
BNL Oracle database services status and future plans

Carlos Fernando Gamboa

RACF Facility

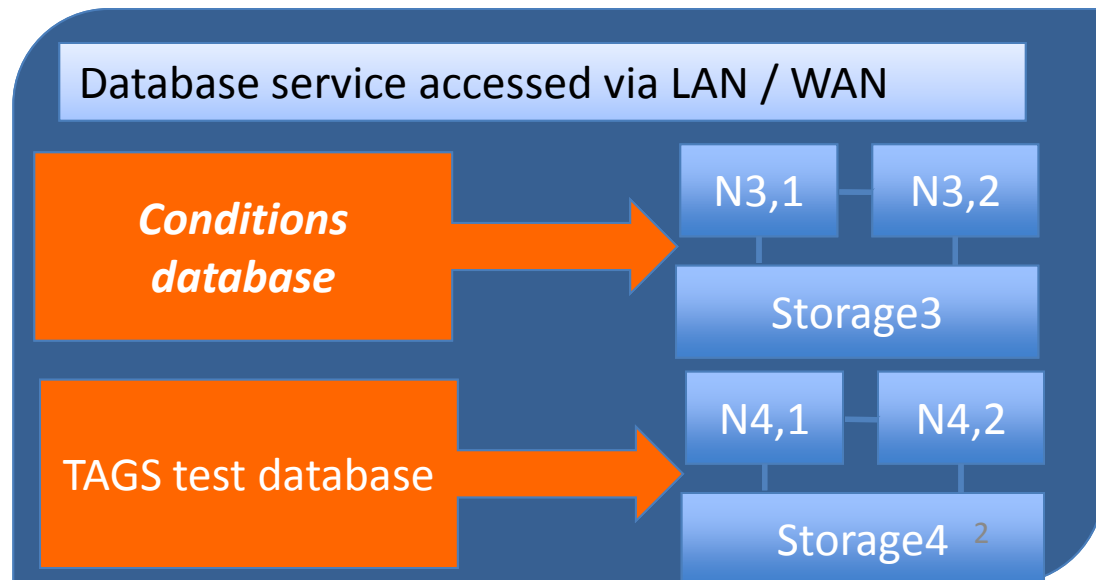
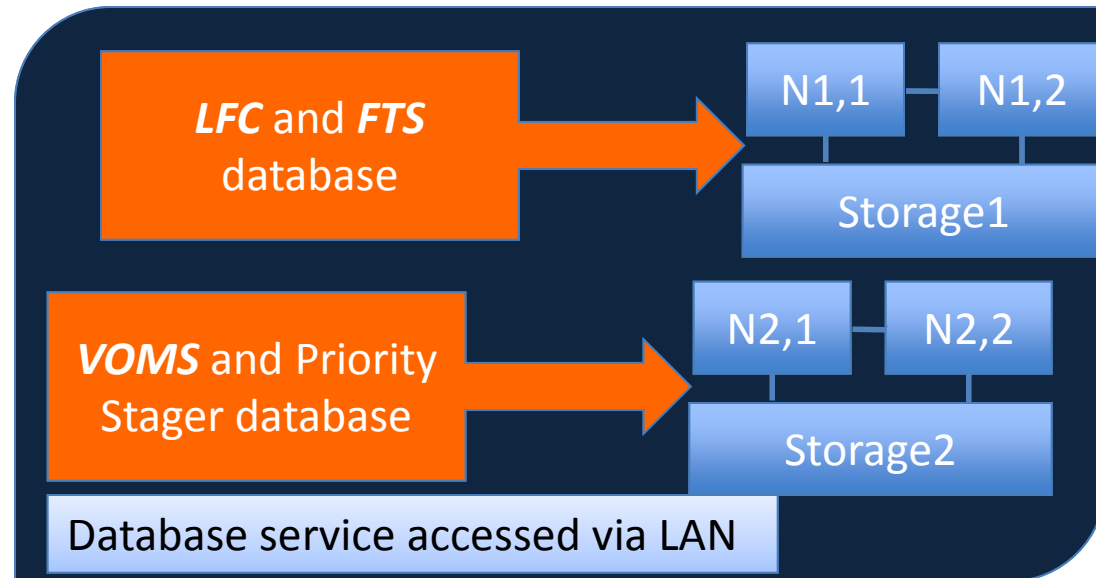
Brookhaven National Laboratory, US

Distributed Database Operations Workshop
CERN, Geneva. November 2010

Topology Oracle Database services hosted at BNL

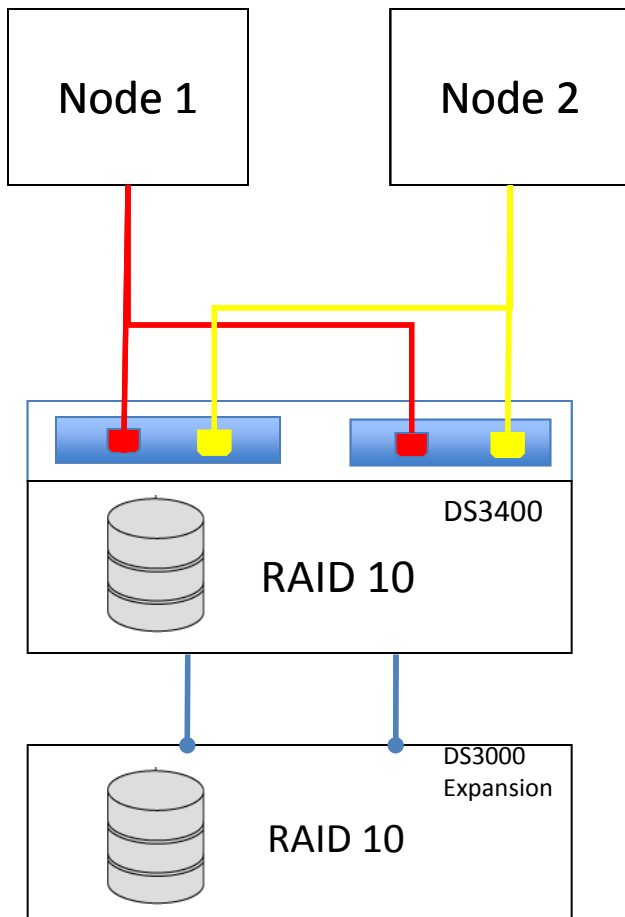
Independent clusters set per application service.

- Dual nodes, Direct Attach Storage (DAS)
- Storage distribution adjusted to application needs
 - Hardware RAID levels
 - Storage and spindles
- Flexible architecture that allows to increase nodes and storage per application needs.
- Homogenous software stack deployed:
Real Application Cluster 10gR2.
 - Database server
 - Clusterware
 - ASM file system



BNL ORACLE Cluster 2 nodes, DAS configuration

General cluster hardware topology



IBM 3550/3650 Server description:

- 2 dual – 2 quad core 3GHz, 64 bits Architecture
- RAM 16GB -32GB

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, size 300 GB/disk or 450GB/disk

IBM DS3000 storage expansion

- 12 SAS disks 15krpm, size 300 GB/disk to 450GB/disk

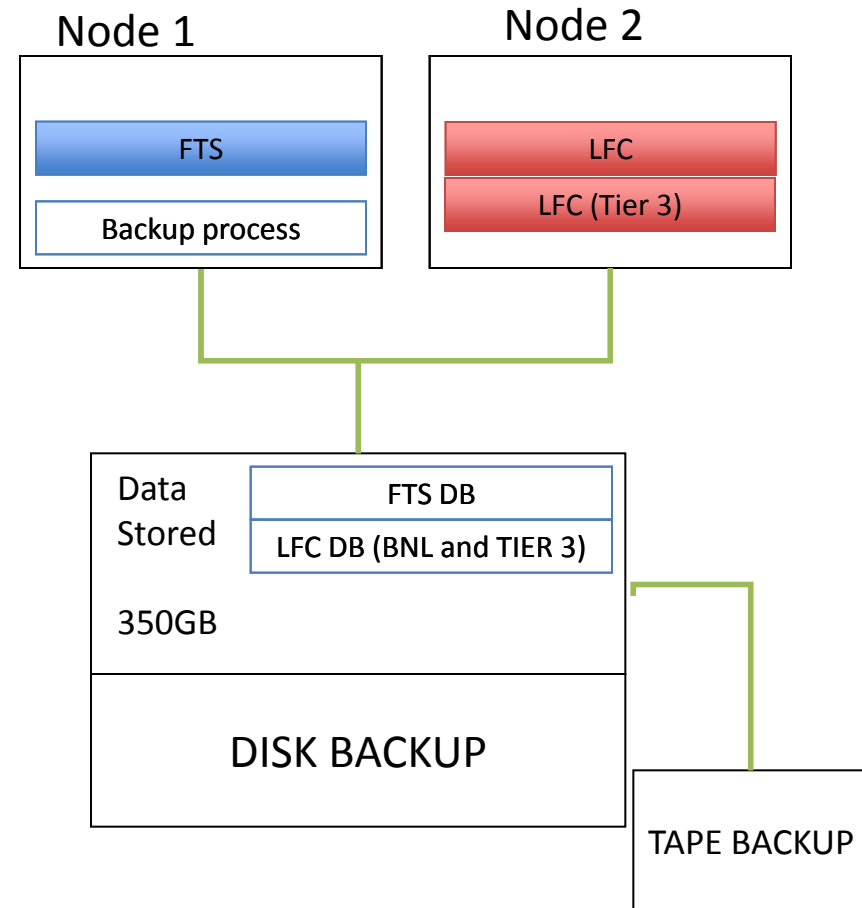
Monitor tools

- Oracle Enterprise Manager Grid Control
- Nagios
- Ganglia

Distribution of database services per production cluster

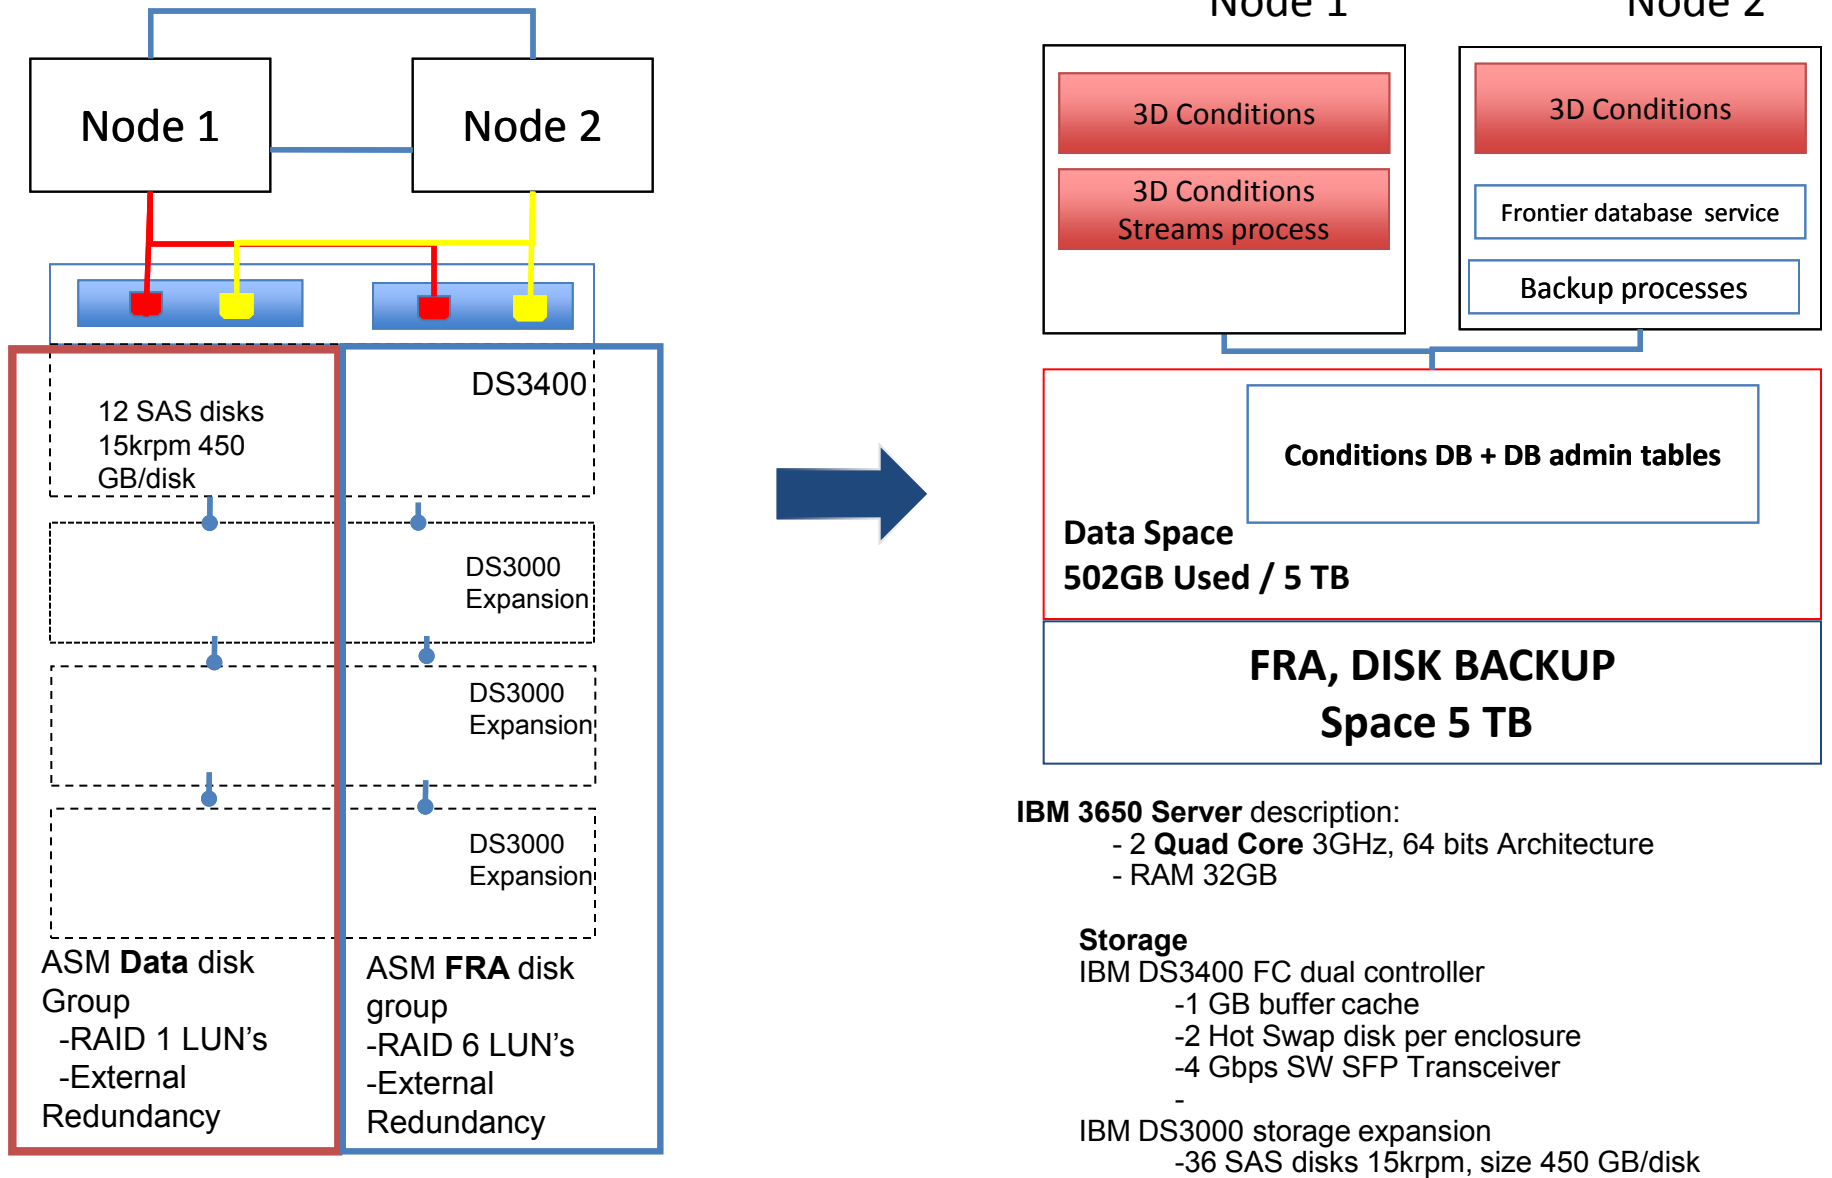
LFC and FTS database

- Dedicated to host BNL, US Tier 3 LFC and FTS data.
- 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.



Distribution of database services per production cluster

Conditions Database



BNL Oracle services status

Operational issues

PSU April 2010 and ORA-7445:

- 2 occurrences, no performance or service degradation issues observed.

-Patch 6196748 applied: Following Oracle SR 3-182235985 recommendation and based on testing efforts by CERN IT (Persistency and Database group) and Atlas DBA group.

Savannah Bug 16836 (<https://savannah.cern.ch/task/?16836>)

Streams inconsistency at BNL affecting 3 schemas.

Entire report was presented in Database Administration Matters 03/01/2010.

<http://indico.cern.ch/conferenceOtherViews.py?view=standard&confId=86629>

BNL Oracle services status

Future plans

New hardware acquired for FTS, LFC and VOMS database IBM DS3500

- 2 IBM 3650 4 cores (3.46GHz) / node.
- Storage IBM DS3500.
 - Up to 4 servers, cluster topology (DAS).
 - Increase on storage connectivity speed between nodes and storage (8Gb HBA, transceiver set).
 - 36 x 600GB disks (~21.6TB RAW space) 15K RPM.
 - 10Gb Ethernet/backups?
- Database service will be migrated using Dataguard.

BNL Oracle services status

Future plans

TAGS test database migration to 11GR2.

Backup slides

Summary production cluster

Software/hardware

BNL Oracle services status

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		16GB		
2	FTS	1	IBM 3650	2 dual Core Intel Xenon Processor 5160 3GHz	8GB	QLogic 4Gb FC Dual-Port PCI-X	1000Gb/s
	LFC	1					
3	Conditions DB	2	IBM 3650	2 quad Core Xeon X5450 3GHz	16GB		

Production head nodes summary

BNL Oracle services status

RAC #	Oracle service	Total RAW space	Total SPACE after RAID 10	Manufacturer Model	Disk type, Speed, Size	Storage Controllers	Redundancy	IOPS per disk measured (ORION VERSION 11.1.0.7.0)
1	TAGS test DB	6TB	2.8 TB	IBM DS3400, DS3000	SAS, 12 Disks 15K rpm 12 Disks 10K rpm 300GB			
2	FTS	6TB	2.8 TB	IBM DS3400, DS3000	SAS, 24 Disks 15K 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.
3	Conditions DB	~21.6TB	5TB	IBM DS3400 3 DS300	SAS, 48 Disks 15K rpm, 450GB			
4	VOMS, Priority Stager	3TB	1.4TB	IBM DS3400	SAS, 12 Disks 15K rpm, 300GB			

Production storage summary

BNL Oracle services status

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- 89.0.31.1 release 4	TAGS Test, VOMS, Priority Stager	10.2.0.4 PSU4	1.4TB	1.4TB	4 GB	
RHEL WS 2.6.18- 1.0.31.1 release 4	FTS and LFC				4GB	2.0.4
RHEL 5 Server 2.6.18- 194.17.1	Conditions DB	10.2.0.4 PSU4	5TB	5TB	9 GB ASMM (disabled)	