

WG on Storage Federations

First meeting

Fabrizio Furano
furano@cern.ch

- The official motivation for the WG:
 - *Federated access to data is one of the main recommendations of the Data Management and Storage TEG*
 - *For the large scale implementation it is important that several technical details and policies are clarified, agreed and documented, so that implementation and setup can progress*

- The official mandate
 - *Clarify in detail which forms of federated data access need to be supported*
 - “fail over” for individual jobs
 - “repair mechanism” of placed data
 - who initiates the repair?
 - how is metadata handled ?
 - pro-active repair?
 - access control and trust mechanism between sites
 - use of catalogues
 - etc.
 - *Suggest a deployment and operations model for the additional components*
 - redirectors
 - xrootd plugins for non-native systems
 - experiment specific configurations/plugins
 - configuration of federations
 - discovery of federations
 - etc.

- Federating storage can be seen as giving seamless direct access to the globality
 - So far the XROOTD clustering through WAN is the solution that came up naturally
 - Others are possible, e.g. HTTP/DAV or the ALICE schema
- Then, this federation can be used for several purposes
 - user/job direct access, fail-over, healing, “staging from”, dynamic caches, etc...
- I also see an intense activity in refining workflows for xrootd federations
 - Things like “go to the SE, then ev. to the regional redirector, then ...”
 - Will be nice to see them consolidated
 - Will be interesting to understand which ones are functional requirements and which ones are induced by the tech framework in use

- Our interests:
 - Making the basic concept better known, make it become a more “normal” thing
 - Get one step closer to “consolidate” what a WLCG site has to do for his SE to join a federation
 - In order to “consolidate” we must first understand
 - Try to reduce the differences among VOs
- A way to start is to document the current status and the directions of the ongoing efforts
 - May be a milestone

- The WG is composed by:
 - Tech providers
 - One representative each
 - Experiments
 - One (two?) representative each
 - Sites
 - Two (three?) site managers, representing the requirements of the community, with its main flavors
- Common requirement is that all participants have a background on the topic

- ALICE
 - Costin Grigoras
- ATLAS
 - Douglas Benjamin
 - Rob Gardner
- CMS
 - Brian Bockelman
 - Giacinto Donvito
- LHCb
 - Philippe Charpentier

- DPM
 - Fabrizio Furano
- dCache
 - Paul Millar
- XROOTD
 - Andy Hanushevsky (ev. substituted by Wei Yang)
- STORM
 - Michele Dibenedetto
- EOS
 - Dirk Duellmann

- Need to satisfactorily represent their flavors without overcrowding
- Preliminary list:
 - Luca dell'Agnello (INFN-CNAF)
 - Andy Hanushevsky/Wei Yang (SLAC)
 - Jeff Templon (NIKHEF)
- Giacinto, Brian and Doug also have a relevant history dealing with Tier2/3s

- The mailing list
wlcg-wg-storagefederations@cern.ch
- The TWiki page
<https://twiki.cern.ch/twiki/bin/view/LCG/WLCGStorageFederationsWG>
- First milestone: would be very positive to summarize to the Federated data stores workshop in Lyon
 - A good candidate may be the “document current status and wishes” referenced before



The interested people

	dCache	Xrootd	DPM	STORM	EOS	ATLAS	CMS	ALICE	LHCb	Tier1	Tier2/3
<i>Costin Grigoras</i>								✓			
<i>Doug Benjamin</i>						✓					✓
<i>Rob Gardner</i>						✓					
<i>Brian Bockelman</i>							✓				✓
<i>Giacinto Donvito</i>							✓				✓
<i>Philippe Charpentier</i>									✓		
<i>Fabrizio Furano</i>			✓								
<i>Paul Millar</i>	✓										
<i>Andy Hanushevsky</i>		✓									
<i>Michele Dibenedetto</i>				✓							
<i>Dirk Duellmann</i>					✓						
<i>Luca dell'Agnello</i>										✓	
<i>Jeff Templon</i>										✓	
<i>Stuart Wakefield</i>							✓				

- Phone meetings, indicative duration is 60m
- If we go roughly weekly, I propose to have 8 of them
 - This fits our timescale of 3 months
- The next slides just contain a first proposal of such a timescale.
 - I am interested in having feedback and refining it

- Two meetings (total!) around the initiatives of the experiments: CMS, ATLAS, ALICE, LHCb
- Would be interesting to see advertised
 - Their ultimate goal in using the technology
 - Current usage of the tech: status, interests, caveats (e.g. name translations)...
 - Sites that are participating, which ones are joining
 - A message to the sites that should install their federated endpoints in the next year
 - What should they do to join
 - What level of effort is expected (e.g. manpower, debugging, babysitting, ...) for the site
 - Which are the parts that sites could share among experiments

- One meeting to discuss the usages of such federations
 - *“fail over” for individual jobs*
 - *“repair mechanism” of placed data*
 - *who initiates the repair?*
 - *how is metadata handled ?*
 - *pro-active repair?*
 - *access control and trust mechanism between sites*
 - *use of catalogues*
- Can these be matched among the experiments?
- What should the admin of a remote site do?
- Who will design and maintain each instrumentation?

- Two meetings (total!) around each storage federations technology we know about
 - Xrootd framework
 - Dynamic Federations (with WebDAV/HTTP frontend)
- A description of the set of consolidated features with respect to the federations tech
 - Where's the config? Central or in the endpoints?
 - Discovery, resilience to downtimes
 - Indicative steps for their tech to be deployed at a site
 - A collaborating one, AND one which does not know anything and will just look at it from time to time
 - Bundles available? How the config of a site is handled
- Would be nice to know about:
 - Why is it so cool?
 - What is the vision about making it really powerful and ubiquitous?
 - Are there contacts with other sciences/VOs setting up federations?
 - Is there an idea on a future interactive access (tools?)

- One session about monitoring
- I see different views on this
 - One seems closer to the “delayed accounting” idea, the other is closer to being “realtime” information
- The topic is very complex, and involves nearly everybody
 - What’s the status
 - Which features are requested (to whom?)
 - Mandatory (why?)
 - Nice to have
 - What is acceptable by the sites (e.g. open ports)
- Which info is produced
- How the info is delivered
- What collects it
 - E.g. What’s the plan to feed the EGI or WLCG infrastructures

- Sites day

- One talk from each site representative (Jeff, Luca, Giacinto, WEI?)
 - Experience so far with the existing federations
 - Wishes with respect to the tech providers/experiments
 - Would they install a bundled thing or will they work on their own completely custom thing? Doing what?
 - Message to the techs, to the exps. What are the requirements of the sites similar to yours ?

- Summary day
 - Putting together the message of the WG
 - What do you want to be said ?
 - What about storage-federated WLCG sites?

- Proposal
 - We fix now the next date
 - next Wed 08 Aug? Or the 22th ?
 - We start a Doodle poll for the subsequent 2 appointments

Thank you

Questions?