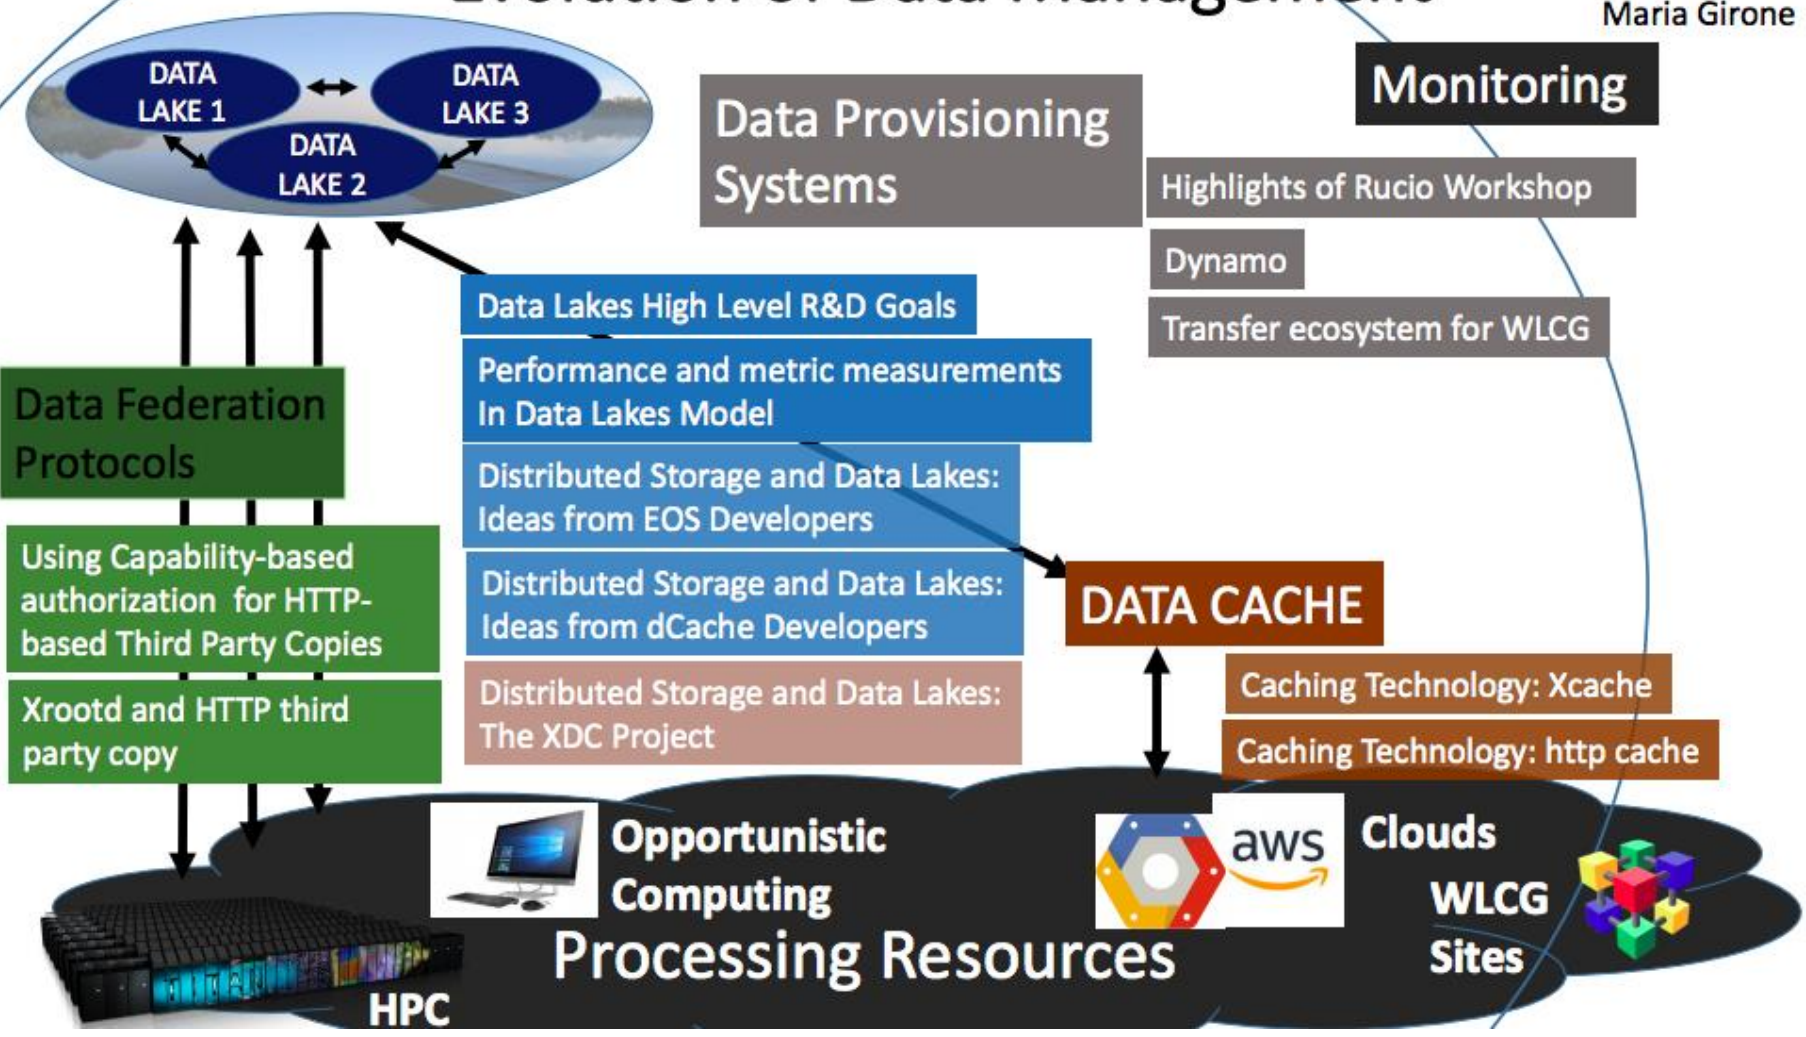


# Data Management plenary and parallel sessions: summary and outlook

Maria Girone, Simone Campana  
CERN

# Evolution of Data Management

Simone Campana  
Maria Girone



What common projects are identified for this area that will advance?

- How will these be followed up?
- Are the timelines well defined?
- Are the right stakeholders already involved? Should other experiments or organisations be contacted?
- Are any resource needs covered?

Are there projects that are not currently advancing?

- What would allow them to do so?

Links with communities outside HEP

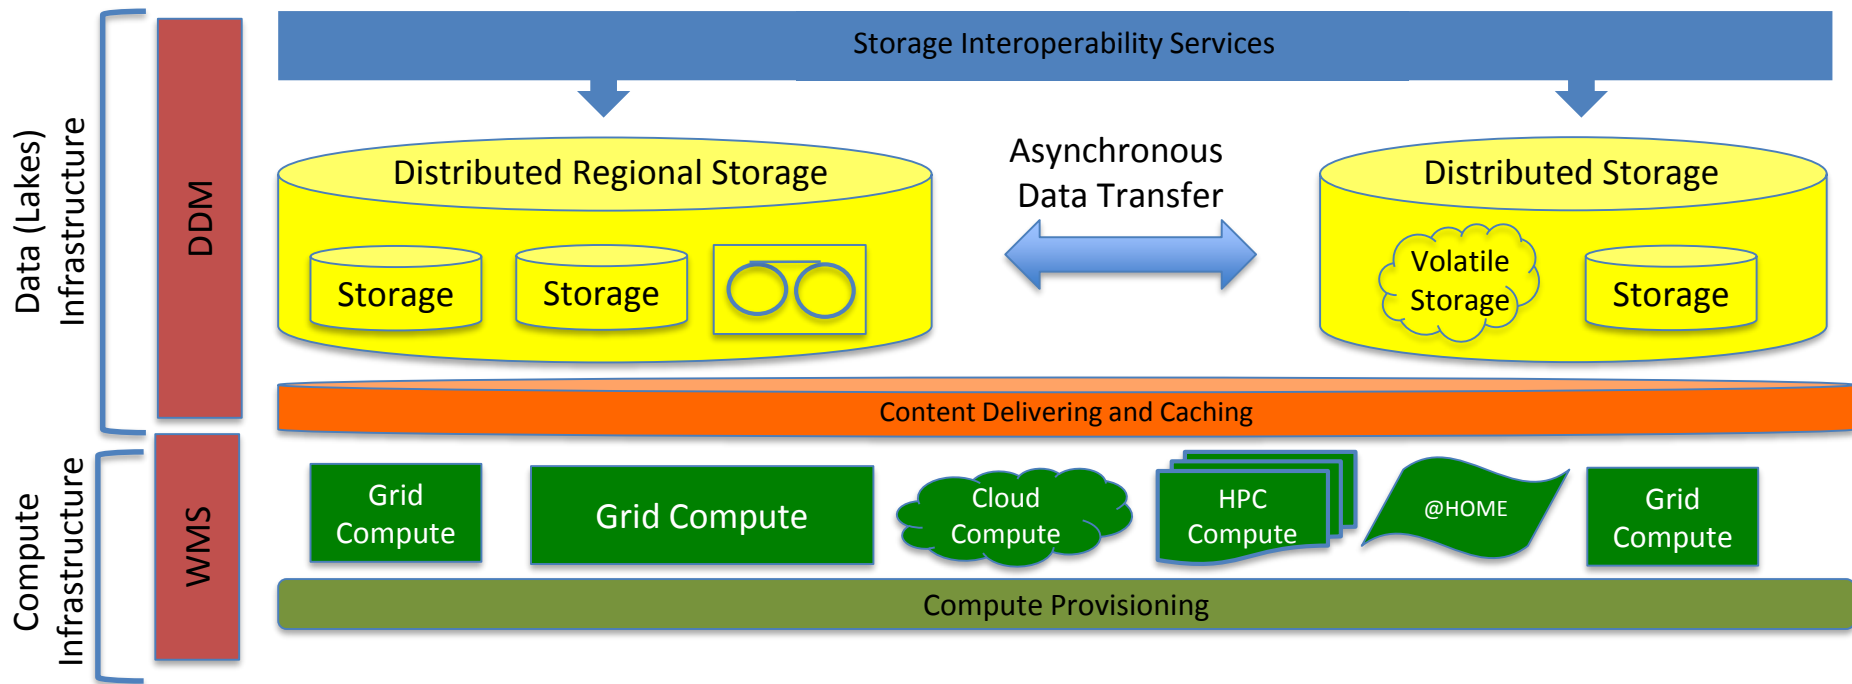
- Are any established? Are any links needed that we don't have?

Are we making best use of community expertise in this area?

- How best do we advance collaborations that will be of benefit more widely and improve our software and maximise the effectiveness of effort

Best ways to work collaboratively from here are...?

# Evolution of Data Organization, Management and Access



# DOMA Q&A

- **Data Management and Provisioning Systems**
  - Pursue commonality in experiment data management systems. Rucio being a very promising candidate as common solution between ATLAS and CMS
  - Implement common functionalities in a common middleware layer (see FTS for example)
- **Geographically distributed storage services: evolve existing technologies (dCache, EOS, ..) to “advanced data federations”**
  - Performance, impact of latency/bandwidth, QoS, notification mechanism, deployment model, operation model
- **Content Delivery and Caching technologies**
  - Protocols, granularity, methodology
- **Data Federation Protocols**
  - From short/medium term gridFTP replacement to longer term token based system

# Trends from Workload Management session



## Thoughts on Storage Budget

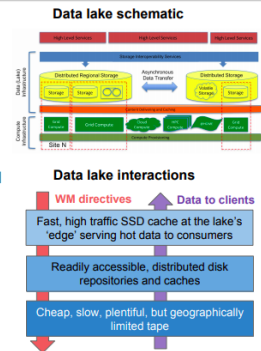


- The tape budget is something we can not shrink significantly
    - “virtual data” makes no sense given the high CPU cost of reconstruction.
  - Disk is ~x5 more expensive than tape.
- => The best bet for reducing storage budget is to rely more on tape and less on disk.**

## Data lakes and workload management



- Our sites are linked with (ever higher) high-bandwidth networking
  - We can expect **~100x bandwidth growth** by HL-LHC
- **Data lakes:** integrated consolidation of distributed storage (and compute) facilities, leveraging high-bandwidth networks
- Data lake encompasses facilities with several levels of storage
  - **Tape**, at a relatively limited number of sites
  - **Standard disk**, at large storage repositories and smaller caches
  - Fast SSD **'edge cache'** for the hottest data
  - Should be able to **place data optimally** based on (dynamic) need
- Workload management knows the hot popular data in use
  - Use that knowledge to drive preparing data in the lake, asynchronously to the processing, e.g.
    - tape staging in a **carousel workflow**
    - placing hot data in SSD cache **'close'** to available CPU
    - **transforming/marshaling data** optimally for client delivery
  - Requires APIs supporting WM directives
- **Instead of 1.8 replicas on disk today, WM + data lake manages dynamic availability of actively used data with replica count <<1**



BROOKHAVEN

T. Wenaus 2018-03-27



End to End content delivery  
Hot/Warm/Cold storage



# Next steps

- Create a Data Organization Management Access evolution project:
  - keep track of developments and advancements in all DOMA areas
  - provide a forum to discuss ideas
  - foster interoperability of solutions
  - an umbrella for stakeholders, national initiatives, EU projects, already existing working groups
- Stakeholders: experiments, middleware developers and storage providers, facilities