Scientific data lake for High Luminosity LHC project and other data-intensive particle and astroparticle physics experiments

A.Alekseev, *A.Kiryanov*, A.Klimentov, T.Korchuganova, V.Mitsyn, D.Oleynik, S.Smirnov and A.Zarochentsev

5th International Conference on Particle Physics and Astrophysics

October 6th, 2020

Outline

- HEP Computing (LHC Run1 and Run2)
- HL-LHC timeline and challenges
- HEP R&D projects
- ESCAPE - project to exploit synergies
- Data Lake concept, prototype, timeline
- More challenges ahead

HEP distributed computing software tools

Workflow Management:
"translates" physicist requests into production tasks. **ProdSys2**

Workload Management:
submission and scheduling of jobs & tasks.
PanDA/ALIEN/DIRAC

Monitoring and Analytics:
production jobs & tasks, shares, users. **BigPanDA**, *Kibana, MonaLisa, Perfsonar*

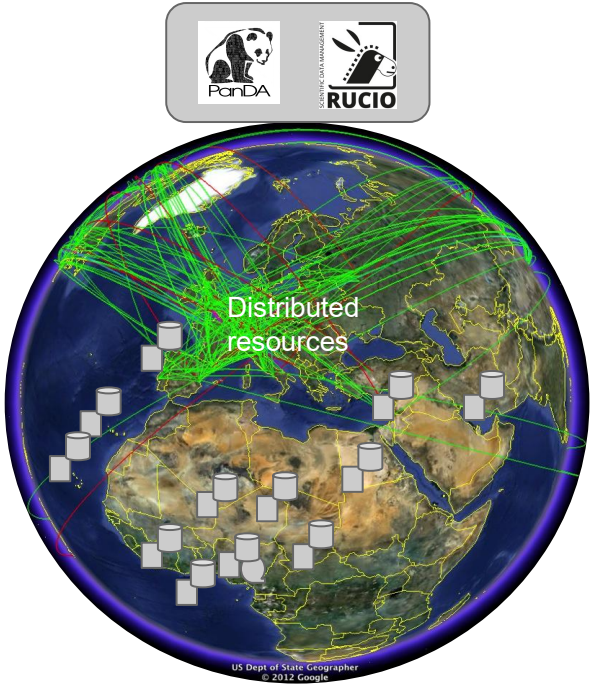

LHC Experiments adopted Grid paradigm
~200 Grid centres in 40+ countries
1M CPU cores (x86)
1EB of storage (50:50 disk/tape)
Global networking (1—100 Gbps)

Data Management:
bookkeeping and distribution of files & datasets. **Rucio**

Information System:
Queues and resources description. **AGIS/CRIC**

Databases:
Conditions and data processing (ORACLE, mySQL, PostgresSQL)

Towards HL-LHC

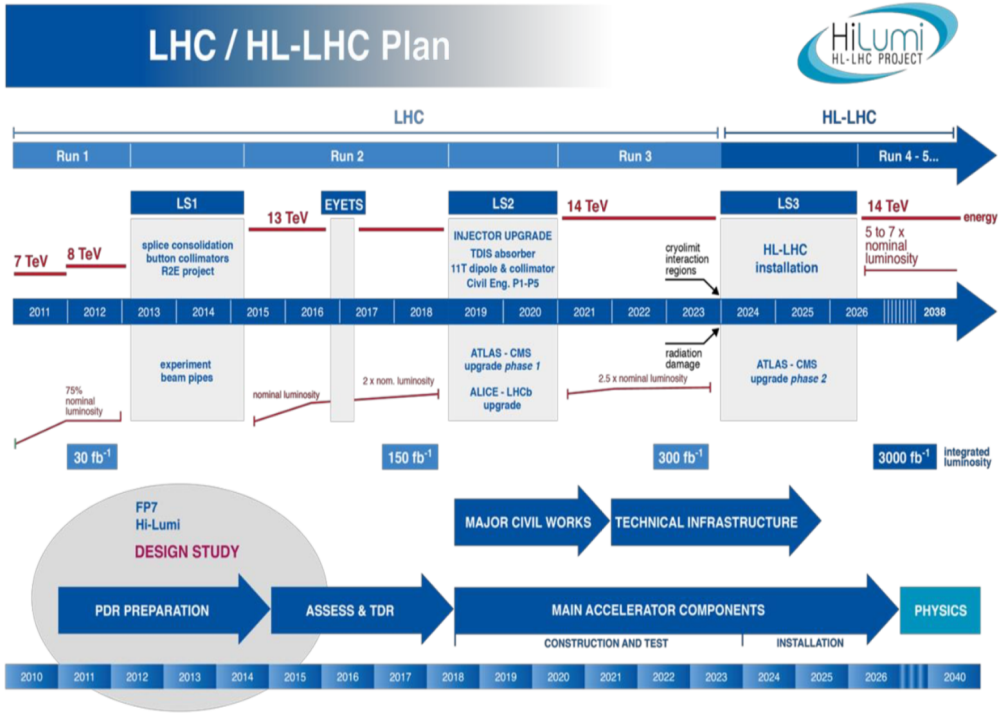

Hot Press

- Extend Run 3 by one year – such that 2024 is included in Run 3 and LS3 starts in 2025;
- Extend LS2 by two months – such that the experimental caverns will be closed on May 1st 2021
- Extend LS2 in well-justified circumstances by up to additional six weeks – ATLAS and LHCb
- Consider to drop one of the ion runs after 2021 and attach it to the scheduled ion period in 2024 – to allow cool-down before LS3
- Decide on the final beam energy after the magnet training at the end of the extended LS2

The High Luminosity LHC Challenge

Growth in CPU Needed

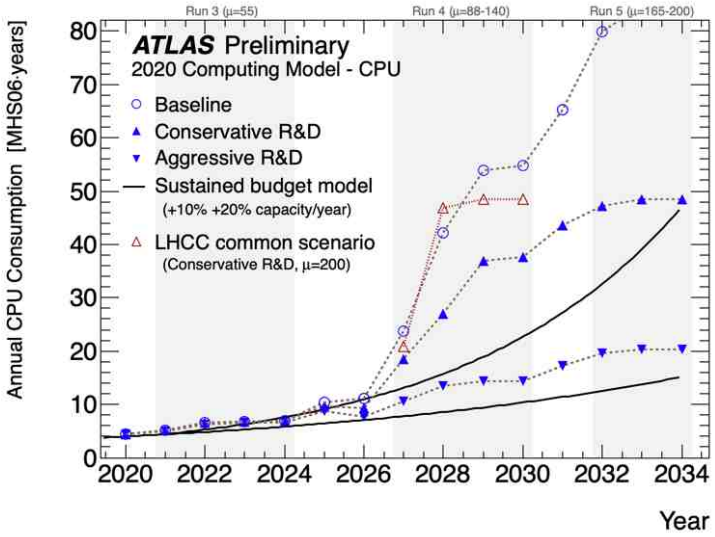

Growth in Disk Storage Needed

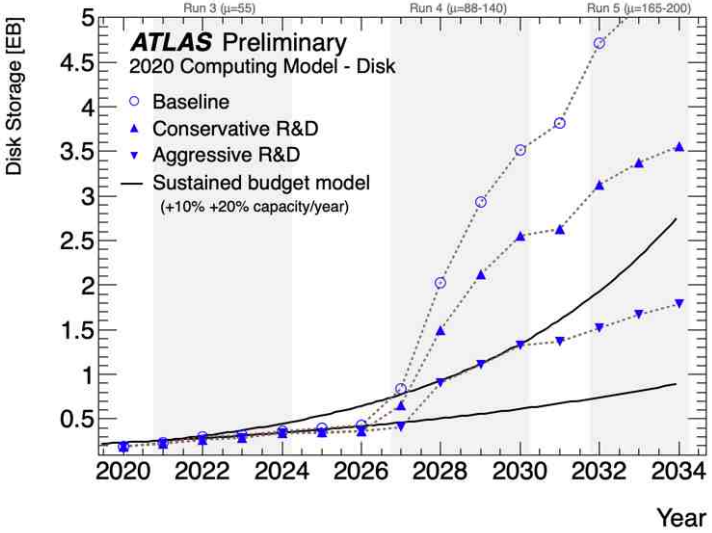

- High Luminosity LHC will be a multi-exabyte challenge where the envisaged Storage and Compute needs are a factor 10 to 100 above the expected technology evolution.
- LHC experiments have successfully integrated HPC facilities into its distributed computing system. “Opportunistic storage” basically does not exist for LHC experiments.

The HEP community needs to evolve current computing and data organization models in order to introduce changes in the way it uses and manages the infrastructure, focused on optimizations to bring performance and efficiency not forgetting simplification of operations.

Astronomy and Particle Physics

- Many common topics in computing and data handling – The ESCAPE project to exploit synergies
  - Builds on communities' complementary excellences in data stewardship:
    - Astronomy Virtual Observatory infrastructure
    - HENP expertise in Exabyte-scale data management and large-scale distributed computing
  - Builds on existing inter-RI synergies, intersections; overlapping competence and authority of national stakeholders
  - Recognizes that ESCAPE communities will be Exascale data generators, early adopters of ICT and data management innovations, push state-of-the-art
  - Both Observatory- and Facility- operations require global, open access to data, long term curation, and sustainability

Data Lake – 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-Cold" storage 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 transfer 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 effective and active in distributed computing context.

Data Lake Concept

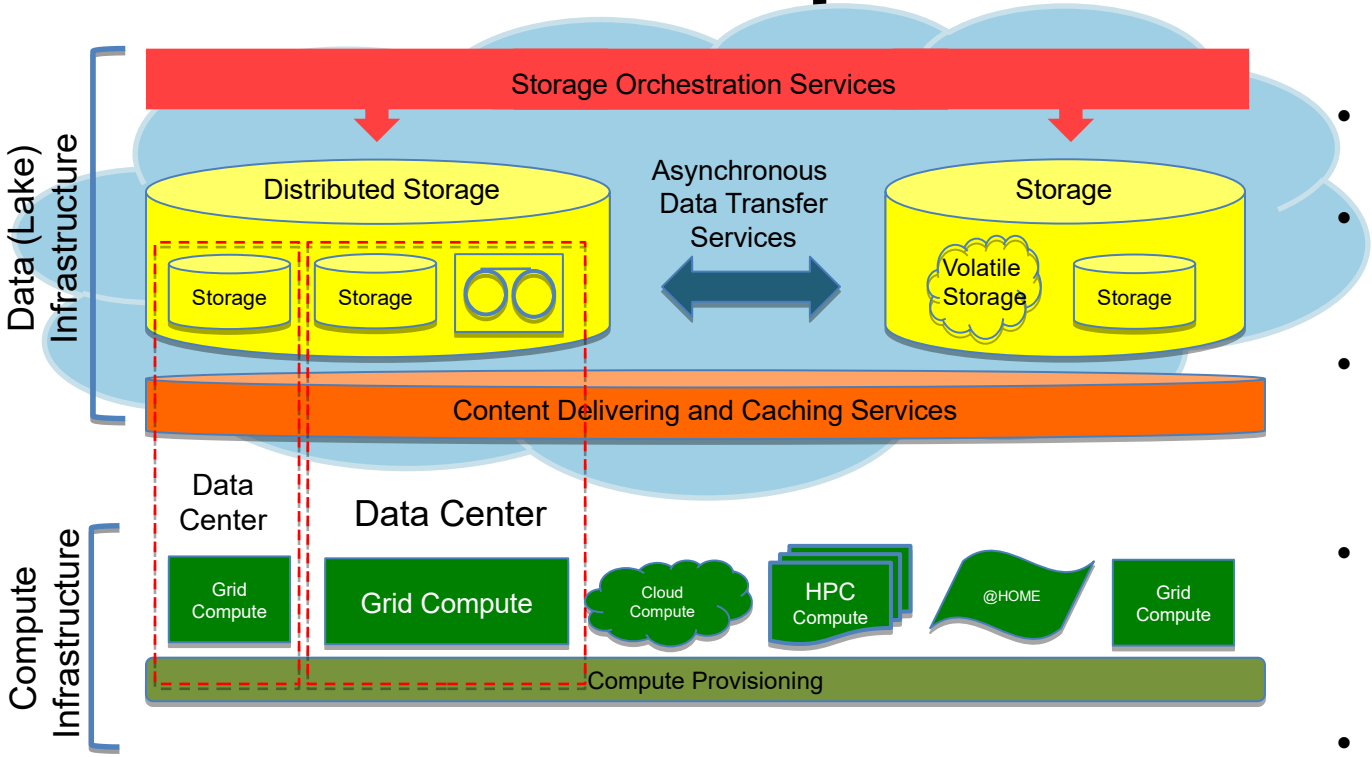

- Federation through token-based AAI
- Policy-driven data replication and distribution
- Distributed storage for reliability, accessibility, sustainability
- Serving data, remote, cached, streaming, to heterogeneous compute facilities
- Hide complexity – transparent access to data

Data Lake Building Blocks

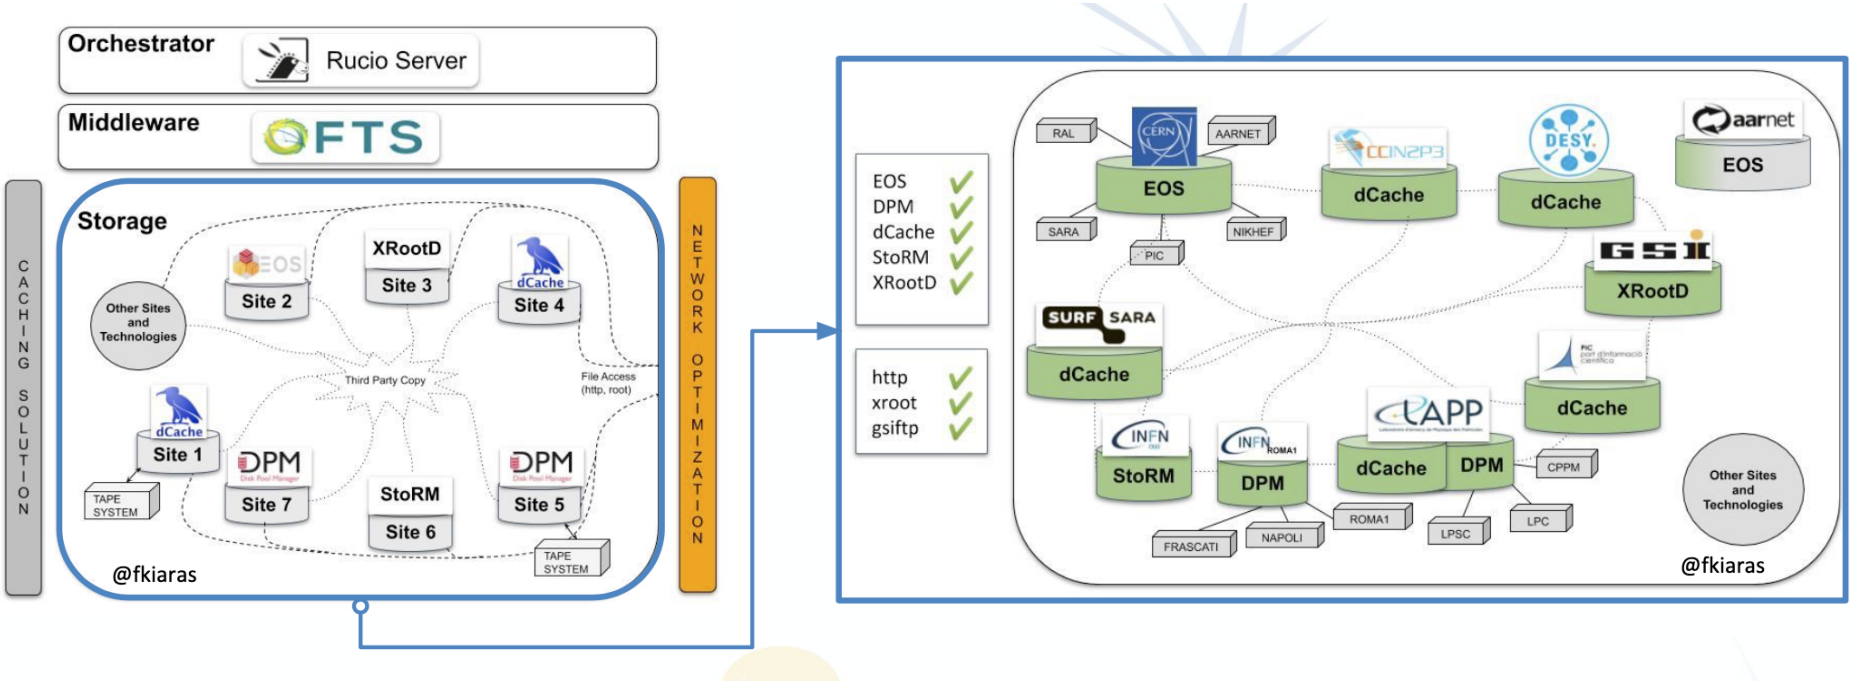

Data Lake Proof of Concept Phase

- Now – mid 2021 : Proof of Concept and prototyping phase
  - Data in/out; Content delivery and latency hiding mechanisms (data caching?)
  - HEP core SW access and deployment
  - AAI mechanism(s)
  - Identify workflows (data (re)processing, end-user analysis)
  - Identify network requirements
  - What about new (non-x86) architectures ?

Data lake prototype in Russia

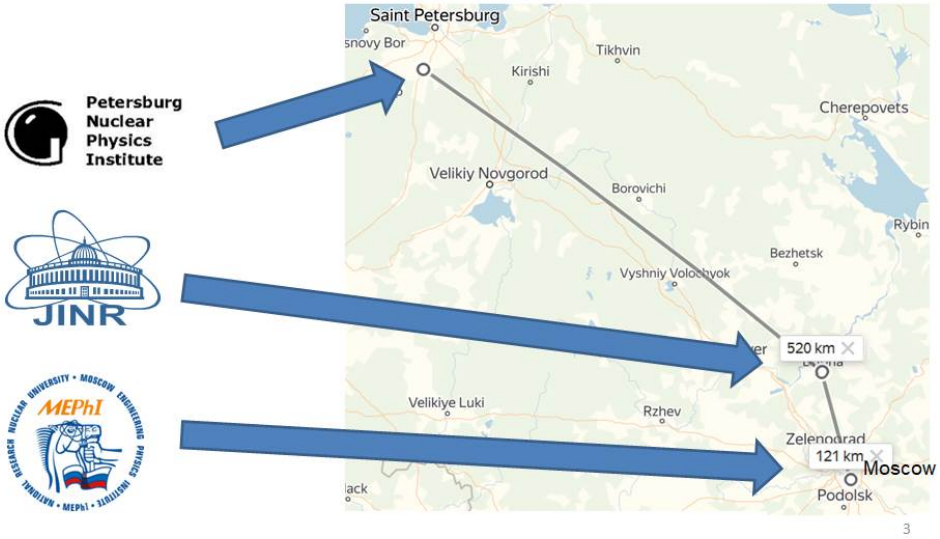

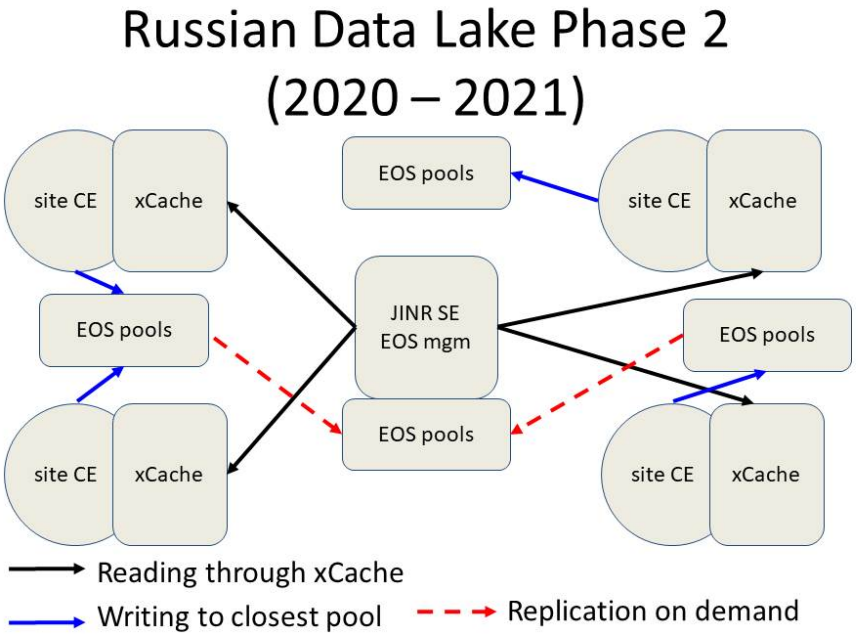

More Challenges Ahead

- We need to bring together Astronomy, Astrophysics, Astro-Particle, High Energy and Nuclear Physics communities
  - Common interests in Exabyte-scale FAIR data management and open science
- We successfully and quickly passed “a pilot project phase” to prototype a federated storage
- We need to address data lake project in National and International context
- We need to integrate academic and commercial data lakes with Grid infrastructure and supercomputers

Thanks

- Thank you to all who contributed materials, discussions and -- most of all – work for Data Lake prototyping and conceptual design
- Especially…
  - Stephane Jezequel, Simone Campana, Xavier Espinal, Ian Bird

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)