

Oracle for administrative, technical and Tier-0 mass storage services

openlab Major Review Meeting

29 September 2009

Lucia Moreno Lopez,
Carlos Garcia Fernandez (presenter),
Ruben Domingo Gaspar Aparicio,
Eric Grancher,
Juan Manuel Guijarro,
Anton Topurov



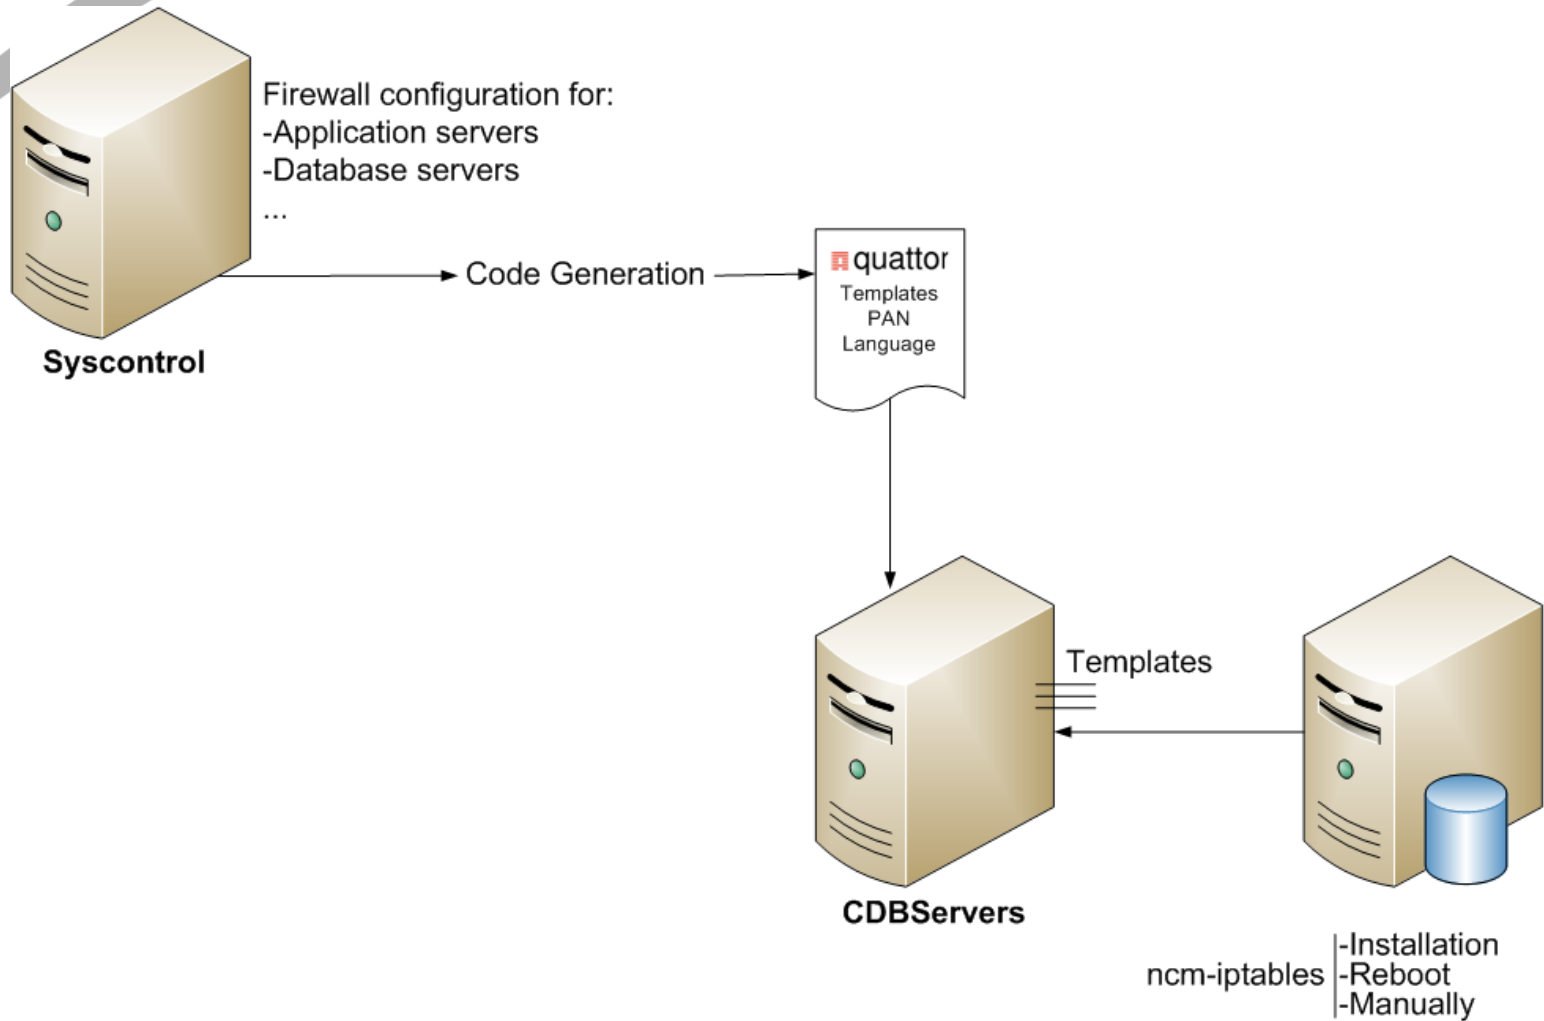
Main areas of activity

- **Security**
- Oracle Database Virtualization
- Joint Software Testing Program (mostly 11.2 database for the past 6 months)
- Monitoring
 - Web applications monitoring using EM / beacons
 - Database policies

- **CERN ELFms** (Extremely Large Fabric management system)
- **Quattor** (System installation and configuration toolsuite)
- **Syscontrol** (Management and Inventory)

- Use OS firewall (iptables) to protect our data. Another security layer.
- Need to build rules for RACs which run complex environments and are very sensitive, using several network interfaces → Extensive testing
- Need to change Quattor ncm-iptables component
- Also provides a log of “inappropriate” accesses

Firewall Automatic Generation




Note: Not auto propagated with every change

Main areas of activity

- Security
- **Oracle Database Virtualization**
- Joint Software Testing Program (mostly 11.2 database for the past 6 months)
- Monitoring
 - Web applications monitoring using EM / beacons
 - Database policies

- Need to integrate OracleVM within CERN ELFms
- Central Linux installation (PXE) + Quattor



- Manual installation:
OVM + packages \approx 1-3 h

- Whole automatic installation \approx
20 min

- Integration of OracleVM as host already done.
- Integrate the installation of the guest hosts using CERN ELFms
- Installation of the Oracle Database on guest hosts using Quattor

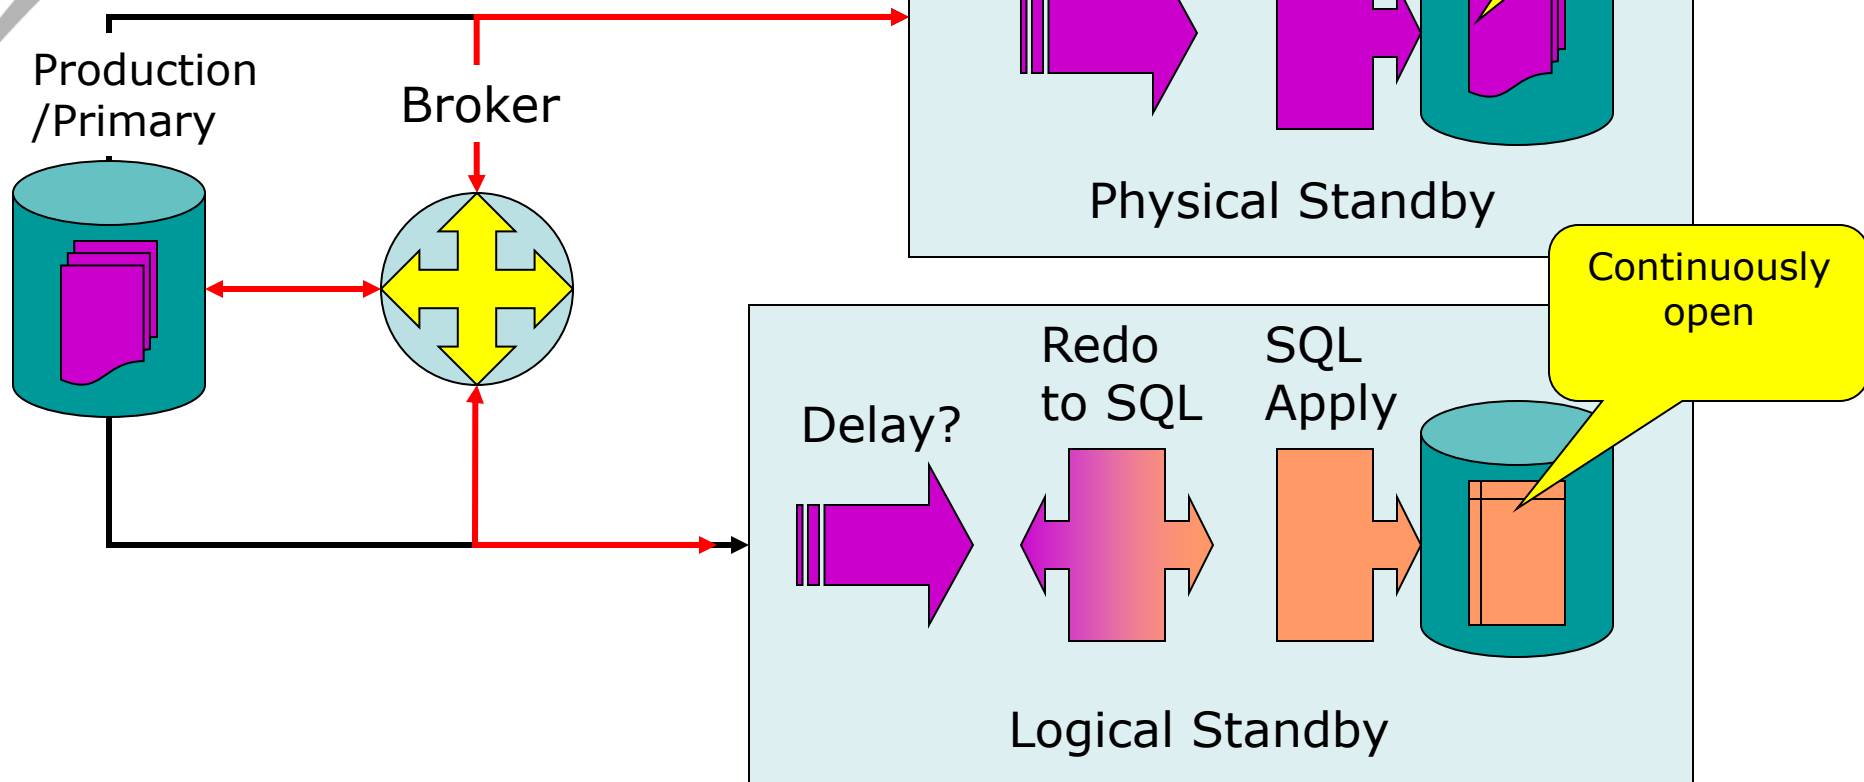
- Special thanks to OracleVM experts from Oracle:
 - Chris Barclay
 - Adam Hawley
 - Madhup Gulatifor their help

Main areas of activity

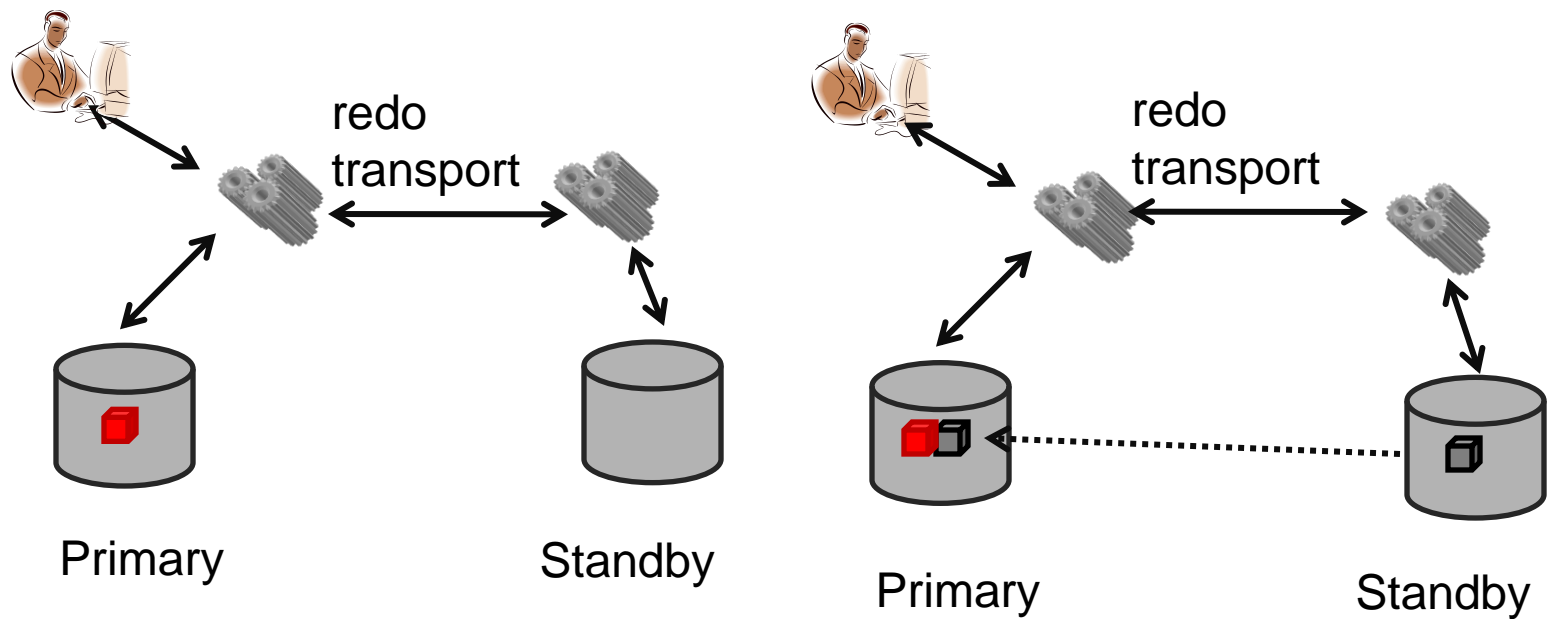
- Security
- Oracle Database Virtualization
- **Joint Software Testing Program (mostly 11.2 database for the past 6 months)**
- Monitoring
 - Web applications monitoring using EM / beacons
 - Database policies

- What is Data Guard?
 - It's a technology that establishes and maintains secondary "standby databases" as alternative/supplementary repositories to production "primary databases".
 - Physical standby (Redo Apply)
 - Replicates the exact contents of its primary database across the Oracle Net network layer.
 - Logical standby (SQL Apply)
 - Convert the redo generated at the Primary database into data and SQL and then re-apply those SQL transactions on the Logical standby.




Data Guard in a



- High Interest on this technology
 - Disaster recovery solution
 - Offload primary database:
 - Backups can be done on the physical standby
 - Physical Standby can be used for read-only queries (11g only) i.e. reporting, data warehouse apps, lengthy queries,...
 - Several use cases already identified for the “active dataguard” including CASTOR name server...
 - Hardware migration solution



User completes its query. It doesn't get:
ORA-01578: ORACLE data block
corrupted

- Access to binaries via Oracle Joint Software testing program
- High interest on new features
 - Already 2 applications designed based on 11g new features
- Features (full report provided to Oracle)
 - Active duplication ☺
 - Automatic block recovery
 - Primary  Standby ☺
 - Primary  Standby  in progress
 - DUPLICATE Without Connection to Target Database ☺
 - Configurable Real-Time Query Apply Lag Limit ☺
 - Direct NFS better performance vs Kernel mounts ☺

Main areas of activity

- Security
- Oracle Database Virtualization
- Joint Software Testing Program (mostly 11.2 database for the past 6 months)
- **Monitoring**
 - Web applications monitoring using EM / beacons
 - Database policies

OEM – Service tests and beacons

- OEM has been **migrated** to version **10.2.0.5** following the intensive pre-release testing and feedback (with some highlight on the benefits)
- The **virtualization console** has been intensively tested in the pre-release. This work was mentioned in the **Oracle press release**.
- **Enable** the use of **web transactions monitoring**.

OEM – Service tests and beacons

- **AIM:** We need to extend the monitoring provided by Lemon in aspects like the web transactions.
 - Decided to take profit of **OEM** Services
 - Provided feedback in the **CAB** (Customer Advisory Board)
 - Our requirements were included in the last release of OEM (10.2.0.5)
 - Satisfied with the tests we plan to deploy them widely.
-

- **Steps to follow to create an “EM service”:**
 - Creation of an “OEM service beacon” (remote piece of software that plays back recorded Web transactions or service tests for various protocols)
 - Creation of an “OEM system” for each CERN service to be monitored
 - Creation of an “OEM service” for each OEM system
 - Creation of a set of “OEM service tests” for each CERN service
-

- 1st Step: Creation of an OEM beacon

ORACLE Enterprise Manager 10g [Setup](#) [Preferences](#) [Help](#) [Logout](#)


Grid Control [Home](#) **[Targets](#)** [Deployments](#) [Alerts](#) [Compliance](#) [Jobs](#) [Reports](#)

[Hosts](#) | [Databases](#) | [Middleware](#) | [Web Applications](#) | [Services](#) | [Systems](#) | [Groups](#) | [Virtual Servers](#) | [All Targets](#)

Beacon: **J2EEPS beacon** Page Refreshed Sep 21, 2009 5:01:08 PM MEST [Refresh](#)

[Home](#) [Watch Lists](#) [Monitored Targets](#)

General



Status **Up** [Black Out](#)

Availability (%) **100**
(Last 24 Hours)

Host [ithp08.cern.ch](#)

Properties

Proxy Host

Proxy Port

Do Not Proxy For Domain

Alerts

Metric	Severity	Message	Alert Triggered	Last Value	Last Checked
No Alerts found.					

Host Alerts

Metric	Severity	Message	Alert Triggered	Last Value	Last Checked
No Alerts found.					

Related Links

[Past Changes](#)

[Home](#) [Watch Lists](#) [Monitored Targets](#)

OEM – Service tests and beacons

■ 2nd Step: Creation of an OEM system

ORACLE Enterprise Manager 10g [Setup](#) [Preferences](#) [Help](#) [Logout](#)
Grid Control [Home](#) **[Targets](#)** [Deployments](#) [Alerts](#) [Compliance](#) [Jobs](#) [Reports](#)
[Hosts](#) | [Databases](#) | [Middleware](#) | [Web Applications](#) | [Services](#) | [Systems](#) | [Groups](#) | [Virtual Servers](#) | [All Targets](#)

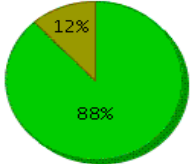
System: J2EEPS system Page Refreshed Sep 21, 2009 5:02:10 PM CEST [Refresh](#) [Launch Dashboard](#)

[Home](#) [Charts](#) [Administration](#) [Components](#) [Topology](#)

General

Owner **LMORENOL**
Problem Jobs **0**
Last 7 days.

[Status History](#)



■ Up(104)
■ Metric Collection Error(14)

Alerts

[Alert History](#)

Severity	Current	Last 24 hours
✗	0	0
⚠	0	0
Total	0	0

Blackouts

[Create](#)

Status	Submitted to the Group	Submitted to any Member
Scheduled	0	0
Active	0	0

Services

Name	Type	Status	Performance Alerts	Usage Alerts	Policy Violations
J2EEPS service	Generic Service		0 0	0 0	0 0 0

Policy Violations

Severity	Current	Last 24 Hours		Distinct Policies Violated
		Cleared	New	
✗	61	0	0	1
⚠	0	0	0	0
i	0	0	0	0
Total	61	0	0	1

[Policy Trend Overview](#)

Configuration Changes

Configuration changes detected for the last 7 days

Security Policy Violations

Severity	Current	Last 24 Hours		Distinct Policies Violated
		Cleared	New	
✗	61	0	0	1
⚠	0	0	0	0
i	0	0	0	0
Total	61	0	0	1

[Security At a Glance](#)

Critical Patch Advisories for Oracle Homes

OEM – Service tests and beacons

■ 3rd Step: Creation of an OEM service

ORACLE Enterprise Manager 10g Setup Preferences Help Logout
Grid Control Home Targets Deployments Alerts Compliance Jobs Reports

Hosts | Databases | Middleware | Web Applications | Services | Systems | Groups | Virtual Servers | All Targets

Generic Service: J2EEPS service

Home Charts **Test Performance** System Topology Monitoring Configuration

Page Refreshed Sep 21, 2009 5:04:32 PM CEST

General

Status **Up** Black Out

Up Since **Sep 21, 2009 4:33:44 PM**

Last Calculated **Sep 21, 2009 4:58:22 PM**

Availability (%) **70.2**
(Last 24 Hours)

Performance ✓

Usage ✓

Actual Service Level (%) **74.6968**
(Last 24 Hours)

Expected Service Level (%) **85.0000**

Key Component Summary

System **J2EEPS system** [\(Topology\)](#)

Status ⬆ **14** ⬆ **104**

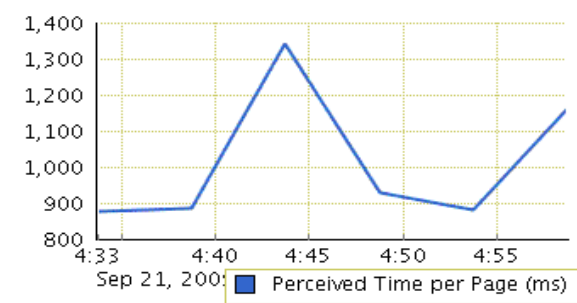
Alerts ✗ **0** ⚠ **0**

All Service Alerts

View All Service Alerts

Target Name	Target Type	Severity	Alert Triggered	Message
(No alerts)				

Performance



■ Perceived Time per Page (ms)

Key Test Summary

Test	Test Type	Status	Alerts
containers hosting using apps	Web Transaction	⬆	0 0

OEM – Service tests and beacons

- 4th Step: Creation of OEM service tests to monitor a CERN service

ORACLE Enterprise Manager 10g
Grid Control

Home **Targets** Deployments Alerts Compliance Jobs Reports

Hosts | Databases | Middleware | Web Applications | Services | Systems | Groups | Virtual Servers | All Targets

[Generic Service: J2EEPS service](#) >

Service Tests and Beacons

Specify the service tests and beacons to monitor the availability and performance of your service. OK

Service Tests

All enabled service tests are run by the beacons shown below.

Verify Service Test Enable Disable Remove View Edit | Test Type CalDAV Add

Select	Service Test	Test Type	Enabled	Status	Key Service Test
<input checked="" type="radio"/>	Shibboleth j2eeeps	Web Transaction	Yes	⬆	<input type="checkbox"/>
<input type="radio"/>	containers hosting using apps	Web Transaction	Yes	⬆	<input checked="" type="checkbox"/>
<input type="radio"/>	httpd at j2eeeps.cern.ch - service proxy	Web Transaction	Yes	⬆	<input type="checkbox"/>
<input type="radio"/>	jmxmc on j2eeeps01 - jmx monitor of the user tomcats	Web Transaction	Yes	⬆	<input type="checkbox"/>

Beacons

Specify the beacons that will execute the service tests.

Remove Edit | Add Create

Select	Name	Version	Operating System	Status	Key Beacon
<input checked="" type="radio"/>	J2EEPS beacon	10.2.0.5.0	Linux	⬆	<input checked="" type="checkbox"/>

Tip

This page allows you to add, remove, and edit service tests and beacons for your service, but availability is defined on a separate page.

To change the availability definition for your service, click "Availability Definition" under Related Links.

- Performed EM service tests – **Test types:**
 - Web Transactions:
 - ✓ SSO: CERN authentication
 - ✓ Form authentication: AIS login
 - ✓ Realm: Tomcat Realm
 - JDBC SQL Timing
 - LDAP
-

OEM – Service tests and beacons

■ Example of service test – Type Web Transaction

ORACLE® Enterprise Manager 10g [Setup](#) [Preferences](#) [Help](#) [Logout](#)
Grid Control [Home](#) **[Targets](#)** [Deployments](#) [Alerts](#) [Compliance](#) [Jobs](#) [Reports](#)
[Hosts](#) | [Databases](#) | [Middleware](#) | [Web Applications](#) | [Services](#) | [Systems](#) | [Groups](#) | [Virtual Servers](#) | [All Targets](#)
[Generic Service: cernauthen_service](#) > [Service Tests and Beacons](#) > [Edit Service Test: cernauth_serviceTest](#) >
Edit Step: Find application - J2EE Public Service [Cancel](#) [Continue](#)

▼ Input Parameters

Nonsensitive Parameters	
Name	Value
V_CTL00\$CONTENTPLACEHOLDER1	Imorenol

Example: Name = USRJ, Value = jsmith. During playback, "[USRJ]" in the URL or POST Data fields will be substituted with "jsmith".

Sensitive Parameters	
Name	Value
V_CTL00\$CONTENTPLACEHOLDER1

Values are masked. Reentry of values for Sensitive Parameters is required when the URL is changed or a new sensitive parameter reference is added.

▼ Request

* URL:

HTTP Method:

POST Data:

HTTP Headers:

Recorded Query Parameter:

This is used by version 10.2.0.3 or earlier beacons.


Request Mode:

▼ Actions

DHTML Script			
type	<input type="button" value="newValue"/>	newValue="[V_CTL00\$CONTENTPLACEHOLDER1\$XTFORMSLOGIN]" name="ctl00\$C	<input type="button" value="+"/> <input type="button" value="X"/>
type	<input type="button" value="newValue"/>	newValue="[V_CTL00\$CONTENTPLACEHOLDER1\$XTFORMSPASSWORD]" name="c	<input type="button" value="+"/> <input type="button" value="X"/>
mouseDown	<input type="button" value="name"/>	name="ctl00\$ContentPlaceHolder1\$btnFormsLogin"	<input type="button" value="+"/> <input type="button" value="X"/>
waitForPageToLoad	<input type="button" value=""/>		<input type="button" value="+"/> <input type="button" value="X"/>
waitForPageToLoad	<input type="button" value=""/>		<input type="button" value="+"/> <input type="button" value="X"/>

This section is for information purposes only. They are only used in Browser Simulation Mode.

■ Example of service test – Type LDAP

ORACLE Enterprise Manager 10g 
Grid Control

[Home](#) **[Targets](#)** [Deployments](#) [Alerts](#) [Compliance](#) [Jobs](#) [Reports](#)

[Hosts](#) | [Databases](#) | [Middleware](#) | [Web Applications](#) | [Services](#) | [Systems](#) | [Groups](#) | [Virtual Servers](#) | [All Targets](#)

[Generic Service: J2EEPS.service](#) > [Service Tests and Beacons](#) >

Edit Service Test: central database cluster

Test Type **JDBC SQL Timing**

Name **central database cluster** Description

* Collection Frequency (minutes)

Test Parameters

▼ Database Connection

* Connection String

* Class String

▼ Credentials

* Username

* Password

Role

▼ Query

* Statement

■ Example of service test – Type JDBC SQL Timing

ORACLE Enterprise Manager 10g [Grid Control](#) [Home](#) [Targets](#) [Deployments](#) [Alerts](#) [Compliance](#) [Jobs](#) [Reports](#) [Setup](#) [Preferences](#) [Help](#) [Logout](#)

[Hosts](#) | [Databases](#) | [Middleware](#) | [Web Applications](#) | [Services](#) | [Systems](#) | [Groups](#) | [Virtual Servers](#) | [All Targets](#)

[Generic Service: J2EEPS service](#) > [Service Tests and Beacons](#) >

Edit Service Test: ldap connection to CERNs Active Directory

[Cancel](#) [Verify Service Test](#) [OK](#)

Test Type **LDAP**

Name **ldap connection to CERNs Active Directory** Description

* Collection Frequency (minutes)

Test Parameters

* LDAP Host Address

* LDAP Port

* LDAP User Name

* LDAP Password

Total Number of Retries

Retry Interval (minutes)

* LDAP Search Filter

* LDAP Search Base

* LDAP Compare Attribute Name

* LDAP Compare Attribute Value

* LDAP Timeout (seconds)

Request Type

Authentication Mode

Main areas of activity

- Security
- Oracle Database Virtualization
- Joint Software Testing Program (mostly 11.2 database for the past 6 months)
- **Monitoring**
 - Web applications monitoring using EM / beacons
 - **Database policies**



Increasing productivity at CERN with Enterprise Manager Grid Control

Manuel Guijarro

Lucia Moreno

Chris Lambert

CERN

Oracle OpenWorld – October 2009

Manuel.Guijarro@cern.ch



Snapshots of OOW presentation



1. Introduction to OOW: Overview of the OOW system, its components, and its role in the CERN environment.

2. Agenda: Overview of the presentation topics, including Introduction to OOW, OOW Architecture, OOW Usage, OOW Performance, OOW Security, OOW Integration, and OOW Future.

3. Introduction to OOW: Detailed overview of the OOW system, its components, and its role in the CERN environment.

4. OOW Architecture: Diagram illustrating the OOW architecture, showing the interaction between the OOW system and the CERN environment.

5. OOW gets ready ...: Image showing the OOW system interface, highlighting the 'Ready' status.

6. OOW Computing Challenge: Image showing the OOW system interface, highlighting the 'Computing Challenge' section.

7. OOW Applications: Image showing the OOW system interface, highlighting the 'Applications' section.

8. Our OOW and OOW Architecture: Diagram illustrating the OOW architecture, showing the interaction between the OOW system and the CERN environment.

9. Our Challenge: Text slide discussing the challenges faced by the OOW system, including scalability, performance, and security.

10. OOW Control Solution: Diagram illustrating the OOW control solution, showing the interaction between the OOW system and the CERN environment.

11. User OOW Policies Usage: Text slide discussing the usage of OOW policies, including the role of the OOW system in enforcing policies.

12. New Policy - Using User OOW Policies: Text slide discussing the usage of OOW policies, including the role of the OOW system in enforcing policies.

13. New Policy - Using User OOW Policies: Text slide discussing the usage of OOW policies, including the role of the OOW system in enforcing policies.

14. New Policy - Using User OOW Policies: Text slide discussing the usage of OOW policies, including the role of the OOW system in enforcing policies.

15. New Policy - Using User OOW Policies: Text slide discussing the usage of OOW policies, including the role of the OOW system in enforcing policies.

16. New Policy - Using User OOW Policies: Text slide discussing the usage of OOW policies, including the role of the OOW system in enforcing policies.

17. New Policy - Using User OOW Policies: Text slide discussing the usage of OOW policies, including the role of the OOW system in enforcing policies.

18. New Policy - Using User OOW Policies: Text slide discussing the usage of OOW policies, including the role of the OOW system in enforcing policies.

19. New Policy - Using User OOW Policies: Text slide discussing the usage of OOW policies, including the role of the OOW system in enforcing policies.

20. New Policy - Using User OOW Policies: Text slide discussing the usage of OOW policies, including the role of the OOW system in enforcing policies.

21. New Policy - Using User OOW Policies: Text slide discussing the usage of OOW policies, including the role of the OOW system in enforcing policies.

22. New Policy - Using User OOW Policies: Text slide discussing the usage of OOW policies, including the role of the OOW system in enforcing policies.

23. New Policy - Using User OOW Policies: Text slide discussing the usage of OOW policies, including the role of the OOW system in enforcing policies.

24. New Policy - Using User OOW Policies: Text slide discussing the usage of OOW policies, including the role of the OOW system in enforcing policies.

25. New Policy - Using User OOW Policies: Text slide discussing the usage of OOW policies, including the role of the OOW system in enforcing policies.

❖❖ *Our Challenge*

- **Minimize cost** of monitoring growing architecture
- Provide timely, **standardized access** to meaningful information
- Enable **pro-active management & problem avoidance**
- Identify and **remove** configuration **exceptions**

❖❖ *Conclusion: Increased productivity*



- Furthermore... Grid control **reduces manpower** needs by:
 - Providing centralized access to meaningful information
 - Enforcing compliance with our standards
 - Decreasing time consumed by daily operations
 - Reducing downtime by pro-active monitoring
 - Assisting DBAs in their tuning and performance improvement tasks
 - ...and all with little additional effort even for a constantly expanding IT infrastructure

- Beta project around Oracle Weblogic
- Database virtualisation deployment
- 11.2 database deployment start

- 11.2 testing result
- OOW conference
- UKOUG conference
- Phone conferences on Oracle VM and EM
- Swiss Oracle User Group on Oracle 11
- CERN Computer Newsletter article
- Oracle Press release about OEM management to OVM server

Supporting Quotes

- "Two of the main roadblocks of customer adoption of virtualization technologies has been the high cost and complexity of managing virtual environments and a lack of tools to provide an application perspective," said Richard Sarwal, Oracle senior vice president Product Development. "With added management capabilities for Oracle VM, Oracle Enterprise Manager has taken a major step forward in helping our customers drive down the cost of managing applications in virtual environments."
- "CERN has been searching for a comprehensive virtualization management solution for some time. We recently evaluated Oracle Enterprise Manager's Oracle VM Management Pack and we were impressed with its virtualization dashboard. It simplifies administrators' jobs by providing an integrated solution for monitoring and administration of virtual servers as well as tracking configuration relationships. We especially liked the ability to provision servers within minutes using templates and live-migrate VMs. These management capabilities can greatly simplify virtual infrastructure management, reduce IT operations overhead, provide cost savings, and reduce risk," said Anton Topurov, Database Fellow, openlab, CERN.
- "With the latest additions for Oracle VM deployment and management, Oracle Enterprise Manager will enable us to manage both our existing physical environments and new virtualized infrastructure from a single console. The virtualized Oracle grid infrastructure will allow us to maximize the use of server resources while maintaining high levels of system availability and performance," said April Sims, Database Administrator, Southern Utah University.

Technical brief

openlab tests Oracle VM for database service virtualization

Why database virtualization
There is a combination of factors that leads to increased interest for virtualization of database servers: number of servers, power usage, cost, new platforms with multi/many cores, high availability and ease of management. Database servers often do not use the full CPU/server capacity. This can be for reasons such as separate test-and-development servers or because spare capacity has to be kept available in case of increased load or for high availability.

In addition, several database services can often not be merged onto a single database or physical "machine" as they need to be on different versions of the database software or operating system. We also need to "protect" each of the database instances from the others so that memory can be guaranteed for a given instance.

Having a smaller number of servers will obviously reduce the cost of hardware, maintenance and electrical power. The power usage is especially important due to the constraints on "critical power" in computer centres as well as the impact on a cooling system.




Fig. 1. Fighting underutilization by combining several clustered databases on two servers.




Fig. 2. Architecture: Oracle VM Manager and Oracle VM Server are installed on separate servers with communication provided by the management server and Oracle VM Agent.

thus providing scalability and high availability. High availability is the main driving force behind implementing RAC in less workload-intensive infrastructure databases, which also leads to the idea of running RAC in a virtualized environment, thus utilizing more of the hardware without compromising the high availability features (see figure 1). The concept of RAC on VM was proved to be successful by openlab summer student Maria Lertner in 2007. She created a cluster database inside a single physical host with Red Hat Enterprise Linux as a base operating system. The set-up was quite stable and the performance loss

due to the virtualization layer was at an acceptable level.

In November 2007, Oracle announced its virtualization solution: Oracle VM, which consists of three parts: Oracle VM Server, Oracle VM Manager and Oracle VM Agent (see figure 2).

Based on Red Hat Enterprise Linux (RHEL) and open-source Xen, Oracle VM Server is the base virtualization software to be installed as a bare metal server. The Oracle VM Agent is the application, installed together with the Oracle VM Server and is responsible for communication with Oracle VM Manager Web-based GUI.

Oracle then certified this single instance (November 2007) and RAC (September 2008) databases in the Oracle VM based

virtualized environment, which was the step we were waiting for.

It was decided to continue with our tests and compare RAC performance on Oracle VM versus the native Xen set-up we tested before. Together with openlab summer student Andrii Dumitru in 2008, we created two set-ups on identical hardware, using Oracle VM and pure Xen, included in RHEL 5. The stability of both set-ups proved to be good and the performance of the Oracle VM set-up was to 20% better than on pure Xen, depending on the workload. Tests of live VM migration with an active database showed only a few seconds downtime with no session-state loss. This is an outstanding addition to the high availability features initially provided by RAC.

Following these successful tests, we are now working closely with IT-IO to integrate this Oracle VM solution into the CERN fabric management system – E-Data.

In summary, these Oracle VM features, their ease of use and the official certification have directed our choice towards Oracle VM as a virtualization solution for our virtualized cluster databases for the future.

Useful links

CERN openlab: <http://cern.ch/openlab> (with openlab summer student reports in Technical Documents section)
Oracle VM: www.oracle.com/technology/oracle-vm

Anton Topurov and Eric Gracich, IT-DES