



yooyoo360.com



ALICE Wuhan Site Report

Daicui Zhou

Zhongbao Yin

Hengtong Ding

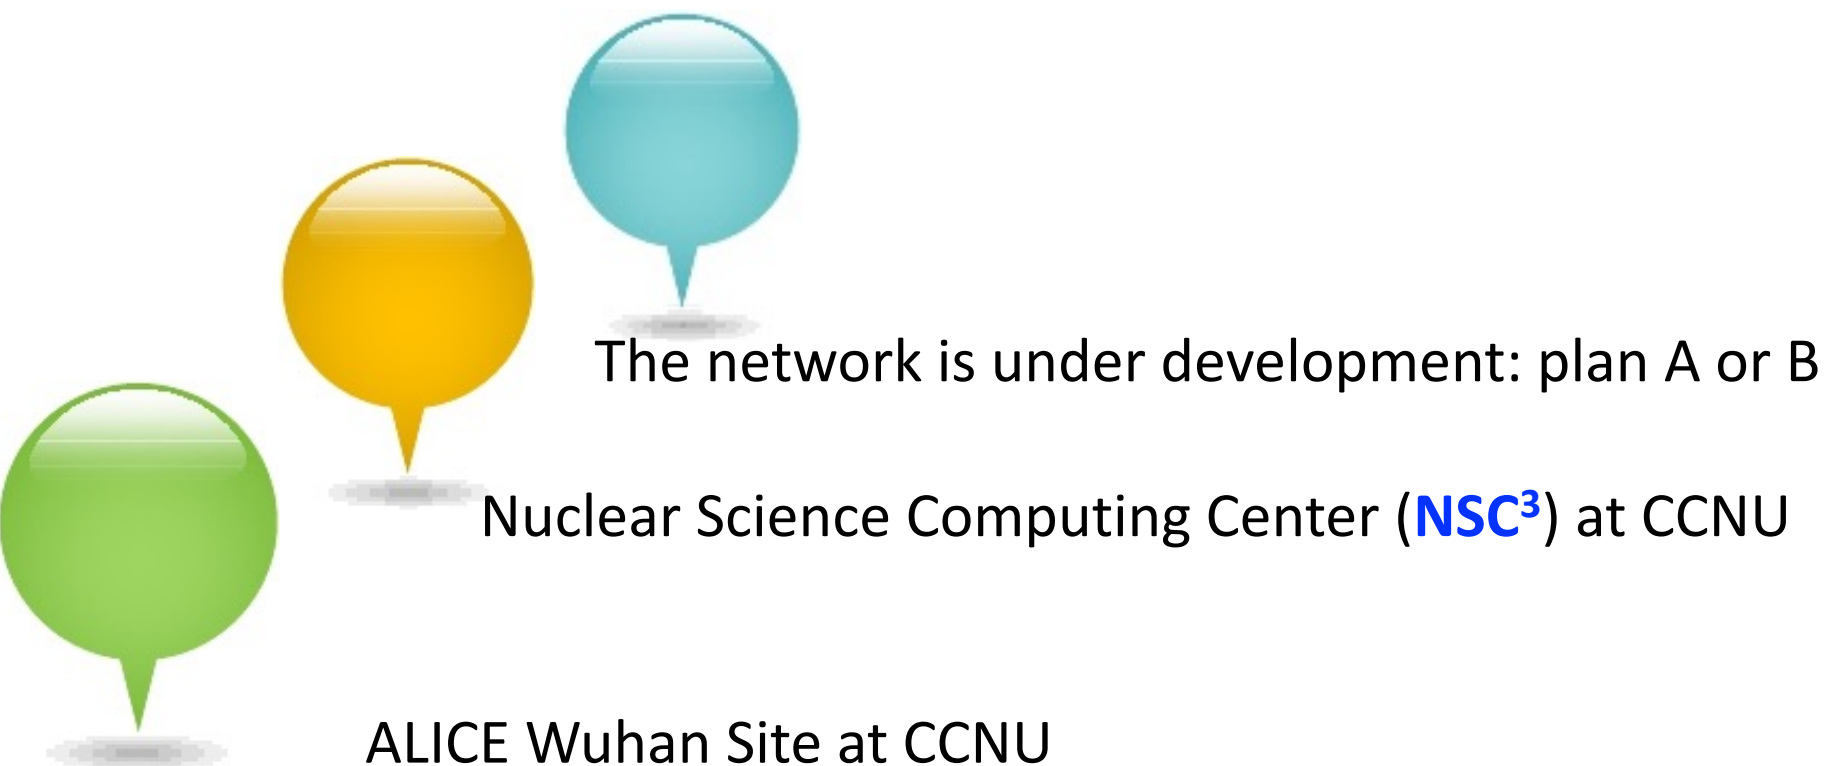
Qiye Shou

Jianlin Zhu

Hua Pei

Nov. 21-24, ATCF2022/Bangkok, Thailand

Current Status of Wuhan Site



A map view of ALICE computing sites



MonALISA Repository for ALICE



My jobs | My home dir | Catalogue browser | LEGO Trains | Administration Section | ALICE Reports | Alert XML Feed | Firefox Toolbar | MonALISA GUI


ALICE Repository

- ALICE Repository
- Google Map
- Shifter's dashboard
- Run Condition Table
- Production Overview
- Production info
- Job Information
- SE Information
- Services
- Network Traffic
- FTD Transfers
- CAF Monitoring
- SHUTTLE
- Build system
- HepSpec
- Dynamic charts

close all

This page: bookmark, URL

Active jobs trend



Active jobs trend

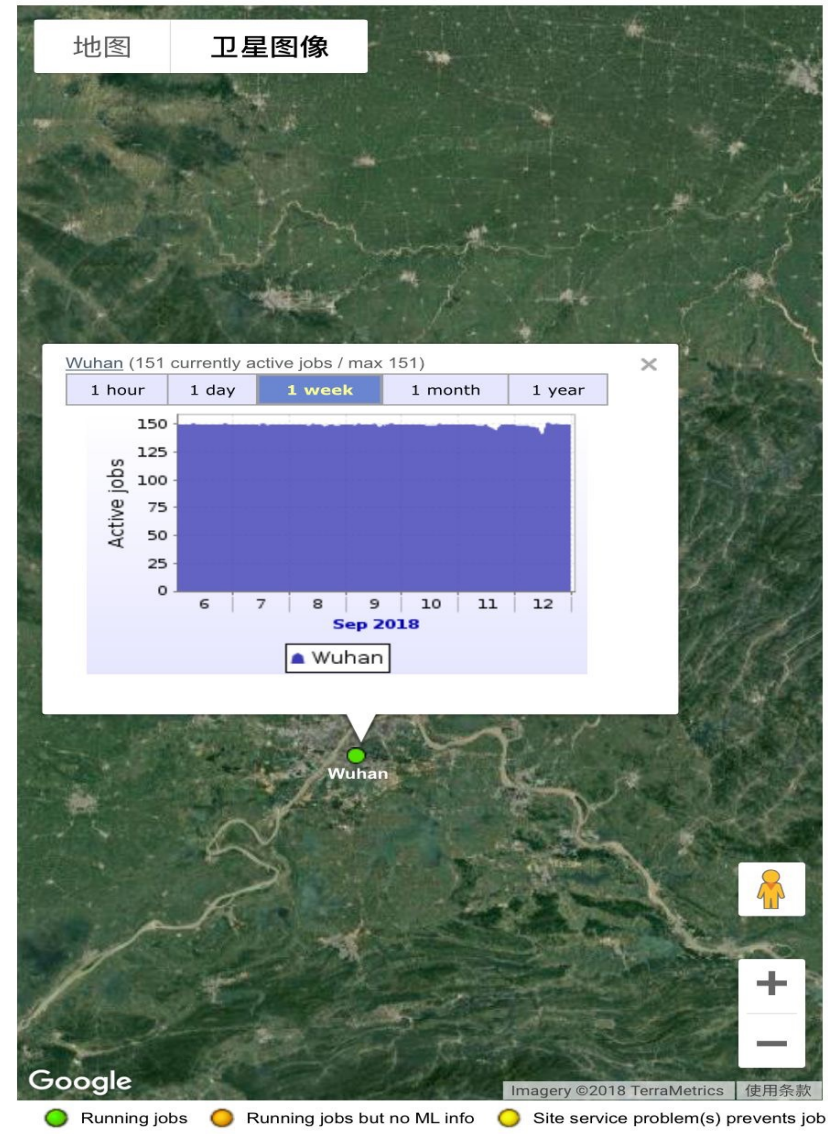


● Running jobs ● Running jobs but no ML info ● Site service problem(s) prevents job execution ● No jobs match the site resources ● ML service down & no running jobs

Map options


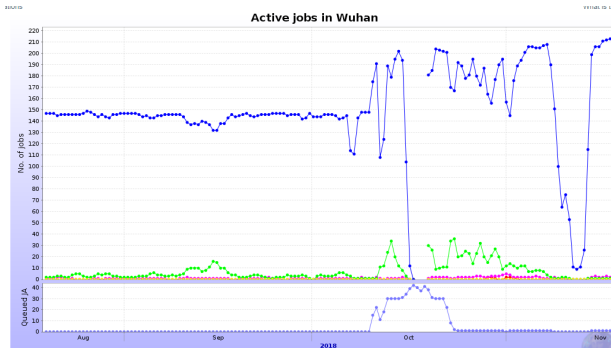
Previously available resources

- "Old" computers reported 2018
 - VOBOX: 8 cores, 16GB RAM, 120GB disk
 - CREAM-CE: 8 cores, 16GB RAM, 120GB disk
 - WNs: each WN has 24 cores, 50GB RAM, 280GB disk.
Total 240 cores
(CPU purchased by 2015)
- SE: ThinkServer RD530, 80TB



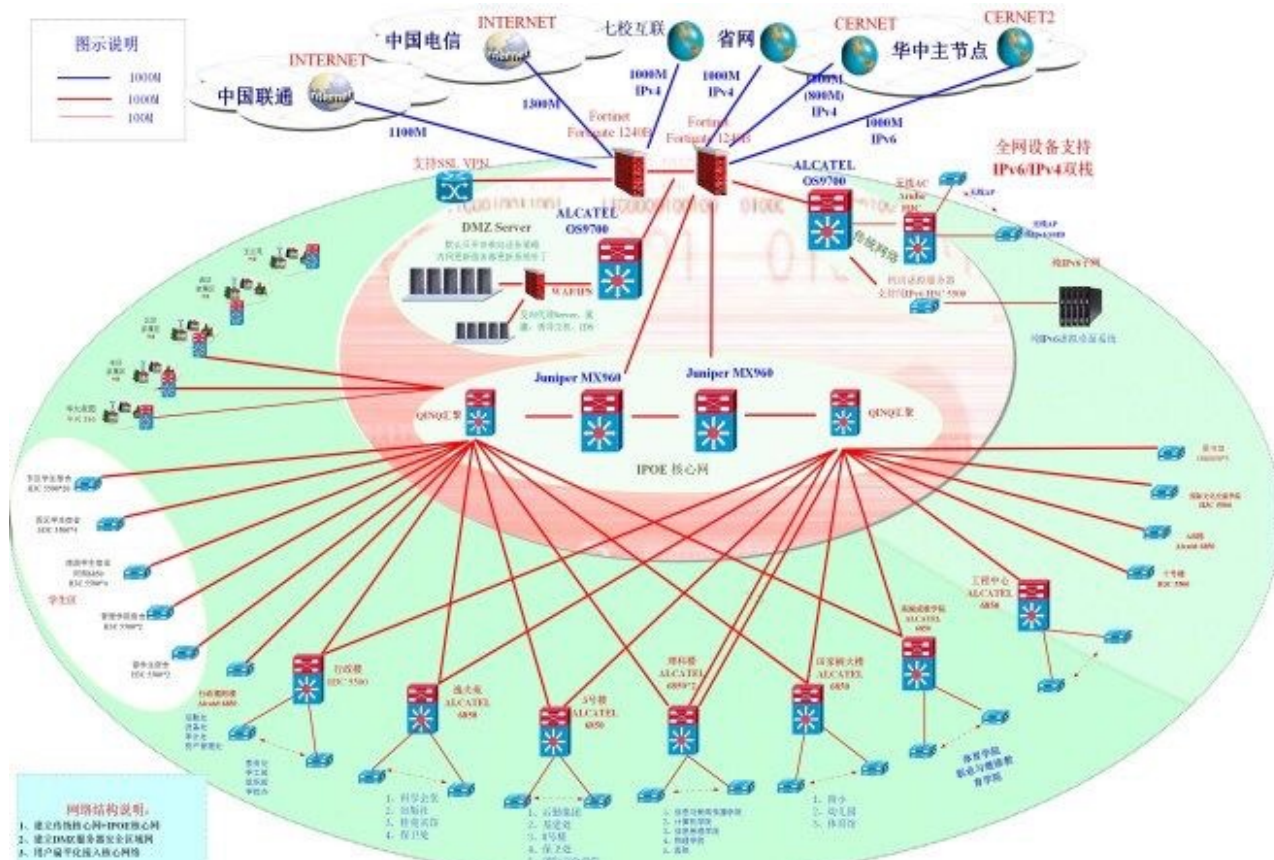
Machine status and jobs activity

- In brief, the existing computers are **too old**
- And has been **limited storage space**
- plus, the network needs improvement in both **bandwidth and latency**
 - according to previous discussion, latency is probably more urgent for user experience
- **Overall, the jobs activity is far from optimistic**



Network Topology of CCNU

- Improvement of Tier 2 network depends on the university
- However, the connection of campus network to "the Internet" is limited, and relies on unexpected resources, i.e. treaty with neighboring universities and bargain with major ISPs



Data flow in/out of Wuhan Tier 2

- Improvement of bandwidth and latency depends on the campus network

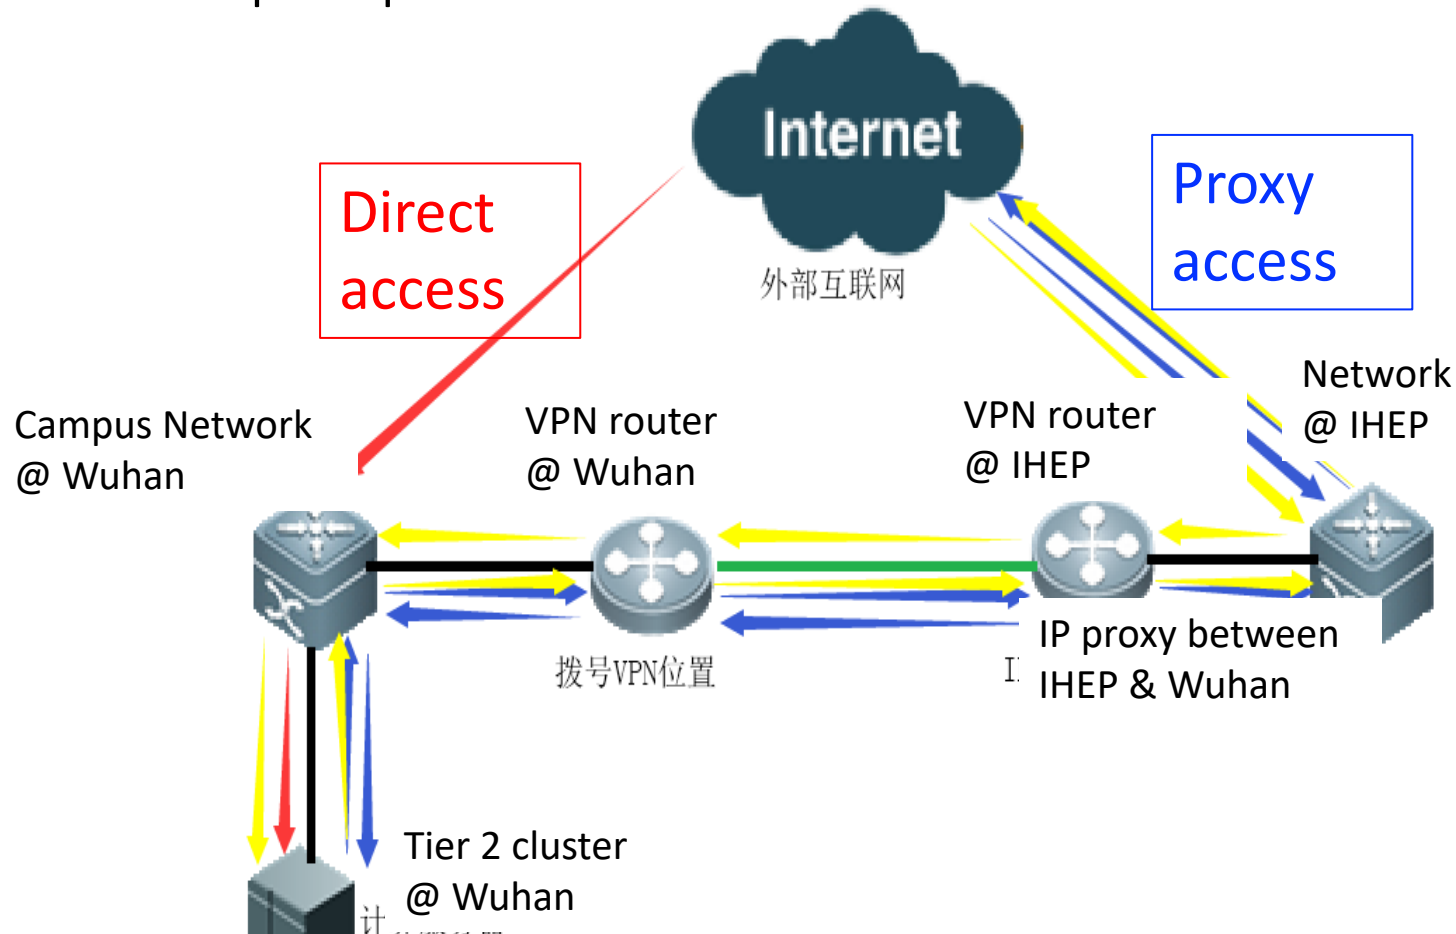
- However, the connection of campus network to "the Internet" is limited, and relies on unexpected resources, i.e. treaty with neighboring universities and bargain with major ISPs

Alternative View: Draft Map

IN from								OUT to							
No.	ID	Site	When	Speed (Mbps)	Hops	RTT (ms)	Streams	No.	ID	Site	When	Speed (Mbps)	Hops	RTT (ms)	Streams
1.	1610123	KISTI-CREAM	09 Aug 2013 10:09	16.78			1	1.	3268629	KISTI_GSDC	15 Nov 2018 20:35	494.95	20	88.28	1
2.	917325	UCT_CERN_RC	15 Jul 2011 21:59	10.24	27	474.07	1	2.	3268004	Tsukuba	15 Nov 2018 04:12	343.95	26	110.90	1
3.	633953	Dortmund	23 Nov 2010 17:54	3.07			1	3.	3270333	ORNL	yesterday 16:48	192.95	25	280.84	1
4.	3265164	Altaria	11 Nov 2018 23:58				1	4.	3268779	RRC_KI_T1	16 Nov 2018 00:34	167.78			1
5.	3076919	Athens	09 Apr 2018 04:38				1	5.	3267505	GRIF_IRFU	14 Nov 2018 15:06	159.39			1
6.	2341574	Bandung	20 Nov 2015 12:41				1	6.	3268349	NIHAM	15 Nov 2018 13:16	142.61	32	312.92	1
7.	3267643	Bari	14 Nov 2018 18:42				1	7.	3267117	Prague	14 Nov 2018 04:47	142.61	30	416.34	1
8.	3264333	Birmingham	11 Nov 2018 01:52				1	8.	3266301	UiB	13 Nov 2018 06:50	142.61	27	285.30	1
9.	3266516	BITP	13 Nov 2018 12:45				1	9.	3267037	HPCS	14 Nov 2018 02:42	134.22			1
			05 Jun					10.	3263502	PAKGRID	10 Nov 2018	134.22	22	254.09	1

Direct connection to IHEP?

- VPN construction between **CCNU** and Institute of High Energy Physics Chinese Academy of Science
- Also coupled with new computer purchase (in next slides)



"Old" resources of whole Department of Physics

However...

Existing computing resources not sufficient for all research purposes:

- LQCD: QCD phase structure, Transport coefficients...
- Hydro simulations
- STAR/ALICE/LHCb data analyses



Batch 1 (2008) 10 servers, 10 * 2 Intel Xeron E5-2620 (2.0GHz, 6 core)CPU, 48GB RAM

Batch 2 (2011) 32 servers, 10 * 2 Intel Xeron E5620 (2.4GHz, 4 core)CPU, 24GB RAM;
22 * 2 Intel Xeron X5650 (2.66GHz, 6 core)CPU, 16GB RAM

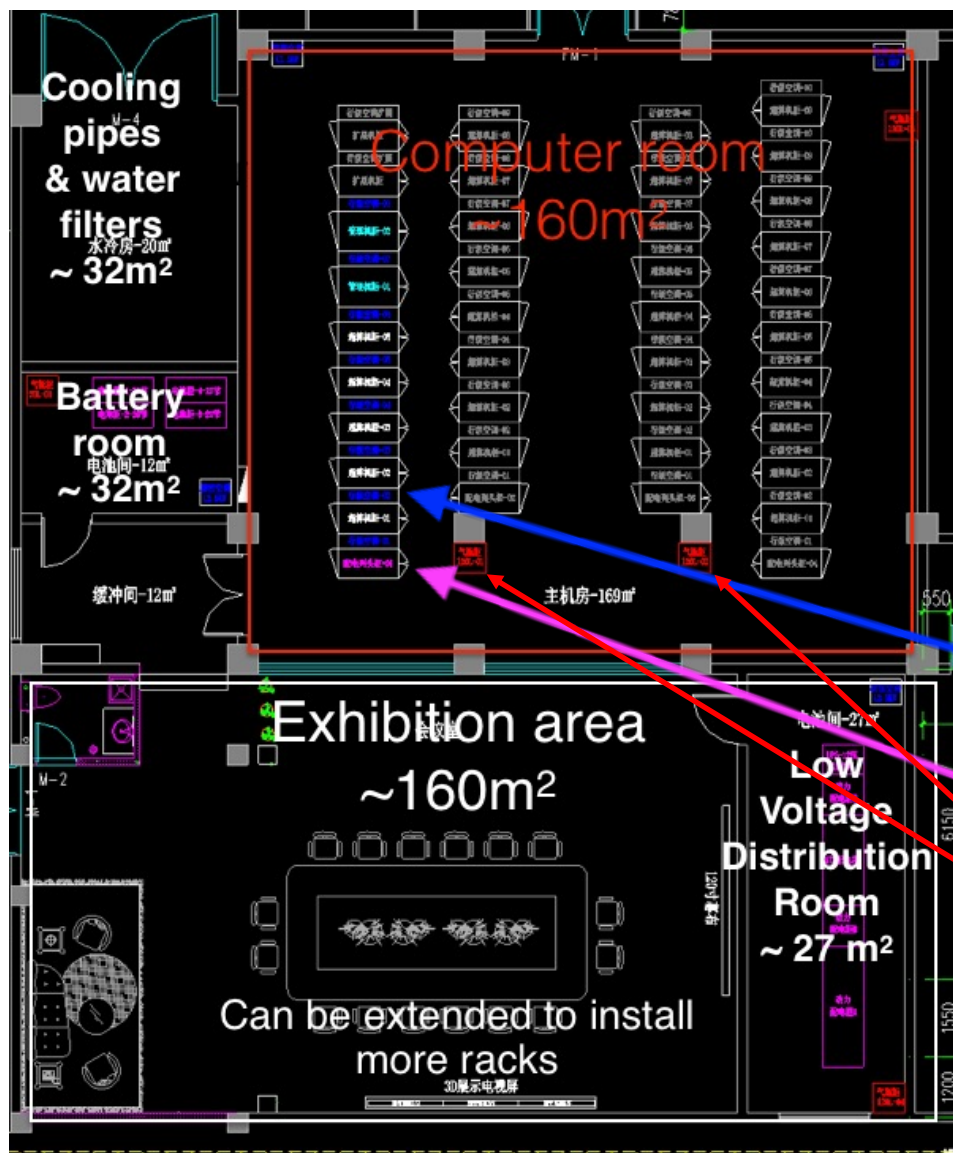
Batch 3 (2012) 55 servers, 42 * 2 Intel Xeron E5-2620 (2.0GHz 6 cores)CPU, 48GB RAM;
13 * 2 Intel Xeron E5410 (2.33GHz, 4 core)CPU, 16GB RAM

Batch 4 (2014) 12 servers, 12 * 2 Intel Xeron E5-2640 (2.50GHz 6 cores)CPU, 48GB RAM

Layout of NSC³



Nuclear Science
Computing Center at CCNU



- ✓ Total area ~ 360 m²
Computer room ~ 160 m²
- ✓ Capacity of 36 Racks
7 racks installed at June 2018
9 racks installed at October 2019
20 racks installed at October 2022
(not yet fully equipped with computers)

Liquid Cooling Package

Electricity rack

Configurations of the GPU cluster

Originally designed for theorists, the computers were **balanced between CPU and GPU**:

In one GPU node:
8 Volta 100 GPU cards
2 skylake-SP (12 cores)
256 GB DDR memory
2 HBA cards

GPU node interconnected
via EDR Infiniband

In total: 18 GPU nodes,
144 GPUs
432 (288 for computing) CPU cores
500 + 500 TB storage



NSC³ prospect in the next years

- Budget cut in computing happened at Wuhan in the last few years
- Fortunately, a new three-year computing budget for the Department of Physics was approved by the university at August 2022
- It includes a 10-million Yuan (~1.3 million US dollars) package, **but that package is for everything: computers, network hardware, ISP, etc.**
- And it is for everybody in the department, so a balance between different users is necessary
- An agreement for the year 2022, about 3-million Yuan budget has been reached this month:
- It focus on CPU purchase by the strong support of HEP experimentalists, which fits the need of ALICE Tier 2 (and possible LHCb Tier2 at Wuhan)

The current purchase plan

- **Computing nodes**

26 computing servers, each with 2 Xeon CPUs, each CPU \geq 32 cores, each core with 6GB RAM (by average)
each server with two 200GB IB card for network

- **Managing nodes**

2 gateway machines

- **Network**

1 IB switch with 5 800G modules and 20 200G modules

1 10-Giga Management switch with 48 10G SFP+fiber and 6 40G QSFP+fiber

- **Software**

Job submission and usage statistics (for billing)

But please keep in mind:

these are for share (e.g. no storage purchase in plan)

and not designed for off-campus usage (e.g. network)


Future Plan



Nuclear Science
Computing Center at CCNU

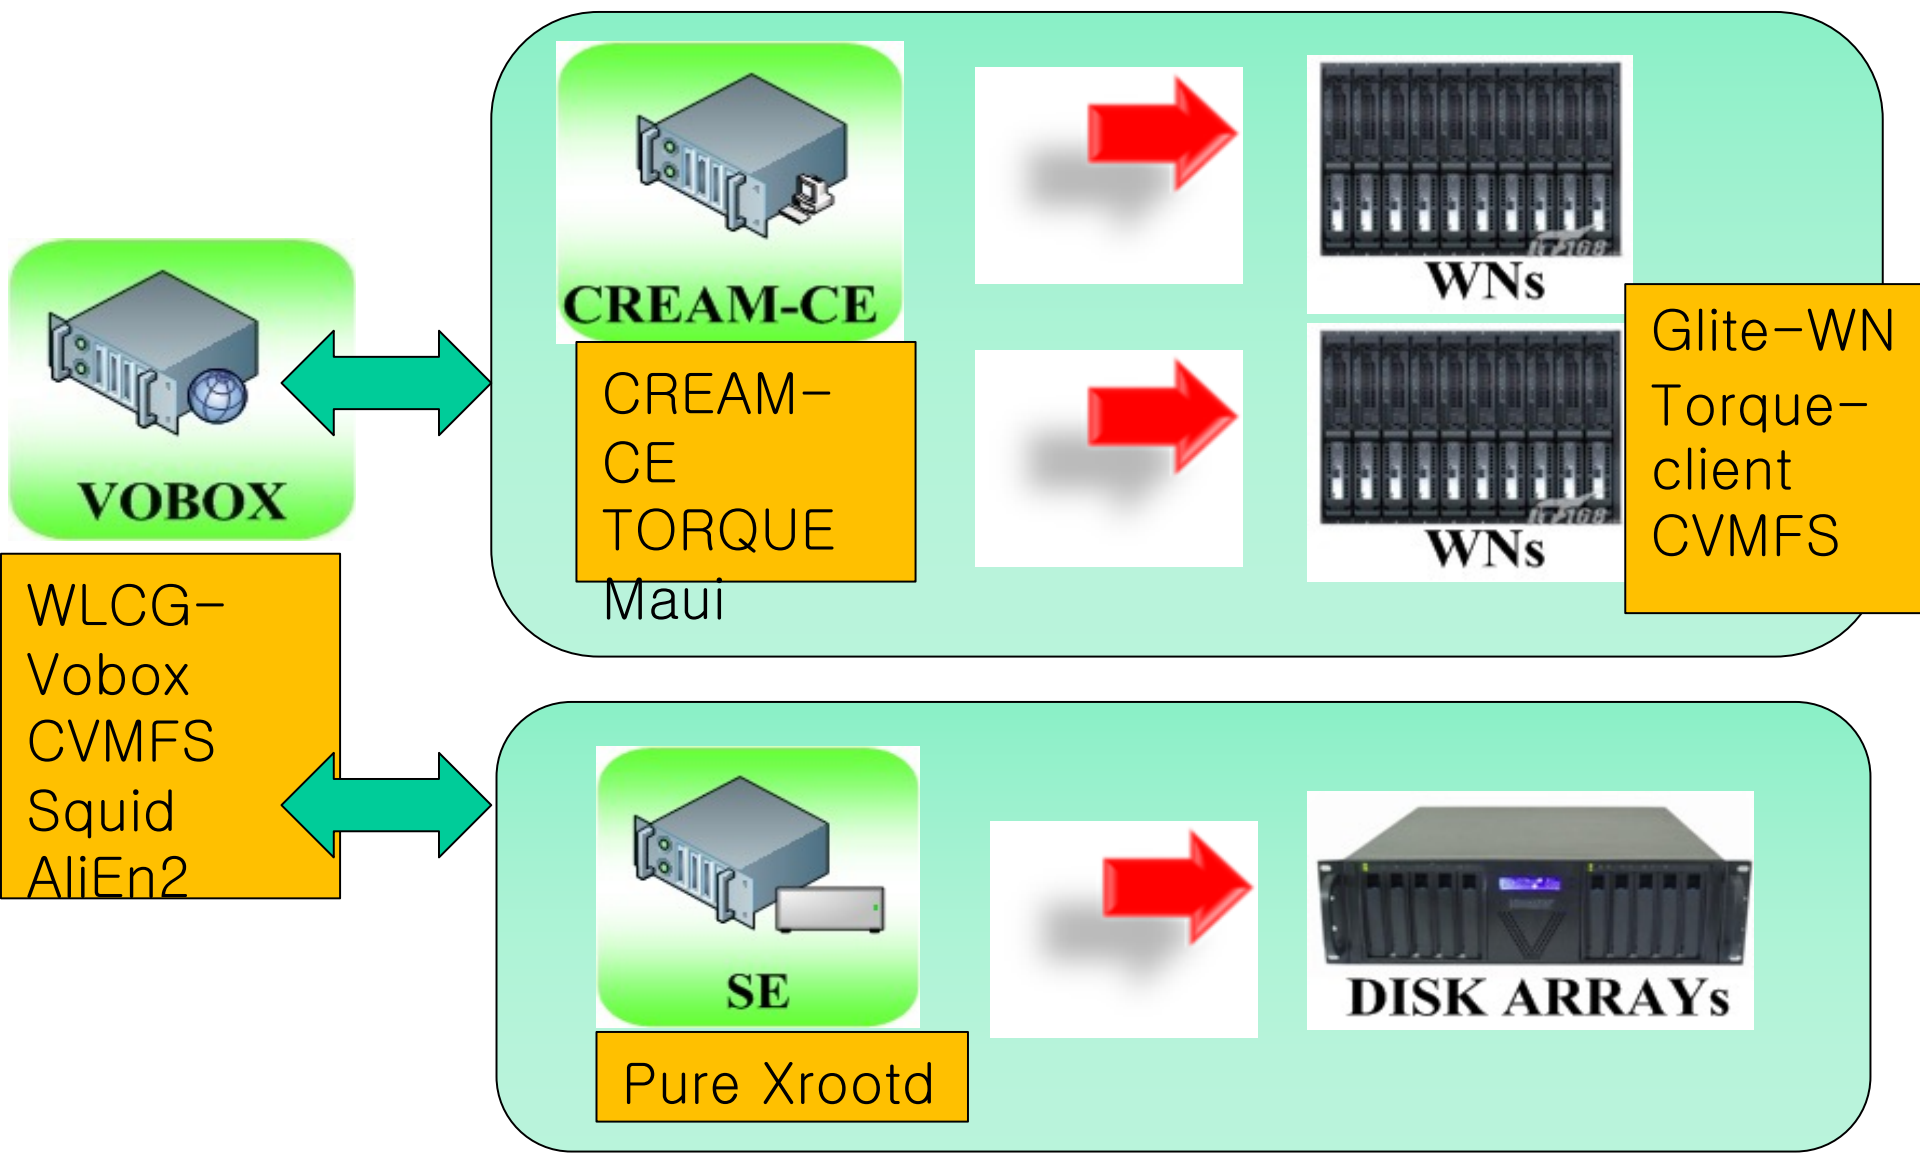
- Computer Grids joint management with IHEP
- Not only renting network from IHEP to Wuhan, but possible **shift of (partial) newly purchased computers to IHEP site**
- But that decision can't be made by ALICE group at Wuhan: Moving properties worth one million US dollars off-campus needs improvement of university management
- On the other hand, ALICE group at Fudan university (Shanghai) has proposed to share the budget of ALICE Tier 2, **including storage and other accessories**
- The expected (much) improved performance at IHEP than Wuhan site will be a good reason for Fudan university to shift their hard drives to IHEP
 - **while Wuhan network can hardly be better than Shanghai**

Thank you for your attention and suggestion!



Backup

Topology of cluster



Machine status

Select site: Wuhan

MonALISA information Version: 13.11.04 (JDK 1.8.0_92)
Running on: 172.16.30.1
Administrator: Jianlin Zhu <Jianlin.Zhu@cern.ch>

Service health NTP: SYNC, offset: 0.009s

Services status ClusterMonitor: OK
AliEn: v2-19.395 PackMan: n/a
CE: OK
CE info: At the moment we are busy (we ...
Max running jobs:
Max queued jobs:

Proxies status AliEn proxy: OK (1 day, 23:27)
Delegated proxy: OK (1 day, 23:59)
Proxy server: OK (95 days, 08:41)
Proxy of the machine: OK (20:30)

Current jobs status Assigned: 0
Running: 216
Saving: 0

Accounting (last 24h) Success jobs: 150 (profile)
Error jobs: 486 + 10 expired
kSI2k units: 670 / pledged

Site averages Active nodes: 10
(last 24h) Average kSI2k/core: 2.872

Storages status

Name	Status	Size	Used	Free	Usage	No of files	Type	ADD test
ALICE::Wuhan::SE								

VoBox health


CPU: 8x 1596MHz
Mem usage: 10.57% of 15.56 GB
Processes: 266
Sockets: 57 TCP / 26 UDP
Uptime: 29 days, 15:29

CPU usage (last 1h avg) Load: 0.076
User: 0.526%
System: 0.14%
IOWait: 0.067%
Idle: 99.26%

Int: 0%
Soft int: 0.003%
Nice: 0%
Steal: 0%

AliEn LDAP var	VoBox path	Size	Used	Free	Use%
TMP	/home/sgmali01/ALICE/tmp	70.65 GB	4.574 GB	62.49 GB	7%
LOG	/home/sgmali01/ALICE/alien-logs	70.65 GB	4.574 GB	62.49 GB	7%
CACHE	/home/sgmali01/ALICE/cache	70.65 GB	4.574 GB	62.49 GB	7%

Active jobs in Wuhan

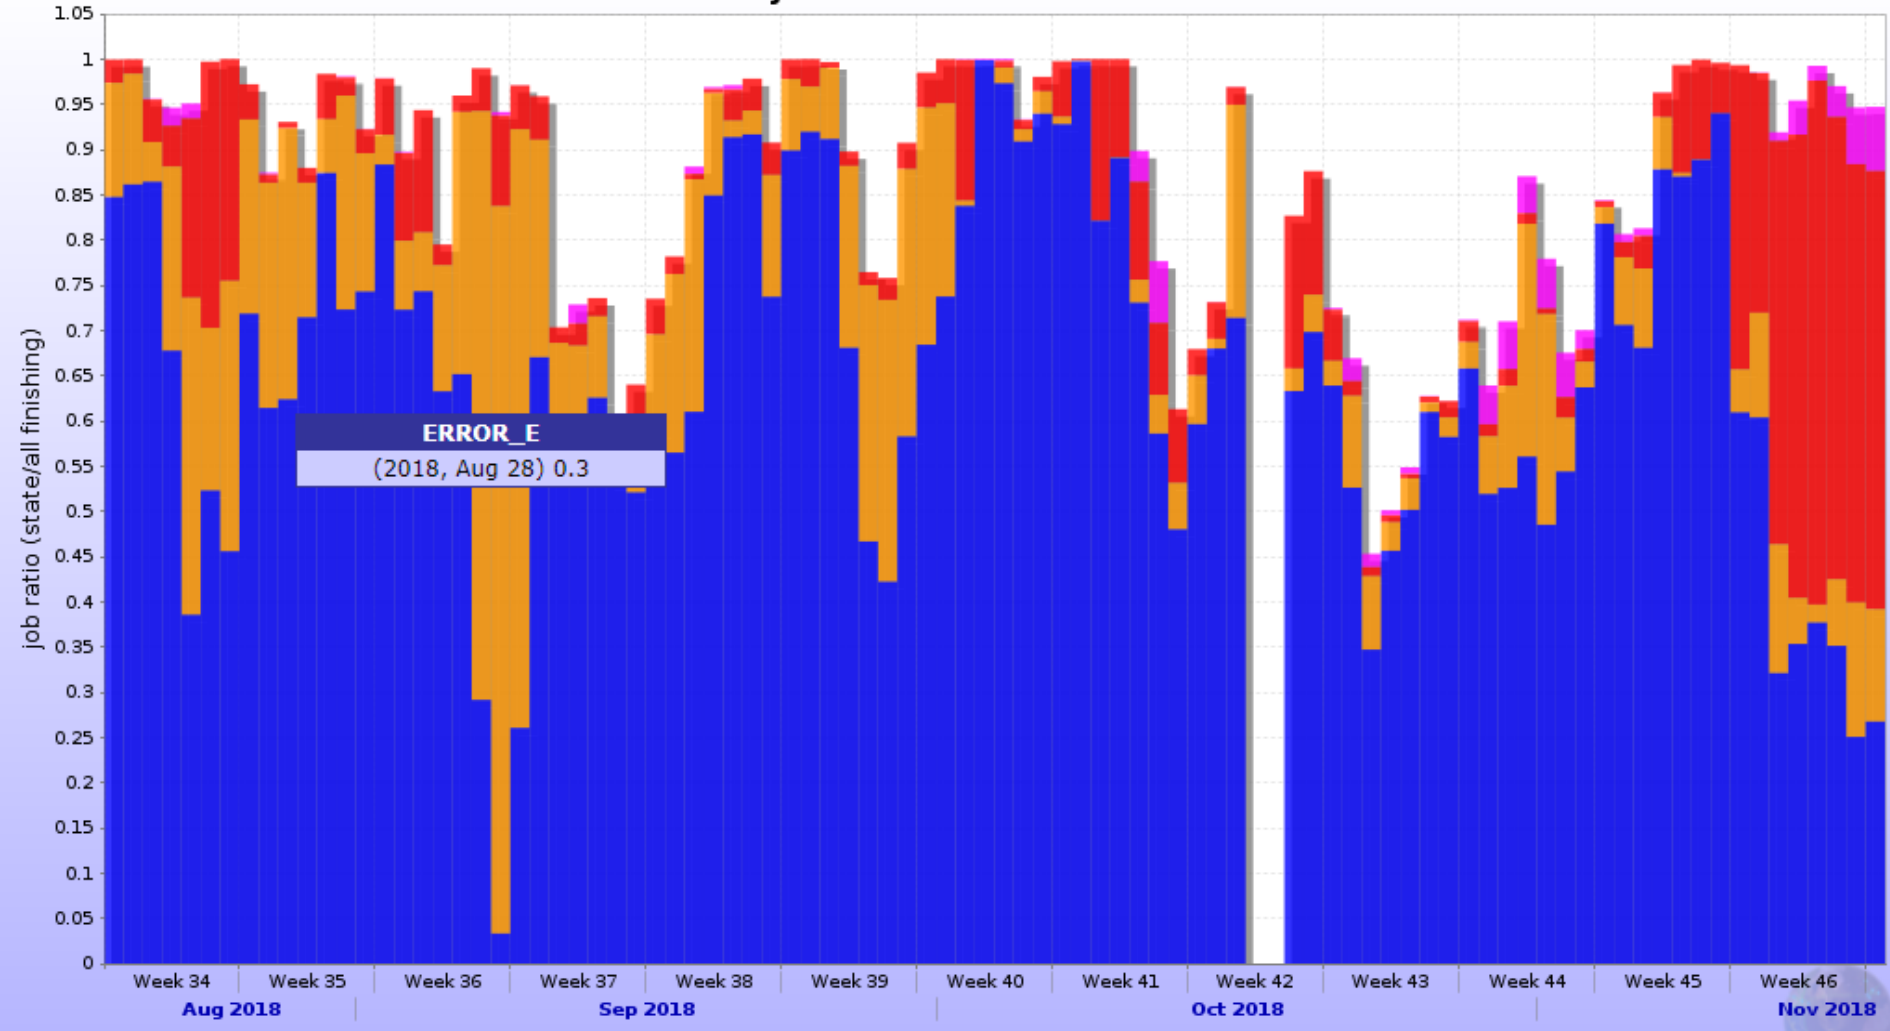


Job state ratio at Wuhan

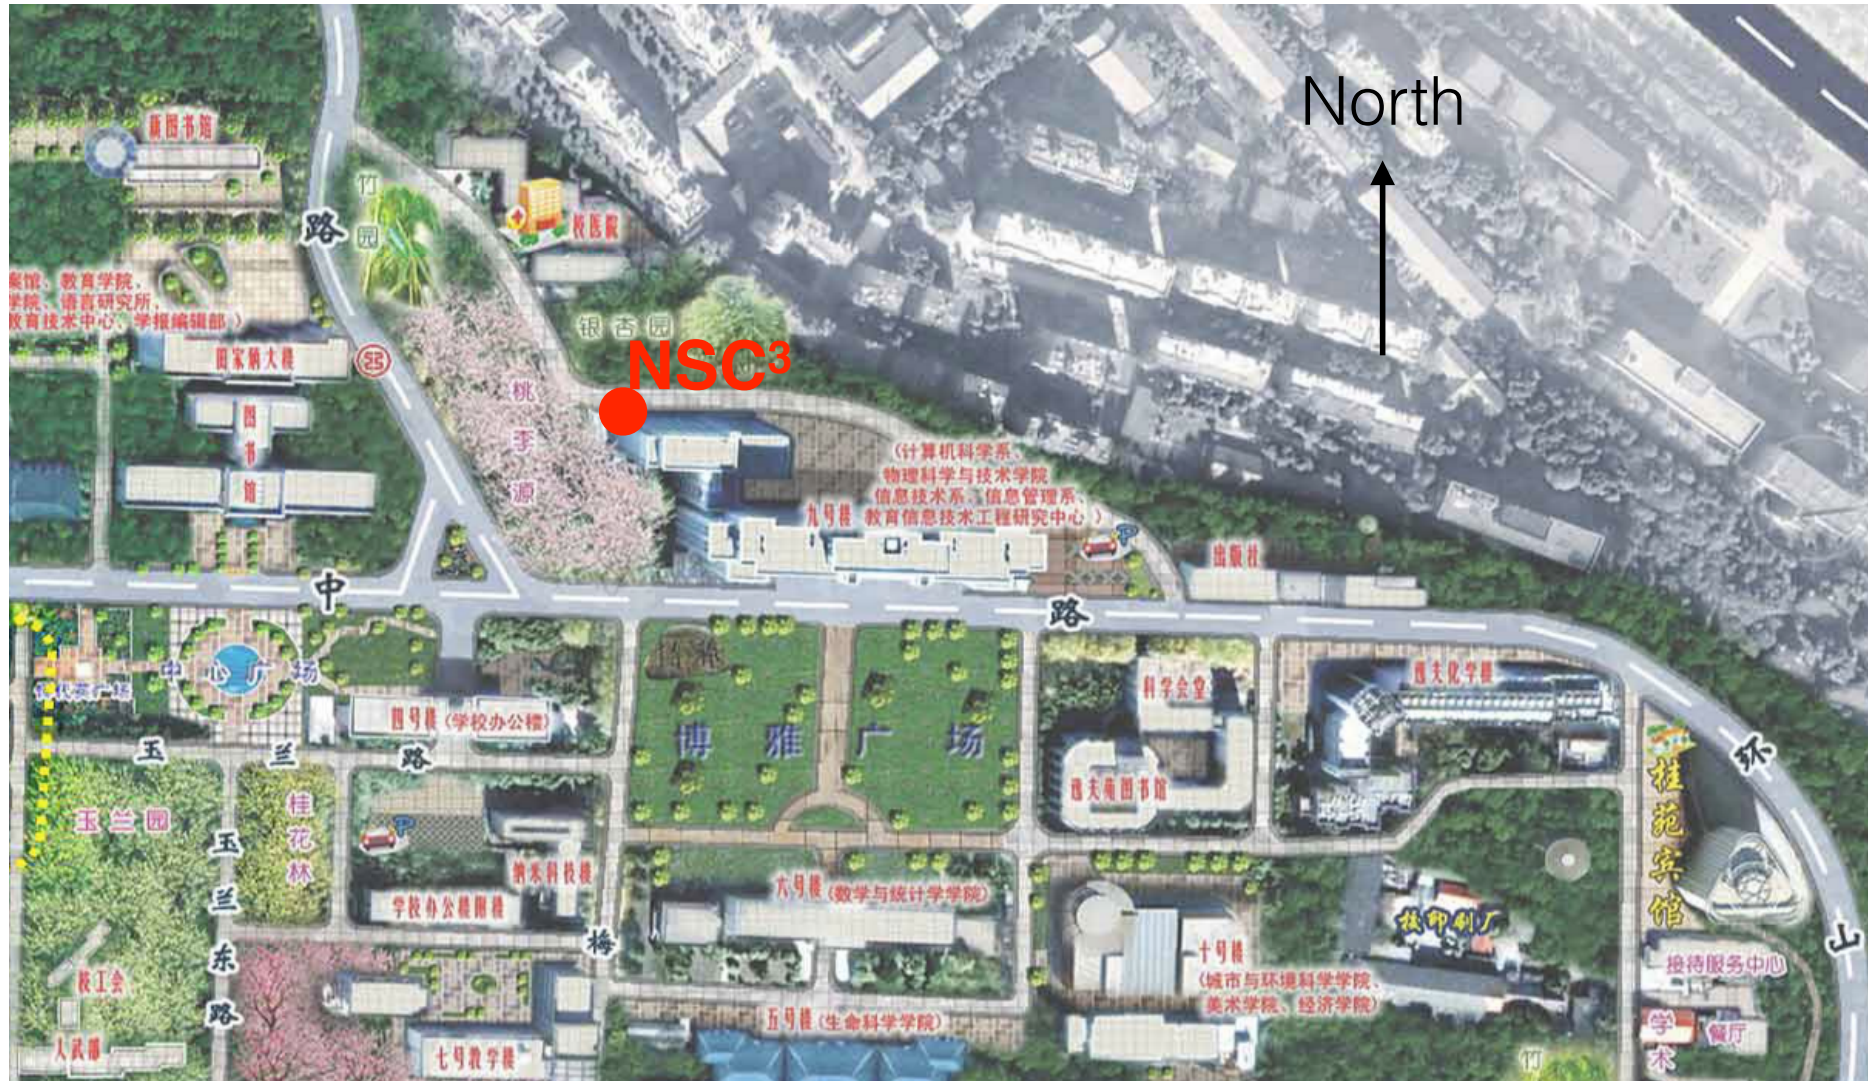
ons

what is this

Final job state ratio at Wuhan



Location of NSC³



Data flow in/out of Wuhan Tier 2

<Wuhan>

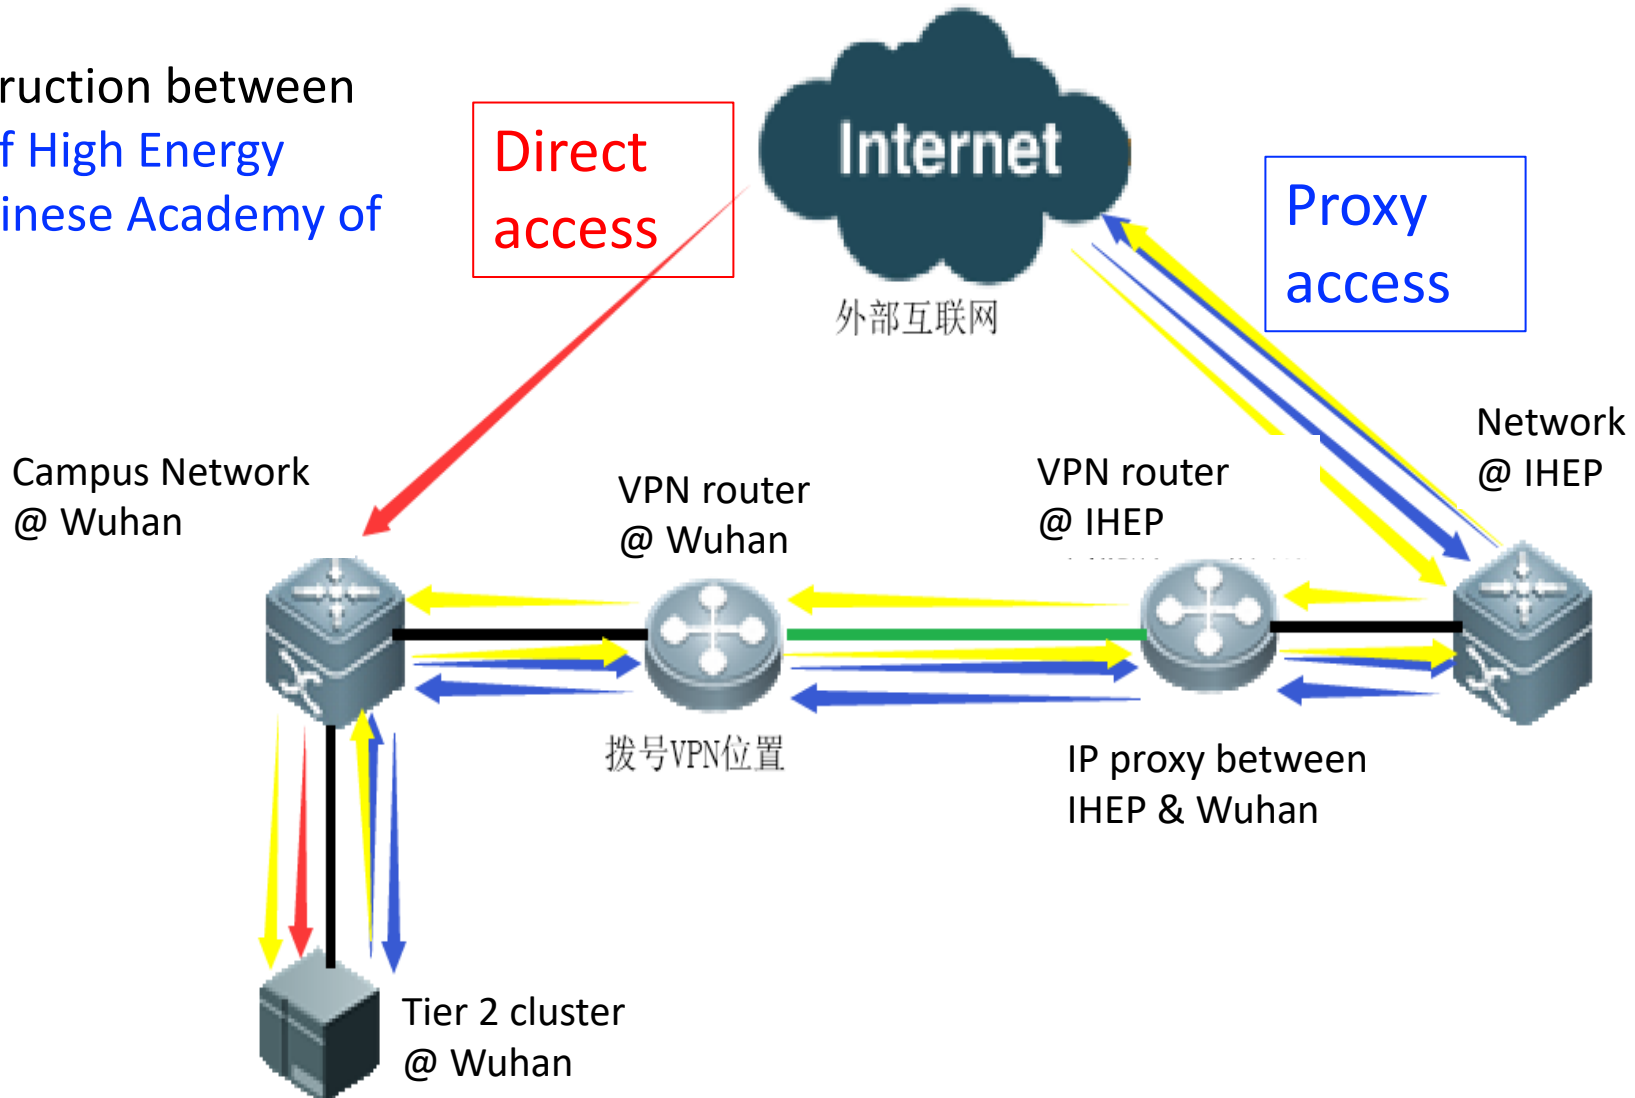
Alternative views: [Chart](#) | [Map](#)

IN from							
No.	ID	Site	When	Speed (Mbps)	Hops	RTT (ms)	Streams
1.	1610123	KISTI-CREAM	09 Aug 2013 10:09	16.78			1
2.	917325	UCT_CERN_RC	15 Jul 2011 21:59	10.24	27	474.07	1
3.	633953	Dortmund	23 Nov 2010 17:54	3.07			1
4.	3265164	Altaria	11 Nov 2018 23:58				1
5.	3076919	Athens	09 Apr 2018 04:38				1
6.	2341574	Bandung	20 Nov 2015 12:41				1
7.	3267643	Bari	14 Nov 2018 18:42				1
8.	3264333	Birmingham	11 Nov 2018 01:52				1
9.	3266516	BITP	13 Nov 2018 12:45				1
			05 Jun				

OUT to							
No.	ID	Site	When	Speed (Mbps)	Hops	RTT (ms)	Streams
1.	3268629	KISTI_GSDC	15 Nov 2018 20:35	494.95	20	88.28	1
2.	3268004	Tsukuba	15 Nov 2018 04:12	343.95	26	110.90	1
3.	3270333	ORNL	yesterday 16:48	192.95	25	280.84	1
4.	3268779	RRC_KI_T1	16 Nov 2018 00:34	167.78			1
5.	3267505	GRIF_IRFU	14 Nov 2018 15:06	159.39			1
6.	3268349	NIHAM	15 Nov 2018 13:16	142.61	32	312.92	1
7.	3267117	Prague	14 Nov 2018 04:47	142.61	30	416.34	1
8.	3266301	UiB	13 Nov 2018 06:50	142.61	27	285.30	1
9.	3267037	HPCS	14 Nov 2018 02:42	134.22			1
10.	3263502	PAKGRID	10 Nov 2018	134.22	22	254.09	1

Bandwidth dependence on routes

VPN construction between
Institute of High Energy
Physics Chinese Academy of
Science and **CCNU**



Big space to improve

However...

Existing computing resources not sufficient for all research purposes:

- LQCD: QCD phase structure, Transport coefficients...
- Hydro simulations
- STAR/ALICE/LHCb data analyses



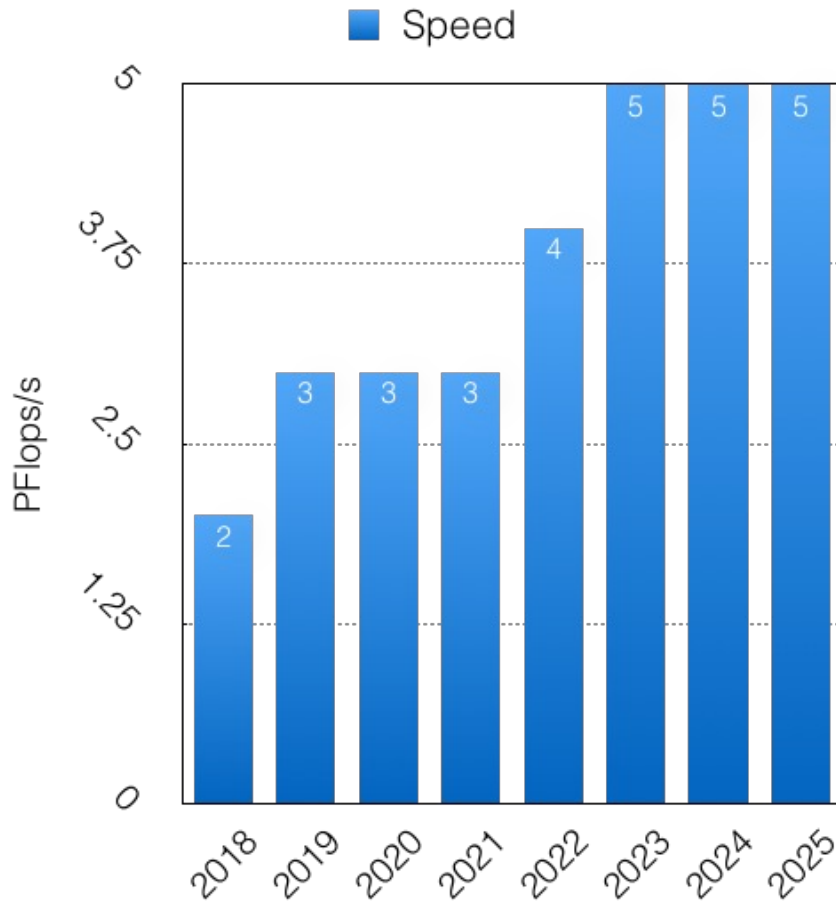
Batch 1 (2008) 10 servers, 10 * 2 Intel Xeron E5-2620 (2.0GHz, 6 core)CPU, 48GB RAM

Batch 2 (2011) 32 servers, 10 * 2 Intel Xeron E5620 (2.4GHz, 4 core)CPU, 24GB RAM;
22 * 2 Intel Xeron X5650 (2.66GHz, 6 core)CPU, 16GB RAM

Batch 3 (2012) 55 servers, 42 * 2 Intel Xeron E5-2620 (2.0GHz 6 cores)CPU, 48GB RAM;
13 * 2 Intel Xeron E5410 (2.33GHz, 4 core)CPU, 16GB RAM

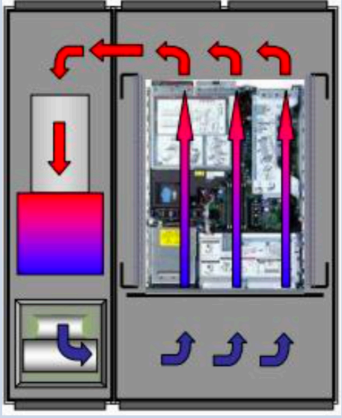
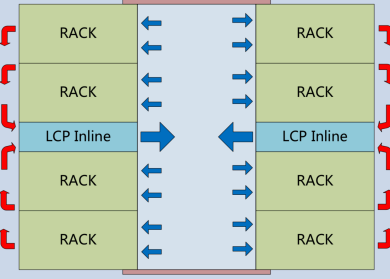
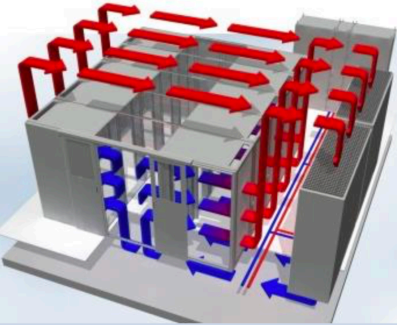
Batch 4 (2014) 12 servers, 12 * 2 Intel Xeron E5-2640 (2.50GHz 6 cores)CPU, 48GB RAM

NSC³ prospect in the next years




- LQCD: QCD phase structure, Transport coefficients...
Hydro simulations
STAR/ALICE/LHCb data analyses
...
- Programs to be determined by a committee
- Domestic and international collaborations are welcome

Cooling strategies

Strategies of cooling	Cool each rack	Cool each two rows of racks	Cool the room
Layout			
Cooling target	IT devices within a rack	IT devices with 1 or 2 rows of racks	whole computer room
Blow distance	0.1~0.8m	1.2~3m	3~15m
Cooling efficiency	High	Medium	Low
Application	High Performance, Green computing	Cloud computing, mixed	Traditional computing room

Temperature distribution in a typical year of Wuhan



There are 4516 hours yearly in Wuhan that $T < 18^{\circ}\text{C}$

51% of the year free cooling is available

Water Cooling

