

# Belle II distributed computing



Pavel Krokovny  
Novosibirsk State University and  
Budker Institute of Nuclear Physics




## Outline:

- Belle II experiment
- Computing model
- Resource estimation
- Production system
- Computing in BINP / NSU



# The Super B factory project

	KEKB		SuperKEKB	
Luminosity:	$2.1 \times 10^{34}$	→	$8 \times 10^{35}$	(x 40)
Total Data:	$1 \text{ ab}^{-1}$	→	$50 \text{ ab}^{-1}$	(x 50)
Detector:	Belle	→	Belle II	




## Belle II Collaboration:

- 26 countries,
- 97 institutes,
- ~600 scientists



# Luminosity prospect



Experiment	Event size	Rate @ Storage	Rate @ Storage
	[kB]	[event/sec]	[MB/sec]
Belle II	300	6,000	1,800
ALICE (Pb-Pb)	50,000	100	4,000
ALICE (p-p)	2,000	100	200
ATLAS	1,500	600	700
CMS	1,500	150	225 (<~1000)
LHCb	55	4,500	250

(@ max. luminosity)

(LHC experiments: as seen in 2011/2012 runs)



# Software system

- A “framework” system with dynamic module loading, parallel processing, Python steering, and ROOT I/O
- Full detector simulation with Geant4
- C++ 11 and gcc 4.7
- Supporting major Linux distributions: SL, Fedora, Ubuntu, etc
- Code management: Subversion
- Formatting tool: astyle
- Building: scons and buildbot system
- Documentation: Doxygen, Twiki
- Issue tracking: Redmine



# Computing model

The BELLE II Computing model has to accomplish, in a distributed environment, the following main tasks:

- RAW data processing
- Monte Carlo Production
- Physics analysis
- Data Storage and Data Archiving





Resource Estimation:

- Event size for RAW (mDST) data is 300 (40) Kb
- RAW data (MC) processing: 45 (90) HepSPEC\*s / event



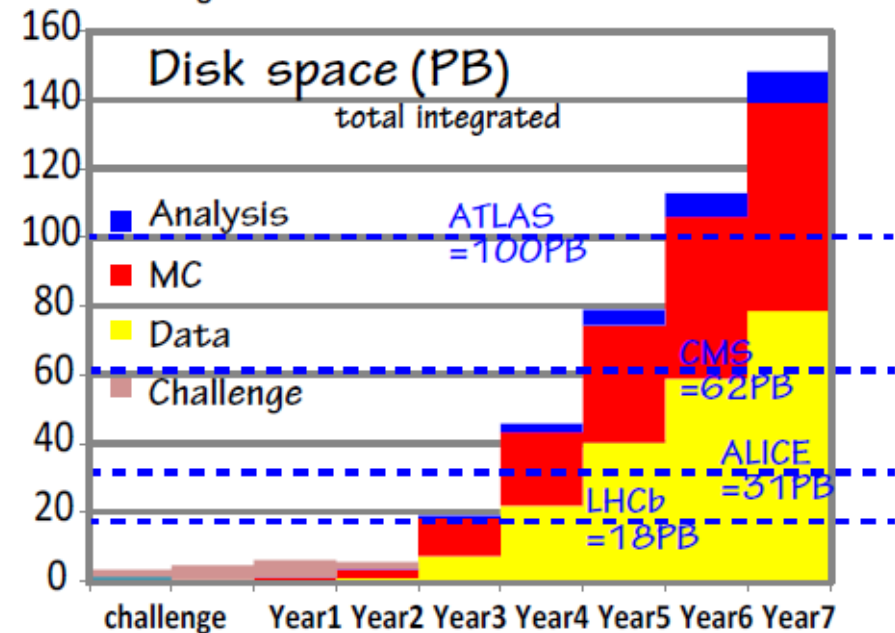
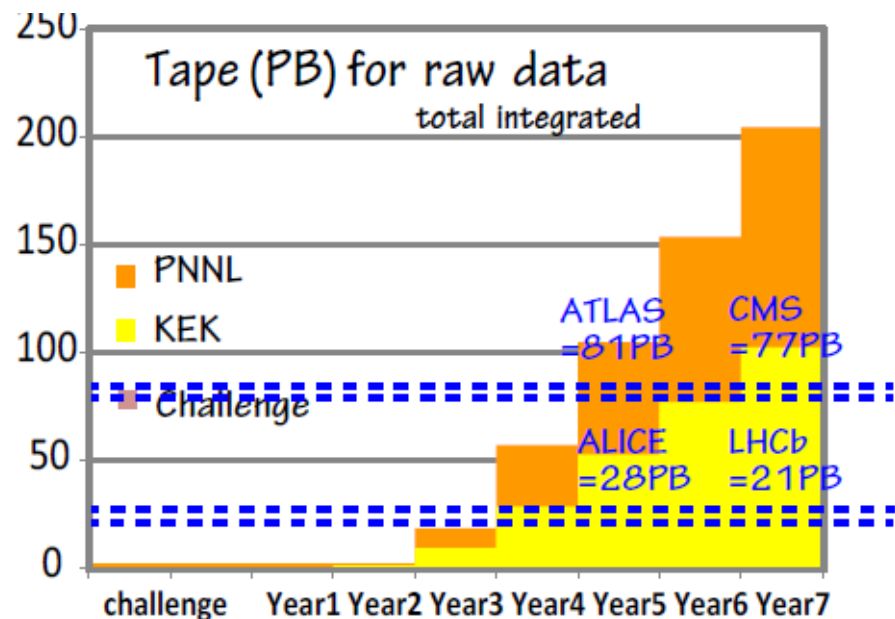
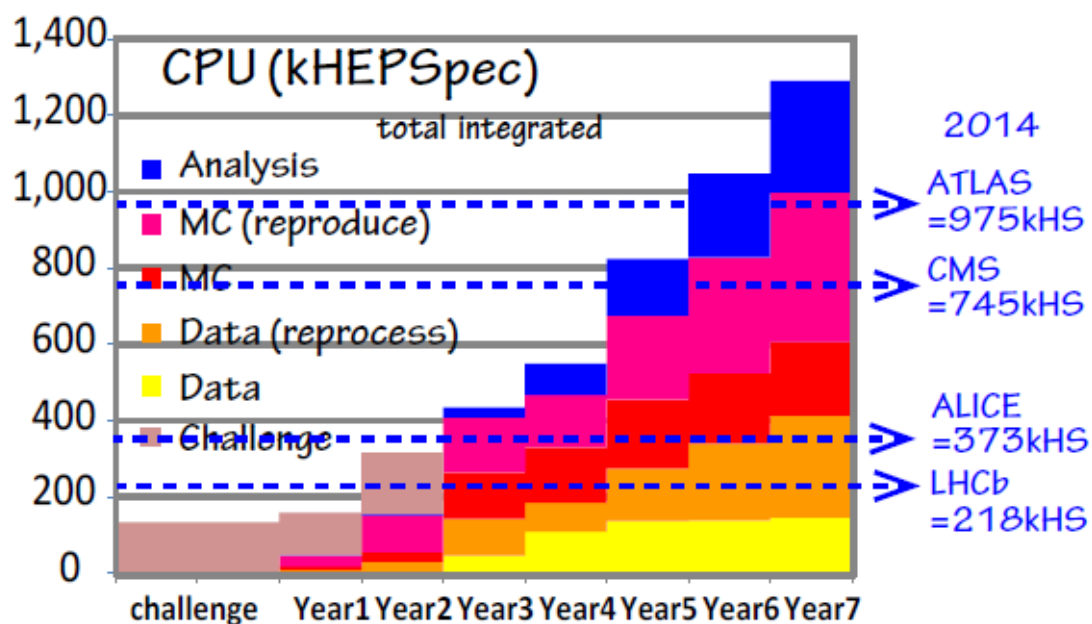



# Required hardware resources

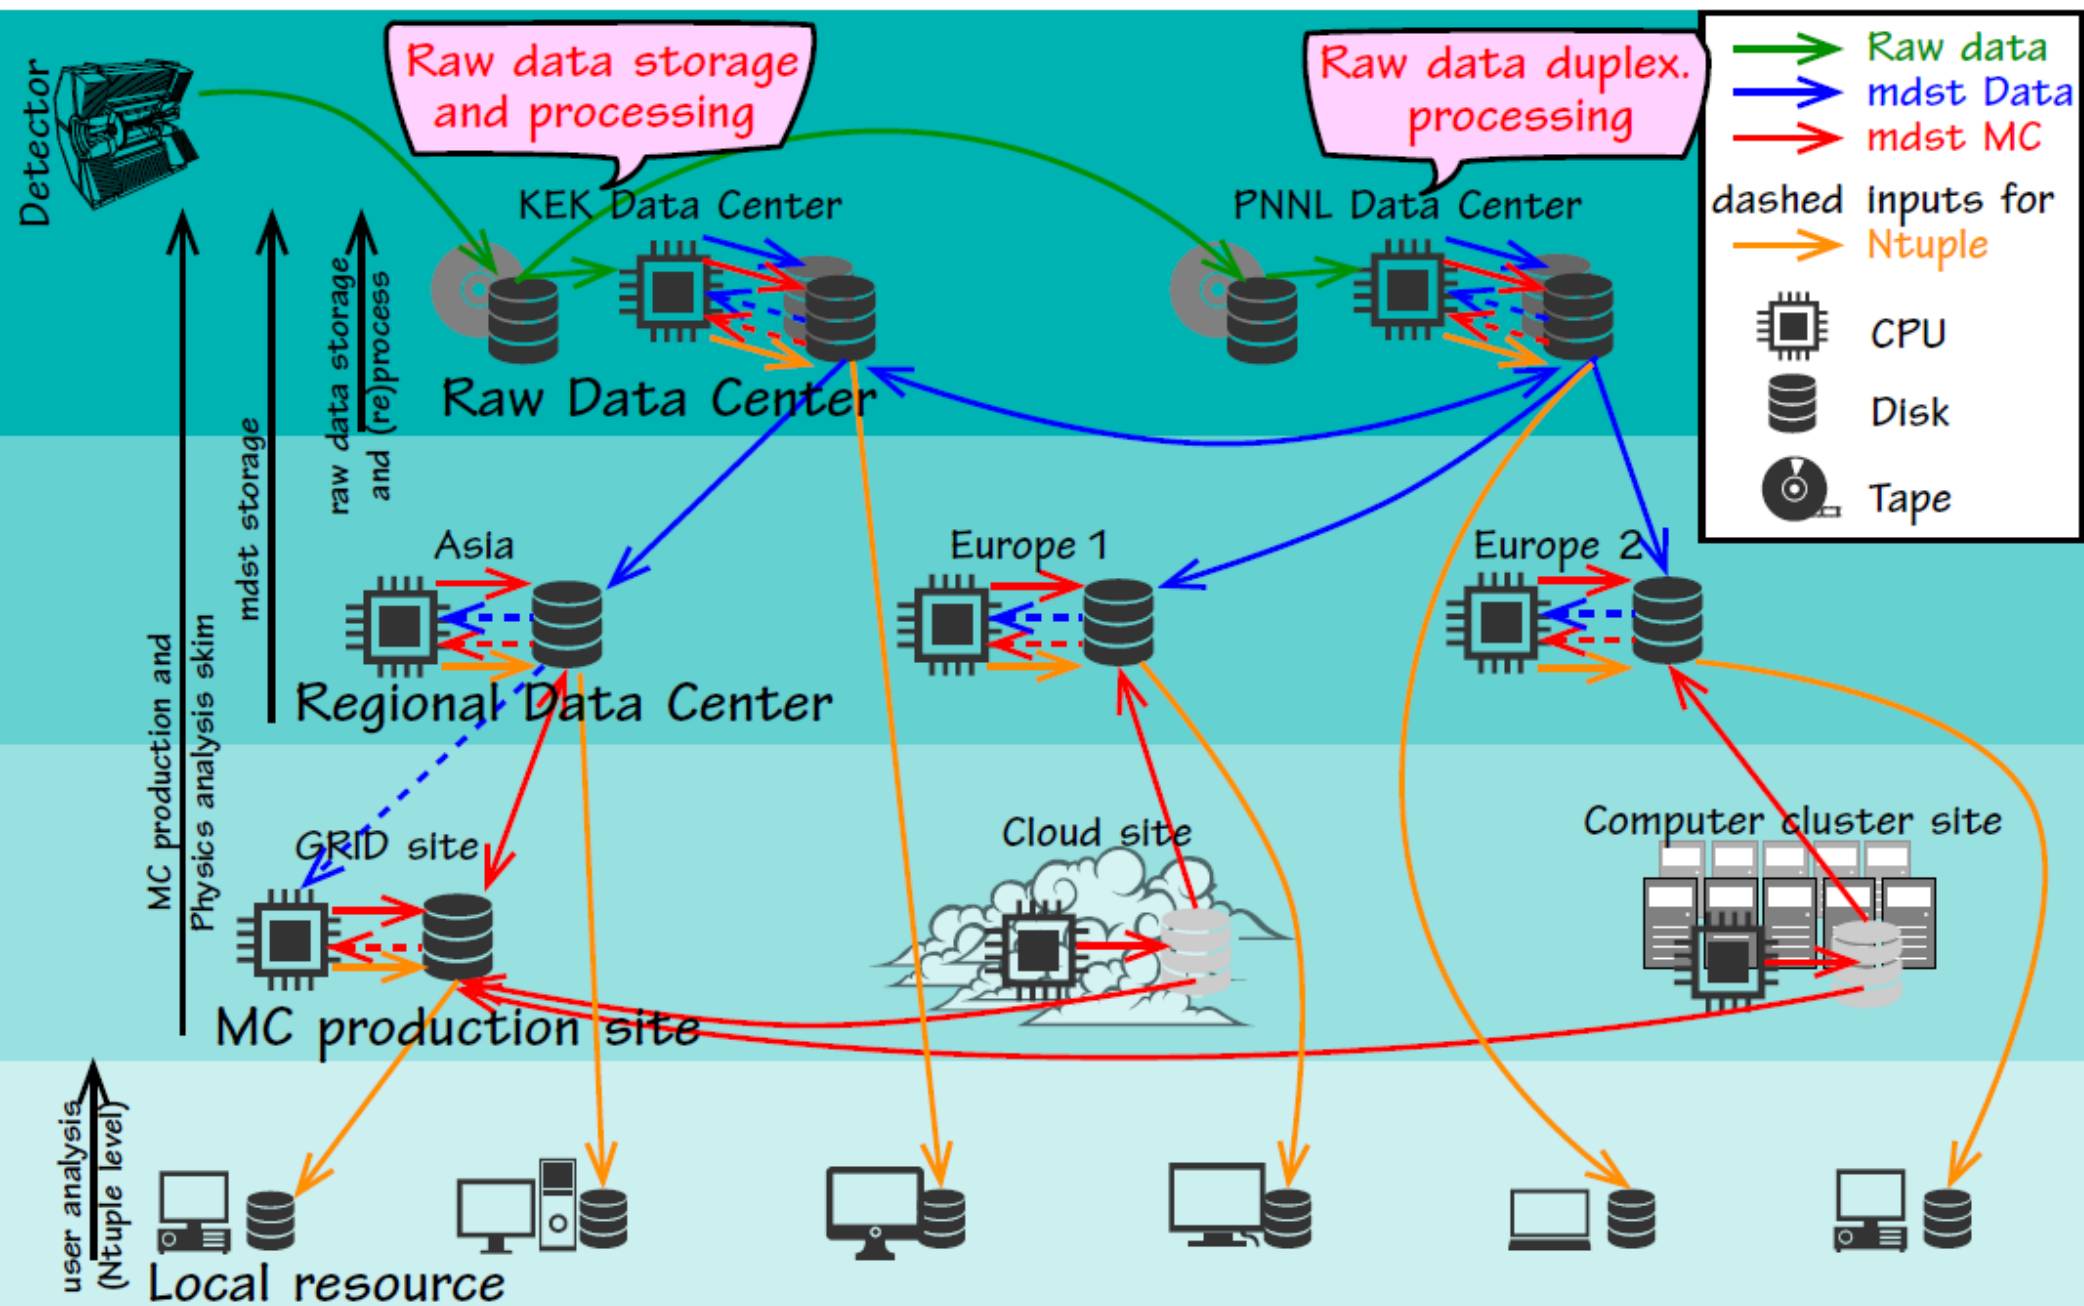




# Comparison of hardware resources

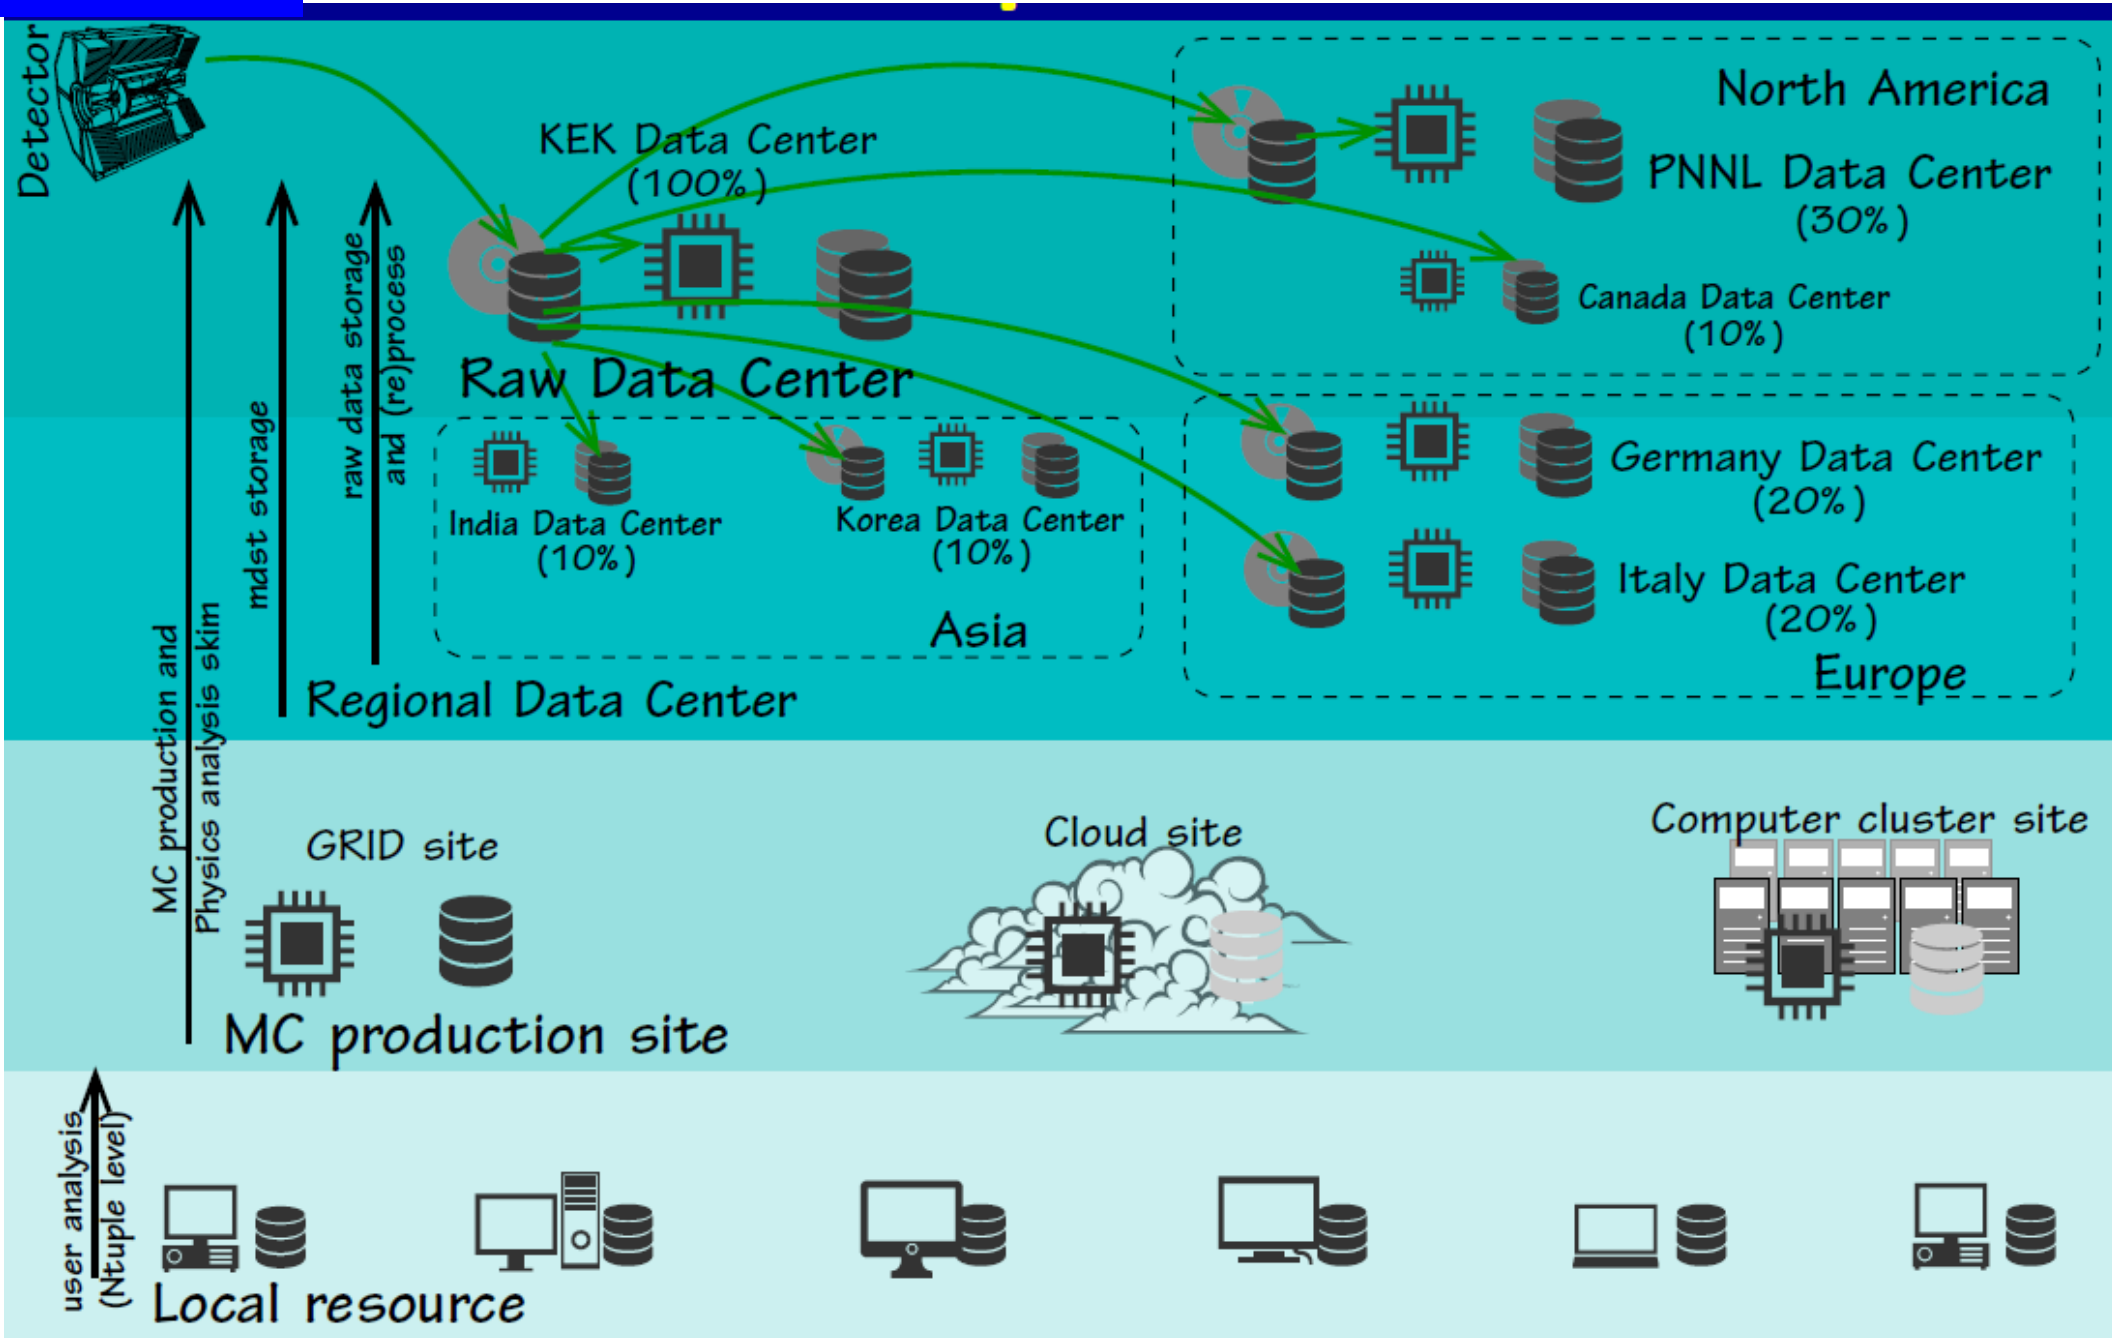


# Initial computing model (first 3 years)





# Raw data storage (2nd stage)






# Current status

Normalized CPU usage by Site

76 Weeks from Week 52 of 2012 to Week 23 of 2014



LCG.DESY.de	18.6%	LCG.CNAF.it	3.6%	OSG.FNAL.us	1.8%
LCG.KIT.de	11.9%	DIRAC.PNNL.us	2.9%	LCG.ULAKBIM.tr	1.2%
LCG.KEK2.jp	11.6%	LCG.SiGNET.si	2.8%	OSG.PNNL.us	1.0%
LCG.KMI.jp	5.3%	LCG.Pisa.it	2.8%	DIRAC.KrakowCloud.pl	0.9%
LCG.UA-ISMA.ua	4.8%	DIRAC.UVic.ca	2.4%	LCG.McGill.ca	0.6%
LCG.CYFRONET.pl	4.4%	LCG.KISTI.kr	2.3%	LCG.Torino.it	0.6%
LCG.CESNET.cz	4.3%	OSG.Nebraska.us	2.3%	LCG.MPPMU.de	0.5%
DIRAC.BINP.ru	4.0%	LCG.Frascati.it	2.2%	LCG.Legnaro.it	0.5%
LCG.Napoli.it	3.9%	LCG.Melbourne.au	2.0%	plus 12 more	

Generated on 2014-05-17 13:07:26 UTC

15 countries/regions

27 sites (+ 2 non-Belle II sites)

HEPHY (Vienna) and MPPMU (Munich)  
joined recently  
GRID, Cloud, local cluster  
is available

First official release of MC samples

BB generic decay/continuum  
tau pair

(corresponding to  $100\text{fb}^{-1}$  w/ and w/o BG)

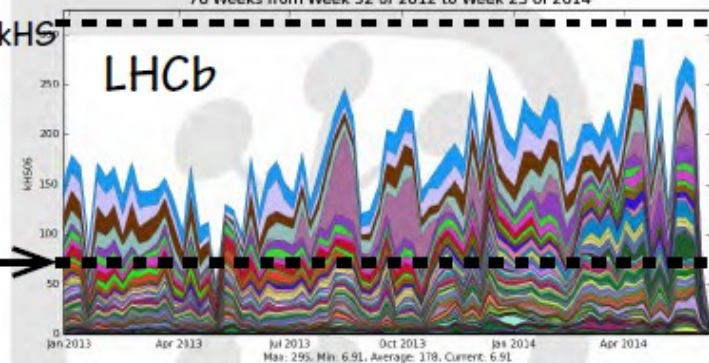
Trans-pacific / trans-atlantic  
network data transfer challenge

300 kHS  
120 sites

Belle II now  
70 kHS

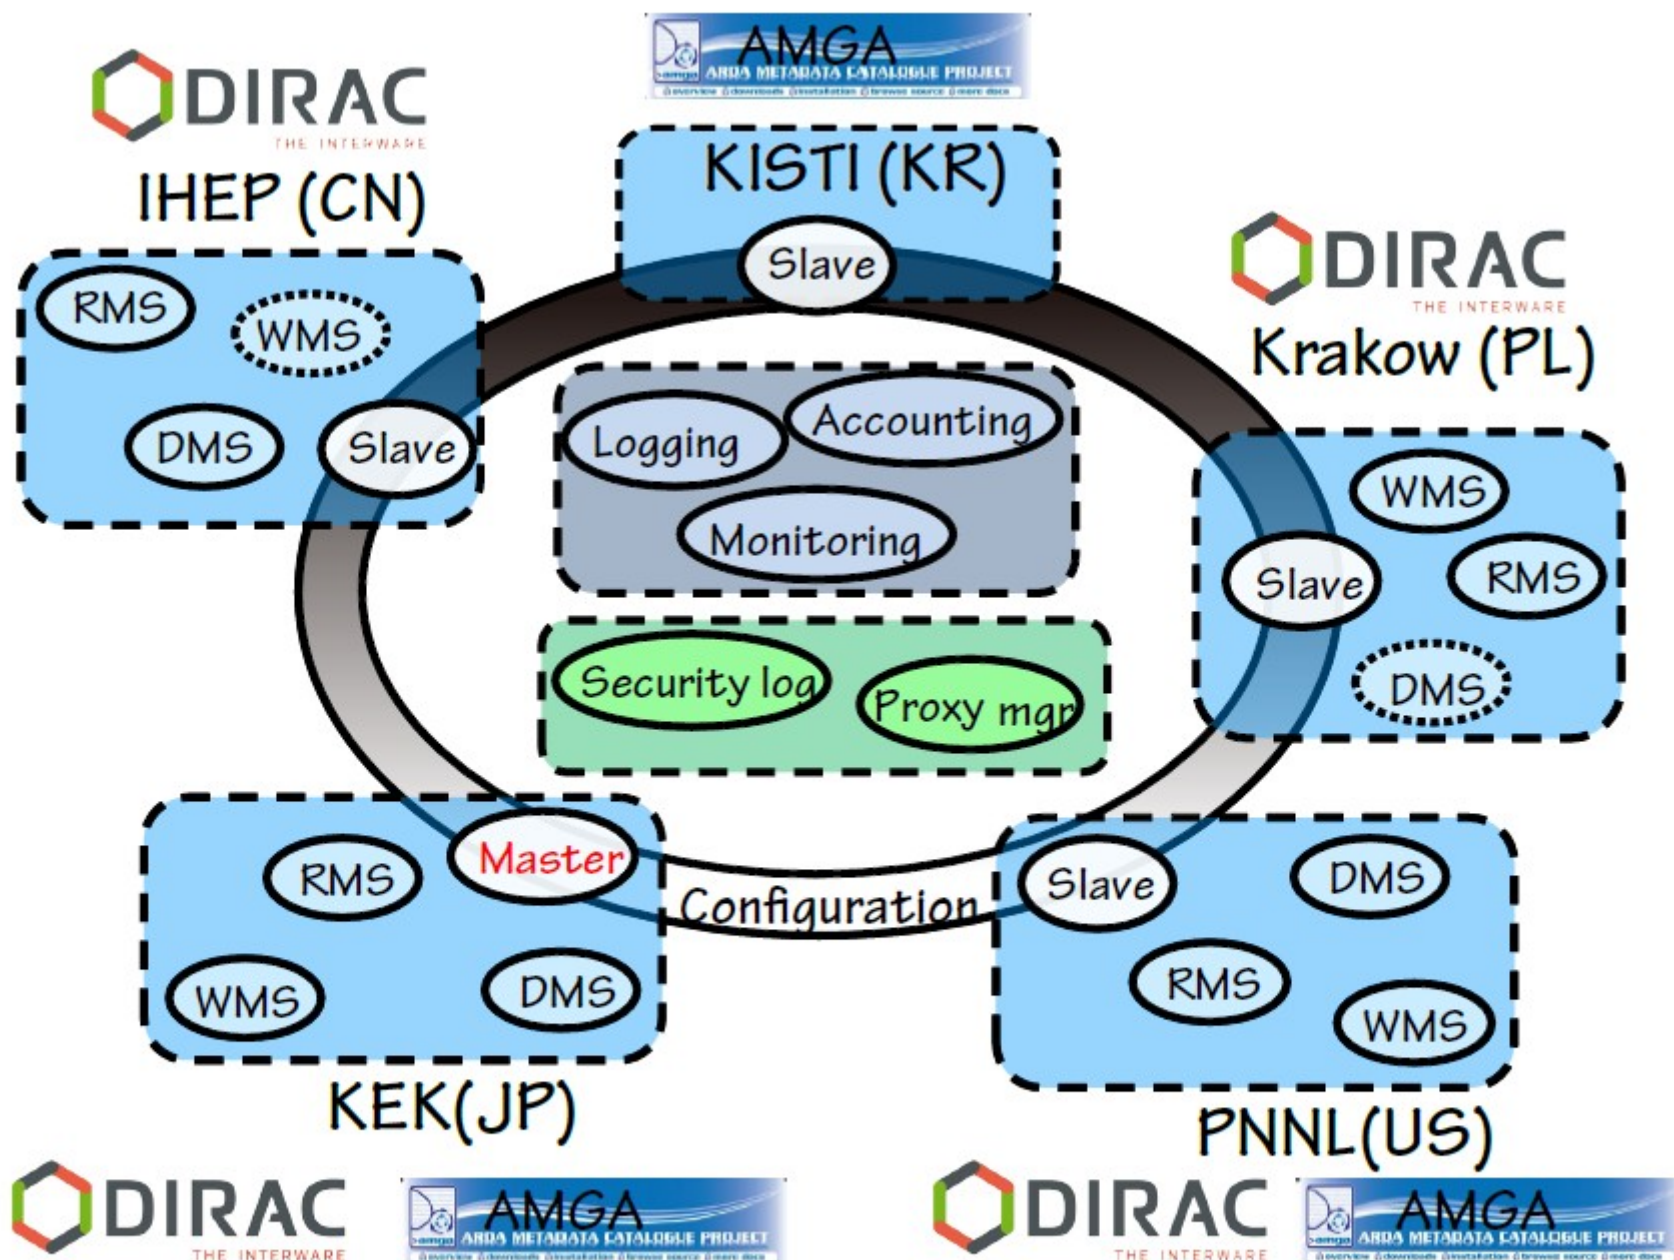
Normalized CPU usage by Site

76 Weeks from Week 52 of 2012 to Week 23 of 2014

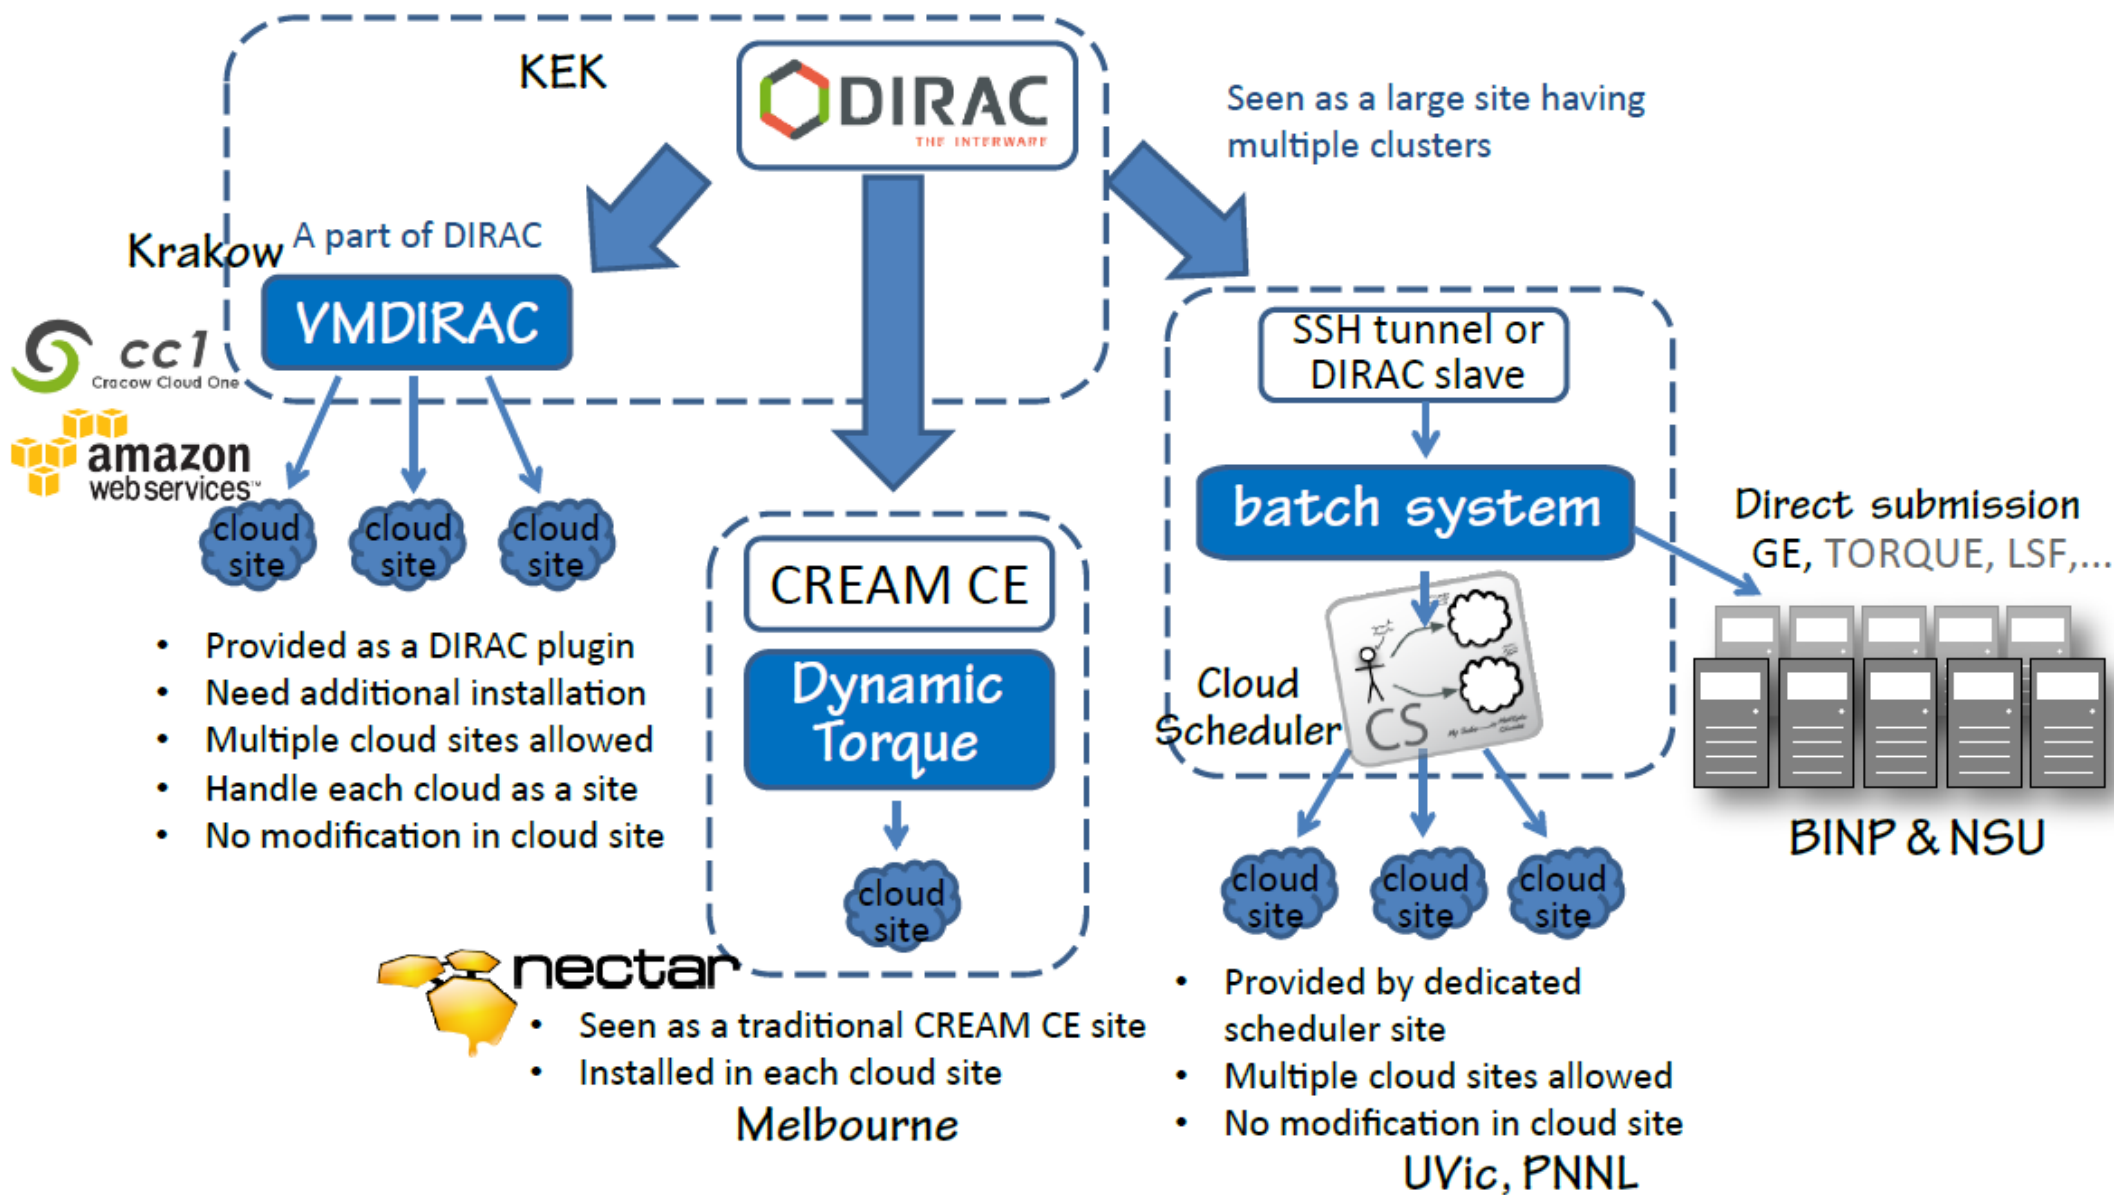




# Production system (core services)



# Production system (cloud usage)





# Computing in BINP & NSU



## Experiments

- BINP: CMD-3, SND, KEDR
- International collaborations: Belle II, ATLAS, LHCb

## Computing resources

- BINP: 512 CPU cores, 6 Tflops
- Novosibirsk State University: 2432 cores, 29 Tflops
- Supercomputer Center: 30 + 85 Tflops (CPU + GPU)

## We want:

- keep the specific computing environment and user's experience
- be like a normal SC user

## The solution is:

- run HEP tasks inside virtual machines,
- run VMs inside supercomputer's batch system jobs.



# Key features of the virtualized infrastructure

- Virtualization by KVM
  - included in modern Linux distributions,
  - quite stable,
  - does not require modified Linux kernel.
- VM disk images are located on SC's file system and accessible via InfiniBand.
- Local snapshots of VM images are used on physical nodes thus leaving master images unchanged.
- Input/output data are located at BINP and accessed by VMs via NFS.
- VMs are just regular batch tasks at a supercomputer.
- VMs are started automatically on user's demands.

# Batch System Integration Mechanisms

NSU

NUSC

PBS Pro Batch System

*STAGE 1  
Job submission  
and automated  
VM group  
deployment  
sequence*

NSC/SCN

Orchestration  
Services

**SGE Batch  
System**

**Belle II Detector  
User Group**

Orchestration  
Services

**SGE Batch  
System**

**SND Detector  
User Group**

Orchestration  
Services

**SGE Batch  
System**

**ATLAS Data  
Analysis Group**

BINP/GCF

BINP




# Batch System Integration Mechanisms

NSU

NUSC

PBS Pro Batch System

**STAGE 1**  
*Job submission  
and automated  
VM group  
deployment  
sequence*



*Group of VMs  
of particular  
type created on  
demand*

*Set of computing  
nodes with KVM,  
IPoIB & HT support  
enabled on  
demand*

NSC/SCN

Orchestration  
Services

**SGE Batch  
System**

**Belle II Detector  
User Group**

Orchestration  
Services

**SGE Batch  
System**

**SND Detector  
User Group**

Orchestration  
Services

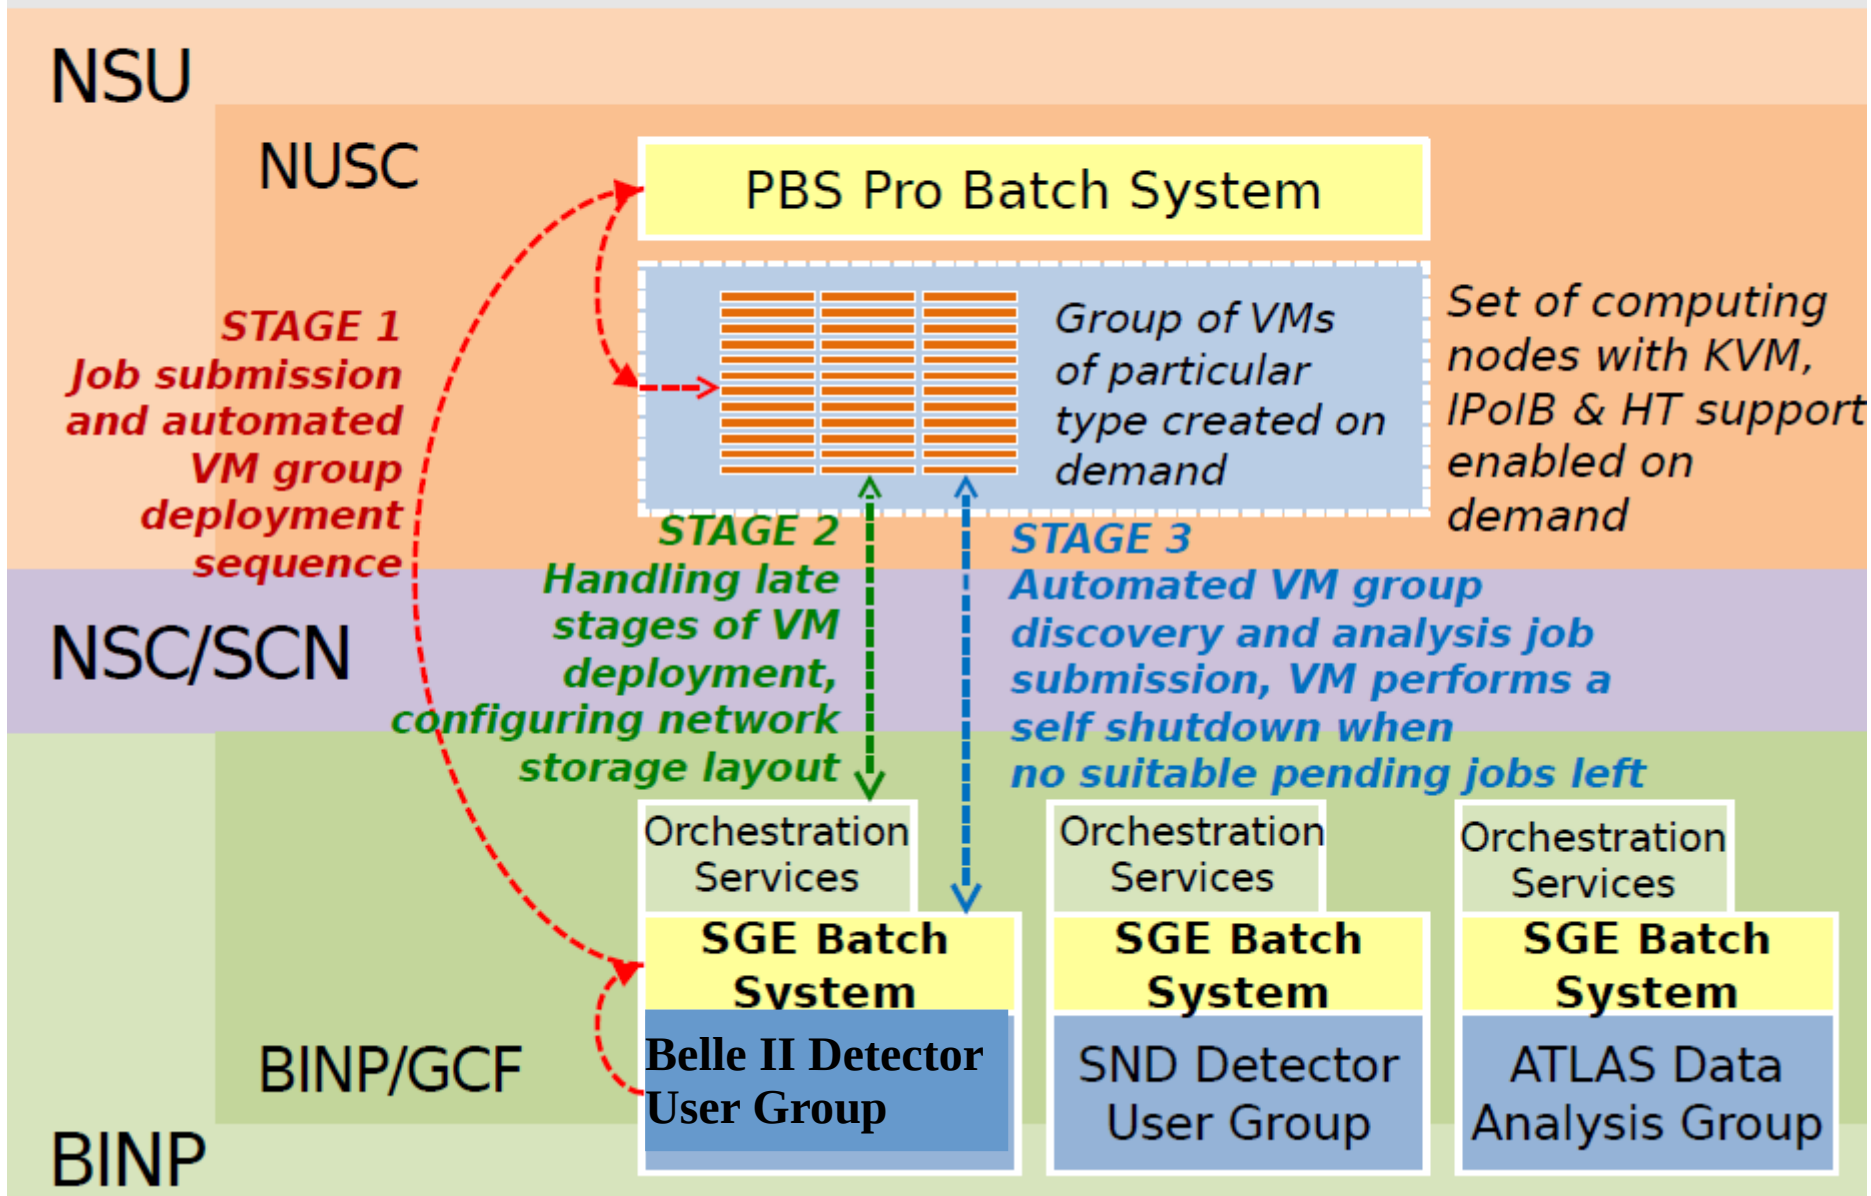
**SGE Batch  
System**

**ATLAS Data  
Analysis Group**

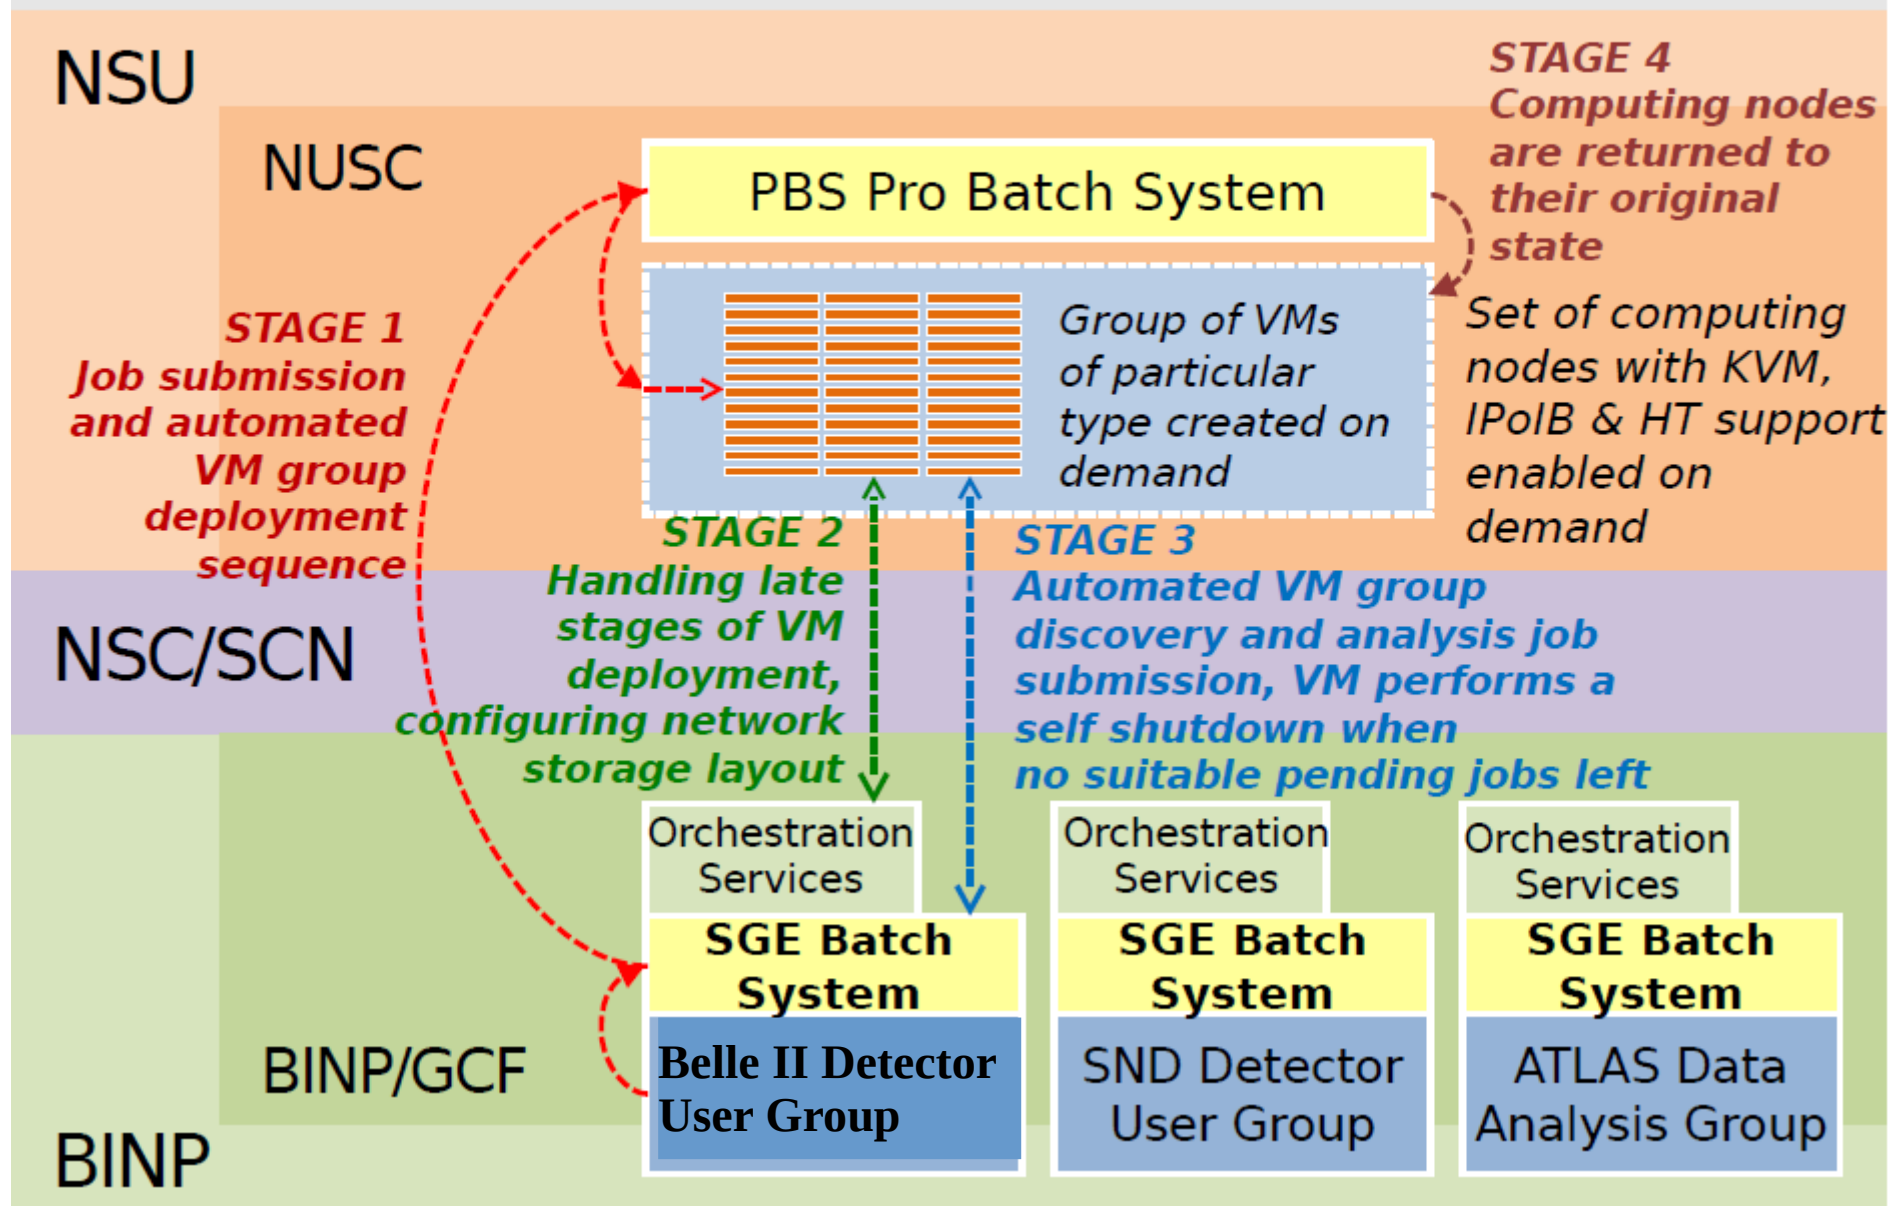
BINP/GCF

BINP

# Batch System Integration Mechanisms



# Batch System Integration Mechanisms






# Summary

- Belle II experiment has successful experience in distributed computing using GRID, Cloud and local clusters.
- Current CPU power is about 70 kHS (25% of LHCb) and used for MC production.
- BINP successfully uses Virtual Machines technique to run up to 1K jobs simultaneously for 6 experiments on 3 supercomputer sites.



# Backup

# SuperKEKB accelerator


- Redesign the lattices of both rings to reduce the emittance
- Replace short dipoles with longer ones (LER)



- Electron cloud mitigation for LER beam pipes



- Belle II
- New Interaction region
- New superconducting final focusing quads near the IP



- Add / modify RF systems for higher beam current



- New positron source

Intensity up



- Low emittance positrons to inject

- Damping ring



- Low emittance gun

- Low emittance electrons to inject

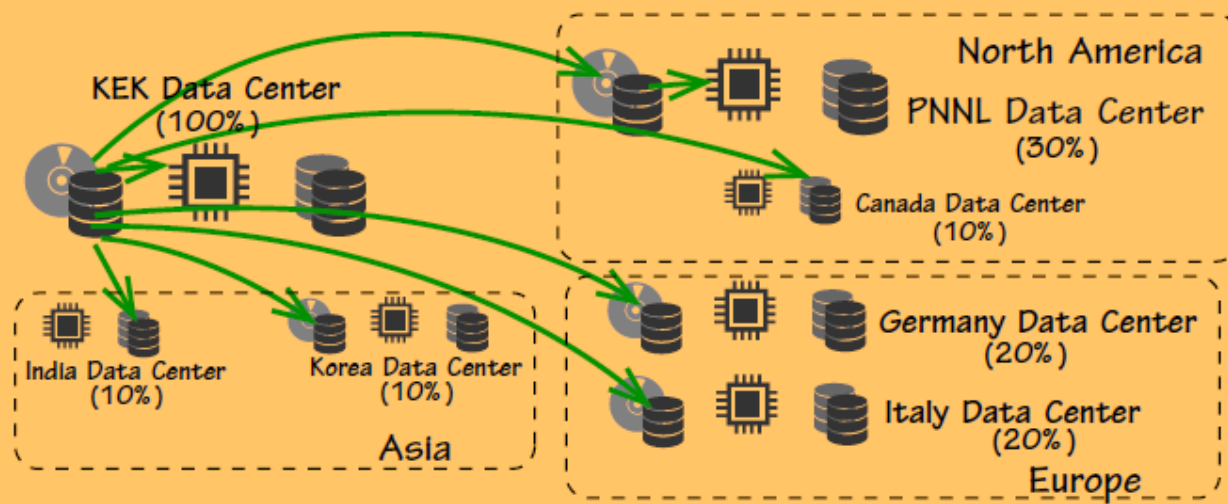


# Raw data distribution



until Year 3



*Scenario 1*  
(copy from KEK)



*Scenario 2*  
(2step copy, KEK → PNNL → Europe)







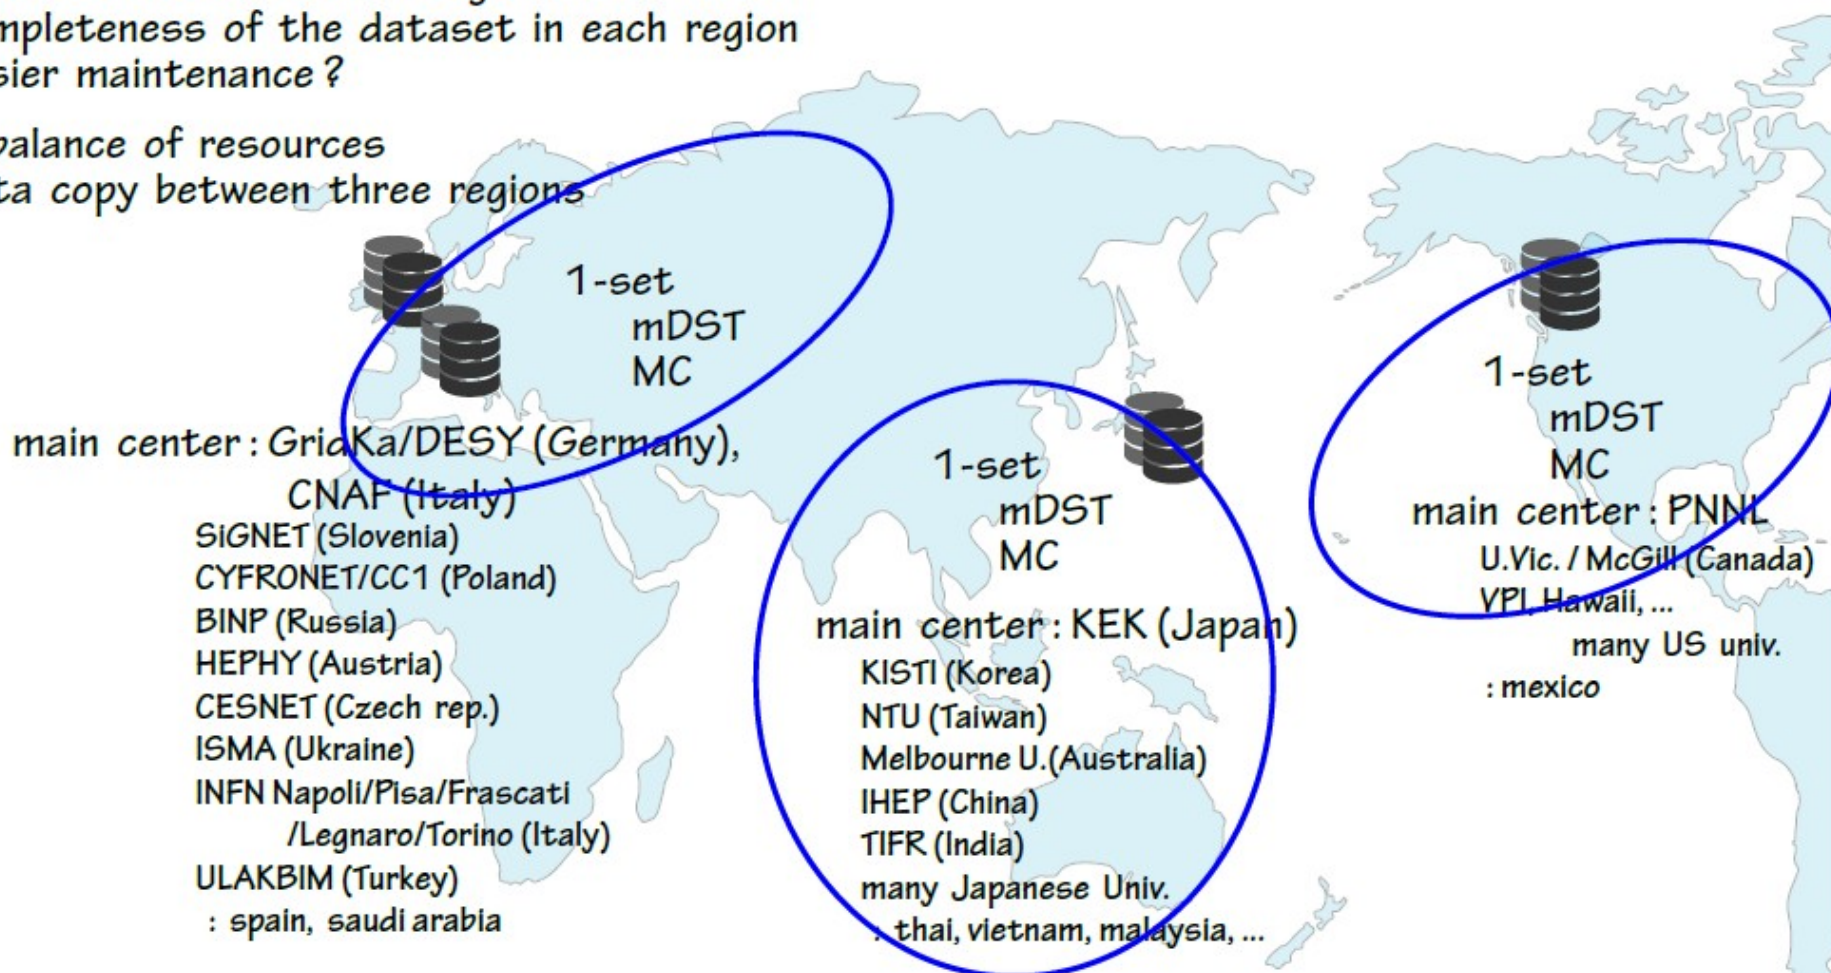
# mDST & MC data distribution

mDST (data) is copied in Asia, Europe, and USA

For the MC data seems to be natural to be the similar structure

better network? in each region  
completeness of the dataset in each region  
easier maintenance?

unbalance of resources  
data copy between three regions







# Network

## Current Connectivity

### Trans-Pacific

10G : Tokyo - LA  
10G : Tokyo - NY  
10G : Osaka - Washington

### Trans-Atlantic

3 x 10G : NY - Amsterdam  
3 x 10G : Washington - Frankfurt  
ANA-100G NY - Amsterdam

### Trans-Asia

2.5G : Madrid-Mumbai  
2.5G : Singapore-Mumbai  
10G : Japan-Singapore

## “Planned” Connectivity

### Trans-Pacific

SINET5  
100G link to US  
in 2016

### Trans-Atlantic

EEX (ESNet Extension to Europe)  
2 x 100G : NY - London  
100G : Washington - Geneva  
40G : Boston - Amsterdam

### Trans-Asia

10G : Mumbai - GEANT  
SINET ?

# Deployment of the virtualized computing infrastructure

Deployment stages:

- Virtualizing of an experimental group's computing environment.
- Tests of virtual machines locally on BINP resources.
- Transferring the VMs to a supercomputer and running them under remote batch system's control.
- Integration of local and remote batch systems.

Finally we have dynamical virtualized computing cluster. Physicists use computing resources in a conventional way.