

ATLAS & DOMA



Few practical thoughts on Distributed Storage & stuff like that

- We are defining the scenarios we want to study to measure the *effectiveness* of distributed storages
 - *Effectiveness*: whatever this means! (i.e. defining/finding the important metrics is a challenge per-se!)
 - Basically trying to figure out how *far* storages can be and still be seen as "one storage" to the experiment. Which workflows? Which data access methods? etc.
- A lot of tools already in place:
 - For now require experts to configure and understand things
 - And a lot of tiny details to get things right (working), different dimensions.
 - E.g. Storage topology -
 - AGIS allows complex config of storages, with LAN and WAN capabilities, firewalls, redirectors, etc: flexibility is good but doc/instructions/usability not optimal yet... many different configs/deployments: results will be most probably very much site(s) dependent.
 - E.g. Precise testing and measurements:
 - HammerCloud - we see this as the possible key tool to allow sites and group of sites to optimize their config.
 - multi-dimensional config: in 2008(or so) when HC was born was a much easier task!
 - ATLAS is working to give sites (group of sites) the capability to optimize by themselves.
- Monitoring and Analytics:
 - One of the (many) fundamental bit we need to get it right from the beginning.

- We see this as a medium/long term activity:
 - The activity is already ongoing and, in general, can be bursty (few experts meeting together and fixing/understanding the various bit and pieces).
 - Order of magnitude, we think 6 months is reasonable timeline to start having first solid numbers.
 - N.b.: we need experts very well aware of all the details: another important part where we have to work on is to be able to "open" to newcomers.
- We are trying to find the right person(s) (with deep ATLAS knowledge) to constantly push this activity and help also in the organization