WLCG Network and Transfer Metrics WG After One Year

Shawn McKee, Marian Babik

GDB
4th November 2015

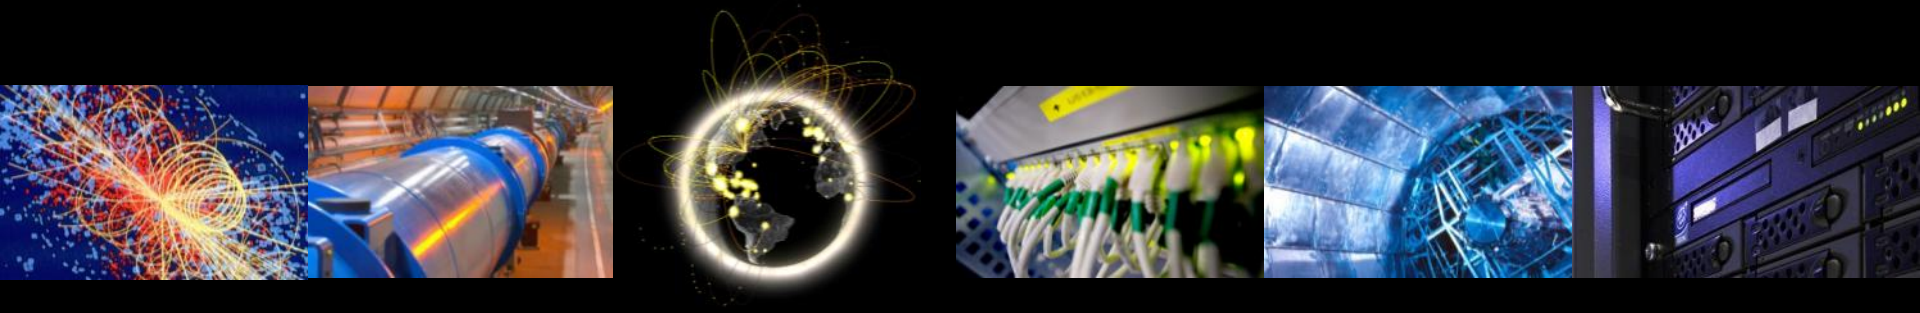

Review of Working Group

- After its first year, it is appropriate to review the WCLG Network and Transfer Metrics working group progress
- Started in Fall 2014, it brings together network & transfer experts
  - Follows up on the WLCG perfSONAR Task Force goals
- **Mandate**
  - Ensure all relevant **network** and **transfer metrics** are identified, collected and published
  - Ensure sites and experiments can better understand and fix networking issues
  - Enable use of network-aware tools to improve transfer efficiency and optimize experiment workflows
- **Membership**
  - WLCG perSONAR support unit (regional experts), WLCG experiments, FTS, Panda, PhEDEx, FAX, Network experts (ESNet, LHCOPN, LHCONE)

https://twiki.cern.ch/twiki/bin/view/LCG/NetworkTransferMetrics

Overview of Work Areas

- The working group has met on a monthly basis and made progress in a number of areas:
  - Defining important use-cases
  - Understanding slow transfers
  - Providing a uniform way to define and access network metrics
  - Integrating network metrics with the experiments
  - Coordinating responses to network issues
  - Baselining existing links; commissioning new links
- We will cover these in the following slides

Defining Important Use-Cases

- The working group has assembled feedback about the use-cases foreseen for network and transfer metrics into a google document
  - Participation from ATLAS, CMS, LHCb and Alice, and various middleware and application providers
  - Details at https://docs.google.com/document/d/1OcggHH4DM6vhe1ydJteYazwc0WgOLcHU31H1X5vbq3I/edit and are summarized later in this presentation
- We also have CHEP 2015 paper, detailing the working group, to be published soon. See abstract https://indico.cern.ch/event/304944/contribution/407.pdf

Understanding Slow Transfers

- Slow transfers can have many causes in complex end-to-end topologies
  - Differentiating “network” issues from application, end-site and end-host issues can be difficult
- A core task for the working group was perfSONAR deployment and commissioning to provide a source of network metrics to enable effective problem debugging
  - 245 active WLCG/OSG instances; 278 registered
  - WG tracks issues in infrastructure and networks
  - Supports problem debugging and resolution
- Uses perfSONAR metrics to identify “network” causes for slow transfers; conversely can rule-out network issues where perfSONAR metrics show good networks.

Providing Network Metrics

- Working closely with the Open Science Grid (OSG) Network Area, the group has provided access to perfSONAR network metrics
  - OSG Network Datastore is used to gather and store ALL WLCG/OSG perfSONAR metrics in a common API (Esmond).
  - CERN hosts a Active Message Queue that OSG "publishes" all metrics to. End users can subscribe to any data they choose
- The working has prototyped a proximity service to handle topology-related requests: closest perfSONAR server to storage element and vice-versa; closest perfSONAR to arbitrary end-host using GeoIP (eventually could use route data).
  - Needs testing and further development but functional now:
  - http://proximity.cern.ch/api/0.3/geoip/nearest?sonar=psum02.aglt2.org&count=10
  - http://proximity.cern.ch/api/0.3/geoip/matrix?vo=atlas
- Collaborating with ESnet on proximity service follow-on, an open-source configuration interface project and the deployment/debugging of perfSONAR toolkit.

Integrating Network Metrics

- It is important to be able to utilize the network and transfer metrics we gather to support higher level services that rely upon the network
- One example is the FTS study which used traceroute information from perfSONAR and FTS log details to analyze FTS performance versus various parameters (See later slides)
- Ongoing pilot projects with ATLAS and LHCb providing pipelines of metrics for the experiments
  - LHCb is processing perfSONAR stream and correlates it with the network and transfer metrics in DIRAC
  - ATLAS processing perfSONAR data to create a network "cost-matrix" for use by PANDA with additional use cases in scheduled transfers and dynamic data access

FTS Transfer Study Details (1/2)

- Saul Youssef undertook a study of FTS transfer data augmented with perfSONAR traceroutes to analyze how transfers are working and if there are specific problematic links in our networks.
  - The analysis uses the average transfer rate and assigns that rate to all the hops along the network transfer path for a given time-window
  - The sum of all transfers in a time-window for a specific network hop can be used to infer a lower bound on its maximum transfer rate
  - In this way hops with low maximum transfer rate are identified and can be checked to see if there are problems that need addressing

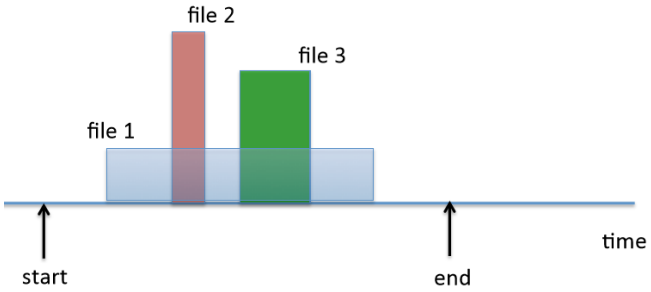

Transfer rate bounded >= size(file1+2+3)/(end-start)

- The study also identified a number of issues
  - Number of streams may not be optimized
  - Lots of SRM overhead impacting performance (esp. small files)
  - CMS consistently getting better throughput than ATLAS

FTS Transfer Study (2/2)

- The details of the analysis and results are available at http://egg.bu.edu/LHC%7Binf:LHC%7D/gadget:FTS/section:history/fts/index.htm
- This has become a collaboration among a number of people (including ATLAS, CMS, LHCb) and most recently the FTS developers / operators
- This collaboration is now exploring the impact of a new way to optimize FTS transfers:
  - Because of SRM overhead it is better to multiplex more transfers at lower number of streams each as opposed to using more streams and a smaller number of transfers
  - Each end-site has only so many resources (Storage I/O, memory and network bandwidth). The algorithm needs a tunable limit (MAXBUFFER) which may vary by site.
  - The number of TCP streams for best performance varies. Lower numbers may be better. In some cases 1 seemed optimal. The error rate (timeouts) with 1 TCP stream is only slightly higher than the rate with multi-TCP streams.
- Looking forward to seeing how this new strategy works

LHCb DIRAC bridge

- Federico Stagni, Henryk Giemza and Stefan Roiser are working on the perfSONAR-DIRAC bridge
- LHCbDIRAC provides information on transfer metrics such as
  - number of files transferred
  - “quality”, i.e. amount of successful/failed transfers
  - throughput
- Main use case is to correlate with perfSONAR and use it to better debug transfer issues
  - prototype is being developed
- Another one is to commission new links, determine capabilities of the new network links

ATLAS Use Cases (Mario Lassnig)

- ATLAS utilizes both scheduled and dynamic transfers
  - **Scheduled** flows for Tier-0 Export, Analysis Output, Production Input/Output, Data Consolidation
  - **Dynamic** for jobs brokered for remote access, local data access failure recovery and the event service
  - Both could benefit from using network and transfer metrics
  - We need better understanding of network infrastructure and characteristics for full performance and rely upon the working group to provide that information
- ATLAS perspective: storage is expensive — make better use of the network
  - Leads to more use of dynamic data access, less movement of files
  - Depending on activity, only small parts of the files are needed — prime use case
- Assumption: current network infrastructure is configured for long flows, big files — is this good for us in the long run?
  - Need a dedicated study how an increase in short flows on our infrastructure will behave
  - Can we do something on our side? Virtual switches with end-to-end traffic policing sound promising

ATLAS Network Metrics Pipeline

- Ilija Vukotic, Kaushik De, Rob Gardner and Jorge Batista are working with our WG to make perfSONAR metrics available to PANDA
- **Pipeline**: OSG Network Datastore -> CERN Active MQ -> Flume -> ES -> PANDA
- Prototype working and analytics being performed in Elastic Search to validate data
- Plan is to create a network source-destination cost-matrix PANDA can use to evaluate options
  - Actual interface details being discussed with PANDA team

Coordinating Network Issue Response

- The working group has created a support unit to coordinate responses to potential network issues
  - Tickets opened in the support group can be triaged to the right destination
  - Many issues are potentially resolvable within the working group
  - Real network issues can be identified and directed to the appropriate network support centers
- Documented at https://twiki.cern.ch/twiki/bin/view/LCG/NetworkTransferMetrics#Network_Performance_Incidents

Baselining Our Networks

- The working group has defined a number of meshes with corresponding tests to measure the performance of our network paths
- WLCG-wide meshes for **latencies**, **traceroutes** and **throughput**
  - perform full mesh tests involving all WLCG sonars
- Network provider meshes
  - **LHCOPN** and **LHCONE** meshes to help debug weak links
- Specific project meshes
  - Dual-stack – compares throughput between IPv4 and IPv6 instances
  - Belle II – covers the additional non-LHC sites in Belle II

Next Steps

- Main focus will be on the pilot projects in collaboration with experiments (ATLAS, CMS, LHCb)
- Integration of the higher-level infrastructure services (such as MadAlert, PuNDIT, etc.)
- Enabling easy integration with data analytics platforms (ES, Hadoop)
- Finalize production deployment of the base infrastructure and follow up
- perfSONAR 3.6 (release around summer 2016)

Summary

- The working group has made very good progress in all of its areas of responsibility
- ATLAS, CMS and LHCb are all involved and should benefit from our pilot projects and the FTS study work.
- We have had very successful (and continuing) collaboration with ESnet and OSG
- Focus during the first year has moved from deployment debugging issues to use of the metrics, which is where we want to be.

QUESTIONS?