

# CNAF Database Service

<i>Barbara Martelli</i>	CNAF-INFN
<i>Elisabetta Vilucchi</i>	CNAF-INFN
<i>Simone Dalla Fina</i>	INFN-Padua



# Overview

- Present Database Service Organization
- CNAF Database Experience
- CNAF – Tier1's HW Resources
- Plans for the Near Future
- CNAF Contribution to LCG 3D Project
- Summary

# Present Database Service Organization

- For each experiment, one person is responsible for installing, maintaining and monitoring all experiment specific applications and services.
- LCG local group is responsible for installing, maintaining and monitoring all LCG applications and services.
- A few Tier1 generic databases are maintained by Database group.
- Database group is composed by 4 people, 3 onsite and one (Simone Dalla Fina) collaborates from Padua.

# CNAF Database Group Experience

- MySQL
- PostgreSQL
- Oracle 9i and Oracle 10g
- High Availability environment built on RedHat Enterprise Cluster Suite.
- Entity-Relationship schema writing.

# MySql Experience

- As almost database based applications on LCG software rely on MySql, CNAF LCG local support group has experienced installation, administration and tuning of MySql databases.
- Experience in development of MySql based applications (e.g. VOMS), exploiting innoDB tables and one-way replication.

# PostgreSQL Experience

- CNAF Hardware resources database built on PostgreSQL.
- Developed an application which contains information about all HW resources and performs automatic configuration of switches, DNS and DHCP servers needed for farms installation.
- Daily tasks such as backups and vacuum are performed with cron jobs.
- During down time due to RDBMS upgrades, patches installation and application upgrades, a secondary database is used, this is not a replica database, it's updated "by hands" restoring the last primary database dump

# Oracle 9i Experience

- CASTOR application, based on Oracle 9i, is running in production environment at CNAF.
- Database group is involved in installation and maintenance tasks.
- Backups
  - Weekly Hot Backup of the entire database
  - Daily full export and Castor tables export
  - Legato NetWorker module for Oracle (which exploits with RMAN) backup on Tape Library STK180 (LTO/9840)

## Oracle 9i Experience (contd)

- Collaboration with CERN in Oracle based RLS component testing.
- Installation of Oracle 9.2.0.4 and Oracle Application Server 9.3.0.1 through CERN installation toolkits.
- Installed 2 biprocessor (PIII 2,4 GHz) with RedHat Advanced Server 2.1 and:
  - Oracle 9i single instance for the RLS database backend, Oracle datafiles reside on a SAN partition (RAID 5) on the first machine
  - Oracle 9i Application Server single instance for the RLS application on the second machine



## High Availability experience

- Deployment of an high available Oracle service on top of RedHat ES 3.0 cluster.
- 2 PIII 2,4 GHz with shared storage in failover configuration.
- Oracle 9i is configured as cluster service.
- The stand-by machine can be the primary node for another service.

# Oracle 10g Experience

- We have recently tried the installation of new release of Oracle 10g.
- Testbed composed by 2 machines at CNAF
  - Dual Intel Xeon 2,4 GHz, HyperThreading enabled, 512 MB RAM, two 60 GB disks, with Scientific Linux
- 1 machine at Padua-INFN
  - Dual PIII 1,3 GHz, 1GB RAM, two 60 GB disks, with Scientific Linux
- Done the set up of one way replication between the machines at CNAF.

# CNAF Hardware Resources

- At present we have available 4 machines on site:
  - Two dual PIII 2,4 GHz, 512 MB RAM with two 60 GB disks
  - Two dual Intel Xeon 2,4 GHz with HyperThreading, 512 MB RAM with two 60 GB disks
- One machine at INFN-Padua
  - Dual PIII 1,3 GHz, 1GB RAM, two 60 GB disks
- Possibility to connect to our SAN to enlarge storage and share it with other servers
- Gigabit interfaces on all machines
- Possibility to add new hardware in the future, if needed

# Plans for the Near Future

- Necessary to create a production environment which has to be high available, fault tolerant and responsive.
- At the same time we need to centralize the administration of some databases which are now scattered on many different servers, this increase the administration time and the probability of mistakes.
- The simplest solution is the creation of a RedHat cluster of DB servers, in high availability configuration, each one being the primary server for one different DB instance.

# CNAF Contribution to LCG 3D Project

- CNAF can give the following type of contribution:
  - Database support: (install bug fix patch, security patch upgrade/migrate, day to day database administration activity).
  - System support: (install, bug fix patch, security patch, upgrade/migrate, day to day system administration activity)
  - Research and test activities on Streams replication and RAC environment.
  - Support other sites on activities such installation of databases, ER schemata writing, and daily administration.

## CNAF Contribution to LCG 3D Project

- By now, services can be guaranteed on working days and working hours.
- One way to ensure a 24x7 service, could be shifts with other sites.
- In the future, we expect to follow the growth of applications needs increasing our effort and possibly the number of people involved in the database group.

# Summary

- CNAF database group has experience on many RDBMS like MySql, PostgreSQL, Oracle.
- Experience also on high availability cluster configurations.
- We are moving our databases to a centralized, high available, fault tolerant configuration, in order to rationalize hardware and human resources.
- We have a lot of interest in exploring replication and cluster tools and techniques in the Oracle environment.

# Contacts

- Barbara Martelli  
[barbara.martelli@cnaif.infn.it](mailto:barbara.martelli@cnaif.infn.it)
- Elisabetta Vilucchi  
[elisabetta.vilucchi@cnaif.infn.it](mailto:elisabetta.vilucchi@cnaif.infn.it)
- Simone Dalla Fina  
[dfina@pd.infn.it](mailto:dfina@pd.infn.it)
- Pier Paolo Ricci  
[pierpaolo.ricci@cnaif.infn.it](mailto:pierpaolo.ricci@cnaif.infn.it)