Status report from Tokyo Tier-2 for the one year operation after whole scale system upgrade

Tomoaki Nakamura

on behalf of ICEPP regional analysis center group

ICEPP, The University of Tokyo

東京大学 THE UNIVERSITY OF TOKYO

ICEPP regional analysis center

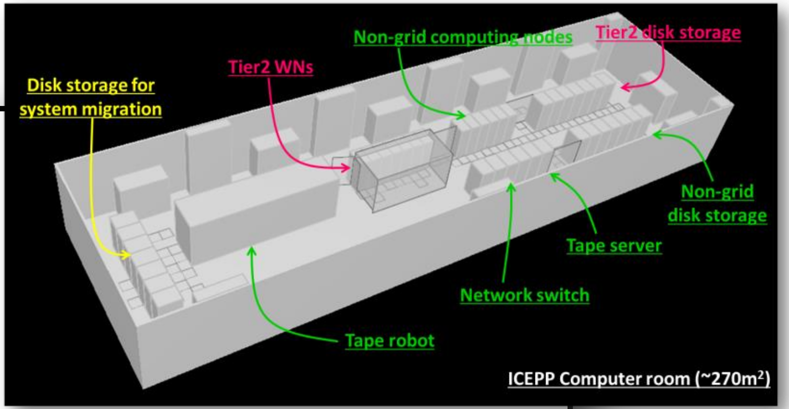

Resource overview

Support only ATLAS VO in WLCG as Tier2.
Provide ATLAS-Japan dedicated resource for analysis.
The first production system for WLCG was deployed in 2007.
Almost of hardware are prepared by three years lease.
System have been upgraded in every three years.
Current system (replaced end of 2012) is the 3rd generation system.

Single VO and Simple and Uniform architecture → 95% availability in 2013.
No major upgrade in terms of hardware for three years operation.
Precise evaluation and optimization of the performance → Improvement of job throughput.
Design for the next system based on the real use case.

Dedicated staff

Tetsuro Mashimo (associate prof.):	fabric operation, procurement, Tier3 support
Nagataka Matsui (technical staff):	fabric operation, Tier3 support
Tomoaki Nakamura (project assistant prof.):	**Tier2 operation**, Tier3 analysis environment
Hiroshi Sakamoto (prof.):	site representative, coordination, ADCoS
Ikuo Ueda (assistant prof.):	ADC coordinator, site contact with ADC
System engineer from company (2FTE):	fabric maintenance, system setup

Removing almost HWs in two days (Dec. 2012)

3rd system (2013 - 2015)

Migration period (Dec. 2012 – Jan 2013)

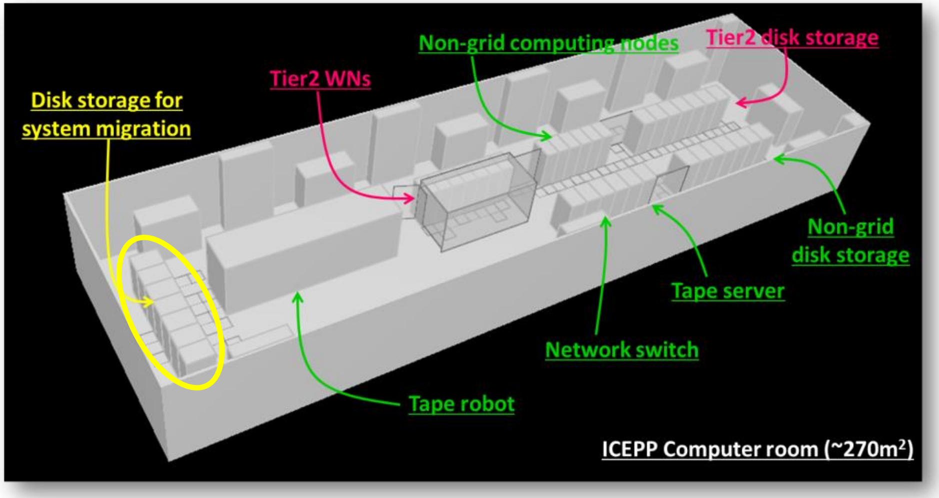

Reduced number of the WNs (32 nodes) and all service instance were operated by using old hardware to minimize the downtime.

1PB of data already stored at Tokyo was copied to the temporal storage a priori so that Grid user could access even in the migration period.

All data was transferred back to the new production storage by retrying "rsync" several times in a few weeks without long downtime.

The number of WNs are gradually increased. Full operation with the new system have been started in Feb. 2013.

Major upgrade (Dec. 2012)

		2nd system (2010-2012)	3rd system (2013-2015)
Computing node	Total	Node: **720 nodes, 5720 cores** (including service nodes) CPU: Intel Xeon X5560 (Nehalem 2.8GHz, 4 cores/CPU)	Node: **624 nodes, 9984 cores** (including service nodes) CPU: Intel Xeon E5-2680 (Sandy Bridge 2.7GHz, 8cores/CPU)
	Non-grid (Tier3)	Node: 96 (496) nodes, 768 (3968) cores Memory: 16GB/node NIC: 1Gbps/node Network BW: 20Gbps/16 nodes Disk: 300GB SAS x 2	Node: 416 nodes, 6656 cores Memory: 16GB/node (to be upgraded) NIC: 10Gbps/node Network BW: 40Gbps/16 nodes Disk: 600GB SAS x 2
	Tier2	Node: **464 (144) nodes, 3712 (1152) cores** Memory: 24GB/node NIC: 10Gbps/node Network BW: 30Gbps/16 nodes Disk: 300GB SAS x 2	Node: **160 nodes, 2560 cores** Memory: 32GB/node (to be upgraded) NIC: 10Gbps/node Network BW: 80Gbps/16 nodes Disk: 600GB SAS x 2
Disk storage	Total	Capacity: **5280TB** (RAID6) Disk Array: 120 units (HDD: 2TB x 24) File Server: 64 nodes (blade) FC: 4Gbps/Disk, 8Gbps/FS	Capacity: **6732TB** (RAID6) + α Disk Array: 102 (3TB x 24) File Server: 102 nodes (1U) FC: 8Gbps/Disk, 8Gbps/FS
	Non-grid (Tier3)	Mainly NFS	Mainly GPFS
	Tier2	DPM: **1.36PB**	DPM: **2.64PB**
Network bandwidth	LAN	10GE ports in switch: 192 Switch inter link: 80Gbps	10GE ports in switch: 352 Switch inter link : 160Gbps
	WAN	ICEPP-UTnet: 10Gbps (+10Gbps) SINIET-USA: 10Gbps x 2 ICEPP-EU: 10Gbps	ICEPP-UTNET: 10Gbps SINET-USA: 10Gbps x 3 ICEPP-EU: 10Gbps (+10Gbps) LHCONE

Fraction of completed ATLAS jobs at Tokyo Tier2

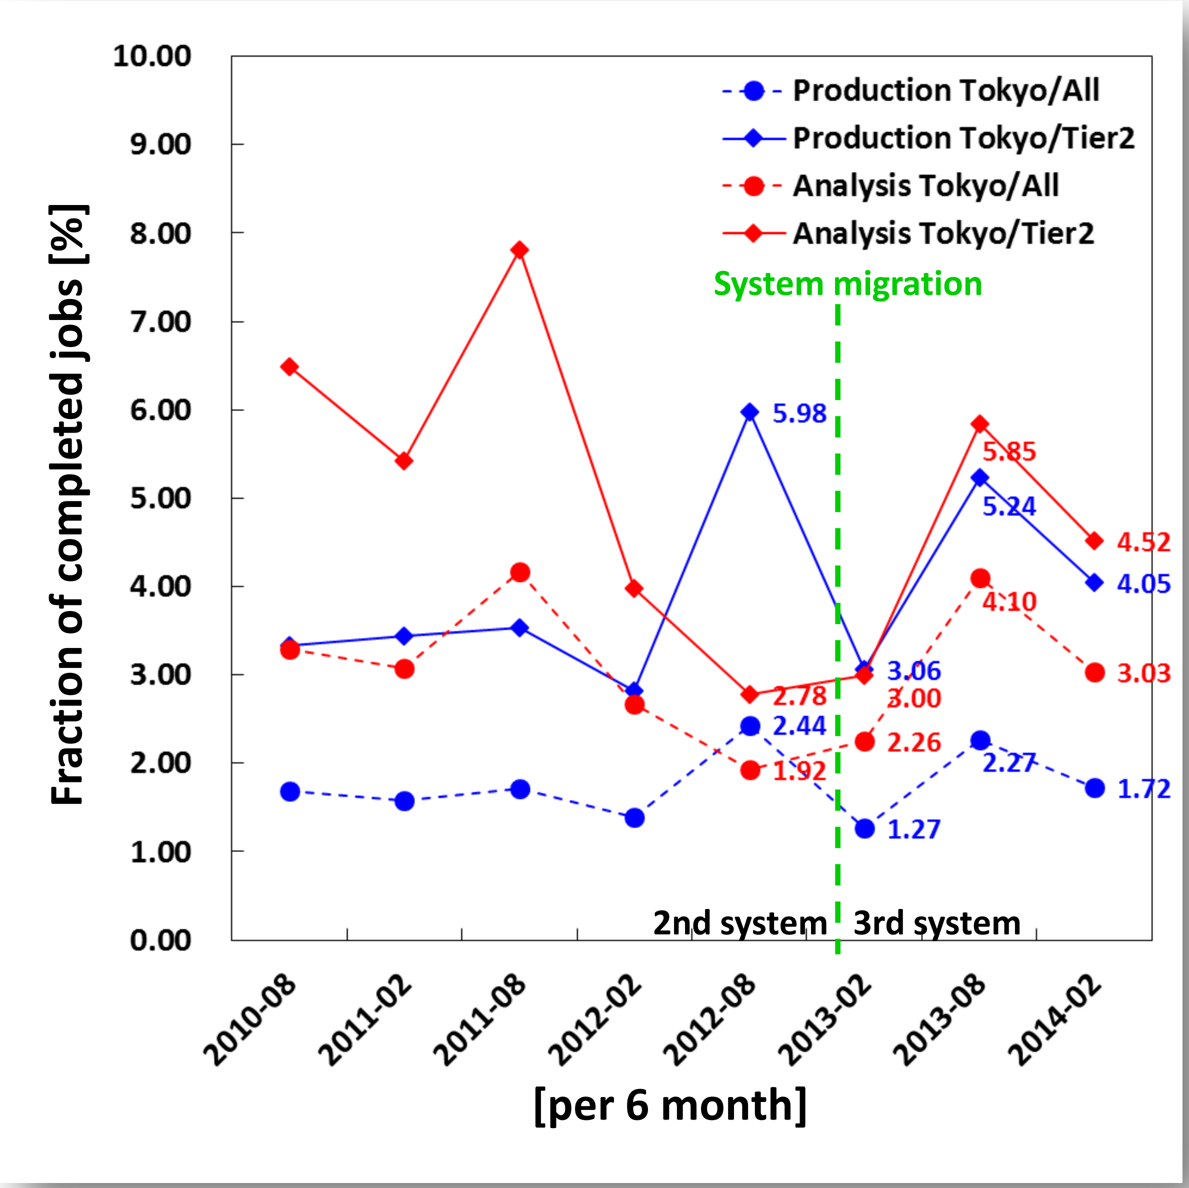

Evolution of disk storage capacity for Tier2

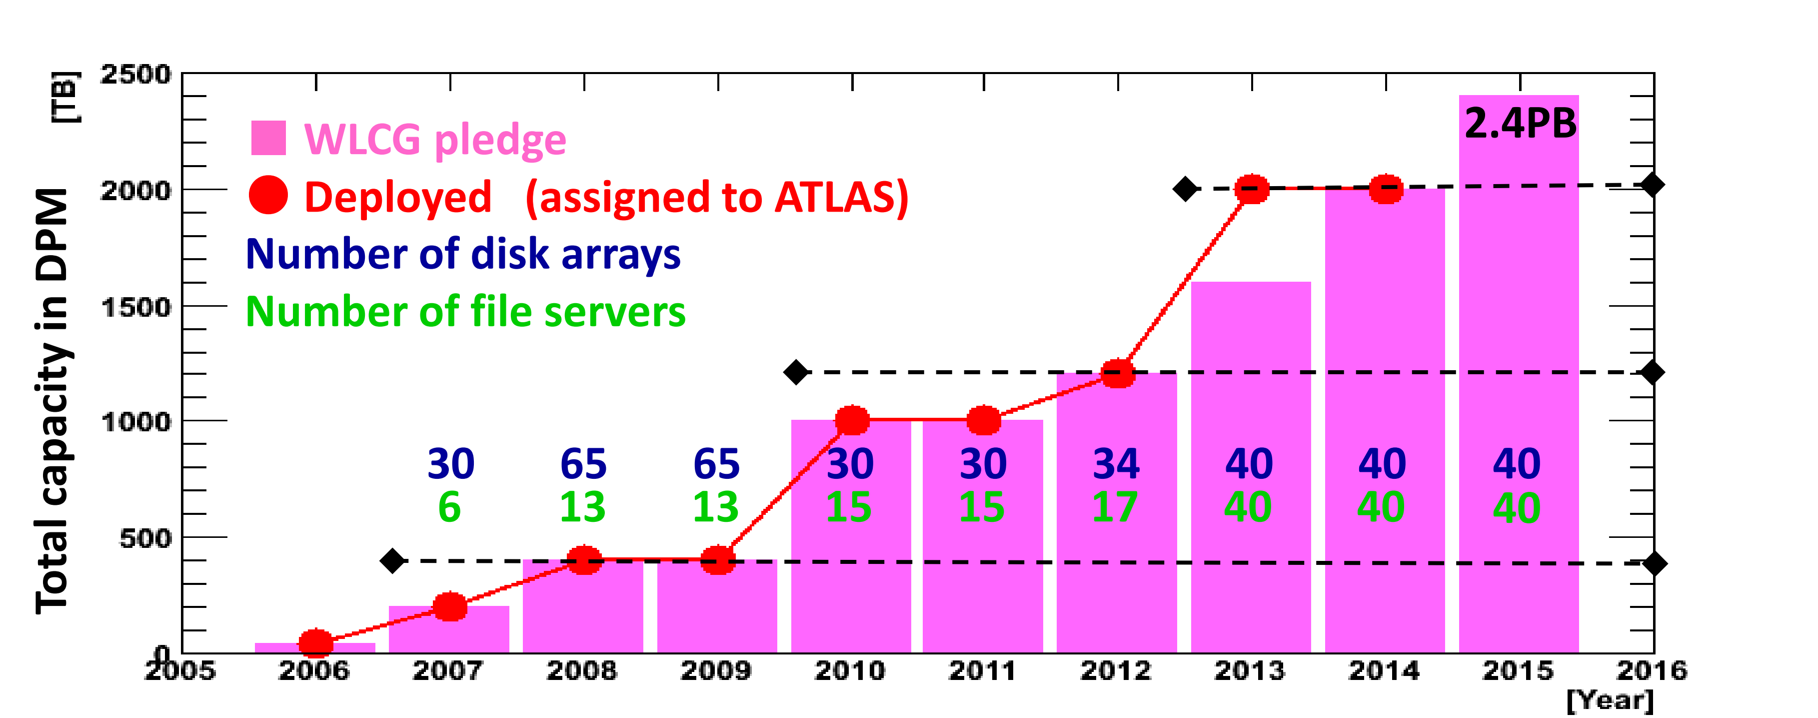

Pilot system for R&D	1st system 2007 - 2009	2nd system 2010 - 2012	3rd system 2013 - 2015	4th system 2016 - 2019
	16x500GB HDD / array	**24x2TB HDD / array**	**24x3TB HDD /array**	5TB HDD?
	5disk arrays / server	**2disk arrays / server**	**1disk array / server**	
	XFS on RAID6	**XFS on RAID6**	**XFS on RAID6**	
	4G-FC via FC switch	**8G-FC via FC switch**	**8G-FC without FC switch**	
	10GE NIC	**10GE NIC**	**10GE NIC**	

ATLAS disk and LocalGroupDisk in DPM

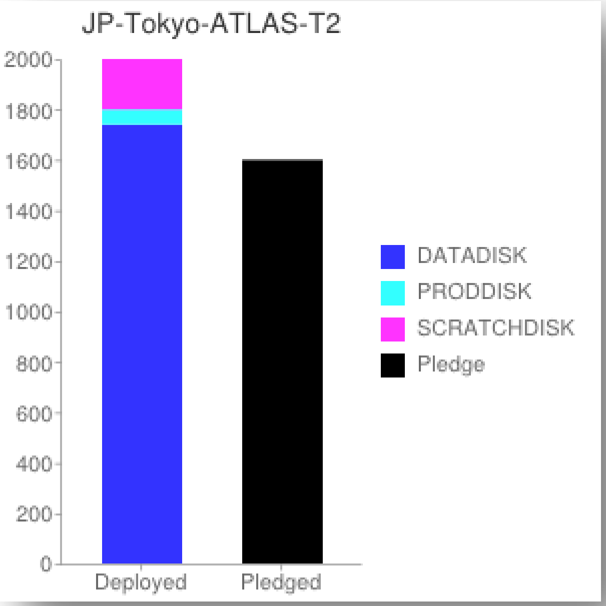

2014's pledge deployed since Feb. 20, 2013

DATADISK: 1740TB (including old GroupDisk, MCDisk, HotDisk)
PRODDISK: 60TB
SCRATCHDISK: 200TB

Total: 2000TB

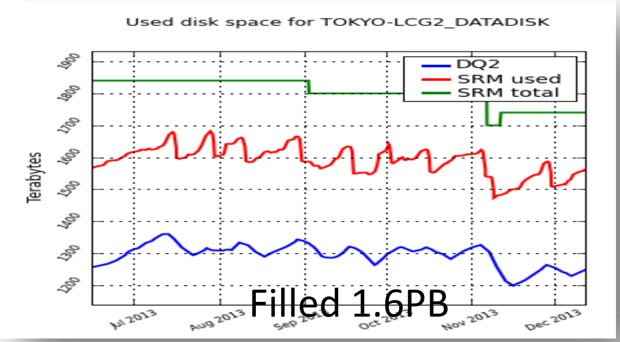

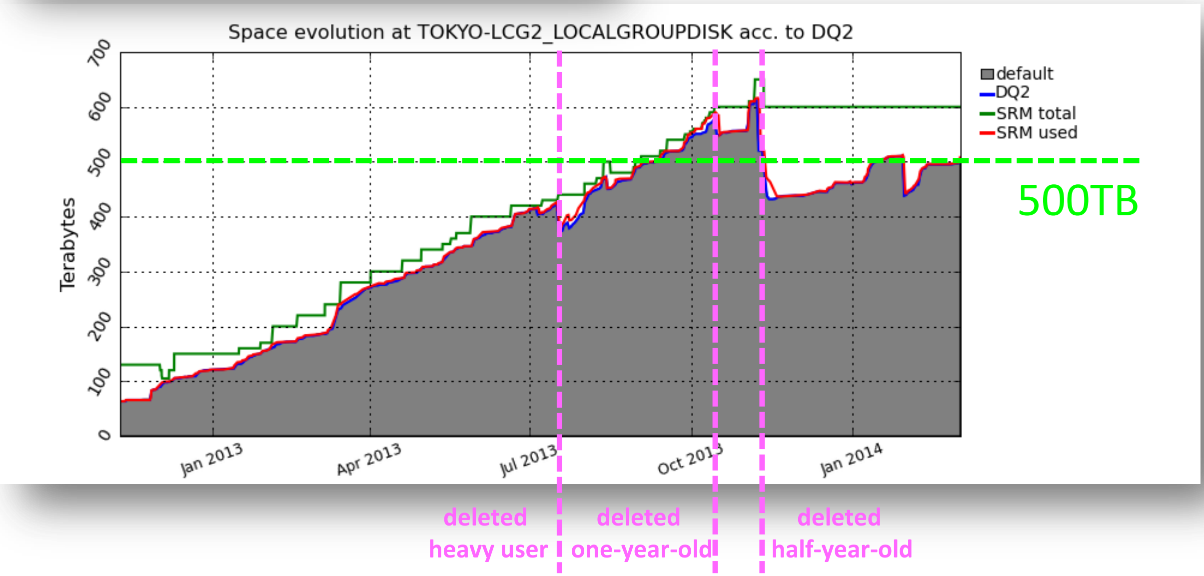

Keep less than 500TB

Manual deletion for several times.

It will be mitigated by the Grid-wide user quota in new dataset catalog.

Rucio

DPM upgrade on scalability and maintainability

DPM is most widely deployed grid storage system
- ~200 sites in 50 regions
- over 300 VOs
- ~45 PB (10 sites with > 1PB)

Scalability

Tokyo is one of the largest sites on DPM capacity with one pool.
10M entries are contained in MySQL-DB for file and directory (size ~50GB).
Performance can be scaled by adding memory (64GB is already added).
But, DB size is still glowing, we will increased the memory to 128GB.

Maintainability

Fusion-IO drive will be attached for DB backup space to reduce the time for maintenance i.e. dump/restore.

- NAND flush memory directly connected via PCI-E
- 1.2TB capacity (MLC)
- I/O: 1.5GB/sec (read), 1.3GB/sec (write)
- IOPS: 270k (read), 800k (write)
- cost: ~12kCHF

Planning to have a redundant configuration to take a daily backup without downtime.

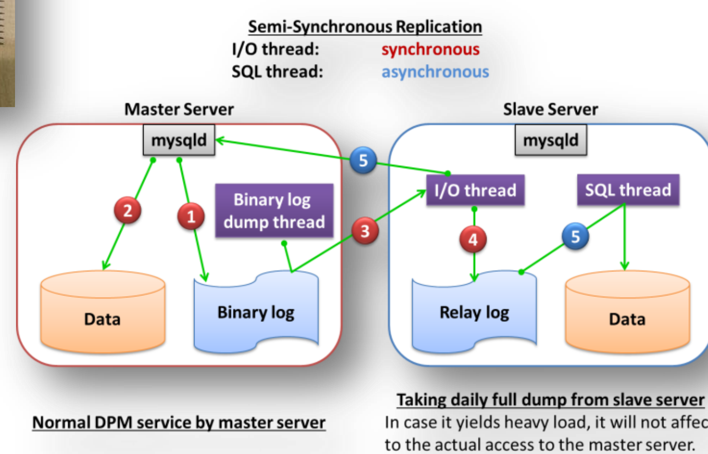

Minor upgrade

CPU performance by the OS upgrade

Migration to SL6 WN was completed at Oct. 2013.

It was done by the rolling transition making TOKYO-SL6 queue to minimize the downtime and risk hedge for the massive miss configuration .

Performance improvement

5% increased as HepSpec06 score

SL5, 32bit compile mode: 17.06±0.02

SL6, 32bit compile mode: 18.03±0.02 (still fastest site in ATLAS)

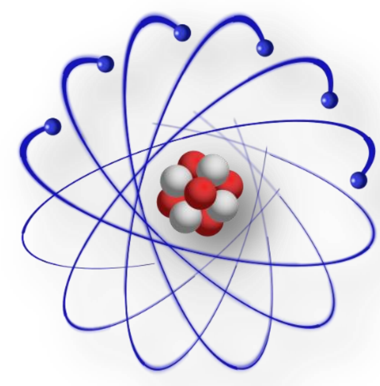

Memory capacity build up

Increased from 2GB/core to 4GB/core for half of the WNs on Nov. 6, 2013

Useful for memory consuming ATLAS production jobs.

Available sites are rare.

ATLAS Grid Information System

Site (OIM/GOCDB): TOKYO-LCG2 **RC:** JP-Tokyo-ATLAS-T2 **Cloud:** FR **State:** ACTIVE **VO:** atlas

CE AGIS name	Endpoint	Status	AGIS status	Queue	Panda Queues
TOKYO-LCG2-CE-lcg-ce02.icepp.jp	lcg-ce02.icepp.jp:8443	production	ACTIVE	atlas default	ANALY_TOKYO_HIMEM TOKYO-LCG2-all-ce-atlas-lcgpbs TOKYO_HIMEM
TOKYO-LCG2-CE-lcg-ce01.icepp.jp	lcg-ce01.icepp.jp:8443	production	ACTIVE	atlas default	ANALY_TOKYO TOKYO-LCG2-all-ce-atlas-lcgpbs

Center network switch configuration

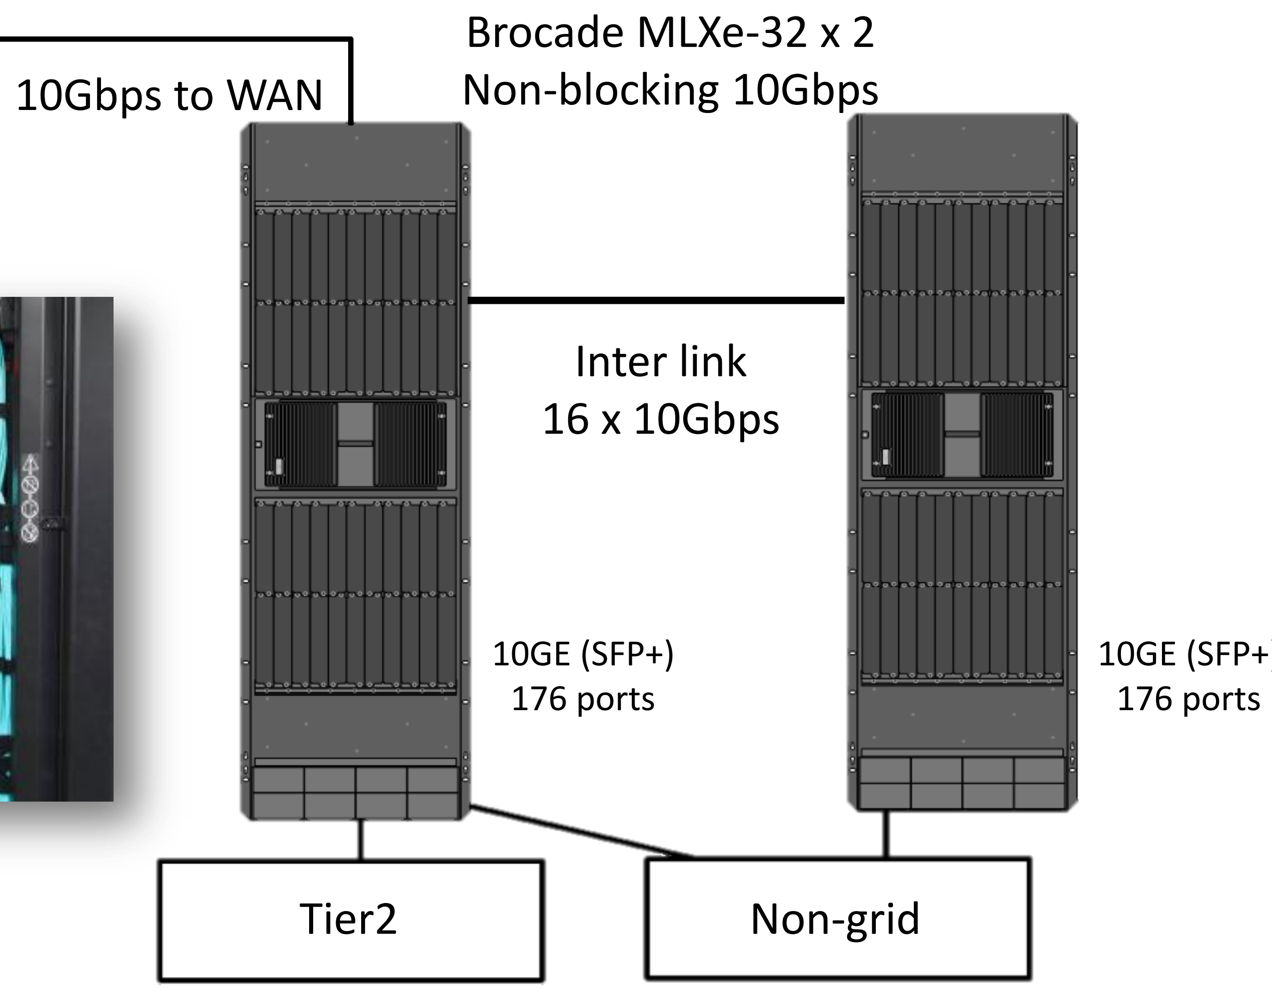

DPM file servers
LCG service nodes
LCG worker nodes

GPFS/NFS file servers
Tape servers
Non-grid service nodes
Non-grid computing nodes

Network bandwidth for WNs

Disk server x40

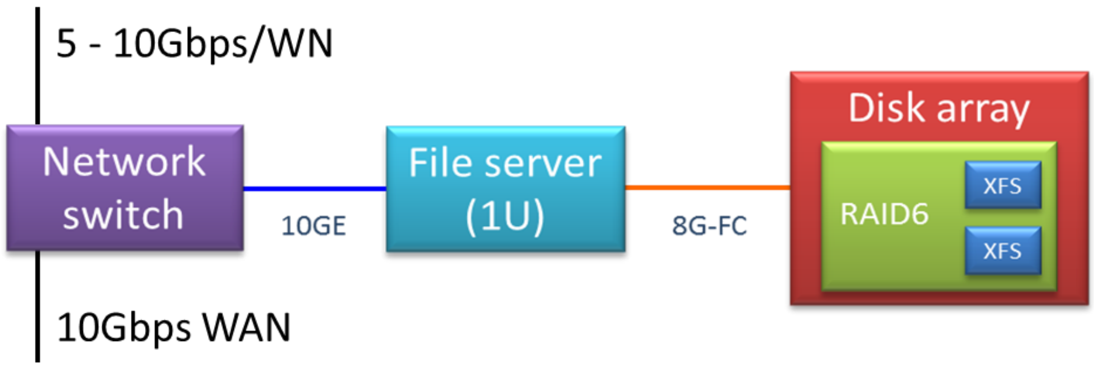

- 66TB x 40 servers
- Total capacity 2.64PB (DPM)
- 10Gbps NIC (for LAN)
- 8G-FC (for disk array)
  500~700MB/sec
  (sequential I/O)

Worker node x160

- 10Gbps pass through module
  SFP+ TwinAx cable
- Rack mount type 10GE switch (8024F)
  10G BASE SR (SFP+)
- Bandwidth
  80Gbps/16nodes
  minimum 5Gbps
  maximum 10Gbps

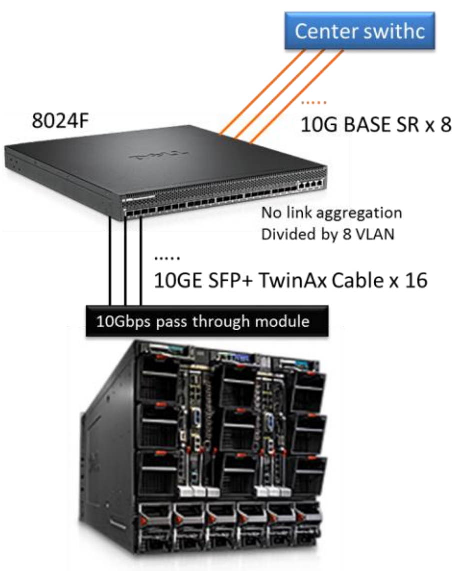

Performance of internal network

(a)

WAN
10Gbps

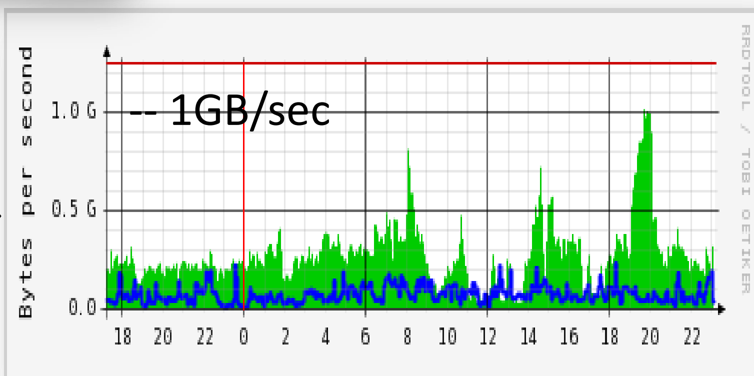

Center network switch
Brocade MLXe-32
non-blocking 10Gbps

(b)

File servers
10GE x 40 nodes
8G-FC with disk array (66TB)

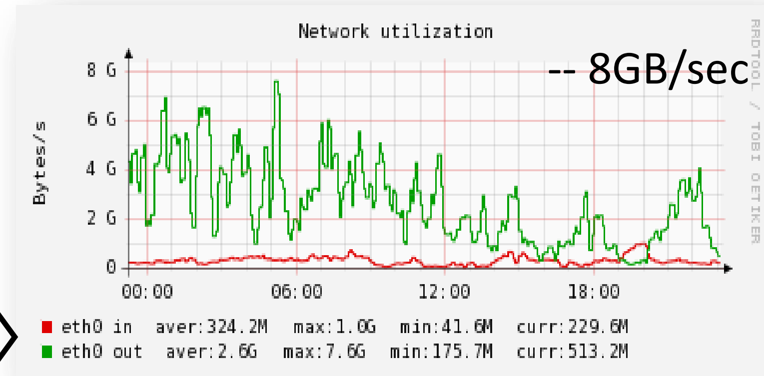

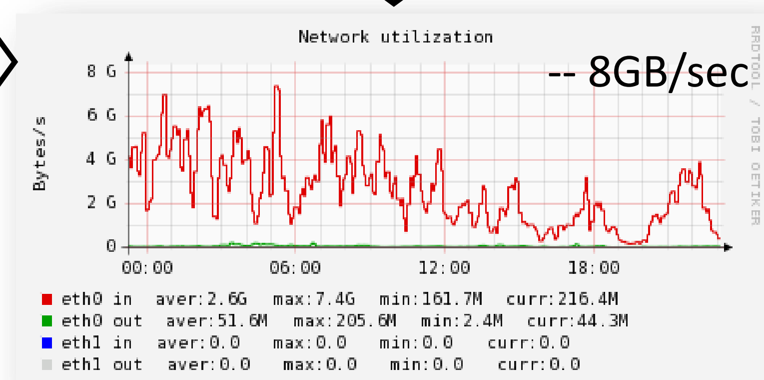

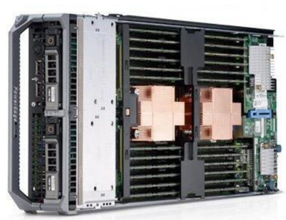

(c)

Worker nodes
10GE x 160 nodes
(minimum 5Gbps/node)

International connection to Tokyo

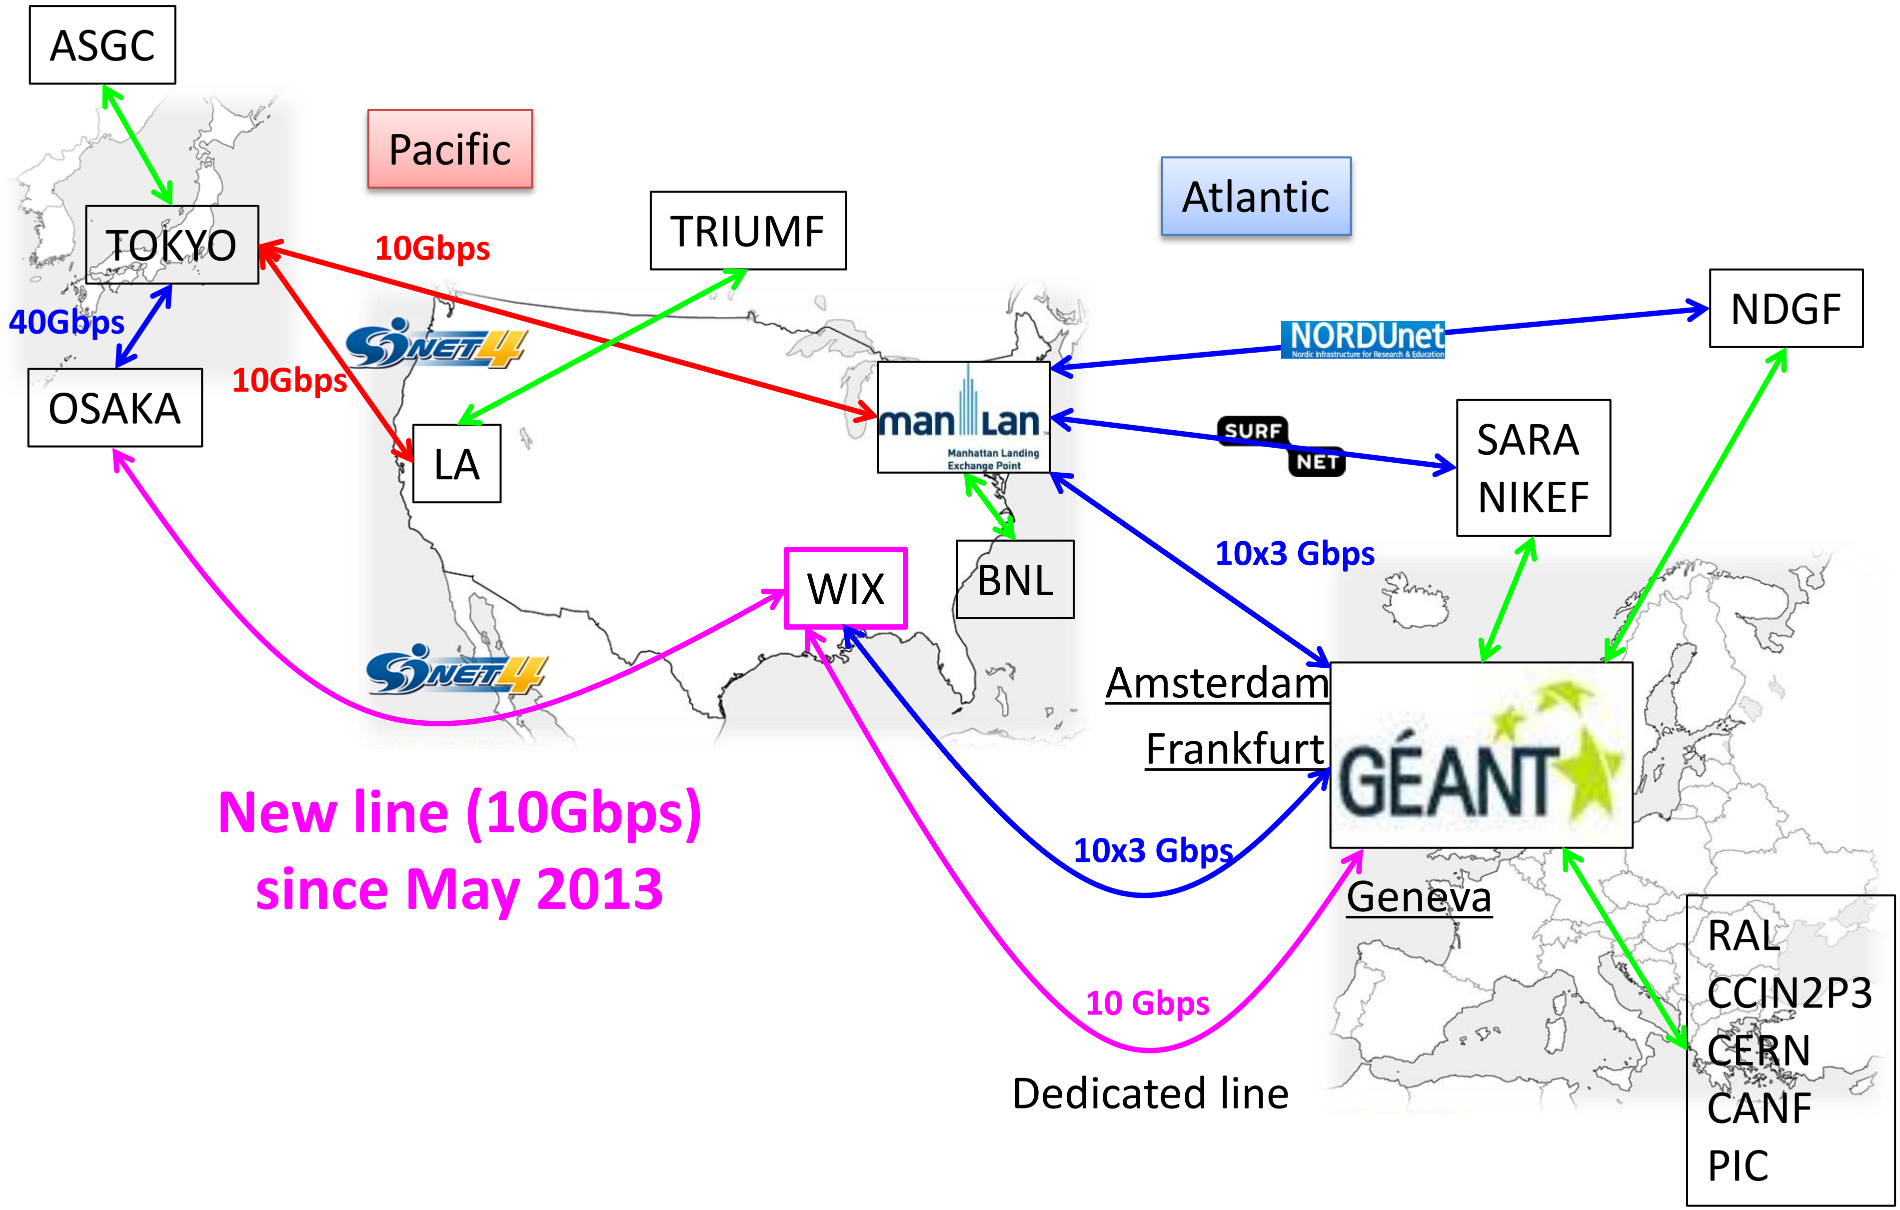

New line (10Gbps) since May 2013

Data transfer throughput (1 day average)

Monitored by network switch

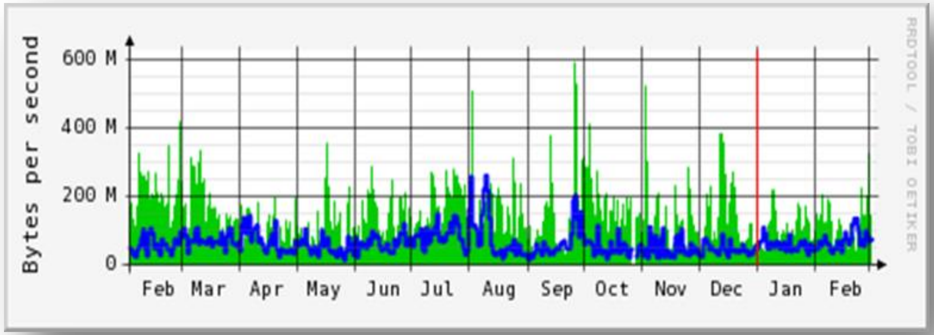

Incoming
Outgoing

Monitored by file servers (extract from grid FTP logs)

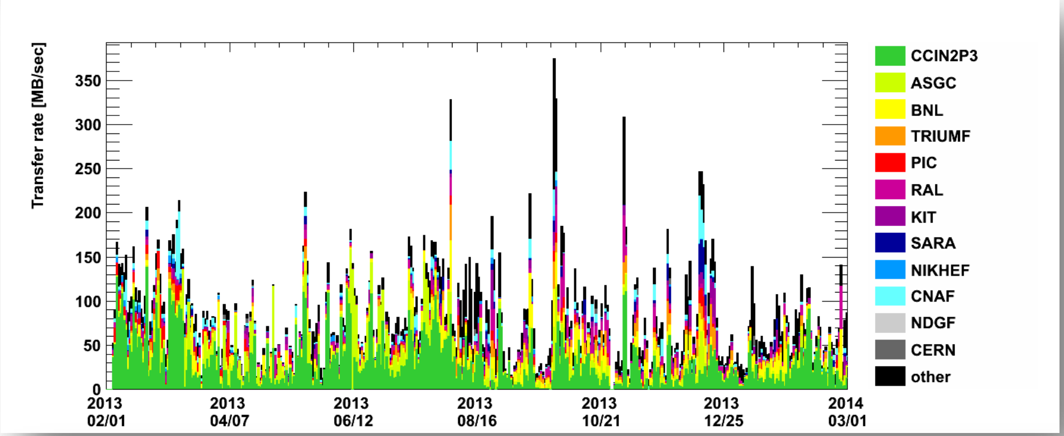

Incoming data

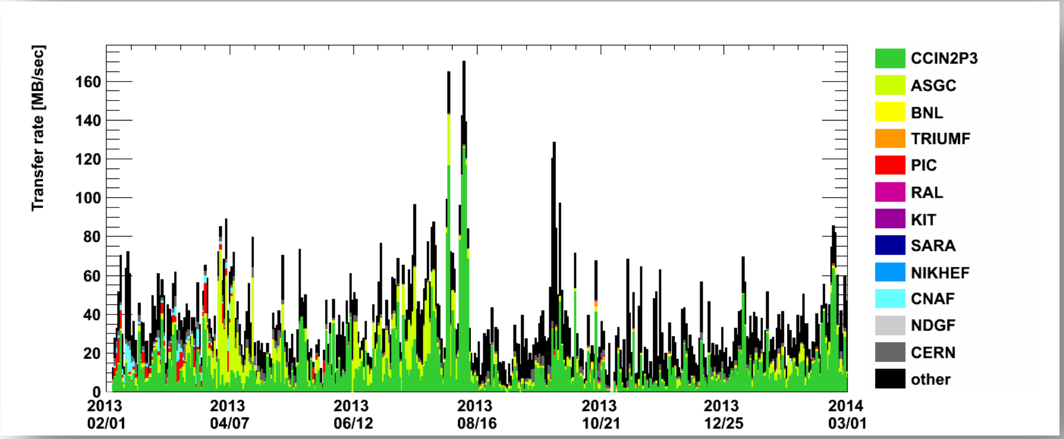

Outgoing data

Data transfer throughput (10 min. average)

Sustained transfer rate

Incoming data: ~100MB/sec in one day average

Outgoing data: ~50MB/sec in one day average

300~400TB of data in Tokyo storage is replaced within one month!

Peak transfer rate

Almost reached to 10Gbps

Need to increase bandwidth and stability!

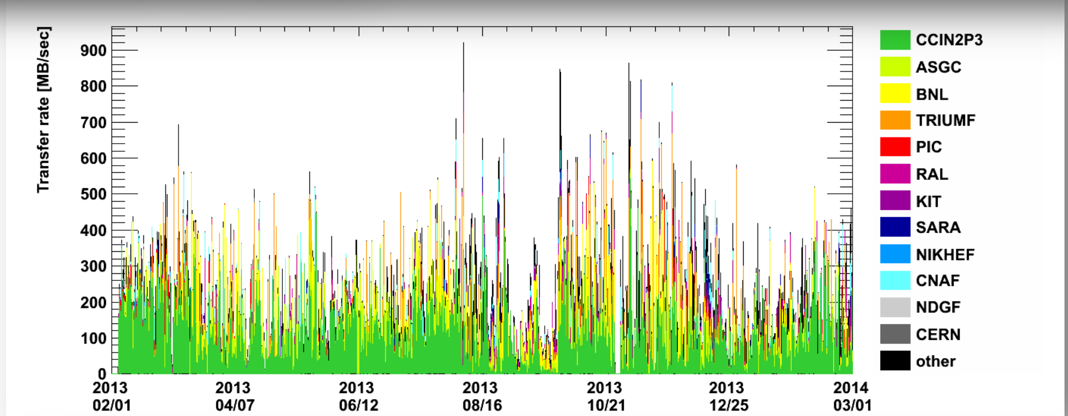

Incoming data

Monitored by file servers (extract from grid FTP logs)

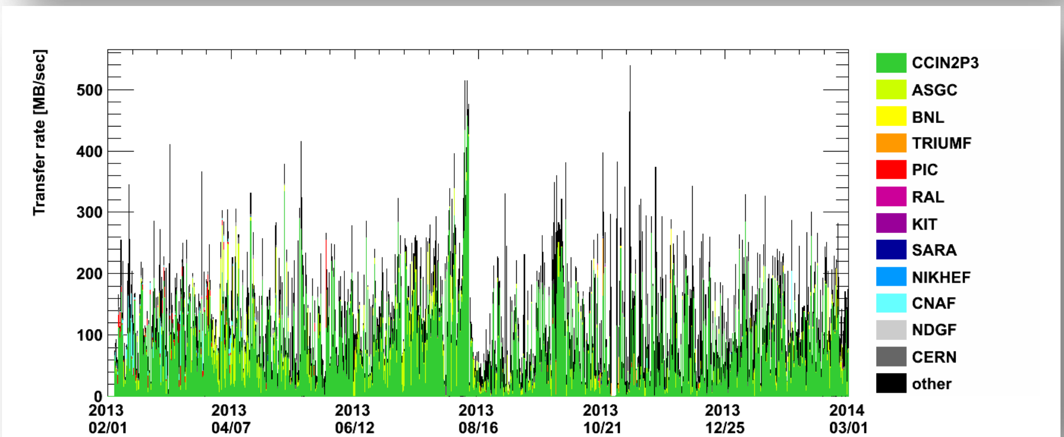

Outgoing data

I/O performance study

The number of CPU cores in the new worker node was increased from 8 cores to 16 cores.

Local I/O performance for the data staging area may become a possible bottleneck.

We have checked the performance by comparing with a special worker node,
which have a SSD for the local storage, in the production situation with real ATLAS jobs.

Normal worker node

HDD:	HGST Ultrastar C10K600, 600GB SAS, 10k rpm
RAID1:	DELL PERC H710P
FS:	ext3
OS:	Scientific Linux 5.8
Sequential I/O:	~150MB/sec
IOPS:	~650 (fio tool)

Special worker node

SSD:	Intel SSD DC S3500 450GB
RAID0:	DELL PERC H710P
FS:	ext3
OS:	Scientific Linux 5.8
Sequential I/O:	~400MB/sec
IOPS:	~40000 (fio tool)

HDD vs. SSD

10sec sampling | 60sec sampling | 3600sec sampling

HDD Service time vs. IOPS

SSD Service time vs. IOPS

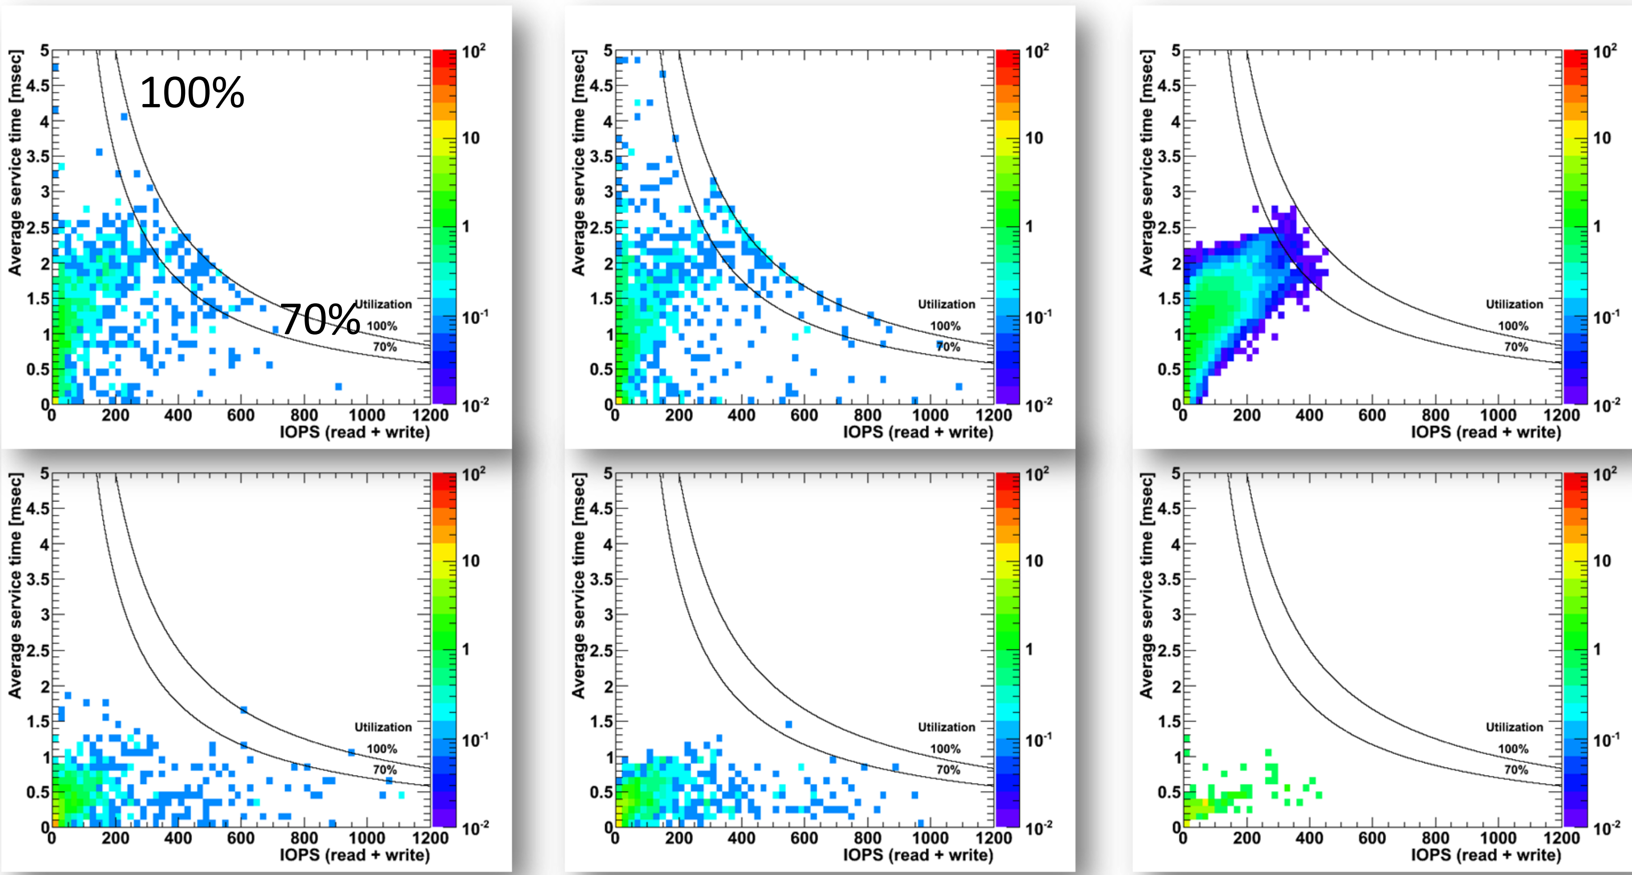

%utilization (iostat) = {IOPS (read) + IOPS (write)} x Service time

Utilization

10sec sampling

60sec sampling

3600sec sampling

Utilization

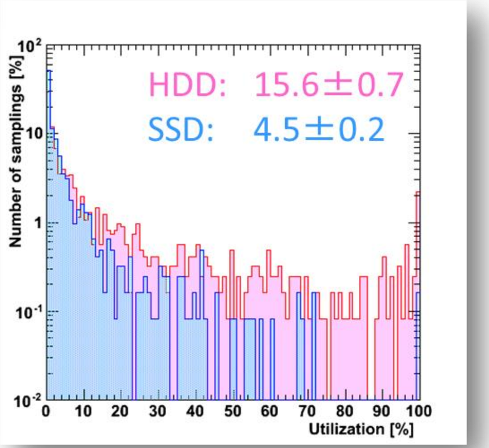

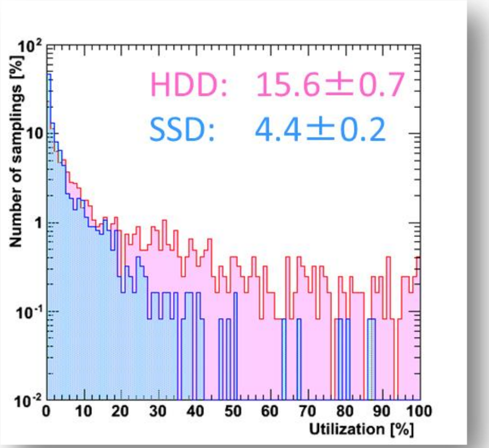

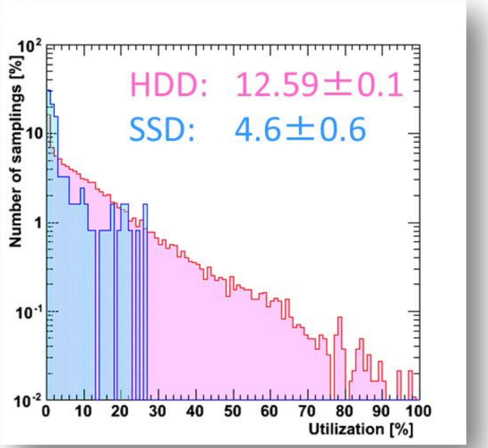

%utilization (iostat) = {IOPS (read) + IOPS (write)} x Service time

Roughly corresponds to the inverse of CPU efficiency (CPU time / Wall time).
Directly affect to the job throughput for the data intensive jobs.

A few percent of I/O request reach 100% of utilization in 10sec sampling.
It leads to the very slow response for the interactive command.

However, it is not a issue by smearing of I/O request in log batch type jobs.

ext3 vs. ext4 on Scientific Linux 6.2

ext3		Number of parallel I/O					
	Block size	1	2	4	8	**16**	32
Read [MB/sec]	64kB	151	90	103	151	**159**	102
Read [MB/sec]	1MB	151	98	113	158	**157**	104
Write [MB/sec]	64kB	120	122	120	118	**116**	110
Write [MB/sec]	1MB	122	120	120	118	**116**	110

ext4		Number of parallel I/O					
	Block size	1	2	4	8	**16**	32
Read [MB/sec]	64kB	169	152	146	139	**154**	116
Read [MB/sec]	1MB	152	131	126	118	**169**	112
Write [MB/sec]	64kB	151	157	154	148	**134**	129
Write [MB/sec]	1MB	143	140	136	129	**117**	136

Measured by using fio tool under Scientific Linux 6.2.

Read access is comparable between ext3 and ext4.
Write access can be expected ~15% improvement.
reservation
delayed write

Utilization (ext3 vs. ext4)

ext3 on SL5

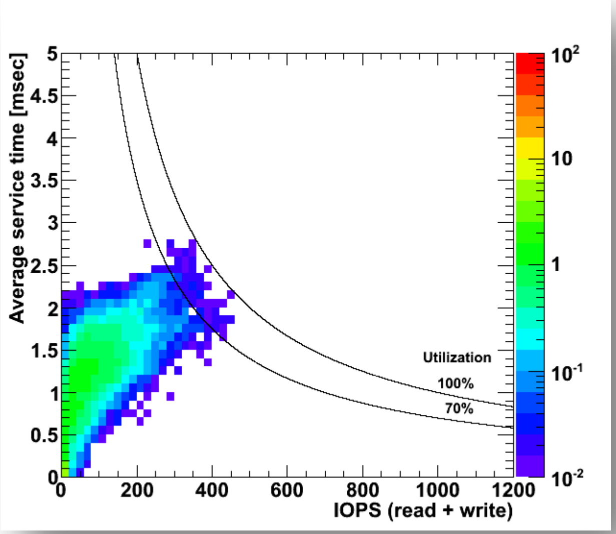

ext4 on SL6

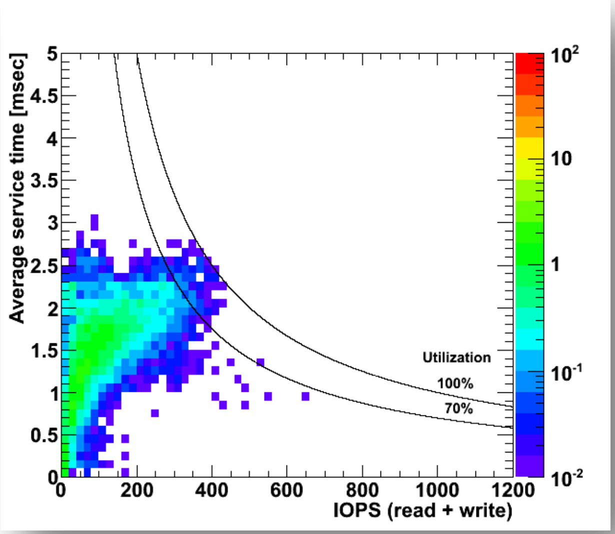

No problem for both filesystem.

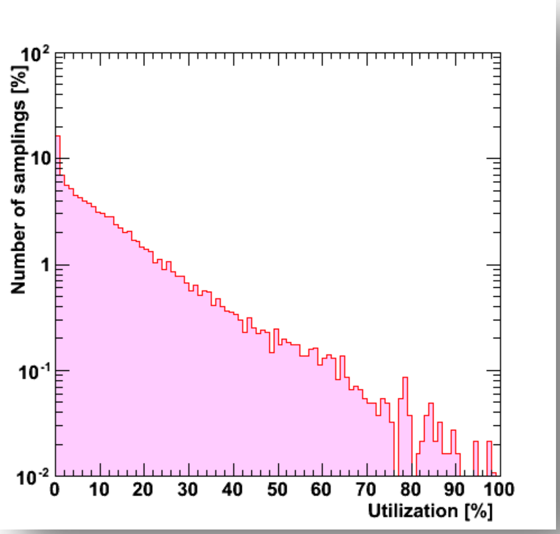

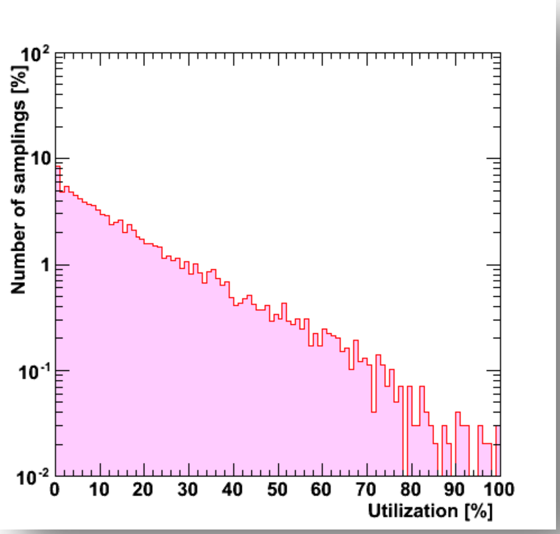

Read performance (ext3 vs. ext4)

ext3 on SL5

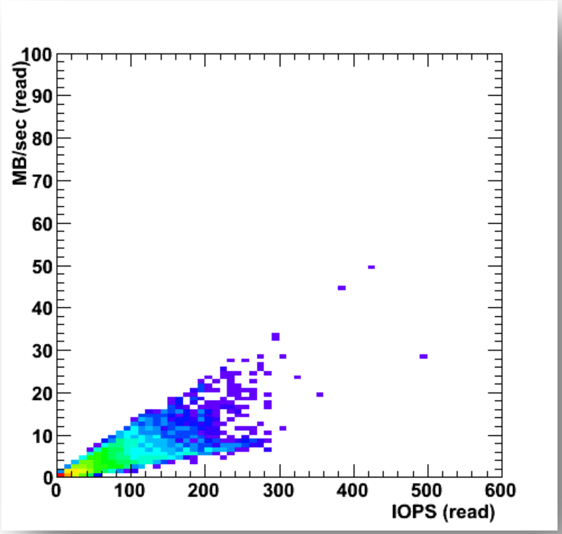

ext4 on SL6

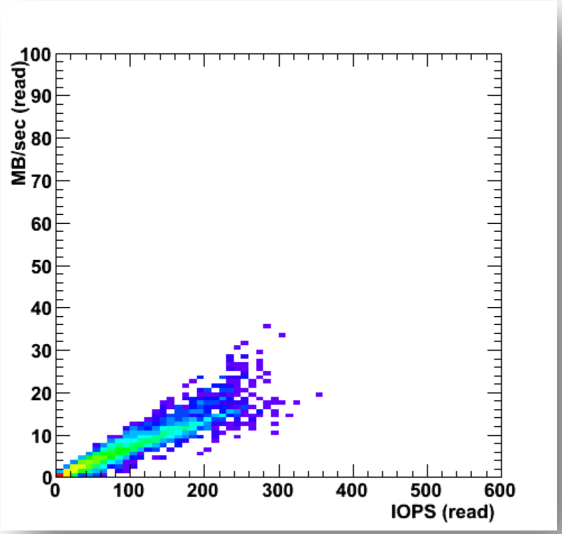

Read performance is also comparable with the ATLAS jobs.

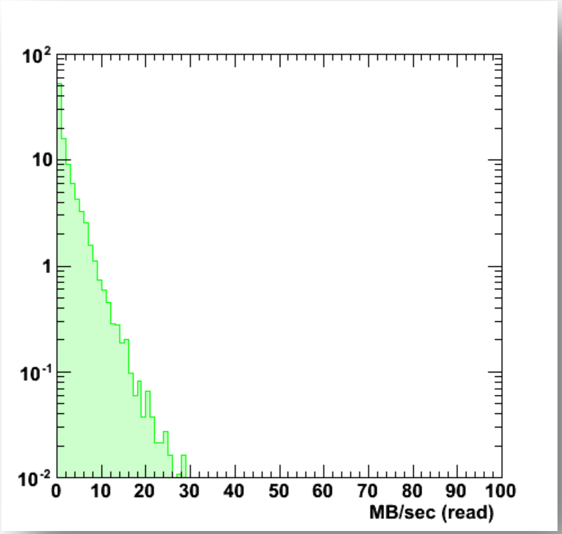

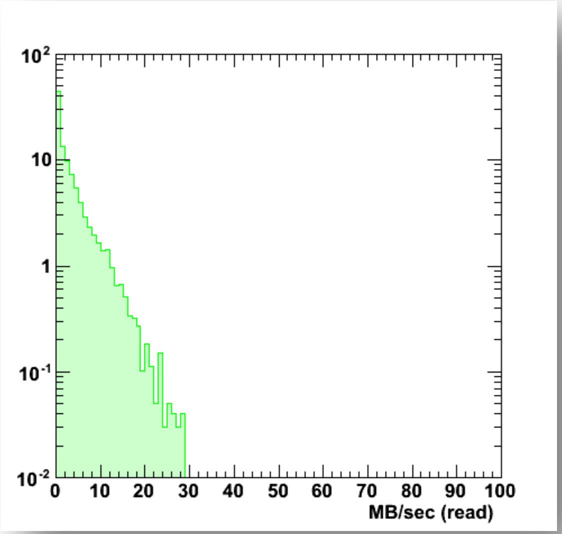

Write performance (ext3 vs. ext4)

ext3 on SL5

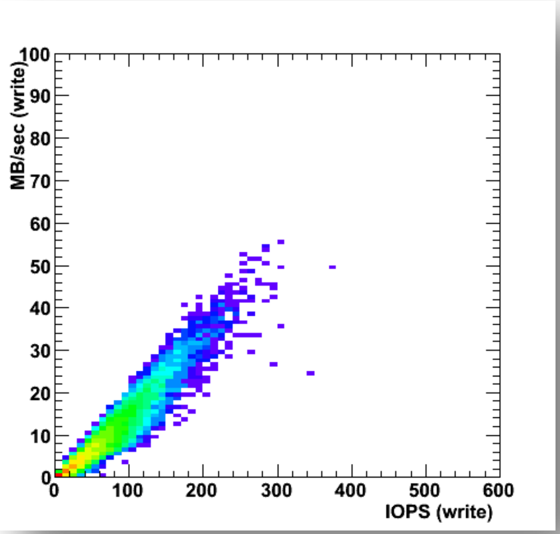

ext4 on SL6

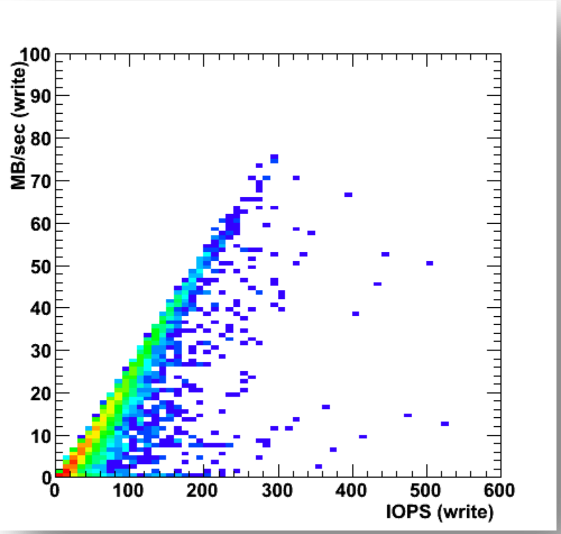

Write performance is improved in the case with ext4 at the peak.

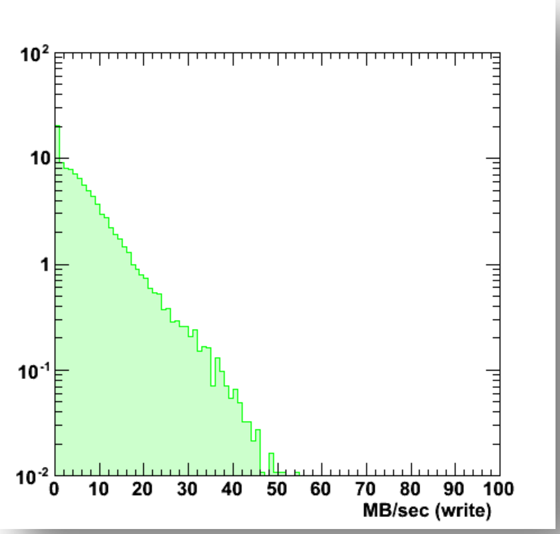

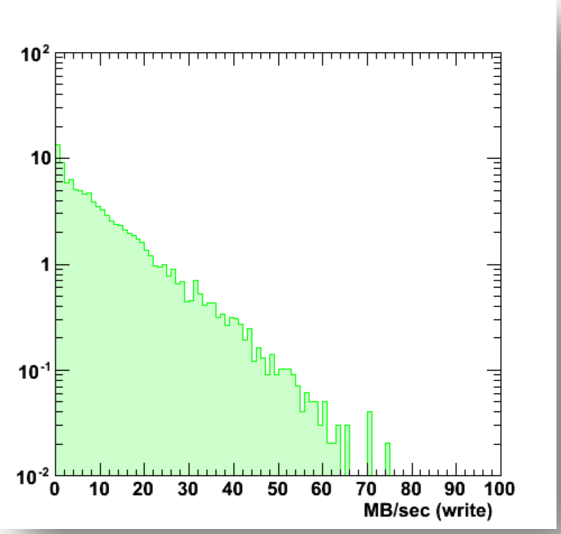

Summary

Operation

- Tokyo Tier2 have been running smoothly by the new system since Feb. 2013.

Performance

- The local I/O performance in our worker node have been studied by a comparison with HDD and SSD at the mixture situation of running real ATLAS production jobs and analysis jobs.
- We confirmed that HDD in the worker node at Tokyo Tier2 center is not a bottleneck for the long batch type jobs at least for the situation of 16 jobs running concurrently.
- Write performance of the local disks for staging is aplenty improved by the ext4, but read performance is not improved so much as compared to ext3.
- The same kind of things should be checked also for the next generation worker node, which have more CPU cores greater than 16 cores, and also for the XFS (will be chosen for RHEL7).
- The improvement of the I/O performance with respect to the direct mount of DPM should be confirmed by ourselves more precisely.

  Job throughput → Staging to local disk vs. Direct I/O from storage

  User's point of view → Jobs will be almost freed from the input file size and the number of input files.

For the next system

- Survey of the new hardware for the next system will be started soon in this year.
- Internal network usage (depend on how to use the storage?) should be checked toward the next system design.