

PanDA@JINR (COMPASS and more)

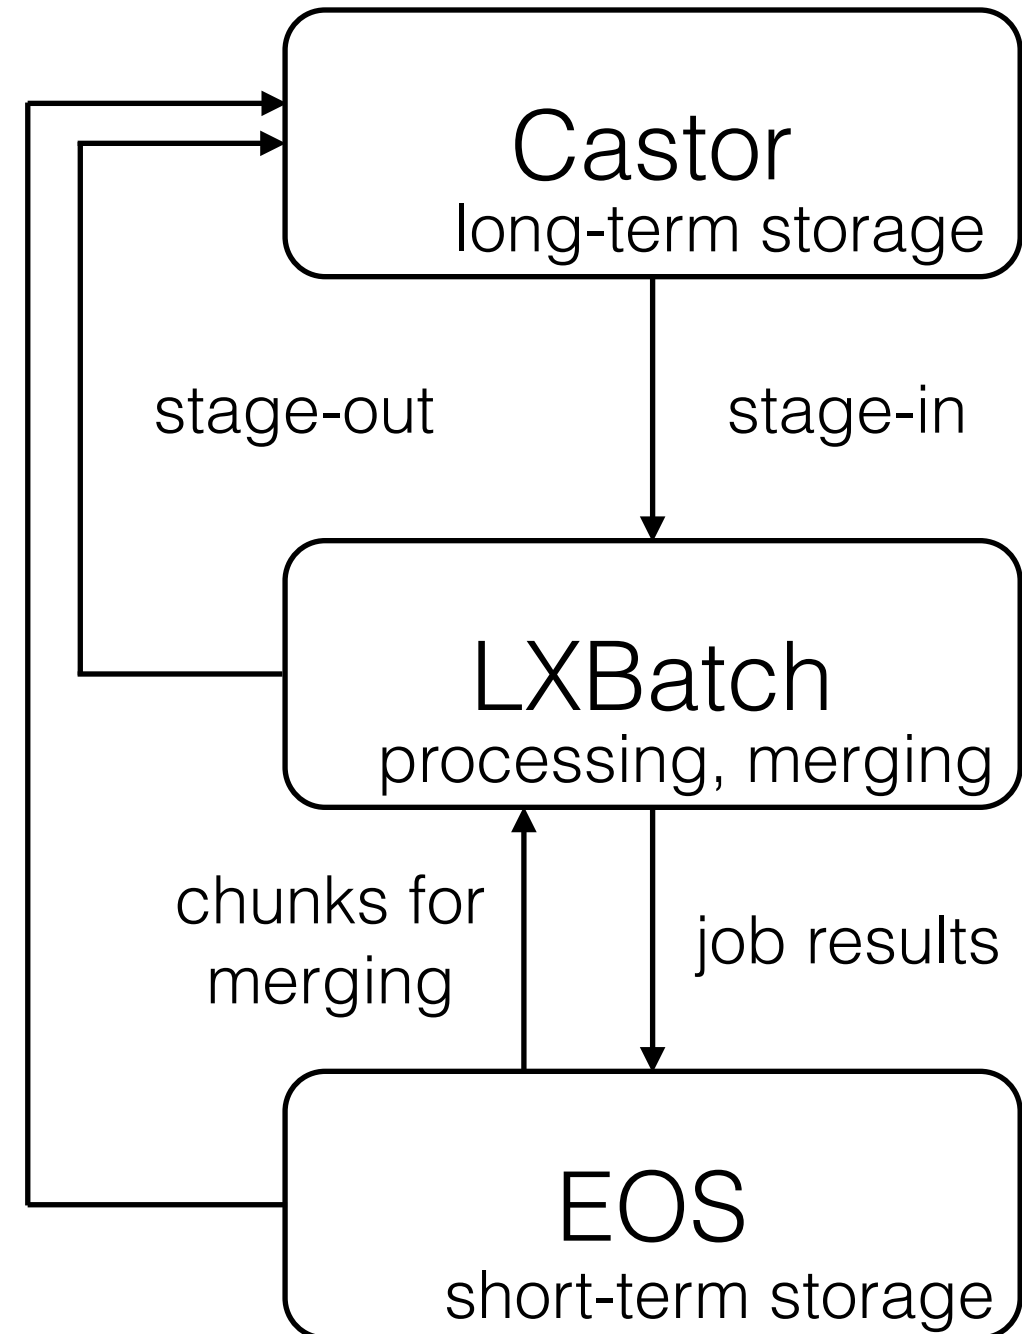
Artem Petrosyan
September 24, 2017

What is COMPASS

- **C**ommon **M**uon **P**roton **A**pparatus for **S**tructure and **S**pectroscopy (COMPASS) is a high-energy physics experiment at a Super Proton Synchrotron (SPS) at CERN
- The purpose of the experiment is the study of hadron structure and hadron spectroscopy with high intensity muon and hadron beams
- First data taking run started in summer 2002 and sessions are continue
- Each data taking session containing from 1.5 to 3 PB of data
- More than 200 physicists from 13 countries and 24 institutes are the analysis user community of COMPASS

COMPASS production dataflow

- All data stored on Castor
- Data is being requested to be copied from tapes to disks before processing (may take ~6 hours)
- Task moves files directly from Castor to lxbatch for processing, several programs are used for processing
- After processing results are being transferred to EOS for merging or short-term storage or directly to Castor for long-term storage
- Merging
- Results are being copied to Castor for long-term storage



Features of “classic” implementation

- We can run jobs on the only one computing resource and LSF will be decommissioned at the end of 2018
- We strictly connected to AFS, local file system, which will be replaced by EOS
- Strictly connected to CASTOR, which will be replaced by EOS
- User jobs and production jobs are sent directly to computing resources and can not be managed (we can not set priority, quota at user’s level)

Motivation

- Move processing to CERN Condor
 - Even more: get ability to switch computing sites, get more resources, any type, not only LSF
 - Enable processing on Blue Waters HPC
- Get rid of self-written code and start using some "common" solution

Action items to enable processing via PanDA

- PanDA instance installation
- Production chain management software preparation
- Grid environment setup
- Production jobs execution by PanDA expert
- Physics validation
- Production by COMPASS production manager

Grid environment

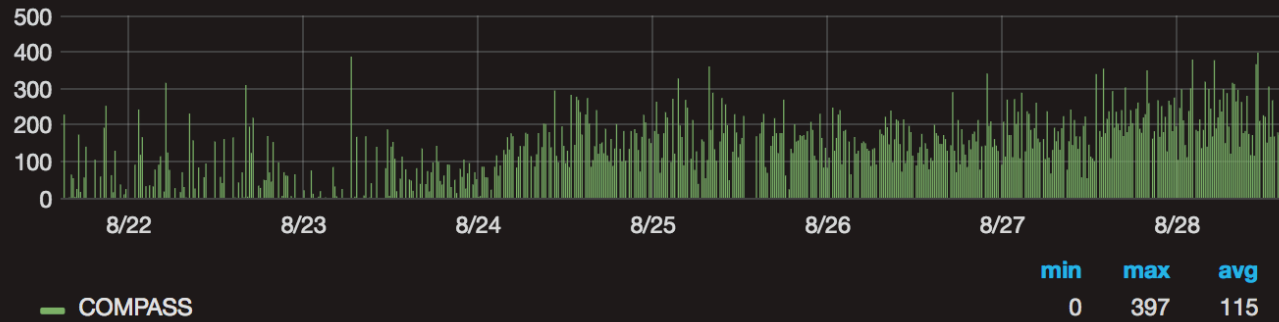
- AFS COMPASS group
 - Production account
- Local batch queue
- EOS directory
- AFS directory to deploy production software
- Virtual organisation
 - Production role
- Computing element
- EOS storage element
- CVMFS

Status

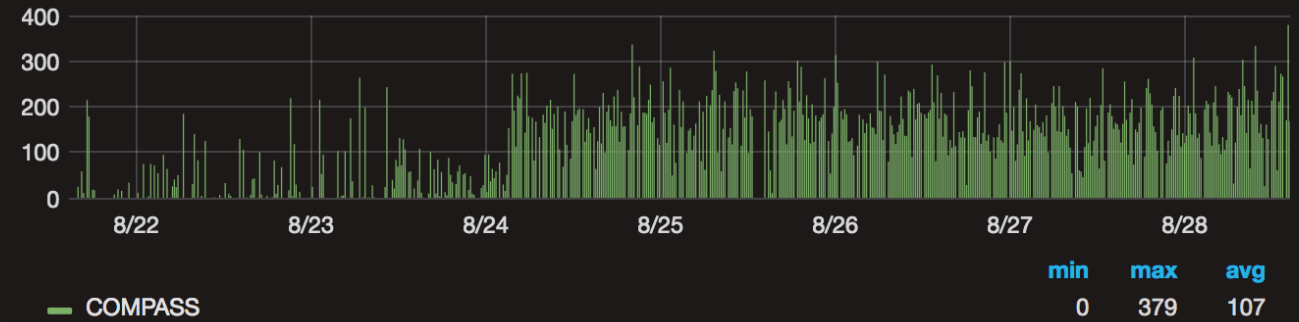
- 3 sites: CERN (Condor), JINR (PBS behind Cream), Trieste (LSF behind Cream)
- 1 storage element: EOS at CERN
- Processing with only one storage allowed to get rid of DDM, files management done by pilot at the stage-out step
- Processing with few computing sites allowed to get rid of ATLAS Grid Information System, network sonars, network closeness monitoring, etc.
- ~12000 simultaneously running jobs
- ~1 million jobs processed during last two months
- PanDA server over MySQL, Monitoring, AutoPilotFactory, Production System deployed in Dubna on production area of our cloud service
- Production system is a brand new one, based on Django framework, prepared to manage COMPASS production chain tasks and jobs — the only thing we had to develop using logic of the previous COMPASS ProdSys
- We also created simple Web UI for schedconfig database

7 days stats


Recent Jobs Started



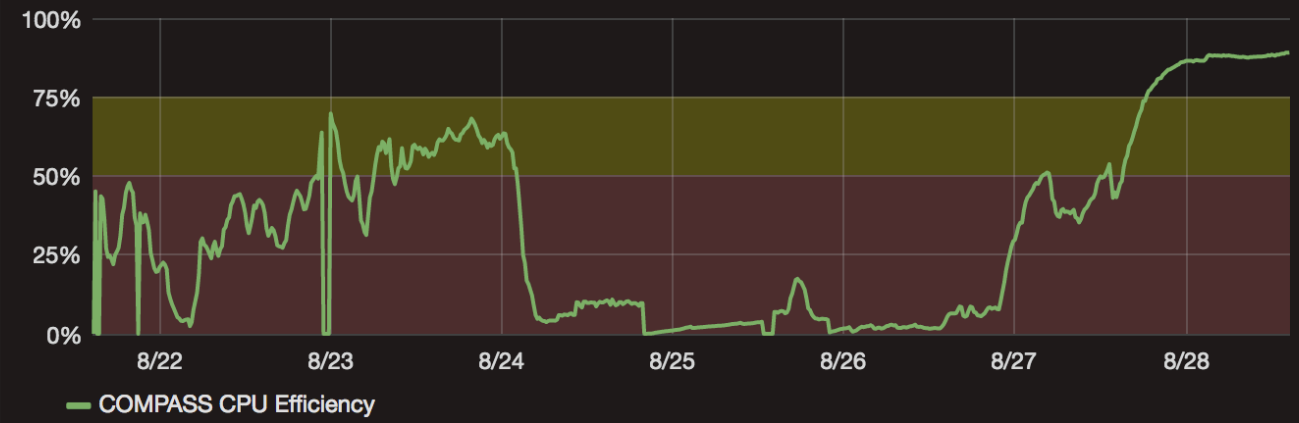
Recent Jobs Submitted



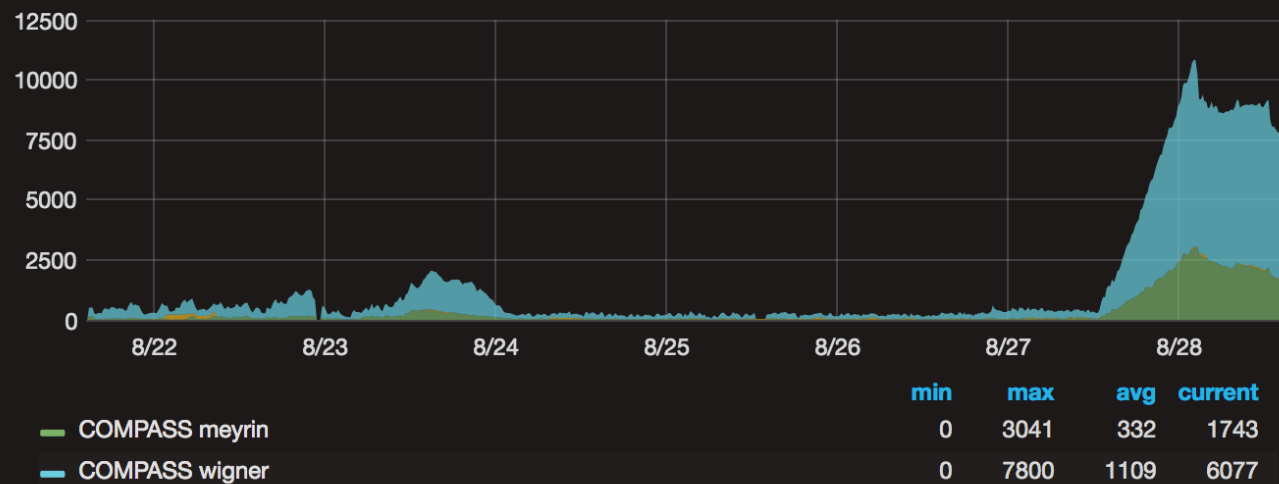
Job Summary



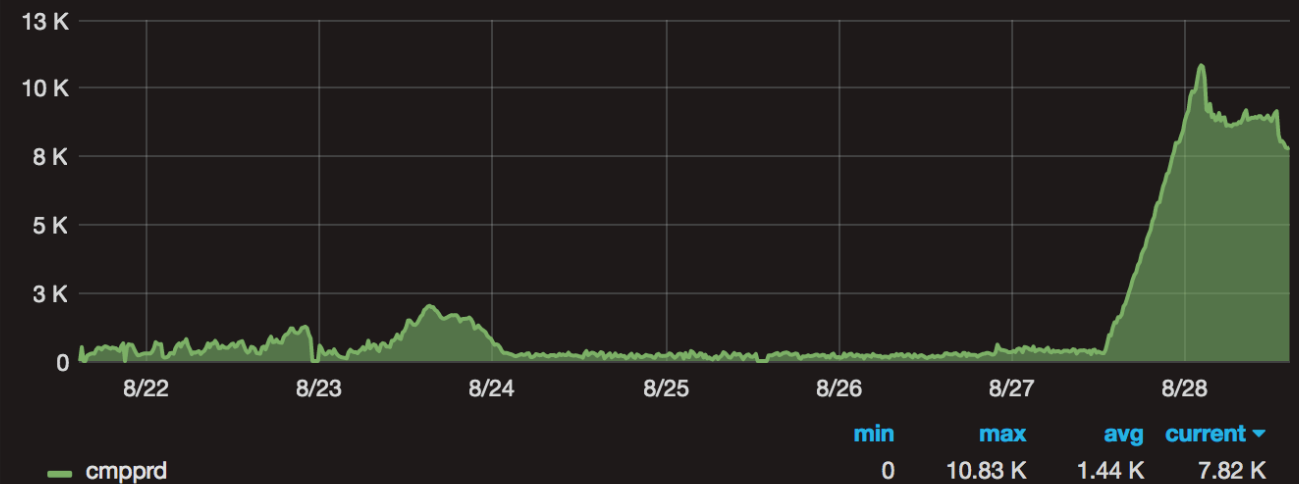
CPU Efficiency



Running by Site



Running Jobs by User




JINR T2 jobs per VO

Resource Centre JINR-LCG2 — Total number of jobs by VO and Month (Official VOs)

Mar 2017	Apr 2017	May 2017	Jun 2017	Jul 2017	Aug 2017
33,069	57,822	37,082	29,131	28,196	26,986
323,132	397,144	366,224	320,417	335,946	308,425
5,079	17,423	54,963	3,277	2,186	1,827
87,329	68,556	48,814	46,711	55,061	66,463
0	0	2	0	0	0
11,253	9,313	36,665	66,805	27,778	33,527
47,090	81,684	64,305	55,729	76,062	51,983
15,674	15,441	13,687	12,989	13,476	13,243
0	2	208	0	198	64,802
522,626	647,385	621,950	535,059	538,903	567,256
8.38%	10.38%	9.97%	8.58%	8.64%	9.10%

Factory view

Factory JINR-pandawms
Version 2.4.9
Last startup 1 day ago
Email artem.petrosyan@jinr.ru
Activity  150
Links [logs](#) [queues.conf](#)

Factory label	last msg
CERN_COMPASS_PROD-ce301-cern-ch	19 mins ago
CERN_COMPASS_PROD-ce302-cern-ch	19 mins ago
CERN_COMPASS_PROD-ce401-cern-ch	19 mins ago
CERN_COMPASS_PROD-ce402-cern-ch	19 mins ago
CERN_COMPASS_PROD-ce403-cern-ch	19 mins ago
CERN_COMPASS_PROD-ce404-cern-ch	19 mins ago
CERN_COMPASS_PROD-ce405-cern-ch	19 mins ago
CERN_COMPASS_PROD-ce406-cern-ch	19 mins ago
CERN_COMPASS_PROD-ce407-cern-ch	19 mins ago
CERN_COMPASS_PROD-ce408-cern-ch	19 mins ago
CERN_COMPASS_PROD-ce503-cern-ch	seconds ago
CERN_COMPASS_PROD-ce504-cern-ch	1 min ago
CERN_COMPASS_PROD-ce505-cern-ch	seconds ago
CERN_COMPASS_PROD-ce506-cern-ch	5 mins ago
CERN_COMPASS_PROD-ce507-cern-ch	1 min ago
CERN_COMPASS_PROD-ce508-cern-ch	5 mins ago
CERN_COMPASS_PROD-condorce01-cern-ch	2 mins ago
CERN_COMPASS_PROD-condorce02-cern-ch	4 mins ago
CNAF_COMPASS_PROD-ce04-lcg-cr-cnaf-infn-it	19 mins ago
JINR_COMPASS_PROD-lcgce12-jinr-ru	4 mins ago
JINR_COMPASS_PROD-lcgce21-jinr-ru	seconds ago
TRIESTE_COMPASS_PROD-cel-ts-infn-it	19 mins ago

COMPASS ProdSys administration

Home › COMPASS ProdSys › Tasks › dvcs2016P09t1PANDAcvmfs

Change task

Name:

dvcs2016P09t1PANDAcvmfs

Type:

test production ↕

Home:

/cvmfs/compass.cern.ch/

Path:

generalprod/testcoral/

Soft:

dvcs2016P09t1PANDAcvmfs

Period:

ProdSlit:

0

PhastVer:

7

Template:

template_mu-.opt ↕

Filelist:

```
/castor/cern.ch/compass/data/2016/raw/W14/cdr12019-275772.raw
/castor/cern.ch/compass/data/2016/raw/W14/cdr12002-275772.raw
/castor/cern.ch/compass/data/2016/raw/W14/cdr12030-275772.raw
/castor/cern.ch/compass/data/2016/raw/W14/cdr12031-275772.raw
/castor/cern.ch/compass/data/2016/raw/W14/cdr11004-275772.raw
/castor/cern.ch/compass/data/2016/raw/W14/cdr11003-275772.raw
```

Select job to change

ADD JOB +

Action: 0 of 100 selected

<input type="checkbox"/>	TASK	FILE	RUN NUMBER	CHUNK NUMBER	PANDA ID	ATTEMPT	STATUS	PANDA ID MERGING	ATTEMPT MERGING	STATUS MERGING
<input type="checkbox"/>	dvcs2016P09t1PANDAcvmfs	/castor/cern.ch/compass/data/2016/raw/W14/cdr13040-275772.raw	275772	13040	588	1	finished	616	5	finished
<input type="checkbox"/>	dvcs2016P09t1PANDAcvmfs	/castor/cern.ch/compass/data/2016/raw/W14/cdr14037-275772.raw	275772	14037	587	1	finished	616	5	finished
<input type="checkbox"/>	dvcs2016P09t1PANDAcvmfs	/castor/cern.ch/compass/data/2016/raw/W14/cdr14036-275772.raw	275772	14036	586	1	finished	616	5	finished
<input type="checkbox"/>	dvcs2016P09t1PANDAcvmfs	/castor/cern.ch/compass/data/2016/raw/W14/cdr13029-275772.raw	275772	13029	585	1	finished	616	5	finished
<input type="checkbox"/>	dvcs2016P09t1PANDAcvmfs	/castor/cern.ch/compass/data/2016/raw/W14/cdr13028-275772.raw	275772	13028	584	1	finished	616	5	finished
<input type="checkbox"/>	dvcs2016P09t1PANDAcvmfs	/castor/cern.ch/compass/data/2016/raw/W14/cdr13015-275772.raw	275772	13015	583	1	finished	616	5	finished
<input type="checkbox"/>	dvcs2016P09t1PANDAcvmfs	/castor/cern.ch/compass/data/2016/raw/W14/cdr13008-275772.raw	275772	13008	582	1	finished	616	5	finished
<input type="checkbox"/>	dvcs2016P09t1PANDAcvmfs	/castor/cern.ch/compass/data/2016/raw/W14/cdr13004-275772.raw	275772	13004	581	1	finished	616	5	finished
<input type="checkbox"/>	dvcs2016P09t1PANDAcvmfs	/castor/cern.ch/compass/data/2016/raw/W14/cdr11001-275772.raw	275772	11001	580	1	finished	616	5	finished
<input type="checkbox"/>	dvcs2016P09t1PANDAcvmfs	/castor/cern.ch/compass/data/2016/raw/W14/cdr11013-275772.raw	275772	11013	579	1	finished	616	5	finished
<input type="checkbox"/>	dvcs2016P09t1PANDAcvmfs	/castor/cern.ch/compass/data/2016/raw/W14/cdr13021-275772.raw	275772	13021	578	1	finished	616	5	finished

Site administration

AUTHENTICATION AND AUTHORIZATION

Groups [+ Add](#) [✎ Change](#)

Users [+ Add](#) [✎ Change](#)

COMPASS PRODSYS

Jobs [+ Add](#) [✎ Change](#)

Tasks [+ Add](#) [✎ Change](#)

COMPASS SCHEDCONFIG

Cloudconfig [+ Add](#) [✎ Change](#)

Schedconfig [+ Add](#) [✎ Change](#)

Siteaccess [+ Add](#) [✎ Change](#)

Users [+ Add](#) [✎ Change](#)

DJANGO_CELERY_BEAT

Crontabs [+ Add](#) [✎ Change](#)

Intervals [+ Add](#) [✎ Change](#)

Periodic tasks [+ Add](#) [✎ Change](#)

DJANGO_CELERY_RESULTS

Task results [+ Add](#) [✎ Change](#)

Recent actions

My actions

[✎ /castor/cern.ch/compass/data/2...
278570.raw](#)

Job

[✎ /castor/cern.ch/compass/data/2...
278679.raw](#)

Job

[✎ /castor/cern.ch/compass/data/2...
278679.raw](#)

Job

[✎ /castor/cern.ch/compass/data/2...
278679.raw](#)

Job

[✎ CN-Elena_Zemlyanichkina__CN-
653827__CN-elenazem__OU-
Users__OU-Organic_Units__DC-
cern__DC-ch](#)

User

[✎ dvcs2017P01t1_mu+_part2](#)

Task

[✘ CN-Artem_Petrosyan__CN-
631691__CN-virthead__OU-
Users__OU-Organic_Units__DC-
cern__DC-ch](#)

User

[✘ Daniil_Malevanniy](#)

User

[✘ Artem_Petrosyan](#)

User

[✎ dvcs2017P01t1_mu+_part2](#)

Task

COMPASS ProdSys administration

Home › COMPASS ProdSys › Jobs

Select job to change

1 result (299593 total)

Action: ----- 0 of 1 selected

<input type="checkbox"/>	TA	E	RUN NUMBER	CHUNK NUMBER	PANDA ID	ATTEMPT	STATUS
<input type="checkbox"/>	dvcs2017P01t1_mu+_part1	/castor/cern.ch/compass/data/2017/raw/W04/cdr12116-278485.raw	278485	12116	1114334	0	defined

1 job

TEMPLATE	STATUS	JOBS	MERGING MDST
template.opt	running	all: 10131, sent: 0, running: 35, failed: 154, check: 2, finished: 9940	sent: 0, running: 0, failed: 0, finished: 3236
template.opt	running	all: 23141, sent: 0, running: 1, failed: 95, check: 0, finished: 23045	sent: 0, running: 0, failed: 0, finished: 15475
template.opt	send	all: 13433, sent: 0, running: 0, failed: 86, check: 8, finished: 13339	sent: 0, running: 19, failed: 113, finished: 9655
template.opt	running	all: 19464, sent: 0, running: 14, failed: 210, check: 0, finished: 19240	sent: 0, running: 65, failed: 0, finished: 11746
template.opt	running	all: 20906, sent: 0, running: 2, failed: 0, check: 0, finished: 20888	sent: 0, running: 0, failed: 0, finished: 18886
template.opt	send	all: 184, sent: 0, running: 0, failed: 0, check: 0, finished: 184	sent: 0, running: 0, failed: 0, finished: 178
template_mu-.opt	done	all: 453, sent: 0, running: 0, failed: 0, check: 0, finished: 453	sent: 0, running: 0, failed: 0, finished: 0
template_mu+.opt	running	all: 26499, sent: 7, running: 47, failed: 171, check: 0, finished: 26239	sent: 0, running: 42, failed: 0, finished: 12425
template_mu-.opt	running	all: 12544, sent: 0, running: 0, failed: 0, check: 0, finished: 12516	sent: 0, running: 44, failed: 0, finished: 9124
template_mu+.opt	running	all: 16681, sent: 0, running: 4, failed: 5, check: 0, finished: 16582	sent: 0, running: 0, failed: 0, finished: 11453

278 jobs selected
User: [Artem Petrosyan](#)
Task ID: 9

Job modification times in this listing range from Sept. 6, 2017, 12:40 a.m. to Sept. 11, 2017, 12:50 p.m.

Job current priorities in this listing range from 1000 to 1000

computingsite	CERN_COMPASS_PROD (278)
destinationse	local (278)
jobstatus	failed (107) finished (171)
prodsourcelabel	prod_test (278)
produsername	Artem Petrosyan (278)
taskid	9 (278)
transformation	merging mdst (278)
vo	vo.compass.cern. (278)

Owner / VO	Task ID	PanDA ID	Transformation	Status	Created	Start	End	Site
Artem Petrosyan / vo.compass.cern.	9	1252607	merging mdst	finished	2017-09-11 11:30	09-11 11:30	09-11 11:39	CERN_COMPASS_PROD
Artem Petrosyan / vo.compass.cern.	9	1252606	merging mdst	finished	2017-09-11 11:30	09-11 11:30	09-11 11:52	CERN_COMPASS_PROD
Artem Petrosyan / vo.compass.cern.	9	1252605	merging mdst	finished	2017-09-11 11:30	09-11 11:30	09-11 11:52	CERN_COMPASS_PROD
Artem Petrosyan / vo.compass.cern.	9	1252604	merging mdst	finished	2017-09-11 11:30	09-11 11:30	09-11 12:46	CERN_COMPASS_PROD
Artem Petrosyan / vo.compass.cern.	9	1252603	merging mdst	finished	2017-09-11 11:30	09-11 11:30	09-11 11:37	CERN_COMPASS_PROD
Artem Petrosyan / vo.compass.cern.	9	1252602	merging mdst	finished	2017-09-11 11:28	09-11 11:28	09-11 11:31	CERN_COMPASS_PROD

Status **finished** indicates that the job has successfully completed.

View the job's [stdout](#), [job outputs](#)

[Download the job cache tarball](#) containing the job execution scripts

View the pilot job's [stdout](#), [stderr](#), [batch log](#)

Job files

Filename (Type)	Size (bytes)	Status	Dataset
275772-13040-0-7_a9b05321-21c2-4d8f-b9c6-4b3a705692eb.job.log.tgz (log)	24292	ready	panda.destDB.a9b05321-21c2-4d8f-b9c6-4b3a705692eb
mDST-275772-13040-0-7.root (output)	43628441	ready	panda.destDB.a9b05321-21c2-4d8f-b9c6-4b3a705692eb
275772-13040-0.root (output)	5752274	ready	panda.destDB.a9b05321-21c2-4d8f-b9c6-4b3a705692eb
testevtdump.raw (output)	2601980	ready	panda.destDB.a9b05321-21c2-4d8f-b9c6-4b3a705692eb
payload_stdout.txt (output)	55306	ready	panda.destDB.a9b05321-21c2-4d8f-b9c6-4b3a705692eb
payload_stderr.txt (output)	777494	ready	panda.destDB.a9b05321-21c2-4d8f-b9c6-4b3a705692eb

Other key job parameters

Job type	prod_test
Payload script (transformation)	cdr13040-275772.raw;
# events	18643
Output destination	local
CPU consumption time (s)	1727
Job metrics	nEvents=18643

PanDA monitor

Job details for PanDA job 616

Owner / VO	PandaID	TaskID	Status	Created	Start	End	Site
Artem Petrosyan / vo.compass.cern.	616	2	finished	2017-06-27 13:14	06-27 13:17	06-27 13:44	CERN_COMPASS_PROD

Job name: dvcs2016P09t1PANDAcvmfs-275772-0-7-merging

Status **finished** indicates that the job has successfully completed.

View the job's [stdout](#), [job outputs](#)

[Download the job cache tarball](#) containing the job execution scripts

View the pilot job's [stdout](#), [stderr](#), [batch log](#)

Job files

Filename (Type)	Size (bytes)	Status	Dataset
dvcs2016P09t1PANDAcvmfs-275772-0-7-merging_e8e9e34d-91f0-478b-9d3a-5fe5985ffbc6.job.log.tgz (log)	20814	ready	dvcs2016P09t1PANDAcvmfs.275772-0-7-merging
mDST-275772-0-7.root.023 (output)	1905658237	ready	dvcs2016P09t1PANDAcvmfs.275772-0-7-merging
payload_stdout.txt (output)	17624	ready	dvcs2016P09t1PANDAcvmfs.275772-0-7-merging

Other key job parameters

Job type	prod_test
Payload script (transformation)	echo \$PWD;
Output destination	local
CPU consumption time (s)	741

Plans

- Migrate all types of data processing to new system
 - Production, merging are there already
 - Daq data decoding is in development
 - MC
 - User analysis
- Prepare COMPASS-specific monitoring
 - Tatiana Korchuganova already at CERN
 - Wish list includes COMPASS-specific items to be presented on the monitoring pages, such as data taking periods, years, runs, etc.
 - VM prepared at JINR
- Enable processing on BlueWaters HPC
 - There is a team adopting COMPASS software to run on HPC, looks like pretty soon it will be ready to be run by PanDA
- Discussing possibility of adding Rucio to organise data

COMPASS issues

- No files catalog, data of different periods stored in different databases: several Oracle schemas, now in MySQL
- Dataflow of MC is not defined
- Software fails without sending exit codes, thus I have to analyse logs to understand whether the job finished correctly or not
- Strong connection to AFS, configuration database contains links to the files at AFS, simply moving to EOS will not help because all the software will have to be adopted to start reading from EOS; upload to CVMFS is more preferable but requires update of the procedure:
 - reduce number of persons who have access to the db
 - organise update sessions: after db update sw update should be performed
- etc, every day I receive new requests, details on the next slide

To do list for COMPASS production via PanDA

<u>task</u>	<u>sub-task</u>	<u>sub-sub-task</u>	<u>details</u>	<u>date start</u>	<u>date end</u>	<u>status</u>	
<u>PanDA</u>	use internal mover	for output files	to EOS			done	
		for input files	from CASTOR/EOS				
Test production	data reconstruction	CORAL+PHAST	reconstruction			done	
			check of all produced chunks : End of Job		2017-09-18	done	
		a possibility to resend jobs by hand				done	
		to have copies of log files in a tar archive	/eos/compass/generalprod/testcoral/PanDALogs/			done	
	error handling	A FATAL ERROR APPEARED	to manual check			done	
		(core dumped) \$CORAL/./phast/coral/coral.exe	to manual check			done	
		Aborted \$CORAL/./phast/coral/coral.exe	to manual check? restart?		2017-09-09		
		char* exception: <u>DaqEventManager::GetEvent(): no event</u>	to manual check			done	
		** TLattice::TLattice: Cannot Allocate 132503040 Bytes in Memory!	job needs to be restarted			done	
		Expected output file testevtdump.raw does not exist	always create testevtdump.raw			2017-09-07	done
		Encountered zero file size for file testevtdump.raw	check size of testevtdump.raw		2017-09-09	2017-09-09	done
		Abnormal job termination. Exception #2 had been caught	job needs to be restarted		2017-09-15		testing
	merging	mDST	with PHAST: phast.exe -m				done
			check : End of Job			2017-09-18	done
			check : compare number of events			2017-09-20	done
a possibility to remerge files					2017-09-08	2017-09-08	done

JINR issues

- Used to work on old hardware, during spring we migrated to brand new hardware installed in the same area of our computing site with T2, which significantly improved stability of cloud services
- I installed Perfsonar service to monitor network connectivity
- There is a noticed JINR cloud network connectivity latency delay and low bandwidth, details on the next slide
 - Yesterday we agreed that new network switch will be installed for cloud service

perfSONAR test results

[Share/open in new window](#)

Source

cloud-perfsonar.jinr.ru
159.93.220.15

[Host info](#) ▾

Destination

perfsonar-bw.cern.ch
128.142.223.246

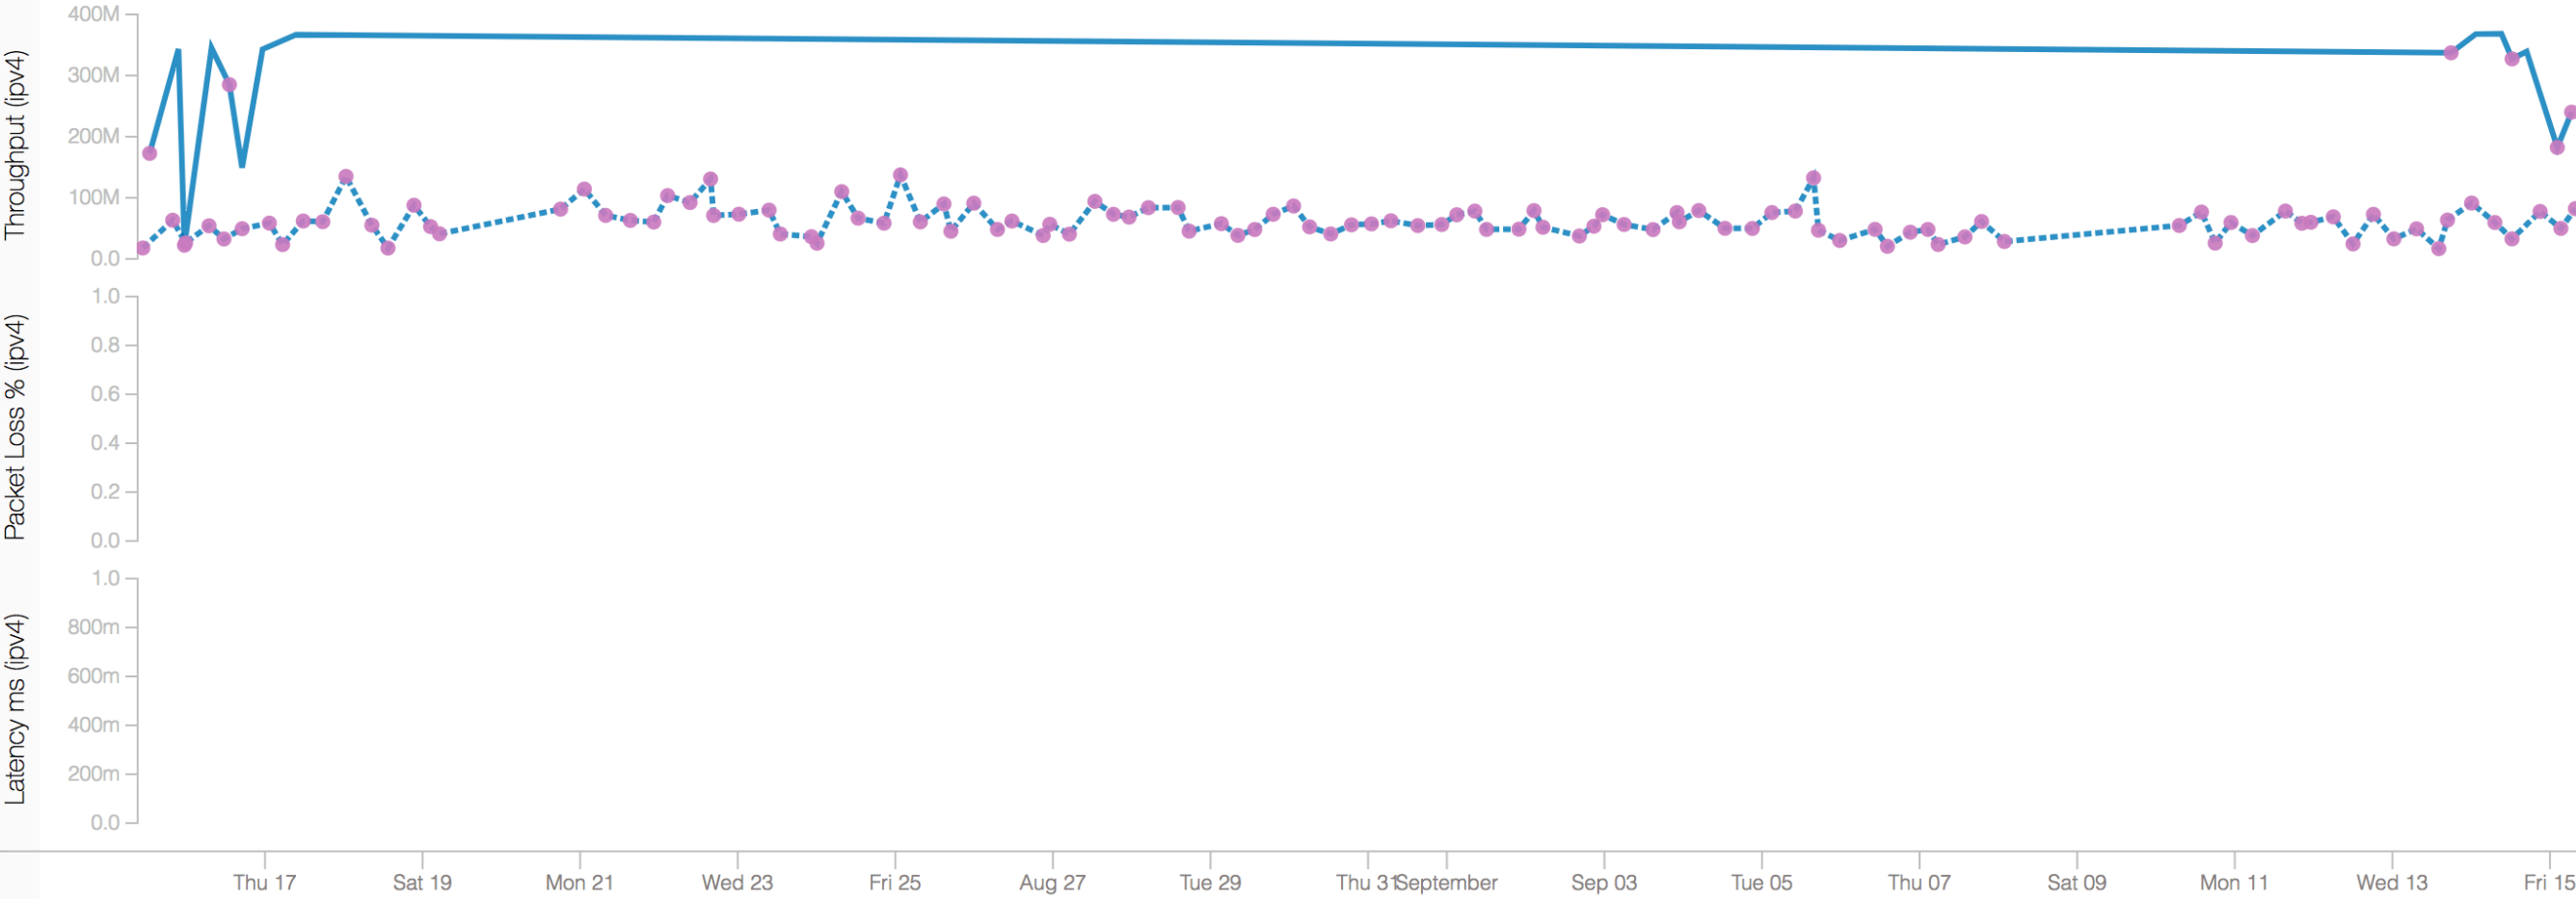
[Host info](#) ▾

Report range

← 1 month →

Tue 08/15/2017 09:12:26 (GMT+2) to Fri 09/15/2017 09:12:26 (GMT+2)

Tput (TCP) Tput (UDP) Loss (UDP) Loss (owamp) Loss (ping) Retrans Latency Latency (ping) Forward Reverse Failures



PanDA issues

- PanDA server
 - No versions, thus no safe update
 - Code in Git does not work in MySQL
- Pilot
 - We need to organise error codes in a way so that they would have a leading part to indicate experiment, screenshot on the next slide
- DB
 - Available db schema for MySQL does not correspond to code base in Git
- ProdSys
 - No generic production system which could be used by any project to define a workflow, each has to start from scratch

```
120 ERR_FILEEXIST = 1221
121 ERR_GETKEYPAIR = 1222
122 ERR_BADALLOC = 1223
123 ERR_ESRECOVERABLE = 1224
124 ERR_ESMERGERECOVERABLE = 1225
125 ERR_GLEXEC = 1226
126 ERR_ESATHENAMPDIED = 1227
127 ERR_ESFATAL = 1228
128 ERR_TEFATAL = 1229
129 ERR_TEBADURL = 1230
130 ERR_TEINVALIDGUID = 1231
131 ERR_TEWRONGGUID = 1232
132 ERR_TEHOSTNAME = 1233
133 ERR_AFATALERRORAPPEARED = 1235
134 ERR_COREDUMPED = 1236
135 ERR_NOEVENTS = 1237
136 ERR_ABNORMALJOBTERMINATION = 1238
137 ERR_ABORTED = 1239
138 ERR_NOENDOFJOB = 1240
139 ERR_CANNOTALLOCATE = 1241
140
141 # internal error codes
142 ERR_DDMREG = 1
143 ERR_FILEONTAPE = 2
144
145 pilotError = {
146     ERR_UNKNOWNERROR : "",
147     ERR_GENERALERROR : "General pilot error, consult batch log",
148     ERR_GETDATAEXC : "Get function can not be called for staging input file",
149     ERR_NOLOCALSPACE : "No space left on local disk",
150     ERR_STAGEINFAILED : "Get error: Staging input file failed",
151     ERR_REPNOTFOUND : "Get error: Replica not found",
152     ERR_LRCREGCONNREF : "LRC registration error: Connection refused",
153     # 1102 : "Expected output file does not exist", # not used, see ERR_MISSINGOUTPUTFILE below
154     ERR_NOSUCHFILE : "No such file or directory",
155     ERR_USERDIRTOOLARGE : "User work directory too large",
156     ERR_LFCADDCSUMFAILED : "Put error: Failed to add file size and checksum to LFC",
157     ERR_STDOUTTOOBIG : "Payload stdout file too big",
```

Other activity

- NA62
 - Initiated by CERN IT, Condor service support
 - First contact this summer with Russian colleagues, after which they introduced me to NA62 community
 - I did a presentation about COMPASS, PanDA at CERN in September
 - Technical meeting with Paul Laycock
 - Agreement is ready that we will test NA62 job submission via PanDA server at JINR
 - I sent him PanDA client package
 - I will define a queue for tests once he sends me details of their VO and computing sites, looks like VO is not even registered at CERN
- NICA
 - No real development from our side at the moment