B) Data Centres and Facilities Technologies (Clouds, Networks, Storage)

Panel: Paul Rossman (Google), D.Kcira (Caltech), B.Holzman (Fermilab), Ian Bird (CERN), K.Tomko (OSC)
Moderator: Oliver Gutsche (Fermilab)

Panel questions:

1. What are the most significant technical improvements and changes that you expect to see in this area over the next 5 years?

2. How do you see the costing model for compute, storage, and networking evolve over the next 5-10 years?

3. What hardware capabilities are we not making effective use of now that we should be trying to better leverage capabilities and to improve our efficiency?

4. How can we make the most effective use of heterogeneous/new resources in HEP (technology changes, computing models evolution)?

5. What are the major changes required in the software stack to data management and distribution at the HL-LHC time scale (~2025) and the challenges in network and storage technologies ?

Live minutes:

Main expected changes

- Bill (Altair): main changes expected are on the SW side that will have to deal with a much more diverse HW and will need to integrate power management in its scheduling decisions as power=money.
- Burt: no major changes expected in the 5 next years, a short time for any major technology change (new memory technology that could be used for cold storage like Memristors will not take over in such a short time). Disks remains the most active market as the demand on this particular market, conversely to compute, is still growing.
- Dorian(Caltech): storage has become bigger and relying more and more on the network: Ceph is a good example and may change the storage in the future. 100G cards in servers coming: need to see what does it mean. SDN is also an important technology that will enable new usages and use cases.
- Ian Bird: the border between facility and distributed computing is becoming weaker. Network is more and more important, as monitoring to understand how the facility is used. Managing the whole memory/storage hierarchy is another challenge that we are just starting to tackle.

A future without facilities

- I. Bird: difficult to imagine this in the next 10 years. For CPUs we know how to provision resources at facilities we don't own but there is an impact on the network cost. We need to find a way to do the accounting of the network costs: we don't have it now
- Burt: the positive effect of trying to understand our costs and have a cost model when using commercial providers is that we have a better idea of our internal costs.
- CPU costs going down and may be almost free (relatively to other components) in 10 years. Storage remains a medium cost but the most expensive resource is network.
- Bill: what is important is to make cost visible to end users. They cannot adapt their behaviour if they don't have the information about the costs.
- Kate (Ohio SC): charging IOPS is a good idea as poor IO patterns are hurting performances not only for the application but also for other users who are impacted too. If users are charged, they will think more at what they do.
- I. Bird: the cost to optimize is the overall cost. No point in optimizing only one component. Need to ensure that we understand the impact of optimizing one component on the others ones (e.g. streaming data to compute resources vs. wasting CPU cycles).
- Miron: the ultimate goal is not to optimize cost but to find the best tradeoff between science and cost. Need the appropriate tools to decide this tradeoff: how we can have a chance to get them.
- Dorian: agree in principle with Miron but in practice this is difficult to attach a specific value to a (fundamental) research done by somebody… Cannot refuse a physics payload because it is inefficient. Decision on what is worth or not must be done elsewhere, in the experiments typically
- Ian B: it's true that we are entering a period where we may not have the resources to execute all the payloads from all experiments but we are not in a position to decide/prioritization. This has to be done at another level.