



# ASGC T1 and associated T2s

Jason Shih

WLCG Collaboration Workshop

Jan 24 2007



# Agenda

- ASGC T1
  - Computing Resource Planning
  - Castor Services
  - Data Transfer
  - High Availability
- APROC/Tier2s



# ASGC Resource Level Targets

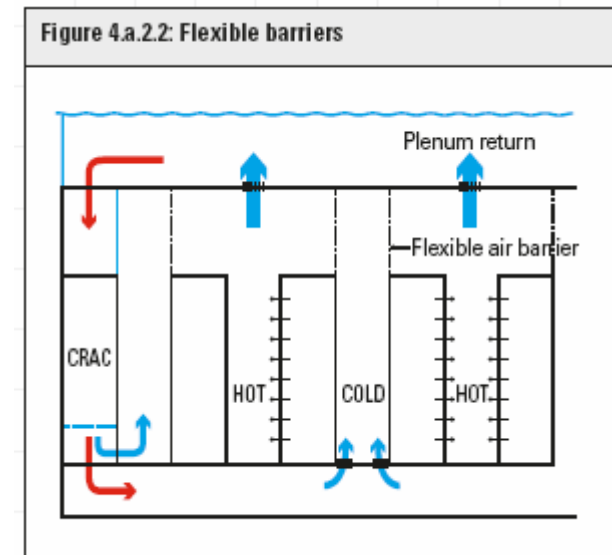
Date	CPU (ksi2k)	Disk (TB)	Tape (TB)
2006	950	400	500
2007	1770	900	800
2008	3400	1500	1300
2009	3600	2400	2000

- 2006 Target
  - Phase 1 complete in Oct., 189KSI2k added
  - Procurement of resources complete end of 2k6
  - Delivery and installation in early 2k7
    - 490kSI2k introduced, including also 360TB raid subsystem



# First thing: Facility Systems


- Second Data Center room being designed
  - Reliable Design
    - **N+1** for critical power and HVAC units
  - Separation of Power and Network
    - Power under raised floor and network in overhead tray
  - Prevent mixing and maximize HVAC efficiency
    - Hot and cold aisle configuration and air barriers
    - Cabling grommets to reduce cool air leakage
- Facility Estimates for 2009
  - CPU 12 racks, 156 kVA, 44 tons of cooling
  - Disk 22 racks, 198 kVA, 56 tons of cooling
  - Total power: 574 kVA
- Data Center Constraints
  - Limited space
  - Limited power
  - Limited cooling





# Hardware Selection I


- General selection considerations include
  - Density
  - Performance
  - Management capabilities
  - Power consumption
- CPU: Blades
  - High density: 70 server/rack
    - Compared to 1U 36 servers/rack
    - **Dual core** servers in recent procurement
  - Management features
    - Remote power management, KVM over IP
    - Reduced power, console and network cabling
  - Power supply efficiencies over **90%**
  - **More Costly?** – switch ports, KVM, RPM
- New expansion will use:
  - Intel Xeon 5130 Woodcrest Dual Core 2.0GHz 4M L2





# Hardware Selection II

- Disk: Direct attached fiber disk-array
  - High density:
    - When combined with blades for as disk servers
    - **750GB** disks or 18TB in 4U chassis
  - Better data protection
    - **RAID6** for protection against 2 disk failures
- Compared to NAS with built in servers comparison
  - Has no flexibility for server hardware
  - Higher price compared among local vendors





# Grid services

- CE split into:
  - NFS
  - Torque Server
  - Remove NAT (Public WN)
- Oracle RAC production service for Grid and Castor (in progress)
- LFC + MySQL
  - Loading issue resolved by client fix in DDM provided by Jean-Philippe
- VOBOX
  - Upgraded to dual core server to address heavy loading
- FTS upgraded to Oracle
- BDII with hot-standby
- 3D Deployment
  - 2 RAC servers
    - OCFS2 file system
    - SAN STK storage
    - Streaming tuning with CERN increased throughput from 4 to +40MB/s rate
  - Frontier server



# Agenda

- ASGC T1
  - Computing Resource Planning
  - Castor Services
  - Data Transfer
  - High Availability
- APROC Services





# H/W Profile (I) - Disk

- Stager/NS:
  - SMP Xeon 3.0, 4GB ECC memory (Fig 1)
- Disk servers:
  - SMP Xeon 3.0, 4GB ECC with x86\_64 (kernel 2.6.9-34.EL.cernsmp) (Fig 1)
- DAS to backend storage
  - max up to 10.5TB for single partition per chassis (Fig 2)
  - Raid 6, 24 bays in same closure, single controller only.
  - Total Capacity 45TB
- Sun/STK raid system shared
  - specific zoning for CASTOR group (*redundant controller*)
  - FLX-B280 – 20TB and FLX-B380 – 30TB (Fig 3)

Fig 3



Fig 1




Fig 2





# H/W Profile (II) – Tape Lib

- IBM 3584 (Fig 1)
  - Base frameset L52 and D52 expansion
  - 700 LTO3 cartridges, 400GB ea.
  - Total capacity 280TB
- Tape drives (Fig 2)
  - 4 x IBM 3592 LTO3 tape drives
  - 80MB/s max throughput
  - Observed performance 20~30MB/s
- Tape Servers (Fig 3)
  - IBM 306m, Pentium dual core, 4GB ECC
  - 3 production use
  - 1 for Castor testbed

Fig 1



Fig 2



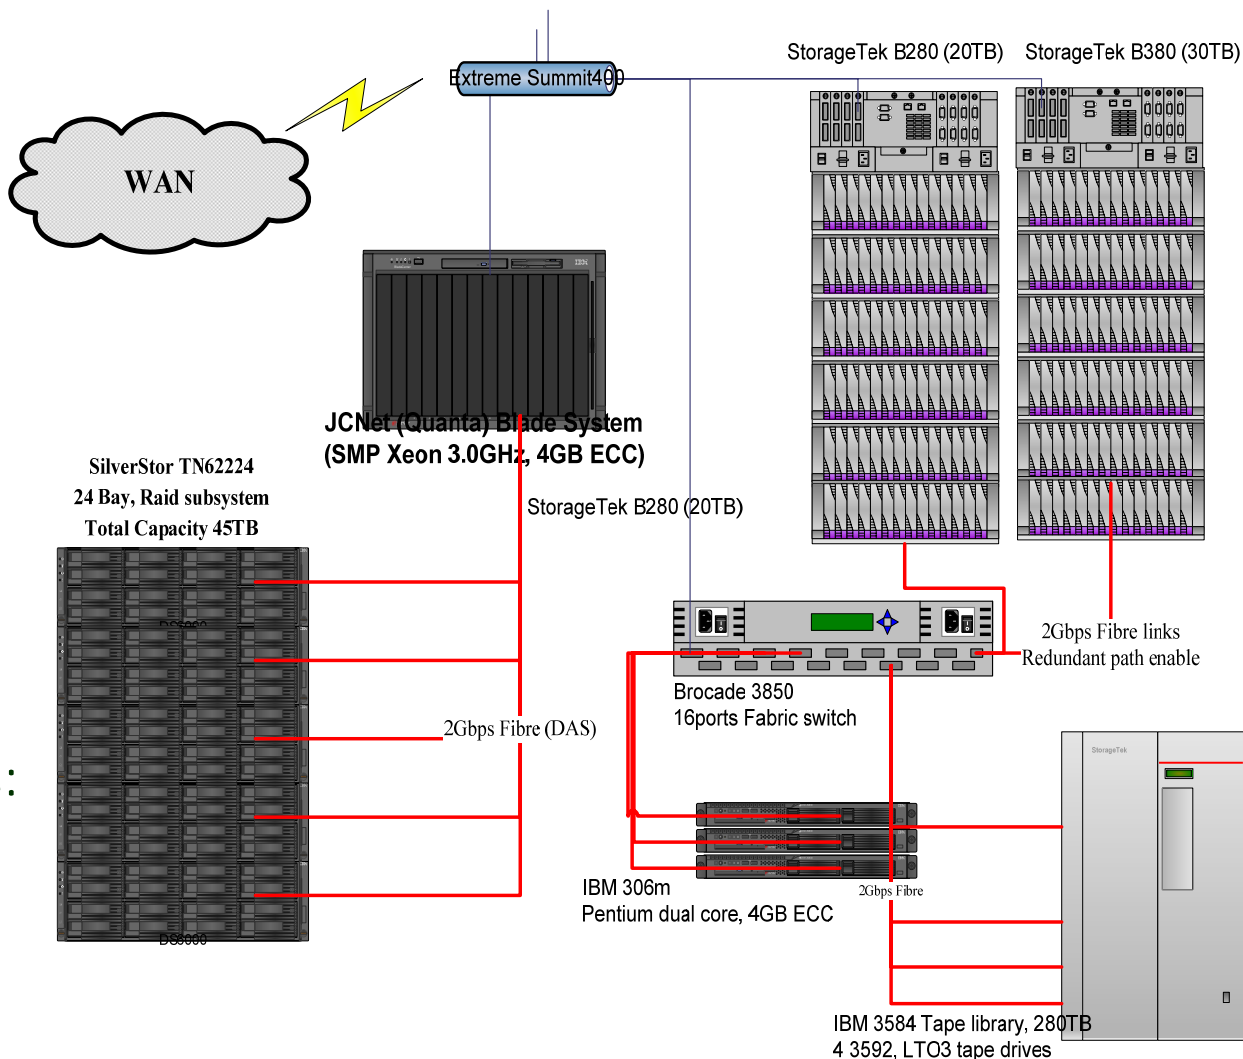
Fig 3





# Current Architecture

- Stager/NS – castor v1
  - MySQL backend
- DNS RR SRM headnodes
- Backend disk servers
  - castor v2 (griftp and rfio)
- Tape servers
  - Base on Castor v2
  - V1 VDQM use direct communication with rtcg bypassing v2 components:
    - Migrator
    - MigHunter
    - Err Handlers





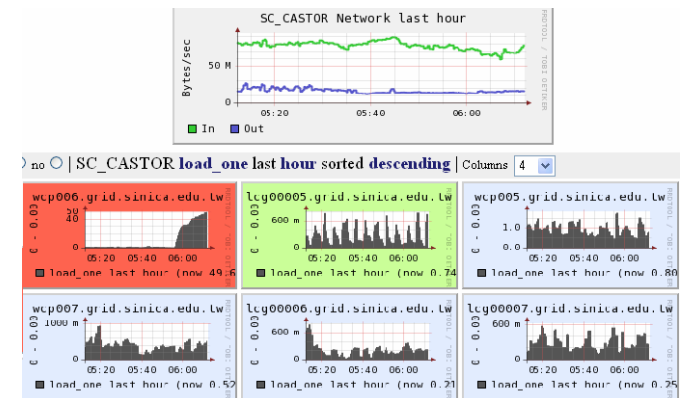
# Castor: Addressed Issues

- Upgraded server hardware to SMP Xeon
  - Stager/NS split from SRM
  - Smoother and higher transfer rates
- Disk server stack overflow causing instabilities
  - Increasing stack size from 4 to 8k but only reduced frequency of issues
  - Migrated disk servers to **EMT64** to fully resolve this issue
- STK disk arrays observed high IO wait and low throughput
  - New single controller arrays used to migrate data from STK
  - Recently isolated this problem to the small **RAID segment size** of 64k
    - Increasing to 256k was sufficient to obtain over 100MB/s write performance with low IO Wait





# Remaining Issues

- High disk server loading when **SRM pushmode** in Castor SRMv1
  - Expected to be resolved with Castor SRMv2
- Scalability limitations of Castor v1
  - Stager catalogue limited to 100k entries
  - Lack of **scheduling** and resource sharing capabilities
- Migration from v1 to v2
  - Tape integration testing with v2
  - Pure V2 implementation – stager/name server
  - Oracle RAC deployment for NS,VDQM and Stager
  - Multiple instances (stager) serving different VO
  - Migration of disk servers into new production v2 env
  - Performance and stress testing
- Castor is quite complex and has a steep learning curve
  - Many thanks to Castor team for their support!
  - Need **dedicated** support for deployment and operations issues






# Accounting ASGC C1 MSS



- disk to disk nominal rate
  - currently ASGC have reach 100+ MB/s static throughput
  - Round robin SRM headnodes associate with 11 disk servers, each provide ~8-12 MB/s constantly
  - debugging kernel/castor s/w issues early stage of SC4 (reduction to 25% only, w/o further tuning)





# Agenda

- ASGC T1
  - Computing Resource Planning
  - Castor Services
  - Data Transfer
  - High Availability
- APROC Services



# Nominal Tier0 – Tier1 Data Rates


<i>Tier1 Centre</i>	<i>ALICE</i>	<i>ATLAS</i>	<i>CMS</i>	<i>LHCb</i>	<i>Target</i>
IN2P3, Lyon	9%	13%	10%	27%	200
GridKA, Germany	20%	10%	8%	10%	200
CNAF, Italy	7%	7%	13%	11%	200
FNAL, USA	-	-	28%	-	200
BNL, USA	-	22%	-	-	200
RAL, UK	-	7%	3%	15%	150
NIKHEF, NL	(3%)	13%	-	23%	150
<b>ASGC, Taipei</b>	-	<b>8%</b>	<b>10%</b>	-	<b>100</b>
PIC, Spain	-	4% (5)	6% (5)	6.5%	100
Nordic Data Grid Facility	-	6%	-	-	50
TRIUMF, Canada	-	4%	-	-	50
<b>TOTAL</b>					<b>1600MB/s</b>





# ASGC Transfer Rate: Oct 06

CMS cleaned 20TB data on Oct 29th. All eleven servers had free space. Then, the rate ramped up to 110Mb/s.






**We are not running out of disk space, but disk servers !**

**ASGC has 11 servers now. 3 servers were added in the middle of CSA06.**

**Server capacity : 10T x3, 7.5T x2, 5T x1, 3T x2, 2T x3**



# CMS: Inbound Rate T0-T1




CMS can hit 100MB/s when Atlas was not transferring

Contention issue may be related current Castor architecture - (single stager)



# ASGC T1 and TW-FTT T2 data transfers



Tier1 inbound transfer:  
230TB data transferred


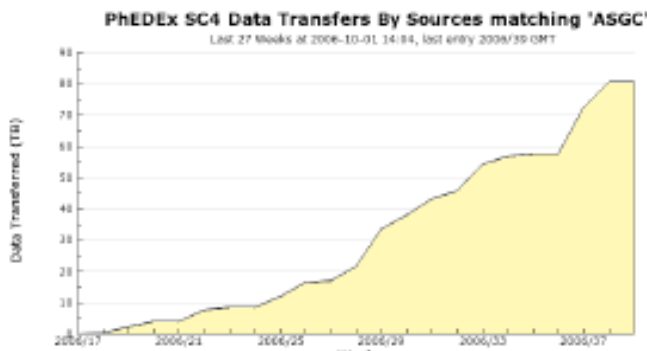


Tier1 outbound transfer:  
80TB data transferred


PhEDEx SC4 Data Transfers By Destinations matching 'Taiwan'



(up) Tier2 inbound transfer: 45TB data transferred




PhEDEx Prod Data Transfers By Sources matching 'ASGC'



(down) Tier2 outbound transfer:  
12TB data transferred



# Transfer Quality in SA06





# FTS T2 Testing I

- Purpose
  - Test and exercise FTS service at ASGC
  - Test T1-T2 network connectivity
  - Test functionality of T2 SRM endpoints and SEs
  - Troubleshoot and optimize T1-T2 transfer rates
- Three phase tested
  - Functionality test
    - 1MB file transferred to each SRM endpoint
  - Performance test
    - 100 1GB files transferred
  - Stability test
    - Test sustained transfer rates for over 10 hours (1000's 1GB files)
    - Sites tested individually
    - Tested with FTS with Oracle backend



# FTS T2 Testing II

- **Functional test:** sites tested
  - FTT, IPAS, NCUHEP, NIUCC, THUHEP
  - AU, KEK, KNU, Beijing, Tokyo-LCG2

- **Performance test:** Average throughput MB/s

	<b>AU</b>	<b>Beijing</b>	<b>IPAS</b>	<b>KEK</b>	<b>KNU</b>	<b>NIU</b>	<b>Tokyo</b>
<b>Rate</b>	3.2	16.3	36.9	9.8	36.7	4.3	40.2

- **Stability tests:** Average throughput MB/s

	<b>AU</b>	<b>IPAS</b>	<b>KNU</b>	<b>Tokyo</b>
<b>Rate</b>	15.2	72.0	28.4	47.7

- T2 -> T1 Testing:
  - KNU 31.8 MB/s
  - IPAS: 82.6 MB/s

Increase # of files with reduced # streams



# FTS T2 Testing Plans

- T1-> T2 Stability test with remaining T2s
  - 20MB/s target for each site
- T2 -> T1 Stability test
  - 10MB/s target for each site
- T1->T1 testing in progress
- Continuous background transfer
  - After Castor V2 is deployed in production
  - Dedicated stager for dteam
  - Enable site admins to control FTS channel
  - Distribute cronjob to clean files



# Agenda

- ASGC T1
  - Computing Resource Planning
  - Castor Services
  - Data Transfer
  - High Availability
- APROC Services





# Tier1 MoU Availability Targets

3.6 days/year

Service	Maximum delay in responding to operational problems			Average availability measured on an annual basis	
	Service interruption	Degradation of the capacity of the service by more than 50%	Degradation of the capacity of the service by more than 20%	During accelerator operation	At all other times
Acceptance of data from the Tier-0 Centre during accelerator operation	12 hours	12 hours	24 hours	99%	n/a
Networking service to the Tier-0 Centre during accelerator operation	12 hours	24 hours	48 hours	98%	n/a
Data-intensive analysis services, including networking to Tier-0, Tier-1 Centres outside accelerator operation	24 hours	48 hours	48 hours	n/a	98%
All other services – prime service hours <sup>[1]</sup>	2 hour	2 hour	4 hours	98%	98%
All other services – outside prime service hours	24 hours	48 hours	48 hours	97%	97%

<sup>[1]</sup> Prime service hours for Tier1 Centres: 08:00-18:00 in the time zone of the Tier1 Centre during the working week of the centre, except public holidays and other scheduled centre closures.



# 24x7: Current Coverage

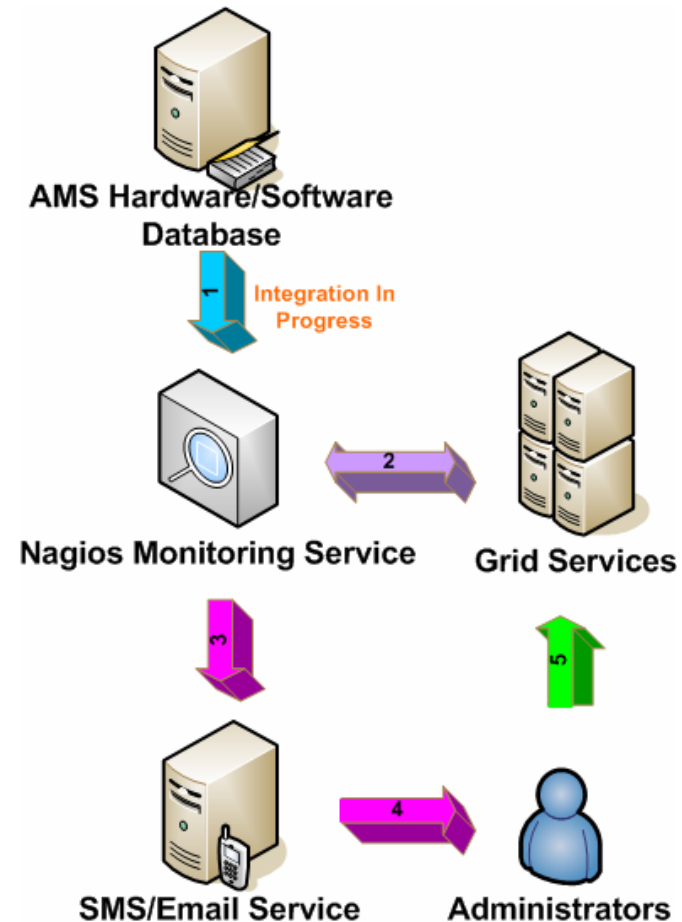


- Grid Administrators
  - Handling of complex Grid services
  - Best effort basis on nights and weekends
- Network Administrators
  - Handling of network faults, coordination of recovery, notification of affected parties
  - On-site 24x7: network administrators
- Facility Operators
  - Handling of power, cooling, security and environment issues
  - On-site 24x7: technicians
- 2007 Plans: Extend Grid Service coverage
  - Start rotation for 24x7 on-call Grid administrator coverage
  - Hire operators to extend on-site coverage to 16x7



# Monitoring and Notification

- Planned integration of AMS DB
- Nagios plugins developed
  - CE
  - LFC
  - VOMS
  - Storage
  - IT services
  - OS
- Notification via Email
  - SMS transmission device currently being tested





# Future Plans

- Increase monitoring coverage to
  - Experiment services hosted at T1
  - Worker Node environment testing
  - IS, Castor, FTS and Oracle plugins
- Test and deploy HA solutions for critical Grid components.



# Availability Results

- Based on SAM tests on CE, SE and SRM services
- Availability from May to Sept: **89%**
  - One of two sites to reach 88% target
  - Still much more effort needed to reach 99%

	CERN-PROD	FZK-LCG2	IN2P3-CC	INFN-T1	RAL-LCG2	SARA-MATRIX	TRIUMF-LCG2	Taiwan-LCG2	US-CMS	PIC	BNL	Avg all site
May	89%	85%	83%	89%	68%	58%	77%	87%	68%	61%	n/a	77%
Jun	92%	15%	89%	62%	76%	49%	88%	75%	64%	88%	n/a	70%
Jul	90%	54%	87%	31%	73%	84%	80%	98%	20%	87%	n/a	70%
Aug	95%	69%	94%	69%	59%	83%	87%	97%	4%	88%	n/a	74%
Sep	89%	57%	83%	88%	69%	92%	66%	85%	26%	90%	n/a	74%
Oct	96%	54%	85%	85%	77%	74%	80%	89%	55%		n/a	76%
Nov	90%	85%	70%	94%	89%	84%	87%	95%	77%	79%	56%	82%
Dev	93%	63%	37%	84%	89%	82%	91%	95%	79%	90%	70%	79%
Avg.	92%	61%	78%	75%	75%	76%	82%	90%	49%	83%	63%	75%



SD: LAN Refactoring and

Service software/hardware upgrade



# ASGC T1 Accounting (2k6 - 2k7)

Normalised CPU time [units 1K.SI2K.Hours] by TIER1 and VO																
TIER1	alice	atlas	biomed	cms	dteam	egeode	esr	fusion	geant4	lhcb	magic	ops	planck	Other VOs	Total	%
AsiaPacific	0	448,502	56,591	464,715	29	0	0	0	0	0	0	16	0	17,500	987,353	3.98%
BNL	0	4,130,425	0	0	0	0	0	0	0	0	0	0	0	0	4,130,425	16.65%
CERN	371,898	766,193	0	689,122	469	0	0	0	6,445	990,168	0	27	0	6,186	2,830,508	11.41%
FNAL	0	0	0	1,206,806	0	0	0	0	0	0	0	0	0	0	1,206,806	4.86%
FZK	257,662	391,340	0	427,820	475	0	0	0	0	471,660	561	18	0	33,295	1,582,831	6.38%
IN2P3	250,150	1,011,389	210,442	149,714	154	0	905	0	24	115,680	0	0	106	321,604	2,060,168	8.30%
INFN-T1	291,876	877,263	0	368,020	36	0	0	0	0	1,301,439	0	39	0	827,871	3,666,544	14.78%
PIC	0	660,517	68,809	309,019	65	0	0	0	0	448,556	0	15	0	636	1,487,617	6.00%
RAL	101,552	1,038,255	0	842,369	84	0	991	1,849	1,621	1,566,347	0	0	0	638,633	4,191,701	16.90%
SARA/NIKHEF	192,364	839,489	266,843	19,683	466	0	687	0	573	700,233	0	0	0	384,036	2,404,374	9.69%
TRIUMF	0	258,128	0	190	30	0	0	0	0	517	0	14	0	0	258,879	1.04%
<b>Total</b>	<b>1,465,502</b>	<b>10,421,501</b>	<b>602,685</b>	<b>4,477,458</b>	<b>1,808</b>	<b>0</b>	<b>2,583</b>	<b>1,849</b>	<b>8,663</b>	<b>5,594,600</b>	<b>561</b>	<b>129</b>	<b>106</b>	<b>2,229,761</b>	<b>24,807,206</b>	
<b>Percentage</b>	<b>5.91%</b>	<b>42.01%</b>	<b>2.43%</b>	<b>18.05%</b>	<b>0.01%</b>	<b>0.00%</b>	<b>0.01%</b>	<b>0.01%</b>	<b>0.03%</b>	<b>22.55%</b>	<b>0.00%</b>	<b>0.00%</b>	<b>0.00%</b>	<b>8.99%</b>		





# T1 staffing in 2009

Scaling assumed	0.3	doubling of hardware would require 30% more staff
<b>Tier-1 Operations</b>		
CPU Service	2.0	Hardware infrastructure for operations, installation, monitoring, and scheduling of all of the CPU resources. Staff will be responsible for the security, integrity and performance of the operating system. (see incident response below)
Data Services: Disk	1.0	All data storage components including those required to ensure reliable, robust, well-managed and guarantee the integrity of the data. Both Disk and Tape services are included plus any HSM capability needed to manage the temporary staging of data between tape and disk. Support for Oracle and similar databases is included in this (at RAL)
Data Services: Tape	1.0	
Data Services: Database	1.0	
Core Services Operations	1.0	Consoles, user file systems, monitoring, software development/deployment and conditions databases.
Operations	4.0	Machine room environment (safety, power, cooling, deliveries, installation), hardware diagnostics and repair, automation and monitoring, deploying Grid fabric management tools, tape movement and robot intervention.
Networking	1.5	Administering the centre's LAN and a share of local WAN access.
Deployment	2.0	The Tier-1 Centre will play a leading role in UK deployment, managed as part of the Grid Deployment Team to provide support for the whole of GridPP.
Experiment Interface	1.5	This is outward-looking Tier-1 effort that is intended to complement dedicated experiment software support.
Management	1.5	Includes both internal management of the Centre and its team and finances, and the outward-facing interactions with LCG, EGEE, and the other Tier centres.
Other	1.0	Any remaining effort not conveniently allocated to another line. E.g. at RAL this is an "Incident Response Unit" which will be developed to meet the MoU service requirements. This constitutes the extra staffing required to enable flexible time working will enable system staff cover from 0800-1800 Mon-Fri excluding UK public holidays.
total	17.5	



# Milestones of ASGC, 2k6

- **Jan-Feb**
  - FS of disk servers reformatted, from ext3 to xfs (performance)
  - Castor v2 gridftp pkg adoped
- **Feb-Mar – kernel upgrade from 2.4 to 2.6 (tcp buffer size, bdflush issue).**
- **May – CASTOR-SC relocated, disk server migrated into dual core servers (IPMI).**
- **Jul –**
  - migrate FTS backend DB to Oracle (continuous dead lock error with MySQL)
- **Aug –**
  - Castor2 testbed implemented
  - Migrate 3D service to new dual core servers (remaining h/w error cause bad performance of db streaming)
- **Sep –**
  - Disk server migrate from dual core server to SMP Xeon blade system, reinstalling OS with x86\_64 for the long exist XFS kernel bug
  - Stager/NS split from SRM headnodes, transfer quality improved
  - 11 disk servers installed for CSA06 (total capacity 65TB)
- **Oct –**
  - nagios plug-in to help monitoring disk servers status and service availability (rfio and gridftp)
  - Installation of 60TB raid subsystem in C1 fabrics
  - 288KSI2k blade system install (old blade system offline for MA)
- **Procurement of CPU/DISK complete end of Dec (2k6 pledged nr)**
  - Delivered earlier of Jan, plan to have new resources online – mid of Feb
- **Dec - Prod C2 design delivered, review by FIO experts**
- **Jan 2k7 –**
  - Prod C2 implementation (plan to be complete mid of Feb)
  - 3D RAC cluster migrate from dual core to SMP blades system (remote pwr mgmt: IPMI fails sometime if kernel panic), considering also the migration of backend raid subsystem (STK B280 to SilverStor)





# Agenda

- ASGC T1
  - Computing Resource Planning
  - Castor Services
  - Data Transfer
  - High Availability
- APROC Services



# APROC Introduction

- APROC Mission
  - Provide deployment support facilitating Grid **expansion**
  - Maximize the **availability** of Grid services
- Supports EGEE sites in **Asia Pacific** since April 2005
  - 18 production sites
  - 8 sites deployed and certified this year
  - 2 new sites, KISTI, and HKU, pass site certification earlier of 2007
  - 1 sites in certification process
    - Mongolia: (MAS IPT) Mongolian Academy of Sciences
- Operates ASGCCA that provide certificates for AP EGEE/LCG sites without domestic CA.
- EGEE Operations
  - CIC-on-duty: EGEE global operations
  - Monitoring tool development: GStat and GGUS Search
  - VO services: VOMS and LFC
  - TPM: Front line user support (Q4 2006)
  - OSCT: Incident Response duty (Dec 2006)



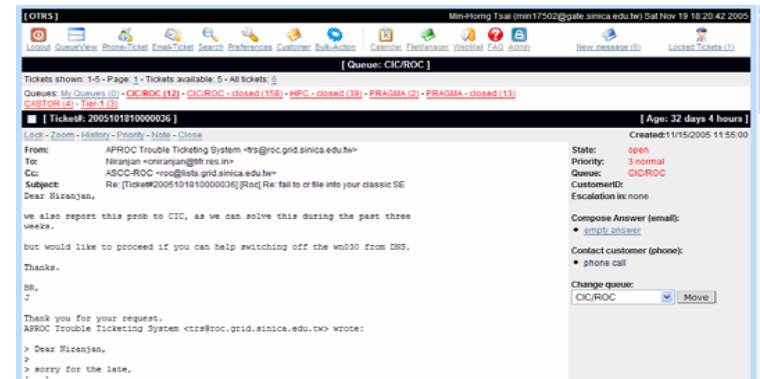
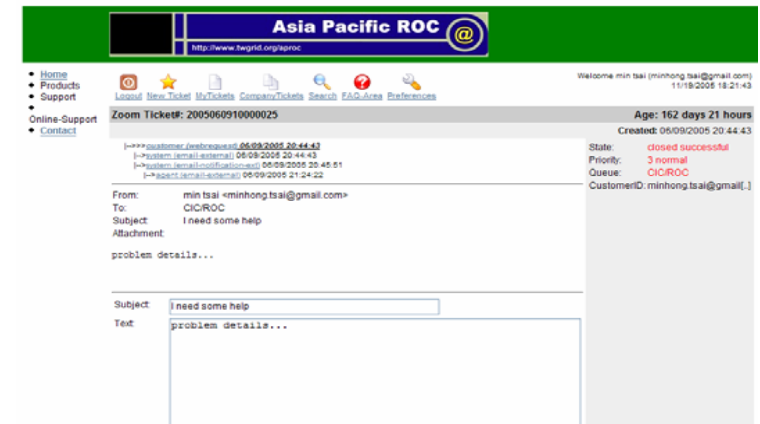
# Site Deployment Services

- Deployment consulting
  - Directing to important references
    - Tutorial DVDs (Chinese)
  - Site architecture
  - Hardware requirements
- Middleware installation support
  - Configuration
  - Troubleshooting
- Site certification
  - Functionality testing
  - Official EGEE infrastructure registration



# Operations Support Services

- Operations Support
  - Monitoring
  - Diagnosis and troubleshooting
  - Problem tracking via OTRS ticketing system
- M/W release deployment support
  - Pre-Production site operations
  - Certification testbed
  - Supplementary release notes
- Security Coordination
  - Security release announcement, instructions and follow-up
- Documentation: APROC Portal and wiki
  - <http://www.twgrid.org/aproc>
  - <http://list.grid.sinica.edu.tw/apwiki>
    - **Troubleshooting Guides (New)**
- Site communication and support channels
  - Phone, Email, OTRS Ticketing System
  - Monthly meeting with AsiaPacific sites over VRVS





# Application Startup

- Initial startup: APESCI VO
  - Provided for new communities to test and develop Grid applications
  - Acts as incubator VO for fast access to Grid resources
  - Centralized services already running
    - Resource Broker, LFC and VOMS services
- Next step: Production VO
  - Discuss with NA4 to join existing VO and collaborate
  - Create a new VO
    - APROC can also help host LFC and VOMS for the new VO



# Regional Monitoring Tools

- SAM for Grid monitoring
- GStat for IS monitoring
- Smokeping for network quality

The screenshot shows a monitoring interface with a 'Tests Displayed' list on the left, a central table of test results, and a table of site status at the bottom.

show	stat	description	sum
<input type="checkbox"/>	NA	no status available	0
<input checked="" type="checkbox"/>	OK	normal status	13
<input checked="" type="checkbox"/>	INFO	useful information	0
<input checked="" type="checkbox"/>	NOTE	important information	0
<input checked="" type="checkbox"/>	WARN	subject may fail soon	0
<input checked="" type="checkbox"/>	ERROR	subject has failed and problem is localized	4
<input checked="" type="checkbox"/>	CRIT	subject has failed and problem is fatal	0
<input checked="" type="checkbox"/>	MAINT	subject is under maintenance	0

No	SiteName	NodeName	Status	cpu	mem	disk	net
AsiaPacific							
1	Australia-UNIMELB-LCG2	lcr-compute.lpc.unimelb.edu.au	ERROR	error	3.0.2	ok	ok
2	GOG-Singapore	melon.naps.nsl.or.sg	OK	ok	3.0.2	ok	ok
3	IN-DAE-VECC-01	gridce01.tier2.kol.res.in	OK	ok	3.0.2	ok	ok
4	JP-KEK-CRC-01	dg10.cc.kek.jp	OK	ok	3.0.2	ok	ok

GStat: 01:51:21 11/30/06 GMT

home alert table service regional service metrics links grid eps test baltic eds exchima esmed secmid gbls nimid

AsiaPacific CERN CentralEurope France Germany/Switzerland Italy NorthernEurope Russia SouthEasternEurope SouthWesternEurope UKI

INDIACMS-TIER	PAKGRID-LCG2	JP-KEK-CRC-01	NCP-LCG2	TW-NTCU-HPC-01	TW-THU-HPC
Taiwan-NCUCC-LCG2	Australia-UNIMELB-LCG2	GOG-Singapore	IN-DAE-VECC-01	JP-KEK-CRC-02	LCG KNU
TOKYO-LCG2	TW-FTT	TW-NCUHEP	TW-NTU-EECS-01	Taiwan-IPAS-LCG2	Taiwan-LCG2

Color Legend  
 GSTAT: OK INFO NOTE WARN ERROR CRIT MAINT OFF  
 SFT: OK NonCrit Crit JobSub JobList Match Sched Down

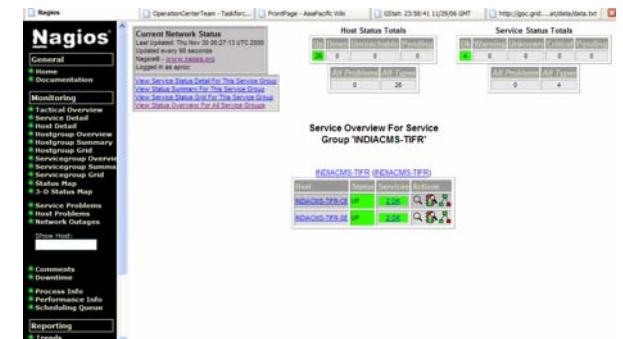
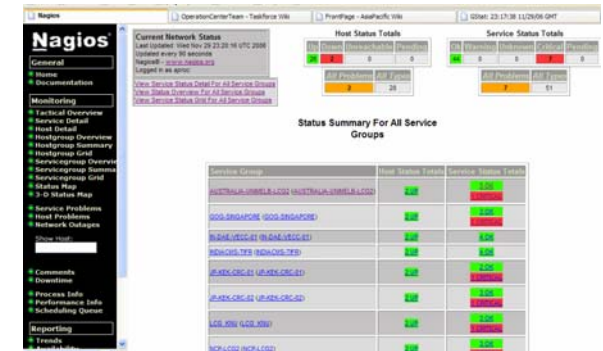
No	Site Reports	GIIS Host	node	cpu	mem	disk	net	version	totalCPU	freeCPU	runJob	waitJob	isAvailable	U	TB
1	Australia-UNIMELB-LCG2	lcr-compute.lpc.unimelb.edu.au	ok	ok	ok	ok	ok	GLITE-3.0.039	30	9	928	1.22	0.35		
2	IN-DAE-VECC-01	gridce01.tier2.kol.res.in	ok	ok	ok	ok	ok	GLITE-3.0.026	26	0	0	3.55	0		
3	INDIACMS-TIER	ce.indiares.in	ok	ok	warn	error	error	GLITE-3.0.030	30	0	0				
4	JP-KEK-CRC-01	dg10.cc.kek.jp	ok	ok	ok	ok	ok	GLITE-3.0.014	11	3	0	1.23	0.14		
5	JP-KEK-CRC-02	h03.cc.kek.jp	ok	ok	ok	ok	ok	GLITE-3.0.046	33	12	0	1.57	1.10		
6	TOKYO-LCG2	dgce01.supp.jp	ok	ok	ok	ok	ok	GLITE-3.0.0174	12	162	5	6.07	1.78		
7	LCG KNU	ehost05.knu.ac.kr	ok	ok	ok	ok	ok	LCG-2.7.0.144	42	2	11	0.11	6.38		





# Nagios Regional Monitoring

- Nagios is being setup to supplement current monitoring tools
- Benefits
  - Tests run at faster frequency
    - 5-10 minutes
    - **Faster** response to faults
  - Add customized plugins
    - Run **low level** tests for faster isolation of problems
    - Tests that may **not be available** in global monitoring tools yet
    - Ability to run tests **on the target host** via NRPE





# APROC resource centers

- Latest events
  - KISTI start certification since late Nov. and become production sites in earlier of Jan.
  - HKU, start certification in earlier Dec., and stable at earlier Jan. Continue effort will try integrating with local batch pool composted by > 180 processors.
  - Mongolia being certified

No	SiteName	NodeName	Status	ops				
				js	ver	ca	bi	csk
<b>AsiaPacific</b>								
1	<a href="#">Australia-UNIMELB-LCG2</a>	lcg-compute.hpc.unimelb.edu.au	OK	<a href="#">ok</a>	<a href="#">3.0.2</a>	<a href="#">ok</a>	<a href="#">ok</a>	<a href="#">ok</a>
2	<a href="#">GOG-Singapore</a>	melon.ngpp.ngp.org.sg	OK	<a href="#">ok</a>	<a href="#">3.0.2</a>	<a href="#">ok</a>	<a href="#">ok</a>	<a href="#">ok</a>
3	<a href="#">HK-HKU-CC-01</a>	ce.grid.hku.hk	OK	<a href="#">ok</a>	<a href="#">3.0.2</a>	<a href="#">ok</a>	<a href="#">ok</a>	<a href="#">ok</a>
4	<a href="#">IN-DAE-VECC-01</a>	gridce01.tier2-kol.res.in	OK	<a href="#">ok</a>	<a href="#">3.0.2</a>	<a href="#">ok</a>	<a href="#">ok</a>	<a href="#">ok</a>
5	<a href="#">INDIACMS-TIFR</a>	ce.indiacms.res.in	OK	<a href="#">ok</a>	<a href="#">3.0.2</a>	<a href="#">ok</a>	<a href="#">ok</a>	<a href="#">ok</a>
6	<a href="#">JP-KEK-CRC-01</a>	dg10.cc.kek.jp	OK	<a href="#">ok</a>	<a href="#">3.0.2</a>	<a href="#">ok</a>	<a href="#">ok</a>	<a href="#">ok</a>
7	<a href="#">JP-KEK-CRC-02</a>	rls02.cc.kek.jp	OK	<a href="#">ok</a>	<a href="#">3.0.2</a>	<a href="#">ok</a>	<a href="#">ok</a>	<a href="#">ok</a>
8	<a href="#">KR-KISTI-GCRT-01</a>	venus.gridcenter.or.kr	OK	<a href="#">ok</a>	<a href="#">3.0.2</a>	<a href="#">ok</a>	<a href="#">ok</a>	<a href="#">ok</a>
9	<a href="#">LCG_KNU</a>	cluster50.knu.ac.kr	OK	<a href="#">ok</a>	<a href="#">2.7.0</a>	<a href="#">ok</a>	<a href="#">ok</a>	<a href="#">ok</a>
10	<a href="#">NCP-LCG2</a>	pcncp04.ncp.edu.pk	OK	<a href="#">ok</a>	<a href="#">3.0.2</a>	<a href="#">ok</a>	<a href="#">ok</a>	<a href="#">ok</a>
11	<a href="#">PAKGRID-LCG2</a>	CE.pakgrid.org.pk	ERROR	<a href="#">error</a>	<a href="#">3.0.2</a>	<a href="#">ok</a>	<a href="#">ok</a>	<a href="#">ok</a>
12	<a href="#">TOKYO-LCG2</a>	dgce0.icepp.jp	OK	<a href="#">ok</a>	<a href="#">3.0.2</a>	<a href="#">ok</a>	<a href="#">ok</a>	<a href="#">ok</a>
13	<a href="#">TW-FTT</a>	f-ce01.grid.sinica.edu.tw	OK	<a href="#">ok</a>	<a href="#">3.0.2</a>	<a href="#">ok</a>	<a href="#">ok</a>	<a href="#">ok</a>
14	<a href="#">TW-NCUHEP</a>	grid01.phy.ncu.edu.tw	OK	<a href="#">ok</a>	<a href="#">3.0.1</a>	<a href="#">ok</a>	<a href="#">ok</a>	<a href="#">ok</a>
15	<a href="#">TW-NIU-EECS-01</a>	niugce.grid.niu.edu.tw	OK	<a href="#">ok</a>	<a href="#">3.0.2</a>	<a href="#">ok</a>	<a href="#">ok</a>	<a href="#">ok</a>
16	<a href="#">TW-NTCU-HPC-01</a>	host001.hpc.ntcu.edu.tw	OK	<a href="#">ok</a>	<a href="#">3.0.2</a>	<a href="#">ok</a>	<a href="#">ok</a>	<a href="#">ok</a>
17	<a href="#">TW-THU-HPC</a>	ce.hpc.csie.thu.edu.tw	OK	<a href="#">ok</a>	<a href="#">2.7.0</a>	<a href="#">ok</a>	<a href="#">ok</a>	<a href="#">ok</a>
18	<a href="#">Taiwan-IPAS-LCG2</a>	atlasce.phys.sinica.edu.tw	OK	<a href="#">ok</a>	<a href="#">3.0.2</a>	<a href="#">ok</a>	<a href="#">ok</a>	<a href="#">ok</a>
19	<a href="#">Taiwan-LCG2</a>	lcg00125.grid.sinica.edu.tw	OK	<a href="#">ok</a>	<a href="#">3.0.1</a>	<a href="#">ok</a>	<a href="#">ok</a>	<a href="#">ok</a>
20	<a href="#">Taiwan-LCG2</a>	quanta.grid.sinica.edu.tw	OK	<a href="#">ok</a>	<a href="#">3.0.0</a>	<a href="#">ok</a>	<a href="#">ok</a>	<a href="#">ok</a>
21	<a href="#">Taiwan-NCUCC-LCG2</a>	ce.cc.ncu.edu.tw	OK	<a href="#">ok</a>	<a href="#">3.0.2</a>	<a href="#">ok</a>	<a href="#">ok</a>	<a href="#">ok</a>






# APROC Availability II

- Daily snapshots of SFT results of region
  - Availability increased to **70-80%** range from 60-70% a half year ago
- CT mostly replica management failure
  - Sensitive to Information System access/performance
  - Request that data management clients can **failover** to secondary BDII
- Network Issues
  - Often the root cause of CT, JL and JS
  - Pakgrid set up local top-level BDII
    - Increase default update timeout and breath time

Remove  
Slow BDII

JS from  
SSH  
upgrade





# Summary

- New year end procurement to meet MOU levels
- Castor improvements made
  - Focus on Castor2 production deployment
- +100MB/s data transfer rate achieved
  - Achieving 200MB/s may also rely on Castor2 deployment
  - Few T2 sites can sustain 20MB/s, testing will continue
- 88% availability target achieved
  - Increase monitoring coverage and notification system
  - Start on-call rotation to provide 24x7 services
- AP site availability has improved to over 70%
- APROC provides operations resources, services and tools
- Hope to hear feedback from site in AP this afternoon



# Thanks You for Your Attention!

- Thanks to efforts from:
  - Experiment Support
    - Chia-Ming Kou, Dominique (CMS, NCUHEP)
    - Suijian Zhou, Zhijun Liang (Atlas, AS IPAS)
  - T1/APROC
    - Jason Shih (T1)                      Dave Wei (3D)
    - Felix Lee (Fabric)                  Joanna Huang (Fabric)
    - Aries Hong (WAN)                  Hung-Che Jen (T1/CASTOR)
    - Jinny Chien (ROC/CIC)              Shu-Ting Liao (ROC/CIC)
    - Yi-Ping Wu (3D)                      Min Tsai (ROC/CIC)



# FAQ for new site

- **Where to start? How to start?**  
<http://www.twgrid.org/aproc/join/>
- **A suitable layout for your site** (Server Requirements, Hardware spac .)  
<http://www.twgrid.org/aproc/join/>
- **CA question**  
<http://ca.grid.sinica.edu.tw/general/faq.html>
- **Installation and Configuration prob** (How to modify site-info.def ?)  
[http://lists.grid.sinica.edu.tw/apwiki/Grid\\_Administrator\\_Tutorial\\_Hands-on\\_Instructions](http://lists.grid.sinica.edu.tw/apwiki/Grid_Administrator_Tutorial_Hands-on_Instructions)