Oracle® Collaboration Suite

Firewall and Load Balancer Architecture

Release 2 (9.0.4)

Part No. B15609-01

November 2004

This document discusses the use of firewall and load balancer components with Oracle Collaboration Suite Release 2. This document focuses primarily on the use of firewalls and load balancers with the Oracle Collaboration Suite Infrastructure, using the Oracle Application Server Containers for J2EE (OC4J) component as an example. Similar considerations apply when using firewalls and load balancers with other components of Oracle Collaboration Suite, such as the Oracle Email server.

This document discusses the following topics:

Introduction

Successful deployment of Internet applications requires integration of Oracle products with firewalls, load balancers, and Secure Sockets Layer (SSL) accelerator products. This document outlines possible configurations of Oracle Collaboration Suite Release 2 components when used with firewalls, load balancers, and SSL-accelerating devices. It also explains criteria for selecting different architectures and components for applications deployed on the Internet. This document attempts to cater to varying customer requirements.

ORACLE

Overview and Disclaimers

The focus of this document is to detail architectures and methods that can be used to deploy Oracle Collaboration Suite Release 2 components with firewalls and load balancers for Internet-accessible applications. In this discussion, specific Oracle and vendor products from Radware, F5, Check Point, Cisco, Sonic Wall, and Nortel are referenced. Oracle performs tests with assorted load balancers, firewalls, and so on, and often the results of such tests have resulted in fixes and enhancements in these products. Oracle believes that these fixes and enhancements are all available from the vendors discussed in either currently shipping products or in upgrades available from the vendor. All Oracle issues have been remedied in initial versions of the Oracle Collaboration Suite Release 2.

While Oracle has performed tests to determine acceptability of these firewall and load balancing products, it does not guarantee these products or recommend them versus other products that may or may not have been tested at Oracle's laboratory. Also, while Oracle believes that these configurations will satisfy a large percentage of Oracle's customer base, Oracle is not making claims regarding the suitability of these products for specific customer situations. Lack of reference to certain vendor's equipment should not be taken to imply that Oracle would not recommend them. If vendors would like their products tested by Oracle in future releases, then they should contact the Oracle product management team.

Overview of General Architecture

Oracle Collaboration Suite components are designed and tested to ensure they work properly in Internet environments. This section provides information that can be used to configure networks, which are enabled for Internet access, with Oracle products.

Prior to discussing specific, tested configurations, it is important to discuss the general recommendations for network architectures suitable for the deployment of Internet-accessible applications. Generally, the common architecture comprising of the Internet, firewall, demilitarized zone (DMZ), firewall, and intranet components is recommended. A simplified representation of this architecture is provided in the following figure.

Figure 1 Traditional DMZ View

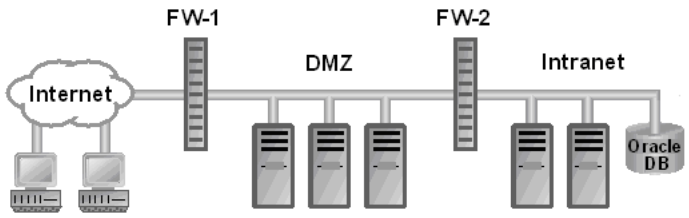

In this document, the DMZ consists of parts of the customer's network that are between the customer's intranet and the Internet. Sometimes, this network zone is called the perimeter network. The DMZ need not always be a simple one-segment LAN as is implied by Figure 1.

Several recommendations are made concerning DMZ when used for deploying Oracle products:

1. Hardware on DMZ is connected through switched connections and not bussed connections. With switched connections, only the sending and receiving devices can view the exchanged messages. With bussed connections, devices that are not involved in the legal exchange of messages also can view the exchanged messages.
2. The firewall between the Internet and the DMZ does not allow incoming Internet traffic that has sender addresses of DMZ hardware.
3. The firewall between the DMZ and the intranet only allows messages coming from the DMZ to the intranet that have DMZ sender addresses.
4. Even though Figure 1 shows two firewalls, F1 and F2, this setup can actually be implemented with a single or many components of hardware in such a network connection scenario.
5. There exists recommended hardware, as described in Figure 2, that can provide switched connections between DMZ hardware and transit rules that might differ for each pair of attached devices, including Internet and intranet devices. The CSS 11000 series of products from Cisco or WSD from Radware and Fireproof series of products are examples of such hardware. These rules can specify that only certain protocols and ports can be used for message interchange between a specific pair of devices. For example, such a configuration might allow System A to send HTTP messages only to System B on port 8080. While such hardware is recommended, use of multiple firewalls is also acceptable.

The following figure shows this recommended architecture.

Figure 2 Switched Connection DMZ

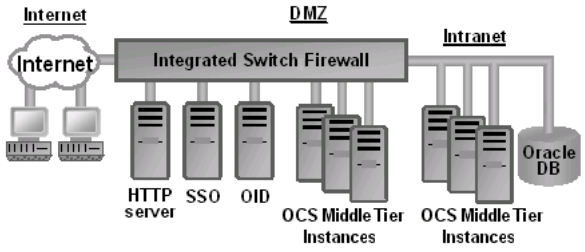

As in Figure 1, Figure 2 also represents the architecture comprising the Internet, firewall, DMZ, firewall, and intranet components. However in this case the Integrated Switch Firewall provides switching rules for each pair of connected components (Internet, intranet,Oracle HTTP Server, OracleAS Single Sign-On Server, and so on). Note that the Integrated Switch Firewall can consist of a single or multiple numbers of devices.

Firewalls

Firewalls are the main defense technology for sites providing Internet access. Different firewall products vary considerably in features and performance. Suitable firewall usage can protect against many common vulnerabilities by prohibiting Internet access to services such as File Transfer Protocol (FTP) and remote shell (`rsh`), especially if these services run on Internet servers.

Firewalls are devices that restrict access between different LAN segments for security. Firewalls perform this function by analyzing traffic and can restrict

communication based on IP address, port, protocol used, protocol transitions, and message content. For example, Check Point Firewall-1 products, that are tested by Oracle, provide a software solution that includes a feature called **Stateful Inspection**, which can restrict access based on illegal Internet protocol transitions. Cisco's PIX, tested by Oracle, is an example of an integrated hardware-software firewall solution.

Some firewalls are software products that are loaded onto client or server computers. These products might be useful to some extent, but are inadequate for corporate firewalls that should always be deployed on separate systems than those deploying application or infrastructure software.

Load Balancers

Load-balancing hardware is used to provide both scalability, by spreading load across multiple processors, and fault tolerance in case of processor failures.

Load balancers have two essential functions. The first is to balance the load of traffic across multiple servers, which results in better scalability. In high traffic situations, a load balancer can prove to be very important. An example of a tested load balancer used in this way is F5's BIG-IP product.

The second function of load balancers is to provide fault tolerance for servers. A load balancer ensures that a single failing server does not result in the loss of a critical resource. It accomplishes this by routing new requests to an alternate server if one server fails.

Typically, load balancers are able to route traffic in both situations where the Infrastructure keeps application state and where the traffic is stateless. In the case of stateless communication, the load balancer can route traffic to any managed server because no server needs to maintain state to be able to correctly process the message. This is generally more efficient because requests can go to the least busy server. However, a stateless operation often puts an unacceptable burden on application writers. Many Oracle products require that the Infrastructure maintain application state.

The term, *sticky* or *persistent* transaction is often used to denote a transaction that should be routed to a particular, load balancer-managed hardware containing an intermediate application transaction state.

For transactions where Infrastructure keeps state, load balancers switch incoming messages to the server maintaining the state. Switching criteria are determined by analyzing cookies, headers, or other attributes. Sometimes only a single server might contain the state. In this case, a processor failure results in the failure of all transactions that have state in the failed processor. All these transactions must be restarted.

In some situations, there are preferred processors but all processors can obtain the state. When failures occur in these situations, a redirect due to failure results in successful processing. However, this might create an added overhead for transactions that had state in failed processors.

SSL Accelerator Appliances

In many sites, SSL key exchange operations can dominate CPU usage. For such sites, the use of an SSL accelerator device can result in significant cost reduction and improved performance.

As a part of this project, devices that provide accelerated SSL handling are included for the following reasons:

1. SSL accelerators are often part of, or included with, load-balancing equipment, such as the BIG-IP SSL accelerator.
2. SSL accelerators are often necessary to improve scalability or reduce costs. An example of such an accelerator is the Sonic Wall Accelerator when used in conjunction with Cisco equipment.
3. SSL Accelerators can be deployed as part of a scalable solution with load-balancing equipment. Such a solution provides added performance and scalability. An example of such a solution is the Radware Web Server Director and CT-100 SSL Accelerator.

Using HTTPS extensively improves security. Where HTTPS use is limited by performance considerations, SSL accelerators should be considered.

There are different types of SSL accelerators. One type is a math coprocessor that offloads expensive cryptographic operations from general purpose CPUs (none tested). A second type is a standalone device that converts HTTPS-to-HTTP protocols. That is to say, it converts incoming HTTPS protocols to HTTP. As the SSL processing of the HTTPS protocol can consume a large percentage, or even most, of a CPUs time, offloading SSL processing may result in a significant reduction in the number of CPUs required to support a workload. Such reduction can result in both cost savings and improved performance, especially in cases where caching can be improved by reducing the number of caching servers.

A problem occurs with HTTPS-to-HTTP converter appliances when client-side X.509 certificates are used. This is because these appliances terminate the SSL session and there is no standard way to provide the client-side X.509 certificate information with the forwarded message. If client-side certificates are only used to allow and deny access to a site or virtual host, then this solution is acceptable. However, if the application or other infrastructure items need certificate information, then custom solutions are required.

As client-side certificates are infrequently used at this time, this consideration is not important for most sites. Customers interested in the use of X.509 client-side certificates with such devices can contact Oracle or appliance providers, because progress toward standard supported solutions is being made.

Deployment and Test Models

Testing was performed at Oracle labs using various Oracle Collaboration Suite and vendor components.

The following section details specific recommendations for the deployment of the core Oracle Collaboration Suite components. In the next section, specific tested configurations of other Oracle Collaboration Suite components are presented. Various configurations have been tested in Oracle labs including the following hardware:

- Check Point Firewall - 1 NG (Installed on Dell GX - 1)
- Cisco CSS 11050
- Cisco Catalyst 6506 W/Content switching module blade
- Cisco PIX 520

- F5 Networks BIG-IP 520 and 540
- Nortel Alteon ACEdirector
- SonicWall SSL-R3
- Radware WSD, CID Fireproof, Linkproof and CT-100

The following observations were made when the tests were performed:

- Persistence: For the Radware, F5, Nortel ACEdirector, and the Cisco CSS 11050, active HTTP cookies inserted into the header were used. The expiration on these cookies was set to zero. This makes them session cookies.
- Balance method: Each of the devices tested has several different traffic-balancing algorithms. Only Round Robin rotation was used for these tests.
- Proxies: Both the Sonic Wall SSL-R3 and the embedded F5 HTTPS accelerator function as proxies with similar application configuration rules.
- Pools: Load-balancing pools consisted of two or three balanced devices.
- Load: Testing was done at light (10 simulated clients) and moderate (more than 50 simulated clients) loads. These tests were not designed to determine maximum loads.

Core Components Associated with Oracle Collaboration Suite Release 2

There are many acceptable configurations for the core components associated with Oracle Collaboration Suite including Oracle Application Server Web Cache, Oracle HTTP Server, Oracle Internet Directory, Oracle Application Server Single Sign-On, a variety of Middle Tier protocol servers, and OC4J. These configurations vary depending on the components used, degree of scalability required, security requirements, and high-availability requirements. Representative and tested configurations are provided in the following sections. Note that these configurations assume connections as described in the earlier sections of this document.

Figure 3 shows a recommended configuration for Oracle Collaboration Suite core components for situations not having special high-availability or scalability requirements. In Figure 3, all Oracle Collaboration Suite components reside in the DMZ with the exception of one Oracle Collaboration Suite Middle Tier Instance.

Figure 3 Core Components Associated with Oracle Collaboration Suite

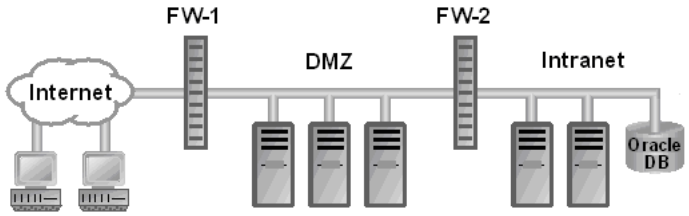

One rule for providing acceptable intranet security requires that all incoming messages first be processed by devices on the DMZ before they are forwarded to the intranet. This is for fault containment reasons. When attacks compromise

DMZ attached devices, damage is constrained to DMZ devices rather than the entire intranet. Successful attacks of intranet devices would be far difficult to constrain.

The reason for recommending switched connections over bussed connections for the DMZ is that protocols between these devices might not be encrypted. For example, the Apache Jserv Protocol (AJP) between Oracle HTTP Server and OC4J is not encrypted. So with switched connections, a successful intrusion into one DMZ device does not allow that device to read another DMZ communication. The reason for using switching equipment that can enforce security restrictions between each connection, is that it can limit the damage during successful intrusions.

High availability is very important for many applications. Load-balancing products are important for providing high availability for Internet applications.

Note: For high availability, critical resources such as OracleAS Single Sign-On and Oracle Internet Directory might need to be deployed on the same hardware. In this case, reliability improves because less hardware is required for successful operation after redundant CPUs are ignored.

In Figure 3, the Oracle Collaboration Suite Middle Tier Instances reside both on the DMZ and on the intranet. Switching equipment with pair-wise rules or multiple firewalls is assumed to be existing. For example, in the order, Internet, firewall, Oracle HTTP Server/ OracleAS Single Sign-On, firewall, Oracle Collaboration Suite Middle Tier Instance/Oracle Internet Directory, firewall, and intranet.

OracleAS Web Cache

OracleAS Web Cache with HTTPS accelerators can significantly improve the performance of many applications.

When OracleAS Web Cache is added, it is placed as the front end for Oracle HTTP Servers. These connections are switched connections, so there is no concept of front or back. However, if switching equipment with pair-wise rules is used, then OracleAS Web Cache should be placed as the front end for the Oracle HTTP Server.

When OracleAS Web Cache is configured with a load balancer, the address of the load balancer is used when configuring URLs, through which OracleAS Web Cache will be accessed. Note also that other components in the network such as OracleAS Single Sign-On, Oracle HTTP Server, and Oracle Internet Directory may have associated hardware load balancers.

As mentioned earlier, it is important to provide high-availability configurations of OracleAS Single Sign-On and Oracle Internet Directory because these are potentially used by many application subsystems and infrastructure sets at a particular site. Therefore, even where scale issues are not important for OracleAS Single Sign-On or Oracle Internet Directory, load balancers should be considered for these components for high-availability reasons.

OracleAS Single Sign-On

Figure 4 represents a recommended load-balanced configuration for OracleAS Single Sign-On Servers. In this configuration, the OracleAS Single Sign-On Server is attached to the DMZ because of access required by Internet devices.

Figure 4 High Availability OracleAS Single Sign-On

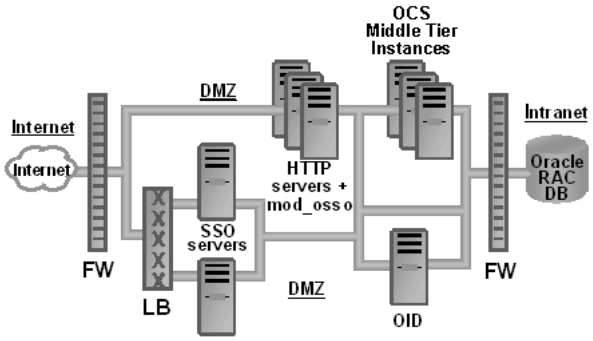

This load-balanced configuration for OracleAS Single Sign-On Server does not include multiple databases, so an Oracle Real Application Cluster configured database is recommended for improved availability. Oracle Internet Directory servers are not in a high-availability configuration in this figure, but such configurations are described later in the document. OracleAS Single Sign-On was tested with `mod_osso` in Oracle HTTP Server and also with OracleAS Portal as described in the following sections.

HTTPS-to-HTTP Appliances

HTTPS-to-HTTP appliances can either route to separate Oracle HTTP Server hardware for secure transactions or to separate ports on hardware that provides both secure and nonsecure transaction access. These appliances are recommended where performance of HTTPS processing is important. Figure 5 provides an architecture where the accelerators are used in front of the Oracle HTTP Servers.

Figure 5 HTTPS-to-HTTP Appliance

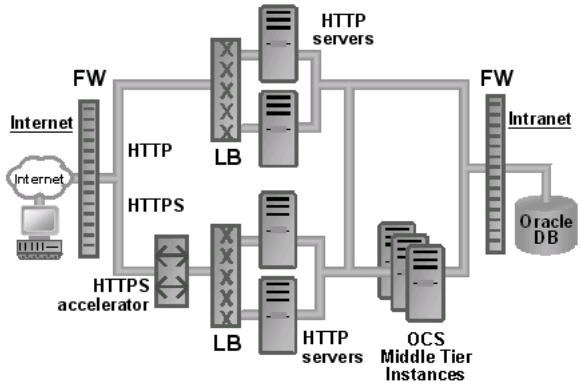

In this figure, the HTTPS accelerator is shown as a separate device from the load balancers. This would be a tested configuration used with Sonic Wall. In the case

of F5 tested configuration, the HTTPS-to-HTTP accelerator is part of the load balancer. Note that neither HTTP processor runs SSL.

An alternative architecture could route the output of the HTTPS accelerator to the normal HTTP servers, but use a different port address, so that the server can distinguish messages that are actually secured by SSL from those that are not.

Oracle Application Server Portal

OracleAS Portal is a key Oracle technology that is used by many applications and leveraged by user applications. Firewalls and load balancers are important to providing secure and well-performing user interfaces for Portal applications.

OracleAS Portal was extensively tested for main-line runtime facilities and provisioning of users, DB Explorer, Portlet Builder, Content Management and Page User Interface, and Providers. Portal deployment is detailed in Figure 6 and requires OracleAS Single Sign-On, Oracle Internet Directory, Oracle HTTP Server, and OC4J. OracleAS Portal's Parallel Page Engine (PPE) runs on OC4J. OracleAS Single Sign-On is used for login access, and it calls Oracle Internet Directory for user information. Here `mod_osso` is not used by OracleAS Portal (refer to Figure 4) because it provides its own `mod_osso` equivalent through a Single Sign-On tool kit.

Figure 6 OracleAS Portal

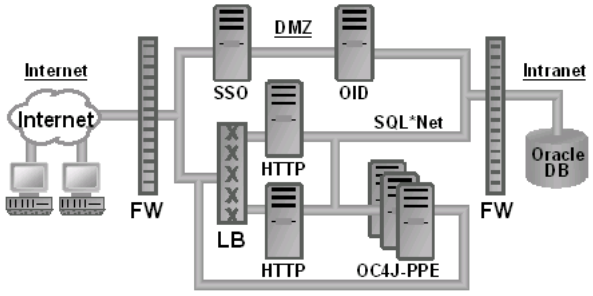

Figure 6 represents the OracleAS Portal Architecture tested. Oracle HTTP Servers are load-balanced and forward OracleAS Portal requests to the PPEs running in OC4J. The PPE forwards these requests to the Providers through HTTP and makes database requests through `mod_plsql` to the database. The PPE is configured to access the Oracle HTTP Servers through the load balancer URL.

Oracle Application Server Enterprise Manager

OracleAS Enterprise Manager was used to test configuration and management of the products for this document. These products include Oracle Internet Directory, Oracle HTTP Server, OC4J, Jserv, and OracleAS Web Cache. It worked well for this purpose.

Oracle Internet Directory

Oracle Internet Directory is a key core technology where lack of scalability can degrade numerous applications and poor availability can completely stop business processing.

The scalability and availability provided by load balancers for Oracle Internet Directory is often the key for acceptable application deployment.

Figure 7 represents a high-availability Oracle Internet Directory configuration. This configuration was not tested, but it is included here because many clients will want such configurations for high availability and load balancing.

Figure 7 High-Availability Oracle Internet Directory Configuration

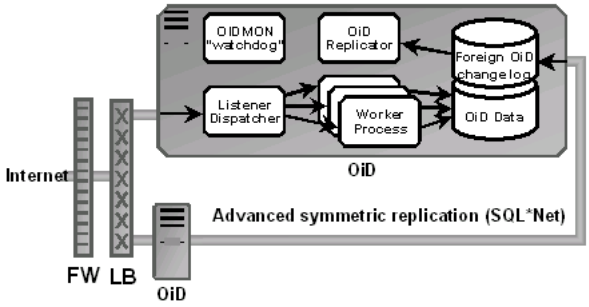

Note that Oracle Internet Directory provides for separate databases for each Oracle Internet Directory instance. These separate databases can be synchronized with each other using Advanced Symmetric replication and Oracle Internet Directory facilities for processing replicated data. This provides excellent fault-tolerance features. Also, this configuration can be used for worldwide deployment to ensure fast access from different parts of our planet.

Summary

Oracle has configured and tested recommended configurations for the core Oracle Collaboration Suite components with a number of popular firewalls, load balancers, and HTTPS-to-HTTP appliances. Based on the results of these tests, Oracle believes that Oracle Collaboration Suite components work well with these products, the use of which is mandatory for Internet deployment of Oracle products.

Tests have found problems with both Oracle and vendor products. These problems have all been remedied, and the remedies shipped in initial production versions of Oracle Collaboration Suite Release 2. It is believed that currently shipping products of the vendors referenced in this document also include the corresponding remedies, or such remedies are available from the product vendors.

Future Work

It is planned that integration testing of Oracle Collaboration Suite components with firewalls, load balancers, and HTTPS-to-HTTP appliances will be extended to provide greater coverage of current and future versions of Oracle products as well as to include additional vendors products.

Additional Information

Further discussion regarding the Internet-DMZ-intranet architecture is discussed in the Best Practices for HTTP Security document, which can be found at

http://technet.oracle.com/products/ias/techlisting.html

This paper also includes important best practices for other areas of HTTP security.

Acknowledgements

Oracle wishes to thank F5, Check Point, Cisco, F5, Nortel/Alteon, Sonic Wall, and Radware for the donated use of their products.

Documentation Accessibility

Our goal is to make Oracle products, services, and supporting documentation accessible, with good usability, to the disabled community. To that end, our documentation includes features that make information available to users of assistive technology. This documentation is available in HTML format, and contains markup to facilitate access by the disabled community. Standards will continue to evolve over time, and Oracle is actively engaged with other market-leading technology vendors to address technical obstacles so that our documentation can be accessible to all of our customers. For additional information, visit the Oracle Accessibility Program Web site at

http://www.oracle.com/accessibility/