

Database Competence Centre

openlab Major Review Meeting 2011

13th October 2011

Zbigniew Baranowski
Andrei Dumitru
Carlos Garcia Fernandez
Luigi Gallerani
Mariusz Piorkowski
Anton Topurov



- Oracle VM at CERN
 - Deployment status update
 - Intel Single Root I/O Virtualization (SR-IOV) 
- Enterprise Manager
 - Enterprise Manager 12 beta testing overview
- Replication Technologies update
 - 11g deployment plans
 - Validations with production data
- Outreach

OracleVM at CERN




- **3** pools of OracleVM Servers running **53** VMs fully hardware virtualized
- Oracle clusters (**RACs**) of **11g** installed in virtual machines
- CERN CC integration Quattor, Lemon, SLS LANDB, syscontrol
- **OracleVM Manager** with **redundant configuration**, **OVM repository on our production DB**

- **5** Pools of OracleVM Servers running **102** VMs using **HW, HW+PVdriver, PV**
- OVM 2.2: 2 Pool Production, 1 dev, 2 test
- **HighAvailability** stress test **successful**, auto recovery of the VM after multiple server power cut
- OVM manager on 2 virtual machines with DNS load balancer!
- **EM12 Cloud** test beta program
- **OVM 3.0** Test in SafeHost with Intel 10GB Ethernet card

- What is Intel Single Root I/O Virtualization (SR-IOV)?

SR-IOV is a PCI device virtualization technology that allows a single PCI device to appear as multiple PCI devices on the physical PCI bus



- **The aim of the test, is to measure the benefit of using SR-IOV with CERN database workloads on ORACLE VM 3.0**
 - Assigning one or more virtual functions to a virtual machine allows the virtual machine to directly exploit the hardware without any mediation by the hypervisor.
 - This means less CPU usage, reduced latency and increased bandwidth!

- New Intel SR-IOV 10GB Ethernet Cards installed on new Dell machines
- Oracle VM 3.0 server installed last week @SafeHost
- Deep analysis on 8 different virtual OS and solution of NFS I/O problem
- SR-IOV on OVM 3.0 test completed for December

Oracle Enterprise Manager 12 Beta @CERN



CERN
openlab



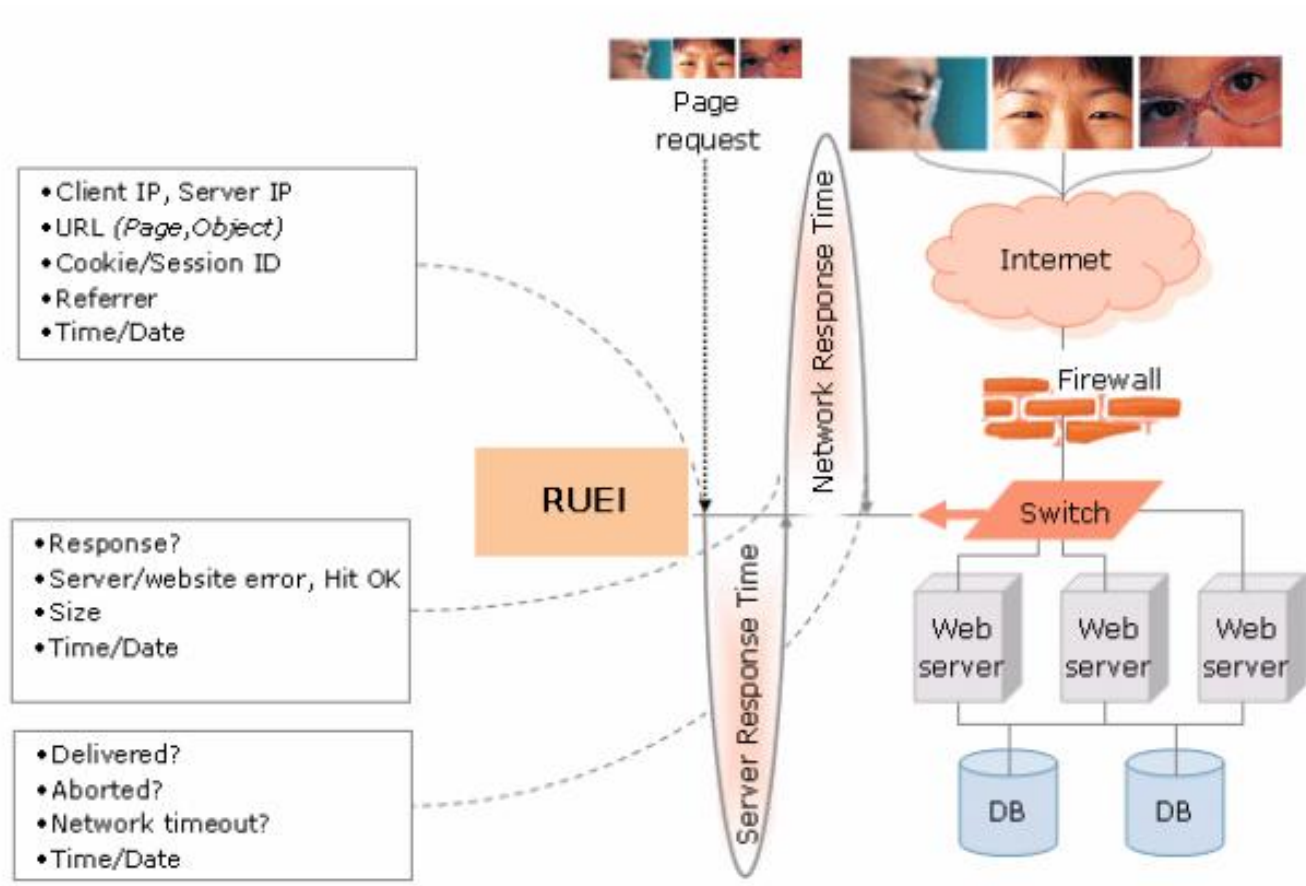
EM 12 new features interesting for CERN

- Manageable upgrade
- SSH Keys and LDAP based authentication
- Admin Groups
- AWR/ASH/Emergency ADDM
- Active Reports
- EM Cloud Management instead of OVM Manager
- Consolidation planner
- Bare metal provisioning


- Testing in 5 tracks:
 - Install
 - Upgrade (two system upgrade)
 - Database Management and framework
 - Application Replay
 - Cloud Management

Application Replay - product overview

- What is AR?
 - realistic testing of any part of the application stack from application server down to disk
 - capture a workload on the production system from network traffic
 - replay it on a test system with the exact timing, concurrency, and transaction characteristics of the original workload
 - Based on **Oracle Real User Experience Insight (RUEI)**
- What is RUEI?
 - end-to-end monitoring based on network protocol analysis
 - offers wide variety of online applications performance reporting
 - deployment requires network appliance or special configuration on switches (copy port)



Oracle EM12 Cloud Management



- Really good product with great GUI
- Lots of functionalities appreciated by CERN
- Modularity and 2-system upgrade are important
- Looking forward to get the 1st patchset and go to production
- Great collaboration with Oracle teams

- Andrew Bulloch:
“CERN has a very good reputation inside Oracle development for delivering high quality work and feedback and I'm sure in the EM Beta test their reputation will only be further enhanced with your efforts so a hearty thanks from us...”

Replication technologies update



- Streams replacement with ADG
 - CMS (online – offline)
 - ALICE (online – offline)

- No changes (Streams11g) for
 - ATLAS
 - online - offline (ADG foreseen in 2013)
 - offline - T1s
 - LHCb
 - online – offline
 - offline – T1s
 - COMPASS (online – offline)

- Streams
 - Downstream capture (ATLAS, LHCb)
 - 11g to 11g (default configuration)
 - 11g to 10g (intermediate state during upgrade)
 - Cascading downstream capture (in combination with Dataguard)
 - Long term replication test (over 3 months) with production data
 - Validation of various patches
- CMS active standby
 - Tested with CMS critical data sets (Conditions) and client applications
 - Good replication performance (avg. 38MB/s, max 82 MB/s)
 - Good stability
- Feedback provided to Oracle development teams

- Goal
 - Validation of technologies with CERN's data
 - Direct performance comparison
- Tested technologies
 - Streams 11g
 - GoldenGate 11g
 - Active DataGuard 11g
- Workload from production
 - LHCb LHC file catalog
- On same hardware and software configuration

- Results sent to Oracle development teams
- Further optimizations suggested by development teams
 - Streams: GoldenGate XStream optimization
 - ADG: patching + tuning
- Tests repeated after implementation of the improvements
- Much more better performance observed!



Monitor	Maps	DBs	Streams	Standbys	Plots	Errors	Availability	History	Reports	
All	All-Prod	ALICE	ATLAS	CMS	COMPASS	LHCb	OTHERS	TEST	TEST11G2	TEST_ATLAS

DATAGUARDS SETUP									
#	Redo log stream	Standby State	Type	Data latency	File transport lag	Apply rate	Recovery status		
1	ADCR.CERN.CH=>ADCR_STDBY.CERN.CH	MOUNTED	PHYSICAL STANDBY / MAXIMUM PERFORMANCE	22 sec	0 sec	8.7 MB/s	1 MRP(s) up		
2	ALIONR.CERN.CH=>ALIONR_STDBY.CERN.CH	MOUNTED	PHYSICAL STANDBY / MAXIMUM PERFORMANCE	18 sec	0 sec	791.9 kB/s	1 MRP(s) up		
3	ATLR.CERN.CH=>ATLR_STDBY.CERN.CH	MOUNTED	PHYSICAL STANDBY / MAXIMUM PERFORMANCE	26 sec	1 sec	821.6 kB/s	1 MRP(s) up		
4	ATONR.CERN.CH=>ATONR_STDBY.CERN.CH	MOUNTED	PHYSICAL STANDBY / MAXIMUM PERFORMANCE	29 sec	1 sec	617.3 kB/s	1 MRP(s) up		
5	CMSONR.CERN.CH=>CMSONR_STDBY.CERN.CH	MOUNTED	PHYSICAL STANDBY / MAXIMUM PERFORMANCE	32 sec	1 sec	653.9 kB/s	1 MRP(s) up		
6	CMSR.CERN.CH=>CMSR_STDBY.CERN.CH	MOUNTED	PHYSICAL STANDBY / MAXIMUM PERFORMANCE	25 sec	0 sec	982.0 kB/s	1 MRP(s) up		
7	CMSTEST.CERN.CH=>CMSTEST_STDBY.CERN.CH	READ ONLY	/ MAXIMUM PERFORMANCE	8 day(s) 22 hr 14 min 52 sec	8 day(s) 4 hr 54 min 15 sec	-	DOWN		
8	LCGR.CERN.CH=>LCGR_STDBY.CERN.CH	MOUNTED	PHYSICAL STANDBY / MAXIMUM PERFORMANCE	17 min 56 sec	2 min 33 sec	53.9 MB/s	1 MRP(s) up		
9	LHCBONR.CERN.CH=>LHCBONR_STDBY.CERN.CH	MOUNTED	PHYSICAL STANDBY / MAXIMUM PERFORMANCE	26 sec	2 sec	206.0 kB/s	1 MRP(s) up		
10	LHCBR.CERN.CH=>LHCBR_STDBY.CERN.CH	MOUNTED	PHYSICAL STANDBY / MAXIMUM PERFORMANCE	19 sec	5 sec	256.9 kB/s	1 MRP(s) up		
11	TEST11G.CERN.CH=>TEST11G_STDBY.CERN.CH	READ ONLY WITH APPLY	PHYSICAL STANDBY / MAXIMUM PERFORMANCE	3 hr 6 min 22 sec	3 hr 6 min 22 sec	0 B/s	1 MRP(s) up		

07-10-11 14:43:09 auto refresh

Links: [3D OEM](#), [3D TWIKI](#)

Contact: [Pdb Service](#)

Monitor	Maps	DBs	Streams	Standbys	Plots	Errors	Availability	History	Reports
Info	Overview	Stats							

ATLR_STDBY.CERN.CH

Instances

Id	Name	CPU Usage	CPU Load	Stream Pool Size	Stream Pool Free	Stream Pool Utilization	Redo Size	Redo Generated/s	Bytes Read/s	Bytes Written/s	PGA size	Logons /s	Current logons
1	atlr1	1%	0.41	592.0 MB	592.0 MB	0%	0 B	0 B	2.0 MB	2.5 MB	432.5 MB	1	37
2	atlr2	1%	0.42	592.0 MB	592.0 MB	0%	0 B	0 B	43.1 MB	1.2 MB	558.0 MB	1	41

Standby stats

Primary database [Recovery Status](#)

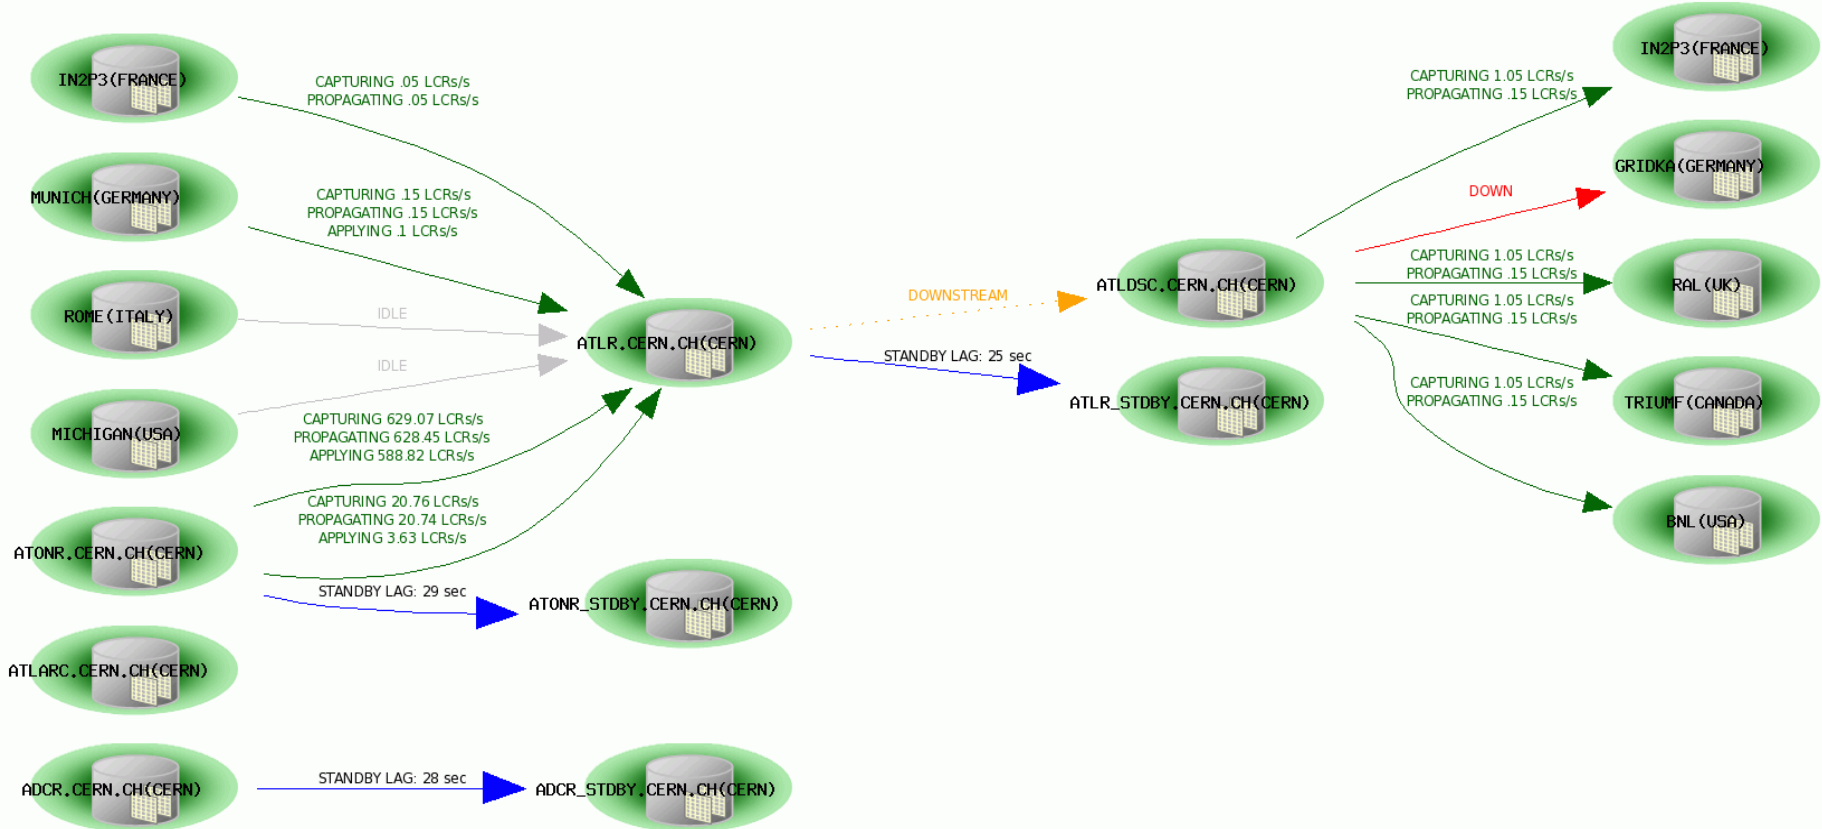
[ATLR.CERN.CH](#) Up (1 processes)

Apply Lag	Transport Lag	Applied redo volume	Current redo apply speed	Redo apply active speed	Redo apply average speed
30 sec	3 sec	4168.3 GB	1.2 MB/s	12.6 MB/s	1.7 MB/s

07-10-11 14:45:41 auto refresh

Monitor **Maps** DBs Streams Standbys Plots Errors Availability History Reports
 ALICE **ATLAS** CMS COMPASS LHCb OTHERS TEST TEST11G2 TEST_ATLAS

TOPOLOGY | auto refresh | show standbys | hide details



Outreach



CERN
openlab

- “Overview of Database Technologies” **Eric Grancher**, Computing and Astroparticle Physics 2nd ASPERA Workshop, Barcelona 30-31 May 2011
- “CERN IT-DB Deployment, Status, Outlook” **Luca Canali**, ESA-GAIA DB Workshop, ISDC, March 2011
- “Oracle at CERN” **Eric Grancher**, CERN openlab summer students programme 2011, CERN, August 2011
- “AWR and ASH in 3-D: Performance Analysis Tools No DBA Has Seen Before” **Anton Topurov, Debaditya Chatterjee, John Beresniewicz**, Oracle Open World 2011, San Francisco 2-6 September 2011
- “Successful Oracle Database 10.2 to 11.2 Migration with Oracle Real Application Testing” **Katarzyna Dziejniewicz-Wojcik, Prabhaker Gongloor, Eric Grancher**, Oracle Open World 2011, San Francisco 2-6 September 2011

