Virtualized HPC infrastructure of the Novosibirsk Scientific Center

for HEP data analysis

D. Maximov

on behalf of NSC/SCN consortium

International Workshop on Antiproton Physics and Technology at FAIR

Budker INP, Novosibirsk, Russia

19 November 2015

Contents

BINP/GCF — History

- started in 2004
- initial goal: participate in LCG project
- currently: a gateway to the NSC computing resources for BINP experimental groups.

Supercomputing Network of the Novosibirsk Scientific Center

Isolated 10 Gbps network connecting main computing resources of Akademgorodok

Organizations involved:

- Institute of Computational Technologies (ICT SB RAS)
- Novosibirsk State University (NSU)
- Institute of Computational Mathematics and Mathematical Geophysics (ICM&MG SB RAS)
- Budker Institute of Nuclear Physics (BINP SB RAS)

Expansion perspectives: other NSC institutes, Tomsk State University

Supercomputing Network of the NSC

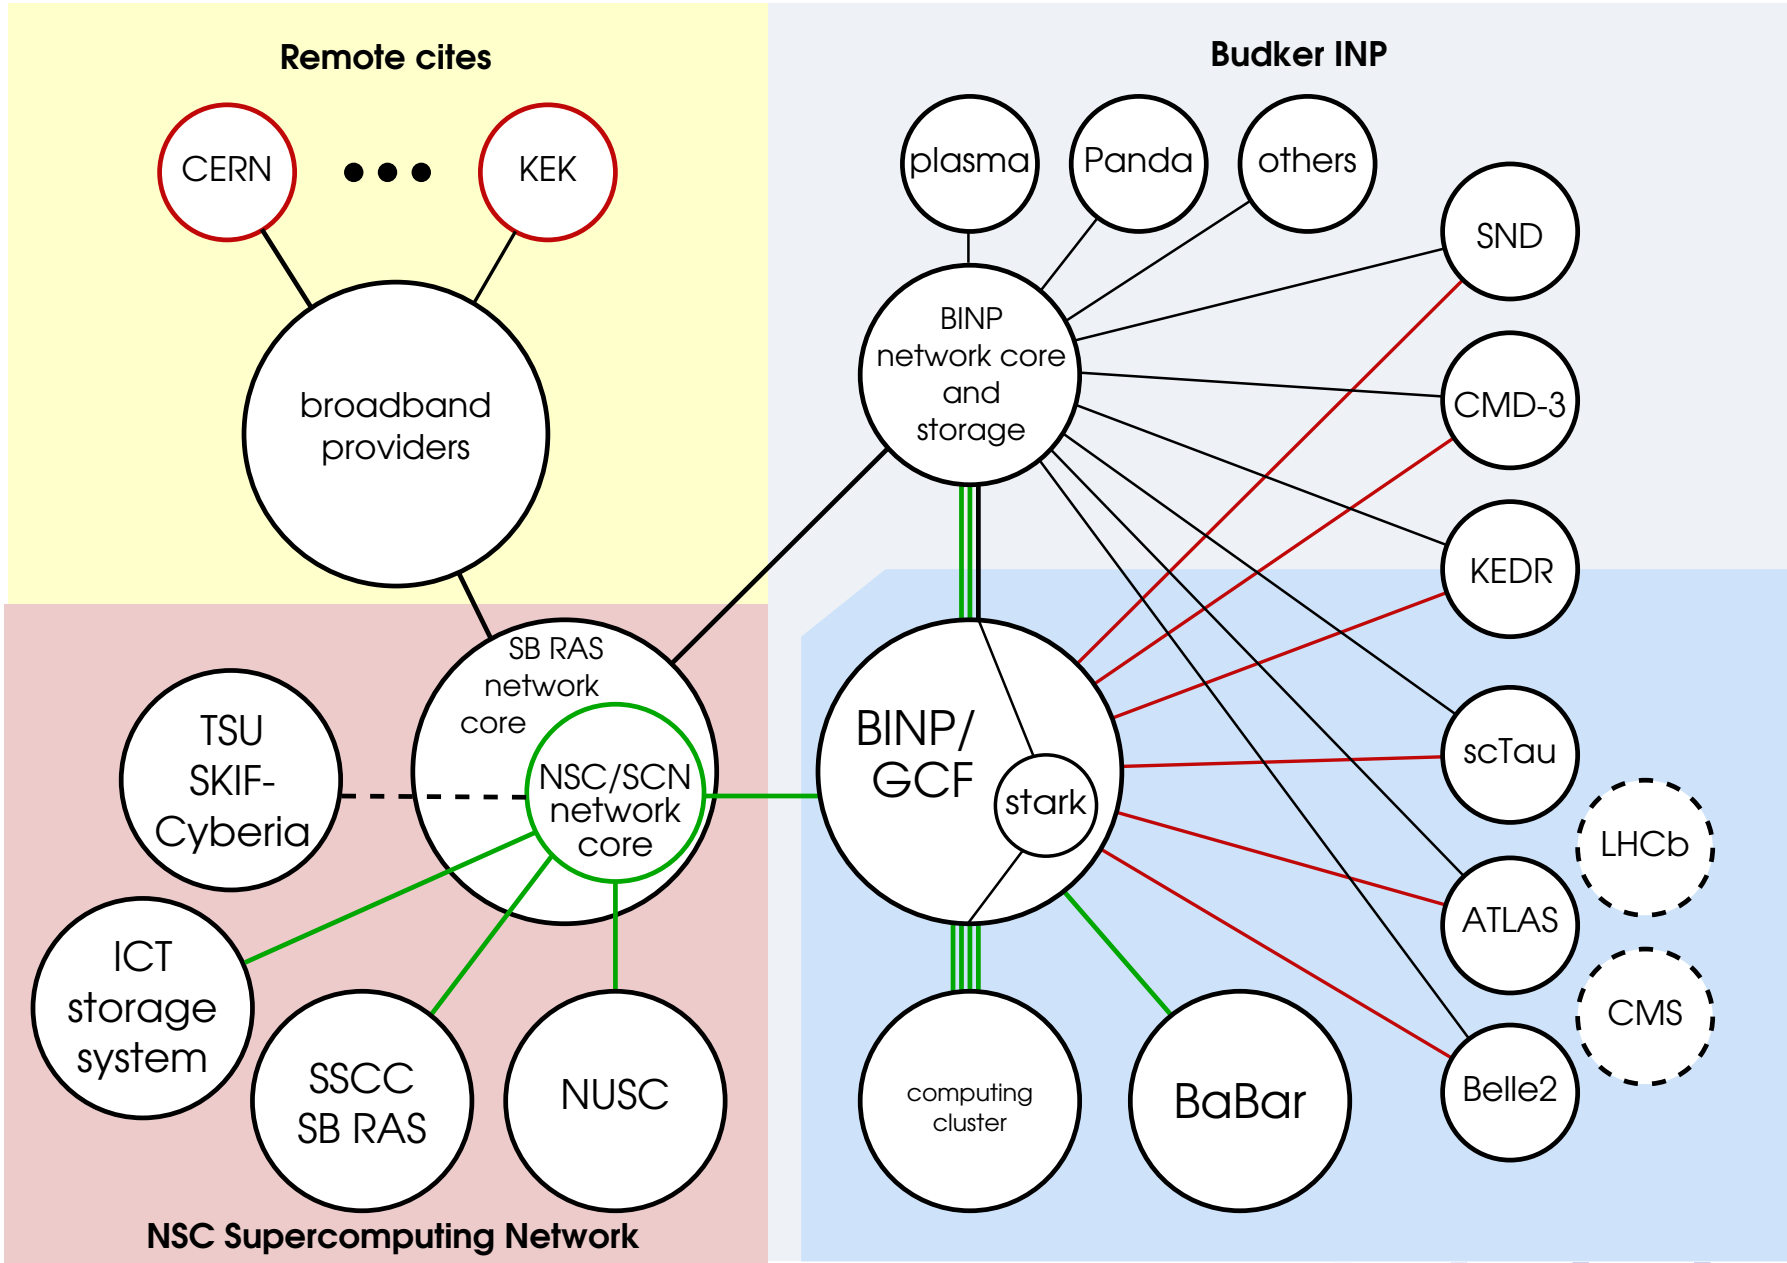

Novosibirsk State University (NSU) Supercomputer Center (NUSC)

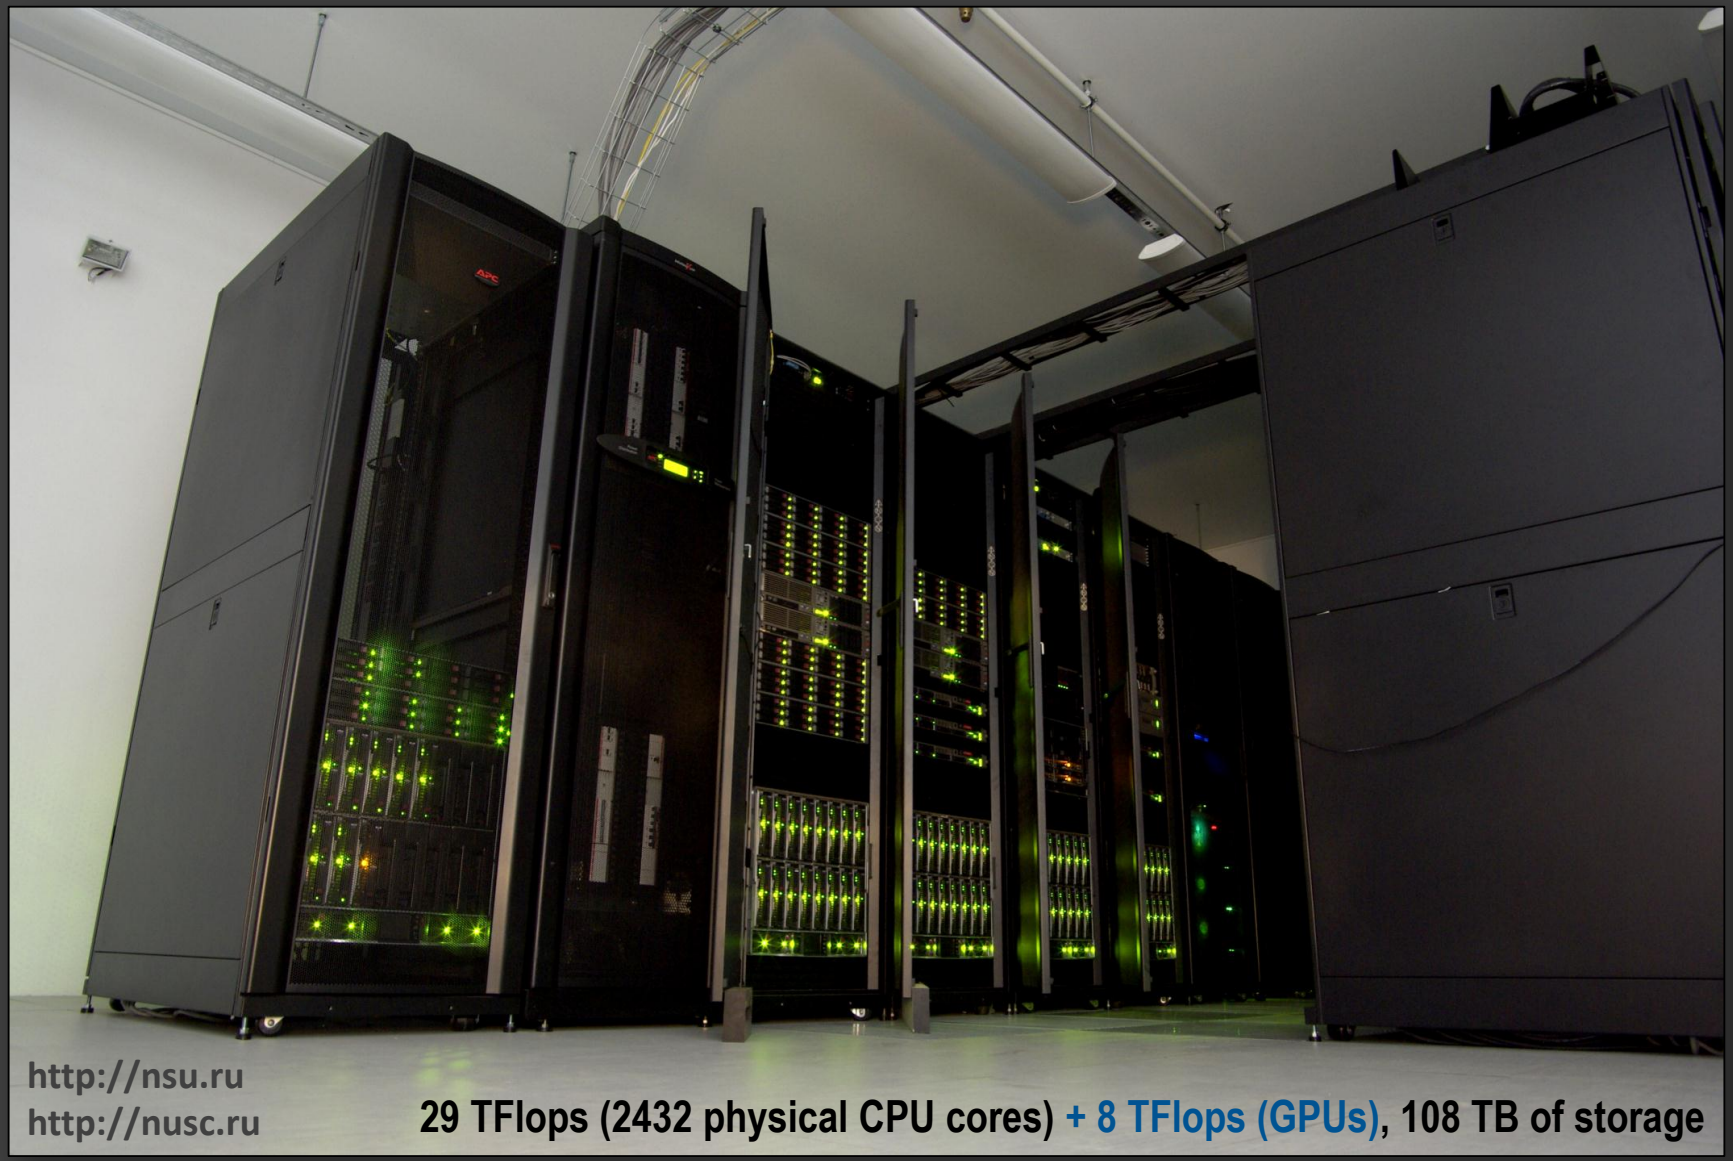

http://nsu.ru
http://nusc.ru

29 TFlops (2432 physical CPU cores) + 8 TFlops (GPUs), 108 TB of storage

Siberian Supercomputer Center (SSCC) at the Institute of Computational Mathematics & Mathematical Geophysics (ICM&MG)

SSCC was created in 2001 in order to provide computing resources for SB RAS research organizations and the external users (including the ones from industry)

30 + 85 TFlops of combined computing performance since 2011Q4 (CPU + GPU)

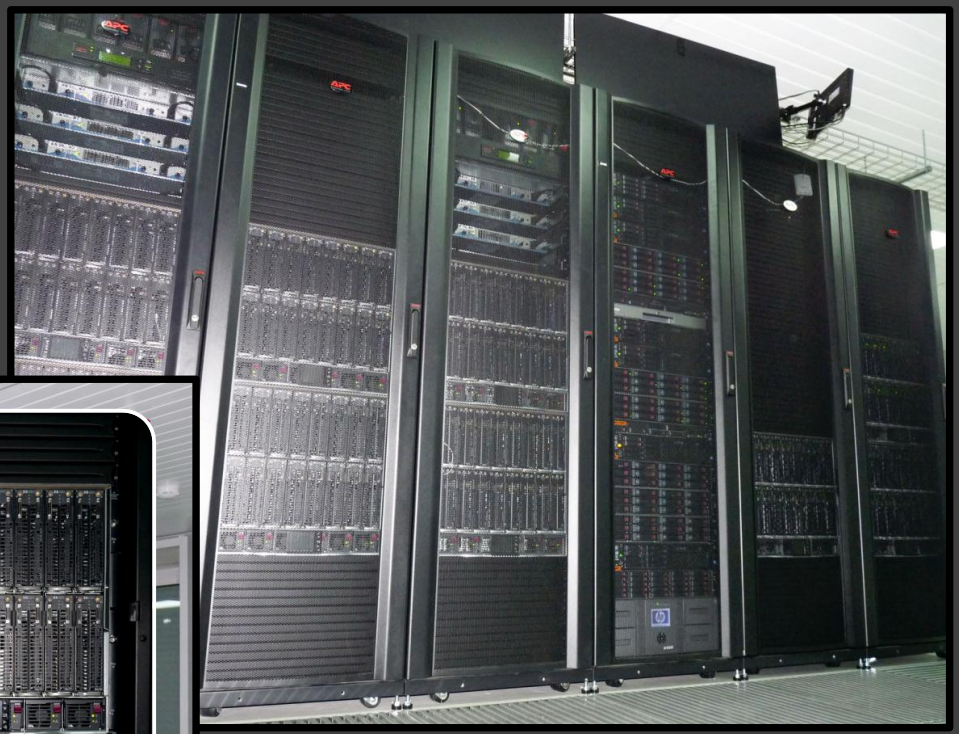

◀ NKS-30T ▲

120 TB of storage

2x 70 sq.m of raised floor space

Up to 140kVA power input & 120kW of heat removal capacity (combined)

CPU/Cores: 128/512, 5 TFlops Peak

Key properties of the HEP computing environment

- Each experiment has unique computing environment
  - wide variety OS and standard packages versions
  - a lot of specifically developed software
- Software can be easily parallelized by data
- Mostly non interactive programs, executed via some batch system.

How to glue HEP and HPC together?

We want:

on the HEP side:

- keep the specific computing environment and user's experience

on the supercomputer side:

- be like a normal SC user

The answer is:

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

Batch System Integration Mechanisms

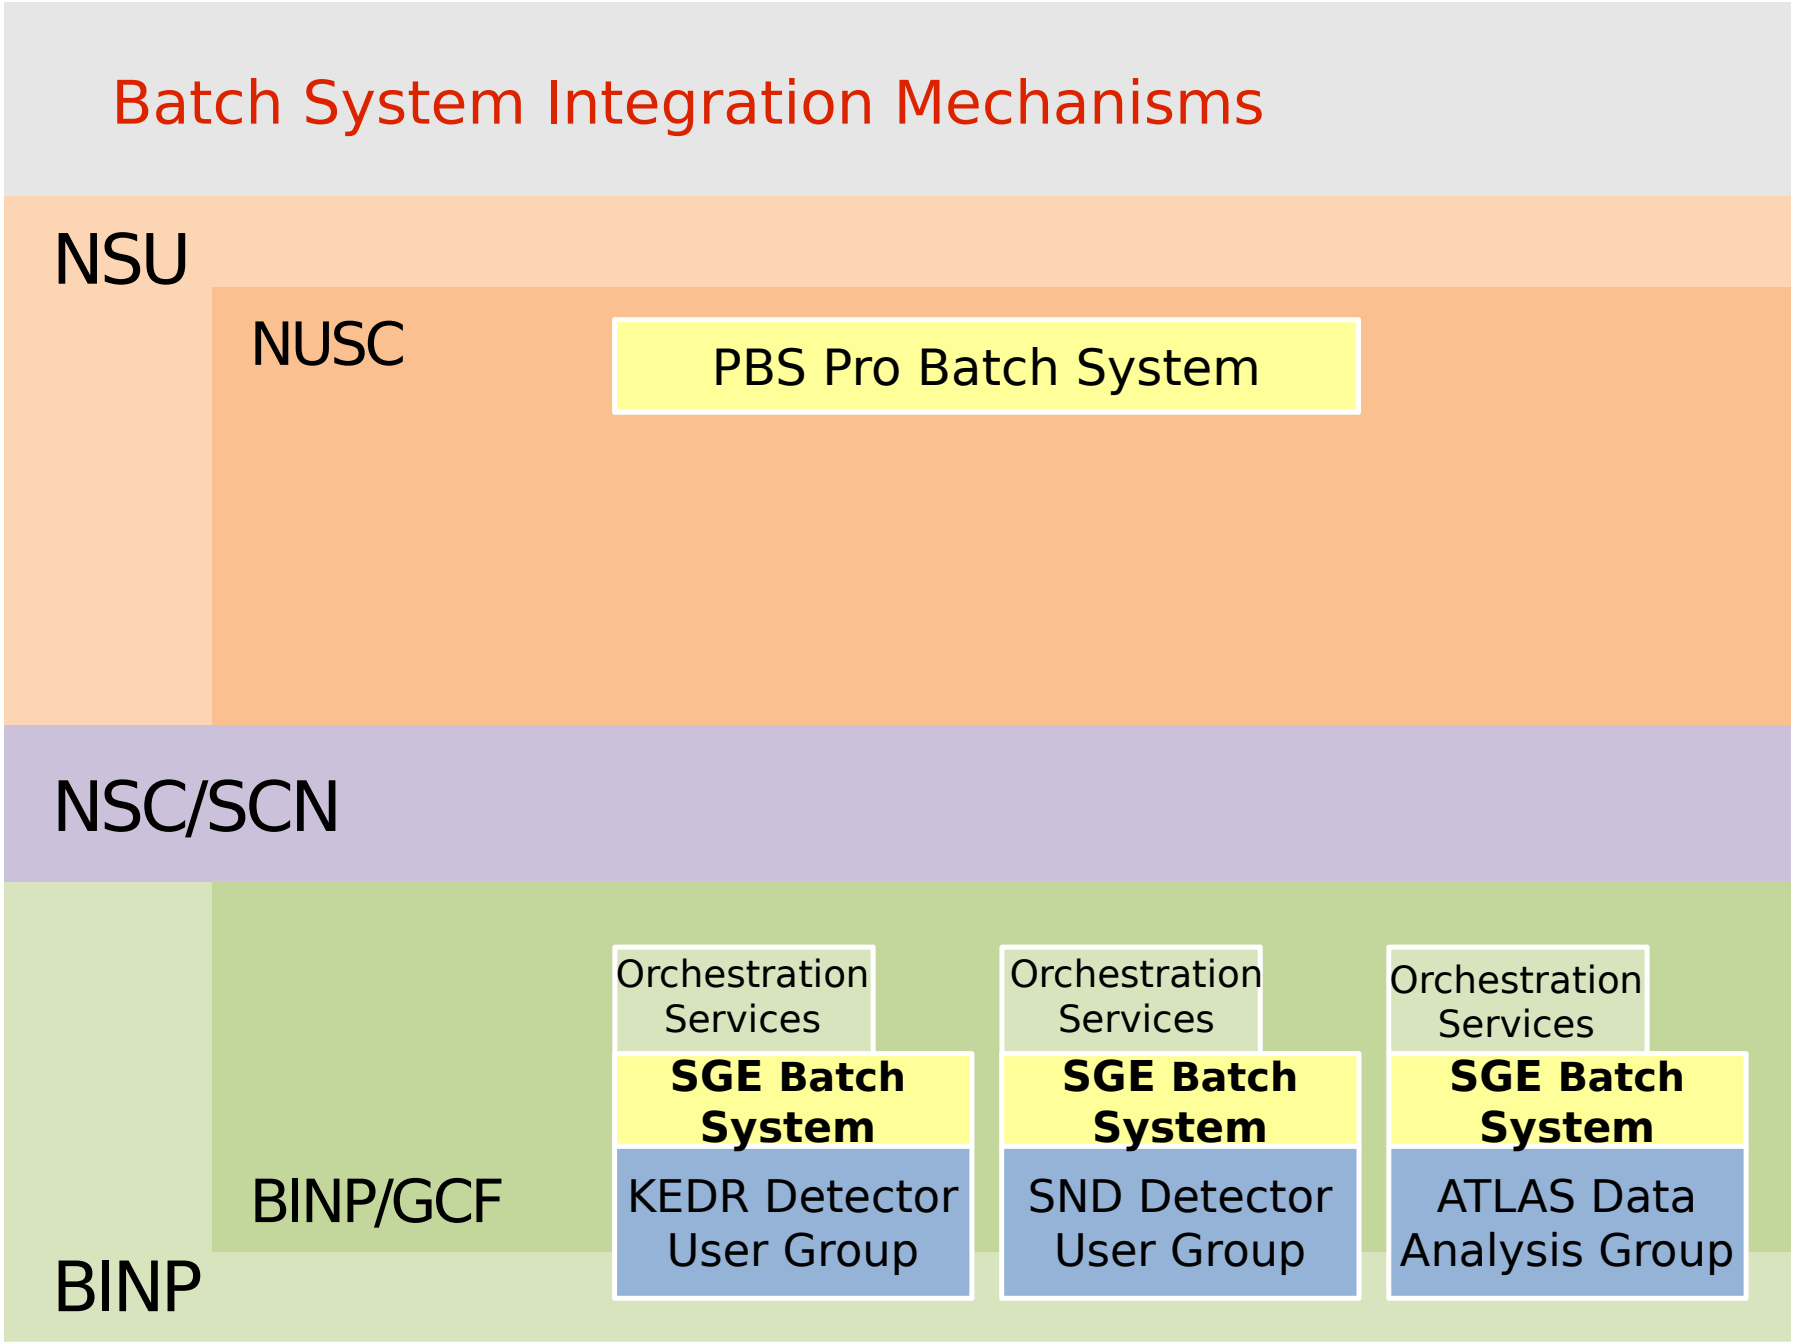

Batch System Integration Mechanisms

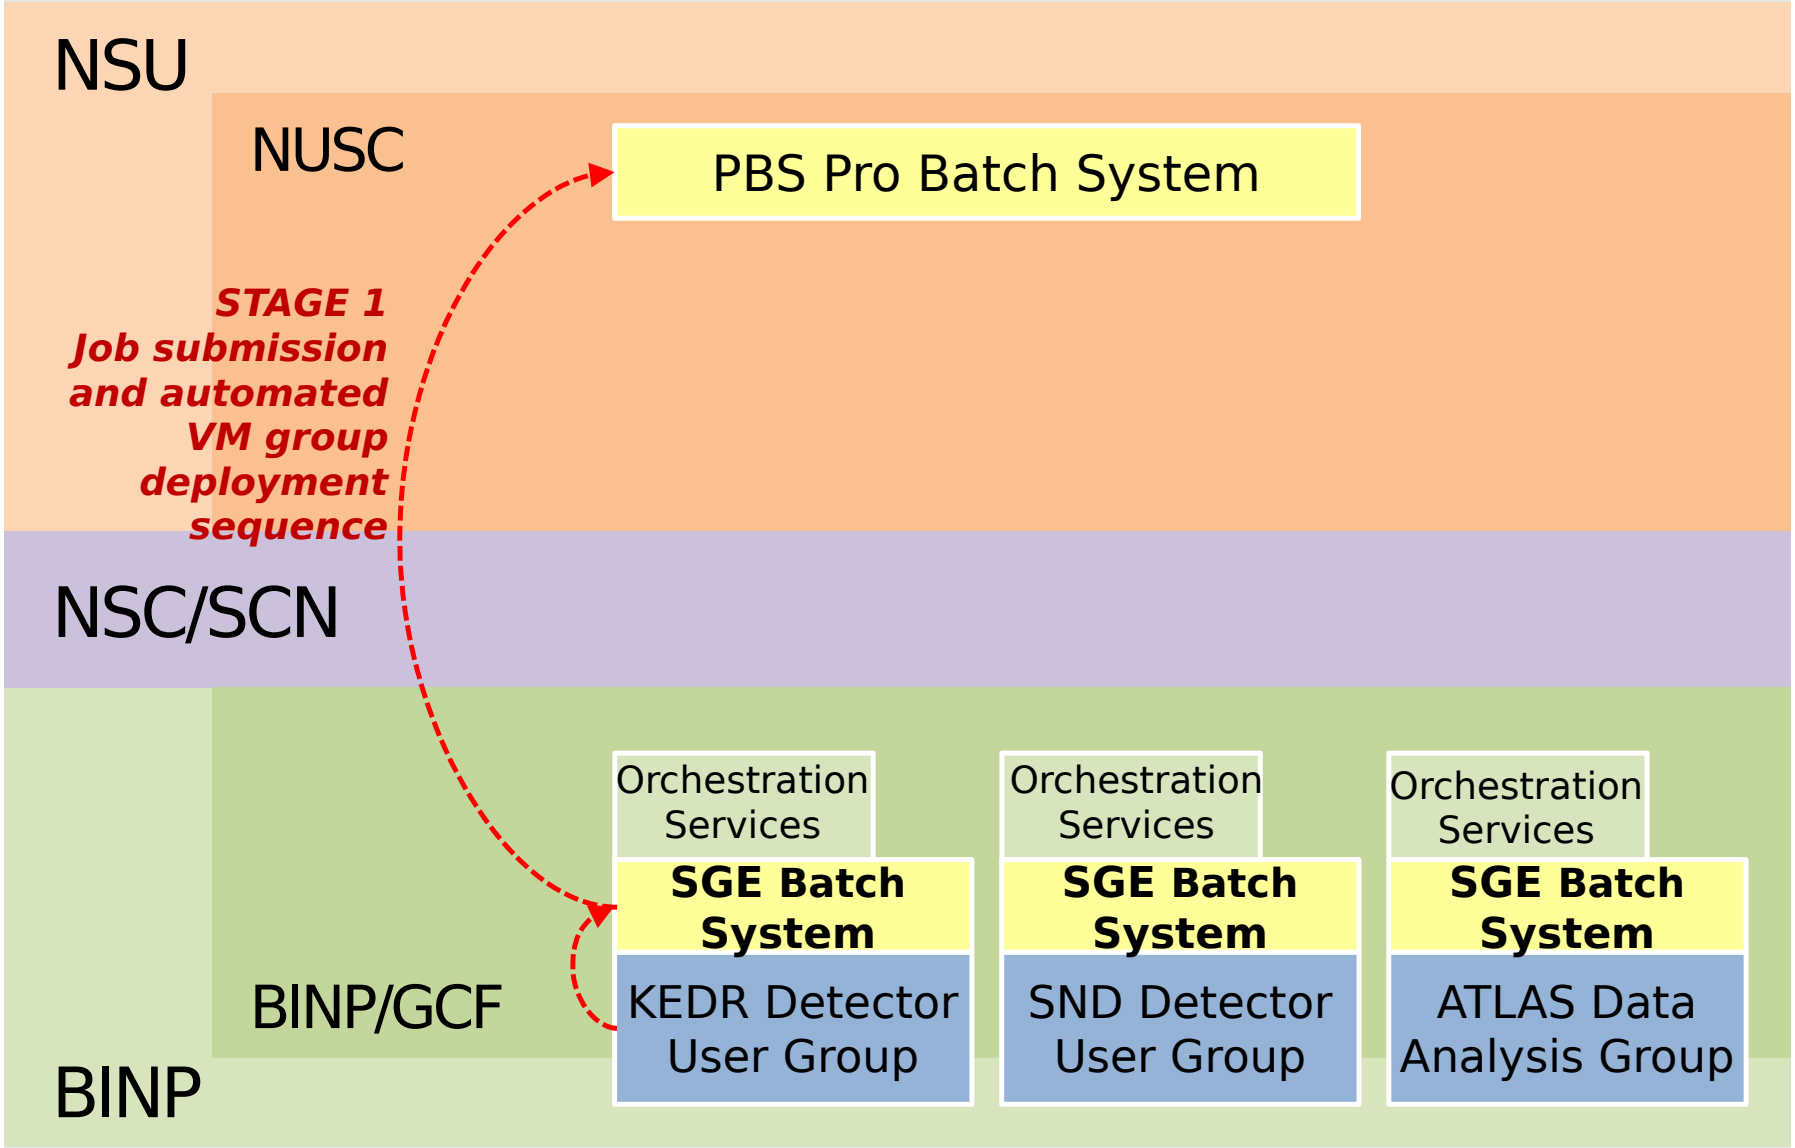

Batch System Integration Mechanisms

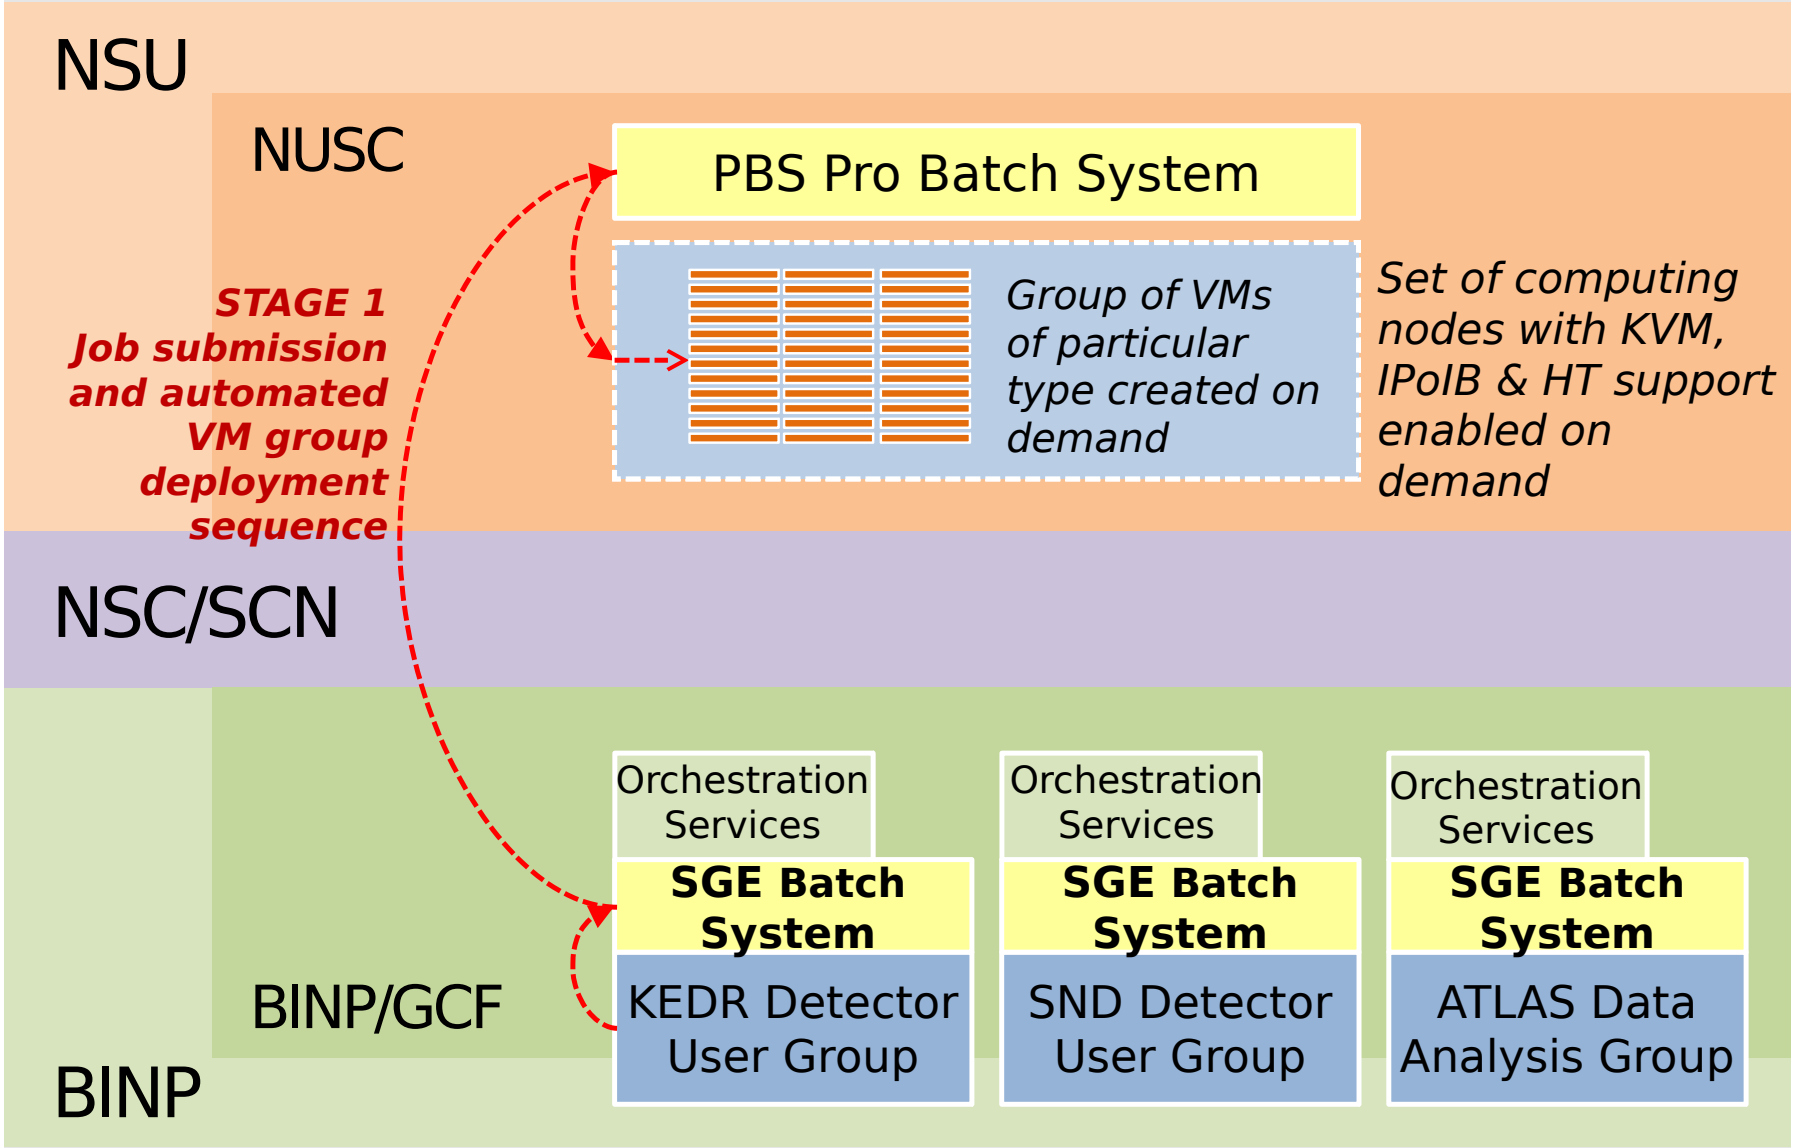

Batch System Integration Mechanisms

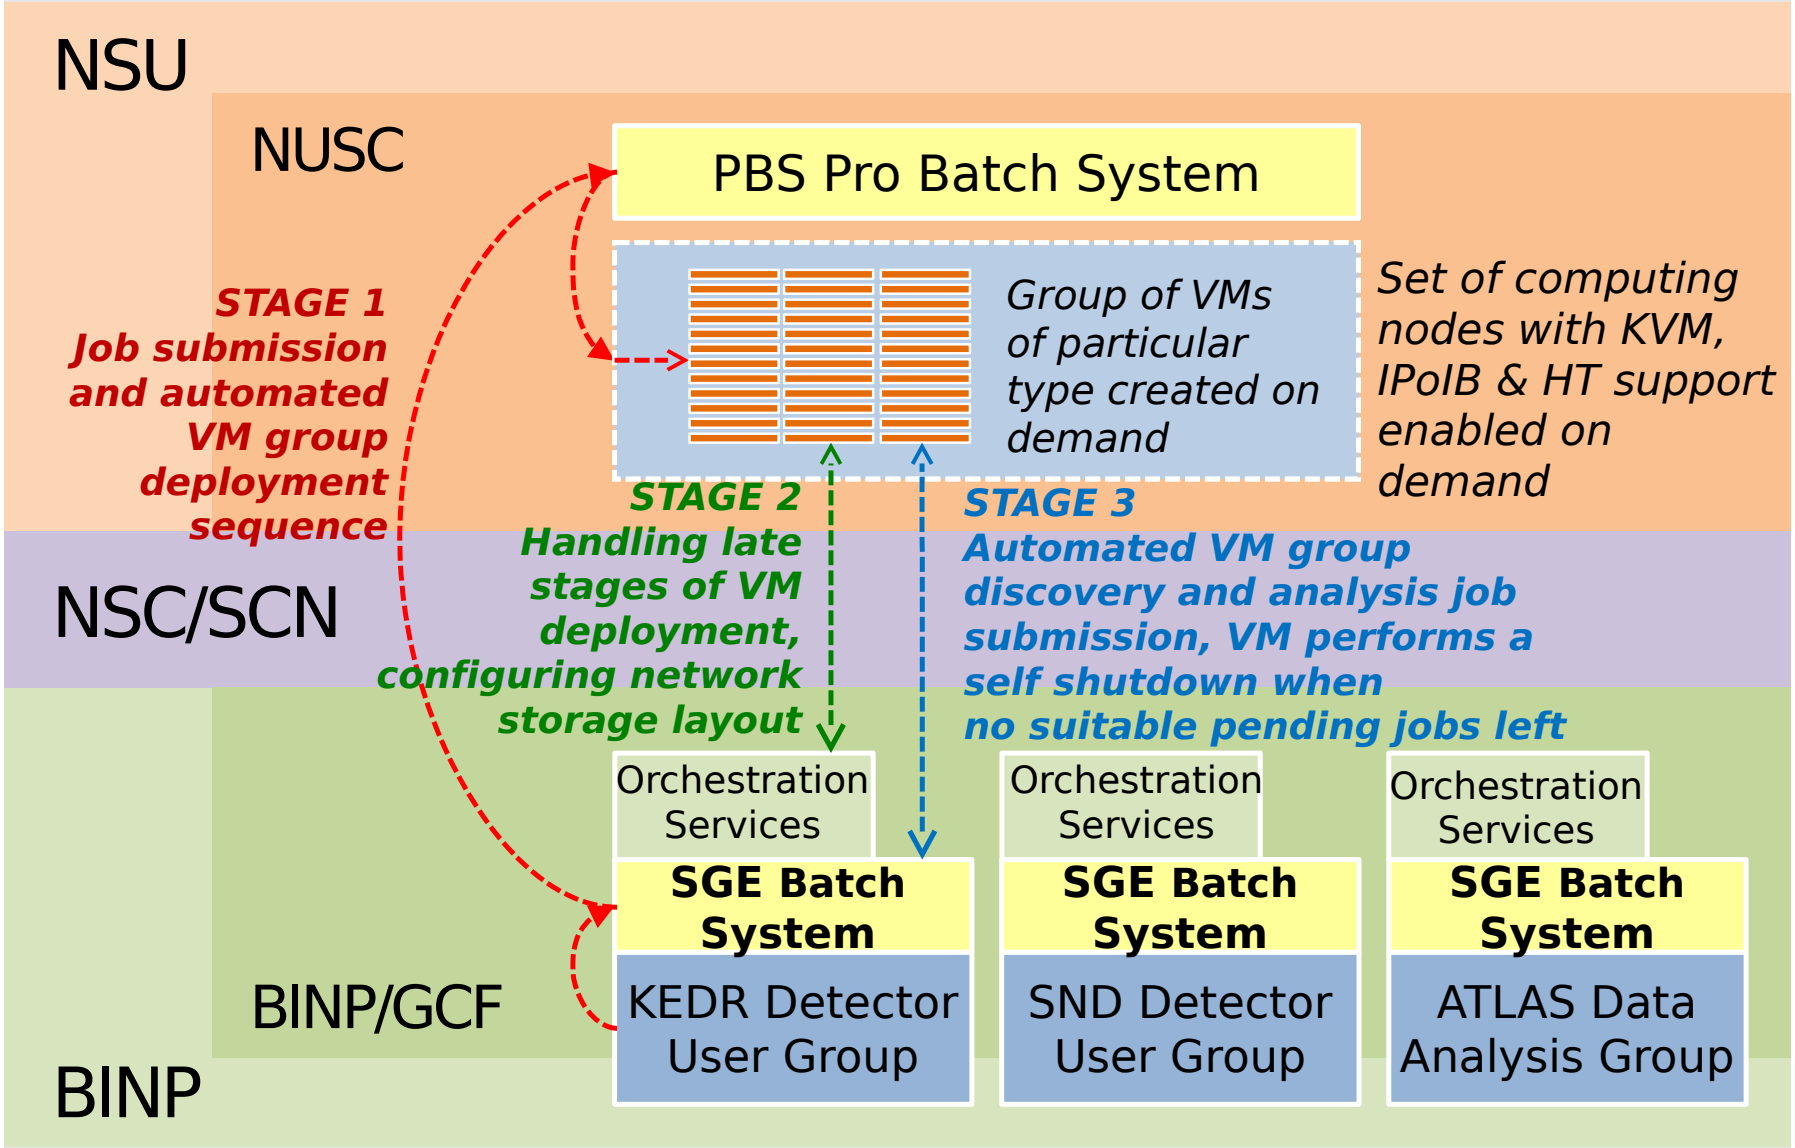

Batch System Integration Mechanisms

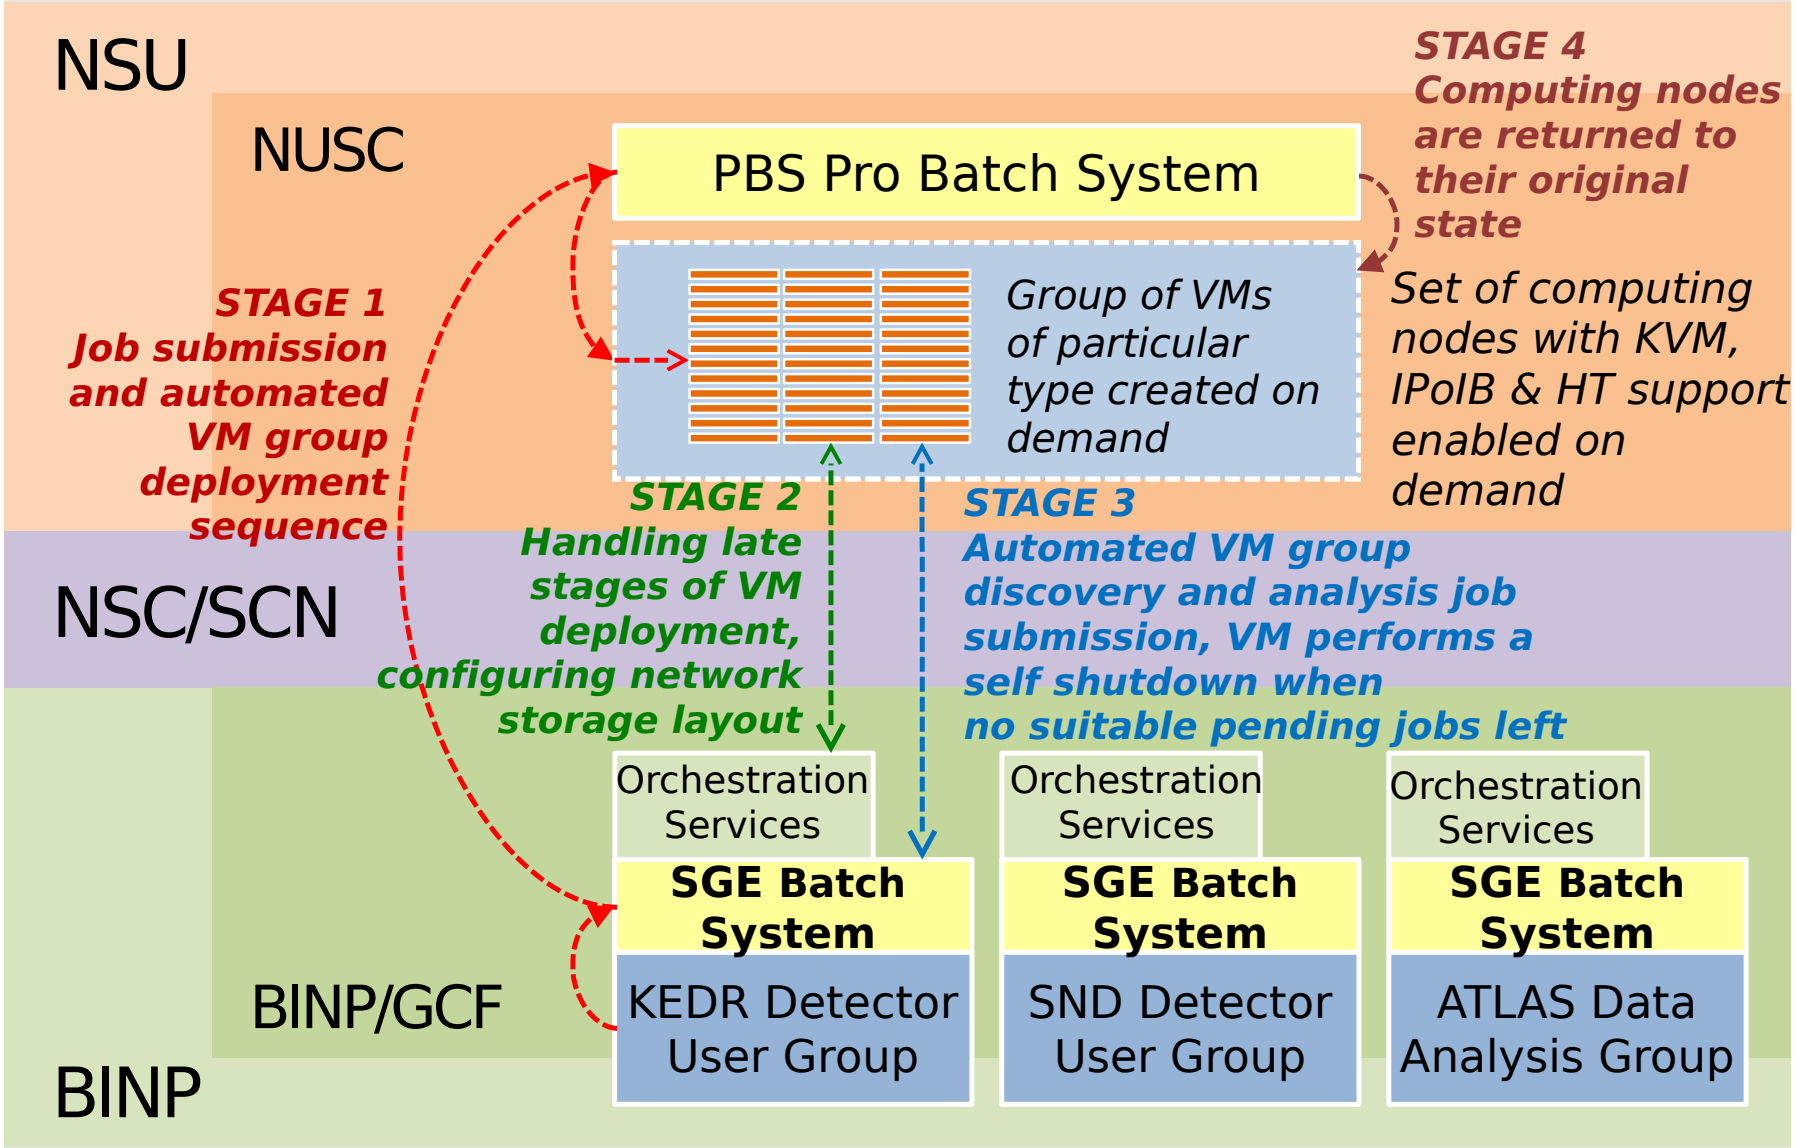

Virtualized computing infrastructure

In this way we have *dynamical virtualized computing cluster* (DVCC).
Physicists use computing resources in a conventional way.

Experiments at Budker INP integrated with GCF

High Energy Physics

- Local
  - KEDR
  - CMD-3
  - SND
  - Super-$c\tau$ (planned)
- External
  - ATLAS
  - Belle2
  - BaBar

Other activities

- Plasma & accelerator physics, engineering calculation...

Virtualized infrastructure: what we've learnt so far

HEP data analysis could be successfully performed using the virtualized HPC infrastructure of the Novosibirsk Scientific Center

- Long term VM stability obtained (up to a month at NUSC, up to year at BINP)
- Most of the underlying implementation details are completely hidden from the users.
- No changes were required for experimental group's software and/or its execution environment.

Virtualized infrastructure: what we've learnt so far (2)

Main benefits:

- Ability to use free capacity of supercomputer sites in order to run much more simple (from HPC point of view) single threaded HEP software.
- Ability to freeze software of an experimental group and its execution environment and (exactly!) reproduce it when needed.
- This scheme could be easily extended for other experimental groups and computing centers.

Milestones

- Initial deployment of GCF at BINP in 2004.
- DVCC middleware development started in 2010.
- KEDR runs production jobs at NUSC since 2011Q1.
- Other BINP groups joined the activity in 2012Q1.
  (Started by ATLAS)
- SSCC connected in August 2012, in production since 2013Q1.
- Belle2 experiment joined in 2014Q2.
  Works through the DIRAC system.

Usage of the local cluster

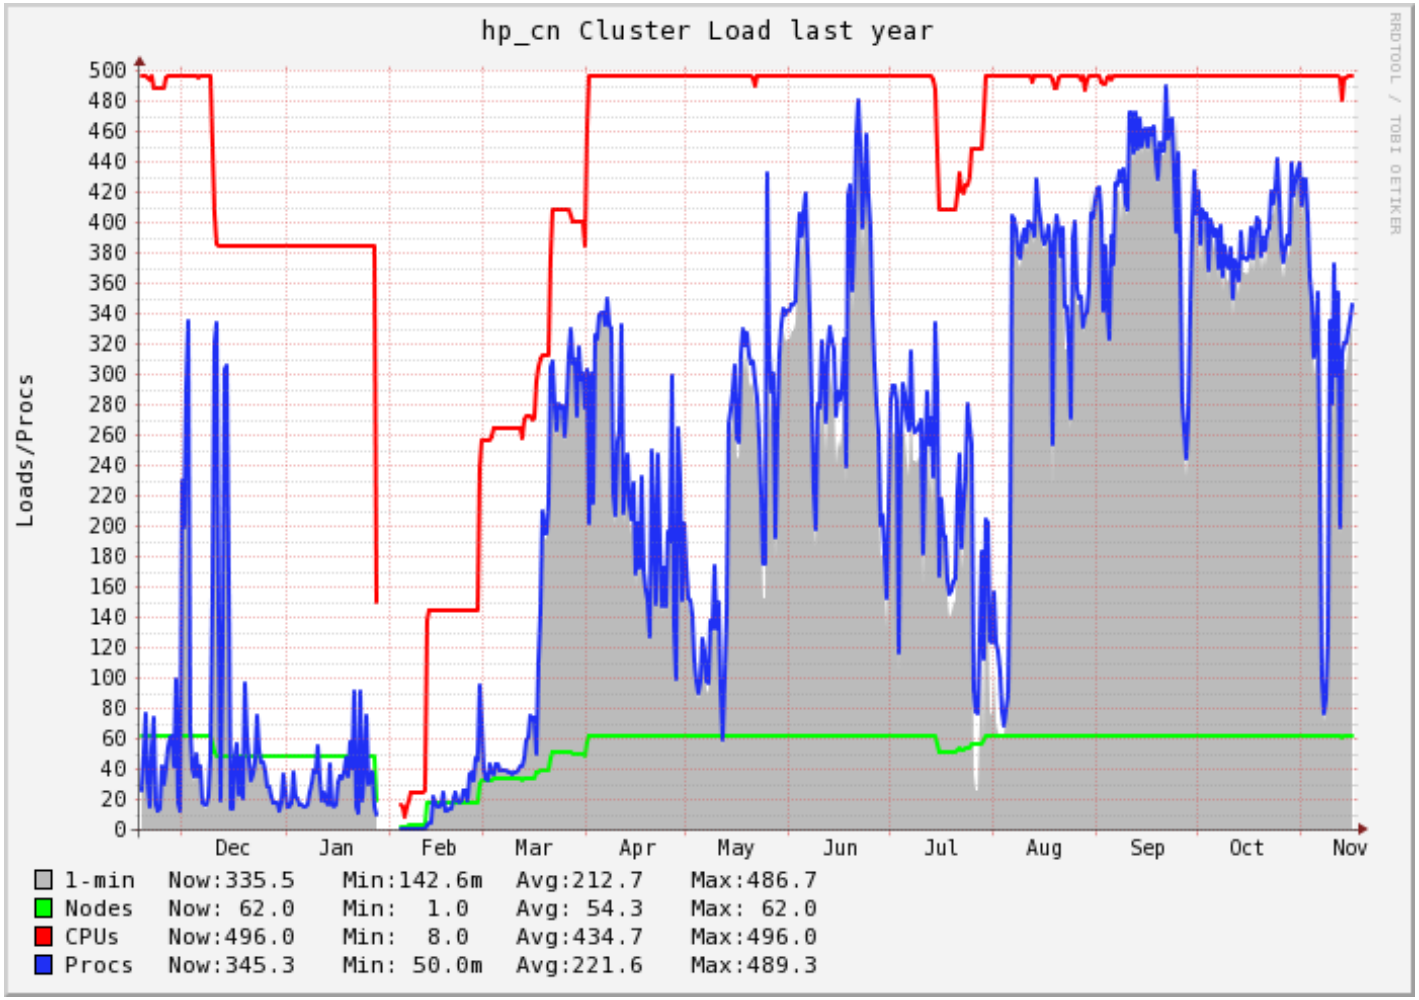

Average Utilization (last year) — 49%.

Participation in Belle2

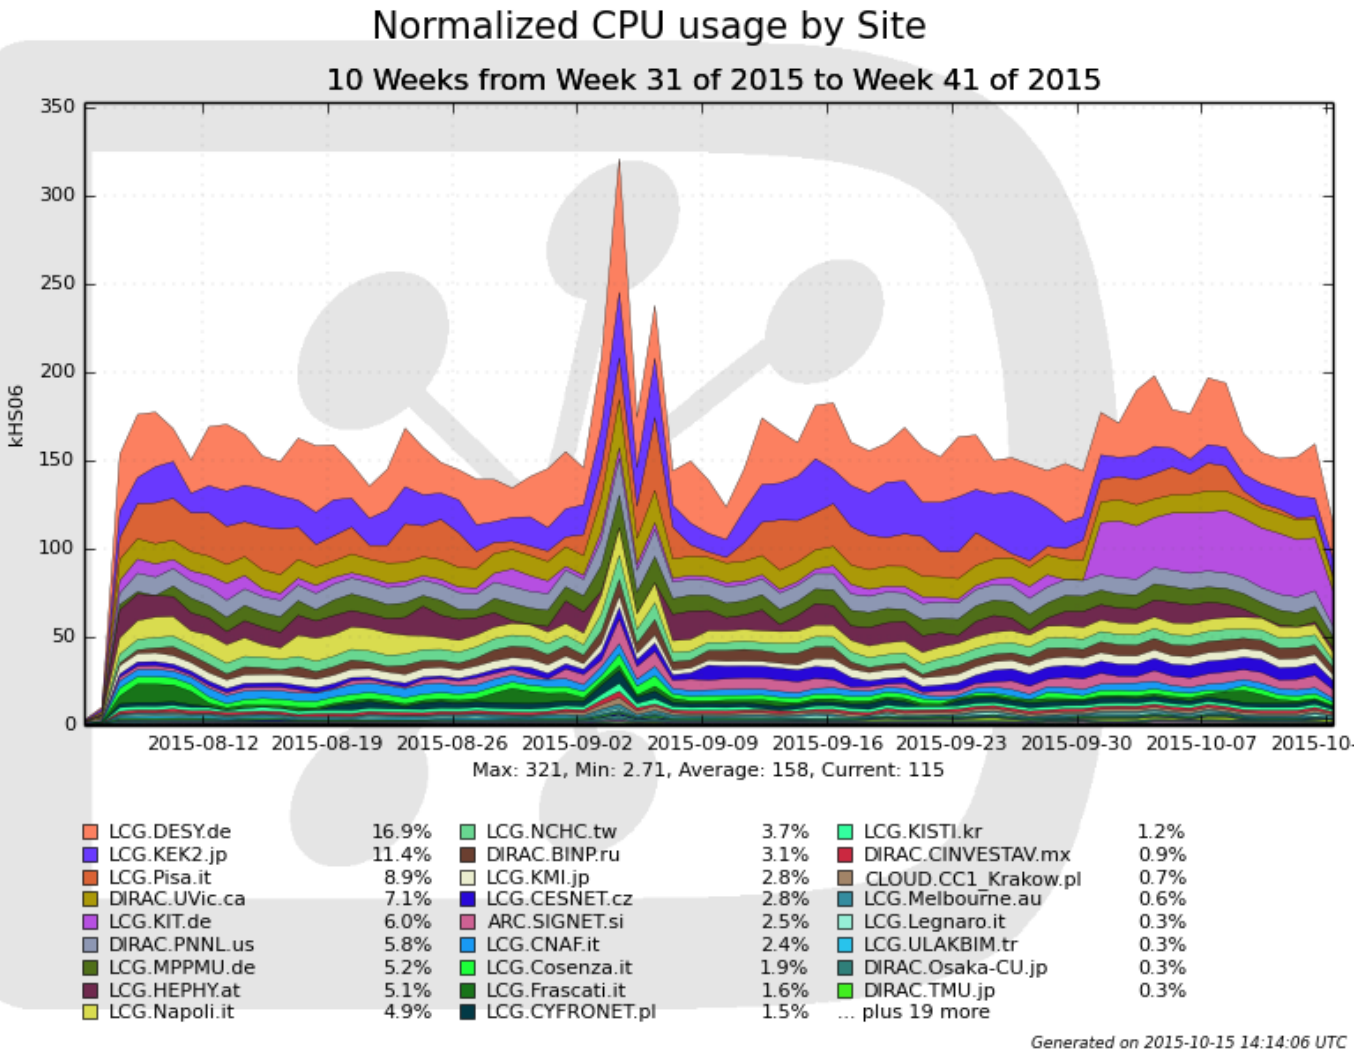

Produced — about 3% CPU-hour for simulation.

Conclusion

Results

- NSC/SCN resources are successfully used for HEP data processing in BINP making analysis cycles 10-100 times faster.
- Our experience could be applied for other computing centers and HEP experiments.

Plans

- Install new computing and storage hardware
- Extend the number of user groups
- Further development of the DVCC middleware.
- Deploy special storage for BaBar experiment.
- Access to others computing centers
- Present our resources to LCG

Thank you for attention!

Key features of the virtualized infrastructure

- Virtualization by KVM
- VM disk images are located on SC's file system
- Input/output data are located at BINP/GCF and accessed by VMs over the network.
- VMs are just regular batch tasks at a supercomputer.
- VMs are started automatically on user's demands.

Results from KEDR detector

Measurement of main parameters of the $\psi(2s)$ resonance (http://arxiv.org/abs/1109.4215)

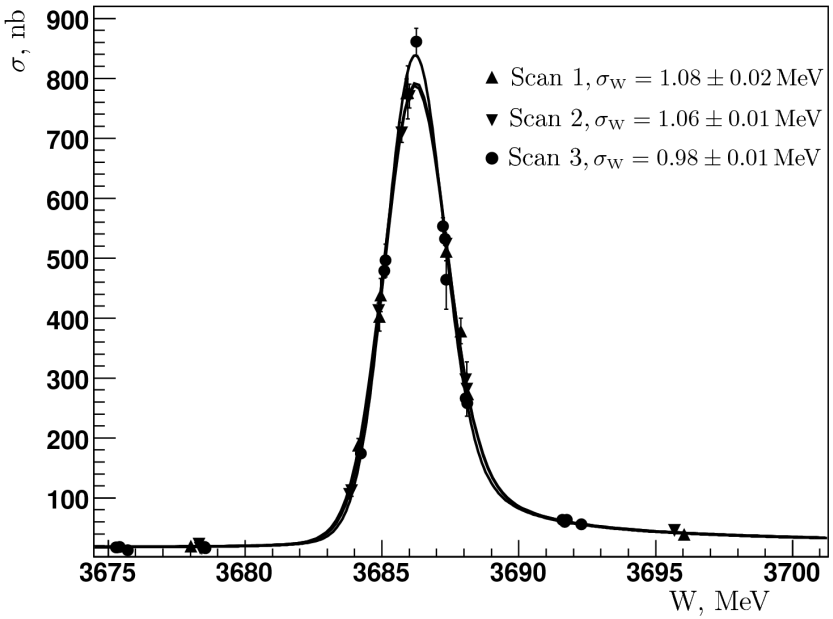

The multihadron cross section as a function of the c.m. energy for three scans in the $\psi(2s)$ region

Measurement of $\psi(3770)$ parameters (http://arxiv.org/abs/1109.4205)

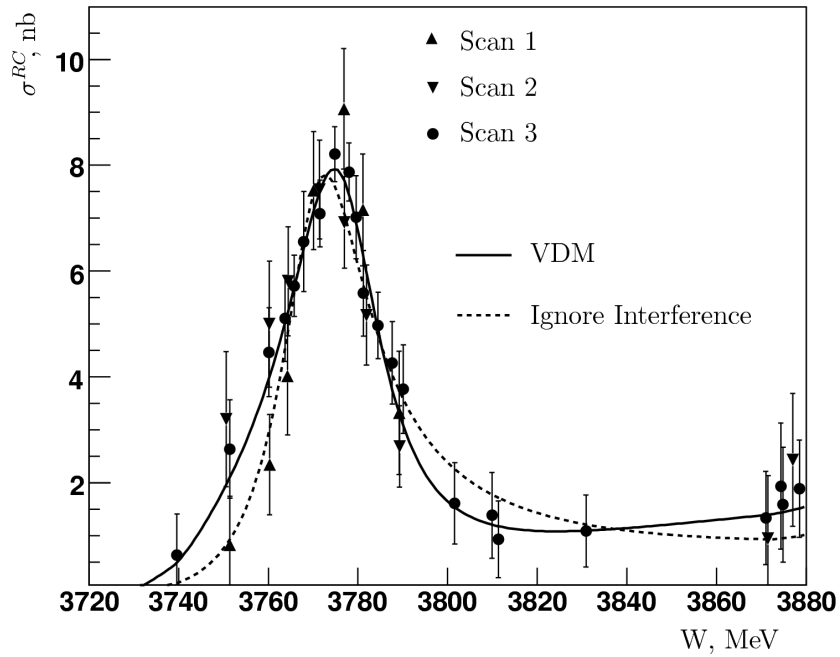

Cross section of $e^+e^- \to$ hadrons vs. c.m. energy in the vicinity of $\psi(3770)$

ATLAS analysis

G. Aad et al. (ATLAS Collaboration) *"Search for heavy neutrinos and right-handed W-bosons in events with two leptons and jets in pp collisions at sqrt(s)=7 TeV with the ATLAS detector"*, Eur.Phys.J. C72 (2012) 2056.

Experimentally observed and expected from MC simulation 95% CL limits on Majorana neutrino mass M_N and right-handed W-boson M_{W_R}.

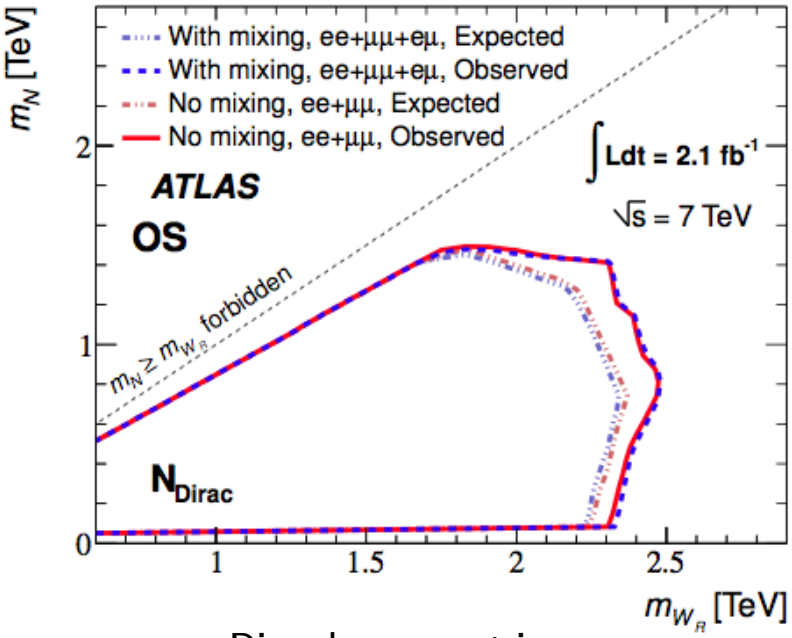

Dirac heavy neutrino

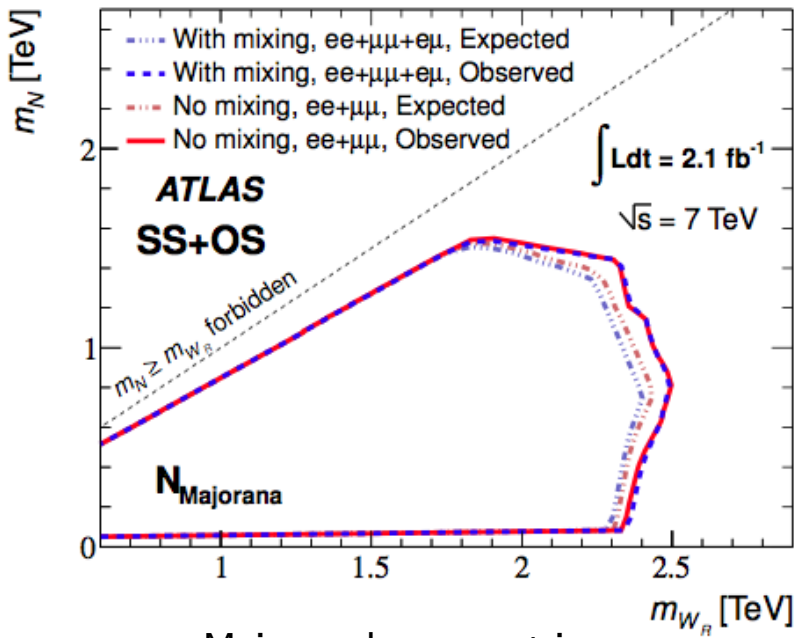

Majorana heavy neutrino