Status of ALICE Tier-2 at Hiroshima

ALICE T1/T2 Workshop
at Univ. Tsukuba
on March 6, 2014

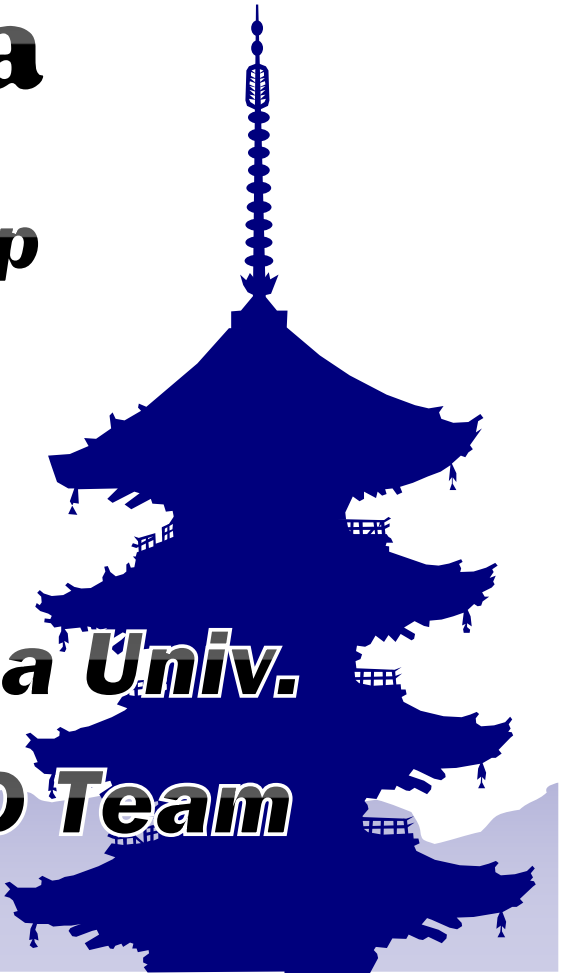

Toru SUGITATE, Hiroshima Univ.
for the ALICE-Japan GRID Team

Hiroshima Tier-2 for ALICE

広島大学

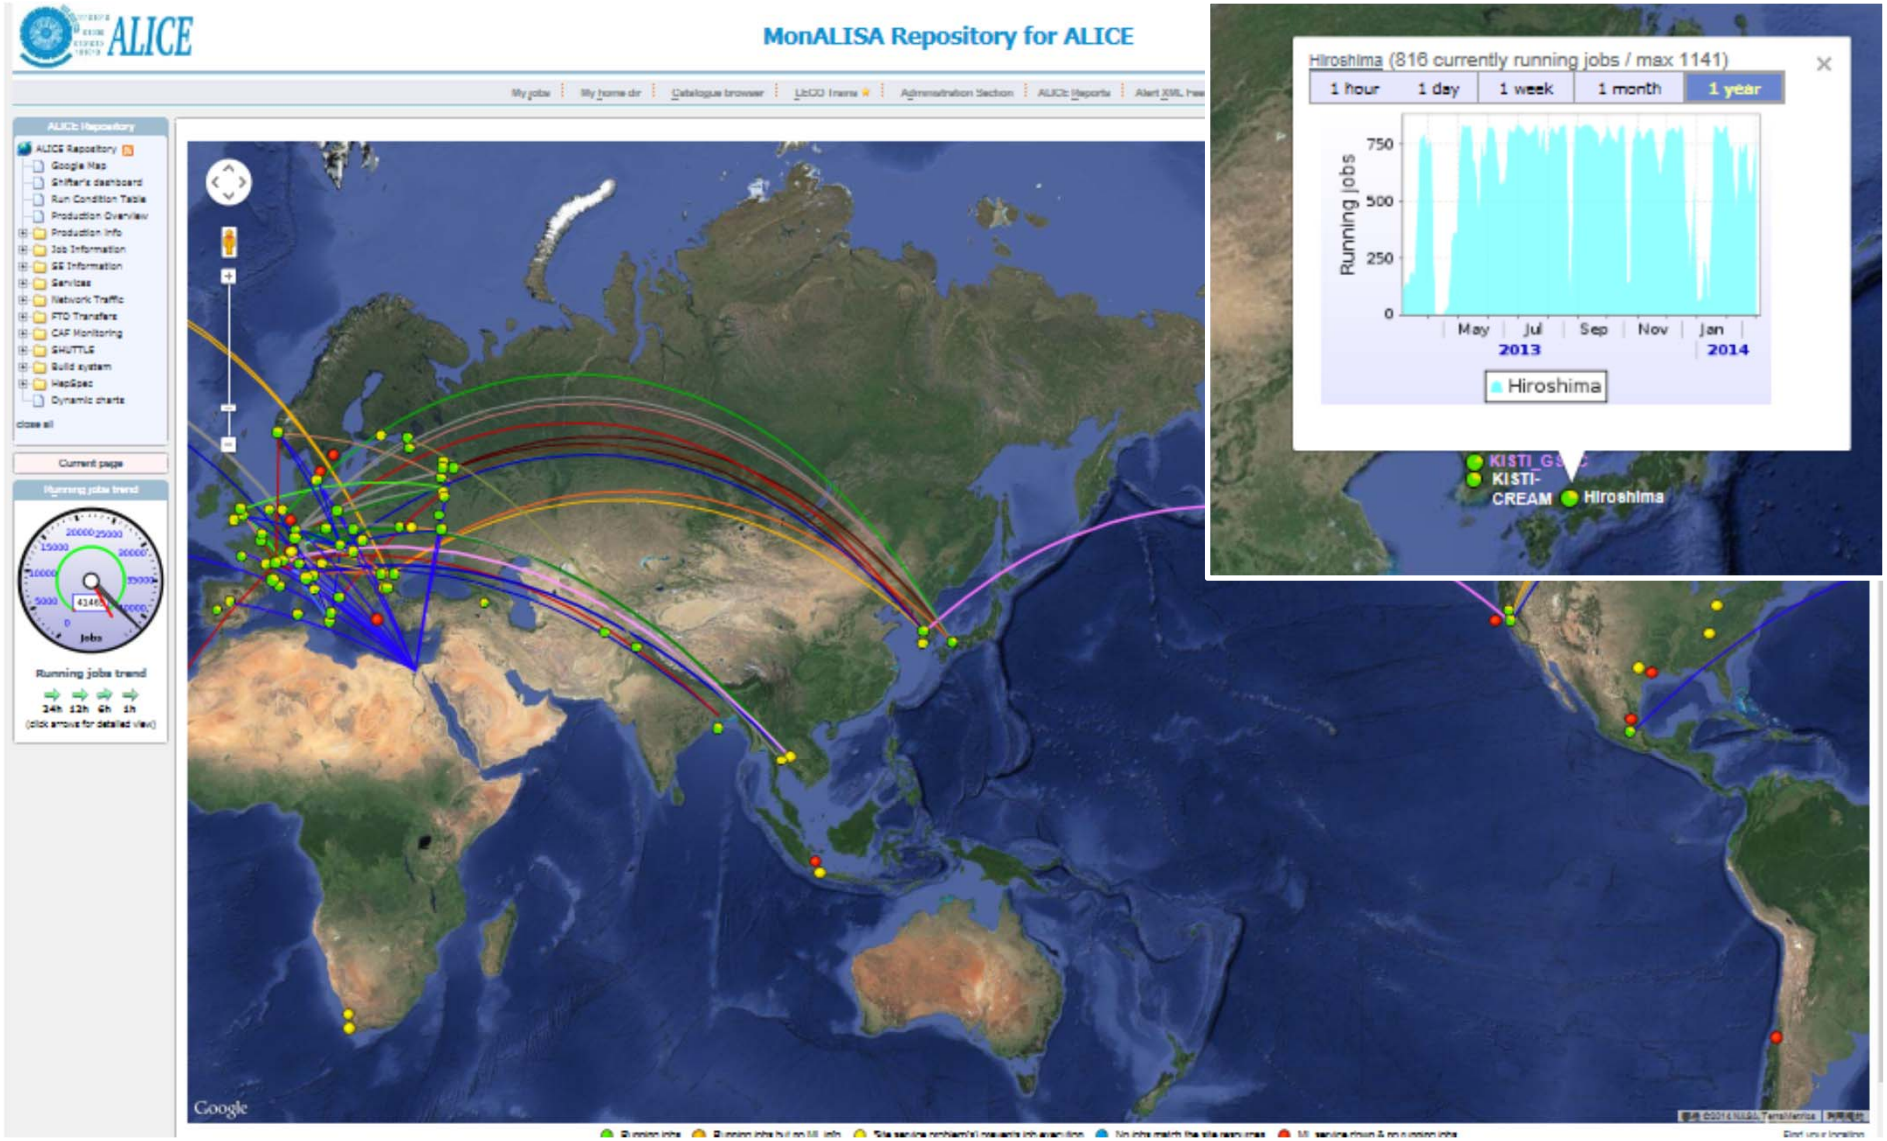

A Little Info. about Us

3 major sightseeing spots:
松島 :Matsushima
天橋立 :Amanohashidate
宮島 :Miyajima

3 major sake breweries:
灘 :Nada
伏見 :Fushimi
西条 :Saijo

Saijo Campus of Hiroshima Univ.

広島大学

ALICE Tier-2 at Hiroshima

- The ALICE LCG site “JP-HIROSHIMA-WLCG” with grid middleware EMI-2 on SL6.4… as stable as possible.
- GRID service; APEL, sBDII, CREAM-CE, XROOTD, DPM-SE, VOBOX… as compact as possible.
- WN resources; 1164 Xeon-cores in total
  Xeon5355(4cores@2.6GHz) x 2cpu x 32boxes &
  Xeon5365(4cores@3.0GHz) x 2cpu x 20blades &
  Xeon5570(4cores@2.9GHz) x 2cpu x 26blades &
  Xeon5670(6cores@2.9GHz) x 2cpu x 3blades &
  Xeon5660(6cores@2.8GHz) x 2cpu x 42blades
- Storage cap; 408TB disks on 6 servers and no MS
- Around 2/3 resource deployed in the ALICE GRID
- The rest in a local cluster
- Network B/W: 1Gbps on 40Gbps-SINET4 in Japan
- WLCG support by ASGC in Taiwan
- Responsible by Prof. Toru Sugitate
- Operated by TS and 中宮義英, and remote technical supports by a part-time SE of 創夢（株）in Tokyo

Configuration since Feb. 2012

- Secure/Robust subnets
- 10Gbps Internal connection (part.)
- Space/Energy saving

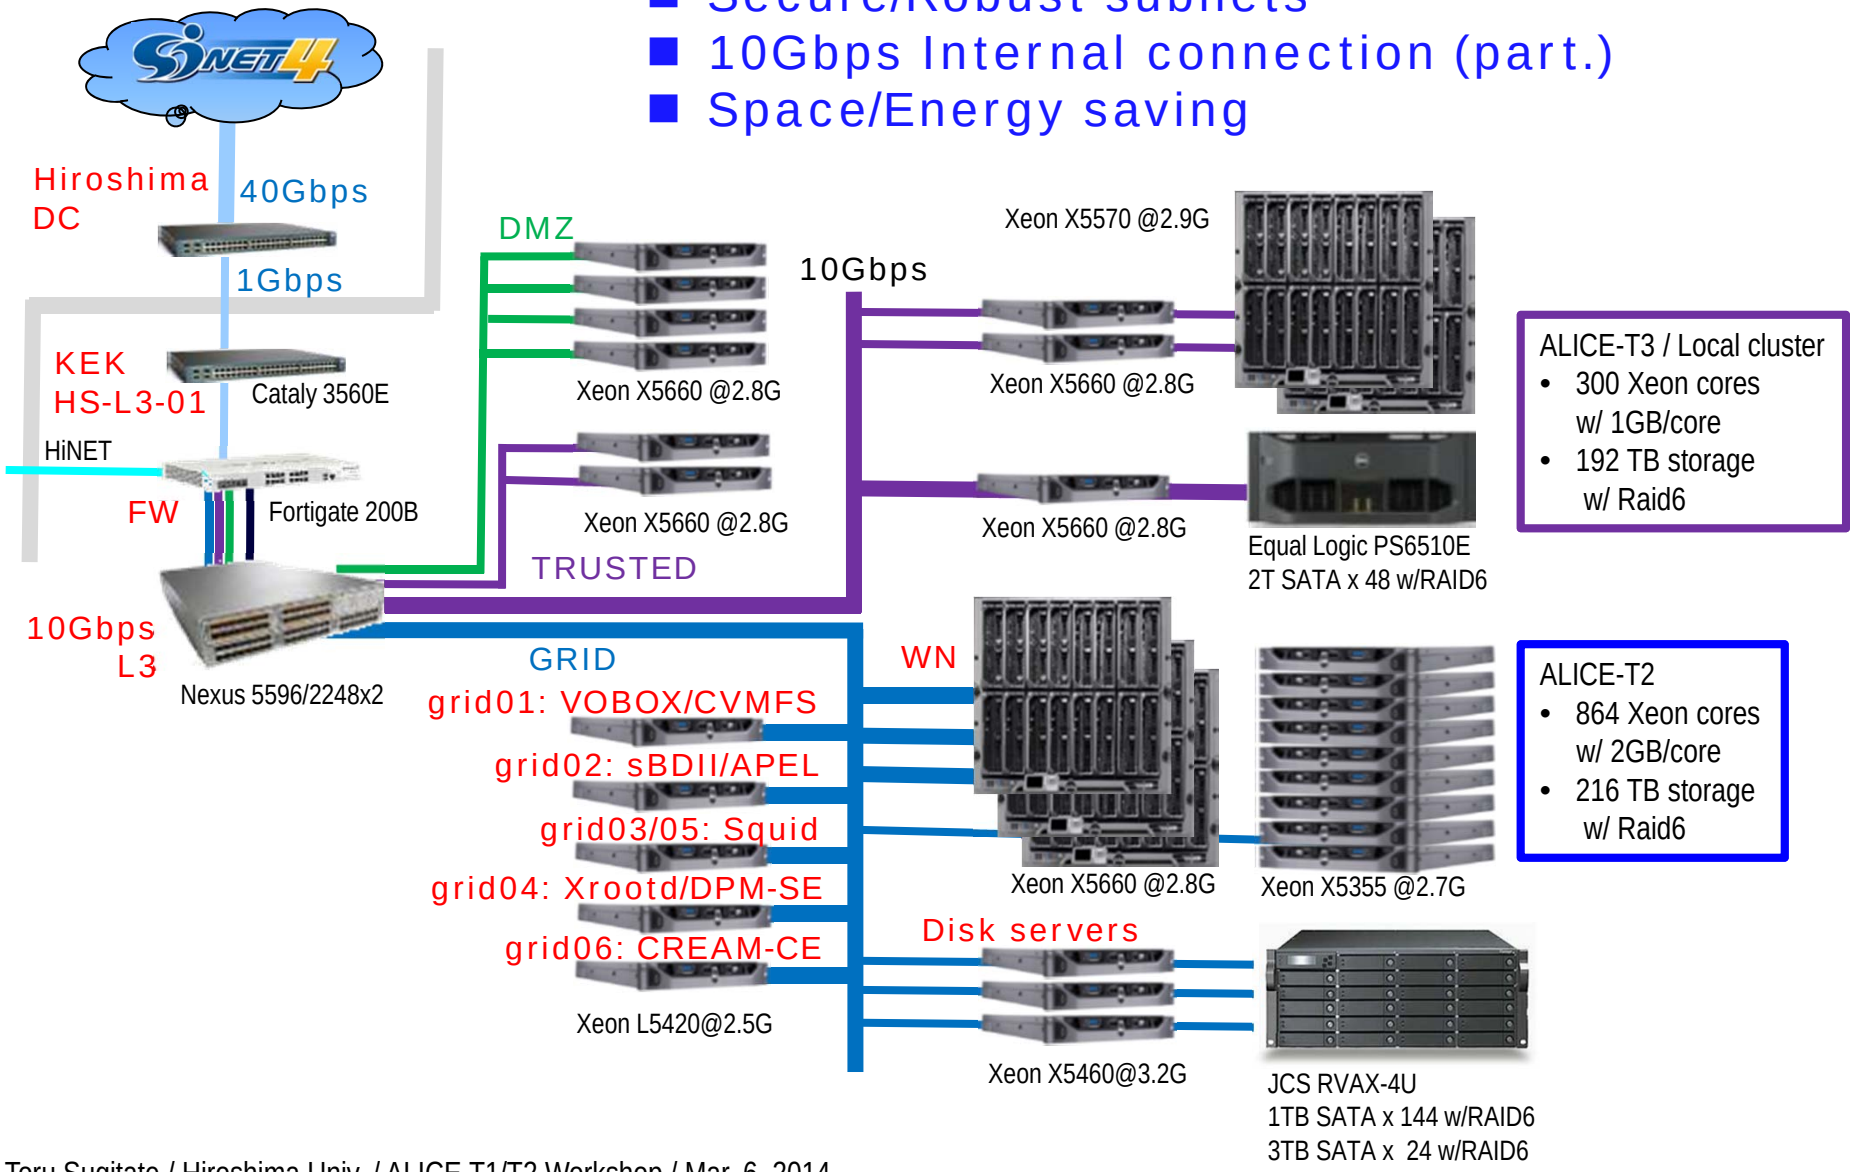

Major upgrade records in last years

In operation since 2009

2012

- Feb major replacement of equipment (every 60 mon.) and resume the services under a new network configuration
- Jun deploy 42 WN's (252 Xenon cores) on SLC 5.7
- Nov migrate to EMI-1 from gLite3.2

2013

- Feb deploy a new disk server (nfs08: 72TB)
- Apr migrate to EMI-2 on SL 6.4 from SLC 5.7
  update to Xrootd 3.2.6
- May update to EMI-2 Update 12
- Aug install CVMFS 2.1.15 and update VOBOX

2014 and beyond

soon

- migrate to EMI-3 from EMI-2

later

- network upgrade to 10Gbps
- move to LHCONE

Daily Score as of Today

Select site: Hiroshima »

MonALISA Information Version: 13.11.04 (JDK 1.7.0_45)
Running on: grid01.hepl.hiroshima-u.ac.jp
Administrator: Toru Sugitate,Hiroshima <sugitate@hiroshima-u.ac.jp,wlcg-hiro@ml.hepl.hiroshima-u.ac.jp>

Service health NTP: SYNC, offset: 0s

Services status
AliEn: v2-19.239
ClusterMonitor: OK
PackMan: n/a
CE: OK
CE info: At the moment we are busy (we ...
Max running jobs: 1000
Max queued jobs: 50

Proxies status
AliEn proxy: OK (1 day, 23:25)
Delegated proxy: OK (2 days, 00:00)
Proxy server: OK (27 days, 10:24)
Proxy of the machine: OK (18:17)

Current jobs status
Assigned: 0
Running: 827
Saving: 0

Accounting (last 24h)
Success jobs: 5449 (profile)
Failed jobs: 0
Error jobs: 293
kSI2k units: 2022 / pledged

Site averages (last 24h)
Active nodes: 86
Average kSI2k/node: 2.433

Storages status

Name	Status	Size	Used	Free	Usage	No of files	Type	ADD test
ALICE::Hiroshima::SE	OK	177.3 TB	60.24%	70.49 TB	106.8 TB	2.898 M	FILE	OK

VoBox health
CPUs: 24x 2792MHz
Mem usage: 28.19% of 23.45 GB
Processes: 503
Sockets: 133 TCP / 27 UDP
Uptime: 136 days, 18:18

CPU usage (last 1h avg)
Load: ***0.118***
User: 0.567%
System: 0.12%
IOWait: 0.003%
Idle: 99.31%
Int: 0%
Soft Int: 0.001%
Nice: 0%
Steal: 0%

AliEn LDAP var	VoBox path	Size	Used	Free	Use%
TMP	/home/sgmali01/ALICE/tmp	501.9 GB	6.421 GB	470 GB	2%
LOG	/home/sgmali01/ALICE/alien-logs	501.9 GB	6.421 GB	470 GB	2%
CACHE	/home/sgmali01/ALICE/cache	501.9 GB	6.421 GB	470 GB	2%

ALICE jobs as of Today

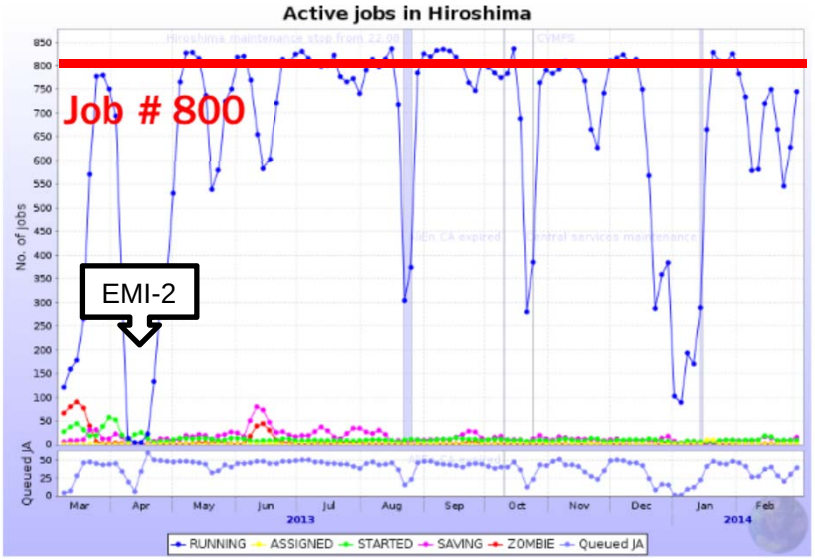

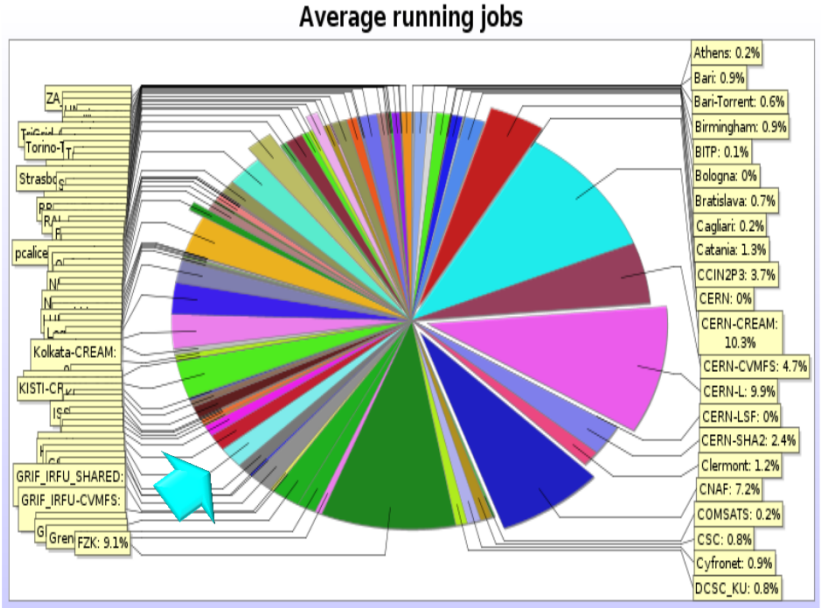

- 864 Xeon-cores in CRAEM-CE.
- Stably accepting over 800 jobs and process around 5000 jobs a day,
- corresponding to about 2% contribution to the entire ALICE jobs.

Job status

Service	VOBox Address	Jobs Running	Started	Saving	Zombie	Total (R+S+S+Z)
34. FZK	alice-kit.gridka.de	3323	5	11	0	
12. CCIN2P3	ccwlcgalice02.in2p3.fr	3155	1	1	473	
24. CNAF	ui01-alice.cr.cnaf.infn.it	2885	9	6	10	
18. CERN-CVMFS	voalice13.cern.ch	2525	19	8	27	
17. CERN-CREAM	voalice12.cern.ch	2286	27	2	172	
19. CERN-L	voalice14.cern.ch	2076	25	2	162	
39. GRIF_IRFU	node09.datagrid.cea.fr	1930	0	14	1	
77. NIHAM	hgate.nipne.ro	1411	1	2	0	
89. Prague	alicebox.farm.particle.cz	1275	21	6	2	
68. LBL	palicevo1.nersc.gov	1132	1	3	28	
38. GRIF_IPNO	ipnvobox.in2p3.fr	983	1	6	0	
108. Torino	alibox2.to.infn.it	972	0	1	0	
4. Bari	vobox-alice.ba.infn.it	952	2	1	51	
71. LLNL	glcc37.ucllnl.org	864	0	0	0	
43. GSI_2	lxcealice01.gsi.de	855	4	37	38	
49. Hiroshima	grid01.hepl.hiroshima-u.ac.jp	827	1	2	0	
62. KISTI_GSDC	vobox11.sdfarm.kr	779	3	2	0	
88. Poznan	vobox.reef.man.poznan.pl	731	8	1	0	
69. Legnaro	vobox-alice.lnl.infn.it	731	0	2	9	
21. CERN-SHA2	voalice11.cern.ch	728	1	18	17	
11. Catania	vobox.ct.infn.it	721	1	1	0	
52. IPNL	lyogrid08.in2p3.fr	701	2	1	0	
95. RRC_KI_T1	rhole.t1.grid.kiae.ru	671	1	2	13	
117. UiB	alien.bccs.uib.no	666	3	1	0	
65. Kolkata-CREAM	grid01.tier2-kol.res.in	660	1	1	0	

Internal Network Traffics

- All traffic filtered with a FW/FortiGate200B.
- In-/Out-bound traffic rates at 0.2-0.5 Gbps.
- Xrootd storage of 177 TB in 4 disk servers.
- A large amount of traffic between WN's and global IP servers routed in the L3 SW.

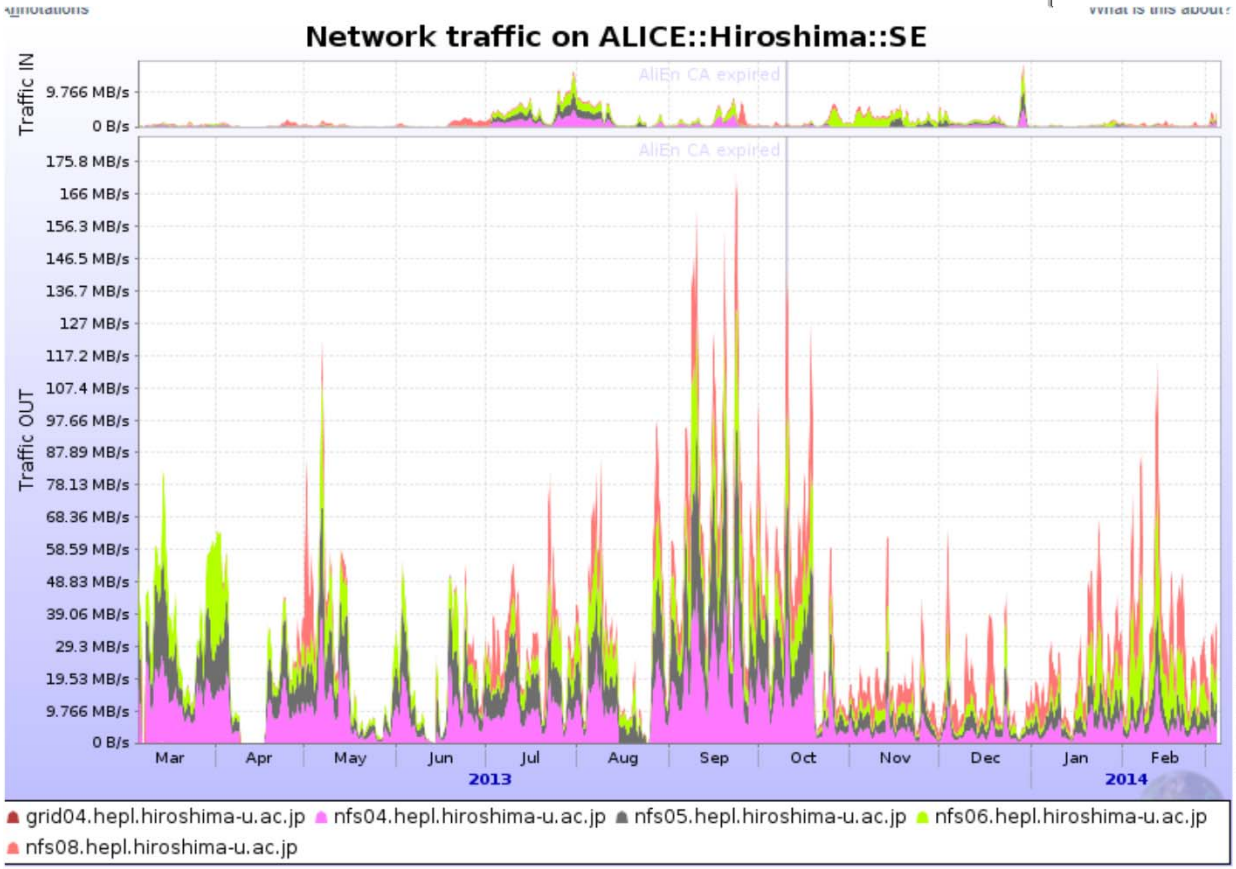

Network Connection to WAN

Hiroshima »

Chart view »

Links: FDT, Kernel parameters tuning

‹Hiroshima›

IN from

No.	ID	Site	Speed (Mbps)	Hops	RTT (ms)	Streams
1.	1253777	CCIN2P31	167.61			1
2.	778391	GSI-SCLAB	58.72			1
3.	1768037	KISTI_GSDC	58.72			1
4.	1768396	LLNL	58.72			1
5.	1648090	RAL_ARC	58.72	17	302.58	1
6.	1384961	GSI-SGE	50.33			1
7.	1765774	NIKHEF	50.33			1
8.	1266811	PDC	50.33			1
9.	1769353	Prague	50.33			1
10.	1769336	ISMA	41.95	18	354.07	1
11.	1767461	Legnaro	41.95			1
12.	1766677	NECTEC	41.95			1
13.	1766472	Torino	41.95			1
14.	1765939	Yerevan	41.90			1
15.	918809	UCT_CERN_RC	33.82			1
16.	1766406	Birmingham	33.56			1
17.	1766290	JINR	33.56			1
18.	1767488	UiB	33.56			1
19.	1610554	Wuhan	33.56			1
20.	1766338	Bari	25.17			1
21.	1766456	Catania	25.17			1
22.	1765494	FZK	25.17			1
23.	1764903	GRIF_IPNO	25.17			1

OUT to

No.	ID	Site	Speed (Mbps)	Hops	RTT (ms)	Streams
1.	1767532	GRIF_IPNO	167.78			1
2.	1266070	PDC	150.85	16	313.72	1
3.	1767468	Legnaro	134.22	21	311.05	1
4.	1767244	Grenoble	92.28			1
5.	1767351	Subatech	92.28	20	325.21	1
6.	1764852	UNAM_T1	92.28	20	251.25	1
7.	1769122	DCSC_KU	83.89	15	304.71	1
8.	1767421	Madrid	83.89			1
9.	1766344	CERN-CREAM	75.50	24	298.70	1
10.	1767270	Bratislava	67.11	18	319.73	1
11.	1765465	CERN-SHA2	67.11	24	310.97	1
12.	1767207	GRIF_IRFU	67.11	17	294.64	1
13.	1764719	JINR	50.33	16	327.68	1
14.	1768912	ITEP	41.95	15	252.91	1
15.	491471	Houston	34.89	15	171.57	1
16.	1766417	COMSATS	33.56	23	347.85	1
17.	1769214	SARA	33.56	14	302.60	1
18.	1766618	Bari	25.17	20	336.51	1
19.	1766660	Catania	25.17	20	327.10	1
20.	1765581	CNAF	25.17	20	312.40	1
21.	1769382	ISMA	25.17	21	358.20	1
22.	1764981	ISS_LCG	25.17	24	334.83	1
23.	1764734	KNU	25.17	20	337.18	1

port13 -- hepfw

System: hepfw in
Maintainer:
Description: port13
ifType: ethernetCsmacd (6)
ifName: port13
Max Speed: 125.0 MBytes/s
Ip: 202.13.220.249 (hepfw.hepl.hiroshima-u.ac.jp)

The statistics were last updated **Wednesday, 5 March 2014 at 20:45**, at which time **'hepfw'** had been up for **192 days, 12:13:54**.

`Daily' Graph (5 Minute Average)

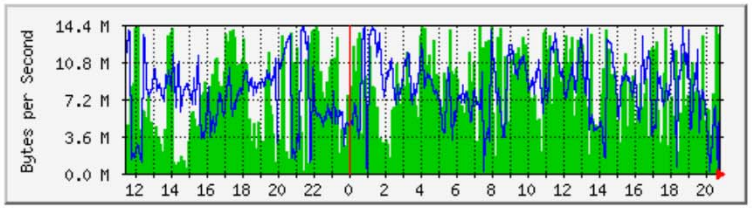

	Max	Average	Current
In	14.2 MB/s (11.4%)	7113.0 kB/s (5.7%)	2740.6 kB/s (2.2%)
Out	14.2 MB/s (11.4%)	7887.6 kB/s (6.3%)	9639.1 kB/s (7.7%)

`Yearly' Graph (1 Day Average)

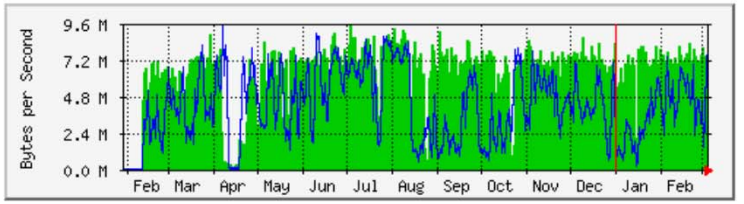

	Max	Average	Current
In	9422.7 kB/s (7.5%)	6479.5 kB/s (5.2%)	6403.7 kB/s (5.1%)
Out	9456.6 kB/s (7.6%)	4287.7 kB/s (3.4%)	7326.9 kB/s (5.9%)

Routes to the World

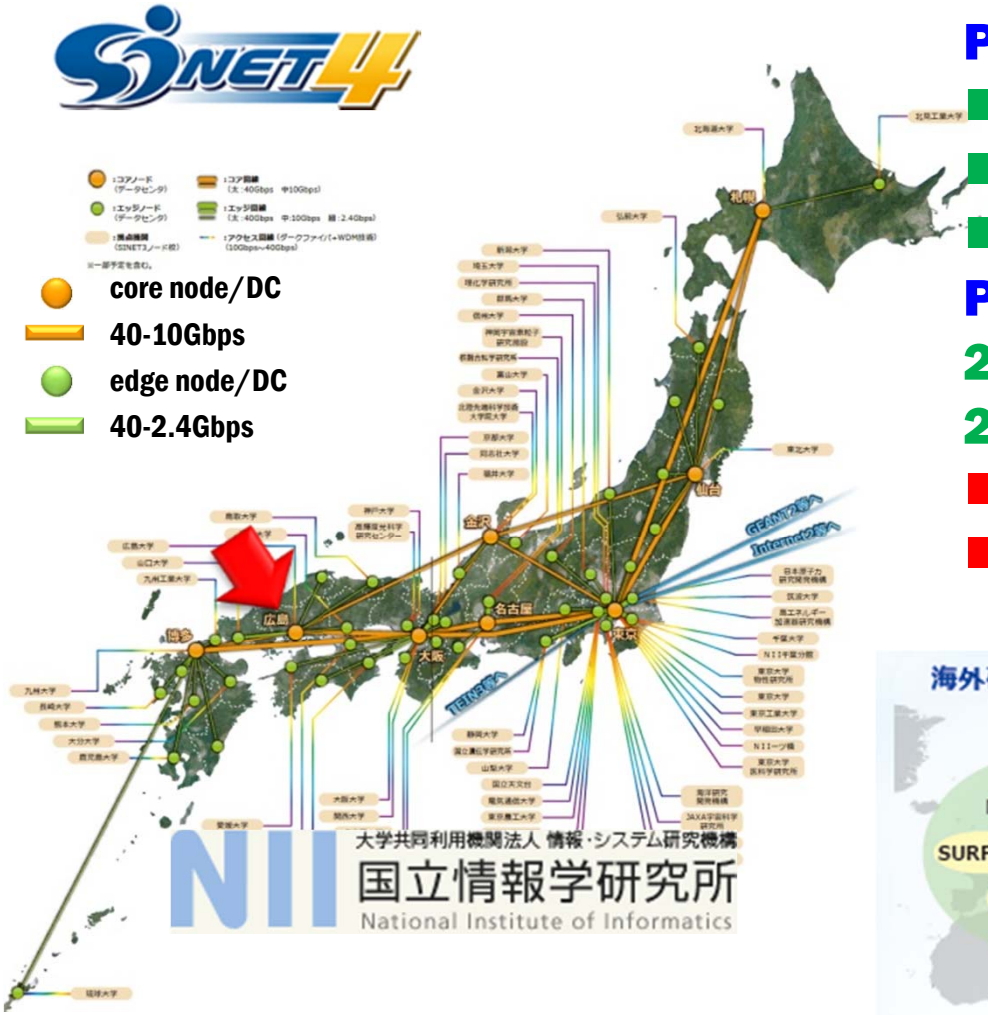

Present: SINET-4

- Hiroshima DC: 40Gbps Core node
- Hiroshima T2 to DC: 1Gbps-MPLS
- International connection via SINET

Plans

2015: Campus LAN upgrade

2016: SINET-4 >> SINET-X

- Hiroshima T2 to DC: 10Gbps
- Approach to LHCONE

Important to develop routing inside Asian communities!

Summary

- **Hiroshima Tier-2 /3**
  - Accepts over 800 jobs stably, and process around 5000 jobs a day,
  - corresponding to about 2% contribution to the entire ALICE jobs.
  - Supporting local/non-local users at local CPU Cluster / "Tier-3".
  - Trace network and tune up speed to increase productivity, but…
  - Declare a 10Gbps connection to the DC in SINET-X.
- **LCG operation**
  - CVMFS done; "27th site out of 52." by Predrag.
  - Thanks strong supports by LCG/EGI/EMI/GGUS folks.
    - SLC4 >> SLC5 >> SL6
    - gLite 3.1 >> gLite 3.2 >> EMI-1 >> EMI-2
    - >> EMI-3 coming soon
- **For Asian community**
  - ***Asian ALICE Analysis Facility (A3F)*** to strengthen physics productivity through the Asian community.
    - Need ESD/AOD in Asia – *thanks to KISTI-T1! ^^)*
    - Develop routing inside Asian countries.
    - Strengthen human network in the communities.
- ***CHEP2015*** and ***QM2016*** will be held in Japan.

Please join us.