An Oracle White Paper
March 2014

Oracle Traffic Director on Exalogic – Configuring VIPs

Executive Overview

Oracle Traffic Director is a powerful Layer-7 Application Delivery Controller (aka software load balancer), capable of large-scale deployments and of flexible enough to handle the multitude of various request types that might go through it. It adds a number of key capabilities, such as intelligent routing, compression, and content caching, among others. Couple with Exalogic's low-latency and high throughput Infiniband networking, the result is an ideal environment for hosting Fusion Middleware applications. Oracle Traffic Director can be set up to handle all incoming requests from the outside world, as well as internal traffic going from component to component internally. When properly set up, it can dramatically reduce traffic on a customer's network, by keeping much of the internal application "chatter" entirely within an Exalogic system. This paper covers the various techniques used to set up Oracle Traffic Director to achieve success.

Introduction

This paper is to be used in conjunction with existing product documentation. The topologies presented here are representations of the recommended topology, as suggested by the documentation. What we add is the Exalogic part – how to apply the concepts using a Virtualized Exalogic rack. This paper will showcase 3 common scenarios where Oracle Traffic Director would be recommended, as well as two common variations of how these would be implemented on Exalogic. The use of Oracle Traffic Director to terminate SSL – one of its primary use cases – is orthogonal to this topic, and is thus omitted. Nothing depicted here prevents SSL from being added into any of these scenarios.

Before the scenarios are introduced, some background material is covered. This is mostly because we need a mapping between the concepts proposed in the Oracle Traffic Director documentation, and those within Exalogic used to create a working topology.

In terms of usage, this paper assumes that the primary use case for Oracle Traffic Director will be as a software load balancer for Oracle Fusion Middleware. Oracle Traffic Director can be for simple JEE applications running inside WebLogic Server, as well as complex SOA Suite components, or Oracle Applications, which require components to communicate with each other. Because Oracle Traffic Director must run in a highly available manner, in order to prevent it being a single point of failure, we define Virtual IP addresses for use. These VIPs

are the entry point to one or more instances of Oracle Traffic Director, which then does the load balancing, for the purpose of hiding the back end servers from the outside.

Because Oracle Traffic Director is typically set up for high availability (active or passive), the use of OTD Failover Groups can be part of the setup as well. It is thus important to fully understand the various networks available on Exalogic, when to apply these to a particular VM, and how to configure the most common topologies where OTD is of benefit.

The introduction below will start with the OTD topology, as presented in the documentation:

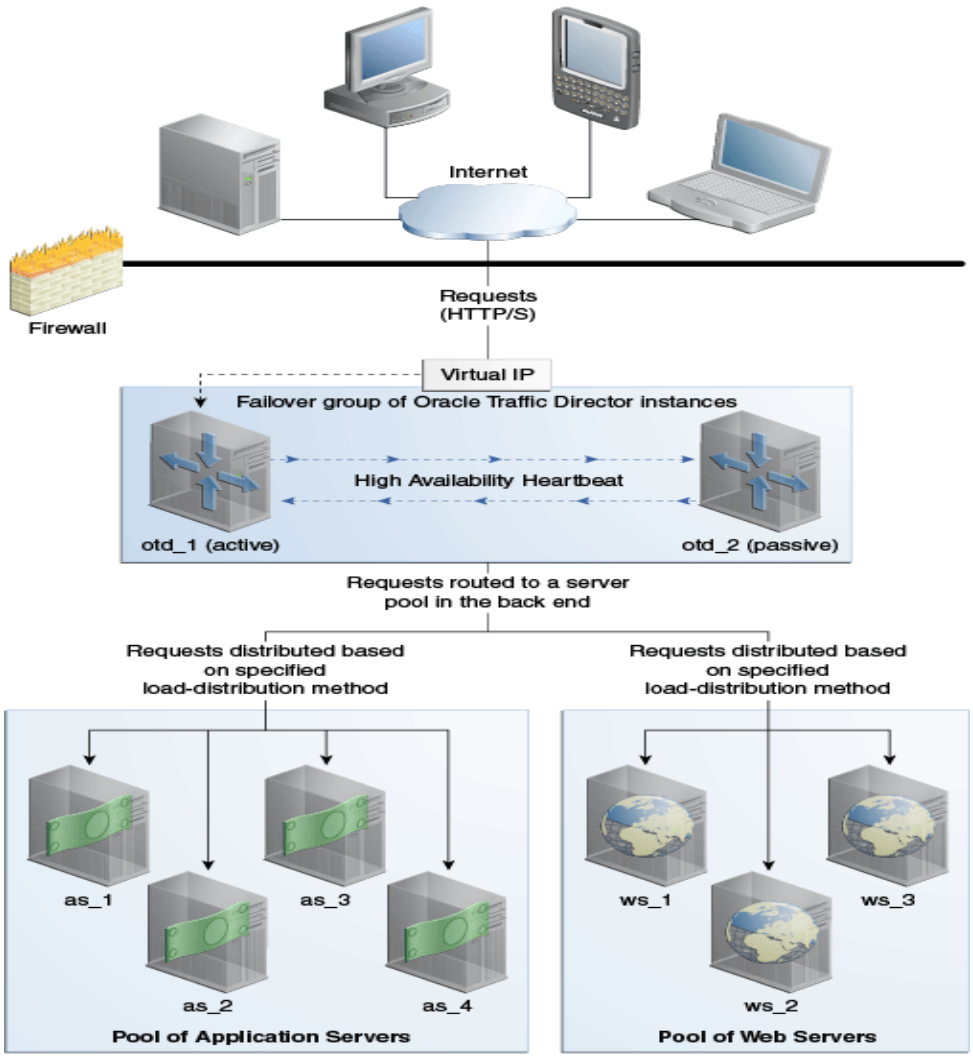

Figure 1 -- Oracle Traffic Director Topology

Having a look at the general topology presented above, it is fairly easy to map the various OTD, web, and application servers (depicted as machines) to Virtual Machines running inside an Exalogic rack. Each component can be represented as a v-server, based on built-in or custom templates and types. However, it is not quite as easy to map other Exalogic cloud concepts, such as Accounts, networks, Distribution Groups, Infiniband, Private V-Nets, and so on and forth. This paper will attempt to fill those voids.

Exalogic Terminology

A (very) brief description of some of the Exalogic Elastic Cloud terms that will be used in this paper appears in this section. By no means is this list comprehensive.

- **V-server** – A virtual machine. In the Exalogic world, a V-server is a given a resource limit (CPU/Memory/Storage), and is granted access to one or more networks. It runs on a single Exalogic compute node at a time.
- **Distribution Group** – This is an anti-affinity concept within Exalogic. By provisioning more than 1 V-server into the same distribution group, Exalogic will guarantee that the V-servers are physically running on separate compute nodes.
- **Virtual Data Center (VDC)** – Exalogic's top level container for all virtualized resources within an Exalogic rack. The VDC contains all networks and cloud accounts created on the entire rack.
- **VDC Cloud Account** – An account is NOT a username/password combination. Instead, it is an allocation of resources and access to networks. Accounts are how Exalogic divvies up resources by group, line of business, or function. Users (one or more) have access to accounts, and can thus share responsibilities within a private cloud scenario.
- **Failover Group** – Oracle Traffic Director mechanism for surfacing a Virtual IP (VIP). This VIP is able to "float" between OTD instances in an active/passive configuration. If the primary OTD instance becomes unavailable, the VIP seamlessly moves to the backup instance.
- **"Public", "Client" or EoIB Network** – A network that has external connectivity outside of the Exalogic rack. Typically this will be the network that v-servers use to get access to resources outside of Exalogic, over the 10Gbe pipe. Each Exalogic rack can host numerous public networks, which allows segregation of applications and application components. NOTE: the "public" part does not necessarily connote "public internet". It simply means outside Exalogic.
- **Private v-net** – A network created entirely within the scope of an account. This network can span ONLY V-servers within a particular account. It provides an internal network, utilizing Infiniband, for which V-servers can communicate. It cannot be accessed from outside Exalogic, and cannot traverse accounts.
- **Infiniband Partitions** – Because Infiniband is a fabric-based technology, this means that its collective bandwidth is pooled. The concept of partitions allows for separation of traffic, without loss of overall capacity. By using partitions (Exalogic does natively), we are able to define multiple network "swim lanes" that cannot see each other's traffic, without having to "slice up" the overall 40Gbps. Partitions are not explicitly created by end users, but rather implicitly by means of creating "Private V-Nets" and by using the in-built Infiniband networks.
- **IPoIB-vserver-shared-storage** – This network is available to be provisioned to each Virtual Machine. Having this network provisioned allows a VM to access the ZFS-based shared storage inside Exalogic. This is important for OTD, as the product installation and instance homes can and

should be stored outside of the VM root disk. This is the case for all Oracle Fusion Middleware products, and OTD is no exception.

- **IPoIB-default** – One of the networks available to the Exalogic rack is the default, which is what would share with an Exadata rack or other Exalogic racks. This is a special type of network that does not participate in Infiniband partitions. It is more or less the "global" partition, thus addresses must be unique on it. Having an address on this network implies that a VM can access global resources, such as Exadata or physical components which have and address on this network.

Shared Versus Multiple Cloud Account

Exalogic's main method of divvying up compute and memory resources is via cloud accounts. By leveraging cloud accounts, different groups of users, lines of business, or application owners can have their own resource pools from which to create v-servers and create networks. Typically the use of a shared or multiple cloud accounts is driven by one of 3 reasons:

1) **Size or organization and the number of administrators** – For example, if one group controls Exalogic, and does all of the deployment tasks, it may make sense to have a single cloud account, to which all resources are allocated.

2) **Number of environments** – In other cases, particularly with Exalogic hosting multiple non-production lifecycle stages of multiple applications, it might be advantageous to have OTD exist as a shared resource. This is advantageous in that multiple copies of the OTD product and configurations do not need to be maintained. Further, application owners need not understand OTD, and can focus on their particular application.

3) **Separation of duties** – Most commonly, the administration and maintenance of OTD may lie with a different group within an organization. They might be responsible for maintaining assets that are closer to the clients, such as DMZs, firewalls, hardware load balancers, and so on. This group isn't typically expected to deal much with the applications themselves, the day to day care and feeding, and the scaling up/out of the application tier.

By changing our model of cloud accounts, the pictures will change, as private v-net traffic cannot traverse the cloud account boundary. Therefore, we'll show a variation for each main scenario, where multiple Exalogic cloud accounts are used. Rather than having each cloud account own and maintain its own copy of OTD, the model will show a shared cloud account for OTD, and allow us to have as many application cloud accounts as necessary.

OTD Topology Overview

Since Oracle Traffic Director could be tasked with load balancing requests from both clients inside or outside of the Exalogic frame, we must follow some general rules when deploying on virtualized Exalogic. The reason is that Exalogic allows multiple network connections, and the correct one to use will depend on the use case. If combined with standard Exalogic best practices, such as the use of the

dedicated shared storage network and use of private v-nets, an OTD v-server will be a provisioned for a minimum of 3 networks. Additionally, one configuration of OTD could be used to support multiple topologies simultaneously. In other words, the scenarios presented here are not mutually exclusive.

Scenario 1 – OTD as Reverse Proxy to External Clients

In this scenario, connections come in over the external (client) EoIB network. OTD listens on this external address in a standalone configuration, or a virtual IP (VIP) in configurations utilizing a Failover Group. If an external hardware-based load balancer is available, it can be used to round-robin requests to each of the available OTD instances directly. If it is not available, OTD's concept of Failover Groups will be applicable. Both of these will be discussed and shown in a diagram. Either way, external clients have, and need, no knowledge of where the back end (Origin) servers are running, and these could be within Exalogic or outside.

Additionally, this paper will include a variation of each aforementioned case with a shared cloud account, as well as with multiple cloud accounts.

Scenario 2 – OTD for Internal Load Balancing

This scenario is similar to the above case, with Exalogic hosting both OTD and the v-servers that run the application that is serving requests. However, in this case, the application components need to communicate with one another, and use load balancing and other OTD functionality as well. Absent of Exalogic, these inter-component calls would go outside each server, to F5, and back into the load-balanced endpoints. But because OTD can serve this purpose, and because OTD and the end application are both on Exalogic, the opportunity for using Infiniband exists. The benefits of higher bandwidth, lower latency, and less traffic going outside Exalogic can thus be realized. For this scenario, v-servers will utilize an Infiniband network, assigned to multiple v-servers. Traffic that goes between origin servers and OTD, as well as internal cluster traffic, will utilize Infiniband networking.

In the exact same fashion as scenario 1, we'll make the distinction between shared and multiple cloud accounts and show the picture both ways.

Scenario 3 – Combined Topology with Exadata

This final scenario will present a combination of both scenarios above, and show communication with an Exadata machine as well. Nothing is fundamentally different here, but this will show a "big picture" of deploying OTD in conjunction with Fusion Middleware and Exadata. In particular, it will highlight the networks to be provisioned for each v-server on the Exalogic side.

OTD Topology Details

In this section, we'll provide some detailed topology diagrams, along with some screen shots and how-to instructions. It is advised to closely follow the Oracle Traffic Director documentation, using the pictures and recommendations below as specific additional data points for an implementation of OTD on virtualized Exalogic.

Scenario 1a – OTD with External Load Balancer

This scenario typically is the simplest path to take when an external load balancer (such as F5) currently points directly to the middleware endpoints (Weblogic or other application servers). OTD is inserted as an intermediary, thus providing more control to the application team over how requests are balanced, logged, and handled. OTD runs on a pair of v-servers (or more as necessary), and all requests come through OTD's EoIB listeners. In this scenario, the load balancer owns the VIP address, and is the only known endpoint for external clients. All traffic comes to the load balancer first, is directed in round-robin fashion to our pair of OTD instances, and is then sent on to the actual origin servers.

The benefit here is two-fold:

1) Clients don't need any updates to their endpoints. They continue to send traffic to the load balancer, and let it do its job.
2) More control is delegated to the OTD administrator. This can help with reducing the effects on the network team when managing aspects of the application. OTD natively supports many of these functions.[1]

Also as stated, this scenario will include both cases for OTD installation – shared and multiple cloud accounts. It is important to note here that both are equally valid, depending on the situation. The differences between the two cases will be highlighted in the implementation notes, following each diagram.

Case 1 – Shared Cloud Account

[1] For a basic list of OTD functionality, see this page:
http://docs.oracle.com/cd/E23389_01/doc.11116/e21036/get_started002.htm

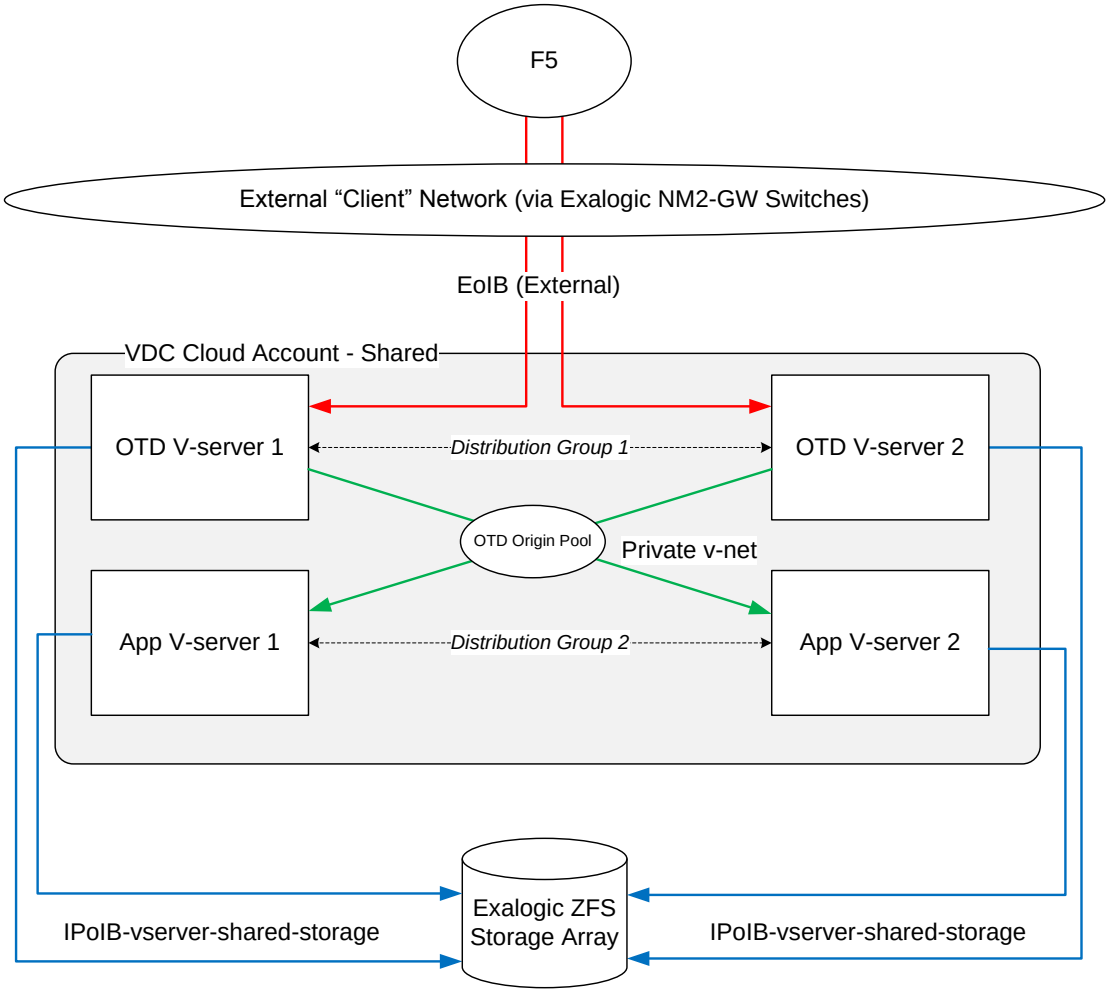

Figure 2 - F5 with OTD as Reverse Proxy to Origin Servers (shared cloud account)

Implementation notes

1. All binaries and instance data for OTD will be stored on the ZFS storage device. Therefore, OTD V-servers are provisioned with the IPoIB-vserver-shared-storage network. While OTD binaries and instance data can be stored on the local v-server's individual root disks, it is always recommended to use shared storage on Exalogic when possible.

2. Each OTD v-server in this scenario has an external connection via the EoIB network (red). The OTD instances on each OTD v-server listen on this address, and these are exposed to the load balancer as its endpoints. In the diagram, (2) EoIB addresses will be provisioned to the cloud account.

3. The private v-net depicted in green is created before provisioning any v-server. All v-servers in this scenario are provisioned with an address on this network, and this address should be used for:

    - Origin server (Application VMs) listen addresses

- OTD's internal communication
- WebLogic or application level clustering

4. Since each v-server has multiple networks, each defines multiple hostnames. Typically the public-facing (EoIB) interface is the primary hostname, whereas a private v-net might have "-priv" appended to it. OTD would be set up to use the private v-net hostnames for its internal communication.
5. Distribution Groups are depicted for both OTD and Application v-servers. For OTD, this would most likely be of size (2). For application v-servers, the size will depend on the application. The drawing depicts (2) application v-servers, but in reality there are likely more.
6. Access to an EoIB network is not always required for application v-servers. If not provisioned, all access to back end servers accessible web consoles would go through OTD only. Remember that Exalogic supports multiple EoIB networks, so if an application v-server was to be connected to EoIB, it would likely be a different VLAN than the one provisioned for OTD. See the OTD product documentation[2] for notes on setting up administrative VIP addresses, should there be no EoIB access assigned to application v-servers.

Case 2 – Multiple Cloud Accounts

Because the private v-net concept does not span Exalogic VDC cloud accounts, we employ a different tactic here. We still would like to access origin servers using Infiniband, and we want to keep all internal traffic inside of the Exalogic rack as possible.

The drawing below depicts using the IPoIB-default network, which will be open to all v-servers that attach to that network. As shown, the OTD v-servers are owned by a cloud account called "OTD", depicted as a shared area in the background. The application cloud accounts are depicted in this fashion as well. The IPoIB-default network, which needs to be made available[3] to all v-servers, is provisioned for each v-server.

This case is functionally the same as the previous, with the exception of IPoIB-default. Even though OTD to v-server traffic does not traverse private v-nets any more, the private v-nets are still created in each account, for the purpose of internal cluster traffic.

[2] http://docs.oracle.com/cd/E18476_01/doc.220/e47690/create_domain_im.htm#BABDBAFF

[3] IPoIB-default is not listed on the Exalogic Configurator spreadsheet, but is created at installation time. It will be the same subnet as the physical Exalogic components' Infiniband addresses, and that of Exadata, if connected. To "make it available", it will need to be allowed as an available network to each cloud account.

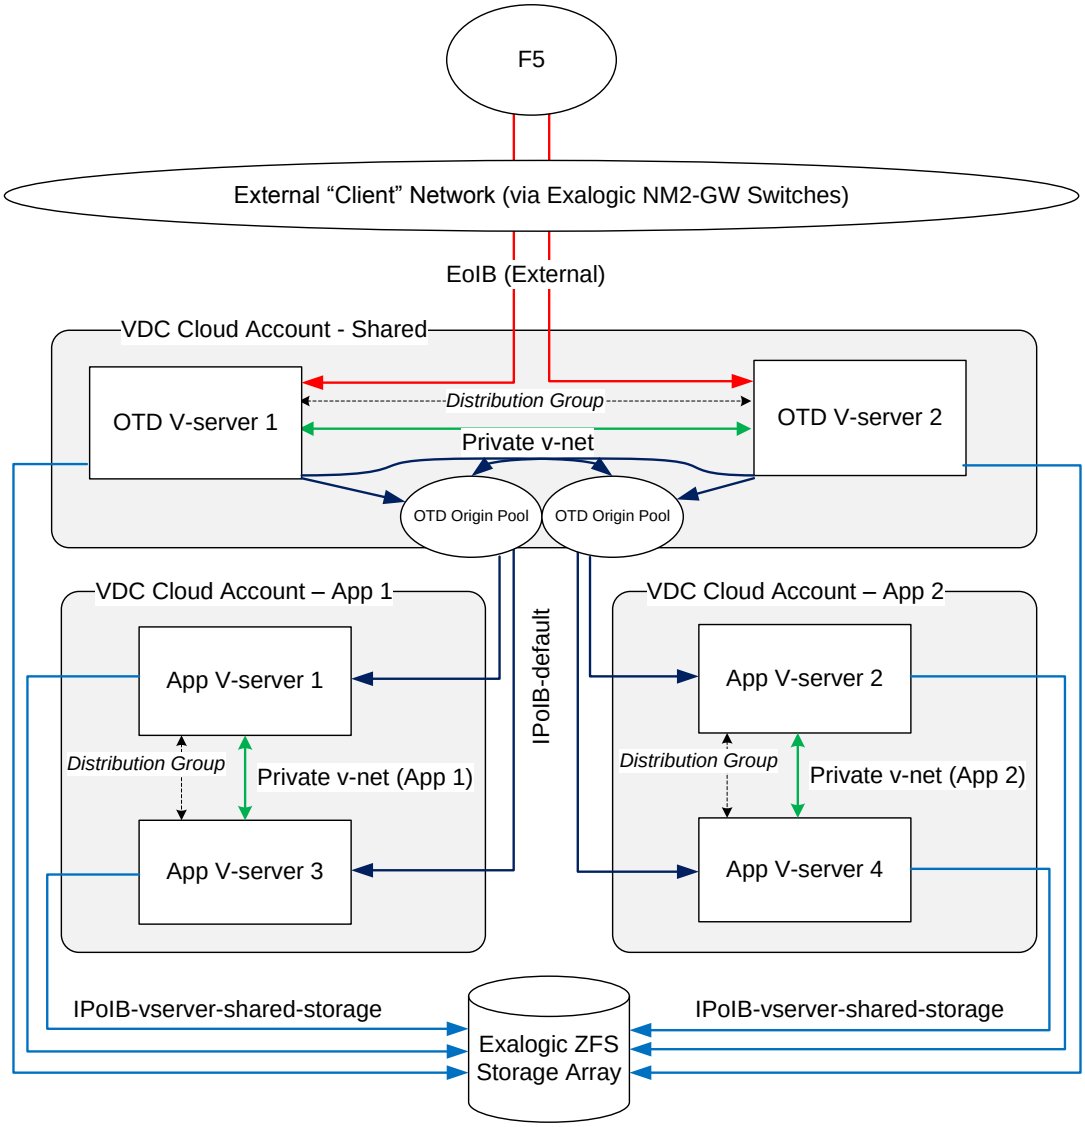

Figure 3 - External VIP with Separate Cloud Accounts

Implementation Notes

1. Cloud accounts are now provisioned for each application, and for OTD. Each cloud account will need an allotment of addresses from IPoIB-default, equal to the number of v-servers likely to be created.

2. As with the shared cloud account diagram, (2) EoIB addresses are provisioned to the OTD cloud account, and one each is assigned to the v-servers.

3. As with the private v-net case, each v-server has multiple networks, and thus multiple hostnames. The IPoIB-default hostnames will be used for this case to define application endpoints for the origin pool.

4. Applications should be configured to listen on the IPoIB-default address or hostname assigned to the v-server, so that access from OTD exists. For WebLogic Server, this would amount to setting the listen-address or using a network channel.
5. Private v-nets are created, one per account, for the purpose of cluster communication.
6. Distribution groups are created per account, and are selected when v-servers are created.
7. OTD Origin pools are configured in a similar fashion to the previous picture, with the main difference being the use of IPoIB-default addresses instead of private v-net ones.

Scenario 1b – OTD with External Failover Group VIP

If F5 or another hardware load balancer is not involved, and there is not another means to direct traffic to 2 or more OTD instances, then OTD can surface its own VIP via a failover group. This will take an additional IP address from the EoIB range, without assigning it directly to one of the OTD v-servers. Instead, the address "floats" between the instances, and acts in an active/passive capacity. Requests come to this VIP and are routed directly to the active instance, unless that instance becomes unavailable. In this scenario, OTD becomes the load balancer, and is solely responsible for directing traffic to a farm of application servers.

As with the other scenarios, a diagram of a shared cloud account and of multiple clouds accounts will be shown, with differences highlighted.

Case 1 – Shared Cloud Account

The picture for using an EoIB VIP on OTD (shared cloud account) looks as follows:

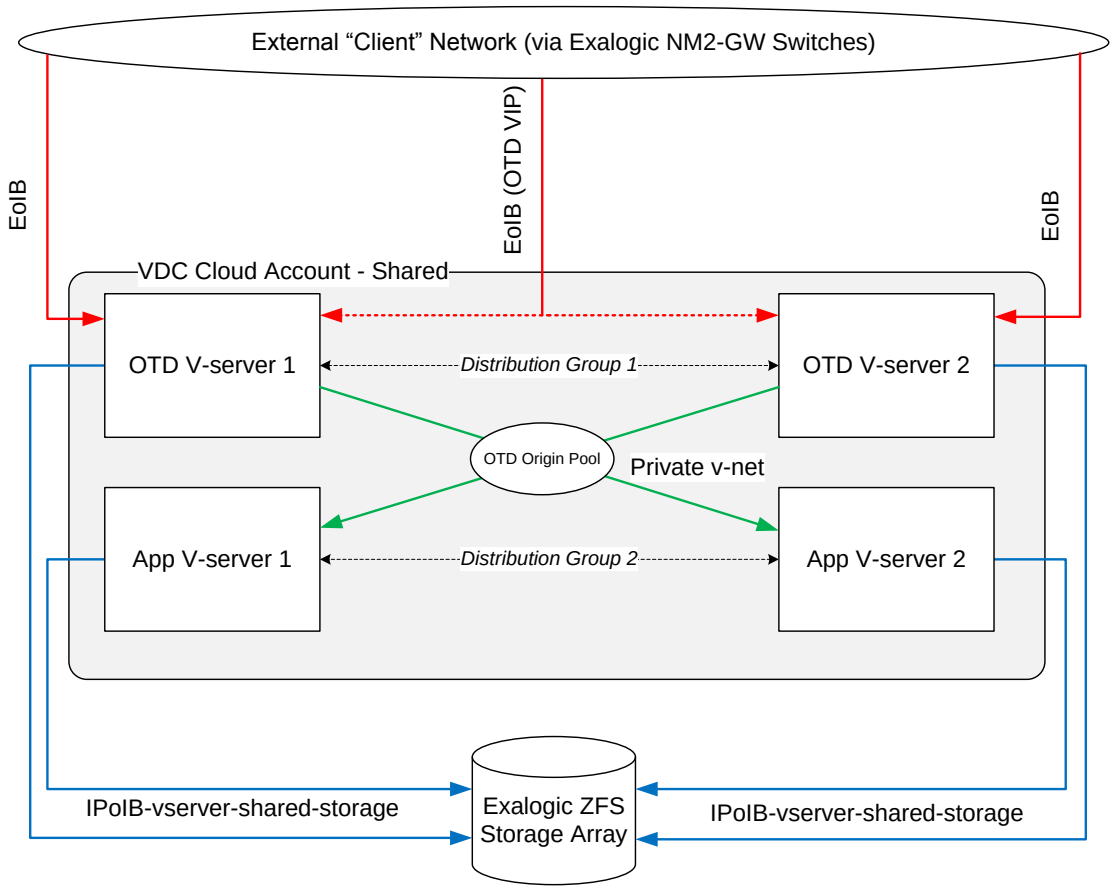

Figure 4 - OTD surfaces external VIP instead of load balancer

As with scenario 1a, the notes for how this is implemented follow the drawing. These will be the same as the prior scenario, with the addition of the failover group notes and diagram.

Implementation Notes

1. All binaries and instance data for OTD will be stored on the ZFS storage device. Therefore, OTD V-servers are provisioned with the IPoIB-vserver-shared-storage network. While OTD binaries and instance data can be stored on the local v-server's individual root disks, it is always recommended to use shared storage on Exalogic when possible.
2. Each OTD v-server in this scenario has an external connection via the EoIB network. The EoIB VIP address (dotted red) is not allocated to a particular v-server, but is done via the Failover Group setup in OTD.
3. The private v-net depicted in green is created before provisioning any v-server. All v-servers in this scenario are provisioned with an address on this network, and this address should be used for:
    - Origin server (Application VMs) listen addresses
    - OTD's internal communication

- WebLogic or application level clustering

4. For OTD Failover Groups (depicted by dotted red line above), it is recommended to "reserve" the selected VIP address(es) once they are assigned. This is done via the Exalogic Control's "Network" section, by editing the EoIB range. For example, the following shows excluding (3) IP addresses from a managed range on an EoIB network, for the purpose of later using those IPs as part of OTD Failover Groups.

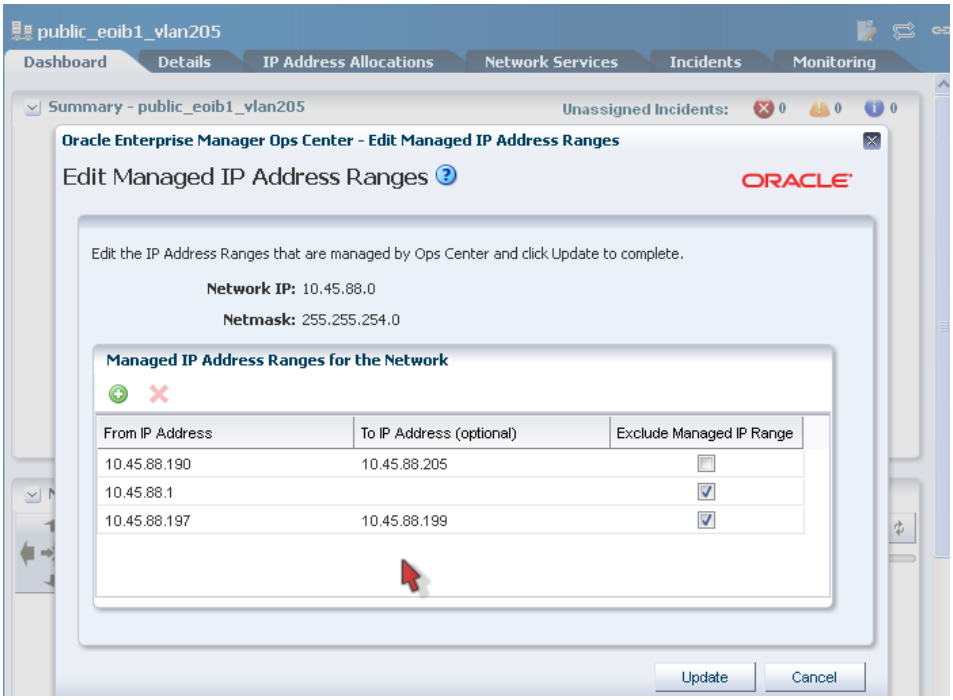

Figure 5 - Reserving IPs to be used for OTD VIP (10.45.88.197 – 199)

Case 2 – Multiple Cloud Accounts

To round out scenario 1, the following drawing represents the addition of multiple cloud accounts to the EoIB VIP example:

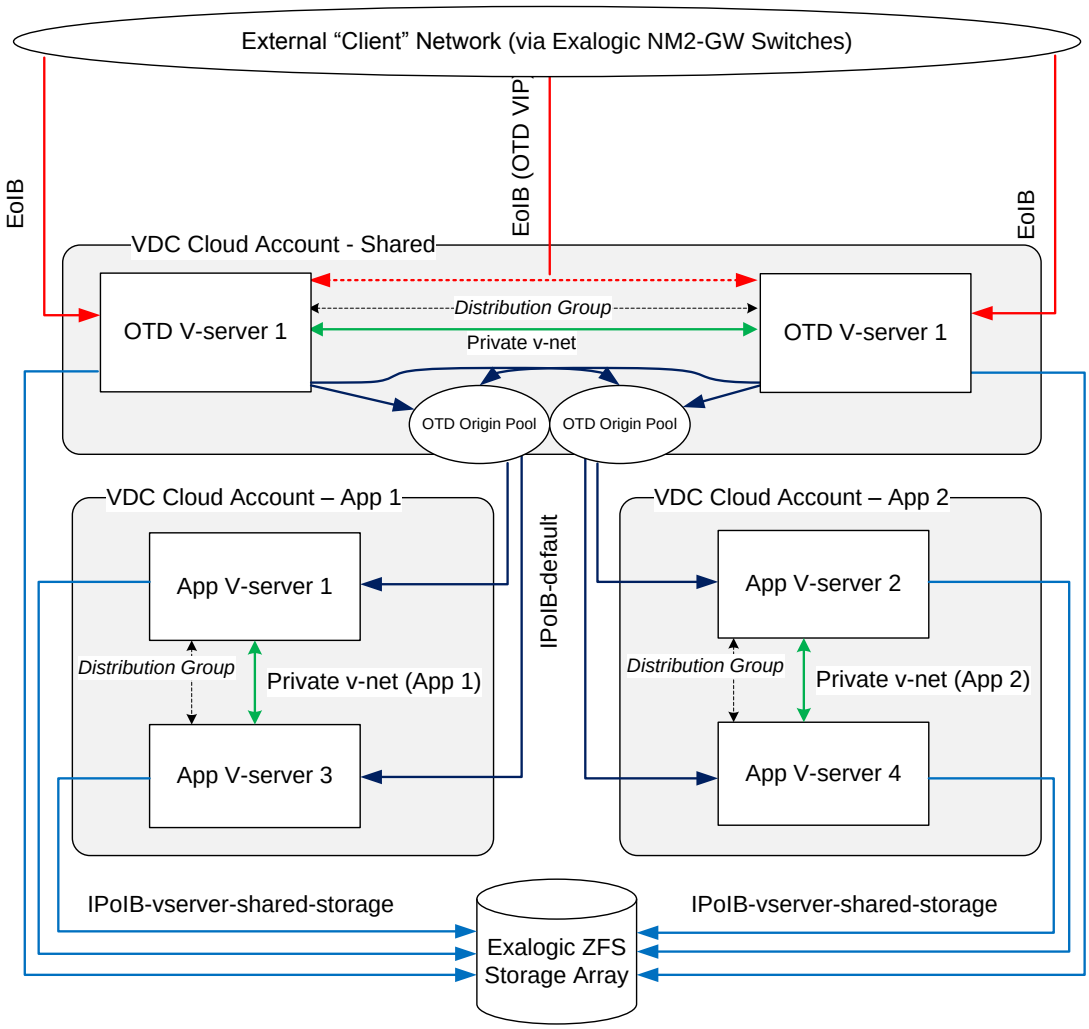

Figure 6 - OTD VIP with Multiple Cloud Accounts

Implementation Notes

1. Each OTD v-server in this scenario has an external connection via the EoIB network. The EoIB VIP address is not allocated to a particular v-server, but is done via the Failover Group setup in OTD.

2. As with the shared account, private v-nets are created here too, but the difference is that they are only used for OTD and application clustering within each account. Since they cannot traverse Cloud Accounts, we use IPoIB-default instead, as that can be used here.

Scenario 2a – OTD with Internal VIP (Shared Cloud Account)

Just as Oracle Traffic Director can proxy in requests from external (to Exalogic) clients, it can also load balance internally. A prime example of this is 2 WebLogic clusters, hosted on the same Exalogic rack. An application may have various components that treat each other as disparate, even though all are running within an Exalogic environment. For maximum performance, it is important to keep all intra-component calls on the Infiniband network. If component A called component B, running on a different v-server or cluster, using an EoIB address, that communication would "leave the frame", and thus be sub-optimal. Using private v-nets and OTD together will eliminate this problem, and allow for taking advantage of the full set of OTD features, without traffic leaving the Exalogic frame.

The following diagram represents an OTD VIP with a private v-net address. EoIB has been removed from the picture to focus on the relevant internal components of the topology. It may still exist for purposes of access and administration, but is not depicted. The cloud account is shared, and is depicted as the shaded region. By using the same cloud account, we can utilize a private v-net for all internal communication.

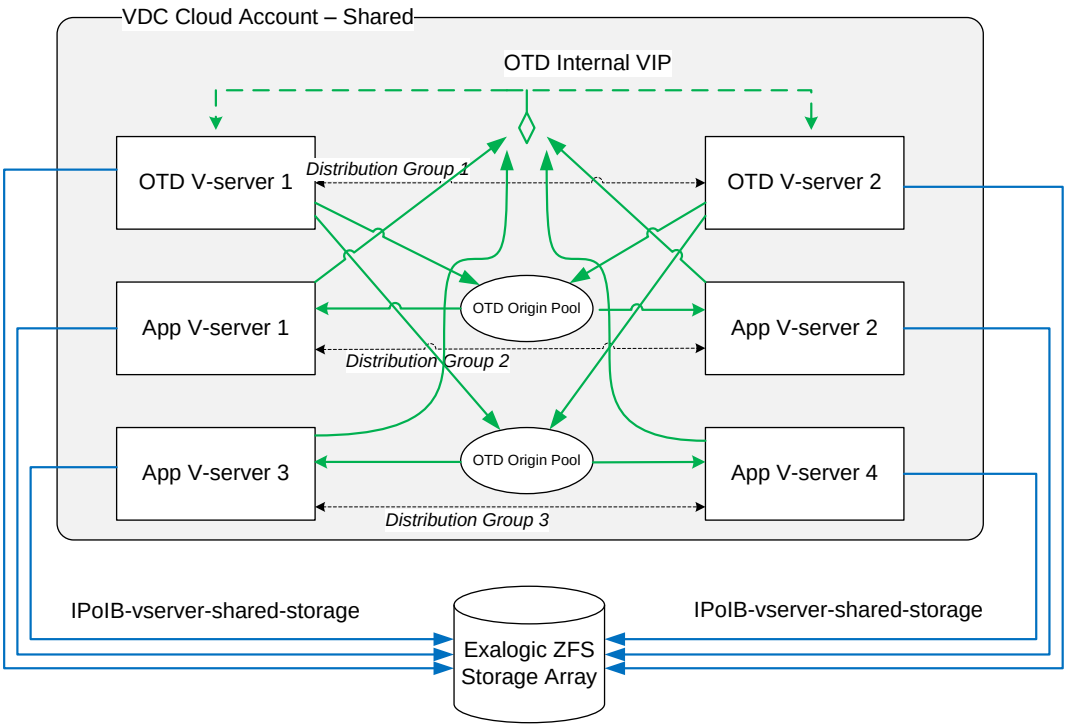

Figure 7 - OTD using internal VIP with shared cloud account

Implementation Notes

1. The most important thing to note here is the private v-net, depicted in green in the middle of the drawing. This network is the centerpiece of the scenario, and is used for:
    - OTD Origin servers (Application v-servers 1-4)
    - All WLS cluster communication
    - OTD Internal administration
    - OTD VIP for Failover Group
2. ZFS storage, distribution group usage, and network provisioning are the same as the previous scenario.
3. Addresses to be used for OTD VIPs internally can be carved out of the private v-net after it is created. This is done via the same network section of Exalogic Control as with scenario 1b, using the Network section with Exalogic Control.

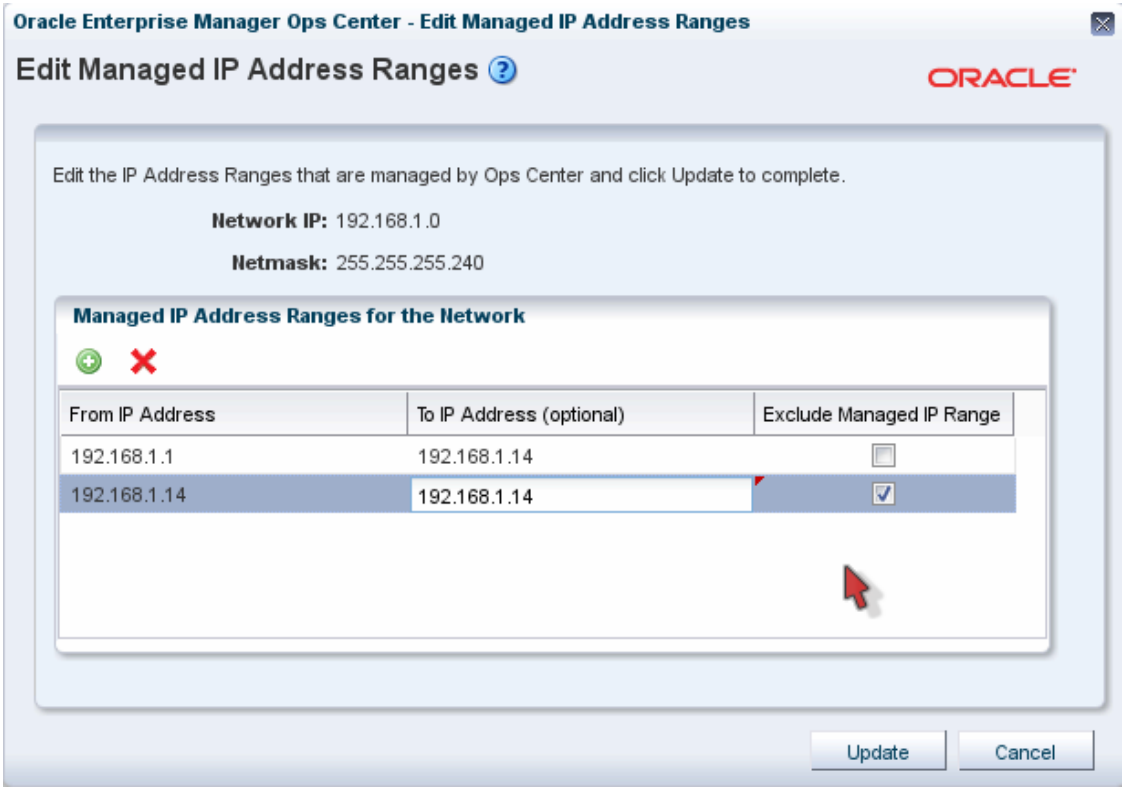

Figure 8 - Reserving a Private v-net VIP address

Scenario 2b – OTD with Internal VIP (Multiple Cloud Accounts)

The internal VIP scenario poses the same decision about separation of duties as do the other scenarios. For any or all of the reasons discussed in the introduction, the decision could be made to keep OTD in its own account. In the case of the interval VIP, the likelihood that Exalogic hosts completely disparate applications that still need to communicate exists. Each application might maintain its own OTD configuration within the shared OTD cloud account, and the benefits from separation of duties can still be realized.

Just like the multiple cloud account diagrams in scenario 1, we use the IPoIB-default network to cross the cloud account boundary. The OTD Failover Group VIP can be created using an IP address in this space, and v-servers owned by any account can access that address, so long as they too are provisioned with access to IPoIB-default.

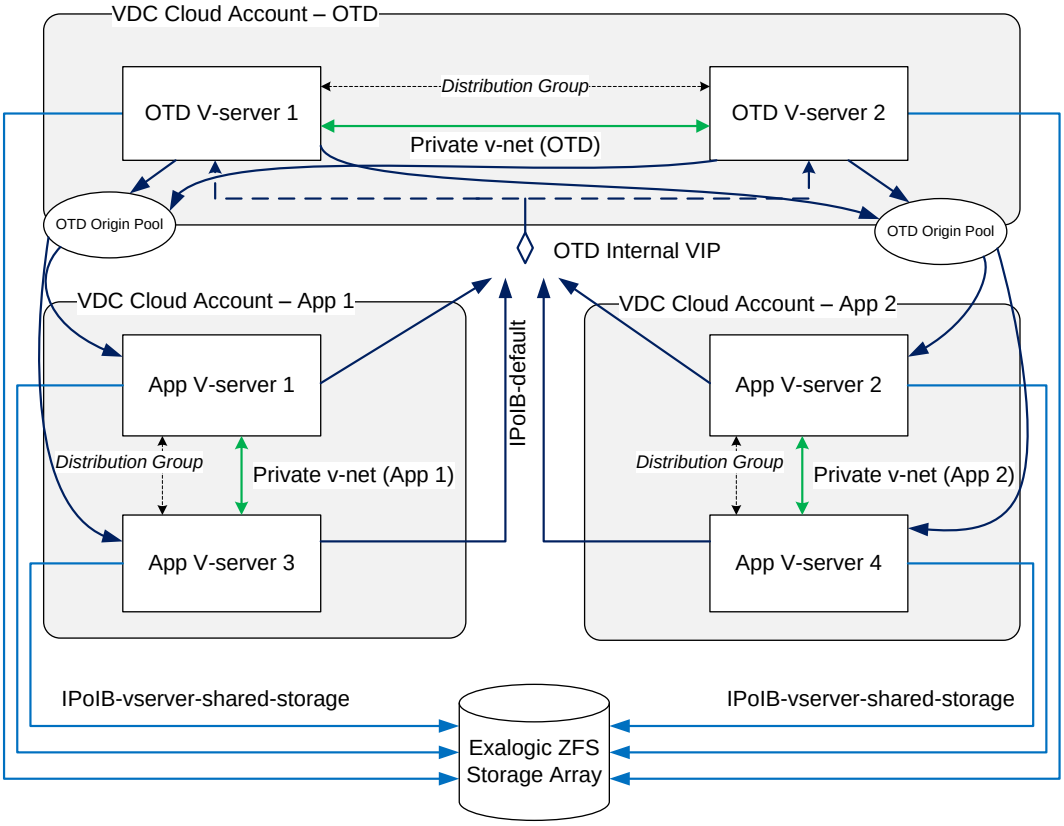

Figure 9 – OTD Internal VIP utilizing multiple cloud accounts

Implementation Notes

1. Both OTD and Application v-servers are provisioned with both IPoIB-default and a private v-net. Because the accounts that own the v-servers differ, so too does the private v-net. Because each private v-net is completely isolated from one another, there is no risk of OTD cluster traffic and application cluster traffic mixing.
2. All traffic between OTD and the application tier goes over the IPoIB-default network, depicted in dark blue. This is capable of traversing cloud accounts. There are actually 2 separate flows

depicted above, one for the OTD → Origin Servers, and one with Applications contacting each other over OTD's internal VIP.

3. Private v-nets are provisioned for each cloud account and v-server contained within, and these are used for internal cluster communications for each product.
4. As with the previous cases, all binaries and instance data for OTD is stored on ZFS. This requires all v-servers to be provisioned for the IPoIB-vserver-shared-storage network as well.
5. It should be noted that since physical devices (compute nodes, ZFS storage, Exadata if present) have addresses on this subnet, security best practices to control access to physical devices should be followed.
6. Similar to scenario 1b, the cloud administrator can (and should) reserve the IP address(es) for OTD VIPs using Exalogic Control. This is found under "Networks" instead of "VDC Management", and is carried out by the "Cloud Admin" user on Exalogic.

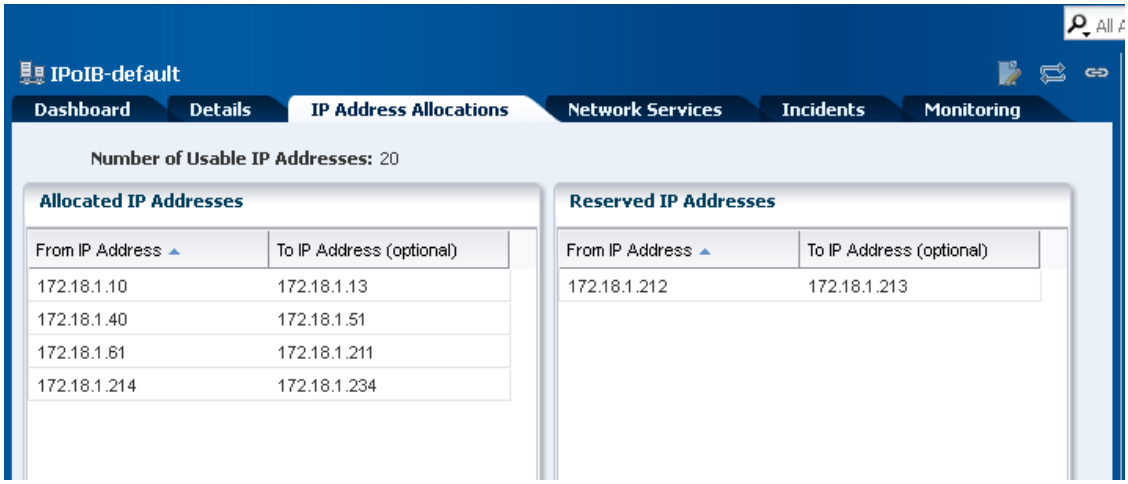

Figure 10 – View of 2 reserved IP Addresses for VIP on IPoIB-default

Scenario 3 – Combined Exadata Topology

As stated in the introduction, a combined scenario with external and internal VIPs often exists. In the following diagram, we combine scenario 1a (external with F5) with scenario 2 (internal VIP). External clients will access a WebLogic cluster via F5 and OTD, while internal clients may use an OTD VIP for load balancing across the same or different WebLogic cluster. Combined scenarios often come into play when using Oracle SOA Suite, or other tiered products that use WebLogic as their base. Such products can surface a number of different access points, running the gamut of JEE technologies, and often requiring many (v-) servers to host components. In these cases, it makes sense to use OTD for both types of communication. And as mentioned earlier, the different OTD configurations are hosted on a shared set of v-servers, and use IPoIB-default to communicate with other v-servers inside Exalogic. Additionally, an Exadata rack is also pictured, as the underlying RDBMS for the applications. The diagram shows access to Exadata via the IPoIB-default network. This completes the picture, and shows how the different networks within Exalogic are used for a complex 3-tier

application. OTD serves as the "gatekeeper", and direct access to application and database components is prevented.

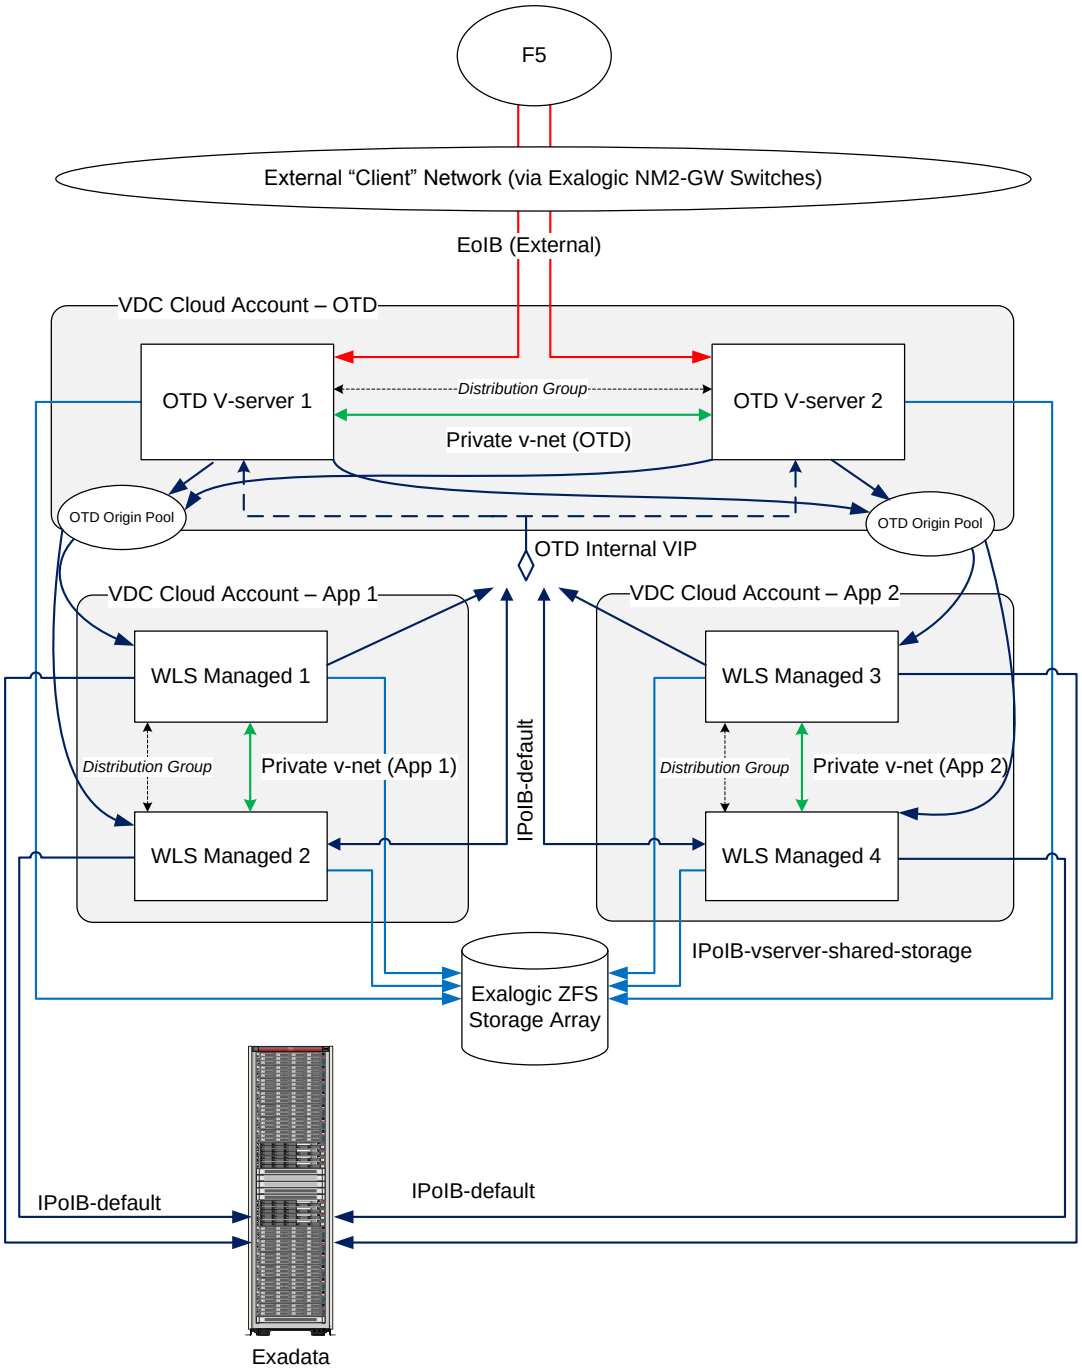

Figure 11 - Combined topology, including Exadata

Implementation Notes

1. As with either previous scenario, the v-servers comprising the application tier could also have EoIB access (either the same or different EoIB network as OTD), depending on whether it is required for the applications. This connection is not shown in the diagram.
2. Each application v-server is provisioned with 3 networks:
    - Private v-net for internal cluster communication
    - IPoIB-vserver-shared-storage for access to ZFS storage
    - IPoIB-default for access to Exadata
3. Additionally, OTD v-servers are provisioned EoIB as well. This allows OTD access from and to the outside world. Access to origin servers hosted on Exalogic goes over IPoIB-default (dark blue).
4. Exadata will host IBVIP listeners on IPoIB-default, for the purpose of application servers on Exalogic being able to utilize the Infiniband connection.

- The application tier v-servers are changed to WebLogic managed servers here. WebLogic can utilize Active GridLink and SDP over the IPoIB-default network.

OTD Installation and Failover Groups

While this paper assumes that the reader is knowledgeable in OTD installation, and has covered the appropriate documentation, there are additional things of note when finalizing the configurations above, and starting to use OTD.

As a general refresher, the process of getting to a working OTD configuration on a virtualized Exalogic should look like this:

1. **Provision Cloud Accounts** – If a separate Cloud Account is desired, do this first. While doing this, pay close attention to network address limits and Exalogic CPU/memory resources. A typical OTD-only account might get 4 vCpu, 16G memory, 2 or 3 EoIB addresses (3 if an external VIP is used), and the ability to create 1 private v-net. This would allow the creation of the OTD assets and networks for any of the scenarios in this paper.
2. **Provision Network** – Choose a scenario and assign IP address(es) to be used for OTD VIPs. Perform steps to exclude VIP addresses within Exalogic Networks, create private v-nets, and allocate static addresses as necessary.
3. **Create Distribution Groups** – One per set of v-servers that will have OTD, WebLogic, or other application tier servers.
4. **Provision V-servers** – Create v-servers from existing (or specialized) v-server types and Exalogic templates.
5. **Software Installation** – Create and mount ZFS shares for OTD, install OTD software, and create instances on each v-server. Pay particular attention to using the correct network for OTD internal communication.

6. **OTD Configuration** – Create a basic configuration, which will consist of a listener, virtual server, and origin pool. The listener's IP or hostname should be consistent with the networks shown in the scenarios in this paper.[4]
7. **OTD Failover Group** – Create and test a failover group using a VIP that has been selected. Test from outside Exalogic (scenario 1b) or from a v-server other than OTD (scenario 2).

Implementation Notes

- OTD will attempt to choose the correct network, based on the address chosen for a VIP. Using the output of "/sbin/ip addr" and the Exalogic Control's Network tab for a v-server, it will be possible to double-check that the correct network has been chosen.
- Even though effort is made here to choose a unique VIP address, and to prevent other v-servers from using it, it is worth doing a ping test on that VIP before starting to create the failover group.
- The "network prefix" requested in step 1 of the Failover Group wizard way need to change, based on the network chosen for the IP address of the VIP. For example, if the EoIB network was created as 10.45.88.0/23, then use 23 for the network prefix. A private v-net with a size of 14 would use a network prefix of 28. OTD may give an error message that contains the correct value, should this be input incorrectly.
- Ensure that the default keepalived service is disabled on any OTD v-server. It may be on by default, and the keepalived process will interfere with that created by OTD for the failover group's internal implementation. Use "chkconfig –list|grep keep" to view the status.
- Ensure that the "Router ID" is unique across all OTD instances. You may need to manually set it to a value other than 255, if you know that other OTD instances exist in your organization. This information should be managed centrally, to ensure uniqueness.

Conclusions

As can be seen from the scenarios discussed in this paper, Oracle Traffic Director can be flexible, just as Exalogic strives to be. Under different circumstances, the recommended topology may change. It is important to have as many facts up front, so that the planning process can go smoothly. This paper introduces several of the most common topologies where OTD can be employed.

There are many other potential use cases and scenarios for which Oracle Traffic Director can be used for, and this paper is simply a starting point. The recommendation is to lay out the topology ahead of time, do the proper planning, and then begin implementation.

---

[4] Specifying a listen address that may not be present (such as a floating VIP address that is only active on 1 v-server at a time), will require the change of a Linux configuration. See here for more details: http://exablurb.blogspot.co.uk/2013/08/limiting-otd-to-listen-only-on-vip.html

Oracle Traffic Director on Exalogic – Configuring VIPs using Internal and External Networks
March 2014
Author: Andrew R. Gregory

Oracle Corporation
World Headquarters
500 Oracle Parkway
Redwood Shores, CA 94065
U.S.A.

Worldwide Inquiries:
Phone: +1.650.506.7000
Fax: +1.650.506.7200

oracle.com

Oracle is committed to developing practices and products that help protect the environment

Hardware and Software, Engineered to Work Together