

SPD Offline Computing Environment

Artem Petrosyan, MLIT, JINR
July 3, 2025

Introduction

- The SPD (Spin Physics Detector) experiment at the NICA collider at JINR is evolving as an international collaboration of physicists from various institutes working together to achieve their stated goals in the field of studying the spin structure of nucleons
- In order to provide the participants of the experiment, both inside and outside JINR, equal opportunities in the field of working with data and computing, it is necessary to create a distributed computing environment of the experiment with the following characteristics: a single entry point for users, data management mechanisms, equitable distribution of available computing resources, security, scalability

Current list of participants

Participants (MOUs signed)

 <p>Joint Institute for Nuclear Research (JINR)  Dubna, Russia  A. Guskov, V. Ladygin</p>	 <p>National Research Nuclear University MEPhI  Moscow, Russia  P. Teterin</p>
 <p>Budker Institute of Nuclear Physics of the Russian Academy of Sciences  Novosibirsk, Russia  A. Barnyakov</p>	 <p>Institute for Nuclear Problems of BSU  Minsk, Belarus  A. Lobko</p>
 <p>Tomsk State University  Tomsk, Russia  S.Filimonov, I. Shreyber</p>	 <p>Skobeltsyn Institute of Nuclear Physics of the Moscow State University  Moscow, Russia  A. Berezhnoy</p>
 <p>Samara National Research University  Samara, Russia  V. Saleev</p>	 <p>Petersburg Nuclear Physics Institute (NRC KI – PNPI)  Gatchina, Russia  V. Kim</p>
 <p>Peter the Great St. Petersburg Polytechnic University (SPbPU)  St. Petersburg, Russia  Ya. Berdnikov</p>	 <p>National Science Laboratory  Yerevan, Armenia  N. Ivanov</p>
 <p>University of Belgrade  Belgrade, Serbia  D. Maletic</p>	 <p>Lebedev Physical Institute of the Russian Academy of Sciences  Moscow, Russia  V. Andreev</p>
 <p>Institute of Nuclear Physics  Almaty, Kazakhstan  S. Sakhiyev</p>	 <p>Belgorod National Research University  Belgorod, Russia  A. Kubankin</p>
 <p>Institute for Nuclear Research RAS  Troitsk, Russia  E. Usenko</p>	 <p>St. Petersburg State University  St. Petersburg, Russia  V. Vechernin</p>
 <p>iThemba LABS  Somerset West, South Africa  N. Stodart</p>	

Participants

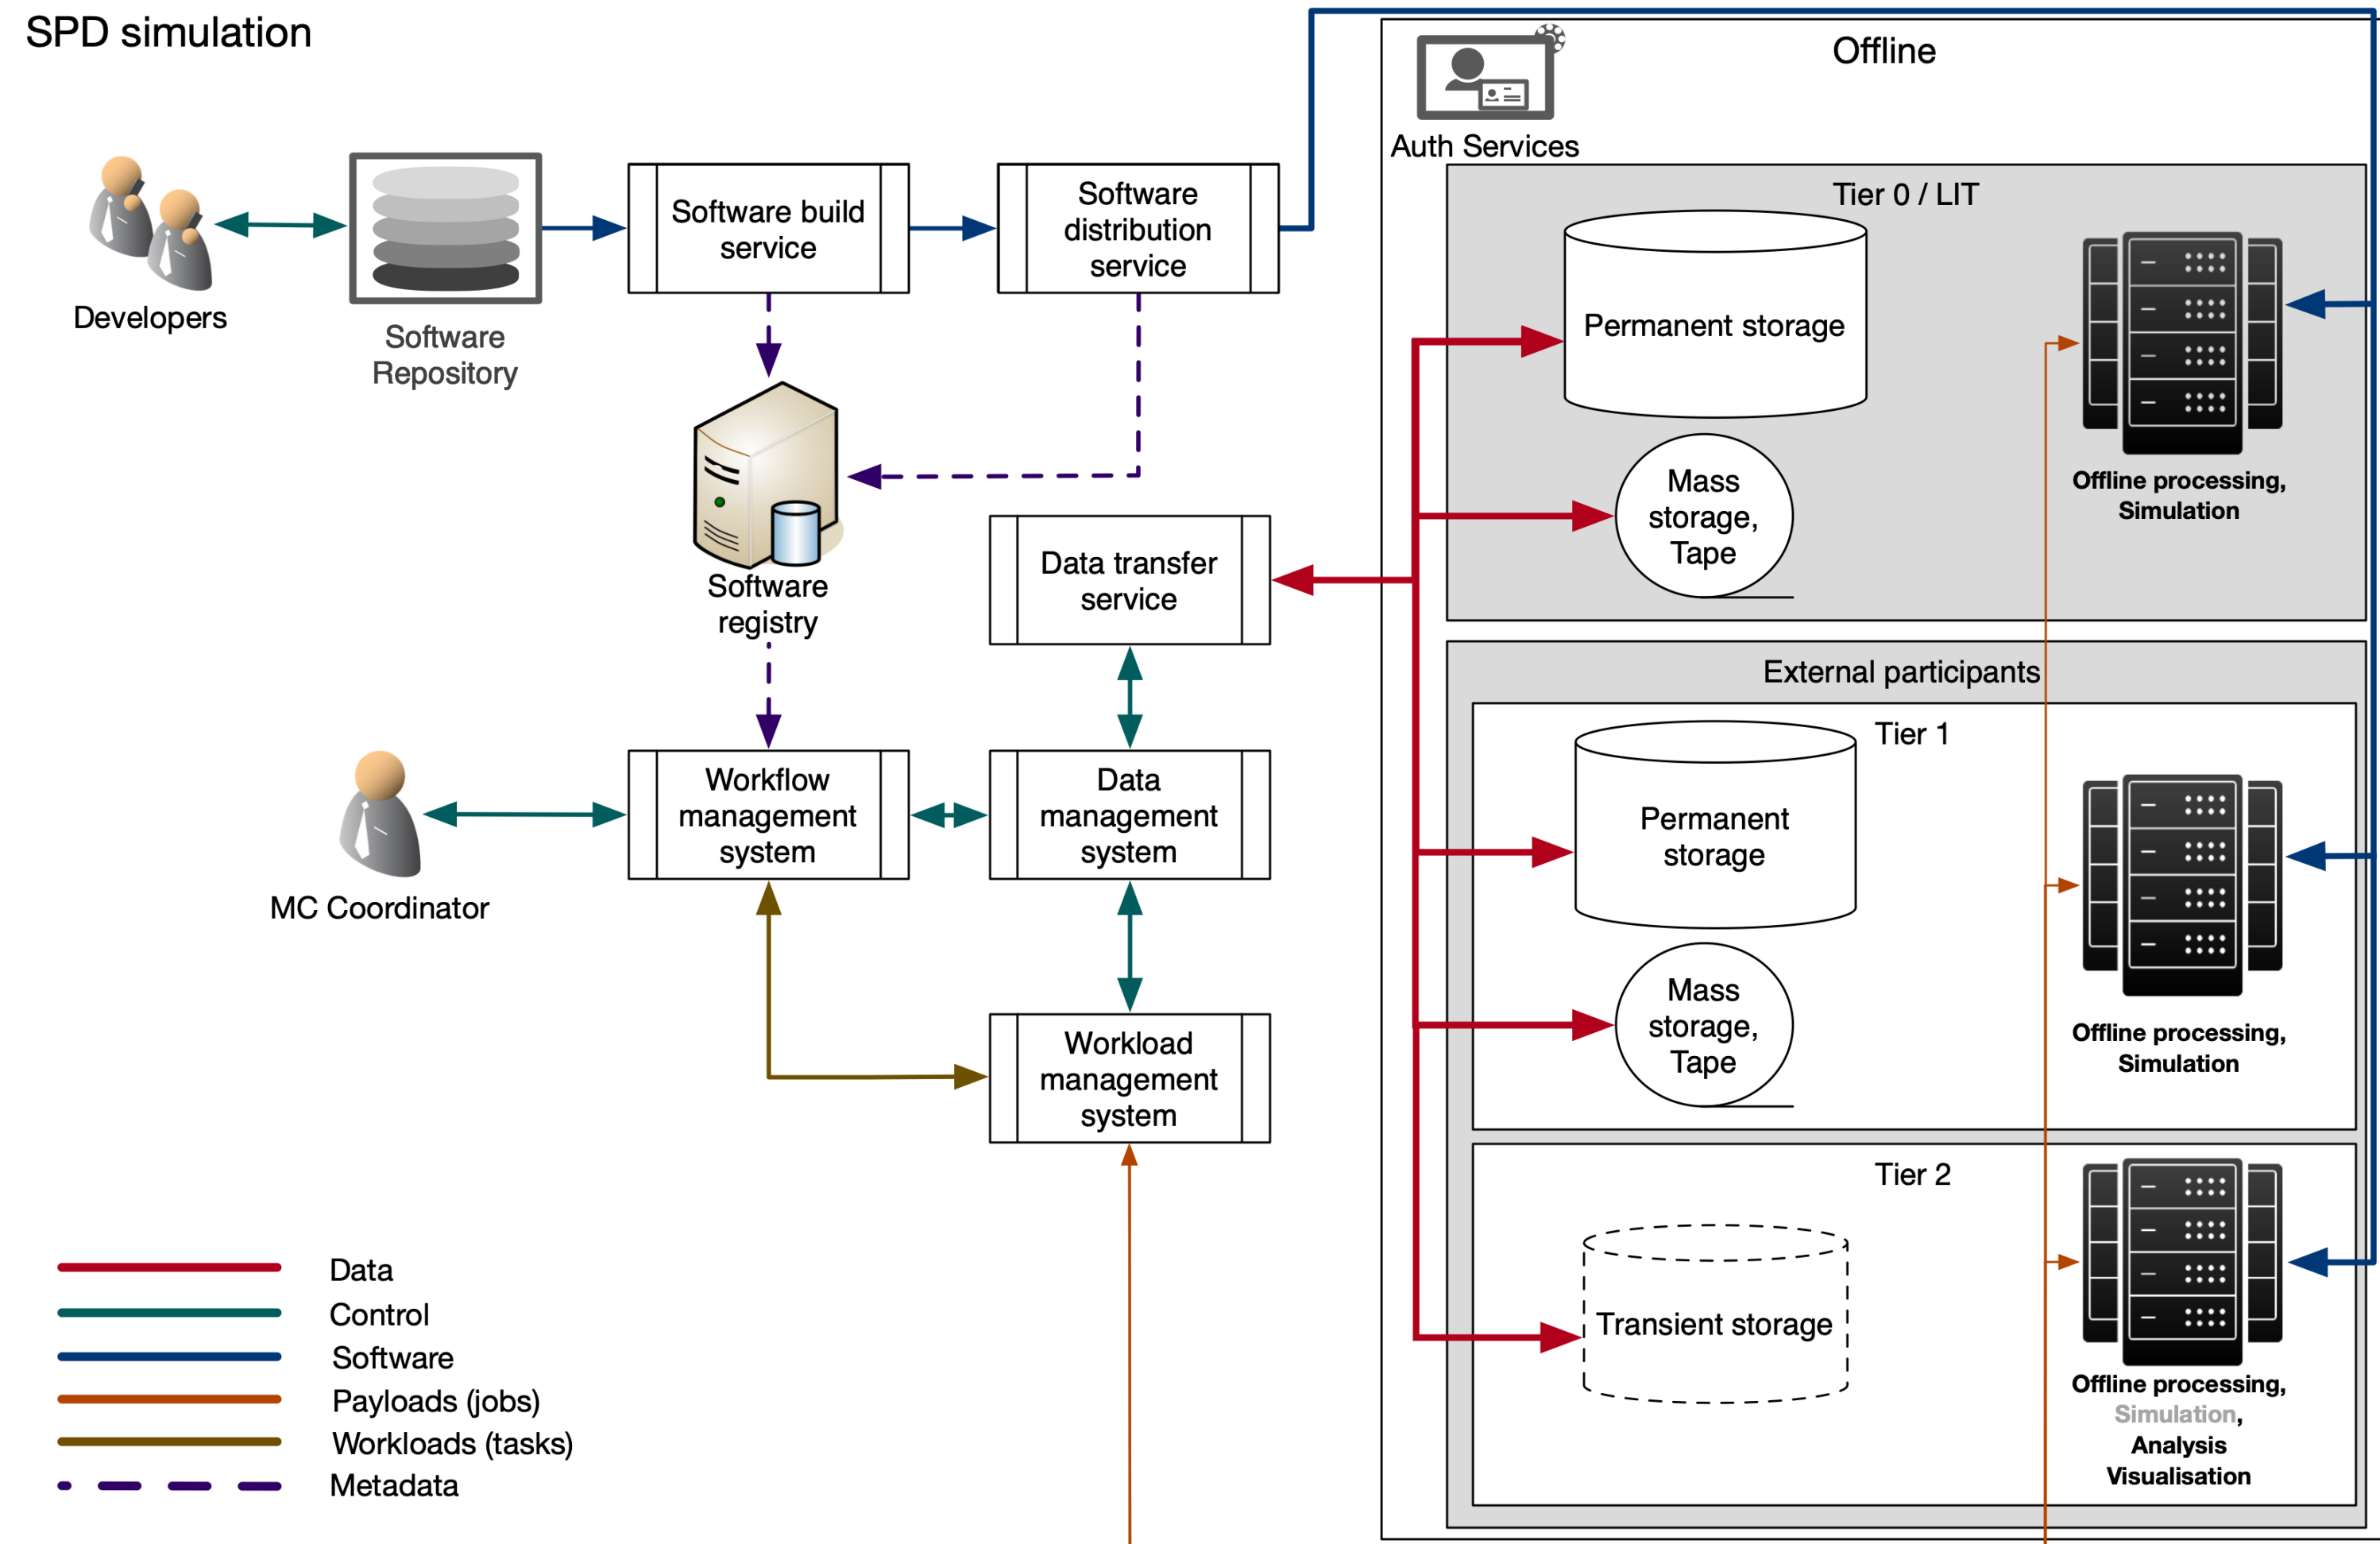
 <p>National Research Center Kurchatov Institute  Moscow, Russia  I. Alexeev</p>	 <p>Higher Institute of Technologies and Applied Sciences (InSTEC)  Havana, Cuba  K. Shtejer</p>
 <p>Cairo University  Cairo, Egypt  R. El-Kholy</p>	 <p>Higher School of Economics  Moscow, Russia  F. Ratnikov</p>
 <p>Tsinghua University  Beijing, China  Y. Wang</p>	 <p>Institute of applied physics of the NAS of Belarus  Minsk, Belarus  R. Shulyakovsky</p>
 <p>CTEPP, UNAB  Santiago, Chile  S. Kuleshov</p>	 <p>SAPHIR  Santiago, Chile  S. Kuleshov</p>
 <p>China Institute of Atomic Energy  Beijing, China  X. Li</p>	 <p>Gomel State Technical University  Gomel, Belarus  V. Andreev</p>
 <p>B.I. Stepanov Institute of Physics of the National Academy of Sciences of Belarus  Minsk, Belarus  Yu. Kulchitsky</p>	 <p>National University of Science and Technology  Moscow, Russia  M. Gorshenkov</p>
 <p>Institute for High Energy Physics  Protvino, Russia  S. Golovnya</p>	

Model of external participant

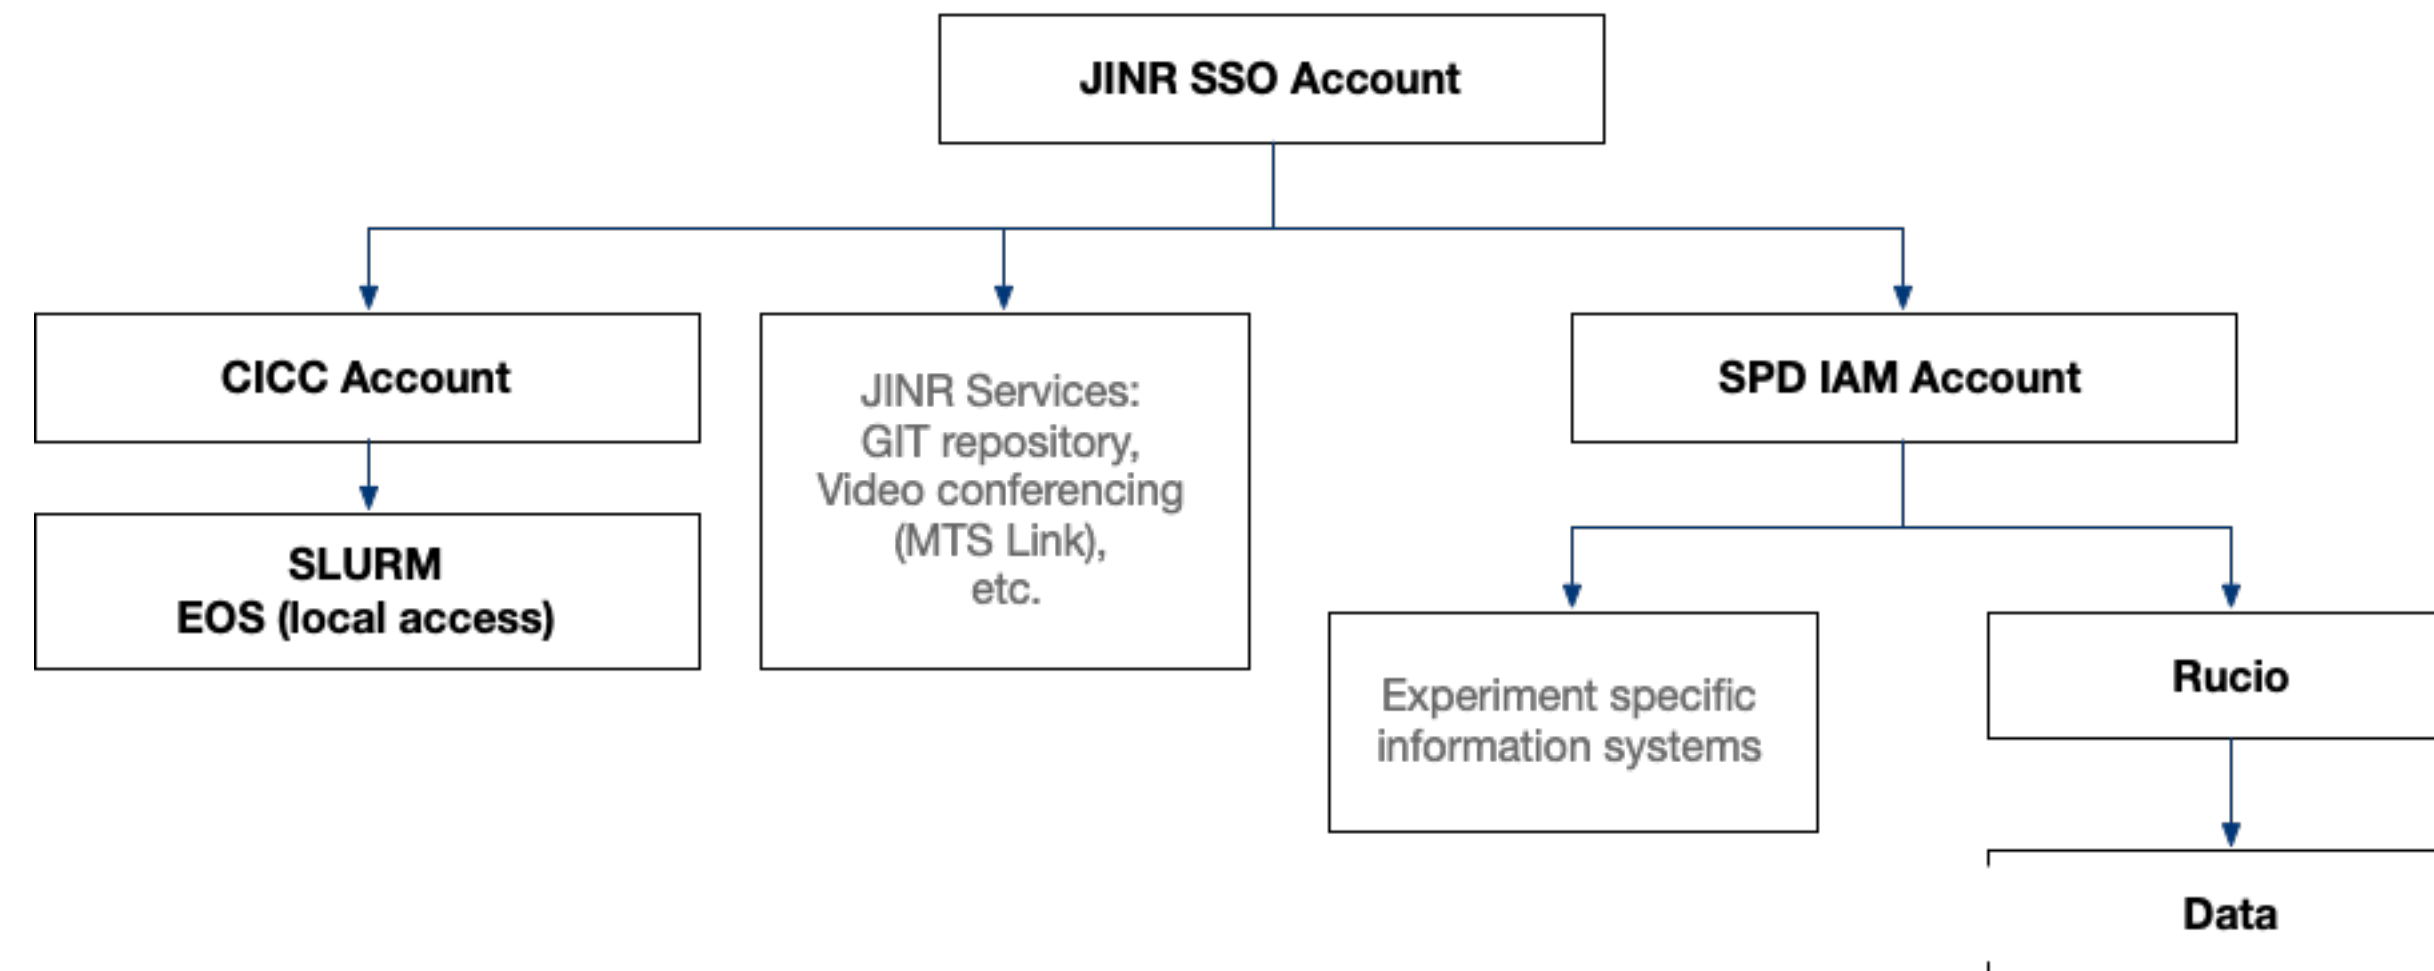
- Physics group
 - Without computing site
 - Not going to setup any computing site
 - Need resources for physics analysis
 - We provide assistance in connection to the JINR computing infrastructure
 - Would like to setup a computing site
 - We provide connection methodology, assistance in documents preparation, consultations with suppliers, etc.
 - Computing site
 - We provide instructions how to connect to our computing environment
- With computing site
 - Not going to participate in the SPD computing
 - Would like to participate in the SPD computing
 - Without experience in grid computing and would like to start using grid
 - We help to setup a grid site
 - With experience in grid computing
 - We provide assistance in connection of the existing site to our computing environment

How we foresee an offline computing environment services for the SPD

- Authentication system: JINR SSO
- Authorization system: IAM
- Information system: CRIC
- Software distribution service: CVMFS
- Data management system: Rucio
- Data transfer service: FTS
- Workload management system: PanDA
- Workflow management system: PanDA/Control Panel



Authentication services



- At the moment we have ~150 users from the JINR and expecting to have at least the same amount of external participants from different institutions
- All users have to have a SSO account (authentication backend for collaboration members) in order to use CICC services
 - Internal have it basing on their user contracts, also have to be in the SPD group in the LDAP
 - External have to become an associate member (sign an MoU, send a letter to the Director of the JINR with a list of persons who is going to use JINR's digital services)
- To start using VO services of the SPD users have to apply for account in the SPD IAM (authorization system only for collaboration members)
- We do not store passwords and any other sensitive information in the db of our identity and access management service, we fully rely on the JINR's SSO which we access via API of Keycloak

Identity and access management

- Address: spd-iam.jinr.ru
- We moved all internal operations between middleware services based on X.509 proxy from the VOMS service to the SPD IAM service
- lxui.jinr.ru, CICC computing and storage resources were configured to work with the SPD IAM as VOMS provider
- A rpm which helps adding IAM VOMS configuration to any computing site was prepared
- An integration between SSO and IAM is now finished, users can log in to IAM (and all underlying services) using JINR SSO account (use JINR SSO button at the SPD IAM login page)
- We have finished transition from VOMS to IAM and we have one entry point for all our computing services — the SPD IAM
- IAM is a single source of info about users and their rights in the distributed computing environment of the SPD experiment including personal quotas in Rucio and EOS
- Middleware services like Rucio and PanDA must be registered in the IAM as well



Welcome to **SPD**

Sign in with your SPD credentials

 Username

 Password

Sign in

[Forgot your password?](#)

Or sign in with

Your X.509 certificate


JINR SSO

Your institutional account

Not a member?

Apply for an account

Clients/services in the IAM

Application		
rucio-auth-client-c1 > more information openid profile offline_access	Authorized: July 23, 2024 Last accessed: November 22, 2024 Expires: Never	Revoke
spd-rucio-auth-client > more information openid profile	Authorized: August 12, 2024 Last accessed: 25 days ago Expires: Never	Revoke
prodsysv1 > more information phone openid profile offline_access rucio email wlcg wlcg.groups	Authorized: October 28, 2024 Last accessed: 5 months ago Expires: Never	Revoke
spd-rucio-auth-client > more information openid profile offline_access	Authorized: December 6, 2024 Last accessed: 25 days ago Expires: Never	Revoke
 cric-sso > more information openid email profile	Authorized: December 17, 2024 Last accessed: 7 hours ago Expires: Never	Revoke
test-panda-client > more information scim iam phone openid profile offline_access groups rucio email wlcg wlcg.groups	Authorized: 5 months ago Last accessed: 22 days ago Expires: Never	Revoke
dev PanDA Monitoring > more information address phone openid email profile	Authorized: 2 months ago Last accessed: 2 months ago Expires: Never	Revoke
dev PanDA Monitoring > more information address phone openid profile scim:read iam:admin.read email	Authorized: 2 months ago Last accessed: 2 months ago Expires: Never	Revoke
dev PanDA Monitoring > more information address phone openid profile offline_access groups scim:read iam:admin.read email wlcg wlcg.groups	Authorized: a month ago Last accessed: a month ago Expires: Never	Revoke

User tokens in the IAM

Tokens

Access Tokens1

Refresh Tokens34

Filter tokens by client...

Filter tokens by user...

#	Client	User	Expires	
1	test-panda-client 384d82d3-8268-4599-bbbc-f25fd3e2e37b	virthead ff4f28c9-71e6-4660-999e-0514193c967c	17 hours from now	<div>✕ Revoke</div>
2	test-panda-client 384d82d3-8268-4599-bbbc-f25fd3e2e37b	virthead ff4f28c9-71e6-4660-999e-0514193c967c	18 hours from now	<div>✕ Revoke</div>
3	dev PanDA Monitoring 37d4261c-a167-4120-aa14-b5b429a6af00	virthead ff4f28c9-71e6-4660-999e-0514193c967c	19 hours from now	<div>✕ Revoke</div>
4	dev PanDA Monitoring 37d4261c-a167-4120-aa14-b5b429a6af00	virthead ff4f28c9-71e6-4660-999e-0514193c967c	19 hours from now	<div>✕ Revoke</div>
5	spd-rucio-auth-client 9dff2cbf-7fb7-40cf-86a1-25e9a866321b	elenazem 5c559362-8bad-4e8d-bbae-0f77c1215910	tomorrow	<div>✕ Revoke</div>
6	test-panda-client 384d82d3-8268-4599-bbbc-f25fd3e2e37b	virthead ff4f28c9-71e6-4660-999e-0514193c967c	2 days from now	<div>✕ Revoke</div>
7	test-panda-client 384d82d3-8268-4599-bbbc-f25fd3e2e37b	elenazem 5c559362-8bad-4e8d-bbae-0f77c1215910	2 days from now	<div>✕ Revoke</div>
8	test-panda-client 384d82d3-8268-4599-bbbc-f25fd3e2e37b	monakov a050f241-8585-4ed5-912b-83e37caa7ad5	3 days from now	<div>✕ Revoke</div>
9	test-panda-client 384d82d3-8268-4599-bbbc-f25fd3e2e37b	virthead ff4f28c9-71e6-4660-999e-0514193c967c	3 days from now	<div>✕ Revoke</div>
10	test-panda-client 384d82d3-8268-4599-bbbc-f25fd3e2e37b	monakov a050f241-8585-4ed5-912b-83e37caa7ad5	3 days from now	<div>✕ Revoke</div>

1


2

3


4

Information system 1/2

cosmo ▾
🏠
Core ▾
Core API ▾
NICA ▾
NICA API ▾
Admin ▾
Logs ▾
🔑 Help ▾
👤 artem.petrosyan@ji... ▾
🔑
🔴



NICA CRIC



TOPOLOGY NAVIGATION

Quickly browse table views for basic topology objects.

- [Federations](#)
- [RC Sites](#)
- [Experiment Sites](#)
- [Services](#)

- [PanDA Sites](#)
(Compute Units)
- [PanDA Queues](#)
(Compute Resources)
- [PanDA Queue parameters](#)

- [Storage Units](#)
- [Storage Resources](#)
(DDMEndpoints, RSEs)
- [Storage Resource parameters](#)
- [Storage Protocols](#)

DOWNTIMES

Browse site downtimes and object exclusion features.

- [Downtimes List](#)
- [Downtime Calendar](#)

- [PanDA Queue Status](#)
- [PanDA Queue Status History](#)
- [PanDA Queue availability](#)

- [DDMEndpoint Status](#)
- [DDMEndpoint Status History](#)

- [Exclusion Probes](#)

API Export

- [Downtime JSON](#) (by RCSites)
- [Downtime JSON](#) (by VOSite)
- [Probes JSON](#)

JSON API

List of mostly used API for data export.

- [Federations JSON](#)
- [RCSites JSON export](#)
- [RCSite Services](#)

- [NICA Sites JSON](#)
- [PandaQueues JSON](#)
- [PandaQueue SW tags](#)
- [DDMEndpoints JSON](#)

Object Exclusion API Export

- [PandaQueueStatus JSON](#)
- [DDMEndpointStatus JSON](#)
- [PandaQueueStatus History](#)
- [DDMEndpointStatus History](#)

[API index](#)
🔧
[Help/Report an issue](#)

OPERATIONS

Regular operations and data modification forms.

- [Request ADMIN privileges](#)
- [Change HC probe](#)
(PandaQueue status)
- [PQ settings for HC](#)
- [StorageResource bulk update](#)

- [Crons List](#)
- [Change Logs](#)

- CRIC stores all info about CE, SE and their relations and exports data to all underlying systems

Information system 2/2

cosmo

Home

Core

Core API

NICA

NICA API

Admin

Logs

Help

artem.petrosyan@ji...

Key

Power

Layout

Link

Export

Print

Columns14/20

Filter

Reload

NICA Site list

200

spd	filter by NICA Site	ACTIVE	filter	filter by Site	filter by Country	filter by Storage Units	filter by PanDA Sites	filter by	filter by	filter by	filter by c	filter by	filter by Data polic
VO	NICA Site	State	Tier	Site	Country	Storage Units	PanDA Sites	ADC notify	Auto proxy	core power	core energy	cloud	Data policy
spd	JINR-SPD	ACTIVE	T0	JINR	Russian Federation	SPD-JINR-DATA	JINR-SPD-PS	✓	✗	10	0	RU	
spd	PNPI-SPD	ACTIVE	T1	PNPI	Russian Federation	SPD-PNPI-DATA	PNPI-SPD-PS	✓	✗	6	0	RU	
spd	SSAU-SPD	ACTIVE	T2	SSAU	Russian Federation			✗	✗	10	0	RU	
VO	NICA Site	State	Tier	Site	Country	Storage Units	PanDA Sites	ADC notify	Auto proxy	core power	core energy	cloud	Data policy

Showing 1 to 3 of 3 entries

Previous

1

Next

cosmo

Home

Core

Core API

NICA

NICA API

Admin

Logs

Help

artem.petrosyan@ji...

Key

Power

Layout

Link

Export

Print

+ new RSE

Columns13/18

Filter

Reload

DDMEndpoint list

200

filter by DDMEndpoint	filter by Experiment si	filter by Storage Unit	filter	filter by Ty	filter by Endpoint	ACTIVE	filter by Resource	filter by	filter by d	filter by	filter by	filter by Space metho
DDMEndpoint	Experiment site	Storage Unit	Tier	Type	Endpoint	State	Resource	cache	determ	volat	mkdir	Space method
JINR EOS2 DATADISK	JINR-SPD	SPD-JINR-DATA	T0	DATADISK	production/	ACTIVE	SPDDATA@JINR SPD EOS	✗	✓	✗	✗	rucio
PNPI SPD DATADISK	PNPI-SPD	SPD-PNPI-DATA	T1	DATADISK	datadisk/rucio/	ACTIVE	SPDDATA@PNPI EOS	✗	✓	✗	✗	rucio
DDMEndpoint	Experiment site	Storage Unit	Tier	Type	Endpoint	State	Resource	cache	determ	volat	mkdir	Space method


Showing 1 to 2 of 2 entries

Previous

1

Next

Distributed data management




- Distributed data managed by Rucio
- Tested thoroughly to be sure that it can handle our needs
- Several daemons were developed to integrate Rucio with the SPD IAM (users, groups, quotas) and CRIC (storages configuration)
- An export module was developed to deliver storages usage info to PanDA
- Rucio clients installed at the CVMFS and available at lxui.jinr.ru
- Since this year, SPD has its own EOS in JINR
- An EOS instance in PNPI also used to store results of central MC productions
- As a side effect, the SPD Rucio was used for JUNO dCache and S3 tests

PanDA IAM integration

- We finished configuring a JWT based authentication in PanDA
- Users can submit tasks via command line client or (preferable) via the Control Panel
- During task submission, in order to identify themselves, users being redirected to the SPD IAM
- The same identity is now used to log in to the Control Panel and to submit a task, it allows us to set up an end-to-end accounting
- PanDA supports auto-registration, so, unlike Rucio, there is no need to develop any identity import service

```
(test) [virthead@vm221-128 task]$ export PANDA_URL=http://vm221-128.jinr.ru:25080/server/panda
(test) [virthead@vm221-128 task]$ export PANDA_URL_SSL=https://vm221-128.jinr.ru:25443/server/panda
(test) [virthead@vm221-128 task]$ export PANDA_CONFIG_ROOT=/home/virthead/pandaclient
(test) [virthead@vm221-128 task]$ export PANDA_USE_NATIVE_HTTPLIB=1
(test) [virthead@vm221-128 task]$ export PANDA_AUTH=oidc
(test) [virthead@vm221-128 task]$ export PANDA_VERIFY_HOST=off
(test) [virthead@vm221-128 task]$ export SSL_CERT_DIR=/etc/grid-security/certificates/
(test) [virthead@vm221-128 task]$ export PANDA_AUTH_V0=spd.nica.jinr:production
(test) [virthead@vm221-128 task]$ python testOpenCharm_simu.py
INFO : Please go to https://spd-iam.jinr.ru/device?user_code=Y1XBDB and sign in. Waiting until authentication is completed
INFO : Ready to get ID token?
[y/n]
y
INFO : All set
(0, [True, 330, None])
```

 INDIGO IAM for SPD virthead

Approval Required for *test-panda-client*

test-panda-client

Access to:

- log in using your identity
- basic profile information
- email address
- offline access

Do you authorize " test-panda-client "?

Authorize Deny

	jeditaskid [PK] bigint	taskname character varying (256)	status character varying (64)	username character varying (128)
266	346	PROD2025-009.SIM.1	done	Elena Zemlyanichkina
267	347	PROD2025-009.RECO.1	finished	Elena Zemlyanichkina
268	348	MC2025_S1-003-SIM.1	aborted	Artem Petrosyan
269	349	MC2025_S1-003-SIM.1	aborted	Artem Petrosyan
270	350	MC2025_S1-003-SIM.1	aborted	Artem Petrosyan
271	351	MC2025_S1-003-SIM.3	aborted	Artem Petrosyan
272	352	MC2025_S1-003-SIM.3	failed	Artem Petrosyan
273	353	MC2025_S1-003-SIM.4	done	Artem Petrosyan
274	354	PROD2025-010.SIM	done	Elena Zemlyanichkina
275	355	PROD2025-010.RECO	finished	Elena Zemlyanichkina

Example of the MC task definition 1/2

- Step 1: simulation
- User defines an output dataset name
- Desired total number of events and events per job
- The system divides the total number of events by the number events per job and generates the required number of jobs
- User can specify either a specific computing queue or a cloud; in the second case, the jobs will be distributed among the queues of the specified cloud
- Jobs execution is performed in the container

```
TaskName = '2024.27GeV.test-MB.2st.DSSD.simu'
DatasetName = '2024.MC.27GeV.test-minbias.00001.SIMUL.0'
#DatasetName = 'jeditest.000023.simu'
```

```
taskParamMap = {}
```

```
taskParamMap['nEventsPerJob'] = 4000
taskParamMap['nEvents'] = 20000000
taskParamMap['noInput'] = True
taskParamMap['skipScout'] = True
taskParamMap['taskName'] = TaskName
taskParamMap['userName'] = 'Artem Petrosyan'
taskParamMap['vo'] = 'spd.nica.jinr'
taskParamMap['taskPriority'] = 900
taskParamMap['architecture'] = 'x86_64'
taskParamMap['transUses'] = 'A'
taskParamMap['transHome'] = None
taskParamMap['transPath'] = 'https://159.93.221.125:8080/spd_simu_VA_transform.sh'
taskParamMap['processingType'] = 'step1'
taskParamMap['prodSourceLabel'] = 'managed'
taskParamMap['taskType'] = 'test'
taskParamMap['workingGroup'] = 'spd.nica.jinr'
taskParamMap['cloud'] = 'JINR'
taskParamMap['ramCount'] = 1900
```

```
outDatasetNameLog = '{0}.log'.format(DatasetName)
outDatasetNameS = '{0}.S'.format(DatasetName)
outDatasetNameP = '{0}.P'.format(DatasetName)
```

```
taskParamMap['log'] = {'dataset': outDatasetNameLog,
                       'type': 'template',
                       'param_type': 'log',
                       'token': 'DATADISK',
                       'value': '{0}.${{SN}}.log.tgz'.format(DatasetName)}
```

```
taskParamMap['jobParameters'] = [
    {'type': 'constant',
     'value': "'singularity run --bind /cvmfs/spd.jinr.ru/production/MC/2024.27GeV.test-MB.2st.DSSD:/prod -H  
./:/WORKDIR  
/cvmfs/spd.jinr.ru/images/spdroot-4.1.6.sif spdroot.py -b -q \'/prod/simu.C({0}, '''.format(taskParamMap['nEventsPerJob'])  
}',  
    . . . . .
```


Example of the MC task definition 2/2

- Step 2: reconstruction
- User defines a name of the input dataset, in this example there are two input datasets of the same size (have the same number of files)
- Sets a name of the output dataset
- Set how many jobs needs to be created per each file in the dataset
- At the job generation stage, the workload management system communicates with the data management service, reads the size (number of files) of the dataset and generates the appropriate number of jobs
- The input files will be staged-in from the storage closest to the computing node

```
scope = '2024'
inDatasetName = '2024.MC.27GeV.test-minbias.00001.SIMUL.0'
outDatasetName = '2024.MC.27GeV.test-minbias.00001.RECO.2'

inDatasetNameS = '{0}.S'.format(inDatasetName)
inDatasetNameP = '{0}.P'.format(inDatasetName)
outDatasetNameR = '{0}.R'.format(outDatasetName)
outDatasetNameLog = '{0}.log'.format(outDatasetName)

taskParamMap = {}

taskParamMap['nFilesPerJob'] = 1
taskParamMap['nEventsPerJob'] = 4000
taskParamMap['noInput'] = False
taskParamMap['taskName'] = TaskName
taskParamMap['userName'] = 'Artem Petrosyan'
taskParamMap['vo'] = 'spd.nica.jinr'
taskParamMap['taskPriority'] = 900
taskParamMap['architecture'] = 'x86_64'
taskParamMap['transUses'] = 'A'
taskParamMap['transHome'] = None
taskParamMap['transPath'] = 'https://159.93.221.125:8080/spd_simu_VA_transform.sh'
taskParamMap['processingType'] = 'step2'
taskParamMap['prodSourceLabel'] = 'managed'
taskParamMap['taskType'] = 'test'
taskParamMap['workingGroup'] = 'spd.nica.jinr'
taskParamMap['cloud'] = 'JINR'
taskParamMap['ramCount'] = 1900

taskParamMap['log'] = {'dataset': outDatasetNameLog,
                      'type': 'template',
                      'param_type': 'log',
                      'token': 'DATADISK',
                      'value': '{0}.${{SN}}.log.tgz'.format(outDatasetName)}

taskParamMap['jobParameters'] = [
    {'type': 'constant',
     'value': ''singulariry run --bind /cvmfs/spd.jinr.ru/production/MC/2024.27GeV.test-MB.2st.DSSD:/prod -H
./:/WORKDIR /cvmfs/spd.jinr.ru/images/spdroot-4.1.6.1.sif spdroot.py -b -q \'/prod/reco.C({0}, ''
.format(taskParamMap['nEventsPerJob'])
    },
```

Control panel of the production manager 1/2

Simulation

Task Creation

Task name:

Output dataset name:
[Naming convention here](#)

Total events:

Events per job:

Cloud:

Data disk:

Skip scout: ☐

Offset:

Path to execution files:
smth like -> /cvmfs/spd.jinr.ru/production/MC/minbias-P8-spdroot417-dev.10GeV.V01

Path to container:
smth like -> /cvmfs/spd.jinr.ru/images/spdroot-dev-4.1.7.sif

Create task

Reconstruction

Task Creation

Task name:

Input dataset name:
[Naming convention here](#), note that no extension expected

Output dataset name:
[Naming convention here](#)

Files per job:

Cloud:

Data disk:

Skip scout: ☐

Offset:


Path to execution files:
smth like -> /cvmfs/spd.jinr.ru/production/MC/minbias-P8-spdroot417-dev.10GeV.V01

Path to container:
smth like -> /cvmfs/spd.jinr.ru/images/spdroot-dev-4.1.7.sif

Create task

- An application allowing users to easily define a MC chain processing via Web UI was put into pre-production in late 2024
- It is integrated with the SPD IAM, and, thus, allows to pass user information to PanDA




Control panel of the production manager 2/2



- The web ui allows one to define each processing step individually, as well as set the entire chain at once
- During several last months a couple dozens of productions were processed basing on requests done by our production manager, Elena Zemlyanichkina

Central production stats

- Successfully processed about 300k jobs across 55 tasks



- Total output datasets volume – more than 425 TB

Task ID	Task name ↑ ↓	Status	Start date	End date	Walltime	Total events	Events per job	Total jobs	Out DS size, GB	Out Log size, GB
364	PROD2025-014.SIM	done	03 May 2025	04 May 2025	24612	40000000	4000	10000	18358.86	1.96
363	PROD2025-013.RECO	finished	02 May 2025	04 May 2025	20	None	None	9992	18536.20	5.06
362	PROD2025-013.SIM.2	done	30 Apr 2025	02 May 2025	24899	40000000	4000	10000	18357.95	1.93
359	PROD2025-012.RECO	finished	28 Apr 2025	29 Apr 2025	24	None	None	9993	18546.74	5.08
358	PROD2025-012.SIM	done	25 Apr 2025	26 Apr 2025	23316	40000000	4000	10000	18360.80	1.89
357	PROD2025-011.RECO	done	23 Apr 2025	24 Apr 2025	10	None	None	1250	2319.57	0.63
356	PROD2025-011.SIM	done	22 Apr 2025	23 Apr 2025	22496	5000000	4000	1250	2295.55	0.24
355	PROD2025-010.RECO	finished	18 Apr 2025	18 Apr 2025	37	None	None	1244	287.74	0.49
354	PROD2025-010.SIM	done	17 Apr 2025	17 Apr 2025	0	5000000	4000	1250	259.39	0.13
353	MC2025_S1-003-SIM.4	done	17 Apr 2025	17 Apr 2025	0	1000	100	10	0.21	0.00

Productions requests spreadsheet

Production name/ID	Status	Description					Software type/version	Short description (for datasets naming)	Number of events	Events per file	Initial seed	Processing type
		Stage	Collision type	Geometry type	Energy	Polarization						
PROD2025-010	Done	S1	dd	Micromegas, TS, ECal, RS, BBC, ZDC (sketch)	4 GeV	UU	spdroot-dev-4.1.7.1	dd-minbias-FTF-spdroot4171-dev	5 000 000	4000	1-1250	reco
PROD2025-011	Done	S2	pp	DSSD, TS, TOF, ECal, FARICH, RS, BBC, ZDC (sketch)	27 GeV	UU	spdroot-dev-4.1.7.2	minbias-P8-spdroot4172-dev test	5 000 000	4000	1-1250	simu
												reco
PROD2025-012	Done	S2	pp	DSSD, TS, TOF, ECal, FARICH, RS, BBC, ZDC (sketch)	27 GeV	UU	spdroot-dev-4.1.7.2	minbias-P8-spdroot4172-dev test	40 000 000	4000	1-10000	simu
												reco
PROD2025-013	Done	S2	pp	DSSD, TS, TOF, ECal, FARICH, RS, BBC, ZDC (sketch)	27 GeV	UU	spdroot-dev-4.1.7.2	minbias-P8-spdroot4172-dev test	40 000 000	4000	10001-20000	simu
												reco
PROD2025-014	Runnig	S2	pp	DSSD, TS, TOF, ECal, FARICH, RS, BBC, ZDC (sketch)	27 GeV	UU	spdroot-dev-4.1.7.2	minbias-P8-spdroot4172-dev test	40 000 000	4000	20001-30000	simu
												reco
PROD2025-015	Ready	S2	pp	DSSD, TS, TOF, ECal, FARICH, RS, BBC, ZDC (sketch)	27 GeV	UU	spdroot-dev-4.1.7.2	minbias-P8-spdroot4172-dev test	40 000 000	4000	30001-40000	simu
												reco
PROD2025-016	Ready	S2	pp	DSSD, TS, TOF, ECal, FARICH, RS, BBC, ZDC (sketch)	27 GeV	UU	spdroot-dev-4.1.7.2	minbias-P8-spdroot4172-dev test	40 000 000	4000	40001-50000	simu
												reco

- At the first we agreed to keep all production requests in the Google doc spreadsheet

Productions requests db 1/2

SPD Data processing																
<div> Home Requests Config API artem.petrosyan@ji... Key Power </div>																
<div> Export Columns 17/23 Filter Reload Processing Request list 100 </div>																
Request	campaign	status	# procs	swproject	version	stage	C	E	P	Events	EF	Tag	S	S	Geometry	description
17 PROD2025-001	SPD MC 2025	DONE	0	spdroot-dev	4.1.7.0	S1	pp	10	UU	5000000	4000	minbias-P8-spdroot417-dev test	1	1250	Micromegas, TS, ECal, RS, BBC, ZDC (sketch)	PRE-PRODUCTION obsolete
16 PROD2025-002	SPD MC 2025	DONE	0	spdroot-dev	4.1.7.0	S1	pp	10	UU	20000000	4000	minbias-P8-spdroot417-dev	1	5000	Micromegas, TS, ECal, RS, BBC, ZDC (sketch)	/eos/nica/spd/users/elenazem/productions/PROD2025-002_recofiles.txt PROD2025-002_paramfiles.txt
15 PROD2025-003	SPD MC 2025	DONE	0	spdroot-dev	4.1.7.0	S1	pp	10	UU	20000000	4000	minbias-P8-spdroot417-dev	5001	10000	Micromegas, TS, ECal, RS, BBC, ZDC (sketch)	/eos/nica/spd/users/elenazem/productions/PROD2025-003_recofiles.txt PROD2025-003_paramfiles.txt
14 PROD2025-004	SPD MC 2025	DONE	0	spdroot-dev	4.1.7.0	S1	pp	10	UU	40000000	4000	minbias-P8-spdroot417-dev	1	10000	Micromegas, TS, ECal, RS, BBC, ZDC (sketch)	/eos/nica/spd/users/elenazem/productions/PROD2025-004_recofiles.txt PROD2025-004_paramfiles.txt
13 PROD2025-005	SPD MC 2025	DONE	0	spdroot-dev	4.1.7.1	S1	pp	5	UU	5000000	4000	minbias-FTF-spdroot4171-dev	1	1250	Micromegas, TS, ECal, RS, BBC, ZDC (sketch)	/eos/nica/spd/users/elenazem/productions/PROD2025-005_recofiles.txt PROD2025-005_paramfiles.txt
12 PROD2025-006	SPD MC 2025	DONE	0	spdroot-dev	4.1.7.1	S1	pp	10	UU	5000000	4000	minbias-FTF-spdroot4171-dev	1	1250	Micromegas, TS, ECal, RS, BBC, ZDC (sketch)	/eos/nica/spd/users/elenazem/productions/PROD2025-006_recofiles.txt PROD2025-006_paramfiles.txt

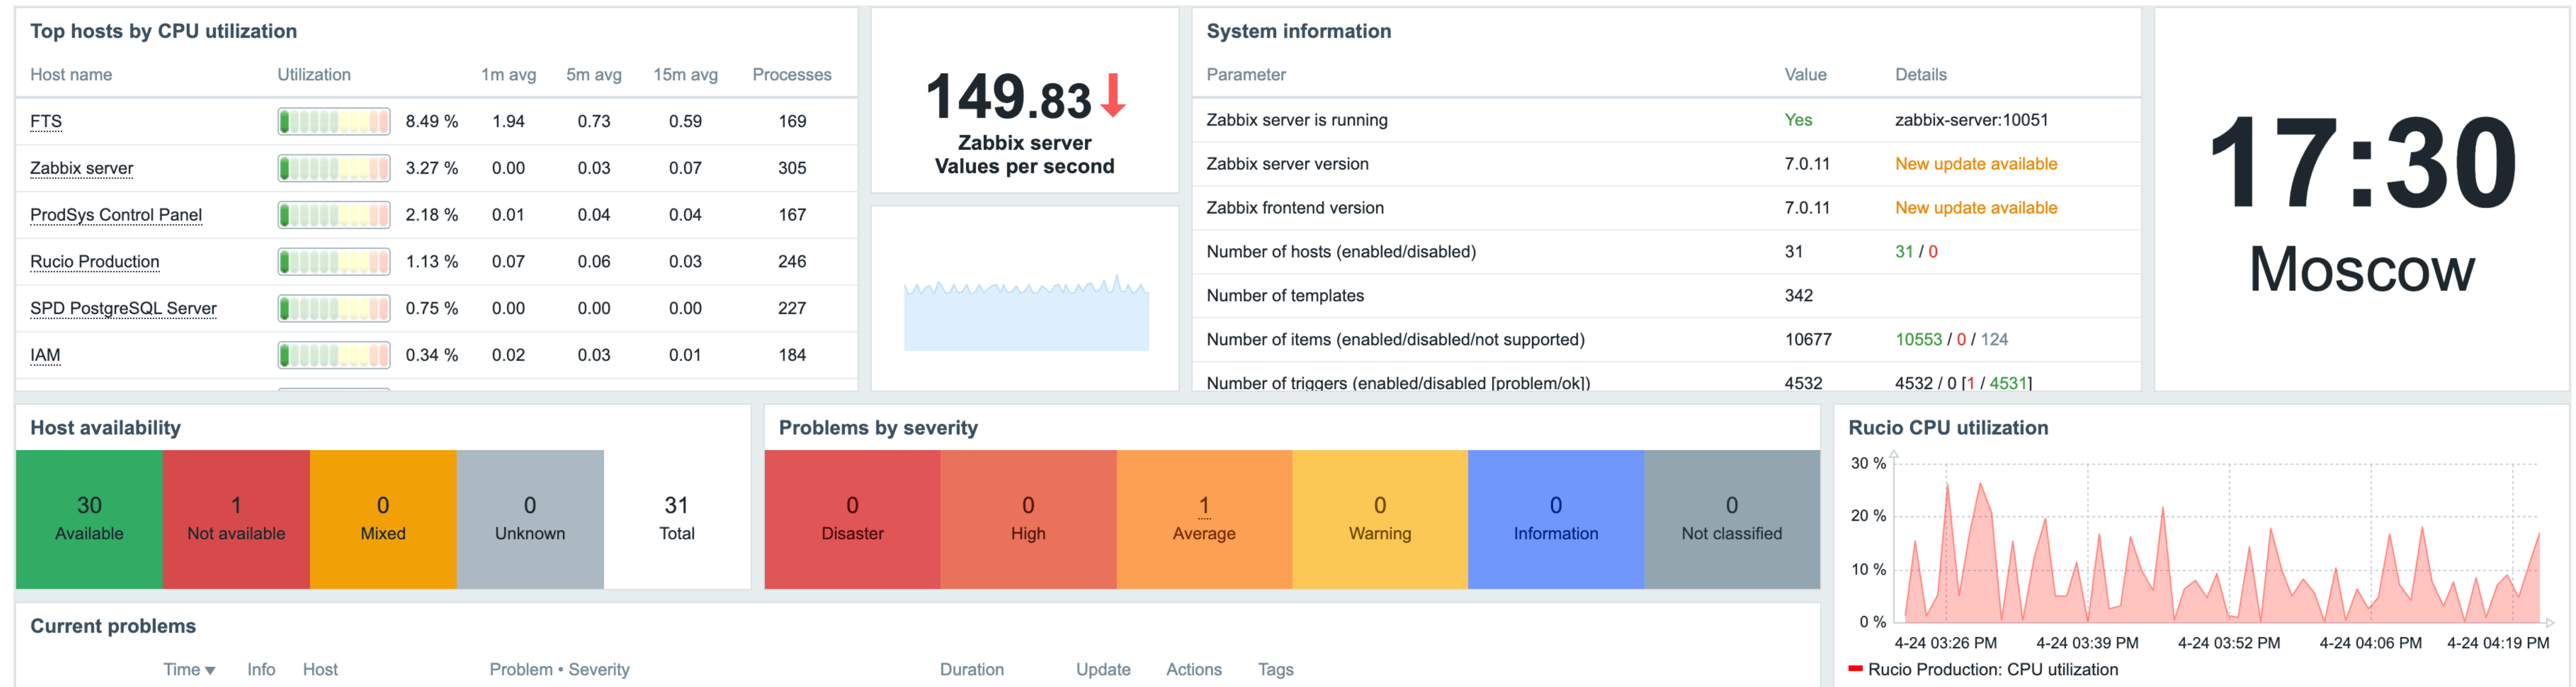
- At the second we decided to create a special application to manage requests

Productions requests db 2/2

SPD Data processing								
<div> Home Requests Config API artem.petrosyan@ji... </div> <div> Export Columns 8/10 Filter Reload Processing Request list 100 </div>								
ID	Request	type	script	dsn_pattern	input_datasets	output_datasets	description	
3	PROD2025-016	simu	/cvmfs/spd.jinr.ru/production/MC/minbias-P8-spdroot4172-dev.27GeV.V02/simu.C	MC2025_S2:minbias-P8-spdroot4172-dev.27GeV-UU.PROD2025-016.SIM*		MC2025_S1:MC2025_S1.minbias-P8-spdroot4172-dev.27GeV-UU.PROD2025-016.SIM.1.S, MC2025_S1:MC2025_S1.minbias-P8-spdroot4172-dev.27GeV-UU.PROD2025-016.SIM.1.P, MC2025_S1:MC2025_S1.minbias-P8-spdroot4172-dev.27GeV-UU.PROD2025-016.SIM.1.log	P8, FARICH PID for pi/K/p fixed, additional to PROD2025-012	
4	PROD2025-016	reco	/cvmfs/spd.jinr.ru/production/MC/minbias-P8-spdroot4172-dev.27GeV.V02/reco.C	MC2025_S2:minbias-P8-spdroot4172-dev.27GeV-UU.PROD2025-016.RECO*		MC2025_S1:MC2025_S1.minbias-P8-spdroot4172-dev.27GeV-UU.PROD2025-016.RECO.1.log, MC2025_S1:MC2025_S1.minbias-P8-spdroot4172-dev.27GeV-UU.PROD2025-016.RECO.1.R	eos/nica/spd/users/elenazem/productions/PROD2025-016_recofiles.txt PROD2025-016_paramfiles.txt	
1	PROD2025-017	reco	/cvmfs/spd.jinr.ru/production/MC/minbias-P8-spdroot4172-dev.27GeV.V02_D0sig/reco.C	MC2025_S2:minbias-P8-spdroot4172-dev.27GeV-UU.PROD2025-017.RECO*		MC2025_S1:MC2025_S1.minbias-P8-spdroot4172-dev.27GeV-UU.PROD2025-017.RECO.1.log, MC2025_S1:MC2025_S1.minbias-P8-spdroot4172-dev.27GeV-UU.PROD2025-017.RECO.1.R	/eos/nica/spd/users/elenazem/productions/PROD2025-017_recofiles.txt PROD2025-017_paramfiles.txt	
2	PROD2025-017	simu	/cvmfs/spd.jinr.ru/production/MC/minbias-	MC2025_S2:minbias-		MC2025_S1:MC2025_S1.minbias-P8-	P8, FARICH PID for pi/K/p fixed, open-charm signal	

- Since this is a database it can be used for some analysis, extended, integrated with other services, etc.

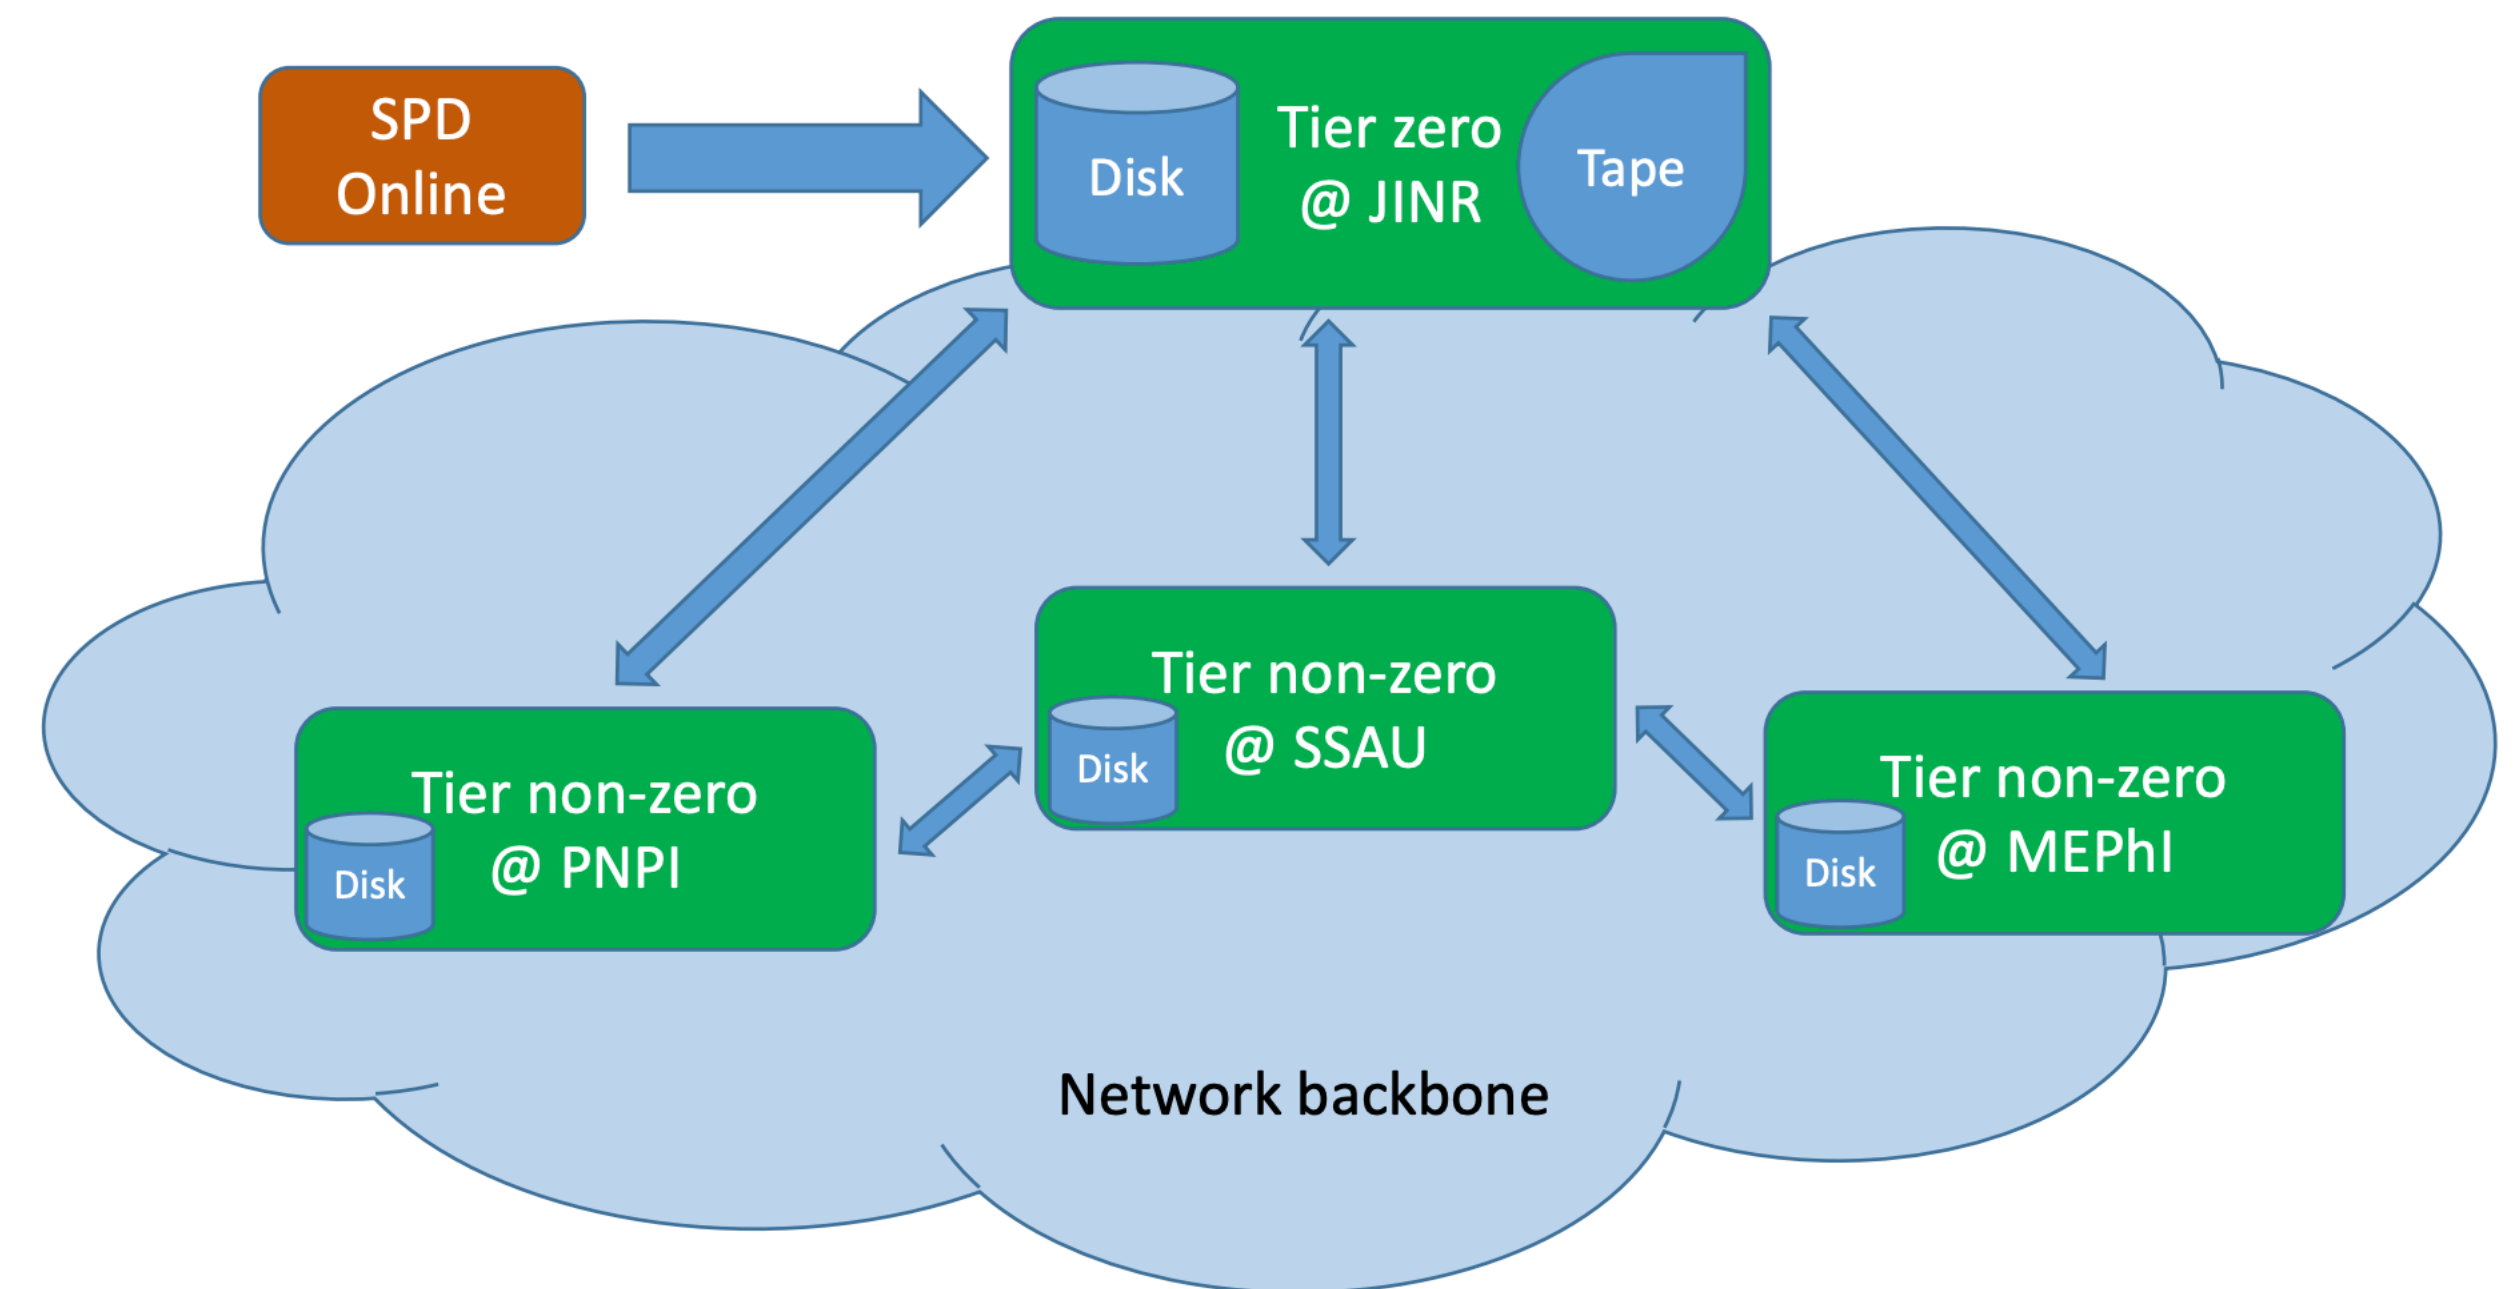
Services monitoring



- We deployed an instance of Zabbix in order to enable monitoring of our growing infrastructure
- At the moment it is very basic setup but we expect to have an integrated monitoring with panels from service-specific metrics, not only OS metrics like CPU utilization, etc.

What is currently deployed in total

- JINR
 - Production system services (prod and devel instances)
 - Computing (2200 cores), usually up to 1500
 - Storage (7.2 PB raw with 27% redundancy = 5.3 PB)
 - Monitoring (somewhat)
- PNPI
 - Computing (288 cores)
 - Storage (190 TB redundant)
 - Monitoring
- SSAU
 - Computing (256 cores)
 - Storage is on the way (240 TB raw with 17% redundancy = 200 TB)
- MEPhI
 - Ongoing negotiations



What is yet missing

- Monitoring coverage for core services
 - Zabbix configuration is ongoing
 - We need to monitor health of services, not just servers
 - Visualization (dashboards, etc.)
- Periodic infrastructure tests
 - PerfSONAR dashboard
 - Job submission tests
 - Worker node health tests
 - Data transfer tests
- Full transition to tokens
 - EOS-side configuration
 - ARC CE-side configuration

JINR certification authority

- Address: ca.jinr.ru
- Finally in production, thanks to the LIT network operation service
- JINR CA is online CA, certificates are issued immediately, which allowed us to optimize our operations a lot
- Any user with JINR SSO account can request user certificate
- LIT CICC computing and storage resources were configured to support JINR CA certificates
- YUM repo with the rpm was prepared to ease installation process
- Can be used as a JINR-centric certification authority for all JINR hosted experiments

📄 Петросян Артем Шмавовнович - Petrosyan Artem Shmavonovich

Мои действительные Host сертификаты: 14

Серийный номер	Имя сертификата
335280838662608440995013	CN=cric.jinr.ru,OU=hosts,OU=GRID,O=JINR,C=RU
330556999940689005096270	CN=vm221-125.jinr.ru,OU=hosts,OU=GRID,O=JINR,C=RU
325831556579249782268863	CN=vm221-128.jinr.ru,OU=hosts,OU=GRID,O=JINR,C=RU
32119164890131102487796	CN=spd-iam.jinr.ru,OU=hosts,OU=GRID,O=JINR,C=RU
382510806361247941477048	CN=10-220-18-146.jinr.ru,OU=hosts,OU=GRID,O=JINR,C=RU
368331973398402739946155	CN=spd-rucio.jinr.ru,OU=hosts,OU=GRID,O=JINR,C=RU
476943422706006476611360	CN=vm221-122.jinr.ru,OU=hosts,OU=GRID,O=JINR,C=RU
491123808586845481041803	CN=spd-fts.jinr.ru,OU=hosts,OU=GRID,O=JINR,C=RU
524181107038056702859852	CN=spd.ssau.ru,OU=hosts,OU=GRID,O=JINR,C=RU
604447717501248574804455	CN=10-220-18-77.jinr.ru,OU=hosts,OU=GRID,O=JINR,C=RU
609173767607613757762429	CN=cric-dev.jinr.ru,OU=hosts,OU=GRID,O=JINR,C=RU
628061244881506981510336	CN=10-220-16-10.jinr.ru,OU=hosts,OU=GRID,O=JINR,C=RU
632795307202027113861823	CN=spd-cric.jinr.ru,OU=hosts,OU=GRID,O=JINR,C=RU
637506996081636867403667	CN=vm221-126.jinr.ru,OU=hosts,OU=GRID,O=JINR,C=RU

Conclusions

- The results of the six-month operation of the created distributed computing environment by users confirm the correctness of the chosen approach:
 - One point of entry allows to manage users and their access rights, as well as applications which are allowed to work in the environment
 - It has become easier to organize centralized mass productions using the web interface
 - The calculations themselves are performed more transparently and users get data for analysis much quicker
 - The results are stored at the common storages and available to all users
 - The system copes well with the load
 - The addition or replacement of computing and storage elements takes place without any/much negative consequences

Side effects

- Results of common efforts that can be (and partially have already) used beyond the SPD collaboration:
 - JINR CA
 - Keycloak as OIDC API to JINR SSO
 - Usage Rucio and FTS as the basis for storages tests of JUNO
 - Arrival of new members of the distributed computing community: Samara
 - Bringing back old members of the distributed computing community: MEPhI

Next steps

- User interfaces development
 - UI/UX improvement of Control panel
 - New features like tasks cloning mechanism in the Control Panel
 - Monitoring and analysis tools
- Integration with the new applied software framework Sampo, transition from SPDRoot to Sampo
- More operation activity along with improvement of the system
 - Finish transition to tokens at services level
- Add more automation for the routine procedures
 - Import users from the SSO database to the IAM automatically basing on LDAP groups
- Users support
 - Finish users data migration from the central JINR EOS to the dedicated one
- Documentation

Thank you!