

22nd International Conference on Computing in High Energy and Nuclear Physics, Hosted by SLAC and LBNL, Fall 2016

## CHEP 2016: Track 4 Highlights

Wahid Bhimji (LBL)  
Patrick Fuhrmann (DESY)  
Elizabeth Gallas (Oxford)  
Maria Girone (CERN)



This talk includes some “Highlights” (not a Summary)

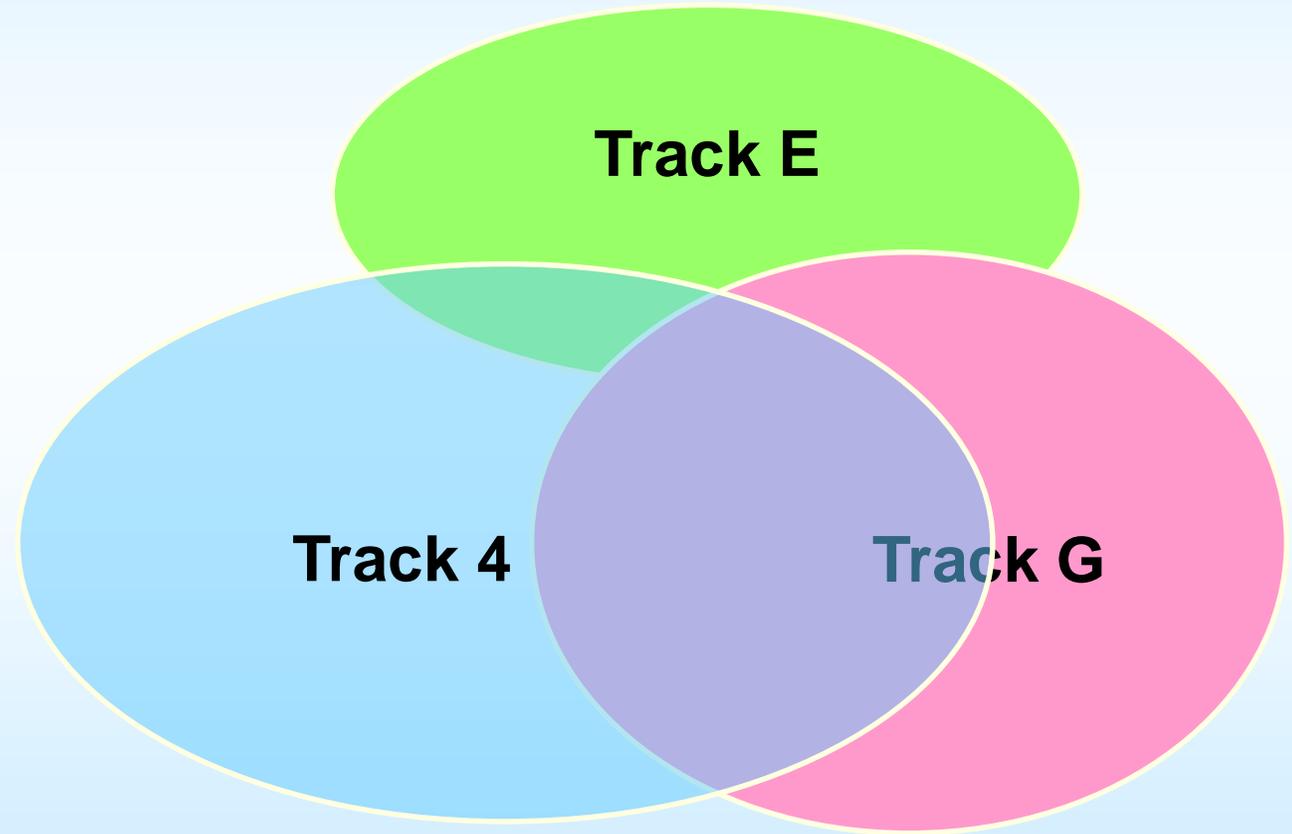
And only 15 minutes !

No time to cover in detail really any presentations

- Thanks to all speakers, both oral and poster
  - For many valuable contributions describing your work and for engaging discussions
- Thanks to my co-convenors
  - Apologies: since they had really no time to review these slides
  - On the other hand, they can deny any responsibility !



- Track 4 includes ...
  - data management; storage management; local I/O and data access; mass storage systems; file systems; object dictionaries; event stores; metadata and supporting infrastructure; databases; access patterns and caching strategies;
- Abstracts: count 85
  - Oral: count 31 ... in 6 Sessions:
    - Storage Middleware
    - File-systems and Cloud Storage
    - Wider HEP and Beyond
    - Experiment Frameworks (1)
    - Experiment Frameworks (2)
    - Data Transfer, Caching, Federation
  - Poster: count 41 ... in 2 sessions:
    - Track A: 23 posters
    - Track B: 17 posters
- Many contributions in other tracks cross-over into Track 4 scope



1. Systems, technologies, and considerations:
  - AFS, DNS, SRM, DPM, NFS, ZFS, XROOT\*, dCache, CernVM, CVMFS, CASTOR, EOS, HPC, Ceph\*, caching, messaging, Object Stores, federations, grid and cloud, storage and processing, networks and data transfer
2. Experiments ... using #1:
  - Belle II, BESIII, AMS, ALICE, ATLAS, CMS, protoDUNE, LHCb, MICE, XFEL
3. Also topics on enabling/improving quality of scientific results, usability

- Metadata provisioning:

- Finding stuff
- An event, by itself, is not complete

- Analytics of processing:

- tuning, forecasting



→ So what's new ... growing, evolving, maturing ... old (retiring?)

→ Cycle time of any entity: determined by its limitations in an evolving climate

→ Performance, Scalability, Connectivity, Flexibility, Usability, Cost



### Common schemes in storage middleware

- Some scaling out into 100 PByte domain
- Make storage cloud-like
  - S3 (http based) cloud protocol defined by Amazon
  - Assigning quality of storage
- Using Ceph as a storage backend:
  - Simplifies Operations and Maintenance
  - Building of object stores
    - Middleware: using for managing files
    - Experiments:
      - using to communicate information and
      - avoiding namespace lookups (reference by an object ID)

Many talks Monday  
and many posters

Important to choose  
non-proprietary standards  
to serve the broader HEP  
and non-HEP communities

- “Storage Quality-of-Service in Cloud-based Scientific Environments: A Standardization Approach”
  - **Quantifies** striking a balance between latency, bandwidth, durability, and cost for a storage implementation appropriate for the use case
- “Unified data access to e-Infrastructure, Cloud and personal storage within INDIGO-DataCloud”
  - give users a unified interface for managing and accessing storage resources provided by different storage providers and to enable them to treat all that space as a single virtual file system with standard interfaces for accessing and transfer, including CDMI and POSIX
- “SciDAC-Data, A Project to Enabling Data Driven Modelling of Exascale Computing”
  - exploit decades of information (analysis patterns and data organization) collected by the Fermilab Data Center to develop realistic models of HEP analysis workflows and data processing.

- **“Taking HEP data management outside of HEP”**
  - Explores use of SAM (FNAL’s data management system) to genome data
    - This metadata based catalog with replication tracking and framework for organizing analysis seems well suited and is being deployed for the project
- Non-HEP Community solutions
  - Many creative and viable solutions to their unique challenges
    - Photon science (XFEL)
    - European projects supporting data federations
  - Some attendees coming to learn from HEP solutions

## “Design of the ProtoDUNE experiment data management infrastructure” (Deep Underground Neutrino Experiment)

- protoDUNE: beam test of a large-scale DUNE prototype at CERN
  - Test beam data on a scale of LHC Run 1 experiments
    - Designing DAQ, data handling and mechanisms for data distribution

## “jade: An End-To-End Data Transfer and Catalog Tool”

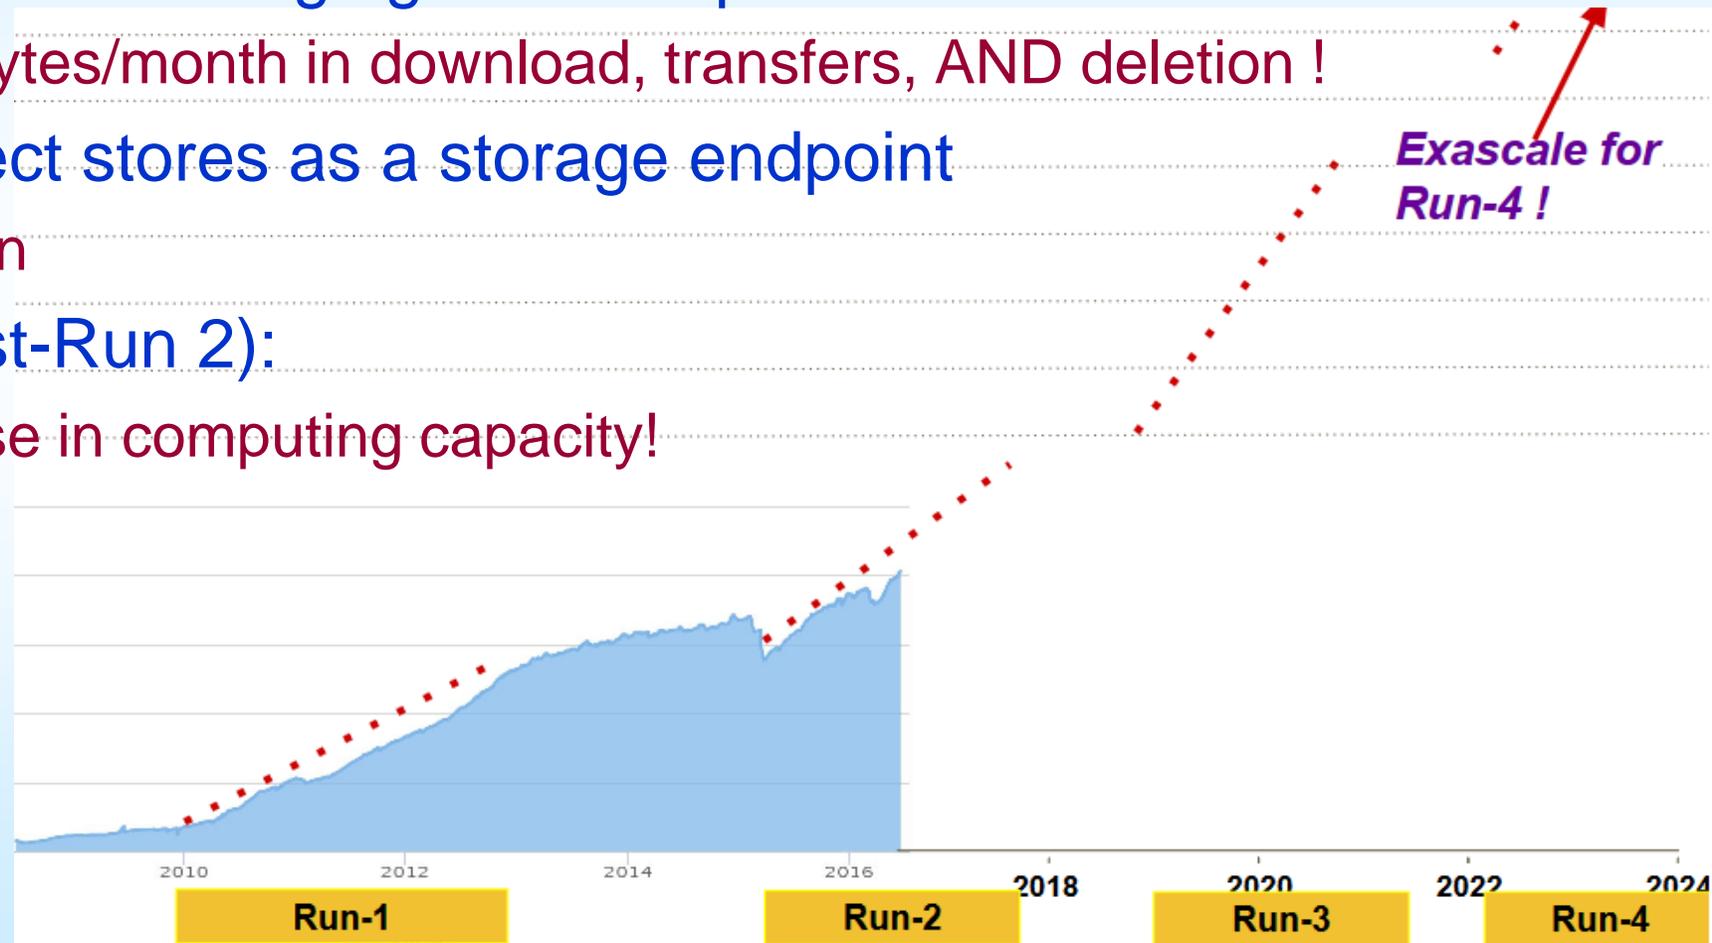
- Addresses the unique challenges in archive and transfer software for the IceCube Neutrino Observatory at the Geographic South Pole
  - Slides include “lessons learned” which apply to experiments large and small

## “Data Management & Database Framework for the MICE Experiment”

- Muon Ionization Cooling Experiment at RAL (UK)
  - Some data handling improvements;
  - Much development on their Configuration Database implementation

# “Experiences w/the new ATLAS Distributed Data Management System” (Rucio)

- working well for Run 2: managing 1B file replicas on over 100 sites
  - Handling many PBytes/month in download, transfers, AND deletion !
- Rucio: can use object stores as a storage endpoint
  - 5 sites in production
- Looking ahead (post-Run 2):
  - need O(10) increase in computing capacity!



- Many contributions using analytics ... a few examples:
  - **“Using machine learning algorithms to forecast network and system load metrics for ATLAS Distributed Computing”**
    - Offers potential for improving:
      - Data rebalancing, Hot replication, Placement and Source selection, Robustness
  - **“GRID Storage Optimization in Transparent and User-Friendly Way for LHCb datasets”**
    - Varied their data popularity algorithm to find a better optimized data replication and distribution strategy
  - **“Dynamo - The dynamic data management system for the distributed CMS computing system”**
    - 2<sup>nd</sup> v of Dynamo: includes components for data deletion and distribution
  - **“Integrating Prediction, Provenance, and Optimization into High Energy Workflows”** (Belle II)
    - Simulate Belle II’s unique anticipated workflow
      - Exploring in advance, then optimize at run time

## Data transfer: Exploring improving efficiency

- “Next Generation high performance, multi-dimensional scalable data transfer”
  - For LCLS from SLAC → NERSC, In collaboration w/ commercial vendors

## Caching ... many contributions touch on it in many areas

- By whole files or at the data block level and for metadata/conditions
  - Exploring if/when caching is effective ... to increase throughput w/o stress on Networks
  - Including dedicated facilities/tests

## “HTTP as a Data Access Protocol: Trials with XrootD in CMS' AAA Project”

- Demonstrate ability of using HTTP data federations from CMS software  
→ necessitated by different protocols to commonly used federations

## “Accessing Data Federations with CVMFS”

- Federations contain a lot of data, but no Posix access to files (which users like)
  - CVMFS provides a Posix namespace which is familiar to users
- This contribution explores providing this, including authentication aspects

- Sofia predictions (CHEP'18) ? Extrapolate from current developments (some of which I haven't mentioned yet)
  - Improvements in **data and storage** management and access
    - With increased use and refinements of latest technologies
  - More detailed **analytics** on transactions (Big Data)
    - Any process (or process set) can be refined
    - Trends more readily recognized
  - More progress in **Conditions** data storage, delivery, and monitoring
    - Payload standardization and optimization
    - Updating technologies to simplify data delivery
  - Expanding use of **metadata**
    - Can we break out of the “event processing loop model” ?
- Final nails in a few coffins imminent
  - SRM
  - AFS
  - ... enter your candidates here ...