

Creating Federated Data Stores For The LHC

jeudi 13 septembre 2012 - vendredi 14 septembre 2012

CC-IN2P3

Programme Scientifique

<center>Motivation</center>

Currently, three LHC experiments have deployed or are ready to deploy a federated model for experimental data access (i.e., ALICE, ATLAS, and CMS). While each experiment has taken a similar approach (i.e. xrootd), the implementations are sufficiently different to warrant a collaborative discussion of the specific needs that spurred a particular direction.

<center>Purpose</center>

Through a collaborative dialogue, this meeting hopes to expedite adoption and support of federated data access models in LHC by leveraging each other's experience.

<center>Points of discussion</center>

Federated Data access models:

Mechanisms:

Copy on demand.

Real time access.

Partial caching.

Pros and cons to each mechanism:

Requirements to make each model work well.

Firewalls.

Federated Name Space:

Providing a uniform global logical name space:

Challenges and solutions.

Federated Monitoring:

Metrics:

Necessary information.

Desired information.

Level of detail.

Visualization:

Architectures.

Incorporating local information:

Mechanisms.

Useful metrics.

Federated Security:

Requirements for different access modes:

Client/Server vs Server/Server access.

Emerging mechanisms:

DOE's InCommon Federated Trust Model, Shibboleth, etc.

Federated Support Structure.

<center>Meeting Format</center>

We envision a 20-30 minute technical and practical presentation in each topic area from each group (experiment or computing group) developing or deploying a federated data access model and related services. The presentation should include a description of what worked and did not work and why; as well as challenges yet to be addressed.