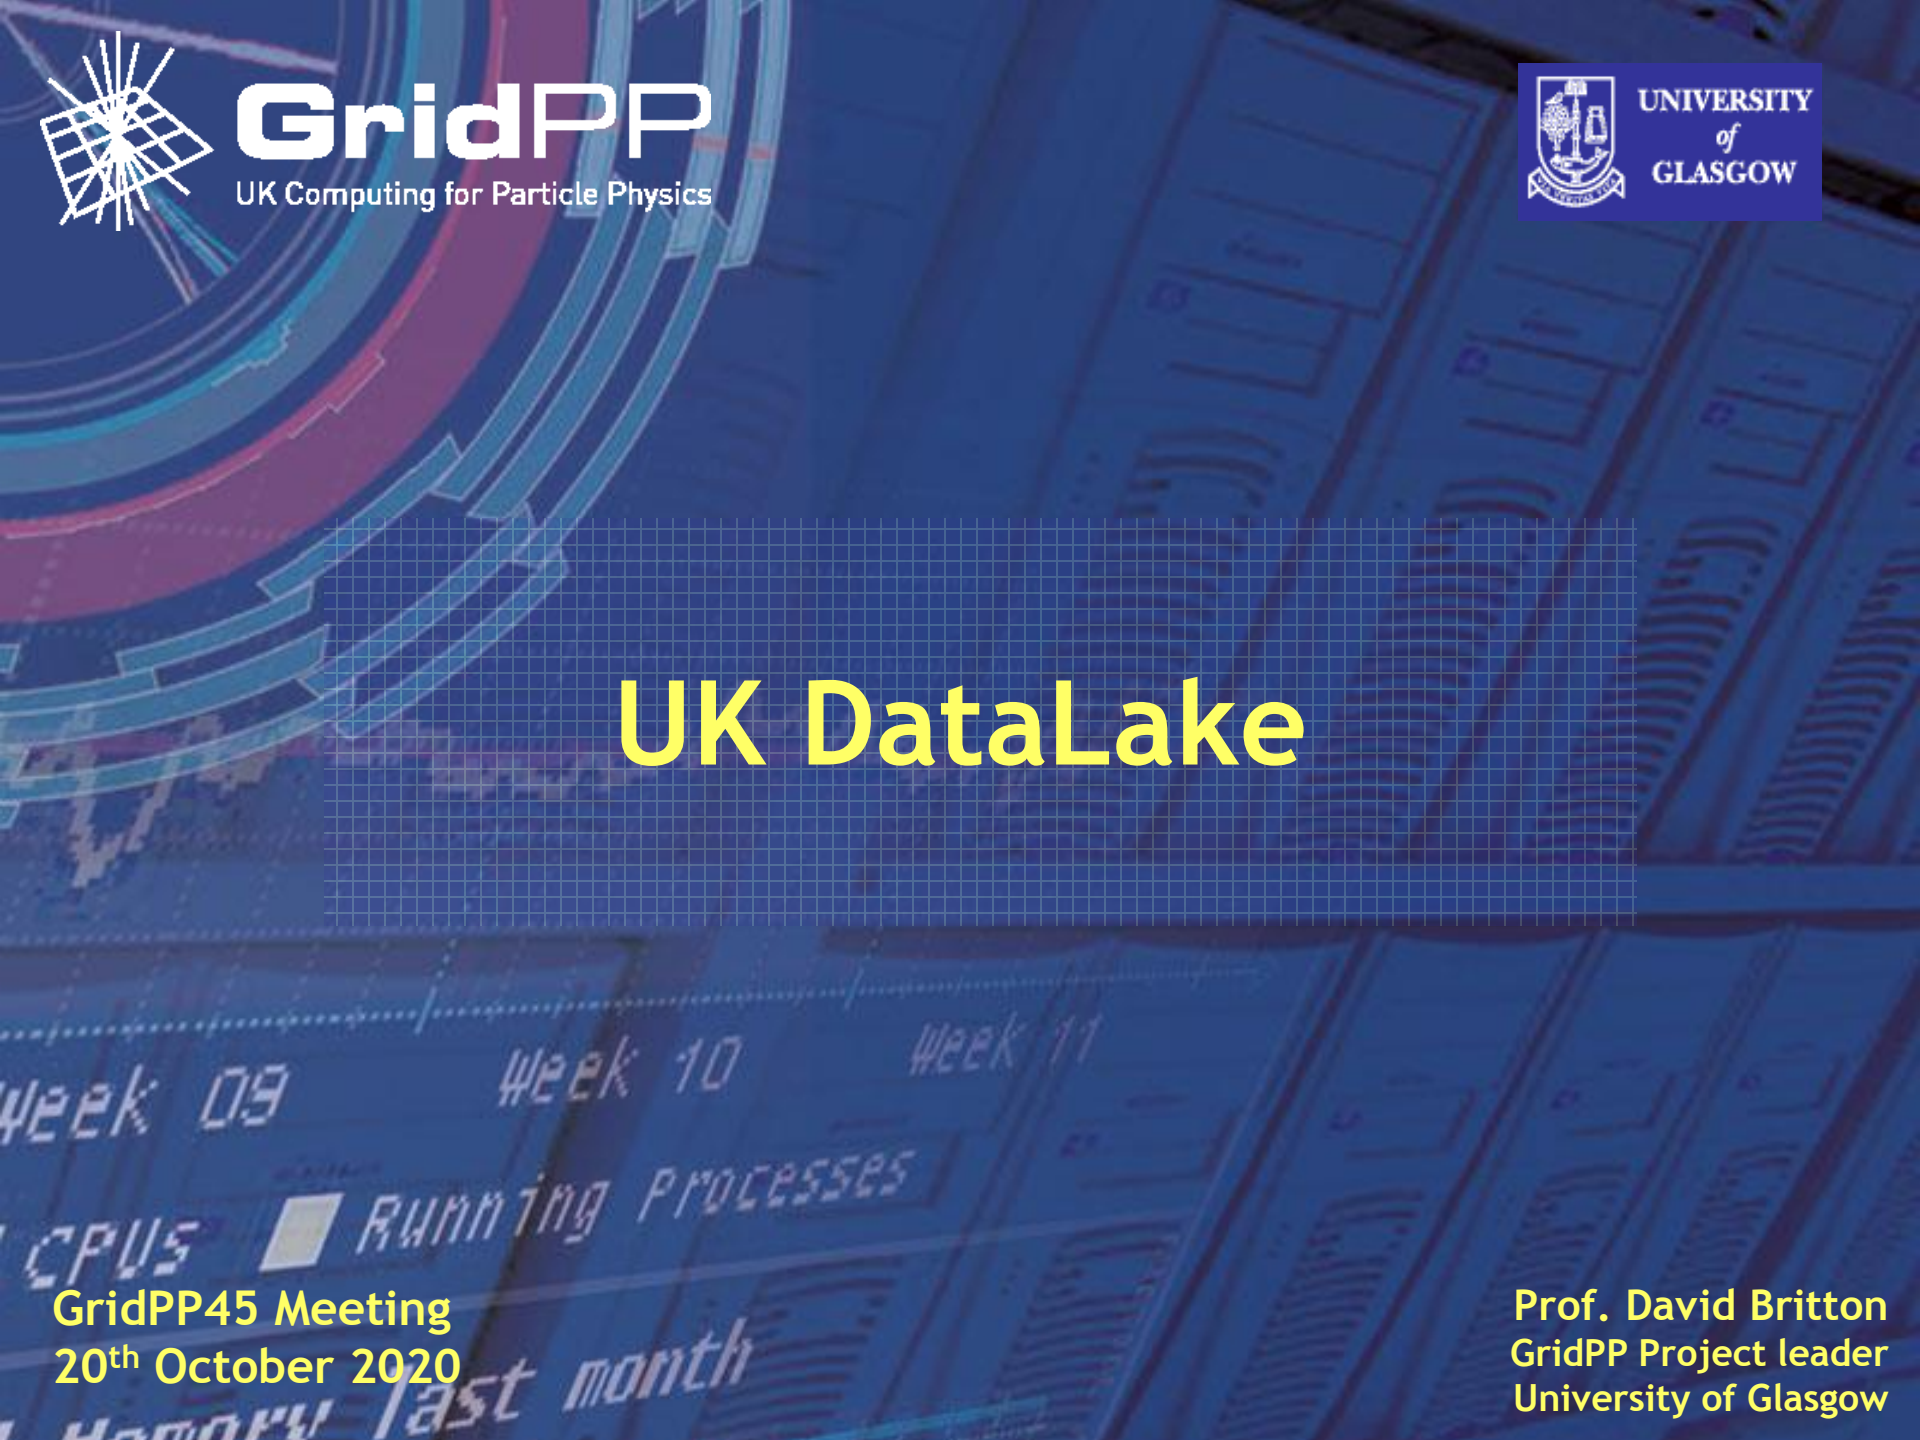
GridPP
UK Computing for Particle Physics
UNIVERSITY of GLASGOW
UK DataLake
Week 09
Week 10
Week 11
CPUs
Running Processes
last month
GridPP45 Meeting
20th October 2020
Prof. David Britton
GridPP Project leader
University of Glasgow

Data Lakes: Why?

- Cost and volume of storage required for HL-LHC.
  - Cannot be delivered within flat-cash budgets
- WLCG storage currently does not expose any notion QoS
  - Experiments end up replicating more data than they really need
- WLCG storage is currently fragmented in more than one hundred endpoints of very different size and with very different level of support.
  - Fewer, larger endpoints could offer the same/better service at a lower (overall) operational cost.
- Vision is to reduce complexity (therefore cost) for the experiments by reducing end-points; providing QoS classes; enabling policy-driven, automated data-movement and replication.

UK Storage Sites

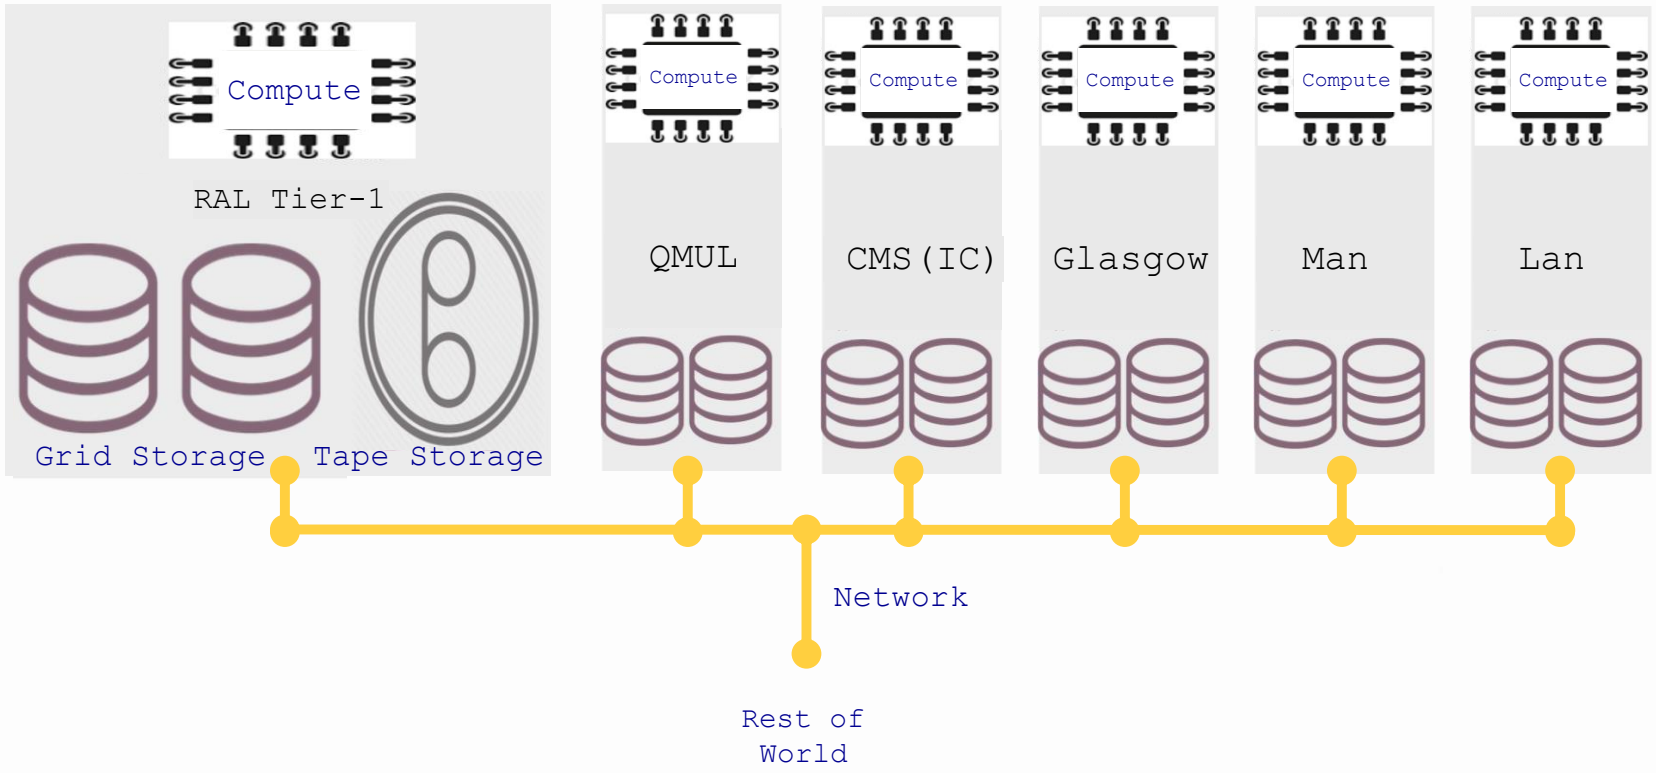

Federating Data

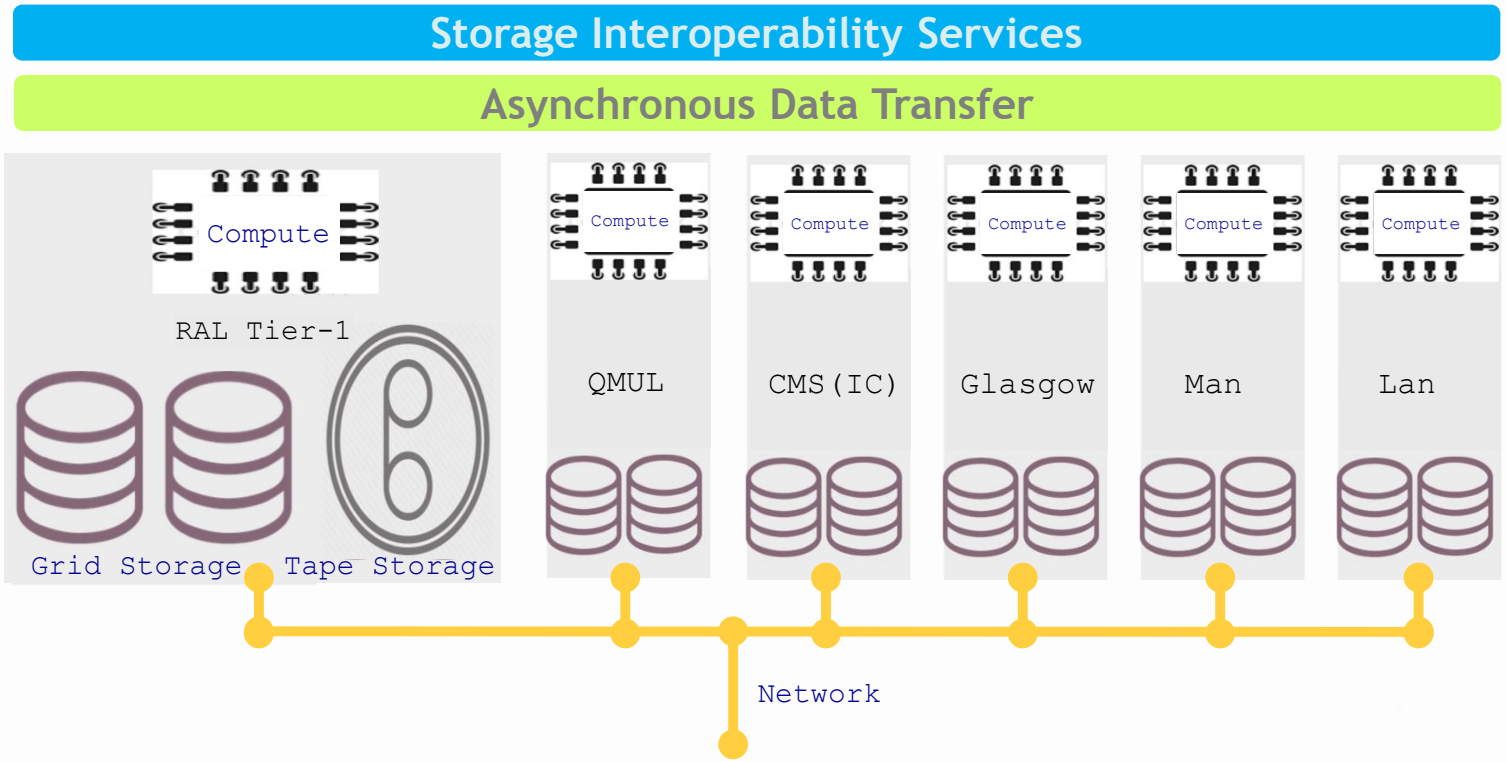

Storage endpoints of perhaps different QOS; data-exchange managed through WAN transfer protocols and orchestrated by storage interoperability and high level services.

Global DM System federates and defines Data Lake including Quotas, ACLs, Replication Policy…. enabling different QoS, Data lifecycles, distributed redundancy, etc.

Data Lake (core)

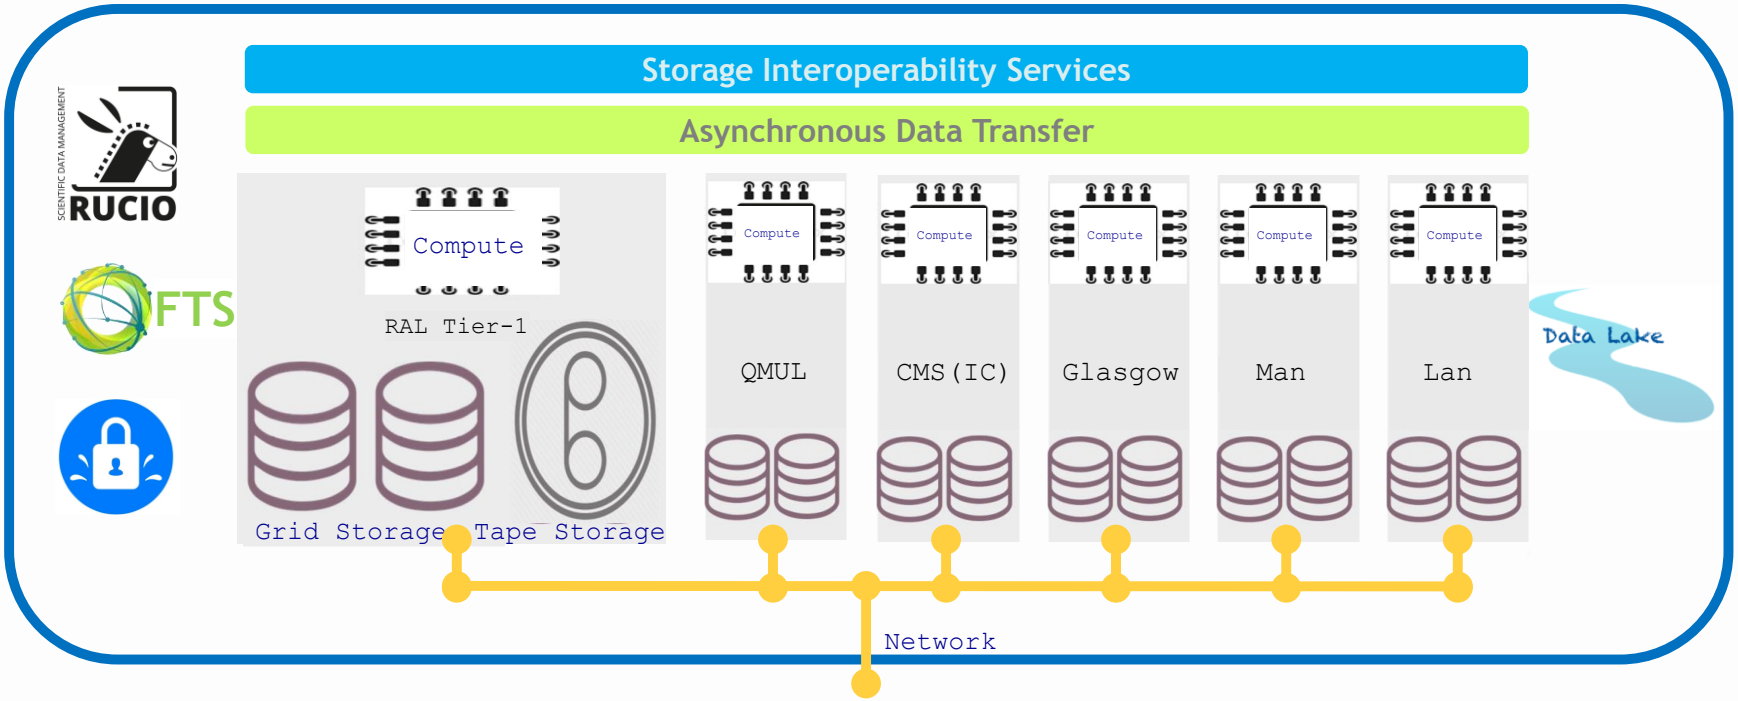

Increasingly, the experiments define data Policy and Requirements (QoS) but the infrastructure implements (it's like moving to a declarative programming paradigm: Say what you want to do, not how you want to do it.)

Data Lake (core)

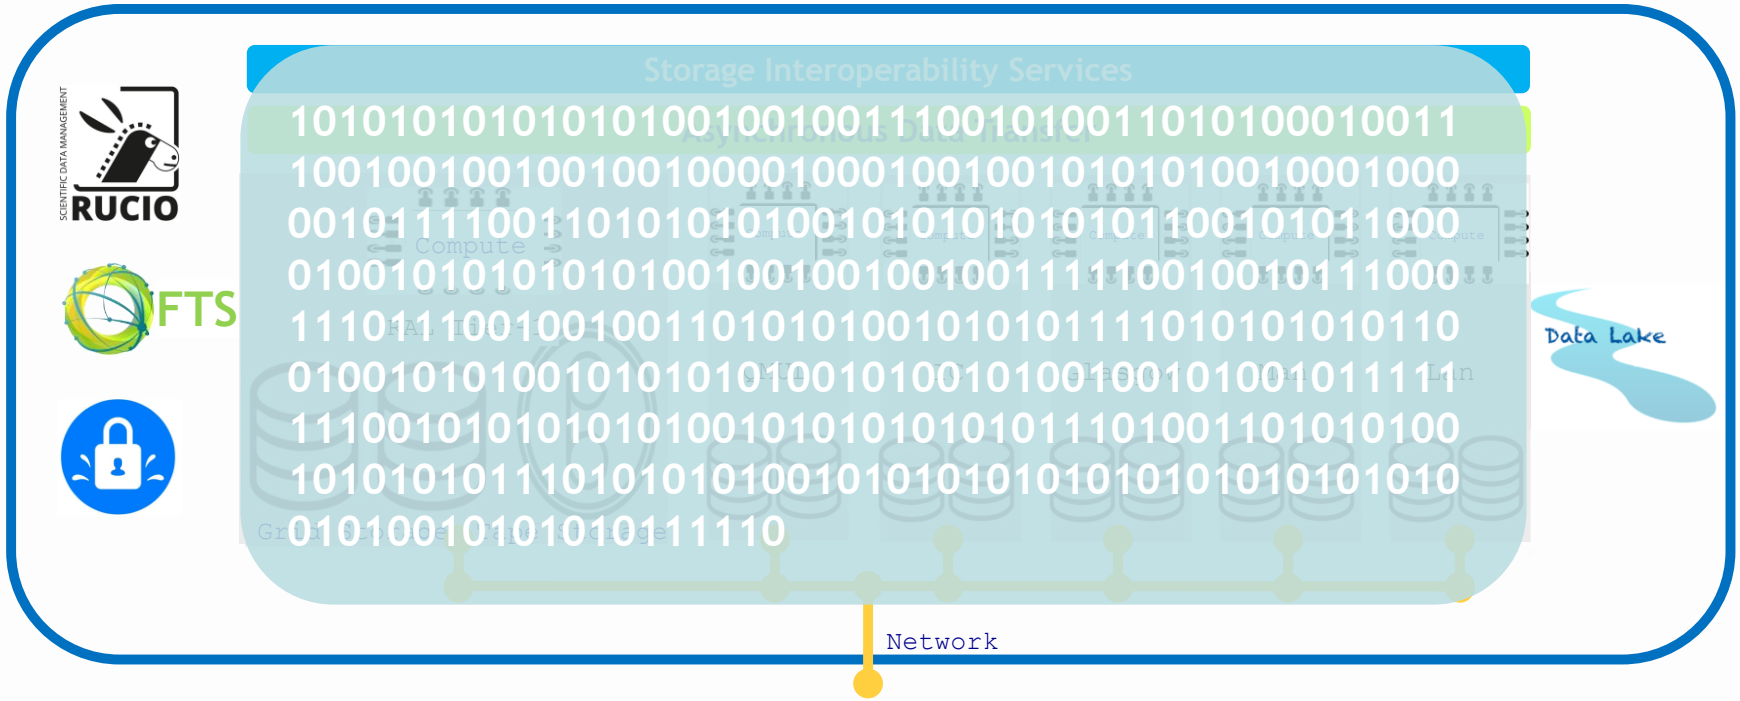

Increasingly, the experiments define data Policy and Requirements (QoS) but the infrastructure implements (it's like moving to a declarative programming paradigm: Say what you want to do, not how you want to do it.)

Data Lake (extending)

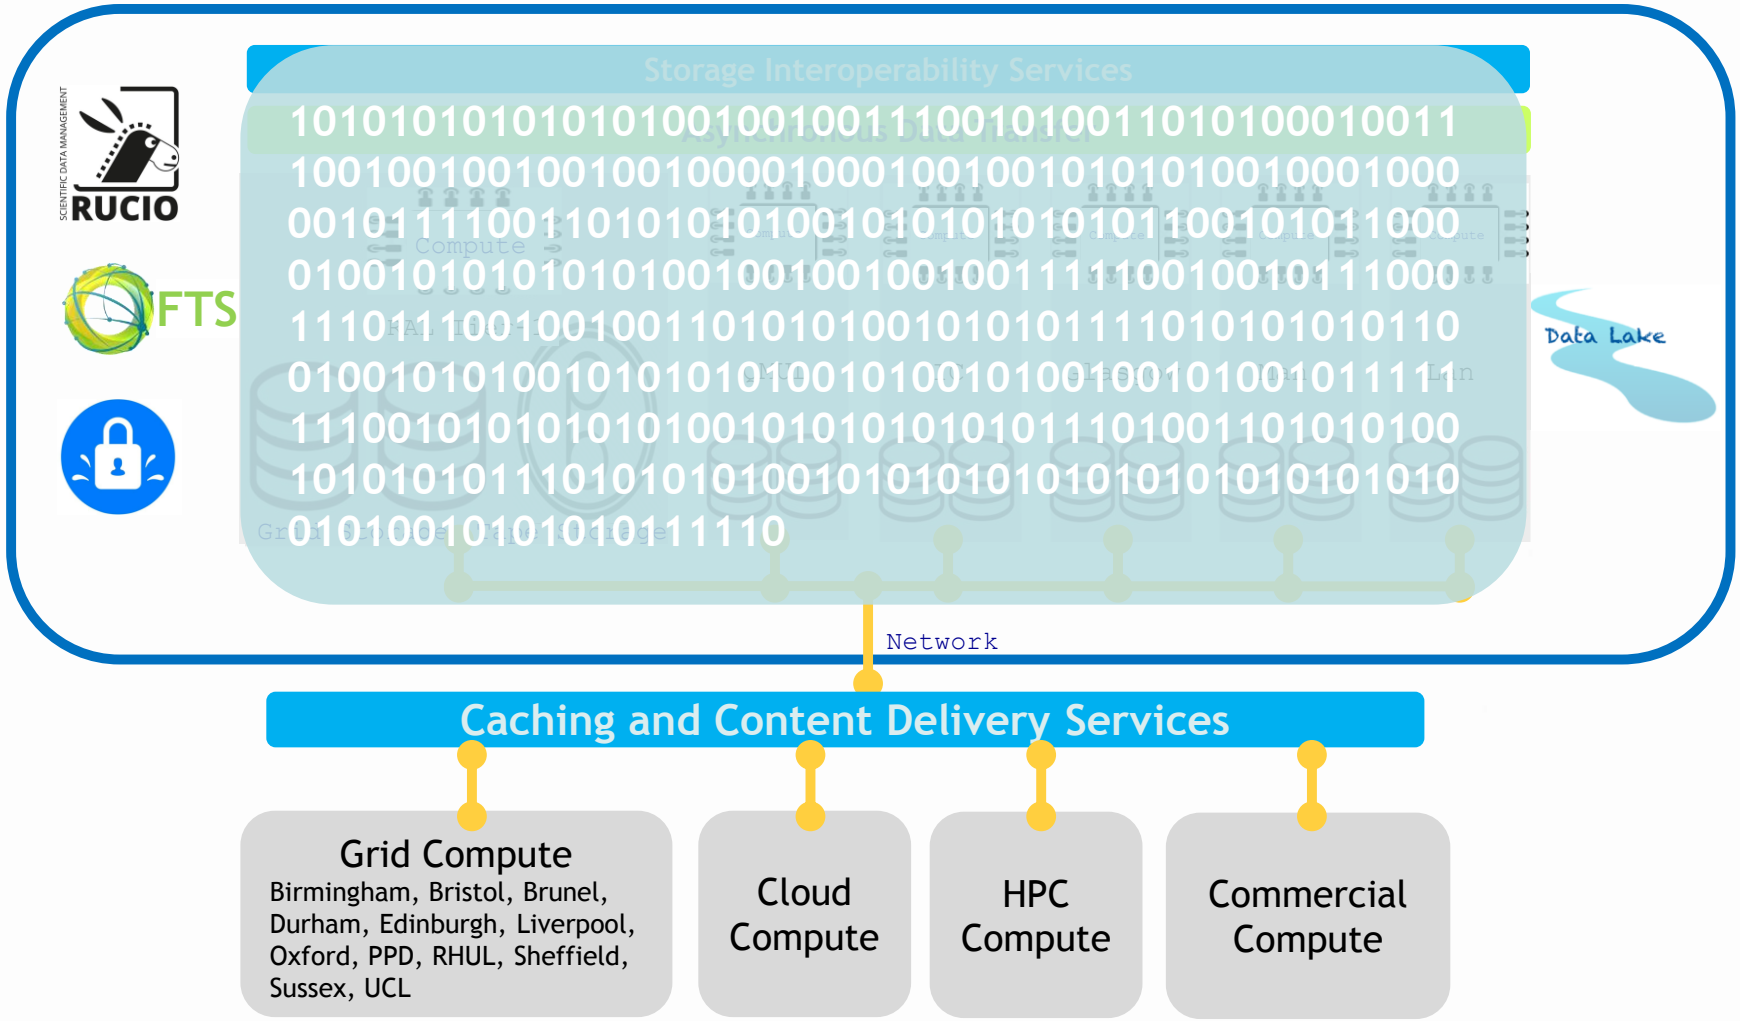

So what do sites look like?

- Do core sites look much difference from now?
  - Large amounts of storage in a scalable/supported system.
  - Excellent network connectivity, growing as needed.
  - Rucio and other data services.
  - QoS implementation

So what do sites look like?

- Do we yet know how to define the non-core sites that plug into the datalake core?
  - Caching and/or remote access?
  - Network requirements or limitations?
  - Attached to Tier-1 or a particular Tier-2, or just “the lake”?
  - What is the future of, or beyond, DPM?
  - Staff requirements – consolidation of storage is about allowing overall limited effort to overall deliver more.
  - How is CPU presented?
    - Recommended single, supported, CE technology?
    - Defined role for VAC?
    - SLATE and other initiatives?
  - Balance with local site interests, obligations, constraints and responsibilities?

Discussion Points

- Are we developing expertise in all the technologies we need (Rucio, Caching, IAM, FTS, TPC…. we should list…)
- Metrics for success:
  - Throughput for different workloads
  - Reduced cost of storage by automated use of QoS classes
  - Reduced effort for experiments
    - Far fewer storage end-points
    - Automation of some data-management processes
  - Deliver (much) more resource with the same overall level of effort?
- What milestones would we envisage?