


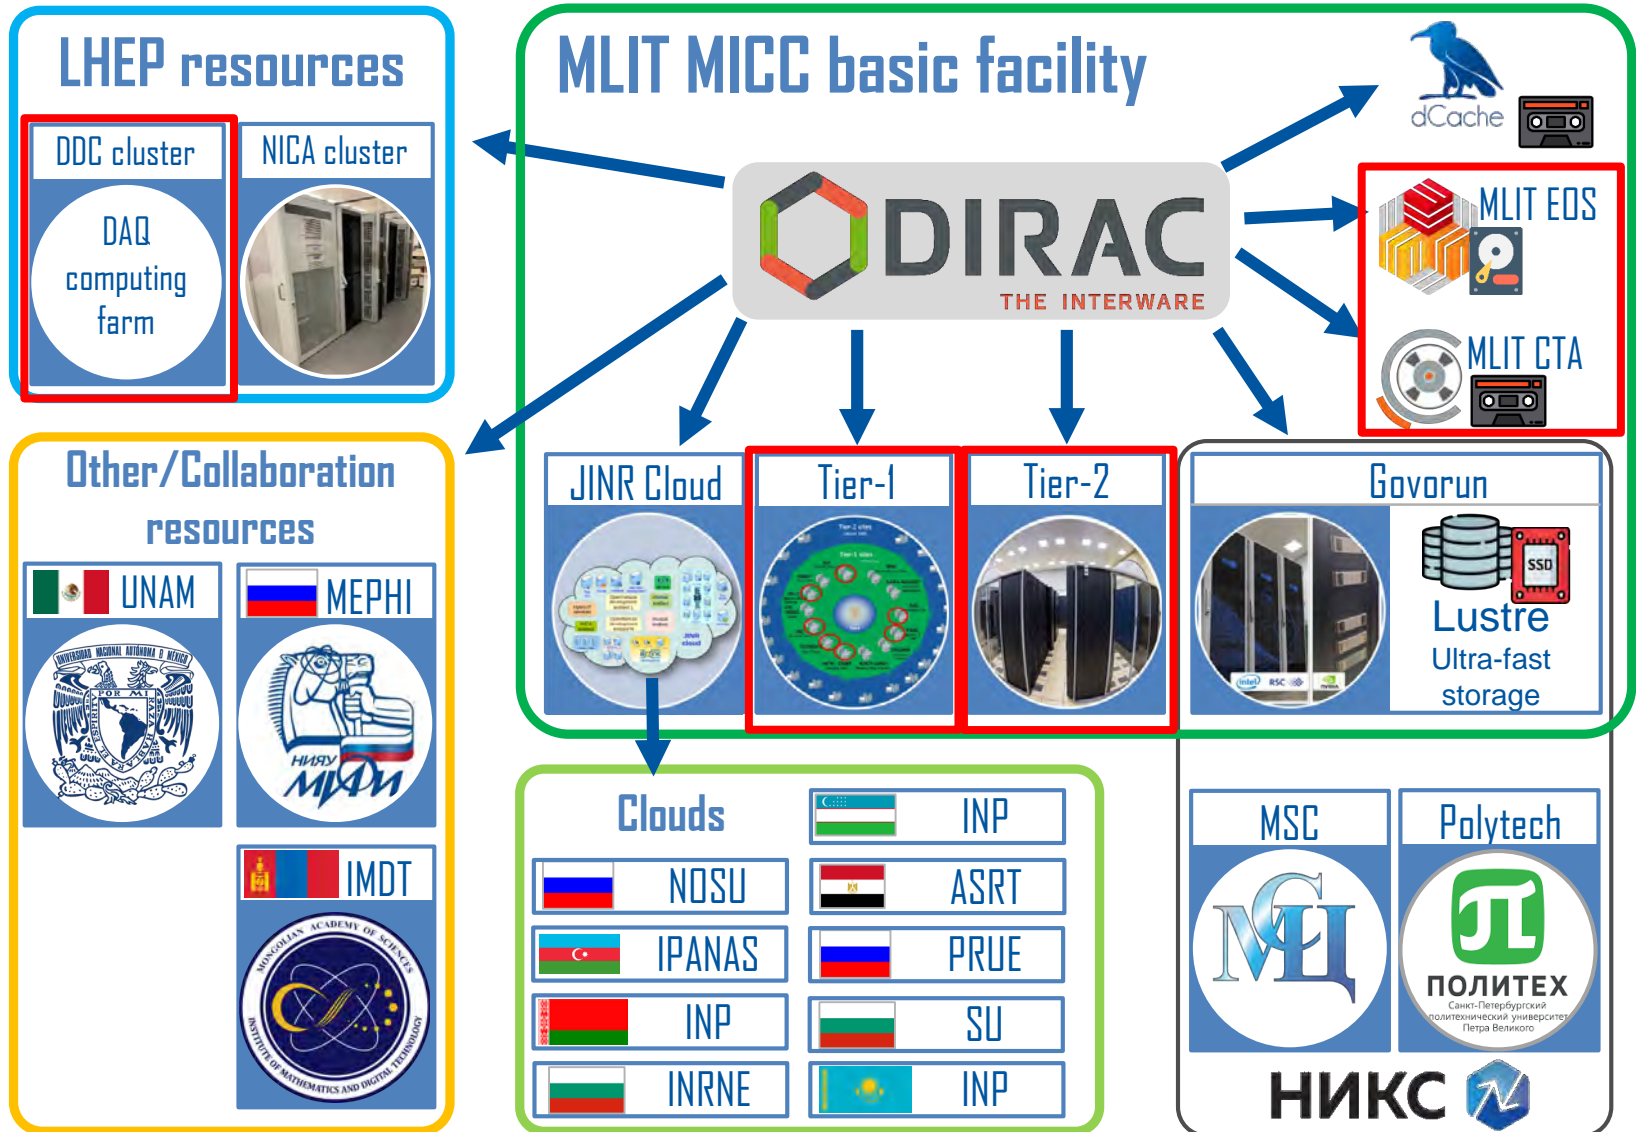
16th BM@N collaboration meeting

Processing of Run9 data on a DIRAC distributed infrastructure

Konstantin Gertsenberger, Igor Pelevanyuk
LHEP MLIT



DIRAC in JINR



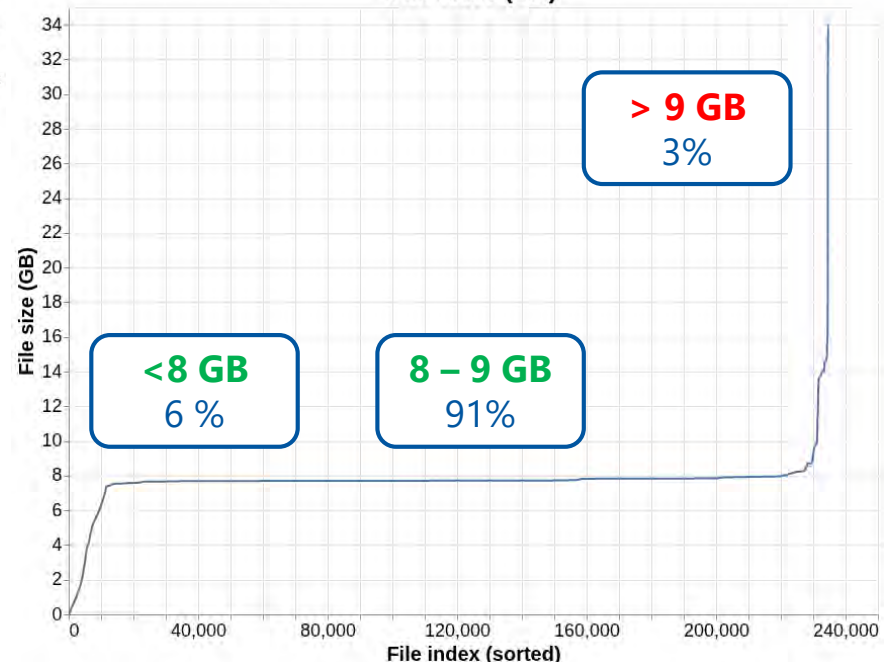
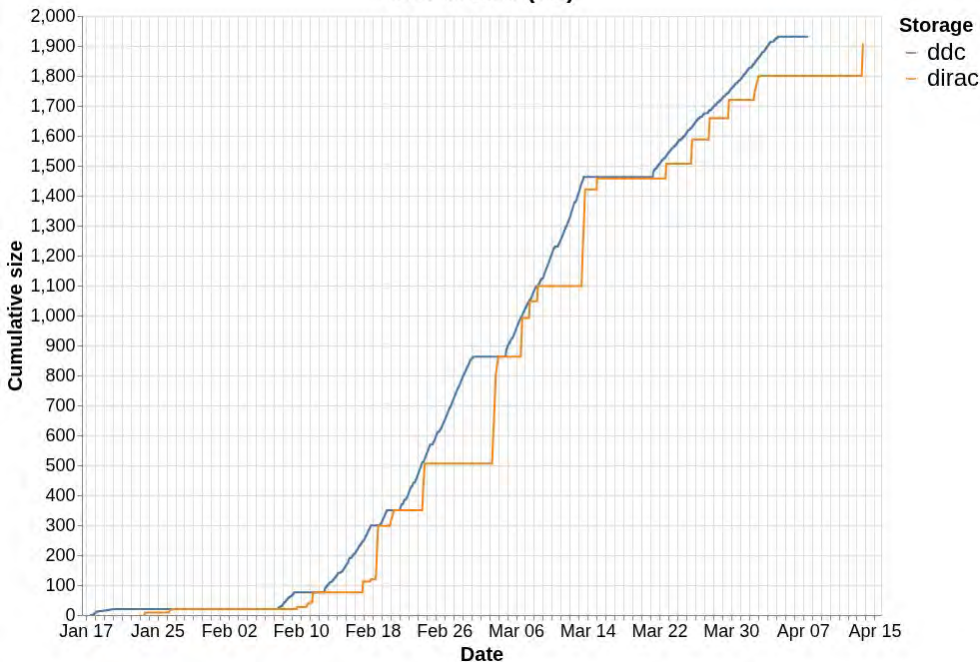
Run 9 in distributed infrastructure

some statistics related to raw data

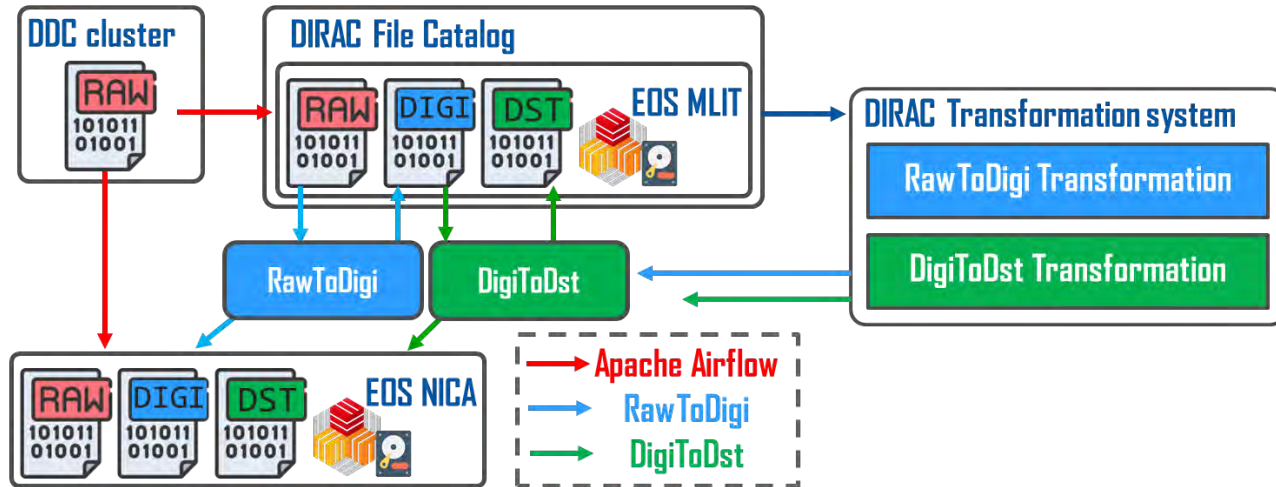
	DIRAC	DDC	Run8
Files - RAW	232349	234616	30003
Size	1906 TB	1925 TB	417 TB
Avg RAW size	8,2	8,2	13,9

Size of files (TB)

File sizes (GB)




Automation ("Triggers+Observer")



This schema was tested during data taking but showed problems while being overloaded with waiting jobs. Full scale use – postponed.


Cascade job initialization was applied. Transfer job that copied data from DDC to EOS MLIT initiated RawToDigi. RawToDigi initiated DigiToDst.



This approach proved to be simple, flexible and controllable for the price of adding a delay between appearance of Raw files and initiation of processing.

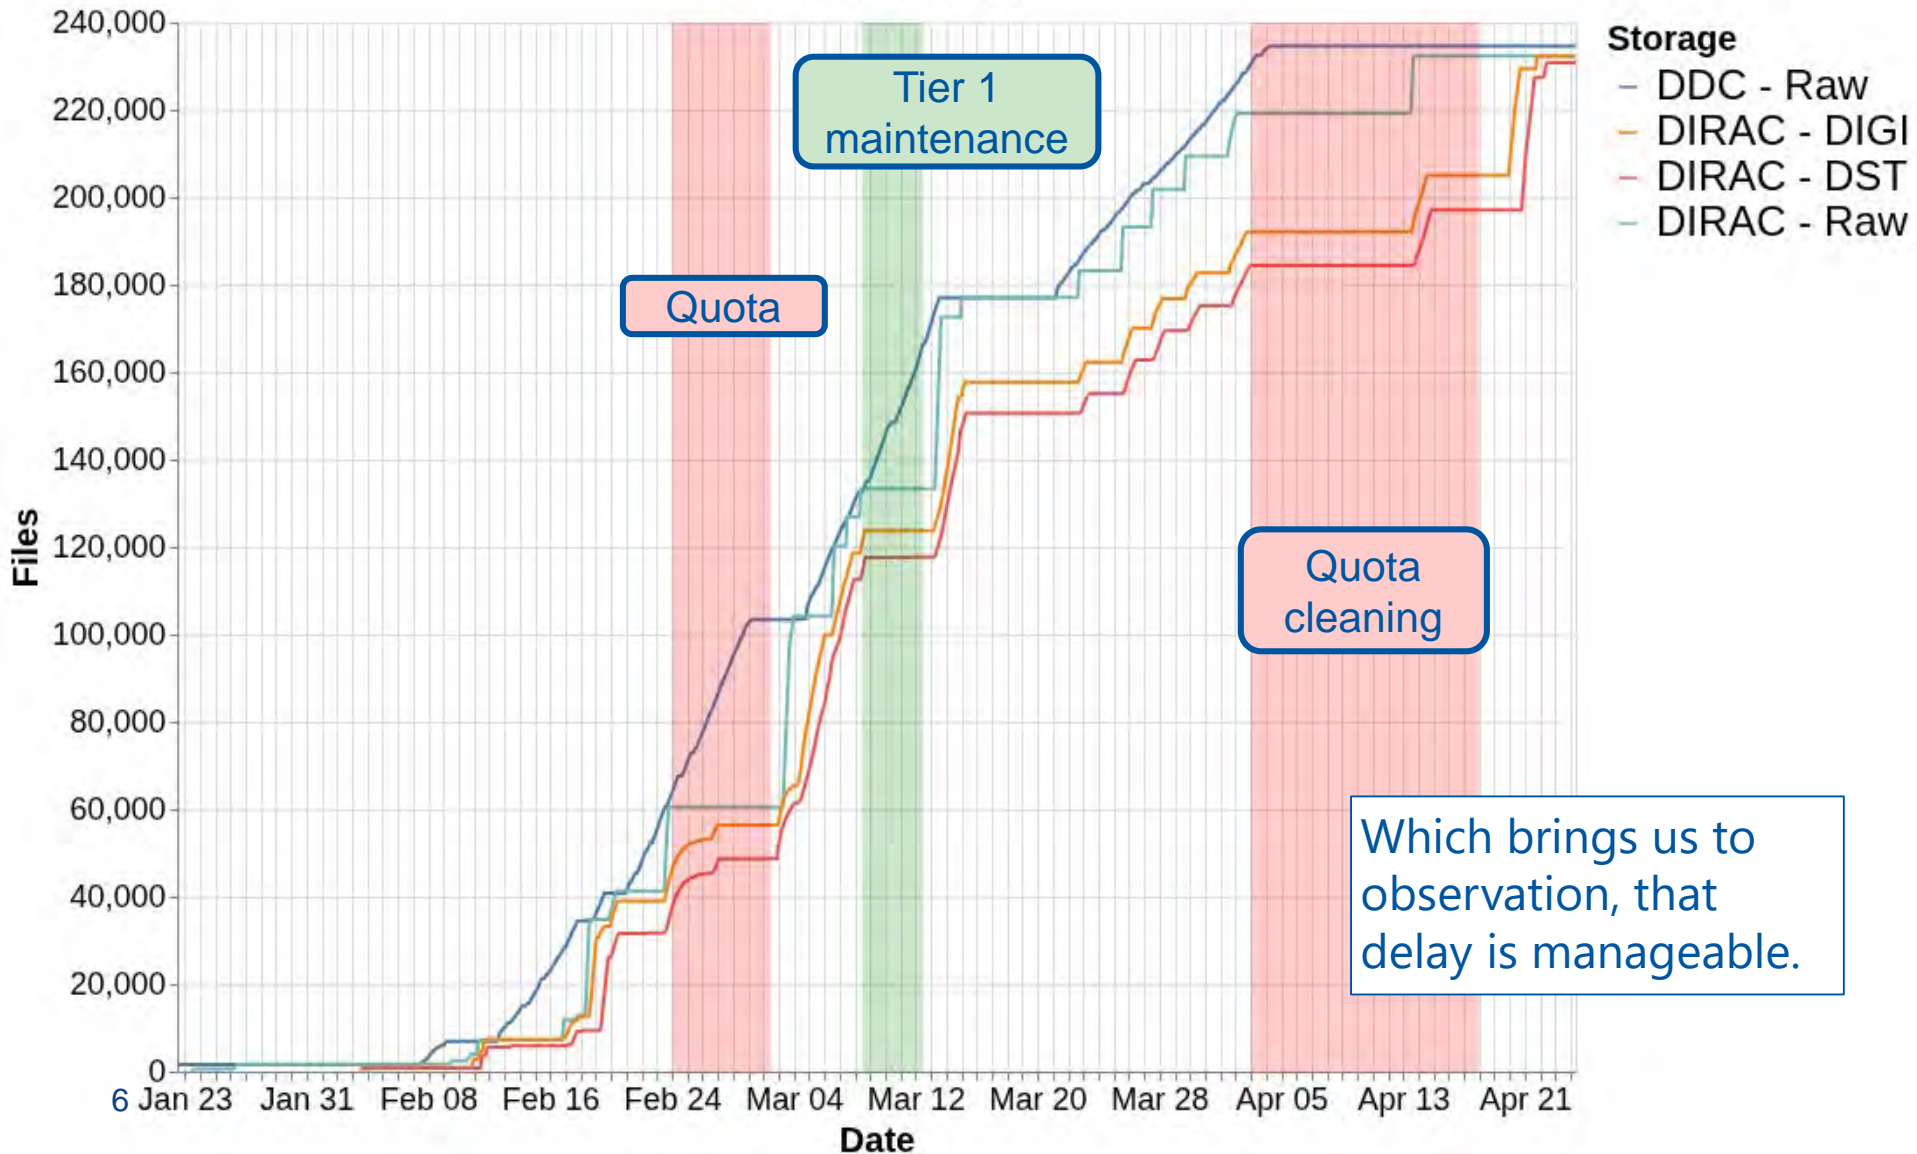
Delay

Amount of files



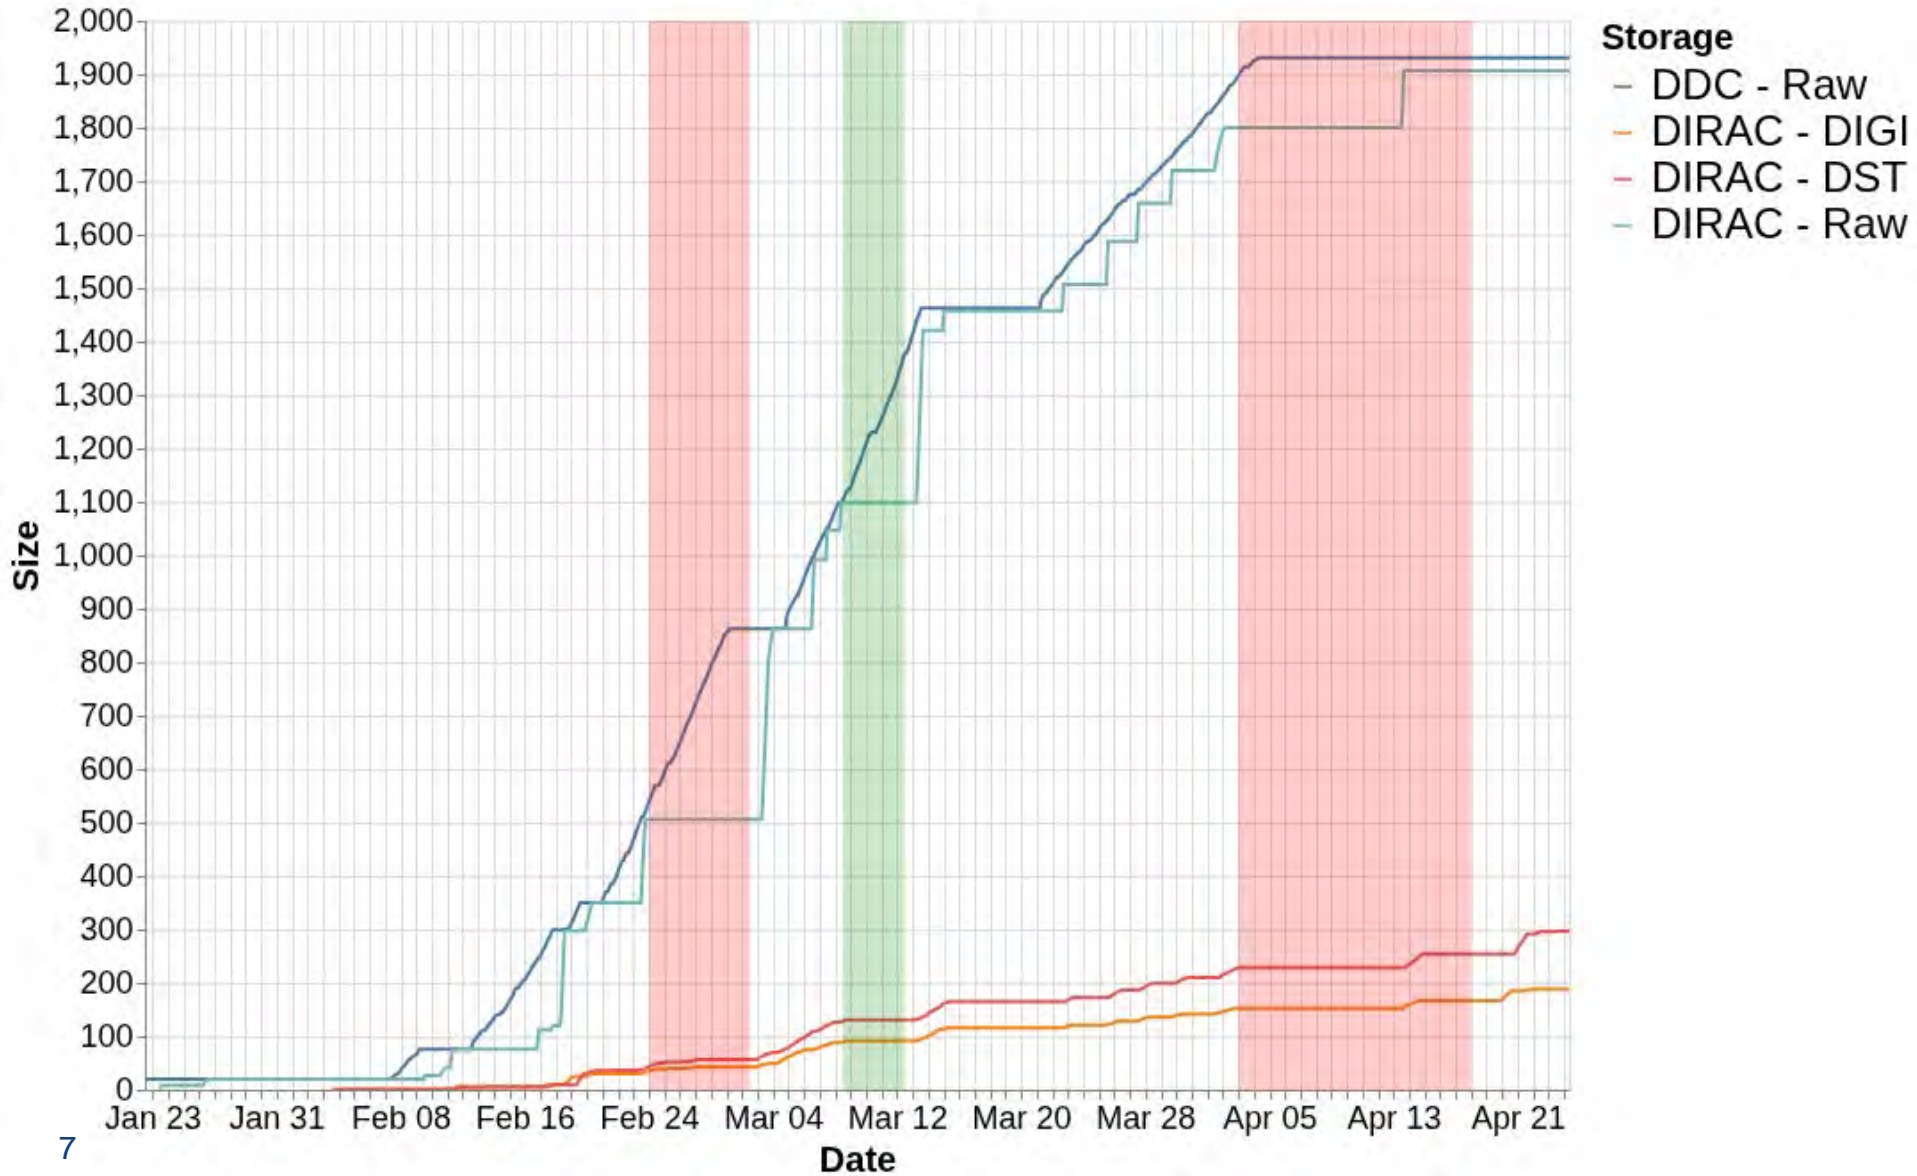
Delay

Amount of files



Size

Size of files




Size summary



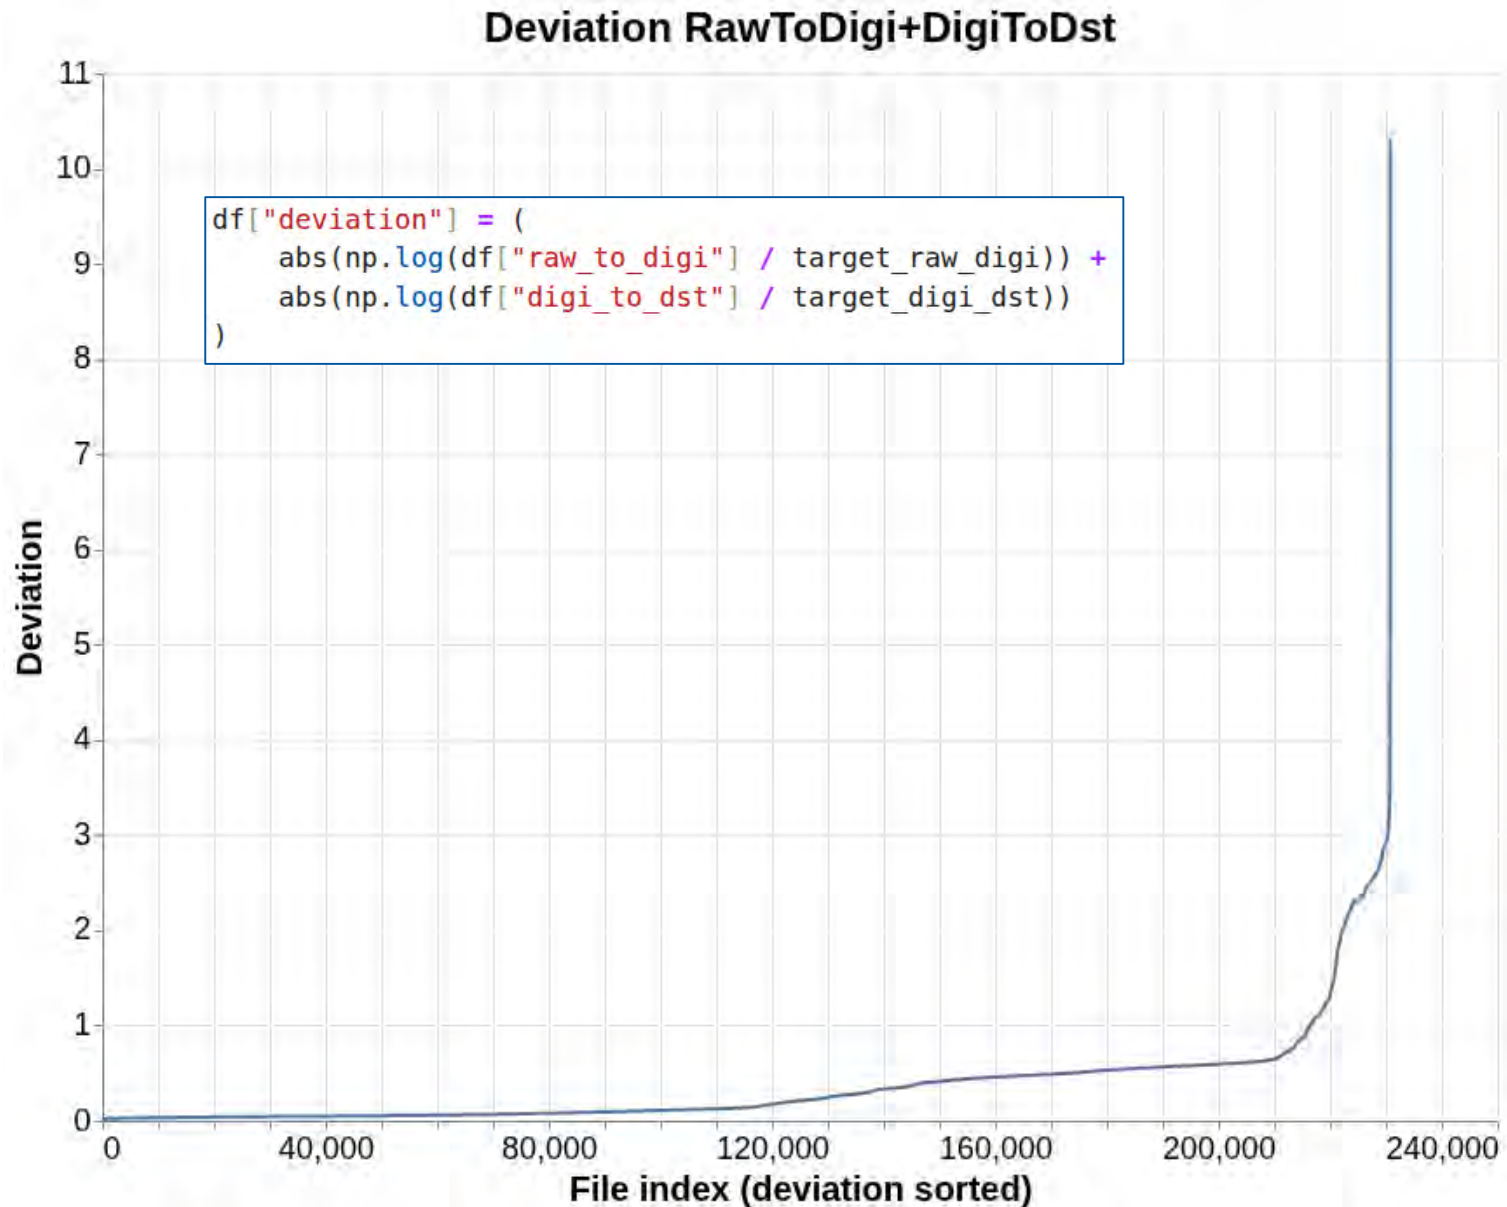
	Run9	Run8
RAW files	232349	30003
RAW size	(~1900 TB)	417 TB
RAW avg file size	8,2 GB	13,9 GB
DIGI files	232241	29999
DIGI size	188 TB	44,2 TB
DIGI avg file size	0,81 GB	1,47 GB
DST Files	230879	29962
DST size	297 TB	67.5 TB
DST avg file size	1,29 GB	2,25 GB
Raw ToDigi ratio	~ 0,09	~ 0,16
DigiToDst ratio	~ 1,56	~ 1,53
Raw ToDst ratio	~ 0,157	~ 0,162
Raw ToBoth ratio	~ 0,256	~ 0,268

Size summary

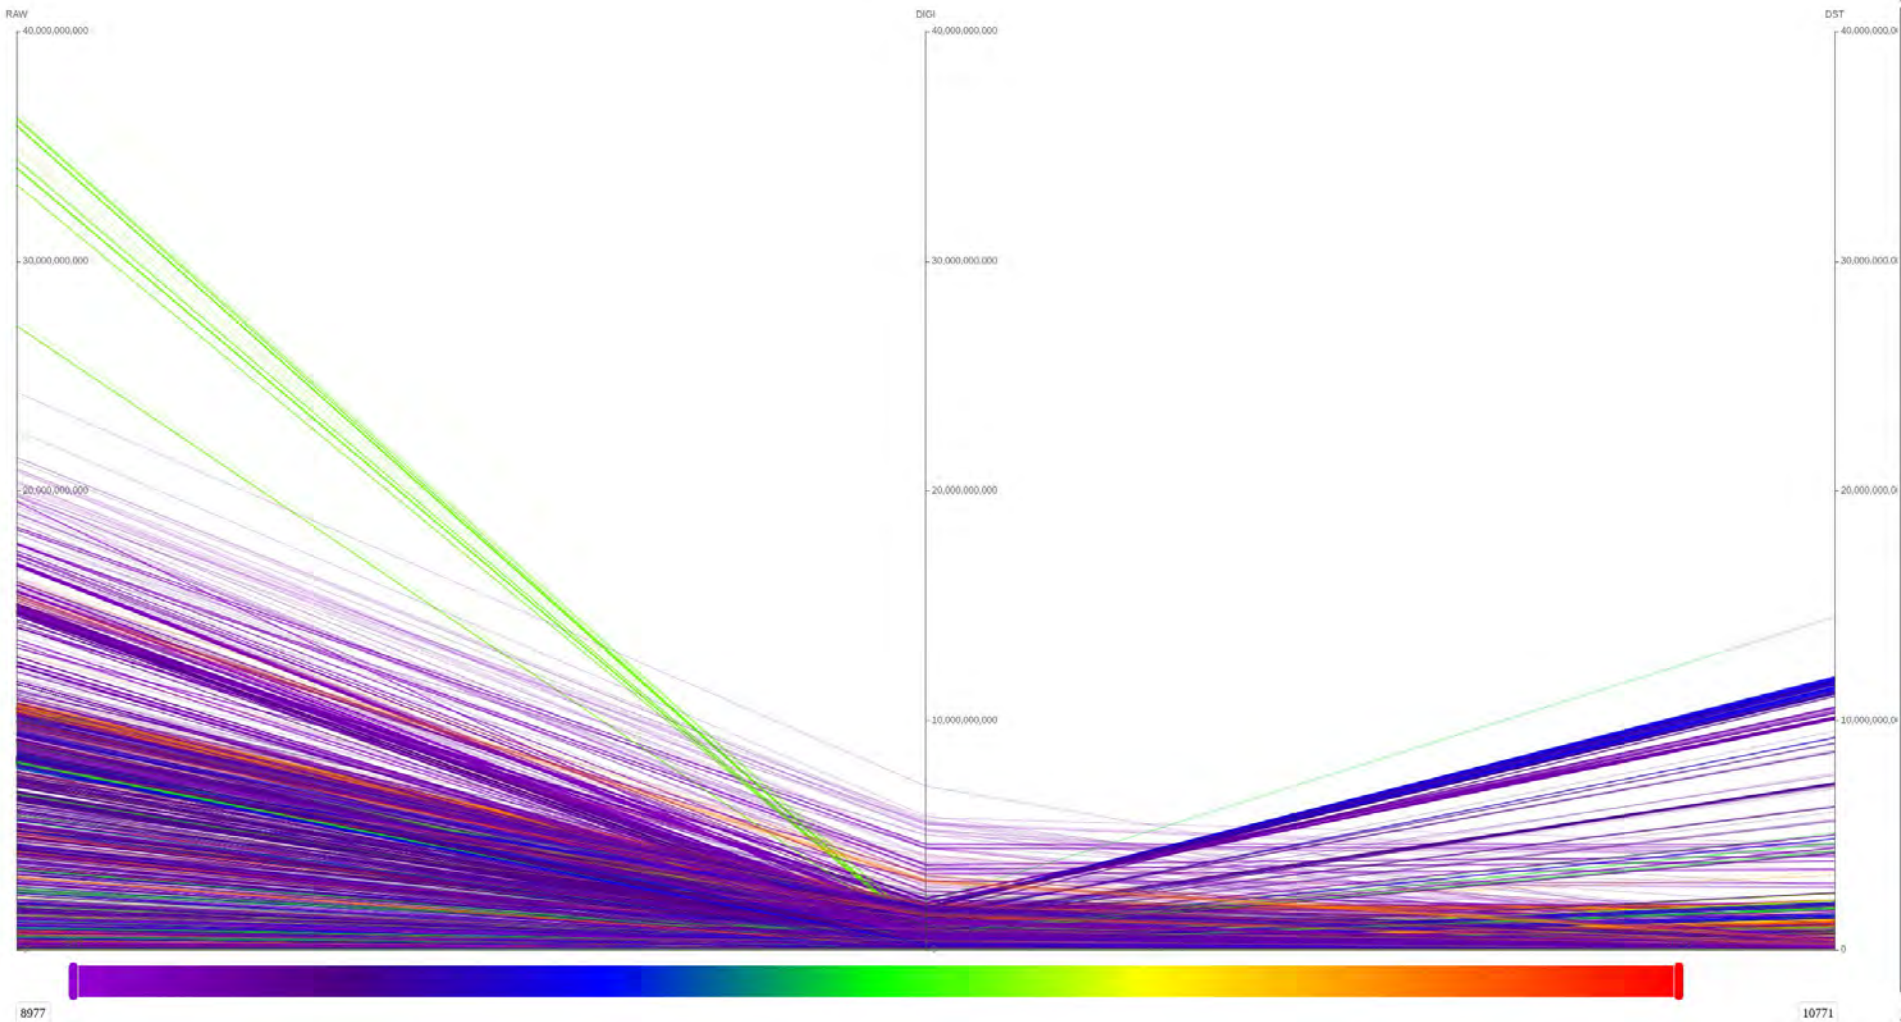
ALL File sizes (GB)



RawToDigi/DigiToDst deviation




Deviation > 1 (~6%)




Job distribution plot

23 Jan 2026, 13:24:15 - 22 Apr 2026, 08:03:44

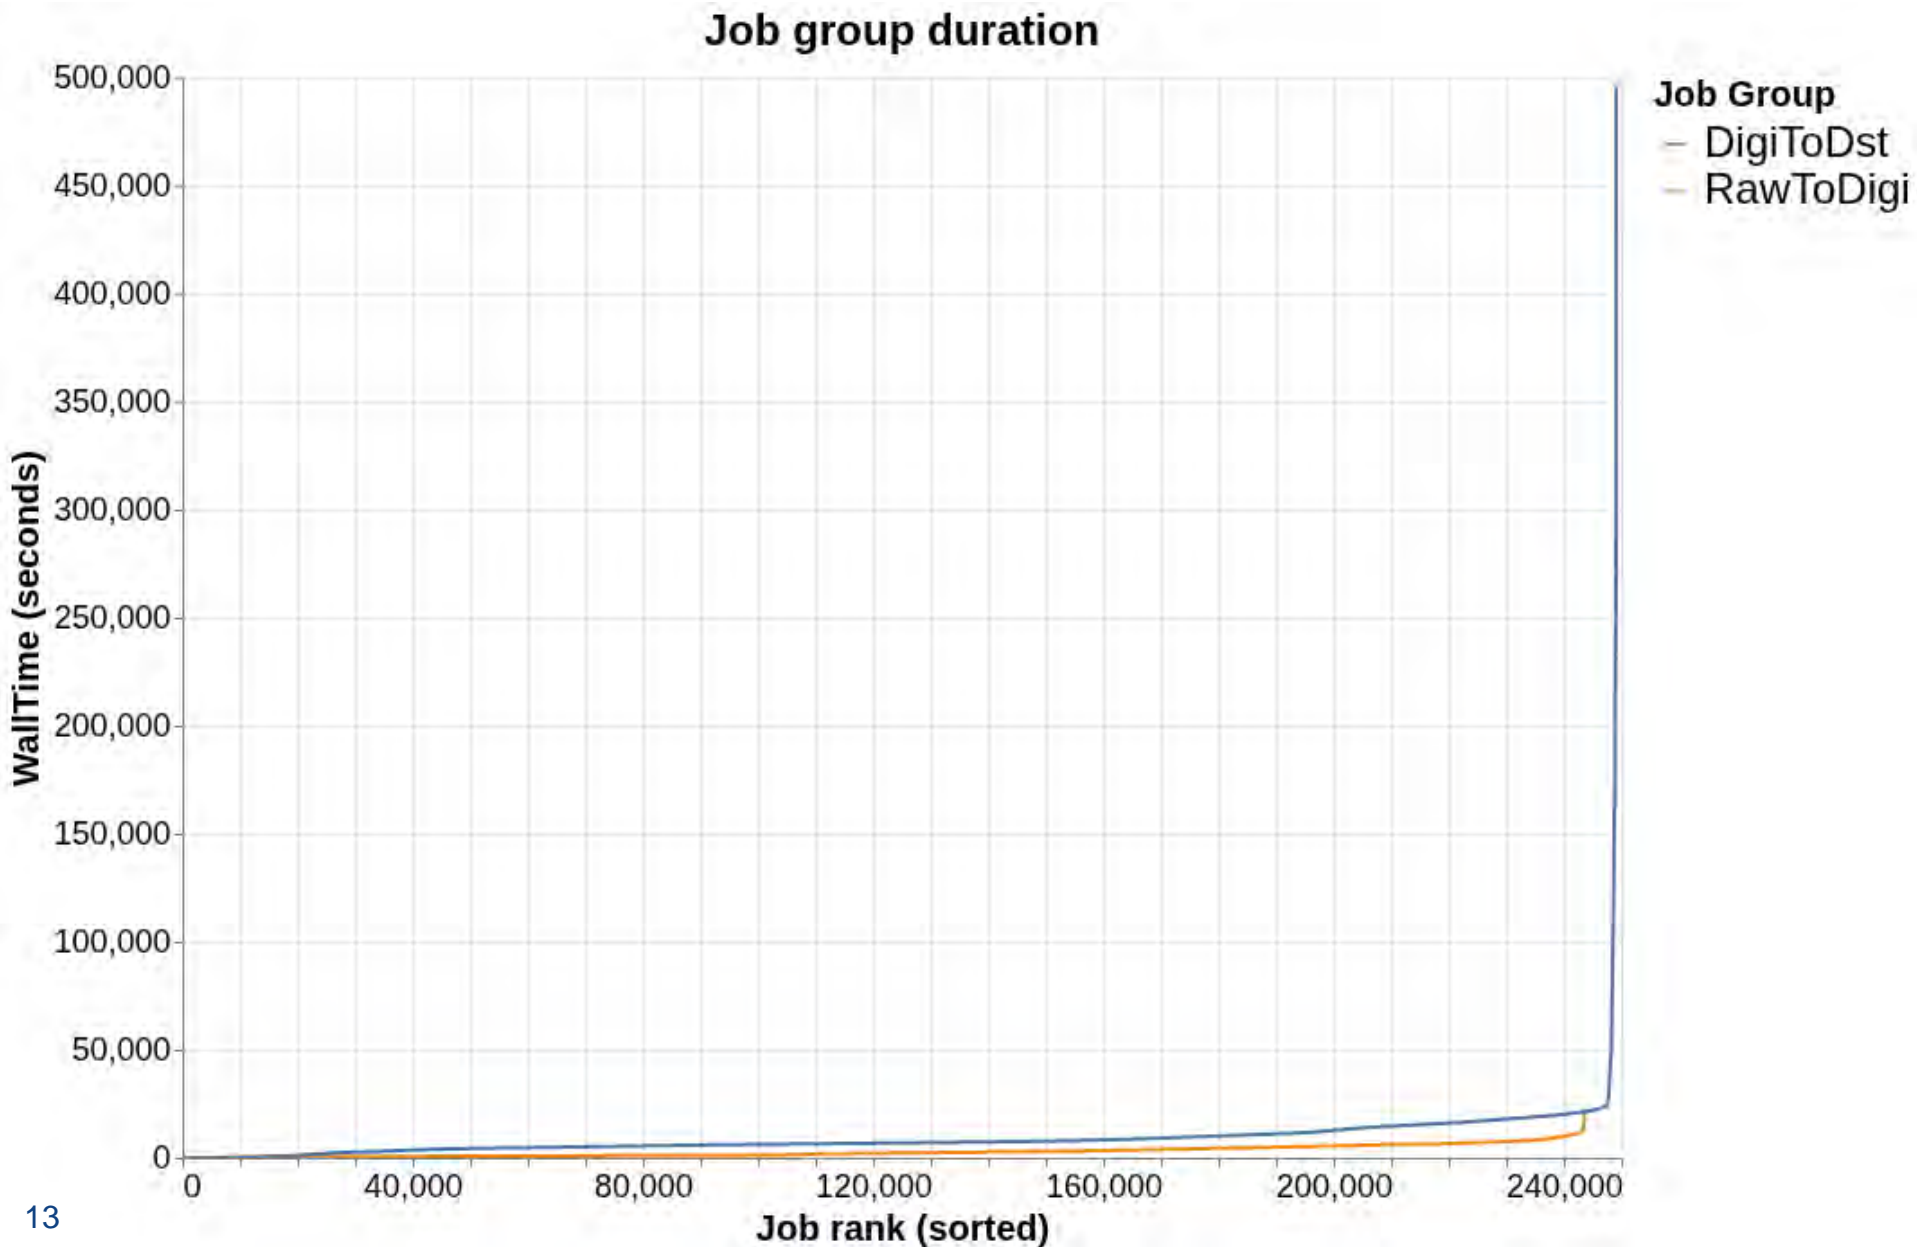
Count of points: 492675



Total jobs per hour



Jobs duration




Super long DigiToDst job



More than 5 days of execution. RAM usage almost 6 GB.
DigiToDst: run_Top_9862_ev4_p45.root

Jobs duration – selected 99%

Job group duration



Super long RawToDigi job

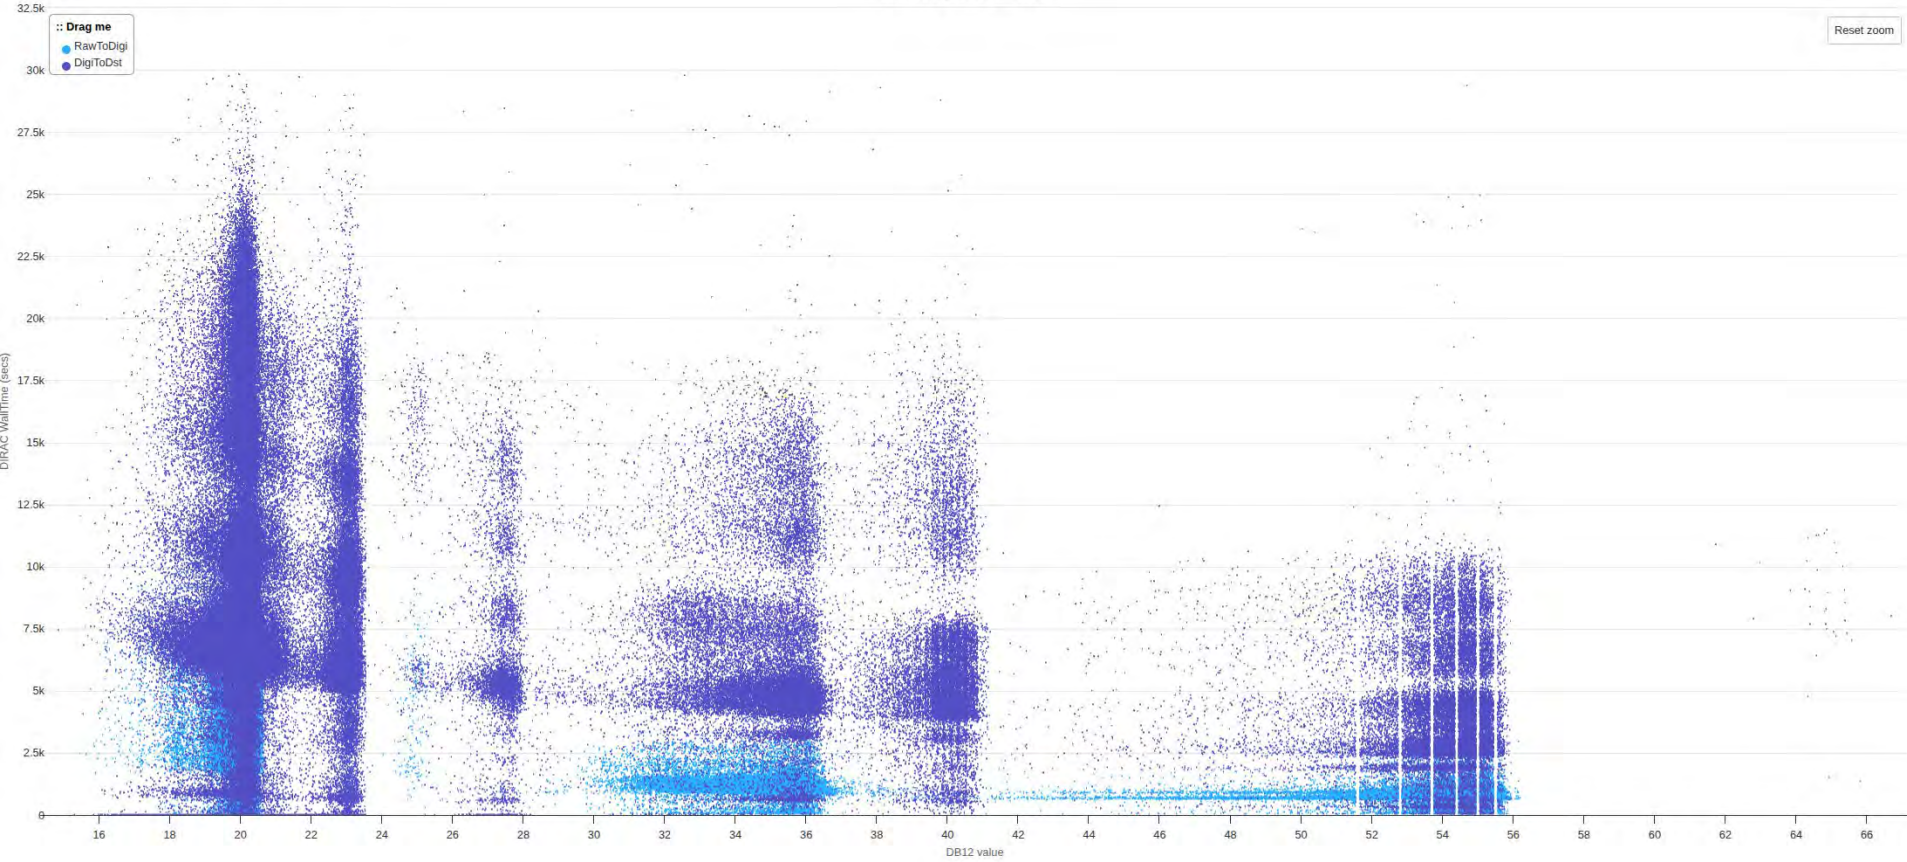


Around 6 hours of execution.
RawToDigi: run_Top_10489_ev0_p0.root


Job distribution plot ~ 99% of jobs

23 Jan 2026, 13:24:15 - 25 Apr 2026, 06:40:16

Count of points: 491411



Total jobs per hour



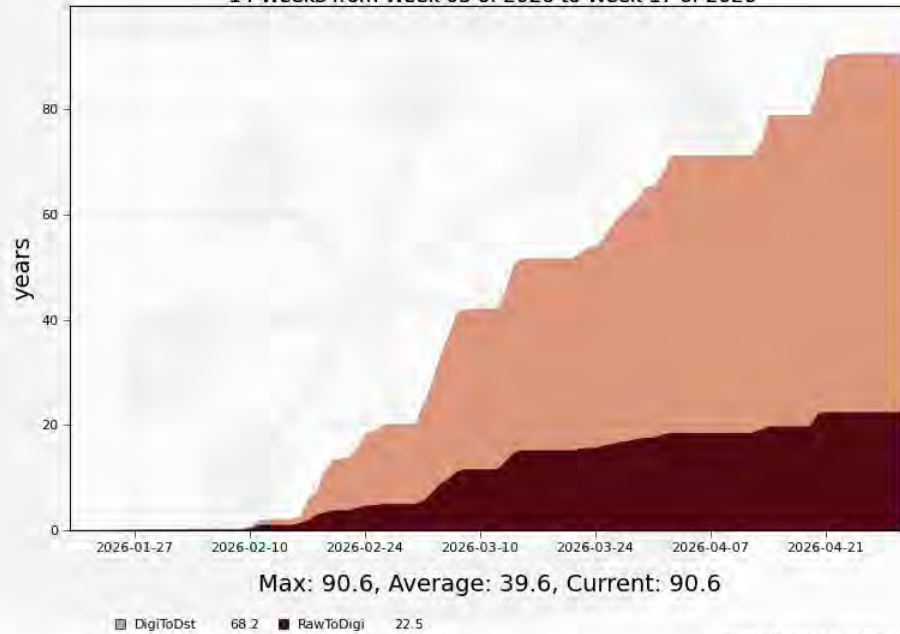
Time: 20 Apr 2026, 21:00:00
Jobs count: 3409

Reset zoom

Time to perform full production

Cumulative wall time by JobGroup


14 Weeks from Week 03 of 2026 to Week 17 of 2026



Generated on 2026-05-13 09:59:47 UTC

Cumulative wall time by Site

14 Weeks from Week 03 of 2026 to Week 17 of 2026



Generated on 2026-05-13 10:04:01 UTC

RawToDigi – 5,5 days on Tier1 (1500 cores)

DigiToDst – 8 days on Tier1, Tier2, DDC (3500 cores)

Jobs duration summary

		Run9	Run8
Duration	Raw ToDigi total	22,5 years	3,8 years
	DigiToDst total	68,2 years	53,2 years
	Raw ToDigi 1 event	0,25 s	0,20 s
	DigiToDst 1 event	0,75 s	2,8 s
	Raw ToDigi 1 file	3050 s	4000 s
	DigiToDst 1 file	9300 s	56000 s
Net Stream	Raw ToDigi Input stream	2,69 MB/s	3,47 MB/s
	Raw ToDigi Output stream	0,26 MB/s	0,36 MB/s
	Raw ToDigi total stream	2,95 MB/s	3,83 MB/s
	DigiToDst Input stream	88 KB/s	26 KB/s
	DigiToDst Output stream	139 KB/s	40 KB/s
	DigiToDst total stream	227 KB/s	66 KB/s

Directory structure

On DDC

```
.../run9/9623/run_Top_9623_ev2_p20.root
```

Better performance in run9 directory on Ceph and EOS.

Looks like it is critically important. But to get all list of files users still have to use:

```
ls -la ./run9/*
```

or

```
find ./run9 -type f -depth 2
```

On DIRAC

```
.../run9/run_Top_9623_ev2_p20.root
```

Optimized to keep hundred thousands file records in a single directory.


DIRAC FileCatalog can handle files in additional directories.

But **within DIRAC FileCatalog it is impossible to "move" files**. Only copy and remove.

Makes it incredibly more complex to build and visualize directories size and structure.


Done: RAW files renamed.
TODO: backup RAW on tapes

BM@N directories structure now




- 1
- 2
- 3
- 4
- 5
- 6
- 7
- 8
- 9
- 10
- 11
- 12
- 13







Total growth up to 260% in 2 month.
 Reduced slightly by removing of test data and old productions of Run8.



MLIT EOS users



Raw space consumed

 /eos/nica/bmn/	6.5 PB
 /eos/nica/mpd/	2.4 PB
 Other	1.9 PB
 /eos/baikalgvd/	1.9 PB
 /eos/dayabay/	622.3 TB
 /eos/meg/	407.3 TB

That means that yet another production and **BM@N occupies more than a half** of the space on general purpose EOS in MLIT!

Compression?

Lets take random RAW file with size 8.265.863.164 and try different compression/decompression

Type	Compression time	New size	Compression strength	Decompression time
zst	66 sec	4635380746	0.44	10 sec
gz	1044 sec	4506895508	0.46	58 sec
xz	7578 sec	3521371564	0.58	155 sec

	Compression		Decompression	
Type	Disk read	Disk write	Disk read	Disk write
zst	125 MB/s	70 MB/s	463 MB/s	826 MB/s
gz	8 MB/s	4,3 MB/s	78 MB/s	143 MB/s
xz	1,1 MB/s	0.465 MB/s	23 MB/s	53 MB/s

TODO

Move from VOMS to IAM-VOMS:

Ongoing. Mostly done. Some minor issues. Done with Artem Petrosyan and Aleksandr Baranov.

Backup on CTA tapes:

Requires rock-solid RAW files' structure.

Perform reprocessing of Run9 data:

Once software ready and request created.

List of participants

DIRAC: Igor Pelevanyuk, Igor Zhironkin

BM@N: Konstantin Gertsenberger, Ilia Romanov

Responsible for resources:

Tier-1, Tier-2, EOS: Valery Mitsyn , Artem Petrosyan,
Aleksandr Baranov, Andrey Kiryanov

CTA Tape library: Vladimir Trofimov

Govorun: Dmitry Podgainy, Dmitry Belyakov, Aleksandr
Kokorev, Maxim Zuev


NICA cluster: Ivan Slepov

DDC cluster: Ilia Slepnev

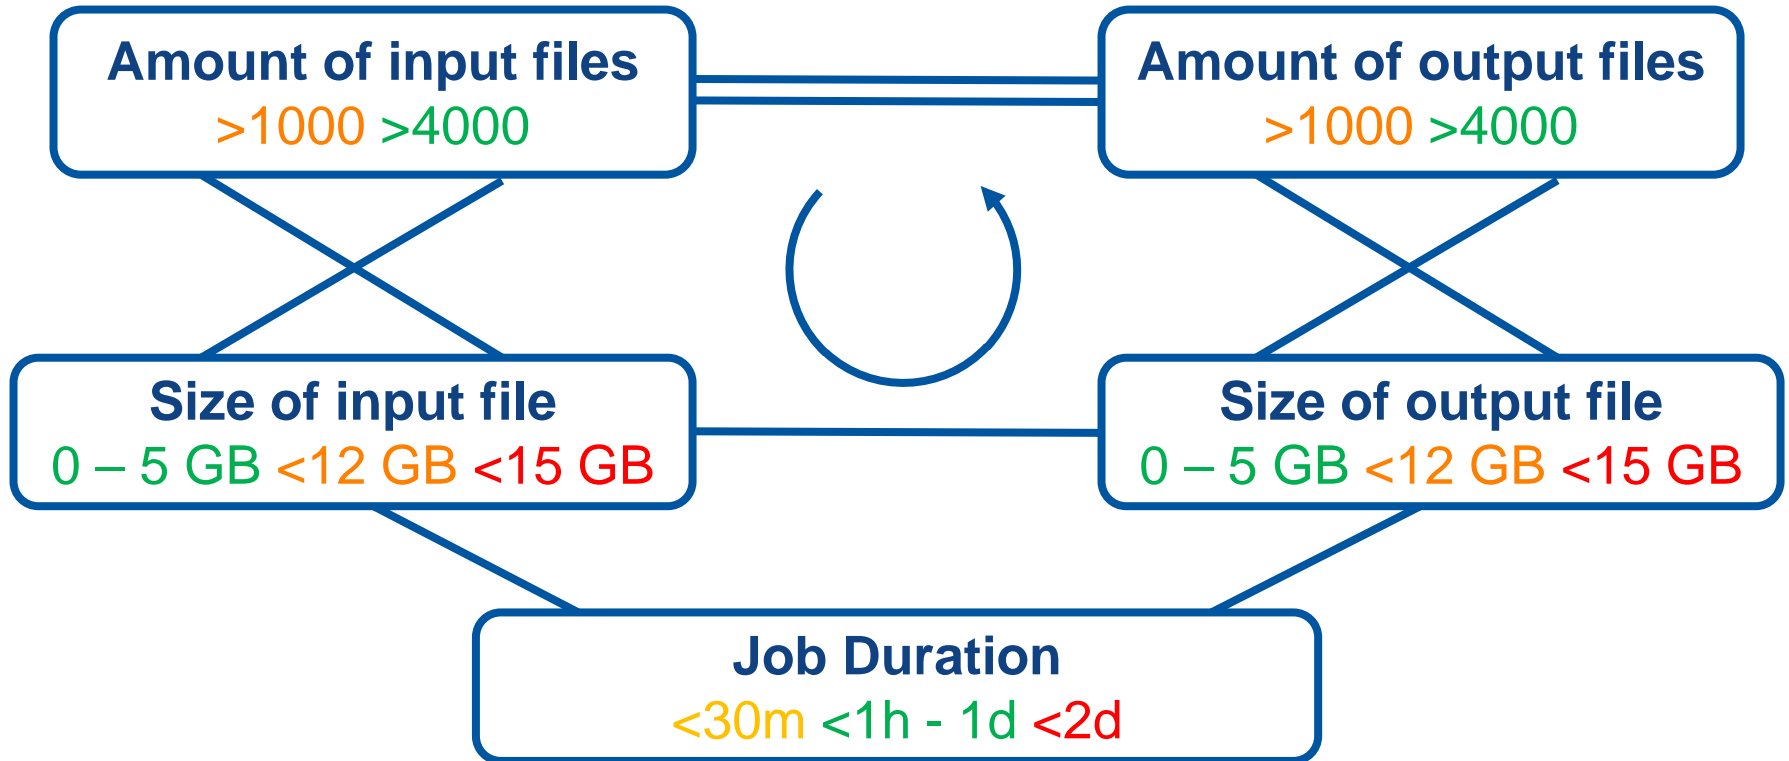
Network: Andrey Dolbilov

Conclusion for 16th BM@N Collaboration

- We successfully processed data from Run9.
- The scale of BM@N exceeded 1 PB. Data handling gonna be a burden now.
- Communication with MLIT about storage and planning growth is essential!
- **For those who perform data analysis!** DST data of one production is around 230k files and 300 TB. Event selection is not gonna be as easy as before. In case of issues, communicate with Konstantin Gertsenberger.




For distributed processing



Comfortable size of RAW files is in range 7-12 GB.
Ideal size considering RAW **Job Duration** is 10 GB.

DIRAC directory structure



Directory structure in DIRAC is now similar to NCX directory structure.

Still requires some work to do, to make it absolutely the same.

DIRAC directory structure

Special metadata applied for run9 data directories. This is crucial for further automatic processing!

```
bmn.nica.jinr
├── exp
│   ├── raw {type: raw}
│   │   └── run9 {run_period: 9}
│   ├── digi {type: digi}
│   │   └── run9 {run_period: 9}
│   │       └── 26.02.0
│   └── dst {type: dst}
│       └── run9 {run_period: 9}
│           └── 26.02.0
```

With that schema it is possible to select files by metadata:

```
$ dirac-dms-find-lfns type=raw run_period=9
/bmn.nica.jinr/exp/raw/run9/file_1_v0_p0.data
/bmn.nica.jinr/exp/raw/run9/file_2_v0_p0.data
/bmn.nica.jinr/exp/raw/run9/file_3_v0_p0.data
```

...