

Digital Research Network Infrastructure

Dave Britton and Pete Clarke, on behalf of GridPP.

This is a proposal to upgrade the network connectivity at institutes that host GridPP Tier-2 sites, and the RAL Tier-1 center. This would benefit all researchers at these institutes, many of whom share network links with HEP, by building a solid foundation on which future growth of digital research could be based. Such an investment would protect the effectiveness of investments made in other network areas.

Context

- The LHC has been a great success and is in the middle of an initial data-taking phase before the first scheduled long shutdown starting in 2013. The luminosity achieved has surpassed expectation and the amount of data collected has been spectacular. The worldwide LHC computing Grid (wLCG) has smoothly handled loads well above the design levels. The second phase of data taking is expected to run from October 2014 until 2017, followed by a second shutdown that will provide for further upgrades.
- wLCG is a tiered Grid Computing infrastructure based on the Monarc simulations of the late 1990's. However, these simulations vastly underestimated the growth in network bandwidth that would occur over the subsequent decade. We are now in a position where networking has become a much more valuable, and more powerful, resource than originally envisaged. The LHC experiments are adapting their computing models to capitalize on this resource, which in turn will allow more efficient use of storage and computational resources. For the international connections between Tier1s wLCG has already put in place an optical private network but the more general connections (Tier-1 to Tier-2 and Tier-2 to Tier-2) rely on the general purpose national research and education networks (JANET in the UK) shared with all university and research traffic.
- UK science (and STFC science in particular) has benefited from an excellent network during the era of JANET5. It is vital that in the future the network component of the e-Infrastructure remains commensurate with the goals for UK excellence and competitiveness in science. JANET6 planning to support research, made manifest through various consultations with the science community, gives us every confidence that JANET6 will provide the core infrastructure needed to ensure that this is so.
- In order for a distributed computing e-Infrastructure to realise the full potential of an excellent core network it is essential that there is no "last mile" bottleneck. This refers explicitly to the network infrastructure connecting hardware resources out across the campus and the JANET Regional Infrastructure (JRI).
- From the perspective of GridPP, one part of this "last mile" is the internal network within the various GridPP computing sites and is an area that can, and is being, addressed. As the capacity of disk storage continues to increase it is becoming increasingly challenging to stream data to and from disk fast enough. This puts high demands on the internal network configuration of computing sites. GridPP is currently working with sites to future-proof their internal network configuration for the 2nd phase of LHC running (2014-2017).
- However, the majority of this "last mile", from GridPP's perspective, is the connection between the Tier1 and Tier2 centres out across campus to JANET, which is over an infrastructure complementary to both the centrally funded JANET network and to the

GridPP funded resource at the site. From a GridPP perspective, we define the “middle-mile” to be that part of the last-mile not directly under our control.

- Funding for distributed e-Infrastructure is needed to allow this area to be upgraded. The “scale” of current middle-mile connectivity is still that of the last generation, i.e. “a few Gbit/s” and often shared with the rest of the campus. These links are becoming saturated. The “scale” of the future is at least 2x 10Gbit/s links to provide resilience, and preferably dedicated paths for heavy users, such as some GridPP Tier-2 sites (it is guaranteed already for the Tier-1), to the edge of the campus.

Proposal

- The process proposed would be:
 - To engage with each institute to facilitate provision of the Tier-2 connectivity to the JANET regional infrastructure as described, and to work with the STFC network group to future-proof the Tier1 connectivity.
 - To engage with JANET (with whom we have an excellent relationship) to determine how to augment JANET6 access links for all Tier2 sites.
- The detailed equipment list can only be drawn up after iterating with the respective network groups but in general will include:
 - Provision of multiple 10 Gbit/s router and/or switch ports.
 - Provision of high performance switches and fibre connections,
 - Provision of lower level driving equipment such as CWDM capability.
- Costs would need to be established on an institute-by-institute basis as carefully tailored solutions are established. For example, the Glasgow Tier-2 currently uses about 2/3rds of the bandwidth of ClydeNet, whereas at other institutions the issues will be completely different, for example, related to departmental or campus level links. Nevertheless, it is likely that an investment of £50k - £100k at each of 19-sites would go a long way to addressing many of the issues. Thus, a total of £1.5m - £2m would enable substantial progress to address immediate issues. A larger total would allow the future-proofing to be extended to a longer period and might lever additional institutional investment as each local project grows to a critical size.

Other Projects

The proposal above is based on the understanding that the likely scenario is Capital Funds available only within the current financial year. There are additional possibilities should Resource and/or multi-year funding be available:

1. Some network-upgrade scenarios would incur recurrent costs that are unlikely to be classified as Capital. Thus, if Recurrent funding was available over multiple years, these options, which might have a longer-term benefit, could be considered.
2. The NASCEnT proposal, recently turned down for PRD funding, addresses some important future areas. This proposal could be re-presented as a 2-year project, rather than the 1 + 1 year that had to be submitted to PRD to meet the funding limit.
3. The push to enable public-access to research-council funded data could be addressed by a pilot study based on one LHC experiment. It would require, something like a 2-year programme of work for 2-3 FTE plus some associated hardware in order to store an additional copy of the data and allow appropriately controlled access.