

Direct Connect Oracle Exadata and Oracle Database Appliance with Oracle Private Cloud Appliance



Oracle Private Cloud Appliance X9-2

PURPOSE STATEMENT

This document provides an overview of Direct Connect functionality for Exadata/ODA included in Oracle Private Cloud Appliance X9-2, with Software Release 3.0.1. It is intended solely to help you assess the business benefits and get an understanding of the configuration procedure involving Exadata-X9/X8/X7/ODA and PCA-X9-2.

DISCLAIMER

This document in any form, software, or printed matter, contains proprietary information that is the exclusive property of Oracle. Your access to and use of this confidential material is subject to the terms and conditions of your Oracle software license and service agreement, which has been executed and with which you agree to comply. This document and information contained herein may not be disclosed, copied, reproduced, or distributed to anyone outside Oracle without prior written consent of Oracle. This document is not part of your license agreement, nor can it be incorporated into any contractual agreement with Oracle or its subsidiaries or affiliates.

This document is for informational purposes only and is intended solely to assist you in planning for the implementation and upgrade of the product features described. It is not a commitment to deliver any material, code, or functionality, and should not be relied upon in making purchasing decisions. The development, release, and timing of any features or functionality described in this document remains at the sole discretion of Oracle.

TABLE OF CONTENTS

Purpose Statement	2
Disclaimer	2
Overview	4
Logical/Physical Network connectivity	6
Configuring the Oracle Private Compute Appliance	9
Create IAAS – from Customer Enclave UI	12
Configuring the Oracle Exadata Database Machine	18
Configuring the Oracle Database APPLIANCE (ODA)	19
Frequently Asked Questions	22

OVERVIEW

The Oracle Private Cloud Appliance provides a simple out-of-the-box converged infrastructure solution and is a key enabler for Oracle-As-A-Service platform. The Oracle Private Cloud Appliance (PCA X9) is the only OCI compatible Engineered System, providing the fastest and most efficient infrastructure for modern software and business applications.

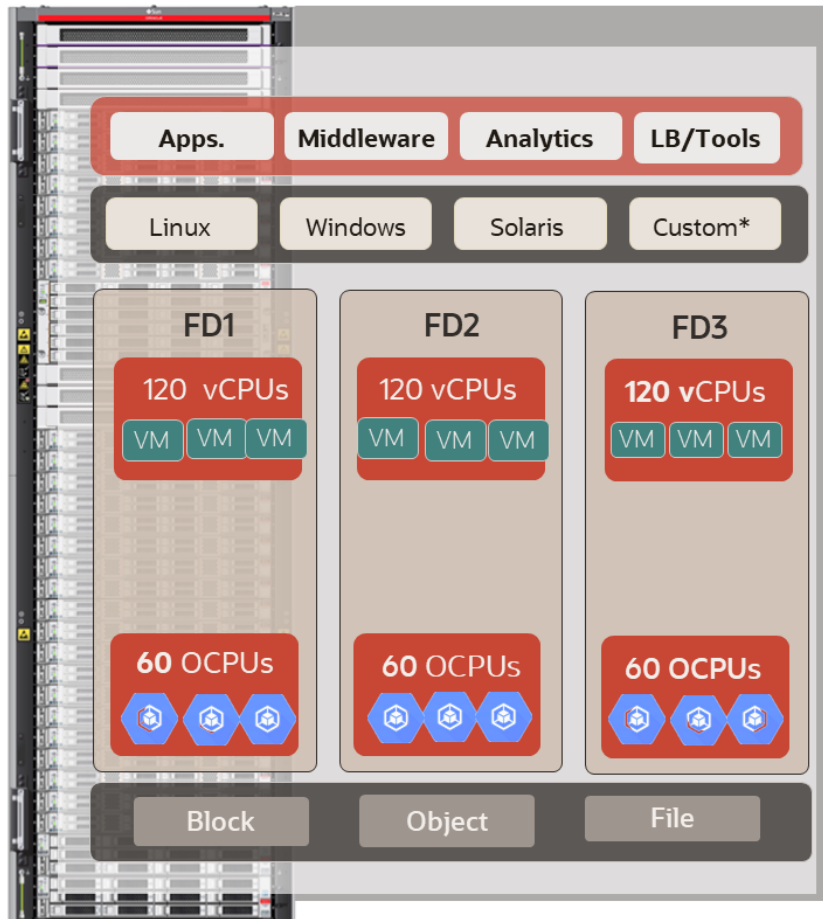


Figure 1: PCA-X9 3 Fault-Domain architecture

The PCA X9 has the same Infrastructure constructs (APIs, SDK etc.) as the OCI. This enables customers to adopt a -- “Develop once and Deploy anywhere, on-premises or on OCI” -- approach to rapidly design and develop high performance applications and middleware.

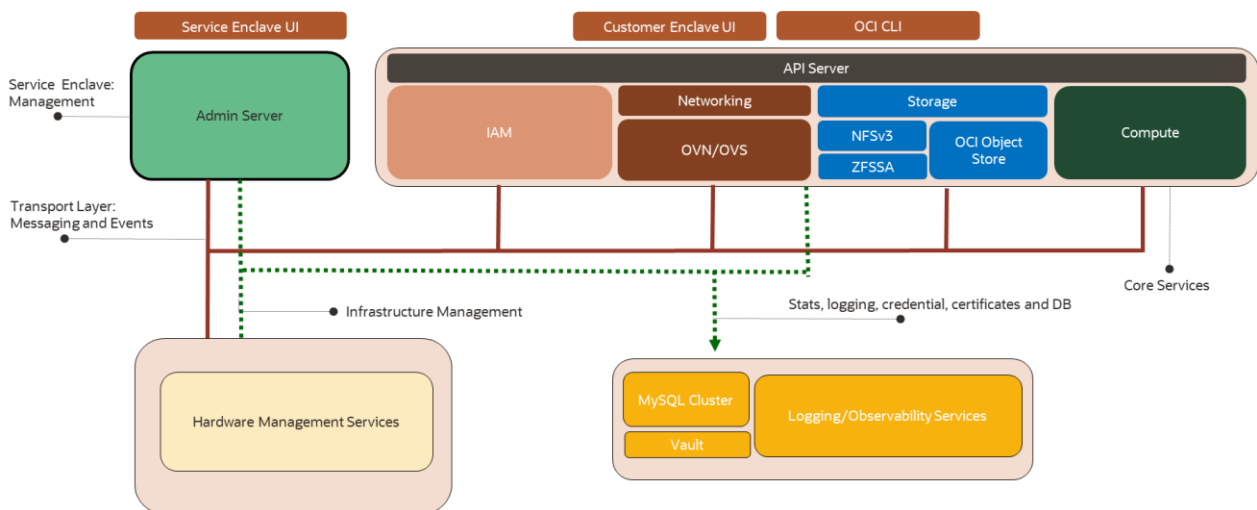


Figure 2: PCA-X9 High-Level Software architecture

Benefits:

- Intelligent and agile infrastructure
- Flexible support for any mix of Oracle and non-Oracle workloads
- Automation speeds install
- Built for rapid and cost-effective private cloud deployment
- Zero Downtime Rolling Upgrades
 - Continuous Integration/Continuous Development (CI/CD) methodology
- Software Define Networking – over 100Gb/s Ethernet backplane
- Migrate legacy Exalogic/BDA/Exalytics workloads to PCA
- No infrastructure license cost

The 100Gb Ethernet network inside Oracle Private Cloud Appliance, is implemented as spine and leaf topology. Each Compute Node has a connection to both leaf switches. Each leaf, is in turn, connected to both the spine switches. Each spine switch has connectivity to the Storage Nodes, as well as a group of ports for external connectivity.

Ports 1-4 on both the spine switches are reserved for the Layer3 based uplink to customer's datacenter. Each of the four customer uplink ports, numbered 1 through 4, may be configured in several ways. Each port may be used as:

- A single 100Gb Ethernet port
- A single 40Gb Ethernet port
- Four 25Gb Ethernet sub ports
- Four 10Gb Ethernet sub ports

Ports 7-10 on both the spine switches are reserved for connecting to database-servers from Exadata/ODA for the Client Network. Each of the four ports is **configured as 4x 25Gb Ethernet breakout** ports to support a total of 16 database-servers which can come from any number of Exadata/ODA platforms.

The Oracle Exadata X8M product line introduces a new and improved networking infrastructure built on high-speed, low latency 100Gb Ethernet optimized for RDMA over Converged Ethernet (RoCE). Connectivity to Applications and Middleware tier is provided via physical Ethernet interfaces on the database-servers. Each Oracle Exadata Database Machine X8M-2 database server consists of the following network components and interfaces:

- 1 embedded 1 GbE (eth0 RJ-45)
- 2 embedded 10GBaseT ports (RJ-45) or **2 embedded 10GbE/25GbE ports (SFP+/SFP28)**
- 1 dual-port QSFP28 100Gb/s RDMA Network Fabric Card
- 1 embedded Ethernet port for Integrated Lights Out Manager (ILOM) remote management
- **1 dual-port 25GbE PCIe 3.0 network card** with Broadcom BCM57414 10Gb/25Gb Ethernet Controller

Note: The SFP28 modules for the 25GbE PCIe 3.0 network cards are purchased separately.

Similarly, Oracle Database Appliance X9-2L/ X8-2S/X8-2M provides support for the following NIC which can be used for connecting to the PCA-X9 Spine switches.

- **Oracle Dual Port 25Gb Ethernet Adapter**

This document focuses on network attachment and configuration steps needed to create secure, private, and high-bandwidth connectivity between the VMs/Containers, comprising the apps/middleware tier running on Oracle Private Cloud Appliance X9, and the database tier running on the Oracle Exadata /ODA.

Directly connecting the database-servers in the Exadata-X8M/X9M, ODA rack to the spine switches in the PCA-X9 rack is ideal for migrating Oracle Exalogic workloads to PCA. This configuration will eliminate the need to acquire and manage dedicated top of rack switches. This option will also provide superior ROI, tighter integration, lowest latency for database transactions and provisions the maximum available bandwidth.

Important Considerations:

- PCA-X9 Spine switch configuration for Direct Connect requires 25Gbps links from the database-servers
- PCA-X9 Spine switch configuration for Direct Connect creates an exclusive Layer3 network for the attached database-servers
- Spine switches inside PCA rack have vPC configuration to make the 2 spines appear as a single-virtual-switch
- LACP **must** be enabled for the Client-Network interfaces to enable maximum bandwidth and failover
- The connection between PCA-spine-switch and database-servers is setup as a private network
- Access from the customer’s network to the database-servers requires additional routing configuration
 - DefaultGateway/DNS/NTP can be reachable via the Spine switches
- The subnet-mask for Client/Other-network should be chosen to accommodate the ip addresses for the VMs in the Oracle Private Compute Appliance.
 - /22 mask provides 1024 addresses, /23 mask provides 512 addresses

This paper addresses the following topics:

- Logical/Physical network connectivity
- Network configuration on Oracle Private Compute Appliance X9-2
- Network configuration on Oracle Exadata X8M database-servers
- Network configuration on Oracle Database Appliance nodes

LOGICAL/PHYSICAL NETWORK CONNECTIVITY

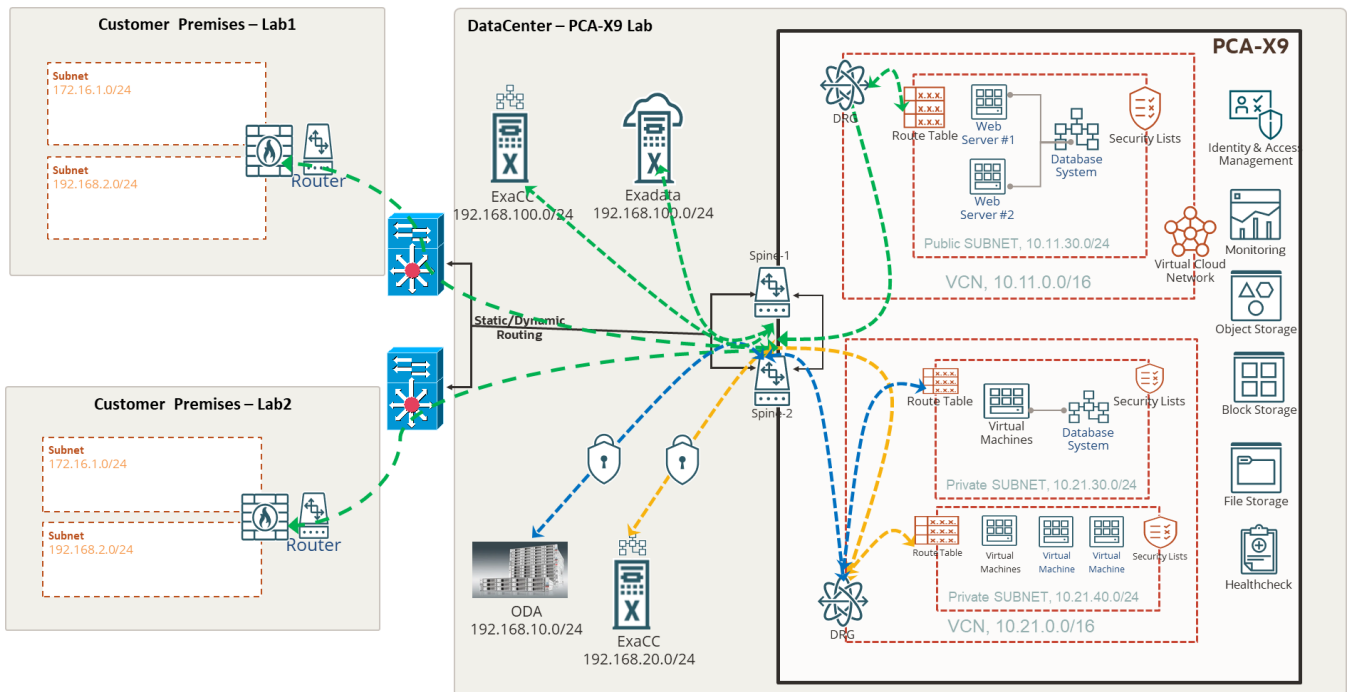


Figure 3: Logical Reference Architecture for connecting Exadata/ODA

The Client-Network is used by the applications to access the database. The database-servers provide 2-pairs of physical interfaces configured by OEDA as bondeth0 and bondeth1 (Exadata) or as btbond1 (ODA), which can be used to create a direct connection to the spine switches in the Oracle Private Cloud Appliance.

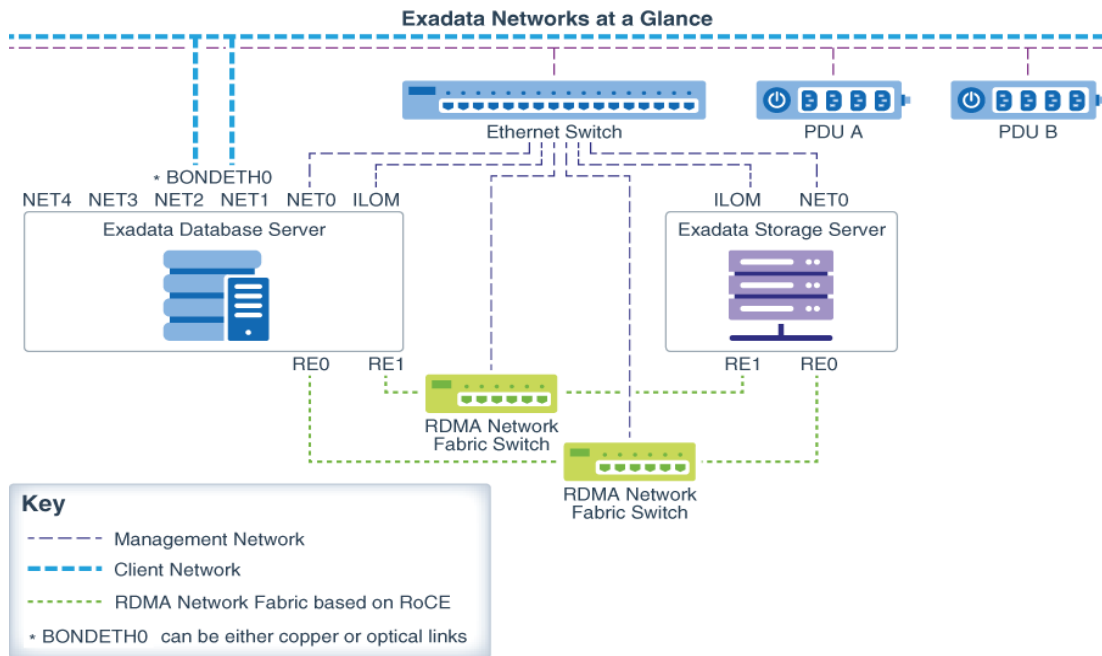


Figure 4: Logical Network diagram of Exadata X8M networks




Figure 5: Available physical interfaces on the database-servers inside Exadata X8M




Figure 6: Available physical interfaces on Oracle Database Appliance X8-2 HA nodes




Figure 7: Physical cable connections between the PCA spine switches and database-servers

PCA-X9 Exadata Network

Sample cable connections for Eighth, Quarter and Half Exadata/ExaCC

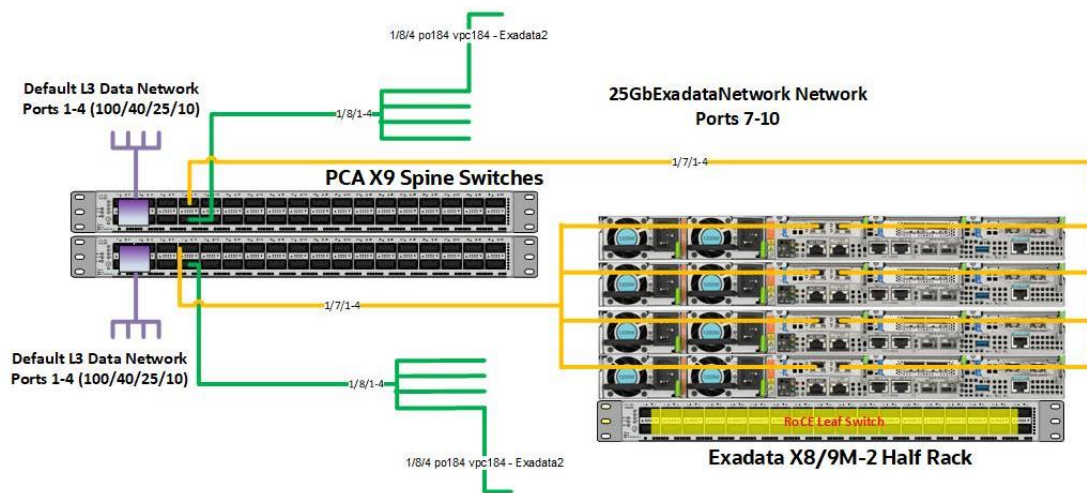


Figure 8: Physical cable connections between the PCA spine switches and database-servers

PCA-X9 Exadata Network

Sample cable connections for ODA X9-2-HA

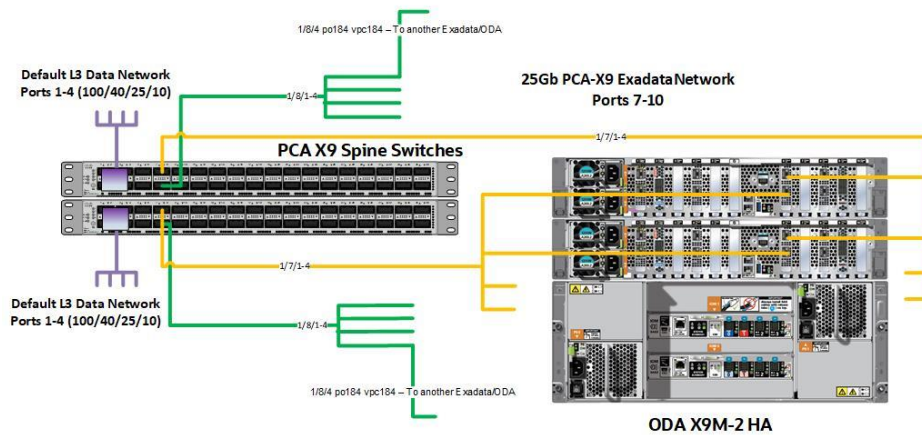


Figure 9: Physical cable connections between the PCA spine switches and ODA nodes (X9-2-HA)

CONFIGURING THE ORACLE PRIVATE COMPUTE APPLIANCE

The Client-Network connection will always be using 25Gb Ethernet cards in the database-servers. Hence, the PCA spine-switch ports 7-10 are pre-configured as 4x25G breakout mode by default, this configuration cannot be changed.

Connecting the Client-network directly to the PCA spine-switches has the benefits of isolating the Client-network from other workloads and providing a greater level of control. Dedicated 50Gbps of bandwidth capacity to each database-server for large workloads and the ability to use jumbo frames, enables higher application efficiency.

Figure 8 depicts the network cabling diagram for a half-rack Exadata-X8M with 4 database-servers. Smaller configurations, like Exadata eighth, and quarter rack or the ODA (Figure 9) will only require connecting port 7 from both switches.

The configuration on PCA is done in multiple steps.

1. Configuration of Exadata Network using PCA-Admin CLI

```
PCA-ADMIN> showallcustomcmds
..
...
exaDataCreateNetwork: ExadataNetwork
exaDataDeleteNetwork: ExadataNetwork
exaDataDisableAccess: Vcn
exaDataEnableAccess: Vcn
exaDataGetNetwork: ExadataNetwork
exaDataListNetwork: ExadataNetwork
```

- a. Create Exadata/ODA Network
 - i. Identify physical ports on Spine switch where database-servers will connect
 - ii. Identify ClientNetwork subnet and reserve 3 IP addresses (one for each Spine, and a VIP)

```
PCA-ADMIN> exaDataCreateNetwork ?
*cidr
vlan
*spine1Ip
*spine2Ip
*spineVip
*ports
advertiseNetwork
```

```
exaDataCreateNetwork cidr="10.80.65.0/24" vlan=2001 spine1Ip="10.80.65.101" spine2Ip="10.80.65.102"
spineVip="10.80.65.1" ports="7/1,7/2" advertiseNetwork=true
```

Note: `advertiseNetwork=true` is needed if you want the database-servers to be accessible from customer's datacenter network, in the case of Dynamic-routing-based config.

- b. Verify successful creation and record the OCID

```
PCA-ADMIN> exaDataListNetwork << No Additional Argument >>
```

exaDataListNetwork

Data:

id	vlan	cidr	spine1Ip	spine2Ip	spineVip	ports
--	----	----	-----	-----	-----	-----
	2001	10.80.65.0/24	10.80.65.101	10.80.65.102	10.80.65.1	7/1,7/2

ocid1.exadata.1913XD300H.broom2.ffv2apzh0dhfljkshxavllqc4I3jnlpvq4wedft7s417tqg6g0ht51nz7szu

2. Create IaaS from the Customer Enclave (UI)

- a. Physical cable connections using QSFP28 to SFP28 breakout cables – refer to Table-2 for the cable and transceiver part-numbers
- b. Create IaaS – [see example in the next section](#)
 - i. Create a DRG first (VMs can only access the Exadata via the DRG)
 - ii. Create VCN(s), Subnets, Route Tables and IGW (if needed)
 - iii. Create DRG-attachment to enable VCN to use DRG
 - iv. Recommendations/Considerations:
 1. If the VM also needs to be accessed from DC, a 2nd VNIC will be needed
 2. Use IGW (Public subnet) for DC access using the Primary VNIC
 3. Use DRG (Private subnet) to access Exadata network
 4. Use separate Route Tables, one for IGW with “0.0.0.0/0” and one for DRG with a “specific” route rule for Exadata subnet
 5. Update Security List(s) as needed to enable ingress traffic
 6. Record the OCIDs of all subnets which would need to access Exadata network

3. Enable access for the required subnets to Exadata Network

- a. Enable communication between Exadata network created in step-1 and VM subnet(s) created in step-2
- b. Repeat the `exaDataEnableAccess` command for each subnet

```
PCA-ADMIN> exaDataEnableAccess ?
```

```
*exadataNetworkId
```

```
*subnetId
```

exaDataenableAccess

```
exadataNetworkId=ocid1.exadata.1913XD300H.broom2.ffv2apzh0dhfljkshxavllqc4l3jnlpvq4wedft7s417tqg6g0ht51nz7szu  
subnetId=ocid1.subnet.1913XD300H.broom2.3ypx4r7hmilpms0rb9dqkx4x5ry3rvqx6zlb0hi0ww6mv2mmy0i8brydnr
```

c. Validate and record final settings

exaDataGetNetwork

```
exadataNetworkId=ocid1.exadata.1913XD300H.broom2.ffv2apzh0dhfljkshxavllqc4l3jnlpvq4wedft7s417tqg6g0ht51nz7szu
```

Data:

```
CIDR = 10.80.65.0/24
```

```
Vlan = 2001
```

```
Spine1Ip = 10.80.65.101
```

```
Spine2Ip = 10.80.65.102
```

```
SpineVip = 10.80.65.1
```

```
Ports = 7/1,7/2
```

```
advertiseNetwork = true
```

```
Access List 1 - Vcn Id =
```

```
ocid1.vcn.1913XD300H.broom2.3qg8a4985mrfwmshtnx3whqxpijy5xpyzzk6s4nomqf9fqq6p1xsZ5l99uz
```

```
Access List 1 - Subnet Ids 1 =
```

```
ocid1.subnet.1913XD300H.broom2.3ypx4r7hmilpms0rb9dqkx4x5ry3rvqx6zlb0hi0ww6mv2mmy0i8brydnr
```

```
Lifecycle State = AVAILABLE
```

CREATE IAAS – FROM CUSTOMER ENCLAVE UI

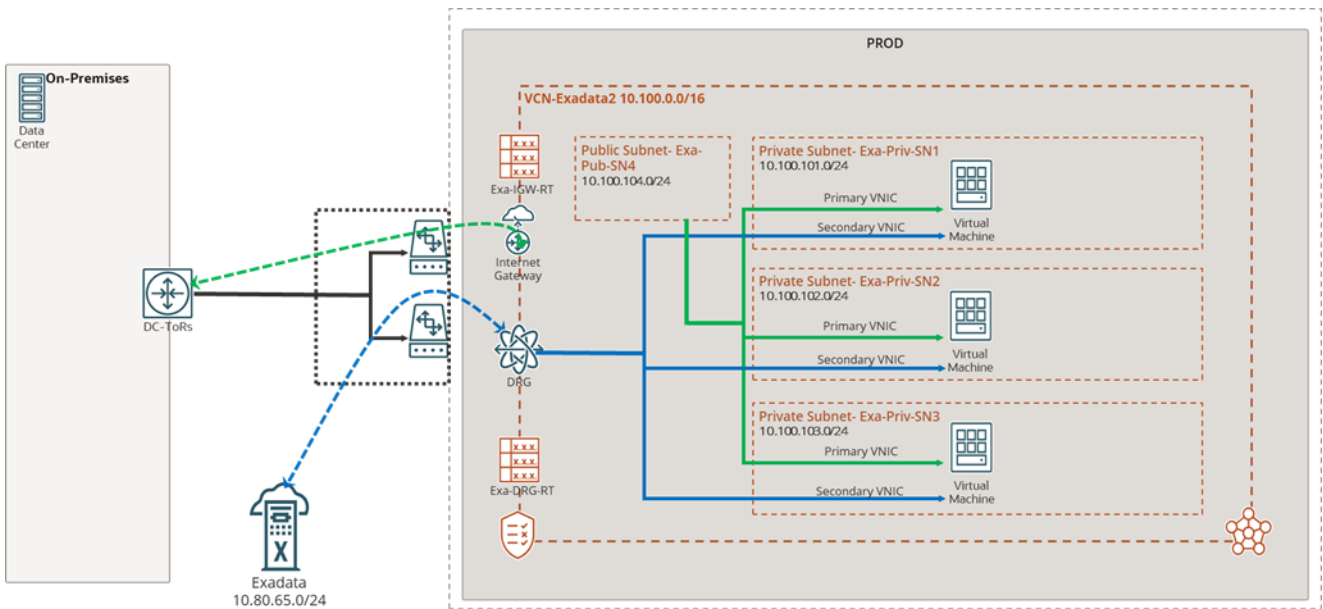
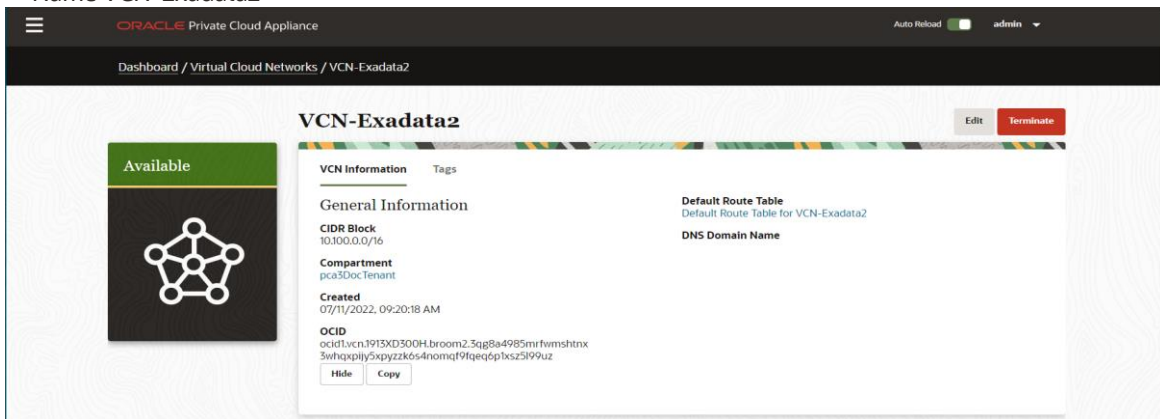


Figure 10: IaaS example with VMs accessing database-servers via DRG

- Each VM has a VNIC with a Public IP address
- Each VM has a secondary VNIC in a Private subnet to access the database-servers

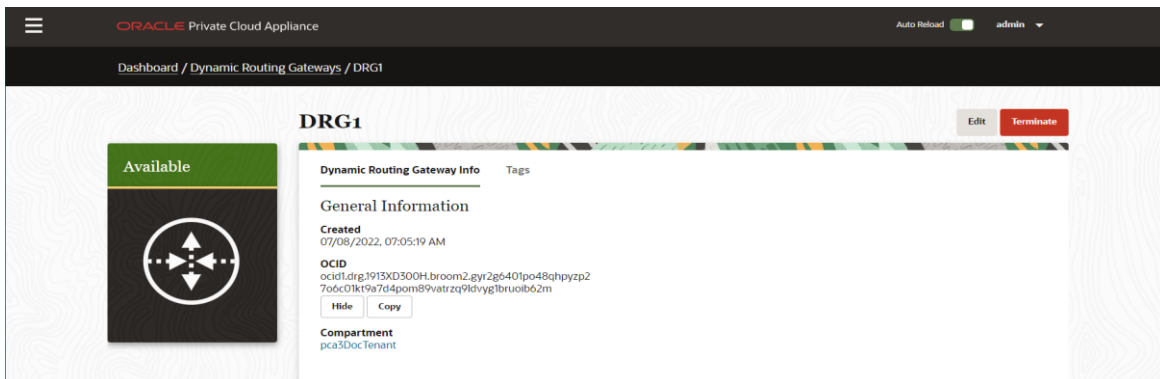
1. Create VCN:

- CIDR 10.100.0.0/16
- Name VCN-Exadata2



2. Create DRG:

- Name DRG1



3. Attach DRG to VCN

VCN-Exadata2 Edit Terminate

VCN Information Tags

General Information

CIDR Block
10.100.0.0/16

Compartment
pca3DocTenant

Created
07/11/2022, 09:20:18 AM

OCID
ocid1.vcn.1913XD300H.broom2.3qg8a4985mrfwmshtnx
3whqxpjiy5xpyzzk6s4nomqf9f9eq6p1xs25i99uz

Hide Copy

Default Route Table
Default Route Table for VCN-Exadata2

DNS Domain Name

Dynamic Routing Gateways Attach Dynamic Routing Gateway

Name	State	Route Table	Created	Actions
DRG1	Attached	-	07/25/2022, 04:41:01 PM	⋮

4. Create IGW:

- Name Ext-IGW-LB

Name	State	Enabled	Created	Actions
Ext-IGW-LB	Available	Yes	11/13/2022, 02:41:00 PM	⋮

5. Create Subnets:

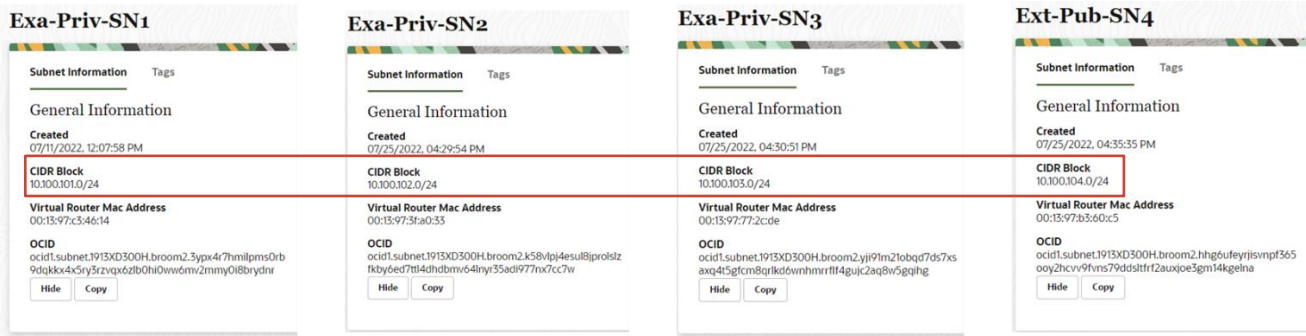
- Exa-Priv-SN1 – 10.100.101.0/24, Exa-Priv-SN2 – 10.100.102.0/24, Exa-Priv-SN3 – 10.100.103.0/24
- Exa-Pub-SN4 – 10.100.104.0/24

Resources

- Subnets (4)
- Route Tables (3)
- Internet Gateways (1)
- Local Peering Gateway (0)
- DHCP Options (1)
- Security Lists (1)
- NAT Gateways (0)
- Network Security Groups (0)

Subnets in the pca3DocTenant compartment Create Subnet

Subnet	State	Created	Actions
Exa-Priv-SN1	Available	07/11/2022, 12:07:58 PM	⋮
Exa-Priv-SN2	Available	07/25/2022, 04:29:54 PM	⋮
Exa-Priv-SN3	Available	07/25/2022, 04:30:51 PM	⋮
Ext-Pub-SN4	Available	07/25/2022, 04:35:35 PM	⋮



6. Create and update Route tables and Security Lists:

Exa-DRG-RT

Route Table Information Tags

General Information

Created
07/11/2022, 12:08:48 PM

Compartment
pca3DocTenant

OCID
ocid1.routetable.1913XD300H.broom2.oe25ovtmu1gmq6302pegmd9ro2c1el0esfwxhs4qg14ggy23tuyvu3wm9sj0

Hide Copy

Route Rules Add Route Rules

Destination	Target Type	Destination Type	Target	Description	Actions
10.80.65.0/24	Dynamic Routing Gateway	CIDR Block	DRG1		⋮

Ext-IGW-RT

Route Table Information Tags

General Information

Created
07/11/2022, 12:09:21 PM

Compartment
pca3DocTenant

OCID
...z49xhgyw72dcj81kykh

Show Full OCID Copy

Route Rules Add Route Rules

Destination	Target Type	Destination Type	Target	Description	Actions
0.0.0.0/0	Internet Gateway	CIDR Block	Ext-IGW-LB		⋮

7. Launch VMs with Primary VNIC in Public subnet

Launch Instance

Shape

Select shape

VM.PCAStandard1.2

Memory (GBs): 32 OCPUs: 2 Disk count/total size (GBs): 0/0

Boot Volume

Specify a custom boot volume size

Subnet

VCN pca3DocTenant (change) Subnet pca3DocTenant (change)

VCN-Exadata2 Ext-Pub-SN4

Public IP Address

Assign Public IP

8. Create Secondary VNIC in Private subnet for each of the VMs

Attached VNICs **Create VNIC Attachment**

Name ↑↓	State ↑↓	VLAN Tag ↑↓	MAC Address ↑↓	Created ↑↓	Actions
LB-Demo-Inst1	■ Attached	0	00:13:97:db:55:c1	06/28/2022, 01:06:30 AM	⋮

Create VNIC Attachment

Subnet

VCN pca3DocTenant (change) Subnet pca3DocTenant (change)

VCN-Exadata2 Exa-Priv-SN1

Skip Source/Destination Check Assign a Public IP Address

Primary IP Information

Private IP Address (Optional)

Must be within 10.0.0.0 to 10.0.255.255. Must not already be in use.

Hostname (Optional)

No spaces. Only letters, numbers, and hyphens. 63 characters max.

Network Security Group

Enable Network Security Groups

Create Attachment Cancel

9. Verify VNICs

- Each VM should have Primary VNIC in Ext-Pub-SN4 and Secondary VNIC in Exa-Priv-SNx

LB-Demo-Inst1 Edit

VNIC Information Tags

General Information

Created
07/26/2022, 09:15:07 AM

OCID
[Redacted]
Hide Copy

Compartment
pca3DocTenant

Subnet
Ext-Pub-SN4

Skip Source/Destination Check
No

MAC Address
00:13:97:0c:d1:fe

VLAN Tag

Primary IP Information

Private IP Address
[Redacted]

Private IP OCID
...21fmxqwi8qunazsh8q
Show Full OCID Copy

Private IP Assigned
07/26/2022, 09:15:07 AM

Fully Qualified Domain Name
lb-demo-inst1

Public IP Address
[Redacted]

Public IP OCID
...5nz43scudnxp6vne1w
Show Full OCID Copy

IP Addresses Assign Secondary Private IP Address

Hostname	Private IP Address	Public IP Address	Created	Actions
lb-demo-inst1	[Redacted]		07/26/2022, 09:15:07 AM	[More]

LB-Demo-Inst1 Edit

VNIC Information Tags

General Information

Created
07/26/2022, 09:26:57 AM

OCID
ocid1.vnic.1913XD300H.broom2.ji7uf4lspajw53fzfspr6yr
pa45wmzocrv4tmbhsf0nj9r1rn64o952unkv7
Hide Copy

Compartment
pca3DocTenant

Subnet
Exa-Priv-SN1

Skip Source/Destination Check
No

MAC Address
00:13:97:d1:a7:8c

VLAN Tag

Primary IP Information

Private IP Address
[Redacted]

Private IP OCID
...1ojcyjr8c116i3l58e7q
Show Full OCID Copy

Private IP Assigned
07/26/2022, 09:26:57 AM

Fully Qualified Domain Name
lb-demo-inst1

Public IP Address
Not Set

Public IP OCID
Show Full OCID Copy

IP Addresses Assign Secondary Private IP Address

Hostname	Private IP Address	Public IP Address	Created	Actions
lb-demo-inst1	[Redacted]		07/26/2022, 09:26:57 AM	[More]

10. Verify VNIC configuration inside each VM

Note: The Secondary VNIC is not enabled automatically, it will need to be configured manually or using the script provided by OCI

[Virtual Network Interface Cards \(VNICs\) \(oracle.com\)](https://docs.oracle.com/en-us/iaas/Content/Resources/Assets/secondary_vnic_all_configure.sh)

https://docs.oracle.com/en-us/iaas/Content/Resources/Assets/secondary_vnic_all_configure.sh

https://docs.oracle.com/en-us/iaas/Content/Resources/Assets/secondary_vnic_windows_configure.ps1

```
[opc@lb-demo-inst1 ~]$ ip addr
```

```
1: lo: <LOOPBACK,UP,LOWER_UP> mtu 65536 qdisc noqueue state UNKNOWN group default qlen 1000
    link/loopback 00:00:00:00:00:00 brd 00:00:00:00:00:00
    inet 127.0.0.1/8 scope host lo
        valid_lft forever preferred_lft forever
    inet6 ::1/128 scope host
        valid_lft forever preferred_lft forever
```

Public

```
2: ens3: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 9000 qdisc pfifo_fast state UP group default qlen 1000
    link/ether 00:13:97:0c:d1:fe brd ff:ff:ff:ff:ff:ff
    inet 10.100.104.2/24 brd 10.100.104.255 scope global dynamic ens3
        valid_lft 3399sec preferred_lft 3399sec
    inet6 fe80::213:97ff:fe0c:d1fe/64 scope link
        valid_lft forever preferred_lft forever
```

Private

```
3: ens6: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 9000 qdisc pfifo_fast state UP group default qlen 1000
    link/ether 00:13:97:d1:a7:8c brd ff:ff:ff:ff:ff:ff
    inet 10.100.101.2/24 scope global ens6
        valid_lft forever preferred_lft forever
    inet6 fe80::213:97ff:fed1:a78c/64 scope link
        valid_lft forever preferred_lft forever
```

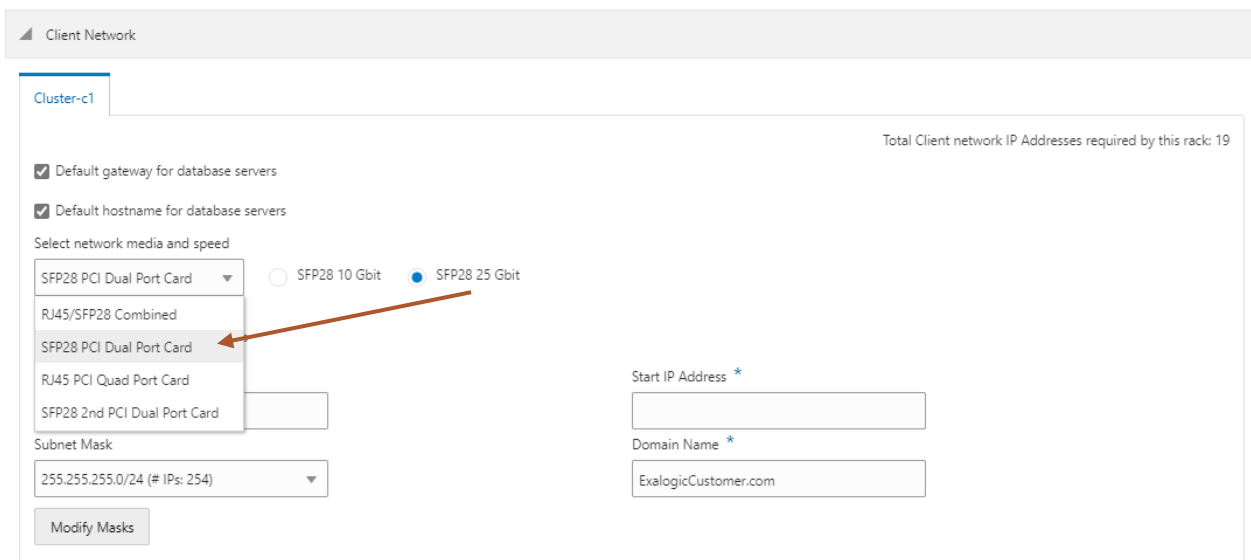
[Return to PCA-Admin shell and execute the final Step to enable subnets to access Exadata Network created in Step1.](#)

CONFIGURING THE ORACLE EXADATA DATABASE MACHINE

The Client-network from Exadata is represented by the green or red 25Gb SFP28 ports in Figure 5.

The Client-network is configured as bondeth0 or bondeth1 interface on each database node. In our example, we will use bondeth1 comprising eth3 and eth4 and should be configured with Linux bonding in active-active/LACP (mode=4). The configuration steps in this solution brief require the use of LACP. Oracle Exadata Deployment Assistant (OEDA) can assist with the configuration of the Client-network.

LACP should be configured to use a policy of layer3+4 and “fast” LACP rate. The default MTU is 1500 bytes but configuring jumbo frames (9000 bytes) is highly recommended. It reduces the level of IP fragmentation and allow for more efficient processing of large streaming workloads with lower CPU overhead. IP networks perform path MTU discovery and auto-negotiate the MTU size between the client (VMs), switches and database-server. Client interfaces using a smaller MTU can still communicate with switches and servers that support jumbo frames.



Optimal settings for bondeth1 are shown in this example:

Client Access Network Configuration

All IP addresses and corresponding host names should be registered in DNS

Starting IP Address : 10.80.65.110
Subnet Mask : 255.255.255.0
Gateway : 10.80.65.1
Client Network Type : SFP28 PCI Dual Port Card Optical Bonded
Client Network Interface : eth3,eth4
Network Speed : 25000

LACP : Enabled, please ensure switches support LACP and configured for LACP and aggregation
BONDING_OPTS="mode=802.3ad miimon=100 downdelay=200 updelay=200 lacp_rate=1 xmit_hash_policy=layer3+4"

```
$ cat /proc/net/bonding/bondeth1
Bonding Mode: IEEE 802.3ad Dynamic link aggregation
Transmit Hash Policy: layer3+4 (1)
MII Status: up
MII Polling Interval (ms): 100
Up Delay (ms): 200
Down Delay (ms): 200
```

```
802.3ad info
LACP rate: fast
Min links: 0
Aggregator selection policy (ad_select): stable
```

Slave Interface: eth3 MII Status: up Speed: 25000 Mbps Duplex: full <...>	Slave Interface: eth4 MII Status: up Speed: 25000 Mbps Duplex: full <...>
---	---

CONFIGURING THE ORACLE DATABASE APPLIANCE (ODA)

Oracle Database Appliance is the easiest and most affordable way for small or medium-size organizations to run Oracle databases and applications and is an ideal platform for remote and edge computing environments. It provides a choice of networking connections for easy integration with current customer infrastructure. The ports needed to connect Client-Network from ODA to PCA-X9, are available from PCIe slots which support Oracle Dual Port 25Gb Ethernet NICs. The following table provides connectivity options from ODA-X8/X9, refer to the ODA documentation to determine the appropriate ports and additional details:

PLATFORM	NETWORK INTERFACES	NETWORK BONDS	REFERENCE
ODA X8-2-HA	BM: [p7p1 p7p2] [p2p1 p2p2] [p10p1 p10p2] VM: [eth3 eth4] [eth7 eth8] [eth11 eth12]	[btbond1] [btbond3] [btbond5] [bond0] [bond2] [bond4]	Network Cabling
ODA X8-2S/2M	BM: [p7p1 p7p2] [p2p1 p2p2] [p10p1 p10p2]	[btbond1] [btbond3] [btbond5]	Network Cabling
ODA X9-2-HA	BM: [p6p1 p6p2] [p4p1 p4p2] [p8p1 p8p2]	[btbond1] [btbond3] [btbond5]	Network Cabling
ODA X9-2S/2L	BM: [p6p1 p6p2] [p4p1 p4p2] [p8p1 p8p2]	[btbond1] [btbond3] [btbond5]	Network Cabling

By default, the bond configuration on ODA nodes uses mode=1 (Active/Standby) during initial provisioning using “*odacli configure-firstnet*”; this mode is incompatible with PCA-X9 DirectConnect feature.

PCA-X9 requires LACP (802.3ad) mode on the peer device. [MOS note ID 2198475.1](#) provides detailed steps to reconfigure the bonding configuration on ODA nodes to use LACP. Bond configuration after successful setup:

```
# cat /proc/net/bonding/btbond1
Ethernet Channel Bonding Driver: v3.7.1 (April 27, 2011)
```

```
Bonding Mode: IEEE 802.3ad Dynamic link aggregation
Transmit Hash Policy: layer3+4 (1)
MII Status: up
MII Polling Interval (ms): 100
Up Delay (ms): 200
Down Delay (ms): 200
```

802.3ad info
 LACP rate: fast
 Min links: 0
 Aggregator selection policy (ad_select): stable
 System priority: 65535
 System MAC address: 00:10:e0:de:fe:1b
 Active Aggregator Info:
 Aggregator ID: 2
 Number of ports: 2
 Actor Key: 15
 Partner Key: 32939
 Partner Mac Address: 00:23:04:ee:be:01

Slave Interface: p6p1 MII Status: up Speed: 25000 Mbps Duplex: full <...>	Slave Interface: p6p2 MII Status: up Speed: 25000 Mbps Duplex: full <...>
---	---

Table 2. Transceivers and cable options

TRANSCEIVERS AND CABLES	X-OPT MKTG PN
QSFP28 100Gb Fiber Short-range Transceiver	7119728
QSFP+ 40Gb Fiber Short-range Transceiver	X2124A
SFP28 25Gb Fiber Short-range Transceiver	7341433
SFP+ 10Gb/1Gb Dual-Speed Fiber Short-range Transceiver	X2129A, 7358257
MPO to 4 LC Optical Splitter OM4 (extended breakout) 10m/20m/50m	7603516 (10m), 7603517 (20m), 7603518 (50m)
MPO to 4 LC Optical Splitter OM3 10m/20m/50m	X2127A-10M, X2127A-20M, X2127A-50M
CBL,1M,CU SPLITTER,QSFP28 to 4xSFP28	7118363
CBL,2M,CU SPLITTER,QSFP28 to 4xSFP28	7118364
CBL,3M,CU SPLITTER,QSFP28 to 4xSFP28	7118365
CBL,5M,CU SPLITTER,QSFP28 to 4xSFP28	7118366
QSFP+ to 4xSFP+ Splitter Cable 1m/3m/5m	X2125A-1M, X2125A-3M, X2125A-5M

Components and Connections – Cables and Transceivers

Copper Breakout Cable



Optical Breakout Cable



LC-LC Coupler



LC-LC Cable



QSFP28 Transceiver



SFP28 Transceiver



LC-LC OM3 MMF Patch Panel



FREQUENTLY ASKED QUESTIONS

1. Can older Exadata racks or ODA platform be connected to PCA-X9 directly?

Yes, the steps shown above will be the same for Exadata-X7, X8 and ODA-X7 which support the Dual-Port-25Gb SFP28 adapters in the database servers.

2. Is Exadata CC supported?

Currently, ExaCC can be connected to PCA-X9 using DirectConnect only if a single VM-cluster will be configured on the ExaCC database servers. During configuration, VLAN tag=1 should be used (see [Create VM Cluster Network](#) section for more details.)

3. Will the default-gateway, DNS and NTP servers be accessible over the Client-Network on the database-server which is connected to the PCA spine switches?

Yes. PCA-X9 provides a Layer-3 network, the Spine switches act as the GW for the Client-Network for Exadata/ODA and provide a route to customer's DNS/NTP servers via the data uplink ports.

4. How many Exadata racks and ODA nodes are supported?

Up to 16 database servers can be connected and these can come from any number of Exadata or ODA. Only 2 physical ports can be part of the same bond interface on the data-base-server.

ODA example: if ODA has 3 Dual-Port-25Gb SFP28 NICs, only 1 NIC can be used for a given Client-Network. The other 2 NICs can be used to create 2 additional Client-Networks ([by repeating Step1 of the PCA-X9 configuration.](#)) This approach provides a flexible, secure, physically isolated, and non-disruptive way to allow application VMs to access the databases running on ODA.

5. Are VLANs supported if customer wants to have Multi-VM Cluster on their ExaCC or ODA?

No, currently the database configuration is expected to use a single VM cluster with access vlan configuration (see [Create VM Cluster Network](#) section for more details.)

CONNECT WITH US

Call +1.800.ORACLE1 or visit [oracle.com](https://www.oracle.com).
Outside North America, find your local office at [oracle.com/contact](https://www.oracle.com/contact).

 blogs.oracle.com

 facebook.com/oracle

 twitter.com/oracle

Copyright © 2022, Oracle and/or its affiliates. All rights reserved. This document is provided for information purposes only, and the contents hereof are subject to change without notice. This document is not warranted to be error-free, nor subject to any other warranties or conditions, whether expressed orally or implied in law, including implied warranties and conditions of merchantability or fitness for a particular purpose. We specifically disclaim any liability with respect to this document, and no contractual obligations are formed either directly or indirectly by this document. This document may not be reproduced or transmitted in any form or by any means, electronic or mechanical, for any purpose, without our prior written permission.

This device has not been authorized as required by the rules of the Federal Communications Commission. This device is not, and may not be, offered for sale or lease, or sold or leased, until authorization is obtained.

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. 0120

Direct Connect Oracle Exadata and Oracle Database Appliance with Oracle Private Cloud Appliance
December 2222
Author: Lalit Bhola

