# Evaluation of the impact of various local data caching configurations at Tier2/Tier3 WLCG sites

Aleksandr Alekseev, Andrey Kiryanov, Alexei Klimentov,
Tatiana Korchuganova, Valery Mitsyn, Danila Oleynik, Sergey Smirnov,
Andrey Zarochentsev

DLC-2020, 08-10.06.2020

# Background

HENP experiments are preparing for HL-LHC era, which will bring an unprecedented volume of scientific data. This data will need to be stored and processed by collaborations, but expected resources growth is nowhere near extrapolated requirements of existing models both in storage volume and compute power.
=> *Computing models need to evolve.*

This evolution includes multiple aspects:

- Optimized data processing, squeezing the maximum from available CPU/GPGPU/FPGA resources
- Optimized data storage, reduction of the number of copies, different data access methods, full utilization of network resources
- Cost optimizations, no high-end expensive RAID setups, no underutilized CPUs on storage servers, no HDDs with 90% free space on the worker nodes
- Deployment optimizations, scalability and containerization with on-demand expansion into the cloud (both community and commercial)
- Operational cost optimization, more standardized solutions, lower requirements on unique Grid expertise

# Data Lake in HL-LHC R&D Computing Projects

*WLCG and experiments have launched R&D projects to address HL-LHC challenges:*

- **Data Lake**. The aim is to consolidate geographically distributed data storage systems connected by fast network with low latency. The Data Lake model as an evolution of the current infrastructure bringing reduction of the storage and operational costs
- **Intelligent Data Delivery Service** (iDDS). The intelligent data delivery system will deliver events as opposed to delivering bytes. This allows an edge service to prepare data for production consumption, the on-disk data format to evolve independently of applications, and decrease the latency between the application and the storage. The first implementation in April-May 2020 for Data carousel and active ML workflows
- **Hot/Cold storage**. Data placement and data migration between “Hot” and “Cold” storages using data popularity information
- **Data format and I/O**. Evaluating new formats (f.e. parquet) and I/O performance for HENP data
- **Third Party Copy**. Improve bulk data transfers between sites and find a viable replacement to the GridFTP protocol
- **Operations Intelligence**. Reduce HEP experiments computing operations effort by exploiting anomaly detection, time series and classification techniques to help the operators in their daily routines, and to improve the overall system efficiency and resource utilization
- **Data Carousel**. Use tape more effectively and actively in distributed computing context

# Rationale

This research is conducted in collaboration with European Data Lake Project, which is part of DOMA initiative.

In this talk we will show some of the possible ways of optimizing remote data access from the worker nodes in somewhat small T2/T3 setups or dynamically scaled containerized deployments for physics analysis payloads.

This kind of deployment implies the necessity of heavy site-remote read-biased data I/O and time slot (t) allocated for analysis job is normally split into three phases (disregard some overhead): input read ($t_1$), compute ($t_2$) and output write ($t_3$). Sometimes analysis payloads can read and write data while performing computation which makes it hard to separate $t_1$ from $t_2$ and $t_2$ from $t_3$, but in any case at least some data needs to be preloaded before computation can start.

Here we will focus on optimizing $t_1$ and thus improving CPU utilization of a compute resource.

# ATLAS Data Popularity in Users Analysis

## DAOD metrics 2/2

*Users Analysis*

data DAOD datasets (deriv only)

- max used : #181
- #Tasks used DAOD datasets as input

| Not used | 1 | 2 | 3 | 4+ |
|---|---|---|---|---|
| 25122 | 11952 | 6573 | 5635 | 52998 |
| **< 7 days** | **1-2 weeks** | **2 -4 weeks** | **1-3 months** | **3+ months** |
| 11318 | 22400 | 65969 | 376831 | 751417 |

MC16 DAOD datasets (deriv only)

- max used : #1170
- #Tasks used DAOD datasets as input

| Not used | 1 | 2 | 3 | 4+ |
|---|---|---|---|---|
| 73721 | 25602 | 16762 | 14019 | 111403 |
| **< 7 days** | **1-2 weeks** | **2 -4 weeks** | **1-3 months** | **3+ months** |
| 400329 | 68483 | 151691 | 581119 | 957728 |

- Delta = [dataset used as input] – [dataset

8/06/19 *ALL Numbers for prun/pathena (credit to T.Maeno)* 7

There's at least one dataset that was accessed 1170 times.

At the average ATLAS datasets consist of 50 files.

=> Each file in this dataset was accessed at least 20 times if data popularity is evenly split between these files.

# Possible caching configurations

There’s no one-size-fits-all solution because of hardware (especially network) differences on different sites. We were testing three pretty obvious scenarios:

1. A single dedicated cache server (poor external network, good internal)
2. A local isolated cache on every worker node (good external network, poor internal)
3. A shared cache between worker nodes (external and internal networks of the same quality) – *requires some sort of service discovery*

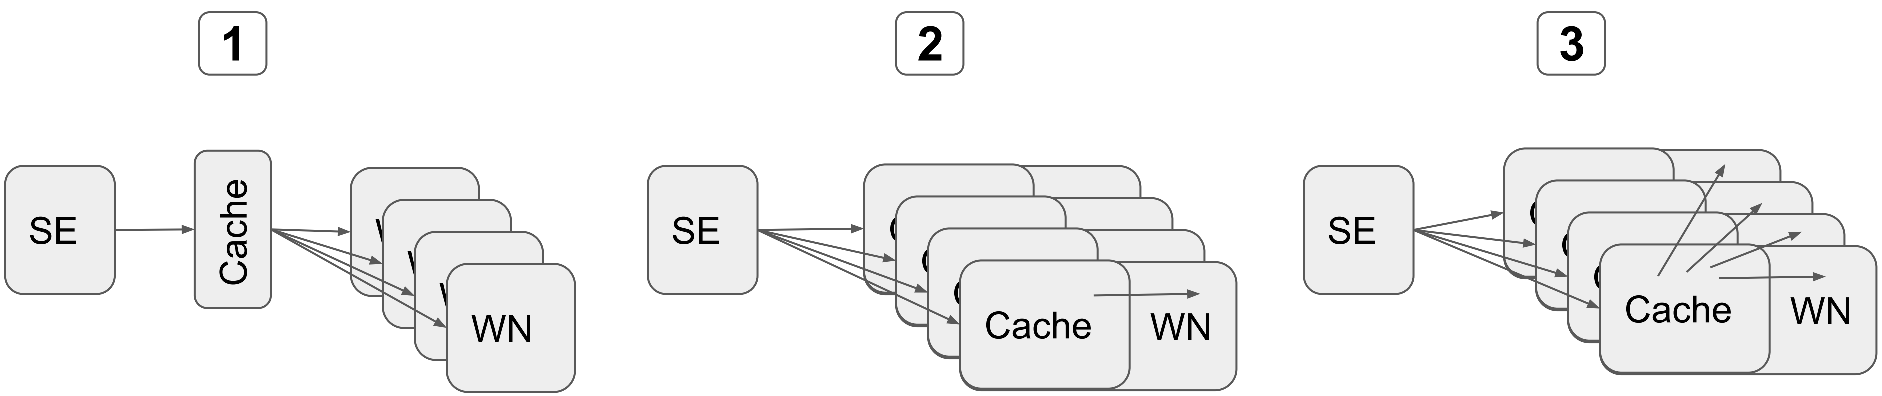

# Russian Data Lake testbed

Petersburg Nuclear Physics Institute — xCache

JINR — dCache

MEPhI — xCache

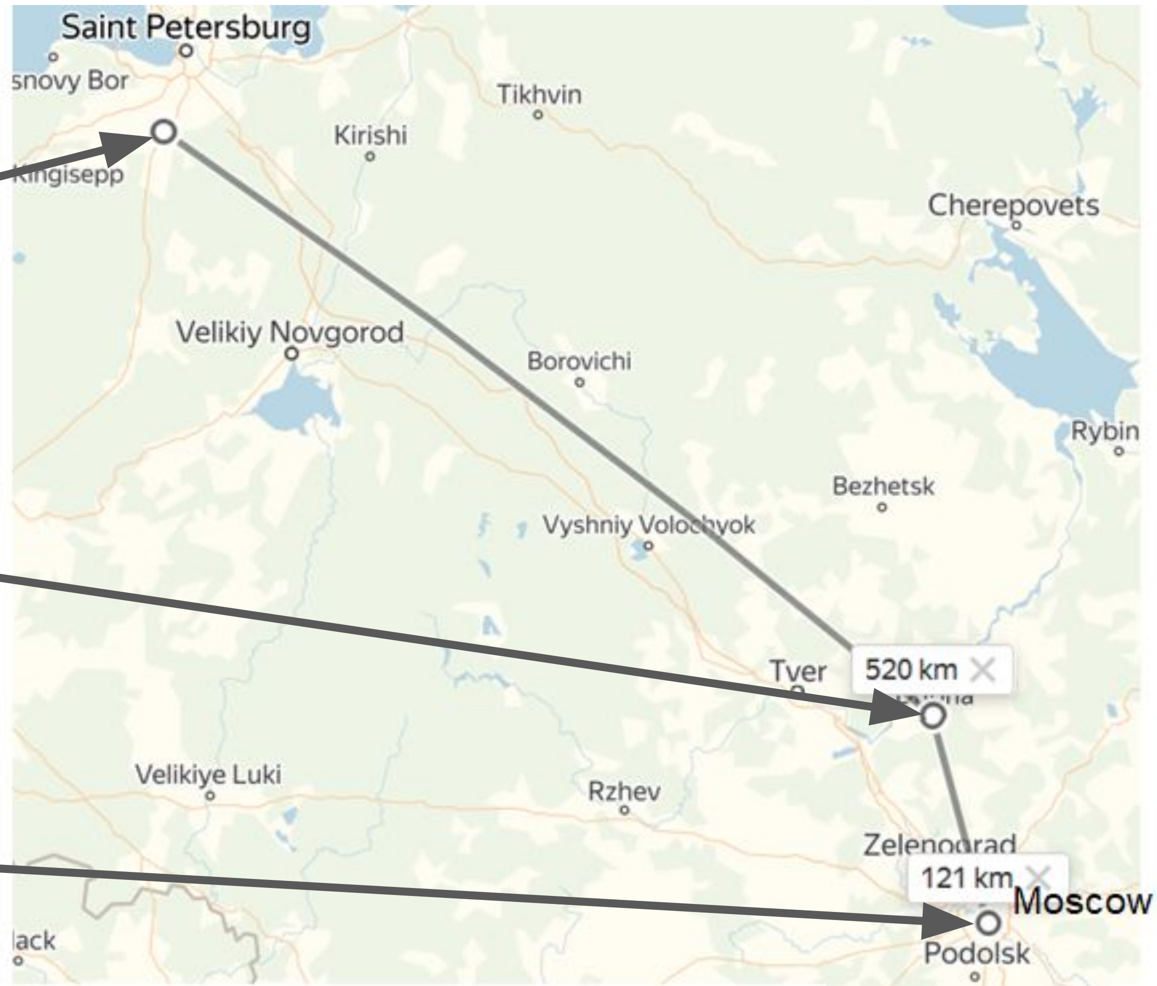

# First tests in 2019 (PNPI and JINR)

*(a slide from our report on HEPiX 2019 Autumn workshop in Amsterdam)*

HammerCloud test results - N20146370 from Template 1099 (copy2scratch)

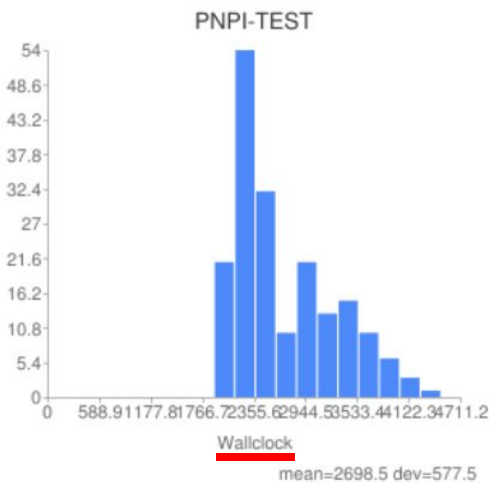


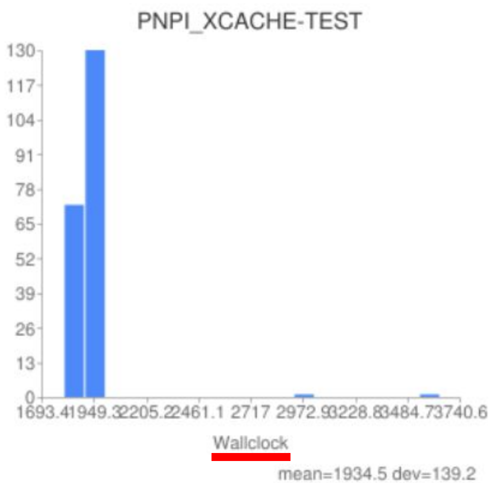


**Wallclock (t)**:

Direct mean time = 2698s ± 577s

xCache mean time = 1934s ± 139s

Difference ~ 770s, ~30%

**Download input files time ($t_1$):**

Direct mean time = 811s ± 574s

xCache mean time = 53s ± 137s

Difference ~ 770s, ~95%

Local (JINR) = 117s ± 17s

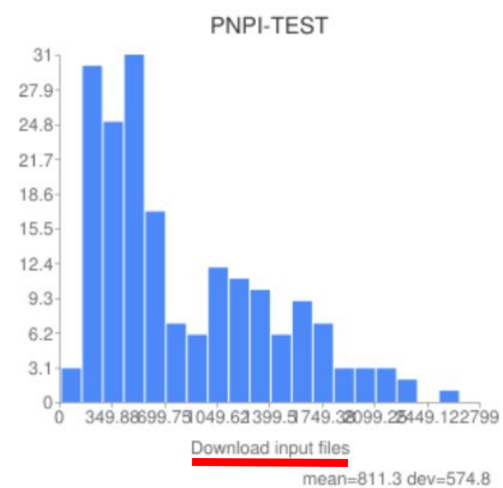


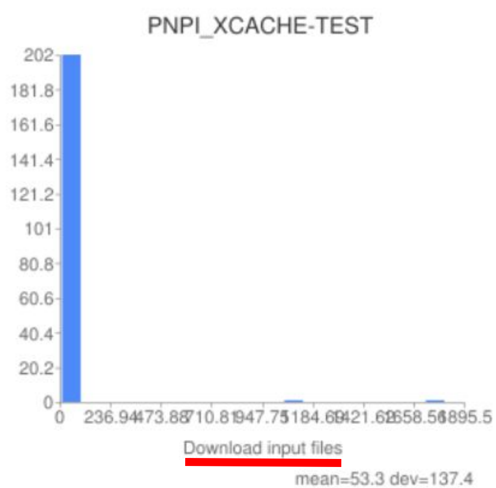


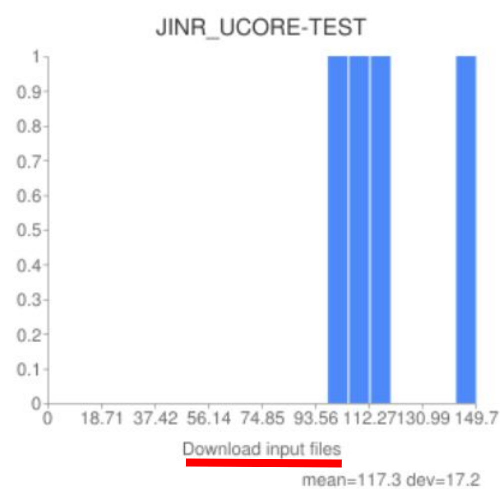

# Testbed changes

- New site - MEPhI (Moscow)
- New Torque with task affinity patches
- New ARC CE deployment
- New monitoring (ELK)
- Addition of node-local tests
- dCache upgrade/downgrade at JINR
- Network backbone reconfiguration in Moscow

Every bullet in this list required quite a lot of effort. New testbed was finally ready for testing in May 2020.

# HammerCloud tests (copy2sctarch) in 2020 at PNPI

Direct read | Cached read

$t_1$

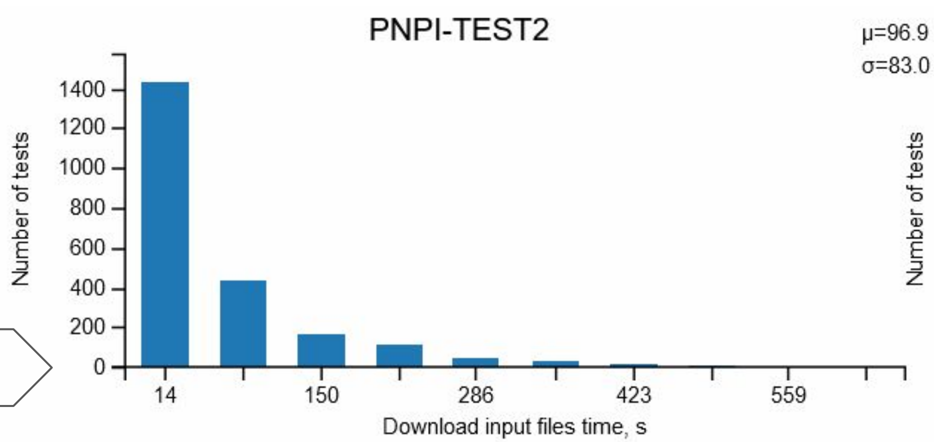


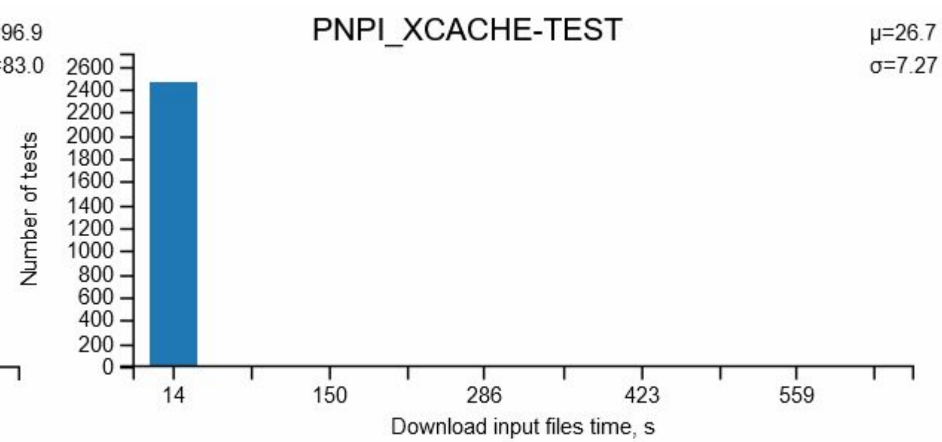


t

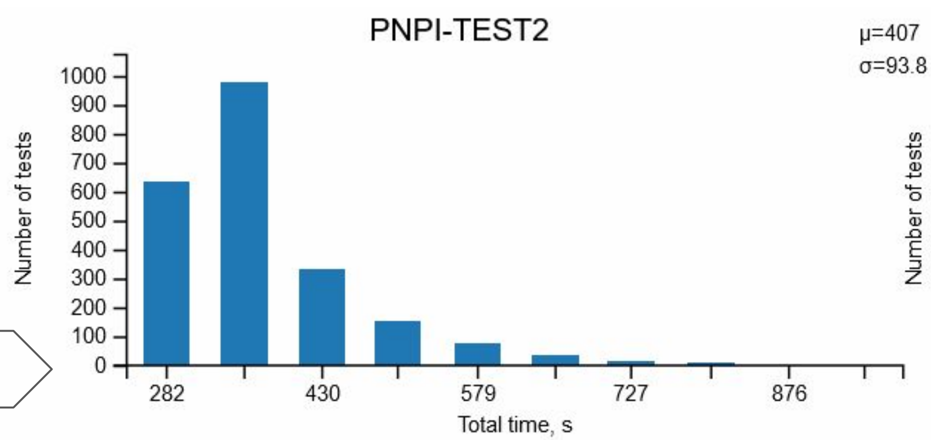


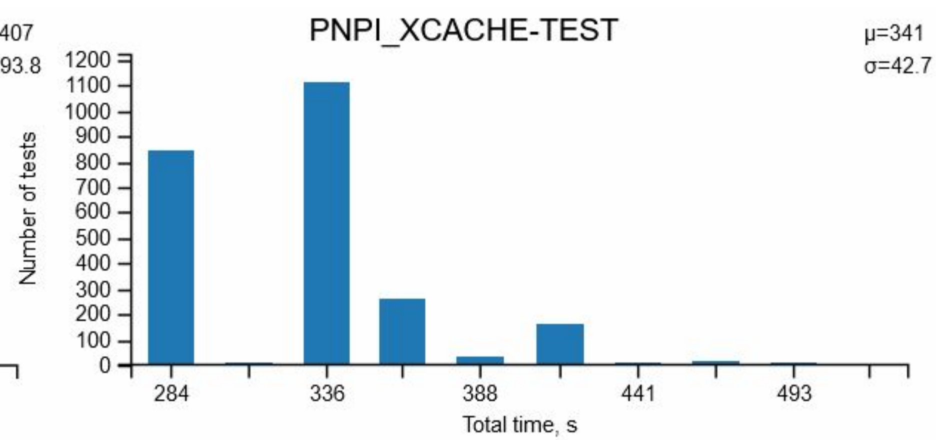

# HammerCloud tests (copy2sctarch) in 2020 at MEPhI

Direct read

Cached read

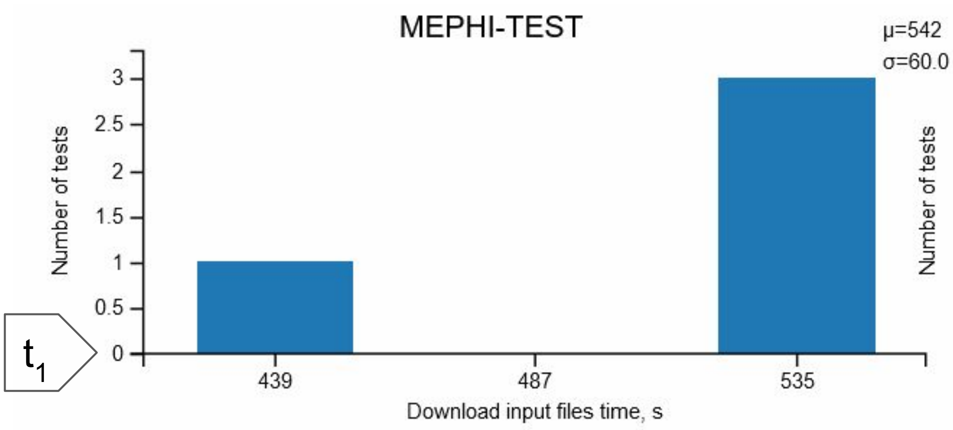


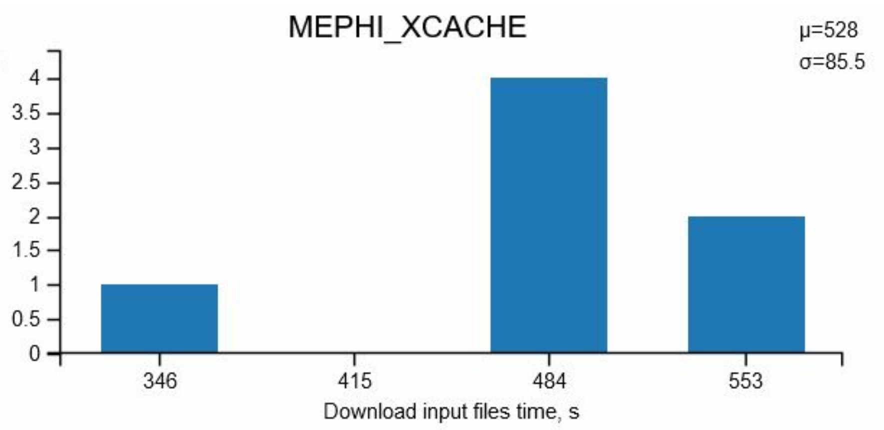


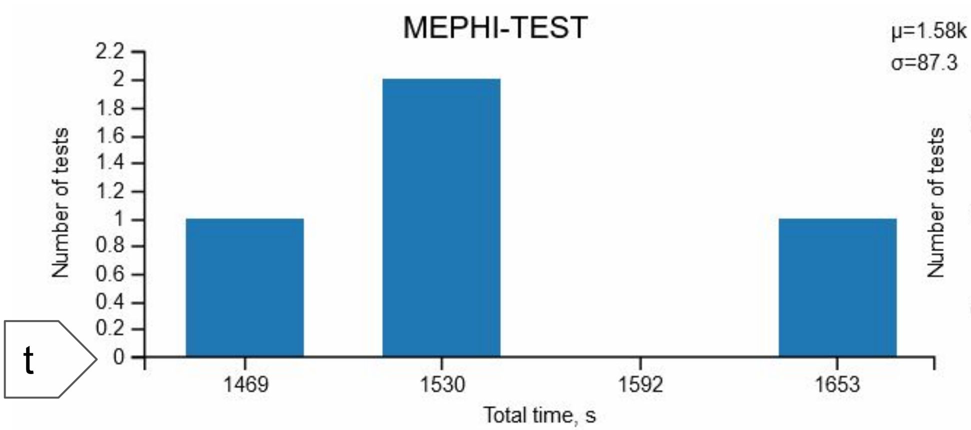


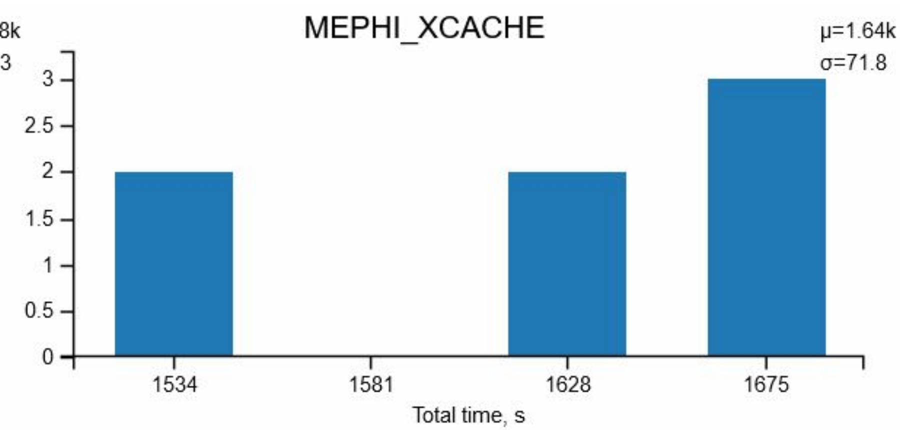

# Synthetic tests (local xcache vs direct access) at PNPI

We submit 20 jobs into the queue, each of which reads one (always the same) file via **local cache**. 12 jobs run in parallel on 4 nodes with independent local caches, each node downloads a file from external storage. This caching scheme shows minimal performance gain in this test.

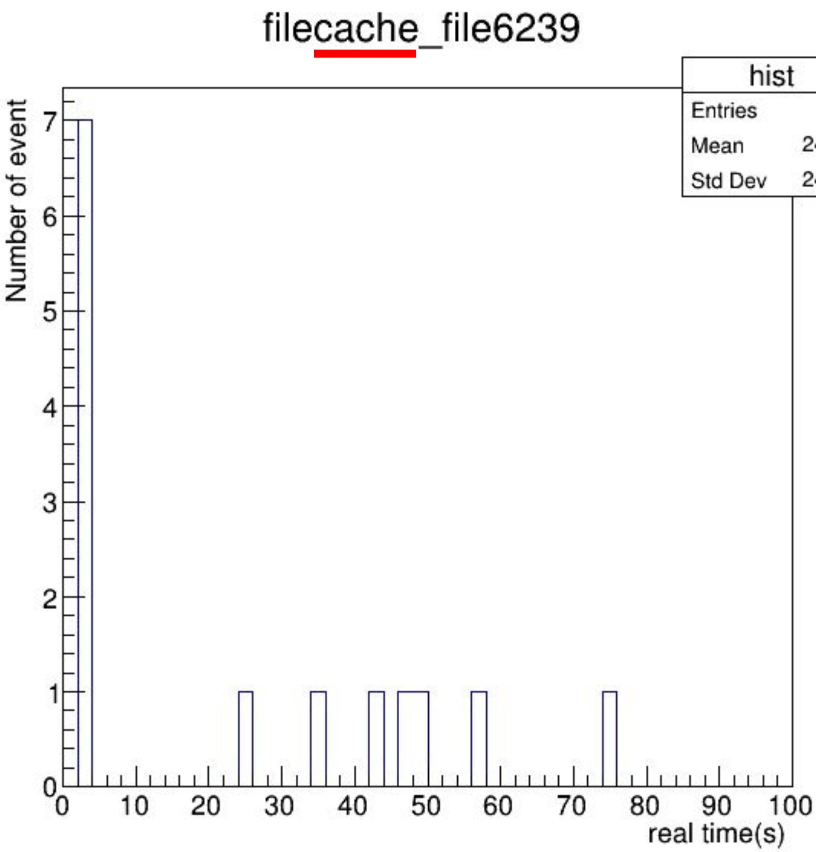


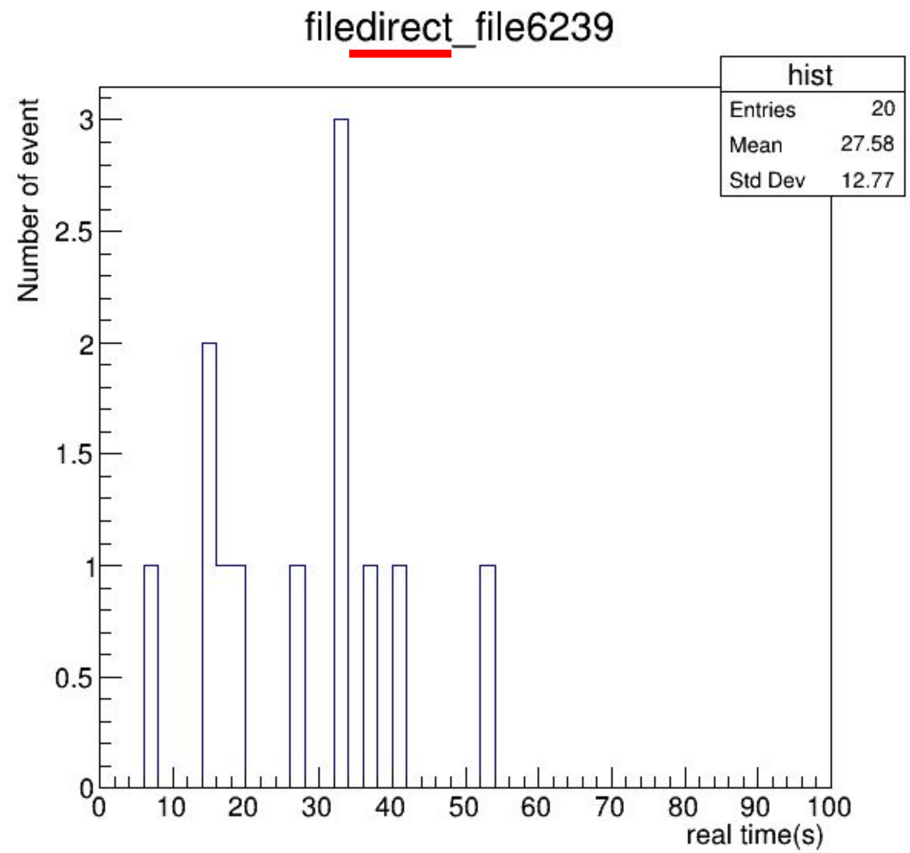

# Synthetic tests (dedicated xcache vs direct access) at PNPI

We submit 20 jobs into the queue, each of which reads one (always the same) file via **dedicated common cache**. 12 jobs run in parallel on 4 nodes, the file is downloaded from external storage only once. This time performance gain is more visible.

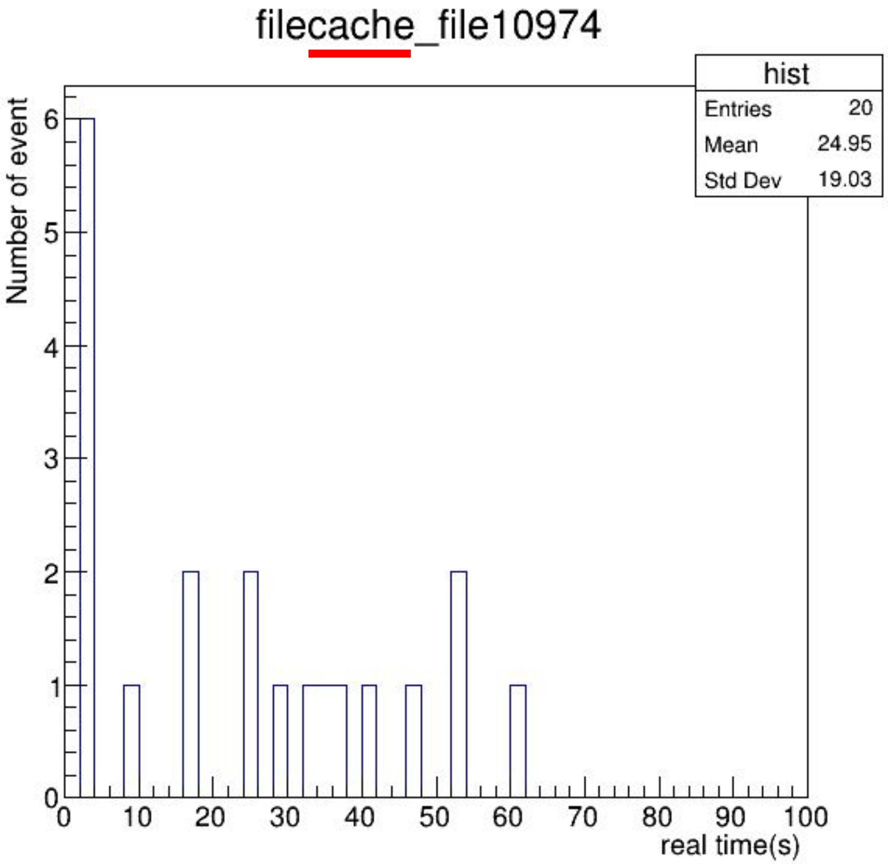


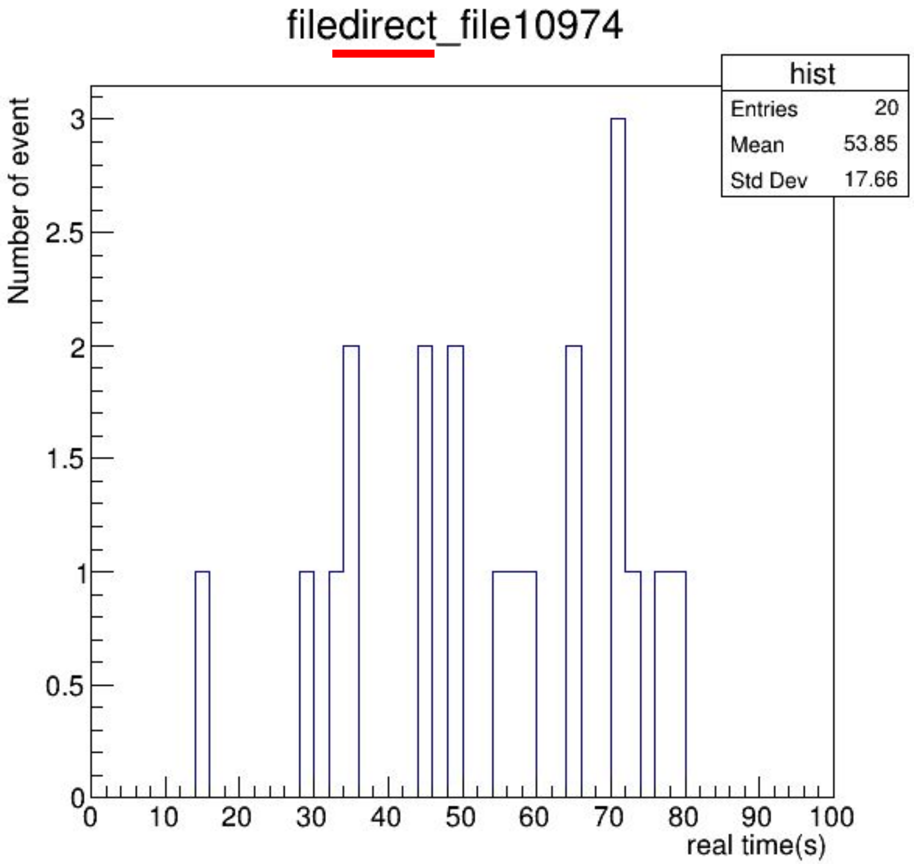

# Synthetic tests (dedicated xcache vs direct access) at MEPhI

We submit 20 jobs into the queue, each of which reads one (always the same) file via **dedicated common cache**. 12 jobs run in parallel on 4 nodes, the file is downloaded from external storage only once. This time performance gain is more visible.

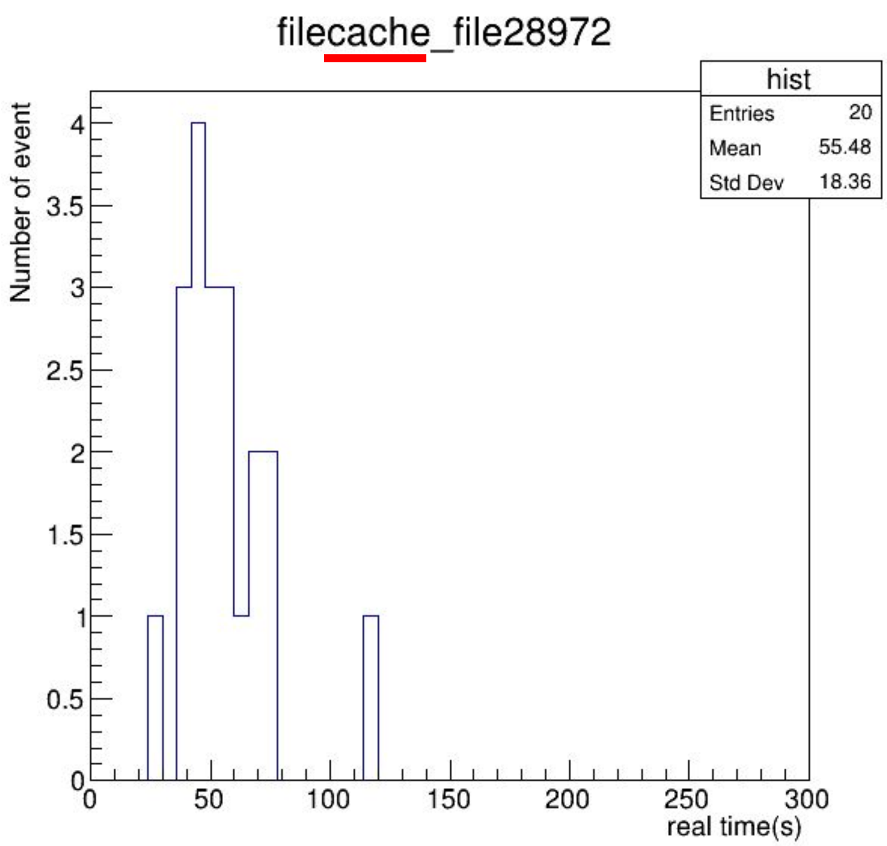


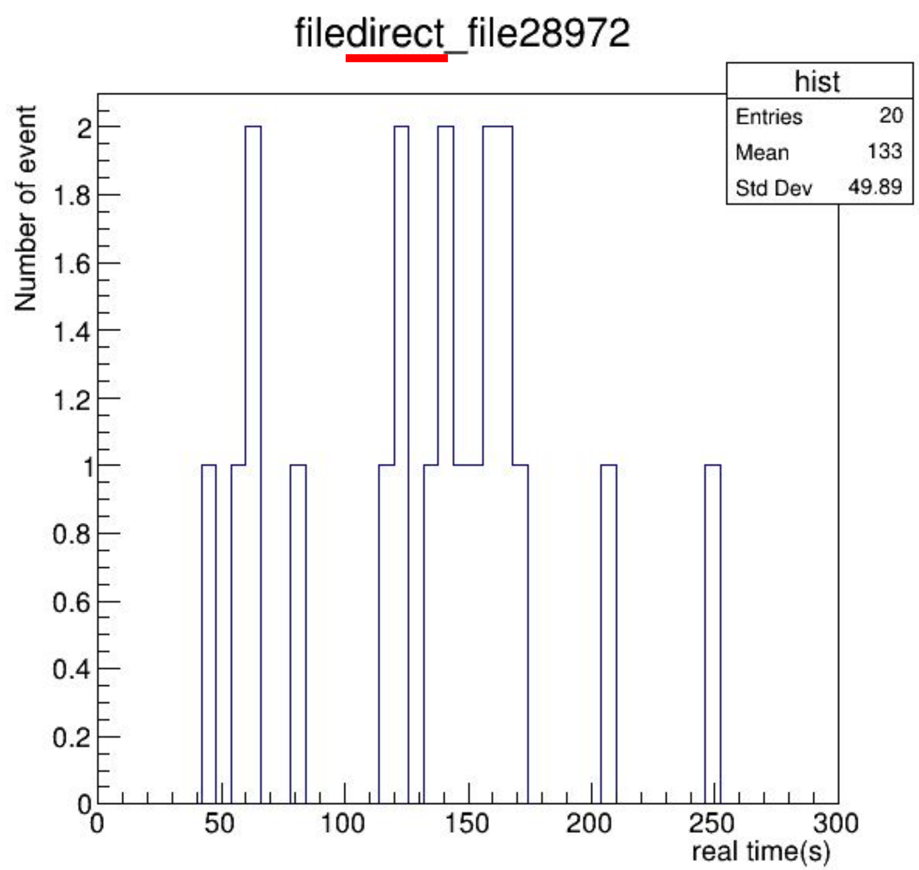

# Conclusions

- We’ve successfully passed “a pilot project phase”
  - We have configured and tested two types of xCache setup
    - dedicated cache
    - local cache
  - We have shown performance benefits of using xCache on smaller sites with
    - synthetic tests
    - real-life ATLAS analysis workloads
- We (as WLCG community) need to address Data Lake topic in a global context. DOMA ACCESS initiative is the first step in this direction. We will work closely with DOMA and ATLAS to define the next steps
- We will be interested to test our setup for other HL-LHC R&Ds, such as Data Carousel, QOS and hot/cold storage, etc.

# Near-term plans

- Set up an extended testbed with some ALICE and CMS sites
- Run multiple workflows and production style tasks (in addition to Hammer Cloud tests)
- Test shared cache scenario
- Evaluate TPC and non-x509 authentication to Data Lake access mode
- Work out xCache authentication for multi-VO setup

# Acknowledgements

It is a collaborative effort in ATLAS and WLCG (Operations, Distributed Computing, Software developers) and sites (JINR, MEPhI, SPbSU and NRC KI - PNPI). Thanks to all.

Thank you to X. Espinal and S. Jezequel for discussions and contribution to HammerCloud tests methodology.

This project in Russia is supported by Russian Science Foundation award No. 19-71-30008 (research is conducted in the Plekhanov Russian University of Economics).

# Thank you!