





New features of the JINR cloud

A. V. Baranov¹, N. A. Balashov¹,
N. A. Kutovskiy¹, A.N. Makhalkin¹,
Ye. Mazhitova^{1,2}, R. N. Semenov^{1,3}





- ¹ Laboratory of Information Technologies, Joint Institute for Nuclear Research
- ² Institute of Nuclear Physics, Almaty, Kazakhstan
- ³ Plekhanov Russian University of Economics, Moscow, Russia

JINR cloud architecture changes

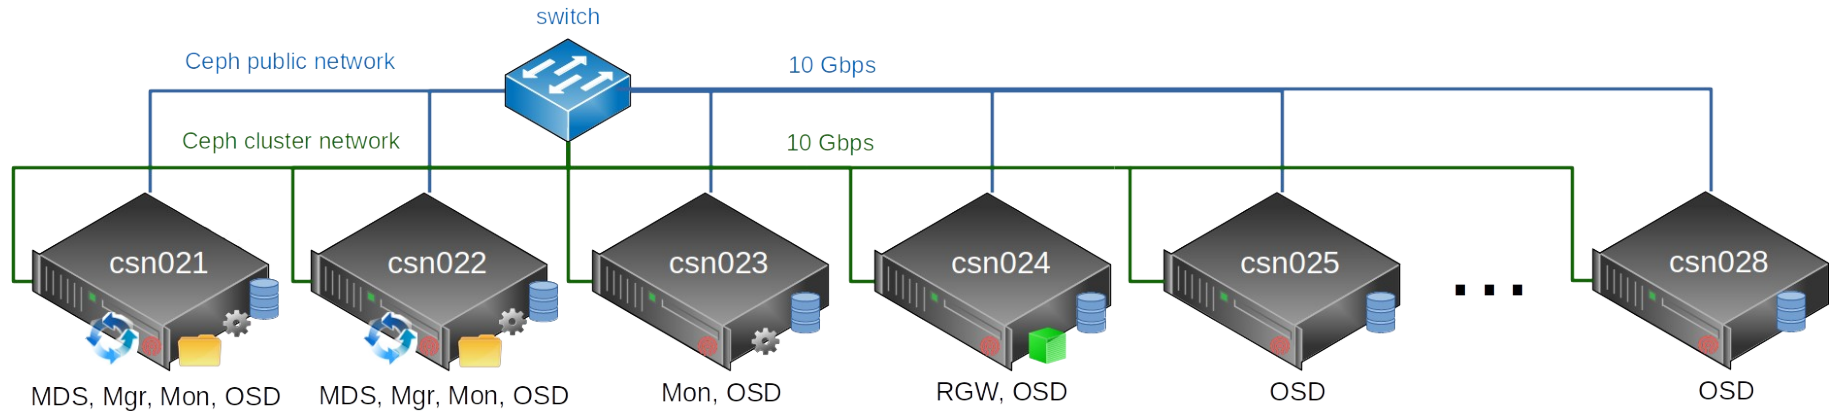


HA setup: 2 FNs, DRBD + Heartbeat
 Distributed storage: LizardFS, 2x replicas



HA setup: 3 FNs, leader elections based on raft consensus algorithm
 Distributed storage: ceph, 3x replicas

Ceph-based software defined storage

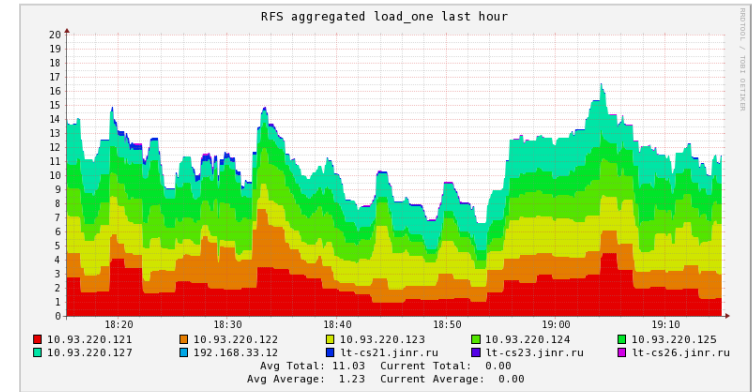


- 12.2.8 Luminous release
- HA setup: 3 MON, 2 MDS, 2 MGR
- Object storage: backups (git, discourse) over S3 protocol
- Block storage: KVM VM images, user data
- CephFS: user data, user software
- Ceph capacity:
 - Total raw disk space: ~896 TB (+128 TB will be added soon, so ~ 1 PB expected)
 - effective ~300 TB (+40TB) due to 3x replication

JINR cloud service monitoring

Limit Results: 100

Host	Service	Status	Last Check
cloud	ONE - MM SCHED	OK	02-10-2015 16:53:36
	ONE - ONED	OK	02-10-2015 16:54:02
	ONE - SUNSTONE	OK	02-10-2015 16:54:24



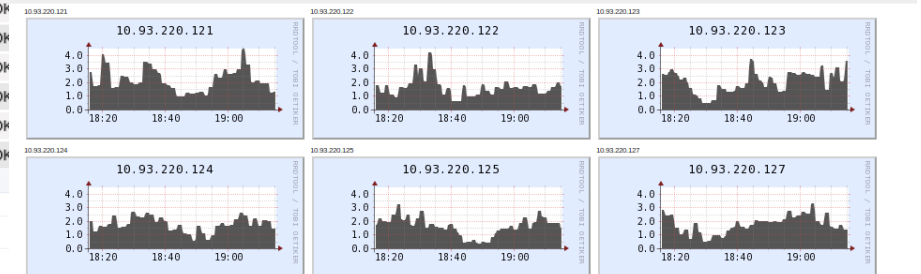
cloud-mon.jinr.ru/nagios/

Apps http://ark.intel.com Phantom 2 Vision

RFS load_one last hour sorted by name

Metric: load_one Show Hosts Scaled: Auto Same None Size: small Columns: 4 (0 = metric + reports)

Show only nodes matching Filter Max graphs to show: all Sorted: ascending descending by name



Nagios

Current Network Status
 Last Updated: Tue Feb 10 16:59:08 MS
 Updated every 90 seconds
 Nagios® Core™ 3.5.1 - www.nagios.org
 Logged in as nagiosadmin

View History For all hosts
 View Notifications For /...
 View Host Status Deta

Duration	Status
01 61d 6h 43m 29s	PING OK
51 27d 3h 52m 52s	PING OK
31 53d 23h 29m 42s	PING OK
01 55d 0h 26m 12s	PING OK
51 54d 1h 15m 32s	PING OK
51 61d 6h 43m 39s	PING OK

General

Home
Documentation

Current Status

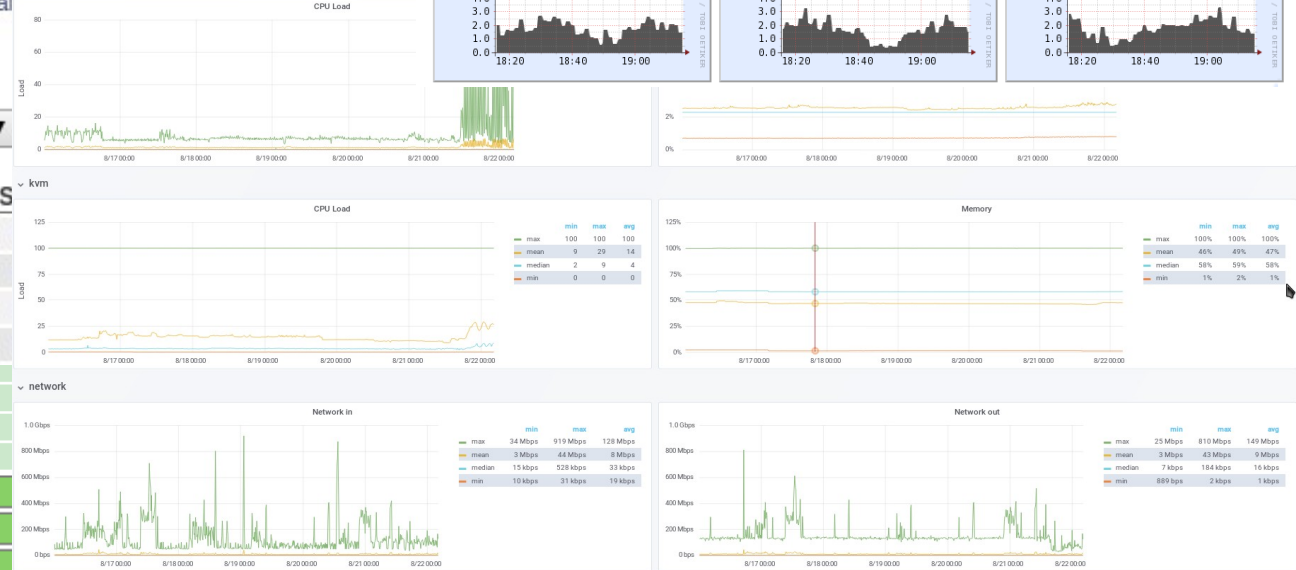
Tactical Overview
 Map
 Hosts
 Services
 Host Groups
 Summary
 Grid
 Service Groups
 Summary
 Grid

Limit Results: 100

Host

Host	Status
cloud	UP

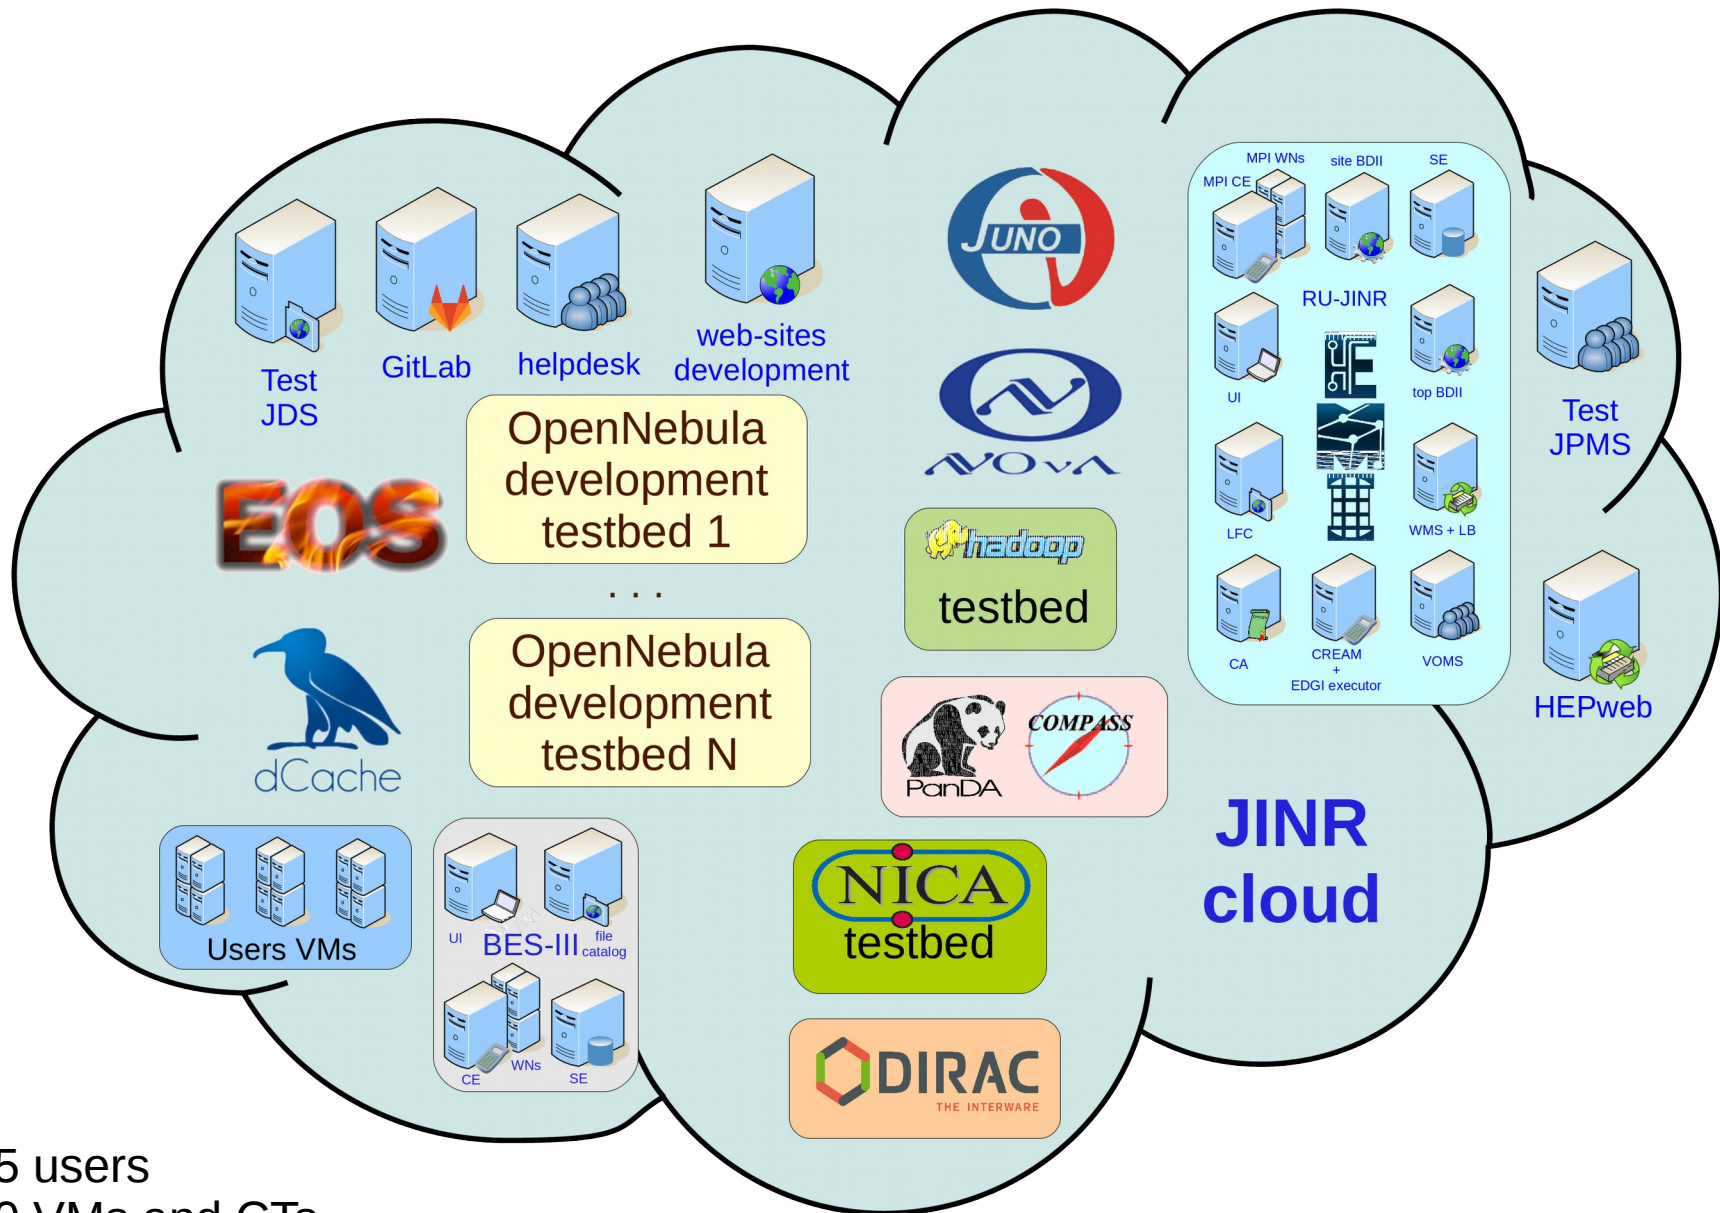
PING	lt-qt1-vm3	UP
RAID	lt-qt1-vm4	UP
RAM Count	lt-vm4	UP
SSH Server	lt-vm5	UP
SWAP		OK
Total Processes		OK
Uname		OK
		02-10-2015 16:56:24



Hardware resources

- Current hardware resources
 - ~80 servers
 - ~1600 CPU cores
 - ~8 TB of RAM
 - ~896 TB (+128 TB will be added soon) of ceph-based SDS raw disk space, effective ~300 TB (+40TB) due to 3x replication


Usage



~115 users

~240 VMs and CTs

Clouds integration (1/2)

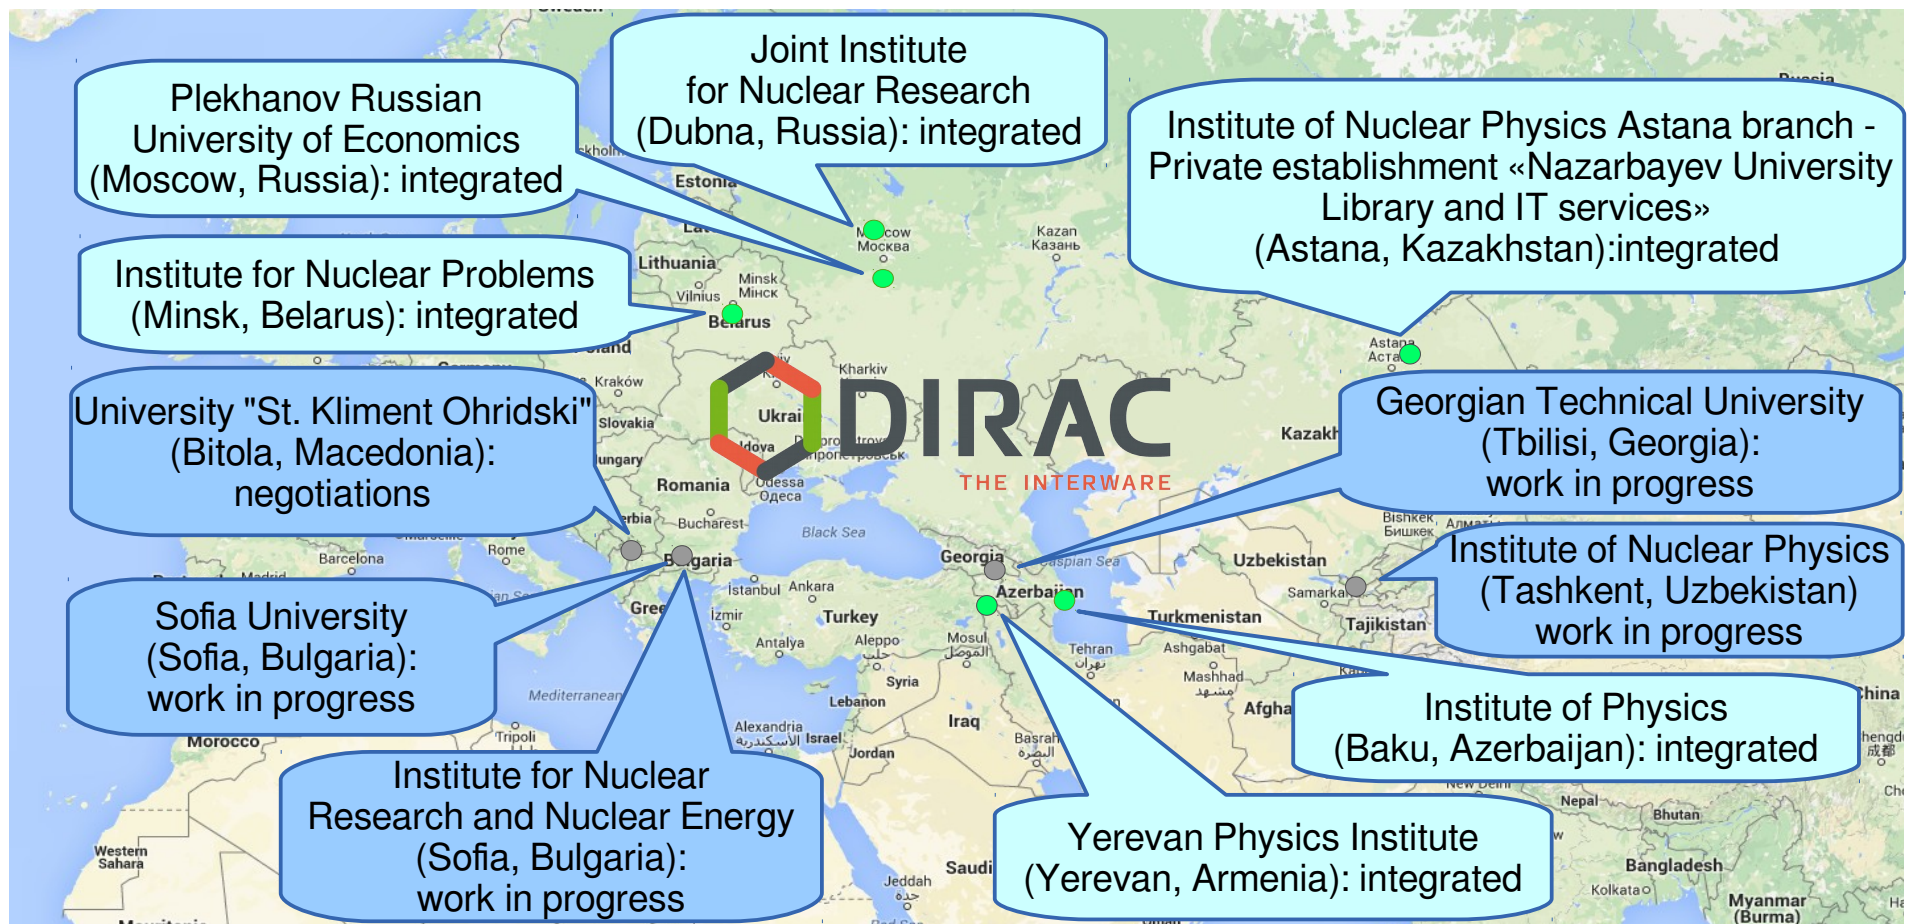


Cloud bursting approach:
- poor scalability
- hard to maintain

Grid middleware (interware) – DIRAC:
- good scalability
- easier to maintain
- DIRAC services are deployed at JINR cloud

Clouds integration (2/2)

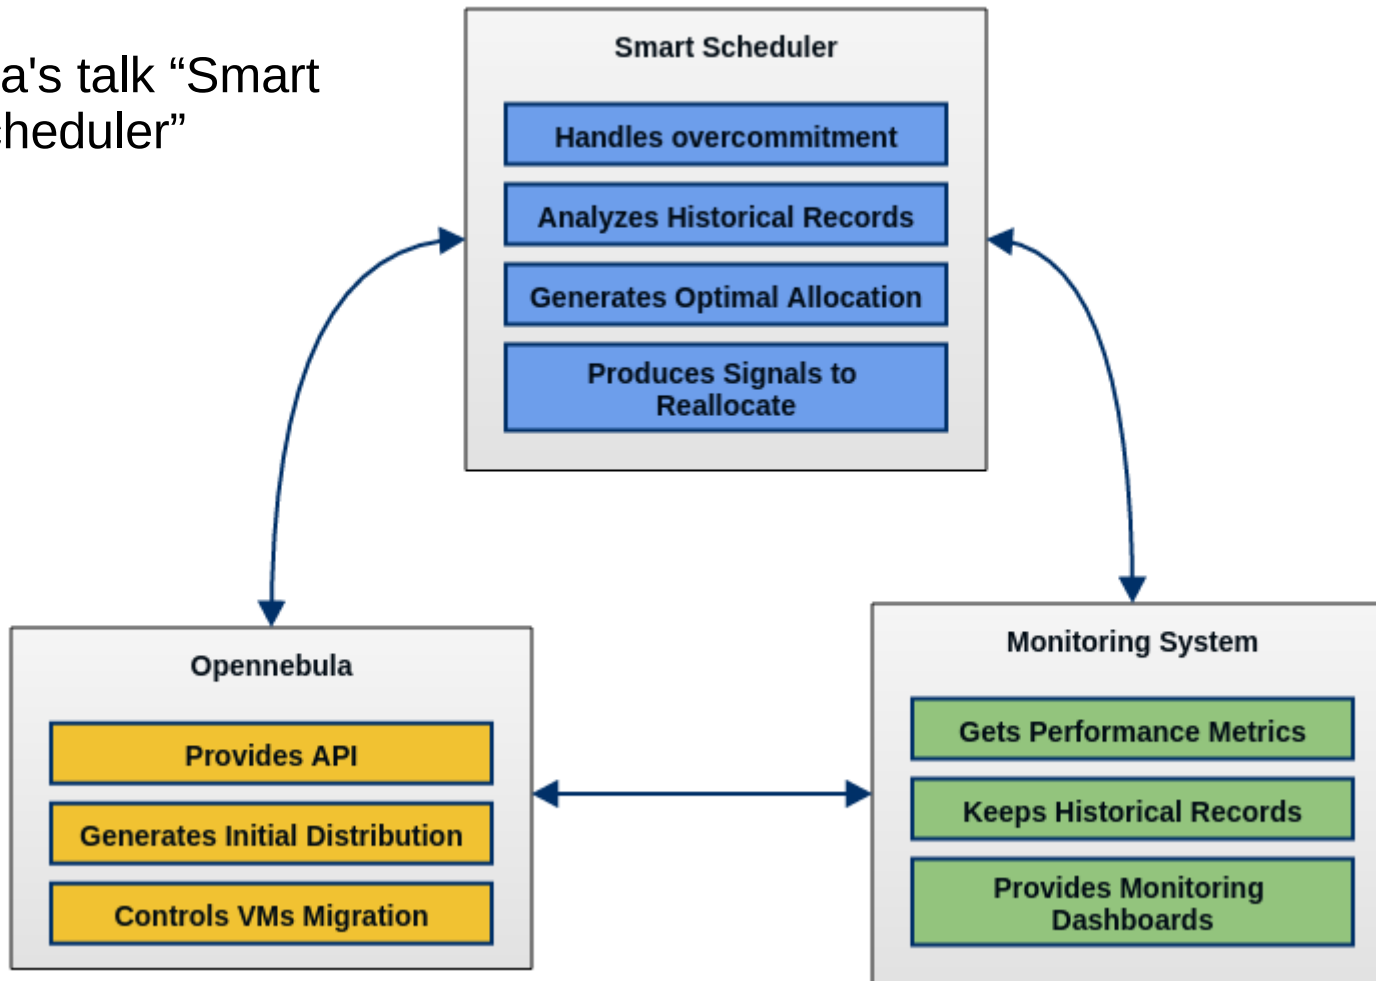
- To join resources for solving common tasks as well as to distribute a peak load across resources of partner organizations



Migration from OCCI to Opennebula XML-RPC endpoint secured with SSL

Smart cloud scheduler

See Nikita's talk "Smart Cloud Scheduler"



The work on the smart cloud scheduler is supported by RFBR grant #15-29-07027

Plans

- Migrate to OpenNebula 5.6
- Switch to Virtuozzo 7
- Put Smart Cloud Scheduler from pre-production into production in JINR cloud
- Adopt a set of experiments' workflow to run into distributed JINR Member State computational (cloud?) infrastructure (DJMSCI)

References

- Web-GUI: <http://cloud.jinr.ru> (authentication is required, accessible from JINR, CERN and Dubna local ISP only)
- Web-portal about JINR cloud infrastructure
 - <http://miccom.jinr.ru> → Cloud service
 - JINR cloud description, quick user and admin guides, contacts, publications, etc
- OpenNebula: <http://opennebula.org>
- Virtualization systems:
 - OpenVZ: <http://openvz.org>
 - KVM: <http://linux-kvm.org>
- Ceph: <http://ceph.com>