Grid and Cloud Computing at IHEP in China

Weidong Li

IHEP, CAS, Beijing

NEC2017 at Budva

09-25-2017

Contents

- Introduction
- DIRAC-based distributed computing
- Progress with IHEPCloud
- Running BOINC jobs on WLCG Tier 2
- Summary

Experiments

BESIII (Beijing Spectrometer III at BEPCII)

DYB (Daya Bay Reactor Neutrino Experiment)

JUNO (Jiangmen Underground Neutrino Observatory)

YBJ (Tibet-ASgamma ARGO-YBJ Experiments)

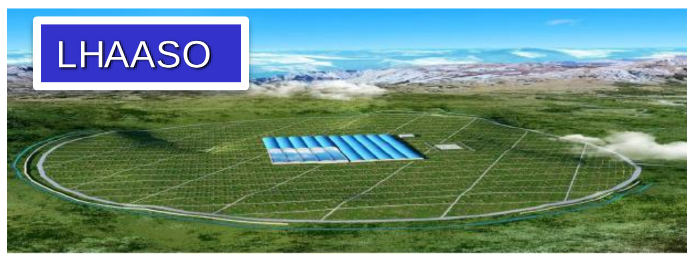

Large High Altitude Air Shower Observatory

Hard X-Ray Moderate Telescope

Circular Electron Positron Collider

Computing resources

- HTCondor clusters
  - HTC jobs: series jobs
  - ~13,500 CPU cores
- SLURM clusters
  - HPC jobs: parallel jobs
  - ~2,750 CPU cores
- Grid site (WLCG)
  - Tier 2 site
  - 1,200 CPU Cores
- The DIRAC-based distributed computing system
  - ~ 3,000 CPU cores
- IHEPCloud based on Openstack
  - ~ 2000 CPU cores

Storage

- Lustre as main disk storage
  - Capacity: 9 PB storage
- Gluster system
  - 734TB storage with replica feature
- DPM & dCache
  - 940TB, With SRM interface
- EOS system
  - 1.2 PB
- HSM using modified CASTOR
  - 2 tape libraries + 2 robots, 26 LTO4 drivers and 1 LTO7 driver
  - Capacity: 5 PB

Network at IHEP

- For office users
  - The largest campus network and bandwidth among all CAS institutes
    - 10G backbone
    - IPv4/IPv6 dual-stack
    - Wireless covered at (>250 APs)
  - Email/web/ services
  - >3000 end users
- For the data center at computing center
  - 160 Gbps (4X40Gbps) for 2-layer switches
  - 2X10 Gbps for storage nodes

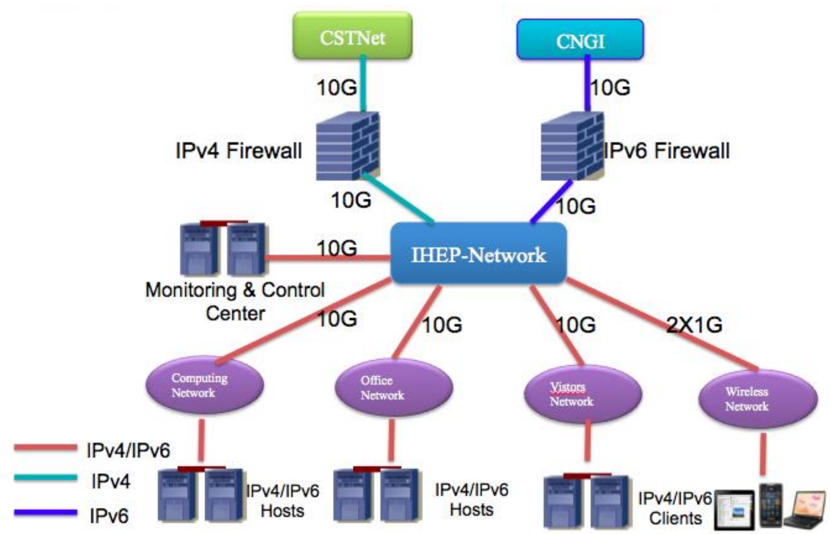

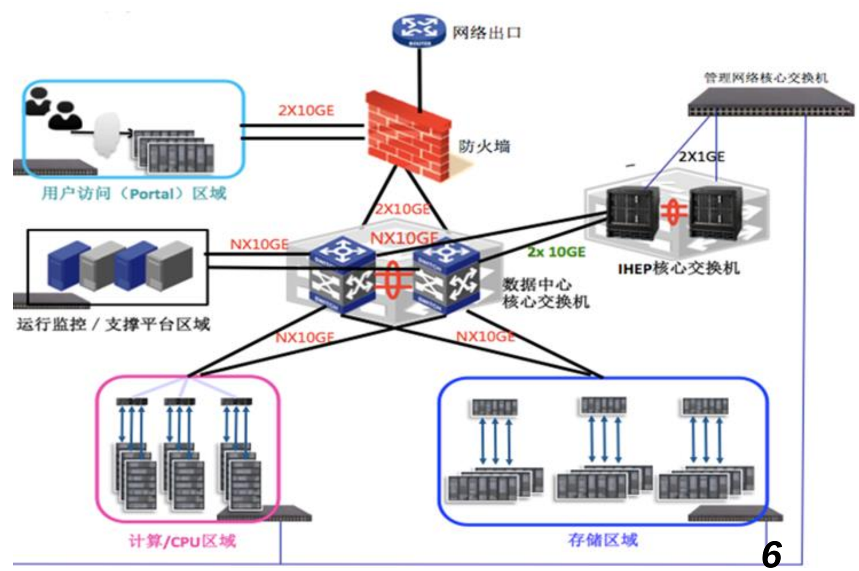

International and domestic links

- Good Internet connections
  - IHEP-Europe: 10 Gbps
  - IHEP-USA: 10 Gbps
  - ~4 PB/year data exchange

- Dedicated Links for three other IHEP sites (two in the future)
  - Shenzhen (Dayabay)
  - Dongguan (CSNS)
  - Tibet (YBJ/ARGO)
  - Kaiping (JUNO)
  - Chengdu (LHAASO)

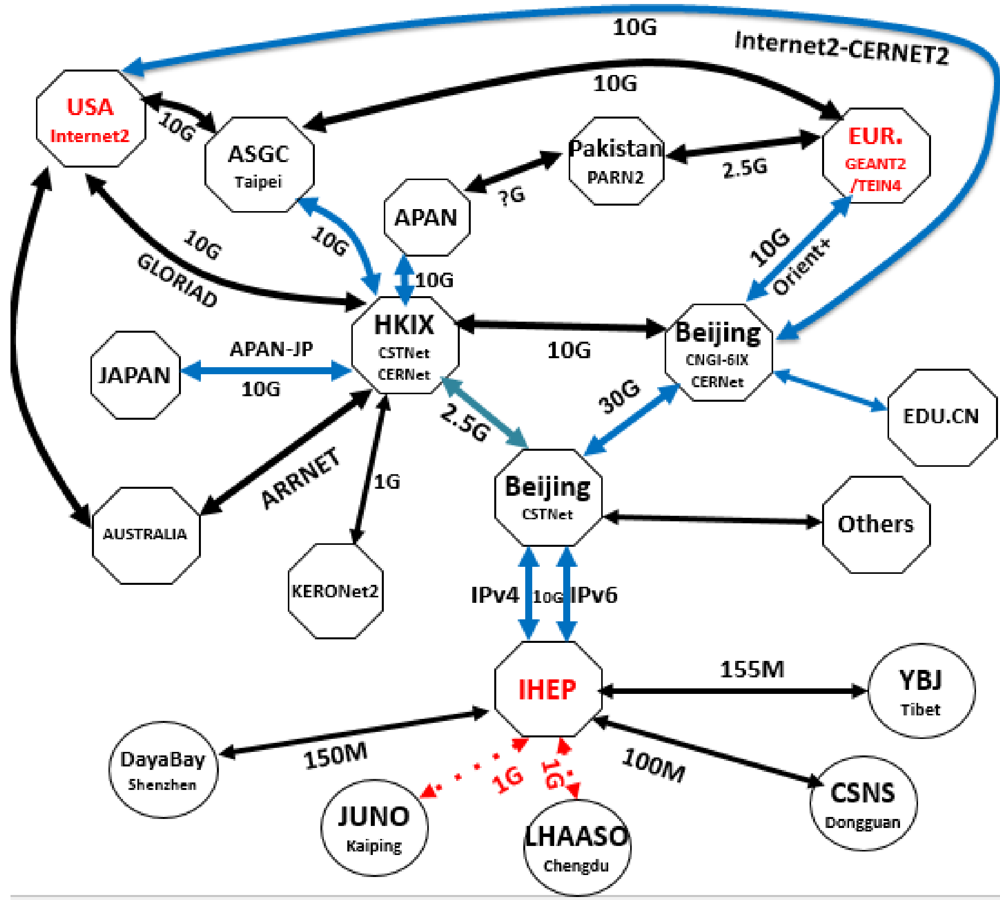

DIRAC-based Distributed Computing

- Have developed the distributed computing system with DIRAC in collaboration with JINR.
  - DIRAC: a general purpose framework for distributed computing
- Currently is being used by several experiments, including BESIII, JUNO, CEPC.
- 15 Sites from USA, Italy, Russia, Turkey, Taiwan, Chinese Universities etc.
- Network: planning to join LHCONE to further improve the network performance.

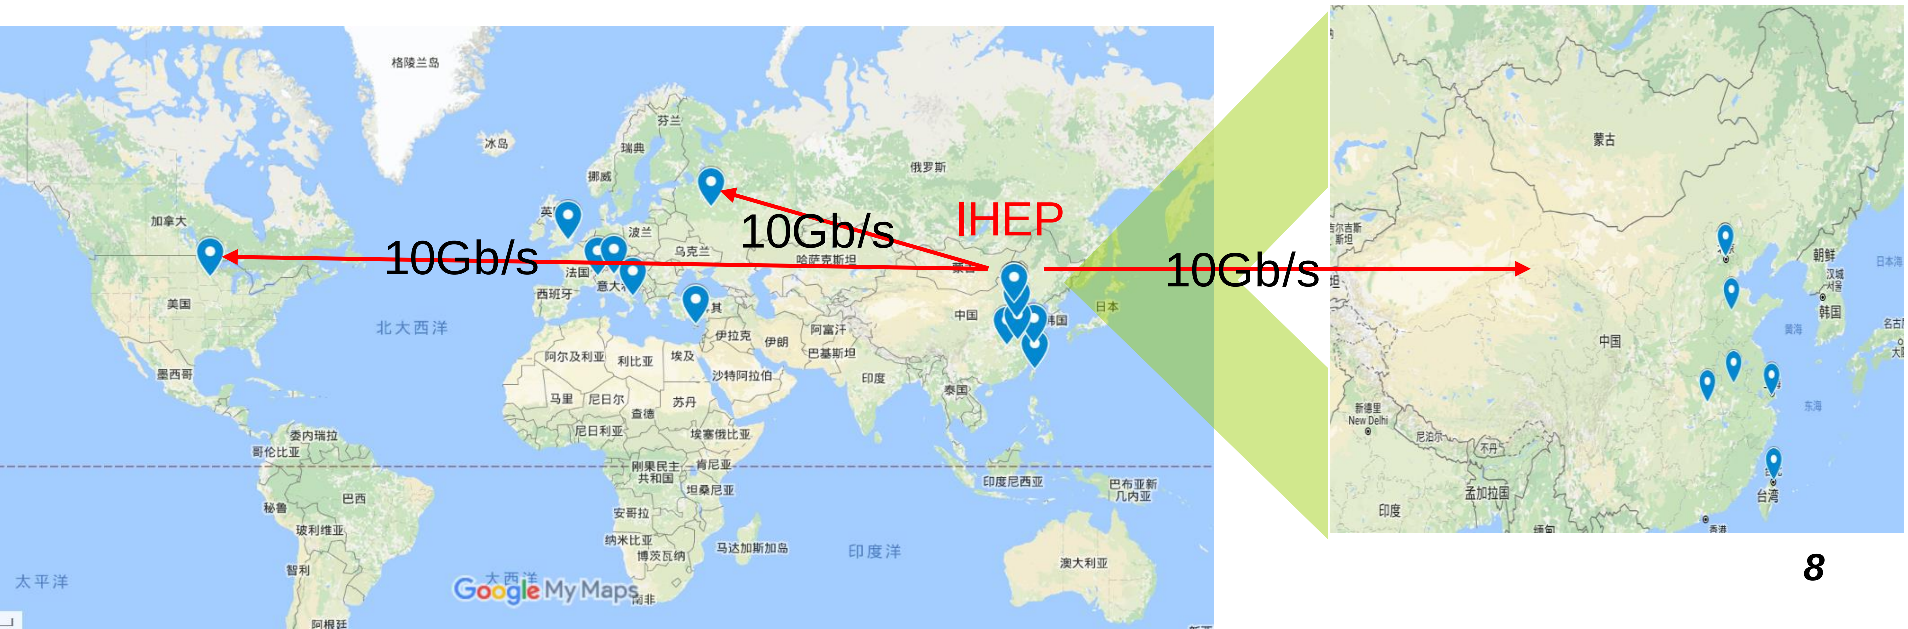

System monitoring (1)

- Monitoring is the key point to assure high availability of the system
  - Especially sites not able to provide manpower or experts for maintenance
- A central monitoring system was designed with two layers to help manage sites and track down problems .
  - System level based on Nagios and Gearman
    - Quickly know health of WNs in the site
  - Application level based on DIRAC
    - Check application environment of a site with SAM test jobs

Site	SiteType		MaskStatus	CE-Test	SE-Test	Storage Usage(%)	Efficiency(%)	Job Usage(%)	WN Status
GRID.QMUL.uk	GRID		Active	Busy		0	92	22.6	Warn
CLOUD.IHEP-OPENS...	CLOUD		Active	Busy	OK	41	100	47	OK
CLOUD.JINR.ru	CLOUD		Active	OK	OK	41	99.5	0	OK
GRID.INFN-ReCas.it	GRID		Active	OK			93.9	0	OK
GRID.INFN-Torino.it	GRID		Active	Bad			0	0	
CLUSTER.IPAS.tw	CLUSTER		Active	Busy	OK	41	81.4	36.8	Warn
CLOUD.IHEP-OPEN...	CLOUD		Active	Busy	OK	41	100	9.2	OK
GRID.JINR.ru	GRID		Active	OK	OK	41	90.1	6.5	Warn
CLOUD.IHEPCLOUD...	CLOUD		Active	Bad	OK	41	0		
CLUSTER.SJTU.cn	CLUSTER		Active	Bad	OK	41	100	0	OK
CLOUD.TORINO.it	CLOUD		Active	Bad			0	0	
CLOUD.INFN-PADO...	CLOUD		Active	Bad			100	0	OK
CLUSTER.NEU.tr	CLUSTER		Active	OK	OK	0	100	0	OK
CLUSTER.USTC.cn	CLUSTER		Active	OK	OK	83.9	99.2	2.5	OK
CLUSTER.UMN.us	CLUSTER		Active	OK	Bad	61	93.3	0	OK

System monitoring (2)

- Collect and report status of network, I/O and memory from remote hosts
- Implementation with Nagios and Gearman
  - Nagios Daemon on WNs Collect info of host and service
  - Errors and warnings are sent to IHEP central system with Gearman
  - Central master can send warning messages through email or Wechat
- With help of the system, IHEP experts can work on problems remotely.
- The System can complete 20,000 checks within 5 minutes.

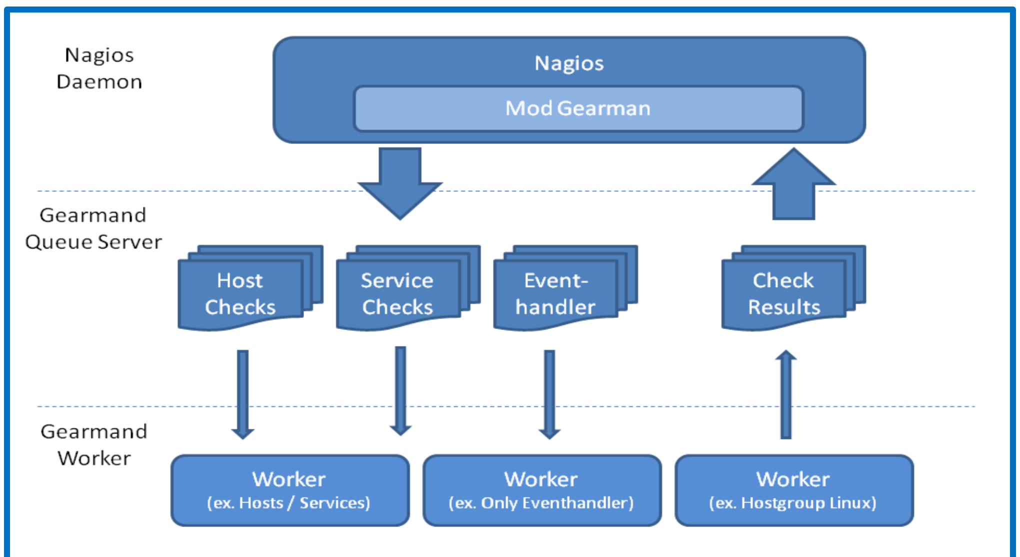

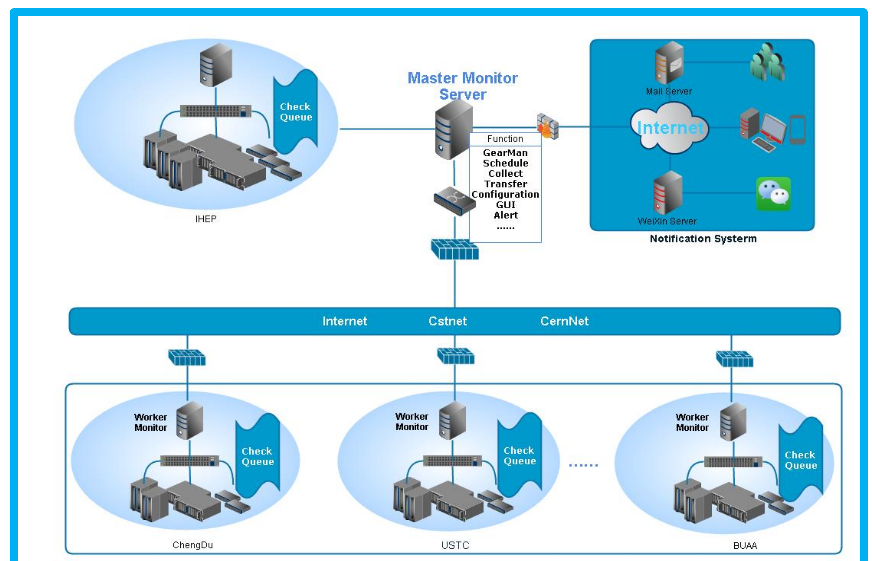

JUNO distributed computing

- The JUNO experiment is expected to take data in 2020
  - Estimated to produce 2 PB data/year for 20 years
- JUNO computing model is being investigated
- IHEP as the main centre
  - Play major role for data processing
  - Maintain central storage for all the data
  - Plan: more than 10,000 CPU cores
- European sub-centres
  - Could be several sites with shared storage
  - Mainly for simulation and analysis
  - Storage for simulated data and data analysis, even raw data as backup
- Be able to make use of opportunistic computing resources

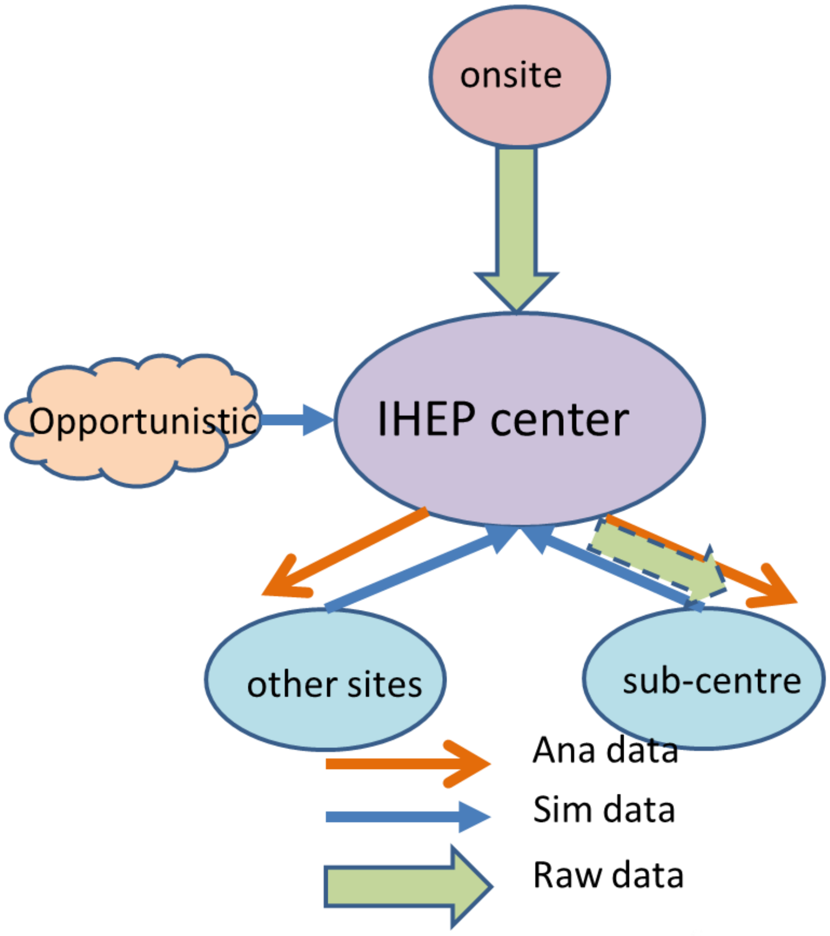

JUNO offline Software

- JUNO offline software has been developed on the underlying framework called SNiPER
- Parallelization was implemented based on TBB (Threading Building Blocks)
  - Enable multi-thread and multi-process applications
  - Take advantage of features of the new hardware architecture
- Event-level parallelization is available in the prototype
  - SNiPER Muster developed for multi-threaded computing
  - SNiPER Task instances can be invoked by TBB tasks

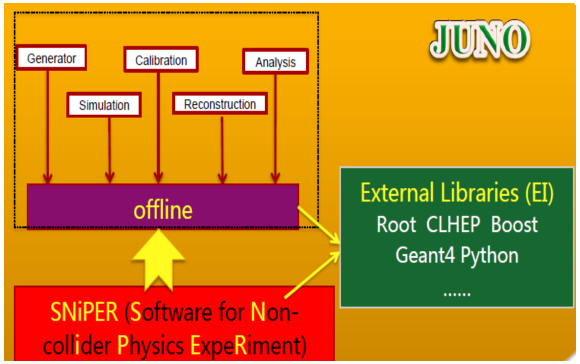

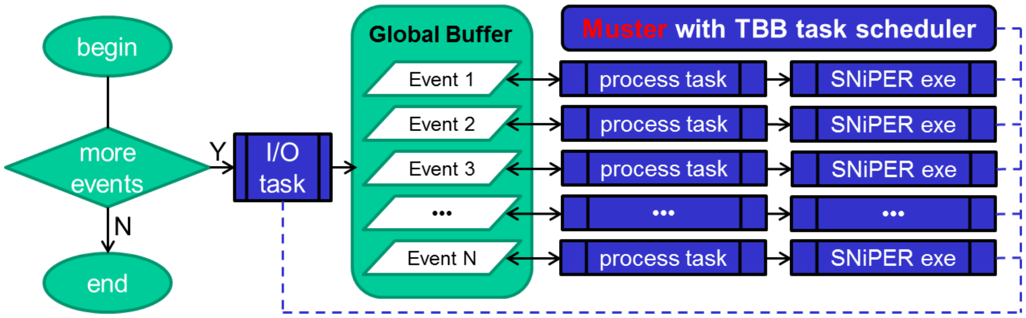

Multi-core supports (1)

- In order to support JUNO’ s multi-core jobs, Multi-core(Mcore) scheduling is being studied
  - Single-core scheduling: One-core pilot pull one-core job (A)
  - multi-core scheduling: Two ways are considered
    - Mcore pilots added to pull Mcore jobs (B)
      - M-core pilots pull M-core jobs
      - Pilot “starving” when matching with mixture of n-core job (n=1, 2...)
    - Standard-size pilots with dynamic slots (C)
      - Standard-size: whole-node, 4-node, 8-node....
      - Pilots pull n-core jobs (n=1, 2.....) until all internal slots are used up

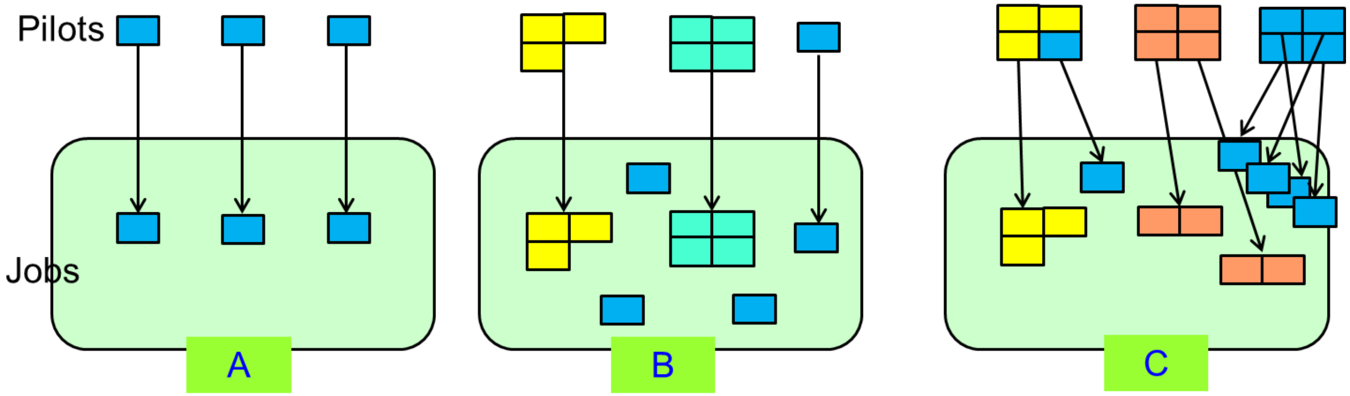

Multi-core supports(2)

- Basic tests have been done on a SLURM testbed
- CPU utilization rate has been measured and compared between single-core and multi-core modes
- With mixture of jobs, the utilization rate of standard-size pilots (C) is higher than that of (B), but still lower than the original single-job case (A) which are 92%
- In (B), the efficiency is low, mostly because some pilots are idle without proper jobs being pulled.
- In (C), some pilots did get jobs, but not fulfill the whole node

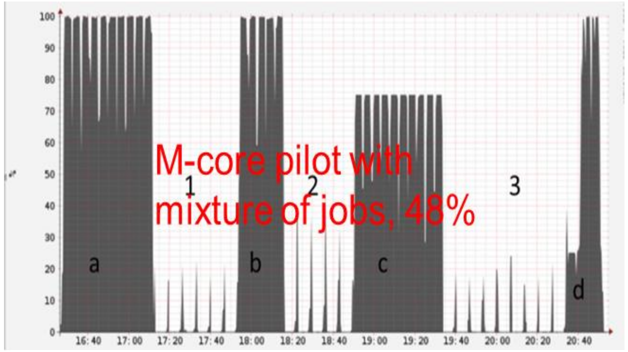

(B)

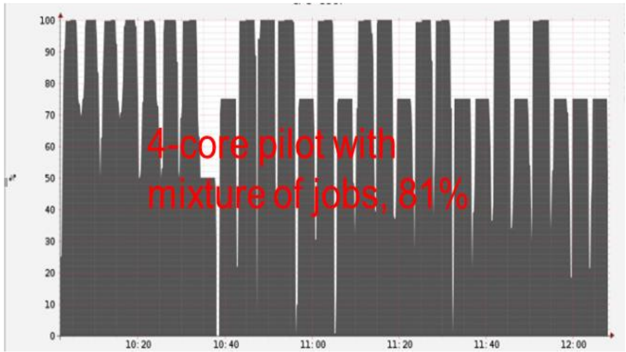

(C)

Virtual computing cluster

- To allocate resources on demand to improve resource utilization
- Resource integration and sharing among different experiments at IHEP
- To meet peak demand

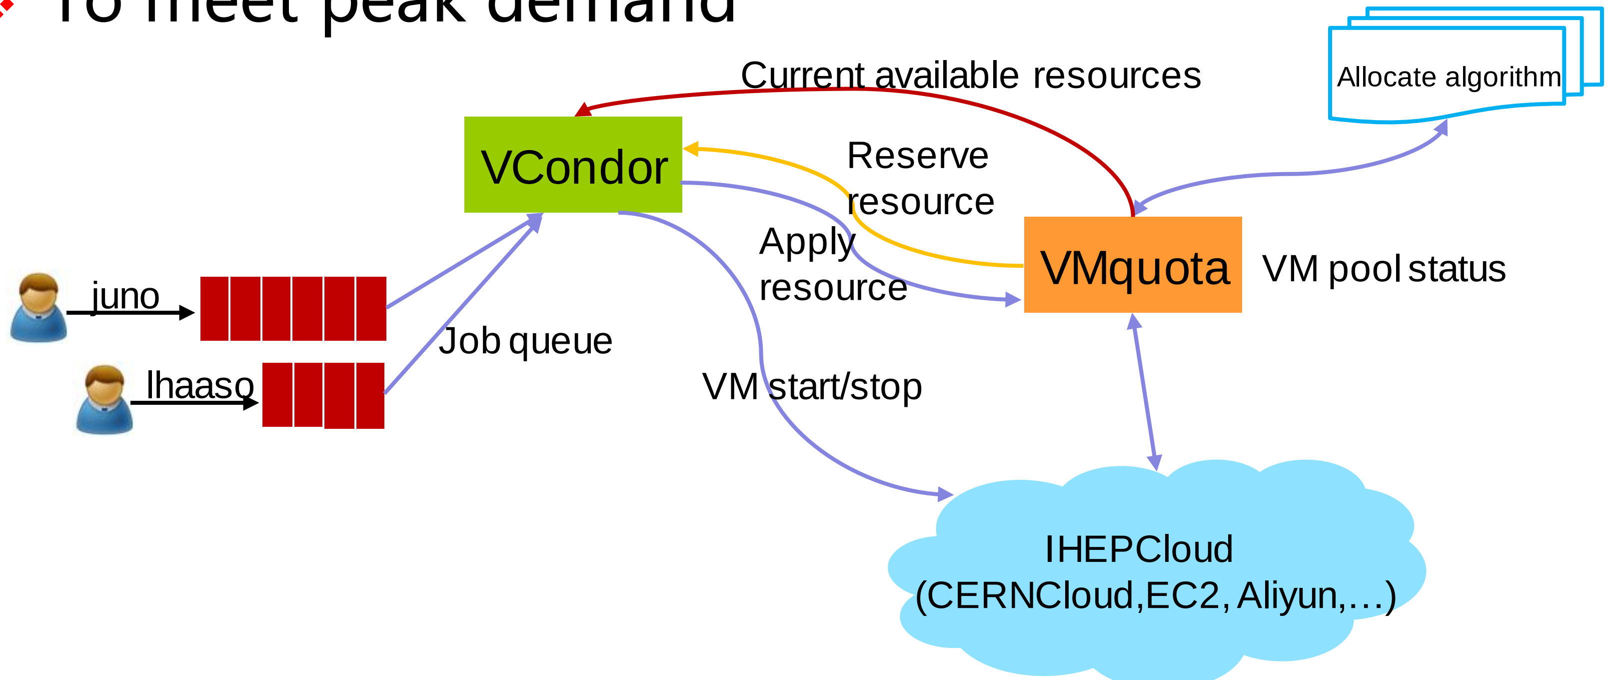

Performance evaluation and optimization (1)

- Evaluations have been done with JUNO M.C. simulaton
  - DetSim, ElecSim, ElecSim without output, PmtRec, EvtRec
- Results showed that
  - Most of Performance Loss (PL) is acceptable under 5%
  - I/O penalty is still a major issue for I/O intensive jobs (ElecSim with large outputs, PL 18.5%)

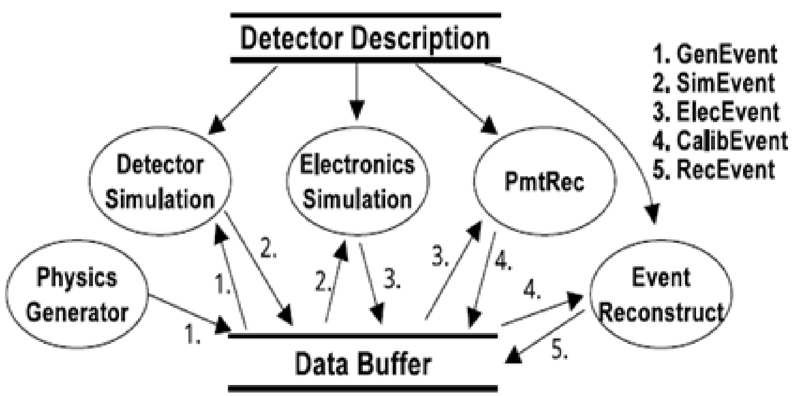

DP	PL
DetSim	5.5%
ElecSim	18.5%
ElecSim* (no out)	1.8%
PmtRec	0.6%
EvtRec	1.8%

Performance evaluation and optimization (2)

- Also found
  - Sim using Geant4 has bigger CPU loss than Rec
  - With increased VM size, PL drops dramatically with full simulation processes due to I/O conjunction in the job initialization
- Many factors contributes to the performance loss such as
  - Hardware, application, KVM parameters, OS etc.
- Automatic test and monitoring tools have been developed to keep watch on performance issues.

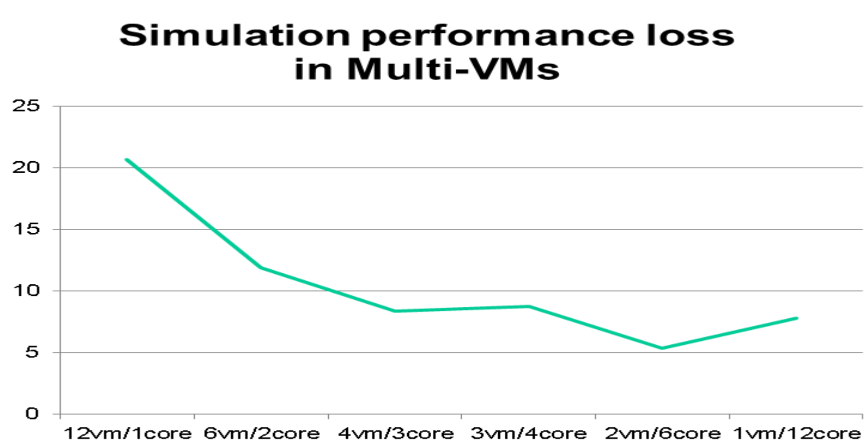

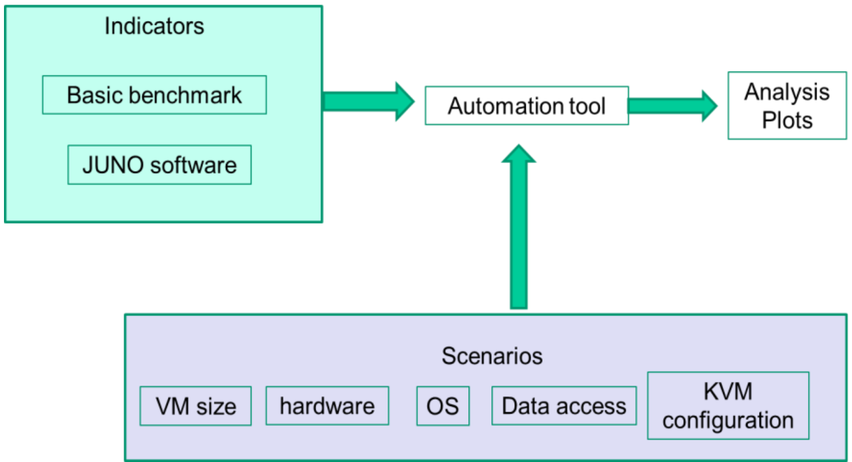

Affinity scheduling

- Affinity definition
  - The relation between VMs: Positive and negative correlation

- Affinity scheduling
  - Divide hosts into different groups
  - Define affinity template for different jobs
  - Apply affinity template to do scheduling of job VMs
  - Use FFD technique to place VMs in consideration of affinity and host resource capacity limitatation.

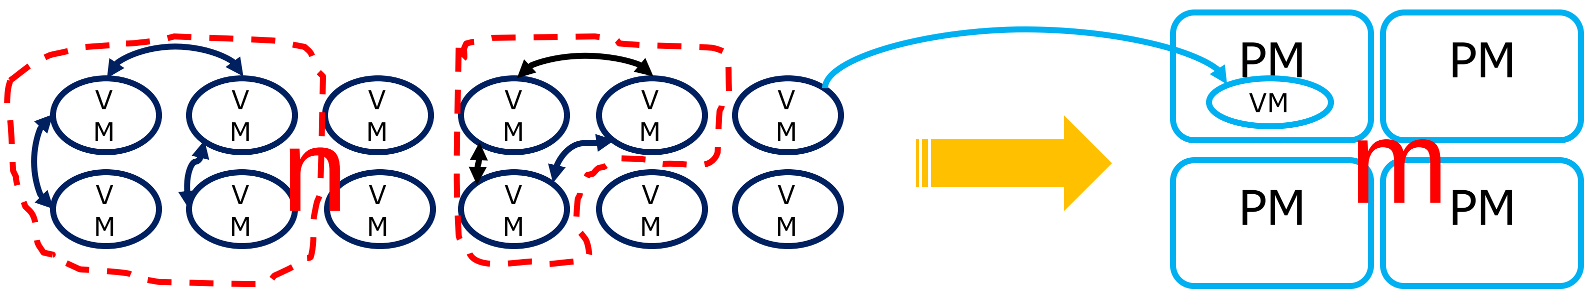

Volunteer Computing

CAS@home

SETI@home

LHC@home

ATLAS@home

BOINC

- ✓BOINC is a commonly used middleware
- ✓harness the **idle** CPU resources on personal computers(desktops, laptops, smart phones)
- ✓ computing tasks run at **a lower priority** (only when CPUs are idle)
- ✓suitable for CPU intensive, low latency computing

CAS@HOME

- First volunteer computing project in China
- Started in 2010, by the computer center, IHEP
- Provides free computing power to scientific computing in China
- Successfully hosted applications from material science and biology.
- Harvests around 2000 cores equivalent computing power

Other harvests from CAS@home

- Have built **a platform** which provides a very considerable amount of computing resources for scientific computing
- Have established **a community** of Chinese volunteers
- Engaged in **the research and development** of BOINC middleware
  - The distribution of heavy virtual machine images
  - Integration of BOINC with the DIRAC interware
  - BOINC scheduling policies
- IHEP is the pioneer in applying volunteer computing to the computing of High Energy Physics experiments.

IHEP Contributions to ATLAS@home

- ATLAS@home is one of the earliest volunteer computing projects running tasks for HEP experiments.
- IHEP involved in the prototype of ATLAS@home and mainly contributed to solve three problems(host virtualization, platform integration, security) which were commonly faced by applying volunteer computing in HEP computing.
- The first ATLAS@home was deployed and maintained at IHEP in 2013, and moved to CERN in 2014.
- IHEP continuously contributes manpower (0.3FTE every year) in maintaining and upgrading ATLAS@home
  - Use virtualization to run ATLAS jobs on heterogeneous computers
  - Event display on BOINC GUI
  - Multiple core jobs
  - Native running (run ALTAS with containers or directly on SLC6)
  - Use ATLAS@home to exploit extra CPU from Grid sites and any clusters
  - BOINC scheduling optimization

Architecture of ATLAS@home

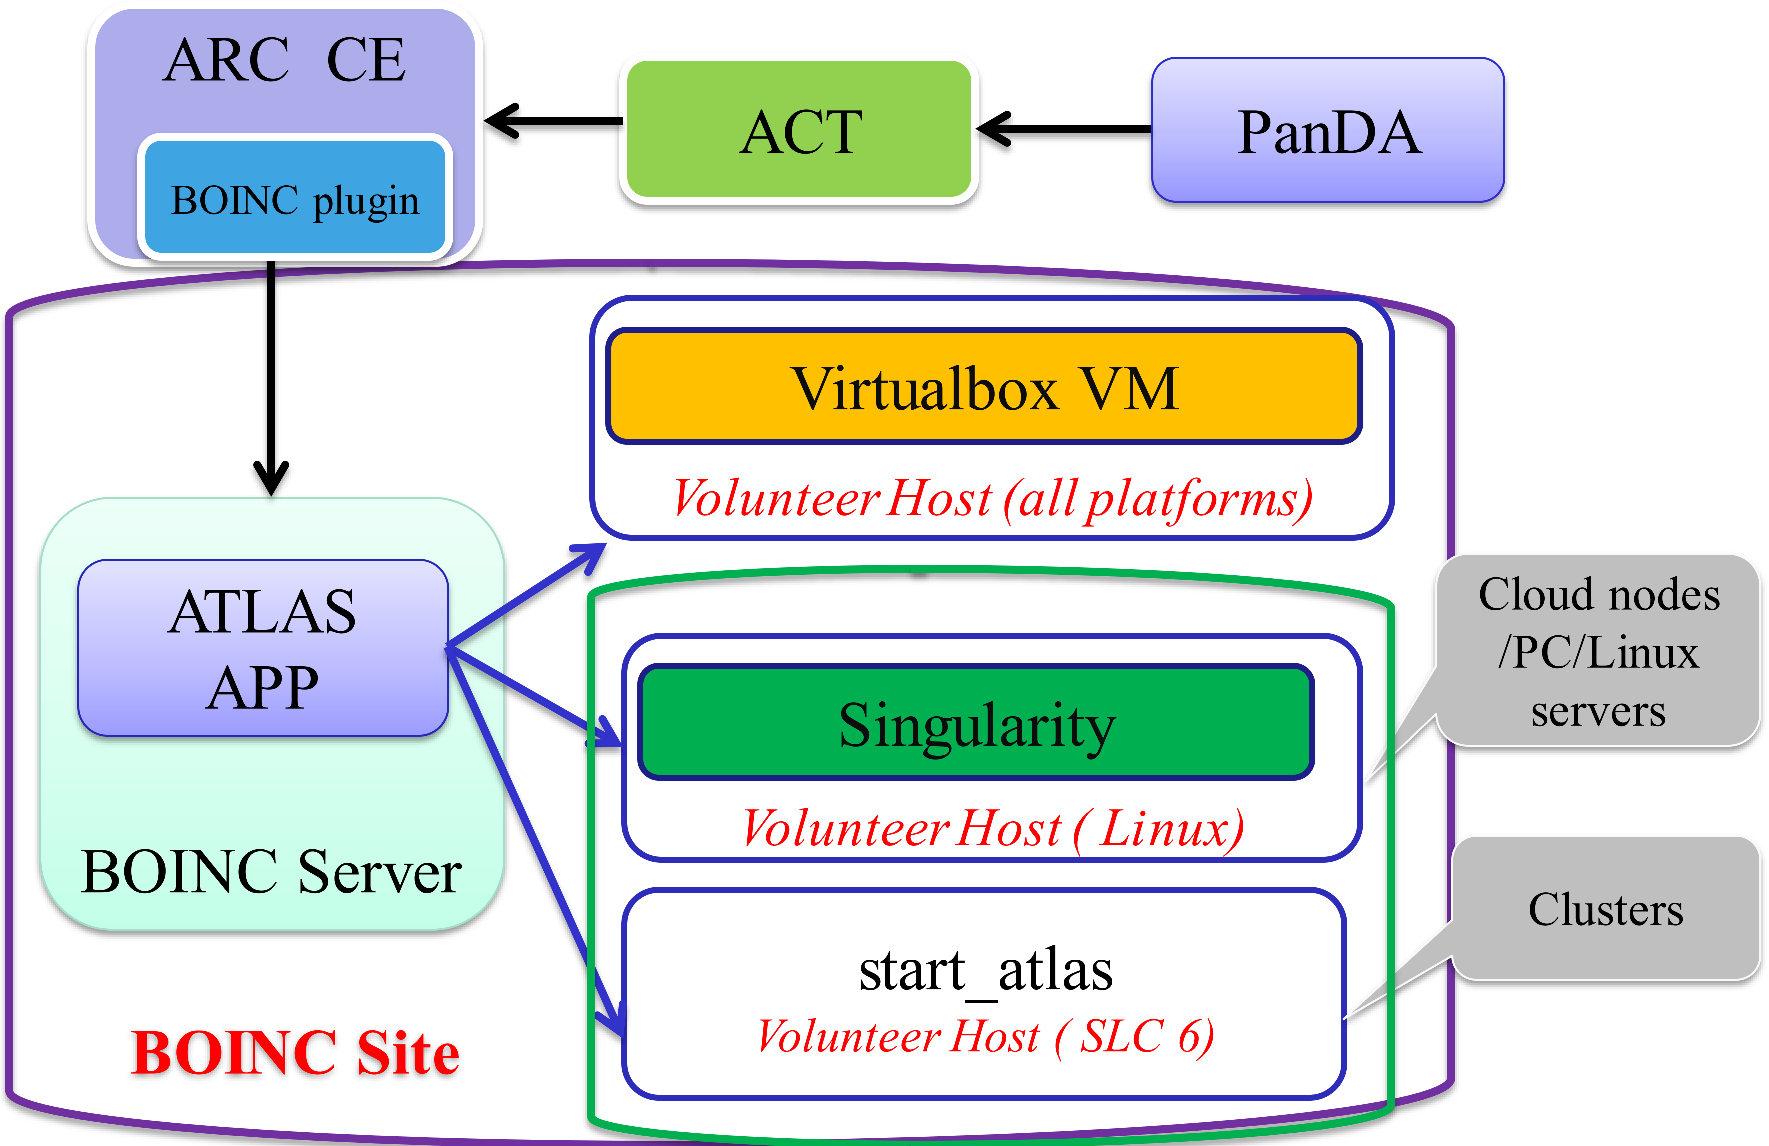

The Whole ATLAS Computing

ATLAS@home(BOINC) provides a very ***visible*** amount of CPU to the ATLAS computing, and its current computing power is equivalent to a grid site with **3000** CPU Cores (core power = 10 HS06)

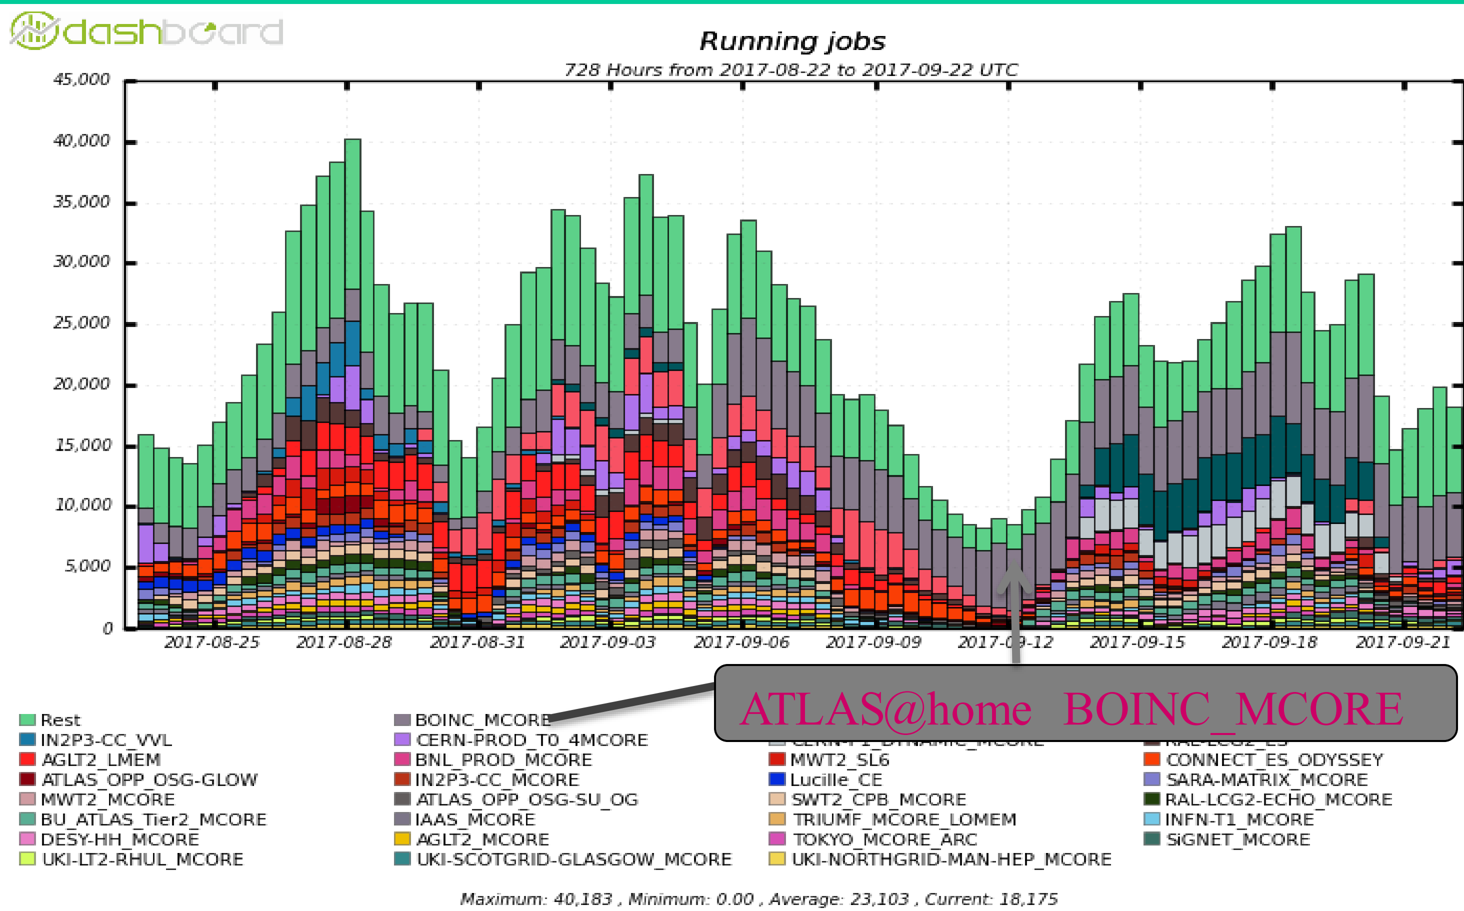

Use ATLAS@home to exploit extra CPU from busy grid sites

- What?
  - Running BOINC native jobs in busy Grid sites, to exploit extra CPU.
- Why?
  - Grid sites or any clusters aren't fully used with traditional ways to use (1 job per job slot)
    - For ATLAS Tier sites, avg. 85% wall utilization, 70% CPU utilization

Site CPU utilization in 100 days

Site	cores	HS06	avg wall utilization	avg cpu utilization	HS06 utilization	avg job eff
BEIJING	634	11780	0.68	0.55	0.63	0.81
TOKYO	6144	111267	0.85	0.72	0.71	0.85
SiGNET	5288	51822	0.88	0.68	0.77	0.77
MWT2	16250	162500	0.83	0.70	0.73	0.84
AGLT2	10224	104414	0.72	0.61	0.63	0.84

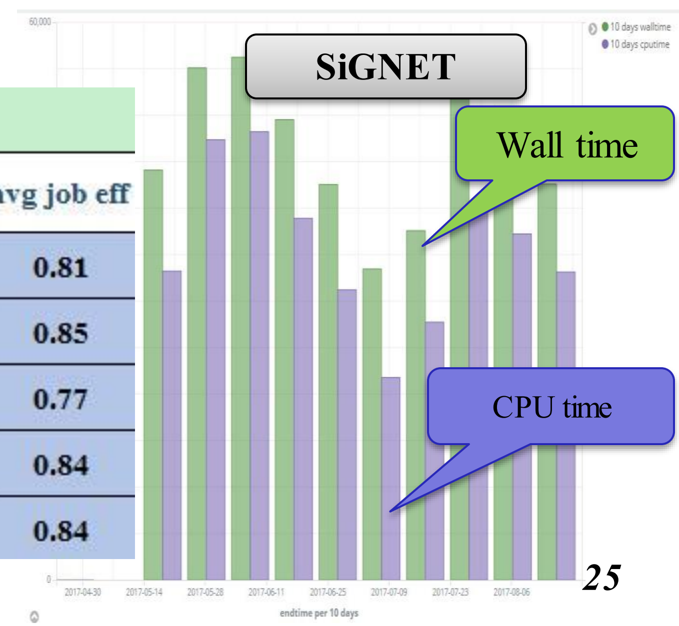

Look into one work node

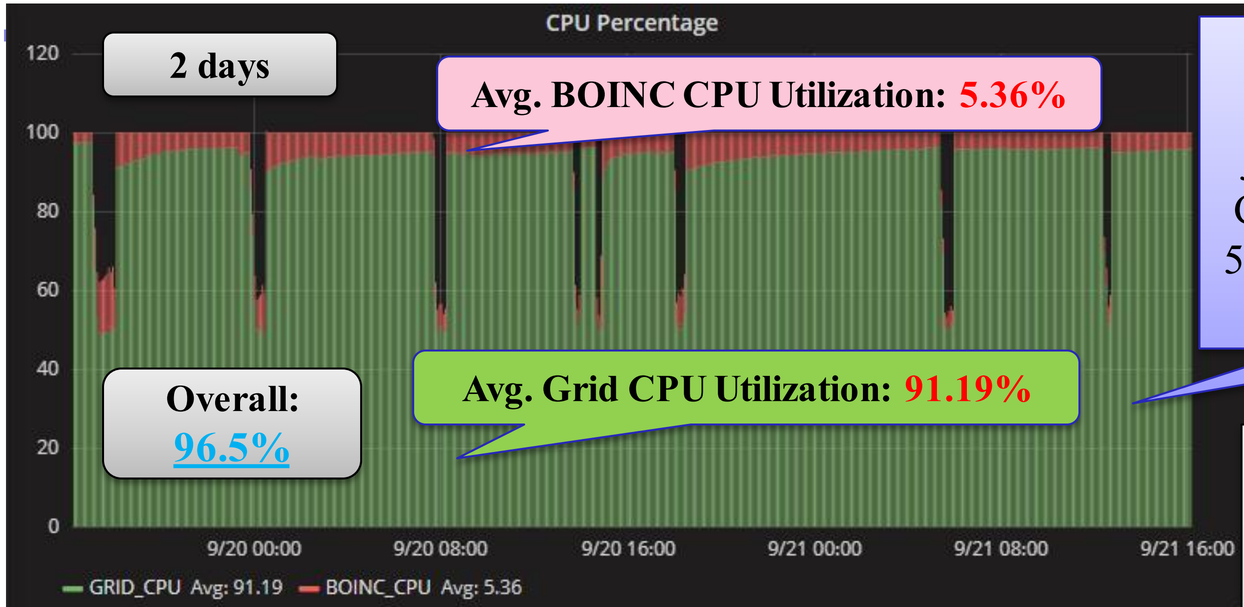

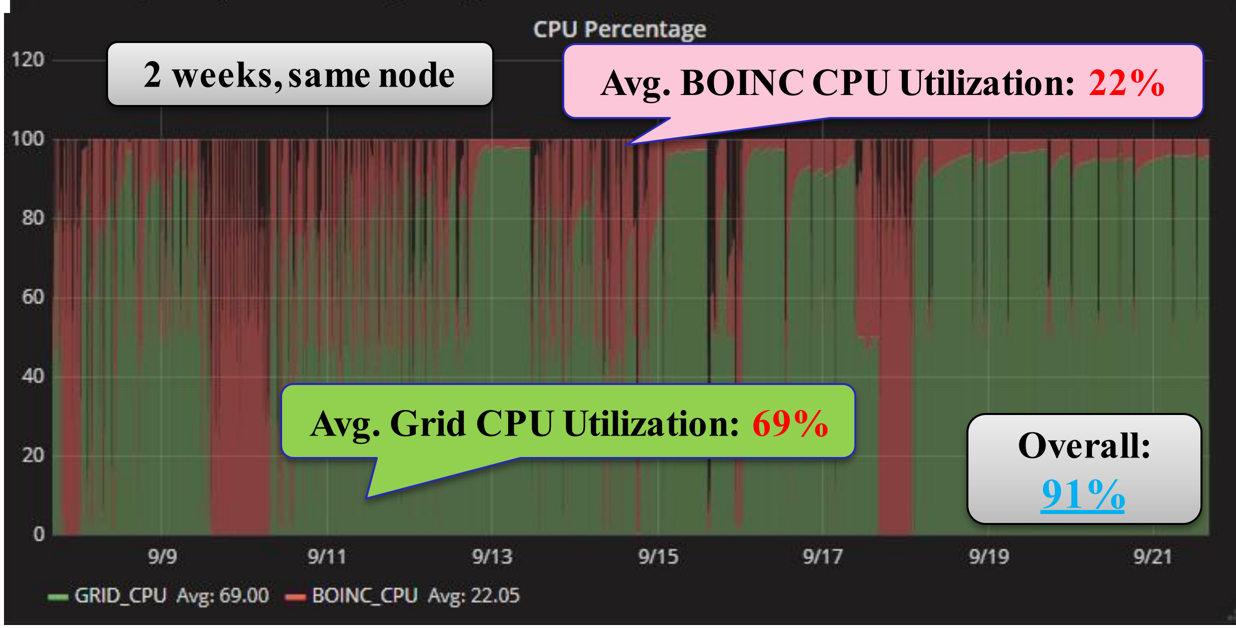

For 2 days, it is very full (Grid jobs use 91% system CPU) and CPU efficient, BOINC gets only 5% , overall 96% CPU utilization

For 2 weeks, it is NOT always full, BOINC gets 22% of system CPU

This node runs multi core production jobs (job cpu_eff>0.92), more cpu efficient than other type of jobs

Running ATLAS@home at BEIJING_LCG2

- How?
  - BOINC acts as a second scheduler, pulling in more jobs (with lowest priority) to increase the CPU utilization of every single work node.
  - Being developed by IHEP and tested at BEIJING_LCG2 site.
- CPU utilization can reach **90%** on all work nodes, and **extra** of **23%** CPU can be exploited from BEIJING_LCG2 Tier 2 site while the site is fully loaded (wall time utilization > 86%)

BEIJING_LCG2 daily utilization in 3 weeks										
	#cores	HS06	Walltime (days)	CPUtime (days)	HS06 (days)	Job CPU eff (%)	Wall Util (%)	CPU Util (%)	HS06 Util (%)	Extra Cores
Grid	420	7560	359	283	5591	78.83	85.48	67.38	73.96	
BOINC	420	7560	367	97	1730	26.43	87.38	23.10	22.88	144.0
Total	420	7560	726	380	7321		172.86	90.48	96.84	

Note: site has 420 dedicated cores for ATLAS, numbers are the average in a consective 3 weeks (starting 2017-08-25), BOINNC processes **0.03 M events/day**

Summary

- Grid and cloud computing technologies were adopted to support various types of HEP experiments in China.
  - Dirac-based grid to integrate resources within an experiment
  - Cloud to promote sharing of resources among different experiments
- Running ATLAS BOINC jobs at BEIJING_LCG2 is being studied to improve CPU utilization.

Thank You！

谢谢