ASGC Site report

Operation Team - Felix Lee, ASGC

Nov. 2010

Outline of ASGC Status Summary

- WLCG Services Status in Taiwan
  - Metrics that ATLAS/CMS/WLCG used
  - Resource plan
  - Improvement
- Data Center.
- Cloud activities in ASGC.

ASGC System Architecture

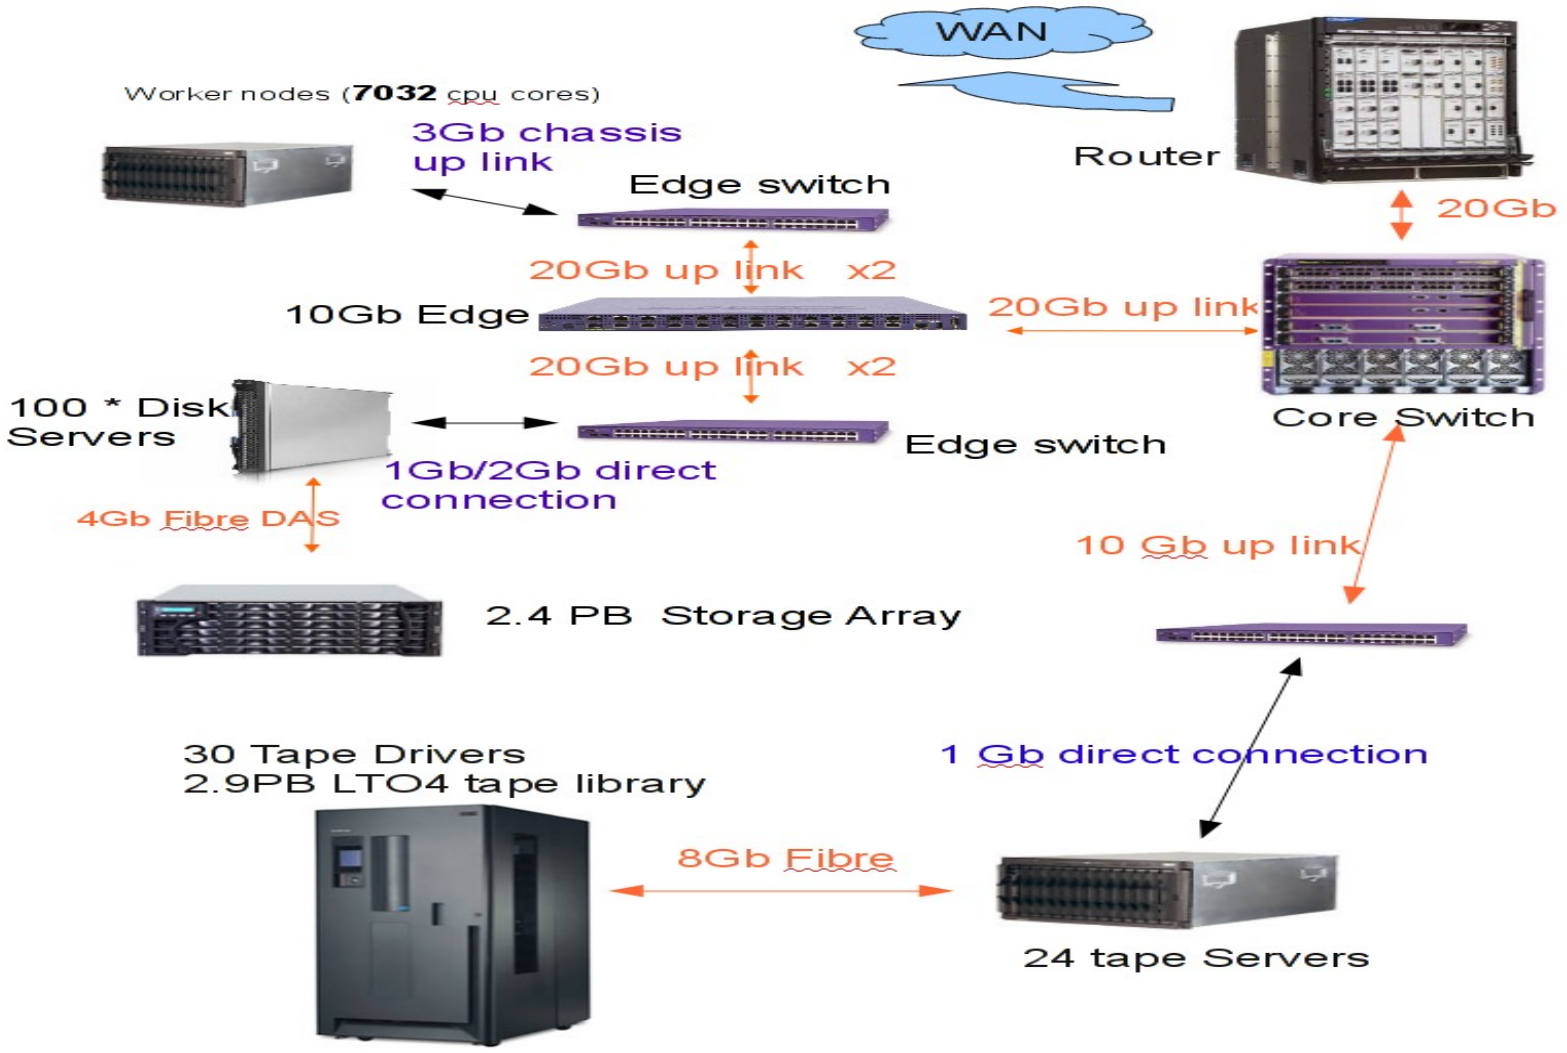

ATLAS Production Jobs at ASGC (year view)

600K jobs finished since 2010-03-01

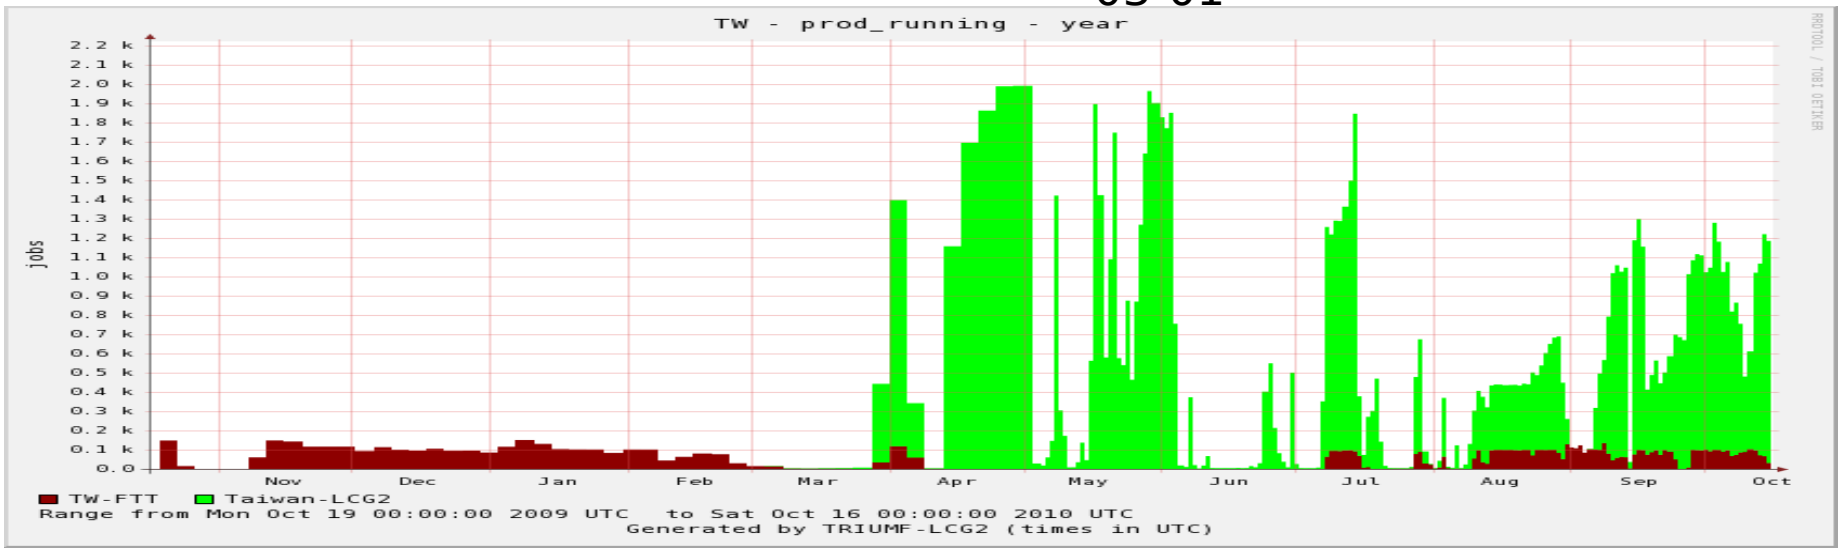

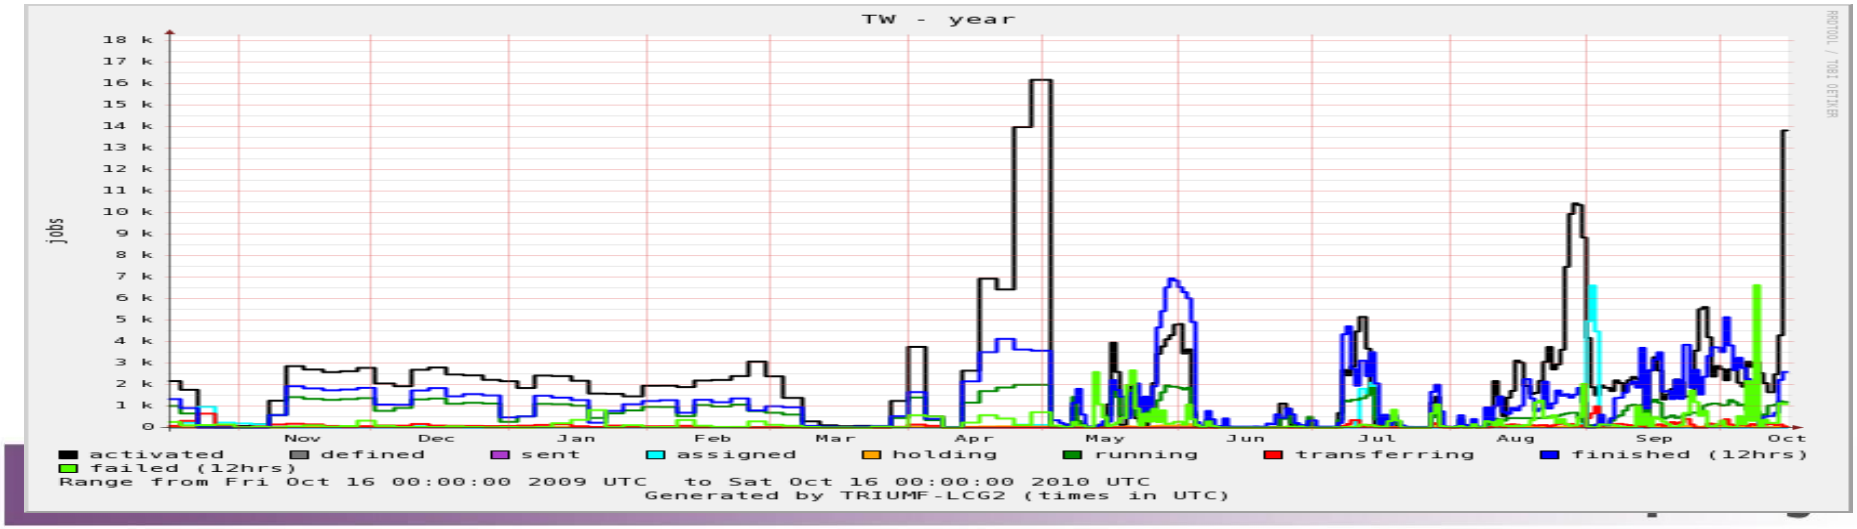

ATLAS Analysis Jobs at ASGC (year view)

631K jobs finished since 2010-03-01

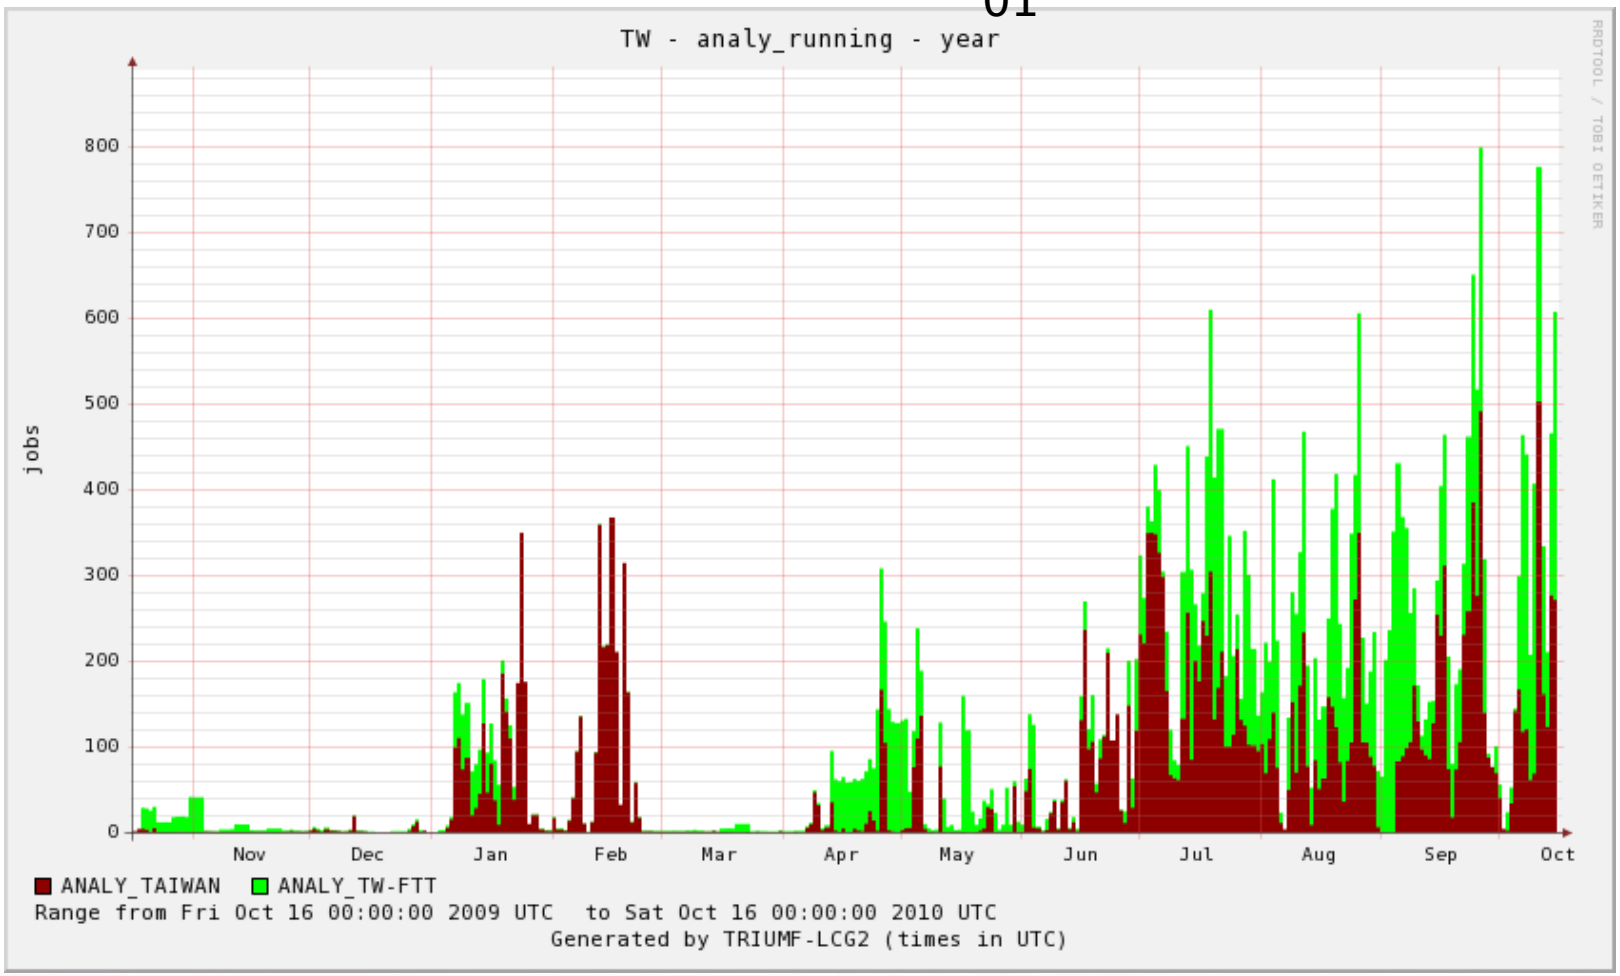

CMS Jobs at ASGC T1

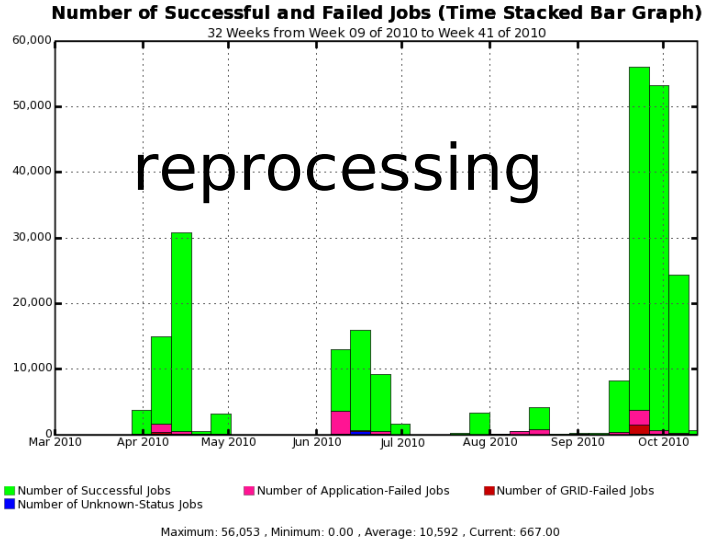

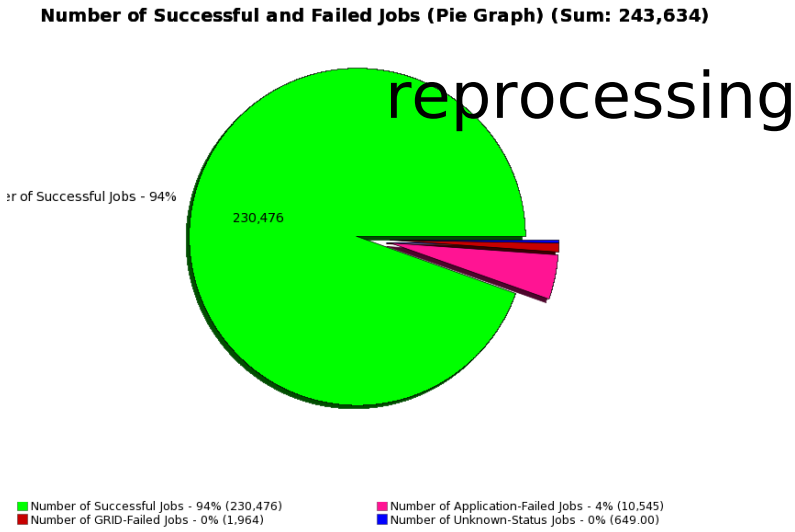

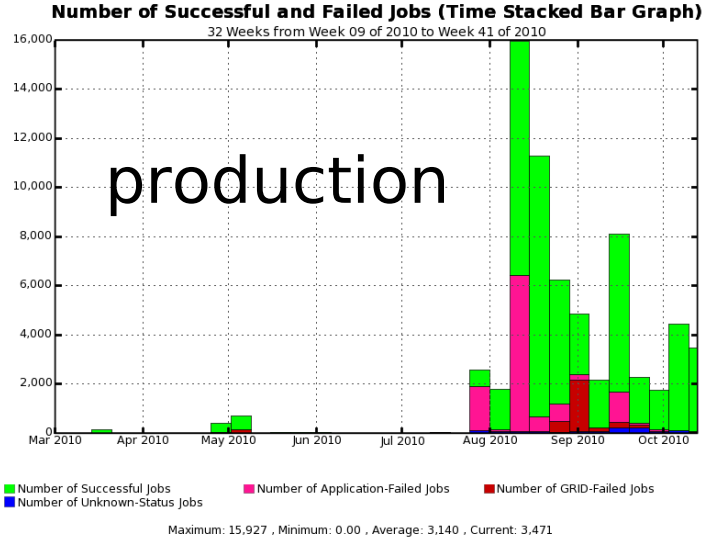

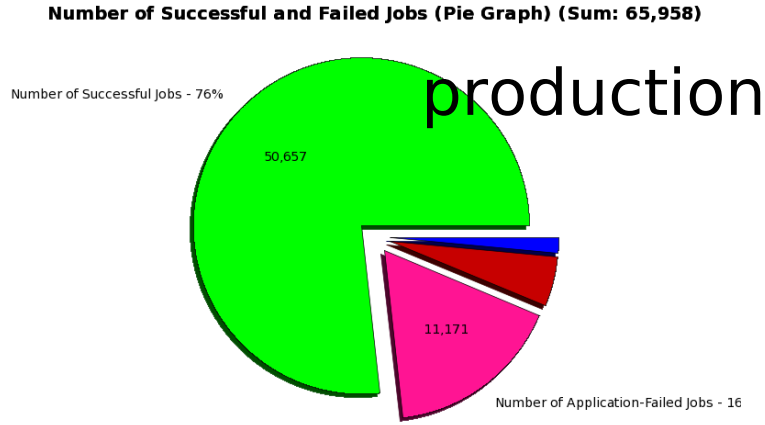

puting

T1 Computing Resources

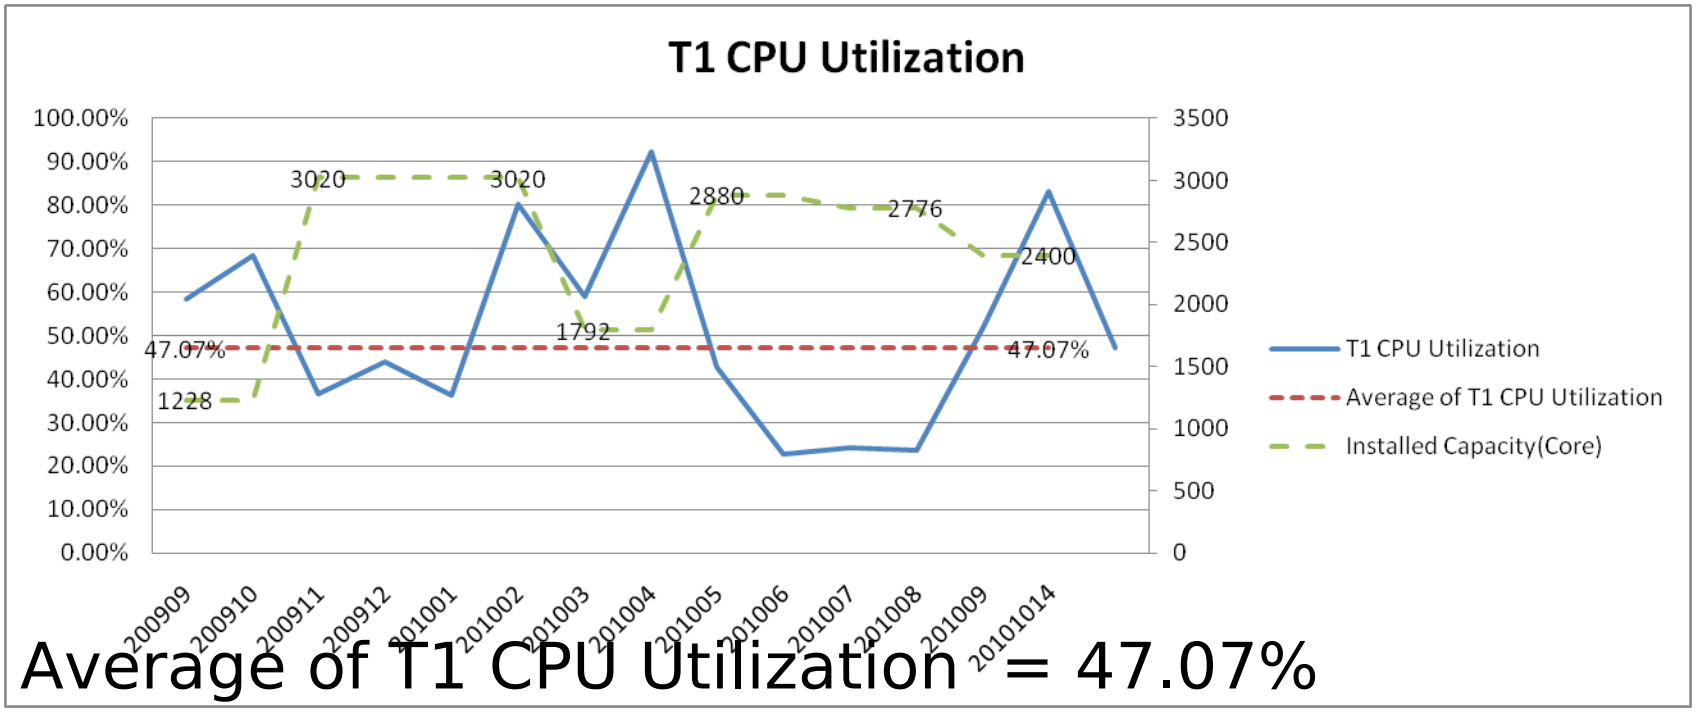

Average of T1 CPU Utilization = 47.07%

- Nov. 2009 - New Quad-Core CPU involved
- Mar. 2010 - Starting to do system migration to SLC5
- May. 2010 - System migration finished
- Sep. 2010 - 376 cores(47 blades) for shared system pre-test
- Oct. 2010 – 800 cores will be added to Tier-1 pool
- Oct. 2010 – Procurement of ~1800 cores is underway

ASGC T1 Computing Resources (II)

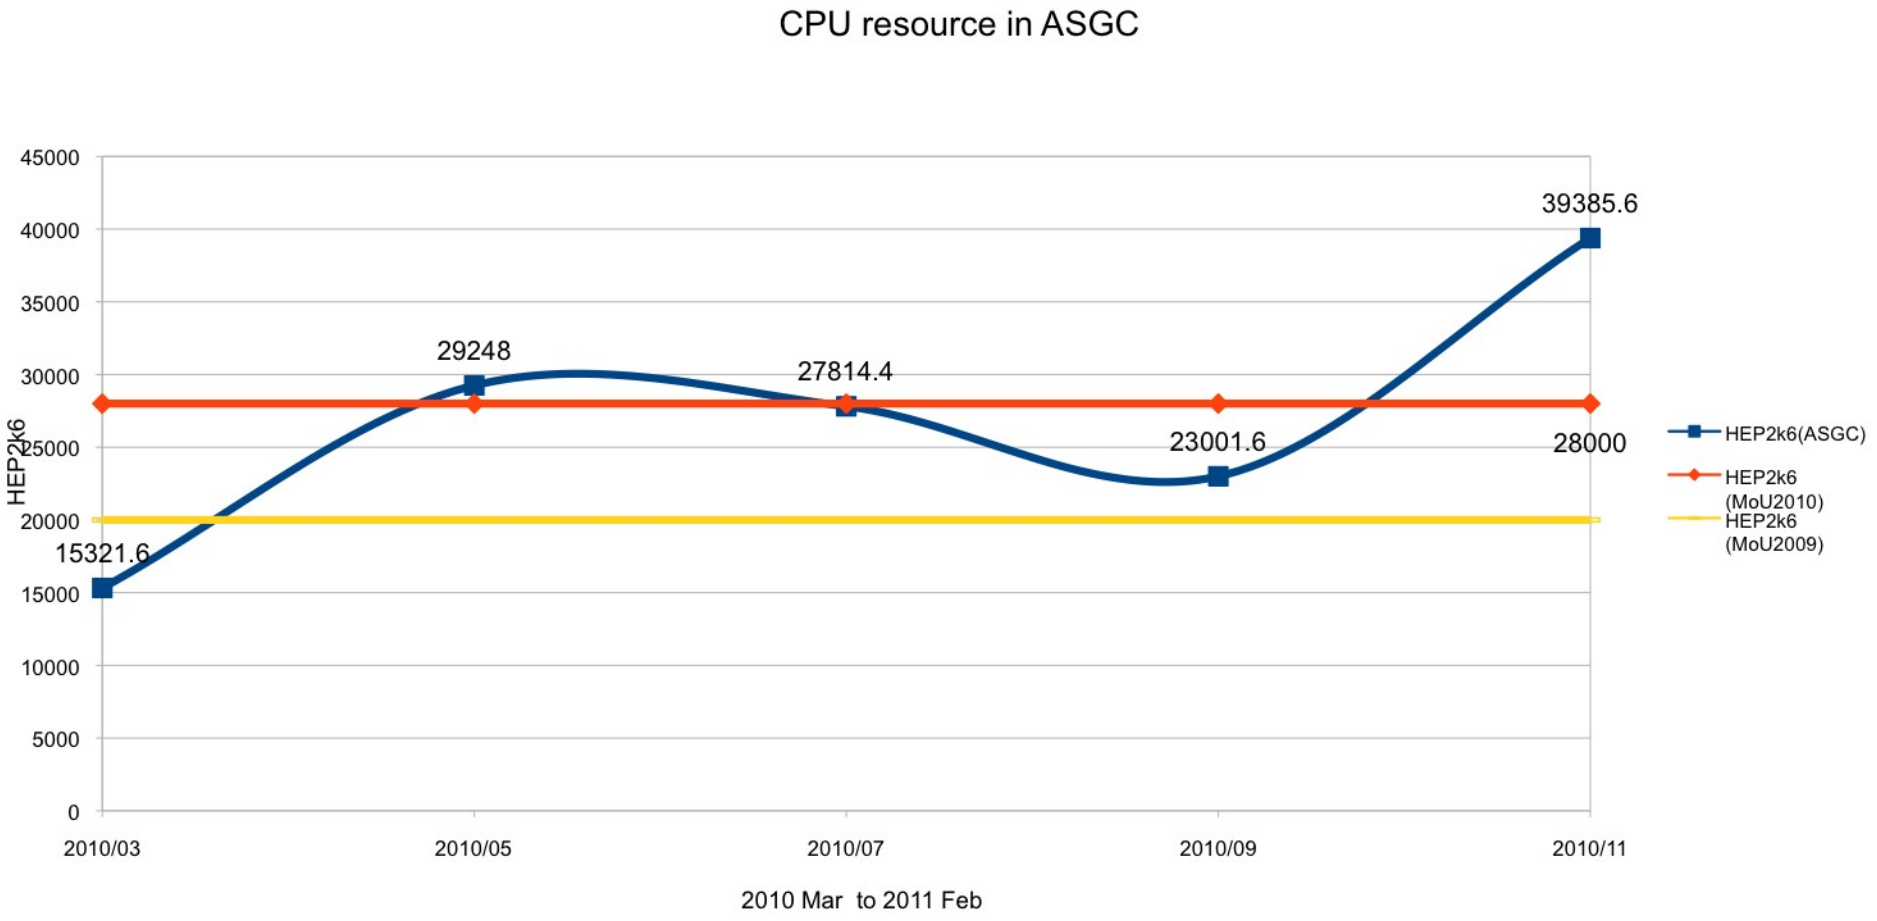

T1 Disk System

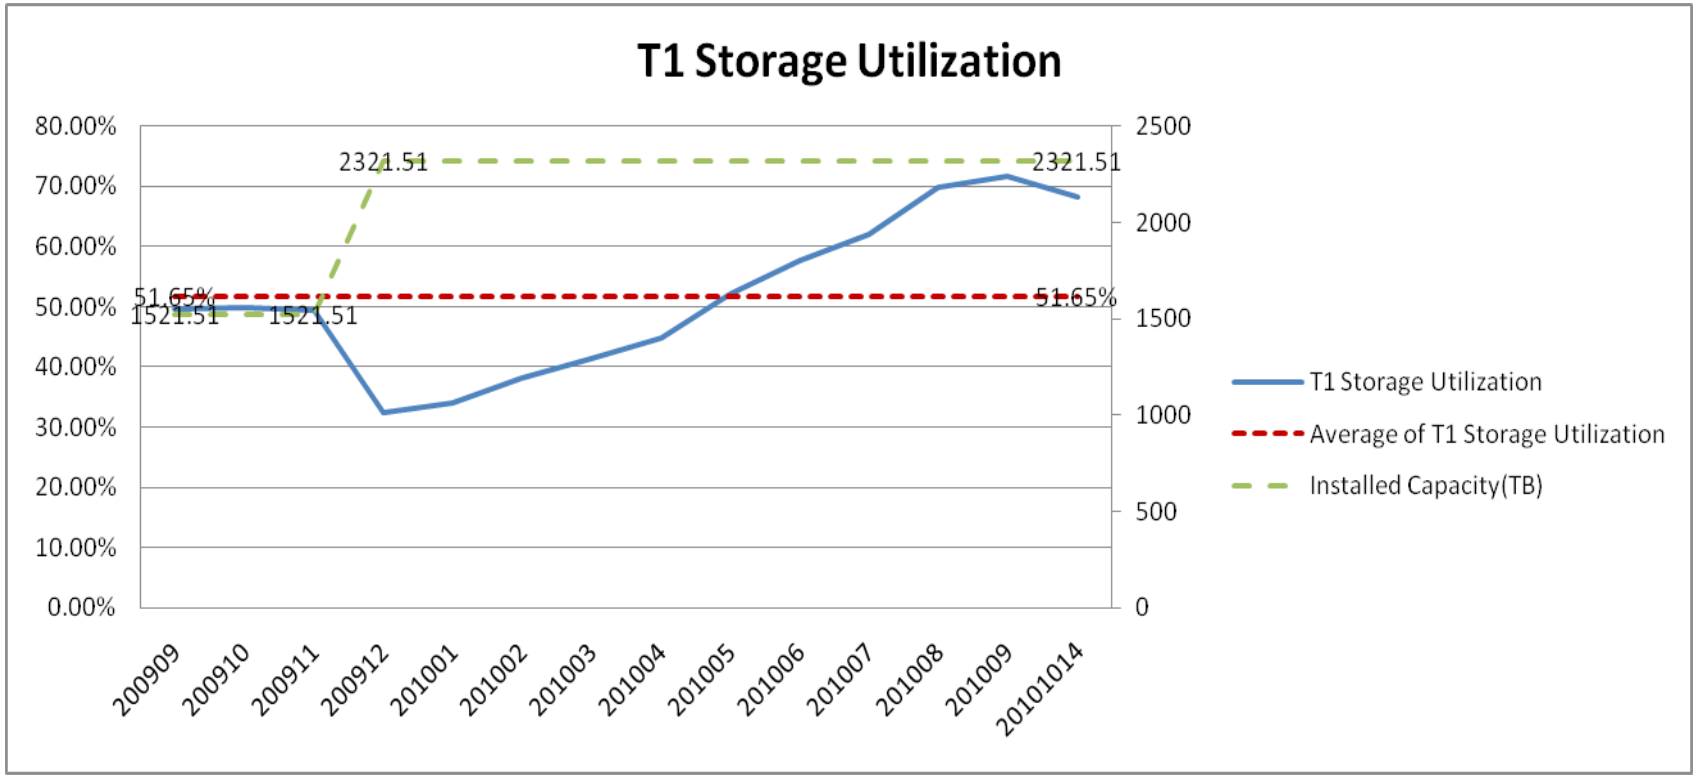

Average of T1 Storage Utilization = 51.65%

- Dec. 2009 – New procured 800TB disk online
- Oct. 2010 – New disk array/server architecture is under evaluation
- Oct. 2010 – Procurement of 2PB disk capacity is underway

T1 Tape System

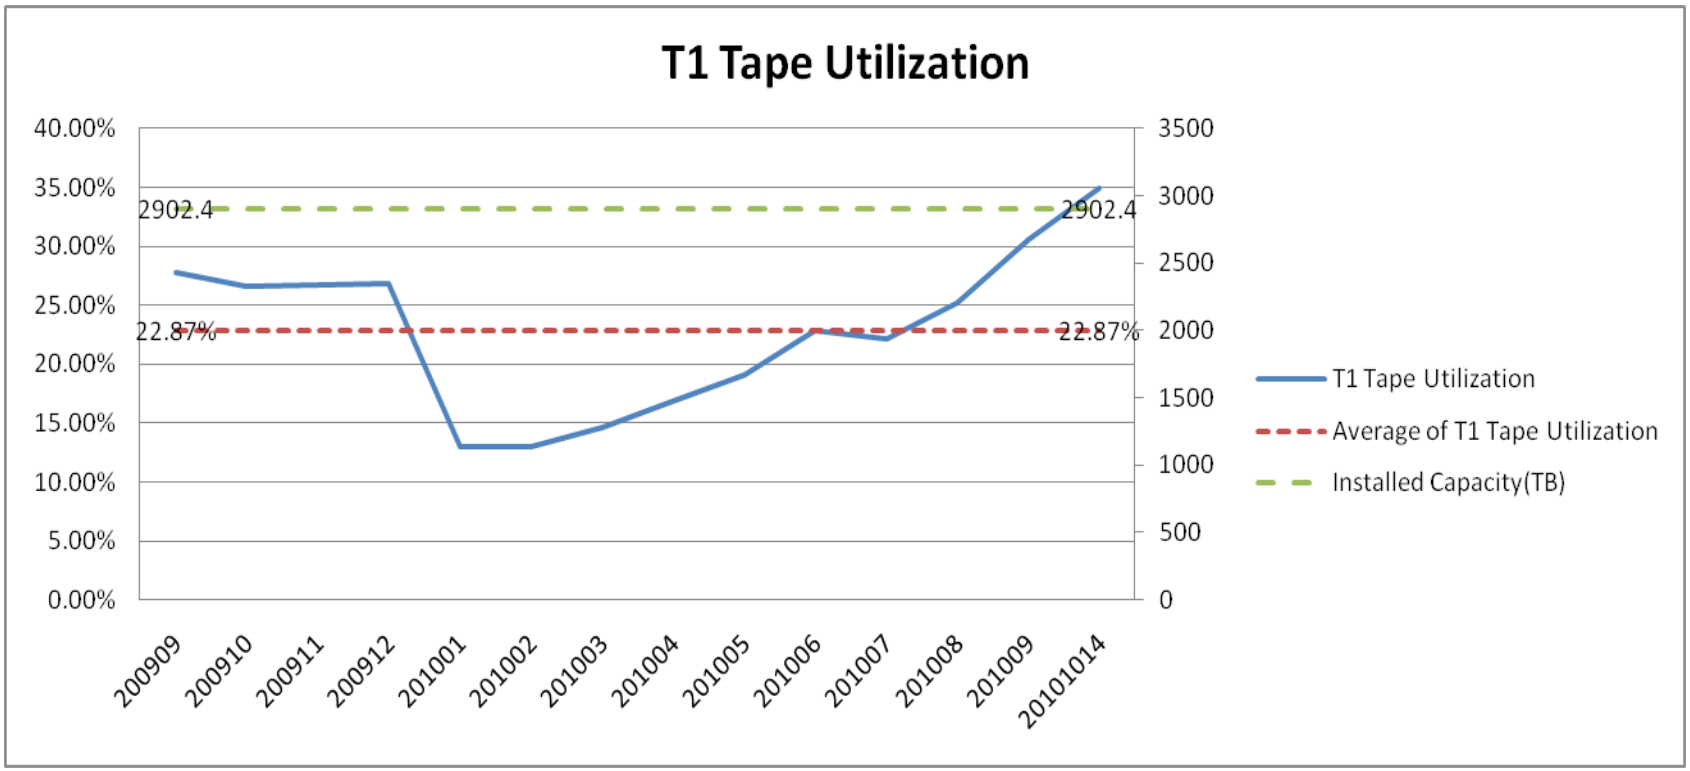

Average of T1 Tape Utilization = 22.87%

- HA by dual online robot were implemented in Mar 2010
- 3,000 x 800GB/LTO4 tapes will be online by Dec. 2010.

E-Science Collab Networking in Asia Pacific Region

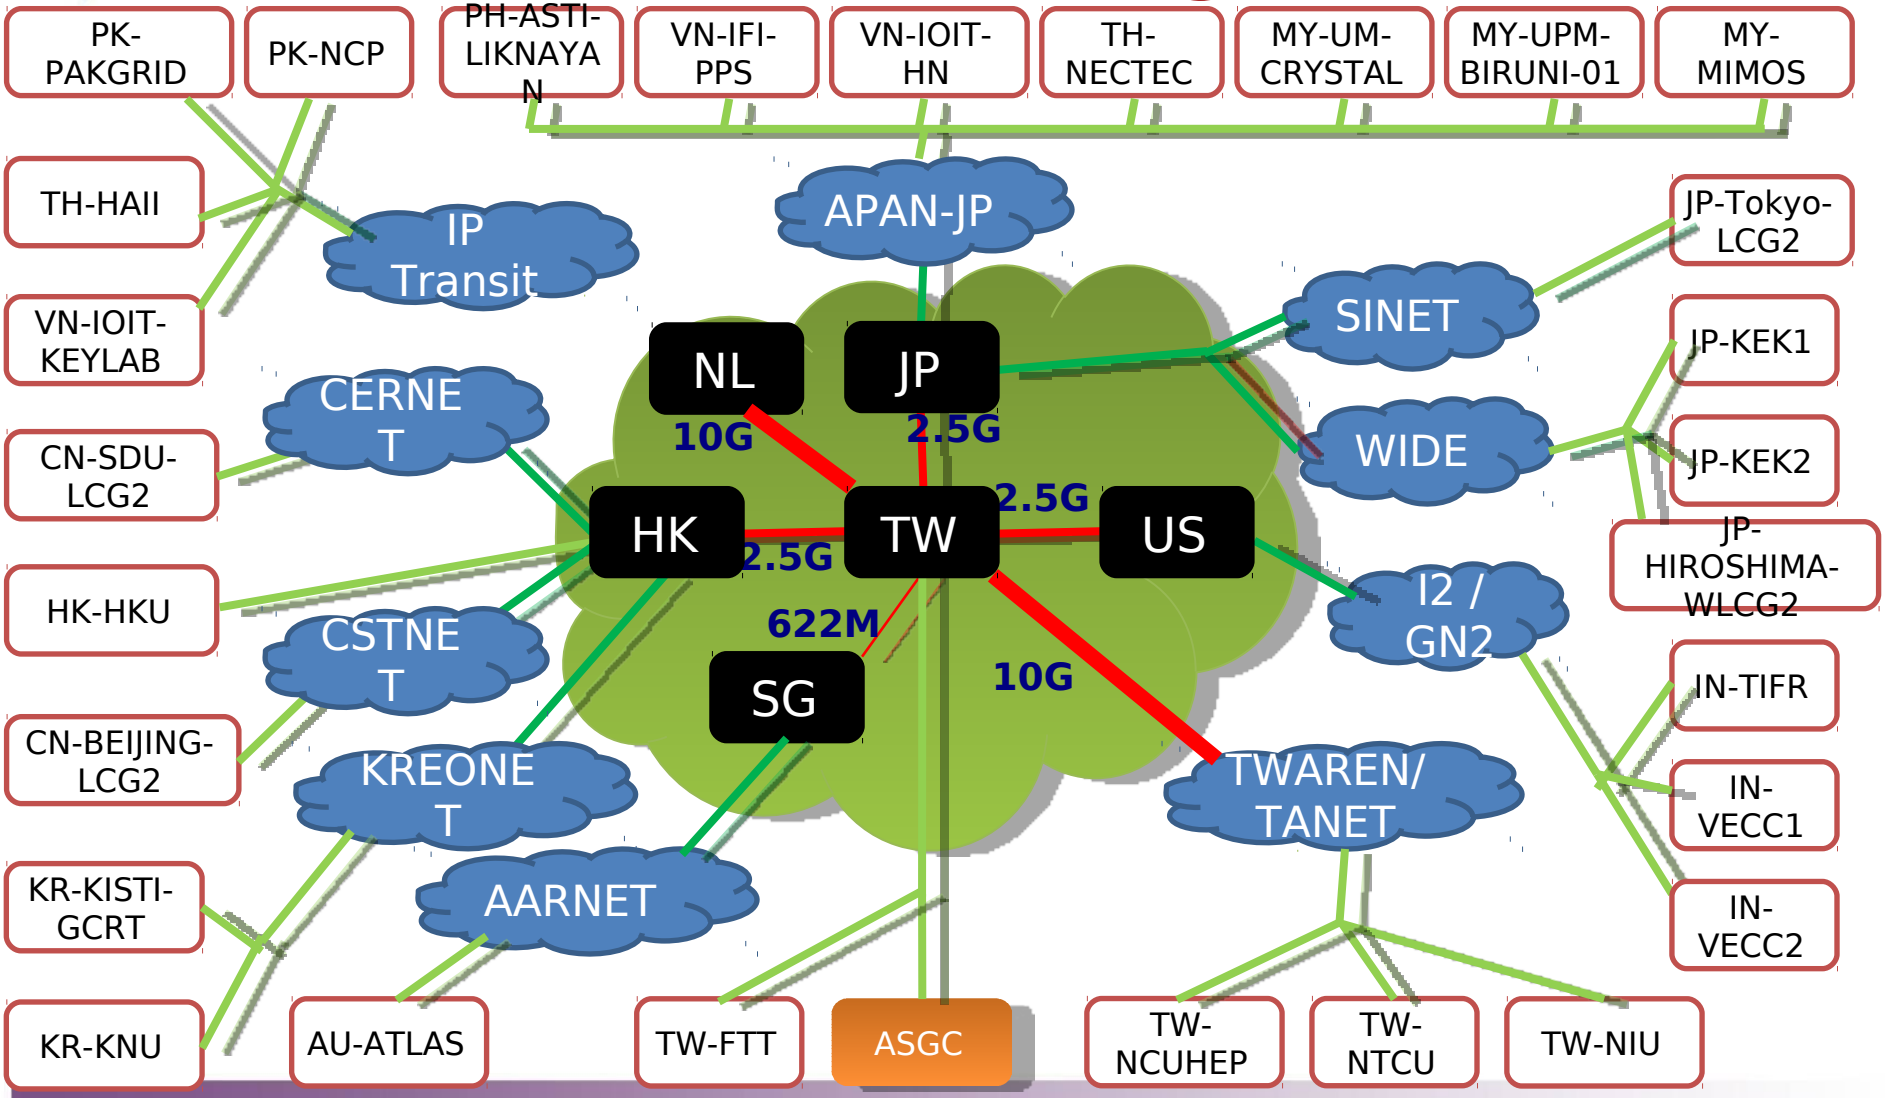

DC Networking

- Blade chassis bandwidth relies on switch module.
  - For current switch module, the lacp seems that is not good enough.

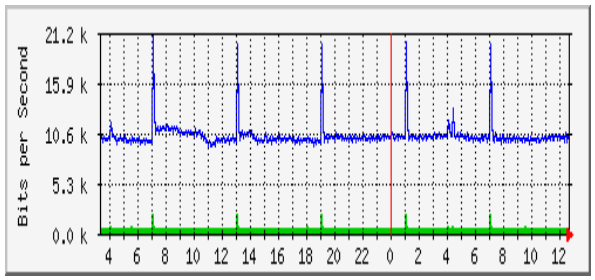

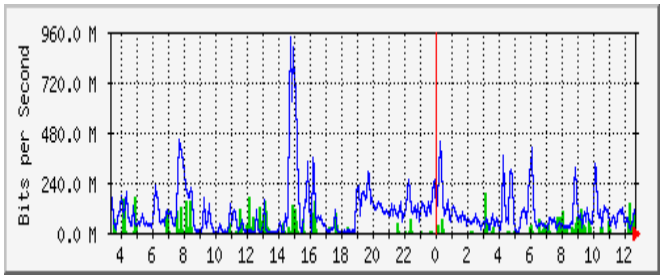

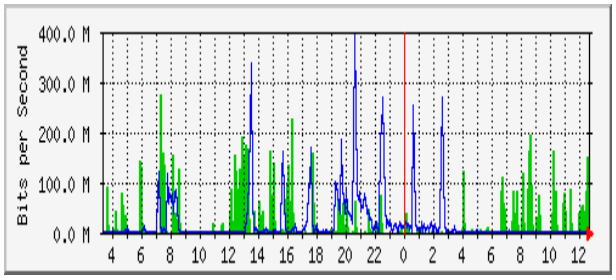

- Upgrading all blade chassis and disk server to be 10GbE.

Tier1 Server Improvement

- System architecture of ASGC
  - Networking (architecture and 10G backbone)
  - Upgrade bandwidth between WN/CE and SE
- Storage system architecture
  - 160TB/server
  - 10 Gb interface
- Castor System – to upgrade to 2.1.9
- Oracle System – Dataguard, TSM backup
- Data Centre Monitoring and Operation
- Monitoring for Intelligent and Preventive Operation
- High Availability of Core Services and software scratch

ASGC Data Center

- **Total Capacity**
  - 2MW, 330 tons AHUs
  - 93 racks
  - ~ 800 m²
- **Resources**
  - 7,032CPU Cores
  - 2,400 TB Disk
  - 2,900 TB Tape
- **Rack Space Usage**
  - AS e-Science: ~63%
  - ASCC: 10%
  - IPAS: 7.0%
  - Free: ~20%

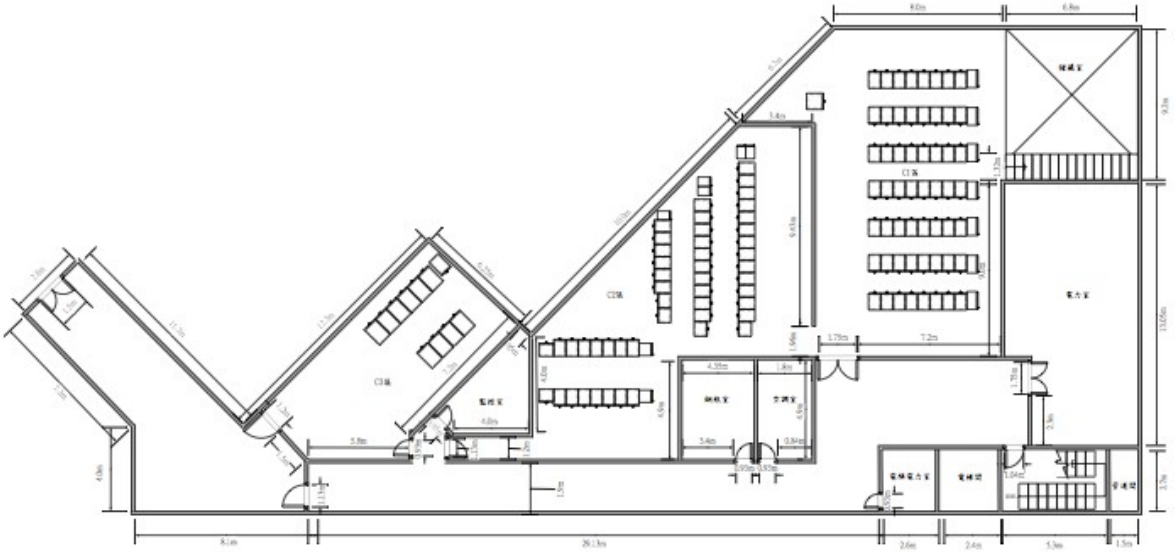

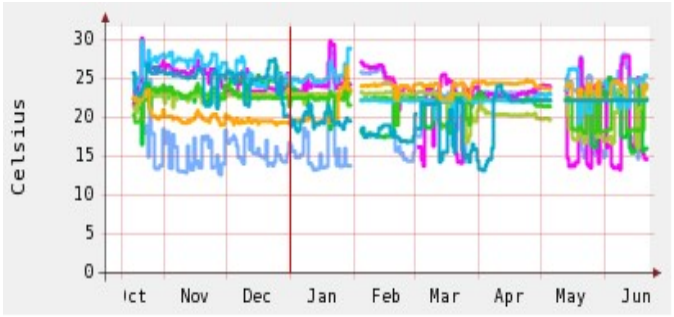

Monitoring the power consumption and temperature of every piece of equipment every 10 seconds.

Data Center

- Power generator will start in 15 sec once there is any power cut (2x400KW generators installed)
- UPS Protection for Core Services – based on data safety requirements
  - Oracle RACs, Disk Arrays and File Servers have already been protected
  - Network devices, LFC, FTS, SRM, BDII, WMS, CE, SE and storage systems will be installed soon
- Power Consumption
  - Power consumption of air conditioners is linearly related to the condenser temperature outside the data center. Analyzing long term data will allow us to build models that could assist us in operating the center intelligently.
- Thermal Management
- Battery Management

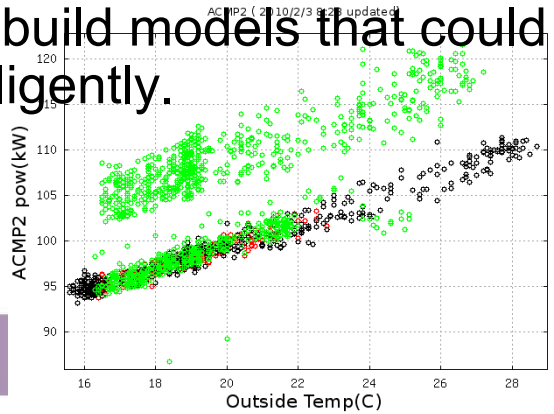

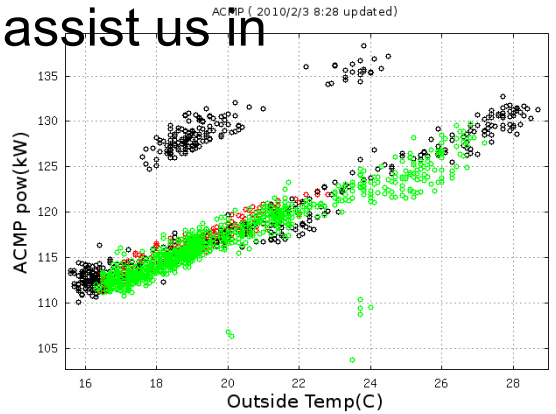

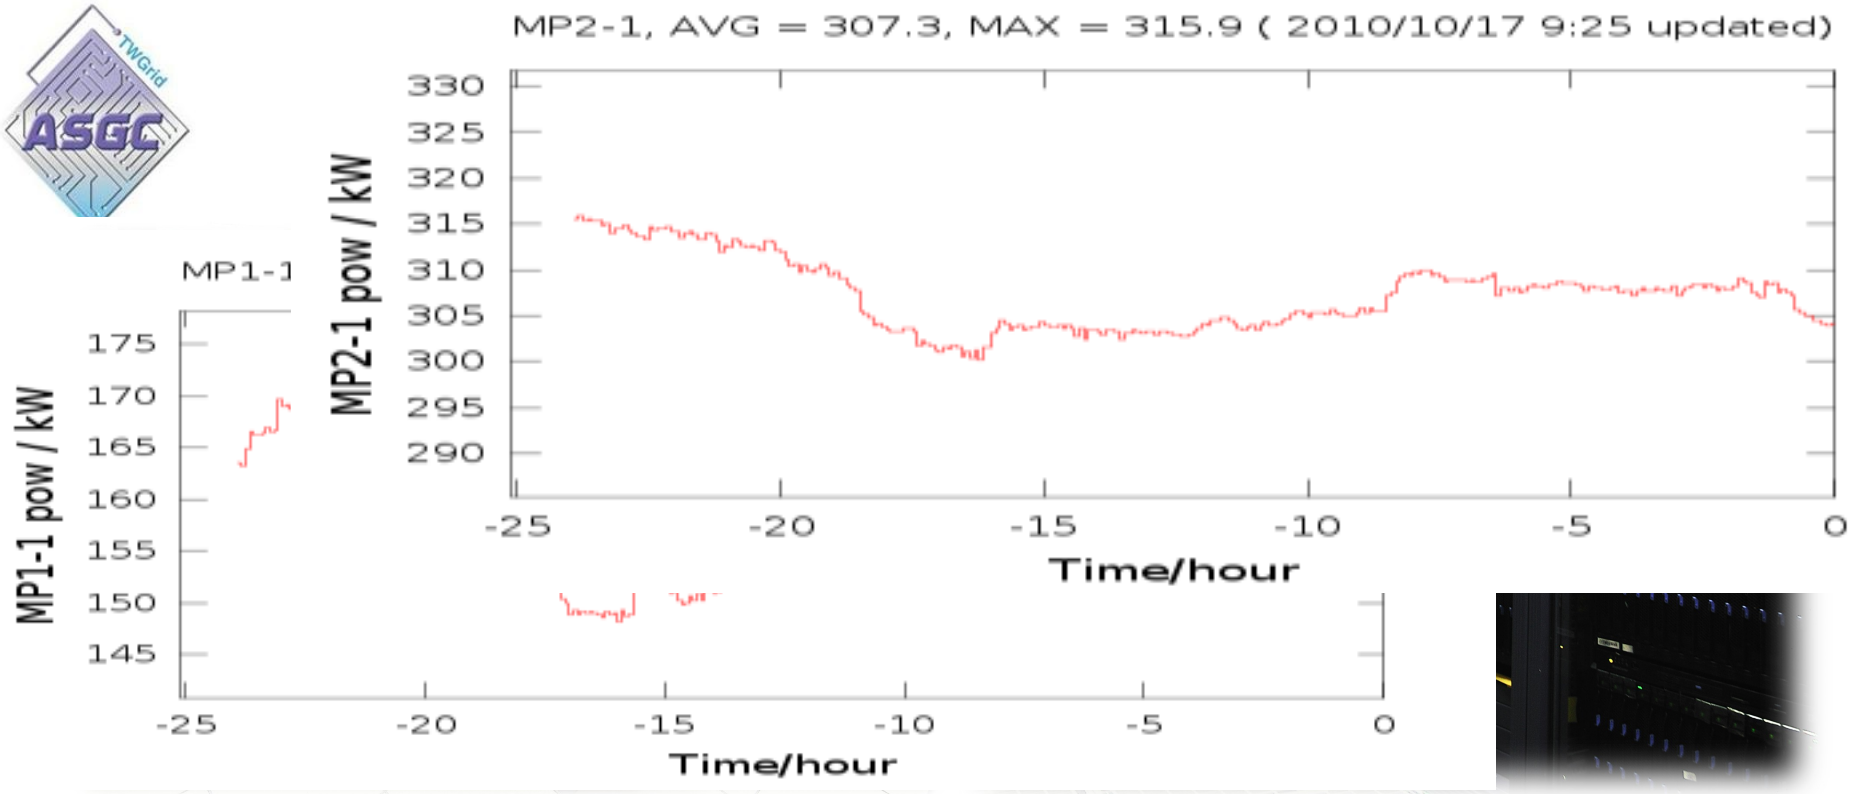
MP2-1, AVG = 307.3, MAX = 315.9 (2010/10/17 9:25 updated)
MP2-1 pow / kW
330
325
320
315
310
305
300
295
290
-25
-20
-15
-10
-5
0
Time/hour
MP1-1
MP1-1 pow / kW
175
170
165
160
155
150
145
-25
-20
-15
-10
-5
0
Time/hour

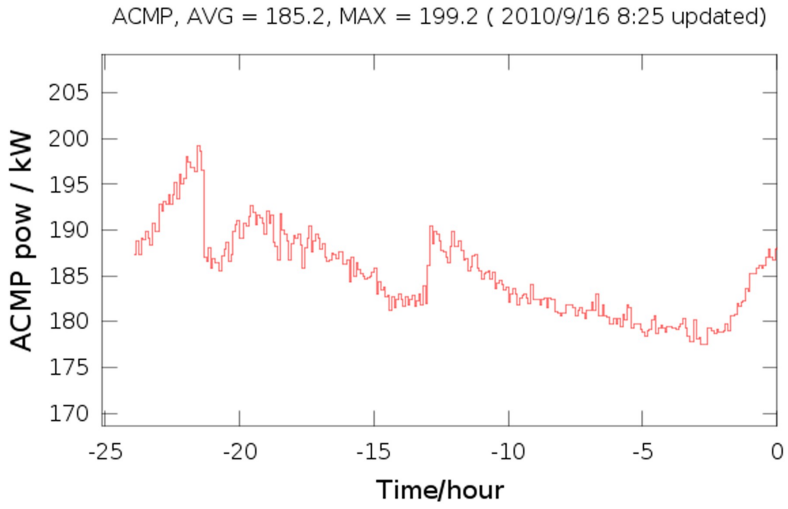
ACMP, AVG = 185.2, MAX = 199.2 (2010/9/16 8:25 updated)
ACMP pow / kW
205
200
195
190
185
180
175
170
-25
-20
-15
-10
-5
0
Time/hour

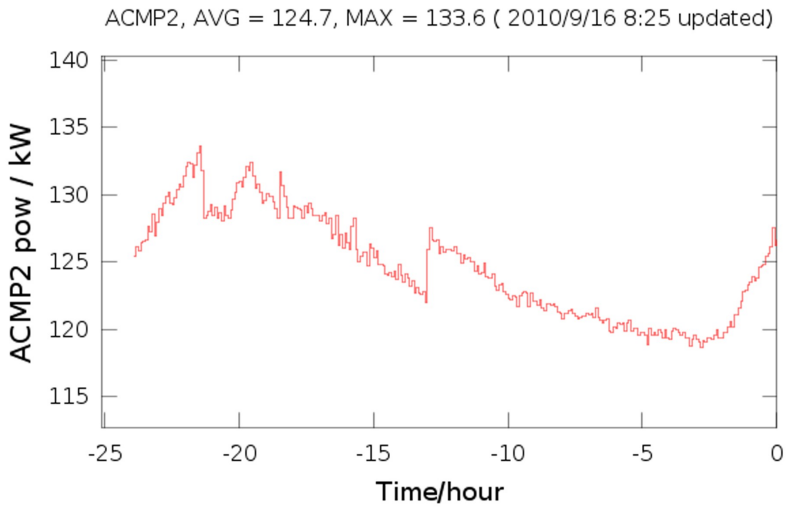
ACMP2, AVG = 124.7, MAX = 133.6 (2010/9/16 8:25 updated)
ACMP2 pow / kW
140
135
130
125
120
115
-25
-20
-15
-10
-5
0
Time/hour

Power System

- No burst power cut but only voltage low
- Have 10 sec monitoring
- Only one power flicker from Jan. 2010, but it's not cleared if the problem comes from TaiPower.
- The new supercomputing in Kobe, UPS is for storage system only, and no power generator in place.
- > 90% incidents were caused by man-made fault.

24x7 Operation

- On-site engineer in place from 4 Jan. 2010.
- Started from DC emergency and incident process, and had been extended to Fabric Services. Will further cover Grid Operation issues after 3Q2010.
- Senior Grid Engineers are on-site from 01:00 – 16:00 UTC, in working days (16x5).
  - By both staffs in Taipei and in CERN.
- On-call support is always available after work hours.
- Service Quality Improvement
  - By training and good FAQ and documentation.
  - Constantly improve proactive monitoring and alarm mechanism

Cloud Activities

ASGC Profile: Support Research Innovation and Innovate ICT

	2002	2004	2005	2006	2008	2009	2010 -
Int'l Collaboration	WLCG TWGrid	EGEE, OSG	WLCG T1 established		EUAsiaGrid		EGI-Inspire, CHAIN, EMI DEGISCO
HEP	ATLAS, CMS		CDF	Belle	WLCG-CCRC	WLCG-STEP09	WLCG in 7TeV
BioInfo, BioMed			mpiBlust-g2	Drug Disc - AF	Drug Disc - AF2	GVSS-DengueF	eHealth, drug-res. bacteria DD
Earth Science				Seismic Data Center		SWave Sim	Seismogram Sim SGT, HazardMap
Env. Change					Carbon Flux		Weather Sim.
Hum. SS	Long-Term Preservation					GISGrid	Social Sim
Grid Infra & Tech	LCG-0 SRB ASGCCA		APROC, GGUS, GStat	Grid Interop - MS CCS	GAP, SRM-SRB	Gstat2, Wkf eng. VC, DesktopG	Data Manager, VMM/Env, DesktopGrid. Smart Center
Dissem	ISGC and Training since 2003			Grid Camp and Internship Program			CHEP

System Architecture

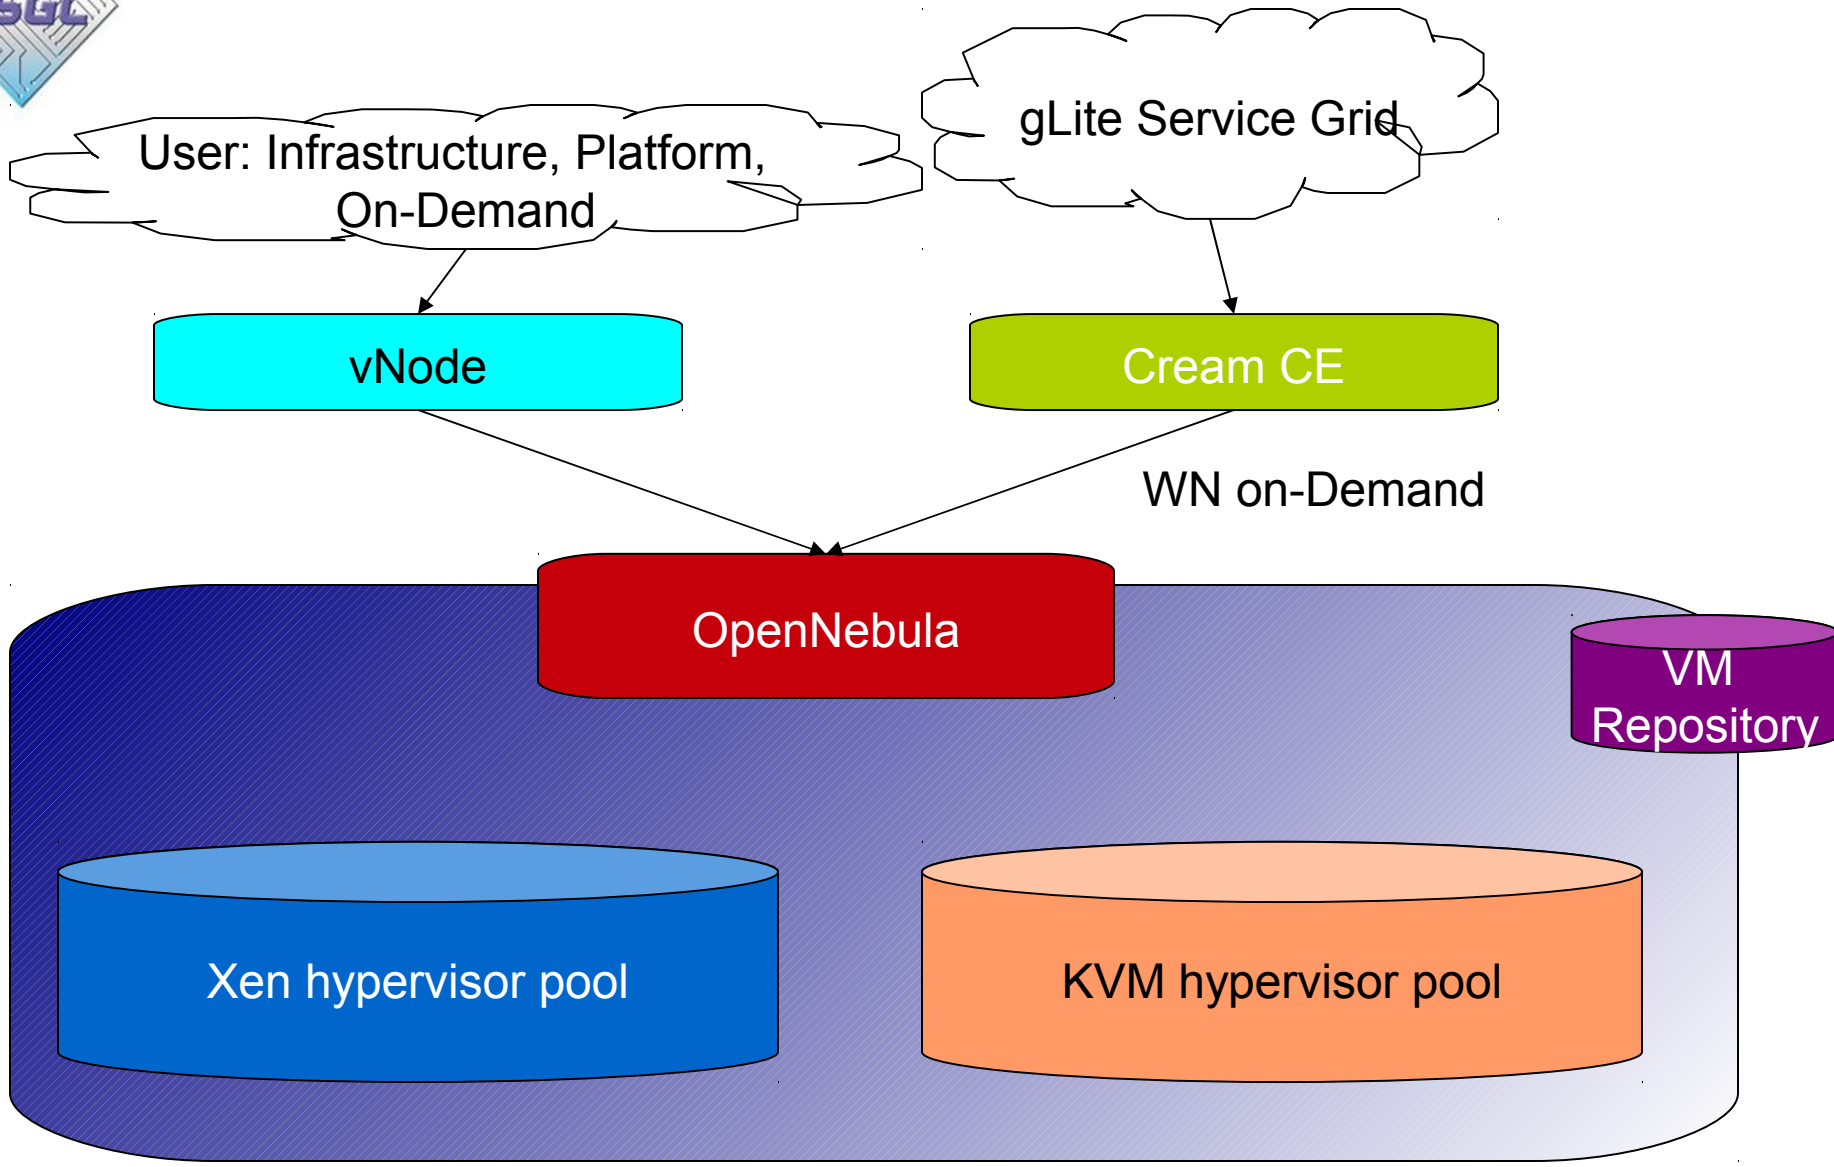

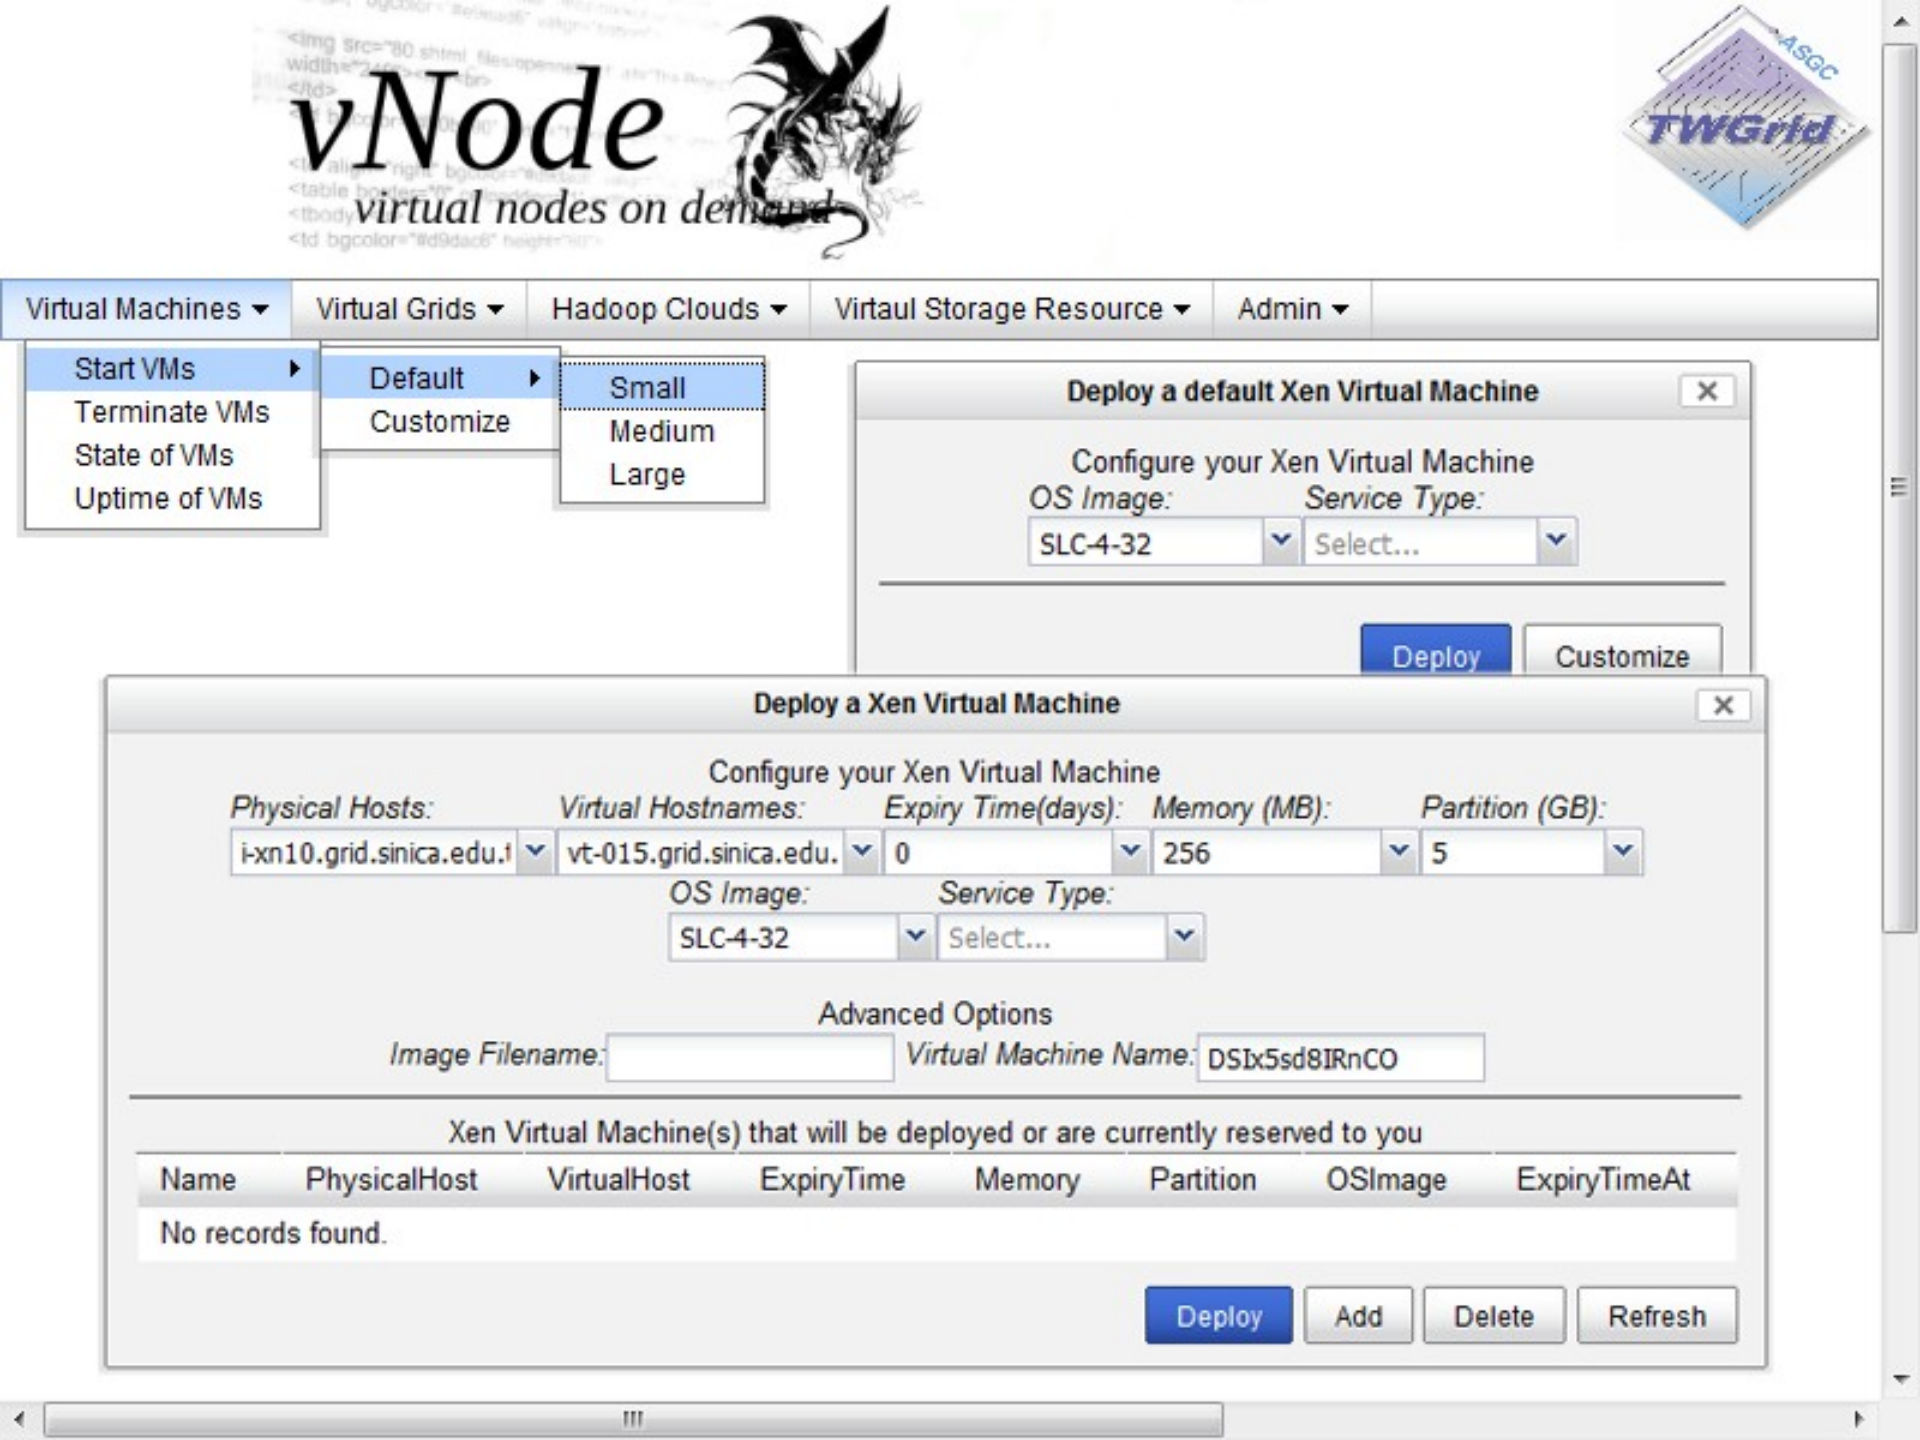
vNode
virtual nodes on demand
ASGC
TWGrid
Virtual Machines
Virtual Grids
Hadoop Clouds
Virtaul Storage Resource
Admin
Start VMs
Terminate VMs
State of VMs
Uptime of VMs
Default
Customize
Small
Medium
Large
Deploy a default Xen Virtual Machine
Configure your Xen Virtual Machine
OS Image:
Service Type:
SLC-4-32
Select...
Deploy
Customize
Deploy a Xen Virtual Machine
Configure your Xen Virtual Machine
Physical Hosts:
Virtual Hostnames:
Expiry Time(days):
Memory (MB):
Partition (GB):
i-xn10.grid.sinica.edu.t
vt-015.grid.sinica.edu.
0
256
5
OS Image:
Service Type:
SLC-4-32
Select...
Advanced Options
Image Filename:
Virtual Machine Name:
DSIx5sd8IRnCO
Xen Virtual Machine(s) that will be deployed or are currently reserved to you
Name
PhysicalHost
VirtualHost
ExpiryTime
Memory
Partition
OSImage
ExpiryTimeAt
No records found.
Deploy
Add
Delete
Refresh

Provisioning of Customized Application Environment

Example: Map Reduce + HDFS

vNode
virtual nodes on demand

Virtual Machines | Virtual Grids | Hadoop Clouds | Virtaul Storage Resource | Admin

Deploy a Xen Hadoop Cloud

Configure Xen Hadoop Cloud

☑ Automatically select the Xen Virtual Machines:
Note: Selecting the option will deploy n number of virtual machines with a default configuration, v
☑ Automatically start virtual site
Note: Under construction
How many virtual Machines: 64

Sitename:
ASGC_Hadoop1
Repository: rfio:///castor/grid.sinica.edu.

Further plan

- WLCG VM catalog service
- EMI Cloud task.
- Add 100 blades for KVM/Xen hypervisor.

Thank you