Dell Integrated Systems for Oracle Business Analytics - Owner's Guide v1.0

Dell Oracle Solutions Engineering
December 2014

Revisions

Date	Description
December 2014	Initial release 1.0

Table of contents

1 Overview - Dell Integrated Systems for Oracle Business Analytics (DISOBA)

Dell Integrated Systems for Oracle Business Analytics (DISOBA) is a fully integrated hardware stack that is purpose built as a high performance Oracle Business analytical solution by integrating Oracle OBIEE and Oracle Database 12c. The following figure shows the 1.0 release of DISOBA that is a 21 rack-units solution.

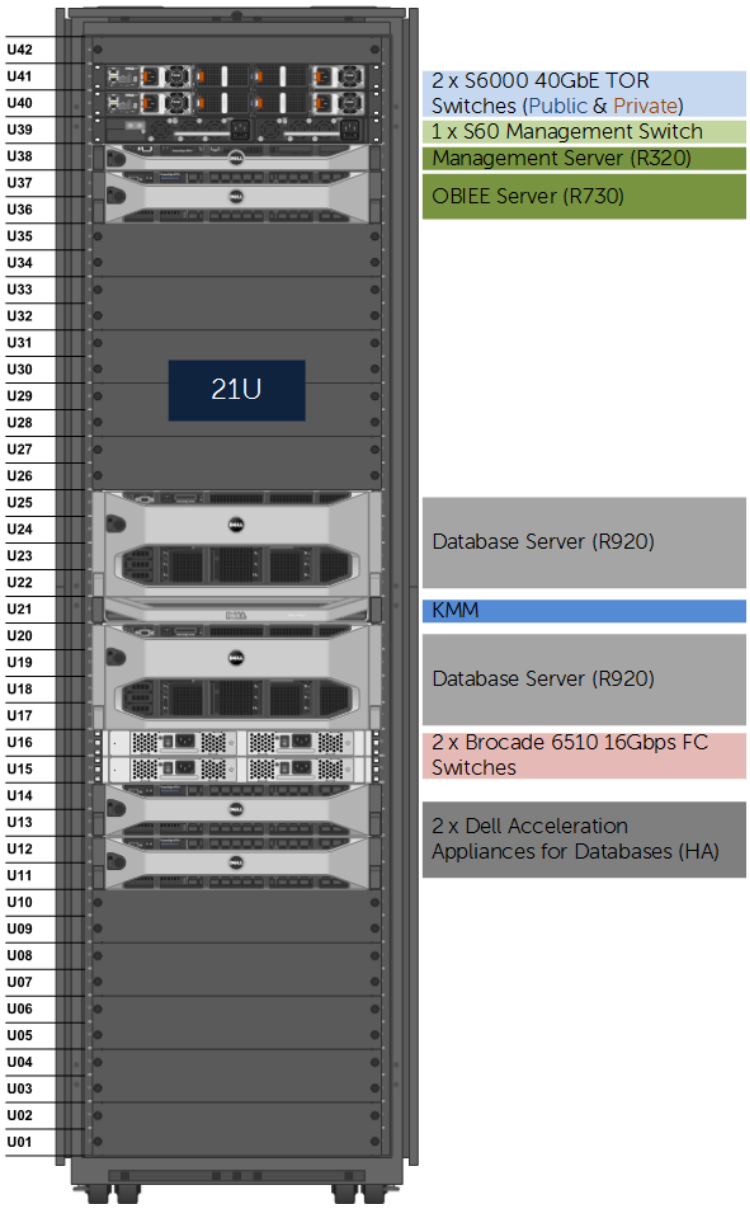

Figure 1 Dell Integrated Systems for Oracle Business Analytics (DISOBA) Hardware Rack

DISOBA provides an out-of-box experience to customers where everything up to the point of the Oracle Business analytics and Database software installation is pre-configured. DISOBA ships with the following pre-integrated, pre-configured, racked, stacked and cabled hardware and software:

- **Business Analytics (OBIEE):** One Dell PowerEdge R730 server pre-installed and configured with Oracle Linux 6.5 running Unbreakable Enterprise Kernel operating system.
- **Database Servers**: Two Dell PowerEdge R920 pre-installed and configured with Oracle Linux 6.5 running Unbreakable Enterprise Kernel operating system.
- **Shared Storage**: Two high performance Dell Acceleration Appliances for Databases (DAAD) storage array pre-configured as a highly available Fibre Channel shared storage.
- **Network Switches**: Two pre-configured high bandwidth Dell Networking S6000 40Gb Ethernet switches used for public and private network traffic.
- **SAN Switches**: Two pre-configured Brocade 6510 16Gbps Fibre Channel switches used for high-speed SAN connectivity.
- **Management**: One Dell PowerEdge R320 server and one Dell Networking S60 1Gb Ethernet switch providing access to the private network inside the rack solution to manage and monitor all the components.

Figure 2 show the logical architectural diagram of DISOBA

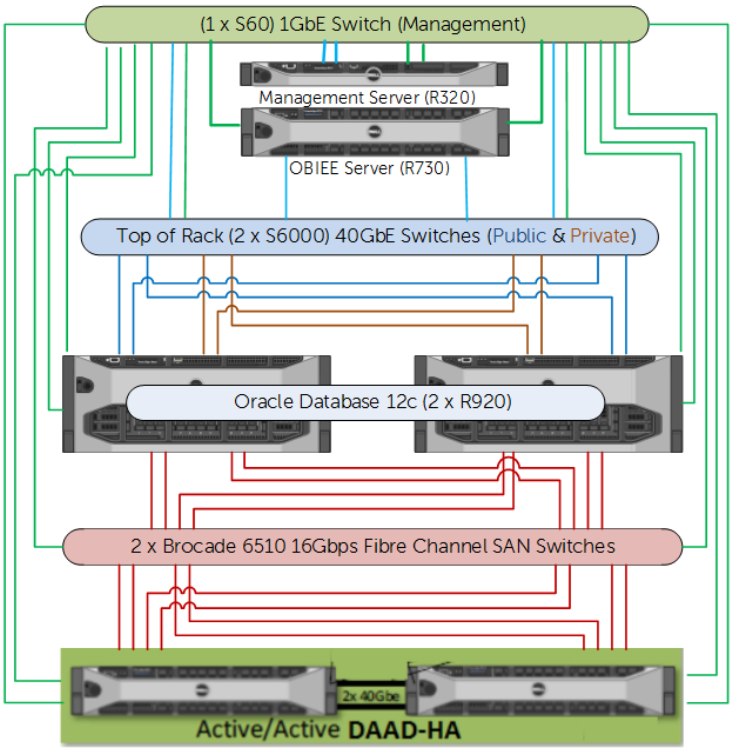

Figure 2 Architectural Diagram of DISOBA

By providing integrated and pre-configured hardware and software solutions, DISOBA provides a fast and easy-to-deploy Oracle Business Analytics integrated with Oracle Database rack solution to the customers.

2 DISOBA Hardware Components

This section provides the hardware details of each of the DISOBA components

2.1 Business Analytics (OBIEE) (R730)

The DISOBA solution rack ships with one Dell PowerEdge R730 server that is used as OBIEE server. The following table shows the details of the hardware components within each of these servers:

Table 1 OBIEE Server (R730) Hardware Details

Database Nodes Hardware Components			
Servers	UP to 1 x 2S Dell PowerEdge R730		
CONFIG	SMALL Config	MEDIUM Config	LARGE Config
Processor	1 x E5-2637 v3 3.5GHz	2 x E5-2637 v3 3.5GHz 4 core (or) 2 x E5-2643 v3 3.4GHz 6core	2 X E5-2667 V3 3.2GHz 8 core (or) 2 x E5-2687W v3 3.1GHz 10 core
Memory	4 x 16G RDIMM @2133 mhz	4 x 16G RDIMM @2133 mhz	8 x 16G RDIMM @2133 mhz
Network Adapters	2 x Mellanox ConnectX-3 DP 10Gb Ethernet Adapters per server		
Network Daughter Card (NDC)	Broadcom 5720 Quad Port 1GbE NDC		
Remote Management	iDRAC7 Enterprise		
RAID Controller	PERC H730P (RAID 1)		
Local Storage	4 x 600GB 15K 6Gbps SAS 2.5in Hot-plug HDDs per server		
Operating System	Oracle Linux UEK Rel 3 - 3.8.13-44.1.1.el6uek.x86_64		

Figure 4 shows the back panel of the R730 chassis with the adapters populated according to Dell's best practice for optimum performance.

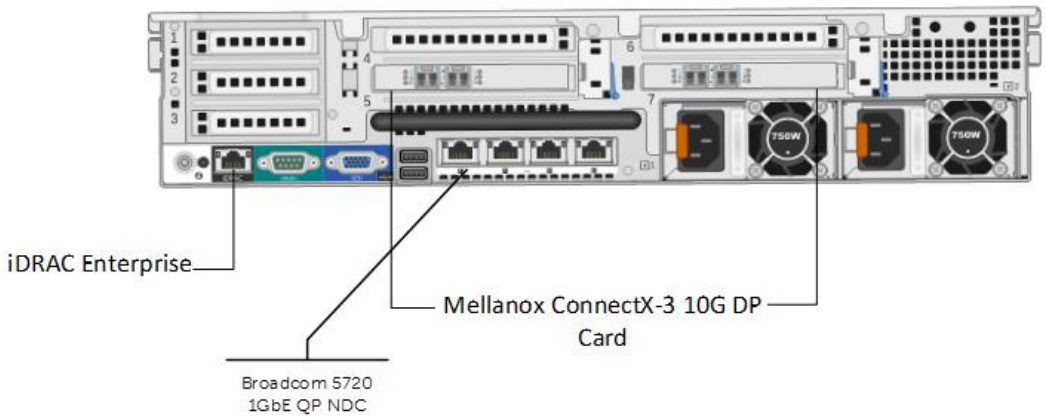

Figure 3 OBIEE Server (R730) PCIe Slot Population

The R730 OBIEE servers is pre-installed and pre-configured with Oracle Linux running Unbreakable Enterprise Kernel. For more details on the OS configuration, refer to section OBIEE Server (R730) Configuration

2.2 Database Servers (R920)

The DISOBA solution rack ships with two Dell PowerEdge R920 servers that are used as Oracle Database 12c nodes. Table 2 shows the details of the hardware components within each of these servers:

Database Nodes Hardware Components	
Servers	2 x 4S Dell PowerEdge R920s
Processor	Up to 4 x Intel Xeon E7-8891 v2 10c 3.2GHz per server
Memory	Up to 2TB using 32GB DDR3 1600MT/s QR LRDIMMs per server
Network Adapters (Oracle Database Public & Private)	2 x Mellanox ConnectX-3 EN DP 40Gb Ethernet Adapters per server
Network Daughter Card (NDC)	Broadcom 5720 Quad Port 1GbE NDC
Remote Management	iDRAC7 Enterprise
HBAs	2 x QLogic 2662 DP 16GB FC Adapters per server
RAID Controller	PERC H730P (RAID 1)
Local Storage	2 x 900GB 10K 6Gbps SAS 2.5in Hot-plug HDDs per server
Operating System	Oracle Linux 6.5 running Unbreakable Enterprise Kernel (UEK) Release 3 Update 3 (3.8.13-44)

Table 2 Database Server (R920) Hardware Details

The following figure shows the rear side of the R920 chassis with the adapters populated according to Dell's best practice for optimum performance.

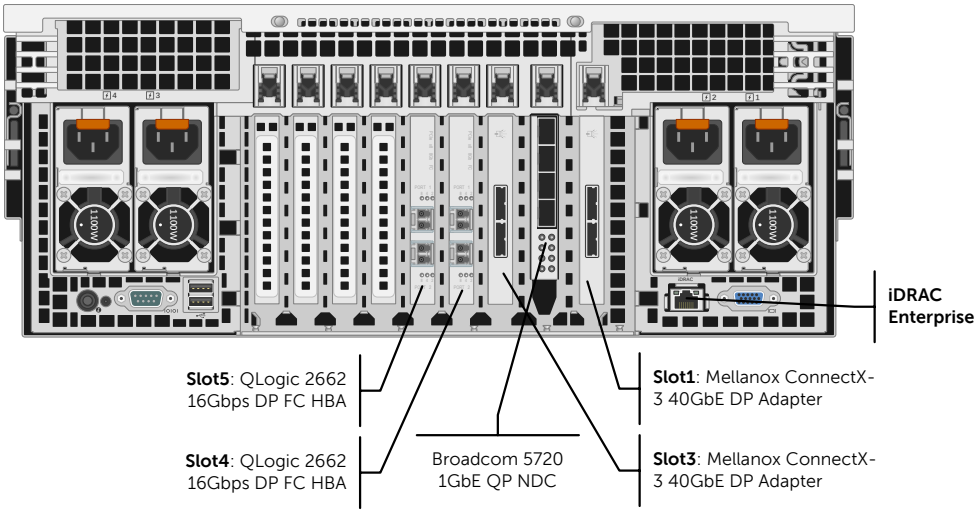

Figure 4 Database Server (R920) PCIe Slot Population

The R920 database servers are pre-installed and pre-configured with Oracle Linux running Unbreakable Enterprise Kernel. For more details on the OS configuration, refer to section Database Nodes (R920) Configuration

2.3 Shared Storage – Dell Acceleration Appliances for Databases (DAAD)

DISOD ships with Dell Acceleration Appliances for Databases (DAAD) storage array as the high performance, low-latency backend shared storage. The following table lists the DAAD storage array components.

DAAD Storage Array Components	
Storage Array Nodes	2 x Dell PowerEdge R720 Servers (HA Mode)
Processors	2 x Intel Xeon E5-2690 v2 per server
Memory	256GB per server
Solid State Disks	4 x Fusion-io ioDrive2 3.0TB MLC Adapters per server
Network Daughter Card (NDC)	Broadcom 5720 Quad Port 1GbE LOM
Remote Management	iDRAC7 Enterprise
HBAs	2 x QLogic 2662 DP 16Gb FC Adapters per server
Ethernet Adapter (HA)	Mellanox ConnectX-3 DP 40GbE Adapter per server
Total Raw Capacity	12TB
Acceleration Software	Fusion-io ION Acceleration Software v2.4.1 (w/ HA option installed)

Table 1 DAAD Storage Array Components

As seen in Figure 5, DISOD consists of two Fibre Channel based DAAD storage array nodes that are connected in a highly available configuration.

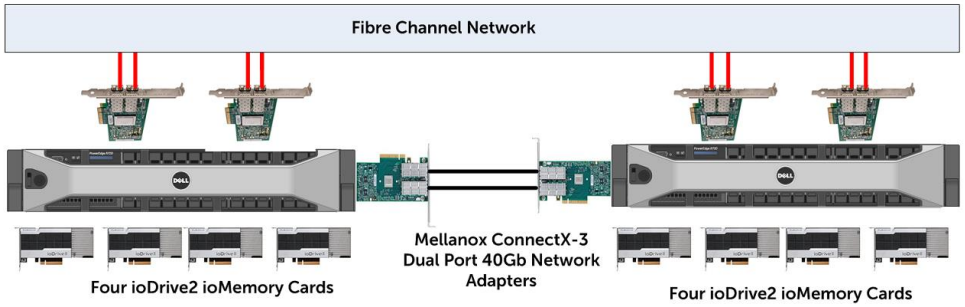

Figure 5 DISOD Shared Storage: Highly Available FC based DAAD Storage Array

For more details on the DAAD storage array configuration for DISOD refer to the section DAAD Storage Array Configuration.

2.4 Network Switches

DISOD solution rack is configured with high bandwidth, low latency, and highly available LAN and SAN switches. The solution also includes a Gigabit switch for management of all

the hardware in the rack. Table 2 shows the model of switches that are used for all the traffic that flows in and out, and within the DISOD Solution rack.

DISOD Network Switches	
LAN Switches (ToR)	2 x Dell Networking S6000 40Gb Ethernet Switches
SAN Switches	2 x Brocade 6510 16GB Fibre Channel Switches
Management Switch	1 x Dell Networking S60 1Gb Ethernet Switch

Table 2 DISOD Network Switches

The following sub-sections describe these switches in more detail.

2.4.1 Dell Networking S6000 Ethernet Switches

Redundant high bandwidth Dell Networking S6000 40GbE switches are setup as Top-of-Rack switches. In order to provide flexibility, the S6000s are pre-configured to provide both 10GbE and 40GbE ports to plug the DISOD in to 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:** Provide private network for the Oracle Real Application Cluster (RAC) database's private interconnect traffic.
- **Remote Management:** Provides connectivity to DISOD's dedicated private management network (S60 switch) for customers and Dell Services/Support to remotely be able to 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 pre-configurations:

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

For more details on the S6000 switch configurations refer to S6000 Switch Configurations.

2.4.2 Brocade 6510 Fibre Channel Switches

In order to sustain the high throughput and performance delivered by the DAAD storage arrays, redundant high bandwidth Brocade 6510 16Gbps Fibre Channel switches are used to connect the database nodes to the shared storage. These switches ship with pre-configured port-based zoning in order to provide out-of-box access to the storage volumes from the database nodes.

Using one of the Fibre Channel switches as an example, the following figure shows the pre-defined port connectivity configuration between 6510 and the database nodes, and between 6510 and the DAAD storage array nodes. The two database nodes are referred to as DB-N1 and DB-N2, and the two DAAD storage array nodes are referred to as DAAD1 and DAAD2. Port range 0-7 are used for connecting to the DAAD storage array nodes and port range 8-15 are used to connect to the database nodes. The remainder of the ports, from 16 through 47, are unconfigured.

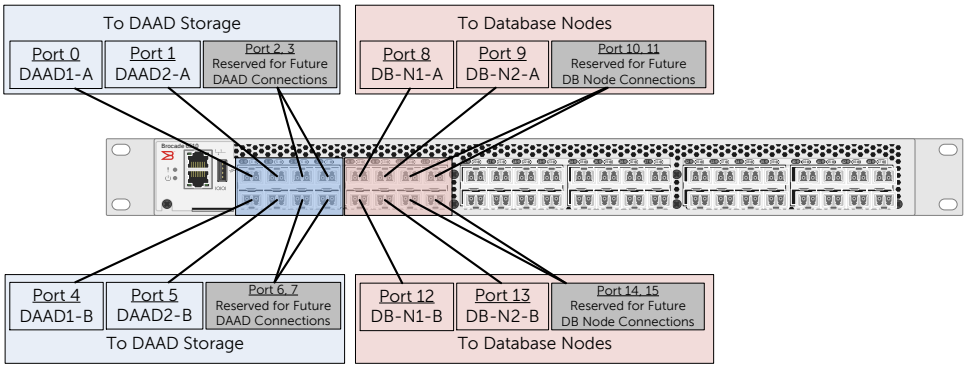

Figure 6 Brocade 6510 Fibre Channel Switch Port Connectivity

Each of the 6510 switches is configured with port-based zoning. As part of this, each initiator port to the database nodes is in its own zone with all the ports from the DAAD storage array. Thus, in the DISOD solution each 6510 is pre-configured with eight zones, four in that are reserved for future database node expansion.

2.4.3 Dell Networking S60 Switch

The Dell Networking S60 1GbE switch is used as the management switch. The ports on the switch are pre-configured and pre-cabled 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.

For more configuration details refer to S60 Switch Configurations.

2.5 Management Server (R320)

Dell PowerEdge R320 server is setup as the management server that allows remote management and monitoring of all the DISOD hardware components. Table 3 shows all the components within the management server.

Management Server Components	
Server	Dell PowerEdge R320 Server
Processor	1 x Intel Xeon E5-2430L v2 6c 2.4GHz
Memory	48 GB (6 x 8GB 1600MT/s RDIMMs)
Network Adapters	1 x Broadcom 5720 DP 1GbE Adapter 1 x Broadcom 5720 DP 1GbE LOM

Remote Management	iDRAC7 Enterprise
RAID Controller	PERC H310 (RAID 1)
Local Storage	2 x 500 GB 7.2K 3Gbps SAS 3.5in Hot-plug Hard Drives Per Server
Operating System	Oracle Linux 6.5 running Red Hat Compatible Kernel

Table 3 Management Server (R320) Components

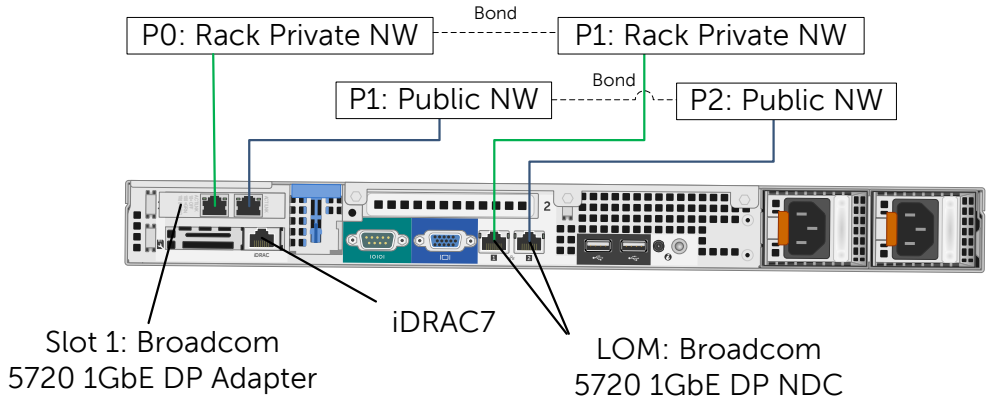

Figure 7 Management Server Network Connectivity

As seen in the above figure, management server provides two redundant and highly available 1Gb network ports that allows:

- Connectivity to the public network to allow user to access the management server remotely, and
- Connectivity to the private rack network for management of other hardware components in the rack.

3 DISOBA Network Configuration

This section provides the hardware and software networking details that DISOBA ships with. It provides details on the following network configuration:

- Private rack network setup
- S6000 switch configurations
- S60 swtich configurations
- Management and monitoring network

3.1 Private Rack Network Setup

In order to provide an out-of-box experience, DISOBA is built with a self-contained private rack network that is pre-configured from the Dell factory before the solution is shipped out. This helps to take out a lot of the hardware and software network configuration effort needed to be done at the customer site. This is achieved by pre-defining the private IP addresses for all the hardware components in the solution rack. Proper care was taken while designing the private rack network to ensure that it is isolated and will not interfere with the customer's public environment.

Table 4 shows the pre-assigned private IP addresses for the various DISOBA hardware components. End users should refer to Table 4 to get the specific IP addresses to be able to remotely connect to the needed DISOBA hardware component via the management server. The name of the hardware component listed in Table 4 starts with the rack unit number that that component is placed in, in the rack. For example, U41-S6000 refers to the S6000 switch that is located at rack unit #41 in the DISOBA solution rack.

NOTE: All the customer network facing adapters - the Oracle public network adapters on the two database nodes, the public adapters (bonded) and the iDRAC network port on the management server, are not pre-configured with an IP address. This is due to the dependency on the customer's network policy. Hence, these adapters need to be configured at the customer's site.

	DISOBA Component	Port	IP Address	Netmask	Network Address	Subnet Class
Switches	U41-S6000	Management	192.168.211.1	255.255.255.224	192.168.211.0	/27
	U40-S6000	Management	192.168.211.2	255.255.255.224	192.168.211.0	/27
	U39-S60	Management	192.168.211.3	255.255.255.224	192.168.211.0	/27
	U16-6510	Management	192.168.211.4	255.255.255.224	192.168.211.0	/27
	U15-6510	Management	192.168.211.5	255.255.255.224	192.168.211.0	/27
	U38-R320 Management Server	iDRAC	Not Configured – to be configured at customer site			
		Private (bond0)	192.168.211.20	255.255.255.224	192.168.211.0	/27
		Public (bond1)	Not Configured – to be configured at customer site			
	U36-R730 OBIEE Server	iDRAC	192.168.211.9	255.255.255.224	192.168.211.0	/27
		Private Mgmt. (em1)*	192.168.211.19	255.255.255.224	192.168.211.0	/27
		Public (bond0)	Not Configured – to be configured at customer site			
Database Nodes	U22-25-DB-N2	iDRAC	192.168.211.11	255.255.255.224	192.168.211.0	/27
		Private Mgmt. (em1)*	192.168.211.21	255.255.255.224	192.168.211.0	/27
		Oracle Private A (eth5)	192.168.211.33	255.255.255.248	192.168.211.32	/29
		Oracle Private B (eth7)	192.168.211.41	255.255.255.248	192.168.211.40	/29
		Oracle Public (bond0)	Not Configured – to be configured at customer site			
	U17-20-DB-N1	iDRAC	192.168.211.12	255.255.255.224	192.168.211.0	/27
		Private Mgmt. (em1)*	192.168.211.22	255.255.255.224	192.168.211.0	/27
		Oracle Private A (eth5)	192.168.211.34	255.255.255.248	192.168.211.32	/29
		Oracle Private B (eth7)	192.168.211.42	255.255.255.248	192.168.211.40	/29
		Oracle Public (bond0)	Not Configured – to be configured at customer site			
DAAD Storage Array	U13-14-DAAD1A	iDRAC	192.168.211.15	255.255.255.224	192.168.211.0	/27
		Private Mgmt. (eth0)*	192.168.211.25	255.255.255.224	192.168.211.0	/27
		HA 1	192.168.211.49	255.255.255.252	192.168.211.48	/30
		HA 2	192.168.211.53	255.255.255.252	192.168.211.52	/30
	U11-12-DAAD1B	iDRAC	192.168.211.16	255.255.255.224	192.168.211.0	/27
		Private Mgmt. (eth0)*	192.168.211.26	255.255.255.224	192.168.211.0	/27
		HA 1	192.168.211.50	255.255.255.252	192.168.211.48	/30
		HA 2	192.168.211.54	255.255.255.252	192.168.211.52	/30

* - Denotes the first NIC port on the LOM that is used for (remote) private management

Table 4 DISOBA Private Rack Network Pre-Assigned IP Addresses

3.2 S6000 Switch Configurations

This section provides the details of the configuration that the S6000 switches ship with.

3.2.1 S6000 Pre-defined Port Functions

Using one of the S6000 switches as an example, the following figure shows the details of the pre-defined port functions and configuration.

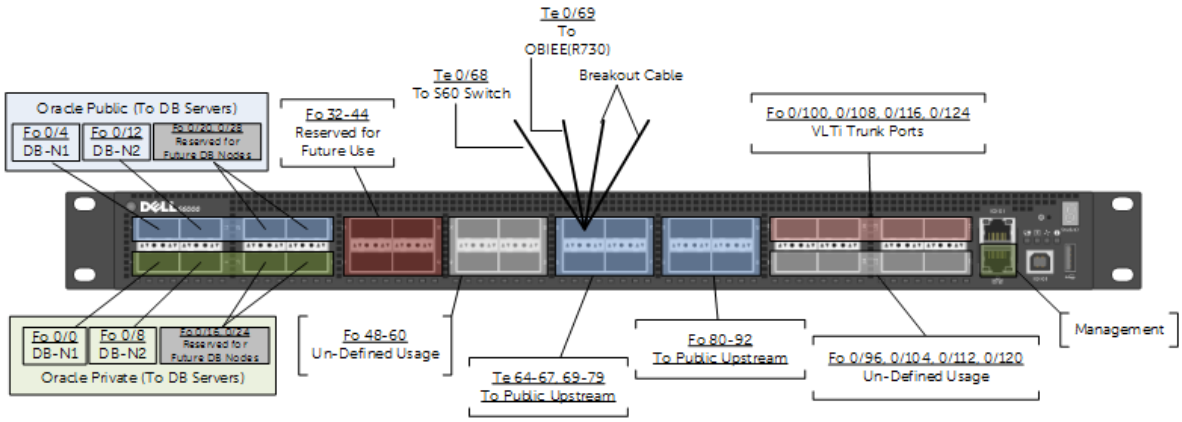

Figure 8 S6000 Port Functions and Configurations

The pre-defined port functions are as follows:

- Ports 0/0, 0/8, 0/18 and 0/24 are configured as FortyGigE ports and are used for Oracle database's private interconnect network. These ports are pre-cabled and connected to the private network ports of the database nodes. 0/18 and 0/24 are not pre-cabled but are reserved for future database node expansion.
- Ports 0/4, 0/12, 0/20 and 0/28 are configured as FortyGigE ports and are used for Oracle database's public network. These ports are pre-cabled and connected to the database nodes' public network ports. 0/20 and 0/28 are not pre-cabled but are reserved for future database node expansion.
- Port range 0/32-44 are configured as FortyGigE ports and are reserved for connectivity to other Dell Integrated Systems.
- Port ranges 0/48-60 and ports 0/96, 0/104, 0/112 and 0/120 are configured as FortyGigE ports with un-defined functionality or usage.
- Port ranges 0/64-79 are configured as TenGigE ports and are used for connecting to customer's 10GbE network environment. The FortyGigE port 0/68 ships with a QSFP+ to 4 x SFP+ 10GbE direct attach breakout cable. The TenGigE port 0/68, 0/69 in that is pre-cabled and connected to the 10GbE module on the S60 switch, Mellanox card on the R730.
- Port ranges 0/80-92 are configured as FortyGigE ports and are used for connecting to customer's 40GbE network environment.
- Ports 0/100, 0/108, 0/116 and 0/124 are configured as FortyGigE ports and are used as VLTi trunk ports. They are connected to the corresponding VLTi trunk ports on the other S6000 switch.

3.2.2 S6000 VLAN and Port Channel Configuration

Though it is typically recommended to have dedicated switches for the Oracle database's public and private traffic, in this pre-integrated solution, the S6000 switches are used for handling both types of traffic. The switches are however pre-configured using Dell and Oracle's best practices by segregating the public and the private traffic using VLANs. Using one of the S6000 switches as an example, Figure 9 shows the VLAN and port channel configuration on the S6000 switches.

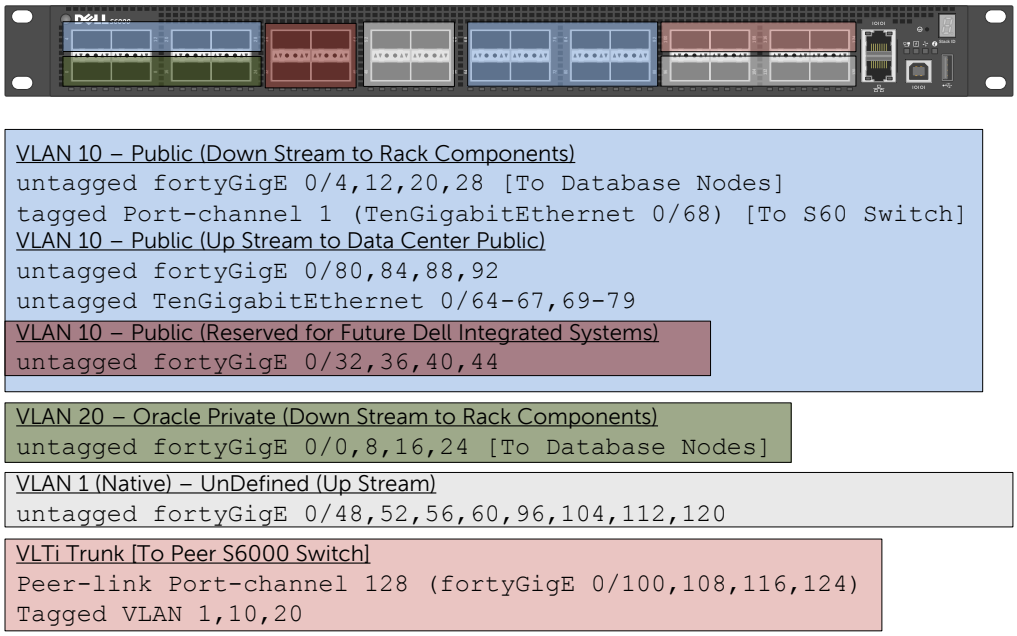

Figure 9 S6000 VLAN and Port Channel Configuration

Thus, the ToR S6000 switches are configured to provide secure and segregrated connectivity for both public and private interconnect traffic.

3.2.3 S6000 VLT Domain Configuration

Virtual Link Trunking (VLT) is a proprietary aggregration 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 links.

Figure 10 shows the VLT Domain configuration across the two S6000 switches implemented in the DISOBA solution rack. Four FortyGigE ports are used for VLTi trunking to provide adequate bandwidth for any customer network connecting to it. The management port is used for the VLT backup link. Port Channel 1 on both the S6000 switches that connect to the S60 management switch are configured as VLT peer LAG to provide 20GB of bandwidth and high availability.

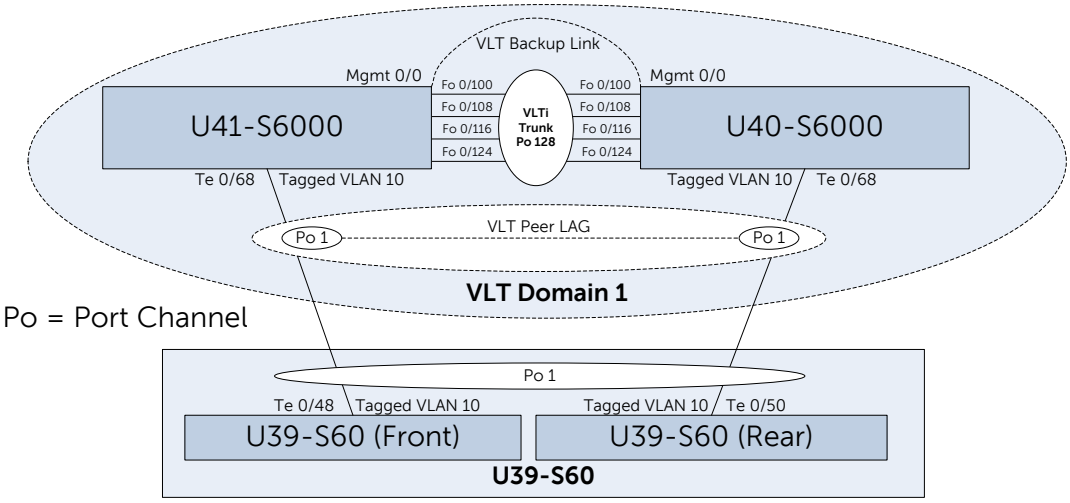

Figure 10 S6000 VLT Domain Configuration

3.2.4 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 across the two individual S6000 switches. However, RSTP on both the S6000 switches are configured with the default bridge priority of 32768.

3.3 S60 Switch Configurations

Figure 11 shows the VLAN and the port channel configuration on the S60 switch.

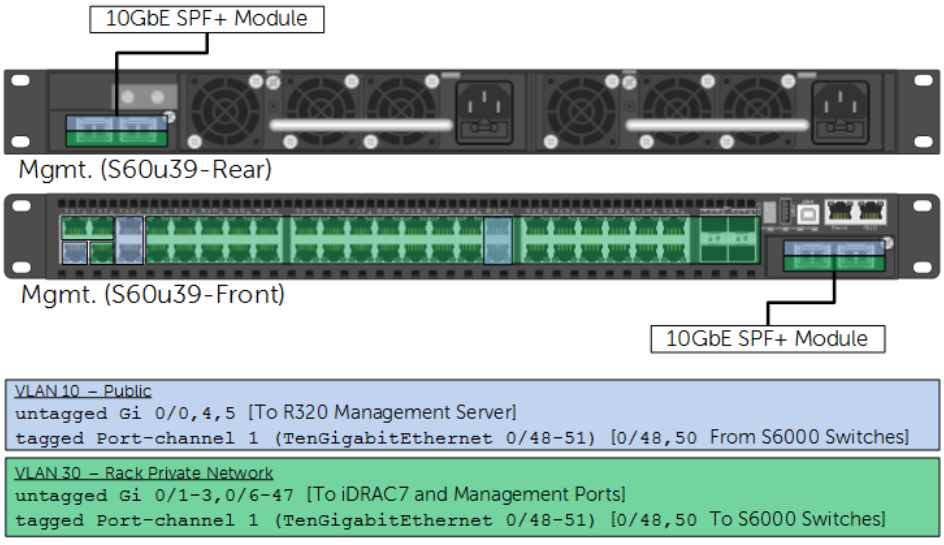

Figure 11 VLAN Configuration on the S60 Management Switch

RSTP is enabled across the switch with bridge priority set very low to 40960.

3.4 Management and Monitoring Network

DISOBA solution rack has a built-in single point of management using a dedicated infrastructure management network and management node. Using the DISOBA management infrastructure, Customers and Dell Services/Support can remotely manage and monitor any of the hardware components.

Figure 12 shows how the DISOBA management infrastructure is setup for providing the remote access to manage and monitor the solution rack.

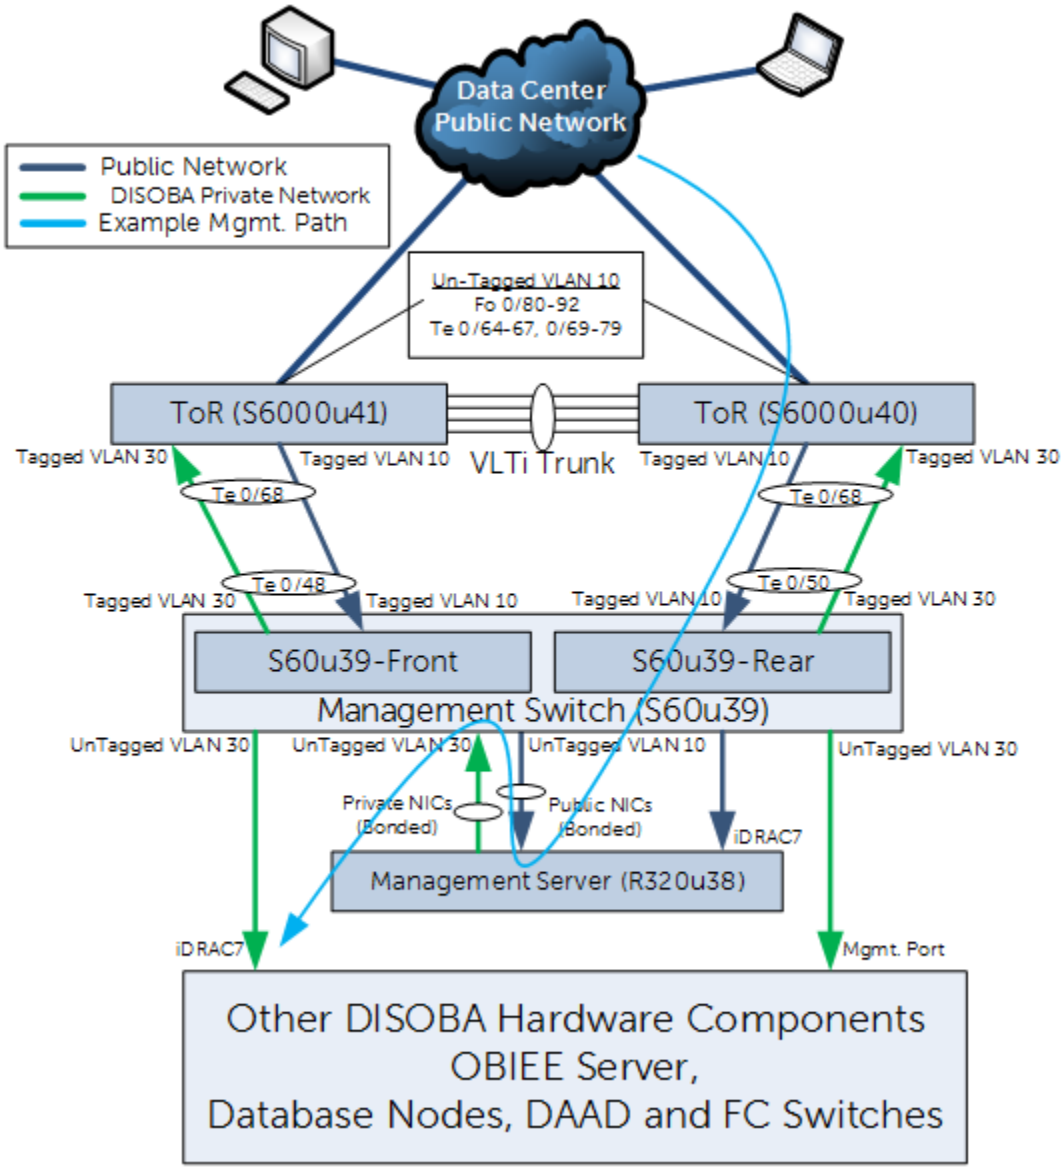

Figure 12 DISOBA Management and Monitoring Network

As seen in Figure 12, the Te 0/68 port on the ToR S6000 switches are connected to the 10GbE modules (ports Te 0/48 and 0/50) on the S60 management switch. And, the 1GbE

ports 0/4 and 0/5 on the S60 management switch is connected to the public network adapters on the R320 management server. This is the route in to the management server from the customer's public environment that is isolated as VLAN 10 network traffic. Once connected to the management server, any of the DISOBA hardware components that are on the isolated private rack network VLAN 30 can be accessed for management or monitoring.

Users can connect to the management server

- Using a browser on their client-machine to log in to the management server's iDRAC for graphical console, or
- Using a terminal on their client-machine to SSH in to the management server's public IP for faster command line access

NOTE: The IP address for the iDRAC7 or the bonded public adapters on the management server need to be configured first at the customer's site before any user can remotely access them.

And, once remotely logged in to the management server (MS), the user can connect to any of the DISOBA components with any of the options listed in Table 5.

DISOBA Component	Graphical (Use Browser on MS)	CLI (Use Terminal on MS)
OBIEE Server (R730)	iDRAC	SSH (private management IP)
Database Nodes* (R920s)	iDRAC	SSH (private management IP)
DAAD Storage Arrays	iDRAC	SSH (private management IP)
Brocade 6510 Switches	Management IP	SSH (management IP)
S6000 Switches	Management IP	Telnet (management IP)
S60 Switch	Management IP	Telnet (management IP)

Table 5 DISOBA Components Remote Management Connectivity Options

***IMPORTANT**: Though the Database Nodes can be accessed directly via its Oracle public network ports, it is strictly recommended not to connect to it for management or monitoring purposes. Oracle public network should be strictly used for database traffic.

Refer to Table 4 for the pre-assigned private IP addresses for each of the DISOBA hardware components on the private rack network.

4 DISOBA Software Configuration

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

- DAAD Storage Array Configuration
- OBIEE Server (R730) Configuration
- Database Nodes (R920s) Configuration
- Management Server (R320) Configuration

4.1 DAAD Storage Array Configuration

This section provides the details of the DAAD storage array that is configured using Dell and Fusion-io's best practices for Oracle Database.

Each DAAD storage array node has four 3.0TB ioMemory cards. With ION HA clustering enabled, the four ioMemory cards (ioDrive1-4) of one DAAD storage array are mirrored with the four in the other node in the HA cluster, as shown in Figure 13.

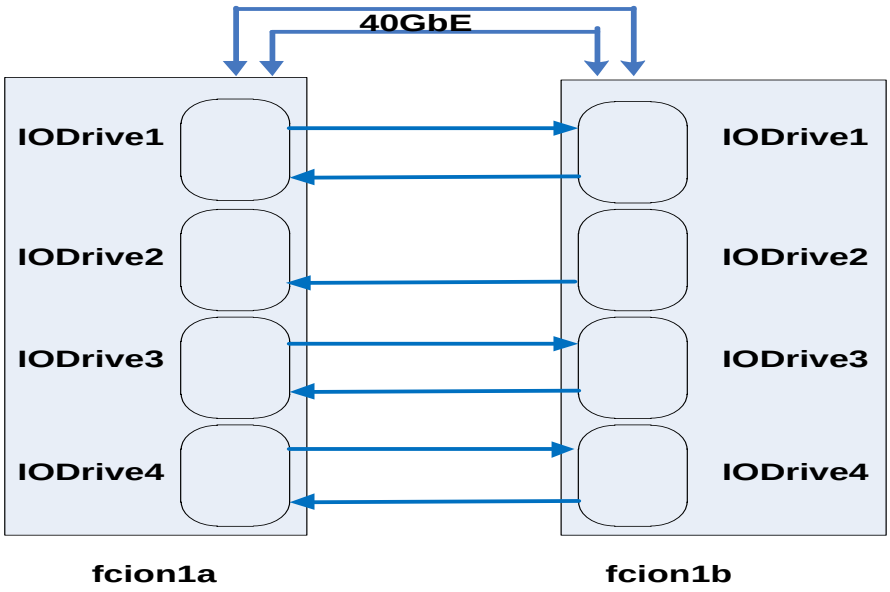

Figure 13 Mirroring between the ioMemory cards between the DAAD Storage Array nodes

A storage pool is created for each ioMemory card, and then two volumes are created for each storage pool. Each volume has a primary storage node and a secondary node. 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. When the database servers update data on the volumes, updates load onto the primary nodes and are then replicated to their mirrored volumes on the secondary node. This design allows

us to balance the workloads evenly over two storage array nodes while maintaining redundancy for high availability.

Each of these eight volumes is created with the same size as they will be presented to the Oracle database nodes to form Oracle ASM disks of an ASM disk group.

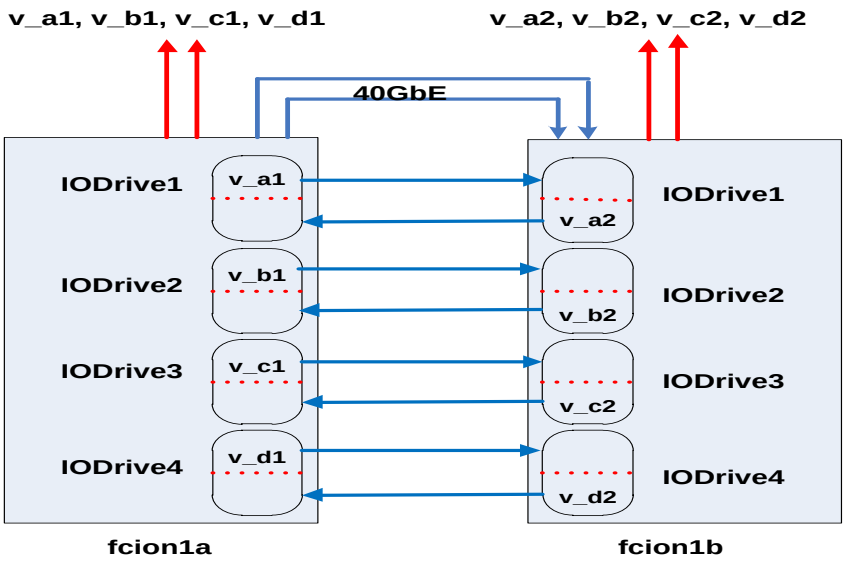

Figure 14 Storage volume setup on the DAAD Storage Array

DAAD storage array uses the default name `fcion1` for the first cluster array. With HA clustering configuration, this array has two nodes: `fcion1a` and `fcion1b`. If there is a need to add additional clusters, this naming convention carries on for each additional cluster array, such as `fcion2` for the second cluster array. With the same naming convention, the node names for the second cluster will be `fcion2a` and `fcion2b`. By default, DISOD is shipped with one DAAD cluster which consists of two DAAD storage array nodes.

8 volumes are created for each DAAD cluster array. These 8 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 will be the cluster name + the local name. For example, the 8 volumes of the first cluster array (`fcion1`) will be:

```
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 of size 1.5TB each, giving a total raw storage capacity of 12TB are mapped to the database nodes for Oracle Database usage. Refer to section *Shared Storage Volume Partitions* for further details on how these eight volumes are configured on the database nodes.

4.2 OBIEE Server (R730) Configuration

This section described the OS pre-configuration that is applied on the OBIEE server to ease the Oracle Business Analytics deployment at the customer. The configuration is done based on Dell and Oracle's best practices.

Database nodes are pre-configured and optimized in the following areas:

1. Hardware configuration
2. Oracle Business Analytics 11g Pre-Installation Setup
3. Oracle network configuration

4.2.1 Hardware Configuration

Before any software is installed, we have taken steps to ensure that the hardware configuration is adequate to support the Oracle environment. This includes making sure we have optimized the BIOS settings, updated appropriate firmware, used the proper RAID configuration, and created properly sized/named partitions. Some details are as follows:

- BIOS matches our supported configuration
- All hardware firmware matches the supported version
- Each R730 has 4 600GB local HDDs in 2 x RAID 1
- File System layout:
  Virtual Disk 1
  - 60GB `/usr/home`
  - 17GB swap (minimum 16GB + 1KB required by Oracle)
  - Rest dedicated to `/u01` to be used for Oracle BI software installation

  Virtual Disk 2
  - 600GB /data – owned by oracle user for application data

4.2.2 Oracle Business Analytics 11g Pre-Installation Setup

The OBIEE comes pre-installed with OL 6.5 running UEK3 (3.8.13-44). To make integration of Dell hardware and Oracle BI software seamless, the Oracle Database 12c's preinstall utility has been customized to eliminate many of the manual steps 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 Oracle recommended module options

- Create special oracle user with specific groups and permissions for Oracle BI

To supplement the Oracle BI 11g preinstall utility, an additional customized deployment utility is used. This utility builds on top of the preinstall utility, and takes additional steps to ensure that the environment is setup properly. Some of the steps include:

- Disable SELinux (optional)
- Create additional users/groups/permissions for Oracle's BI Infrastructure
- Create base directories and permissions for the BI installation
- Install the necessary RPMs needed for Oracle BI 11g

4.2.3 Oracle Network Configuration

In order to minimize the needed steps taken by the customer to setup the Oracle BI environment, DISOBA also ships with some of the network pre-configured on the OBIEE server. The network is configured using Dell and Oracle's best practices in mind. To keep every possible component redundant for high availability, the public networks are each configured in a redundant mode. To further increase redundancy, each of the 10Gb Mellanox Ethernet cards have one link from each of the public networks to the Public switches to reduce yet another single point of failure. By spreading both networks over two cards, we are able to eliminate the possibility of a downed card resulting in lost communication.

NOTE: The public network interface is not pre-assigned with an IP address and requires to be assigned at the customer's site.

4.3 Database Nodes (R920) Configuration

This section describes the OS pre-configuration that is applied on the database nodes to ease the Oracle Database deployment at the customer site. The configuration is done based on Dell and Oracle's best practices.

Database nodes are pre-configured and optimized in the following areas:

4. Hardware configuration
5. Oracle Database 12c pre-installation setup
6. Shared storage volume partitions
7. Disk permissions setup using UDEV rules
8. Device-Mapper multipath
9. ION tuner
10. Oracle network configuration

4.3.1 Hardware Configuration

Before any software is installed, we have taken steps to ensure that the hardware configuration is adequate to support the Oracle environment. This includes making sure we have optimized the BIOS settings, updated appropriate firmware, used the proper RAID configuration, and created properly sized/named partitions. Some details are as follows:

- BIOS matches our supported configuration
- All hardware firmware matches the supported version
- Each R920 has two 900GB local HDDs in RAID 1
- Logical volume configuration:
    - 60GB `/usr/home`
    - 17GB swap (minimum 16GB + 1KB required by Oracle)
    - Rest dedicated to `/u01` to be used for Oracle database software installation

4.3.2 Oracle Database 12c Pre-Installation Setup

The database nodes come pre-installed with OL 6.5 running UEK3 (3.8.13-44). To make integration of Dell hardware and Oracle Database software seamless, the Oracle Database 12c's preinstall utility has been customized to eliminate many of the manual steps 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 Oracle recommended module options
- Create special users with specific groups and permissions for Oracle Database 12c's management

To supplement the Oracle Database 12c's preinstall utility, an additional customized deployment utility is used. This utility builds on top of the preinstall utility, and takes additional steps to ensure that the environment is setup 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

4.3.3 Shared Storage Volume Partitions

With DAAD not being a traditional spinning disk based storage, but being 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, etc. As a result, the eight storage volumes that are pre-configured and assigned to the database nodes are pre-partitioned at the OS level and presented as raw disks these various Oracle database functions.

Each of the eight 1.5 TB DAAD volume gets two partitions - one 15 GB and the other spanning rest of the disk i.e. 1485 GB. Thus, DISOBA solution provides and is configured with the following shared storage at the database node level:

1. Eight 15 GB partitions or disks, for a total of 120 GB, three of which are used for OCR and Voting Disks, and the rest of which are free to use for additional storage.

2. Eight 1485 GB partitions or disks, for a total of 11.88 TB, that are available for data files, FRA, REDO log files, etc.

NOTE: Though DISOBA provides up to 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, giving all-together six copies.

4.3.4 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 pre-configured on the DISOBA database nodes. These ownership and permissions are setup using UDEV rules across eleven disks – 3 disks for OCR/VDs and 8 disks for data/FRA. The UDEV rules are configured in the `/etc/udev/rules.d/99-dell.rules` file.

4.3.5 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 setup, with redundancy verified and using Dell and Fusion-io's best practices, before shipping to the customer.

4.3.6 ION Tuner

The OS on the database nodes are configured with best practices and storage performance optimizations from Fusion-io. Fusion-io has included these optimizations in an rpm package called ION Tuner. This package sets low-level kernel parameters, as well as special UDEV rules, to optimize the performance of the DAAD. A few of the steps taken by this optimization (as per fusionio.com) are as follows:

1. Tune block devices
2. IRQ pinning
3. Multipath optimization

For more details, see: http://www.fusionio.com/white-papers/accelerate-oracle-rac-on-dell-poweredge-r720-servers

4.3.7 Oracle Network Configuration

In order to minimize the needed steps taken by the customer to setup the Oracle environment, DISOBA also ships with some of the network pre-configured on the database nodes. The network is configured using Dell and Oracle's best practices in mind. To keep every possible component redundant for high availability, both the public and private networks are each configured in a redundant mode. To further increase redundancy, each of the 40Gb Mellanox Ethernet cards have one link from each of the public and private networks to reduce yet another single point of failure. By spreading both

networks over two cards, we are able to eliminate the possibility of a downed card resulting in lost communication.

The public interfaces are in Linux bonded mode 1, which is an active-backup configuration. This means that if a single public interface loses connection, there is always an alternate path that takes over to ensure continuous communication.

NOTE: The public network interface is not pre-assigned with an IP address and requires to be assigned at the customer's site.

The private interfaces are not bonded at the OS level since the recommendation is to utilize Oracle's HAIP, 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. The following two private interface network scripts are pre-configured at the OS level:

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

NOTE: Refer to Table 4 for the pre-assigned IP addresses for the private interfaces

During Oracle grid installation select the above two pre-configured network interfaces for the private interconnect.

4.4 Management Server (R320) Configuration

The management server (MS) is pre-installed with Oracle Linux 6.5 running the Red Hat Compatible Kernel. The MS ships with the following pre-configuration:

- As an isolated and contained DHCP and TFTP Server
    - This ensures the IP address for the private rack network components, e.g. iDRAC and private management port on the servers and management ports on the switches is static and consistent across reboots. This is achieved by pre-configuring the `/etc/dhcpd/dhcpd.conf` file with the MAC addresses of these network ports.
    - This allows to easily upgrade the firmware or apply pre-configured switch settings on the S6000 and the S60 switches by booting them in `jumpstart` mode. Though this feature is available, customers are not recommended to upgrade or change the switch settings without consulting with Dell Support.
    - DHCP is configured to always use rack private network so that it does not cause any conflict with customer's DHCP server. DHCP is setup to use `bond0` private adapter on the MS for servicing all DHCP requests.
    - To quickly access and manage DISOBA components, on the management server Desktop, shortcuts to the iDRAC sessions are created for each component.

A DISOBA Rack Hardware Cabling Diagrams

DISOBA components are pre-cabled before they are shipped out to the customer. This section provides the rack cabling diagrams for reference. The diagrams are divided on the following logical basis:

- Cabling Diagram - LAN Infrastructure
- Cabling Diagram - SAN Infrastructure
- Cabling Diagram - Management Infrastructure

A.1 Cabling Diagram - LAN Infrastructure

This section provides, as seen in Figure 15, the cabling diagram of the hardware components that are part of the LAN infrastructure that is used for Oracle Business Analytics, Oracle Database's public and private interconnect traffic.

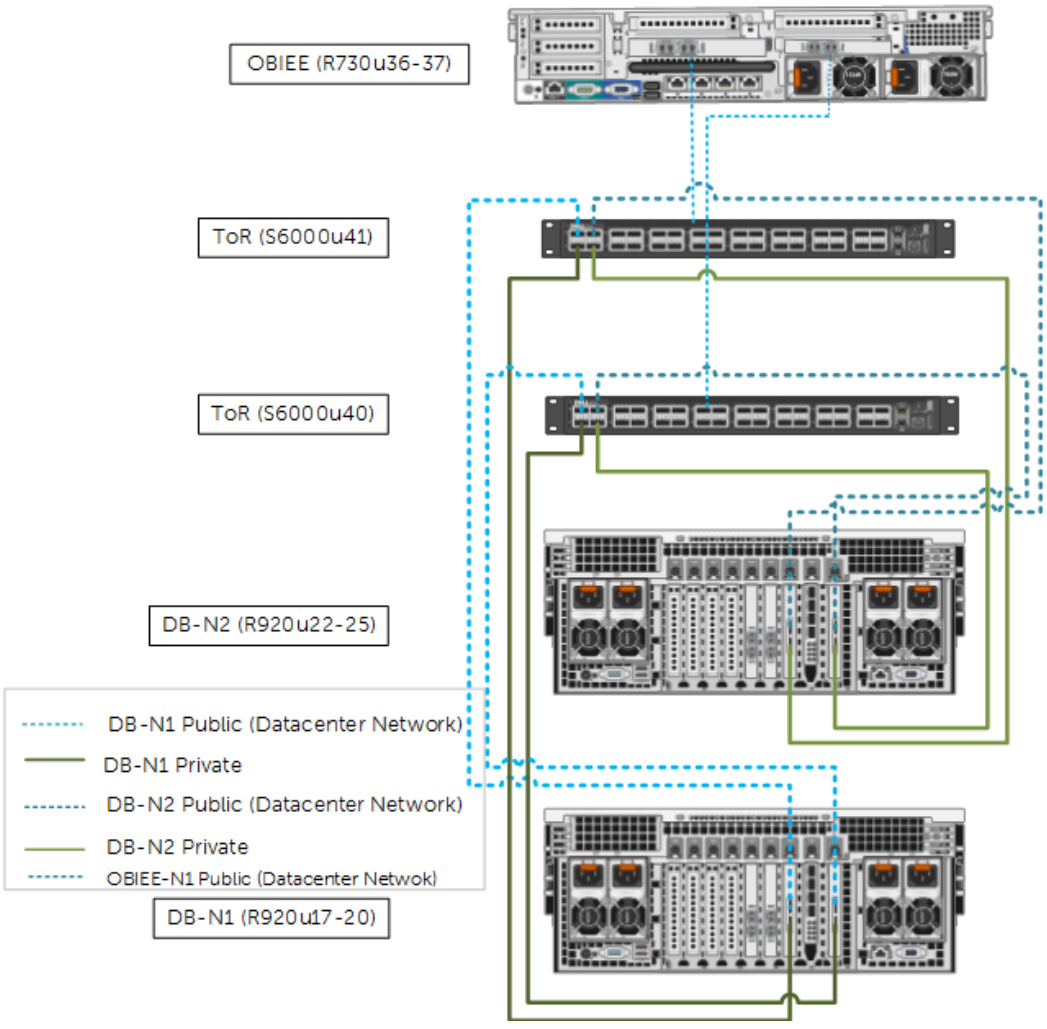

Figure 15 Cabling Diagram – LAN Infrastructure

A.2 Cabling Diagram - SAN Infrastructure

This section provides the cabling diagram for the hardware components that are part of the SAN infrastructure.

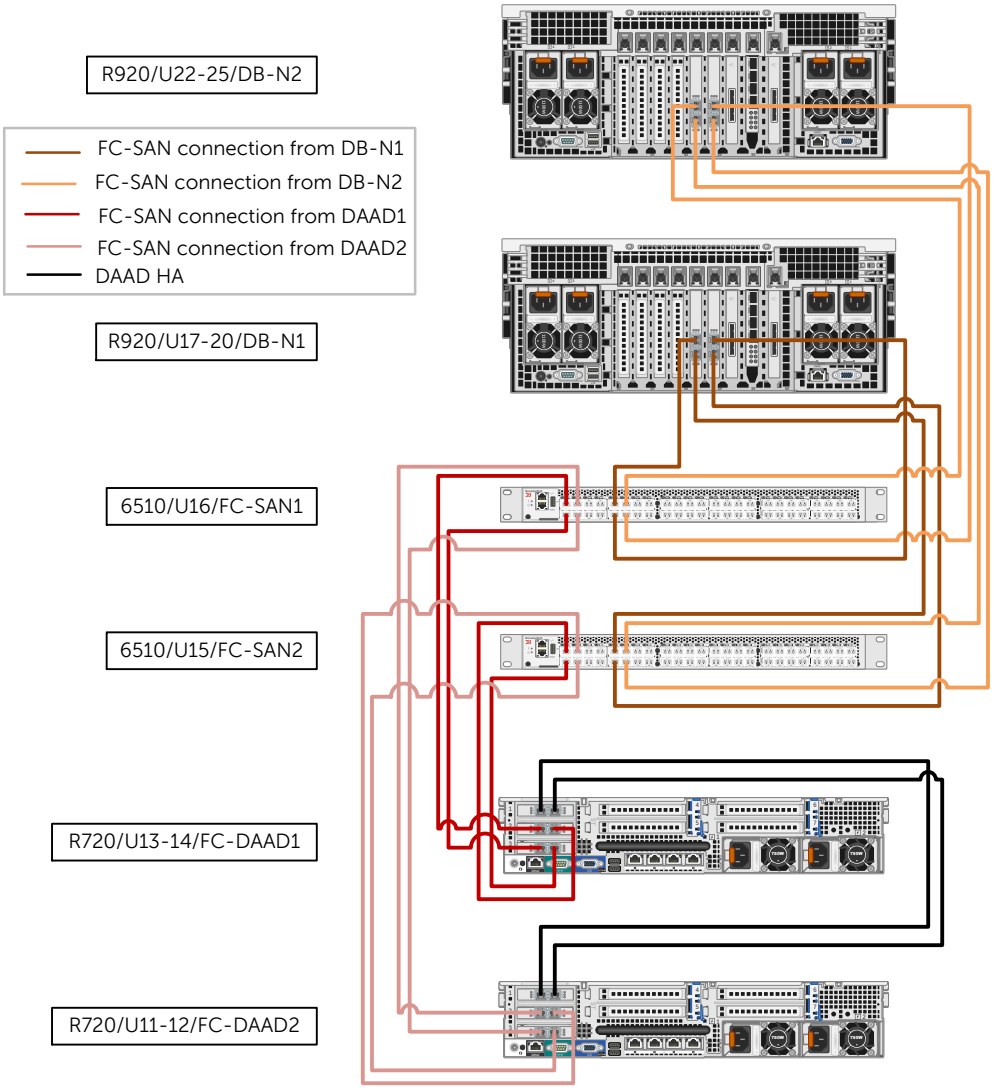

Figure 16 Cabling Diagram – SAN Infrastructure

A.3 Cabling Diagram - Management Infrastructure

This section provides the cabling diagram for the hardware components that are part of the management infrastructure.

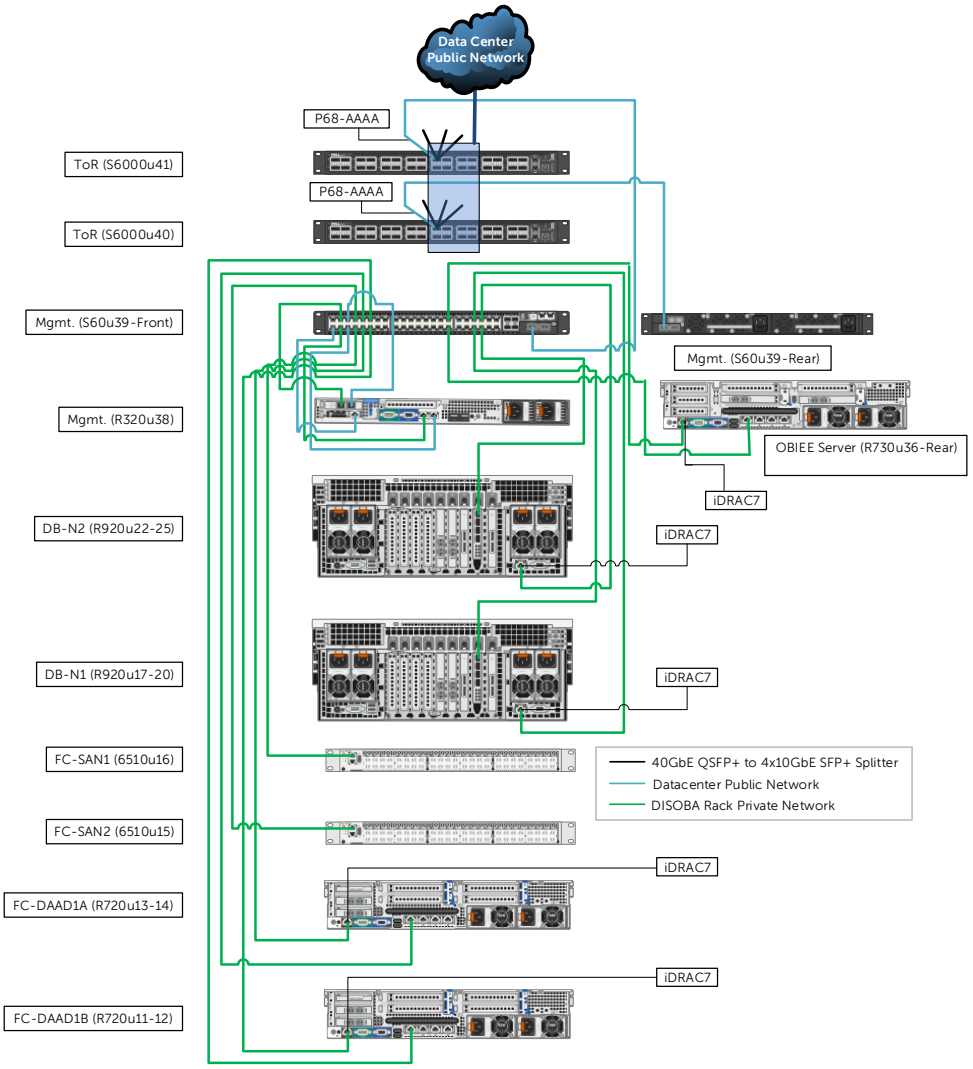

Figure 17 Cabling Diagram – Management Infrastructure

B DISOBA Authentication Settings

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

- `root/oracle` : BI node (R730), Database nodes (R920s) and management server (R320)
- `admin/oracle` : DAAD storage array nodes
- `admin/password`: Management switch (S60), 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 DISOBA solution rack to the customer's network environment.