Production System Status

Artem Petrosyan, MLIT, JINR
XI SPD Collaboration Meeting, Tomsk, Russia
May 20, 2026

Introduction (memory pill)

- SPD, like any modern experiment, requires organized data production for further user analysis

- Taking into account the specifics of the experiment: data volumes, installation on one site, computing resources on another, as well as the possibility of using resources from participating institutions as their contribution to the experiment, we decided to build a system that allows efficient use of remote resources

- When building such a system for the SPD experiment, we use best practices of LHC experiments

Progress since the previous collaboration meeting 1/2

- Computing center in St. Petersburg State University is actively participating in data processing
- Rucio internal monitoring was deployed
- Rucio version upgrade to 38.5.1
- PanDA host upgraded to AlmaLinux9
- PanDA server upgraded to version 0.5.0
- PanDA database migrated to the central DB

Progress since the previous collaboration meeting 2/2

- Automated SpdRoot containers delivery to CernVM FS was organized
- There is ongoing integration between Control panel of the production manager and Requests DB
- Services test suite is being developed
- SPD group quota was removed from the main JINR EOS
- Dedicated EOS is token-ready
- Zabbix service was updated

PanDA upgrade consequences

- After upgrading rotation of empty jobs has finally finished, submission of jobs now is fully controlled by PanDA
- Harvester communicates PanDA how many jobs are in activated status for each queue and submits corresponding amount of pilots, taking into account depth of the queue and amount of running jobs

```
2026-05-15 14:07:18,675 panda.log.worker_adjuster: DEBUG    define_num_workers <site=JINR_SPD_T2_PROD> Processing queue JINR_SPD_T2_PROD job_type ANY resource_
type ANY with static_num_workers {'nReady': 0, 'nRunning': 292, 'nQueue': 8, 'nNewWorkers': 0}
2026-05-15 14:07:18,680 panda.log.worker_adjuster: DEBUG    get_queue_no_pilots_when_no_active_jobs <site=JINR_SPD_T2_PROD> ret_val=False
2026-05-15 14:07:18,680 panda.log.worker_adjuster: DEBUG    define_num_workers <site=JINR_SPD_T2_PROD> available activated panda jobs 9
2026-05-15 14:07:18,680 panda.log.worker_adjuster: DEBUG    get_queue_activate_worker_factor <site=JINR_SPD_T2_PROD> ret_val=1.0
2026-05-15 14:07:18,681 panda.log.worker_adjuster: DEBUG    get_activate_worker_factor <site=JINR_SPD_T2_PROD> ret_val=1.0
2026-05-15 14:07:18,681 panda.log.worker_adjuster: DEBUG    define_num_workers <site=JINR_SPD_T2_PROD> limiting max_queued_workers to min(n_activated=9, queue_
limit=1000)
2026-05-15 14:07:18,681 panda.log.worker_adjuster: DEBUG    define_num_workers <site=JINR_SPD_T2_PROD> setting n_new_workers to 1 in max_queued_workers calcula
tion
2026-05-15 14:07:18,681 panda.log.worker_adjuster: DEBUG    define_num_workers <site=JINR_SPD_T2_PROD> setting n_new_workers to 1 to respect max_workers
2026-05-15 14:07:18,681 panda.log.worker_adjuster: DEBUG    define_num_workers <site=JINR_SPD_T2_PROD> setting n_new_workers to 1 in order to respect universal
 maxNewWorkers
2026-05-15 14:07:18,796 panda.log.worker_adjuster: DEBUG    define_num_workers <site=JINR_SPD_T2_PROD> defined {'JINR_SPD_T2_PROD': {'ANY': {'ANY': {'nReady':
0, 'nRunning': 292, 'nQueue': 8, 'nNewWorkers': 1}}}}
2026-05-15 14:07:19,005 panda.log.worker_adjuster: DEBUG    define_num_workers <site=SPbSU_SPD_PROD> start
2026-05-15 14:07:19,005 panda.log.worker_adjuster: DEBUG    define_num_workers <site=SPbSU_SPD_PROD> static_num_workers: {'SPbSU_SPD_PROD': {'ANY': {'ANY': {'n
Ready': 0, 'nRunning': 241, 'nQueue': 59, 'nNewWorkers': 0}}}}
2026-05-15 14:07:19,010 panda.log.worker_adjuster: DEBUG    define_num_workers <site=SPbSU_SPD_PROD> Processing queue SPbSU_SPD_PROD job_type ANY resource_type
 ANY with static_num_workers {'nReady': 0, 'nRunning': 241, 'nQueue': 59, 'nNewWorkers': 0}
2026-05-15 14:07:19,014 panda.log.worker_adjuster: DEBUG    define_num_workers <site=SPbSU_SPD_PROD> No n_new_workers since n_queue(59) + n_ready(0) + n_runnin
g(241) >= max_workers(300)
2026-05-15 14:07:19,014 panda.log.worker_adjuster: DEBUG    define_num_workers <site=SPbSU_SPD_PROD> setting n_new_workers to 0 in order to respect universal m
axNewWorkers
```

SPD offline computing landscape

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

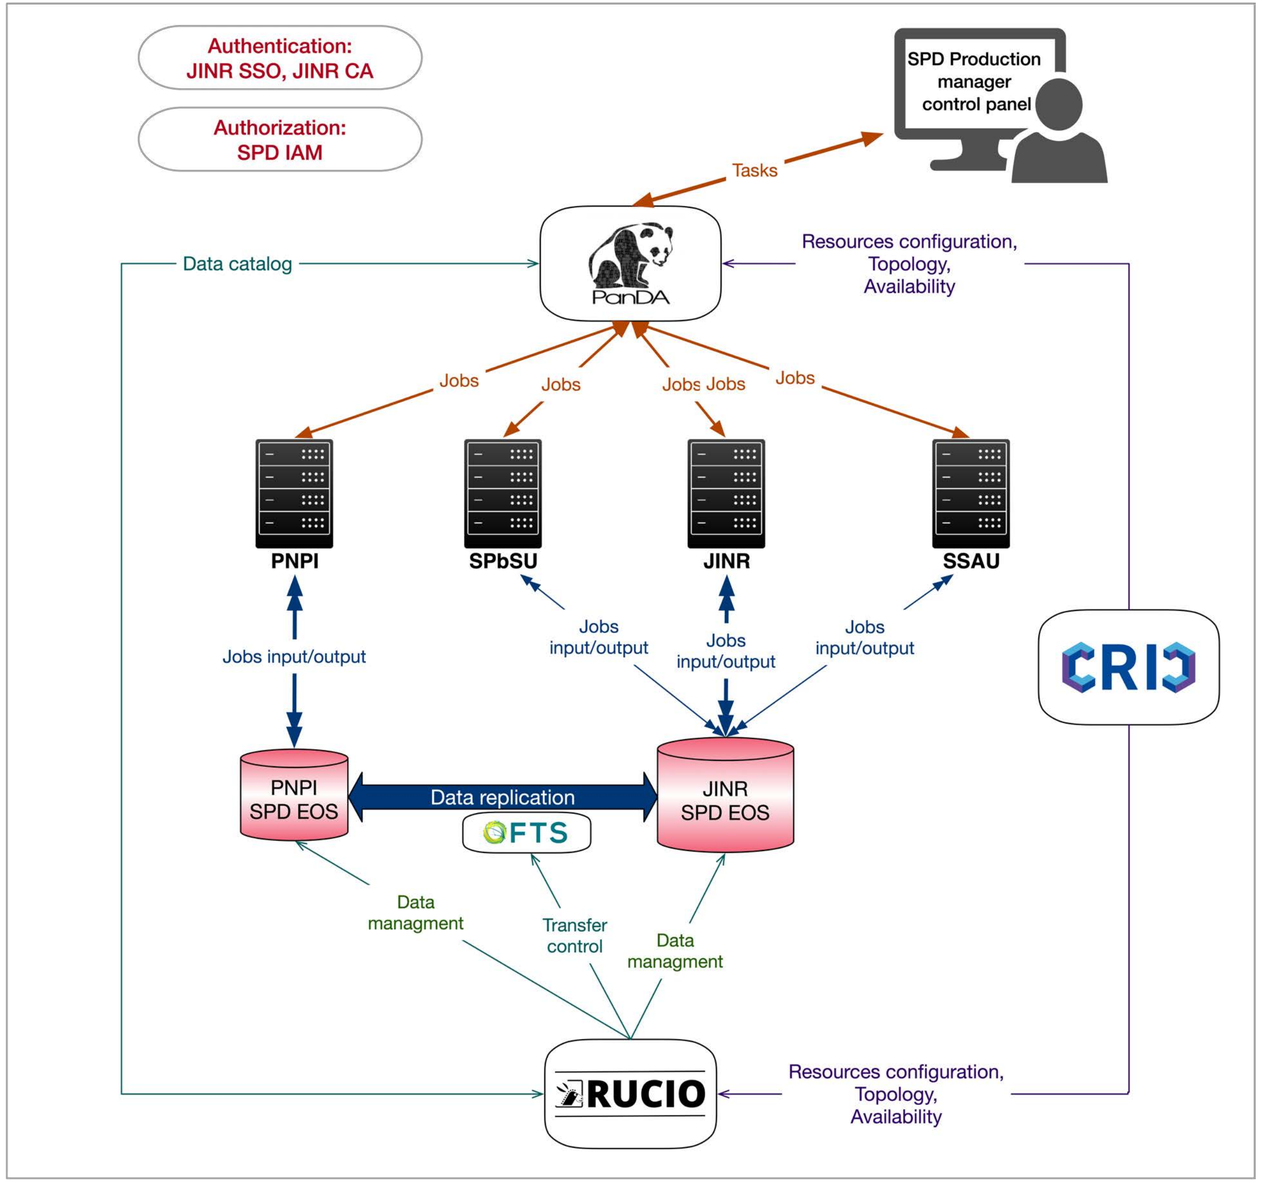

SPD computing cloud

- JINR
  - Computing (2 queues, 3000 cores), but sometimes problematic
  - Storage (7.2 PB raw with 27% redundancy = 5.3 PB)
- PNPI
  - Computing (304 cores)
  - Storage (280 TB redundant)
- SSAU
  - Computing (256 cores)
  - Storage is on the way (240 TB raw with 17% redundancy = 200 TB)
- SPbSU
  - Computing (272 cores)
  - Storage is being configured
- MEPhI
  - Ongoing negotiations

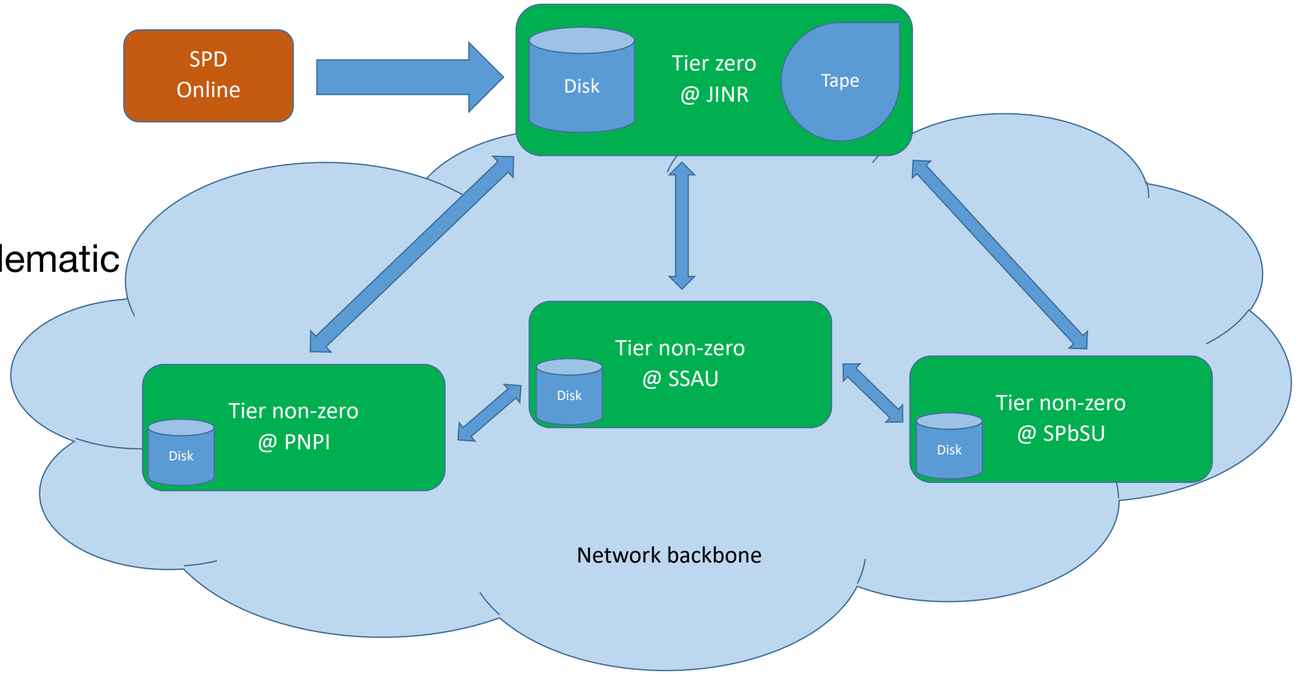

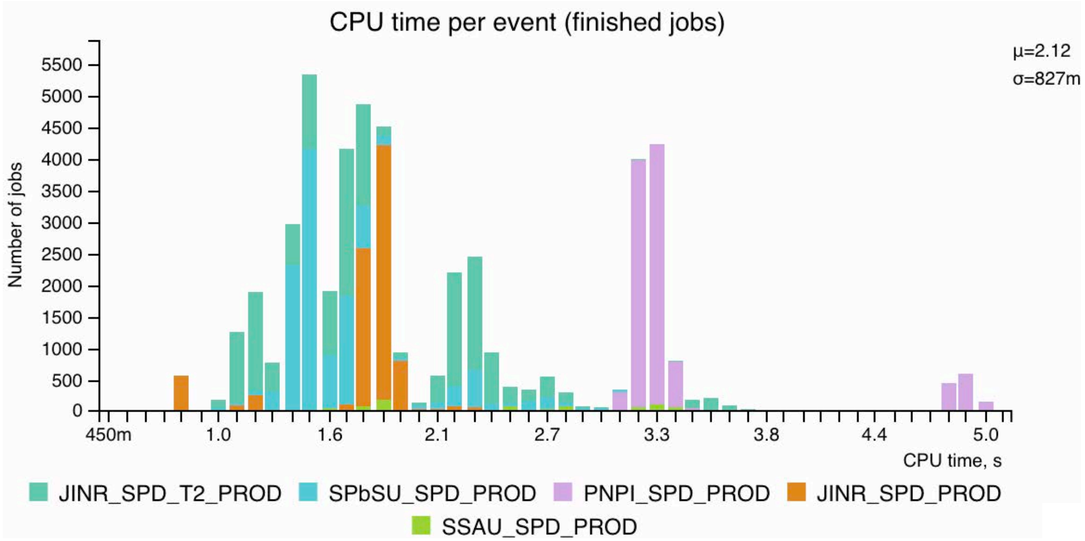

PanDA monitoring

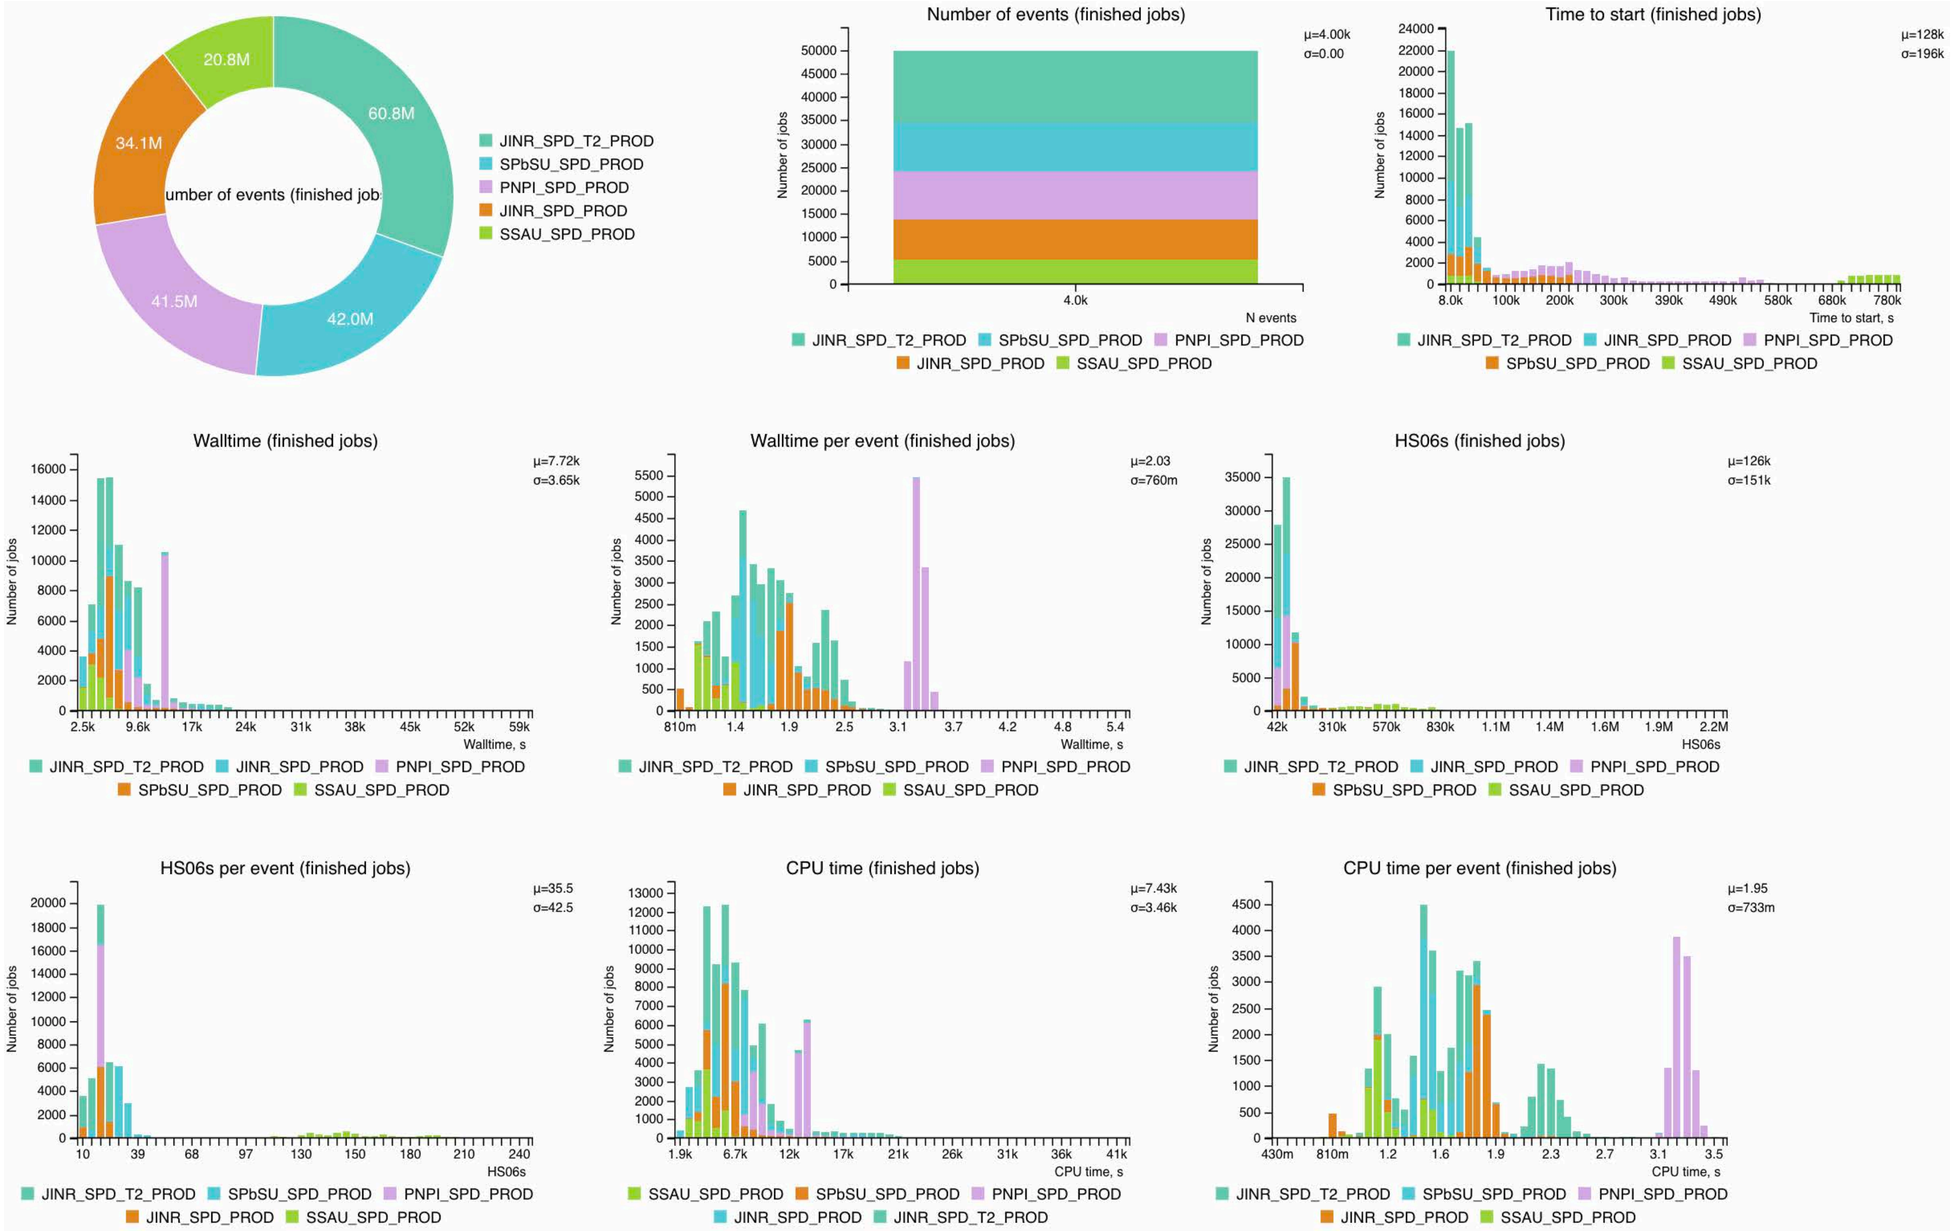

Production statistics

- Successfully processed about 440k jobs across 137 tasks

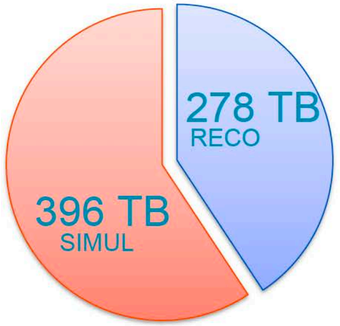

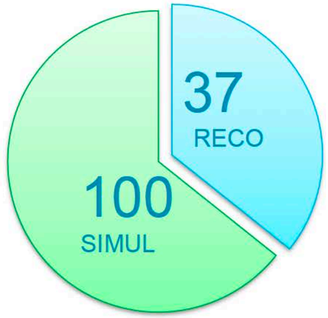

- Total output datasets volume – more than 680 TB

We faced long running productions this spring accompanied by EGI-SVG emails exchange about different vulnerabilities in Linux, Nginx and even in Rucio, and we started thinking about load balancers and high availability configuration of all key systems in order to have ability to upgrade servers without stopping data processing

- Very active productions during last two months
- We collected more data during the last month than ever before and almost reached storage limit in PNPI
- We’re close to 1 PB of data

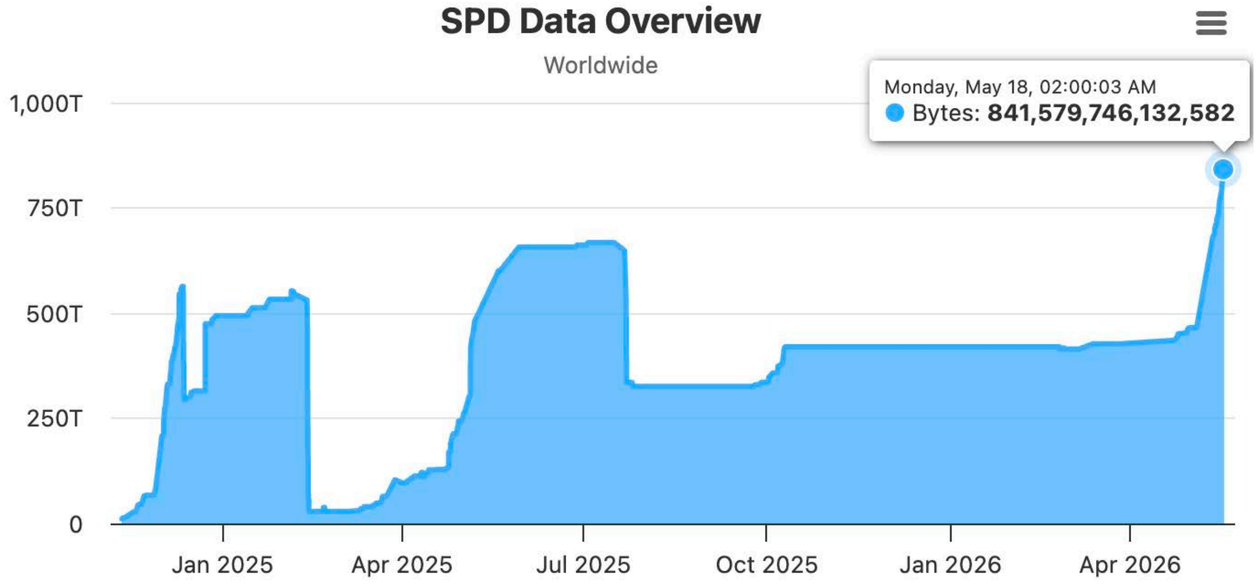

Related talks

- Containers delivery to CVMFS, see talk by Rinat Korotkin
- Services test suite and Rucio monitoring, in the talk by Alexey Konak
- Control panel of the production manager, see talk by Nikita Monakov
- Storage federation proposal, see talk by Andrey Kiryanov

Services test suite

- List of services of computing infrastructure of the SPD experiment is already quite long
- In case of failure it’s hard to quickly identify the source of problems
- In order to have all services and their latest status we decided to develop a test suite which would solve that task
- It’s not just monitoring page, it is proactive testing suite, which runs regular tests, gathers results and presents them at one dashboard with indication about problems and errors
- A notification system can be developed based on this test suite

Monitoring system

The main purposes of creating such monitoring are:

- Collect information about the health and operation of systems in one place
- Visualization of this data for analysis by system administrators

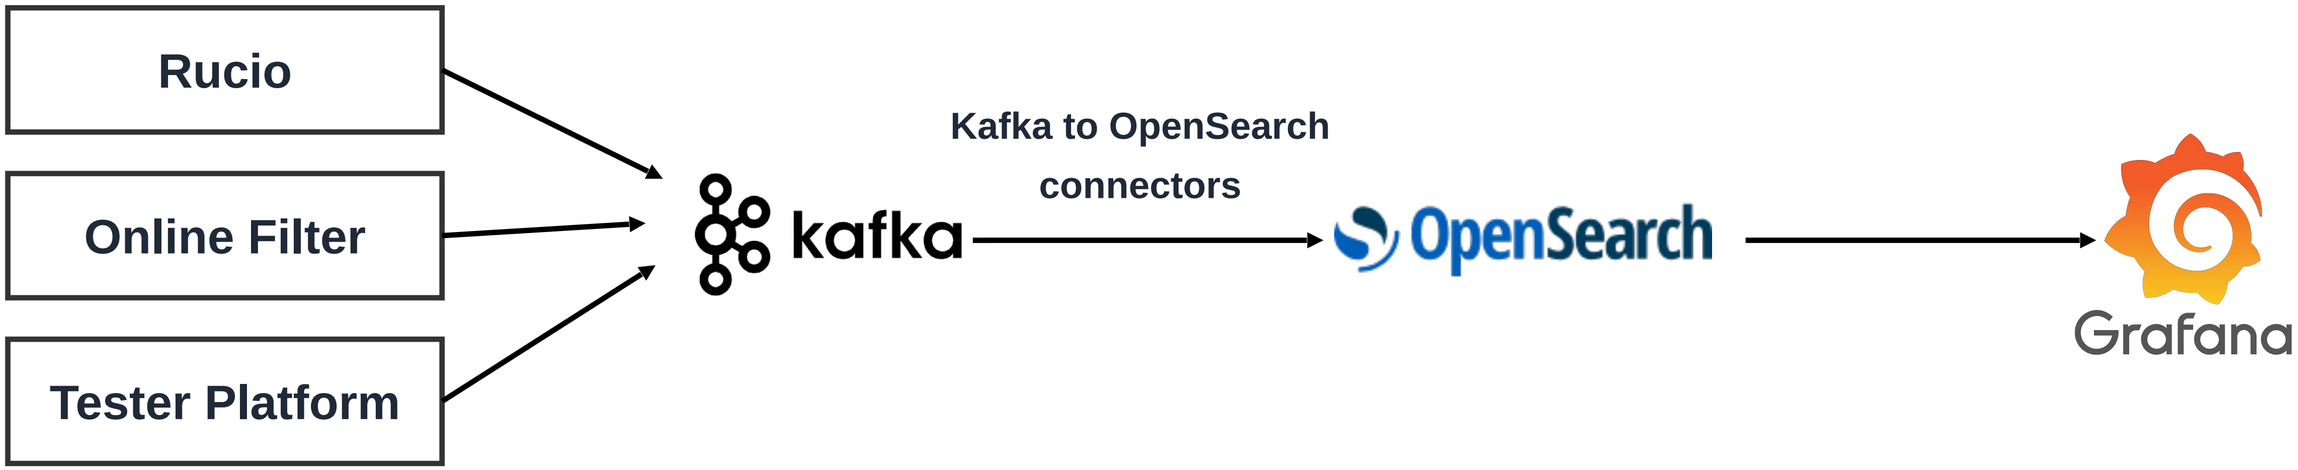

Data sources: collecting traces and metrics from critical experiment nodes and the testing system.

Apache Kafka: provides fault-tolerant reception and buffering of large data streams in real time.

OpenSearch: indexing and long-term storage of monitoring data with the possibility of fast full-text search.

Grafana: building dynamic dashboards and visualizations for operational monitoring of the distributed environment.

Services monitoring

test results

timestamp	site	service...	service_name	result	error_type	mes...	endpoint	check_type	response...
05.05.26 12:24	PNPI	SE-XROOTD	PNPI_EOS_xrootd	success		Alive	root://mss3.pnpi.nw.ru:1094	XRootD Stat	0.726
05.05.26 12:24	JINR	SE-XROOTD	JINR_SPD_EOS-XROOTD	success		Alive	root://eos-spd.jinr.ru:1094	XRootD Stat	0.326
05.05.26 12:24	PNPI	SE-WEBDAV	PNPI_EOS_webdav	success		Alive	davs://mss3.pnpi.nw.ru:8000	WebDAV Stat	0.379
05.05.26 12:24	JINR	SE-WEBDAV	JINR_SPD_EOS-WEBDAV	success		Alive	davs://eos-spd.jinr.ru:8000	WebDAV Stat	0.224
05.05.26 12:24	INPBSU	SE-WEBDAV	INPBSU-SPD-STORM	fail	ServiceNotFound	Failed afte	davs://grid02.hep.by:8443	WebDAV Stat	0.437
05.05.26 12:24	JINR	Rucio	Rucio-prod	success		Alive (Ruci	https://spd-rucio.jinr.ru:443	Rucio Client	0.0217
05.05.26 12:24	JINR	PanDA	PanDA-prod	success		Alive	https://vm221-129.jinr.ru:25443	HTTP	0.0200
05.05.26 12:24	JINR	Harvester	SPD-221-126	success		Alive	https://vm221-126.jinr.ru:8080	HTTP	0.0274
05.05.26 12:24	JINR	FTS	SPD-FTS	success		Alive	https://spd-fts.jinr.ru:8446	HTTP	0.0554
05.05.26 12:24	SSAU	CE	SSAU-CE	fail	Network	Failed afte	https://spd.ssau.ru:443	ARC CE	7.20
05.05.26 12:24	SPbSU	CE	SPbSU-CE	fail	Network	Failed afte	https://alice30.grid.pu.ru:443	ARC CE	6.96
05.05.26 12:24	PNPI	CE	PNPI-CE	success		Alive	https://gt2.pnpi.nw.ru:443	ARC CE	0.387
05.05.26 12:24	JINR	CE	JINR-CE-T2	success		Alive	https://lcgce01.jinr.ru:443	ARC CE	0.267
05.05.26 12:24	JINR	CE	JINR-CE-T1	fail	SSL	Failed afte	https://ce03.jinr-t1.ru:443	ARC CE	2.61
05.05.26 12:24	INPBSU	CE	INPBSU-CE	fail	Network	Failed afte	https://grid04.hep.by:443	ARC CE	7.13

Rucio monitoring

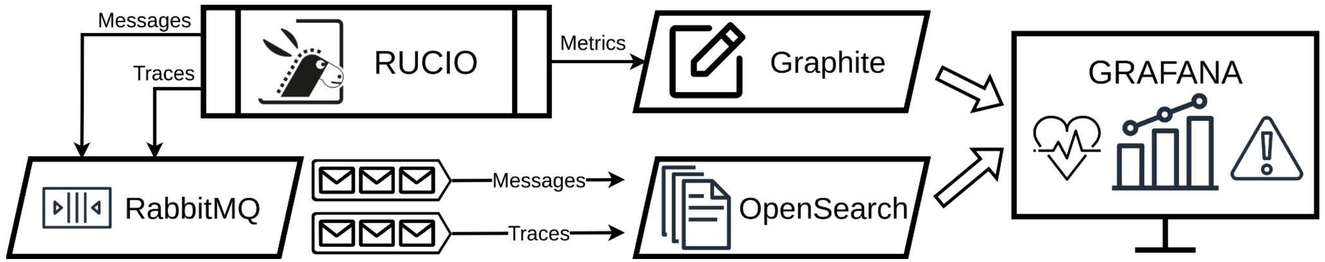

Next steps 1/2

- User interfaces development
  - UI/UX improvement of the Control panel and Production Requests
  - Monitoring and analysis tools
- Integration with the new applied software framework Sampo, transition from SpdRoot to Sampo
- Finish transition to tokens
  - ARC CE-side configuration
- Add more automation for the routine procedures
  - Import users from the SSO database to the IAM automatically basing on LDAP groups
- High availability configuration of the production services
- Migrate from SPD FTS to JINR FTS
- Documentation

Next steps 2/2

- Monitoring coverage for core services
  - Zabbix configuration
  - 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
- Logs processing system

Conclusions

- Production system of the SPD works stable and is being filled by different infrastructure and service elements
- All middleware services deployed at the MLIT cloud platform, and, in common, provides all possible support (many thanks!)
- We are grateful to our colleagues who are sharing their computing resources with us and looking forward to have new participants
- There are activities which are aimed at making the system to be more reliable

Thank you!