iris hep

**Institute for Research and Innovation in Software for High Energy Physics**

# DOMA Team

Brian Bockelman

# Why DOMA?

- What's the point of a DOMA R&D area?
    - We believe data management is one of the key cost drivers for the HL-LHC and our group will make key contributions to reducing this cost. **Advances in technology do not compensate** for the increase in data volumes and complexity.
        - To reduce cost, we must **reduce the data resident on disk**: reduce replicas, reduce size of the dataset, reduce the percentage of data kept on disk (vs tape).
        - We must also **reduce the barriers for user data access**. Human costs, although more difficult to quantify, loom large for the HL-LHC.
- Some of these issues are social and some are technical: we cannot tell the HL-LHC what to do, but rather work together to provide potential solution mechanisms.

# Activities

- DOMA is organized into three areas:
  - **Modeling**: Better understand LHC data usage patterns as a precursor to HL- LHC planning.
  - **Data Organization**: Improve how HEP data is serialized and stored.
  - **Data Access**: Develop capabilities to deliver filtered and transformed event streams to users and analysis systems.
  - **Data Management**: Improve and deploy distributed storage infrastructure spanning multiple physical sites. Improve inter-site transfer protocols and authorization.

# The DOMA Team

## Morgridge Institute for Research

Brian Bockelman

## University of Chicago

Rob Gardner

Andrew Chien

(UChicago has an additional accepted offer starting later this month…)

## University of Illinois Urbana-Champaign

Ben Galwesky

Mark Neubauer

## Princeton

Jim Pivarski

## UC San Diego

Frank Würthwein

Edgar Fajardo

Igor Sfiligoi

## UC Santa Cruz

Carlos Maltzahn

Jeff LeFevre

Aaron Chu

**Note: This just shows those directly supported by IRIS-HEP specifically for DOMA: several others have been contributing and participating!**

# Organizing “Road Map”

(Carefully avoided the words “architecture diagram”)

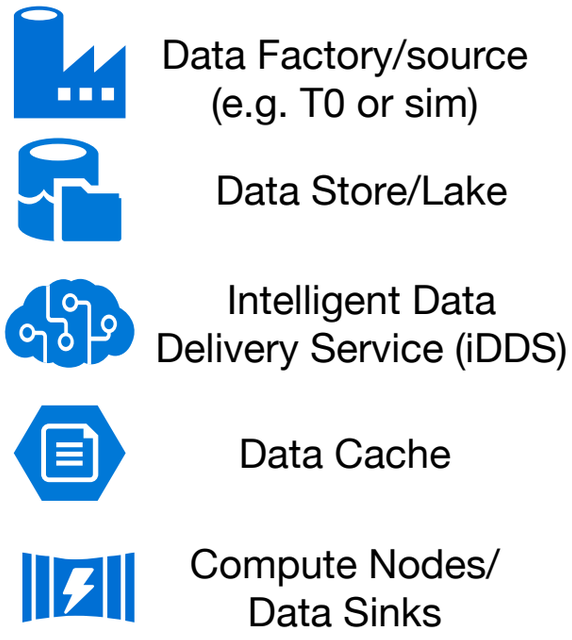


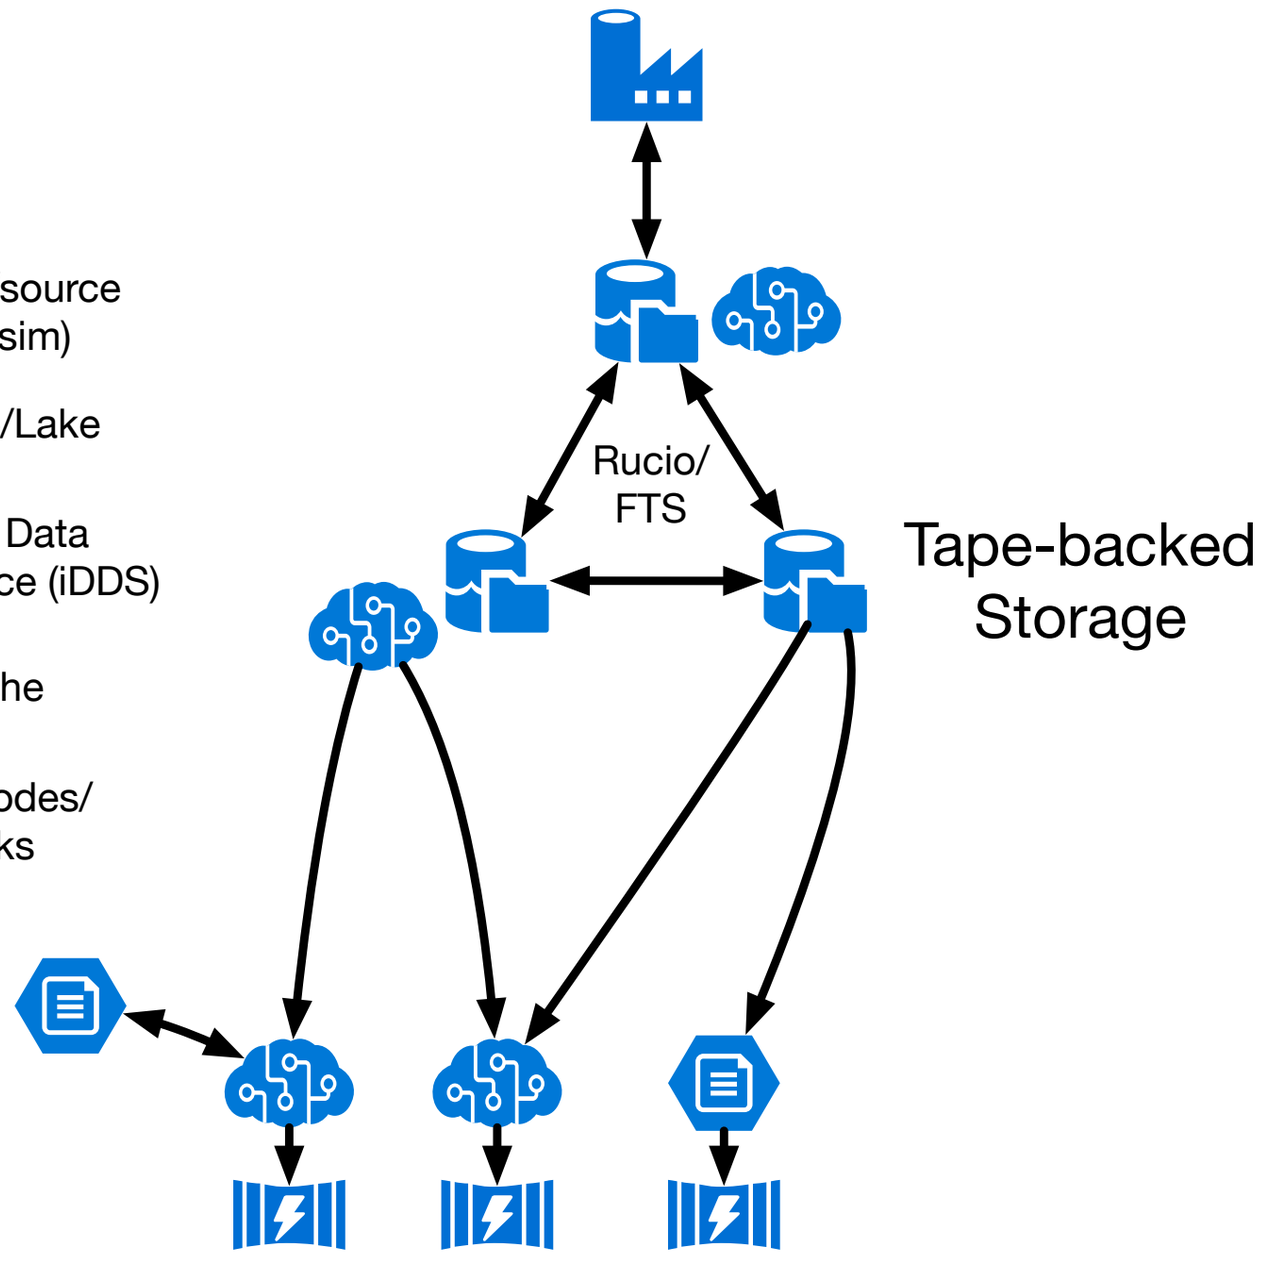


**Notes:**

- We treat computation and data storage as more separate in this model
- There will always be multiple independent data stores that need to be managed.
- Some data stores may look like today’s storage elements (file-based) while others may have an internally optimized format.

# Ongoing Projects

- The DOMA team has been very active across wider community projects such as the WLCG:
    - WLCG DOMA: The access group is co-lead by Würthwein; the third-party-copy group is co-lead by Bockelman.
    - WLCG AAI (Authentication and Authorization Infrastructure): Participation from Bockelman - AAI models have a significant impact on the DOMA area.
- Working to grow alternates to GridFTP protocol for interoperable data transfer.
- Studying alternative data management models for storing HEP data.

# Intelligent Data Delivery Service (iDDS)

- We think of a **cache** as a service that can deliver specific bytes to the user. Responsible for finding the right data (either locally or remotely) and sending it.

- We see a **data delivery service** as something that can deliver the requested data - possibly transforming or pre-conditioning things as needed.

- A cache responds to “deliver bytes 1024 - 2048 of file /foo/bar/baz.root”.
    - A data delivery service responds to “deliver the electrons and muons from events 1 - 10,000 of dataset foo_baz_AOD”.

- The IDDS will be responsible for handling user queries, determining the appropriate data source, transforming the data (potentially working with another iDDS at the source side), deciding whether to store it locally, and finally delivering to the user.

- See one-page “use case paper” attached to the Indico agenda.

**Goal: Setup one-on-one meetings with USATLAS and USCMS to go over these use cases, in preparation for HOW2019**

# Take-Home Message

- Over the last month, the DOMA area has gotten close to filling the open positions.

- IRIS-HEP DOMA has provided leadership and technical to impactful community projects such as WLCG DOMA.

- Gearing up for the IDDS project — looking to schedule meetings with USATLAS and USCMS to gather requirements.

# Backup Slides

# Top-level WBS Items

- **WBS3** (DOMA): Develop systems capable of managing, organizing and accessing Exabyte-scale datasets across national and international cyberinfrastructures that integrates network, storage, and compute investments at all scales - from small colleges to national centers.
  - **W3.1** (Modeling): Better understand LHC data usage patterns as a precursor to HL-LHC planning.
  - **W3.2** (Organization): Improve how HEP data is serialized and stored.
  - **W3.3** (Access): Develop capabilities to deliver filtered and transformed event streams to users and analysis systems.
  - **W3.4** (Management): Improve and deploy distributed storage infrastructure spanning multiple physical sites. Improve inter-site transfer protocols and authorization.

# Deliverables and Milestones

| Type | Time | Description | WBS task | |
|---|---|---|---|---|
| D | 0-3 | Start a process for regular exchanges between ATLAS and CMS on understanding use data from both experiments. | 3.1 | ✓ |
| M | 0-3 | Begin technical work on caches and interoperable data transfer mechanisms | 3.3 | ✓ |
| D | 3-6 | Organize topical meeting on data delivery through CDNs within LHC and comparisons with other users. | 3.1 | IN-PROGRESS |
| D | 3-6 | Initial web presence explaining the DOMA area activities. | 3.* | TODO |
| D | 3-6 | Organize discussions for “data lake” and “IDDS” use cases, stakeholders, and development plans. | 3.3, 3.4 | IN-PROGRESS |
| M | 3-6 | First draft of IDDS design | 3.4 | Post-HOW2019 |
| D | 3-6 | Prepare internal document outlining opportunities for minimizing event data storage sizes. | 3.1, 3.2 | Summer 2019 |

# Deliverables and Milestones

| Type | Time | Description | WBS task |
|---|---|---|---|
| D | 6-12 | Deliver report on LHC data access patterns, usage and intelligent caching approaches for the HL-LHC | 3.1 |
| M | 6-12 | First technical demo of IDDS. | 3.4 |
| D | 6-12 | Report on cache usage on the WLCG and potential use cases & deployment scenarios for the US LHC facilities. | 3.3 |
| D | 9-15 | Prototype of HEP event data on object storage using alternate organization approaches. | 3.2 |
| M | 12-18 | With OSG, help transition 30% of data transfers at one USLHC site to use a non-Globus Toolkit implementation. | 3.4, 7.3 |
| D | 12-18 | Report on feasibility of hardware accelerators for event data. | 3.4 |
| M | 12-18 | IDDS prototype that can automatically transform / optimize data formats. | 3.2, 3.4 |
| M | 18-24 | With SSL, scale test a data lake prototype across multiple sites. | 3.3 |
| D | 18-24 | First version of IDDS at production site. | 3.4 |