


BM@N Run 8 raw data production on distributed infrastructure with DIRAC

Konstantin Gertsenberger, Igor Pelevanyuk

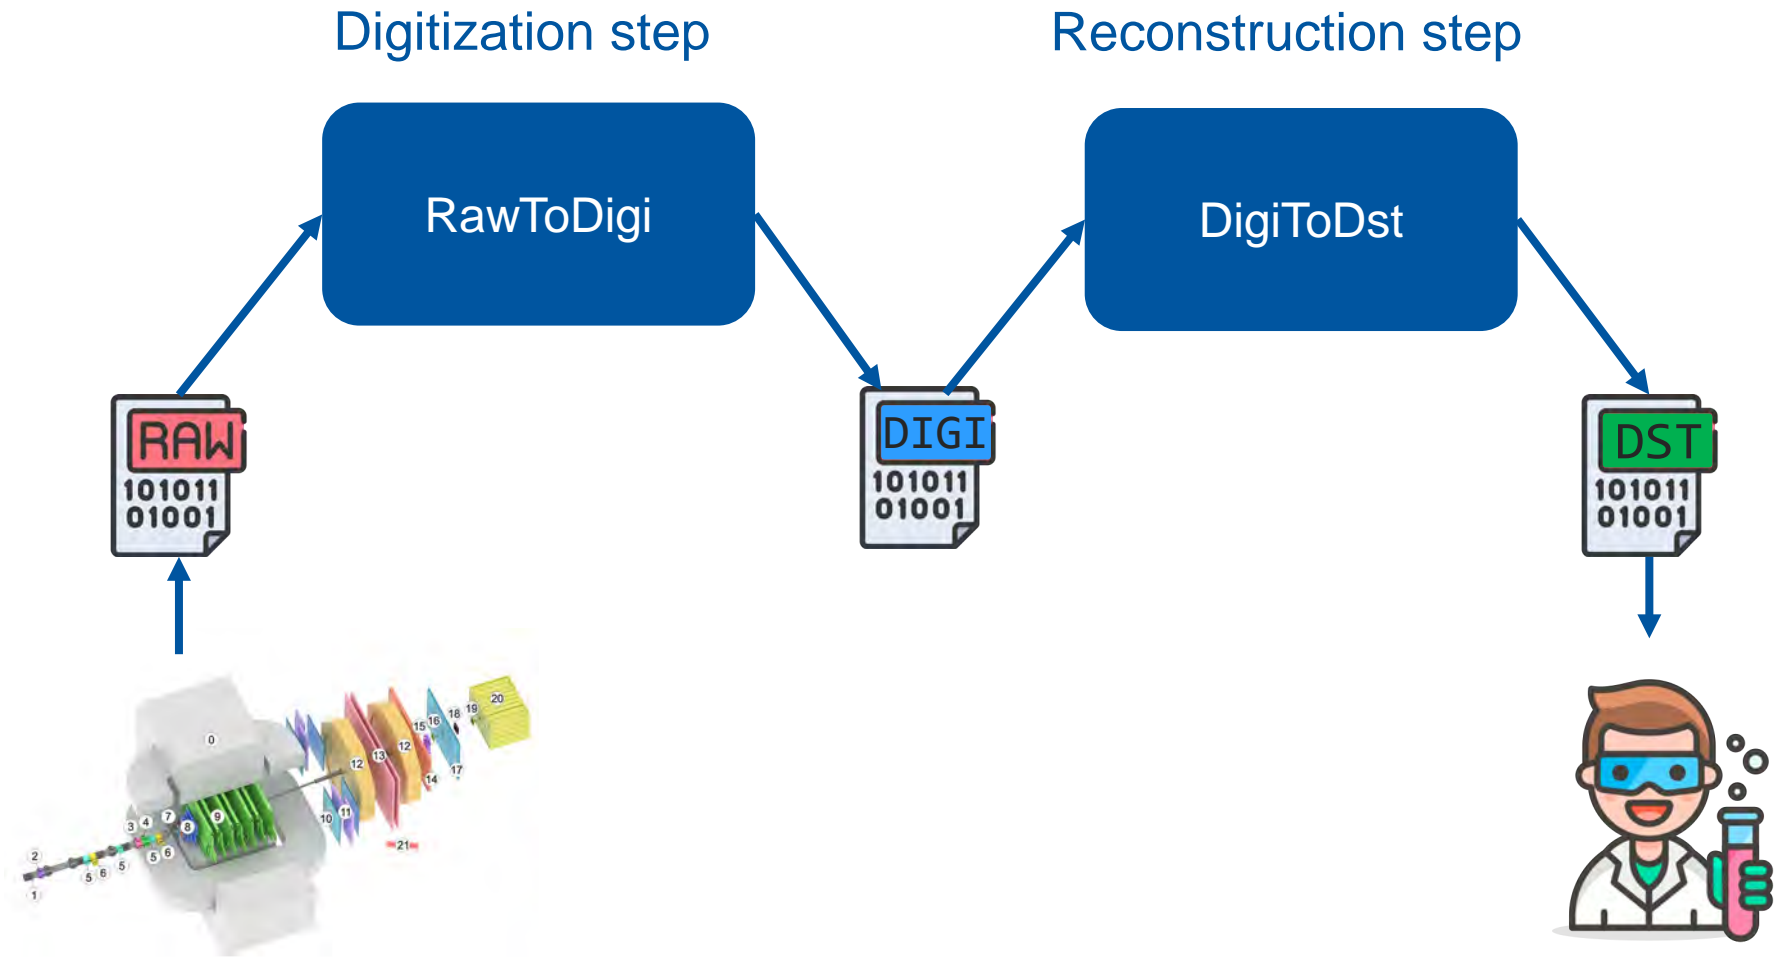
LHEP

MLIT




Status report
12.09.2023

Workflow of production



DIRAC in JINR

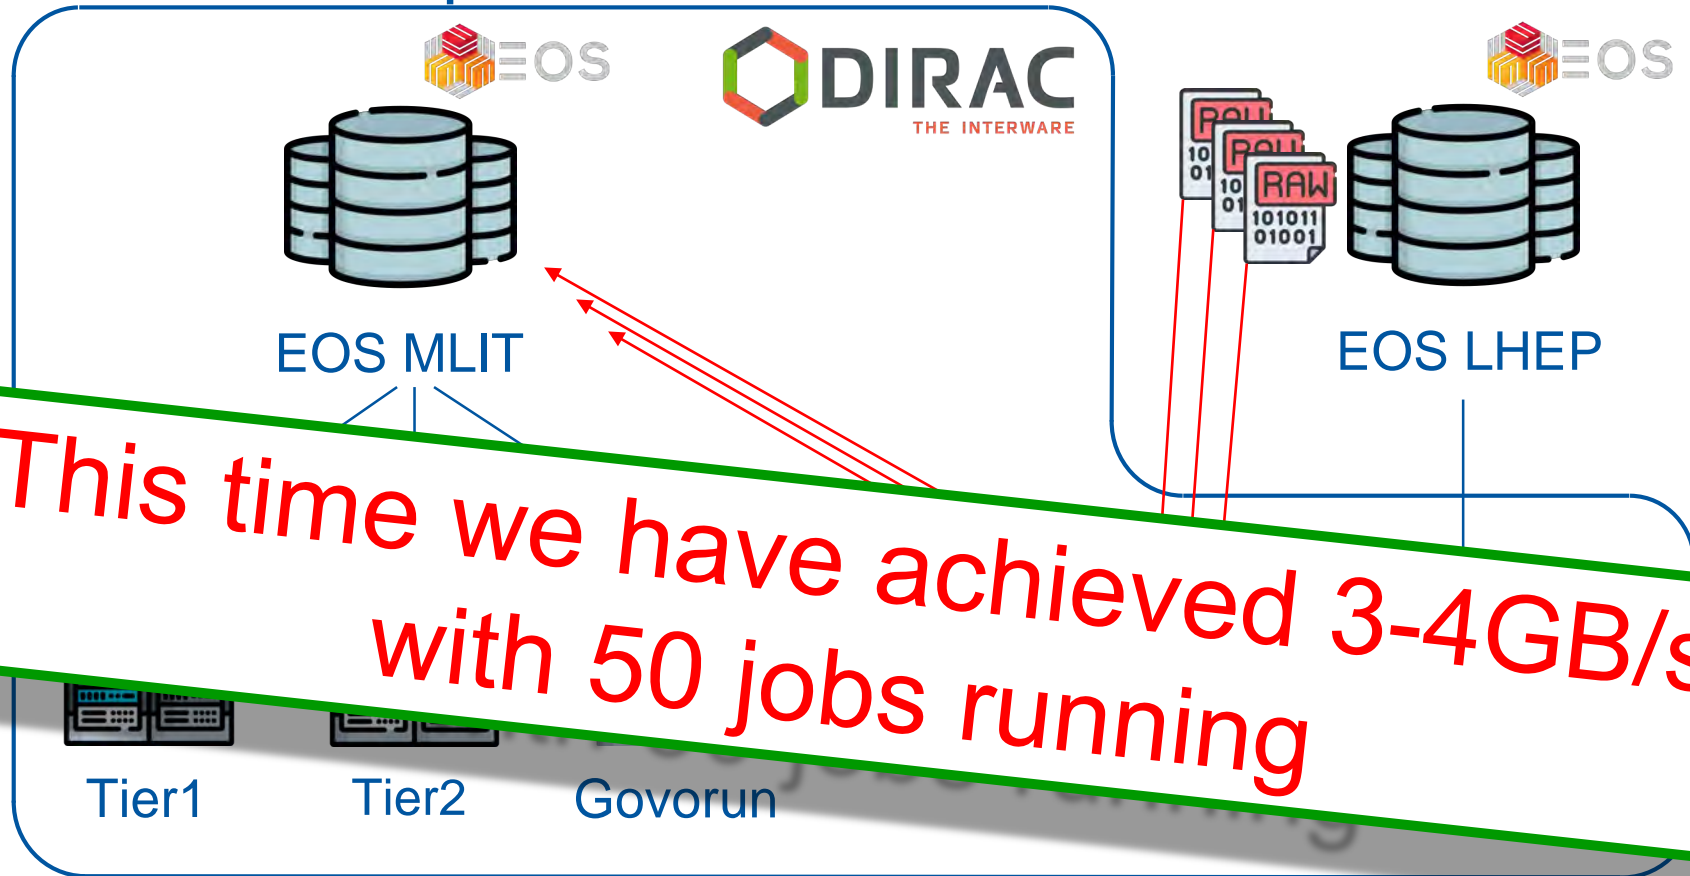


NRCN (National Research Computer Network) is the Russia's largest research and education (R&E) network. May allow execution of jobs submitted to Govorun on a resources of the network. Massive tests with MPD jobs were performed successfully in the beginning of 2022

Summer EOS Failure®

- July-August 2023: Fail occurred due to the bug in the EOS source code. This bug activates only during high load, so it is not observable during initial functionality tests.
- All files larger than 500 MB were in risk.
- Around 3-7 % of files under DIRAC Data Management system were affected(subjectively). This had great effect on BM@N production process
- Partial BM@N data processing chain was performed for half of runs

Step 1: Data transfer

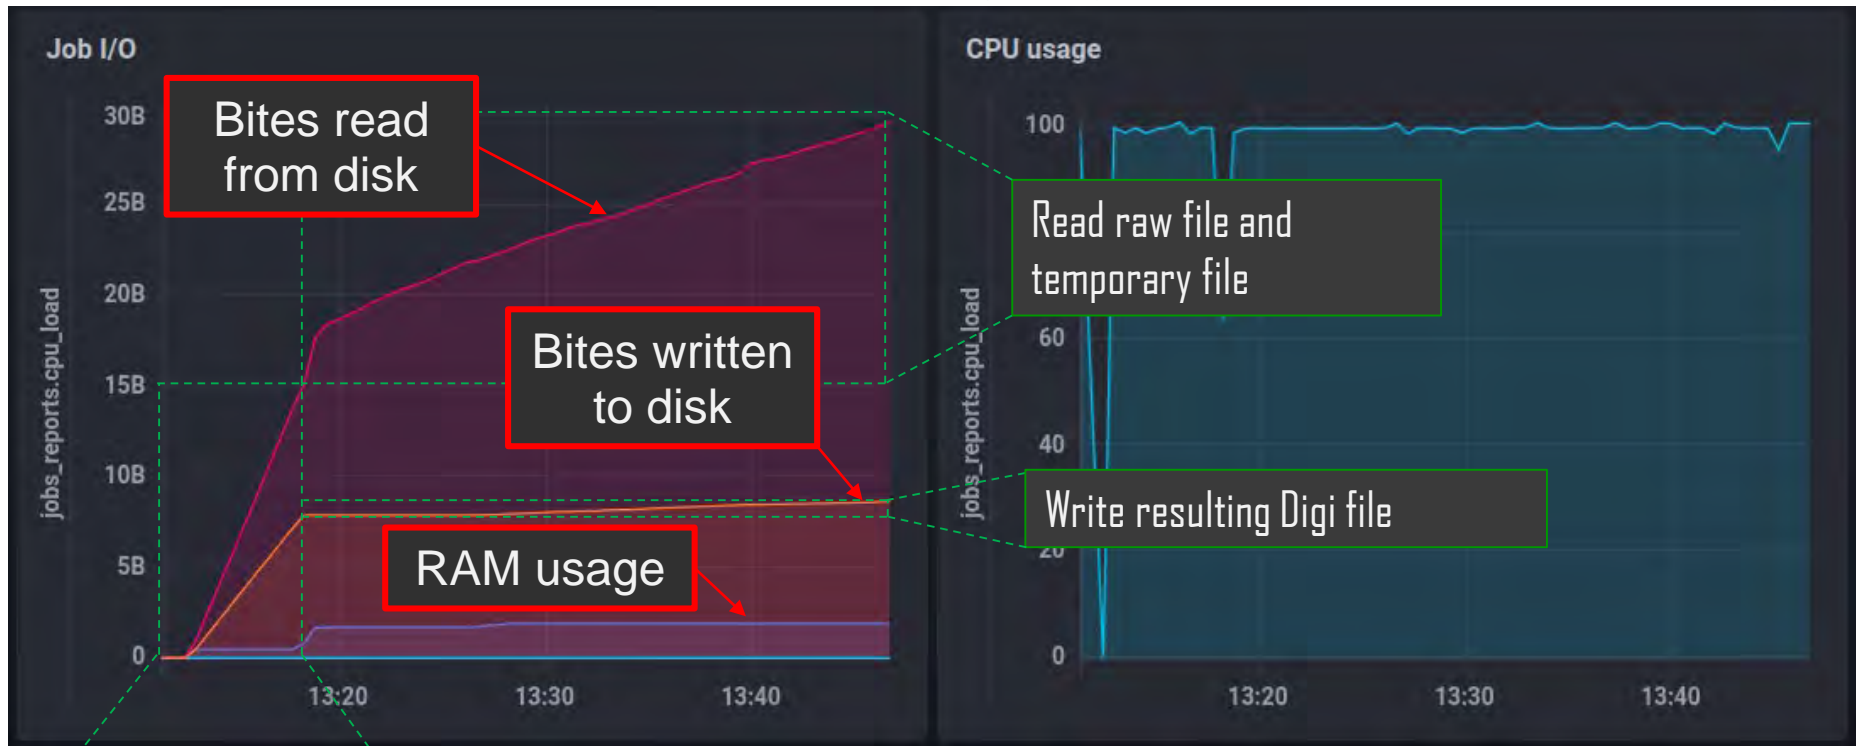


This time we have achieved 3-4GB/s with 50 jobs running

- Single stream of xrootd transfer can not exceed 100MB/s. **Transfer would take ~ 50 days.**
- NCX interface node can sustain not more than 10 streams(1GB/s total). And that would overload its network.

So, 20 independent DIRAC jobs were sent to NICA cluster to perform transfers with one stream each.

Step 2: Raw2Digi job profiling

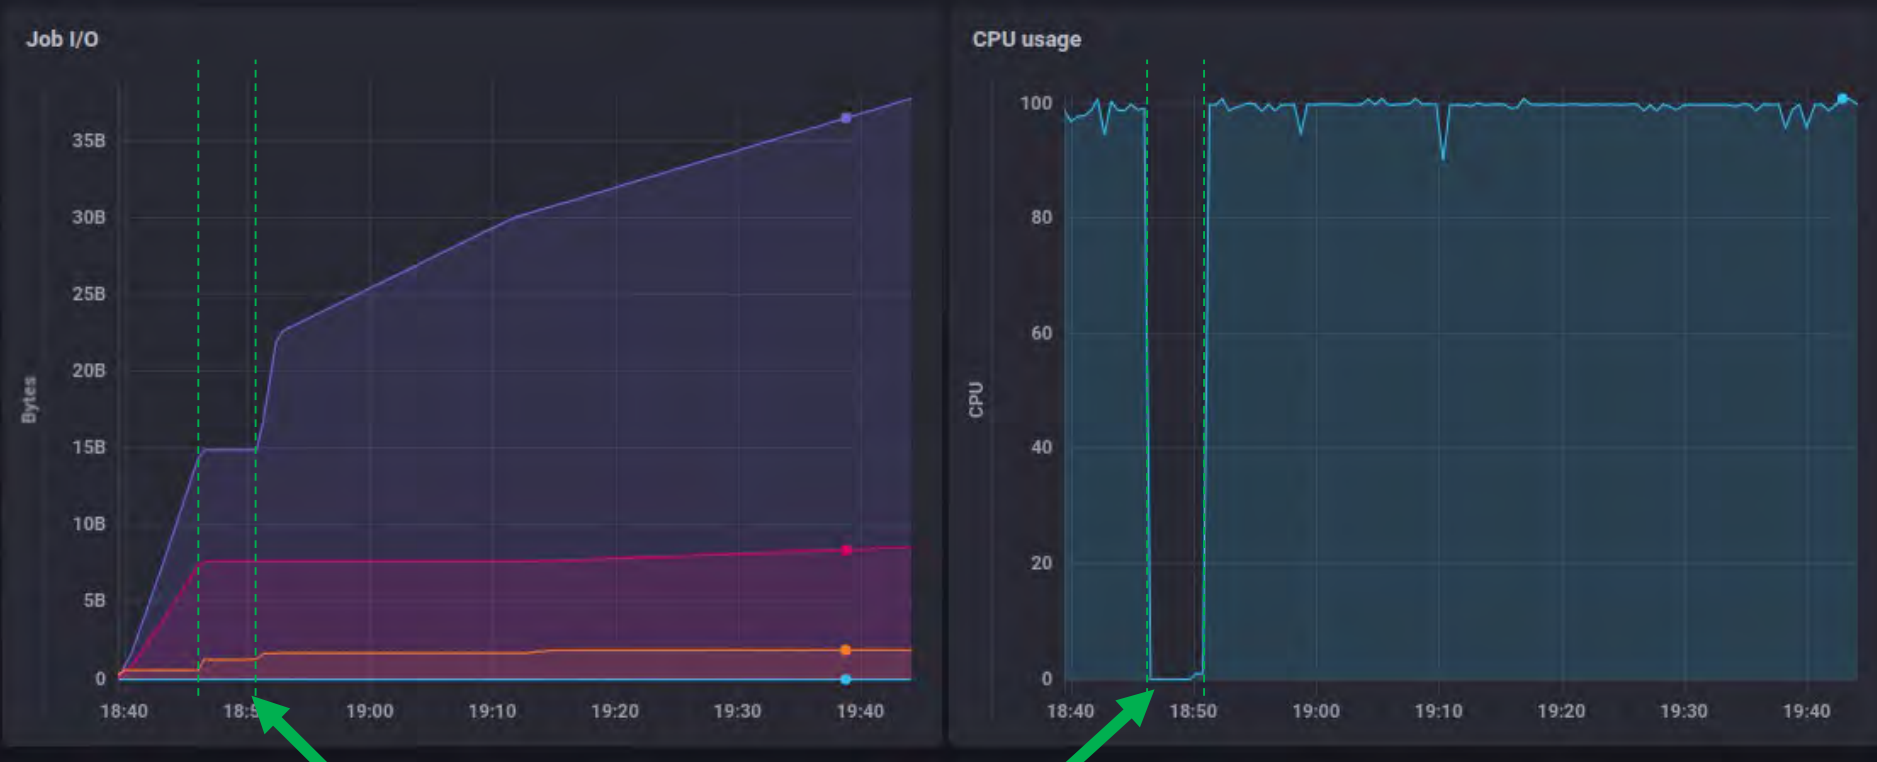


Initial read of 15GB raw file and creation of temporary 8 GB file

Disk usage
Temporary file: +8 GB
Result file: 800MB
Total disk usage per 15 GB job: 25 GB

RAM usage : ~2GB

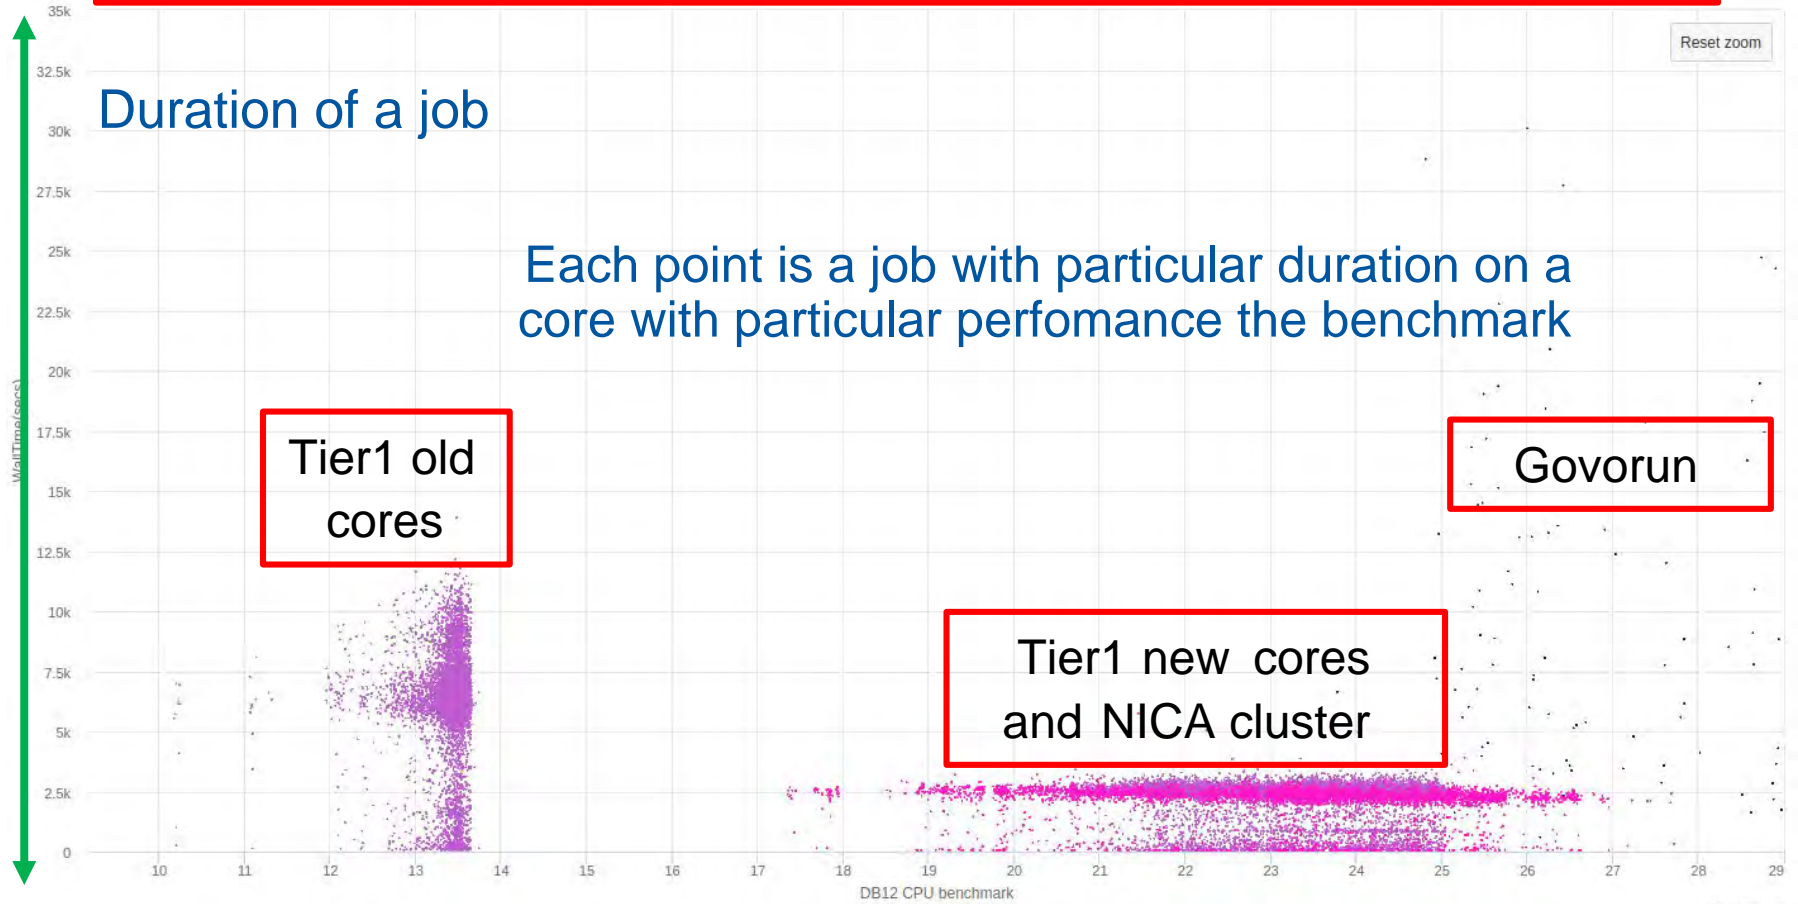
Step 2: Raw2Digi job profiling



New Raw2Digi have strange 5m idle period

Step 3: Massive production Raw2Digi

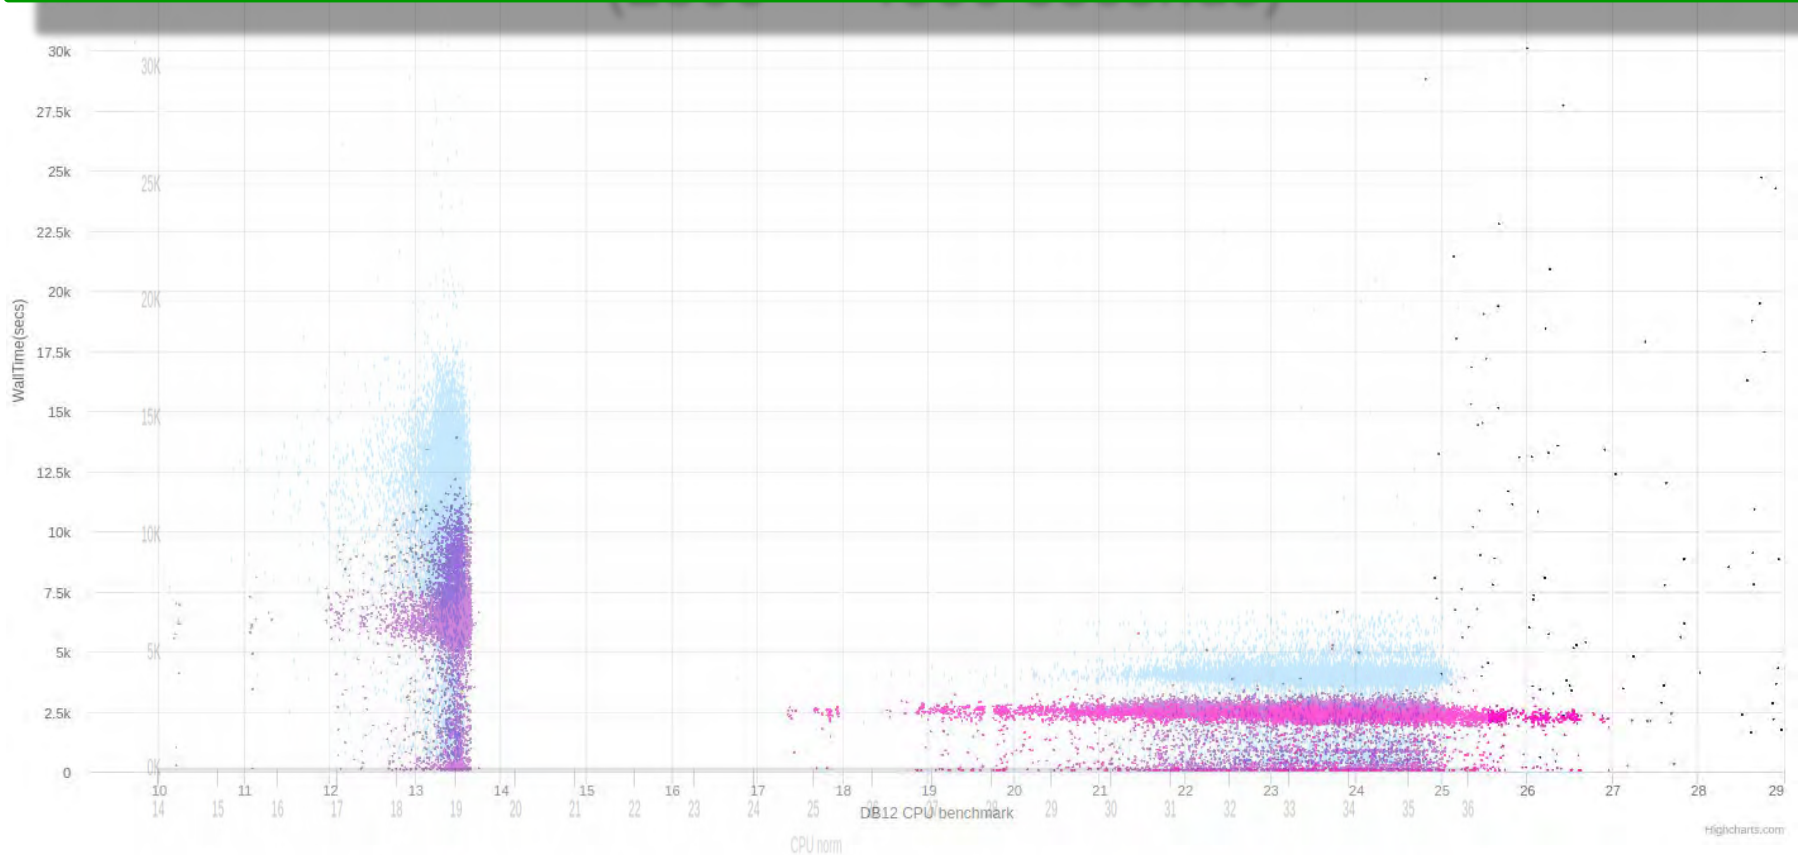
Total duration of Raw2Digi campaign – 35 hours




CPU core performance on benchmarks

Step 3: Massive production Raw2Digi

Average Raw2Digi calculation time increased by 60%
(2500 -> 4000 seconds)



Step 3: Network usage



300 jobs running
4 MB/s per job


NICA Cluster

Tier1

1580 jobs running
4.1 MB/s per job

Maximal transfer speed (Read+Write) with EOS in MLIT – 7.5 GB/s

Step 4: Digi2Dst profiling



Step 4: Digi2Dst VF

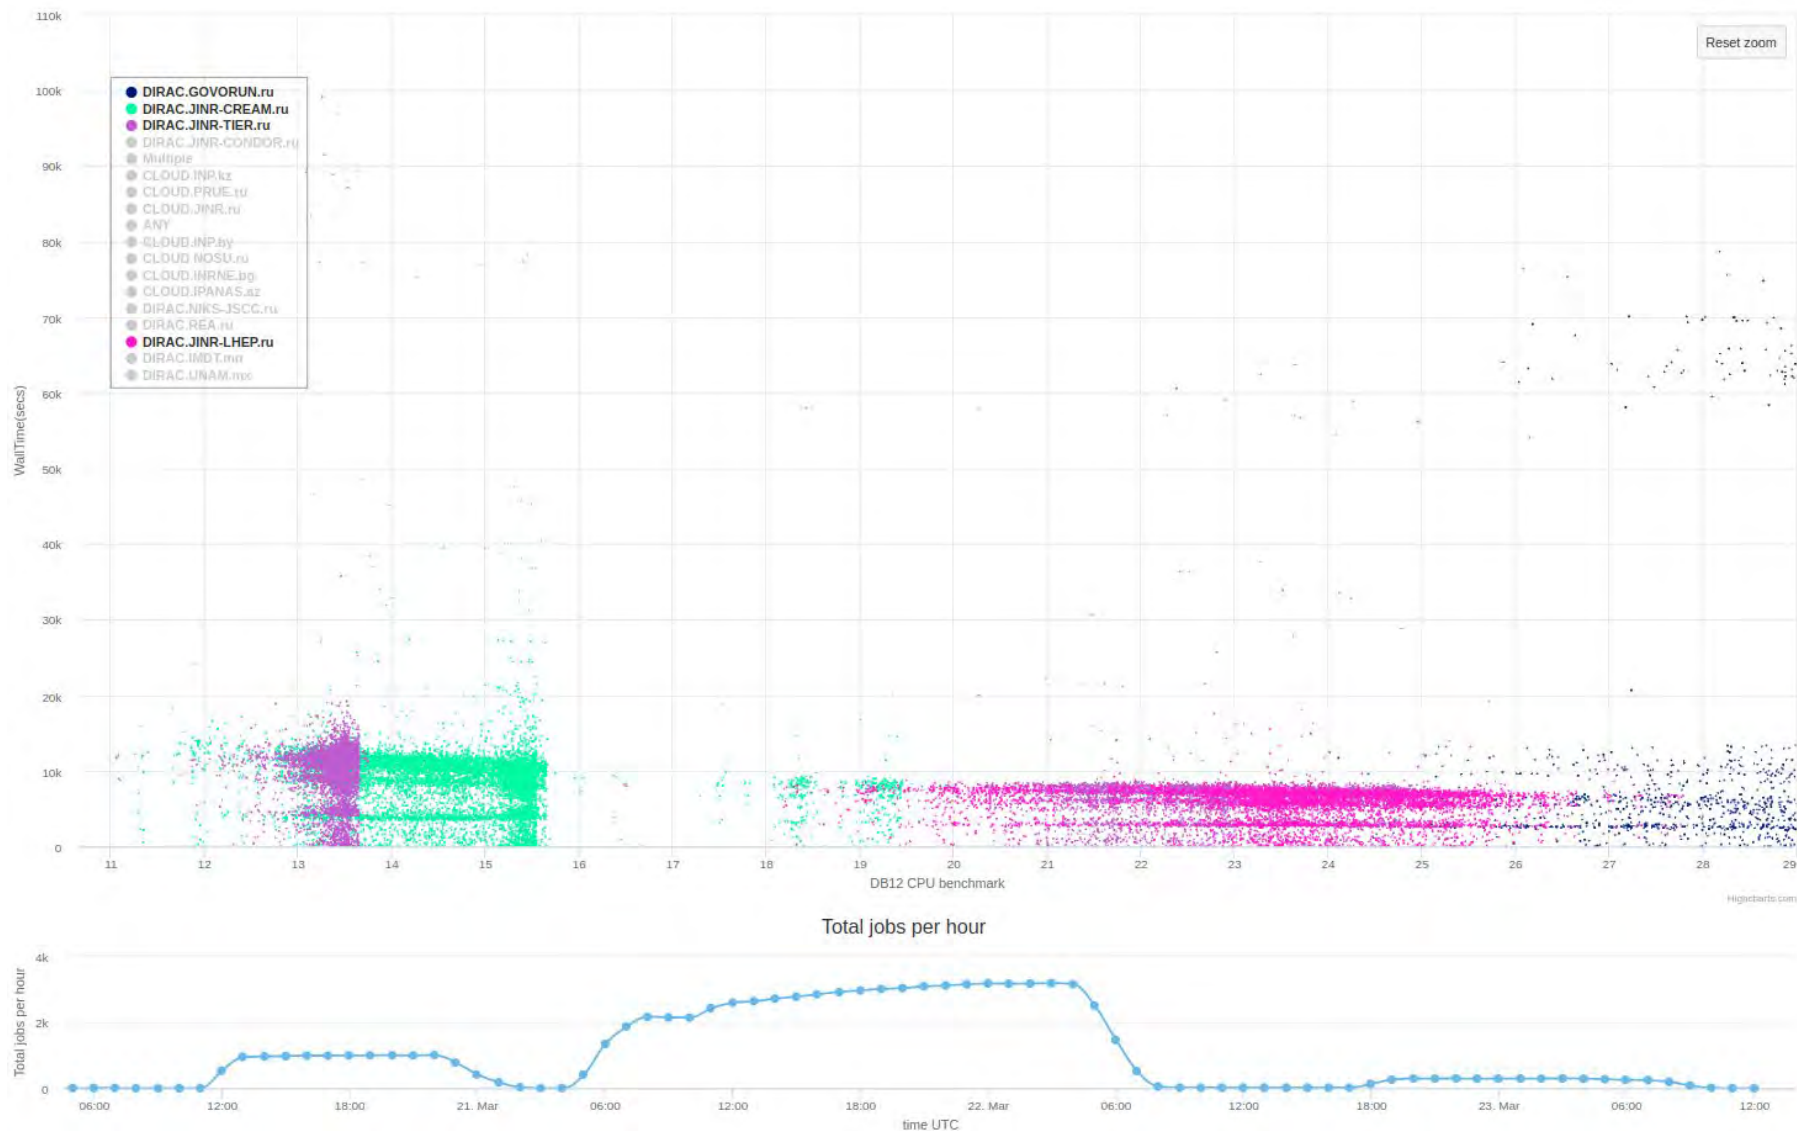
Job I/O



CPU usage



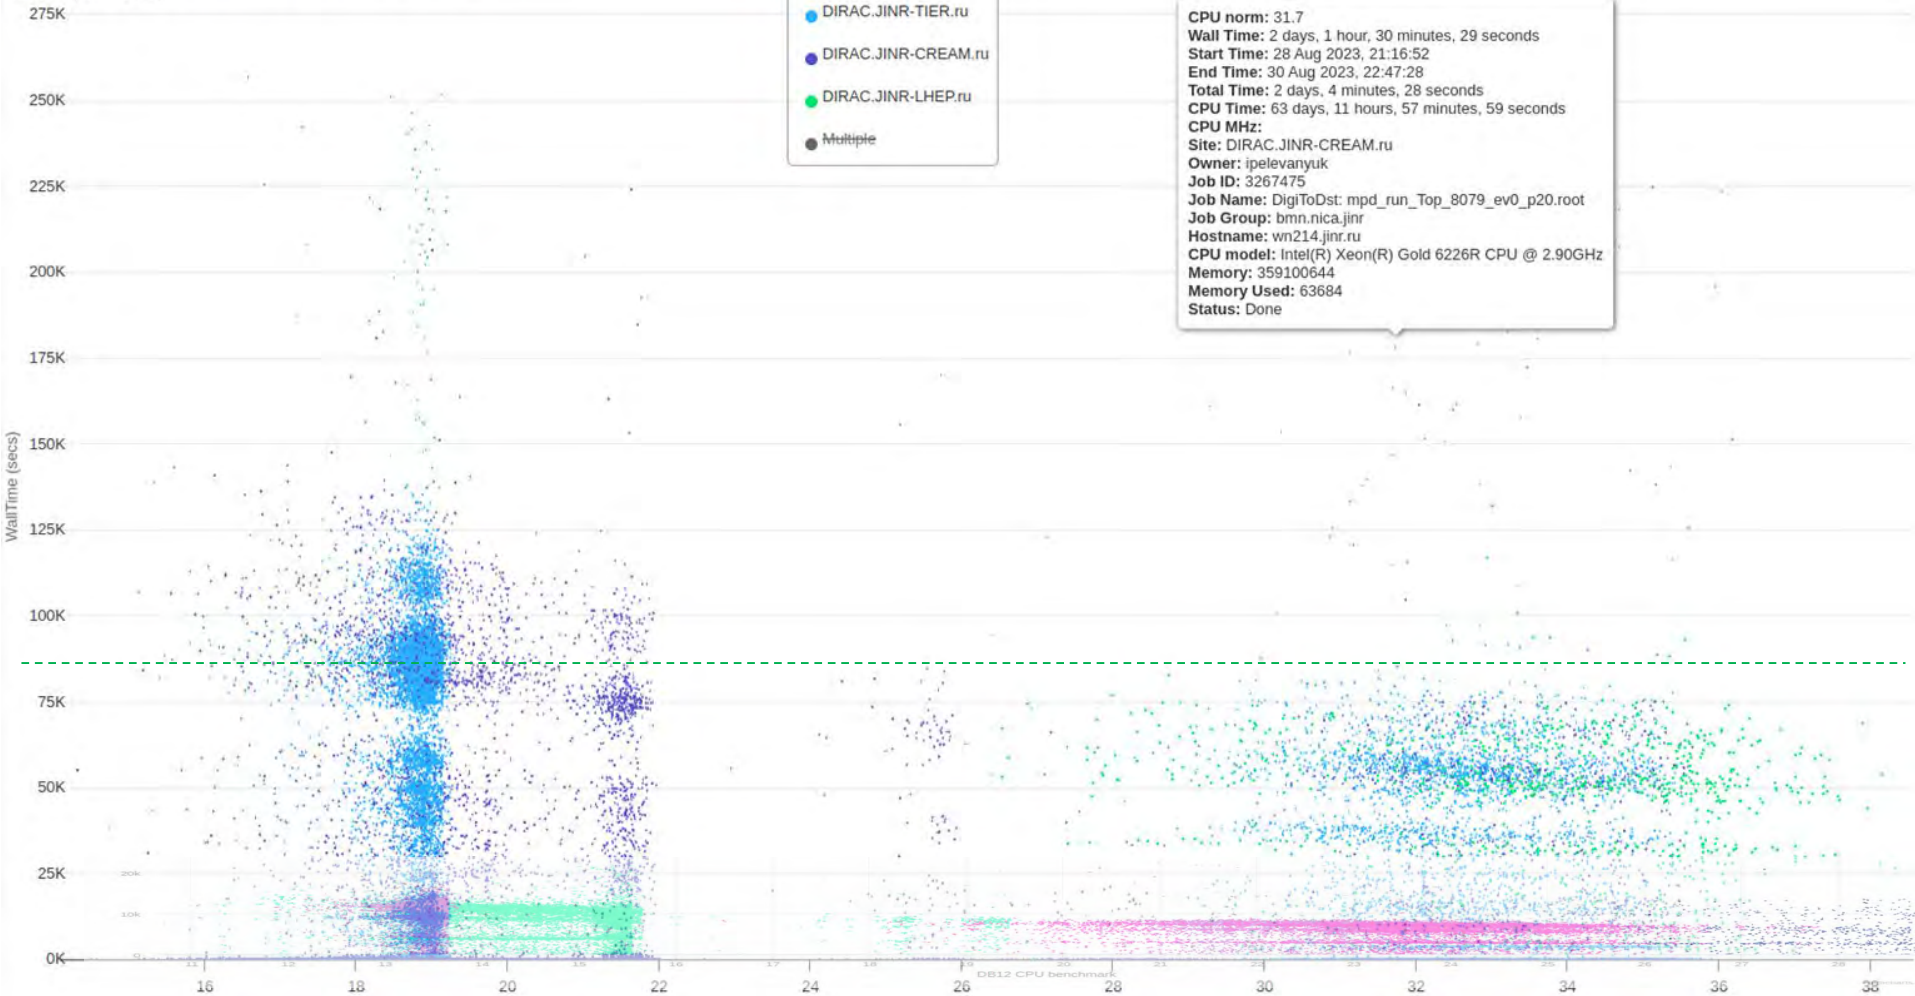
Step 5: Old Massive production Digi2Dst



Step 5: New VF Digi2Dst

28 Aug 2023, 10:00:09 - 7 Sep 2023, 10:51:06


Count of points: 30537



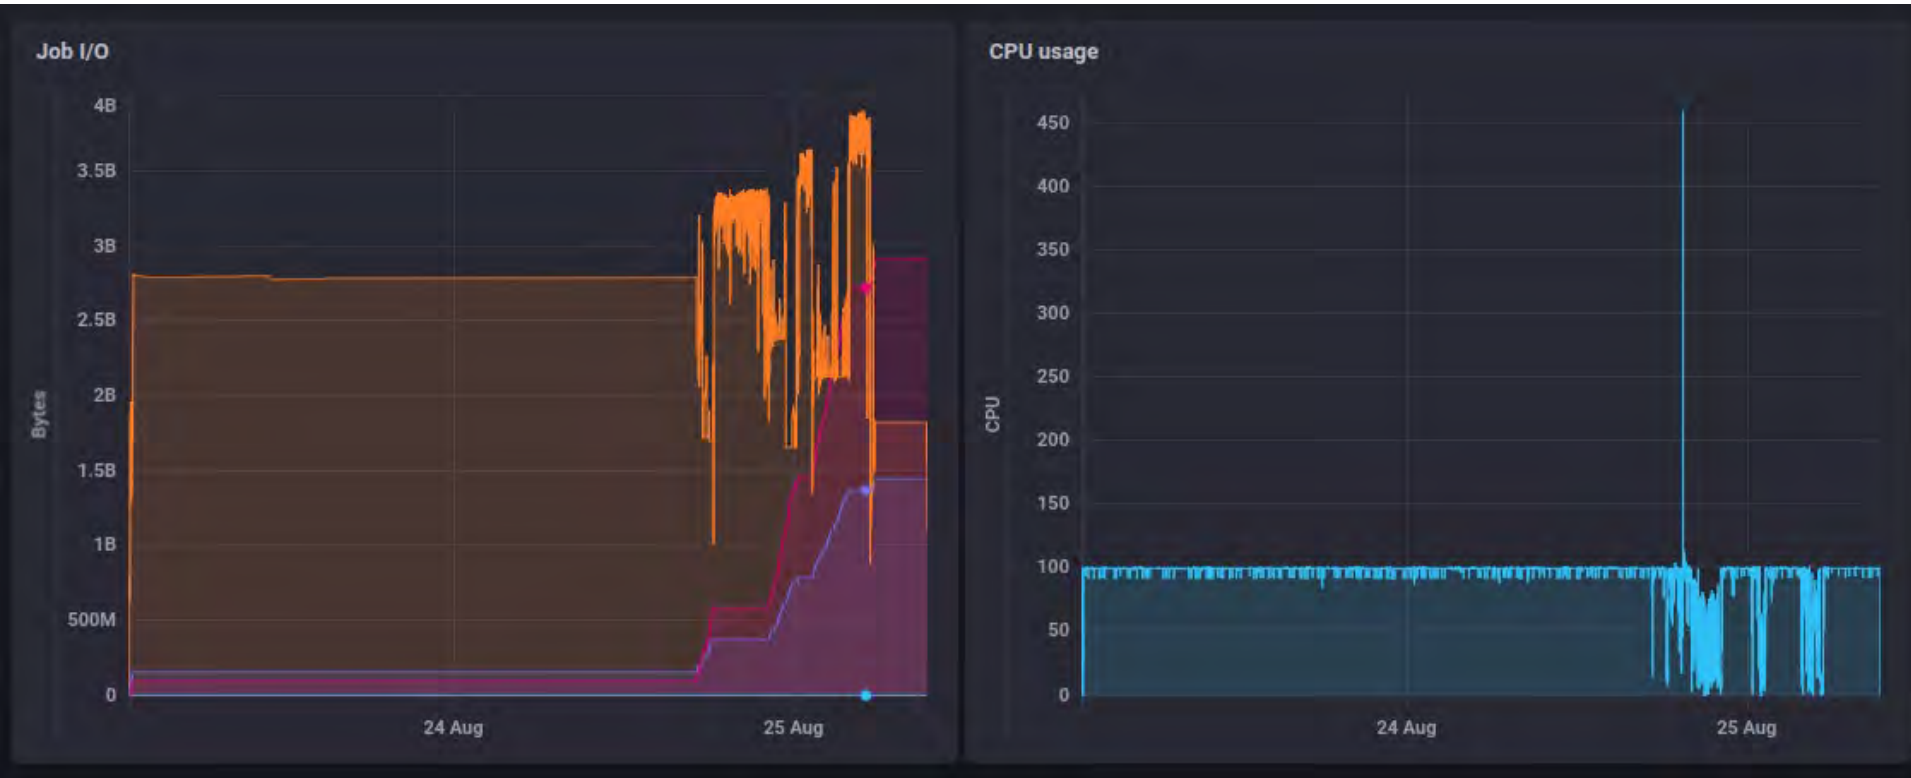
Step 5: New L1 Digi2Dst

21 Aug 2023, 02:53:37 - 23 Aug 2023, 06:53:18

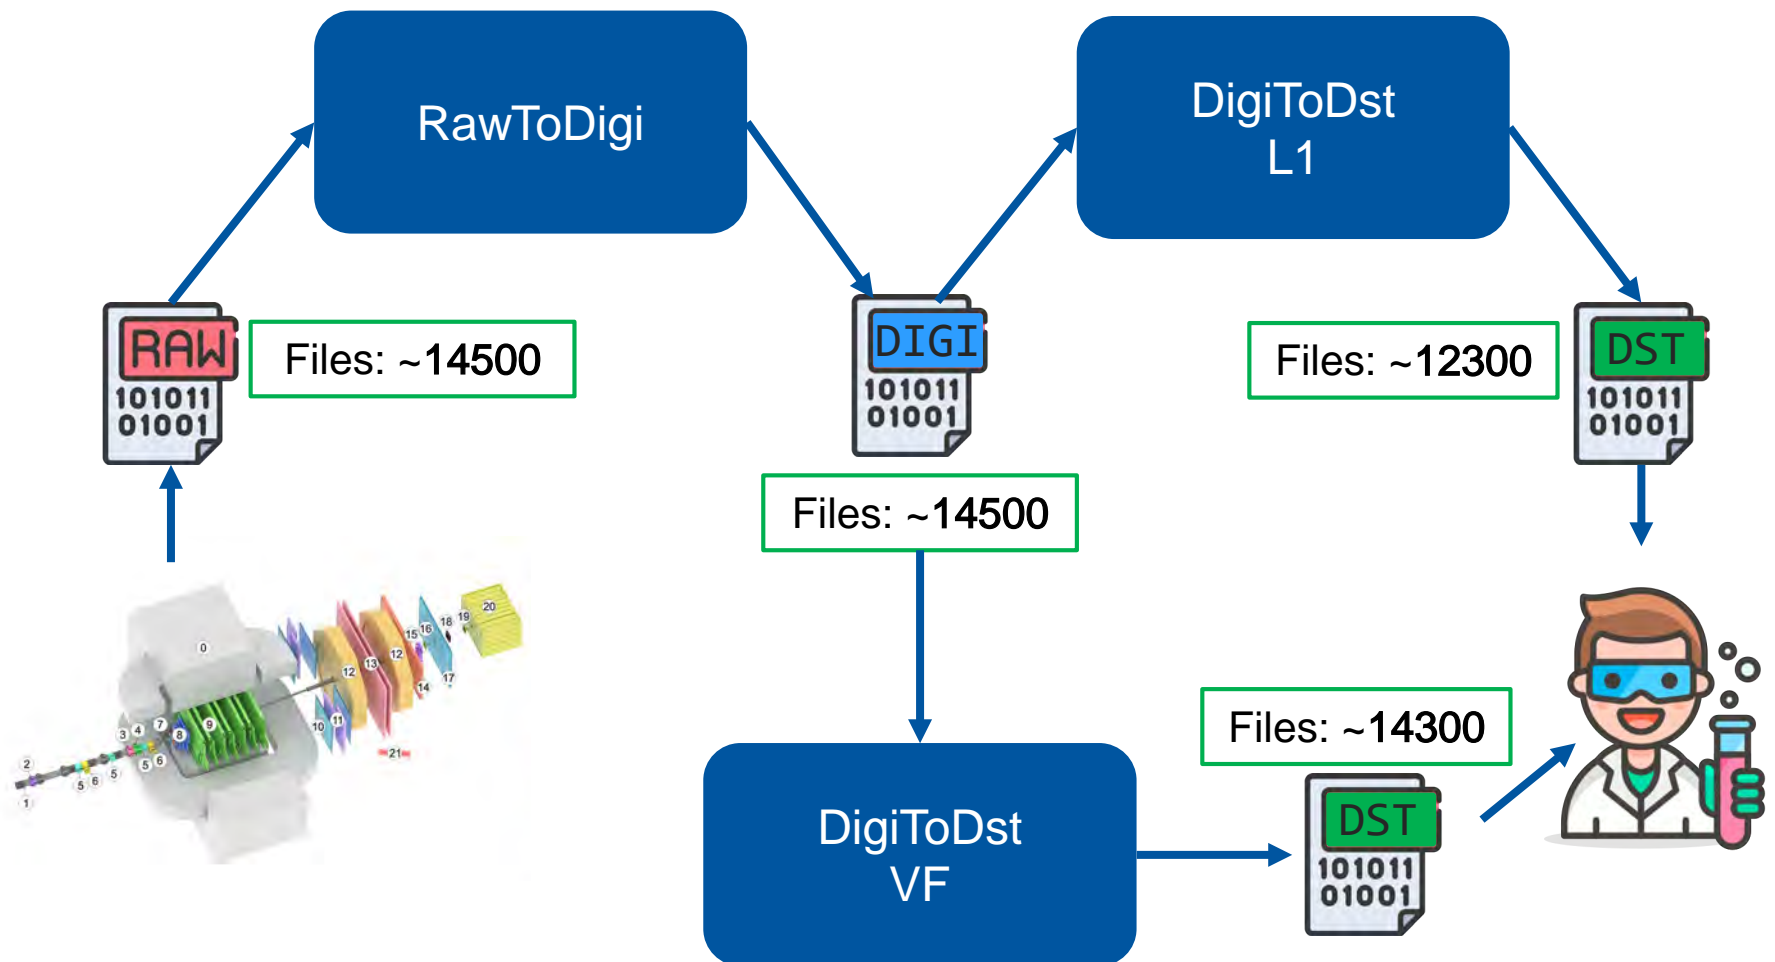
Count of points: 9482




Step 5: New L1 Digi2Dst



Results of reproduction



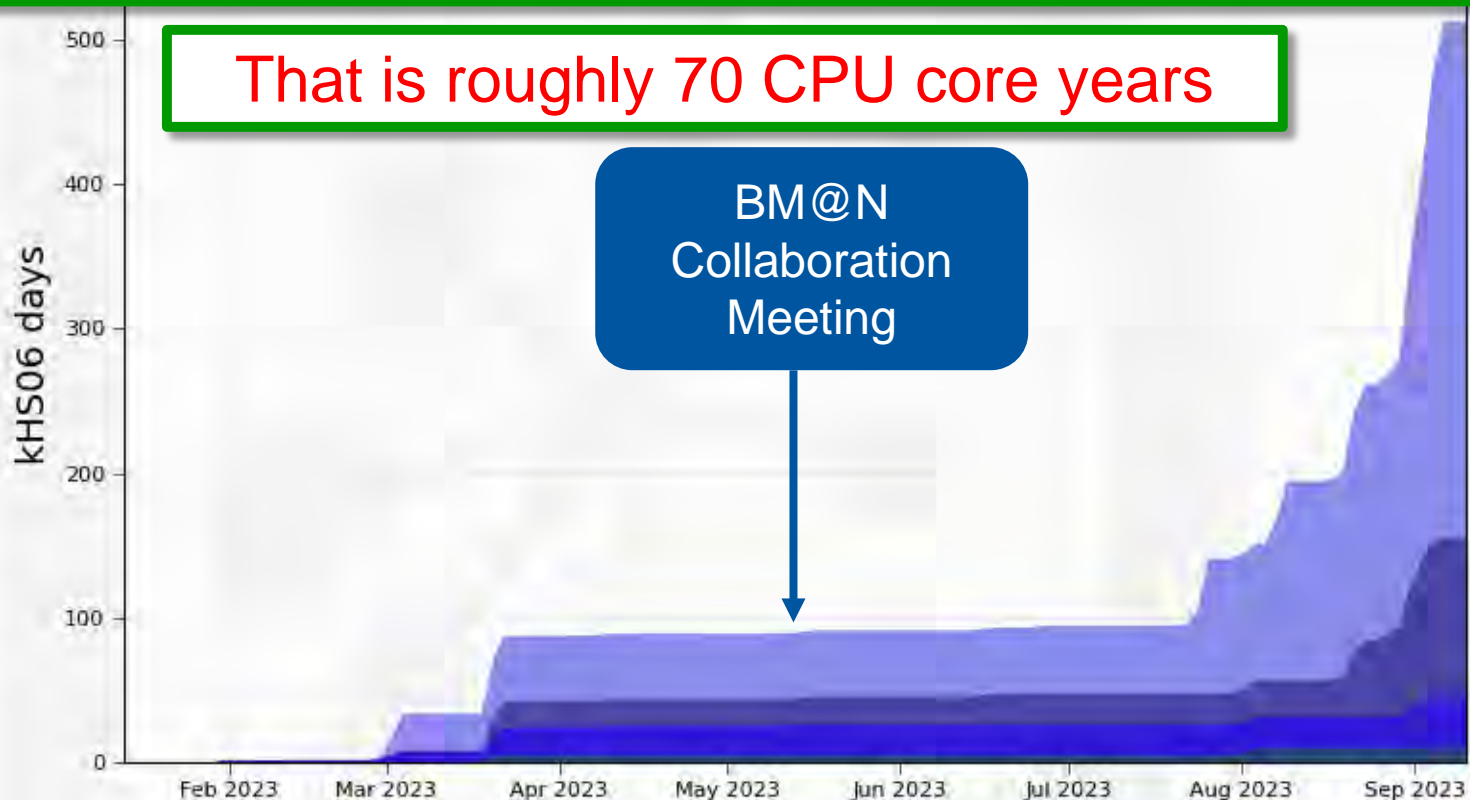
How we have got there



Generated on 2023-09-11 10:58:58 UTC

Consumed resources

We consumed 4 times more CPU resources after meeting comparing to what was consumed before meeting



Max: 513, Average: 103, Current: 513

DIRAC.JINR-TIER.ru	357.4	DIRAC.JINR-LHEP.ru	36.9
DIRAC.JINR-CREAM.ru	109.4	DIRAC.GOVORUN.ru	9.1

Generated on 2023-09-11 11:06:11 UTC

Results

- Summer EOS Failure hit us hard, a lot of time and efforts invested in investigation and working with the issue.
- 5 minor productions and 5 major productions were performed since the last BM@N Collaboration meeting(14-17 May 2023)
- Requirements to the computing capabilities grows. Especially when BM@N and MPD prepare for collaboration meetings
- There was not enough computing resources to perform second extensive check of all BM@N files. So, all results achieved on files that did not suffer from the EOS failure. After special file check it is possible to reupload damaged data and reprocess it.

List of participants

DIRAC: Igor Pelevanyk

BM@N: Konstantin Gertsenberger

Responsible for resources:

Tier-1, Tier-2, EOS: Valery Mitsyn

Govorun: Dmitry Podgainy, Dmitry Belyakov, Aleksandr Kokorev

NICA cluster: Ivan Slepov

Network: Andrey Dolbilov

