Data Access in HEP

Focus areas for data

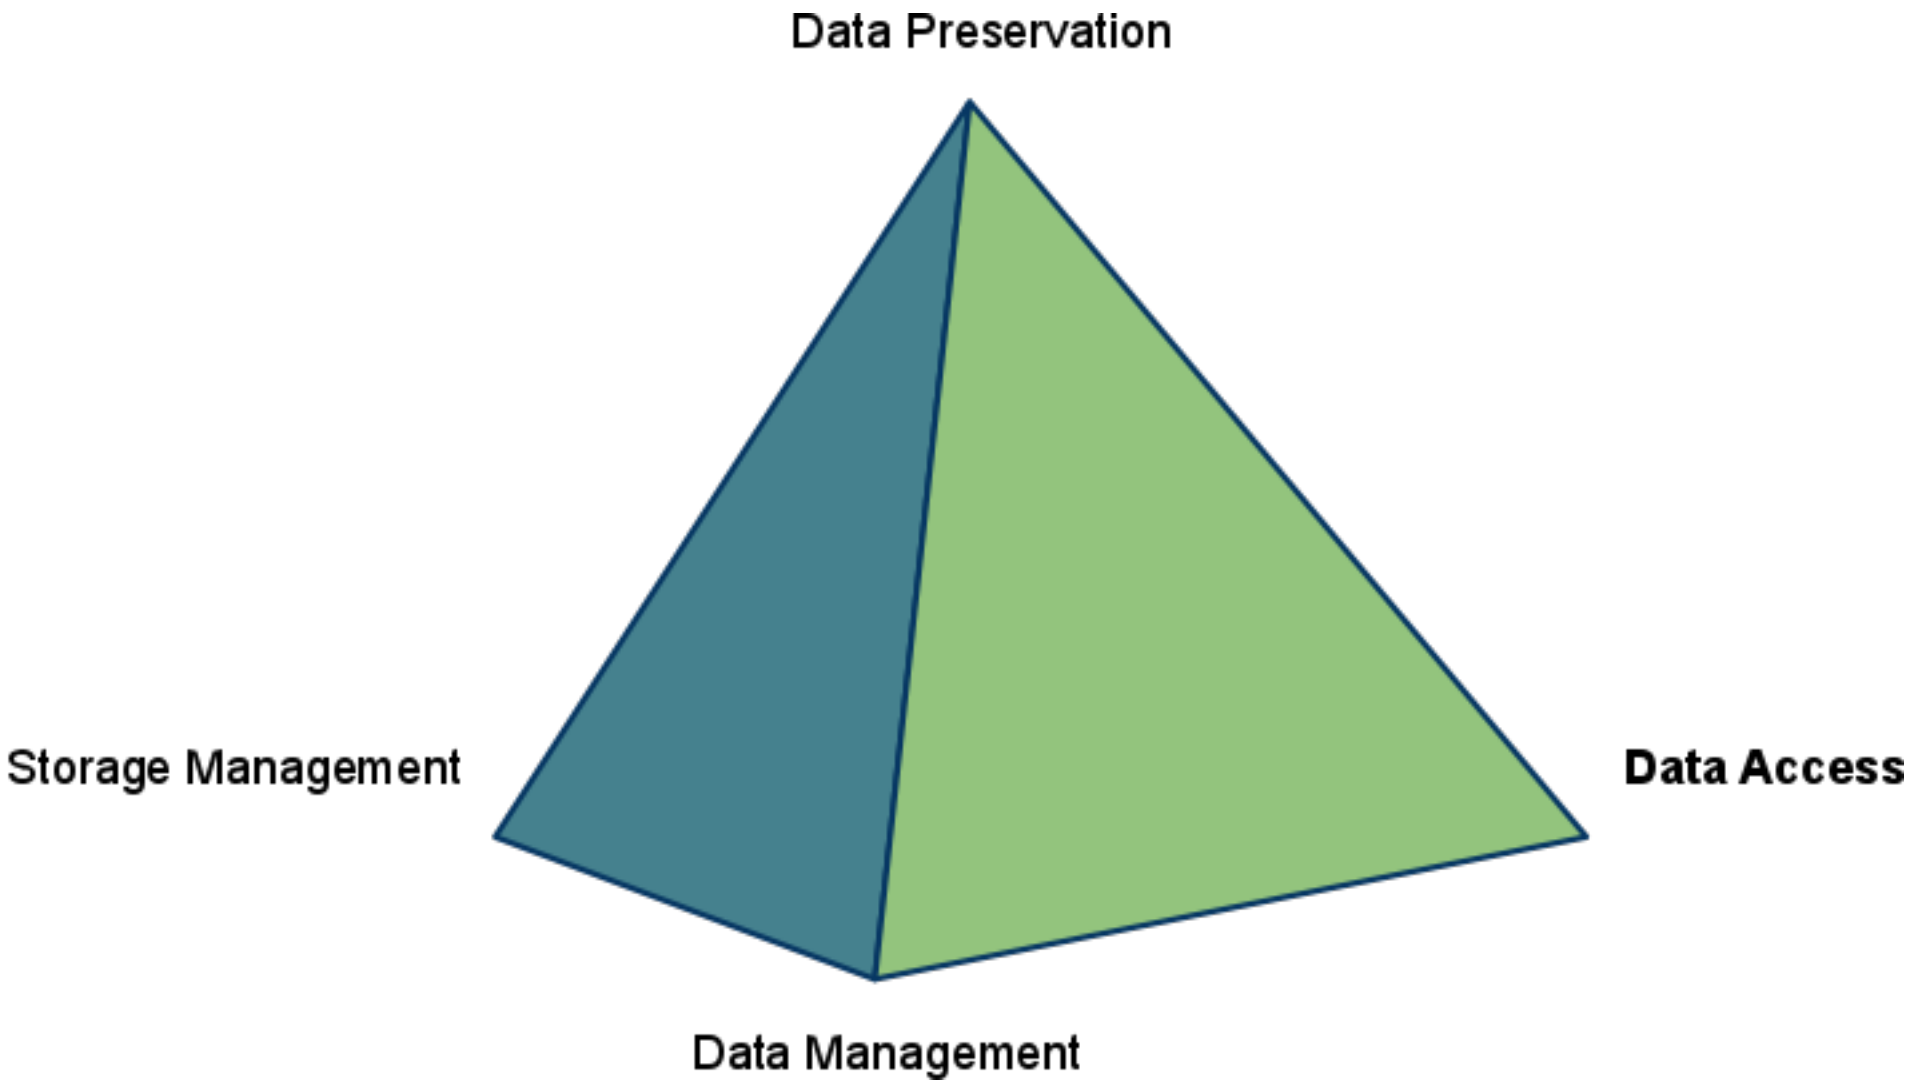

Quick Picture (as seen yesterday)

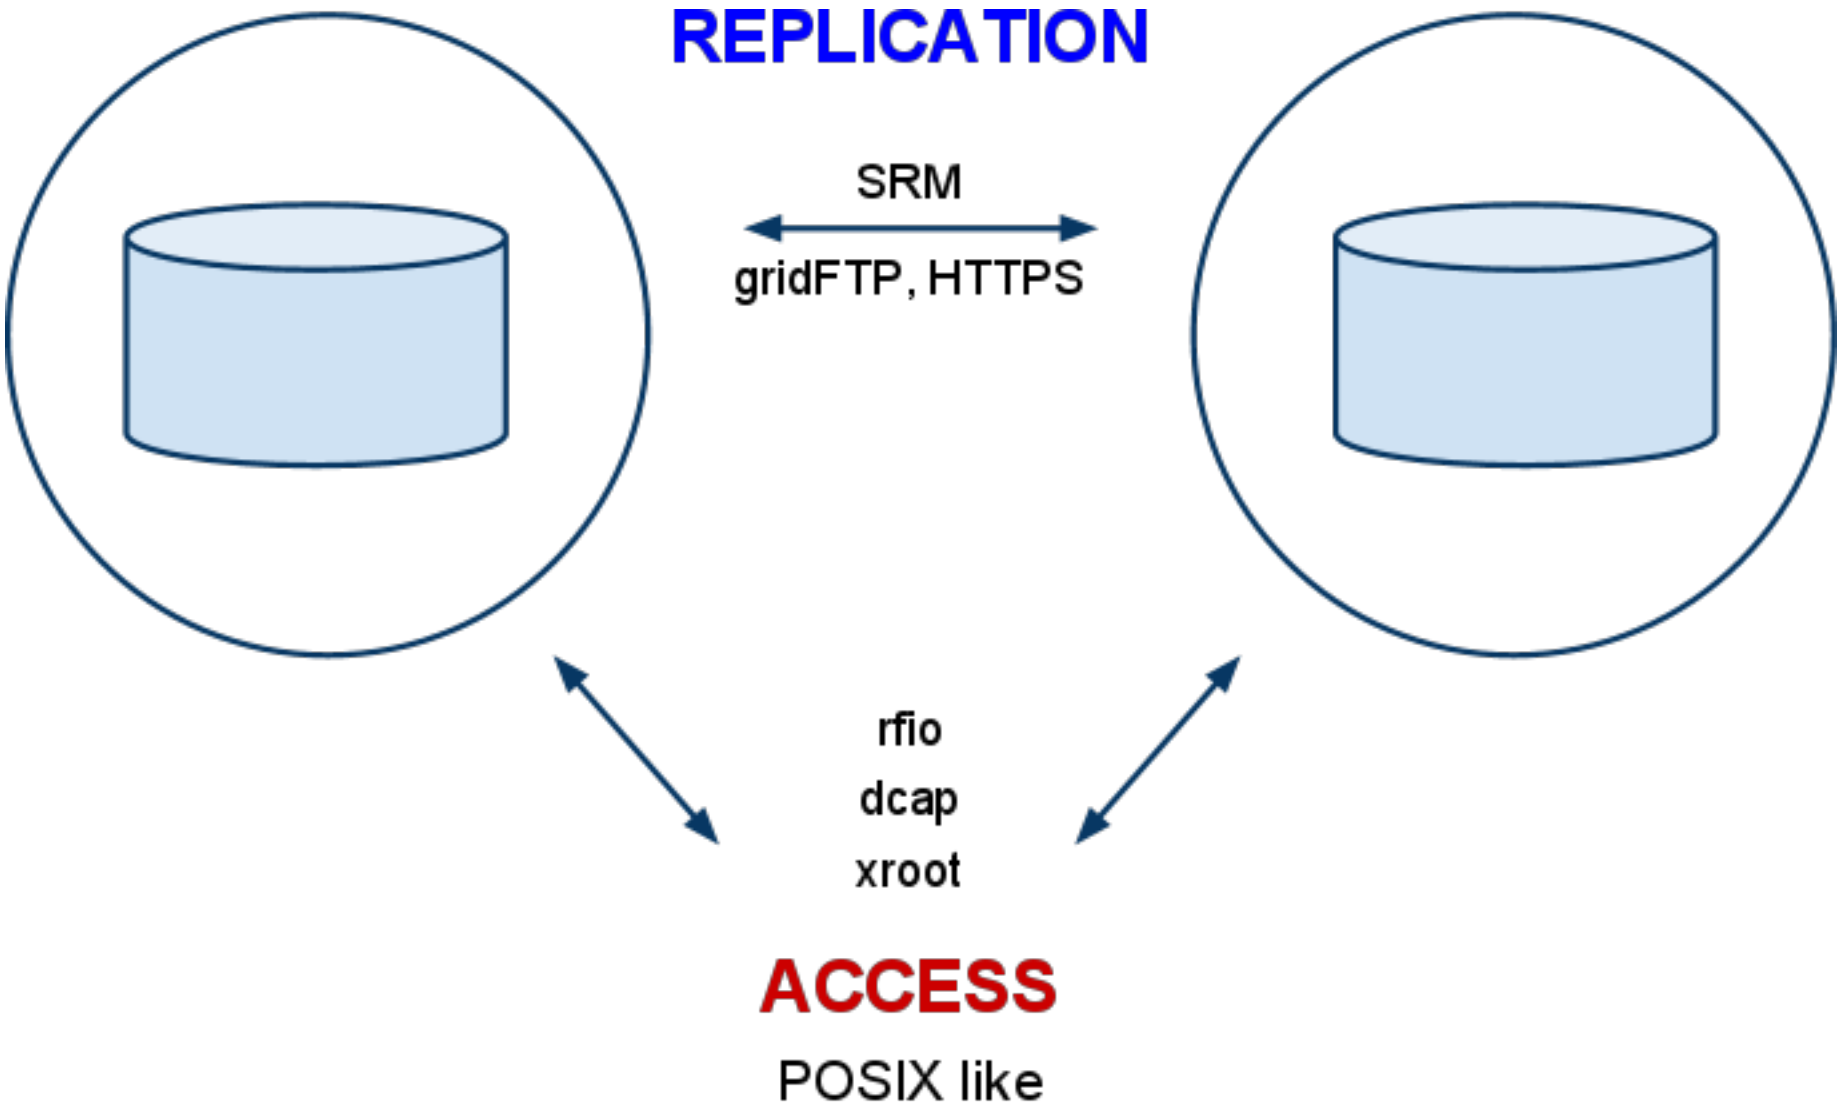

What our I/O protocols provide :-)

- POSIX like access
  - Expected by our analysis libraries
- Strong security
  - Currently PKI (X509) & kerberos (popular internally)
- Scalability
  - Thousands of concurrent clients
  - Accessing local systems, but also over the WAN
  - Distributed data access (separate metadata and I/O requests)
- Performance optimizations
  - Block caching, vector reads

No standard protocol offering all of this... until recently

What our I/O protocols impose :-(

- A protocol zoo
  - We have an heterogeneous environment
  - Every vendor proposes his own
  - Some efforts on providing abstraction libraries
- Pain on our users
  - Systems and experts need this knowledge
  - Complexity and load on operators increases significantly
- User Interface Nodes (entry points)
  - Hard to manage
    - version incompatibilities, library conflicts, ...

A standard for POSIX access... pNFS

- NFS4 gave us
  - Strong security (via GSSAPI)
  - Bulk requests (performance, also over WAN)
- NFS4.1 (pNFS) gives us
  - Split between metadata and I/O requests
  - Support for multiple access "layouts": file, block, object
  - Scalable access to distributed data
    - Place for optimization, benefit from tiered storage
- And above all they give us
  - An industry standard
  - Built-in clients in different operating systems (Linux, Solaris, Windows, ...)
    - And all the optimizations already there for years
  - End of vendor lock-in (?)

Where we will end up

- A simpler, easier system to use and manage
  - Work in progress in some of our systems
- But that won't been enough
  - Our libraries (analysis and other remote access) need to improve their failure handling
    - Code for failure, as praised earlier today
    - True for every community
- And only then...