



RAL Status and Plans

**3D Workshop,
CERN,
26- 27 November 2009**

Carmine Cioffi

Database Administrator and Developer





- **3D**
 - Database configuration and HW spec
 - Storage configuration and HW spec
 - Future plans
- **CASTOR**
 - Database configuration and HW spec
 - Storage configuration and HW spec
 - Schemas Size and Versions
 - Future plans
- **Backup configuration**



- **3 nodes RAC for ATLAS (Ogma)**
 - Red hat 4.8
 - 64 bit
 - 2 Quad Core Xeon(R) E5410 @ 2.33GHz
 - 16 GB
- **2 nodes RAC for LHCb (Lugh)**
 - Red Hat 4.8
 - 64 bit
 - 2 Dual- Core AMD Opteron(tm) 2216
 - 16 GB



- **For both RAC Ogma and Lugh:**
 - **Oracle 10.2.0.4**
 - **Single OCR**
 - **Single Voting Disk**



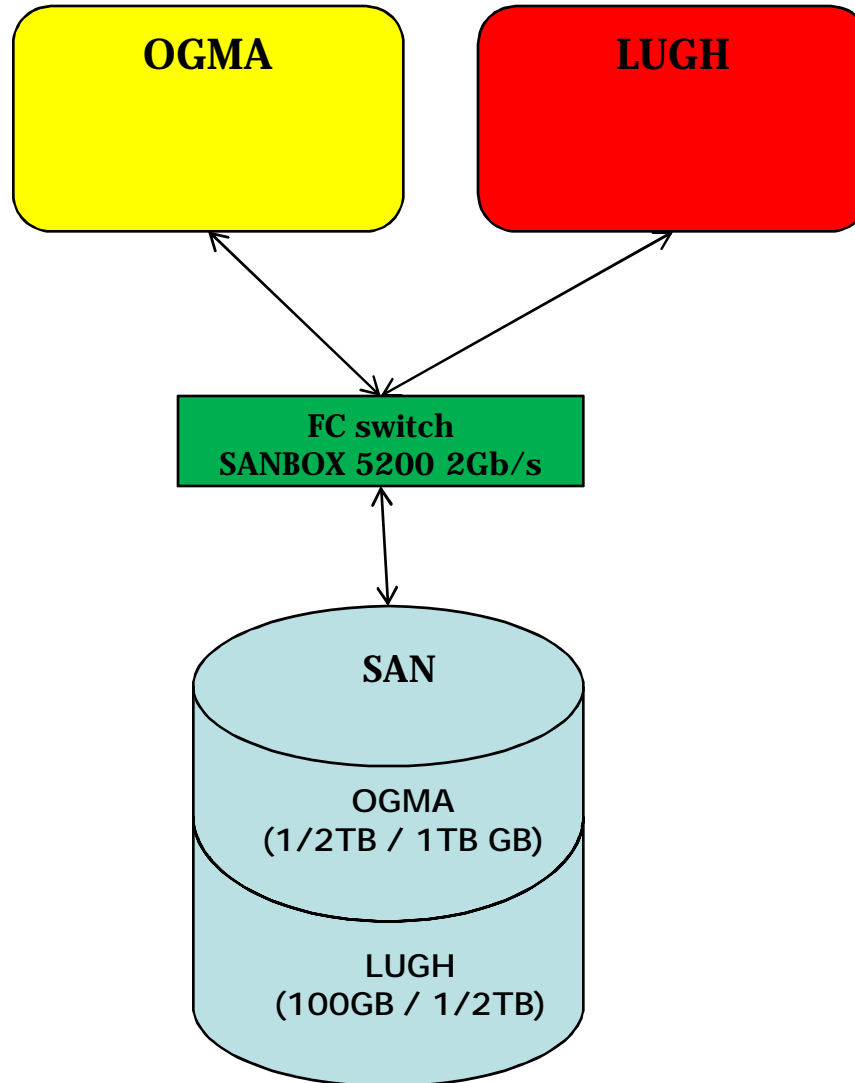
- **Single disk array shared by both databases (Ogma, Lugh):**
 - **Storage (SAN, 2GBps FC):**
 - **Ogma: ~1/2 TB**
 - **Pluto: ~100GB**
 - **Single switch SANBOX 5200 2Gb/s**
 - **16 disks SATA 260GB**
 - **Configured with RAID10**



- **ASM:**
 - **Ogma (ATLAS):**
 - Normal redundancy
 - Single disk group
 - Two failure groups
 - One disk (512G) per failure group
 - **Lugh (LHCb):**
 - Normal redundancy
 - Single disk group
 - Two failure groups
 - One disk (512G) per failure group



3D: Database diagram





- **There will be no changes on:**
 - Number of nodes per RAC
 - Hardware spes
 - Oracle version
- **Deploy on both RAC Ogma and Lugh:**
 - Two OCRs
 - Three Voting Disks



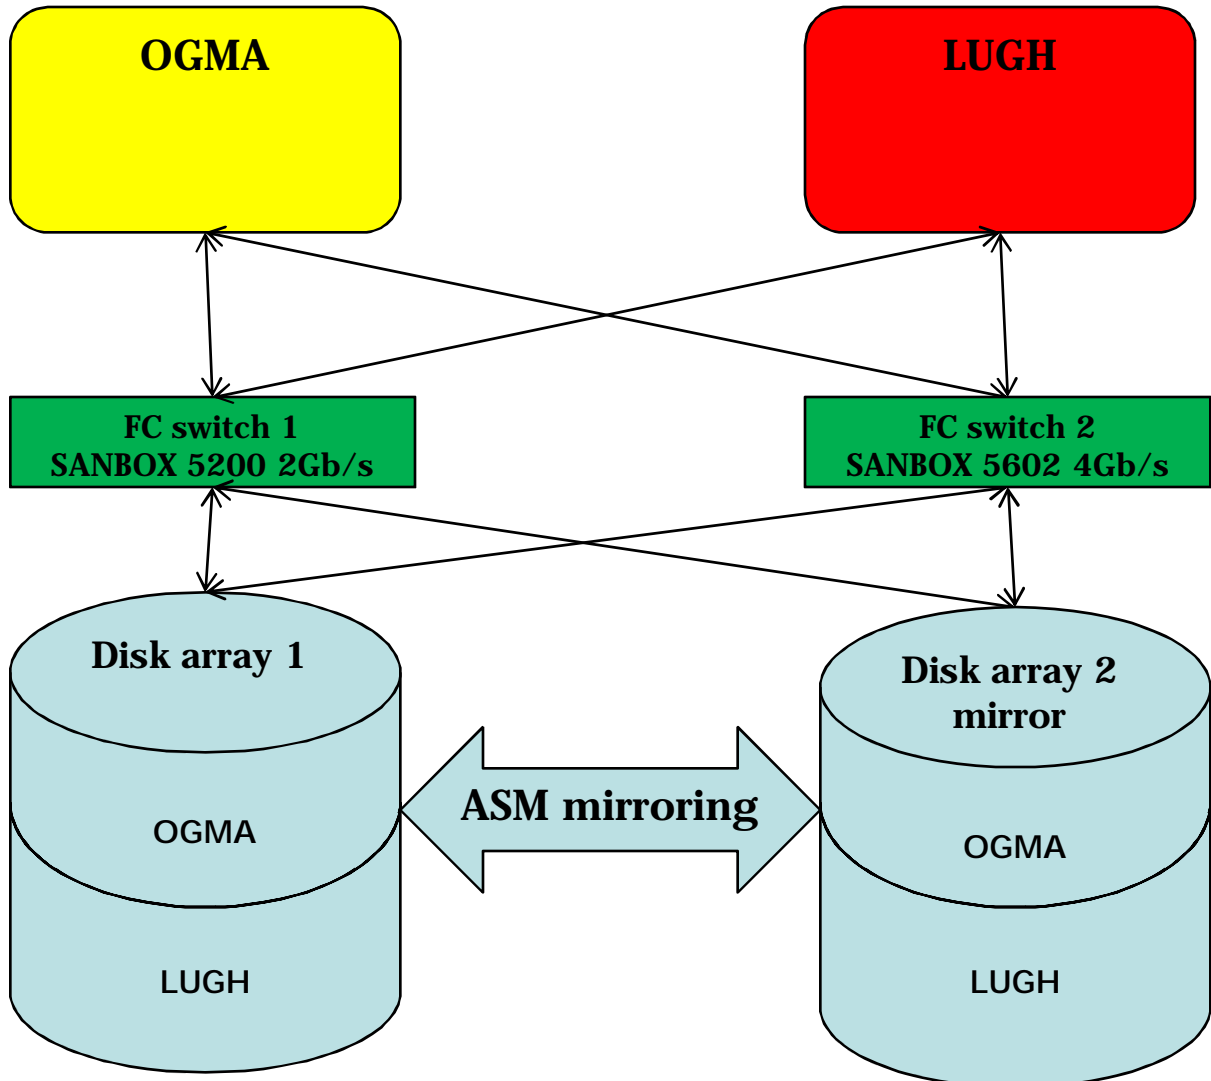
- **Two disk arrays shared by both databases (Ogma, Lugh):**
 - **Storage: SAN 4GBps FC**
 - **Physical disk available:**
 - **Array 1: 16 disks SATA 260GB**
 - **Array 2: 6 disks SATA 550GB**
 - **Arrays with RAID5 configuration**
- **Two switches:**
 - **SANBOX 5200 2Gb/s**
 - **SANBOX 5602 4Gb/s**



- **ASM:**
 - **Ogma (ATLAS):**
 - **Normal redundancy**
 - **Single disk group two failure groups**
 - **Two or more disks per failure group**
 - **Lugh(LHCb):**
 - **Normal redundancy**
 - **Single disk group two failure groups**
 - **One or more disks per failure group**



3D: Database Diagram





- **2 5- nodes RAC (Pluto, Neptune) + one single instance (Uranus)**
 - Red hat 4.8
 - 32 bit
 - Dual Quad Core (Intel Xeon 3Ghz)
 - 4 GB
- **Oracle 10.2.0.4**
- **Single OCR**
- **Single Voting Disk**



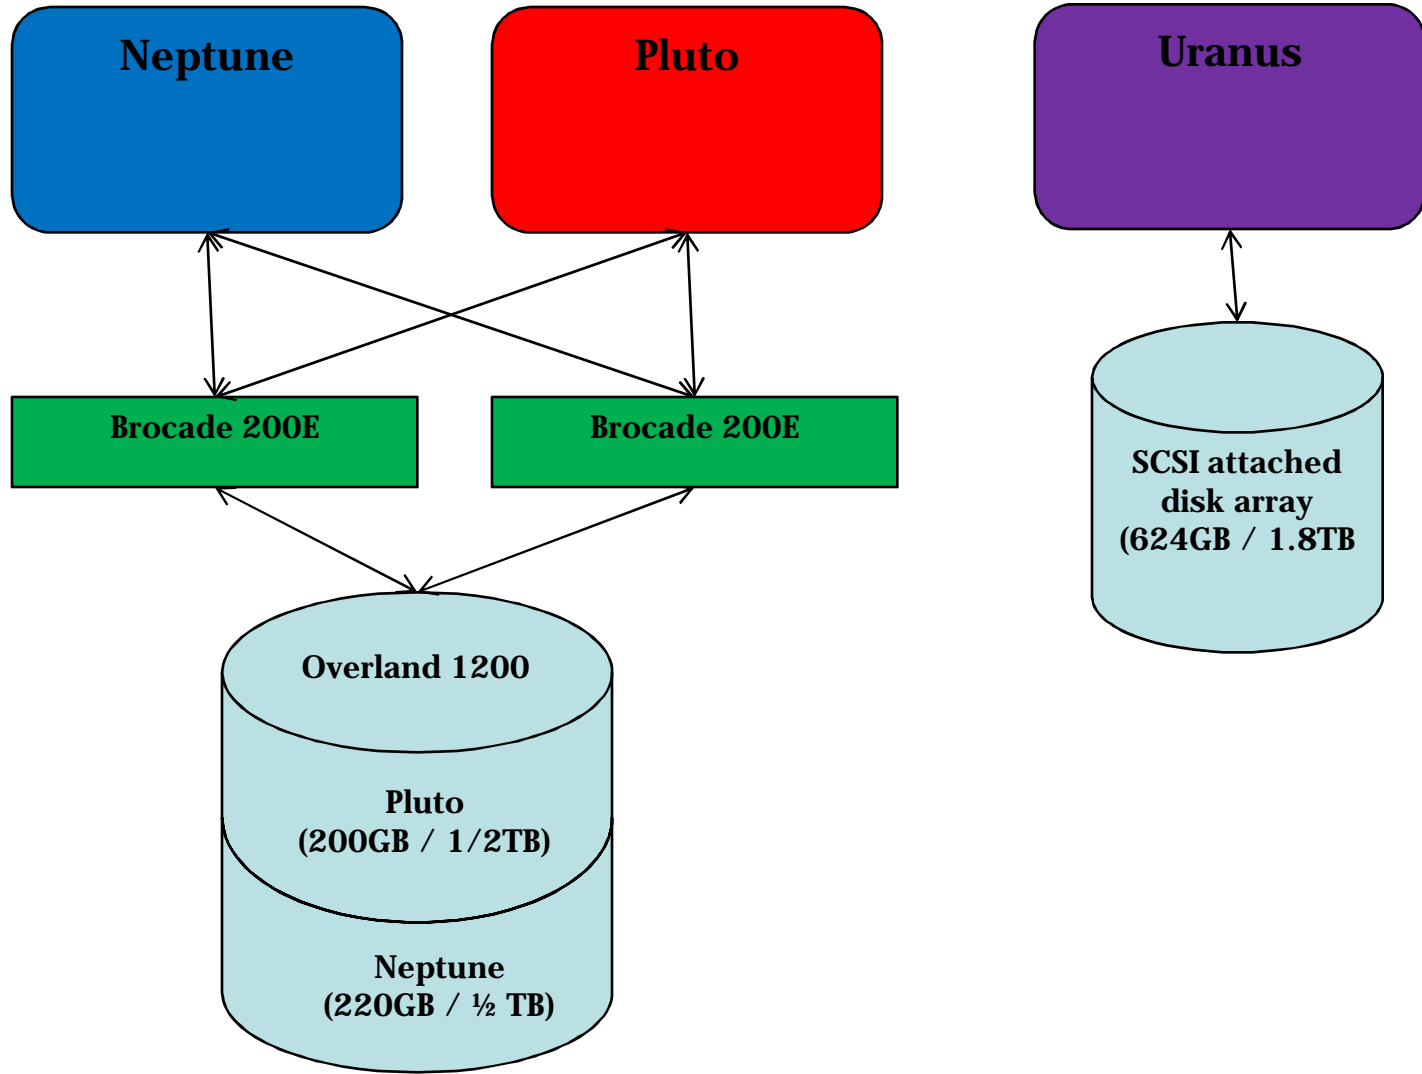
- **Single disk array used by the two RACs:**
- **Storage:**
 - Pluto:~200GB
 - Neptune:~220GB
 - Single instance: 624GB
- **Overland 1200 disk array**
 - Twin controller
 - Twin Fibre Channel ports to each controller
 - 10 SAS disk (300GB each 3TB total gross space)
 - Raid 1(1.5 TB net space)
- **Two Brocade 200E 4Gbit switched**



- **ASM (Pluto, Neptune):**
 - Normal redundancy
 - Single disk group
 - Two failure groups
 - One disk (512G) per failure group



Database Overview





- **There will be no changes on the number of node per RAC, the hardware or Oracle version**
- **Deploy on both RAC Pluto and Neptune:**
 - **Two OCRs**
 - **Three Voting Disks**



- **Two disk arrays shared by both databases (Neptune, Pluto):**
 - **Storage: EMC Clarion**
 - **Physical disk available:**
 - **SAS 300GB Drives**
 - **2TB gross**
 - **RAID5 configuration**
- **Two Brocade 200E 4Gbit switched**



- **ASM (Pluto, Neptune):**
 - **Normal redundancy**
 - **Single disk group two failure groups**
 - **One or more disks per failure group**

Pluto

Schemas	Version	Size
Name Server	n/a	1.8GB
VMGR	n/a	1.7MB
CUPV	n/a	0.2MB
CMS Stager	2_1_7_27_1	1.9GB
Gen Stager	2_1_7_27_1	3.8GB
Repack_219	2_1_9_1	17MB
Repack	2_1_7_27	62MB
Gen SRM	2_8_2	540MB
SRM CMS	2_8_2	1.1GB
VDQM2	2_1_8_3_1	5MB

Neptune

Schemas	Version	Size
Atlas Stager	2_1_7_27_1	18GB
LHCb stager	2_1_7_27_1	1.8GB
SRM Atlas	2_8_2	5.1GB
SRM LHCb	2_8_2	1.2GB



- **Incremental 0 once a week**
- **Incremental 1 the other days of the week**
- **All backups are followed by logical validation**
- **Archived log backup done during the day (for now)**
- **Once we move to the new hardware the archived log will be multiplexed on a shared disk outside ASM**
- **Backup are stored on the local disk**
- **Backup are copied from the local disk to tape and kept for three months**



- **RMAN configuration parameters are:**
 - **CONFIGURE RETENTION POLICY TO RECOVERY WINDOW OF 8 DAYS;**
 - **CONFIGURE BACKUP OPTIMIZATION ON;**
 - **CONFIGURE DEFAULT DEVICE TYPE TO DISK;**
 - **CONFIGURE CONTROLFILE AUTOBACKUP ON;**
 - **CONFIGURE CONTROLFILE AUTOBACKUP FORMAT FOR DEVICE TYPE DISK TO '/oracle_backup/pluto/%F.bak';**
 - **CONFIGURE DEVICE TYPE DISK PARALLELISM 2 BACKUP TYPE TO BACKUPSET;**
 - **CONFIGURE DATAFILE BACKUP COPIES FOR DEVICE TYPE DISK TO 1;**
 - **CONFIGURE ARCHIVELOG BACKUP COPIES FOR DEVICE TYPE DISK TO 1;**
 - **CONFIGURE CHANNEL DEVICE TYPE DISK FORMAT '/oracle_backup/pluto/pluto_%U.bak';**
 - **CONFIGURE MAXSETSIZE TO UNLIMITED;**
 - **CONFIGURE ENCRYPTION FOR DATABASE OFF;**
 - **CONFIGURE ENCRYPTION ALGORITHM 'AES128';**
 - **CONFIGURE ARCHIVELOG DELETION POLICY TO NONE;**
 - **CONFIGURE SNAPSHOT CONTROLFILE NAME TO '/opt/oracle/app/oracle/product/10/db_1/dbs/snapcf_pluto1.f'; # default**



- **Incremental 0:**
 - backup incremental level 0 duration 12:00 database;
 - backup archivelog all delete all input;
 - report obsolete;
 - delete noprompt obsolete;
- **Incremental 1:**
 - backup incremental level 1 duration 12:00 minimize time database;
 - backup archivelog all delete all input;
- **Validation:**
 - restore validate check logical database archivelog all;



**ANY
QUESTIONS?**