

Cloud infrastructure at JINR

N. Balashov, A. Baranov, N. Kutovskiy, R. Semenov
Nikolay.Kutovskiy@jinr.ru

LIT JINR

Motivation

- To provide a modern computing facility for JINR users participating in different national and international scientific projects
- To solve local JINR IT tasks and fulfill JINR commitments in various national and international projects related to the usage of IT technologies
- To satisfy the needs of JINR developers (development, testing and debugging various apps in various environments) and system administrators (testing and studying specifics of installation and operation of new apps or verifying updates)
- All these requires computing and storage resources
- To manage facilities for them in more effective and reliable way a special dedicated infrastructure deployed and managed with help of modern tools – **cloud technologies**.

Implementation: virtualization, servers, cloud access

- Based on OpenNebula software
- Two types of virtualization
 - OpenVZ (OS-level virtualization)
 - KVM (full virtualization)
- Two types of cloud nodes
 - servers with two mirrored disk drives for highly reliable VMs
 - servers with a single disk for educational, research or test VMs
- Cloud access
 - Web-GUI - «Sunstone» with authentication via JINR kerberos credentials (login/password)
 - http-traffic with SSL encryption
 - command-line interface via ssh



```
Scientific Linux release 6.6 (Carbon)
Kernel 2.6.32-042stab103.6 on an x86_64

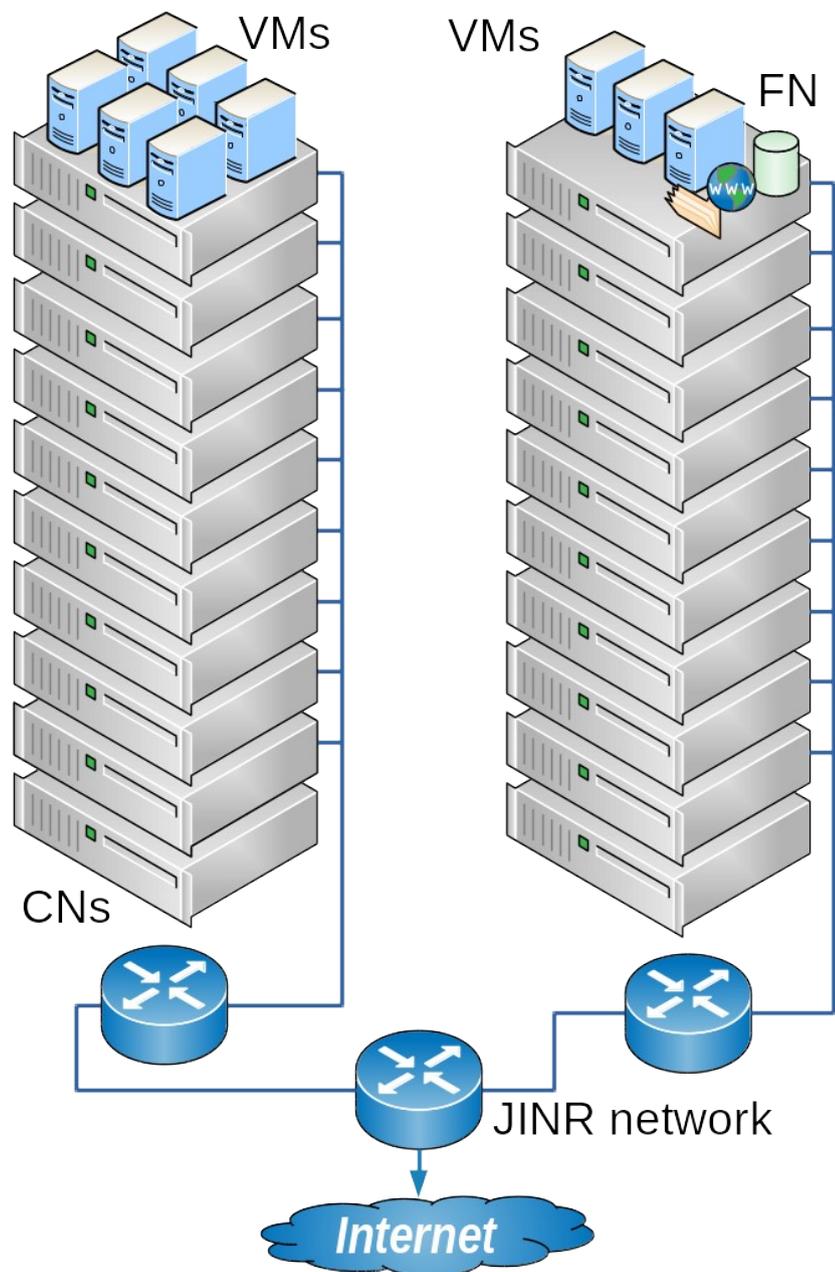
vm168 login: telecast
Password: _
```

A screenshot of the Sunstone web-GUI login page. The header features the JINR logo and the word "Cloud" in a blue box. Below the header, there are two input fields: "Username" with the value "telecast" and "Password" with a masked password. A "Keep me logged in" checkbox is present below the password field. A "Login" button is located to the right of the password field. At the bottom of the page, there is a link for "General information and terms of service" and its Russian equivalent "Общая информация и правила использования сервиса".

Implementation: VMs access, user support, PM

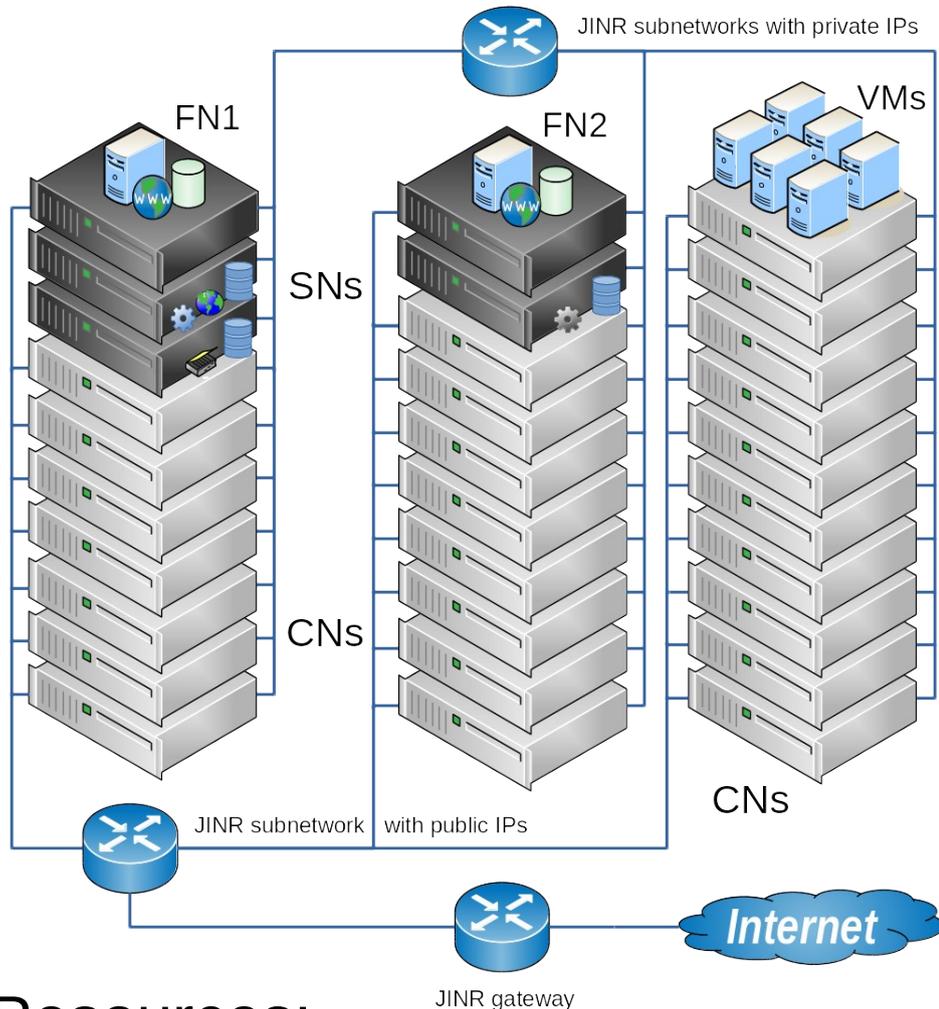
- VMs access
 - via ssh-protocol: root access using {rsa,dsa}-keys or using JINR kerberos credentials and then 'sudo'
 - using VNC in Sunstone with JINR kerberos credentials
- Cloud web-GUI URL: <http://cloud.jinr.ru>
- User support via JINR helpdesk service: <http://helpdesk.jinr.ru>
- JINR cloud project management via JINR project management service (JPMS): <http://pm.jinr.ru>

Implementation: simple setup



- OpenNebula front-end (FN)
 - deployed on single VM
 - core OpenNebula service
 - OpenNebula scheduler
 - MySQL backend
 - User (Web-GUI, CLI) and API (rOCCI) interfaces
- All VM images and data are stored on FN disk
- Resources:
 - Total # of cores: 122
 - Total RAM capacity: 252 GB
 - Datastores size: 500 GB
- Disadvantages:
 - FN is single point of failure
 - Impossible to extend FN disk space «on the fly»

Implementation: HA-setup



Resources:

CPU: + 80 cores

RAM: + 160 GB of RAM

Datastore size: 16TB on DNFS

- OpenNebula front-end deployed on two VM configured in high available (HA) setup
- Each HA-FN has
 - core OpenNebula service
 - OpenNebula scheduler
 - MySQL backend
 - User (Web-GUI, CLI) and API (rOCCI) interfaces
- All VM images and data are stored on distributed network file system (DNFS) based on LizardFS with automatic data replication

JINR cloud service: monitoring

Host Status Details For All Host Groups

Limit Results:

Host	Service	Status	Host	Status	Last Check	Duration	Status Information	
cloud	ONE - MM SCHED	OK	cldwn02	UP	02-10-2015 16:55:01	61d 6h 43m 29s	PING OK - Packet loss = 0%, RTA = 0.31 ms	
	ONE - ONED	OK	cldwn03	UP	02-10-2015 16:54:51	27d 3h 52m 52s	PING OK - Packet loss = 0%, RTA = 0.24 ms	
	ONE - SUNSTONE	OK	cldwn04	UP	02-10-2015 16:54:31	53d 23h 29m 42s	PING OK - Packet loss = 0%, RTA = 0.24 ms	
	PING	OK	cldwn05	UP	02-10-2015 16:56:01	55d 0h 26m 12s	PING OK - Packet loss = 0%, RTA = 0.43 ms	
	RAM Count	OK	cldwn06	UP	02-10-2015 16:51:51	54d 1h 15m 32s	PING OK - Packet loss = 0%, RTA = 0.44 ms	
	SSH Server	OK	cldwn07	UP	02-10-2015 16:54:51	61d 6h 43m 39s	PING OK - Packet loss = 0%, RTA = 0.20 ms	
	SWAP	OK	cldwn08	UP	02-10-2015 16:55:11	61d 6h 43m 29s	PING OK - Packet loss = 0%, RTA = 0.30 ms	
	Total Processes	OK	cldwn09	UP	02-10-2015 16:55:11	61d 6h 43m 29s	PING OK - Packet loss = 0%, RTA = 0.29 ms	
	Uname	OK	cldwn09	UP	02-10-2015 16:55:11	61d 6h 43m 29s	PING OK - Packet loss = 0%, RTA = 0.29 ms	
	localhost	Current Load	OK	cldwn10	UP	02-10-2015 16:55:11	61d 6h 43m 29s	PING OK - Packet loss = 0%, RTA = 0.25 ms
Current Users		OK	cldwn11	UP	02-10-2015 16:55:11	61d 6h 43m 29s	PING OK - Packet loss = 0%, RTA = 0.29 ms	
HTTP		OK	cldwn12	UP	02-10-2015 16:55:11	61d 6h 43m 29s	PING OK - Packet loss = 0%, RTA = 0.25 ms	
PING		OK	cldwn13	UP	02-10-2015 16:55:11	61d 6h 43m 29s	PING OK - Packet loss = 0%, RTA = 0.31 ms	
Root Partition		OK	cldwn14	UP	02-10-2015 16:55:11	61d 6h 43m 29s	PING OK - Packet loss = 0%, RTA = 0.31 ms	
SSH		OK	cldwn15	UP	02-10-2015 16:55:11	61d 6h 43m 29s	PING OK - Packet loss = 0%, RTA = 0.33 ms	
Swap Usage		OK	cldwn19	UP	02-10-2015 16:55:11	61d 6h 43m 29s	PING OK - Packet loss = 0%, RTA = 0.37 ms	
Total Processes		OK	cldwn19	UP	02-10-2015 16:55:11	61d 6h 43m 29s	PING OK - Packet loss = 0%, RTA = 0.06 ms	
lt-qt1-vm1		CPU Count	OK	cloud	UP	02-10-2015 16:55:11	61d 6h 43m 29s	PING OK - Packet loss = 0%, RTA = 0.04 ms
		CPU Model	OK	localhost	UP	02-10-2015 16:55:11	61d 6h 43m 29s	PING OK - Packet loss = 0%, RTA = 0.23 ms
	Current Load	OK	lt-qt1-vm1	UP	02-10-2015 16:55:11	61d 6h 43m 29s	PING OK - Packet loss = 0%, RTA = 0.20 ms	
	Current Users	OK	lt-qt1-vm2	UP	02-10-2015 16:55:11	61d 6h 43m 29s	PING OK - Packet loss = 0%, RTA = 0.29 ms	
	Disk Partition Root	OK	lt-qt1-vm3	UP	02-10-2015 16:55:11	61d 6h 43m 29s	PING OK - Packet loss = 0%, RTA = 0.28 ms	
	PING	OK	lt-qt1-vm4	UP	02-10-2015 16:55:11	61d 6h 43m 29s	PING OK - Packet loss = 0%, RTA = 0.05 ms	
	RAID	OK	lt-vm4	UP	02-10-2015 16:55:11	61d 6h 43m 29s	PING OK - Packet loss = 0%, RTA = 0.22 ms	
	RAM Count	OK	lt-vm4	UP	02-10-2015 16:55:11	61d 6h 43m 29s	PING OK - Packet loss = 0%, RTA = 0.22 ms	
	SSH Server	OK	lt-vm5	UP	02-10-2015 16:55:11	61d 6h 43m 29s	PING OK - Packet loss = 0%, RTA = 0.22 ms	
	SWAP	OK	lt-vm5	UP	02-10-2015 16:55:11	61d 6h 43m 29s	PING OK - Packet loss = 0%, RTA = 0.22 ms	
Total Processes	OK	lt-vm5	UP	02-10-2015 16:56:02	61d 6h 43m 29s	PING OK - Packet loss = 0%, RTA = 0.22 ms		
Uname	OK	lt-vm5	UP	02-10-2015 16:56:24	61d 6h 43m 29s	PING OK - Packet loss = 0%, RTA = 0.22 ms		

cloud-mon.jinr.ru/nagios/

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

Nagios®

Current Network Status
 Last Updated: Tue Feb 10 16:59:08 MS34 ms
 Updated every 90 seconds
 Nagios® Core™ 3.5.1 - www.nagios.org
 Logged in as nagiosadmin

General

- Home
- Documentation

Current Status

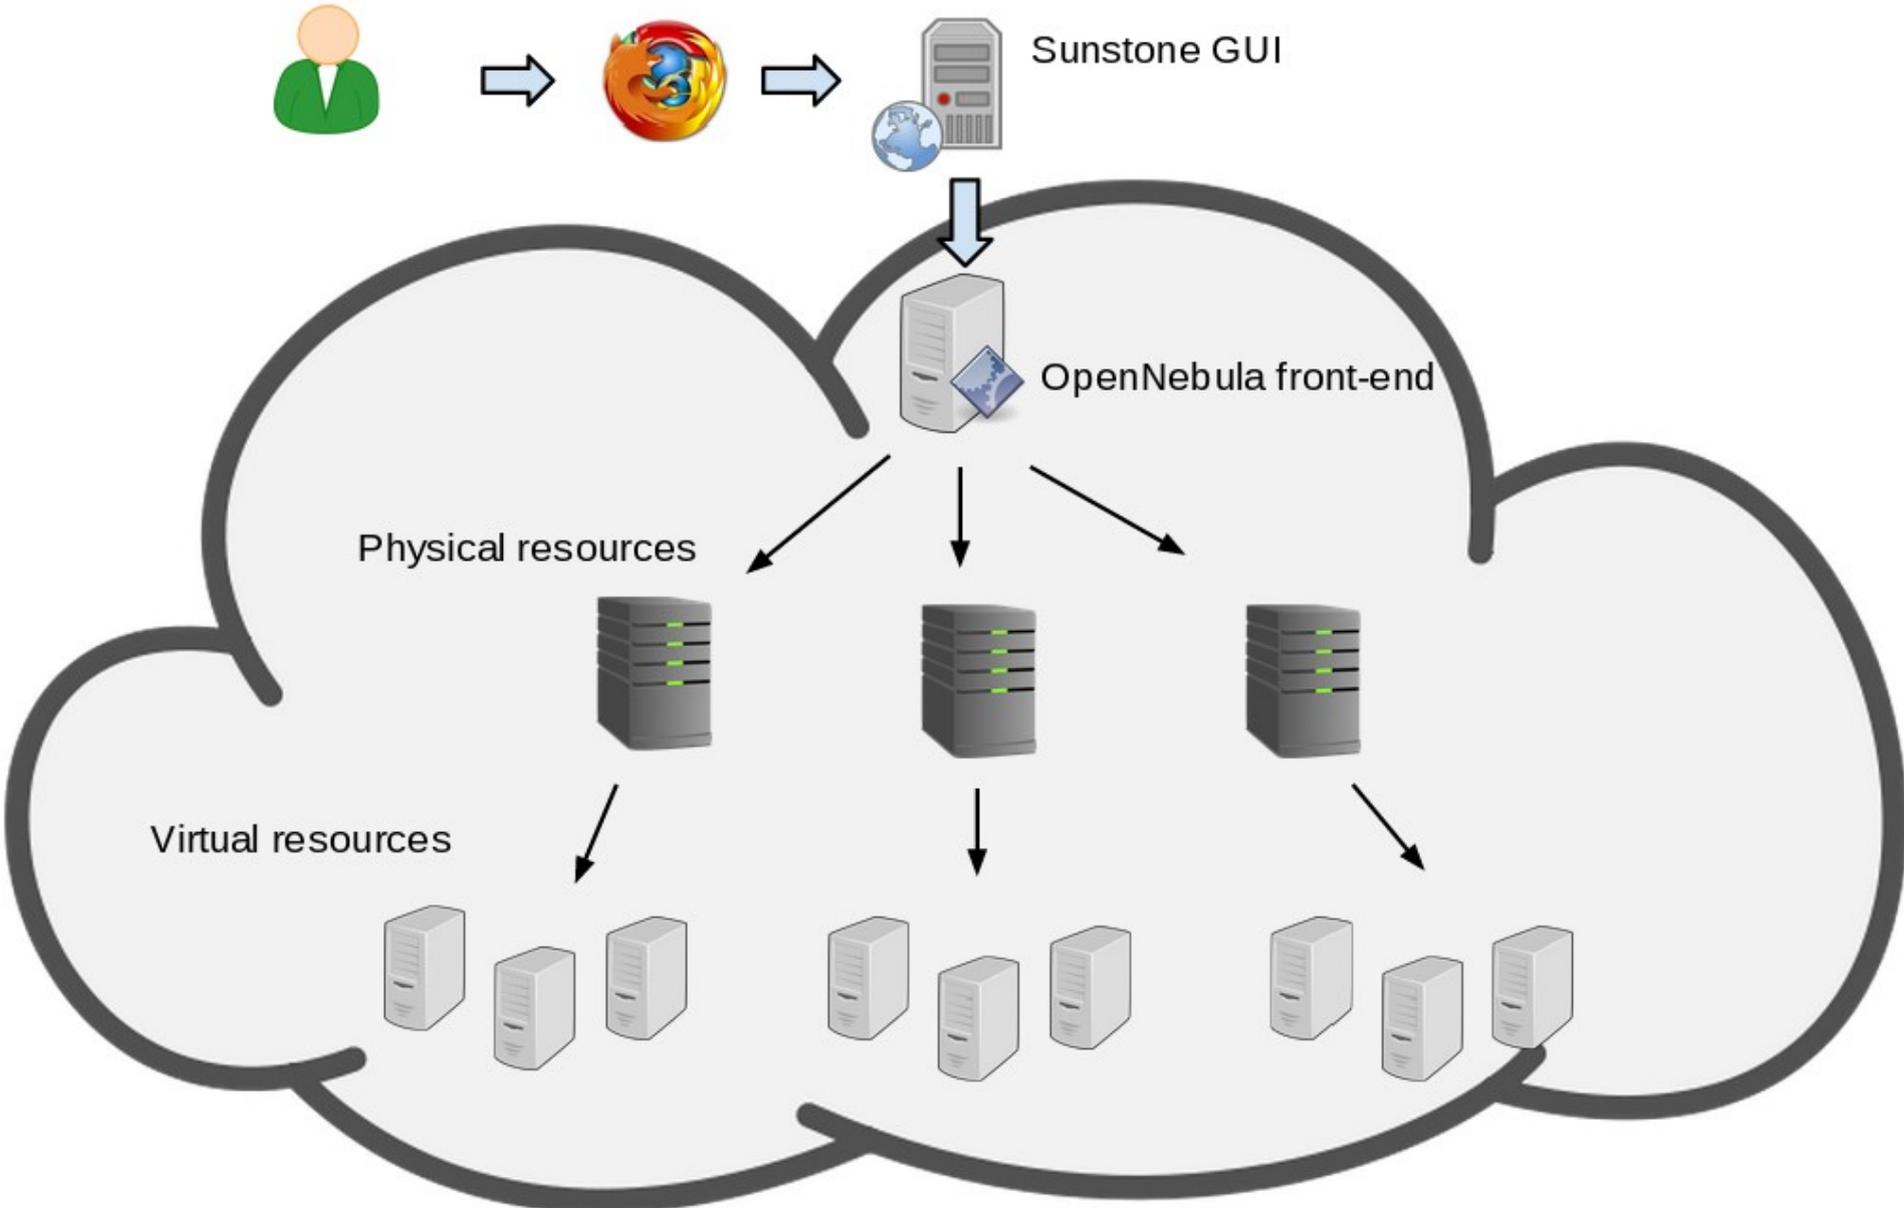
- Tactical Overview
- Map
- Hosts
- Services
- Host Groups
- Service Groups

Limit Results:

Host	Service	Duration
cloud	ONE - MM SCHED	28 ms
	ONE - ONED	05 ms
	ONE - SUNSTONE	22 ms
	PING	

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

Workflow



Web-form for resources request

OpenNebula

 Request resources

 Dashboard

 Virtual Resources

 Infrastructure

 Marketplace

 OneFlow

 Request resources

 Statistics

 Send

General Information

Full name

E-mail

Manager's full name

Manager's e-mail

Laboratory: VBLHEP

Topic number: No topic

Details on the listed topics see [here](#)

Required Resources

CPUs (cores), total per all VMs

RAM (GB), total per all VMs

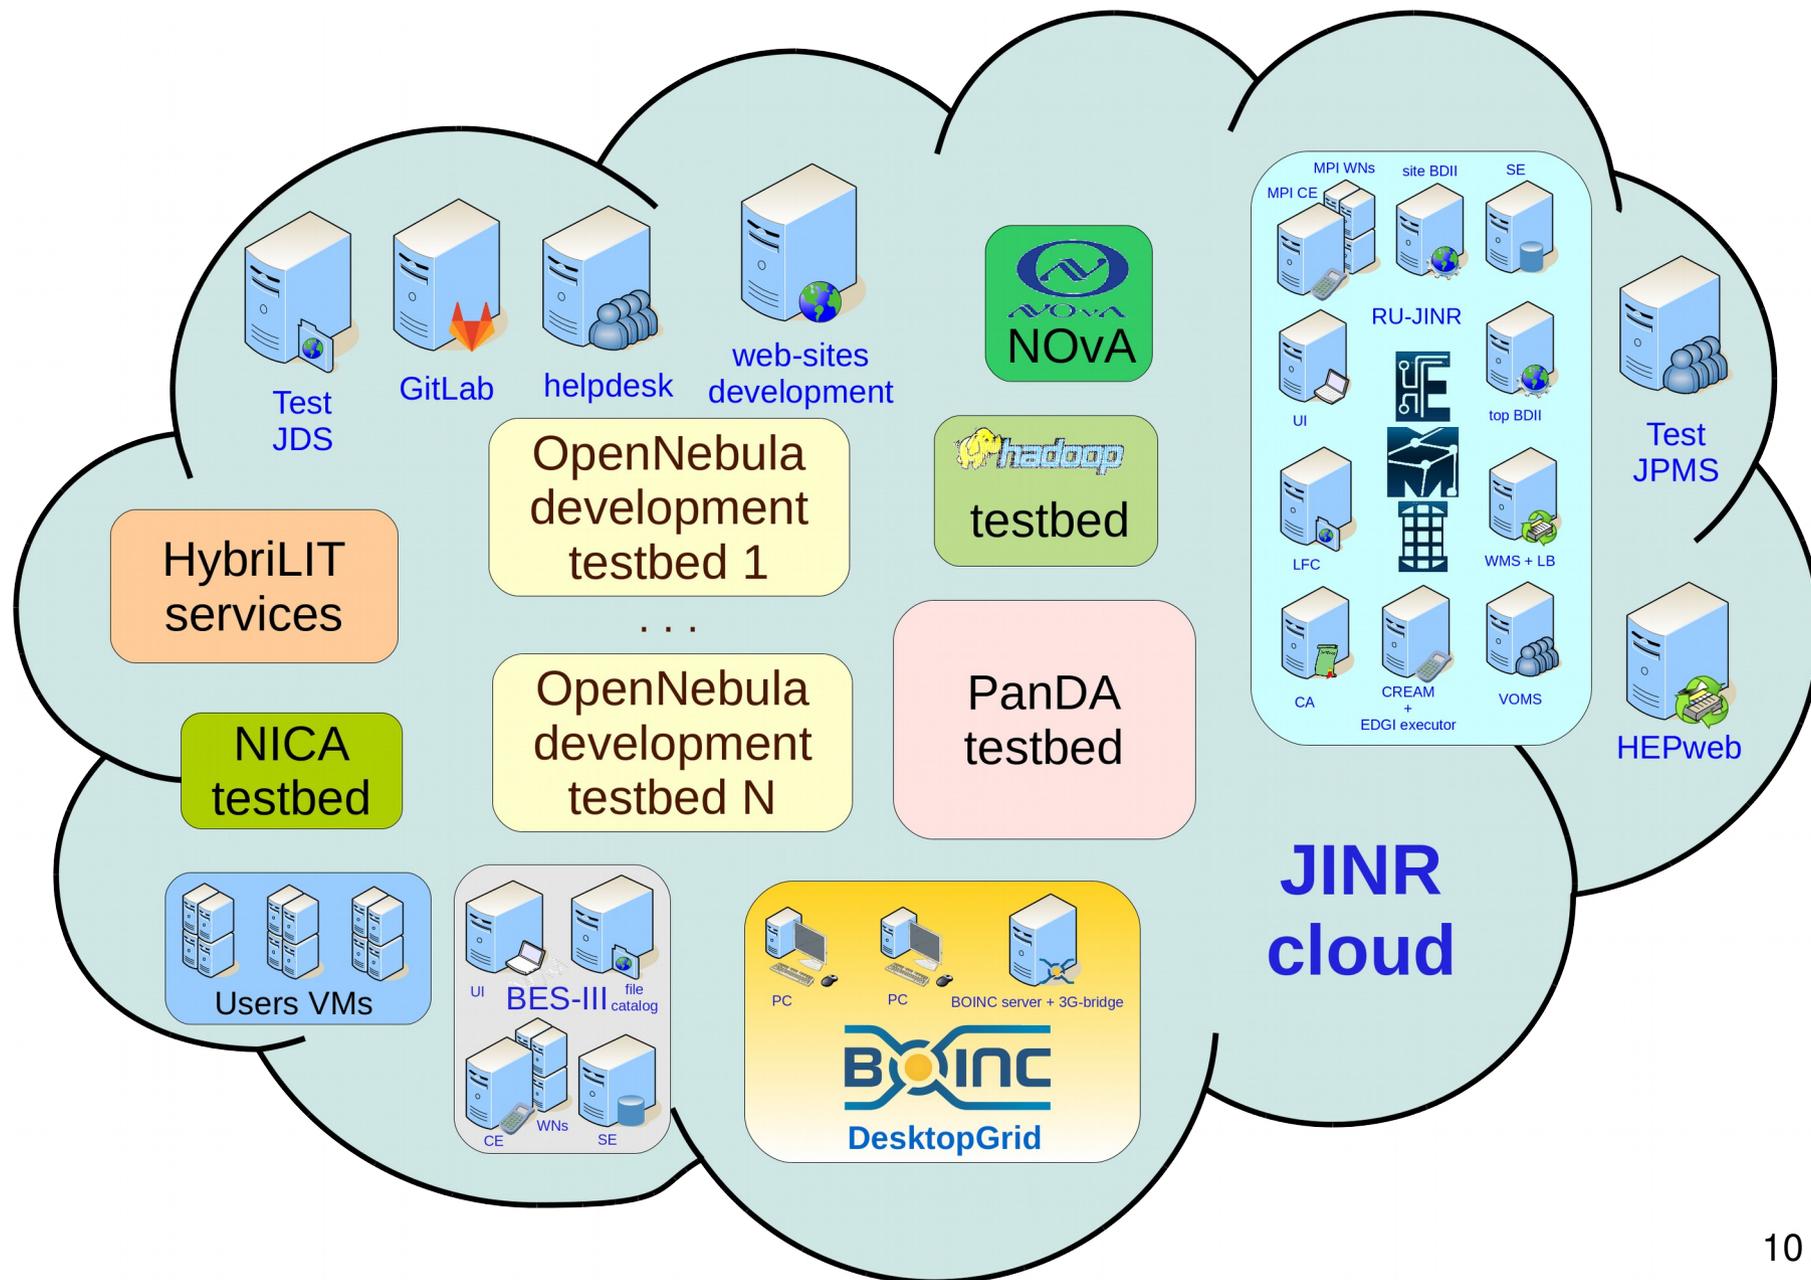
Storage (GB), total per all VMs

Number of virtual machines

OS type: Linux

Comment (purpose of the requested resources or reason for quotas change)

Utilization (1/2)



Utilization (2/2)

- Training and testing grid infrastructure's testbeds:
 - EMI-based testbed
 - ATLAS T3MON + PanDA testbed (monitoring tools development for ATLAS Tier-3 sites, PanDA software development for distributed analysis)
 - DIRAC-based testbed for BES-III experiment (monitoring tools development for BES-III distributed computing infrastructure)
 - DesktopGrid testbed (to estimate the volunteer' computing technology for possible use in solving JINR users' tasks)
- web-service HEPWEB (provides a possibility to use different tools for Monte-Carlo simulation in high-energy physics)
- Helpdesk service (both production and testing/development instances)
- test instances of the JINR document server (JDS), JINR Project Management Service (JPMS)
- Testbeds for OpenVZ driver for OpenNebula:
 - 1 OpenVZ-based VM for opennebula front-end and 2-3 KVM-based VMs as cloud worker nodes with OpenVZ kernel inside

Statistics (1/3)

oneadmin OpenNebula

OpenNebula

Statistics

- Dashboard
- System
- Virtual Resources
- Infrastructure
- Marketplace
- OneFlow

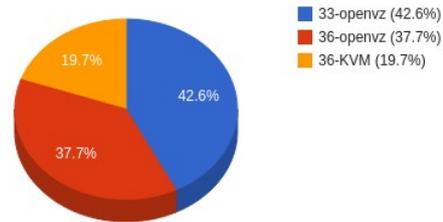
Support
Not connected

Sign in

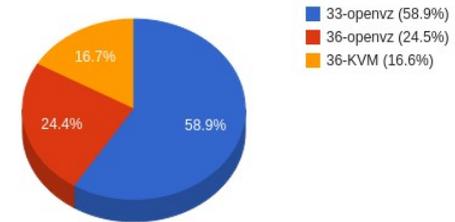
Statistics

Resources distribution over clusters (graphical view)

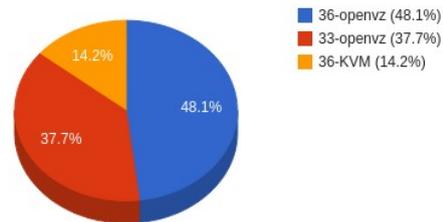
CPU distribution



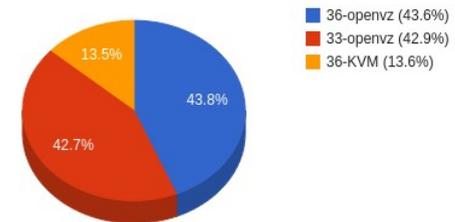
Memory distribution



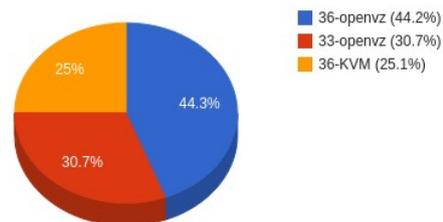
Allocated CPU distribution



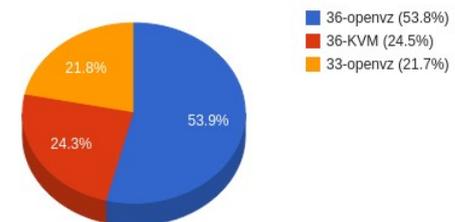
Allocated Memory distribution



Distribution of CPU utilization



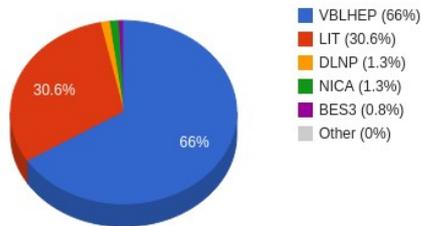
Distribution of Memory utilization



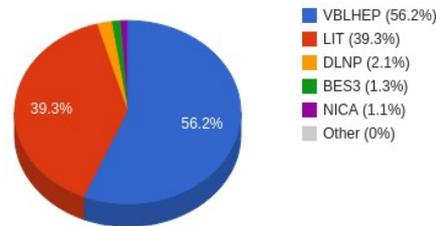
Statistics (2/3)

Resources usage by department (graphical view)

CPU usage by department, core * hours



Memory usage by department, GB * hours



Resources distribution over users

Resources usage by department

Start date: End date:

Lab Name	CPU, Core * hours	RAM, GB * hours
BES3	12275	23243
DLNP	20291	39062
LIT	484089	730749
NICA	20027	20027
TMPK	50.0	50.0
VBLHEP	1044161	1044161

Showing 1 to 6 of 6 entries

Previous **1**

ID	Name	VMs, pcs			CPU, cores			Memory, GB			HDD, GB		
		Used	Available	Allocated	Used	Available	Allocated	Used	Available	Allocated	Used	Available	Allocated
26	evkuz	2	2	4	3	2	5	5	3	8	54.3	115	170
27	alex	0	∞	-1	0	∞	-1	0	∞	-1	-	-	-
29	jemtchou	1	2	3	1	2	3	1	2	3	0.000	60.000	60.000
30	kras	0	∞	-1	0	∞	-1	0	∞	-1	-	-	-
31	svm	0	∞	-1	0	∞	-1	0	∞	-1	-	-	-
32	prosvetov	0	∞	-1	0	∞	-1	0	∞	-1	-	-	-
33	prosvet	0	∞	-1	0	∞	-1	0	∞	-1	-	-	-
34	butorov	0	∞	-1	0	∞	-1	0	∞	-1	-	-	-
35	slepov	1	0	1	4	0	4	4	0	4	0.000	50.000	50.000
36	fira	0	∞	-1	0	∞	-1	0	∞	-1	-	-	-
37	zmey	1	0	1	1	0	1	1	0	1	0.202	9.80	10.00
38	balashov	0	∞	-1	0	∞	-1	0	∞	-1	-	-	-
39	golosc	0	∞	-1	0	∞	-1	0	∞	-1	-	-	-
42	lpelevan	1	1	2	1	1	2	2	1.00	3.00	2.80	17.2	20.000
43	zaporozh	0	∞	-1	0	∞	-1	0	∞	-1	-	-	-
44	miramir	0	∞	-1	0	∞	-1	0	∞	-1	-	-	-
45	abulato	0	2	2	0	2	2	0	2	2	0.000	80.000	80.000
46	bes3	1	4	5	1	4	5	2	8	10	1.13	199	200.000
49	abv	0	∞	-1	0	∞	-1	0	∞	-1	-	-	-
51	roccl	0	∞	-1	0	∞	-1	0	∞	-1	-	-	-
52	dmitryb	1	4	5	1	4	5	1	9	10	0.000	50.000	50.000
53	mvala	3	∞	-1	7	∞	-1	7	∞	-1	2.13	∞	∞
54	strel	1	1	2	2	0	2	2	0	2	0.678	19.3	20.000
55	tmpk	1	0	1	2	0	2	2	0	2	0.000	40.000	40.000
56	virthead	0	5	5	0	10	10	0	20	20	0.000	50.000	50.000
Sum:					111	108	188	144	201	314	234	3275	3380
Unallocated:													
Total:													

Showing 21 to 45 of 72 entries

Previous **1** **2** **3** Next

Statistics (3/3)

Resources usage by department

Start date

01/01/2014

End date

11/10/2015

Get Accounting

Lab Name	CPU, Core * hours	RAM, GB * hours
BES3	12275	23243
DLNP	20291	39062
LIT	484089	730749
NICA	20027	20027
TMPK	50.0	50.0
VBLHEP	1044161	1044161

Showing 1 to 6 of 6 entries

Previous

1

Next

10



User and admin trainings

Organization	Organization location	Training dates	Number of trainees	Training type
Institute of Experimental and Applied Physics, Czech Technical University	Prague, Czech Republic	07-10.07.2015	2	usage
Egyptian scientific organizations	Egypt	05-09.06.2015	3	usage
JINR	Dubna, Russia	26-27.01.2015	11	usage
Gdansk university of technologies	Gdansk, Poland	06.10-12.12.2014	1	usage and administration
National Scientific and Educational Centre of Particle and High Energy Physics of the Belarusian State University	Minsk, Belarus	22-29.09.2014	3	usage and administration

Integration with other clouds

- 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)
 - and Plekhanov Russian University of Economics – PRUE (Moscow, Russia)
 - EGI Federated cloud

JINR cloud team

- **Nikita Balashov**
 - OpenVZ driver for OpenNebula development and support
 - User support
 - documentation
- **Aleksandr Baranov**
 - cloud admin, OpenVZ driver tester
 - cloud users and admins support
 - documentation
- **Nikolay Kutovskiy**
 - coordinator
 - user support
 - documentation
- **Roman Semenov**
 - admin, R&D in cloud storages, cloud storage admin
 - users support
 - documentation

Thank you for your attention!