1-IB2-P9B to R4-U1-P6B	10 meters
R1 IB2 to Rack 5	R1-IB2-P10A to R5-U1-P7B	10 meters
R1 IB2 to Rack 6	R1-IB2-P10B to R6-U1-P8B	10 meters
R1 IB2 to Rack 7	R1-IB2-P11A to R7-U1-P9B	10 meters
R1 IB2 to Rack 8	R1-IB2-P11B to R8-U1-P10B	10 meters

Table 2-27 shows the cable connections for the second spine switch (R2-U1) when cabling eight full racks together:

Table 2-27 Leaf Switch Connections for the Second Rack in a Eight-Rack System

Leaf Switch	Connection	Cable Length
	Connection	Cable Length
R2 IB3 within Rack 2	R2-IB3-P8A to R2-U1-P3A	3 meters
R2 IB3 to Rack 1	R2-IB3-P11B to R1-U1-P10A	5 meters
R2 IB3 to Rack 3	R2-IB3-P8B to R3-U1-P4A	5 meters
R2 IB3 to Rack 4	R2-IB3-P9A to R4-U1-P5A	5 meters
R2 IB3 to Rack 5	R2-IB3-P9B to R5-U1-P6A	10 meters
R2 IB3 to Rack 6	R2-IB3-P10A to R6-U1-P7A	10 meters
R2 IB3 to Rack 7	R2-IB3-P10B to R7-U1-P8A	10 meters
R2 IB3 to Rack 8	R2-IB3-P11A to R8-U1-P9A	10 meters
R2 IB2 within Rack 2	R2-IB2-P8A to R2-U1-P3B	3 meters
R2 IB2 to Rack 1	R2-IB2-P11B to R1-U1-P10B	5 meters
R2 IB2 to Rack 3	R2-IB2-P8B to R3-U1-P4B	5 meters
R2 IB2 to Rack 4	R2-IB2-P9A to R4-U1-P5B	5 meters
R2 IB2 to Rack 5	R2-IB2-P9B to R5-U1-P6B	10 meters
R2 IB2 to Rack 6	R2-IB2-P10A to R6-U1-P7B	10 meters
R2 IB2 to Rack 7	R2-IB2-P10B to R7-U1-P8B	10 meters
R2 IB2 to Rack 8	R2-IB2-P11A to R8-U1-P9B	10 meters

Table 2-28 shows the cable connections for the third spine switch (R3-U1) when cabling eight full racks together:

Table 2-28 Leaf Switch Connections for the Third Rack in a Eight-Rack System

Leaf Switch	Connection	Cable Length
R3 IB3 within Rack 3	R3-IB3-P8A to R3-U1-P3A	3 meters
R3 IB3 to Rack 1	R3-IB3-P11A to R1-U1-P9A	5 meters
R3 IB3 to Rack 2	R3-IB3-P11B to R2-U1-P10A	5 meters



Table 2-28 (Cont.) Leaf Switch Connections for the Third Rack in a Eight-Rack System

Leaf Switch	Connection	Cable Length
R3 IB3 to Rack 4	R3-IB3-P8B to R4-U1-P4A	5 meters
R3 IB3 to Rack 5	R3-IB3-P9A to R5-U1-P5A	5 meters
R3 IB3 to Rack 6	R3-IB3-P9B to R6-U1-P6A	10 meters
R3 IB3 to Rack 7	R3-IB3-P10A to R7-U1-P7A	10 meters
R3 IB3 to Rack 8	R3-IB3-P10B to R8-U1-P8A	10 meters
R3 IB2 within Rack 3	R3-IB2-P8A to R3-U1-P3B	3 meters
R3 IB2 to Rack 1	R3-IB2-P11A to R1-U1-P9B	5 meters
R3 IB2 to Rack 2	R3-IB2-P11B to R2-U1-P10B	5 meters
R3 IB2 to Rack 4	R3-IB2-P8B to R4-U1-P4B	5 meters
R3 IB2 to Rack 5	R3-IB2-P9A to R5-U1-P5B	5 meters
R3 IB2 to Rack 6	R3-IB2-P9B to R6-U1-P6B	10 meters
R3 IB2 to Rack 7	R3-IB2-P10A to R7-U1-P7B	10 meters
R3 IB2 to Rack 8	R3-IB2-P10B to R8-U1-P8B	10 meters

Table 2-29 shows the cable connections for the fourth spine switch (R4-U1) when cabling eight full racks together:

Table 2-29 Leaf Switch Connections for the Fourth Rack in a Eight-Rack System

Leaf Switch	Connection	Cable Length
R4 IB3 within Rack 4	R4-IB3-P8A to R4-U1-P3A	3 meters
R4 IB3 to Rack 1	R4-IB3-P10B to R1-U1-P8A	10 meters
R4 IB3 to Rack 2	R4-IB3-P11A to R2-U1-P9A	5 meters
R4 IB3 to Rack 3	R4-IB3-P11B to R3-U1-P10A	5 meters
R4 IB3 to Rack 5	R4-IB3-P8B to R5-U1-P4A	5 meters
R4 IB3 to Rack 6	R4-IB3-P9A to R6-U1-P5A	5 meters
R4 IB3 to Rack 7	R4-IB3-P9B to R7-U1-P6A	10 meters
R4 IB3 to Rack 8	R4-IB3-P10A to R8-U1-P7A	10 meters
R4 IB2 within Rack 4	R4-IB2-P8A to R4-U1-P3B	3 meters
R4 IB2 to Rack 1	R4-IB2-P10B to R1-U1-P8B	10 meters
R4 IB2 to Rack 2	R4-IB2-P11A to R2-U1-P9B	5 meters
R4 IB2 to Rack 3	R4-IB2-P11B to R3-U1-P10B	5 meters
R4 IB2 to Rack 5	R4-IB2-P8B to R5-U1-P4B	5 meters
R4 IB2 to Rack 6	R4-IB2-P9A to R6-U1-P5B	5 meters
R4 IB2 to Rack 7	R4-IB2-P9B to R7-U1-P6B	10 meters
R4 IB2 to Rack 8	R4-IB2-P10A to R8-U1-P7B	10 meters



Table 2-30 shows the cable connections for the fifth spine switch (R5-U1) when cabling eight full racks together:

Table 2-30 Leaf Switch Connections for the Fifth Rack in a Eight-Rack System

Connection	Cable Length
R5-IB3-P8A to R5-U1-P3A	3 meters
R5-IB3-P10A to R1-U1-P7A	10 meters
R5-IB3-P10B to R2-U1-P8A	10 meters
R5-IB3-P11A to R3-U1-P9A	5 meters
R5-IB3-P11B to R4-U1-P10A	5 meters
R5-IB3-P8B to R6-U1-P4A	5 meters
R5-IB3-P9A to R7-U1-P5A	5 meters
R5-IB3-P9B to R8-U1-P6A	10 meters
R5-IB2-P8A to R5-U1-P3B	3 meters
R5-IB2-P10A to R1-U1-P7B	10 meters
R5-IB2-P10B to R2-U1-P8B	10 meters
R5-IB2-P11A to R3-U1-P9B	5 meters
R5-IB2-P11B to R4-U1-P10B	5 meters
R5-IB2-P8B to R6-U1-P4B	5 meters
R5-IB2-P9A to R7-U1-P5B	5 meters
R5-IB2-P9B to R8-U1-P6B	10 meters
	R5-IB3-P8A to R5-U1-P3A R5-IB3-P10A to R1-U1-P7A R5-IB3-P10B to R2-U1-P8A R5-IB3-P11A to R3-U1-P9A R5-IB3-P11B to R4-U1-P10A R5-IB3-P8B to R6-U1-P4A R5-IB3-P9A to R7-U1-P5A R5-IB3-P9B to R8-U1-P6A R5-IB2-P8A to R5-U1-P3B R5-IB2-P10A to R1-U1-P7B R5-IB2-P10B to R2-U1-P8B R5-IB2-P11B to R4-U1-P10B R5-IB2-P11B to R4-U1-P10B R5-IB2-P8B to R6-U1-P4B R5-IB2-P9A to R7-U1-P5B

Table 2-31 shows the cable connections for the sixth spine switch (R6-U1) when cabling eight full racks together:

Table 2-31 Leaf Switch Connections for the Sixth Rack in a Eight-Rack System

Leaf Switch	Connection	Cable Length
R6 IB3 within Rack 6	R6-IB3-P8A to R6-U1-P3A	3 meters
R6 IB3 to Rack 1	R6-IB3-P9B to R1-U1-P6A	10 meters
R6 IB3 to Rack 2	R6-IB3-P10A to R2-U1-P7A	10 meters
R6 IB3 to Rack 3	R6-IB3-P10B to R3-U1-P8A	10 meters
R6 IB3 to Rack 4	R6-IB3-P11A to R4-U1-P9A	5 meters
R6 IB3 to Rack 5	R6-IB3-P11B to R5-U1-P10A	5 meters
R6 IB3 to Rack 7	R6-IB3-P8B to R7-U1-P4A	5 meters
R6 IB3 to Rack 8	R6-IB3-P9A to R8-U1-P5A	5 meters
R6 IB2 within Rack 6	R6-IB2-P8A to R6-U1-P3B	3 meters
R6 IB2 to Rack 1	R6-IB2-P9B to R1-U1-P6B	10 meters
R6 IB2 to Rack 2	R6-IB3-P10A to R2-U1-P7B	10 meters
R6 IB2 to Rack 3	R6-IB2-P10B to R3-U1-P8B	10 meters
R6 IB2 to Rack 4	R6-IB2-P11A to R4-U1-P9B	5 meters
R6 IB2 to Rack 5	R6-IB2-P11B to R5-U1-P10B	5 meters



Table 2-31 (Cont.) Leaf Switch Connections for the Sixth Rack in a Eight-Rack System

Leaf Switch	Connection	Cable Length
R6 IB2 to Rack 7	R6-IB2-P8B to R7-U1-P4B	5 meters
R6 IB2 to Rack 8	R6-IB2-P9A to R8-U1-P5B	5 meters

Table 2-32 shows the cable connections for the seventh spine switch (R7-U1) when cabling eight full racks together:

Table 2-32 Leaf Switch Connections for the Seventh Rack in a Eight-Rack System

Leaf Switch	Connection	Cable Length
R7 IB3 within Rack 7	R7-IB3-P8A to R7-U1-P3A	3 meters
R7 IB3 to Rack 1	R7-IB3-P9A to R1-U1-P5A	10 meters
R7 IB3 to Rack 2	R7-IB3-P9B to R2-U1-P6A	10 meters
R7 IB3 to Rack 3	R7-IB3-P10A to R3-U1-P7A	10 meters
R7 IB3 to Rack 4	R7-IB3-P10B to R4-U1-P8A	10 meters
R7 IB3 to Rack 5	R7-IB3-P11A to R5-U1-P9A	5 meters
R7 IB3 to Rack 6	R7-IB3-P11B to R6-U1-P10A	5 meters
R7 IB3 to Rack 8	R7-IB3-P8B to R8-U1-P4A	5 meters
R7 IB2 within Rack 7	R7-IB2-P8A to R7-U1-P3B	3 meters
R7 IB2 to Rack 1	R7-IB2-P9A to R1-U1-P5B	10 meters
R7 IB2 to Rack 2	R7-IB2-P9B to R2-U1-P6B	10 meters
R7 IB2 to Rack 3	R7-IB2-P10A to R3-U1-P7B	10 meters
R7 IB2 to Rack 4	R7-IB2-P10B to R4-U1-P8B	10 meters
R7 IB2 to Rack 5	R7-IB2-P11A to R5-U1-P9B	5 meters
R7 IB2 to Rack 6	R7-IB2-P11B to R6-U1-P10B	5 meters
R7 IB2 to Rack 8	R7-IB2-P8B to R8-U1-P4B	5 meters

Table 2-33 shows the cable connections for the eighth spine switch (R8-U1) when cabling eight full racks together:

Table 2-33 Leaf Switch Connections for the Eighth Rack in a Eight-Rack System

Leaf Switch	Connection	Cable Length
R8 IB3 within Rack 8	R8-IB3-P8A to R8-U1-P3A	3 meters
R8 IB3 to Rack 1	R8-IB3-P8B to R1-U1-P4A	10 meters
R8 IB3 to Rack 2	R8-IB3-P9A to R2-U1-P5A	10 meters
R8 IB3 to Rack 3	R8-IB3-P9B to R3-U1-P6A	10 meters
R8 IB3 to Rack 4	R8-IB3-P10A to R4-U1-P7A	10 meters



Table 2-33 (Cont.) Leaf Switch Connections for the Eighth Rack in a Eight-Rack System

Leaf Switch	Connection	Cable Length
R8 IB3 to Rack 5	R8-IB3-P10B to R5-U1-P8A	10 meters
R8 IB3 to Rack 6	R8-IB3-P11A to R6-U1-P9A	5 meters
R8 IB3 to Rack 7	R8-IB3-P11B to R7-U1-P10A	5 meters
R8 IB2 within Rack 8	R8-IB2-P8A to R8-U1-P3B	3 meters
R8 IB2 to Rack 1	R8-IB2-P8B to R1-U1-P4B	10 meters
R8 IB2 to Rack 2	R8-IB2-P9A to R2-U1-P5B	10 meters
R8 IB2 to Rack 3	R8-IB2-P9B to R3-U1-P6B	10 meters
R8 IB2 to Rack 4	R8-IB2-P10A to R4-U1-P7B	10 meters
R8 IB2 to Rack 5	R8-IB2-P10B to R5-U1-P8B	10 meters
R8 IB2 to Rack 6	R8-IB2-P11A to R6-U1-P9B	5 meters
R8 IB2 to Rack 7	R8-IB2-P11B to R7-U1-P10B	5 meters



Index

A	expansion kit
adding Sun Datacenter InfiniBand Switch 36 switch,	when to order, 1-9 extending cellip.ora file, 1-54
1-29	eighth rack in X3-2 or earlier racks, 1-21
addNode.sh script, 1-57	eighth rack in X4-2 or later racks, 1-11
В	F
backslash character, 1-62	files
bend radius requirements, 1-5	add-cluster-nodes.rsp, 1-57 add-db-nodes.rsp, 1-63
С	*
cable management arm, 1-5	<u>H</u>
cabling	health check utility, 1-65
bend radius, 1-5	•
bundling InfiniBand, 1-5	1
estimating cable length, 1-4	<u></u>
InfiniBand precautions, 1-4	installation
multi-rack, 1-42, 2-1	safety guidelines, 1-3
new database servers, 1-34	installing Sun Datacenter InfiniBand Switch 36
new Exadata Storage servers, 1-39	switch, <i>1-29</i>
quarter racks, 2-7	IP addresses
capacity-on-demand, 1-47	cellip.ora, 1-54
cellip.ora file, 1-54	•
checking	1
cell configuration, 1-6	L
cell disk configuration, 1-6	leaf switches, 2-1
firmware version, 1-6	icai switches, 2 1
flash logs, 1-6	
HugePages memory configuration, 1-6	O
image history, 1-6	Oragla Evadata Databasa Mashina HaalthChaak
InfiniBand switch information, 1-6	Oracle Exadata Database Machine HealthCheck
IORM plans, 1-6	utility, 1-65
IP addresses, 1-6	
CMA guidelines, 1-5	S
configuring	
cell alerts, 1-62	spare cables, 2-4
CSI, 1-3, 1-11	spine switches, 2-1
Customer Support Identifier, 1-3, 1-11	
	U
E	-
	utilities
Exachk utility, 1-6	Exachk, 1-6

utilities (continued) health check, 1-65 utilities (continued) oifcfg, 1-46

