US ATLAS

Western Tier 2 Status and Plan

Wei Yang

ATLAS Physics Analysis Retreat
SLAC

March 5, 2007

ATLAS Distributed Data Management system

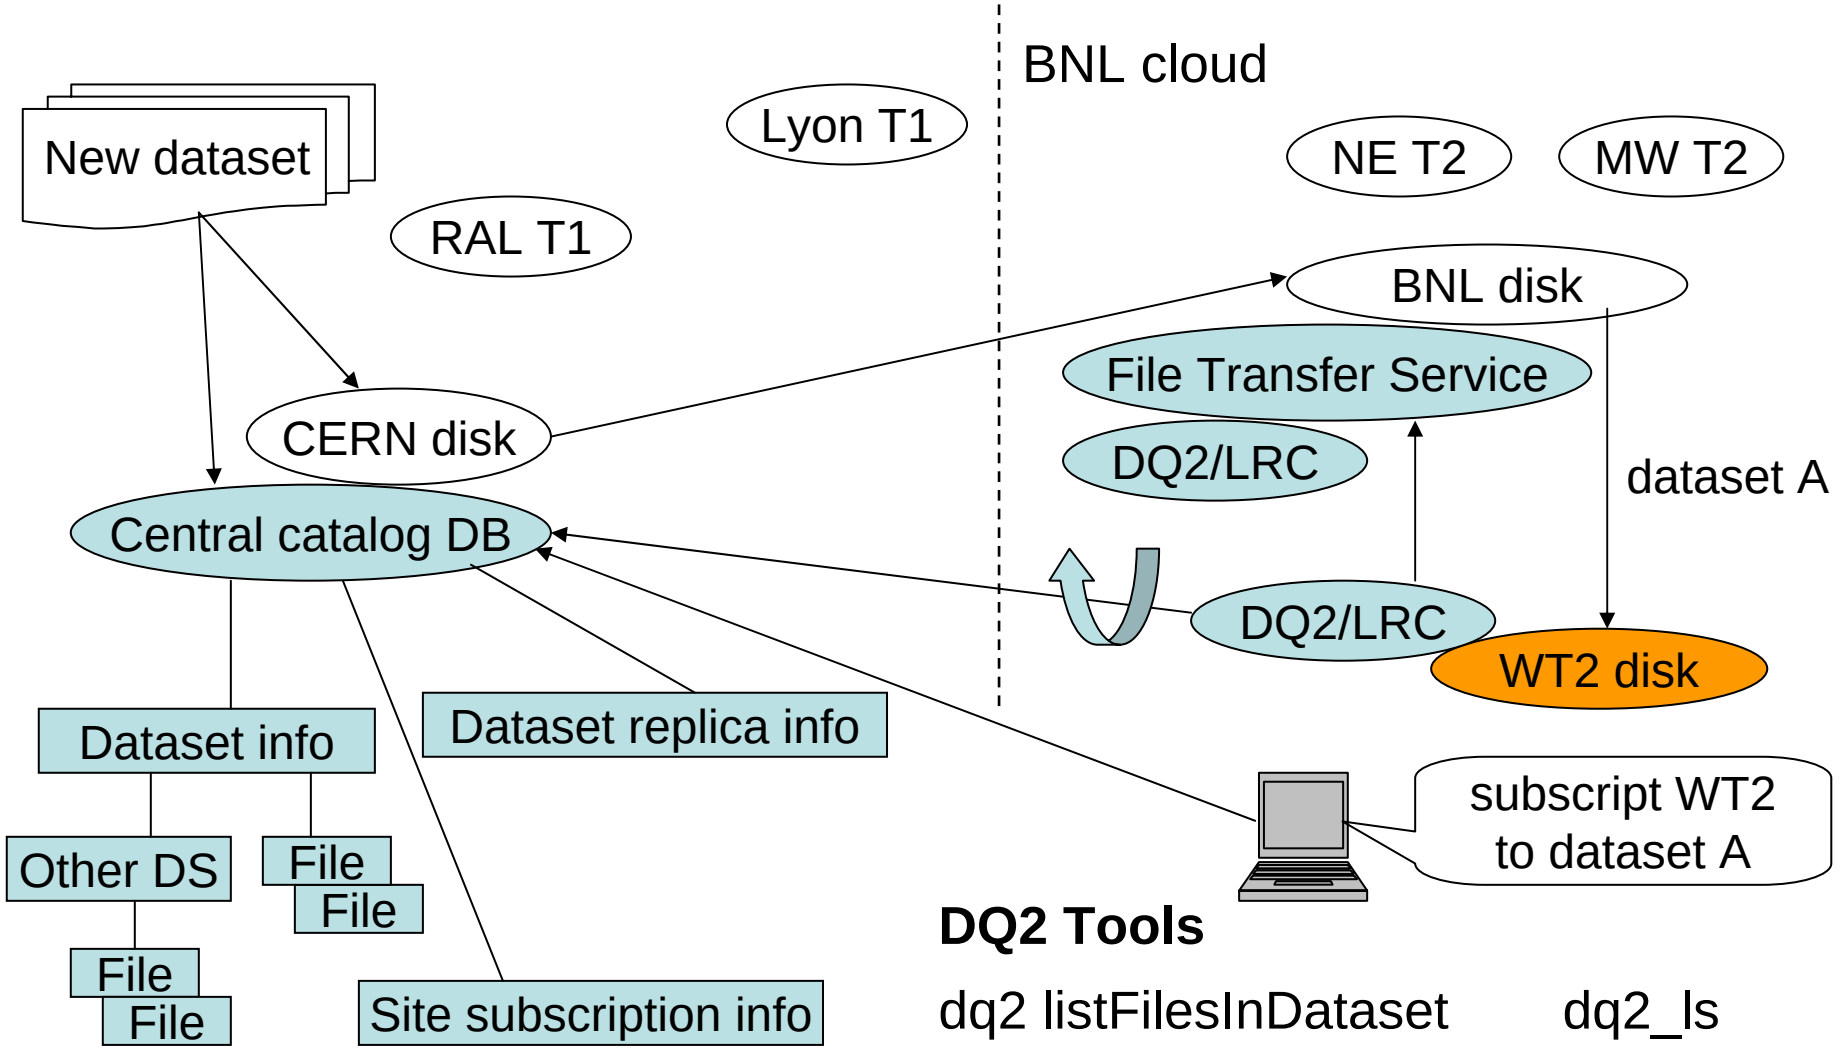

DQ2 Tools

dq2 listFilesInDataset	dq2_ls
dq2 listDatasetReplicas	dq2_get
dq2 listDatasetsInSite	

How local users access datasets

- Use the “dq2” command:
  - Check files in dataset
  - Which site has interesting datasets
  - Go to that site and run jobs

  OR

- dq2 client tools to get dataset to your own storage
  - dq2_ls
  - dq2_get
  - Run jobs on your laptop

Panda

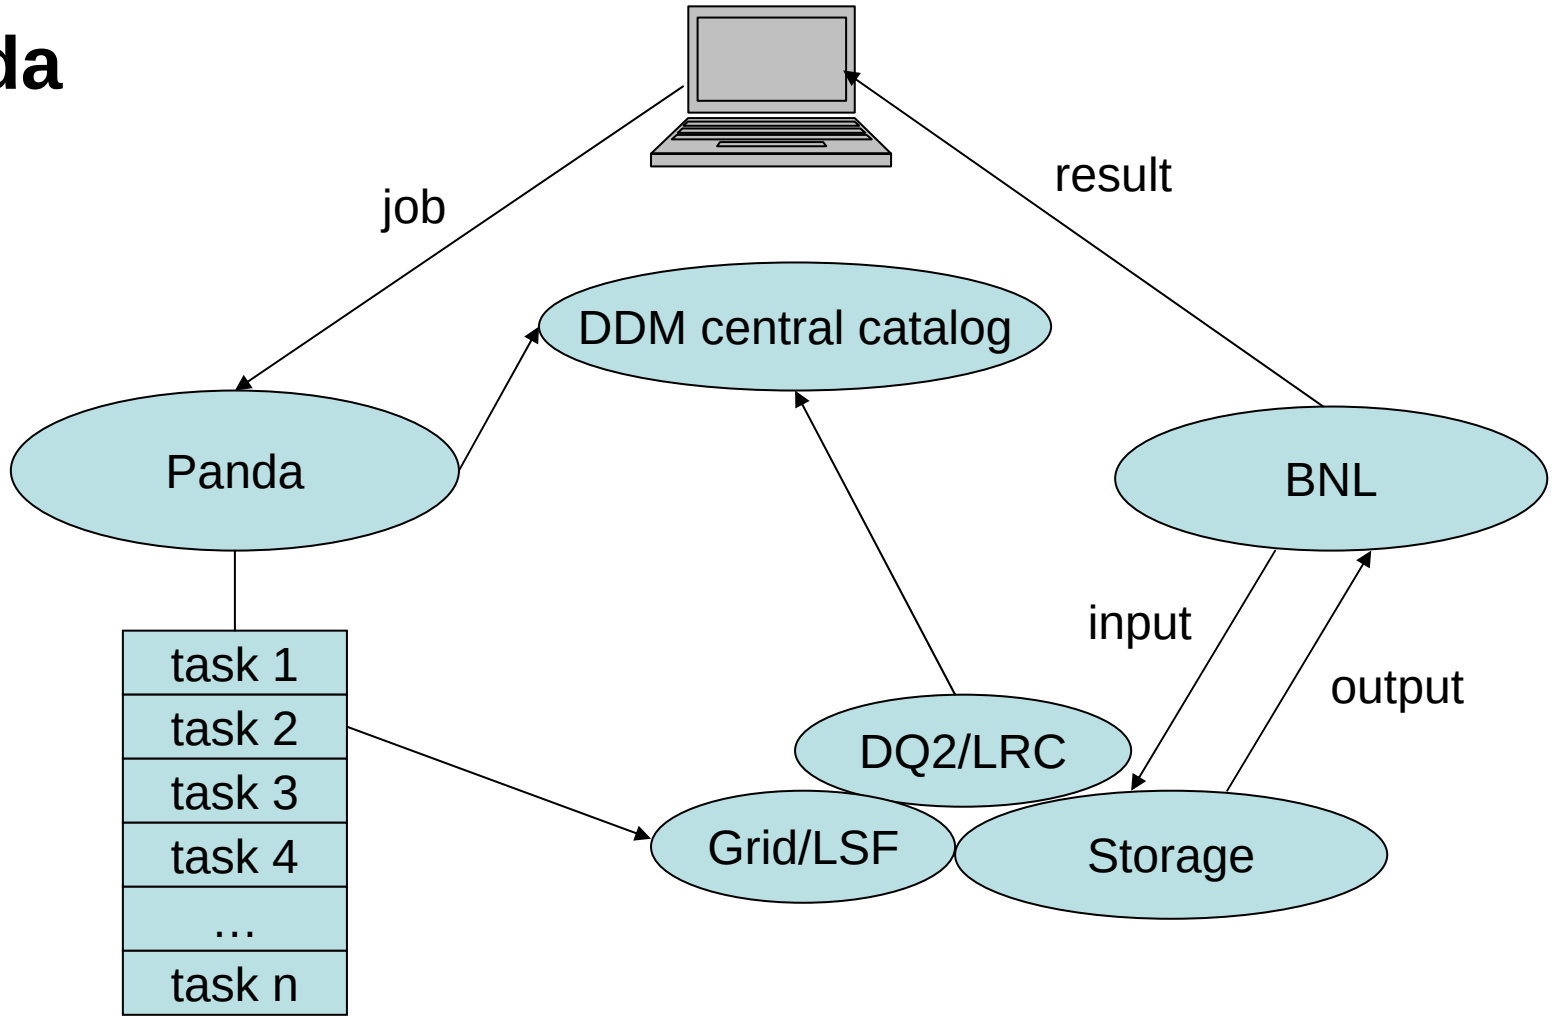

- Panda uses DDM
- Panda is a US ATLAS production system
- Panda supports distributed analysis via grid

Current Resource at Western Tier 2

- 7 % fair share of SLAC “shared” LSF batch resource for local ATLAS users
- 2.4 TB NFS space for user working directories
- AFS space for ATLAS release
- Pacman mirror for ATLAS software

- Grid infrastructure. Gatekeeper/gridftp, GUMS
- DQ2 site service
- 2.3 TB NFS space (NetApp) for DQ2 repository
- 500 GB NFS space for ATLAS production software
- 7 % LSF fair share for production

None from ATLAS funding. ATLAS equipments will arrive in March/April

Architecture

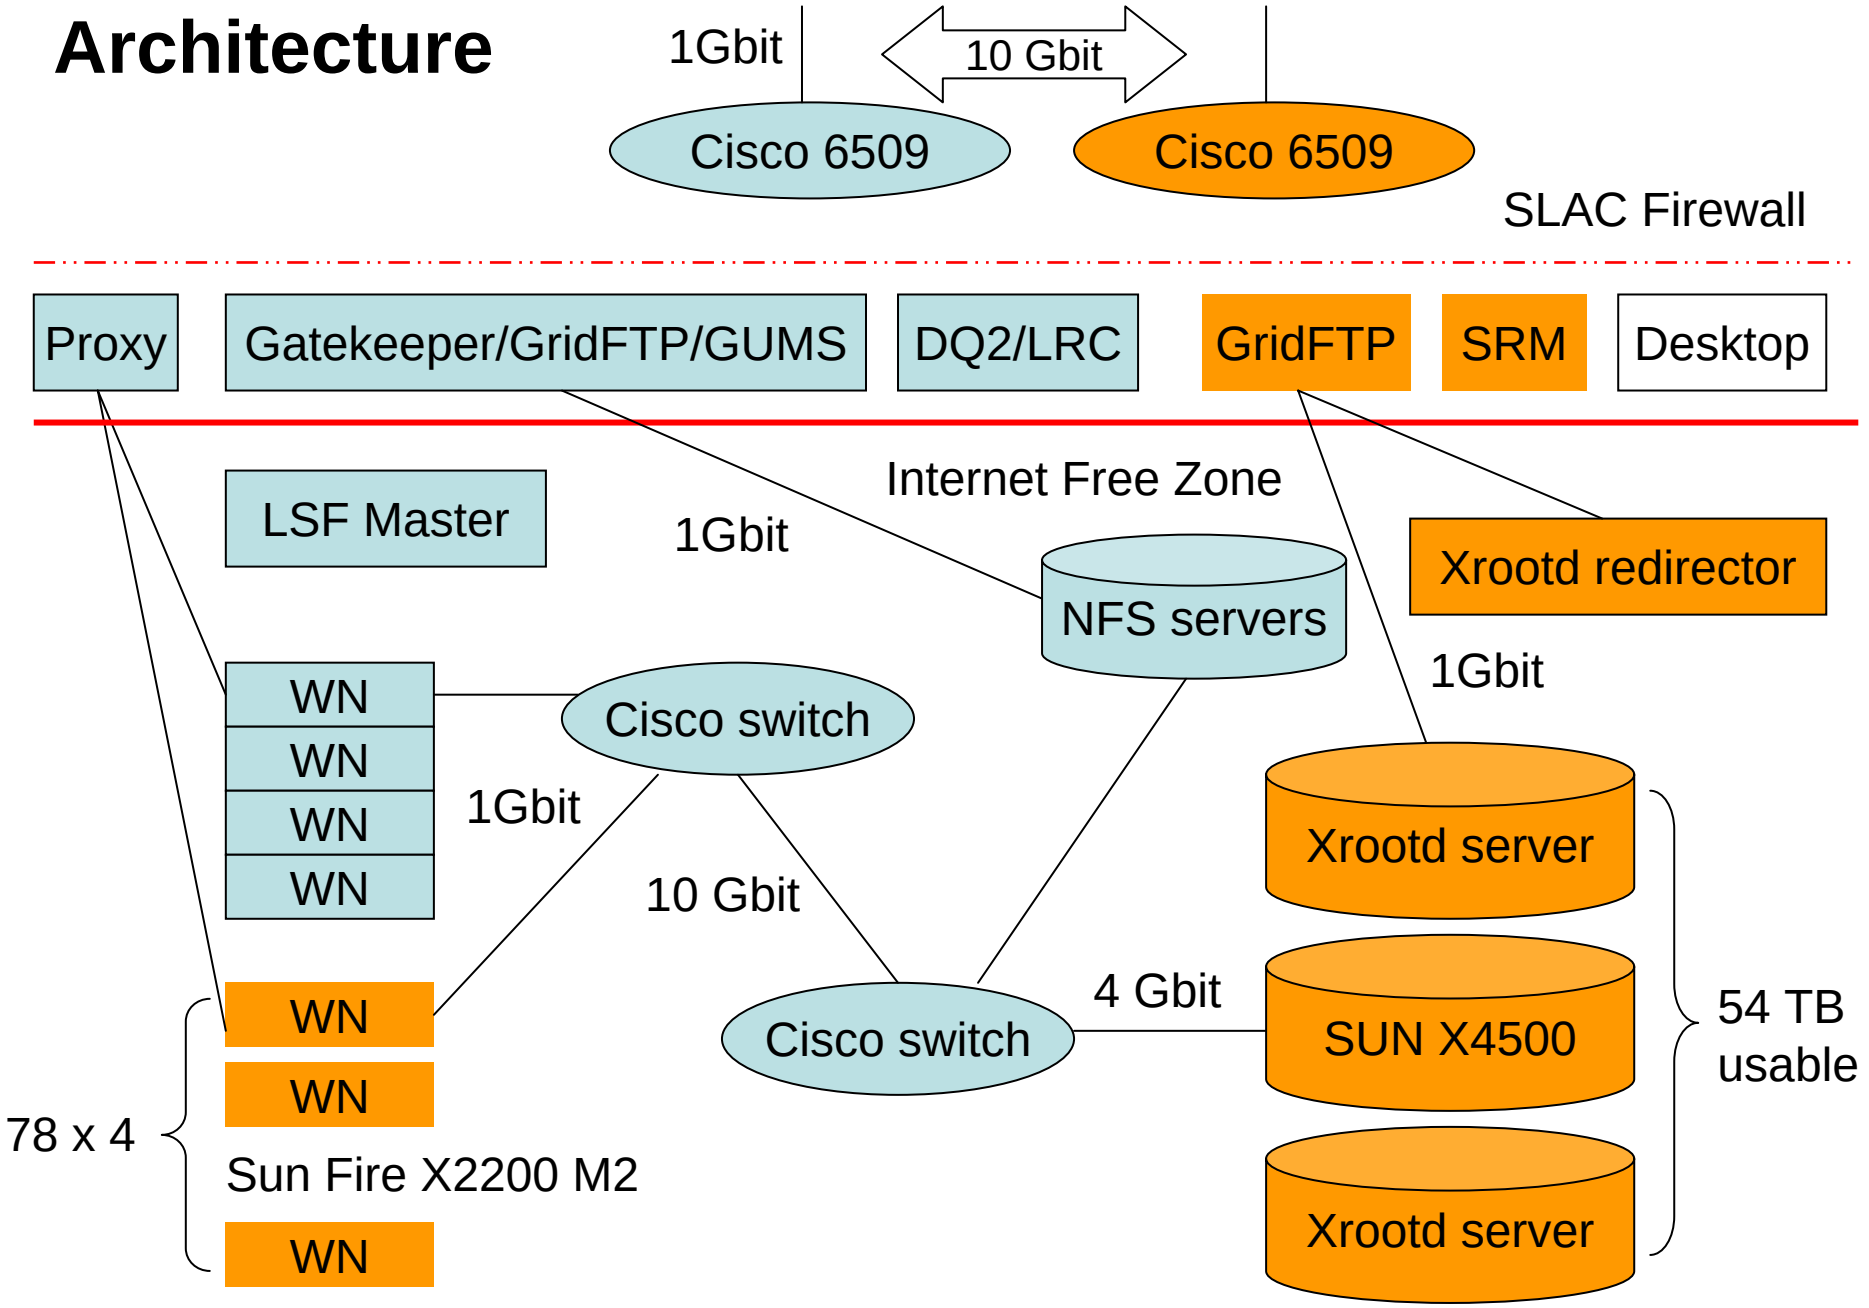

Western Tier 2 Plan

Software:

- RedHat Enterprise

  RHEL 3 32 bit, RHEL 4 32/64 bit

- LSF 6.1

  Many years of experience with LSF

- Xrootd

  Expertise in Xrootd
  Experience of real usage by Babar

Western Tier 2 Plan (cont'd)

Hardware:

- Expect ATLAS equipment to arrive at March/April
- 312 CPU cores in 78 node

  Sun X2200 M2 AMD Opteron 2218, 4 core, 8GB memory, 2 x 250GB
- 54 TB usable storage in 3 Thumper Boxes

  Sun X4500, 18 usable TB (48 SATA drive),
  4 AMD Opteron CPU core, 16GB, 4 x 1Gbit,
  Solaris 10 x86 and ZFS
- 10 Gbit external network

Resource

Revised hardware plan for CPU and storage

	2006 - 2007	2007-
CPU $ ratio	2/3	1/3
Storage $ ratio	1/3	2/3

	Projected	Purchased	Current
CPU (kSI2K)	346	421	
Disk (usable TB)	54	54	5.2

Current Resource Allocation

CPU	Production	50%	7% LSF shares
	Local user	50%	7% LSF shares
Disk	Production data	44%	2.3 TB
	Production software	10%	0.5 TB
	Local user	46%	2.4 TB

Future resource allocation is determined by ATLAS management and Western Tier 2 Advisory Committee

Resource Utilization

78 x 4 x 1.35 kSI2K	**421 kSI2K**	**→**	**6.0 x 10^4 CPU hours / day**
½ for production	**210 kSI2K**	**→**	**3.0 x 10^4 CPU hours / day**
Future resource ↑		↓	*Current resource*
Best day in February	**239 kSI2K**	**←**	**3.4 x 10^4 CPU hours / day**
Average in February	**91 kSI2K**	**←**	**1.3 x 10^4 CPU hours / day**

- CPU hours in SLAC unit
- Utilization is low due to many reasons
- The above numbers don't tell us how to improve

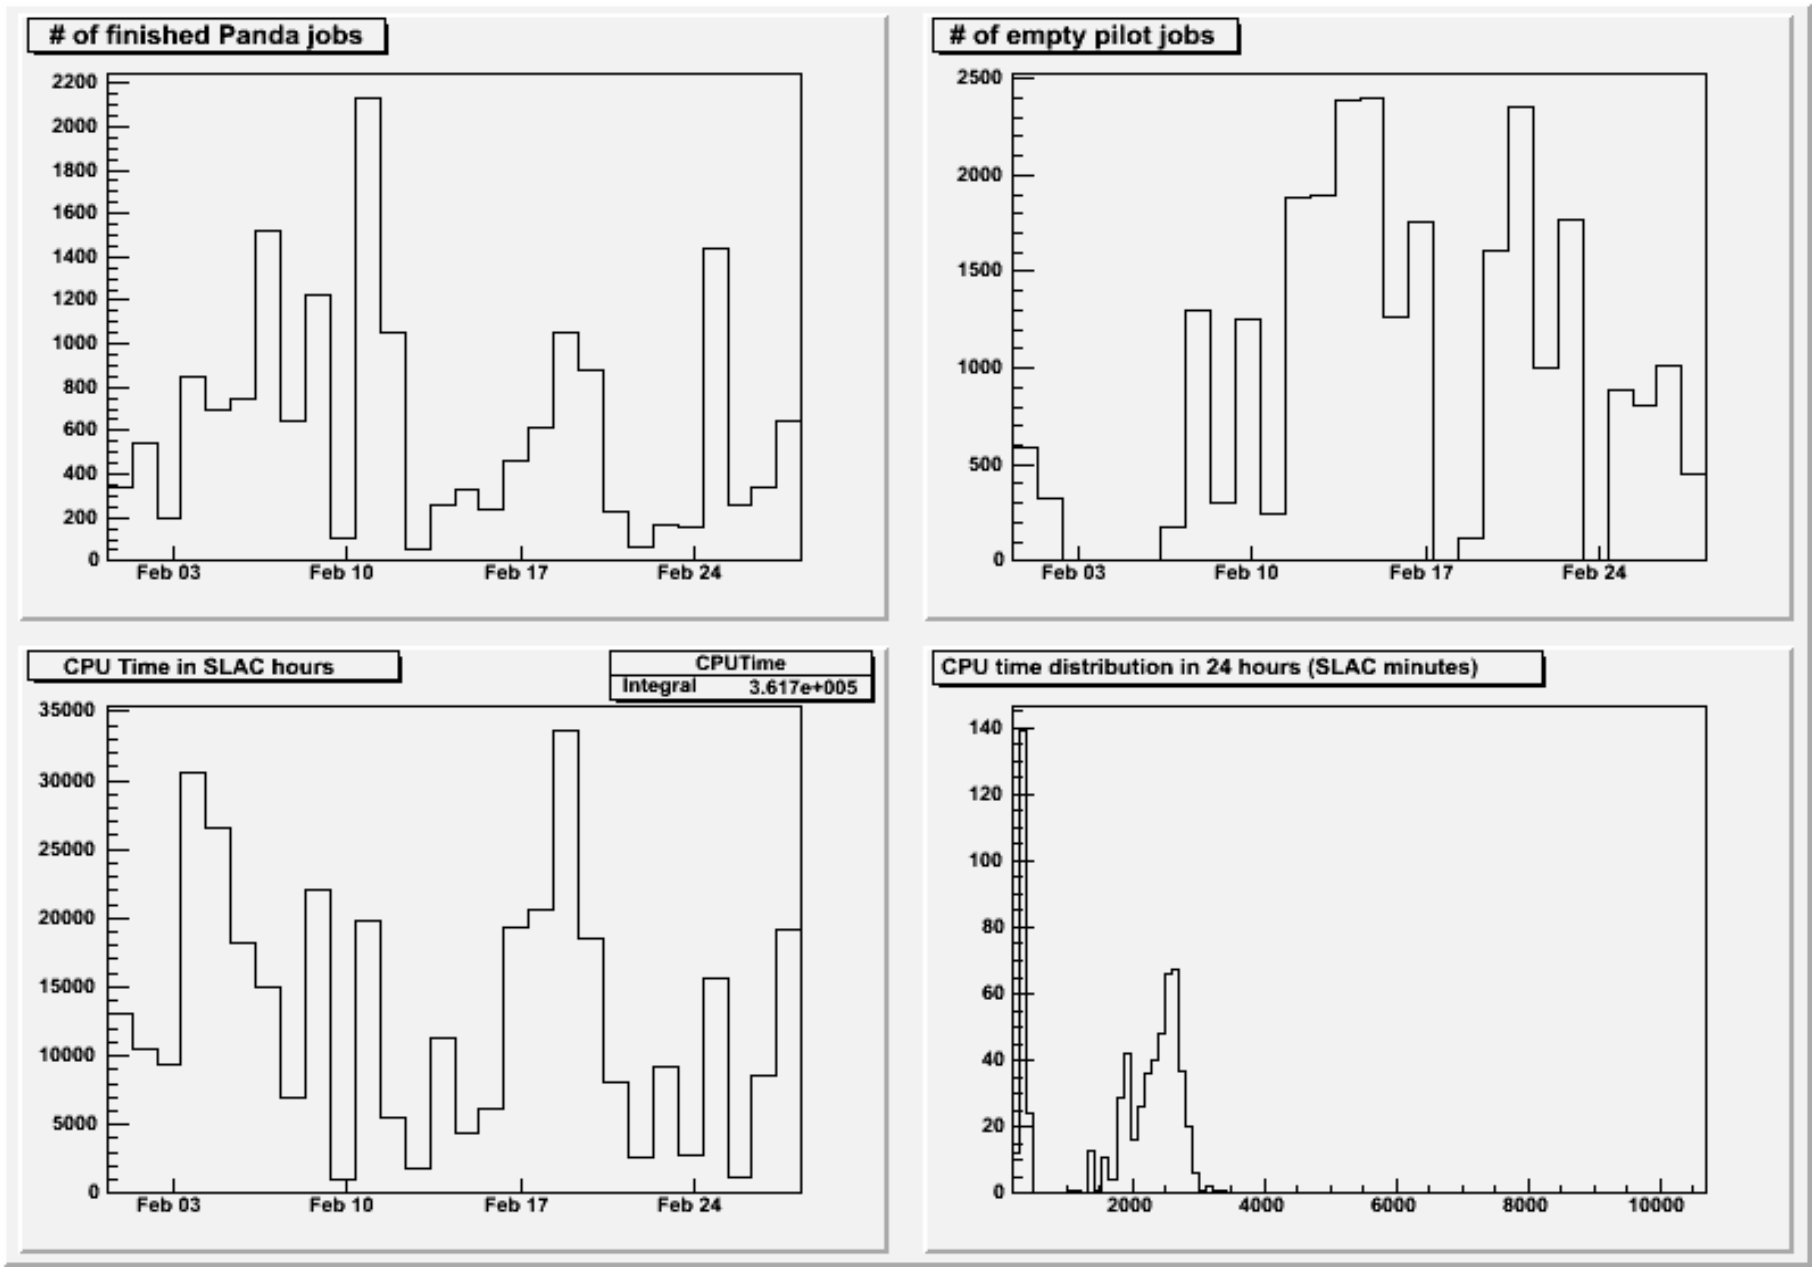

Pictures at left column: **Monthly average is more accurate than single day**

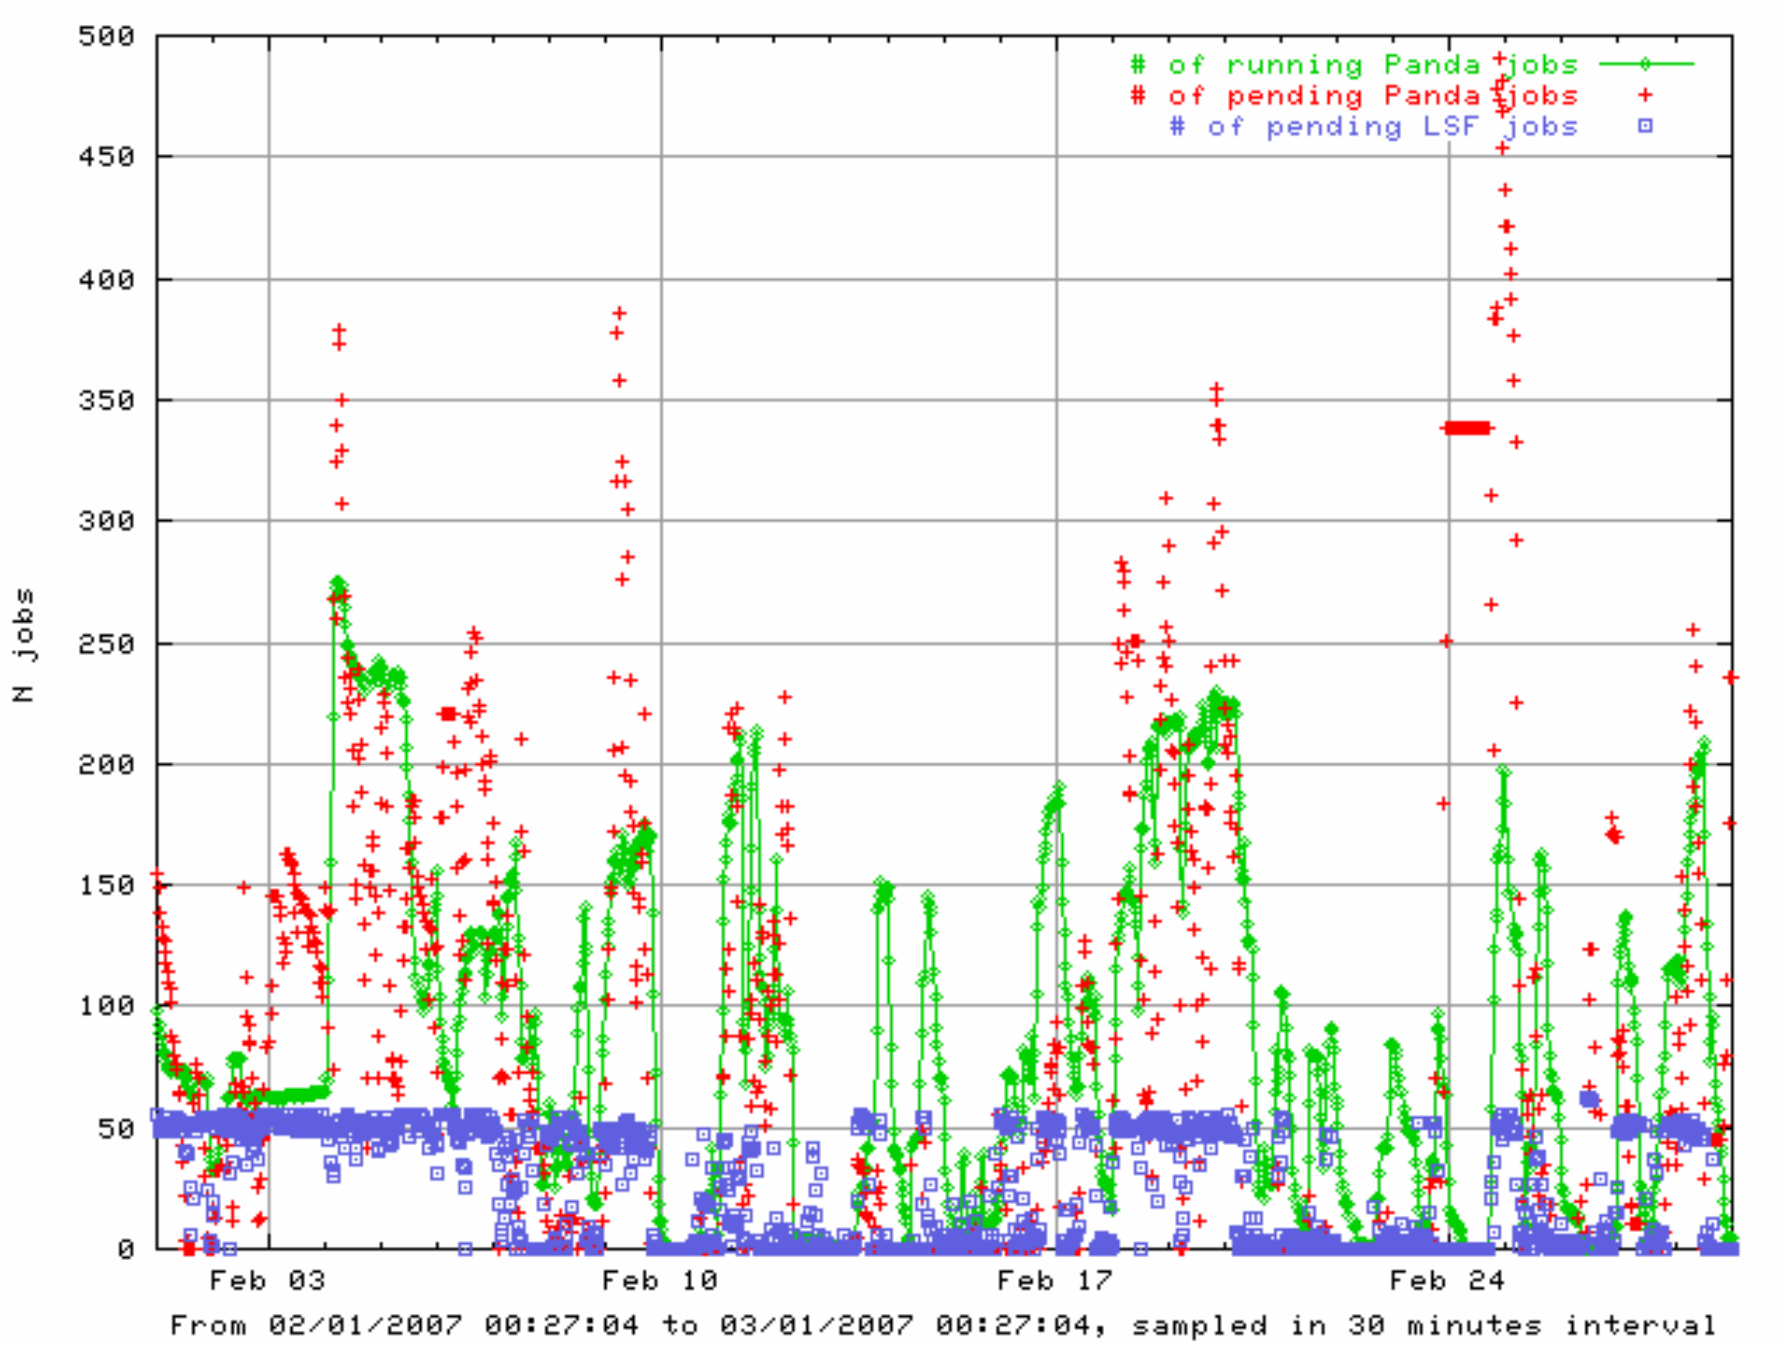
# of running Panda jobs
# of pending Panda jobs
# of pending LSF jobs
N jobs
500
450
400
350
300
250
200
150
100
50
0
Feb 03
Feb 10
Feb 17
Feb 24
From 02/01/2007 00:27:04 to 03/01/2007 00:27:04, sampled in 30 minutes interval

Recent Activities at Western Tier 2

- Developed a GridFTP Data System Interface for Xrootd
- Installed a DQ2 0.2.12 site services with Xrootd backend
- Working with Panda term to use this new DQ2 site and Xrootd storage
  - Can Pathena job ***READ*** root files from Xrootd server directly ?
  - Can local Athena job ***READ*** root files from Xrootd server ?
- Tier 2 hardware selection and purchasing order completed

Issues

- Power
  - 18 TB will be online in April. High priority
  - CPU and the rest storage in racks in June/July, w/o power
  - 10 Gbit external network installed, waiting for power
  - Hope to solve the current power crisis in late summer
- Can dq2 client tools work with xrootd storage ?
  or any other non-UNIX/NFS storage ?
- DDM software is better than before, but still have problems