

From the Transatlantic Networking Workshop to the DAM Jamboree

David Foster

CERN-IT

LHC Network Activities

2005-2010

- LHCOPN designed to ensure T0-T1 data transport and provide capacity for (most) T1-T1 needs.
 - Closed group of parties, CERN + T1's
 - Bounded problem with good estimates of needs
- T2 connectivity uses general purpose IP connectivity.
 - T2-T1, T2-T2
 - Open problem with poor estimates of needs

We have entered a new era

- A workshop on transatlantic connectivity was held 10-11 June 2010 at CERN.
 - <http://indico.cern.ch/conferenceDisplay.py?confId=88883>
 - ~50 Participants representing major stakeholders in R&E networking
 - ESNNet, I2, Dante, NRENs, NSF, DOE, Industry, Major Labs etc
- And revealed the following:
 - Flows are already larger than foreseen at this point in the LHC program, even with lower luminosity
 - Some T2's are very large (not new).
 - All US ATLAS and US CMS T2's have 10G capability.
 - Some T1-T2 flows are quite large (several to 10Gbps)
 - T2-T2 data flows are also starting to become significant.
 - The vision progressively moves away from all hierarchical models to peer-peer
 - True for both CMS and ATLAS
 - For reasons of reduced latency, increased working efficiency
- Expectations of network capability are reaching unrealistic proportions without forward planning.

Problems and Opportunities

- Networking Technology and Bandwidth is not the problem itself, but the flows are scaling up to occupy much of the existing infrastructures.
- This will not be a problem if:
 1. You can clearly define what you need, now and over the next few years.
 2. “You” (or your agency) can pay for it.
 3. You can integrate it into a system, where the end-site facilities and networks operate together in a consistent fashion.

1. Define what you need

- The networking community needs a much better definition of real requirements.
 - Excessive use of general purpose networks will cause operators to take defensive action.
 - International network architectures, with sufficient reliability and capacity to cope with the expected traffic growth and flow patterns need to be designed and implemented.
 - This is not at all trivial; it needs planning and time
 - Experiments must work with the network community to create a definition that will enable implementation of infrastructures to support T1-T2-T3 matrix flows.

2. Pay for it

- Given an agreed architectural plan with capacity and other objectives (e.g. resilience and adaptability to shifting flows)
 - Optimal solutions in terms of costs can be found;
 - The limits of what could be afforded can be understood
 - The funding bodies can plan for the resulting infrastructure, within feasible cost bounds.
 - The sites can budget to connect.
- This requires conviction and excellent justification of the costs from the experiments – In line with the justifications given for other parts of the LHC program
 - and ultimately: Commitment from the funding bodies.

3. Integrate it into a system

- Provisioning capacity and providing it to end sites is manageable
 - If you can define what you need.
 - If you can show how your needs follow from a well-defined comprehensive operational plan, with well-justified costs.
 - And you can pay for it.
- Using it effectively is more of a problem
 - Site/Campus/Regional network issues
 - End system hardware issues
 - End system platform and interface issues (O/S, drivers, NIC, etc)
 - End system application issues
 - Experiments' data movement and job placement software.
 - Real end-end awareness is needed.

We Need

3 Core Interrelated Activities

- Data Architecture Vision
 - Encompassing all substantial data movement operations, with clear throughput or time-to-complete goals
 - Leading to clear (and effective) capacity planning
 - System-wide view enabling key efficiencies: e.g. strategic data placement with some form of Content Distribution Network (not a new concept)
- International Network Architecture
 - Takes time to put in place.
- End-end everything
 - Monitoring
 - Automated Operations
 - Consistent site and network resource-use and/or performance optimization

LHC Networking Activities

2010-2013

First Step

- A small planning group that engages a small number of people to come up with an agreed plan for network bandwidth requirements for the experiments.
 - That network architects can then translate into core infrastructure designs.
 - That funding bodies can understand and plan for; based on a convincing, comprehensive operational model
 - That sites can then understand how to connect to these infrastructures and any policy implications.
- A design group that will understand the architectural and organisational issues around an additional international core infrastructure.
 - The LHCOPN working group is a good start of key stakeholders.
 - Needs to involve and inform T2/T3 stakeholders in some managed way.
- This should be done now to ensure that the infrastructures will be in place for 2013 data taking.
 - In time for a “STEP13” exercise in 2012
- Must work closely with the other core activities.