



ASGC Status report

ASGC/OPS
Jason Shih

Nov 26th 2009
Distributed Database Operations Workshop



Outline

- Architecture and configurations
 - Hardware (server, storage, SAN)
 - Monitoring
 - Backup
- Applications
 - Grid services: CASTOR, LFC, FTS, SRM, 3D
- Performances & license statistics
- Remarks



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 (Infotrend), dual controller
 - 750T & 1T SATA2 drives.
 - Spare: N3600 x 2, 20T per enclosure
- Monitoring
 - Oracle Enterprise Manager
- Backup:
 - RMAN (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

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
<input checked="" type="radio"/>	obk_grid.sinica.edu.tw		0 0	5 0 0	64	.48	92.66	12.44
<input type="radio"/>	oms_grid.sinica.edu.tw		0 2	7 0 0	63	2.54	81.31	36.63
<input type="radio"/>	w-rac01_grid.sinica.edu.tw		0 3	6 0 0	63	3.6	27.79	138.1
<input type="radio"/>	w-rac02_grid.sinica.edu.tw		0 7	6 0 0	63	4.22	54.5	142.91
<input type="radio"/>	w-rac03_grid.sinica.edu.tw		0 7	6 0 0	63	86.07	77.41	1330.87
<input type="radio"/>	w-rac04_grid.sinica.edu.tw		0 6	6 0 0	63	14.43	86.79	180.45
<input type="radio"/>	w-rac05_grid.sinica.edu.tw		0 6	6 0 0	63	7.27	63.74	176.24
<input type="radio"/>	w-rac06_grid.sinica.edu.tw		0 1	6 0 0	63	3.71	96.83	138.09
<input type="radio"/>	w-rac07_grid.sinica.edu.tw		0 0	5 0 0	82		60.68	
<input type="radio"/>	w-rac08_grid.sinica.edu.tw		0 0	5 0 0	82	.47	59.59	22.57
<input type="radio"/>	w-rac09_grid.sinica.edu.tw		0 0	5 0 0	82	50.16	78.46	480.44
<input type="radio"/>	w-rac10_grid.sinica.edu.tw		0 0	5 0 0	82	2	76.31	104.91
<input type="radio"/>	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 (Isnr: 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 (Isnr: 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


General



Shutdown Black Out


Status Up
Instances 2 (↑ 2)
Availability (%) 100
(Last 24 hours)
Cluster Castor_crs
Time Zone GMT
Database Name srmdb
Version 10.2.0.3.0
Oracle Home /u01/app/oracle/product
/10.2.0/db_1

Host CPU



Load 5.55

Active Sessions



Maximum CPU 8



Hardware Profile

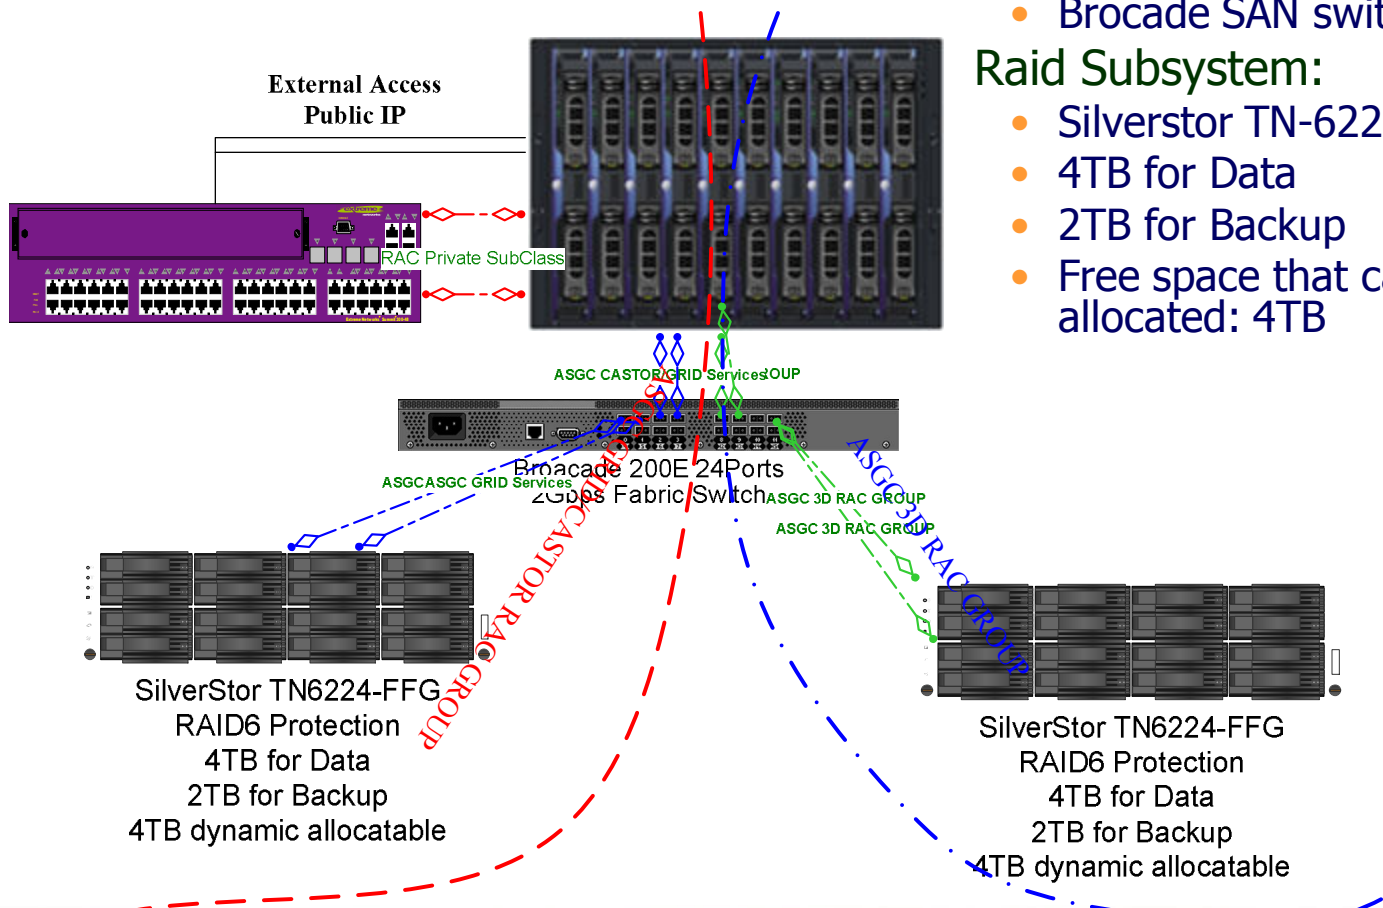
- SAN Storage:
 - Fabric switch:
 - Brocade SAN switch E200, 3850
 - Raid Subsystem:
 - LFC/FTS
 - Silverstor TN-6224S-FFG RAID 6
 - 4TB for Data
 - 2TB for Backup
 - Free space that can be dynamically allocated: 4TB
 - 3D/CASTOR
 - 15TB, Enstor G2430
- Servers (old)
 - Quanta Blade System run EM64T
 - 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





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