



TRIUMF

CANADA'S NATIONAL LABORATORY FOR PARTICLE AND NUCLEAR PHYSICS

*Owned and operated as a joint venture by a consortium of Canadian universities
via a contribution through the National Research Council Canada*

Databases at TRIUMF

Andrew Wong

ORACLE

 PostgreSQL


MySQL

LABORATOIRE NATIONAL CANADIEN POUR LA RECHERCHE EN PHYSIQUE NUCLÉAIRE ET EN PHYSIQUE DES PARTICULES

*Propriété d'un consortium d'universités canadiennes, géré en co-entreprise à partir d'une contribution administrée par le
Conseil national de recherches Canada*

PostgreSQL / MySQL

PostgreSQL 8.3.4 – dCache, size ~ 20GB
(Used Slony replication to upgrade from 8.2 on Jan 7, 2009)

MySQL 4.1 (SL)

Database sizes: Monbox ~4GB / LFC ~10GB / Smallhsm ~ 1GB

Replication – for failover
– LFC & Monbox

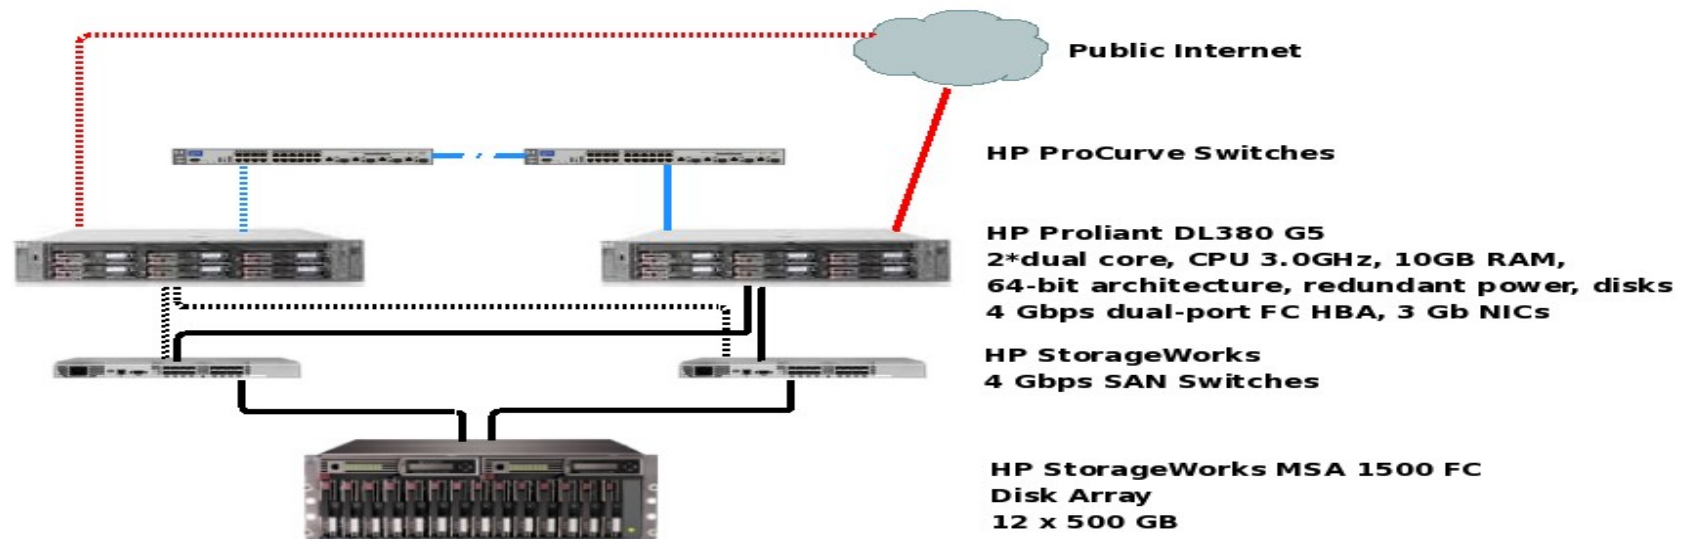
Database Backups:

Database export dumps (depending on database every 2 or 6 hours)

Testing PostgreSQL warm standby



TRIUMF Oracle RAC



High availability configured in 2008

Databases: Conditions, TAGS, FTS



TRIUMF Oracle RAC

Node 1

- 3D Conditions instance
- 3D Conditions Streams process

Node 2

- 3D Conditions instance
- TAGS Testing instance
TRIUMF is volunteering to host ATLAS TAGS
- RMAN backups



Oracle Administration

Backups: Nightly RMAN jobs

Support / Monitoring:

OEM – email alerts (alert.log errors, apply process stopped, etc)

Scripts – CRS, ASM usage

Nagios – ping, ssh, paged if Oracle node is down

Ganglia - <http://gridinfo.triumf.ca/ganglia/>



Monitoring


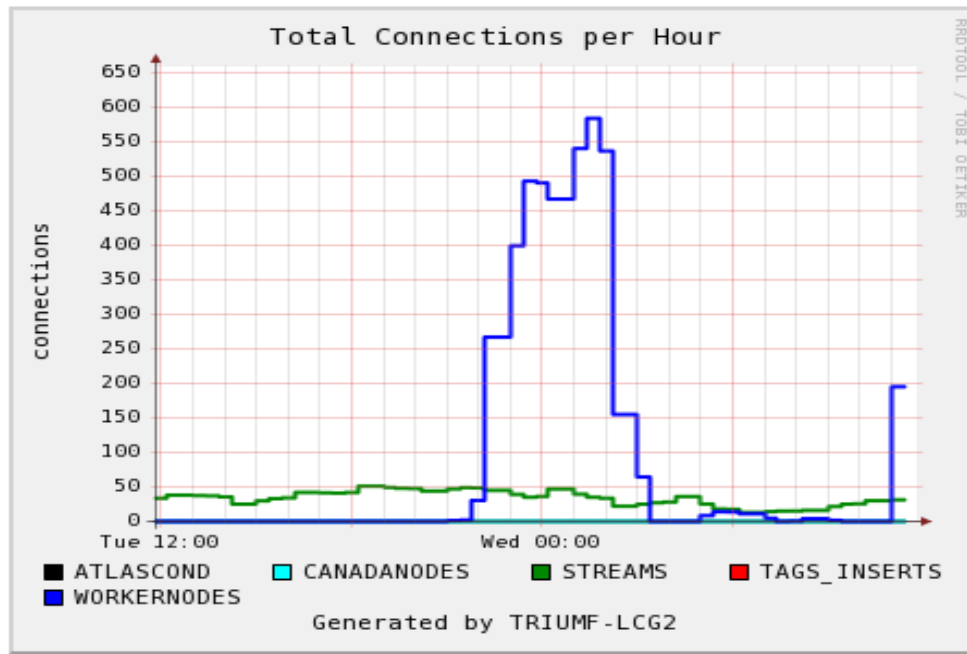
Ganglia – customized using PHP

Oracle Connection Summary

Individual

Cummulative

Last Day



Oracle RAC Storage

HP StorageWorks MSA1500 FC

Disk Array – 12 x 500GB

=====

MSA1500_physicaldrive Status Report on host tr-orac2.triumf.ca:

Tool /usr/sbin/hpacucli ctrl wwn=500508B30092F7B1 physicaldrive all show status

physicaldrive 5:1 (box 5:bay 1, 500.0 GB): OK

physicaldrive 5:2 (box 5:bay 2, 500.0 GB): OK

physicaldrive 5:12 (box 5:bay 12, 500.0 GB, spare): OK

physicaldrive 5:3 (box 5:bay 3, 500.0 GB): OK

physicaldrive 5:4 (box 5:bay 4, 500.0 GB): OK

physicaldrive 5:5 (box 5:bay 5, 500.0 GB): OK

physicaldrive 5:6 (box 5:bay 6, 500.0 GB): OK

physicaldrive 5:7 (box 5:bay 7, 500.0 GB): OK

physicaldrive 5:8 (box 5:bay 8, 500.0 GB): OK

physicaldrive 5:9 (box 5:bay 9, 500.0 GB): OK

physicaldrive 5:10 (box 5:bay 10, 500.0 GB): OK

physicaldrive 5:11 (box 5:bay 11, 500.0 GB): OK



Current RAC Performance

I/O Performance:

From <https://twiki.cern.ch/twiki/bin/view/PSSGroup/OrionTests>

CERN – 128 SATA disks IOPS=12000

PICS – 80 SAS disks IOPS=16000

TRIUMF – 9 SATA disks IOPS=550

Note: ATLAS has significantly improved the conditions DB access, current system can sustain reprocessing activities, nevertheless there is limited performance with concurrent TAGs upload and Conditions DB access (to be addressed in our upcoming RAC upgrade)

More storage for 2009:

Current Max Storage is ~2.2TB

For 2009 will need ~ 24TB (Cool & TAGS)



2009 Plans

For Conditions DB RAC:

Add more storage to current RAC

Upgrade to higher performance drives

Exploring new Solid State Drives (SSD) technology

-> Single SSD IOPS=35,000 versus a typical HDD which has 350 IOPS.

-> As long as it is supported by WLCG middleware



2009 Plans

For TAGS DB: (New Oracle RAC)

2 - 2 x Dual Core Servers
High performance drives
More, more, more storage

For LFC: Will leave on MySQL



Questions ?

