Building Highly Available Database Servers Using Oracle Real Application Clusters

An Oracle White Paper
May, 2002

ORACLE

Building Highly Available Database Servers Using Oracle Real Application Clusters

EXECUTIVE OVERVIEW

High Availability (HA) is becoming a must-have requirement for e-businesses that cannot afford system down time. The Internet makes it easier to reach customers around the world and around the clock. Since companies must always be prepared to serve their customers, this expanded reach also makes it more costly when service is not available, either planned or unplanned.

Oracle Real Application Clusters are the multi-node extension to Oracle database server. They enable e-businesses to build a multi-node database server that are highly available and highly scalable. An e-business does not need to sacrifice scalability or performance for high availability with clustered database.

This white paper presents the architecture that makes Oracle cluster databases highly available while functioning as highly scalable database servers. It starts with the basic availability issues faced by a generic cluster system, followed with a discussion on how Oracle cluster database addresses those issues. The paper then discusses enhancements in Oracle cluster databases which distinguish them from generic cluster systems on availability capabilities. The paper also discusses limitations associated with using cluster systems as HA solutions and gives suggestions on how to address them.

This white paper also addresses practical issues with using Oracle cluster databases, including planning, installation and configuration. It discusses options for applications to use Oracle cluster databases so that the applications can be made more available when building a highly available application system.

This paper also shows why Oracle cluster database, with active instances on all nodes, is also a much better choice for database fail-over solution than those offered at operating system level for generic application fail-over.

This paper demonstrates how Oracle cluster databases are used to save costs and expand into new markets in two typical high availability cases. First, VeriSign uses Oracle cluster databases as the back-end databases for its various Internet based service offerings, including certificate subscription and Internet payment services. Second, Oracle uses Oracle cluster database to consolidate its global email systems that makes it much more robust and performant. Oracle also uses the same architecture for its email hosting business.

Finally, this paper discusses new HA features of the upcoming release.

ARCHITECTURE OF REAL APPLICATION CLUSTERS

Typical Configurations Supported by Real Application Clusters

Real Application Clusters run on top of a hardware cluster. A cluster is a group of independent servers (nodes) that cooperate as a single system.

The primary cluster components are processor nodes, a cluster interconnect, and a shared storage subsystem. The nodes share access to the storage subsystem and resources that manage data, but they do not physically share main memory in their respective nodes. Oracle cluster database combines the memory in the individual nodes to provide a single view of the distributed cache memory for the

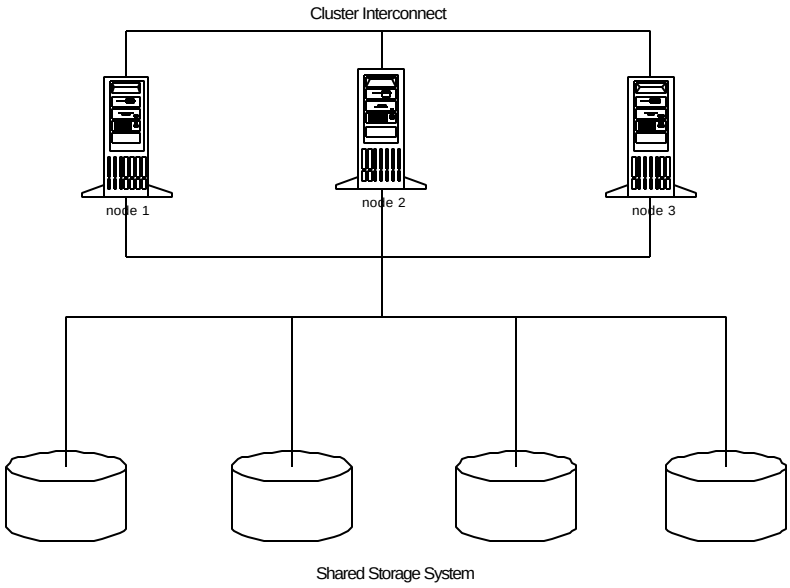

Figure 1. A Cluster Comprises Processor Nodes, The Cluster Interconnect, and A Disk Subsystem

entire database system.

A node can be made up of multiple processors. A common type of node is a Symmetric Multi-Processor (SMP) node. Each node has its own dedicated system memory as well as its own operating system and database instance.

Advantages of Clusters

The benefits of building a database with a cluster of smaller nodes than with a single larger node are-

- Flexibility and cost effectiveness in capacity planning, so that a system can scale to any desired capacity
- Fault tolerance to partial failures within the cluster, especially the node failures.

Scalability

Oracle cluster database gives users the flexibility to add node(s) to the cluster as the demands for capacity increases, scaling system up incrementally to save costs in capital investments and eliminating the need to replace smaller single node systems with larger ones. It makes capacity upgrade process much easier and faster since, in most case, one or more nodes with similar or identical configuration are added to the cluster, compared to using completely new and larger nodes to upgrade systems.

The Cache Fusion technology implemented in Oracle Real Application Clusters enables capacity to be scaled up close to linearly. The Cache Fusion technology is described in details in the Cache Fusion technical white paper.

High Availability

Another main advantage of the cluster architecture is the inherent fault tolerance provided by multiple nodes. Since the physical nodes run independently, the failure of one or more nodes should not affect other nodes in the cluster. In the extreme case, a cluster system can still be available even when all but one node survives, making a system based on cluster highly available. This architecture also allows group of nodes to be taken off-line for maintenance while the rest of the cluster continues to provide services online.

Oracle cluster database takes full advantages of this inherent fault tolerance architecture to provide a highly available database server.

The following sections describe how Oracle cluster database achieves high availability for Oracle databases.

Availability Framework of Cluster Systems

Introduction

To take full advantage of the fault tolerance afforded by cluster architecture, Oracle cluster database enables the Oracle database server to function in the face of various failure scenarios in the cluster. Furthermore, Oracle cluster database is able to recover failed nodes while the database server is online.

Before going into details of the availability framework for cluster systems, we need to understand the differences between availability within a single node and within a cluster. In a single node Oracle database system, availability refers to the ability to survive various application and operation failures within the database instance. In the extreme case of the failure of the node, availability refers to the ability to recover the database to a transaction consistent state as fast as possible. Extensive discussion on this subject is provided in the white paper for Oracle Database High Availability.

For a cluster system, aside from handling failure scenarios in a single node, it needs to handle failure scenarios associated with a node, a group of nodes or

network, while providing required performance. Oracle cluster database builds on top of the fault tolerant capabilities of the single instance Oracle database and enhances the database server to handle failure scenarios unique to a cluster system.

Common Failure Scenarios in a cluster system

The flexibility of a cluster system for each node to function relatively independently does come with some unique problems the cluster system must deal with when failure occurs.

Fault Isolation

The cluster system maintains a consistent system image at all times, especially during failures of individual nodes or the cluster interconnect. The biggest challenge for the cluster system is to be able to quickly and reliably isolate faults and take corresponding actions.

For example, in the case of network failure, a cluster with many nodes can end up as isolated groups of connected nodes. The cluster system must be able to decide that this condition is due to network failure and be able to decide which connected group of nodes (sub-cluster) should continue to operate for the cluster, and which should be temporarily retired from the cluster. This decision is critical to prevent the cluster system from developing the "split brain" syndrome in which different isolated groups of connected nodes (sub-clusters) all claim to represent the whole cluster and all working on the same set of data, unaware of each other's continuing existence.

Recovery

Since a failed node or nodes may contain global information to the whole cluster, the cluster system must be able to re-construct the information as quickly as possible. It can either maintain a hot standby repository for the global information, or re-create the information from the live nodes and the information stored in the storage system. Oracle cluster database's fast recovery is especially crucial to maintain high availability for Oracle database servers. This fast recovery capability is discussed in more details later in the Oracle cluster database recovery section.

IO Fencing

IO fencing refers to the situation that occurs when the left-over write operations from failed database instances (cluster function failed on the nodes, but the nodes are still running at OS level) reach the storage system after the recovery process starts. Since these write operations are no longer in the proper serial order, they can damage the consistency of the stored data. Oracle cluster database utilizes facilities provided by the underlying cluster system to prevent this.

Limitations of the Cluster System's Availability Capability

As highly available as a clustered system is, it is not the solution to all availability problems. A cluster's availability is mostly a system solution that enables a system to function when some of its components fail, such as its instances, physical nodes, or cluster interconnect. To maintain a highly available system, other precautions must be taken.

For example, redundant hardware, such as separate power supplies for each of the sub components or redundant interconnect hardware, can help to avoid a problem caused by a failure of one of these components.

A cluster system cannot prevent failures resulting from operator errors either. Every precaution much be taken to guard against these user errors, just like in the case of a single database instance.

Two types of these additional failures need special attention.

Disaster Recovery

Cluster systems can not be used to guard against disasters that completely shut down a site. Disasters such as earthquakes, fires, power outages or flooding that destroy a physical site can not be guarded using cluster systems at the same site. Solutions for database disaster recovery such as a standby database, either physical or logical, provide the right protection against such failures. For details about Oracle's offerings in this area, please refer to relevant white papers.

Storage System Protection

Since the one of the intrinsic functions of a cluster system is for cluster nodes to utilize a shared storage system, the cluster system itself does not guard against failures of the shared storage system. Even though a plain bunch of disks will work with a cluster system, a more sophisticated storage system that is highly reliable and available should be used as the shared storage sub-systems for cluster systems to guard against storage related failures.

Availability Framework of ORAC

Architecture

Approximately depicted in *Figure 2*,Cluster ware is the component that provides the generic cluster functions at the operating system level. Cluster ware monitors the cluster and handles group membership related events, such as when to include a node into or exclude a node from the cluster when a node starts or fails.

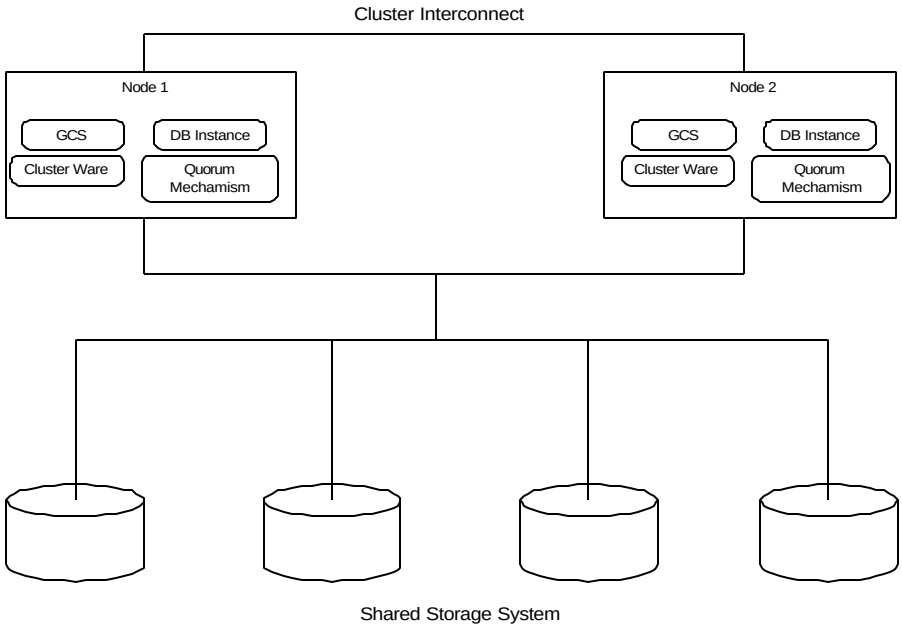

Figure 2. Oracle cluster database software components

The Quorum Mechanism is used by Oracle cluster database to enhance failure detection at database level. For example, cluster ware does not know if a database instance crashes.

Cluster Ware in conjunction with Quorum Mechanism provides reliable fault detection and isolation at both OS level and database level. Fault scenarios are more complicated than a generic cluster system with database as part of it. For example, when a network failure occurs, both Cluster Ware and Quorum Mechanism need to work together to ensure both the right nodes and instances are chosen to act for Oracle cluster database. Another example would be when a database instance fails, Quorum Mechanism and Cluster Ware need to coordinate to make sure no left over writes from the failed instances reach the storage system after the recovery process starts.

GCS (Global Cache Service) insures that a consistent single database image is maintained. It maintains the consistency of the database at the cluster level. Database blocks accessed concurrently by cluster nodes have corresponding GCS

resources to insure the same data block is not updated without coordination by different nodes. When the information is lost during to node failure, it must be reconstructed first before anything else within the database can proceed.

Database Recovery in ORAC environment

The Database Instance relies on all of these components to implement instance recovery for failed instances, in addition to handling normal database operations.

When a database instance or a node in the cluster fails, Oracle cluster database needs to do recovery for the database just as it does for a single instance database. Since other nodes in the cluster is still providing services to the clients, Oracle cluster database makes sure the recovery time is as little as possible through concurrent recovery operations.

The global cache service (GCS) maintains global cache resources status to insure the consistency of database. If a node fails, it needs to rebuild the global cache resource information. Recovery cannot start until the GCS finishes rebuilding the information. Effectively, the whole database is "frozen" during this time. Since the cache resources for database blocks are distributed among the cluster nodes, the time needed for rebuilding GCS information is minimized. Only the cache resources that reside or are mastered by the GCSs on the failed nodes need to be rebuilt or re-mastered. This phase takes only a few seconds on average.

Furthermore, Oracle9*i* uses a two-pass database recovery scheme, where the first pass of redo log scan decides data blocks to recover and then the second pass only accesses the marked blocks to speed up instance crash recovery. Oracle9*i*, can initiate the first pass of this recovery process concurrently with GCS rebuild process. After the first pass, database is made available for service for data blocks that are not impacted by the failure.

Oracle9i also gives you the ability to specify the amount of time a recovery should take, which eliminates potential problems caused by uncertainty about time needed for recovery.

Fail over of TCP/IP Connections

To ensure fast fail over of clients connecting to the failed nodes/instances, database application clients can connect to Oracle cluster database through Oracle Net using load balancing and application fail-over. Client connection load balancing distributes client connections to all nodes of the clusters, alleviating the impact of an individual node failure. With the Transparent Application Failover (TAF) option, Oracle Net will re-connect the failed connections to the fail-over node in the cluster without the client even being awareness of the failure.

Today, most client connections are made using the TCP/IP networking protocol. Since TCP/IP is a reliable communication protocol for wide area network, it ensures every message is either reliably delivered or a failure event is generated.

This can cause problem for Oracle cluster database during node failure. In some condition, if a message is sent to a node that fails and no response is generated back to the client before the failure, the TCP/IP protocol will enforce a time-out period of up to a couple hours before reporting the error condition. This scenario can leave the impression of a hanging client.

To address this problem, Oracle cluster database floats the IP address of the failed nodes to a live node as part of the recovery procedure, ensuring that requests even to the failed node have a destination so that clients do not need to endure long time-outs.

Online Configuration

To minimize the impact of Oracle cluster database configuration on availability, nodes can be added to or taken out of an Oracle database cluster without the database server being shut down.

Though DBCA (database configuration assistant), a new instance can be added to or deleted from an existing Oracle cluster database cluster while the database is online. The GUI-based utility makes a complicated operation of adding and deleting a node online much simpler.

BUILDING HIGHLY AVAILABLE DATABASE SERVERS

Understanding the availability features of Oracle cluster database lays a common ground for understanding the issues involved in constructing highly available Oracle9*i* database servers. This section discusses the practical aspects of building database systems using Oracle cluster database.

The information presented here is intended to serve only as general guidelines. For detailed, step by step procedures, please refer to the Oracle cluster database document set.

Availability Requirements

The most critical steps in constructing a highly available database server are to have a clear objective, to set the right expectations, and to collect the right requirements.

As discussed in previous section, Oracle cluster database is ideal for providing system availability solutions when a site must be operational even if some nodes of the clusters are taken off-line due to maintenance or failures. Using Oracle cluster database to provide scalable database services while making it resilient to partial system failure is the best use for Oracle cluster database. Online stores, general web sites, corporate databases, and most Web based business portal fit into this category.

If the requirement is for the system to survive disasters, Oracle cluster database alone is not enough. A standby backup system that is not physically co-located can

provide extended protection from disasters. One Oracle cluster database system acts as the main server while the other Oracle cluster database system acts as the standby in another location. These systems that are loosely coupled through database standby operations can survive most disastrous events. Considerably more resources are needed for this type of availability requirement. Highly critical applications such as banking systems or airline reservations that are extremely time sensitive may require this type of availability.

Installation Considerations

Cluster Systems

Since cluster systems require additional hardware and software components than a single node, users are advised to pay special attention during installation. In addition to the nodes and the system software required by single instance Oracle database, a cluster requires an interconnect network to link the cluster nodes together and a storage system that is shared accessed by all nodes in the cluster. In addition, most cluster system vendors provide cluster software to manage group membership of the cluster and to monitor the cluster. This adds some complexity to the installation process before installing Oracle cluster database.

To ensure a smooth installation, Oracle works with system vendors to certify cluster platforms as well as storage systems that support Oracle cluster database. Specific certification information is published on www.oracle.com under the Real Application Clusters section. System vendors should be the first line of contact for hardware configuration inquiries.

Peripherals

Considerations should be taken for preparation of the peripherals as well. For example, separate power sources should be used for separate nodes if possible to prevent a single power source failure crashing down the whole system.

We also strongly recommend using a highly available storage system with Oracle cluster database when high availability is the primary concerns for using Oracle cluster database. As pointed out earlier, a cluster system itself does not guard against storage system failure, it relies on the storage system to provide guards against disk failures.

In addition, you should use redundant networking links for the network interconnect when supported by the vendors. This practice can circumvent a hardware-based problem with the interconnect which could disrupt operations.

Client Connections

Connecting clients to Oracle cluster database is the same as connecting to a single Oracle database. No changes are needed to move clients from single database to Oracle cluster database based Oracle databases. This makes migrating to Oracle

cluster database easier. In addition, there are a few connection options that can make the applications more failure resilient.

Connection Fail-over

Connection fail-over uses a low overhead connection option to fail over connections in the event of a failure. When the connection fails during the connect time, applications can fail over the connection to another live Oracle cluster database node transparently. This option prevents applications from trying to connect to failed nodes repeatedly. It does not preserve states for sessions or queries. Applications need to handle the failure recovery during queries or transactions.

Transparent Application Fail-over (TAF)

Applications using the Transparent Application Fail-over option to Oracle cluster database can fail over transparently to other Oracle cluster database instance when the instance they connect to fails for most cases. Currently, session fail-over and SELECT operation fail-over are supported. In progress queries can be picked up where they left off prior to failure at the failed over instance. Transactions in progress during failure needs to be rolled back. A callback function is provided for applications so that they can continue on after fail-over without quitting the applications due to node failure.

ENHANCED FAIL-OVER WITH UNINTERRUPTED CAPACITY

For regular Oracle cluster database operations, when a node in an Oracle cluster database cluster fails, Oracle cluster database continues to operate at a reduced capacity. It only decreases in performance due to the loss of the processing power of the failed node.

For customers who are primarily concerned about database availability, Oracle cluster database can also be configured to function as a much better fail-over cluster than those offered by OS vendors, e.g., FailSafe on Windows platform. Oracle cluster database enables Oracle database servers to provide guaranteed capacity even during node failure. Using this configuration option, RAC Guard, one Oracle cluster database node provides reserved capacity as active standby for the other. When one node fails, the other node takes over immediately without lost of any database server capacity.

This enhanced failover capability provides much better availability than regular fail safe capability offered by OS vendors. Regular fail safe needs to restart all the database related services, including volume groups, network resources, application processes and the database instance from scratch after failure is detected, and database recovery must be completed before service is available. These operations can be time consuming when a database grows larger. Since services are not available when these failover operations are ongoing, the benefits of such

solutions are reduced. Also, they do not allow the standby node to be used for secondary tasks, such as administration and diagnostics.

With Oracle cluster database, the fail-over (standby) node is ready to take requests without restarting database all over again, since an active database instance is already running on the fail-over node. When the main node fails, the fail-over node can take over the operations instantly. This configuration also has the added benefits of using the active standby node to do maintenance works on the database. This is only possible since the active standby shares the same database in the Oracle cluster database cluster. You can not do this with regular fail safe configuration since the database is not available in the standby node.

CASE STUDIES

Oracle cluster database enables users to gain great cost saving benefits and enter new markets. The following two cases demonstrate how Oracle cluster database is used to provide highly available database services in vastly different business applications by two different users: VeriSign and Oracle itself.

VeriSign Internet Trust Services

VeriSign Inc. is the leading provider of internet-based trust services needed by websites, enterprises and individuals to conduct secure communications and e-commerce online. Its core offerings include digital certificates, authentication and payment services. All these internet-based OLTP services are powered by Oracle cluster databases at the back end. VeriSign relies on its stable infrastructure to ensure uninterrupted revenue stream. It puts a high emphasis on availability, scalability, and performance. VeriSign is running multiple mission critical production databases on Oracle cluster databases, which continue to be the strategic platform for future databases deployments.

Before Oracle cluster database implementation, VeriSign was using a system based fail-over solution to deploy Oracle databases. This architecture requires 2 nodes with one being primary and the other being standby. The standby node was inactive without database instance mounted. In case of the primary server or database crash, it would unmount the database storage from the primary node and mount it to the standby node. It will then startup the database on the standby node. It was a viable HA fail-over solution. However, it required database outage during fail-over whether it was planned or not, and it required a fully loaded database server as an idle standby node waiting for the fail-over to happen. The recovery time suffers as well. It was neither a scalable nor economical solution.

Using Oracle cluster database, the production database team was able to bring the VeriSign HA solution to the next level. The new production system can fail over database within a minute. It supports at least four times the database activities in the previous 12 months, and keeps every transaction under 2 seconds. Oracle cluster database was chosen mainly because of its ability to run multiple active

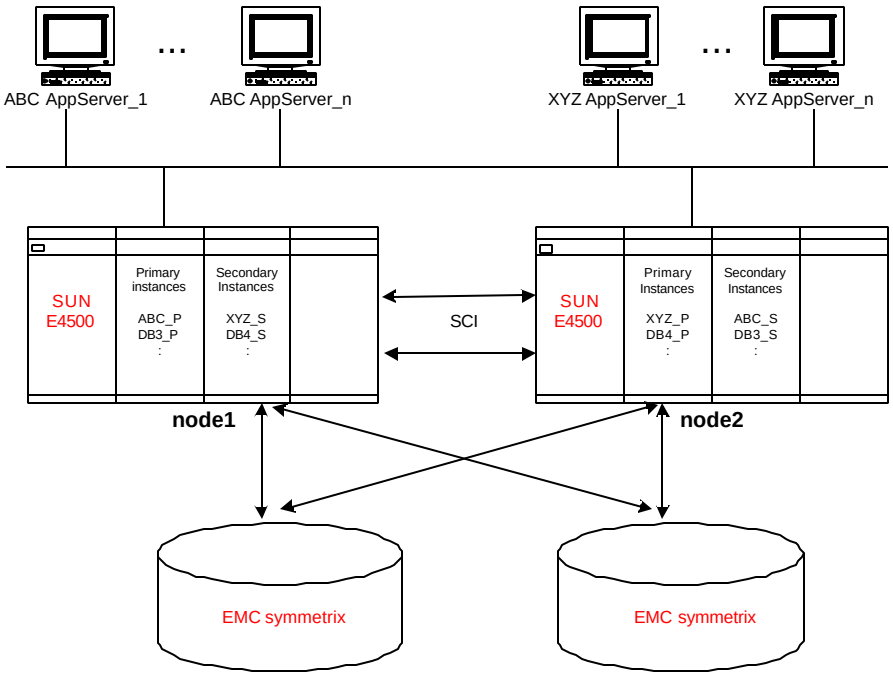

Figure 3. VeriSign Trust Services Architecture

instances for the same data on separate nodes. This completely eliminates any database outage regardless of server failure and instance crash. The database is always up and no fail-over is necessary. Since all the nodes are active, the computing resources of all the high-end cluster servers are fully utilized. Since multiple application databases can be run on the same cluster, a particular application can be designated to access its database on a specific node. This reduces overheads caused by one application using all the nodes in the same cluster. For example, the CERT database application will run on node 1 and use node 2 as an active standby instance while the PAY database application will run on node 2 and use node 1 as an active standby instance. When one of the nodes became inaccessible, the application can simply point to the other node because the database instance will already be there. With this design, VeriSign Trust Service system is able to achieve high availability and scalability without losing any performance. The architecture is shown below.

Oracle Email Services

Utilizing Oracle's own technologies, Oracle Corporation has been able to gain great cost saving benefits and enter into new markets that directly contributed to Oracle's bottom line.

Oracle Email Services handle email communications for all Oracle employees over the world. There are 54,000 named users with a daily traffic load of many hundreds of gigabytes of data. Since email is a global service, it must be up 24 hours a day. The service requires a 3 second response time and 2 minute mean time to recovery.

Oracle consolidated its email servers in middle of 2000. Before the consolidation, the email systems consisted of 97 servers that scattered around over 50 countries. There were 120 Oracle databases with 13 different character sets, different kinds of hardware and software and different email protocols. Email addresses had to be identified with different country codes. There were no backup services when some of them go down. To support the services with all the different configurations and operating environments required 70 system and database administrators, and an equal number of help desk and backup staff.

After the consolidation, Oracle email services use 11 servers, one common UTF8 character set, 5 databases and a set of common protocols such as LDAP, POP3, IMAP5. The service is hosted at one central location with disaster recovery standby at another location. For the first time, email addresses can all use one single domain oracle.com without the country code sub domains. A system diagram is shown at Figure 4.

A simple comparison can show the obvious great cost savings: 97 servers to 11, 120 databases to 5, standardized client software, all inclusive Oracle server software, data center space and staffing.

The server system is built completely using Oracle's own products. The 5 Oracle databases that are used as repositories for the email system are hosted by a two node Oracle cluster database cluster. The Oracle cluster database server is configured for one node to provide standby and for the other to guarantee the capacity of the database in case of node failure.

Aside from satisfying the performance requirements, the databases are required to recover within 60 seconds on average so that the email system can recover within 2 minutes. There is no way for regular fail safe to achieve the required recovery time, given that the sizes of database instances range from several GB to 18 GB, with database sizes up to 1.6 TB. The only option available is to use Oracle cluster database. The Oracle cluster database node acting as standby database has all the necessary resources running and ready to take over while the main node is active. Since the database instance on the failed over node share the same database as the active one, it is ready to switch over as soon as the database recovery allows processing to proceed when the active node fails. Combining Oracle cluster database with Oracle Database's fast start recovery scheme, the Oracle cluster database database server deployed as Oracle Email Service's repository easily meets the recovery time requirements.

Oracle's email hosting business will be utilizing the same architecture to provide fast and reliable email hosting services.

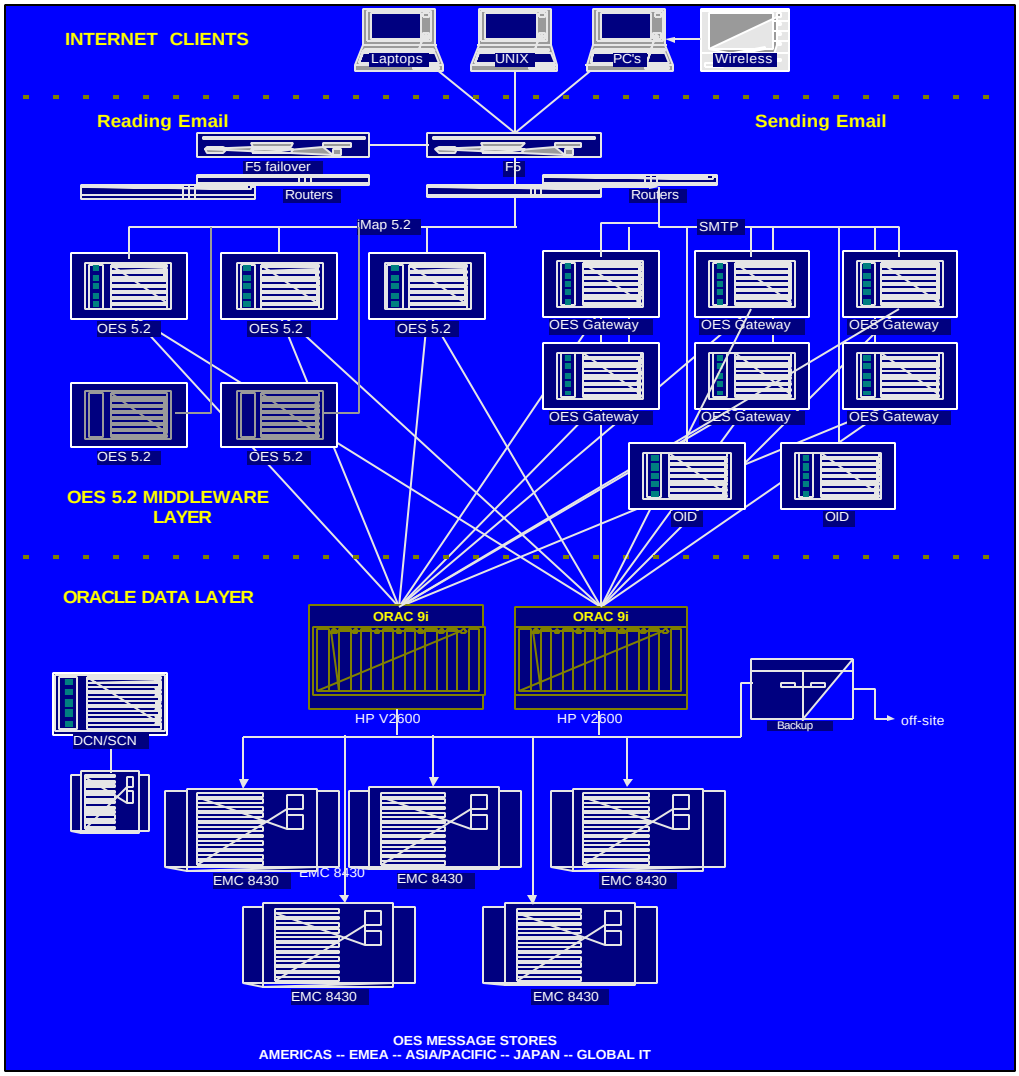

Figure 4. Oracle Email Service

FUTURE ENHANCEMENTS

Oracle plans to build a new HA framework based on Compaq TruCluster's CAA (Cluster Availability Architecture) on all hardware platforms Oracle Real Application Clusters support. These new features will greatly enhance the HA capabilities of Oracle cluster database. The HA framework and its associated features will be available starting with Oracle 9i release 2.

The HA features in Oracle cluster database will enable users to build a highly available Oracle database server with redundant capacity on a cluster greater than 2 nodes. Unlike RAC Guard that requires a separate installation, they will be intrinsic functions of Oracle Real Application Clusters. These HA functions are

automatically enabled with every installation, with configuration options that allow users to tailor to their own HA requirements.

Following are the key features:

Consistent Cluster Ware

Oracle will provide a cluster ware that handles hardware and OS level cluster functions, such as cluster group membership, and cluster monitoring, etc. for use by Oracle cluster database. Cluster wares from vendors will co-exist with Oracle's cluster ware, so that other applications that rely on vendors' cluster ware can still run on the same cluster. Oracle's own cluster ware will allow Oracle cluster database to have a consistent way of dealing with various cluster level events, resulting in a more consistent behavior at the database level.

Cluster Alias and Virtual IP

Cluster alias and/or virtual IP allows applications to see a cluster as one entity rather than as a collection of nodes. This feature makes Oracle cluster database virtual undistinguishable from a single instance database for applications that use them. All the setup and configuration needed are done on the Oracle side, making it much easier for applications to use Oracle cluster database.

Management of Resources

Resources refer to a broad ranges of facilities in the cluster database environment, including virtual IP addresses, processes bound to local nodes, (such as daemon processes, listener processes, EM agent processes), database instances processes, service provider, service attractor, service member, etc. Oracle cluster database treats resources as basic management units. Resources can be composite resources and can have dependency on other resources.

Oracle cluster database automates start, stop, and monitoring of resources and also automates failover, relocation, restart of resources.

Resources can be internal as well as external to allow for users to define resources for application services.

Policy Based HA Management

In addition to a set of pre-defined policy to manage Oracle cluster database, users can define policies on how to utilize resources provided by the cluster database. For example, users can define how to fail over database services in case of instance failure; Users can define whether fail-back is automatic or manual; Users can define whether to keep reserve capacity, and how it is to be used, etc. Users can also define priority of how resources are to be used, whether resources can be preempted, etc.

Service Based Work Load Partitioning

Applications connect to Oracle database through “services” defined in Oracle network. Users can designate which groups of nodes that provide those services, effectively partitioning the work load of the cluster database according to services. Reserve capacity for those services can be designated as well. Similar work loads can be grouped together by binding the services for those work loads together.

Event System

Oracle cluster database provides an event system where related events such as resource failure, status changes, start/completion of failover, failback, etc. can be defined, monitored, notified and associated operations can be taken. Both management tools and applications can utilize the system to respond to events.

Management Facility For Setup And Configuration

The management facilities for the new functions will be integrated in OEM, regular Oracle installation and configuration tools. GUI and command line based facilities will be provided for users to manage resource, resource dependency, events, policy, etc. Lights-out operations will be supported as well.

CONCLUSION

Oracle Real Application Clusters are both highly available and highly scalable. Oracle cluster database achieves both goals without sacrificing one for the other.

This paper has attempted to give readers a clear idea on how Oracle cluster database addresses generic availability issues associated with cluster architecture and how Oracle cluster database is enhanced to provide better availability capabilities.

Real world examples in this paper demonstrated how Oracle cluster database is used in the business world to save costs and help businesses enter new markets with Oracle cluster database as a highly available, highly scalable and cost effective database server platform.

Building Highly Available Database Servers Using Oracle Real Application Clusters
May, 2002
Author: Jack Cai
Contributing Authors: Simon Leung, VeriSign Inc.

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
www.oracle.com

Oracle Corporation provides the software
that powers the internet.