BNL Oracle database services status and future plans

Carlos Fernando Gamboa

Grid Group, RACF Facility

Brookhaven National Laboratory, US

Distributed Database Operations Workshop
CERN, Geneva. November 2009

ATLAS Oracle DB services hosted at BNL

The US ATLAS production Database services at BNL are distributed among 2 independent oracle clusters

RAC 1

Dedicated to serve US ATLAS Conditions Database TAGS database test

-Besides serving BNL worker nodes, this database is also accessed by different Tier 2 and some Tier 3 site worker nodes distributed across the US.

-Temporarily hosts TAGS test data,

-Streams replication process runs on a different node than the RMAN backup jobs.

Database service distribution RAC 1

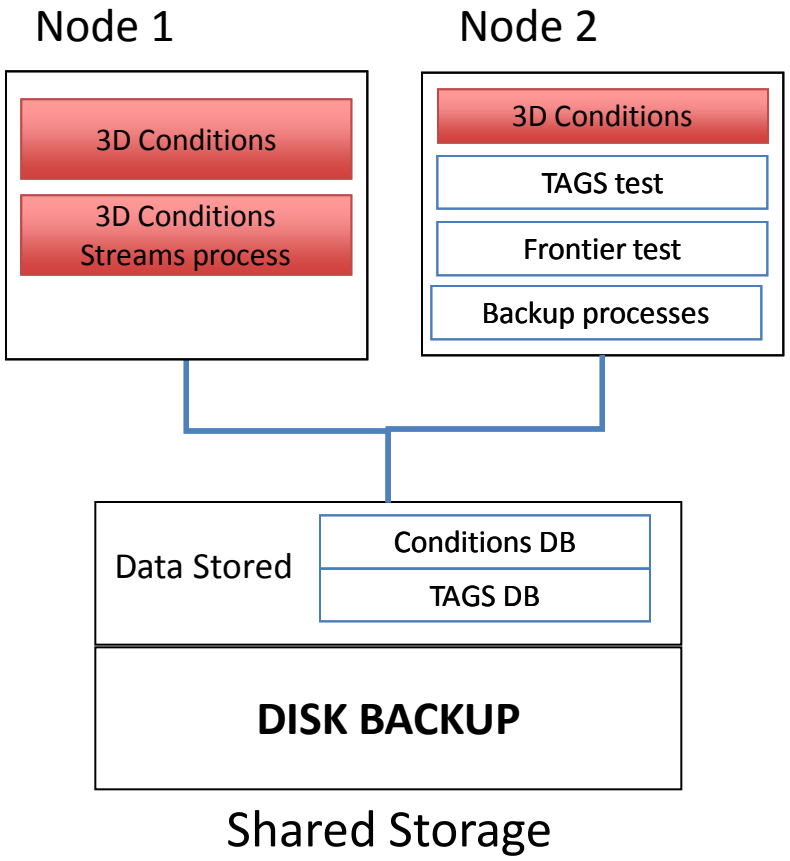

ATLAS Oracle DB services hosted at BNL

RAC 2:

- Dedicated to host LFC and FTS data.
- LFC service for Tier 3 recently enabled.
- Each database service is distributed on only one node. In case of failure, database services will fail over to the surviving node.
- Cluster inside BNL firewall.
- TSM is enabled for tape backups besides the disks backups.

BNL database services are monitored by OEM, Ganglia and Nagios.

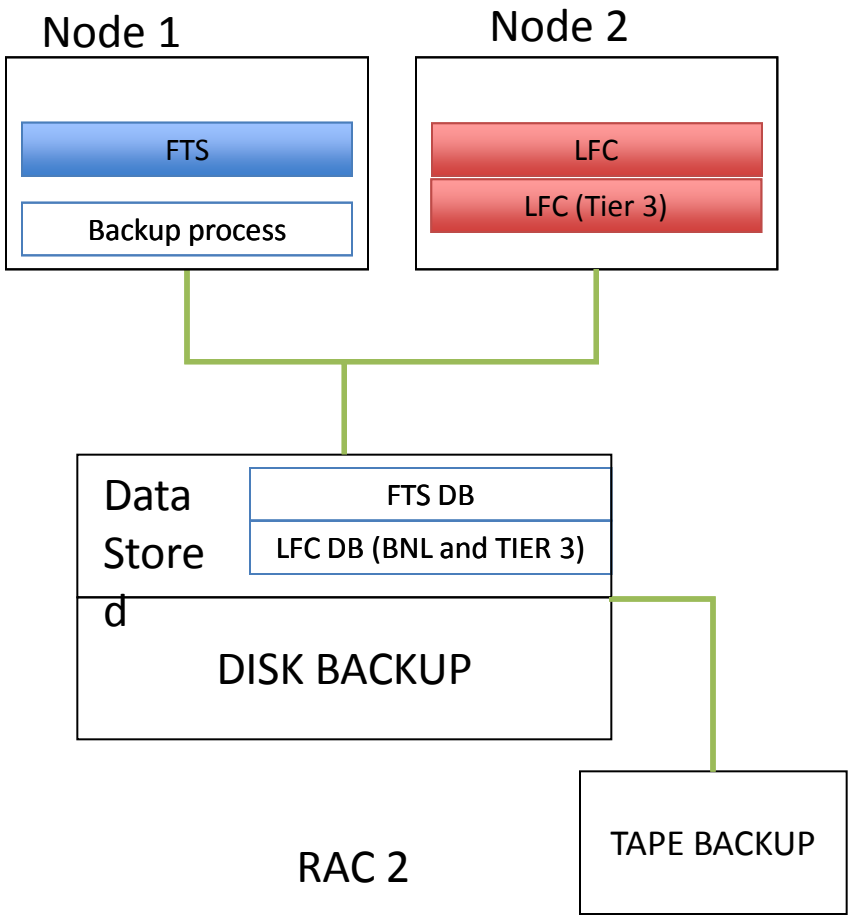

Future plans presented in previous workshop

Monday 20 April 2009 top

09:30->*13:00* **Operations and Service Review** slides

09:30 Welcome and Workshop Goals (20') — Maria Girone (*CERN*) , Luis Diaz (*PIC*)

09:50 WLCG Service Roadmap for 2009 / 2010 (40')
WLCG Collaboration Workshop agenda)
- Summary of recommendations from recent LI
- Scale Test for Experiment Program - STEP 20(
- **Service Issues - Outstanding Oracle bug**
- WLCG Operations Summary - **What do DBAs**
- Remaining challenges and directions... (2011

10:30

10:45 Tier0 and Tier1 Service Review/Plans Round
with contributions by all the sites concerned

Done →

Conditions DB and TAGS DB will be isolated, next week. →

Done →

Future Plans

Atlas Conditions Database

Hardware

-Replace the head nodes with Quad core Nodes of 3GHz.

-To acquire 21.6 TB RAW storage, SAS disks of 450GB and 15krpm distributed among 1 DS3400 and 3 DS3000. Will use the same DAS technology.

-This modular technology allows to be integrated in an SAN configuration when required.

Database Services

-TAGS database tests service will be allocated on a different test hardware

-Frontier test service

As of now, there are no procedures running on the database side for these tests. It is planed to include it to the FroNTier test infrastructure in future development efforts in coordination with CERN IT DB and Atlas DB group and froNtier developers folks.

The goal here is to minimize possible impact on the BNL production service (regular Cond.DB user service and T0->T1 database replication) while testing froNtier technology.

Carlos Fernando Gamboa 10

Future plans presented in previous workshop

Done

Atlas wide policy, Done

Next week, Tests done.

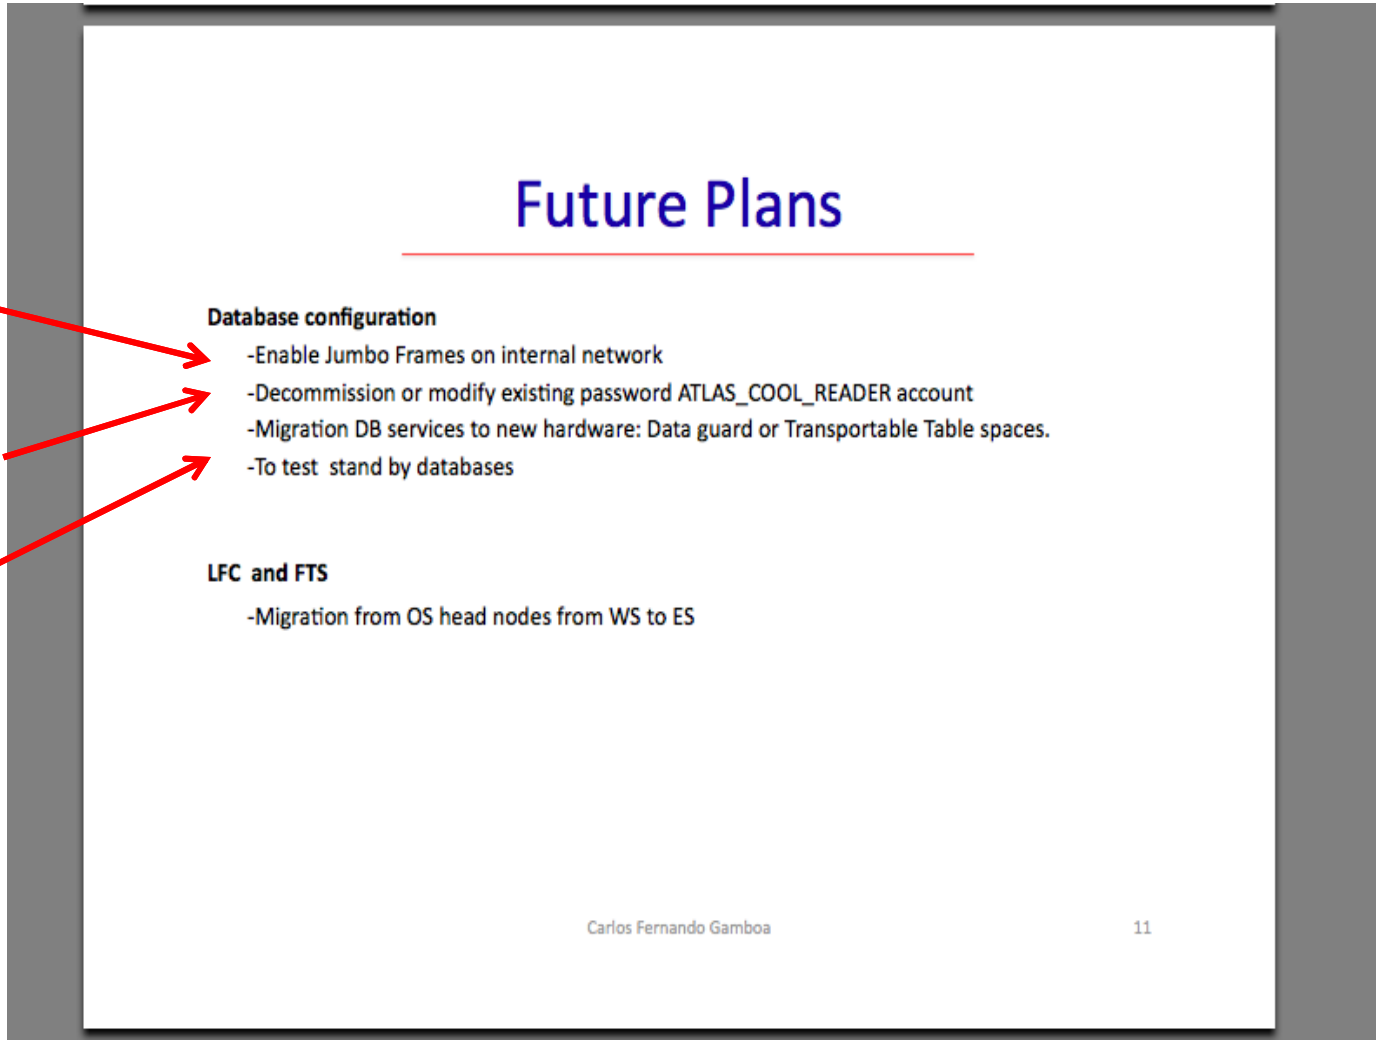

BNL ORACLE Cluster 2 nodes, DAS configuration general cluster hardware topology

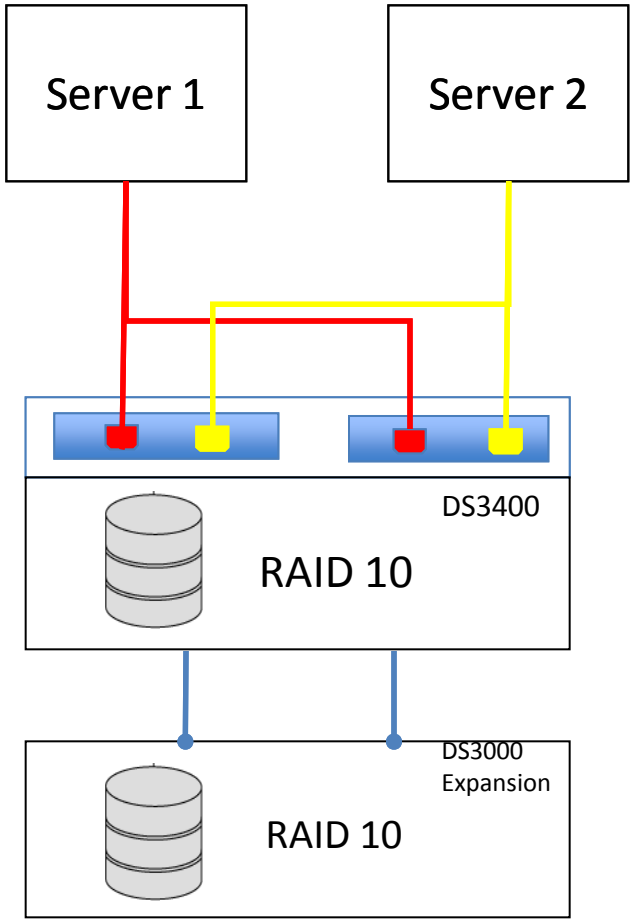

IBM 3550 Server description:

- 2 dual core 3GHz, 64 bits Architecture
- RAM 16GB

Interconnectivity

Server to clients

-NIC 1000Gb/s.

Server to storage

-HBA QLogic 4Gb FC Dual-Port PCI-X

-1M LC-LC Fibre Channel Cable

Storage

IBM DS3400 FC dual controller

-2 Hot Swap disk per enclosure

-4 Gbps SW SFP Transceiver

-12 SAS disks 15krpm and 10krpm, size 300 GB/disk

IBM DS3000 storage expansion

12 SAS disks 15krpm, size 300 GB/disk

New capacity installed / cluster distribution

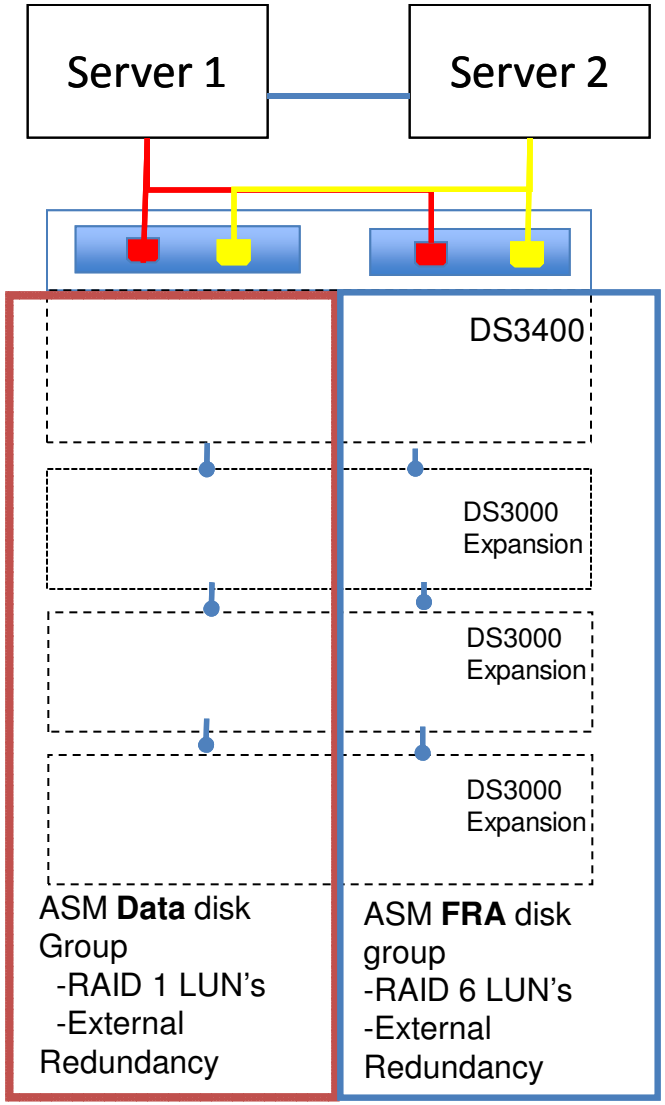

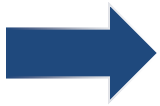

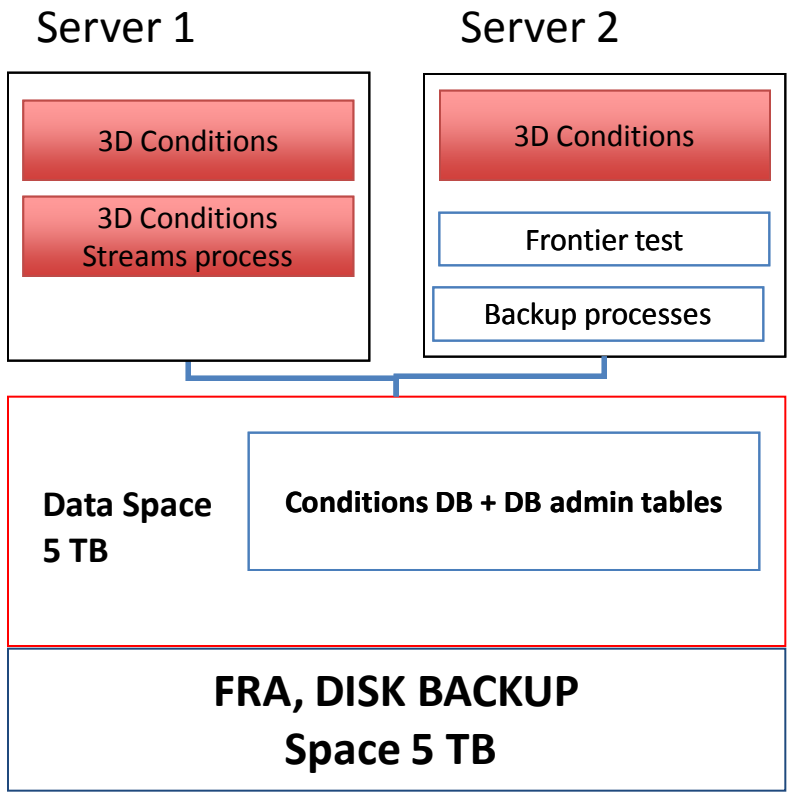

IBM 3650 Server description:
- 2 **Quad Core** 3GHz, 64 bits Architecture
- RAM 16GB - 32GB

Storage
IBM DS3400 FC dual controller
- -1 GB buffer cache
- -2 Hot Swap disk per enclosure
- -4 Gbps SW SFP Transceiver
- -12 SAS disks 15krpm 450 GB/disk

IBM DS3000 storage expansion
- -36 SAS disks 15krpm, size 450 GB/disk

BNL Oracle services status for ATLAS

ATLAS production Oracle services hosted at BNL are distributed among 3 RAC clusters as:

RAC #	Oracle service	Nodes	Manufacture Model	Processor	Memory	HBA	NIC
1	Conditions DB/TAGS	2	IBM 3550	2 dual Core Intel Xenon Processor 5160 3GHz	16GB	QLogic 4Gb FC Dual-Port PCI-X	1000Gb/s
2	FTS	1	IBM 3650		8GB		
	LFC	1					
3	NEW Conditions DB	2	IBM 3650	2 quad Core Xeon X5450 3GHz	16GB		

Production head nodes summary

BNL Oracle services status for ATLAS

RAC #	Oracle service	Total RAW space	Total SPACE after RAID 10	Manufacture Model	Disk type, Speed, Size	Storage Controllers	Redundancy	IOPS per disk measured (ORION VERSION 11.1.0.7.0)
1	Conditions DB	6TB	2.8 TB	IBM DS3400, DS3000	SAS, 12 Disks 15K rpm 12 Disks 10K rpm 300GB	Dual FC controller 4 Gbps SW SFP Transceiver	Hot Swappable SAS disks Dual power supply	~200 IOPS / disk Measured with 5 LUNS RAID 1, 10 disks.
2	FTS	6TB	2.8 TB	IBM DS3400, DS3000	SAS, 24 Disks 15K rpm, 300GB			
3	New conditions Db	~21.6TB	5TB	IBM DS3400 3 DS300	SAS, 48 Disks 15K rpm, 450GB			

Production storage summary

BNL Oracle services status for ATLAS -General database distribution-

OS level 64 Bits	Database	Oracle Database Release	Data ASM disk Group	Backup ASM disk Group	SGA	Oracle ASMlibs
RHEL ES 2.6.9-89.0.11 release 4	Conditions DB / TAGS	10.2.0.4 PSU1	1.4TB	1.4TB	4 GB	2.0.2.4
RHEL WS 2.6.9-89.0.15 release 4	FTS and LFC				4GB	
RHEL 5 Server 2.6.18-164.6.1	Conditions DB	10.2.0.4 PSU2	5TB	5TB	9 GB ASMM (disabled)	

BNL Oracle services status for ATLAS
-Future plans-

- Planning to migrate LFC and FTS BNL services to OS RHEL 5 server / hardware
- Migrate OEM to RHEL5 and hardware.
- Planning to migrate GUMS database service from MySQL to Oracle.
- Continue understanding FRONTIER technology in the ATLAS context.
  Previous summary presented past Atlas Computer & Software week September 2009, http://indico.cern.ch/contributionDisplay.py?contribId=7&confId=50976.
- Continue understanding interactions between user jobs and database.
  http://indico.cern.ch/contributionDisplay.py?contribId=9&confId=50976