

# Brainstorming Workshop Agenda

1. Introduction:
  - Goals of the day, what we need to do, expected outcomes of the workshop; constraints and assumptions that we should take into account; understanding of how to address cost issues (not necessarily just assuming fixed budget)
2. Infrastructure models
  - How to reduce the overall cost of the infrastructure, e.g. by:
    - Reducing the amount of scattered storage
    - How to procure and incorporate large scale compute resources at scale
    - What is the networking model - OPN-style private?
    - The use of other resources: HPC, volunteer, etc. What is the model for distributed sites (Tier 2)?
  - The idea of setting up an early prototype of a data/cloud infrastructure to test ideas (a la “Zephyr” proposal) - suggestion already supported by several Tier 1s
  - What about specialist sites/resources (HPC, GPU etc) for specific workflows
3. Data management tools and infrastructure
  - Can we think about common toolsets for all experiments? What is achievable?
4. Technology Evolution
  - Should we set up a technology tracking activity (as Bernd has been doing), but perhaps now realised together with the concurrency forum and the tech lab activity – to provide a real testing and benchmarking environment (openlab here too).
  - Could be a distributed collaborative effort - in common with HSF/concurrency forum
5. Software topics
  - What is actually useful to do in common?
  - Tools or support for common automated and intelligent build/test/validation services?
  - More ? Common libraries etc? Role of HSF
  - What is usefully shareable?
  - Detailed software performance analysis and tools for analysis
  - All this should be in the context of the HSF. In fact, some activities towards common (lower-level) build tools was started already.
  - Propose a dedicated HSF workshop - making sure the proper audience is present.
  - Do experiments actually want to commit to the HSF?
6. Understanding and Modelling of the distributed infrastructure and computing models
  - How well do we understand our current workflows, their behaviour and resource needs?
    - with respect to storage, remote access, networks, CPU, memory
    - How well do we understand the behaviour of our current infrastructure?
    - What can we do to improve this understanding in an experiment independent way?

- how independent can this be?
- What has been done already in experiments?
- What would be desirable? Ability to model ideas of infrastructure to understand performance, costs, etc.
- What is potentially common across experiments? What is specific?