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

openlab Major Review Meeting 2008

15 September 2008

Andrei Dumitru, Anton Topurov,
Chris Lambert, Lucia Moreno Lopez,
Daniel Lenkes, Eric Grancher

LHC startup and some databases

- IT manages databases for Accelerator control, 3 Real Application Cluster databases.
- Settings Database, settings to drive the accelerators
- Measurement database, relatively fast read-back data from accelerators, quite intensive, ~790GB redo on September 10th.
- Logging database, relatively slow read-back data from accelerators (decision support for post-mortem, off-line analysis), 5TB on September 17th 2008, +10TB per year

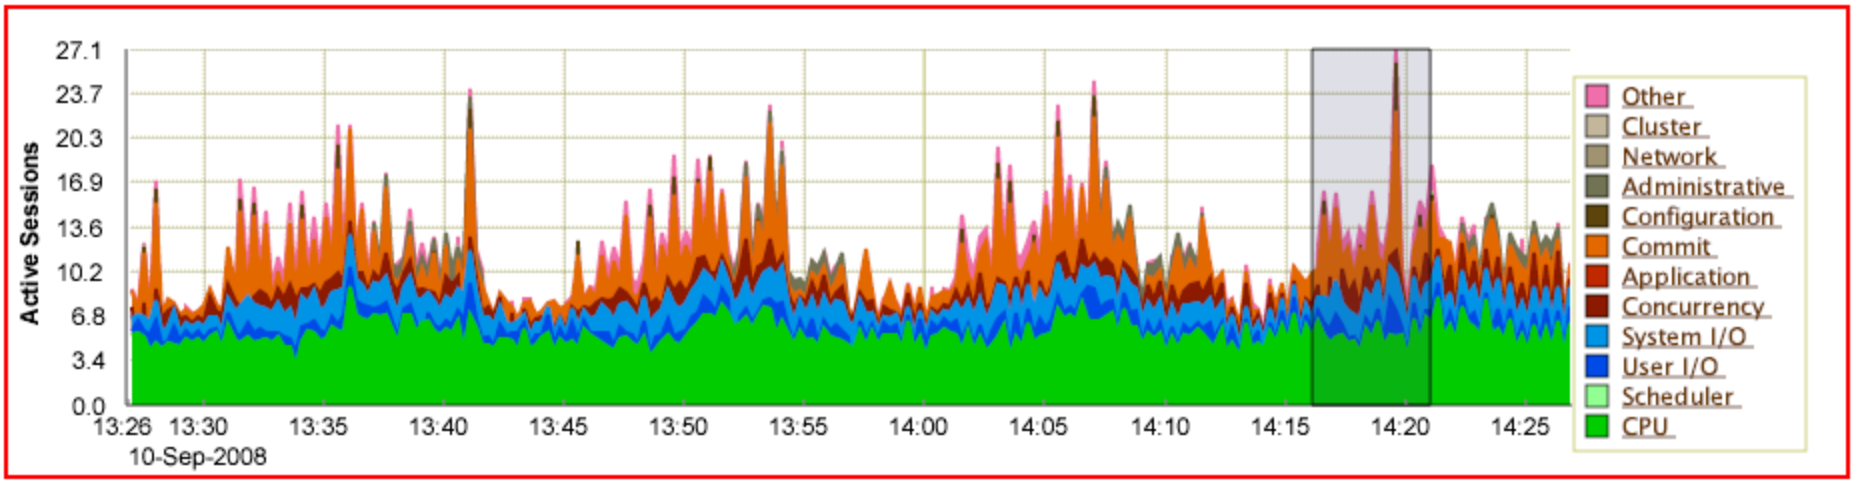

PVSS Real Workload

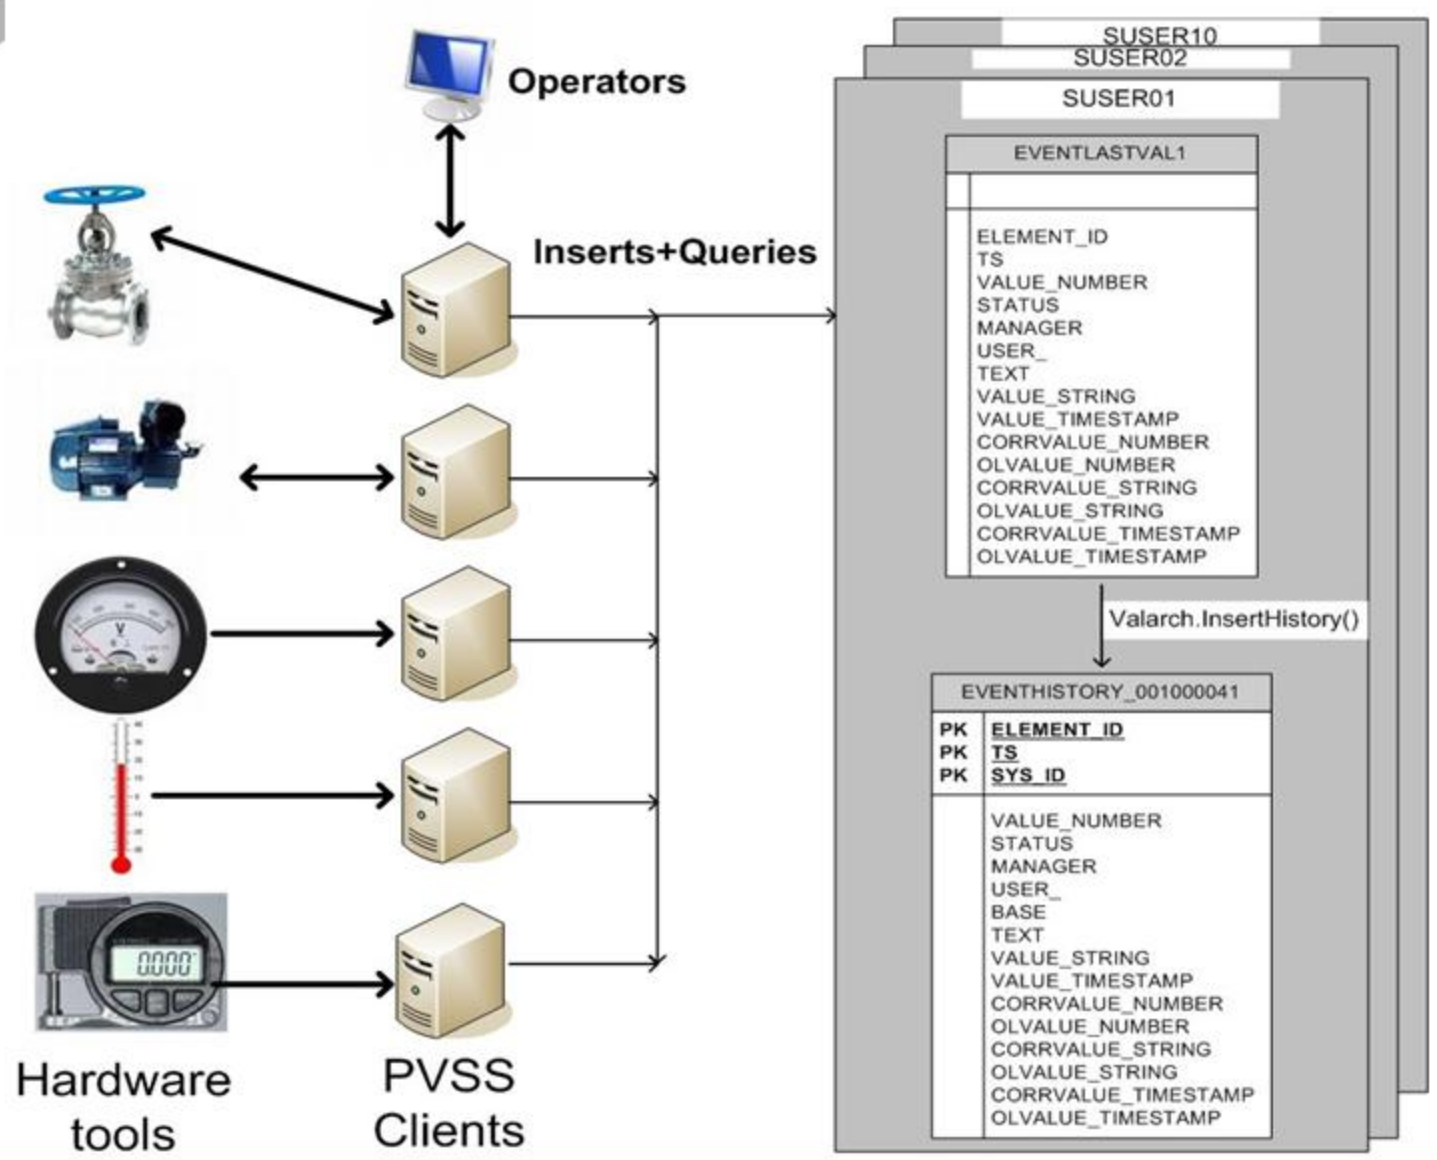

PVSS Workload Simulator

- Created within Swingbench Framework

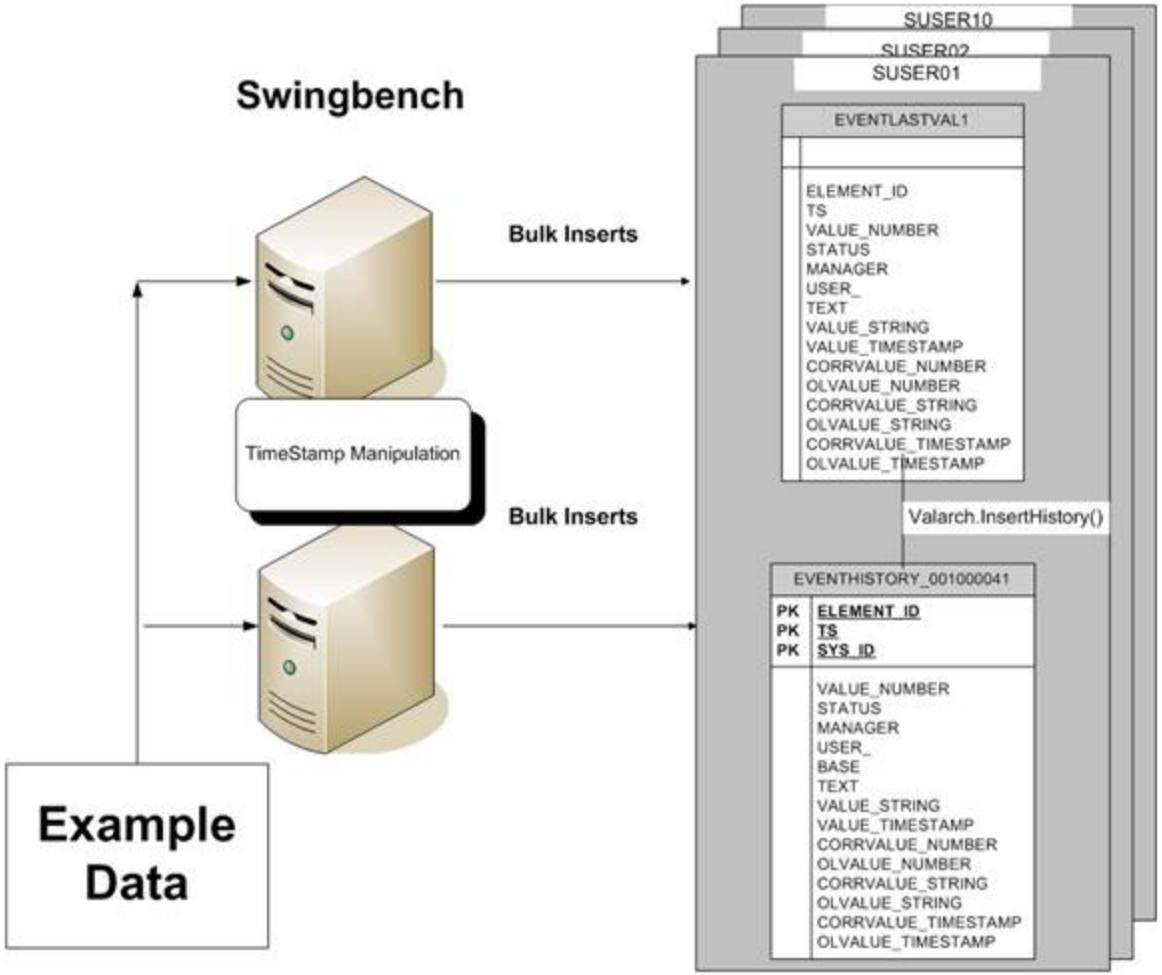

PVSS Workload Generator

- 1000 changes/s as a bulk insert by each session
- Single Swingbench instance runs 15 sessions towards a dedicated schema
- 10 swingbench instances in total needed to generate top load of 150 000 changes/s

Swingbench in action

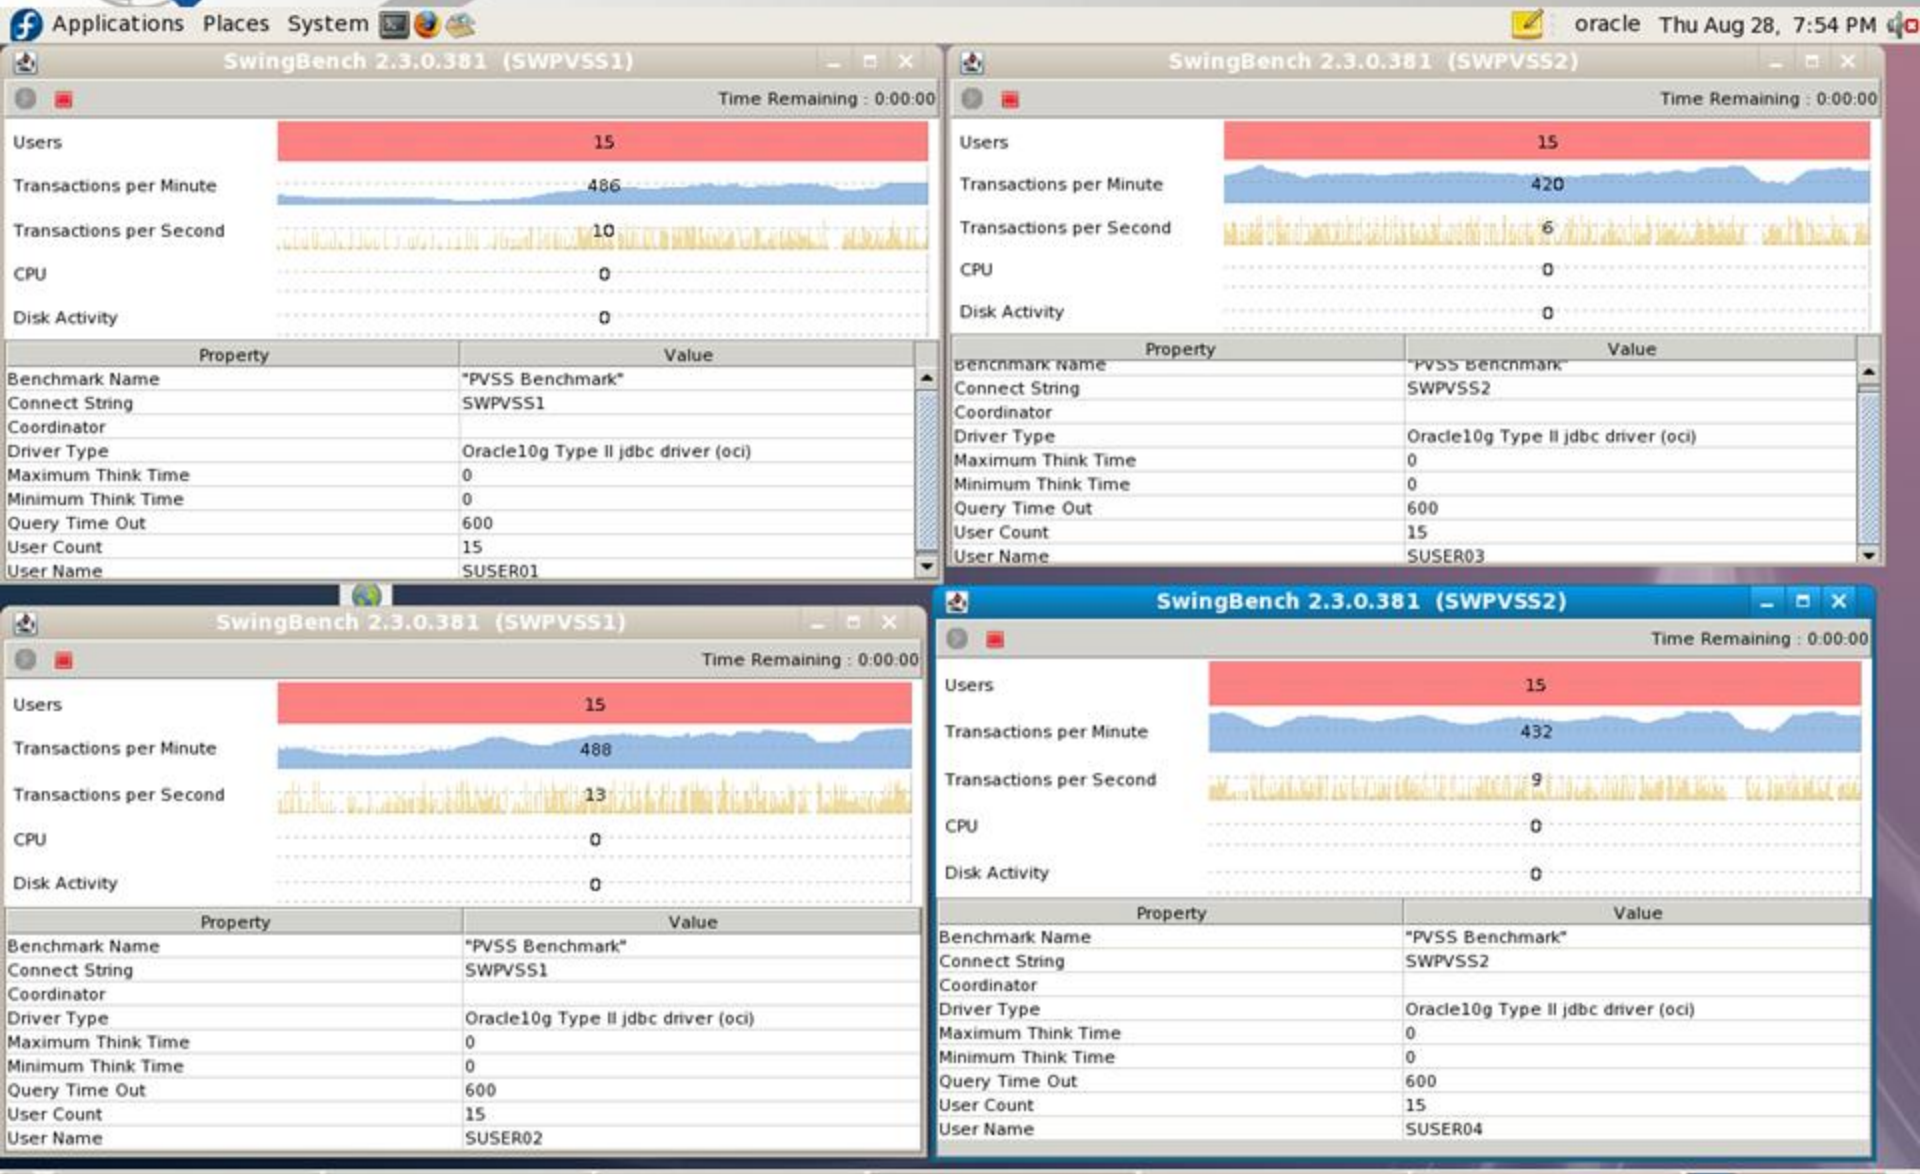

Oracle Database Virtualization

- Continuation of last year summer student project
- Main focus:
  - Test of Oracle RAC on Oracle VM and OEL5
  - Test of Oracle RAC on OEL5 and pure XEN
  - Test of Hardware Virtualization and Paravirtualization
- Work being done by Andrei Dumitru (openlab Summer Student)
- First results show better performance and ease of use for Oracle VM solution

[RAC] Oracle RAC virtualization

Setup 2:

- Same as setup 1
- NFS export from real machine

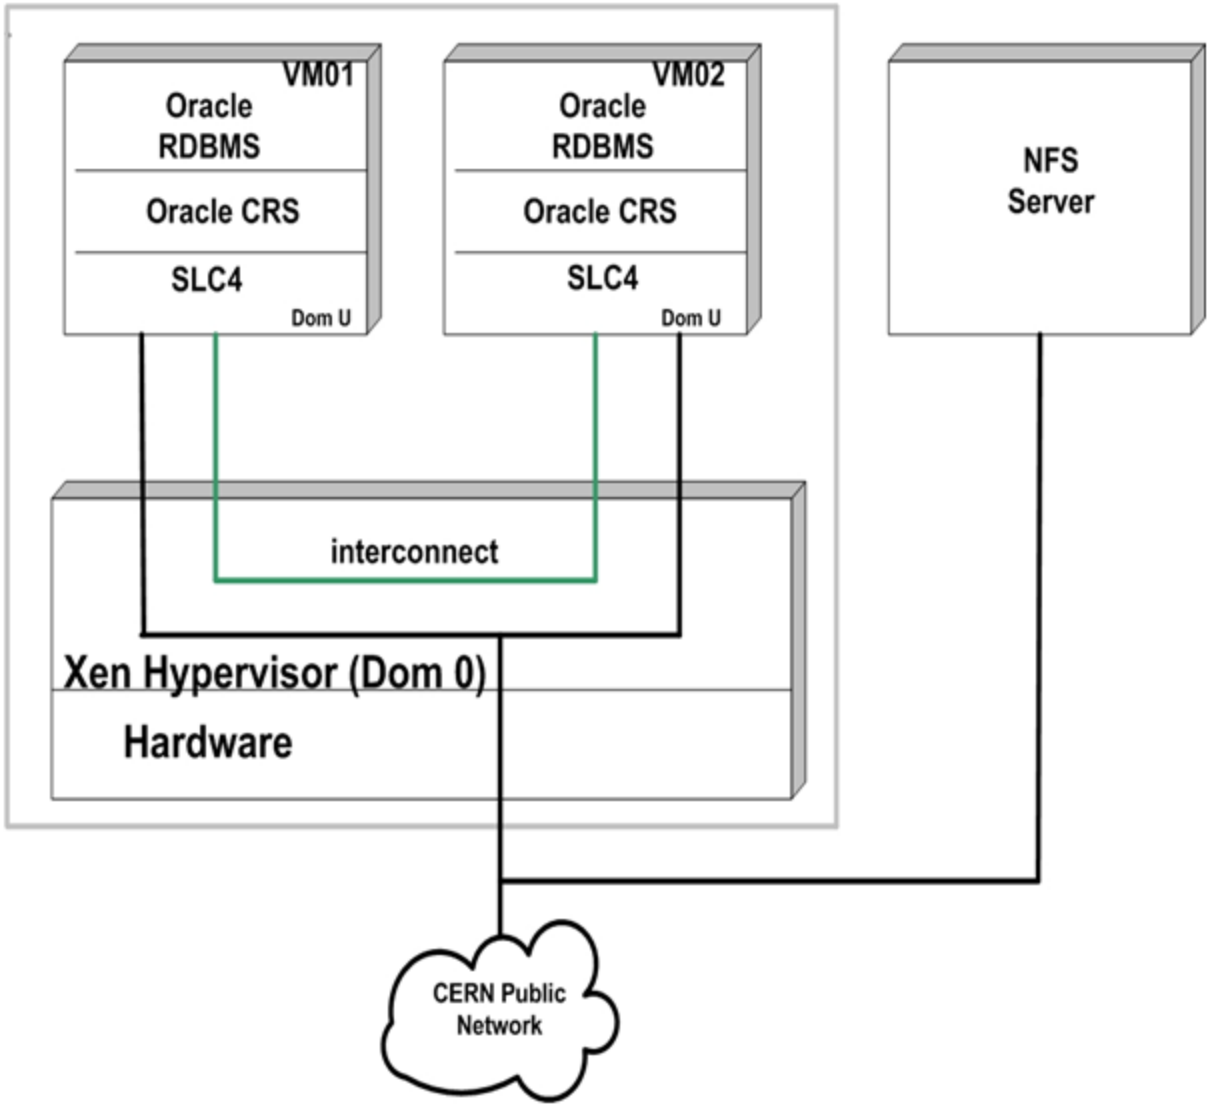

Using EM10g to improve Compliance for Oracle Security

Christopher Lambert

Database and Engineering Services Group

chris.lambert@cern.ch

Changing IT Infrastructure

Legacy Environment	New Environment
Few multiprocessor hosts	Many hosts
Single vendor hardware	Commodity hardware
Solaris	RedHat Linux
Single instance databases	Oracle RAC
Directly attached storage	NetApp NAS storage
Unique configurations	Standard configurations
Individually managed	Manage by exception

Our Grid Control Environment

Agent versions:
- 10.1.0.4, 10.2.0.2
- 10.2.0.3, 10.2.0.4

Solaris and Linux (32&64 bit)

Secure agent upload

Hardware load balancer

OMS version 10.2.0.4

Linux RHEL 3

2 CPU (3.2GHz) - 4Gb RAM

2-Node RAC (load balanced)

RDBMS 10.2.0.3

RedHat Linux 4 (64-bit)

NetApp NAS storage

Total monitored targets: ~ 1500
192 instances (53 RAC dbs)
223 hosts

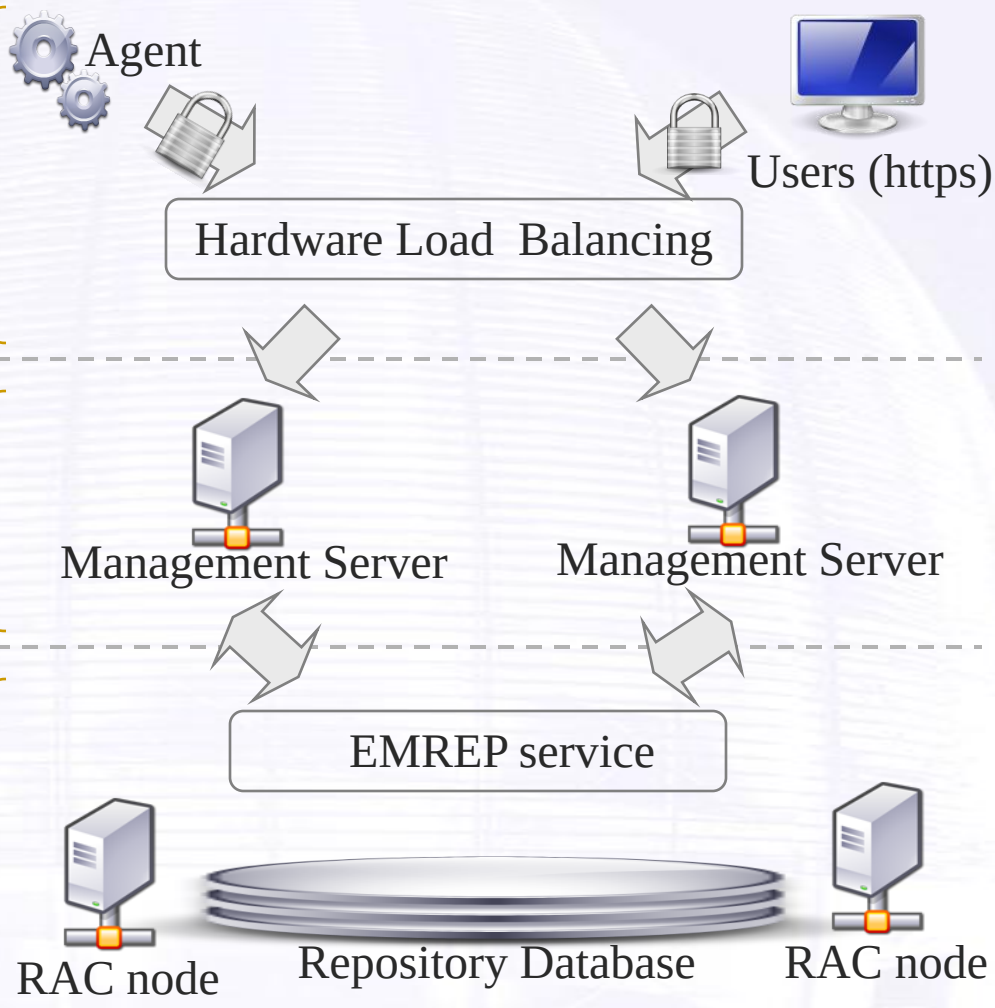

Grid Control Objectives

Minimize cost of monitoring growing architecture

Timely, standardized access to meaningful information

Enable pro-active management & problem avoidance

Identify and remove configuration exceptions

4 Fundamental Questions

Are we monitoring everything we should?

What is our current status?

Where should we focus our efforts?

How can we make this sustainable?

Grid Control Solution

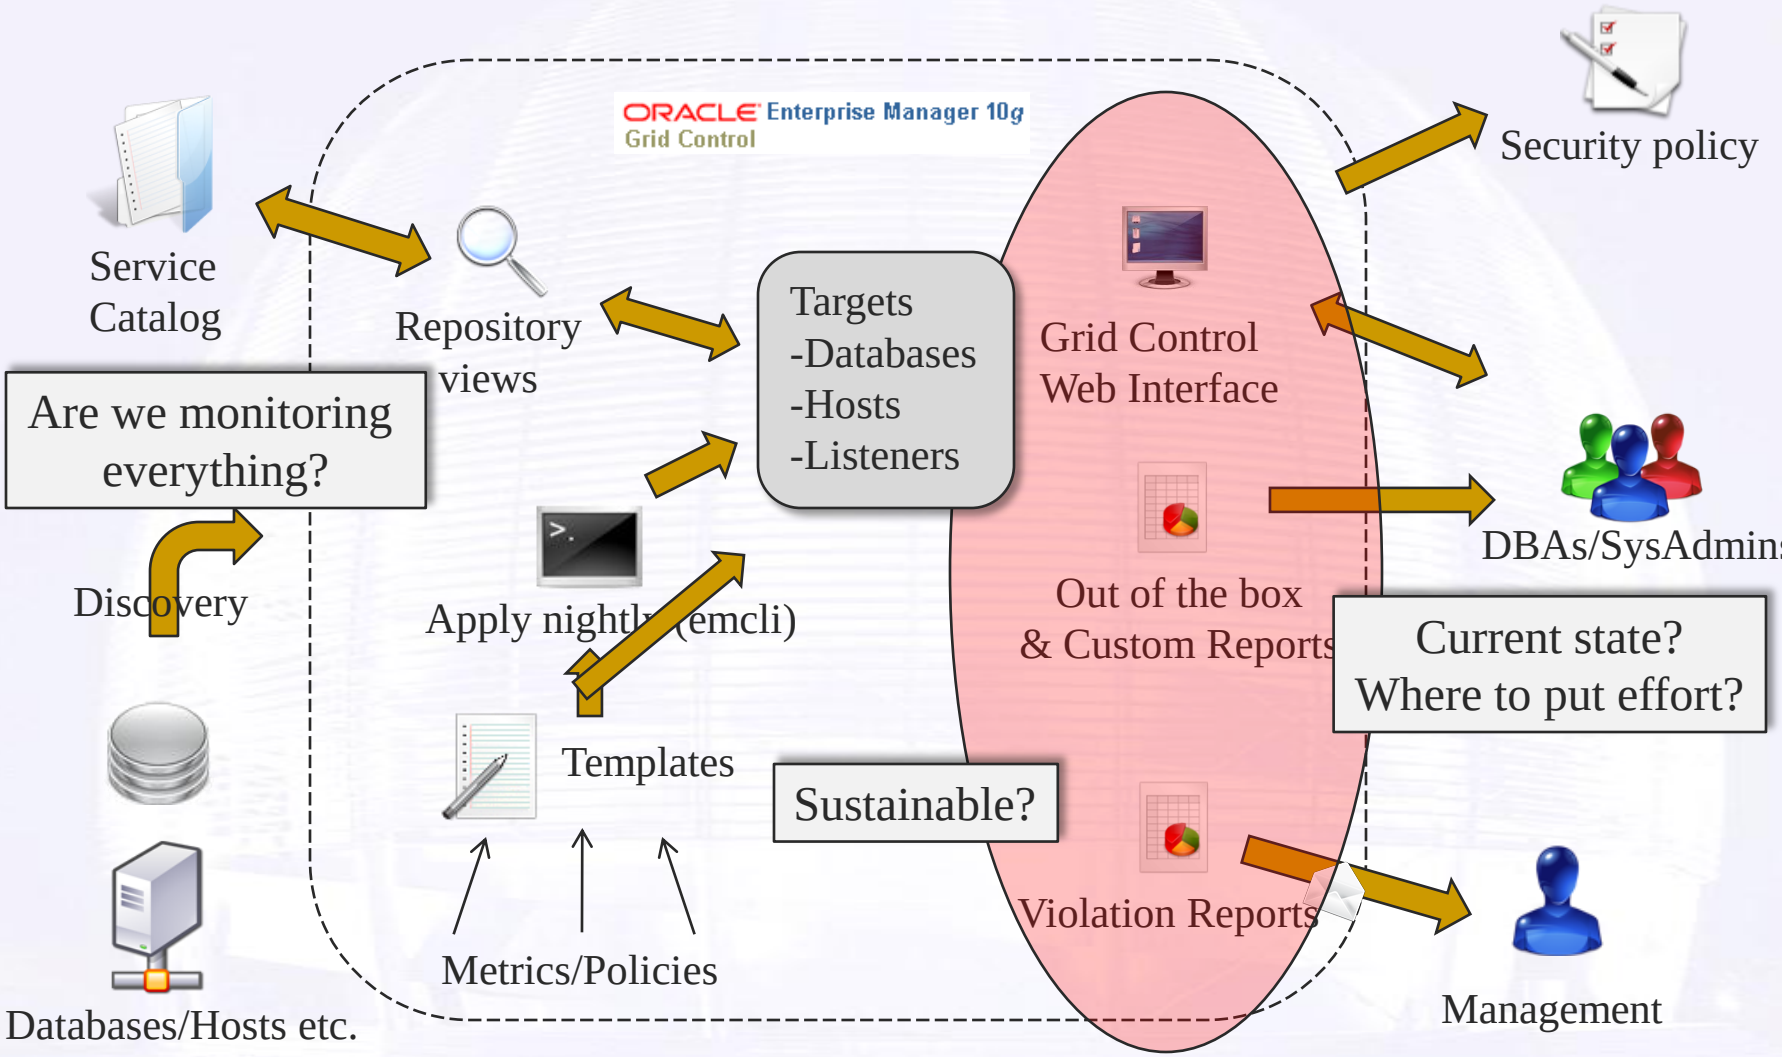

Continual Service Improvement

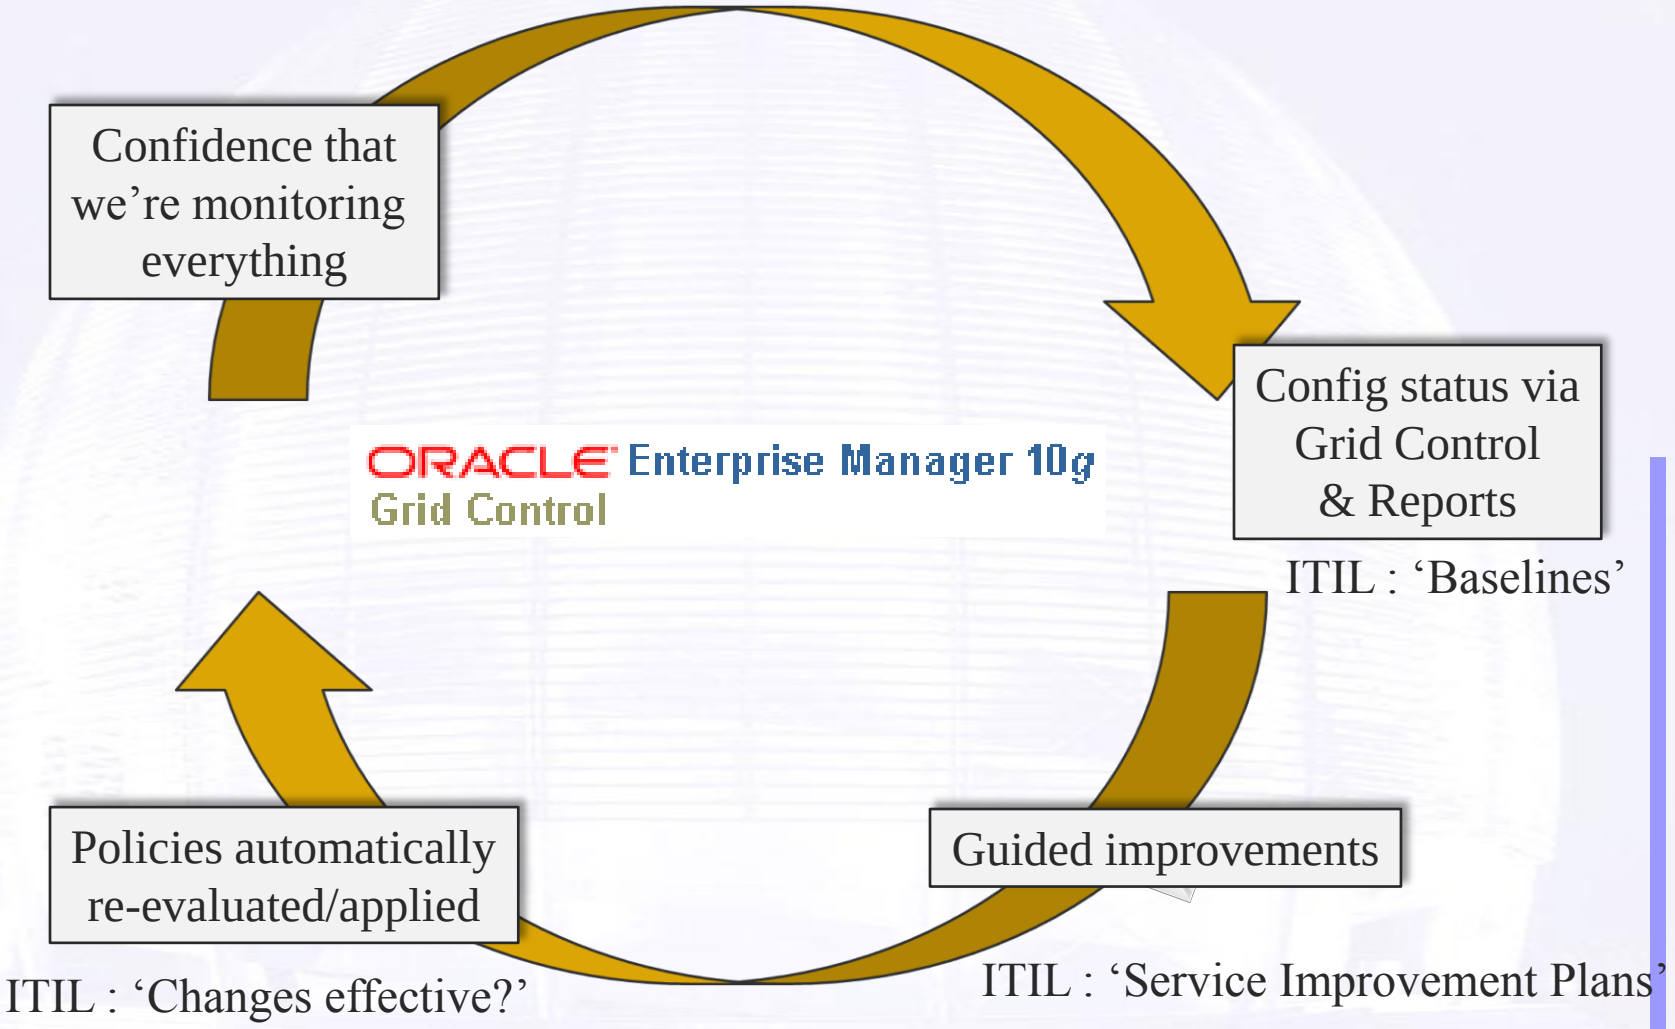

Conclusion

Out-of-the-box + extensible policies
Phasing out home grown solutions

Enables a standard management approach
20+ administrators using the same tool

Scales with our environment
800 -> 1500 targets with little additional effort

Essential tool for enabling compliance and security

Enabling us in our ‘best practice’ efforts