# ASGC Status report

ASGC/OPS
Jason Shih

Nov 26th 2009
Distributed Database Operations Workshop

# Outline

- Architecture and configurations
  - Database preparation and hardware configuration
  - Database service status
  - Hardware topology and CastorDB migration
- Applications
  - Grid services: CASTOR, LFC, FTS, SRM, 3D
- Future plan

# Database preparation

- DB Engine
  - Oracle 10g Rel. 10.2.0.3.0 and 10.2.0.4 mix
  - 11 nodes in total, another 3 obsolete and 1 for TB
  - 5 raid subsystem (15T x 2, 20T x 3)
    - RAID level 6 + HS (Infortrend), dual controller
    - 750T & 1T SATA2 drives.
  - Spare: N3600 x 2, 20T per enclosure
- Monitoring
  - Oracle Enterprise Manager
- Backup:
  - RMAM (disk only atm.)
  - Policy:
    - Incremental level=0 – *every Mon midnight*
    - Differential incremental level=1 - *every week day*
    - restore validate database – *every Sat.*
    - delete obsolete backups – every Sat.
    - Notification – *DB admin list.*
  - Retention policy
  - keep 1 full backups each week for 3 weeks

# Hardware Profile

- SAN Storage:
  - Fabric switch:
    - Brocade SAN switch E200, 3850
  - Raid Subsystem:
    - LFC/FTS
      - Silverstor TN-6224S-FFG RAID 6
      - 4TB for Data
      - 2TB for Backup
      - Free space that can be dynamically allocated: 4TB
    - 3D/CASTOR
      - 15TB, Enstor G2430
- Servers (old)
  - Quanta Blade System run EM64T
  - SMP Intel Xeon 3.0GHz
  - ECC 8GB Physical Memory
- New blade servers
  - 3D/CastorDB
  - IBM Blade server (BCE), DC 5130 2.0GHz
  - FB-DIMM 8GB

# Cluster - *LFC & FTS*

- SAN Storage:
  - Fabric switch:
    - Brocade SAN switch E200

Raid Subsystem:

- Silverstor TN-6224S-FFG RAID 6
- 4TB for Data
- 2TB for Backup
- Free space that can be dynamically allocated: 4TB

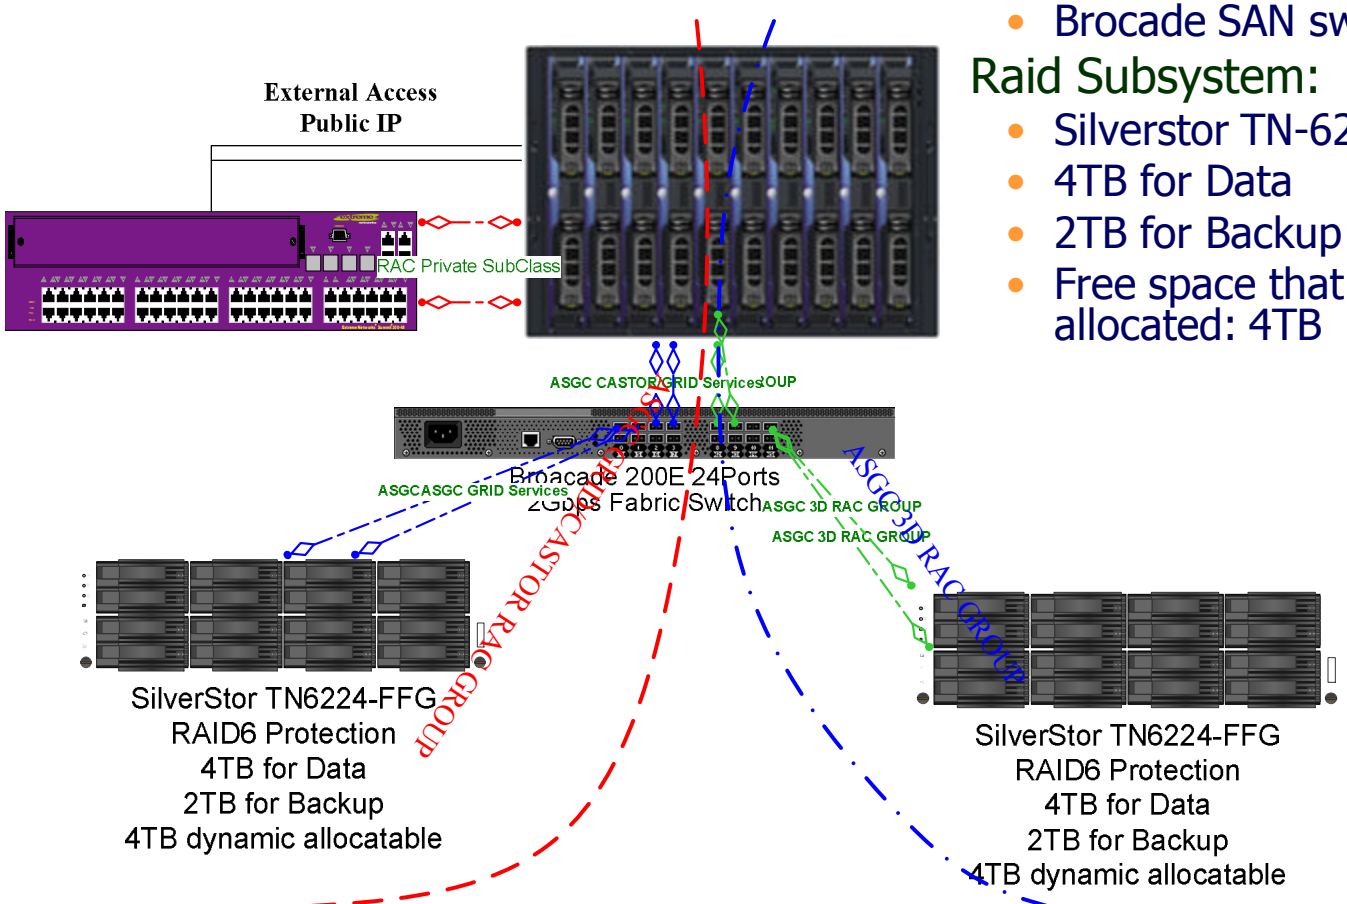

# DB services for WLCG

| DB service | Nr. Nodes | OS/Kernel | FS | 10g Rel. |
|---|---|---|---|---|
| LFC/FTS | 3 | Oracle Unbreakable Linux 4/ 2.6.9-42.0.0.0.1.ELsmp | OCFS | 10.2.0.3 |
| Atlas 3D | 2 | RHEL4u6/ 2.6.9-89.0.15.EL | ASM | 10.2.0.4 |
| CastorDB | 3/2 | RHEL4u6/ 2.6.9-78.0.1.EL | ASM | 10.2.0.4 |

Remove | Configure | Add

| Select | Name | Status | Alerts | Policy Violations | Compliance Score (%) | CPU Util % | Mem Util % | Total IO/sec |
|---|---|---|---|---|---|---|---|---|
| ◉ | obk.grid.sinica.edu.tw | ⇧ | 0 0 | 5 0 0 | 64 | .48 | 92.66 | 12.44 |
| ○ | oms.grid.sinica.edu.tw | ⇧ | 0 2 | 7 0 0 | 63 | 2.54 | 81.31 ⚠ | 36.63 |
| ○ | w-rac01.grid.sinica.edu.tw | ⇧ | 0 3 | 6 0 0 | 63 | 3.6 | 27.79 | 138.1 |
| ○ | w-rac02.grid.sinica.edu.tw | ⇧ | 0 7 | 6 0 0 | 63 | 4.22 | 54.5 | 142.91 |
| ○ | w-rac03.grid.sinica.edu.tw | ⇧ | 0 7 | 6 0 0 | 63 | 86.07 ⚠ | 77.41 | 1330.87 |
| ○ | w-rac04.grid.sinica.edu.tw | ⇧ | 0 6 | 6 0 0 | 63 | 14.43 | 86.79 | 180.45 |
| ○ | w-rac05.grid.sinica.edu.tw | ⇧ | 0 6 | 6 0 0 | 63 | 7.27 | 63.74 | 176.24 |
| ○ | w-rac06.grid.sinica.edu.tw | ⇧ | 0 1 | 6 0 0 | 63 | 3.71 | 96.83 | 138.09 |
| ○ | w-rac07.grid.sinica.edu.tw | ⇧ | 0 0 | 5 0 0 | 82 | | 60.68 | |
| ○ | w-rac08.grid.sinica.edu.tw | ⇧ | 0 0 | 5 0 0 | 82 | .47 | 59.59 | 22.57 |
| ○ | w-rac09.grid.sinica.edu.tw | ⇧ | 0 0 | 5 0 0 | 82 | 50.16 | 78.46 | 480.44 |
| ○ | w-rac10.grid.sinica.edu.tw | ⇧ | 0 0 | 5 0 0 | 82 | 2 | 76.31 | 104.91 |
| ○ | w-rac11.grid.sinica.edu.tw | ⇧ | 0 0 | 5 0 0 | 82 | .94 | 61.98 | 24.14 |

Remove | Configure | Add

# Castor DB migration

- File system
  - OCFS to ASM
- H/W configuration
  - Old: SMP Xeon 3.00GHz, 8GB per nodes
    - 3 nodes serving STG, NS, DLF, VMGR, SRM etc.
    - Quanta Blade server/Single raid subsystem (lsnr: 1521)
  - New: SMP Xeon DC 5130 2.0GHz, 8GB per nodes
    - 5 nodes in total. Two for NS, and 3 for SRM/STG/DLF
    - IBM blade server/two raid subsystem (lsnr: 15501)

# Monitoring

- Nagios probes:
  - Dummy login check for all RAC nodes
  - Oracle deadlocks (per 20min)
    - Alarm trigger if session lock > 10min.
  - Generic NRPE host plugins (CPU load, cache, swap)
- Grid Control
  - Castordb, srmdb, gdsdb

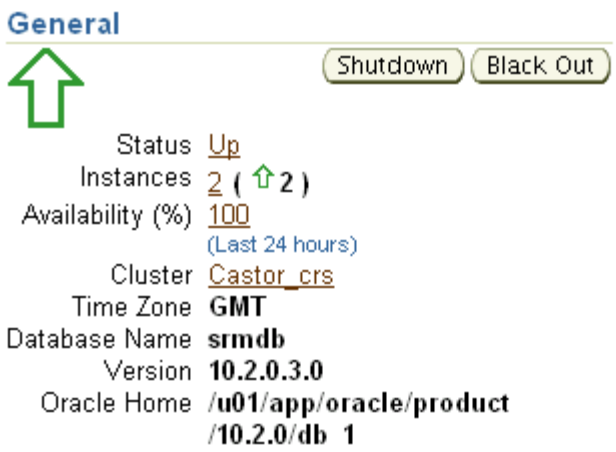


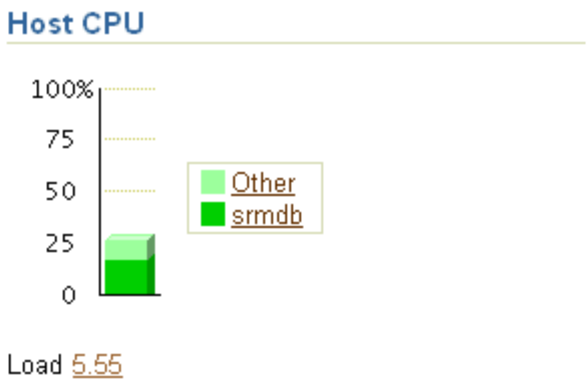


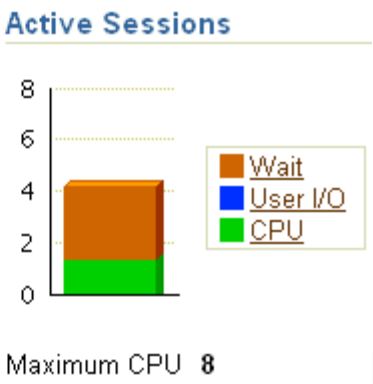

# Future planning

- Tape backup
  - TSM, considered LAN Free SAN backup
- Monitoring service
  - OMS migration
- LFC migration
  - File system and better profile server nodes
  - Est. end of 2009
- Enforcement of recovery exercise
- 24x7 best effort support
  - Type of alarm being handled and proper escalation

# ASGC incident report

ASGC/OPS
Jason Shih

Nov 26th 2009
Distributed Database Operations Workshop

# Outline

- Critical service incident
  - Atlas 3D
  - Castor DB corruption
- Procedures and event logs
- Plan

# Atlas 3D incident - I

- Sept 27th,
  - First observed by Eva ORA-12514 error code, and propagation cannot be re-started normally. verbose exception msg is 'TNS:listener does not currently know of service requested in connect descriptor'

Streams Monitor Error Report
Report date: 2009-09-27 02:21:49
Affected Site: CERN-PROD
Affected Database: ATLDSC.CERN.CH
Process Name: STRM_PROPAGATION_ASGC
Error Time: 27-09-2009 02:21:48
Error Message: ORA-12514: TNS:listener does not currently know of service requested in connect descriptor
ORA-06512: at "SYS.DBMS_AQADM_SYS", line 1087
ORA-06512: at "SYS.DBMS_AQADM_SYS", line 7642
ORA-06512: at "SYS.DBMS_AQADM", line 631
ORA-06512: at line 1

# Atlas 3D incident - II

- 28-Sep-09 07:17 - CERN-PROD : Process error report STRM_PROPAGATION_ASGC@ATLDSC.CERN.CH, service degradation of listener is suspicious.
- 28-Sep-09 13:40 - Eva confirm the same error code "ORA-12528: TNS: listener: all appropriate instances are blocking new connections‟
  - Instance down, err code 'ERROR at line 1: ORA-01507: database not mounted‘
  - 28-Sep-09 14:05 - all lsnr service are blocked, and confirm the db is not open properly except with the fs mounted.

# Atlas 3D incident - events

- 28-Sep-09 - the problem seems affect by the recent power surge and cause hardware degradation of two blade chassis, Consider performing point-in-time restore for the 3D database
  - propagation is disable now and we confirm LCRs activities already via stream monitor page
- 29-Sep-09 14:05, problem persist after point-in-time recovery. In fact, it is not possible to access to some old data so this makes impossible to try any kind of resynchronization using Streams with the current database status.
- 30-Sep-09 - perform point-in-time restore to Sept 21, where we confirm to have full backup
  - 30-Sep-09 13:41 - Eva confirm with some errors when try to access the data and this is why we suspect there is data corrupted (12 entries being selected with owner like 'ATLAS_COOLOFL_DCS' and object_name like '%_F0028_%', and the missing object observed: strmadmin@ASGC3D> select * from ATLAS_COOLOFL_DCS.COMP200_F0028_IOVS_SEQ;
  - ERROR at line 1:
  - ORA-08103: object no longer exists
- 5-Oct-09 13:41, the datafiles are nolonger valid (the point-in-time restored refer to Sept 21 while streams might not be able to recover from a 2 weeks backlog or not having archived log files on disk

# Atlas 3D incident - recovery

```
RMAN> run {
SET UNTIL SCN 6077727611290;
#SET UNTIL TIME "TO_DATE('Sep 21 2009 18:00:00','Mon DD YYYY HH24:MI:SS')";
restore database;
recover database;
}


RMAN-03002: failure of restore command at 09/29/2009 19:21:45
RMAN-20207: UNTIL TIME or RECOVERY WINDOW is before RESETLOGS time
```

```
channel ORA_DISK_1: restored backup piece 1
piece handle=+ATLAS_BACKUP/asgc3d/datafiles/uakpn18k_1_1
tag=TAG20090920T090003
channel ORA_DISK_1: restore complete, elapsed time: 00:00:46
```

```
Tue Sep 29 19:57:47 2009
alter database recover logfile
'+ATLAS_DATA/asgc3d/archivelog/2009_09_28/thread_2_seq_4571.879.698789657'
Tue Sep 29 19:57:47 2009
Media Recovery Log
+ATLAS_DATA/asgc3d/archivelog/2009_09_28/thread_2_seq_4571.879.698789657
Tue Sep 29 19:57:48 2009
Errors in file /u01/app/oracle/admin/asgc3d/bdump/asgc3d1_p000_11562.trc:
ORA-00600: internal error code, arguments: [4552], [2], [0], [], [], [],
[], []
Tue Sep 29 19:57:49 2009
Errors in file /u01/app/oracle/admin/asgc3d/bdump/asgc3d1_p000_11562.trc:
ORA-00308: cannot open archived log
'+ATLAS_DATA/asgc3d/onlinelog/group_15.291.677251137'
```

# Castor DB incident

- DB lock critical alarm
  - Date: Wed Oct 21 00:16:52 UTC 2009
  - Error: OracleLocks CRITICAL - CASTORDB2: 10 are waiting for lock out of the 146 sessions are connected
  - eLog: http://lists.grid.sinica.edu.tw/mailman/private/asgc-db/2009-October/002158.html
    - url: https://nagios.grid.sinica.edu.tw/asgc_ops/cgi-bin/status.cgi?host=w-rac02&servicestatustypes=28

# Castor DB incident - analysis

- Alert log in 2nd instance indicating the block corruption:

```
# tail /u01/app/oracle/admin/castordb/bdump/alert_castordb2.log
Wed Oct 21 10:00:38 2009
Recovery of Online Redo Log: Thread 2 Group 21 Seq 53451 Reading mem 0
  Mem# 0: /u02/oradata/castordb/redo2101.log
Block recovery completed at rba 53451.178442.16, scn 2.3721027505
Wed Oct 21 10:00:39 2009
Corrupt Block Found
        TSN = 10, TSNAME = STAGER_DATA
        RFN = 10, BLK = 1967881, RDBA = 43910921
        OBJN = 84545, OBJD = 84545, OBJECT = SUBREQUEST, SUBOBJECT =
P_STATUS_6
        SEGMENT OWNER = STAGER, SEGMENT TYPE = Table Partition
```

# Castor DB incident – recovery

- Oct 29th – Nov 3rd,
  - Hardware preparation
    - New blade servers, pass-through parts, fabric switch
    - Backend storage
    - 5 nodes in total, split into two clusters
  - Clean cluster preparation
    - collect any additional patches suggested by CASTOR DB team at CERN
  - CERN lead
    - install additional patches and the current security patch
    - start DB and check proper server functioning
    - initiate cloning of binaries to second ASGC cluster (SRM, Stager, DLF)
    - import recovered name server content into ASGC cluster 1
    - bring up cluster 2 and import the other recovered schemata

# Castor DB incident – service restart

- Nov 4th,
  - Recreate dlf schema which is less critical
  - confirm all core services up and running
  - Upgrading SRM to 2.8
    - SRM service resume Nov 5th
- Nov 6th,
  - OP revert the wrong srm config and cause the service degraded start from Fri.
  - Situation clarified:
    - Stager/VO mapping
    - Incorrect CNS host definition
    - Full function since Mon.
      - Exp. confirm with FT and production transfers

# Plan

- Weekly Tape backup
- Spare hardware for standby cluster
- Enforcement of recovery exercises

# Acknowledgment


- Atlas 3D resync
  - Carlos from BNL
  - Eva/Maria from CERN
- CastorDB recovery
  - Eric, Luca, Jacek, Dawid, Nilo, Giuseppe, Maria and Dirk.
  - Castor support team