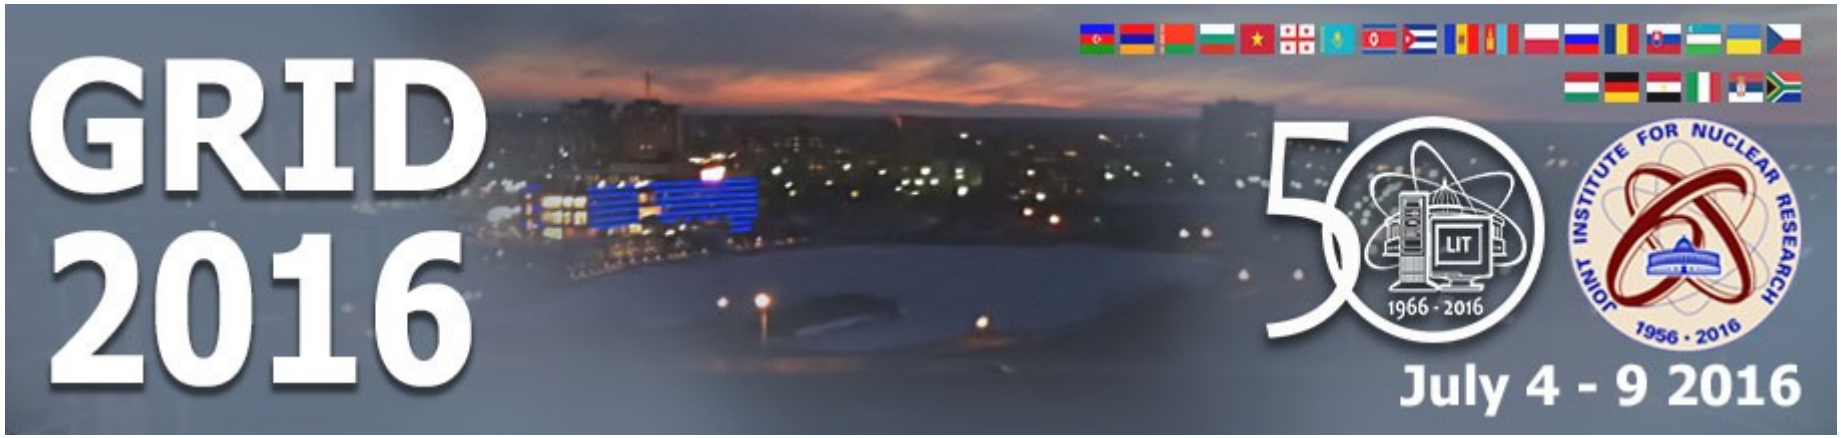

JINR cloud: status report

A. V. Baranov[1], N. A. Balashov[1], N. A. Kutovskiy[1,2], R. N. Semenov[1,2],
K. V. Fedorov[1], A. Kondratyev[1]

[1] Laboratory of Information Technologies, Joint Institute for Nuclear Research
[2] Plekhanov Russian University of Economics

The 7th International Conference "Distributed Computing and Grid-technologies in Science and Education" (GRID'2016)
4-9 July 2016, JINR, Dubna, Russia

JINR cloud architecture

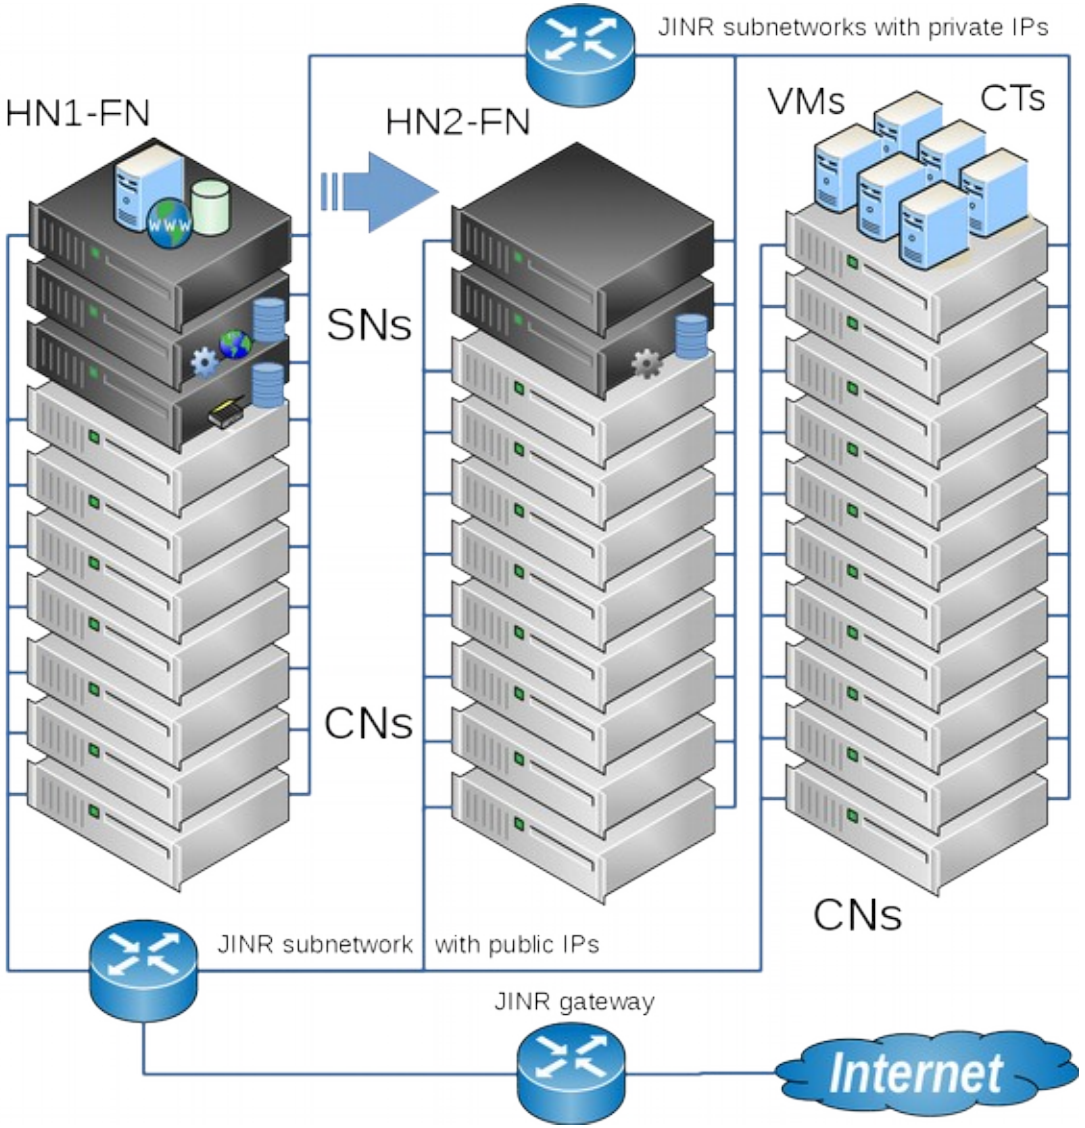

- OpenNebula (v4.12)
  - Core
  - Scheduler
  - MySQL Database
  - Interfaces (web-GUI, CLI, API)
  - OneGate
  - OneFlow
- High availibility&reliability
  - DRDB
  - Heartbeat
  - HN1-FN and HN2-FN are connected to two different UPS
  - DNFS based on LirazdFS

Distributed network file system

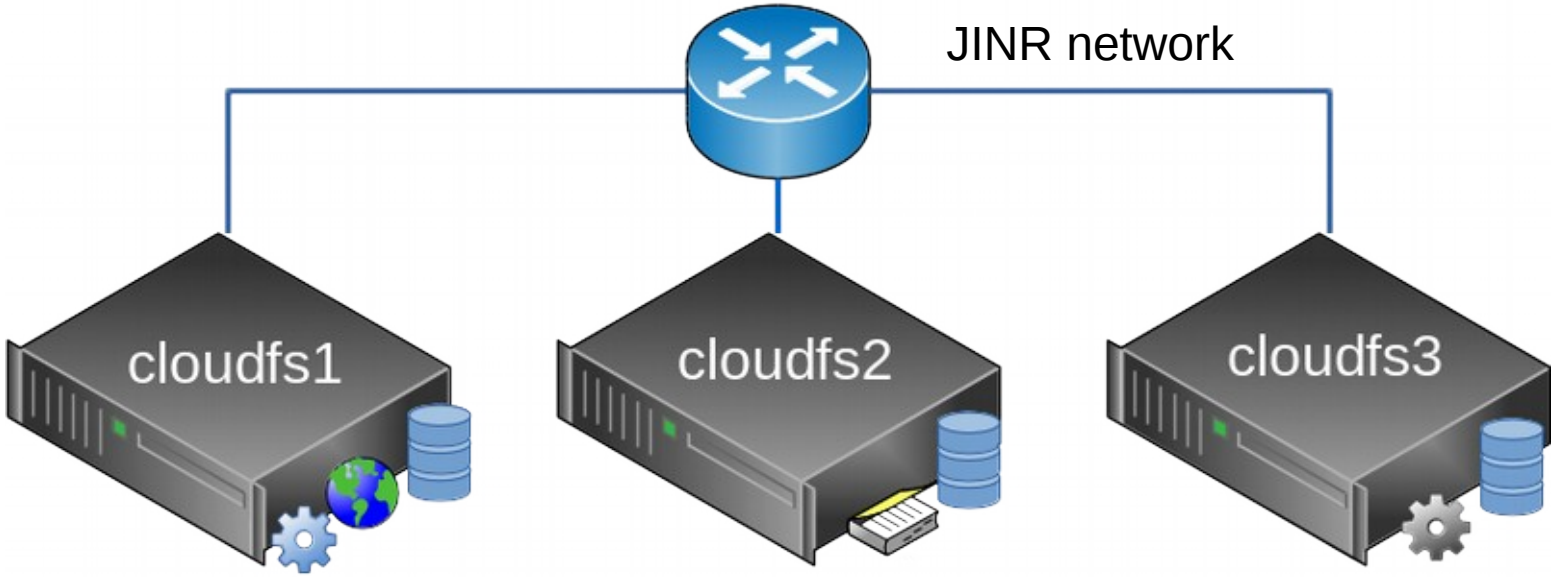

- DNFS is based on LizardFS:
  - SN1: DNFS head node (master), a data server (chunk server) and a graphical web-admin interface (web-based admin GUI);
  - SN2: a duplicate service for logs (metalogger) and a chunk server;
  - SN3: a backup head node DNFS (shadow) and a chunk server.

JINR cloud service monitoring

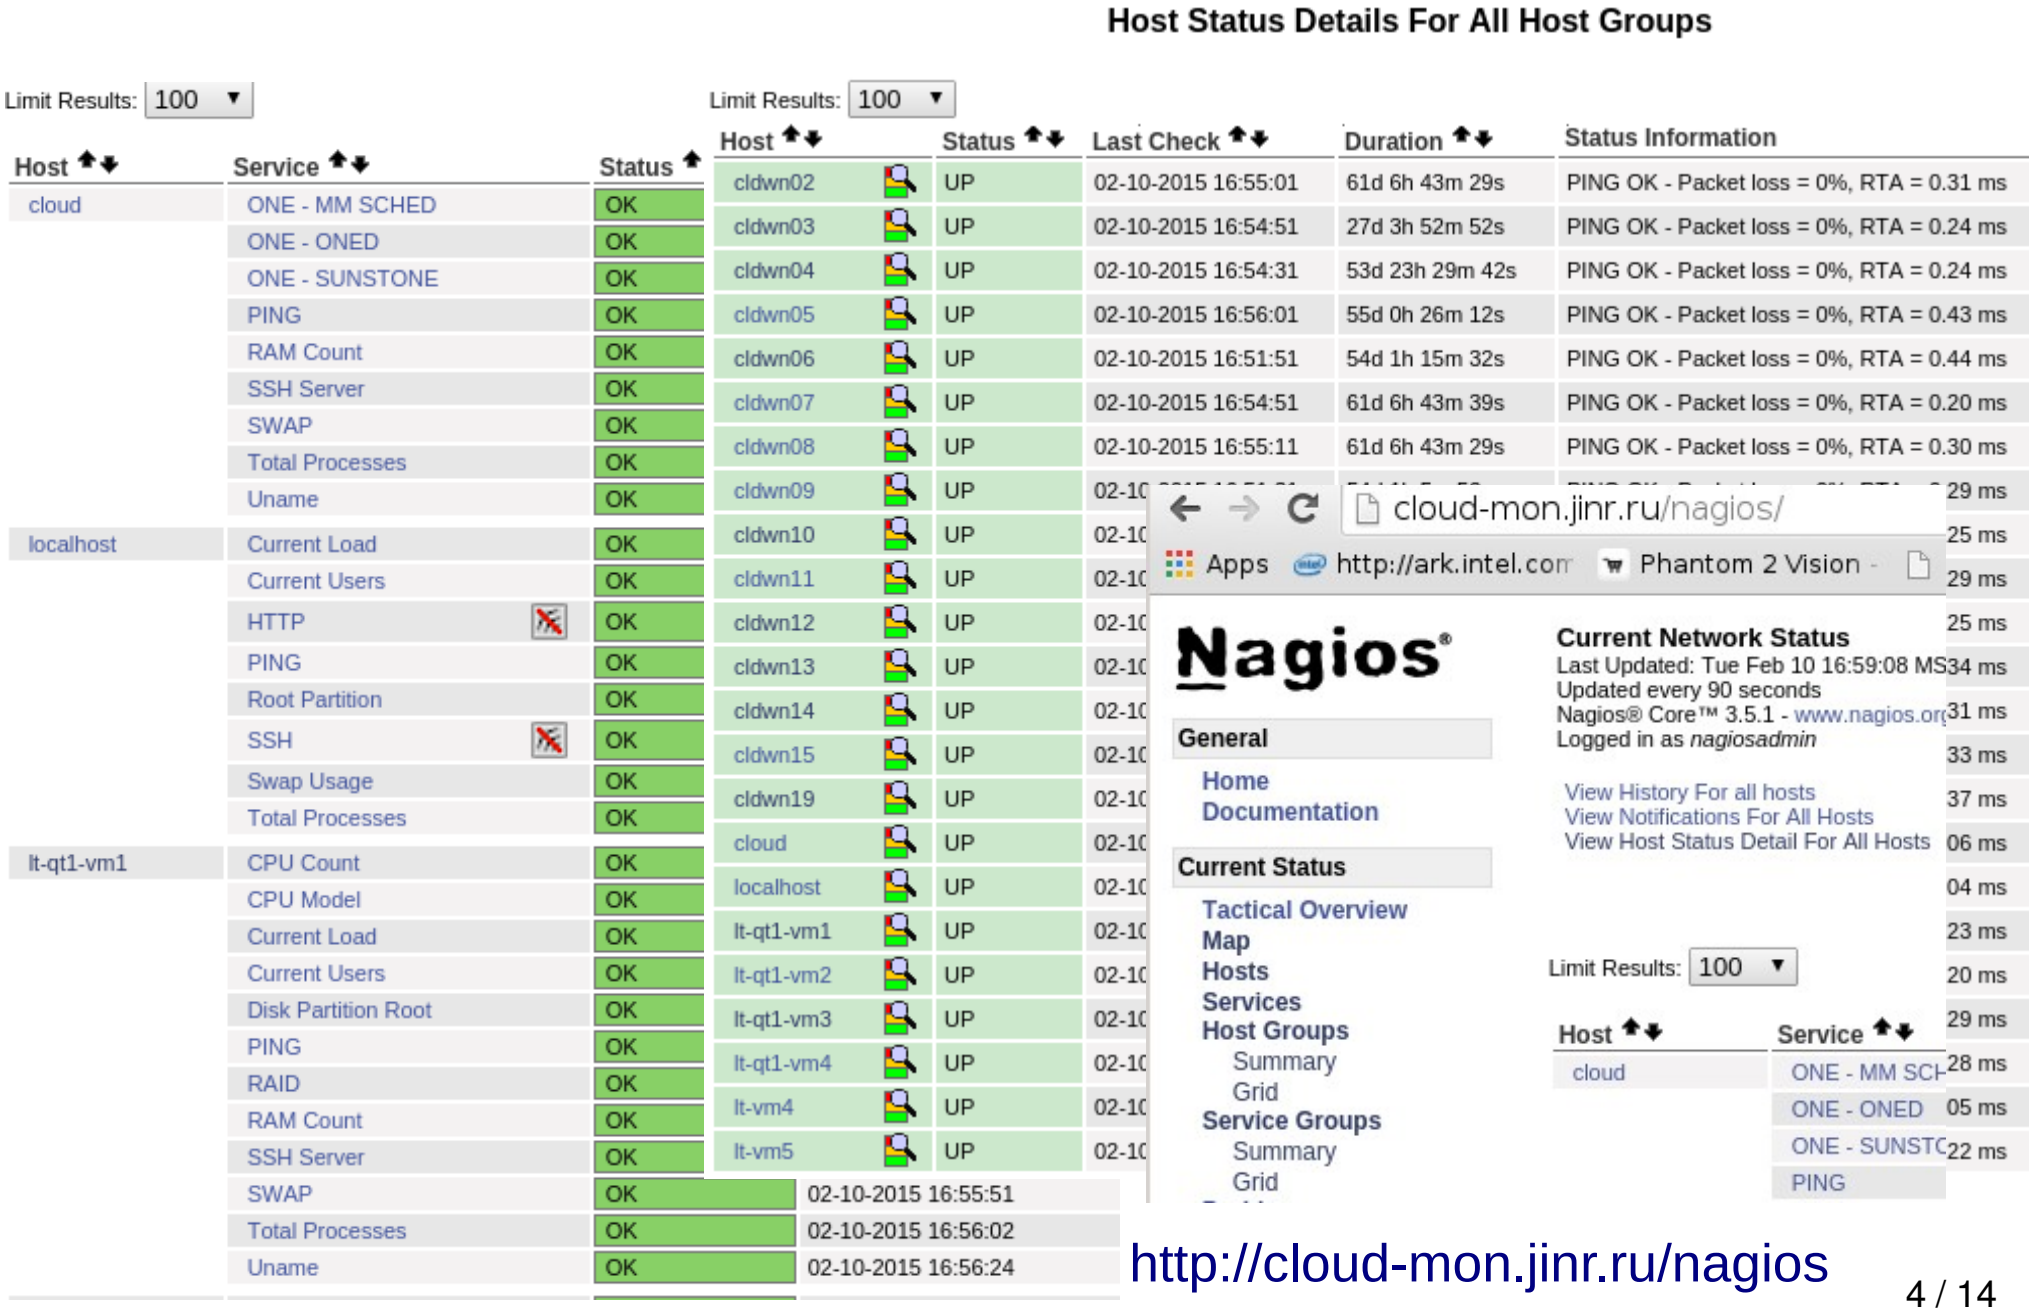

Limit Results: 100

Host	Service	Status
cloud	ONE - MM SCHED	OK
	ONE - ONED	OK
	ONE - SUNSTONE	OK
	PING	OK
	RAM Count	OK
	SSH Server	OK
	SWAP	OK
	Total Processes	OK
	Uname	OK
localhost	Current Load	OK
	Current Users	OK
	HTTP	OK
	PING	OK
	Root Partition	OK
	SSH	OK
	Swap Usage	OK
	Total Processes	OK
lt-qt1-vm1	CPU Count	OK
	CPU Model	OK
	Current Load	OK
	Current Users	OK
	Disk Partition Root	OK
	PING	OK
	RAID	OK
	RAM Count	OK
	SSH Server	OK
	SWAP	OK
	Total Processes	OK
	Uname	OK

02-10-2015 16:55:51
02-10-2015 16:56:02
02-10-2015 16:56:24

Limit Results: 100

Host Status Details For All Host Groups

Host	Status	Last Check	Duration	Status Information
cldwn02	UP	02-10-2015 16:55:01	61d 6h 43m 29s	PING OK - Packet loss = 0%, RTA = 0.31 ms
cldwn03	UP	02-10-2015 16:54:51	27d 3h 52m 52s	PING OK - Packet loss = 0%, RTA = 0.24 ms
cldwn04	UP	02-10-2015 16:54:31	53d 23h 29m 42s	PING OK - Packet loss = 0%, RTA = 0.24 ms
cldwn05	UP	02-10-2015 16:56:01	55d 0h 26m 12s	PING OK - Packet loss = 0%, RTA = 0.43 ms
cldwn06	UP	02-10-2015 16:51:51	54d 1h 15m 32s	PING OK - Packet loss = 0%, RTA = 0.44 ms
cldwn07	UP	02-10-2015 16:54:51	61d 6h 43m 39s	PING OK - Packet loss = 0%, RTA = 0.20 ms
cldwn08	UP	02-10-2015 16:55:11	61d 6h 43m 29s	PING OK - Packet loss = 0%, RTA = 0.30 ms
cldwn09	UP	02-10		29 ms
cldwn10	UP	02-10		25 ms
cldwn11	UP	02-10		29 ms
cldwn12	UP	02-10		25 ms
cldwn13	UP	02-10		34 ms
cldwn14	UP	02-10		31 ms
cldwn15	UP	02-10		33 ms
cldwn19	UP	02-10		37 ms
cloud	UP	02-10		06 ms
localhost	UP	02-10		04 ms
lt-qt1-vm1	UP	02-10		23 ms
lt-qt1-vm2	UP	02-10		20 ms
lt-qt1-vm3	UP	02-10		29 ms
lt-qt1-vm4	UP	02-10		28 ms
lt-vm4	UP	02-10		05 ms
lt-vm5	UP	02-10		22 ms

cloud-mon.jinr.ru/nagios/

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

Nagios®

General
- Home
- Documentation

Current Status
- Tactical Overview
- Map
- Hosts
- Services
- Host Groups
  - Summary
  - Grid
- Service Groups
  - Summary
  - Grid

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 All Hosts
View Host Status Detail For All Hosts

Limit Results: 100

Host	Service
cloud	ONE - MM SCH
	ONE - ONED
	ONE - SUNSTC
	PING

http://cloud-mon.jinr.ru/nagios

Hardware resources

- Current hardware resources
  - ~40 servers (+27 to be connected)
  - ~200 cores (+236),
  - ~400 GB of RAM (+656 GB),
  - ~20 TB of total local disk spaces for VM/CT deployment (+17 TB)
  - ~16 TB of DNFS disk space
- Planned:

	2016	**2017**	**2018**	**2019**
CPU cores	1000	1400	1800	2200
Total RAM, GB	4240	6160	8080	10000
TB of disk space, TB	384	576	768	960

Clusters

- Production cluster
  - Highly reliable services
- Dev cluster
  - Education
  - Development
  - Testing
- Computing cluster
  - Computational resources for users and experiments
  - rOCCI for connecting to experiments' computing infrastructures
    - BESIII
    - NOvA

Usage

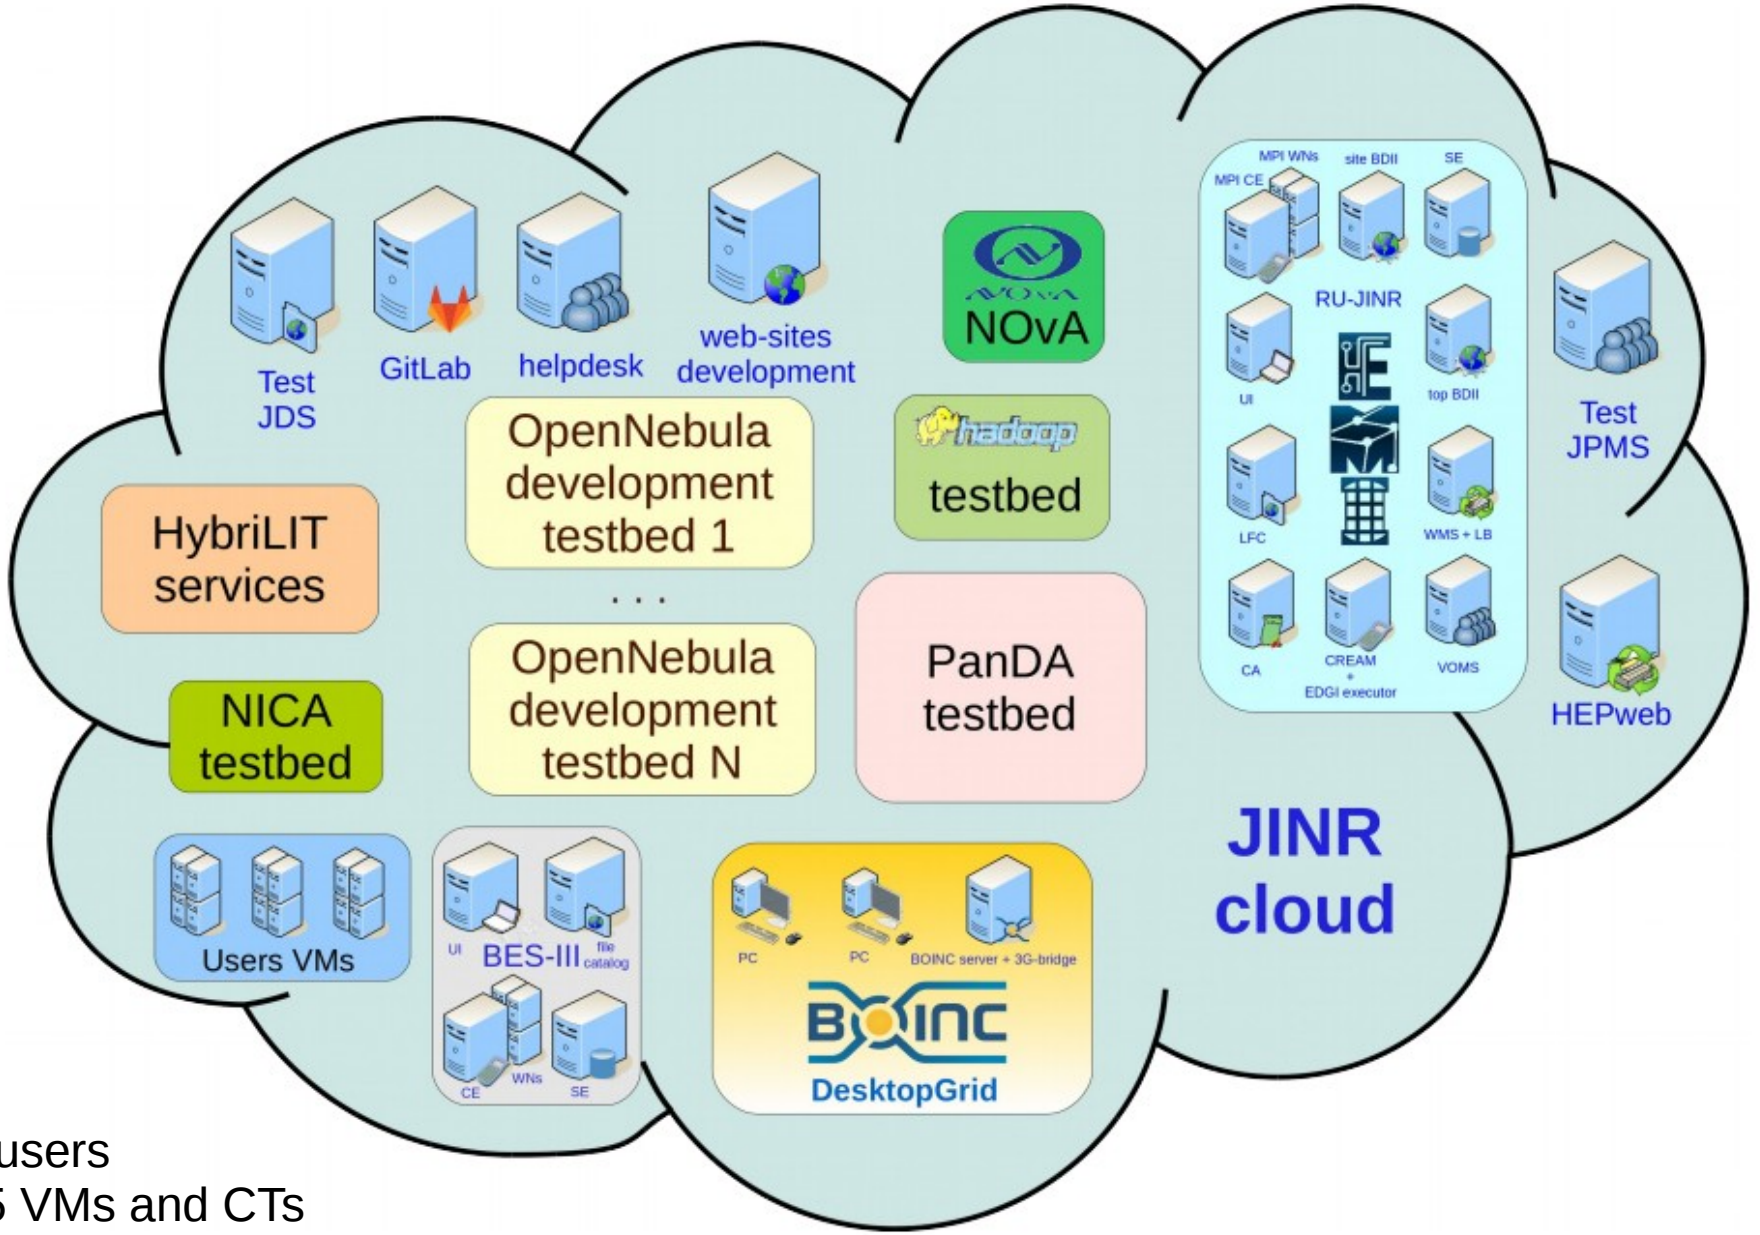

~80 users
~115 VMs and CTs

Clouds integration

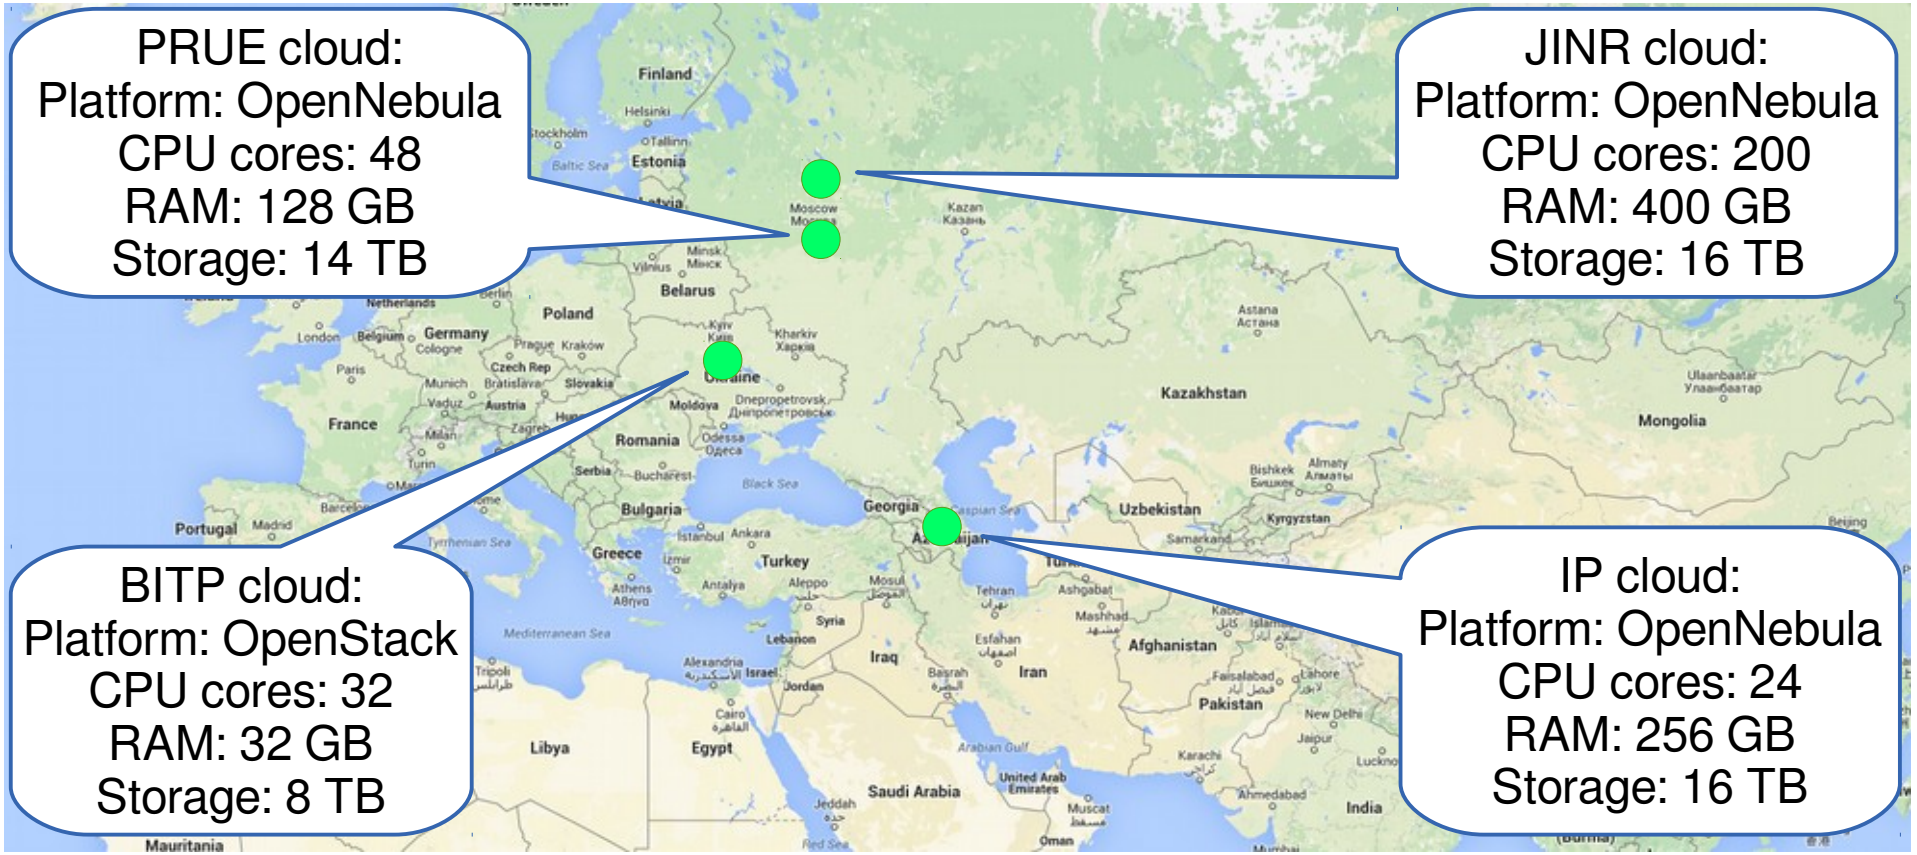

- To join resources for solving common tasks as well as to distribute a peak load across resources of partner organizations
- JINR cloud integration with clouds of partner organizations:
  - Institute of Physics of Azerbaijan National Academy of Sciences – IP (Baku, Azerbaijan)
  - Bogolyubov Institute for Theoretical Physics of the National Academy of Sciences of Ukraine – BITP (Kiev, Ukraine)
  - Plekhanov Russian University of Economics – PRUE (Moscow, Russia)
  - EGI Federated cloud

User and admin trainings

Organization	Organization location	Training dates	Number of trainees	Training type
GRID'2016 school		05.07.16	5	cloud use
Institute of Experimental and Applied Physics, Czech Technical University	Prague, CZ	07-10.07.15	2	cloud use
Egyptian scientific organizations	EG	05-09.06.15	3	cloud use
JINR	Dubna, RU	26-27.01.15	11	cloud use
Gdansk university of technologies	Gdansk, PL	06.10-12.12.14	1	cloud use/admin
NC PHEP BSU	Minsk, BY	22-29.09.14	3	cloud use/admin

Smart cloud scheduler

See Nikita's talk “Smart Cloud Scheduler”

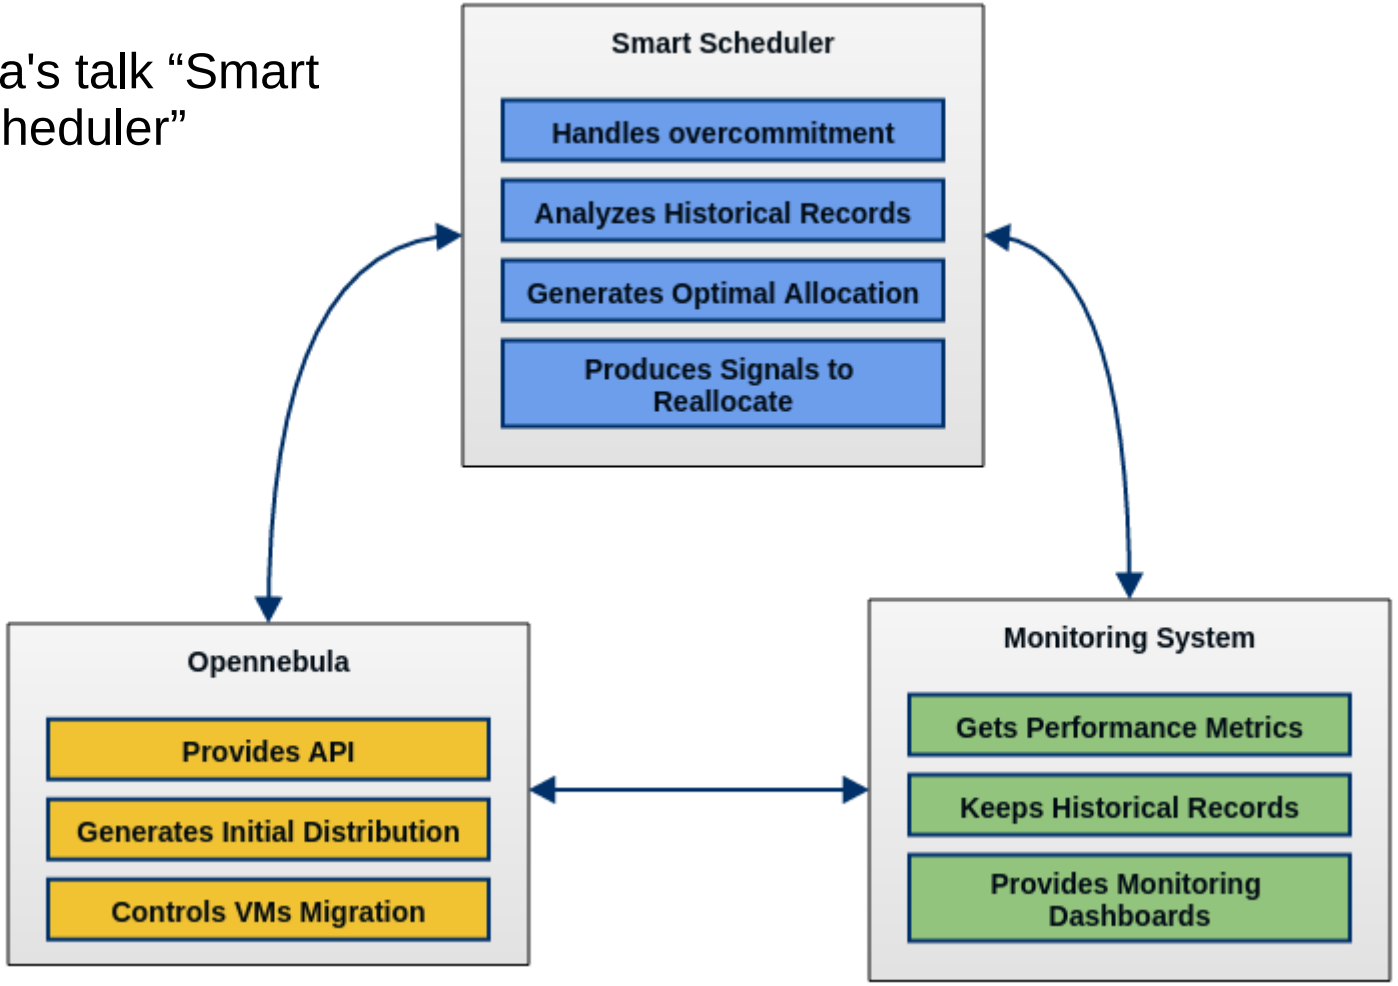

The work on the smart cloud scheduler is supported by RFBR grant #15-29-07027

MPI + CUDA in OpenVZ containers

- One of the modern trends in clouds is to run HPC cluster with GPGPU
  - Access to two Sugon servers were provided:
    - One is in China and the second one is in Moscow testbed of Jet Infosystems company:
      - 1st server config:
        - 2 x Intel(R) Xeon(R) CPU E5-2650 v3 @ 2.30GHz chips (40 cores in total),
        - 128 GB of RAM,
        - 4 x GK110BGL [Tesla K40m]
      - 2nd server config:
        - 2 x Intel(R) Xeon(R) CPU E5-2667 v3 @ 3.20GHz chips (16 cores in total)
        - 500 GB of RAM,
        - 2 x GK110BGL [Tesla K40m]
  - OpenVZ + CUDA (+ MPI): success
  - KVM + CUDA: failure

Plans

- Migrate to OpenNebula 5.0
- Evaluate Virtuozzo 7 and switch to it
- As soon as Smart Cloud Scheduler is ready integrate it into JINR cloud
- As soon as 10Gbps switches arrive and be mounted deploy CEPH storage and migrate to it
- Build HPC segment in the JINR cloud
- JINR cloud as computing backend for HTCondor

JINR cloud team

- Nikita Balashov
  - Custom components for OpenNebula development and support
  - User support and trainings
- Aleksandr Baranov
  - Cloud administration, new components evaluation and testing
  - Cloud users and admins support and trainings
- Nikolay Kutovskiy
  - coordinator
  - User support and trainings
- Roman Semenov
  - Admin, R&D in cloud storages
  - Users support
- Konstantsin Fedorov
  - sysadmin
- Andrey Kondratyev
  - Developer
- Two summer students:
  - Ruslan Gainanov
  - Evgeniy Kotsarev

References

- Web-GUI: http://cloud.jinr.ru (authentication is required, accessible from JINR, CERN and Dubna local ISP only)
- Cloud servers, services, VMs and CTs are monitored with help of Nagios:
  - http://cloud-mon.jinr.ru/nagios (authentication is required)
- 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
- LizardFS: http://lizardfs.com