



Dell Integrated Systems for Oracle Database (DISOD) 2.0

Owner's Guide 1.1

Dell Oracle Solutions Engineering
November 2015

Revisions

Date	Description
August 2015	Document initial release version 1.0
November 2015	Document version 1.1. Replaced S60 with S3048-ON as the Management Switch.

© 2015 Dell Inc. All rights reserved. This product is protected by U.S. and international copyright and intellectual property laws. Dell™ and the Dell logo are trademarks of Dell Inc. in the United States and/or other jurisdictions. All other marks and names mentioned herein may be trademarks of their respective companies.

Contents

1	Overview—Dell Integrated Systems for Oracle Database (DISOD)	4
1.1	DISOD 2.0	5
2	DISOD hardware components	9
2.1	Database servers	9
2.2	Shared storage array	11
2.3	Network switches	12
2.4	Management server	14
3	DISOD network configuration	15
3.1	DISOD private management network	15
3.2	LAN switch configurations	20
3.3	SAN switch configurations	24
3.4	Management switch configurations	26
4	DISOD software configuration	28
4.1	DAAD storage array configuration	28
4.2	Database nodes configuration	31
4.3	Management server configuration	34
5	DISOD management and monitoring	37
A	DISOD rack hardware cabling diagrams	41
A.1	Cabling diagram - LAN infrastructure	41
A.2	Cabling diagram - SAN infrastructure	43
A.3	Cabling diagram - management infrastructure	45
B	DISOD authentication settings	47



1 Overview—Dell Integrated Systems for Oracle Database (DISOD)

Dell Integrated Systems for Oracle Database (DISOD) is a fully integrated hardware stack that is purpose-built as a high performance Oracle Database 12c racked solution. DISOD provides an out-of-box experience to customers, where, everything up to the point of the Oracle Database software installation is preconfigured. By providing integrated and preconfigured hardware and software solution, DISOD provides a fast and easy-to-deploy Oracle Database rack solution to the customers.

The main hardware components of a DISOD rack solution are:

- **Database servers:** Dell PowerEdge servers with preinstalled and preconfigured Linux-based operating system that is Oracle database installation-ready.
- **Shared storage:** All-flash Dell Acceleration Appliance for Databases (DAAD) nodes preconfigured as highly available/high availability (HA) storage arrays and preconfigured with storage volumes required for the Oracle database.
- **Network switches:**
 - **LAN switches:** High bandwidth, low latency, top-of-rack (ToR) Dell Networking Ethernet switches preconfigured for Oracle database's public and private interconnect network traffic.
 - **SAN switches:** High bandwidth SAN switches preconfigured for Oracle database storage network traffic.
 - **Management switch:** Management switch preconfigured and dedicated for the management traffic of the DISOD components.
- **Management server:** Management server preconfigured to enable a one-stop shop management and monitoring of the DISOD hardware components.
- **Console:** KMM (Keyboard/Monitor/Mouse) for local management and monitoring of the DISOD hardware components from within the DISOD solution rack.

The base configuration of a DISOD solution ships with the following preintegrated, preconfigured, racked, stacked, and cabled hardware:

- Two Dell PowerEdge Oracle database servers
- One pair of DAAD storage nodes
- Two ToR Dell Networking Ethernet switches
- Two SAN switches
- One Dell Networking management switch
- One Dell PowerEdge management server, and
- One KMM

The base configuration is also referred to as a 2x2 configuration, where the first '2' represents the number of database nodes, while the second '2' represents the number of individual DAAD nodes. In general, in this document or otherwise, the different DISOD configuration sizes will be referred to with similar representation. For example, a 4x2 configuration refers to a DISOD solution rack that has four database nodes and two individual DAAD storage nodes or one DAAD HA pair.



DISOD is a scale-out solution and is available in multiple configuration sizes in terms of the database compute nodes and storage capacity. However, the scope of this document is to provide the details as a generic design principle that may be applied to various configuration sizes, versus providing the details of each and every configuration sizes supported by the DISOD solution. To know about the different DISOD configurations that are currently supported, contact your Dell Sales team.

This document covers the details of the second release of the DISOD solution, which will be referred to as DISOD 2.0. The subsections and the sections that follow provide the overview, and the hardware and software configuration details of the DISOD 2.0 solution.

1.1 DISOD 2.0

DISOD 2.0 is the second release of the DISOD solution that is built with the latest generation of Dell Enterprise servers and storage. The database nodes and the DAAD storage nodes are based on the latest 13th generation of Dell PowerEdge servers. The second generation of DAAD, that is, DAAD 2.0 that serves as the shared storage in the DISOD 2.0 rack solution is built with newer, faster, and bigger capacity Fusion ioMemory flash adapters than what was used in the DISOD 1.0-DAAD.

Similar to DISOD1.0, DISOD 2.0 also supports only Fibre Channel (FC)-based preintegrated and preconfigured rack solution. The following section provides an overview of the hardware components in the FC-based DISOD 2.0 solution.

1.1.1 FC-based DISOD 2.0

The Fibre Channel-based DISOD 2.0 supports the following main hardware components:

- Database servers
 - Four-socket Dell PowerEdge R930 rack server, or
 - Two-socket Dell PowerEdge R730 rack server
- DAAD shared storage in an HA pair
 - 12.8 TB DAAD 2.0 that uses four 3.2 TB Fusion ioMemory flash adapters
 - 25.6 TB DAAD 2.0 that uses four 6.4 TB Fusion ioMemory flash adapters
- Network switches
 - Two Dell Networking 40 GbE LAN or ToR switches
 - Two Brocade 6510 16 Gbps Fibre Channel SAN switches
 - One Dell Networking S3048-ON 1 GbE management switch
- Management server
 - One Dell PowerEdge R320

The rack layout of a base (2x2) configuration of a Fibre Channel (FC)-based DISOD 2.0 rack solution that uses R730 as the database nodes is shown in Figure 1.



DISOD 2.0 (R730) – Small (2x2)




Figure 1 Fibre Channel-based DISOD 2.0 rack layout – base (2x2) configuration



The architectural diagram of a medium (4x4) FC-based DISOD 2.0 configuration is shown in Figure 2.

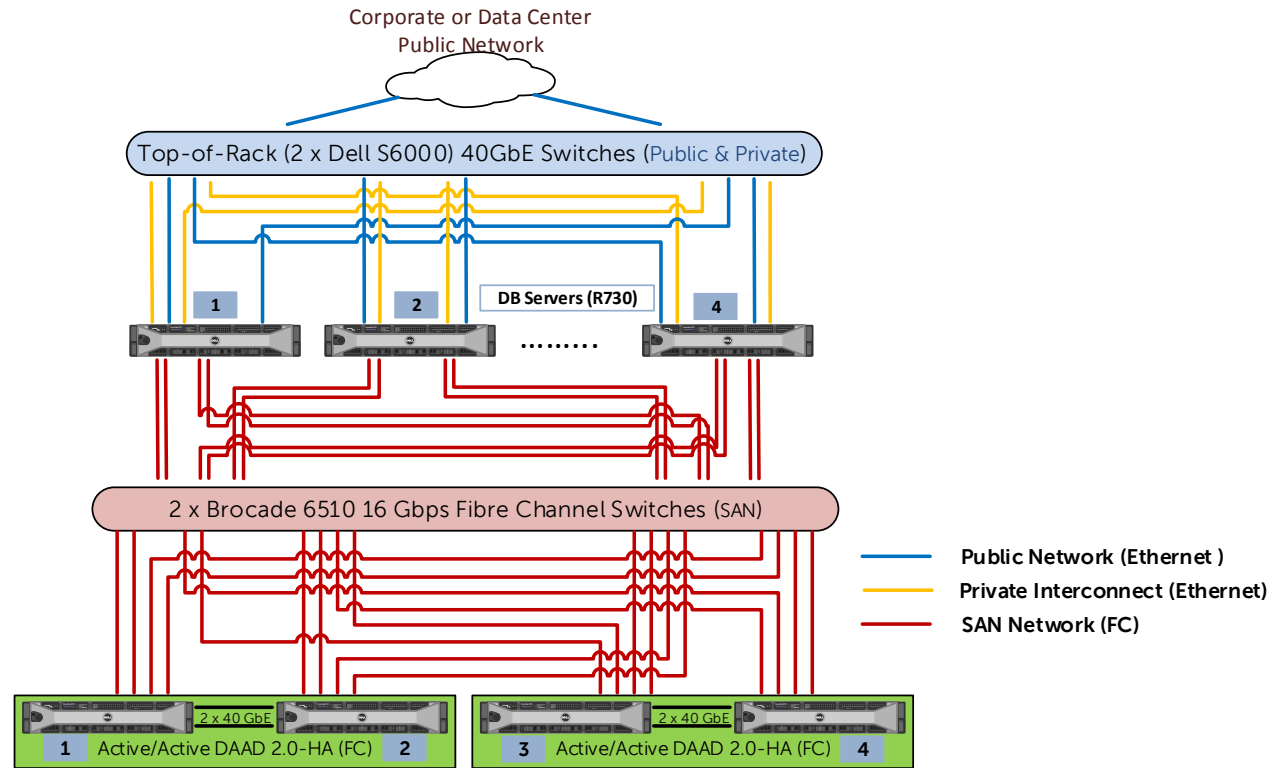


Figure 2 Architectural diagram of DISOD 2.0 – Medium (4x4) FC configuration

The following sections provide the hardware and software configuration details of the individual hardware components in the DISOD 2.0 solution.



2 DISOD hardware components

This section provides information about the main hardware components in the DISOD 2.0 solution.

2.1 Database servers

DISOD 2.0 offers two different servers to be used as Oracle Database nodes:

- Two-socket Dell PowerEdge R730 rack server, or
- Four-socket Dell PowerEdge R930 rack server

The database servers are preinstalled with Oracle Linux 6 running Unbreakable Enterprise Kernel release 3. The OS is preconfigured with the necessary prerequisites that are required for Oracle Database 12c installation. For more information about the OS configuration, see the *4.2 Database nodes configuration* section.

NOTE: Oracle Database software itself is not preinstalled on the database servers.

The following sections provide the hardware details of each of these database servers.

2.1.1 R730 database server

Table 1 provides information about the hardware components used in each of the R730 servers that is supported as the Oracle database compute node in the DISOD 2.0 solution.

Table 1 R730 database server hardware details

Hardware Components – per database server	
Database Server	2S Dell PowerEdge R730
Processor	2 x Intel Xeon E5-2687W v3 10c 3.1 GHz
Memory	Up to 384 GB using 16 GB DDR4 2133 MT/s DR RDIMMs
Network Adapters (Oracle Database Public & Private)	2 x Mellanox ConnectX-3 EN DP 40 Gb Ethernet adapters
Rack Network Daughter Card (rNDC)	Broadcom 5720 quad port 1 GbE rNDC
Remote Management	iDRAC7 Enterprise
HBAs	2 x QLogic 2662 DP 16 GB FC adapter per server
RAID Controller	PERC H730P (RAID 1)
Local Storage	2 x 1.2 TB (RAID 1) + 1 x 1.2 TB hot-spare, 10K 6 Gbps SAS 2.5" hot-plug hard disk drives (HDDs)
Operating System	Oracle Linux 6.6 running Unbreakable Enterprise Kernel (UEK) Release 3 Update 3 (3.8.13-44)



Figure 3 shows the rear side of the R730 chassis with the adapters populated based on the recommended PCIe adapters slot priority of Dell.




Figure 3 R730 database server PCIe slot population

Henceforth, in this document, the DISOD solution that uses R730 servers as the database node will be referenced as R730-based DISOD solution, unless stated otherwise.

2.1.2 R930 database server

Table 2 provides the details of the hardware components used in each of the R930 servers that is supported as the Oracle database compute node in the DISOD 2.0 solution.

Table 2 R930 database server hardware details

Hardware Components – per database server	
Database Server	4S Dell PowerEdge R930
Processor	4 x Intel Xeon E7-8891 v3 10c 2.8 GHz
Memory	Up to 3 TB using 32 GB DDR4 2133 MT/s QR LRDIMMs
Network Adapters (Oracle Database Public & Private)	2 x Mellanox ConnectX-3 EN DP 40 Gb Ethernet adapters
Rack Network Daughter Card (rNDC)	Broadcom 5720 quad port 1 GbE rNDC
Remote Management	iDRAC8 Enterprise
HBAs	2 x QLogic 2662 DP 16 GB FC adapters
RAID Controller	PERC H730P (RAID 1)
Local Storage	2 x 1.2 TB (RAID 1) + 1 x 1.2 TB hot spare, 10K 6 Gbps SAS 2.5" hot-plug HDDs
Operating System	Oracle Linux 6.6 running Unbreakable Enterprise Kernel (UEK) Release 3 Update 3 (3.8.13-44)

Figure 4 shows the rear side of the R930 chassis with the adapters populated based on the recommended PCIe adapters slot priority of Dell.




Figure 4 R930 database server PCIe slot population

Henceforth, in this document, the DISOD solution that uses R930 servers as the database node will be referenced as R930-based DISOD solution, unless stated otherwise.

2.2 Shared storage array

This section describes the shared storage that is preintegrated and preconfigured within the DISOD solution rack. DISOD 2.0 currently supports only Dell Acceleration Appliance for Oracle Databases (DAAD) storage array as the preintegrated shared storage within the solution rack.

The following subsection provides the overview of the DAAD storage array that ships with the DISOD solution rack.

2.2.1 Dell Acceleration Appliance for Oracle Databases (DAAD) storage array

DISOD ships with DAAD storage array as the high performance, low-latency backend shared storage. DISOD does not support standalone DAAD units. As a result, preconfigured DAAD units shipped from factory and future capacity increments are allowed only as DAAD HA pairs.

DISOD 2.0 solution supports the following DAAD:

- Generation: the second generation of DAAD i.e. DAAD 2.0
- Capacities: both the capacity flavors offered with DAAD 2.0
 - 12.8 TB DAAD 2.0 that uses four 3.2 TB Fusion ioMemory flash adapters

- 25.6 TB DAAD 2.0 that uses four 6.4 TB Fusion ioMemory flash adapters
- Fabric: FC for front-end connectivity only

Each HA pair of DAAD in the DISOD solution is precabled to the SAN switches, and is preconfigured with the necessary volumes required for Oracle Database 12c. For more information about the DAAD storage array configuration, see the *4.1 DAAD storage array configuration* section.

For more information about DAAD as a generic storage appliance, click [DAAD](#).

2.3 Network switches

DISOD solution rack ships with the following preintegrated network switches:

- **LAN switches:** High bandwidth, low latency, and highly available LAN switches are used for providing connectivity for customer's applications that are connected to the Oracle databases within the DISOD solution rack. In certain DISOD configurations, it is also used for Oracle database's private interconnect network.
- **SAN switches:** High bandwidth, low latency, and highly available SAN switches are used for providing connectivity between the database nodes and the shared storage within the DISOD solution rack. In certain DISOD configurations, they are also used for Oracle database's private interconnect network.
- **Management switch:** A Gigabit Ethernet switch is used for providing secure and segregated DISOD monitoring and management capability.

The following subsections describe these switches in more detail.

2.3.1 LAN switches

This section describes the LAN switches that are supported, preintegrated, and preconfigured within the DISOD solution rack.

2.3.1.1 Dell Networking S6000 Ethernet switches

Redundant high bandwidth Dell Networking S6000 40 GbE switches are set up as top-of-rack LAN switches. To provide flexibility, the S6000s are preconfigured to provide both 10 GbE and 40 GbE ports to plug the DISOD into the customer's network environment.

The S6000 switches are primarily used for:

- **Oracle Public Network:** Provide public network to connect the customer's Application Servers to the Oracle database's public network.
- **Oracle Private Network (FC-based DISOD only):** Provide private network for the Oracle Real Application Cluster (RAC) database's private interconnect traffic.
- **Remote Management:** Provides remote connectivity to DISOD's dedicated private management network (S3048-ON switch) for customers and Dell Services/Support to remotely monitor and manage the hardware components in the DISOD solution rack.



Both the S6000 switches are configured as Layer-2 switches. The switches ship with the following preconfigurations:

- Rapid Spanning Tree Protocol (RSTP)
- Virtual Link Trunking (VLT)
- Predefined port functions and configuration
- VLAN and Port Channels

For more information about the S6000 switch configurations, see the *3.2.1 S6000 switch configurations* section.

2.3.2 SAN switches

This section describes the switches that are supported, preintegrated, and preconfigured within the DISOD solution rack as the SAN switches.

2.3.2.1 Brocade 6510 Fibre Channel switches

For FC-based DISOD solutions, Brocade 6510 16 Gbps FC switches are used as the SAN switches. To sustain the high throughput and performance delivered by the DAAD storage arrays, redundant and high-bandwidth Brocade 6510 switches are used to connect the database nodes to the shared storage.

For more information about the 6510 switch configurations, see the *3.3.1 Brocade 6510 switch configurations* section.

2.3.3 Management switch

This section describes the switch that is supported, preintegrated, and preconfigured within the DISOD solution rack as the management switch.

2.3.3.1 Dell Networking S3048-ON switch

The Dell Networking S3048-ON switch is a 1GbE enterprise switching platform that supports both an industry hardened OS and open networking, providing freedom to run third-party operating systems (OS). This is configured as the management switch. The ports on the switch are preconfigured and precabled to provide secure and isolated access to the hardware components in the DISOD solution rack. This dedicated management switch also ensures the segregation of the public and the private rack network traffic.

NOTE: For the remainder of this document S3048-ON switch will be referred to as S3048, unless stated otherwise.

For more information about configuration, see the *3.4.1 S3048 switch configurations* section.



2.4 Management server

The Dell PowerEdge R320 server is set up as the management server that allows remote management and monitoring of all the DISOD hardware components. Table 3 lists all the components in the management server.

Table 3 R320 management server components

Management Server Components	
Server	1S Dell PowerEdge R320 server
Processor	1 x Intel Xeon E5-2430L v2 6c 2.4 GHz
Memory	48 GB (6 x 8 GB 1600 MT/s RDIMMs)
Network Adapters	1 x Broadcom 5720 DP 1 GbE adapter 1 x Broadcom 5720 DP 1 GbE LOM
Remote Management	iDRAC7 Enterprise
RAID Controller	PERC H310 (RAID 1)
Local Storage	2 x 500 GB (RAID 1) + 1 x 500 GB hot-spare, 7.2K 3 Gbps SAS 3.5", hot-plug HDDs
Operating System	Oracle Linux 6.6 running Red Hat Compatible Kernel (2.6.32-504.el6.x86_64)

The management server is preinstalled and configured with Microsoft Windows Server 2012 as the base OS. It is also preconfigured with the following two Hyper-V virtual machines (VMs):

- **WinVM:** Preinstalled with Windows Server 2012 as the guest OS. Is primarily created to install the DISOD management software.
- **LinVM:** Preinstalled with Oracle Linux 6.6 running Red Hat Compatible Kernel as the guest OS. Is available to provide a Linux-based environment to access the DISOD hardware components.

Hereafter, in this document, these two VMs will be referred to as WinVM and LinVM, unless stated otherwise. For information about the OS configuration on the management server, see the *4.3 Management server configuration* section.



3 DISOD network configuration

This section provides information about the hardware and software network that DISOD ships with. It provides information about the following network configurations:

- DISOD private network
- LAN switch configurations
- SAN switch configurations
- Management switch configurations
- Management and monitoring network

3.1 DISOD private management network

To provide an out-of-box experience, DISOD is built with a self-contained private management network that is preconfigured before the solution is shipped out to the customers. This helps reduce the hardware and software network configuration efforts required to be done at the customer site. This is achieved by identifying and assigning predefined IP addresses to those network ports that do not have dependency on the customer's public network.

Table 4 lists the preconfigured static IP address ranges for the different DISOD components.

Table 4 DISOD preconfigured static IP address ranges

Network port	DISOD component	IP address range	Max hosts
Network address/Subnet class: 192.168.211.0/26			
Management or private network ports	Switches ¹ , Management Server	192.168.211.1 – 192.168.211.10	10
iDRAC	Database nodes	192.168.211.11 – 192.168.211.20	10
Private mgmt. ports (LOM – em1)	Database nodes	192.168.211.21 – 192.168.211.30	10
iDRAC	DAAD nodes	192.168.211.31 – 192.168.211.40	10
Network address/Subnet class: 192.168.201.0/24			
Oracle interconnect network ports	Database nodes	192.168.201.1 – 192.168.201.10	10
Network address/Subnet class: 192.168.200.0/24			
Oracle redundant interconnect network ports	Database nodes	192.168.200.1 – 192.168.200.10	10
Network address/Subnet class: 192.168.0.0/24 ²			
HA interconnect ports	DAAD nodes	192.168.0.1 – 192.168.0.2	2
Network address/Subnet class: 192.168.1.0/24			
HA redundant interconnect ports	DAAD nodes	192.168.1.1 – 192.168.1.2	2

¹ On the switches, the static management IP is configured on the private VLAN30 network; Physical management ports are not configured.

² The HA network ports on the two DAAD nodes configured in an HA pair are attached directly to each other. Thus, multiple DAAD HA pairs can exist within a single DISOD rack with the same static IP addresses.



Network address/Subnet class: 172.16.211.0/26 ³			
Public network ports	Management Server	172.16.211.8	1
Public mgmt. ports (LOM – eth0)	DAAD nodes	172.16.211.41 – 172.16.211.50	10

Some components within the DISOD solution rack are fixed in quantity, such as the switches and the management server. However, because DISOD is a scalable solution some components such as the database nodes and the DAAD nodes are variable in quantity. As a result, the private IP addresses for the fixed components are predictable and are provided in the following subsections. For information about the private IP addresses used on the switches, see the *3.1.1 Static IP addresses* section, and for the private IP addresses preassigned on the management server, see the *3.1.2 Management server network configuration* section.

The number of database nodes and DAAD nodes within a DISOD solution rack depends on the DISOD size a particular customer is using. Therefore, for the scalable or variable components such as the database nodes and the DAAD nodes, user must refer to Table 4 to determine the IP address for the particular network component or port. The IP addresses are preassigned, starting with the first IP address in the range and incremented with each increasing node number. Refer to section *3.1.3 Static IP addresses of Database and DAAD nodes* that provides a 2x2 DISOD configuration as an example to help with the static IP address design logic for the database and the DAAD nodes.

3.1.1 Static IP addresses of Switches

All the switches in the DISOD solution rack ship with preassigned static IP addresses. Table 5 lists these preassigned static IP addresses.

Table 5 Preassigned switches static IP addresses

DISOD Component	(Virtual) Port	IP Address
S6000-u42	VLAN 30 (Mgmt.)	192.168.211.1
	Physical management port	<Not configured; to be configured at customer site (optional)>
S6000-u41	VLAN 30 (Mgmt.)	192.168.211.2
	Physical management port	< Not configured; to be configured at customer site (optional)>
S3048-u40	VLAN 30 (Mgmt.)	192.168.211.3
	Physical management port	< Not configured; to be configured at customer site (optional)>
6510-u18	Physical management port	192.168.211.4
6510-u17	Physical management port	192.168.211.5

As shown in Table 5, the physical management port on the two S6000 switches and the single S3048 switch are not directly configured. The preconfigured static management IP

³ These routable public IP addresses are temporarily assigned. These should be changed to the customer's public network before plugging the DISOD ToR switches to the customer's public network to prevent any IP address conflicts.



address is assigned to VLAN30, the virtual DISOD private management network. User can choose to configure the physical management port manually at a later time, if need be.

Any direct management of these switches can be done by connecting to the IP addresses provided in Table 5.

3.1.2 Management server network configuration

As seen in Figure 5, the management server (R320) ships with a dual-port 1 GbE LOM and a dual-port 1 GbE PCIe network adapter. The ports are teamed or bonded at the OS level to provide two networks—one used for public network access and the other used for DISOD private management network access.

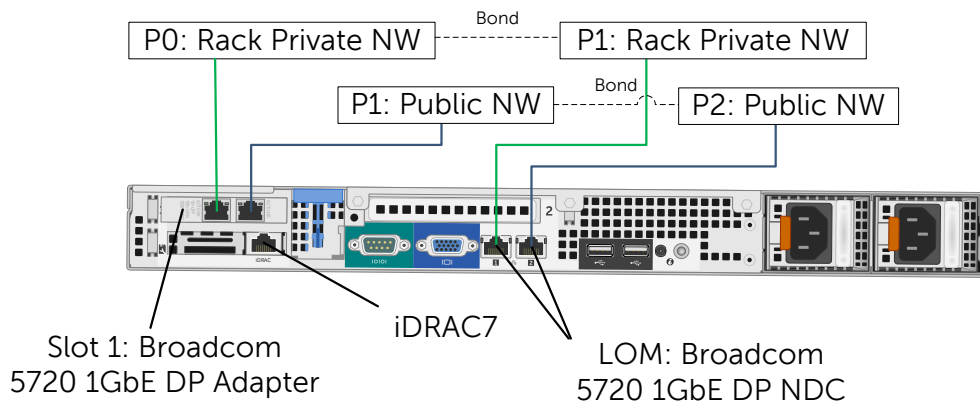


Figure 5 Management server network adapters and ports

Figure 6 shows an overview of the network configuration in the management server—both at the physical OS and at the two VMs' level.

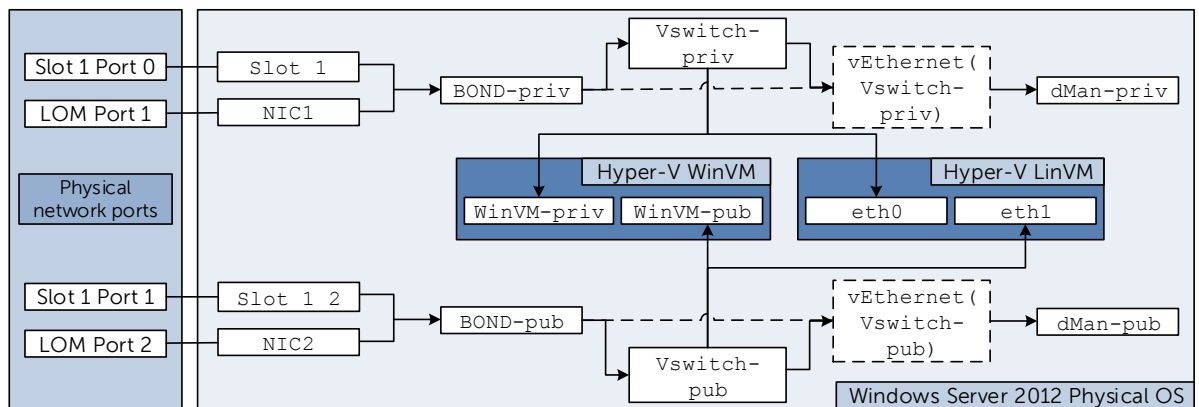


Figure 6 Management Server networking configuration

Network configuration at physical OS level

Table 6 shows the NIC teaming configured at the physical OS level. Teaming is configured using the Windows Server 2012 native teaming software.



Table 6 Management server NIC teaming configuration

Teaming Name	Adapter 1 (LOM)		Adapter 2 (add-on)	
	Physical Port #	OS Name	Physical Port #	OS Name
BOND-priv	1	NIC1	0	SLOT 1
BOND-pub	2	NIC2	1	SLOT 1 2

Table 7 Management server virtual switches configuration

Virtual Switch Name	Teamed Adapter Names
Vswitch-priv	BOND-priv
Vswitch-pub	BOND-pub

After the virtual switches were created, Windows automatically created two more virtual network adapters at the base OS level called vEthernet (Vswitch-priv) and vEthernet (Vswitch-pub). These are technically the same adapter as BOND-priv and BOND-pub. These were renamed with the following two user-friendly names:

- dMan-priv: renamed from vEthernet (Vswitch-priv). This adapter represents the DISOD management (dMan) server's private network adapter at the physical OS level.
- dMan-pub: renamed from vEthernet (Vswitch-pub). This adapter represents the DISOD management server's public network adapter at the physical OS level.

Table 8 lists the pre-assigned IP address for the private management port (dMan-priv). Also, as shown in Table 8, the public management port and the iDRAC port on the management server are not preconfigured because these must be configured with the customer's public network IP address.

Table 8 Management server static IP addresses

Adapter Name	IP Address
dMan-priv	192.168.211.6
dMan-pub	<Not configured; to be configured at customer site>
iDRAC	<Not configured; to be configured at customer site>

Network configuration at VM's guest OS level

Table 9 lists the network configuration of the two Hyper-V VMs in the management server.

Table 9 Management server VM network configuration

Virtual Machine	Virtual Switch	Network Configuration	
WinVM [Windows Server 2012 guest OS]	Vswitch-priv	Private Network	
		Adapter Name	IP Address
		WinVM-priv	192.168.211.7
	Vswitch-pub	Public Network	
		Adapter Name	IP Address
		WinVM-pub	<Not configured; to be configured at customer site>
LinVM [Oracle Linux 6.6 guest OS]	Vswitch-priv	Private Network	
		Adapter Name	IP Address
		eth0	192.168.211.8
	Vswitch-pub	Public Network	



		Adapter Name	IP Address
		eth1*	172.16.211.8

*** IMPORTANT:** eth1 in LinVM is preconfigured with a temporary public routable IP address. This was required to preconfigure the DAAD storage. Customers must ensure to change this IP address to match their public network before plugging in to the DISOD ToR switches to avoid any IP conflicts.

3.1.3 Static IP addresses of Database and DAAD nodes

Table 10 shows an example of the preassigned IP addresses for the different network components on the database nodes and the DAAD nodes within a 2x2 DISOD solution rack. The IP address increases by one with each increasing node number for a particular network component. For example, the private management port (em1) on the database nodes are in the IP address range of 192.168.211.21 to 192.168.211.30—with database node #1 with IP address 192.168.211.21, database node #2 with 192.168.211.22, and so on. Similarly, the iDRAC IP addresses on the DAAD nodes are in the range of 192.168.211.31 to 192.168.211.40—with DAAD node #1 with IP address 192.168.211.31, DAAD node #2 with 192.168.211.32, and so on.

Table 10 Preassigned static IP addresses on the database and DAAD nodes

Node Type	Node ID	Network Component					
		Component	IP Address	Subnet	Mask	Port	
Database Nodes	R730-u19-20 (DB-N1)	iDRAC	192.168.211.11	255.255.255.192	192.168.211.0	/26	11
		Private Mgmt. (em1)	192.168.211.21	255.255.255.192	192.168.211.0	/26	21
		Oracle Private A (eth5)	192.168.201.1	255.255.255.0	192.168.201.0	/24	1
		Oracle Private B (eth7)	192.168.200.1	255.255.255.0	192.168.200.0	/24	1
		Oracle Public (eth4+eth6 bonded)	Not Configured – to be configured at customer site				
	R730-u21-22 (DB-N2)	iDRAC	192.168.211.12	255.255.255.192	192.168.211.0	/26	12
		Private Mgmt. (em1)	192.168.211.22	255.255.255.192	192.168.211.0	/26	22
		Oracle Private A (eth5)	192.168.201.2	255.255.255.0	192.168.201.0	/24	2
		Oracle Private B (eth7)	192.168.200.2	255.255.255.0	192.168.200.0	/24	2
		Oracle Public (eth4+eth6 bonded)	Not Configured – to be configured at customer site				
DAAD Storage Arrays	DAAD1A u15-16	iDRAC	192.168.211.31	255.255.255.192	192.168.211.0	/26	31
		Public Mgmt. (eth0)	172.16.211.31	255.255.255.192	192.168.211.0	/26	31
		HA 1	192.168.0.1	255.255.255.0	192.168.0.0	/24	1
		HA 2	192.168.1.1	255.255.255.0	192.168.1.0	/24	1
	DAAD1A u13-14	iDRAC	192.168.211.32	255.255.255.192	192.168.211.0	/26	32
		Public Mgmt. (eth0)	172.16.211.32	255.255.255.192	192.168.211.0	/26	32
		HA1	192.168.0.2	255.255.255.0	192.168.0.0	/24	2
		HA2	192.168.1.2	255.255.255.0	192.168.1.0	/24	2



3.2 LAN switch configurations

This section describes the switch configuration details of the LAN switches supported by the DISOD solution rack.

3.2.1 S6000 switch configurations

Dell Networking S6000 40 Gb Ethernet switches are used as the ToR or LAN switches in the DISOD solution rack. This section provides information about the switch settings that the S6000 switches ship with. Each S6000 switch is configured by using its own settings file before being shipped out to the customer. The settings that are unique to each switch are listed in Table 11.

Table 11 S6000 switches unique settings

Setting	S6000 at rack unit #42 (Primary)	S6000 at rack unit #41 (Secondary)
Switch name or hostname	S6000u42	S6000u41
VLT domain		
unit-id	0	1
back-up destination	192.168.211.2	192.168.211.1
primary-priority	1	8192
VLAN		
Vlan 30 (DISOD private network)	192.168.211.1/26	192.168.211.2/26

Other settings are common and the following subsections provide more information about the remaining switch configurations.

3.2.1.1 S6000 predefined and preconfigured port functions

Using one of the S6000 switches as an example, Figure 7 shows the details of the predefined port functions and the configuration the switch ships with.

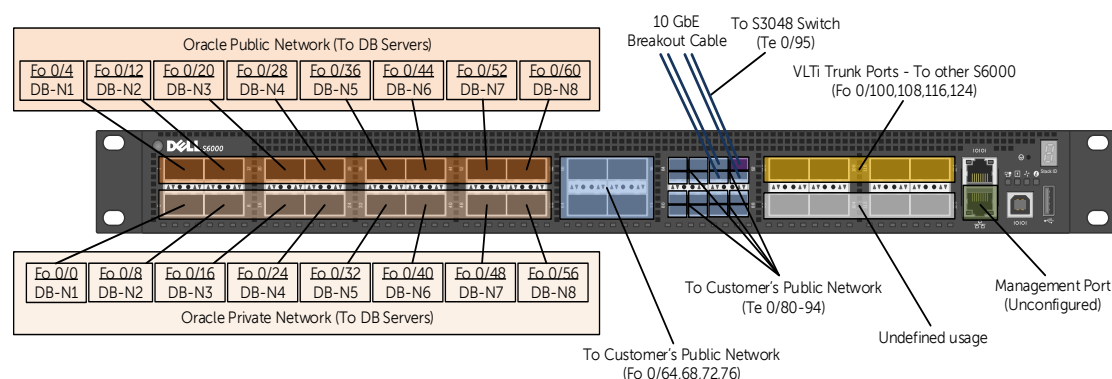


Figure 7 S6000 predefined and preconfigured port functions

The port connectivity is predefined and preconfigured with the following functional logic:

- Database node connectivity:** Ports 0–60 connect to the network ports on the database nodes. These are configured as FortyGigE (Fo) ports.
 - The ports in the same column connect to the network ports on the same database node.

- The database node number increases from left to right
- In each column, top ports (0/4, 12, 20, 28, 36, 44, 52, and 60) connect to the Oracle RAC public network port, and the bottom ports (0/0, 8, 16, 24, 32, 40, 48, and 56) connect to the Oracle RAC private interconnect port on the same database node.
- The ports that are precabled from factory depend on the customer's DISOD configuration size. For example, for DISOD 4x4 configuration, ports 0–28 are precabled while ports 32–60 are reserved for connectivity to expansion database nodes.
- **Customer's network connectivity:** These are the ports that are configured to uplink to the customer's datacenter or public network so that their applications can connect to the Oracle databases in the DISOD solution rack.
 - **FortyGigE connectivity:** Ports 0/64, 68, 72, and 76 are preconfigured as FortyGigE ports to allow for connectivity to the customer's FortyGigE network environment. These ports do not come precabled from factory.
 - **TenGigE connectivity:** FortyGigE ports 0/80, 84, 88, and 92 are configured as TenGigE (Te) ports. Each of these FortyGigE ports are further split in to four logical TenGigE ports, giving rise to sixteen logical TenGigE ports in total, that is numbered from 80–95. A 'FortyGigE QSFP+ to 4xTenGigE SFP+' breakout cable is required to access each of these logical TenGigE ports. Thus, logical ports 80–94 are available to connect to the customer's TenGigE network environment. Ports 80–91 do not come precabled from factory. Logical port 95 is reserved for management connectivity.
- **DISOD management connectivity:** To have a preconfigured access to the DISOD management infrastructure, port 92 comes precabled from factory with the breakout cable. Each SFP+ cable on the breakout cable is labeled AAAA, BBBB, CCCC, and DDDD that maps internally to the four TenGigE ports 92, 93, 94, and 95 respectively. The SFP+ cable labeled DDDD on each S6000 switch that maps to port 95 internally, is directly attached to the TenGigE ports on the front of the S3048 switch.
- **VLTi inter-switch connectivity:** Ports 0/100, 108, 116, and 124 are configured as FortyGigE ports and are used as VLTi trunk ports. They are connected to the corresponding VLTi trunk ports on the peer S6000 switch. These are precabled from factory.
- **Undefined usage:** Ports 0/96, 104, 112, and 120 are configured as FortyGigE ports with undefined functionality or usage. These ports do not come precabled from factory.

Based on the best practices recommended by Dell, the ports dedicated for the Oracle RAC private interconnect connectivity are configured as `portfast` to allow for immediate transition to a forwarding state. This prevents the Oracle Grid infrastructure's Cluster Synchronization Services (CSS) `misscount` parameter from missing any network heartbeats that could lead to database node evictions.



3.2.1.2 S6000 VLAN and Port Channel configuration

Though it is recommended to have dedicated switches for the Oracle database's public and private traffic, in this FC-based preintegrated solution, the S6000 switches are used for handling both types of traffic. In such configurations, it is recommended to segregate these two network traffics within the switch. Hence, the switches are preconfigured by segregating these two network traffics on separate VLANs. Using one of the S6000 switches as an example, Figure 8 shows the VLAN and port channel configuration on the S6000 switches.

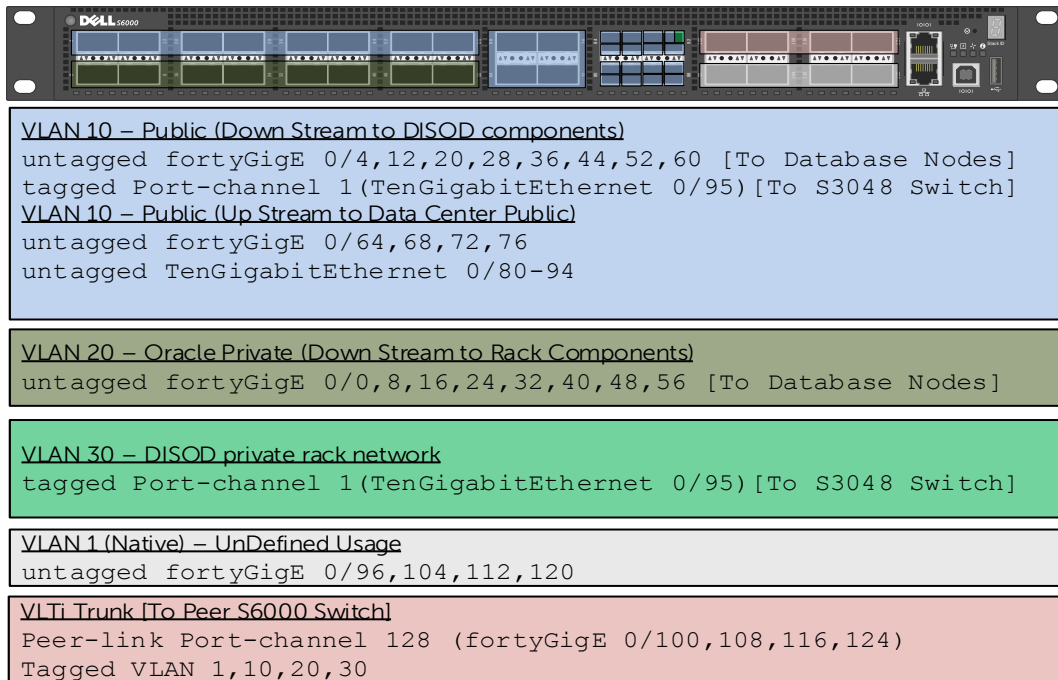


Figure 8 S6000 VLAN and port channel configuration

Four FortyGigE ports (0/64, 68, 72, and 76) are preconfigured for customers to connect and communicate with the public network of Oracle database and to remotely connect with the Management Server. If more than four is required then the undefined ports 0/96, 104, 112, and 120 can be used for the same. Similarly, if forty Gigabit network is not available then customers can choose to use up to 15 preconfigured TenGigE ports (0/80-94) to connect and communicate with the DISOD public network.

VLAN 30 is configured with virtual IP address that is used for managing the switch. Dedicated physical management port on the switch is not configured. Though not recommended, if needed, customers can configure and connect the dedicated management port to their network for direct access and management.

By preconfiguring the Oracle database public, private, and management traffic on distinct VLAN IDs, the DISOD ToR S6000 switches provide a secure and segregated network environment.

3.2.1.3 S6000 VLT domain configuration

Virtual Link Trunking (VLT) is a proprietary aggregation protocol by Dell Force 10 that is available in its enterprise-class network switches. VLT provides a loop-free environment by presenting two physical switches as one logical switch while not losing bandwidth for the devices connecting to it over two separate links.

Figure 9 shows the VLT Domain configuration across the two S6000 switches implemented in the DISOD solution rack, with the S6000 located at rack unit #42 configured as the primary switch (`unit-id 0`) and the S6000 located at rack unit #41 configured as the secondary peer switch (`unit-id 1`). Four FortyGigE ports (0/100, 108, 116, and 124) are used for VLTi trunking and configured on Port Channel 128, to provide adequate bandwidth for any customer network connecting to it. The VLT backup link configured on each S6000 switch is the management IP address of the other peer S6000 switch in the VLT domain. As noted earlier, the management IP address is configured and assigned to VLAN 30 on both the S6000 switches. Port Channel 1 on both the S6000 switches that connect to the S3048 management switch is configured as VLT peer LAG to provide 20 GB of bandwidth and high availability.

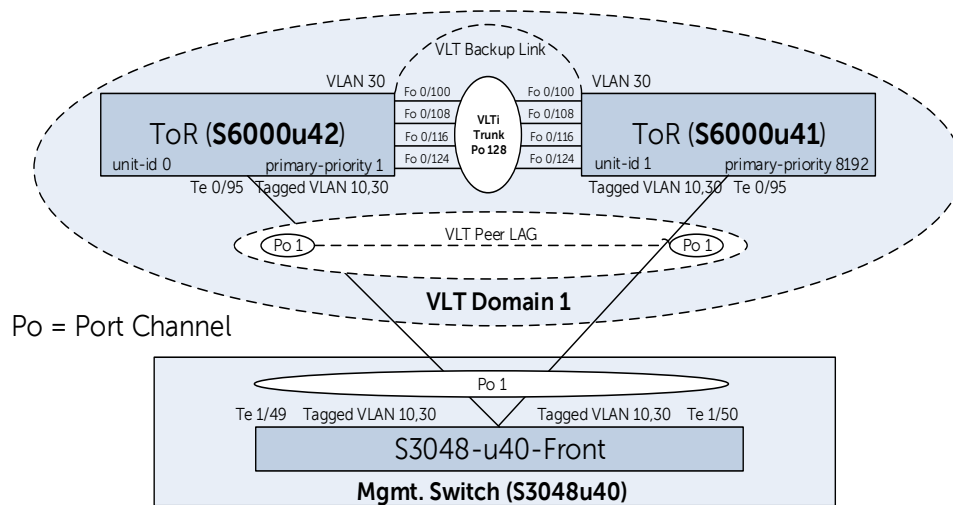


Figure 9 S6000 VLT domain configuration

The VLT domain configuration requires a system MAC address to be set for internal operations. Though the Dell Networking OS automatically generates a system MAC address during the VLT domain configuration, both the S6000 switches in the DISOD solution rack are configured with the same static MAC address in order to speed up the synchronization of the default MAC addresses on the VLT domain on both peer switches when one peer switch reboots. The static MAC address assigned to the two switches is `92:b1:1c:f4:e7:4e` that identifies it as a locally administered address (`x2-xx-xx-xx-xx-xx`) or private MAC address, and prevents any potential conflict with any other universally administered MAC address.



3.2.1.4 S6000 Spanning Tree Protocol (STP) configuration

Though VLT offers a loop-free Layer2 topology, as a best practice to prevent configuration and patching mistakes, Rapid Spanning Tree Protocol (RSTP) is enabled on both the S6000 switches. Enabling RSTP also helps prevent any loops that could occur just before the VLT domain is established. Both the switches are configured with the default S6000 RSTP switch settings. Some of the important ones are:

- Bridge priority on both the S6000 switches are left as-is at the default high value of 32768 to ensure that either of these two DISOD ToR switches do not become the `root` bridge when plugged in to the customer's public network and cause any type of disruption. These priority values can be adjusted or reconfigured later based on the customer's spanning tree topology preferences.
- `PortFast` is enabled on the FortyGigE ports that are used for Oracle database private network to prevent any delay for the Oracle Grid infrastructure's Cluster Synchronization Services (CSS) `misscount` feature that could lead to missed network heartbeats and Oracle node evictions.

3.3 SAN switch configurations

This section describes the switch configuration details of the SAN switches supported by the DISOD solution rack.

3.3.1 Brocade 6510 switch configurations

Brocade 6510 16 Gbps switches are used as the SAN switches in the FC-based DISOD solution. This section provides information about the switch settings that the redundant 6510 FC switches ship with.

The 6510 switches are preconfigured with port-based zoning in order to provide out-of-box access to the storage volumes from the database nodes. Using one of the Brocade 6510 switches in a 4x4 DISOD configuration as an example, Figure 10 shows the predefined port connectivity configuration between 6510 and the database nodes, and between 6510 and the DAAD HA pair nodes.

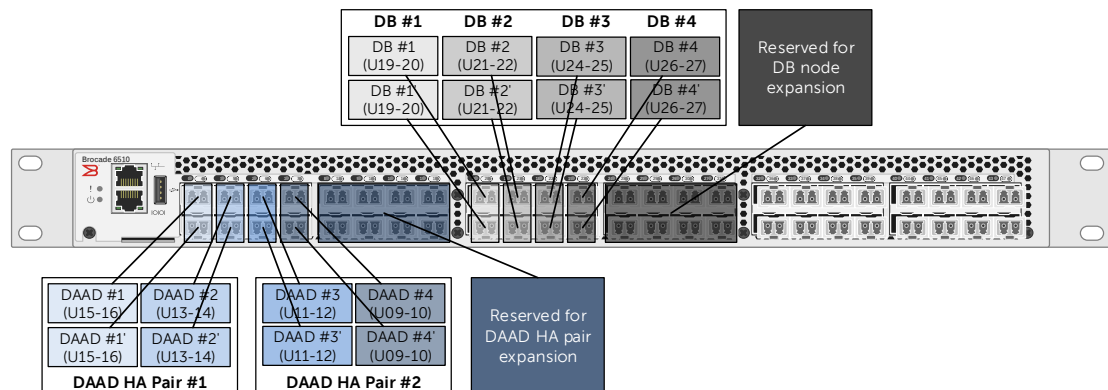


Figure 10 Brocade 6510 FC switch predefined port connectivity



As shown in the example in Figure 10 and in general, the switch configuration design is as follows:

- The two switch ports in the same column connect to two ports on either the same DAAD node or on the same database node
- The DAAD or the database node number increases from left to right
 - Port range 0–15 is designated and used for connecting to the DAAD nodes
 - Port range 16–31 is designated and used for connecting to the database nodes
- Port range 32–47 usage is undefined. Potential use cases:
 - To connect to additional database nodes in other expansion DISOD racks
 - To connect to database backup servers

Each of the 6510 switches is configured with port-based zoning. As part of this, each initiator port on the database nodes along with all the ports from the DAAD storage array are configured in its own zone on the switch.



3.4 Management switch configurations

Dell Networking S3048 1 Gb Ethernet switch is used as the management switch in the DISOD solution rack. This section provides information about the configuration that the S3048 management switch ships with.

3.4.1 S3048 switch configurations

This section provides information about the switch settings that the management switch S3048 ships with.

3.4.1.1 S3048 predefined and preconfigured port functions

Figure 11 shows the details of the predefined port functions and the configuration on the S3048 management switch.

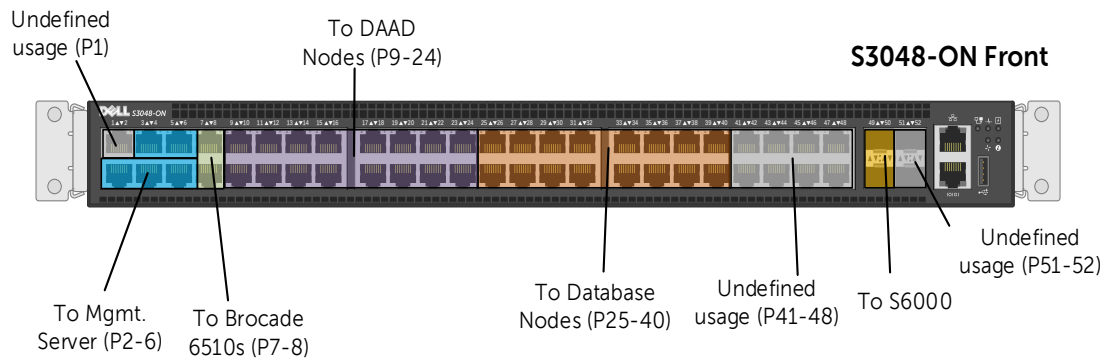


Figure 11 S3048 predefined and preconfigured port functions

The port connectivity is predefined and preconfigured with the following functional logic:

- **Management server connectivity:** Ports 2–6 connect to the network ports on the management server.
- **SAN switches connectivity:** Ports 7–8 that are located in the same column connect to the management port on the two SAN switches.
- **DAAD storage connectivity:** Ports 9–24 connect to the network ports on the DAAD storage nodes.
 - The ports in the same column connect to the network ports on the same DAAD storage node
 - The DAAD node number increases from left to right
 - In each column, top ports connect to the management port (LOM port 1 or `eth0`), and the bottom ports connect to the iDRAC port on the same DAAD node
- **Database node connectivity:** Ports 25–40 connect to the network ports on the database nodes.
 - The ports in the same column connect to the network ports on the same database node
 - The database node number increases from right to left

- In each column, top ports connect to the management port (LOM port 1 or em1), and the bottom ports connect to the iDRAC port on the same database node
- **LAN switches connectivity:** Two 10 GbE SFP+ ports are used to connect to the ToR LAN switches.
 - Port 49 is connected to the top LAN switch
 - Port 50 is connected to the bottom LAN switch
- **Undefined connectivity:** Ports 1, 41–48, and 51–52 connectivity is not predefined.

Refer to appendix section 4.3 for management infrastructure cabling diagram.

3.4.1.2 S3048 VLAN and Port Channel configuration

Figure 12 shows the VLAN and the Port Channel configuration information about the S3048 management switch.

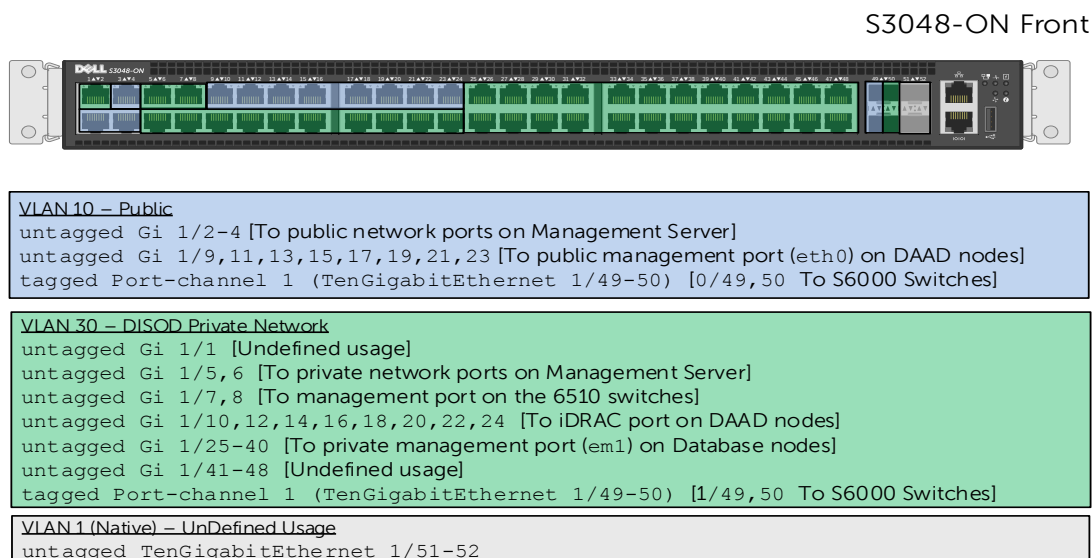


Figure 12 VLAN configuration on the S3048 management switch

The S3048 switch is preconfigured with two segregated VLANs:

- VLAN 10 for the public management traffic coming in to the DISOD solution rack,
- VLAN 30 for the private management traffic to securely access the DISOD hardware components from within the rack. VLAN 30 is also preconfigured with a virtual IP address that is used as the management IP address.

The physical management port on the switch is not configured. Customers can choose to configure this port, if necessary. RSTP is enabled across the switch with a low bridge priority of 40960 to ensure that this switch does not become the root bridge and cause any network disruption when plugged in to the customer's public network. In general, customers are not required to change this bridge priority value.

4 DISOD software configuration

This section describes the software configuration that the following hardware components in the DISOD ships with:

- DAAD storage array configuration
- Database nodes configuration
- Management server configuration

4.1 DAAD storage array configuration

This section provides the details of the DAAD 2.0 storage array that is configured by using Dell's best practices for Oracle Databases. As described in the section *2.2 Shared storage array*, DISOD 2.0 supports both capacity flavors of DAAD 2.0—12.8 TB DAAD 2.0 that uses four 3.2 TB Fusion ioMemory adapters, and 25.6 TB DAAD 2.0 that uses four 6.4 TB Fusion ioMemory adapters. The software configuration details and the best practices described in this section applies to both flavors of DAAD 2.0 storage arrays that are supported with DISOD 2.0.

Each DAAD storage node has four ioMemory flash adapters, each with either 3.2 TB or 6.4 TB disk capacity. As shown in Figure 13, with DAAD ION HA clustering enabled, the four ioMemory adapters (ioMemory1-4) in one DAAD storage node are mirrored with the four ioMemory adapters in the other node of the HA pair.

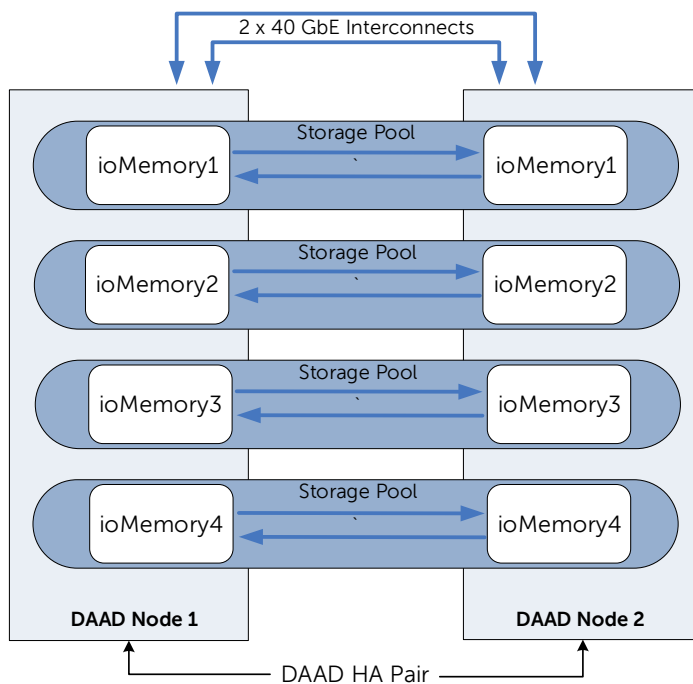


Figure 13 DAAD HA pair - ioMemory mirroring and storage pools

As shown in Figure 13, each storage pool contains two ioMemory cards, one from each of the two DAAD nodes in the HA pair. Altogether, four storage pools are created per DAAD HA pair.

Each volume that is preconfigured is assigned a primary storage node and a secondary storage node. DAAD ION clustering software uses synchronous replication protocol. Local 'write' operations on the primary node are considered complete only after both the local and the remote disk write operations have been confirmed. As a result, loss of a single node is guaranteed not to lead to any data loss. As shown in Figure 14, volumes *v_a1*, *v_b1*, *v_c1*, and *v_d1* use storage node 1 (*fcion1a*) as the primary node and volumes *v_a2*, *v_b2*, *v_c2* and *v_d2* use storage node 2 (*fcion1b*) as the primary node to present the volumes to the database nodes. Volumes *v_a1'*, *v_b1'*, *v_c1'*, *v_d1'*, *v_a2'*, *v_b2'*, *v_c2'* and *v_d2'* are their respective mirrored volumes. By creating two volumes per storage pool and by alternating the primary node between the two DAAD nodes in the HA pair, balances the workloads while maintaining redundancy for high availability.

Each of these eight volumes is created with the same size and are presented to the Oracle database nodes to form Oracle ASM disks of an ASM disk group.

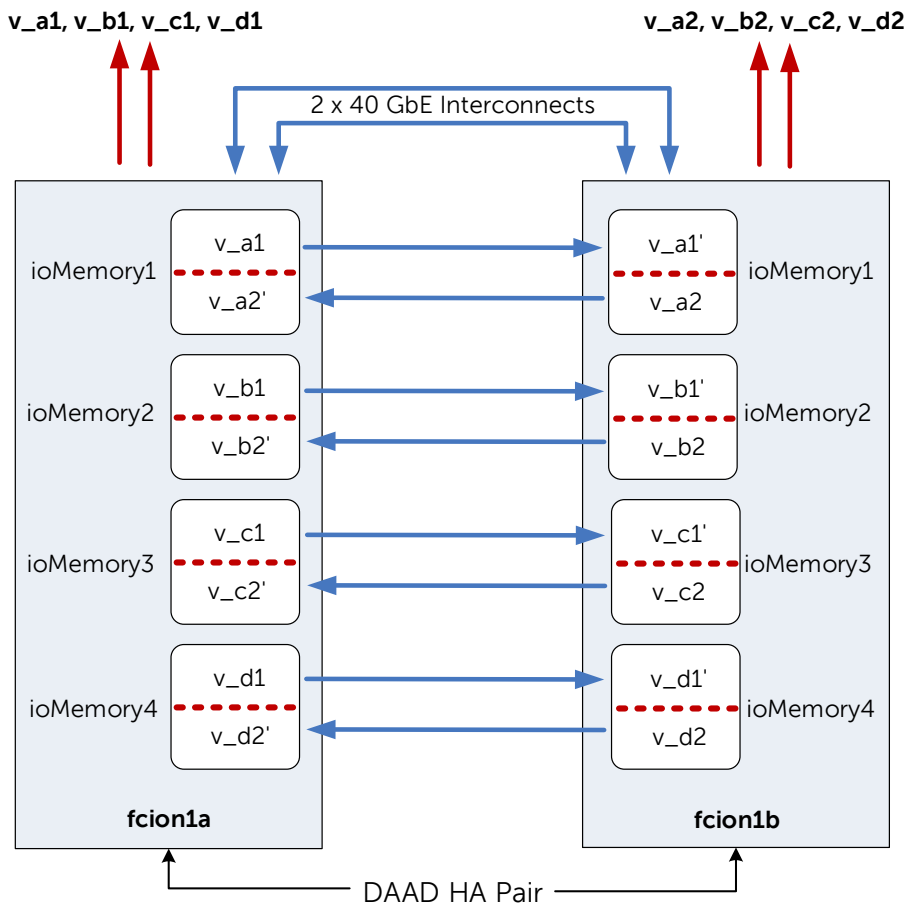


Figure 14 Storage volume setup on the nodes in a DAAD HA pair



Minimum DISOD configuration contains at least one DAAD HA pair because DISOD supports only clustered DAAD nodes. The cluster names of each DAAD HA pair and the individual hostnames of each DAAD node within that HA pair are preconfigured from factory. The first DAAD HA pair is assigned the cluster name `fcion1` and the individual DAAD nodes within that HA pair are assigned the hostnames `fcion1a` and `fcion1b`. Similarly, the second DAAD HA pair is assigned the cluster name `fcion2` and the individual DAAD nodes within that HA pair are assigned the hostnames `fcion2a` and `fcion2b`. This naming convention is followed for all the DAAD clusters or HA pairs in the DISOD solution rack.

Volumes and volume names on the DAAD nodes are preconfigured from factory. Eight volumes are created within each DAAD HA pair. These eight volumes are given the following local names:

```
v_a1
v_a2
v_b1
v_b2
v_c1
v_c2
v_d1
v_d2
```

The full name of each volume is defined as the 'cluster name + the local name'. For example, the eight volumes of the first cluster array (`fcion1`) are preconfigured as:

```
fcion1_v_a1
fcion1_v_a2
fcion1_v_b1
fcion1_v_b2
fcion1_v_c1
fcion1_v_c2
fcion1_v_d1
fcion1_v_d2
```

Eight volumes are created per DAAD HA pair and assigned to the database nodes for Oracle database usage. So, the total number of volumes preconfigured and presented to the database nodes is the number of DAAD HA pairs times eight.

Total number of volumes for Oracle database = Number of DAAD HA pairs × 8

Depending on the capacity flavor of the DAAD nodes within each HA pair, the size of the volumes assigned to the database nodes is as under:

- Each volume is of size 1.5 TB in case of the 12.8 TB-based DAAD HA pairs
- Each volume is of size 3 TB in case of the 25.6 TB-based DAAD HA pairs

For more information about how these volumes are configured on the database nodes, see the *4.2.3 Shared storage volume partitions* section.



4.2 Database nodes configuration

This section describes the OS preconfiguration that is applied on the database nodes to facilitate faster deployment of Oracle Database 12c at the customer site. The configuration is done based on the best practices of Dell and Oracle.

Database nodes are preconfigured and optimized in the following areas:

- Hardware configuration
- Oracle Database 12c preinstallation setup
- Shared storage volume partitions
- Disk permissions setup using UDEV rules
- Device-Mapper multipath
- Oracle network configuration

To view the list of Dell PowerEdge servers that are supported as the database nodes in the DISOD solution rack, see the *2.1 Database servers* section.

4.2.1 Hardware configuration

Before any software application is installed, it was ensured that the hardware configuration is adequate and optimized for Oracle Database environment. This includes setting the recommended BIOS settings for Oracle Databases, appropriate updated firmware, setting the recommended RAID configuration, and creating the required partitions for the OS and for the Oracle Database software. These configuration details are further provided below:

- System Profile and memory configuration is set to 'Performance Optimized' in the BIOS settings
- Installed firmware versions on all the hardware either matches or is newer than what was tested with
- Each database node has two 1.2 TB HDDs in RAID 1 with a global hot-spare (1x1.2TB) configured.
- Local HDD logical volume partition configuration:
 - 60 GB `\/` (root partition for OS)
 - 17 GB swap (minimum 16 GB + 1 KB required by Oracle Database)
 - Rest allocated to `/u01` to be used for Oracle Database software installation

4.2.2 Oracle Database 12c pre-installation setup

The database nodes come preinstalled with OL 6.6 running UEK release 3. To make integration of Dell hardware and Oracle Database software seamless, the preinstall RPM for Oracle Database 12c has been customized to eliminate many of the manual tasks necessary to correctly setup the environment. This utility modifies many settings, including but not limited to:

- Modify kernel parameters and boot settings per Oracle's recommendations
- Add module options as recommended by Oracle
- Create special users with specific groups and permissions required for the management of Oracle Database 12c.



To supplement the Oracle Database 12c preinstall RPM, a customized Dell deployment utility is also used. This utility builds on top of the preinstall RPM, and takes extra steps to ensure that the environment is set up properly. Some of the steps include:

- Disable SELinux (optional)
- Create additional users/groups/permissions for Oracle's Grid infrastructure
- Create base directories and permissions for the Grid/Database installation
- Install the necessary RPMs needed for Oracle Database 12c

The Oracle Database 12c preinstall RPM and the Dell deployment utility are preinstalled on the database servers in the DISOD solution rack. More information regarding these utilities can be viewed at http://en.community.dell.com/techcenter/enterprise-solutions/w/oracle_solutions/4957.oracle-database-deployment-automation

4.2.3 Shared storage volume partitions

DAAD is not a traditional spinning disk-based storage, but is an extremely fast all-flash based storage array. There are no performance implications on putting all Oracle database functions across only two ASM disk groups—one for OCR and Voting Disks (VDs), and the other for data files, Flash Recovery Area (FRA), REDO log files, and so on. As a result, the storage volumes that are preconfigured and assigned to the database nodes are prepartitioned at the OS level and presented as raw disks for these various Oracle Database functions.

Two partitions are created on each storage volume—one of 15 GB disk size and the other spanning rest of the drive. The disk size of the second partition depends on the capacity of DAAD storage array that is used in the backend. For DAAD nodes with a total disk capacity of 12.8 TB, the second partition is of size 1.485 TB; and, for DAAD nodes with a total disk capacity of 25.6 TB, the second partition is of size 2.985 TB. Thus, taking a 2×2 configuration as an example that has one DAAD HA pair, DISOD solution provides and is configured with the following shared storage volumes or partitions at the database node level:

1. Eight 15 GB partitions or disks, for a total of 120 GB, three of which are to be used for OCR and Voting Disks, and the remaining are free to be used as extra storage.
2. Eight partitions or disks of size either 1.485 TB or 2.985 TB, for a total of either 11.88 TB or 23.88 TB, that are available for data files, FRA, REDO log files, and so on.

The total quantity of each of the above two partitions depends on the total number of DAAD HA pairs that are used in the backend. To get the total quantity and the total size of usable storage capacity, multiply the above numbers by the total number of DAAD HA pairs. For example, for two DAAD HA pairs in the backend, a total of sixteen 15 GB partitions will be available, and a total of sixteen second partition that are to be used for data files, REDO log files, and so on will be available.

NOTE: Though DISOD provides at least eight partitions or disks for OCR and VDs, it is recommended to create its ASM diskgroup with Normal Redundancy. DAAD's HA feature



provides three additional mirrored copies residing on separate disks on top of the Normal Redundancy, thus providing a total of six copies.

4.2.4 Disk ownership and permissions setup using UDEV rules

The necessary and appropriate ownership and permissions that is required on the storage disks before Oracle 12c grid and database installation is preconfigured on the DISOD database nodes. These appropriate ownership and permissions are set up by using UDEV rules across all the partitions described in section *4.2.3 Shared storage volume partitions*. UDEV rules are configured in the `/etc/udev/rules.d/99-dell.rules` file.

4.2.5 Device-Mapper multipath

In order to employ high availability with external storage, this integrated system uses the Linux Device-Mapper multipath utility. The advantage of using this is that storage traffic is not only redundant, but is load-balanced across multiple connections, increasing total throughput. Multipathing is already set up by using recommended best practices of Dell, with redundancy verified before shipping to the customer. The multipath is configured in the `/etc/multipath.conf` file.

4.2.6 Oracle network configuration

The network configuration within the database nodes is designed to provide HA, high bandwidth, and load balance for both the Oracle database's public and private interconnect network traffic. Two dual port, 40 GbE Mellanox PCIe adapters are used, with one port from each adapter for the public traffic and the other port from each adapter for the private interconnect traffic. This type of design provides HA at the port and at the adapter level for both the types of Oracle Database traffic. The two ports across the two adapters are also connected to two separate ToR switches to further provide redundancy and load balance at the switch-level, and eliminate any single point of failure in the Oracle Database network.

The two public network interfaces or ports in each of the database nodes are not preconfigured. Dell recommends to configure them at the customer site by using native Oracle Linux bonding driver in active-backup configuration mode. This type of configuration ensures a failover path is available in case the active public interface loses connection and prevents any downtime.

NOTE: In the OS, the two public network interfaces are enumerated as `eth4` and `eth6`.

The two private network interfaces are not bonded at the OS level because Oracle recommends to use the Highly Available IP (HAIP) database feature which places two separate physical links on different subnets, ensuring that there are two distinct paths for data to travel on. HAIP is a feature of Oracle Grid Infrastructure installation, and requires at least two interfaces on separate private subnets to be configured. To help with faster deployment at the customer site, the two private network interfaces in all the database nodes are preconfigured in the following OS interface network scripts:



Oracle Private A: `/etc/sysconfig/network-scripts/ifcfg-eth5`
Oracle Private B: `/etc/sysconfig/network-scripts/ifcfg-eth7`

NOTE: For the preassigned IP addresses for the private interfaces in the database nodes, see the *3.1.3 Static IP addresses of Database and DAAD nodes* section.

During Oracle grid installation, select the above two preconfigured network interfaces for the private interconnect.

4.3 Management server configuration

The management server comes preinstalled with the following setup:

- Windows Server 2012 R2 Standard Edition (SE) as the base physical OS.
- Two Hyper-V VMs:
 - **WinVM:** This is a Windows-based VM that is preinstalled with Windows Server 2012 R2 SE as the guest OS. This VM is created primarily for installing management software at the customer site.
 - **LinVM:** This is a Linux-based VM that is preinstalled with Oracle Linux 6.6 running the Red Hat Compatible Kernel as the guest OS. This VM is primarily used for the configuration of the DISOD components before it is shipped out to the customer. It also provides for a Linux-based environment to remotely access any of the DISOD hardware components, if needed.

Refer to Appendix *B DISOD authentication settings* for the default authentication settings.

NOTE: For information about installation and configuration of the various management software in order to manage and monitor DISOD solution rack, see *Dell Integrated Systems for Oracle Database – Oracle Enterprise Manager Integration Guide*

4.3.1 Base physical OS configuration settings

The base OS on the management server is preconfigured with the following settings:

- Windows Server 2012 R2 64-bit is preinstalled with Hyper-V feature enabled.
- OS partition is set to the maximum HDD available disk space of 497 GB.
- All firewalls (Public, Private, and Domain) are disabled.
- Computer name is set to:

`dMan-<ServiceTag>`

where, `<ServiceTag>` is the unique Service Tag of the management server and `dMan` indicates DISOD Management.

- Two Hyper-V VMs with names WinVM and LinVM are created.
- For network configuration details, see the *3.1.2 Management server network configuration* section.



4.3.2 WinVM configuration settings

WinVM VM on the management server is preconfigured with the following settings:

- As a 'generation 2' VM
- Windows Server 2012 R2 64-bit guest OS
- Four virtual processors and 8 GB dynamic memory
- All firewalls (Public, Private, and Domain) disabled.
- Computer name set to:

WinVM-*<ServiceTag>*

Where, *<ServiceTag>* is the unique Service Tag of the management server

- For network configuration details, see the *3.1.2 Management server network configuration* section.

The primary purpose of this VM is to install and configure the necessary DISOD management software in order to integrate with Oracle Enterprise Manager (OEM). This provides the ability to monitor and manage the DISOD components from within a single OEM console. For information about installation and configuration of various management software, see the *Dell Integrated Systems for Oracle Database - Oracle Enterprise Manager Integration Guide* available at Dell.com/support/home.

NOTE: The virtual memory size is allocated on the basis of requirement for the base DISOD configuration. Customers are recommended to observe the memory usage and change the dynamic memory allocation, as necessary, for their DISOD configuration.

4.3.3 LinVM Configuration Settings

LinVM VM on the management server is preconfigured with the following settings:

- As a 'generation 1' VM.
- Oracle Linux 6.6 running Red Hat Compatible Kernel as guest OS.
- Virtual HDD size of 250 GB.
- Four virtual processors and 32 GB dynamic memory
- Hostname set to:

LinVM-*<ServiceTag>*

Where, *<ServiceTag>* is the unique Service Tag of the management server

- For network configuration details, see the *3.1.2 Management server network configuration* section.
- As an isolated and contained DHCP and TFTP Server:
 - To ensure that the preassigned IP addresses for the private rack network components, for example, iDRAC and private management port on the servers, and management ports on the switches is static and consistent across reboots. This is achieved by preconfiguring



the `/etc/dhcpd/dhcpd.conf` file with the MAC addresses of these network ports.

- To allow easy upgrade of the firmware, or to apply preconfigured switch settings on the S6000 and the S3048 switches, which can be achieved by booting the switches in `jumpstart` mode.
- DHCP is configured to use the DISOD private network so that it does not cause any conflict with the customer's own DHCP server. DHCP on LinVM is configured in the `/etc/sysconfig/dhcpd` file to listen only on the private management adapter `eth0` to limit the DHCP functionality to the DISOD private network.

NOTE: The virtual memory size is allocated based on the requirement for the base DISOD configuration. Customers are recommended to observe the memory usage and change the dynamic memory allocation, as required, for their DISOD configuration.



5 DISOD management and monitoring

DISOD solution rack has a built-in dedicated management infrastructure that consists of the following hardware components:

- Management switch (S3048)
- Management server (R320)
- iDRAC ports on the database servers and the DAAD nodes for GUI-based server accessibility
- Dedicated management port on the database servers (LOM port 1), on the DAAD nodes (LOM port 1), and on the switches for CLI-based hardware accessibility

Using the preconfigured management infrastructure, customers and Dell Services or Dell Support can remotely access, manage, and monitor any of the hardware components within the DISOD solution rack.

Figure 15 shows, as an example, the precabled management infrastructure of a DISOD 2x2 solution.



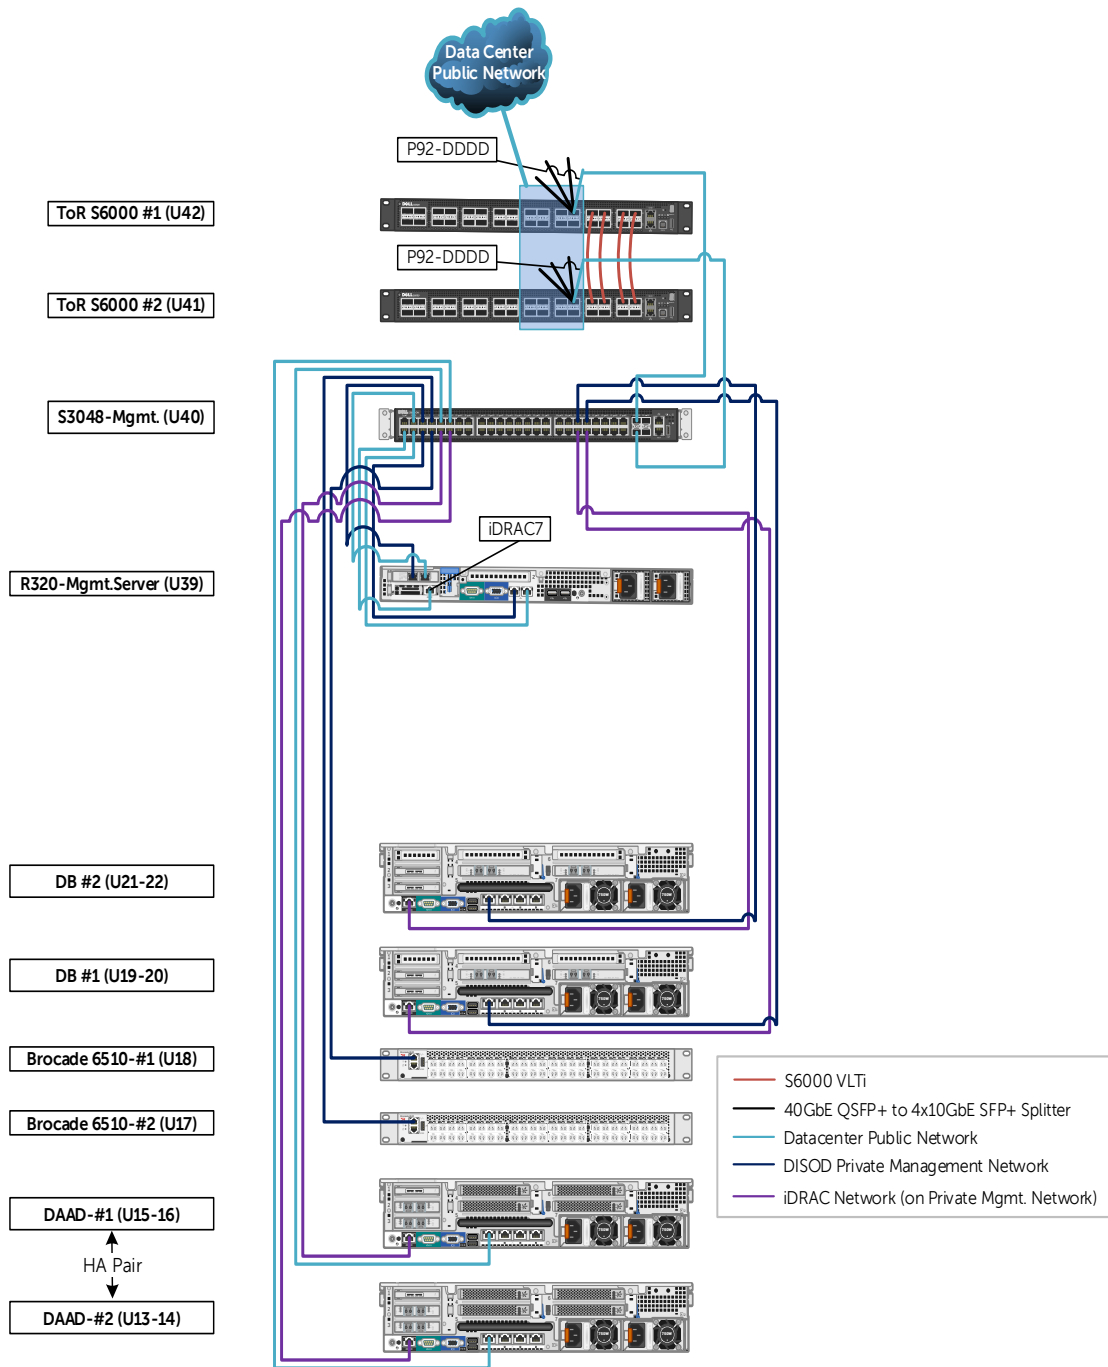


Figure 15 DISOD Management Infrastructure with cabling

The two ToR S6000 switches are connected to the S3048 management switch by using redundant 10 GbE network ports. The management switch, in turn, is connected to two different networks on the management server—redundant 1 GbE public network and redundant 1 GbE private network. Thus, customers can connect remotely to the management server by using the public network that is isolated on VLAN 10 network. After logging in to the management server, any of the DISOD hardware components can be accessed for management and monitoring by using the private network that is isolated on



VLAN 30 network. The DISOD solution thereby provides a secure and segregated access and management solution.

Users can log in to the management server:

- By using a browser on their client-machine to log in to the management server's iDRAC, or
- By using a Linux terminal or an application such as Putty.exe on their client-machine to SSH in to LinVM's public IP address for command line interface (CLI) access, or
- By using Windows Remote Desktop Protocol on their client-machine and connecting to the management server's public IP address.

NOTE: The public IP addresses must be configured first for the above network ports on the management server before users can connect to them remotely.

After remotely logged in to the management server, the user can connect to any of the DISOD components with any of the options listed in Table 12.

Table 12 DISOD components remote management connectivity options

DISOD Component	Graphical (via Browser)	CLI ⁴
Database nodes ⁵	iDRAC	SSH - Private Mgmt. IP
DAAD storage nodes	iDRAC	SSH - Mgmt. IP
Brocade 6510 Switches	N/A	SSH - Mgmt. IP
S6000 Switches	N/A	SSH/Telnet - Mgmt. IP
S3048 Switch	N/A	SSH/Telnet - Mgmt. IP

For the pre-assigned private IP addresses for each of the DISOD hardware components on the private management network, see the *3.1 DISOD private management network* section. The default authentication settings for the DISOD components can be found in the Appendix section *B DISOD authentication settings*.

Though the above method is an available option to log in directly to each of the DISOD components, and manage and monitor them individually, the DISOD management solution goes one step further. The DISOD network has been set up to integrate seamlessly with Oracle Enterprise Manager (OEM) in order to provide DBAs a single console of management and monitoring of the DISOD solution. With the DISOD management solution, the DBAs can not only be able to manage and monitor their Oracle databases, but also manage and monitor the health, inventory, and any alerts generated by

⁴ It is recommended to download and install an application such as Putty to SSH or Telnet in to a DISOD component.

⁵ Though database nodes can be accessed directly by using its Oracle public network ports, it is recommended not to connect to it for management or monitoring purposes. Oracle public network should be used for database traffic only.



the DISOD hardware components, all from a single OEM console. To completely set up the DISOD management system, see the *'Dell Integrated Systems for Oracle Databases Oracle Enterprise Manager Integration Guide* at Dell.comsupport/home.



A DISOD rack hardware cabling diagrams

DISOD components are precabled before they are shipped to the customer. For reference and as an example, this section provides the rack cabling diagram only for a 4x4 R730-based FC DISOD solution, and for a 2x2 R930-based FC DISOD solution. Users can extrapolate the cabling methodology for other DISOD configurations based on the logical design explained in section 3 *DISOD network configuration*. The diagrams in this section are divided on the basis of the following functionalities:

- Cabling diagram - LAN infrastructure
- Cabling diagram - SAN infrastructure
- Cabling diagram - management infrastructure

A.1 Cabling diagram - LAN infrastructure

This section provides the cabling diagram of the hardware components that are part of the LAN infrastructure used for public and private interconnect traffic of Oracle Database.

Figure 16 shows the LAN infrastructure cabling diagram of a 4x4 R730-based FC DISOD solution.

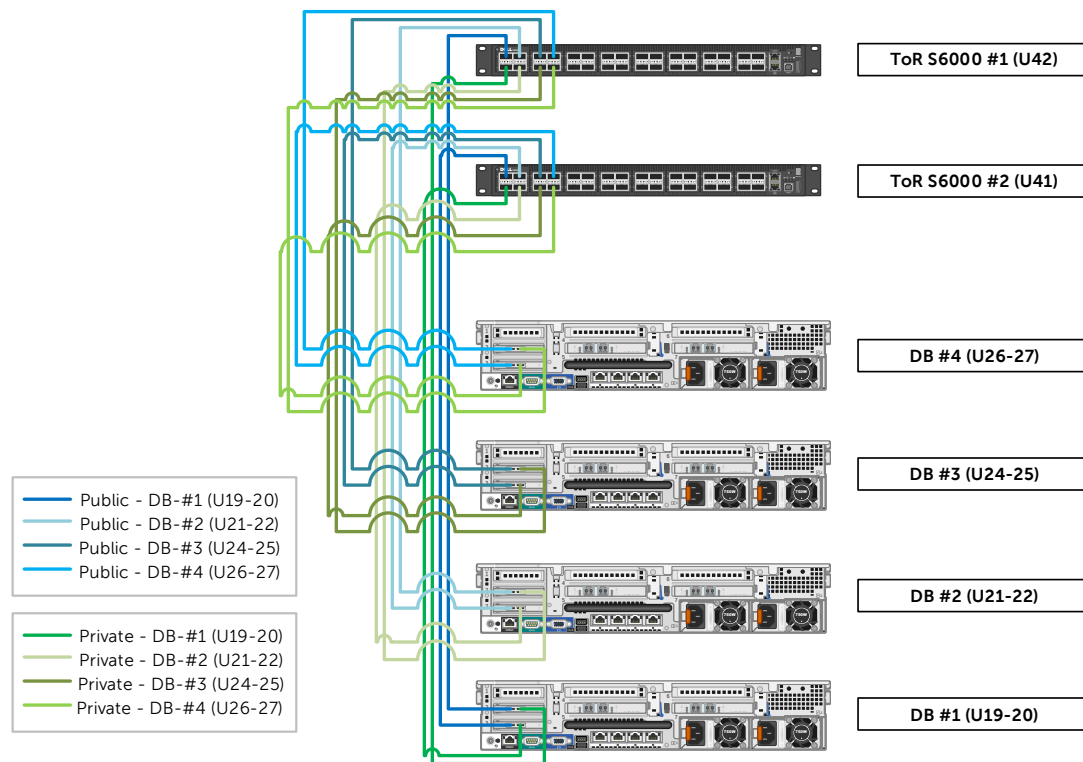


Figure 16 Cabling diagram – LAN infrastructure – 4x4 R730-based FC DISOD solution

Figure 17 shows the LAN infrastructure cabling diagram of a 2x2 R930-based FC DISOD solution.

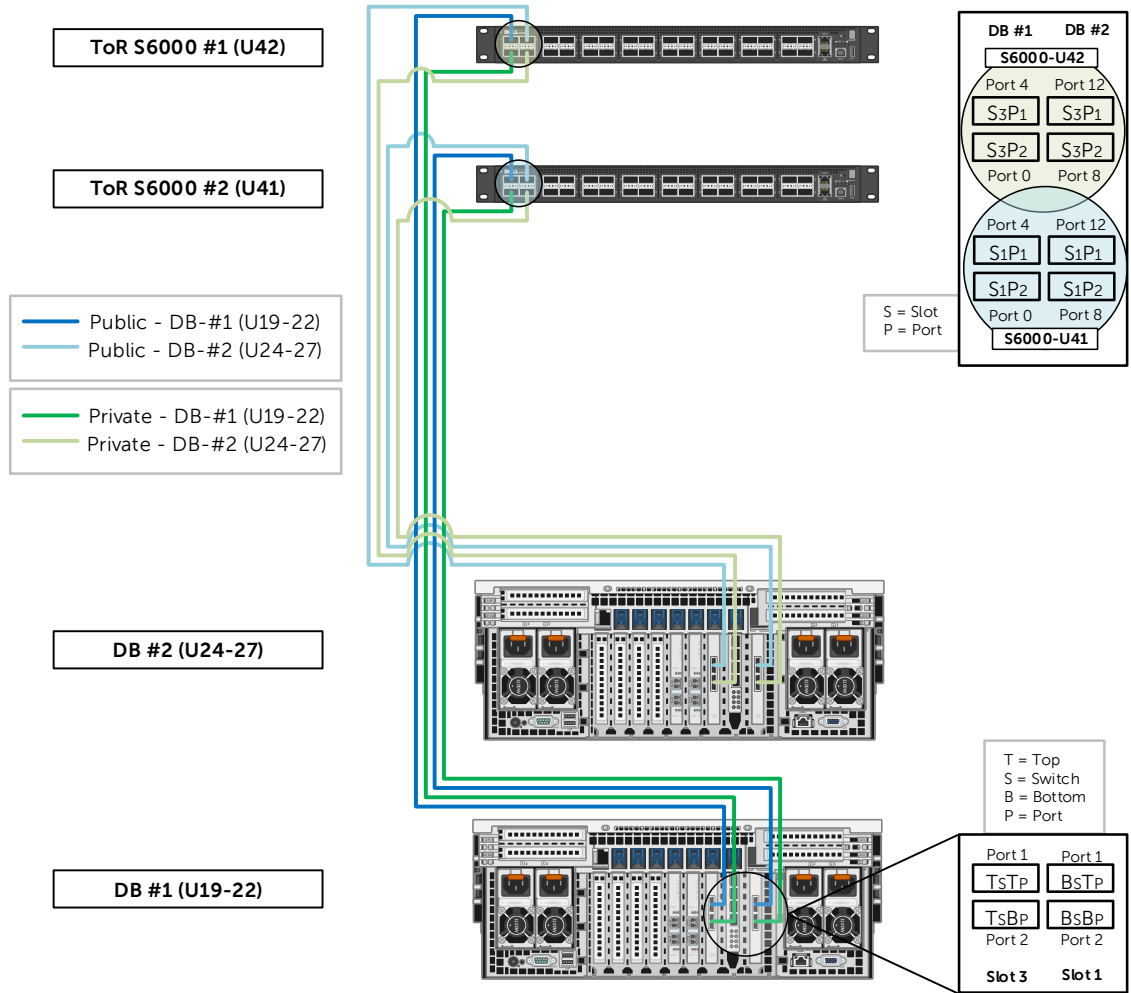


Figure 17 Cabling diagram – LAN infrastructure – 2x2 R930-based FC DISOD solution



A.2 Cabling diagram - SAN infrastructure

This section provides the cabling diagram of the hardware components that are part of the SAN infrastructure, showing connectivity between the database nodes and the SAN switches, and between the DAAD nodes and the SAN switches.

Figure 18 shows the SAN infrastructure cabling diagram of a 4x4 R730-based FC DISOD solution.

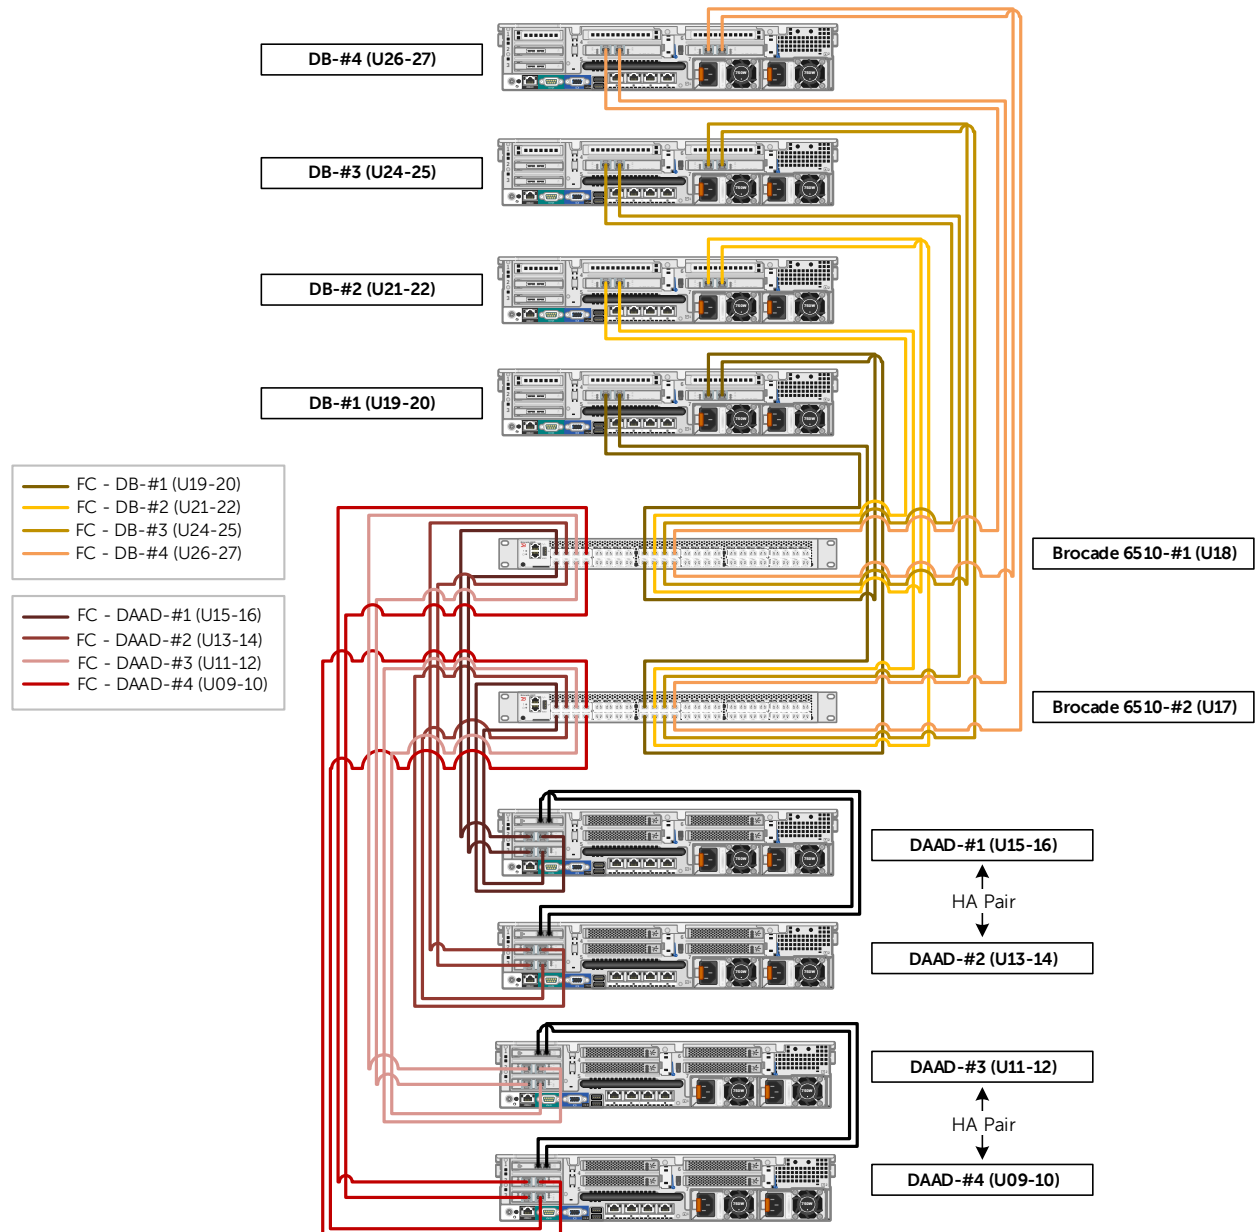


Figure 18 Cabling diagram – SAN infrastructure – 4x4 R730-based FC DISOD solution

Figure 19 shows the SAN infrastructure cabling diagram of a 2x2 R930-based FC DISOD solution.




Figure 19 Cabling diagram – SAN infrastructure – 2x2 R930-based FC DISOD solution

A.3 Cabling diagram - management infrastructure

This section provides the cabling diagram of the hardware components that are part of the management infrastructure, showing the connectivity between the various management ports and the management switches.

Figure 20 shows the management infrastructure cabling diagram of a 4x4 R730-based FC DISOD solution.

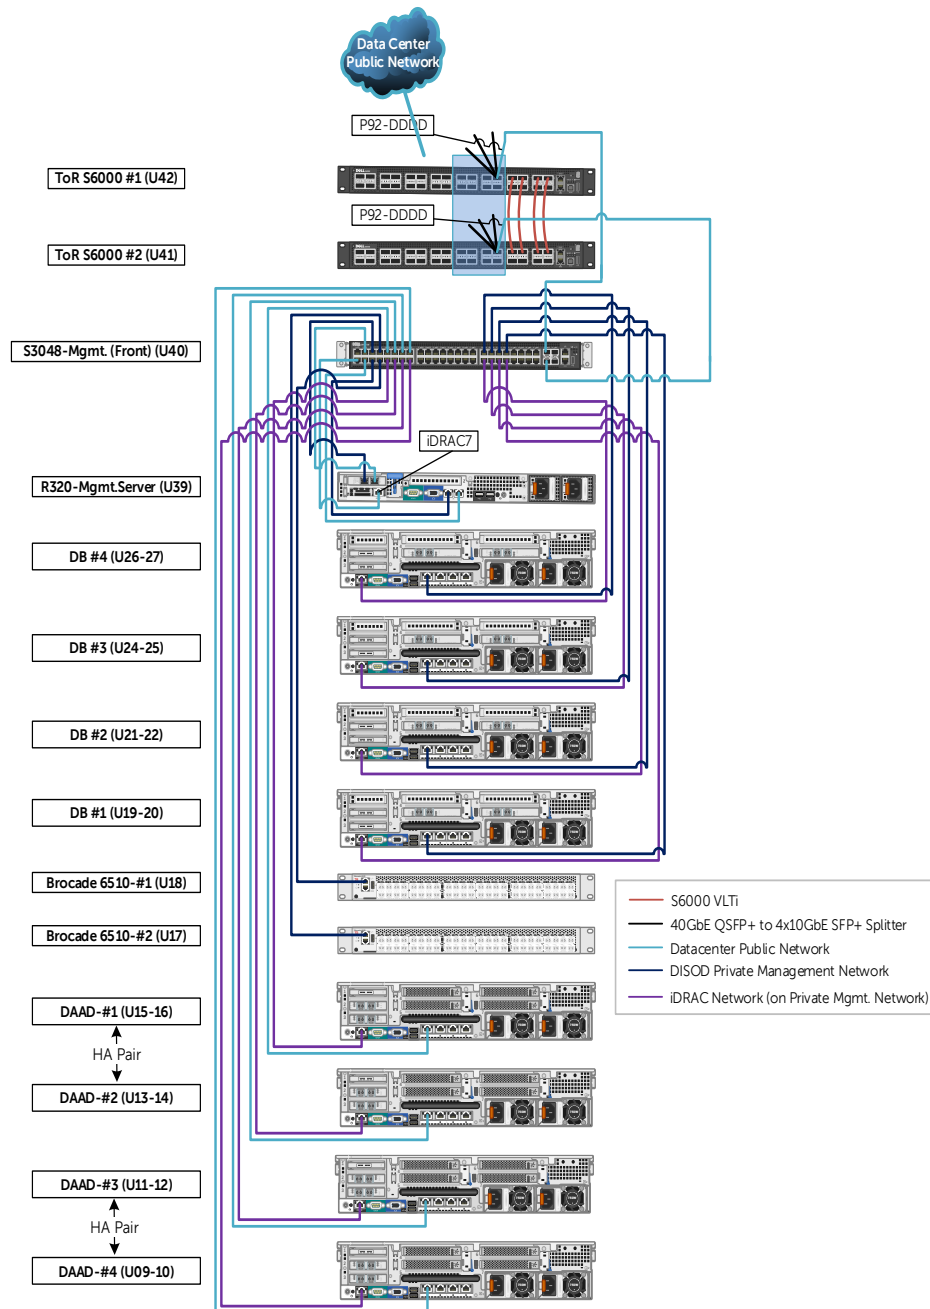


Figure 20 Cabling diagram – management infrastructure – 4x4 R730-based FC DISOD solution

Figure 21 shows the management infrastructure cabling diagram of a 2x2 R930-based FC DISOD solution.

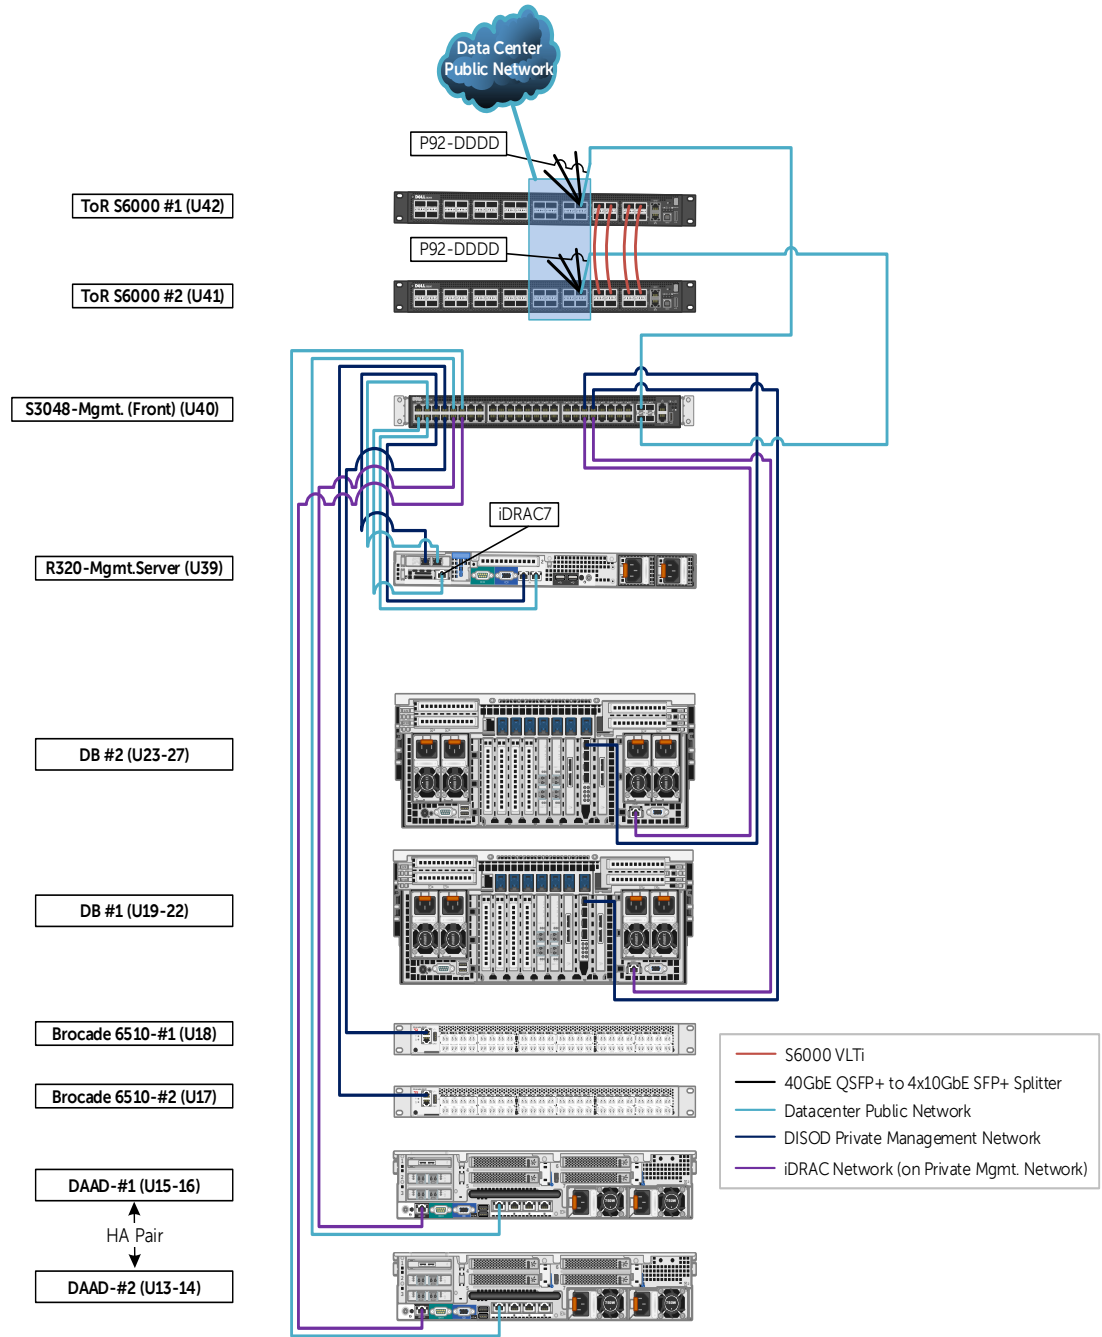


Figure 21 Cabling diagram – management infrastructure – 2x2 R930-based FC DISOD solution



B DISOD authentication settings

The hardware components in the DISOD solution rack ship with the following default authentication settings:

- `root/oracle`: Database nodes and LinVM
- `Administrator/Oracle@123`: Management server base OS and WinVM
- `admin/oracle`: DAAD storage array nodes
- `admin/password`: Management switch (S3048-ON), Top-of-Rack S6000 switches, and the SAN Brocade 6510 switches
- `root/calvin`: iDRAC

It is highly recommended to change these default settings before connecting the DISOD solution rack to the customer's network environment.

