



Maximizing Availability with Oracle Application Server 10g



Simone Bundschu
Hewlett-Packard

© 2004 Hewlett-Packard Development Company, L.P.
The information contained herein is subject to change without notice



Agenda

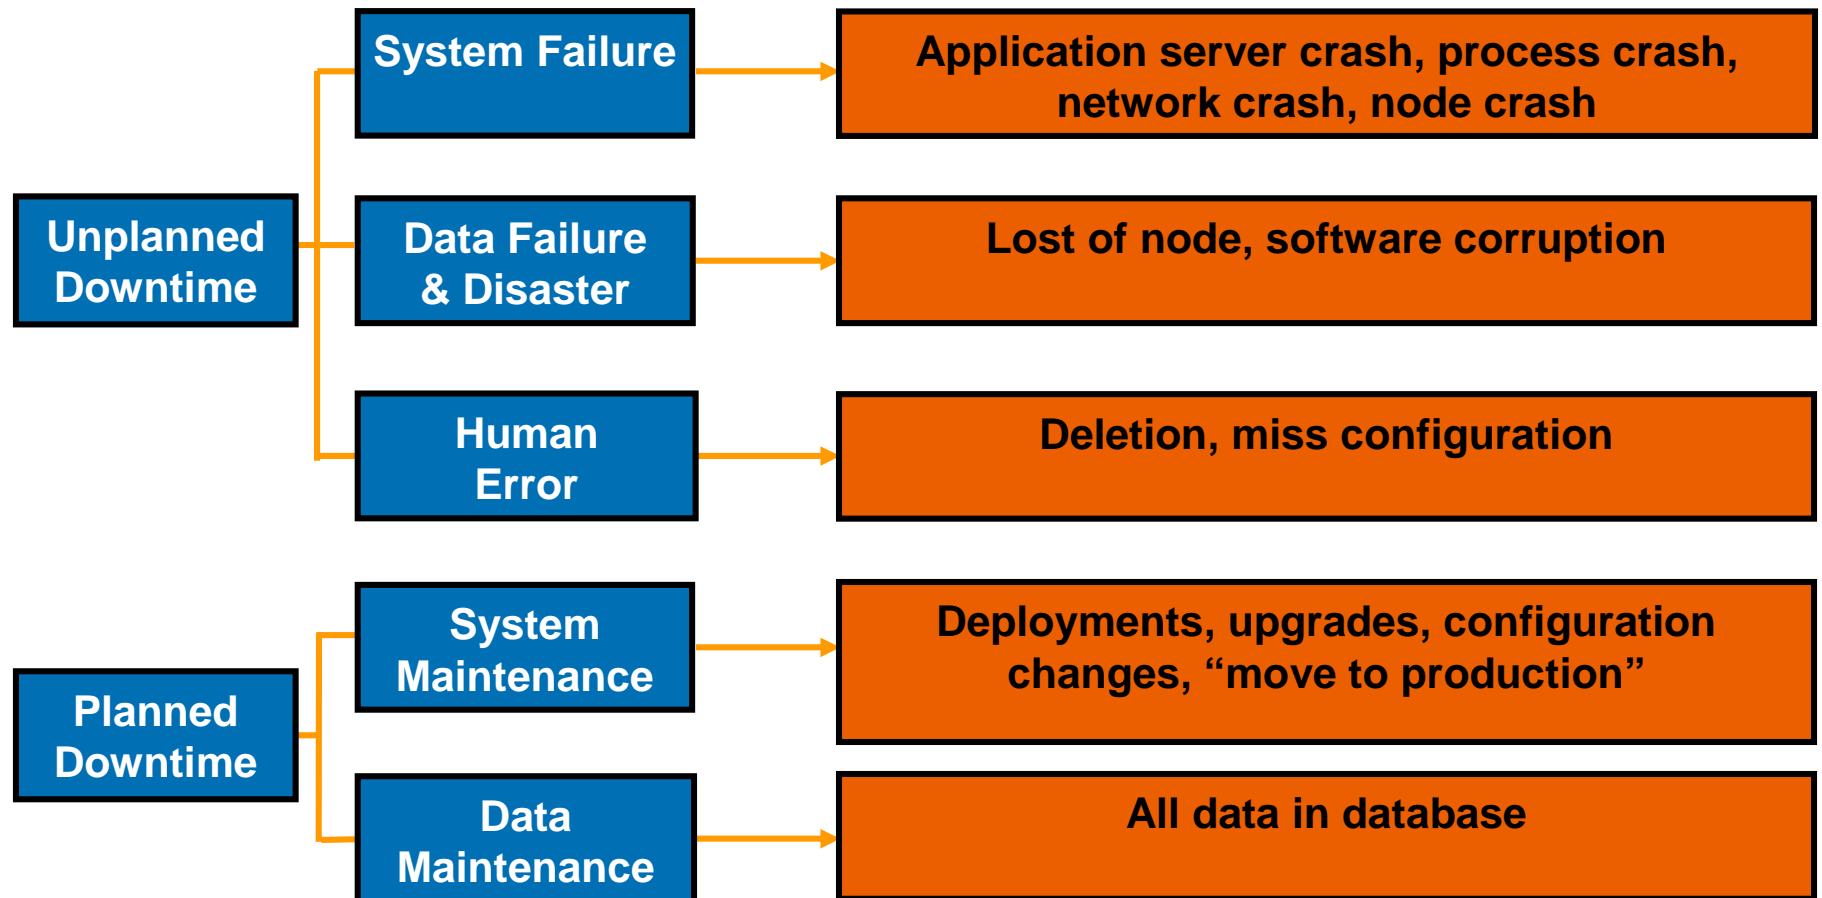
- Introduction
- Oracle Application Server High Availability Features
 - Overall process control: Self-healing
 - All-subtiers clustering:
 - HTTP Server
 - OC4J
 - Connection Pool
 - Portal/Reports
 - High Availability for Metadata and Security Services
 - Backup&Recovery and Disaster Recovery
 - Deployment features
- Sample Configurations/Summary/Q&A

Agenda

- Introduction
- Oracle Application Server High Availability Features
 - Overall process control: Self-healing
 - All-subtiers clustering:
 - HTTP Server
 - OC4J
 - Connection Pool
 - Portal/Reports
 - High Availability for Metadata and Security Services
 - Backup&Recovery and Disaster Recovery
 - Deployment features
- Sample Configurations/Summary/Q&A

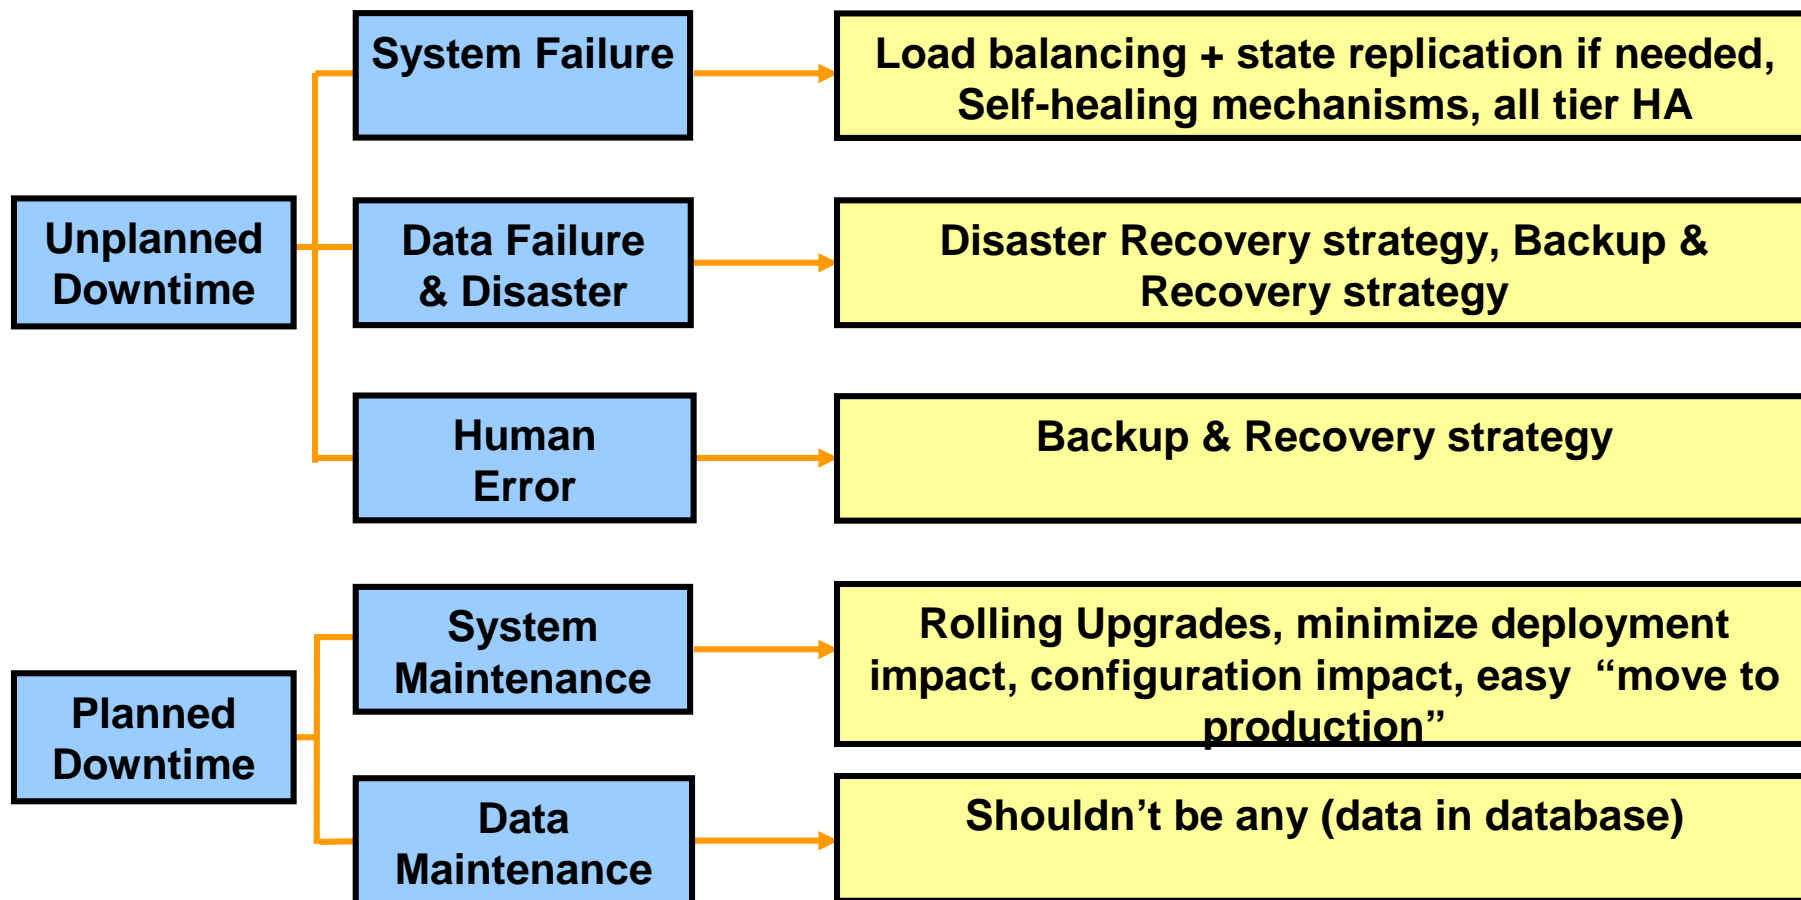


Oracle Application Server High Availability

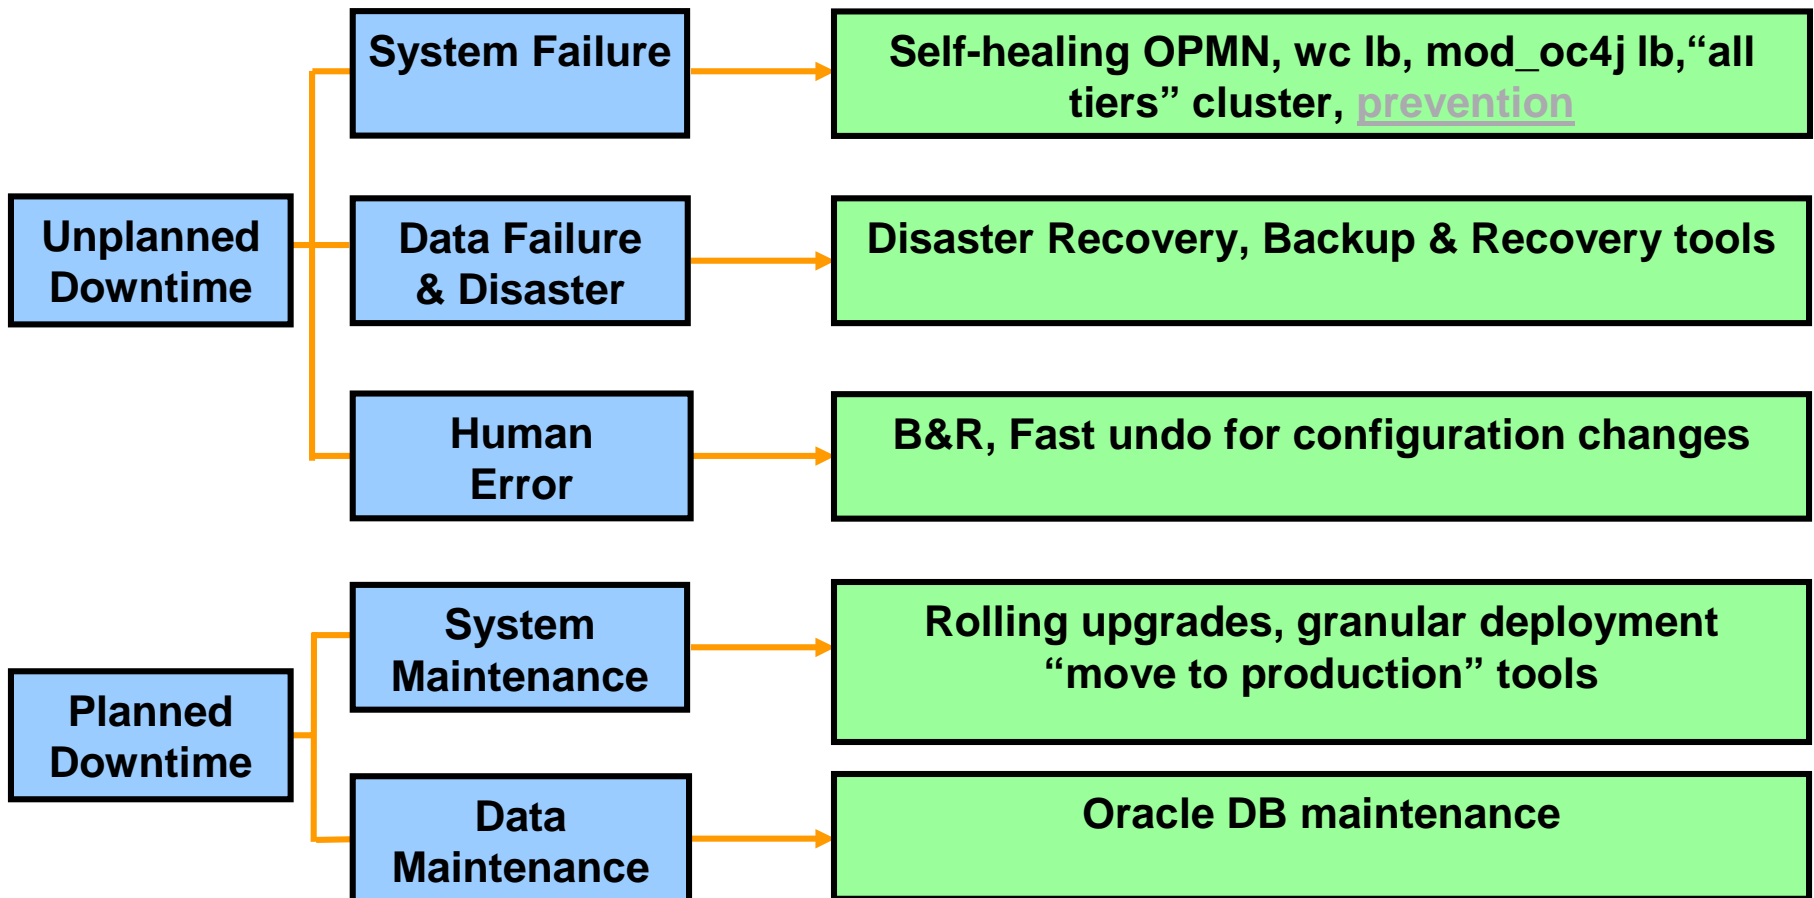




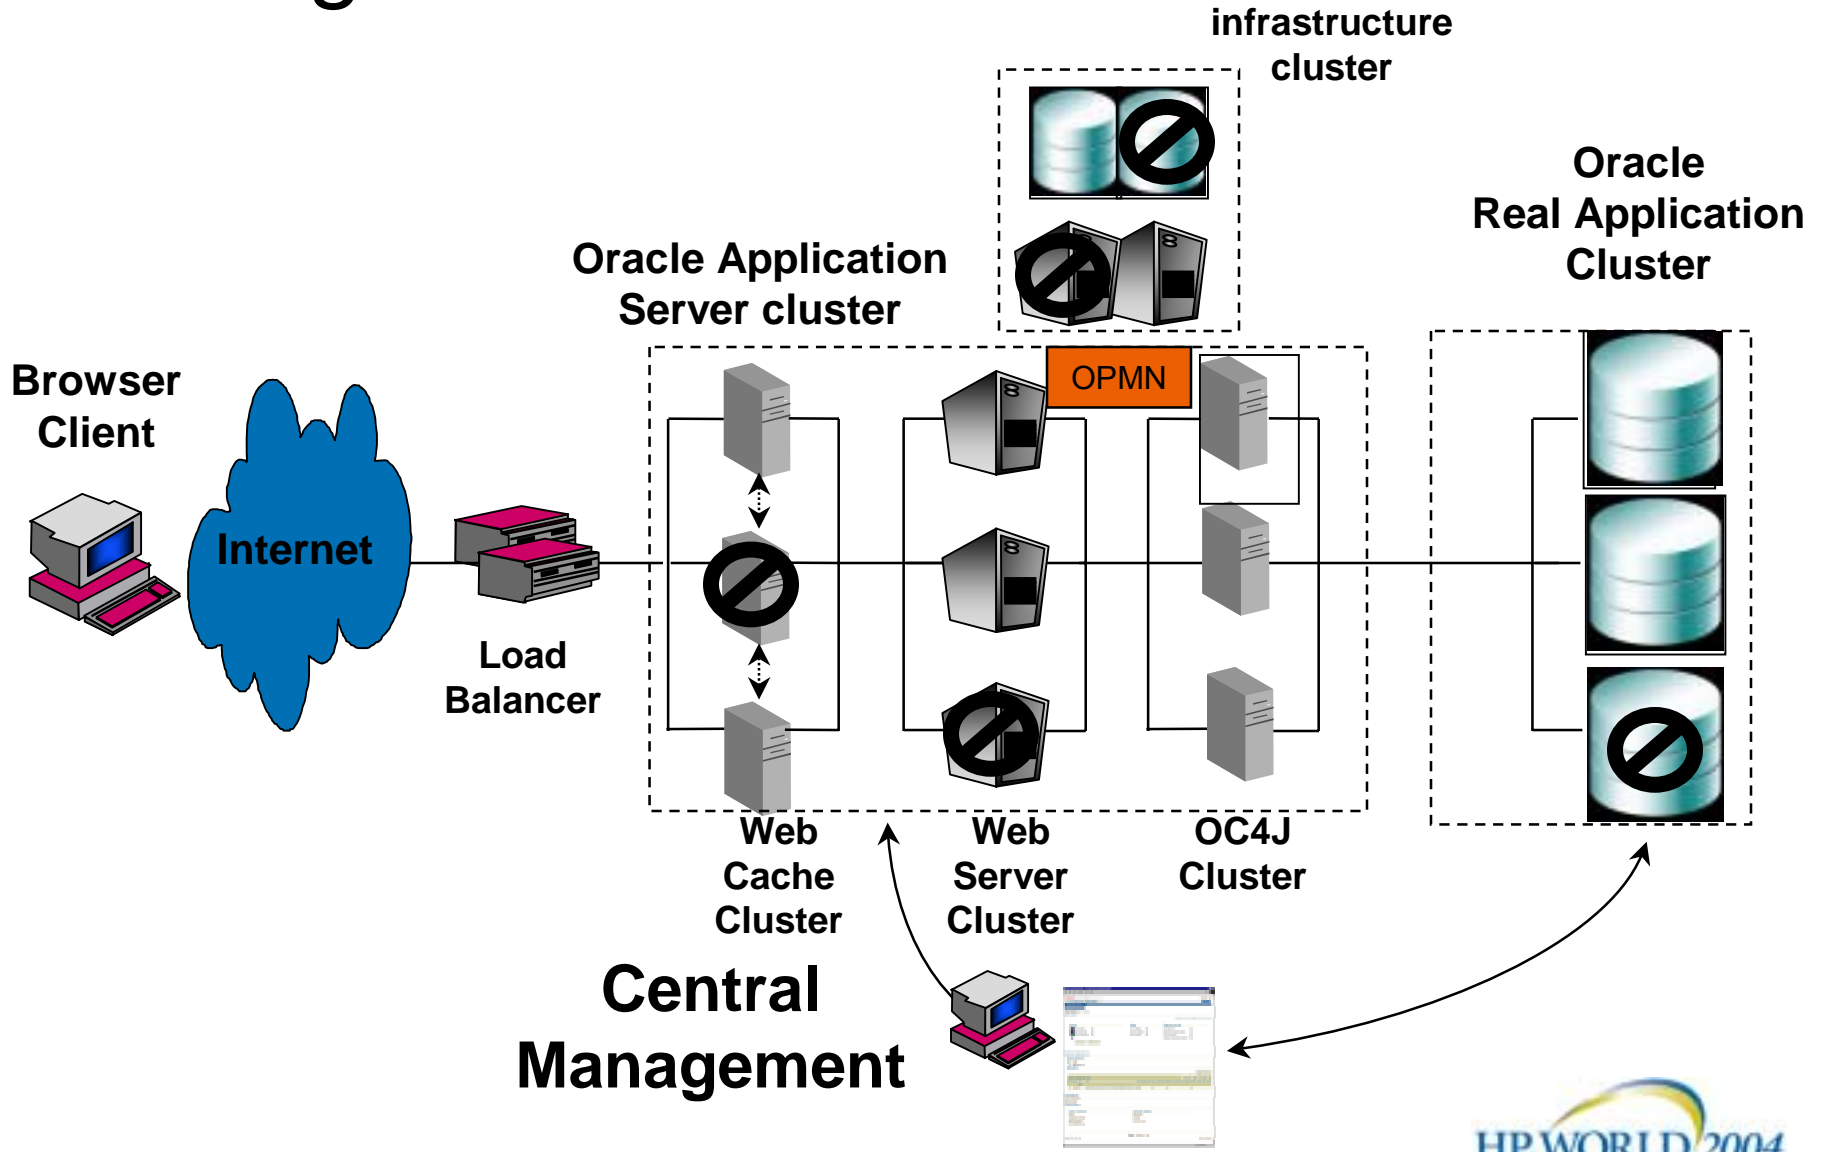
Oracle Application Server High Availability



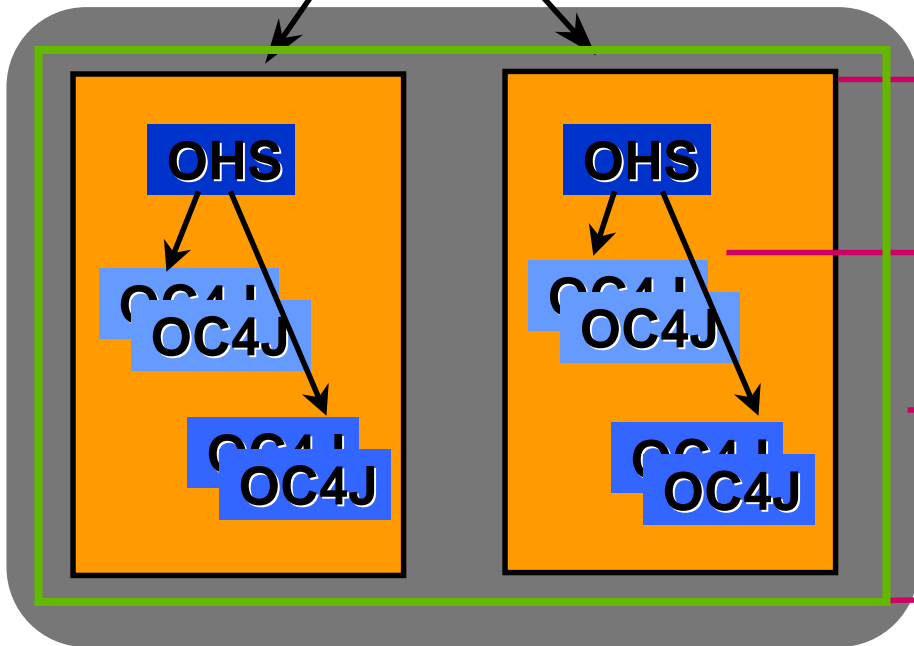

What Oracle Application Server Has to say about High Availability



No Single Point of Failure



Terminology



OracleAS Instance:

Managed App Server processes
(one \$ORACLE_HOME)

Component Instance:

Identical configured components

Farm:

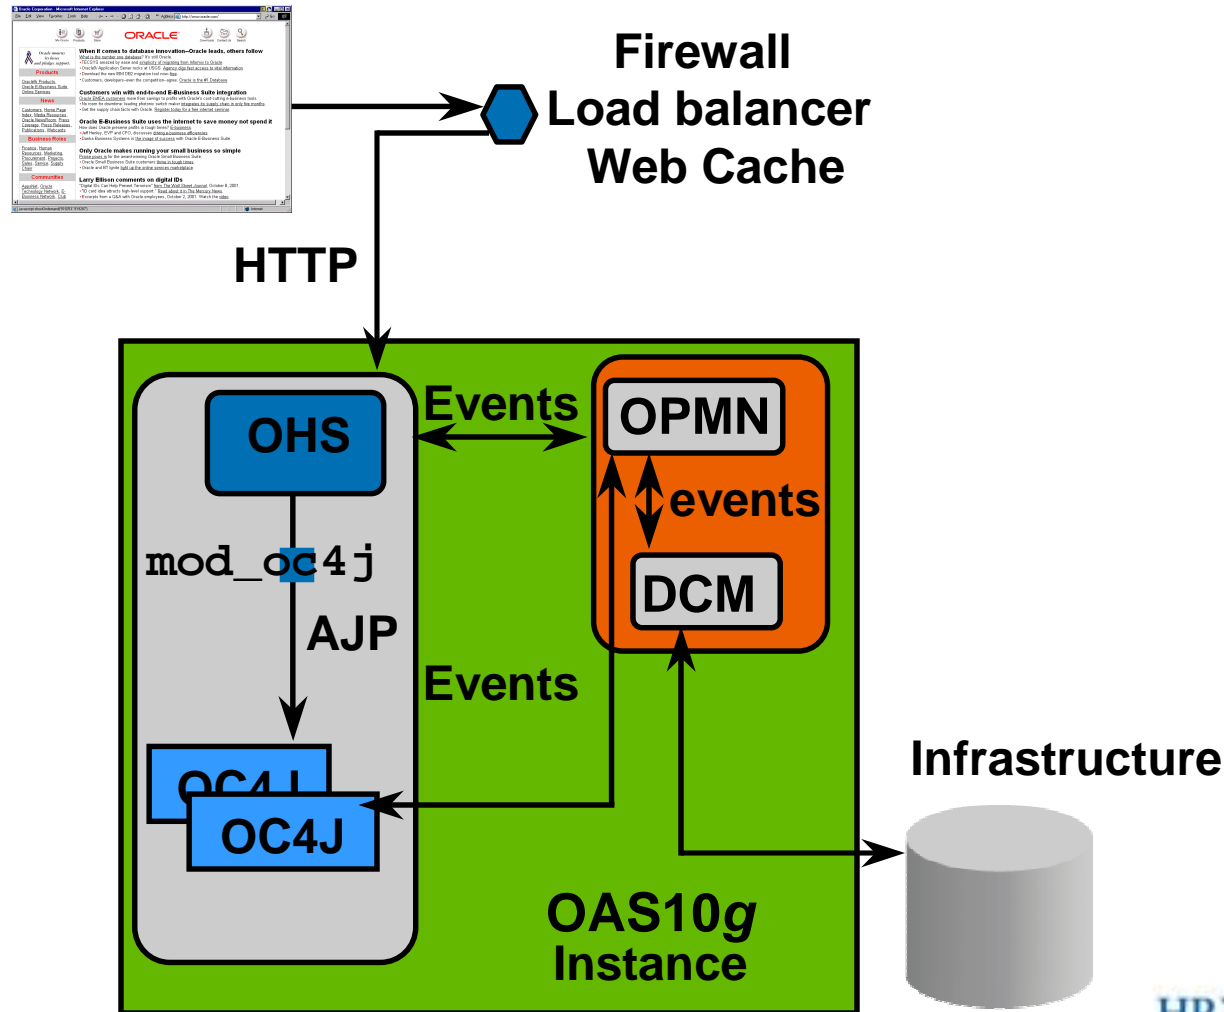
Multiple OracleAS instances
connected to a single
infrastructure

Cluster:

Identical OracleAS instances from
a single farm



Architecture of an Oracle Application Server 10g Instance



Cluster Configuration in EM



The screenshot shows the Oracle Enterprise Manager console interface. The main window title is "Oracle Enterprise Manager - Farm: asdb.de.oracle.com - Mozilla". The address bar shows "http://stulinuxhh1.de.oracle.com:1812/emd/console".

On the left, there is a sidebar with the "Create Cluster" wizard. The "Cluster Name" field contains "904midtierCluster".

The main content area shows a table with the following data:

Select	Name	Status	Instances
<input checked="" type="radio"/>	904midtierCluster	↓	0

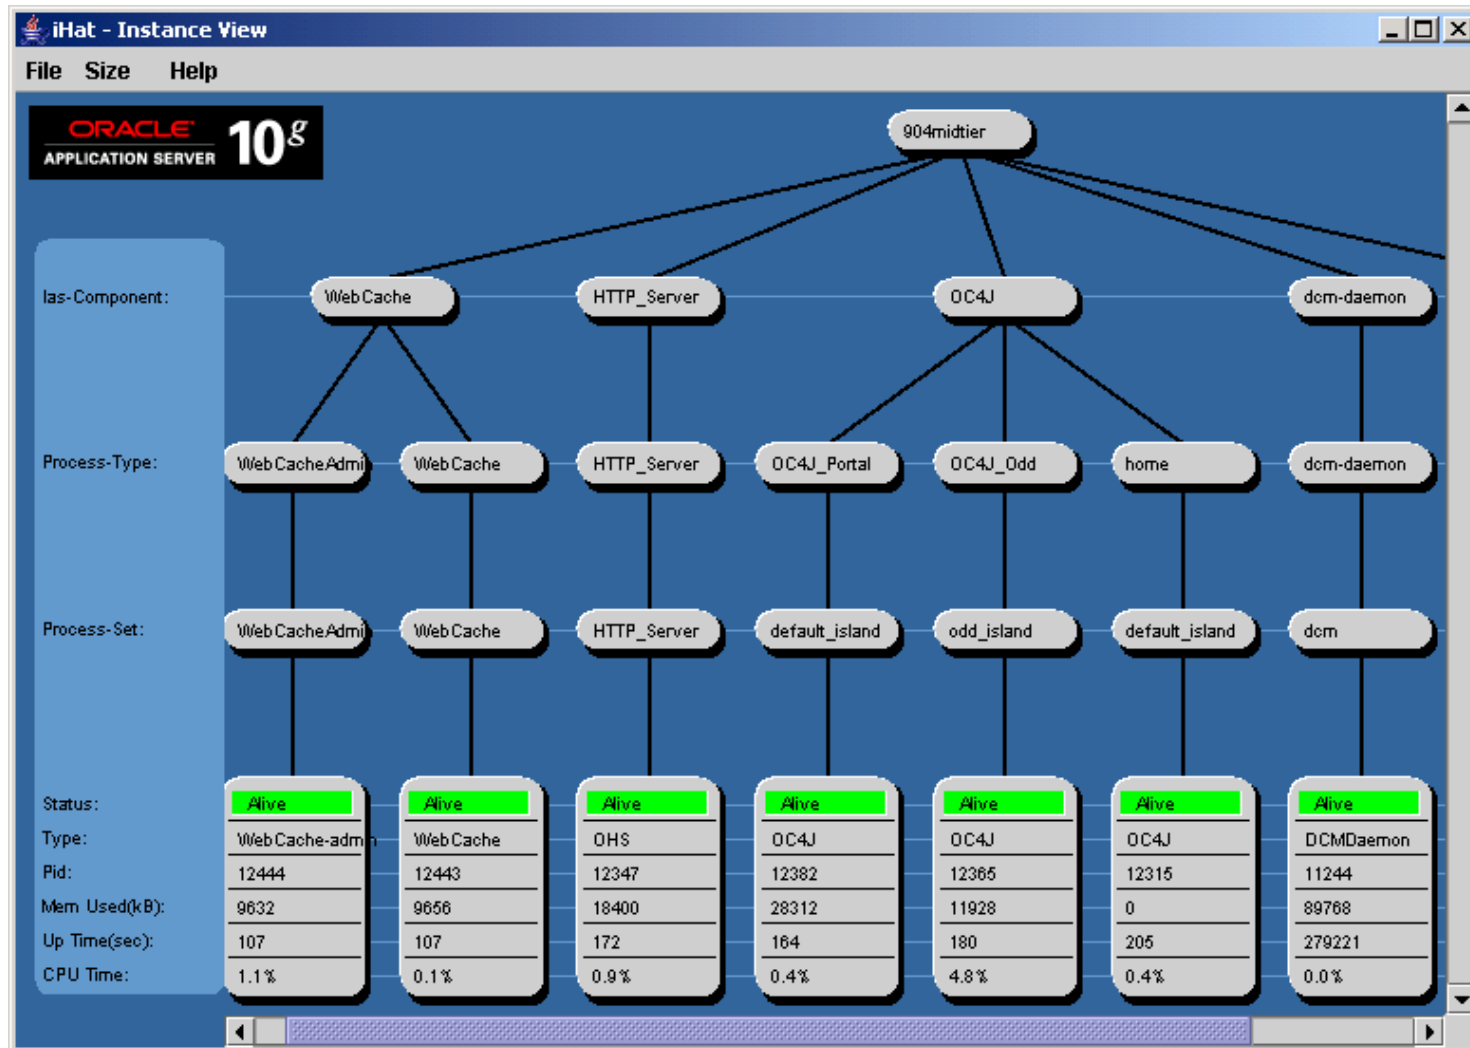
Below this table, there is a section titled "Standalone Instances" with the text "These instances belong to the farm but are not part of any cluster." A "Join Cluster" button is circled in orange.

Below the "Join Cluster" button, there is another table with the following data:

Select	Name	Host	Oracle Home
<input checked="" type="radio"/>	904midtier.sturachh1.de.oracle.com	sturachh1.de.oracle.com	/ufs1/OraHome3/904midtier
<input type="radio"/>	904midtier.sturachh2.de.oracle.com	sturachh2.de.oracle.com	/ufs1/OraHome3/904midtier
<input type="radio"/>	infra904.stulinuxhh1.de.oracle.com	stulinuxhh1.de.oracle.com	/products/oracle/infra904

At the bottom, there are links for "Preferences" and "Help", and a copyright notice: "Copyright © 1996, 2003, Oracle. All rights reserved."

Cluster Monitoring in iHAT





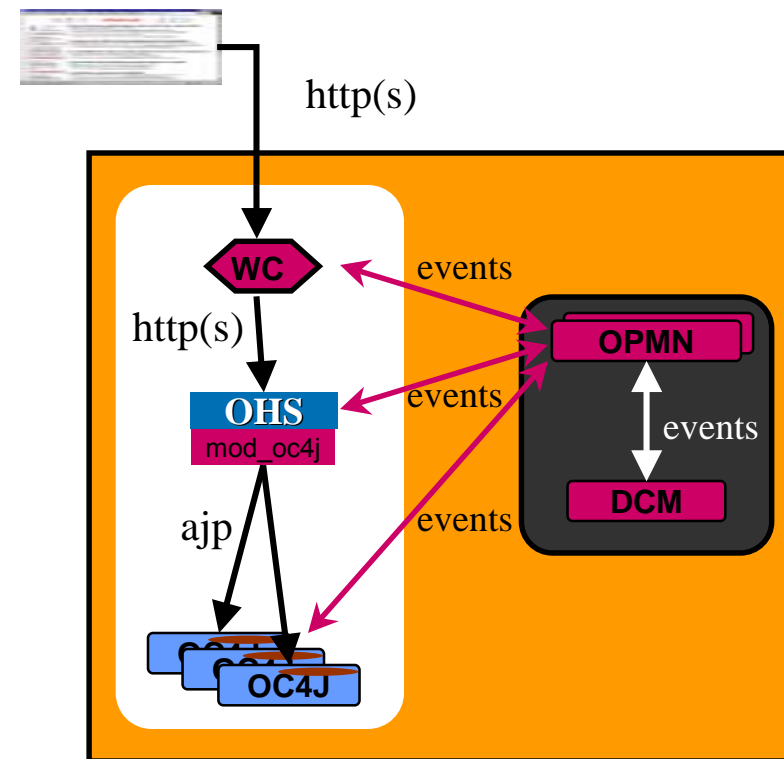
Agenda

- Introduction
- Oracle Application Server High Availability Features
 - Overall process control: Self-healing
 - All-subtiers clustering
 - High Availability for Metadata and Security Services
 - Backup&Recovery and Disaster Recovery
 - Deployment features
- Sample Configurations/Summary/Q&A



Process Monitoring: OPMN

- Monitoring of all Oracle AS processes
- Automatically restarts processes (“self healing”)
 - OPMN-Shadow Process monitors OPMN
- A process failure will cause a notification to dependant components/processes
 - e.g. mod_oc4j / oc4j



Self Healing in OracleAS: OPMN



- OracleAS Process Monitoring and Notification service (OPMN):
 - Process Monitoring (Death detection and Restart)
 - Encloses all processes:
 - HTTP servers
 - J2EE containers
 - LDAP directories
 - Distribute configuration management
 - UDDI registry
 - OracleAS Wireless, OracleAS Reports, OracleAS Discoverer,...
 - Cross-dependencies management

System Failure





Self Healing in OracleAS: OPMN

- Improved starting and stopping of components
- opmnctl is the primary command-line tool for starting and stopping the components
- Increased power and flexibility for configuring the Oracle Application Server
- Event hooks and scripts
- Custom services monitoring
- Operating system-level statistics
- Restore to pre-crash status after "non-graceful" shutdown (such as an OS crash)

System Failure



Agenda

- Introduction
- Oracle Application Server High Availability Features
 - Overall process control: Self-healing
 - All-subtier clustering
 - High Availability for Metadata and Security Services
 - Backup&Recovery and Disaster Recovery
 - Deployment features
- Sample Configurations/Summary/Q&A

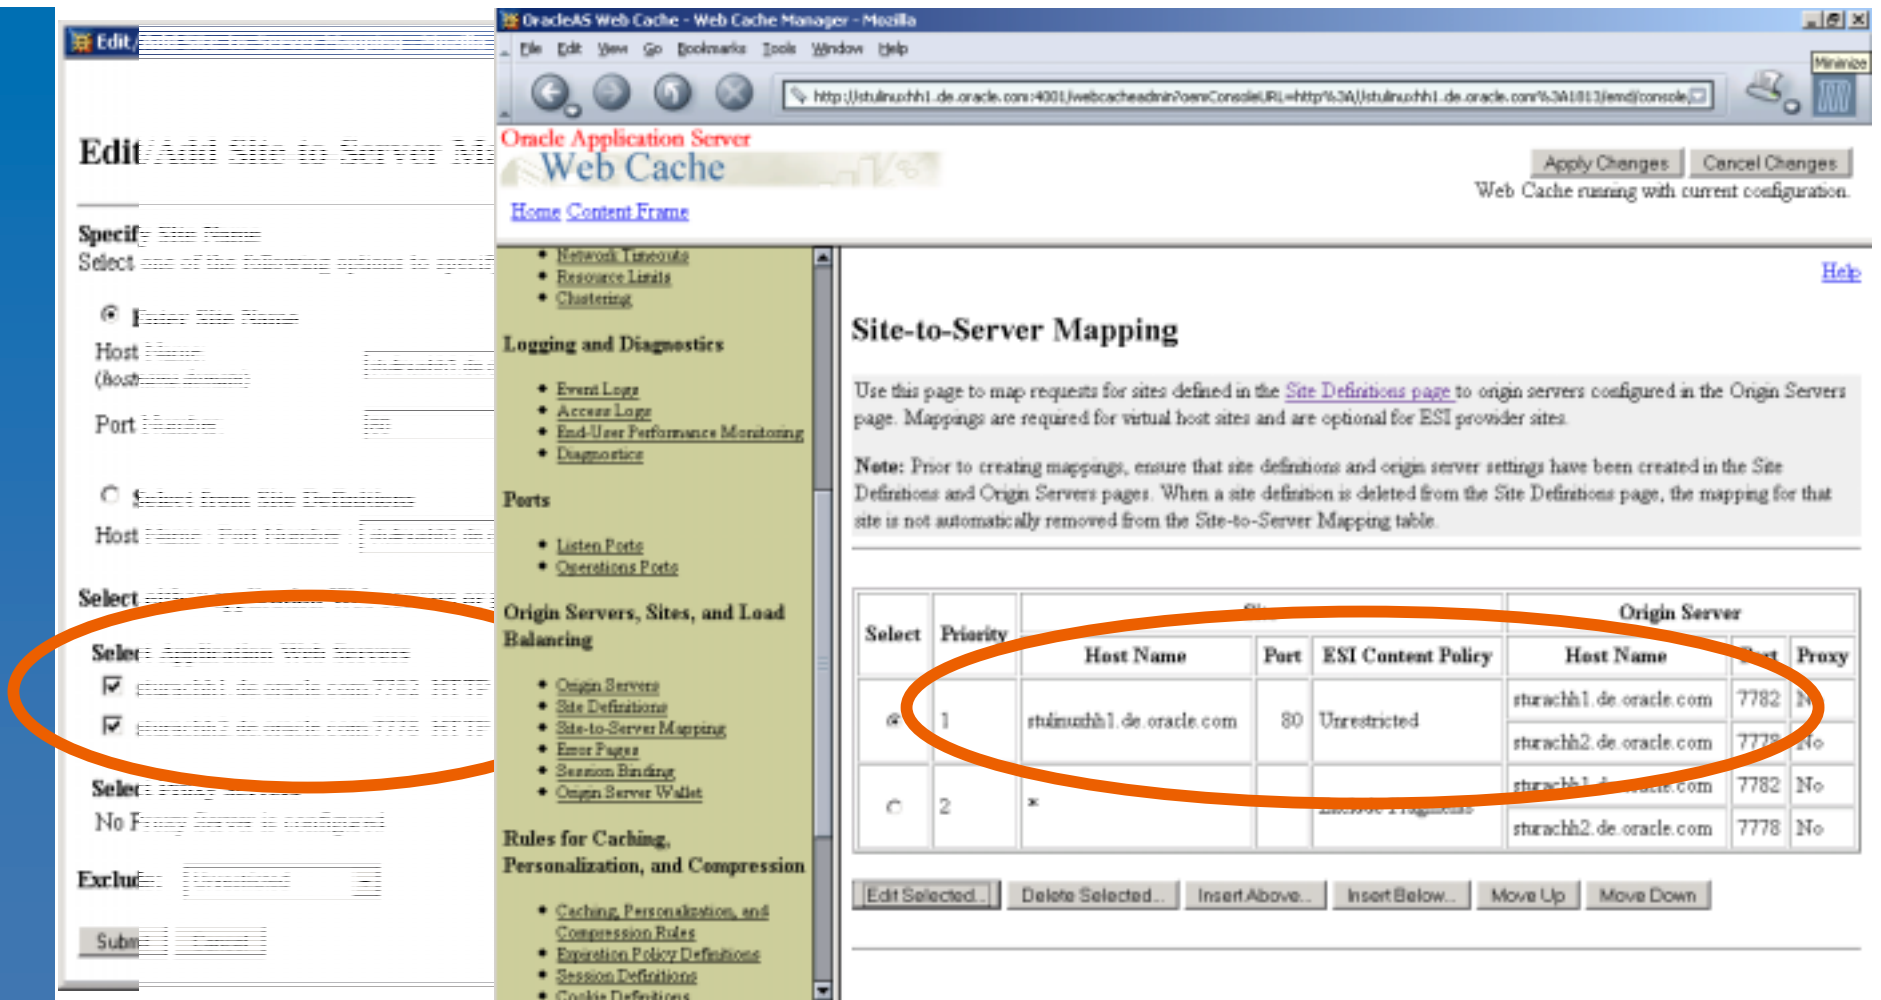
All sub-tier clusters

- Webcache Clusters:
 - Reflection monitoring for state and content miss
- HTTP Session clusters:
 - Island grouping, code agnostic
- Stateful Session Ejb's clusters:
 - On jvm's crash, on each call, code agnostic
- Database Connection state (for instance, open cursors) for jdbc with RAC

Webcache as a Loadbalancer

- webcache load balances to HTTP Server
 - round-robin
- supports session-binding
- recognizes HTTP Server failure
 - detects re-availability of a HTTP-Server
- memory cache
 - can be configured in a webcache cluster over several boxes

Webcache Load-Balancing



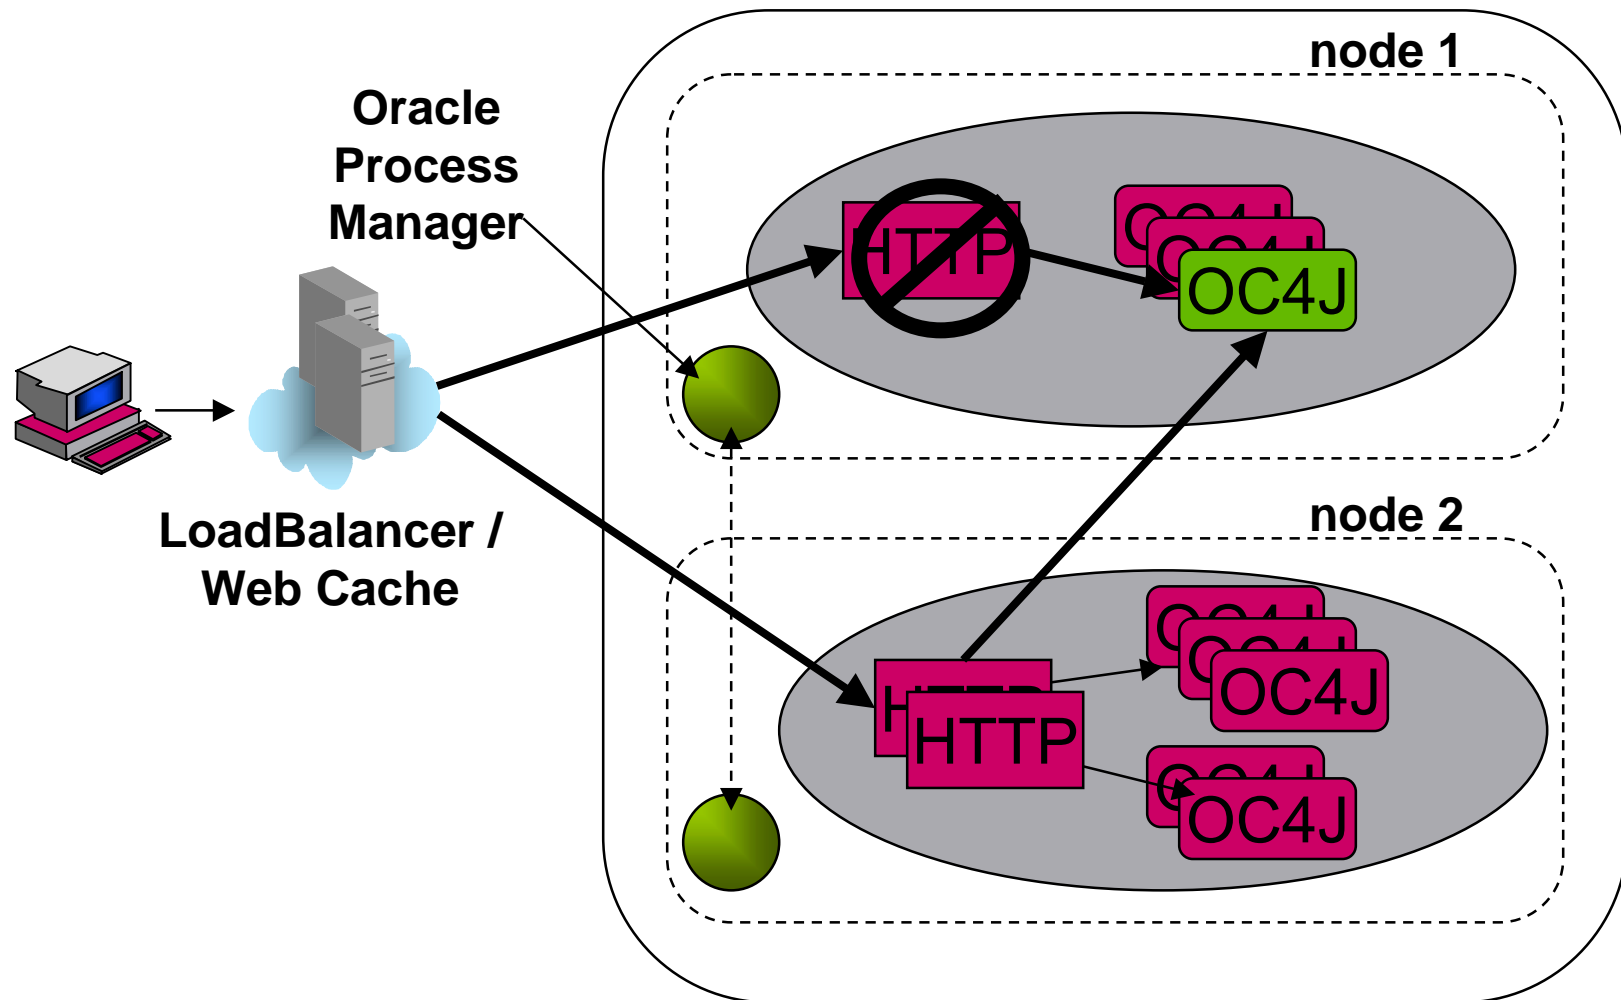
The screenshot shows the Oracle Application Server Web Cache Manager interface. The main window is titled "Oracle Application Server - Web Cache Manager - Mozilla". The URL in the address bar is "http://stulnuzhh1.de.oracle.com:4001/webcacheadmin?oarsConsoleURL=http%3A%2F%2Fstulnuzhh1.de.oracle.com%2F%2Fend/console%2F".

The interface is divided into several sections:

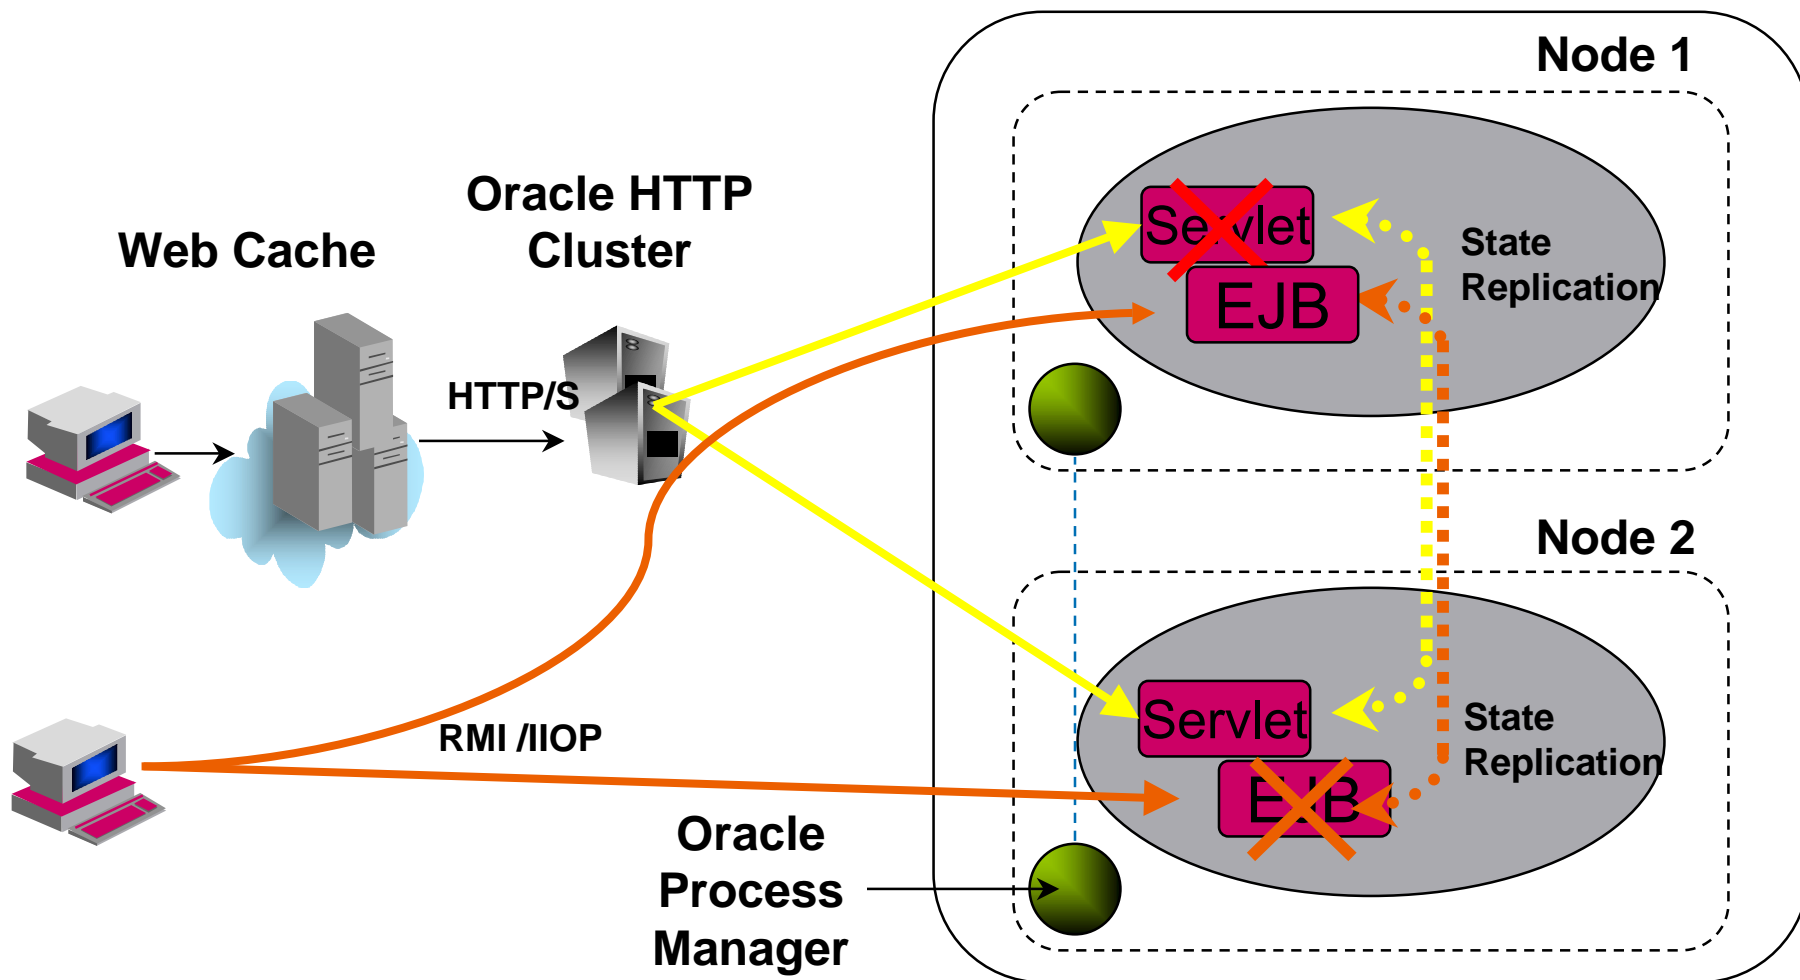
- Left Panel:** "Edit/Add Site to Server Mapping". It has two radio buttons: "Enter Site Name" (selected) and "Select from Site Definitions". Under "Enter Site Name", there are fields for "Host Name" and "Port". Under "Select from Site Definitions", there are fields for "Host Name" and "Port Number". Below these are "Select" buttons for "Application Web Servers" and "Proxy Servers".
- Navigation Menu:** A vertical menu on the left contains:
 - Network Timouts
 - Resource Limits
 - Clustering
 - Logging and Diagnostics
 - Event Logs
 - Access Logs
 - End-User Performance Monitoring
 - Diagnostic
 - Ports
 - Listen Ports
 - Operations Ports
 - Origin Servers, Sites, and Load Balancing
 - Origin Servers
 - Site Definitions
 - Site-to-Server Mapping
 - Error Pages
 - Session Binding
 - Origin Server Wallet
 - Rules for Caching, Personalization, and Compression
 - Caching, Personalization, and Compression Rules
 - Expiration Policy Definitions
 - Session Definitions
 - Cookie Definitions
- Main Content Area:**
 - Web Cache:** Includes "Home Content Frame", "Apply Changes", "Cancel Changes", and "Web Cache running with current configuration."
 - Site-to-Server Mapping:**
 - Text: "Use this page to map requests for sites defined in the Site Definitions page to origin servers configured in the Origin Servers page. Mappings are required for virtual host sites and are optional for ESI provider sites."
 - Note: "Prior to creating mappings, ensure that site definitions and origin server settings have been created in the Site Definitions and Origin Servers pages. When a site definition is deleted from the Site Definitions page, the mapping for that site is not automatically removed from the Site-to-Server Mapping table."
 - Table:** A table with columns: Select, Priority, Host Name, Port, ESI Content Policy, Origin Server (Host Name, Port, Proxy).

Select	Priority	Host Name	Port	ESI Content Policy	Origin Server	Host Name	Port	Proxy
<input checked="" type="radio"/>	1	stulnuzhh1.de.oracle.com	80	Unrestricted	strachh1.de.oracle.com	7782	No	
<input type="radio"/>	2	*			strachh2.de.oracle.com	7778	No	
<input type="radio"/>					strachh1.de.oracle.com	7782	No	
<input type="radio"/>					strachh2.de.oracle.com	7778	No	

OracleAS HTTP Cluster

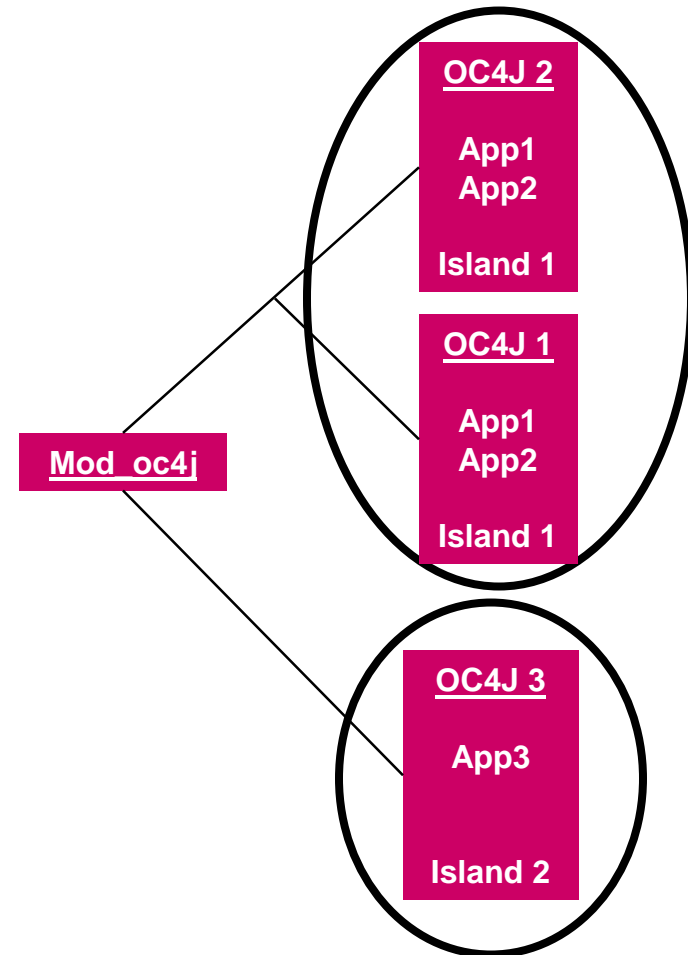


OracleAS OC4J Cluster



Servlet / JSP Clustering

- dynamic node recognition
- automatic load balancing
 - Round Robin
 - Random
 - Metric based
- session replication over IP multicasting
- replication based on island-id
 - Islands enable application specific grouping of cluster nodes



Servlet/JSP Clustering in EM



Oracle Enterprise Manager 10g
Application Server Control

Farm > Application Server: 904midtier.sturachh1.de.oracle.com > OC4J: OC4J Odd > Replication Properties

Replication Properties

Page Refreshed Mar 4, 2004 2:41:56 PM

Web Applications

TIP Session session state replication here will enable session state replication for all web applications. The load-on-startup property will be automatically set to true for all web modules.

Replicate session state

Multicast Host (IP)	230.230.1
Multicast Port	9123

EJB Applications

TIP EJB applications replicate state between all OC4J processes in the OC4J instance.

Replicate State

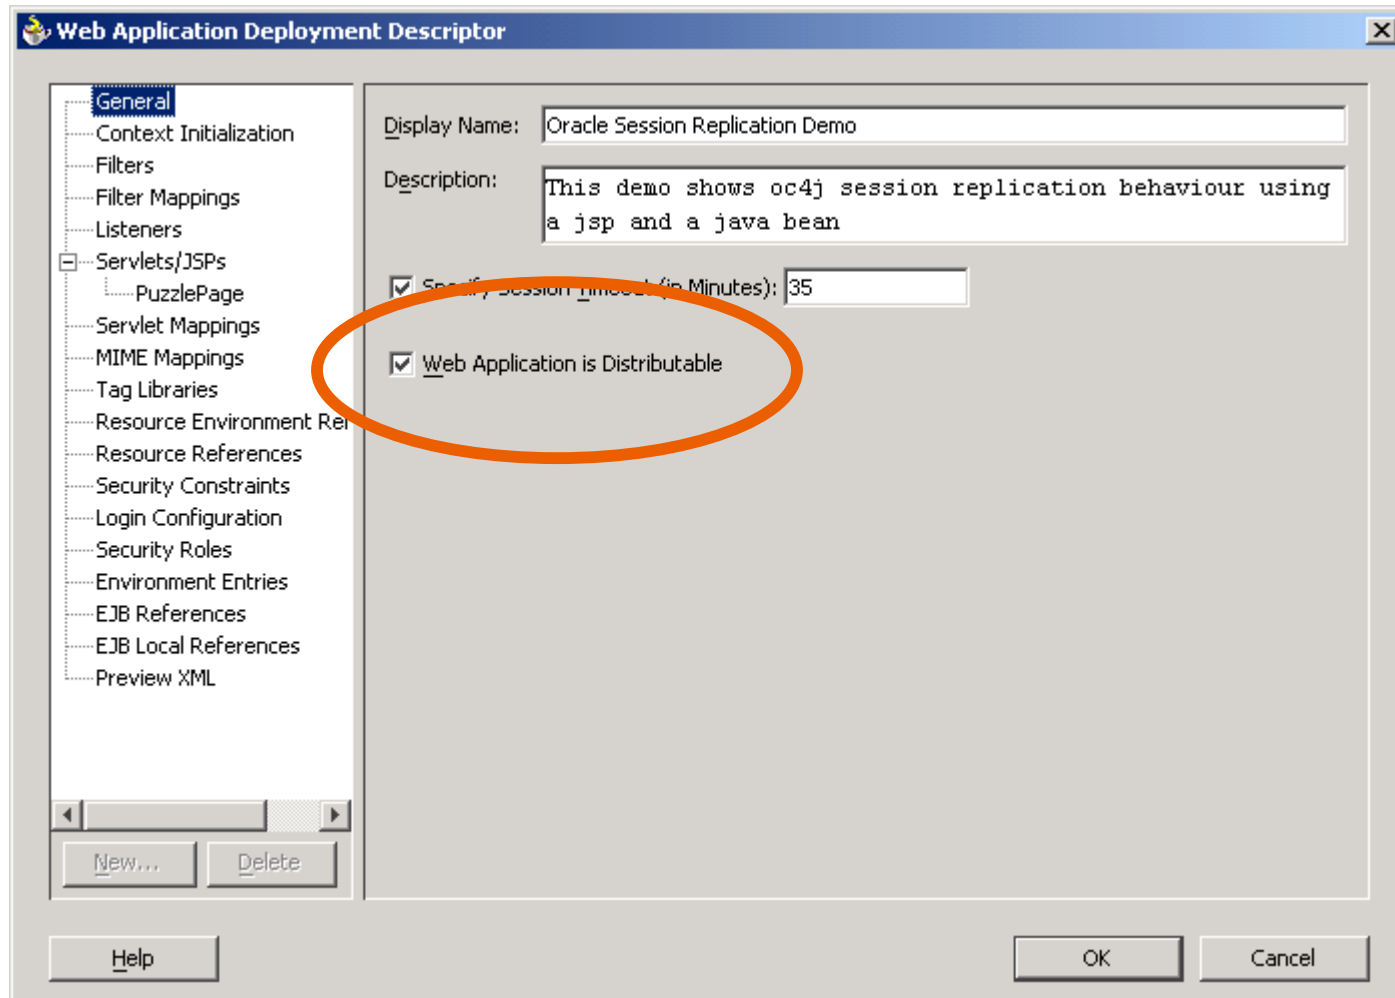
Multicast Host (IP)	
Multicast Port	
Username	
Password	
RMI Server Host	[ALL]

This is usually the name of the machine where the OC4J instance is running.

Revert Apply



Servlet/JSP Clustering in JDeveloper



EJB Clustering

- state replication over IP multicasting
- dynamic and static node recognition
- load-balancing for all EJB types
- stateful session bean fail-over
 - JVM termination
 - End-of-Call
- entity bean fail-over
 - automatic state synchronization

Konfiguration EJB Clustering im EM



Oracle Enterprise Manager 10g
Application Server Control

Farm > Application Server: 904midtier.sturachh1.de.oracle.com > OC4J: OC4J_Odd > Replication Properties

Replication Properties

Page Refreshed Mar 4, 2004 2:49:11 PM

Web Applications

TIP Setting session state replication here will enable session state replication for all web applications. The load-on-startup property will be automatically set to true for all web modules.

Replicate session state

Multicast Host (IP)

Multicast Port

EJB Applications

TIP EJB applications replicate state between all OC4J processes in the OC4J instance.

Replicate State

Multicast Host (IP)

Multicast Port

Username

Password

RMI Server Host

This is usually the name of the machine where the OC4J instance is running.

Database Connection state

- Oracle AS (9.0.4) uses SQL exceptions (Instance down but node/listener still alive) to clean up OC4J connection pools in RAC environments
- Oracle AS 10.1.3 (i.e., oc4j) will use FaN (ONS under the covers), RAC events (10g and up) and JDBC Fast Connection Fail-Over for
 - load-balancing connection requests across JDBC connection Caches on instance UP
 - cleaning up caches (remove orphan connections) on Node/Instance DOWN
- Some components (like portal) will not catch up this functionality until 10.1.4

Load Balancing: “Crash” Prevention



- “The best way not to fall” is to adapt routing to node’s state: let node recover from peaks and route to a more relaxed node
- Most crashes come from overloading...overloading is 90% application’s responsibility
- Load Balancing:
 - From webcache to http servers
 - From http servers to j2ee containers
 - Intra j2ee containers load balancing
 - From j2ee containers to databases

System Failure



“Crash” Prevention



Tiers Involved	Critical? (0-5)	LB mechanisms
Webcache to Http servers	3	Round-robin, weighed round robin
Http servers to j2ee containers	5	Random, Round Robin, Random with Local Affinity, Round Robin with Local Affinity, Random using Routing Weight, Round Robin using Routing Weight, Metrics Based, Metric Based with Local Affinity
Intra j2ee containers	3	Random
J2ee containers to database	5	All oracle Net load balancing mechanisms (RAC to the rescue)

System Failure





All sub-tiers cluster

- and what about Portal, Forms, Reports...?

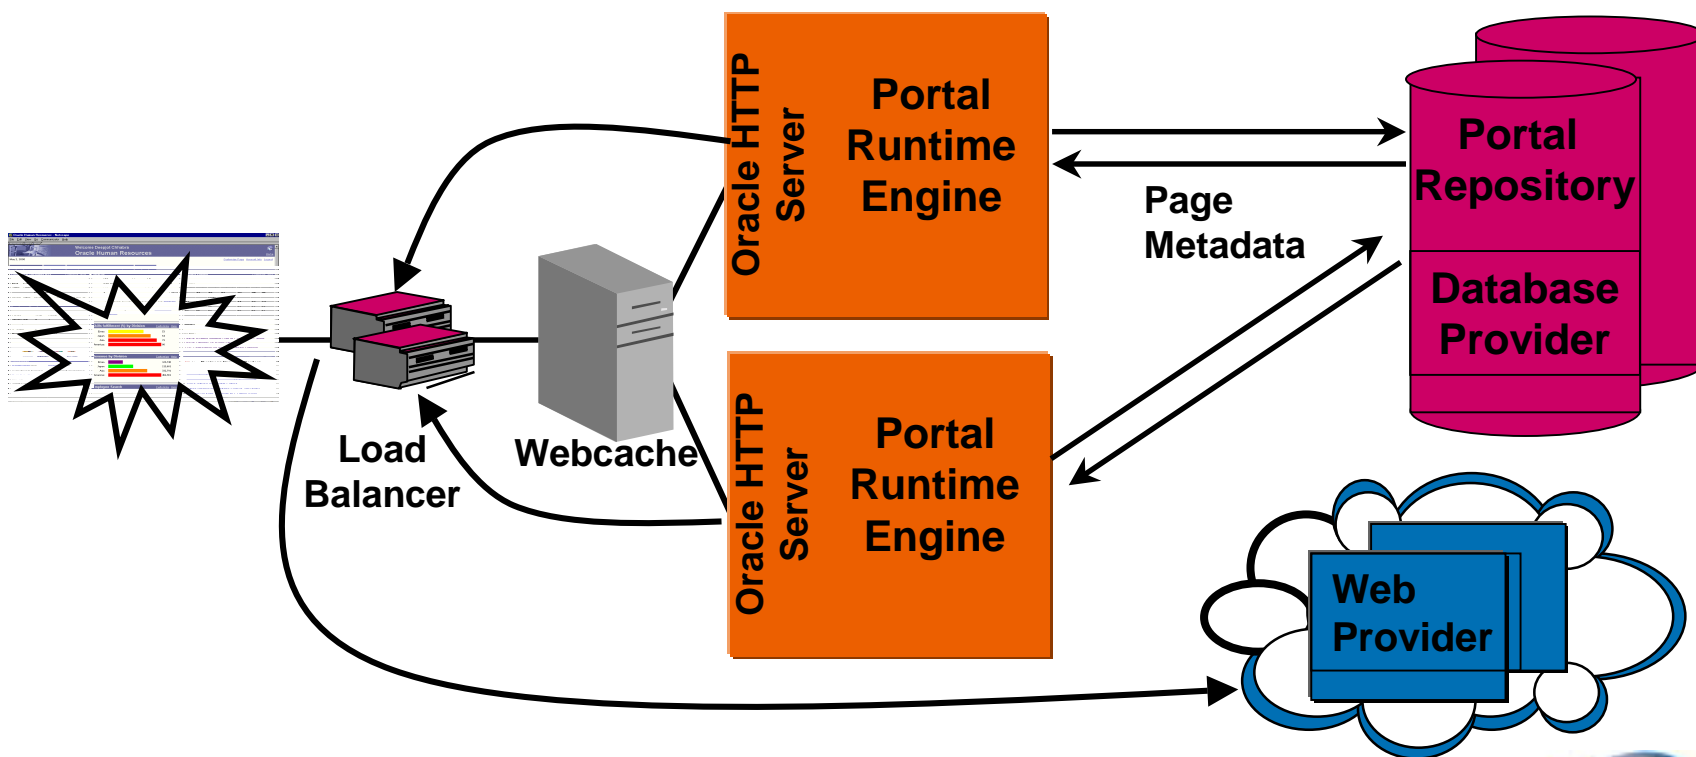


Portal HA

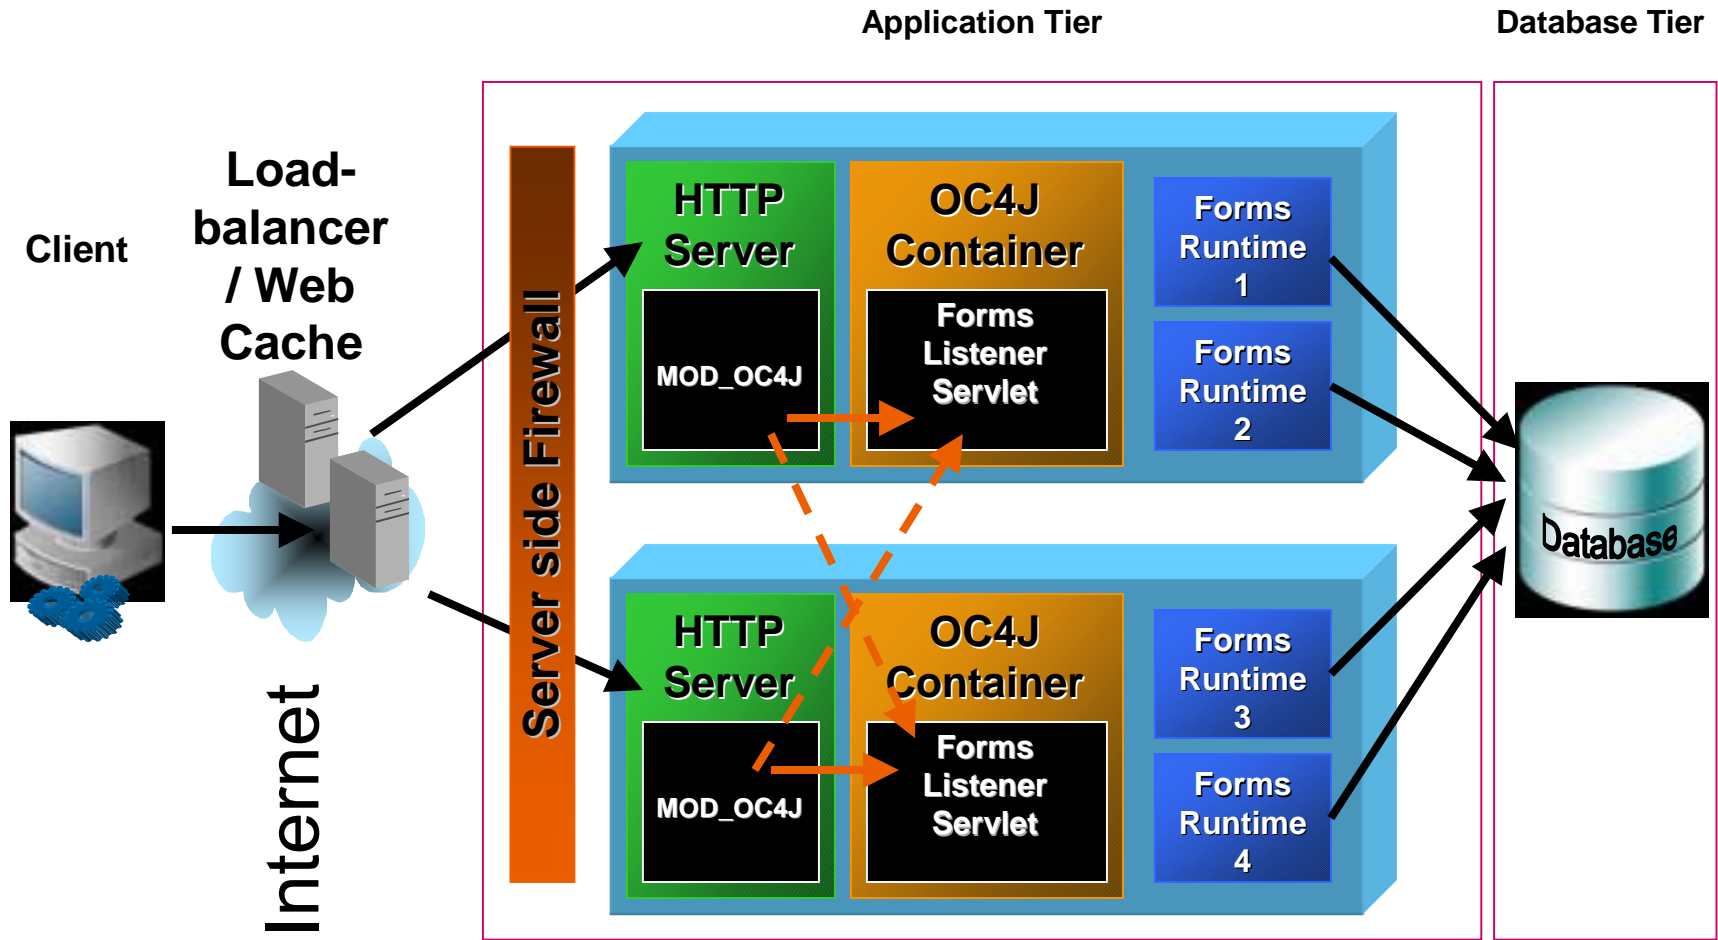
Browser Client

OracleAS Portal Middle Tier Cluster

OracleAS Portal Database Tier RAC



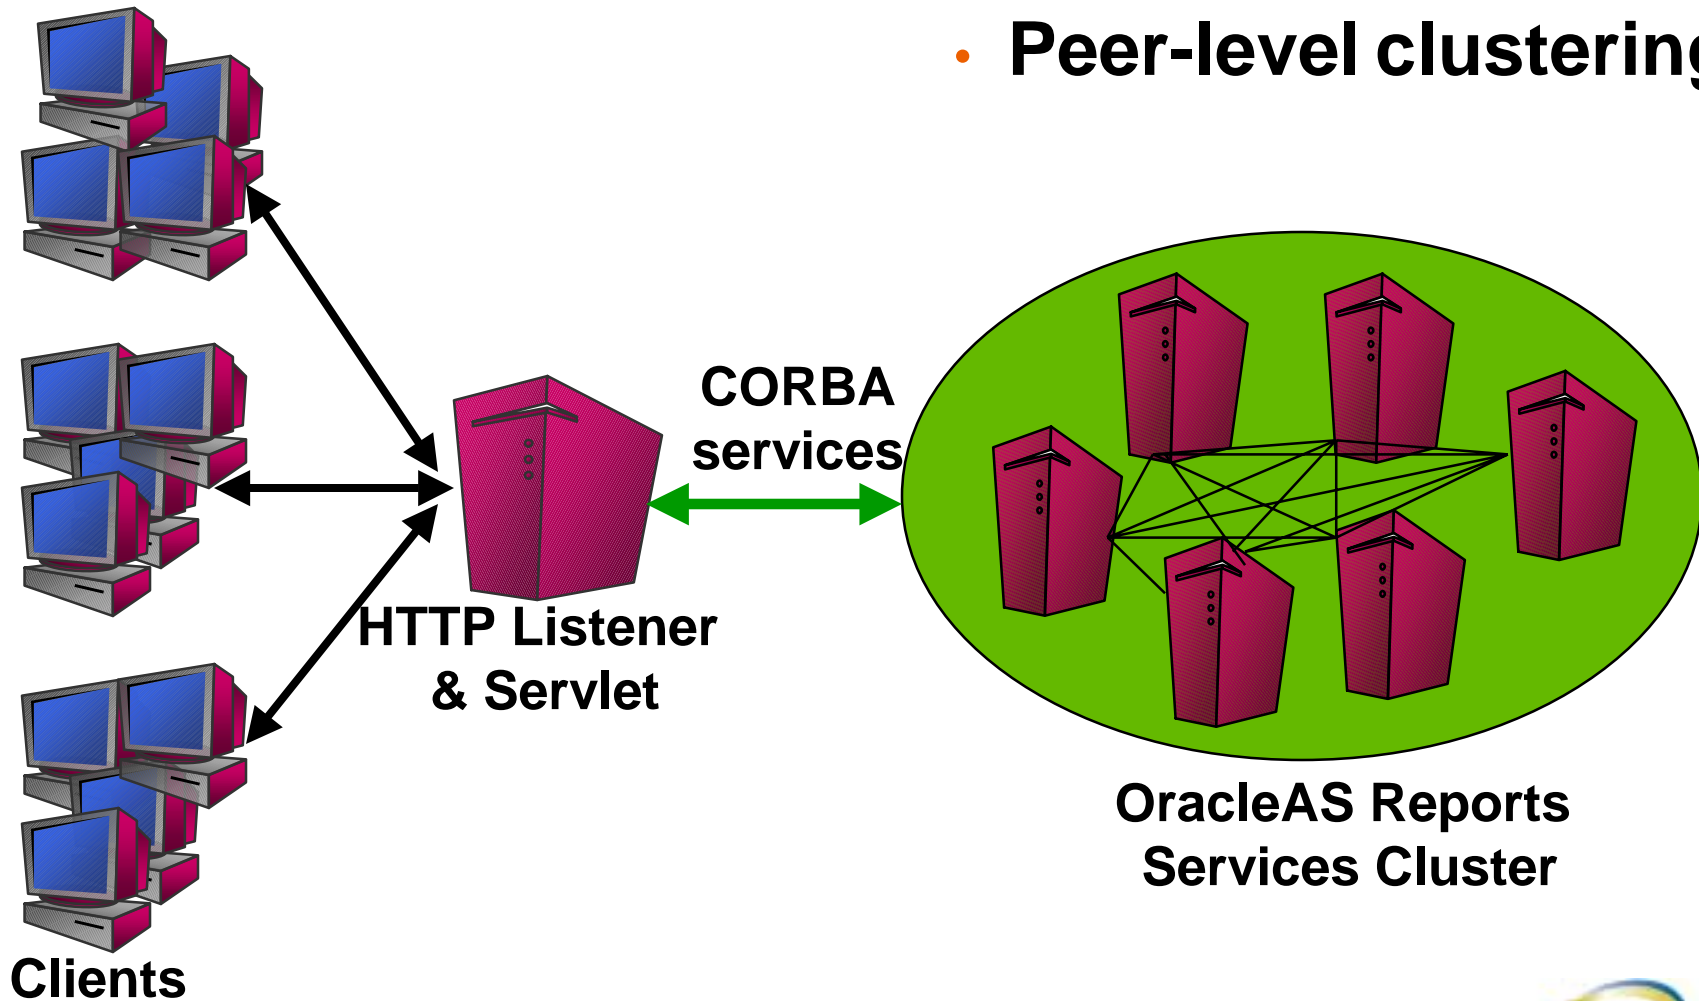
Forms HA



Reports – High Availability



- Peer-level clustering



Agenda

- Introduction
- Oracle Application Server High Availability Features
 - Overall process control: Self-healing
 - All-sub tiers clustering: state replication and load balancing (crash prevention)
 - High Availability for Metadata and Security Services
 - Backup&Recovery and Disaster Recovery
 - Deployment features
- Sample Configurations/Summary/Q&A

Metadata and system services High availability



- Critical Services
 - Identity Management Services – for SSO/OID
 - Metadata Repository Services – for Portal, Wireless, Discoverer, etc.
- HA Options
 - Cold Failover Cluster (CFC)
 - Active Failover Cluster (AFC)

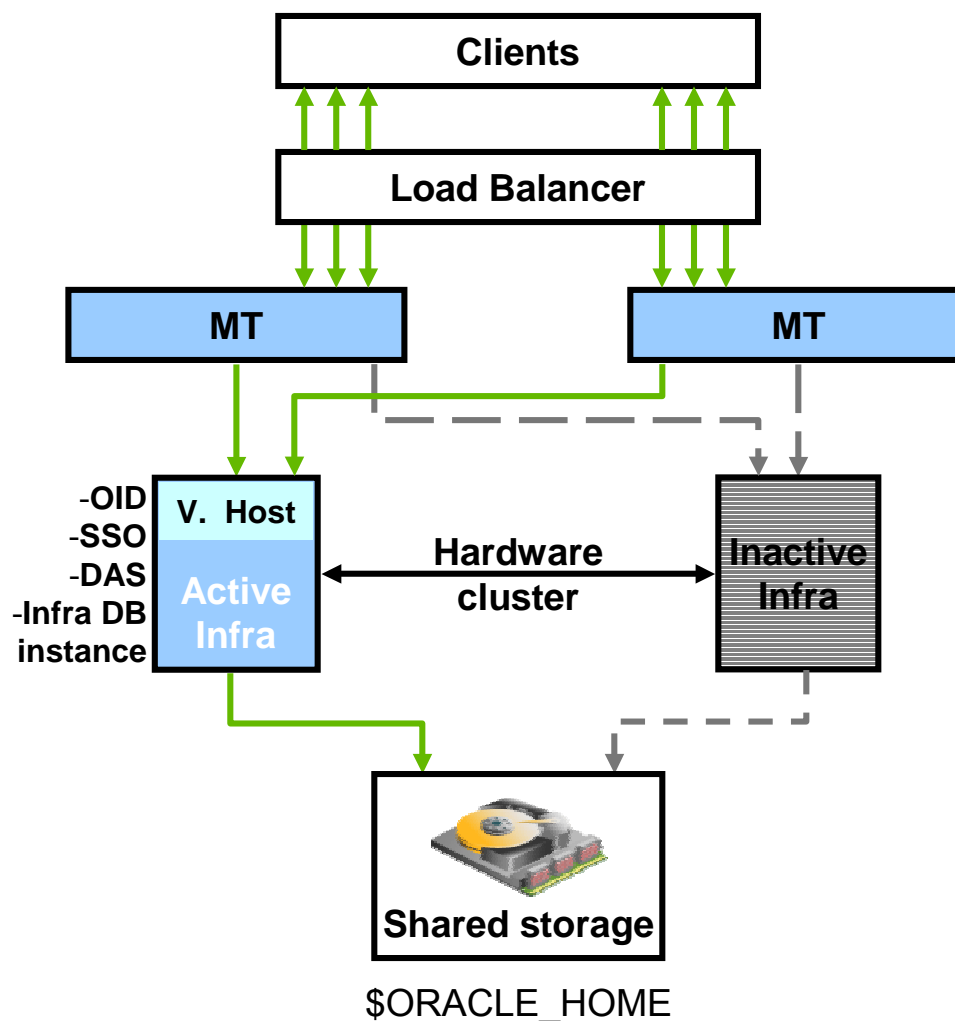
System Failure



Cold Failover Cluster

- Active/Passive setup:
 - One node is “hot” and the other node is “cold”
- Single Install
 - Using virtual hostname & IP (the clusterware associates them with the active node)
 - On shared storage device

Cold Failover Cluster



System Failure

Middle Tiers on Cold Failover Cluster

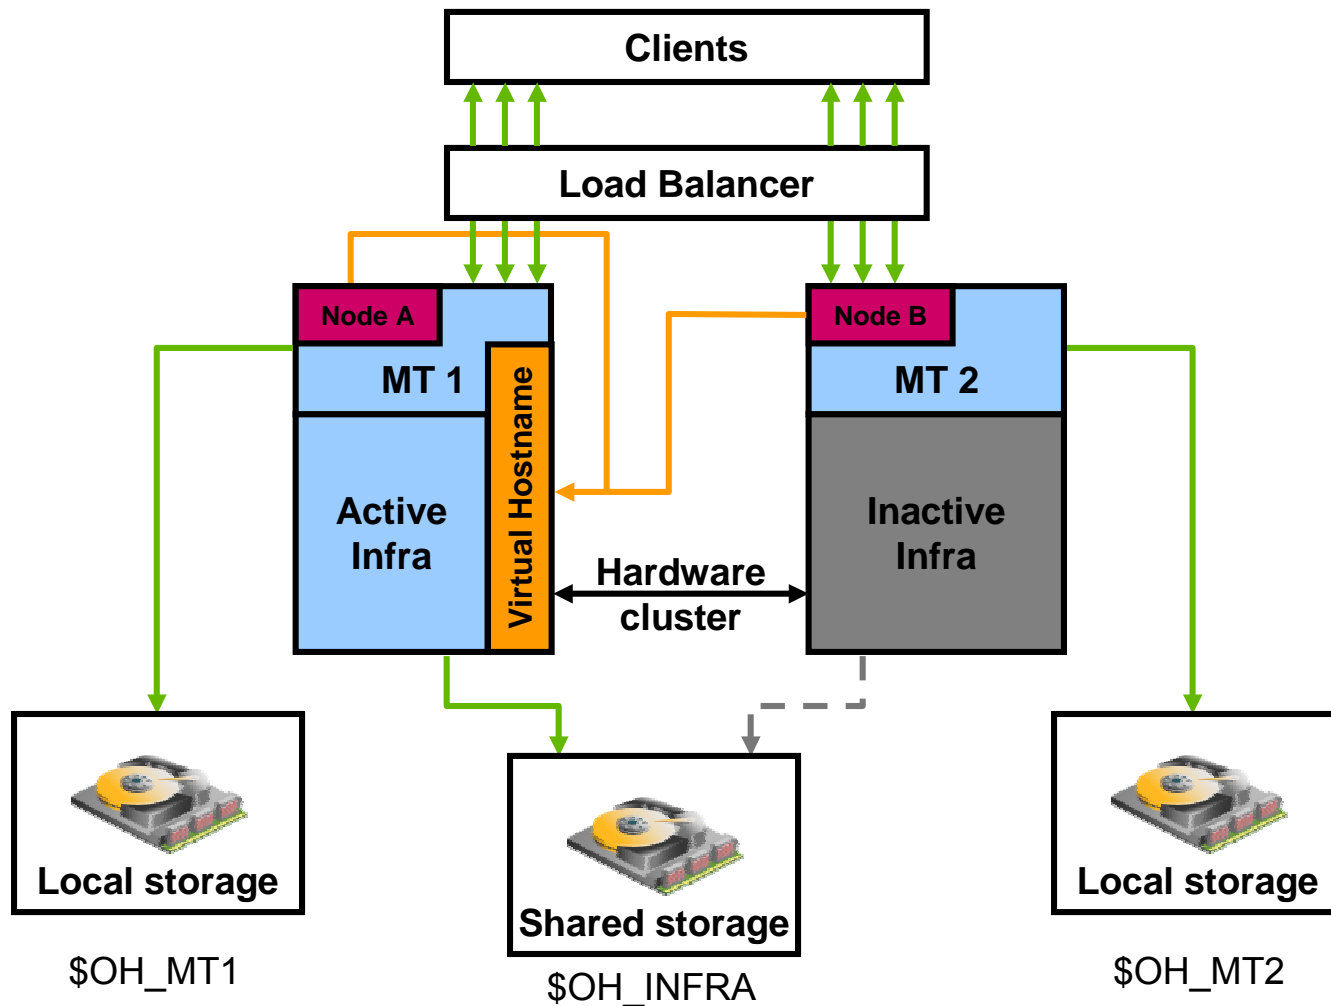


- Install Middle Tiers on CFC nodes
- **No** Failover of Middle Tiers
 - But death detection and restart still available
- Use LBR to route requests across the nodes
- Middle Tiers are installed using physical hostname
 - On local disk drives
- Infrastructure is installed using virtual hostname
 - On shared disk drives

System Failure



Middle Tiers on Cold Failover Cluster



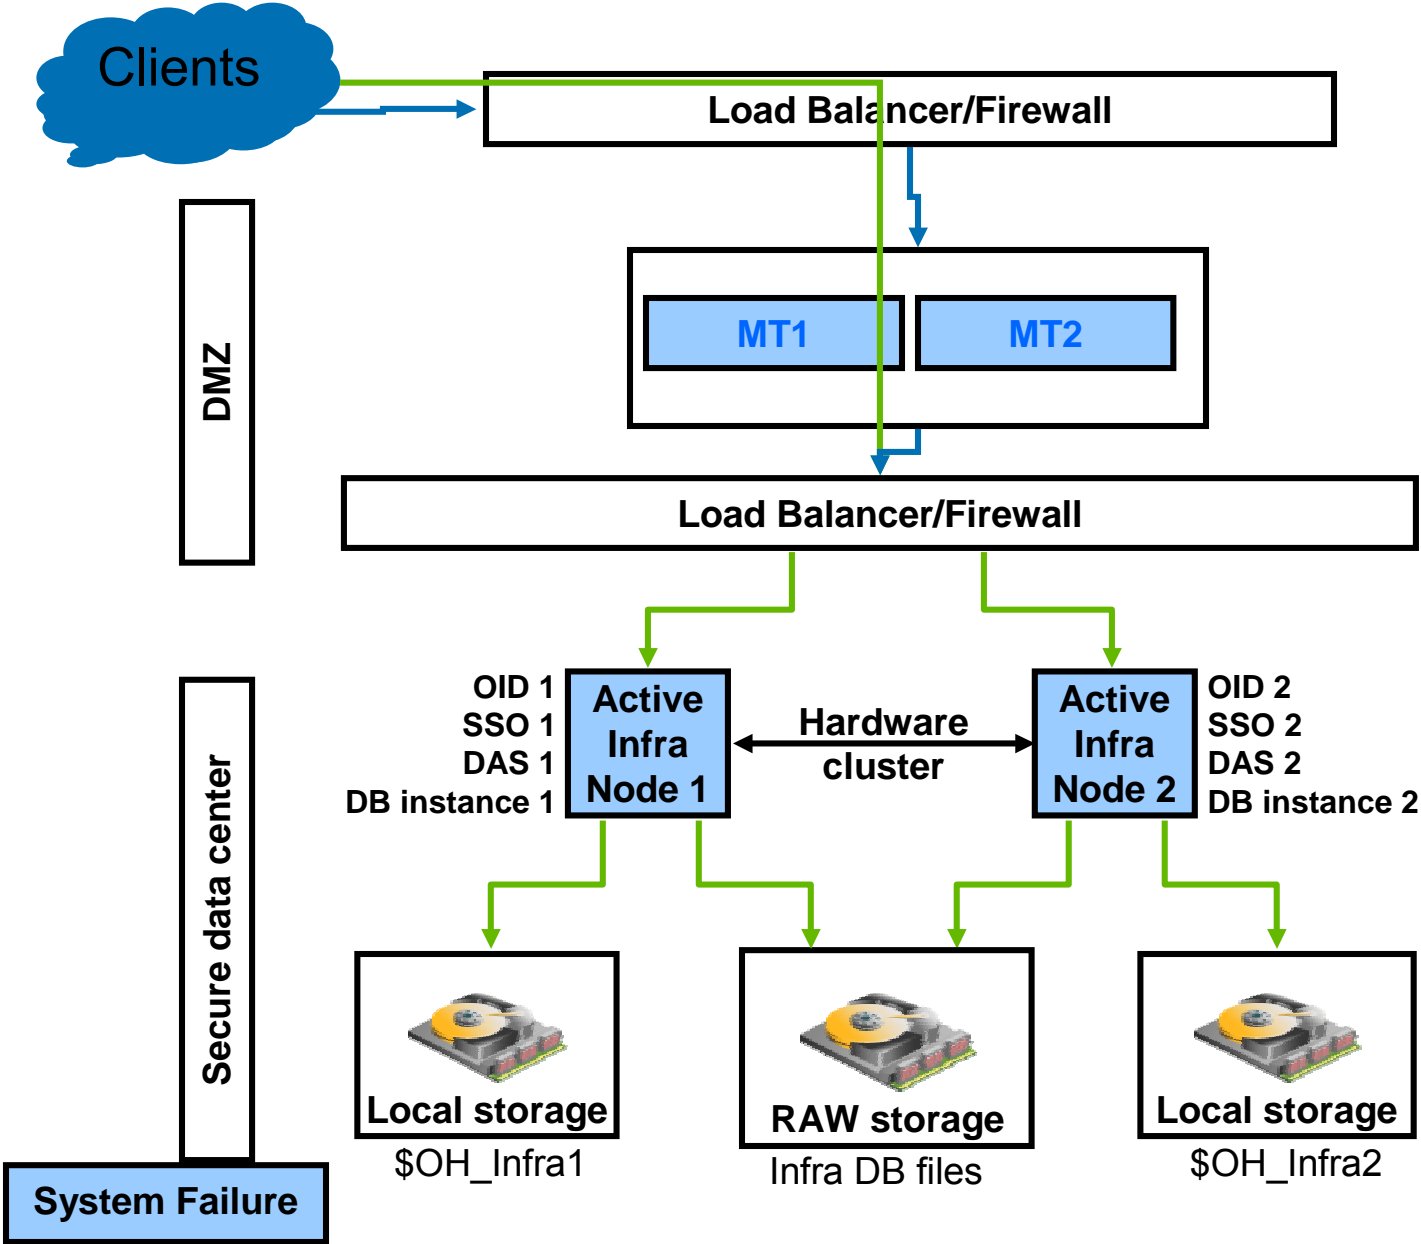
System Failure



Active Failover Cluster

- Active/Active setup: Provides “high 9s” uptime
- Single Install
 - RAC database files on shared RAW disk
 - Infrastructure \$ORACLE_HOME on local storage
 - Requires hardware load balancer for LDAP and HTTP traffic
- Also provides scalability

Active Failover Cluster



Agenda



- Introduction
- Oracle Application Server High Availability Features
 - Overall process control: Self-healing
 - All-subtiers clustering: state replication and load balancing (crash prevention)
 - High Availability for Metadata and Security Services
 - Backup&Recovery and Disaster Recovery
 - Deployment features
- Sample Configurations/Summary/Q&A



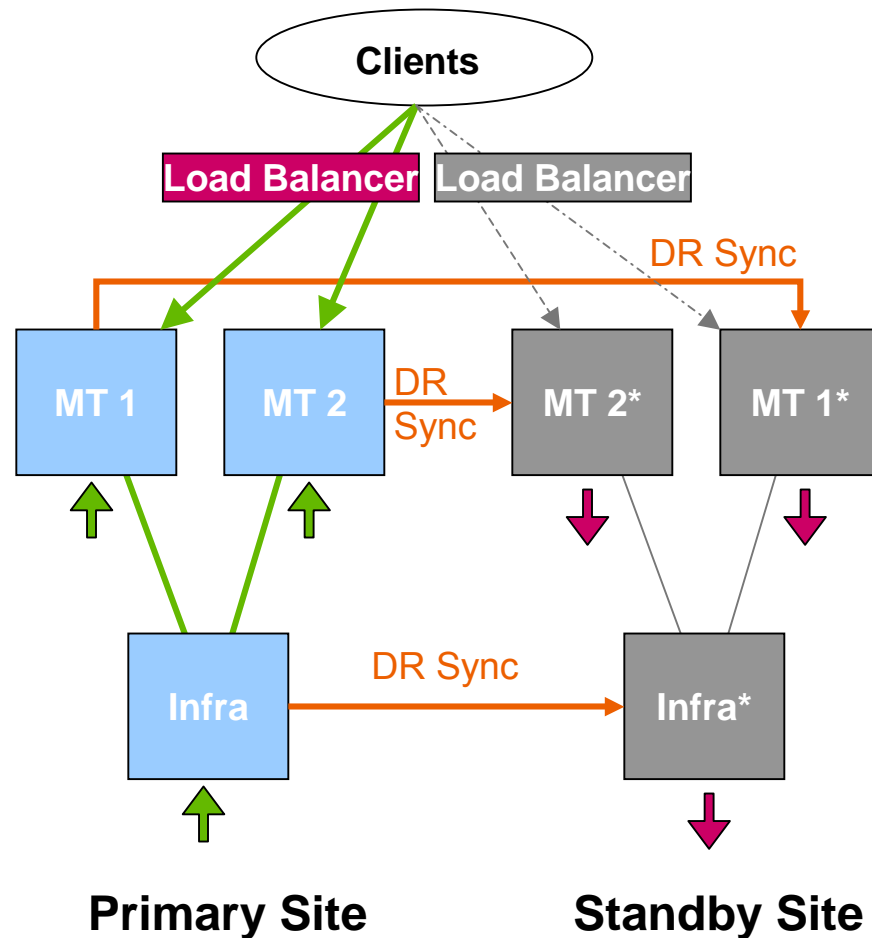
Backup and Recovery

- Backup and Recovery Tool
 - For Middle Tiers as well as Infrastructure
 - Incremental backup and recovery
 - Point-in-time recovery
- DCM Archival and Recovery
 - Revert to a previous configuration
 - Low overhead, quick “undo”
 - Works with DCM-managed components
 - Ideal for quick recovery from human error

Using DCM Archiving

- Archives created using `dcmctl createArchive` have the following properties:
 - Archives contain configuration and application deployment information
 - Archives can be moved from:
 - database repository to database repository
 - file-based to file-based repository
 - database to file-based repository
 - file-based to database repository
 - Can be applied to any compatible instance
 - Can be exported to a file

Disaster Recovery



- Site-to-site DR
- Active/Standby DR using:
 - Virtual hostnames
 - Backup and recovery to synchronize primary & standby sites
 - Data Guard for Infrastructure database

Data failure and Disaster



Metadata Storage and Changes

OracleAS metadata infrastructure storage

- In the Infrastructure database
- In configuration files within the Oracle Home

OracleAS midtier storage

- In configuration files within the Oracle Home

Considerations for Configuration Changes

- Modifies a collection of configuration files
- The component metadata within the infrastructure database.
- Mid-tier components cache and maintain information in the mid-tier
- System Administrators use different utilities to modify configuration
 - Oracle administrations [e.g.. Oracle Grid Control]
 - Jdeveloper deployment
 - text editor

DR requires controlled changes to configurations and application deployment.

Standby Site Synchronization

OracleAS 10g DR utilizes Backup/Restore and Oracle Data Guard to synchronize a standby site.

Oracle Data Guard log transfer and apply is used to sync the infrastructure database

- Managed recover is disabled

Backup/Restore is used to collect/restore all changed Oracle home files

- Full B/R is used for standby site instantiation.
- Incremental B/R is used for standby site synchronization
- Correlate the backup file to the current transaction level of the infrastructure database

User defined transfer techniques are used to transfer and stage backup files to standby.

Agenda

- Introduction
- Oracle Application Server High Availability Features
 - Overall process control: Self-healing
 - All-subtiers clustering: state replication and load balancing (crash prevention)
 - High Availability for Metadata and Security Services
 - Backup&Recovery and Disaster Recovery
 - Deployment features
- Sample Configurations/Summary/Q&A

OC4J Hot Deployment*

	Hot Deployment	Hot Redeployment (Stateless Apps)	Hot Redeployment (Stateful Apps)
Web Tier			
OracleAS OC4J	Yes	Yes	Yes
OC4J standalone	Yes	Yes	Yes
OC4J Clusters (islands)	Yes	Yes	No
EJB Tier			
OracleAS OC4J	Yes	No	No
OC4J Standalone	Yes	No	No
OC4J Clusters (islands)	Yes	No	No

* Application must be deployed as a packaged .ear or .war; no editing of descriptor files permitted

* Certain configuration restrictions apply – see doc set for details

System Maintenance

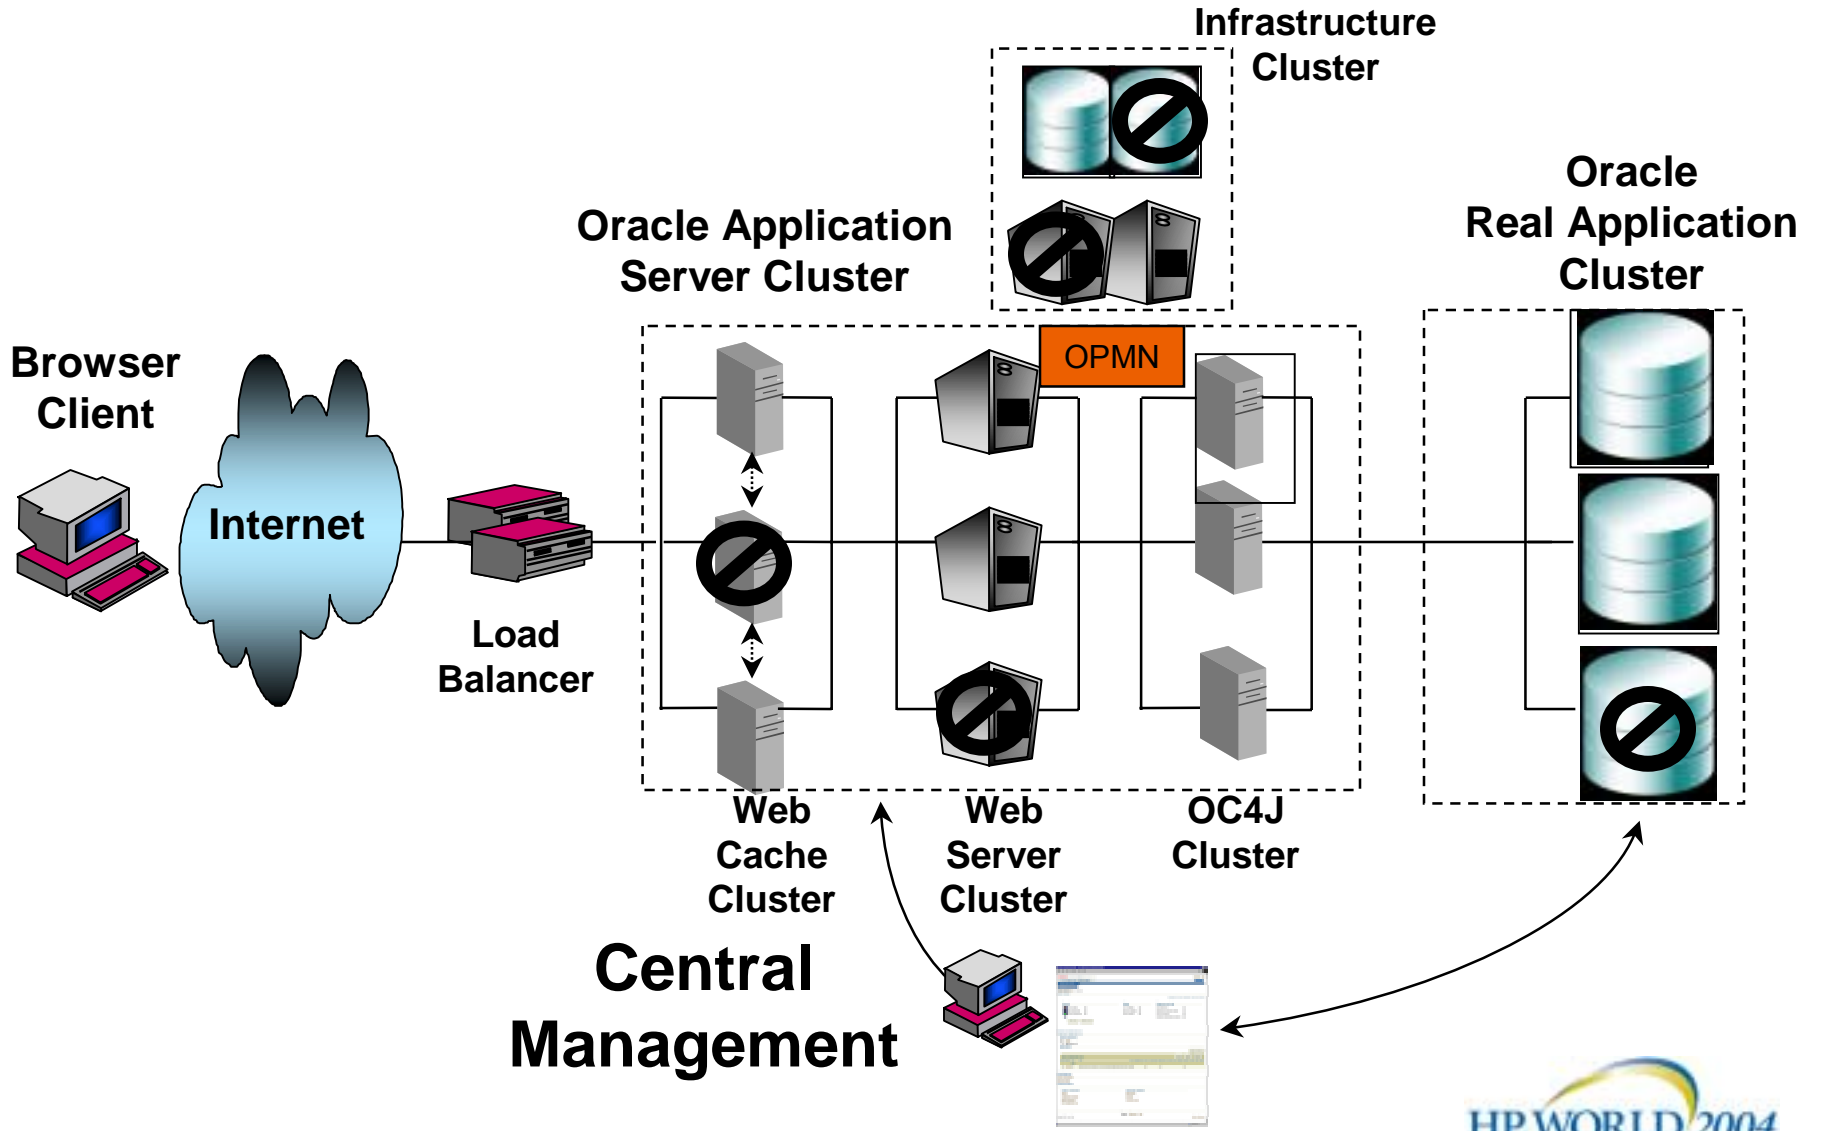
Agenda



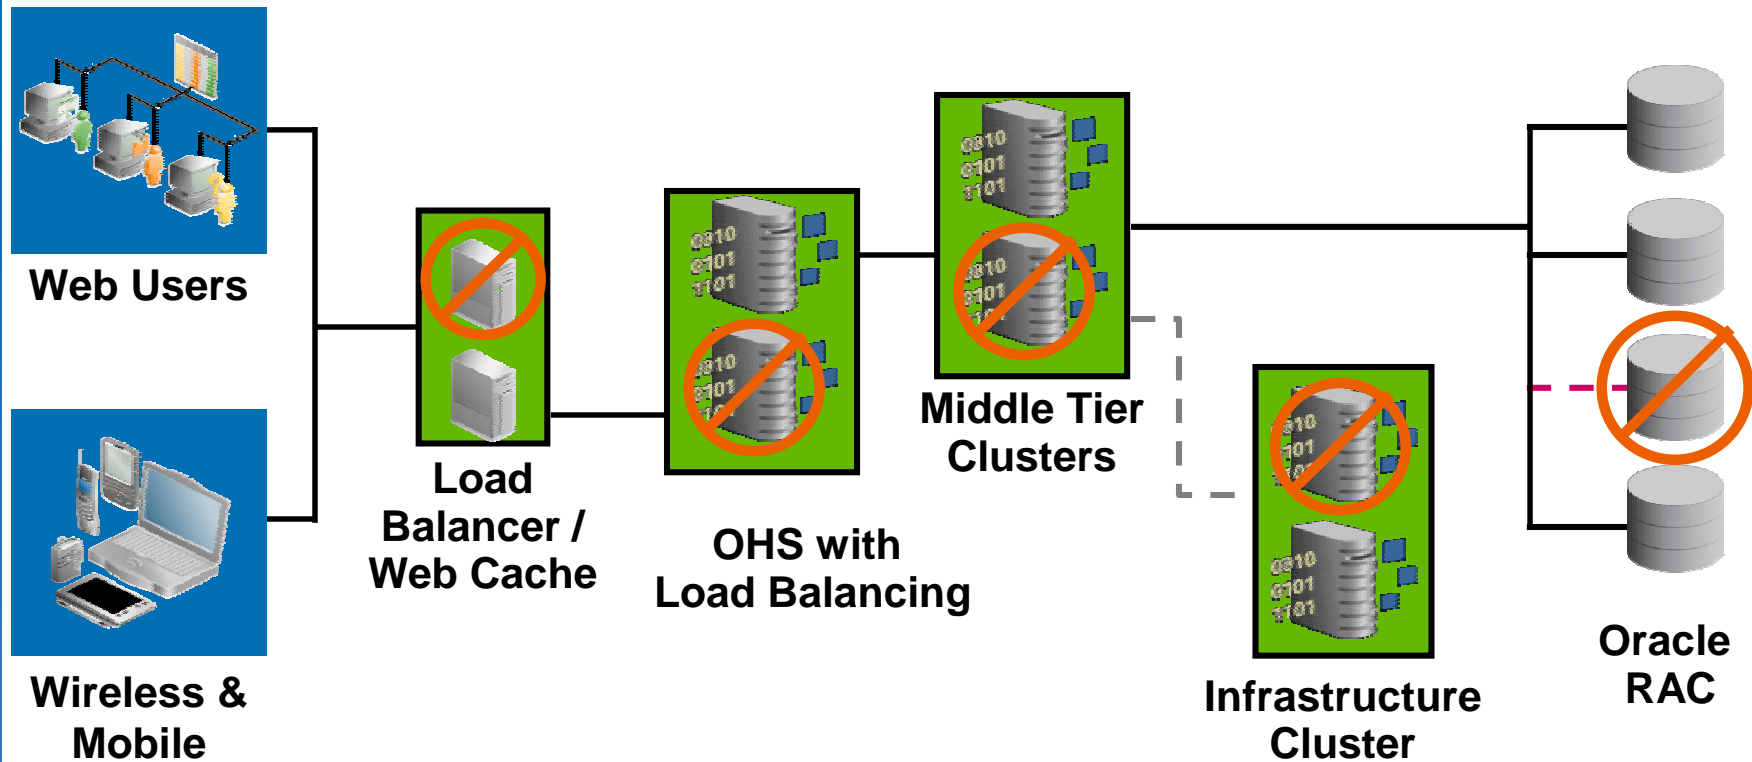
- Introduction
- Oracle Application Server High Availability Features
 - Overall process control: Self-healing
 - All-subtiers clustering: state replication and load balancing (crash prevention)
 - High Availability for Metadata and Security Services
 - Backup&Recovery and Disaster Recovery
 - Deployment features
- Sample Configurations/Summary/Q&A



No Single Point of Failure



HA Architectural Overview



Web Users

Wireless & Mobile

Load Balancer / Web Cache

OHS with Load Balancing

Middle Tier Clusters

Infrastructure Cluster

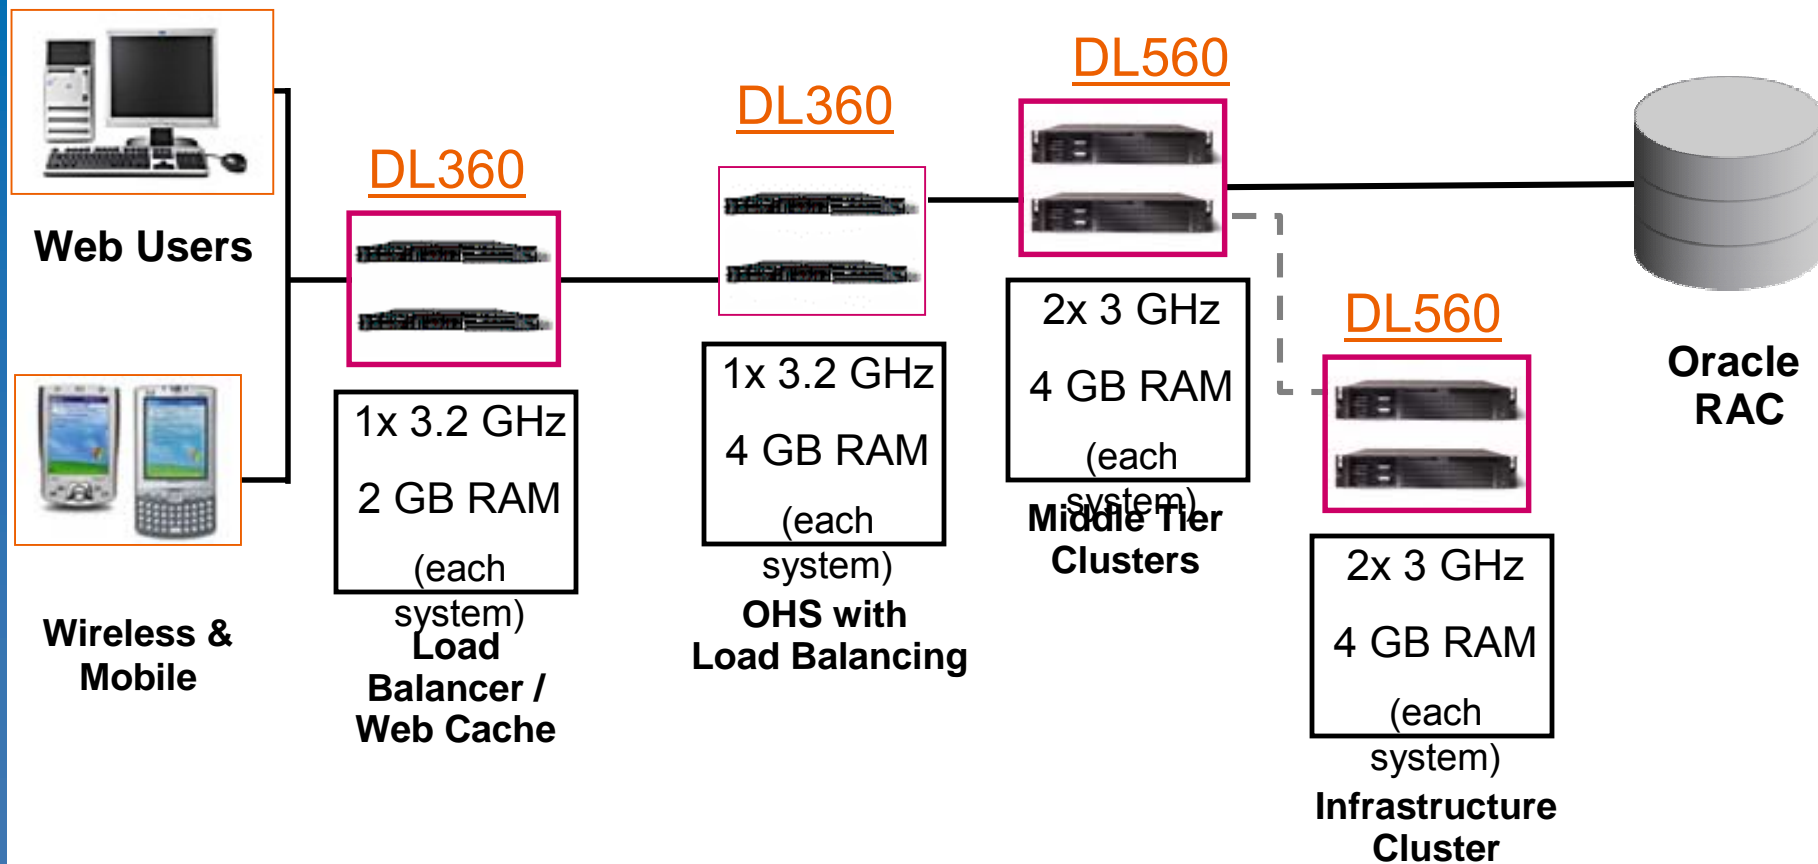
Oracle RAC

Client

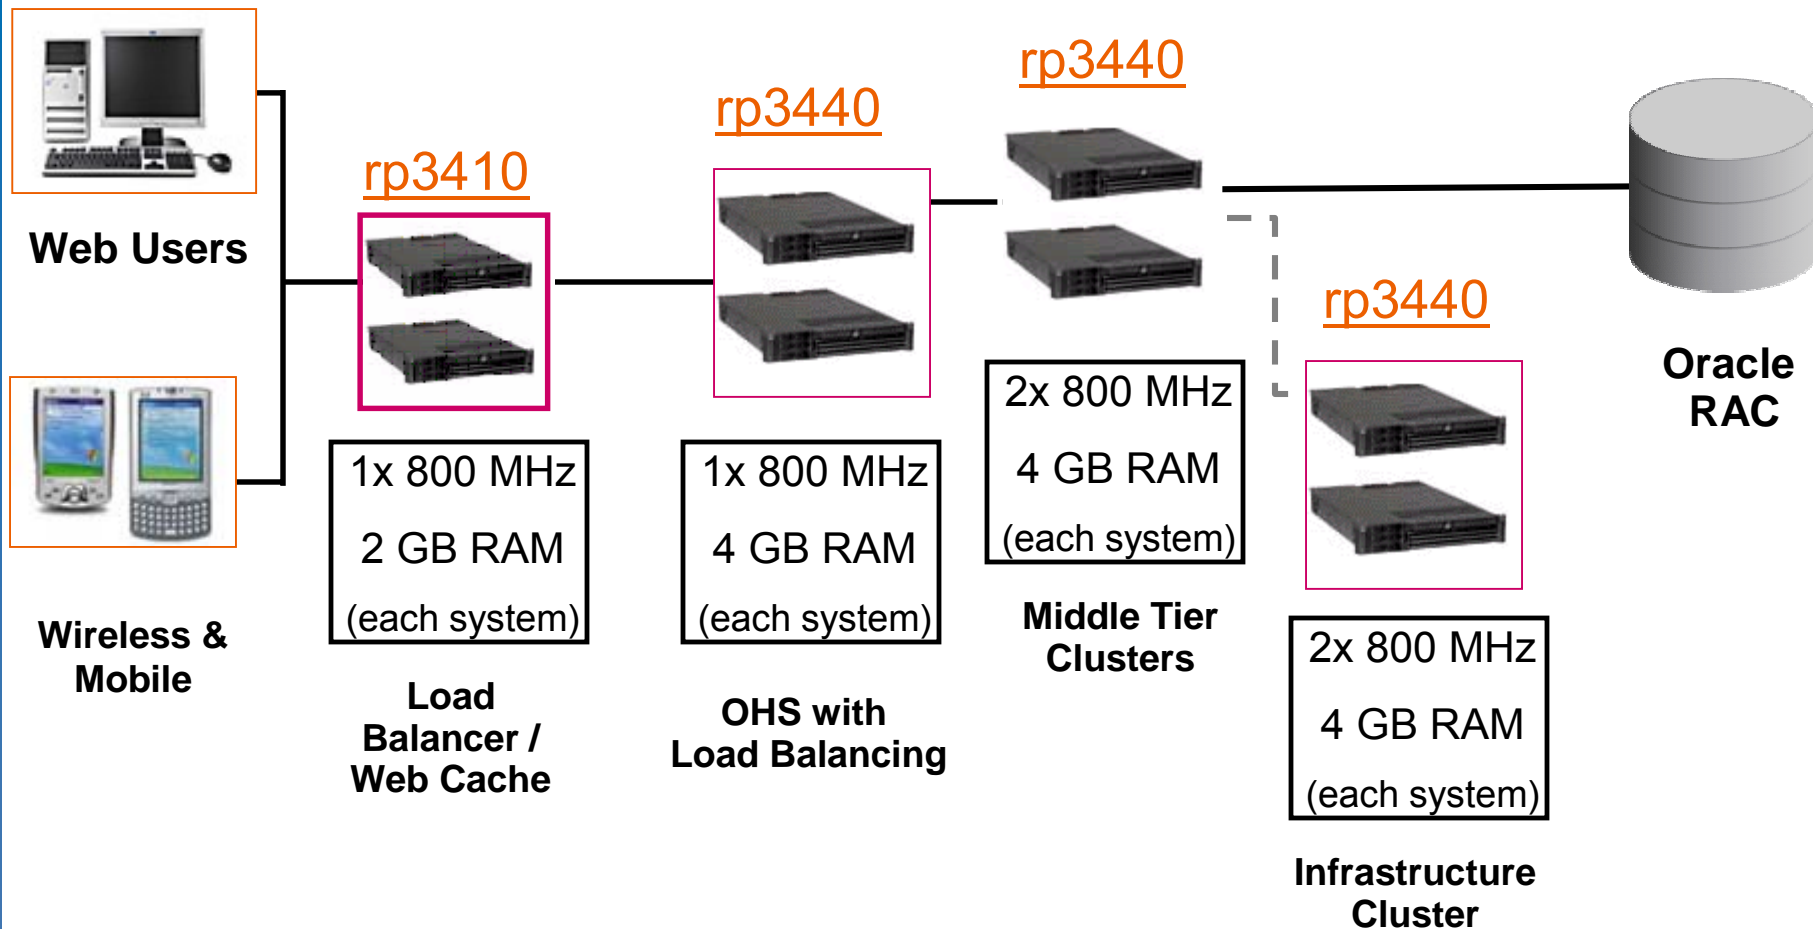
Oracle Application Server 10g

End-to-end High Availability

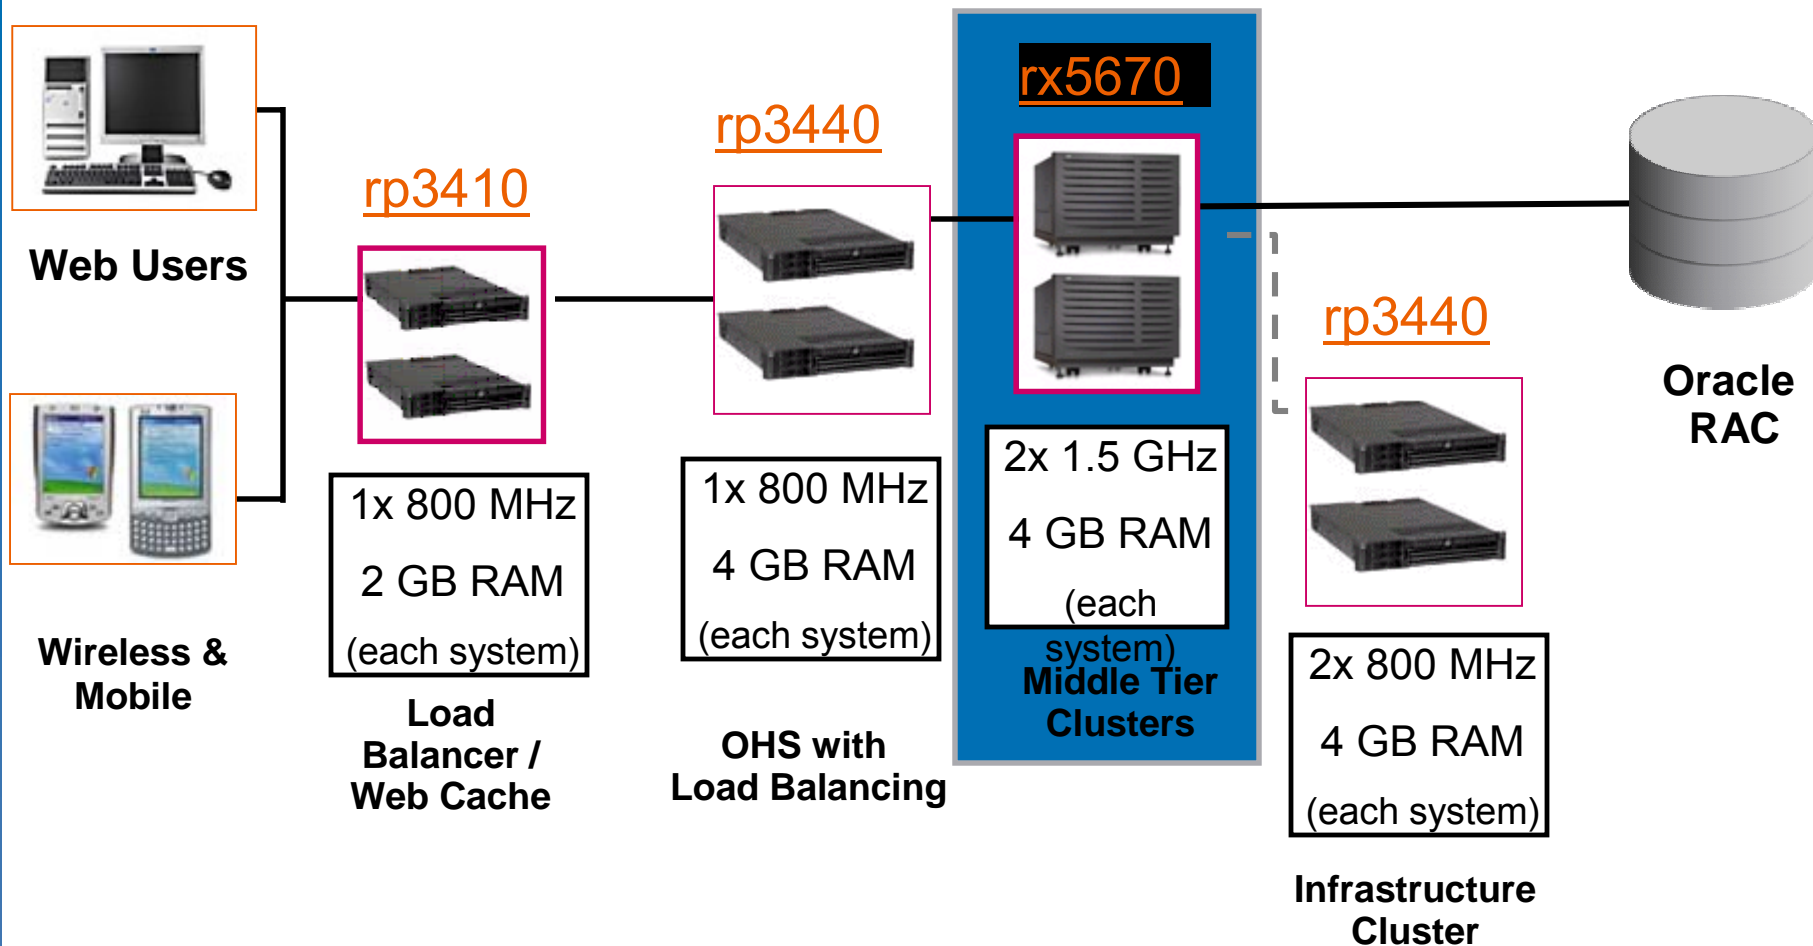
ProLiant (Windows/Linux)



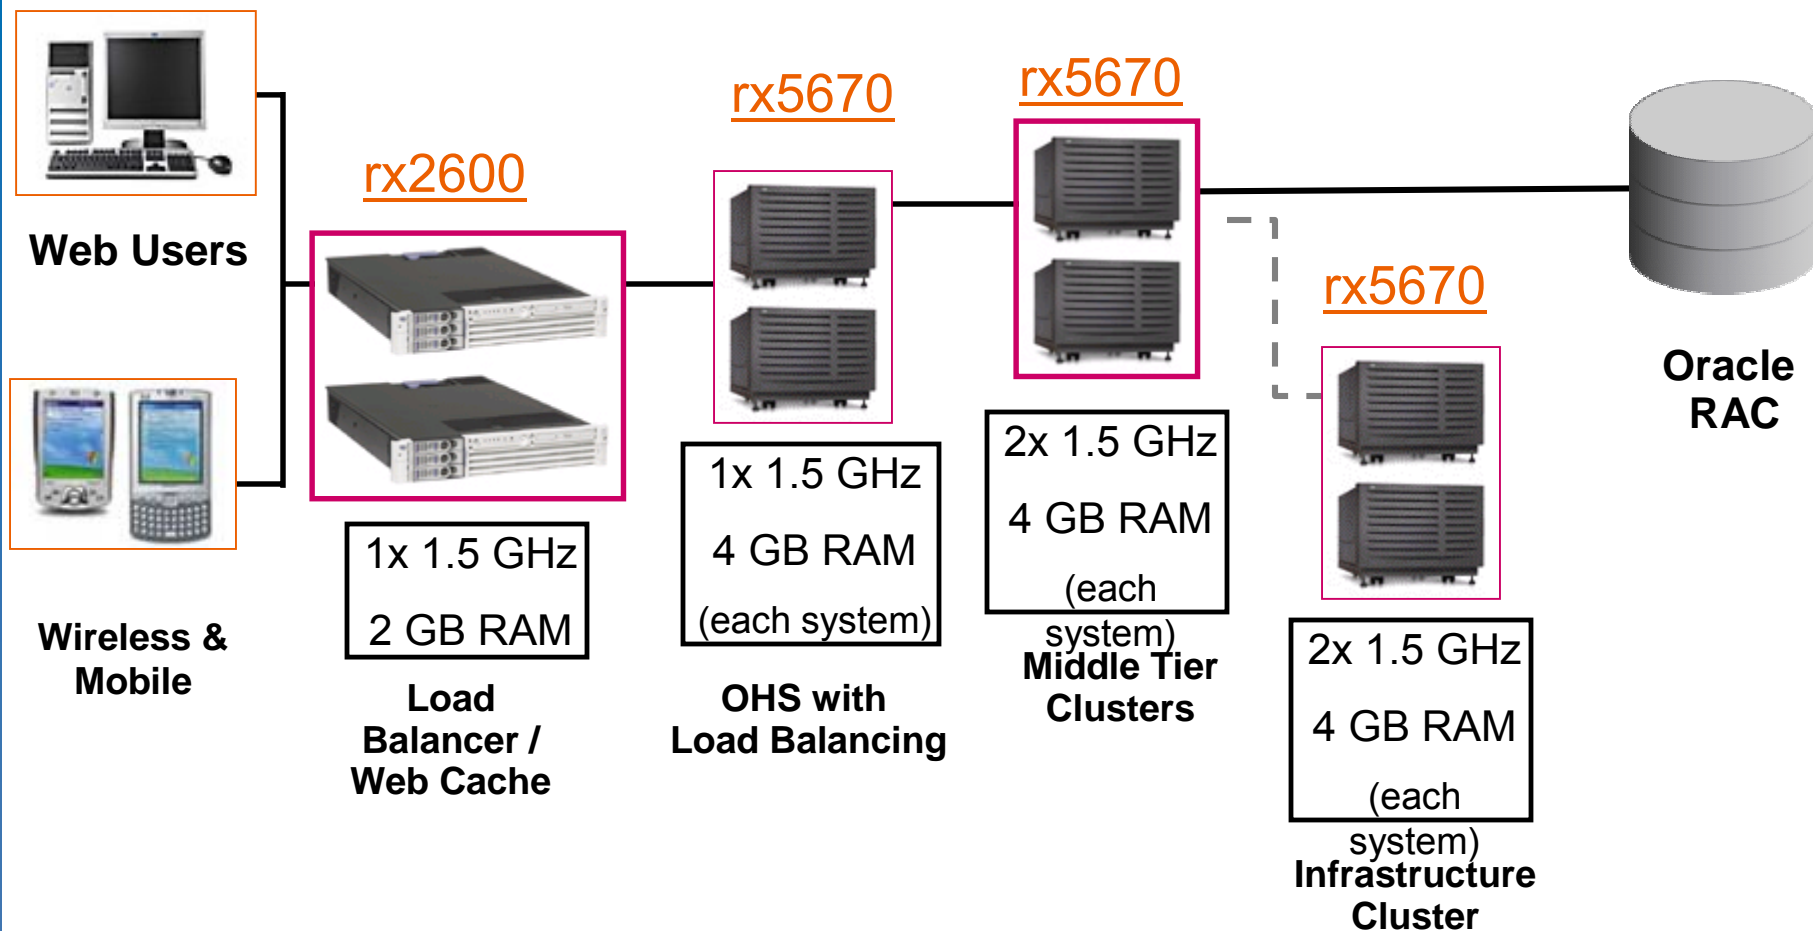
HP9000 (HP-UX)



Split config HP9000/Integrity server



Integrity server (not yet available)



Summary



- No Single Point of Failure: All tiers HA
- Delivering “intelligent HA”: prevention, dynamic resource allocation
- Fast Recovery
 - Self-healing – OPMN
 - Disaster Recovery
- Ease of Configuration
 - EM, DCM Clusters
 - Archival and Recovery
- Rolling Upgrades
- Minimal planned or unplanned downtime
- Determine a high availability solution best suited for customer’s business needs



Resources



- OTN:
 - http://otn.oracle.com/products/ias/hi_av/index.html
- Oracle Application Server Documentation:
 - High Availability Guide
 - Install Guide (contains HA installation options)
 - Admin Guide (contains Backup & Recovery)





HP WORLD 2004

Solutions and Technology Conference & Expo

Co-produced by:



RECOMMENDED TRAINING VENUE FOR THE
HP Certified Professional





Backup slides

Oracle
Application
Server 10g on
Itanium

Itanium Roadmap – Phase 1



Currently only J2EE (Oracle Application Server Containers for J2EE (OC4J) and OracleAS TopLink) are available on Itanium-2 processor based systems.

Oracle Application Server 10g Release 2 (10.1.2) - Phase 1 will be available in the August-October, 2004 timeframe and support additionally the following components:

- Portal
- Wireless
- Business Intelligence
- Business Integration
- Identity Management
- Web Caching





Itanium Roadmap – Phase 2

Oracle Application Server 10g Release 2 (10.1.2) - Phase 2 will be available in the October-December timeframe and updates the Phase 1 components to the following:

- J2EE (Oracle Application Server Containers for J2EE (OC4J) and OracleAS TopLink)
 - Portal
 - Wireless
 - Business Intelligence
 - Business Integration
 - Identity Management
 - Web Caching
 - **Forms and Reports**
- Phase 1**
- Phase 2**





Itanium Roadmap – Available platforms

Oracle Application Server 10g Release 2 (10.1.2) will be available on the following platforms:

Platform	O/S Version
HP-UX	HP-UX 11i v2 (also referred to as HP-UX version 11.23)
Linux	RedHat Linux EL 3.0 SUSE SLES 8.0 and 9.0 (based on market demand)
Windows	Windows 2003 Windows XP (for development purposes only)



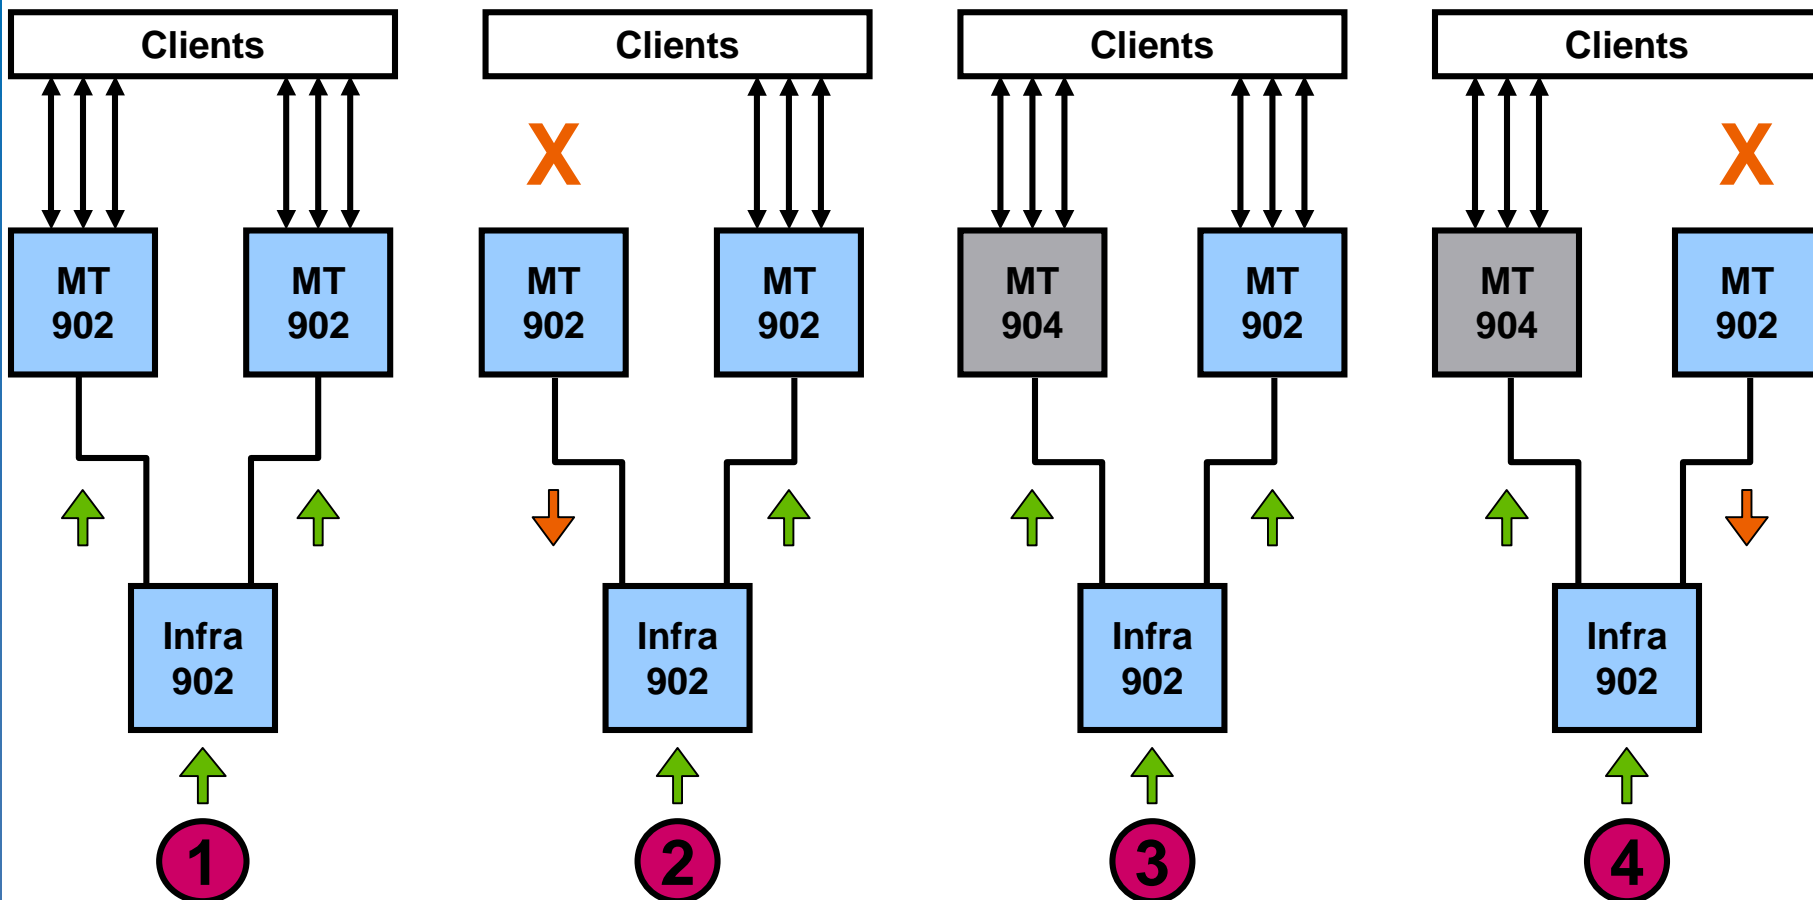
Backup slides

Rolling Upgrades



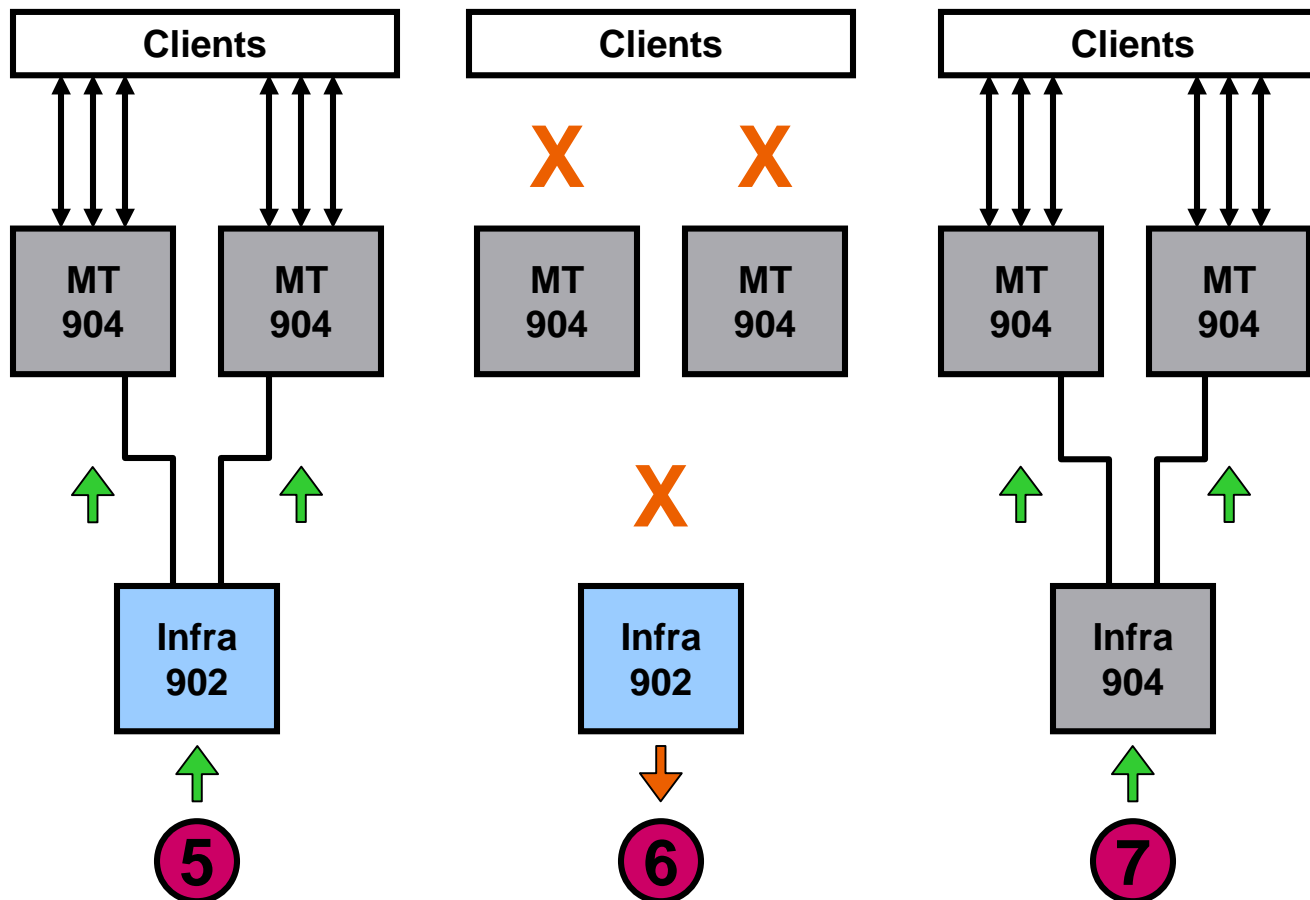
i n v e n t

Rolling Upgrades



System Maintenance

Rolling Upgrades – contd.



System Maintenance



Backup slides

OPMN

List of components integrated with OPMN



Explicit Components

1. Oracle HTTP Server
2. OC4J
3. Webcache
4. OID
5. Wireless
6. Reports
7. Discoverer
8. Process Connect
9. DCM Daemon
10. Log Loader
11. IASPT
12. Custom

Implicit Components

1. BC4J
2. CABO/BALI
3. Forms, SSO/Login Server
4. OAI(HTTP Adapter)
5. OraDAV
6. Personalization Server
7. Portal
8. Soap
9. Syndication Server
10. UDDI Server (web services)
11. Ultrasearch
12. TopLink



The opmn.xml Configuration File Used in Oracle9iAS, Release 2

```
<ias-instance>
  <process-manager>
    <ohs gid="ohs">
      </ohs>
      <oc4j instanceName="home">
        <island id="default_island" numProcs="1"/>
      </oc4j>
    </process-manager>
  </ias-instance>
```

System Failure



The New opmn.xml Configuration File



```
<opmn>
  <process-manager>
    <ias-instance id="10gTrain1">
      <ias-component id="HTTP_Server">
        <process-type id="HTTP_Server" module-id="OHS">
          <process-set id="HTTP_Server" numprocs="1"/>
        </process-type>
      </ias-component>
      <ias-component id="OC4J">
        <process-type id="home" module-id="OC4J">
          <process-set id="default_island"
numprocs="1"/>
        </process-type>
      </ias-component>
    </ias-instance>
  </process-manager>
</opmn>
```

System Failure

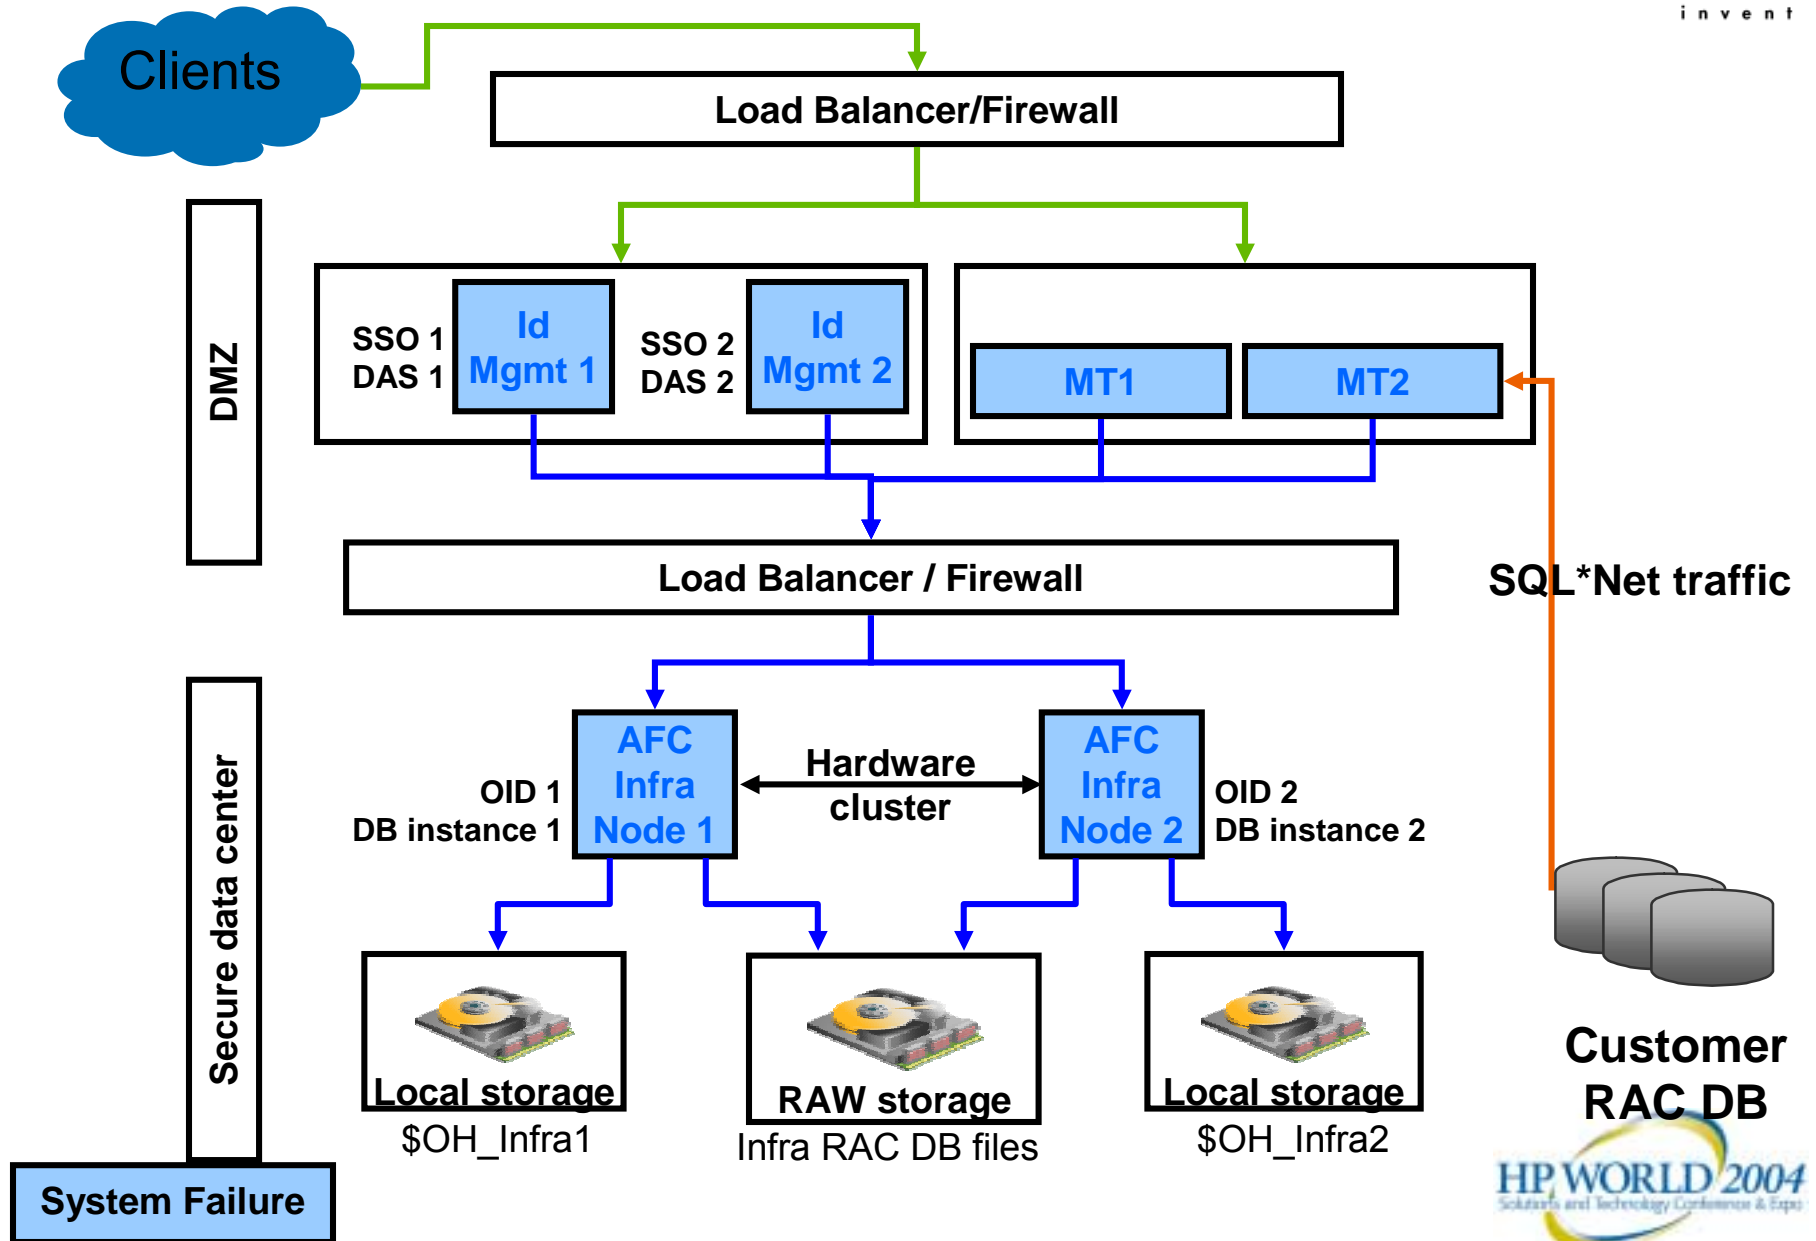




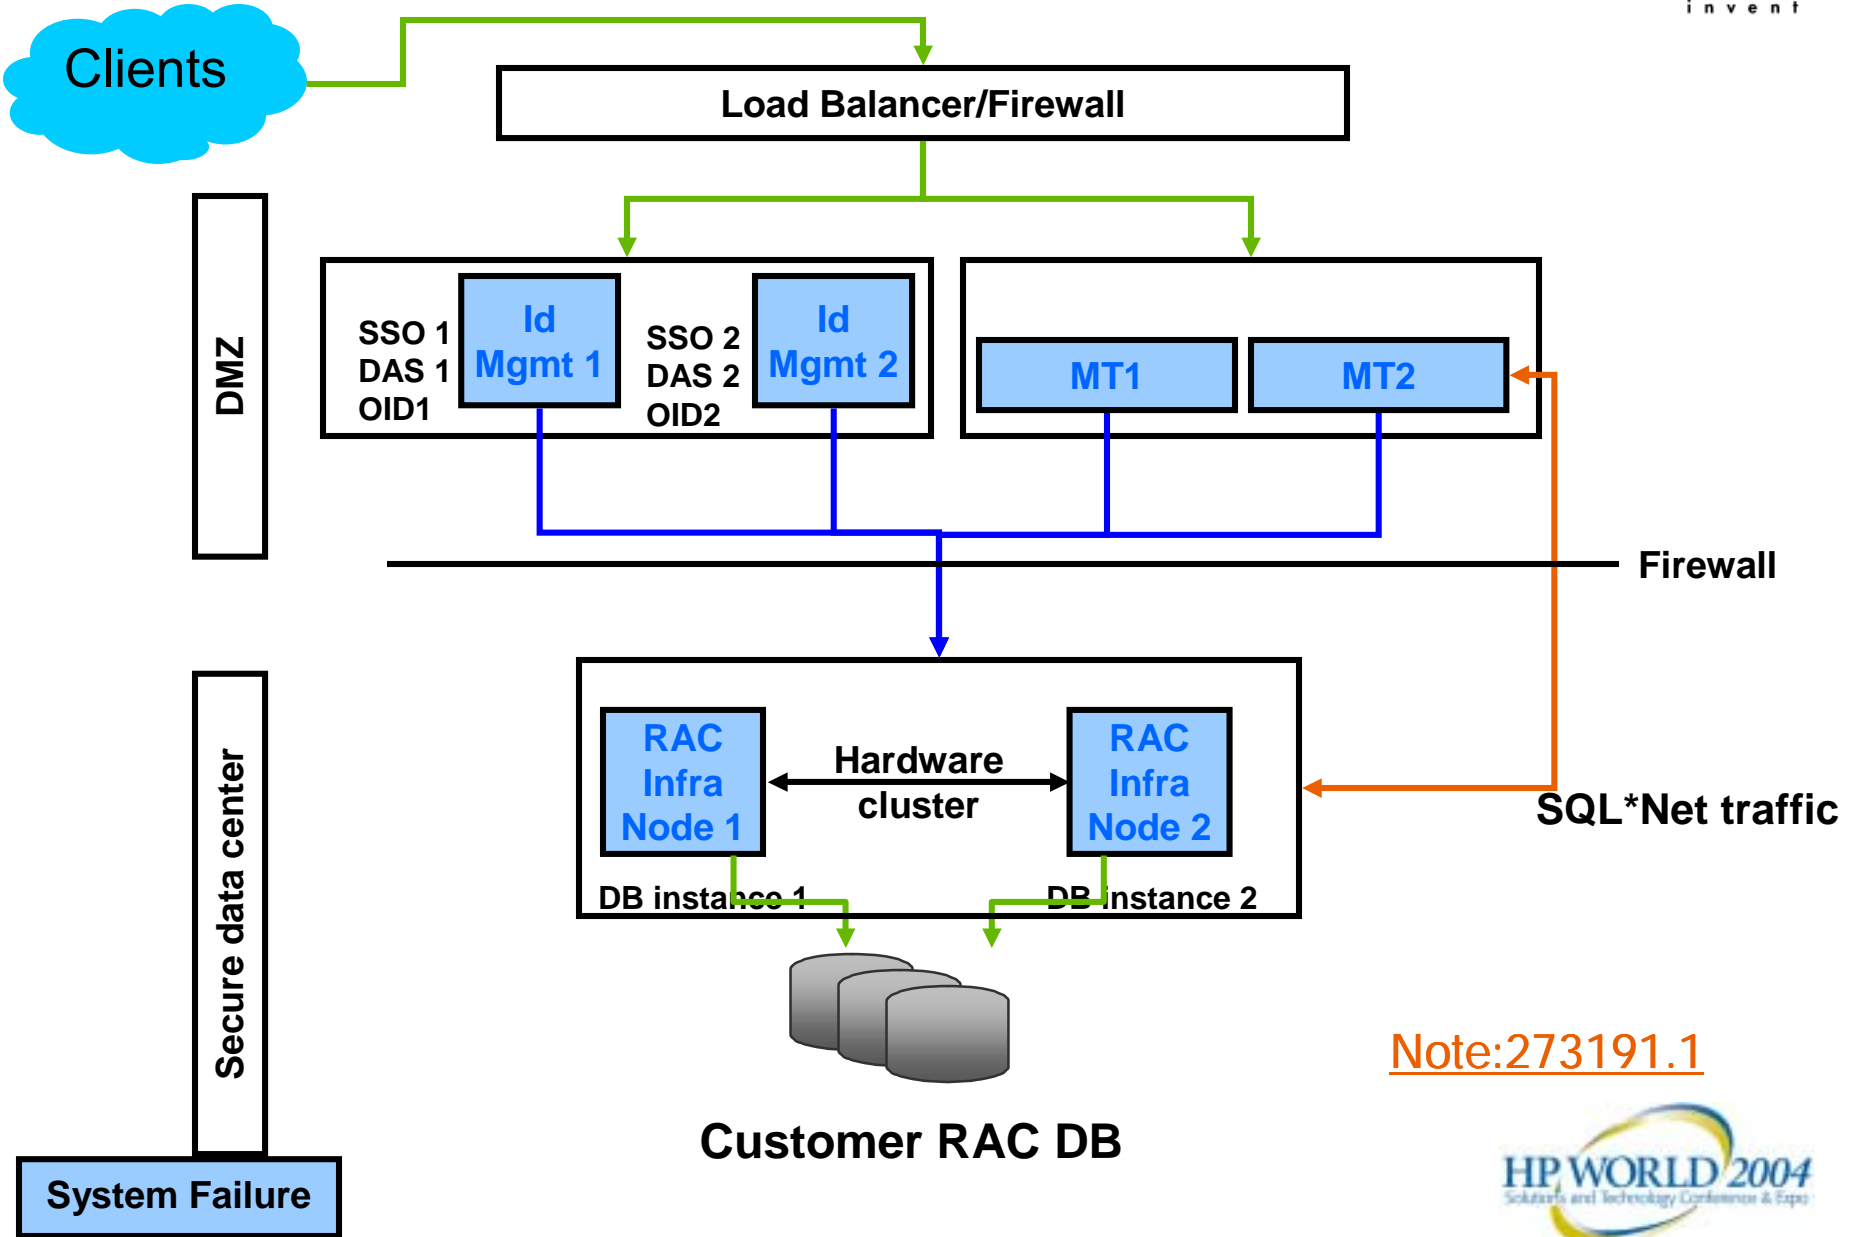
Backup slides

ID Management/ repCa

Distributed ID Mgmt with AFC



Architecture (acknowledged) and report to RAC (I)



[Note:273191.1](#)



IM-HA(CFC) and repCa to RAC (II)

