Storage Management TEG:
status update

Wahid, Daniele

Credits also to Dirk and Brian wherever the topics overlap with the DM-TEG

First steps and general bits

Definition of the membership and balance among areas of expertise

- ~30 members at the start → personal profiles collected
- expanded to >50 people on the e-group (caveat: >~10 observers)
- introductory set of SM-oriented questions (see later: “list 1”)

Organizational bits

- wlcg-teg-storage@cern.ch (open, self subscription enabled)
- https://twiki.cern.ch/twiki/bin/view/LCG/WLCGTEGStorage

Meetings

- Working since roughly one month
- Kick-off meeting on Nov 10th
- Already several meetings since then, also joint with DM-TEG (more later)
  - focus on Qs to experiments (see later: “list 2”)
  - attendance to meetings is >20-25 individuals (on average)
    - not always the same people - although chairs writing minutes could identify a bulk of SM-TEG addicted

Breaking the boundaries among TEGs

SM-TEG and **DM-TEG** approved their mandate and started to work separately in November 2011

- the chairs of each TEG tried to follow the other TEG's meetings at their best
  - not always easy...

It was soon realized that a large overlap exists on many topics

- one topic may indeed "live" within one TEG - as from the original mandate...
- ... but the topic may not be deeply examined without the other TEG's input
  - and, if the target is to form coherent strategies, the two TEGs must work together

Discussions among the chairs, on an overall process optimization

- Good agreement on targets and priorities
  - review the topics as from the mandates
  - decide which is the "owner" TEG in starting/following the work on each topic
  - protect the TEG members from duplicated attendance to meetings on equivalent items
    - Please note that the richness of the TEGs is the membership, not the chairs. And *all* TEGs members are experts and potentially very useful to the task, but also very busy, they help on the TEGs only for a relatively small fraction of their work time, and should be protected (by the chairs) against time wasting in the overall process

SM/DM-TEG topics

Topics in the **SM-TEG** *(as from the original mandate)*:

- [SM.1] Experiment IO/usage patterns (and so performance requirements for storage)
- [SM.2] Requirements on future storage systems, and how storage will evolve
- [SM.3] Separation of archives and disk pools/caches
- [SM.4] Storage system interfaces to Grid (SRM status?)
- [SM.5] Filesystems/protocols (standards?)
- [SM.6] Security/access controls (TBD: to be discussed with the Security TEG)
- [SM.7] (as site-run services:) storage management interfaces, performance measurements, monitoring, etc

- Topics in the **DM-TEG** *(as from the original* mandate)*:

(*) roughly, since DM already presented and re-elaborated the list

- [DM.1] Review of DM demonstrators from Summer 2010
- [DM.2] Dataset management and Data placement (policy-based or dynamic)
- [DM.3] Data federation strategies
- [DM.4] Transfers and WAN access protocols (HTTP, xrootd, gsiftp)
- [DM.5] Data transfer management (FTS)
- [DM.6] Understanding data accessibility and security requirements/needs
- [DM.7] POOL
- [DM.8] ROOT, PROOF frameworks
- [DM.9] Management of namespaces
- [DM.10] Management of catalogues (LFC, future directions)

Reshuffling of {SM,DM}-TEG topics

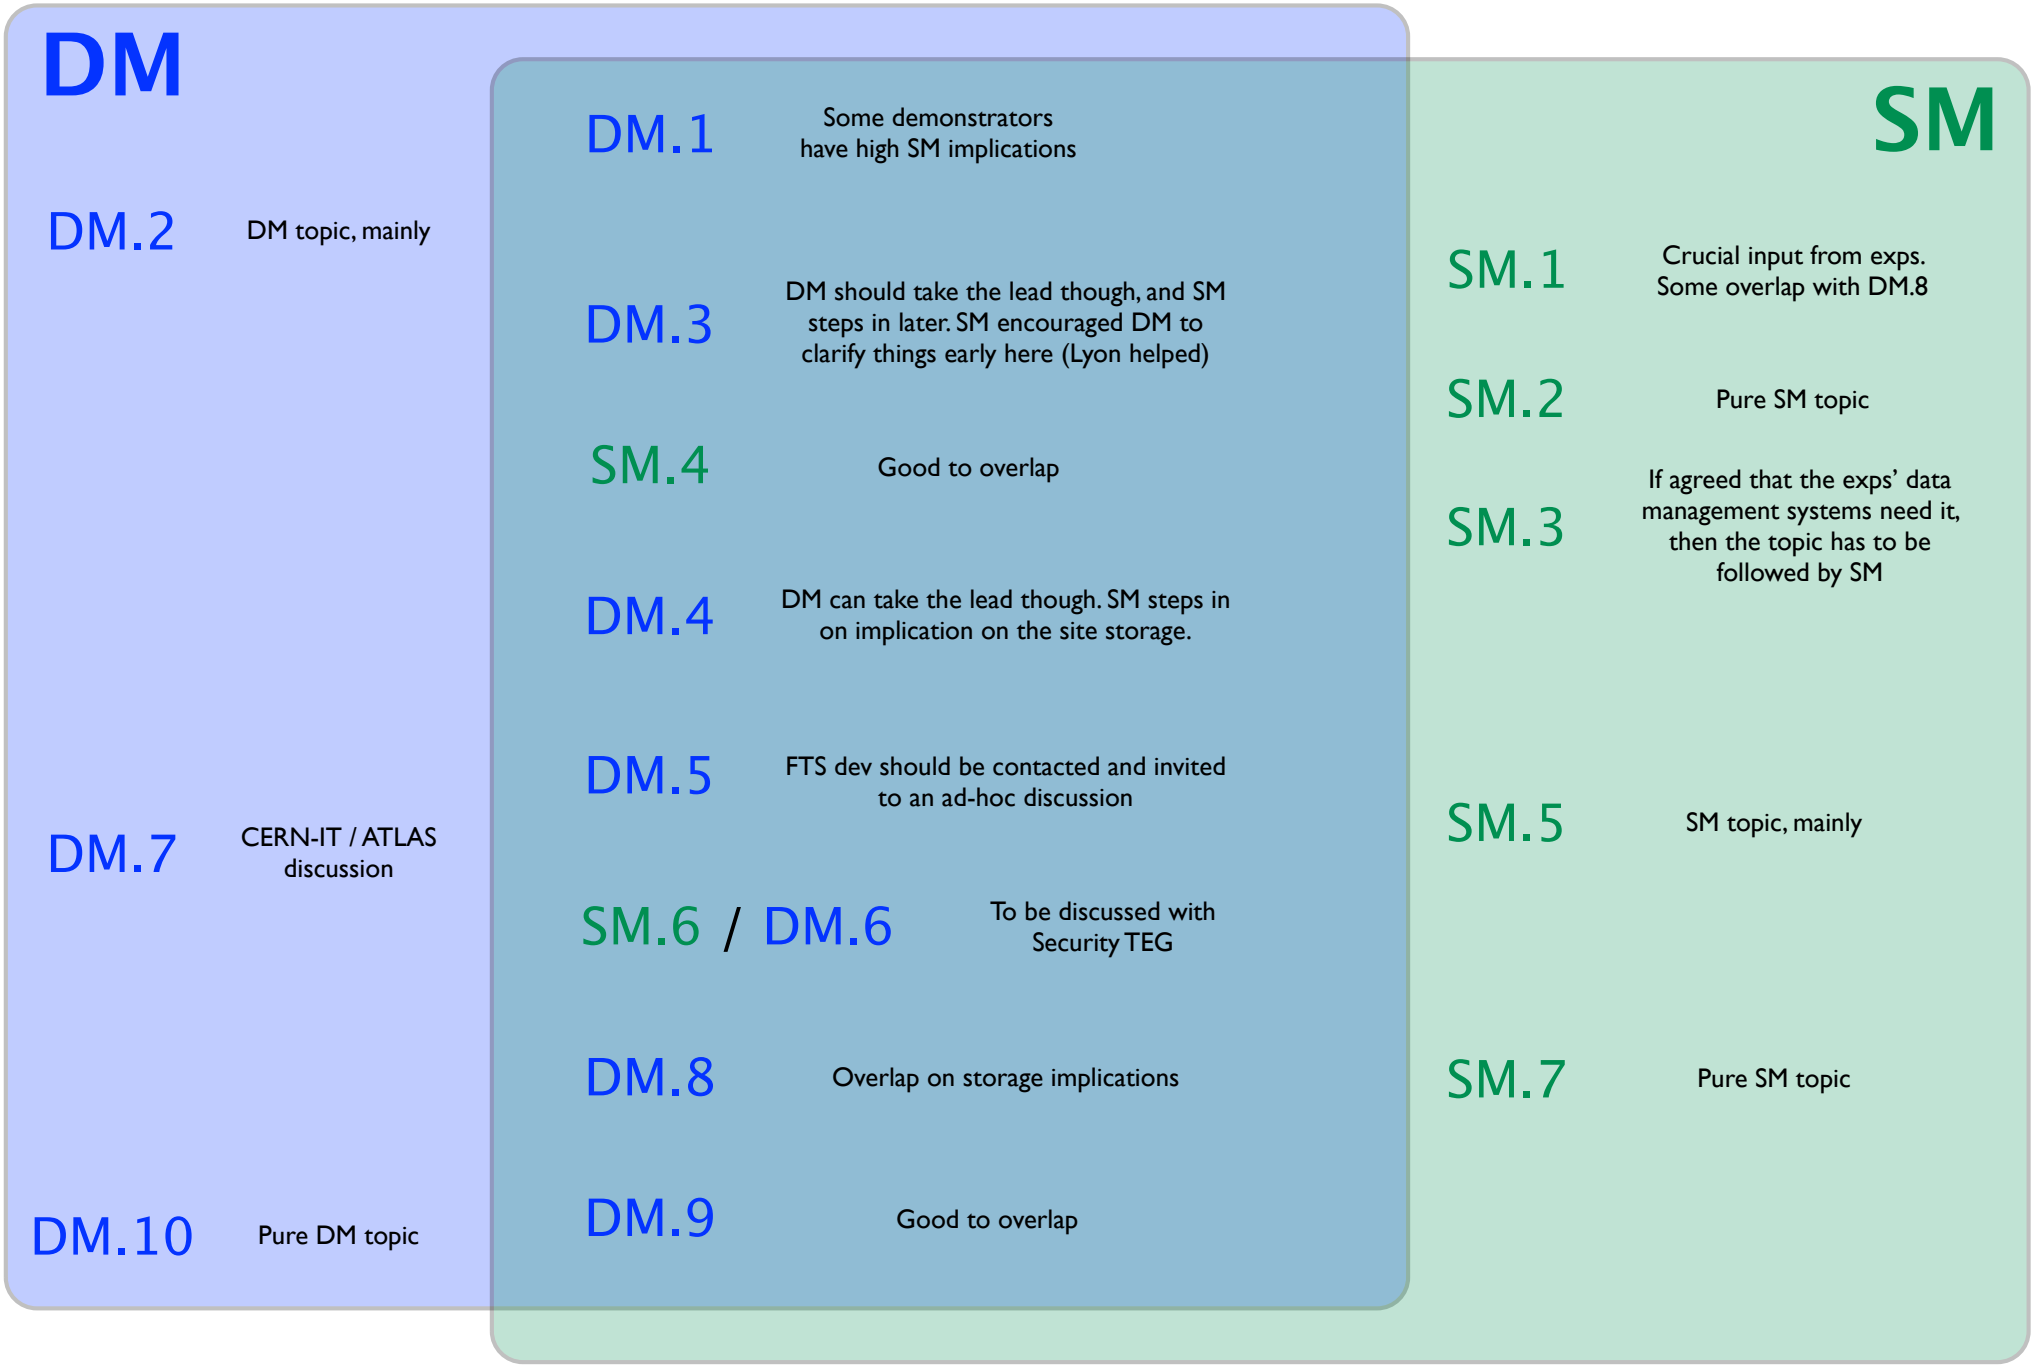

First goal: feedback from experiments

We focussed first on experiences and needs by the experiments

- try to identify commonalities (in requirements, choices, needs, complains)

We are in the information gathering phase (3/4 exps DONE)

- unique set of DM/SM-TEG questions to experiments (see next)
  - joint DM/SM-TEG meetings to listen and discuss the answers by experiments
- Very effective so far:
  - experiments in general very generous in details and elaborated answers
    - precious input! Thanks!
  - large attendance to experiment-specific joint meetings, plenty of discussions
  - minutes available for absent TEGgers
- More info:
  - ALICE: https://twiki.cern.ch/twiki/bin/view/LCG/WLCGTEGDataManagement_ALICE
  - CMS: https://twiki.cern.ch/twiki/bin/view/LCG/AnswersCMS1
  - LHCb: https://twiki.cern.ch/twiki/bin/view/LCG/WLCGTEGDataManagement_LHCb
  - ATLAS: answers in discussion internally - will be presented and discussed at the next TEG meeting

Questions to SM-TEG members

List 1

Asked to all members before the first kick-off meeting:

- Q: “In your view, what are the 3 **main current issues in Storage Management (SM)**?
- Q: “ What is the **greatest future challenge which would greatly impact the SM** sector?”
- Q: “ What is your site/experiment/middleware **currently working on** in SM?”
- Q: “ What are the **big developments that you would like to see from your site/ experiment/storage system** in the next 5 years?”
- Q: “ In your experience and area of competence, what are the (up to) 3 **main successes** in SM so far?”
- Q: “ In your experience and area of competence, what are the (up to) 3 **main failures** or things you would like to see changed in SM so far?”

Questions to Experiments

List 2

Overarching questions

- It has been asserted that we have a working system and do not need a new system, but the current system (cost and complexity) exceeds what is needed. Would the experiment agree with this statement?
- What elements are being used, and how? Which have been added and why?
- Which components failed to deliver (larger) functionality needed? Does that mean we can deprecate and obsolete these components, or was it not truly needed?
- If the middleware dropped complexity and removed functionality, does the experiment have the resources to adopt to the change?
- Where should the complexity and intelligence lie - the experiments or the infrastructure? How do you view the balance now, and how would you like to see this change (if at all)?

State of play: resource use

- What volume of data do experiments plan to move / store (up to the 2017 timescale) (and what do they currently move /store).
- What kind of file access performance are they expecting (any changes expected in next 5 years?) - what WAN data transfer rates?
  - If there is already information provided to the WLCG MB along these lines - could you point me where to find it.

File access: SRM

- Is the archive / disk split an agreed-upon strategy in your experiment? Can HSM be dropped?
- Can we get rid of SRM in front of our storage services? Can we get rid of SRM for disk-only storage?
- How thoroughly does the experiment use space management today?

File access: Efficient Data Placement

- What is the experiment's interest in revising data placement models? What kinds of revisions have gone through?
- What is the experiment's interest in data federations?
- What sort of assumptions does the experiment need to make for data federations work?
- For smaller sites, would caching work for your experiment?
  - (Can refer to Lyon workshop if appropriate for these Federation questions)

File access: WAN protocols

- Do you need gridftp? Can you use http?
- Does your framework already support HTTP-based access?

File access: Clouds

- Could you work directly with cloud file systems (i.e., GET/PUT)? Assume "cloud file systems" implies REST-like APIs, transfer via HTTP, no random access within files, or at least limited byte-range access, no third-party transfer, and possibly no file/directory structure. See Amazon S3 for inspiration.

File access: Local

- Data access protocols - do you need our own special ones (xroot, rfio, dcap)? Do you see that http or WebDav be used here too?
- Could you work directly with clustered file systems at smaller sites?

Security / access controls:

- Security / VOMS - what are your expectations / needs?
- Access control - what is really needed?

Namespace:

- Where do you see your namespace management "philosophy" evolving in the future
- How will we manage a global namespace - whose problem is it? the experiments? Is there a continued need for LFC?

What's done and what's next

- digestion and approval of the mandate
- definition of membership, guarantee a careful balance
- organizational bits (twiki, Indico, etc), plus collection of personal profiles of the members
- definition of (brief) set of general questions on SM personal viewpoints (see "list 1")
- collection and digestion, before the kick-off meeting
- kick-off meeting
- definition of boundaries with DM-TEG, definition of work plans with them
- Lyon workshop on data federation
- definition of a unique list of joint DM/SM questions to LHC exps (see "list 2")
- information gathering phase from **ALICE**, ad-hoc meeting
- information gathering phase from **CMS**, ad-hoc meeting
- information gathering phase from **LHCb**, ad-hoc meeting
- information gathering phase from **ATLAS**, ad-hoc meeting
- digestion of the exps' feedback, and synthesis
- on each sub-topic, the "responsible" TEG elaborates a way to proceed
- potentially large (to be parallelized) homework for different subsets of members
- from each subset: feedback collection, digestion, synthesis
- work on coherence, first steps towards a draft report, including risk management analysis
- aim to have a statement on "where we are" by/at the Amsterdam TEGs week
- post-Amsterdam wrap-up: work on "the future" over Feb-Mar
- ...

Some thoughts and open questions

Some "operational" worries:

- following the activities/meetings of other TEGs is tough
  - challenging for the chairs, impossible for most members
  - the risk that something falls in the cracks and the final outcome is not coherent/complete should not be underestimated
- the Xmas holiday break is a precious time: we should profit of it
  - i.e. homework(s), despite reluctant to ask
- a joint DM/SM meeting in Amsterdam is turning into a deep "TEGs f2f week"
  - different for the original format we had thought of, potentially very fruitful, but...
  - ... also extremely dense, need to prepare it carefully
- we will need time to wrap-up after Amsterdam
  - will we have enough time to be ready to deliver by early Feb?

Areas where we might fail

- we should be more ambitious
  - how can we be so, if asked not to disrupt successful operations?
- hard for TEGs to write a doc that can be inserted into a coherent scheme
  - do we need somebody to start thinking about the deliverables writing phase?