

WG on Storage Federations

07 Aug 2012
Introduction

Fabrizio Furano
furano@cern.ch

- Let's start trying to have a look at what the experiments are doing now
- Interesting things to understand:
 - How the federation is accomplished, what bits make it match the definition* ?
 - Are there many or just one (per experiment)
 - How is the federation used?
 - e.g. seamless direct access, tape-like file healing, caches...?
 - What is the final goal of the exercise?
 - What is its current deployment status?

** A collection of disparate storage resources managed by cooperating but independent administrative domains transparently accessible via a common namespace.*

- What about the involved workflows?
 - Things like “go to the SE, then ev. to the regional redirector, then ...”
- Where are they defined?
 - defaults ones in the xrootd client? wrapper system? instrumentation of the site?
- Are these
 - functional requirements (“we need this and this”)
 - ... or workarounds ? For what?

- What does a site have to do to:
 - join one of these federations?
 - maintain it, keep its corner healthy
 - report on the activity of its federated node(s)
- Does this federation involve instrumenting/customizing the setup at site-level?
 - e.g. custom scripts doing XXX then YYY (file healing... reporting... what?)
 - Who does it?
- Are all the sites configured/instrumented in the same way?
- Can the described approaches share components among experiments?
- What makes a site eager to install your stuff?

- Proposal

- Doodle poll for the next one, spanning 20-31th of Aug
- What about hearing about the point of view of ATLAS and LHCb?
- The subsequent one could be some sort of summary on the usages of a federation (slide 14 of the kickoff). Would you like to say something there?
- Please vote in the poll otherwise we'll never know what to do

Thank you

Questions?