Network and Transfer Metrics WG Meeting

Shawn McKee, Marian Babik

Network and Transfer Metrics Kick-off Meeting
8th September 2014

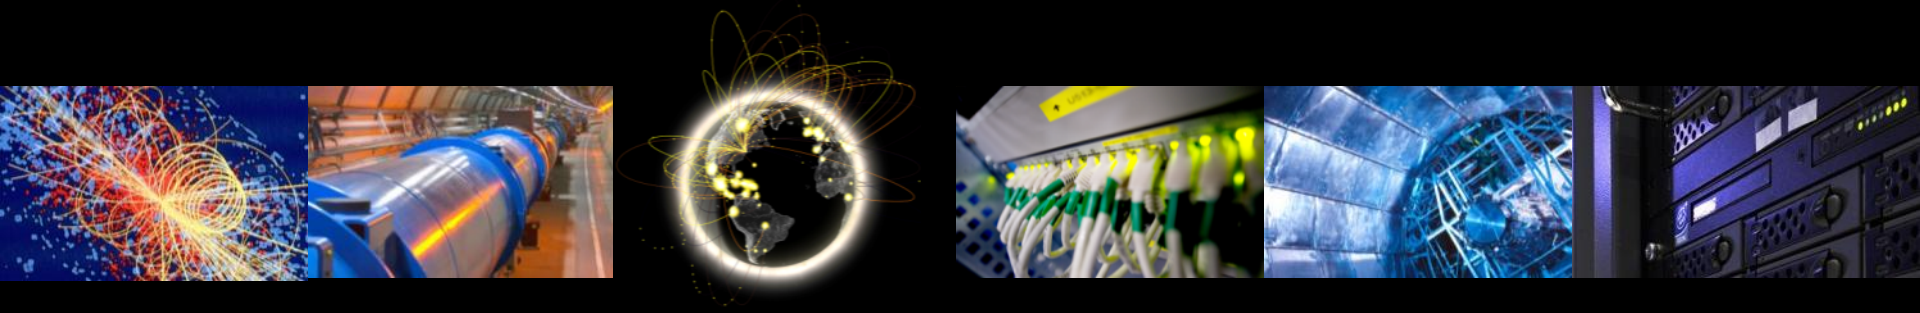

WLCG
Worldwide LHC Computing Grid

Outline

- Introduction
  - Network and transfer metrics WG
  - Mandate, Objectives, Team
  - Organization
- Network monitoring status
  - perfSONAR
- Transfer monitoring status
  - FAX/AAA, FTS, Rucio, PhEDEx
- Review of Proposed Topics/Tasks
  - Meeting schedule

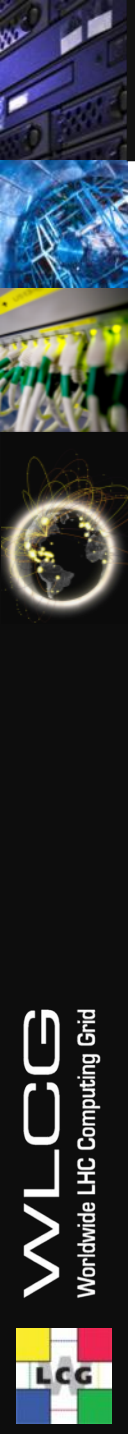

Network and Transfer Metrics WG

Welcome

- Welcome

Shawn McKee (chairperson), Marian Babik (co-chair), ATLAS (Simone Campana), CMS (Tony Wildish), LHCb (Stefan Roiser, Joel Closier), Alice (Latchezar Betev, Costin Grigoras), FAX (Ilija Vukotic), FTS (Mikael Salichos, Oliver Keeble), CMS/PhEDEx (Tony Wildish), Panda (Kaushik De), Rucio (Vincent Garonne)

perfSONAR contacts: US-ATLAS (Shawn McKee), US-CMS (Jorge Alberto Diaz Cruz), UK-ALL (Alessandra Forti, Chris Walker, Duncan Rand), IT-ATLAS (Alessandro de Salvo), IT-CMS (Enrico Mazzoni), CA-ALL (Ian Gable), FR-ALL (Frederique Chollet, Laurent Caillat, Frederic Schaer), TW-ALL (Hsin-Yen Chen, Felix Lee), ND-ALL (Ulf Tigerstedt), DE-ALL (Guenter Duckeck, Andreas Petzold, DE-KIT: Bruno Hoeft, Aurelie Reymund), ES-ALL (Fernando Lopez, Josep Felix), CERN (Stefan Stancu), LHCOPN/LHCONE (John Shade, ESNet: Mike O'Connor), RU-ALL (Victor Kotlyar), ESnet Science Engagement group (Jason Zurawski)

Mailing lists:
wlcg-ops-coord-wg-metrics@cern.ch (ALL)
wlcg-perfsonar-support@cern.ch (perfSONAR contacts)

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

Network and Transfer Metrics

- 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

WG Objectives

- Identify and continuously make available relevant transfer and network metrics
  - Ensure we can consistently publish all the relevant metrics
  - Common metric attributes semantics needed for analytics/correlations
  - Improve our understanding on what metrics are needed and how we can get them
- Document metrics and their use
- Facilitate their integration in the middleware and/or experiment tool chain
  - Work with experiments on their use cases
- Coordinate commissioning and maintenance of WLCG network monitoring
  - Ensure all links continue to be monitored and sites stay correctly configured
  - Verify coverage and optimize test parameters

WG Organization

- Purpose of this meeting:
  - Give an overview of current status in network and transfer metrics
  - Agree on a list of topics/tasks to follow up in future meetings
- Meeting schedule
  - We propose to have topical meetings, called in advance via doodle
- Reporting and Task tracking
  - WLCG ops coordination/report needed every 2 weeks
  - Proposing JIRA for task tracking
  - Volunteers needed to lead specific topics

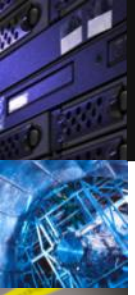


Network Monitoring Status perfSONAR

Motivation

- All distributed, data-intensive projects **critically** depend upon the **network**.
  - Network problems can be hard to diagnose and slow to fix
  - Network problems are **multi-domain**, complicating the process
  - Standardizing on specific tools and methods allows groups to focus resources more effectively and better self-support (as well as benefiting from others work)
  - Performance issues involving the network are complicated by the number of components involved end-to-end. We need the ability to better isolate performance bottlenecks
- WLCG wants to make sure our scientists are able to effectively use the network and **quickly** resolve network issues when and where they occur

Metrics

- **Network Path** - We use perfSONAR's traceroute to track the network path between WLCG sites. Currently 1/hour between ALL WLCG perfSONAR latency instances.
- **Bandwidth** - We use perfSONAR's Iperf tool to measure achievable bandwidth. Depending upon the grouping (mesh) we test every 6 hours both directions between src-dst. We additionally test **all** WLCG pairs 1/week)
- **Latency** – We send 10Hz of one-way delay measurement packets between **all** WLCG sites. The packet statistics (avg,min,max delay) are summarized every minute and any **packet losses** (x/600 packets) are noted (a critical metric)

Deployment Organization

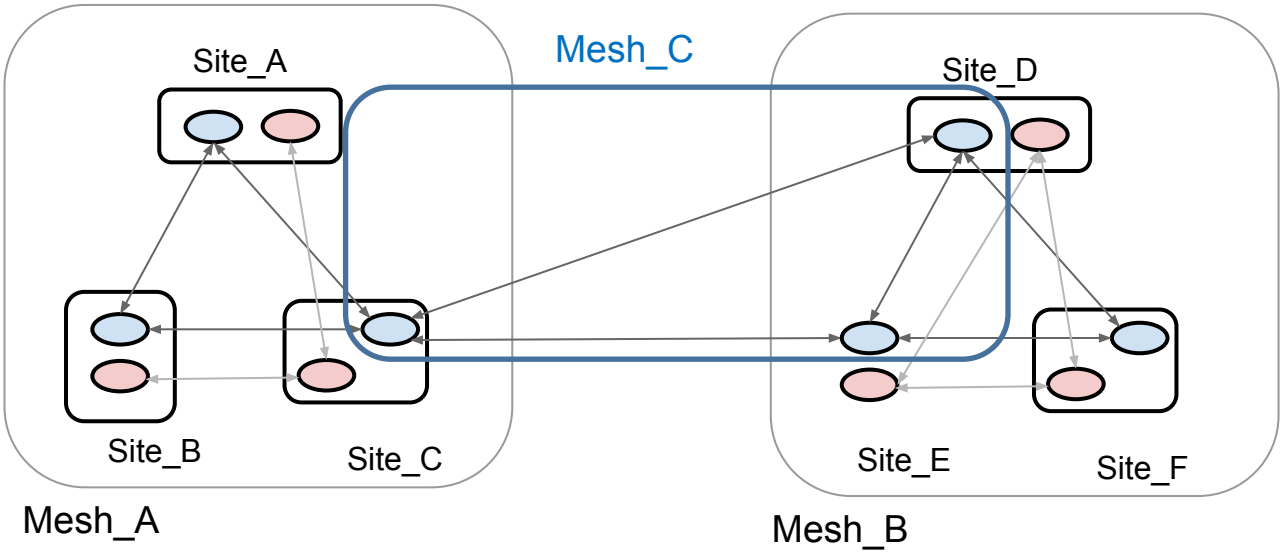

perfSONAR bandwidth

perfSONAR latency

We use groupings of perfSONAR instances (by region, experiment or both) to organize the testing for WLCG. These are referred to as "meshes" and we provide a central mesh-configuration JSON file to control tests, members and test parameters.

We can additionally test "between" meshes by adding subsets of sites to a new mesh (Mesh_C above)

Our working group can customize and tune the type of meshes and the testing details as needed.

perfSONAR network

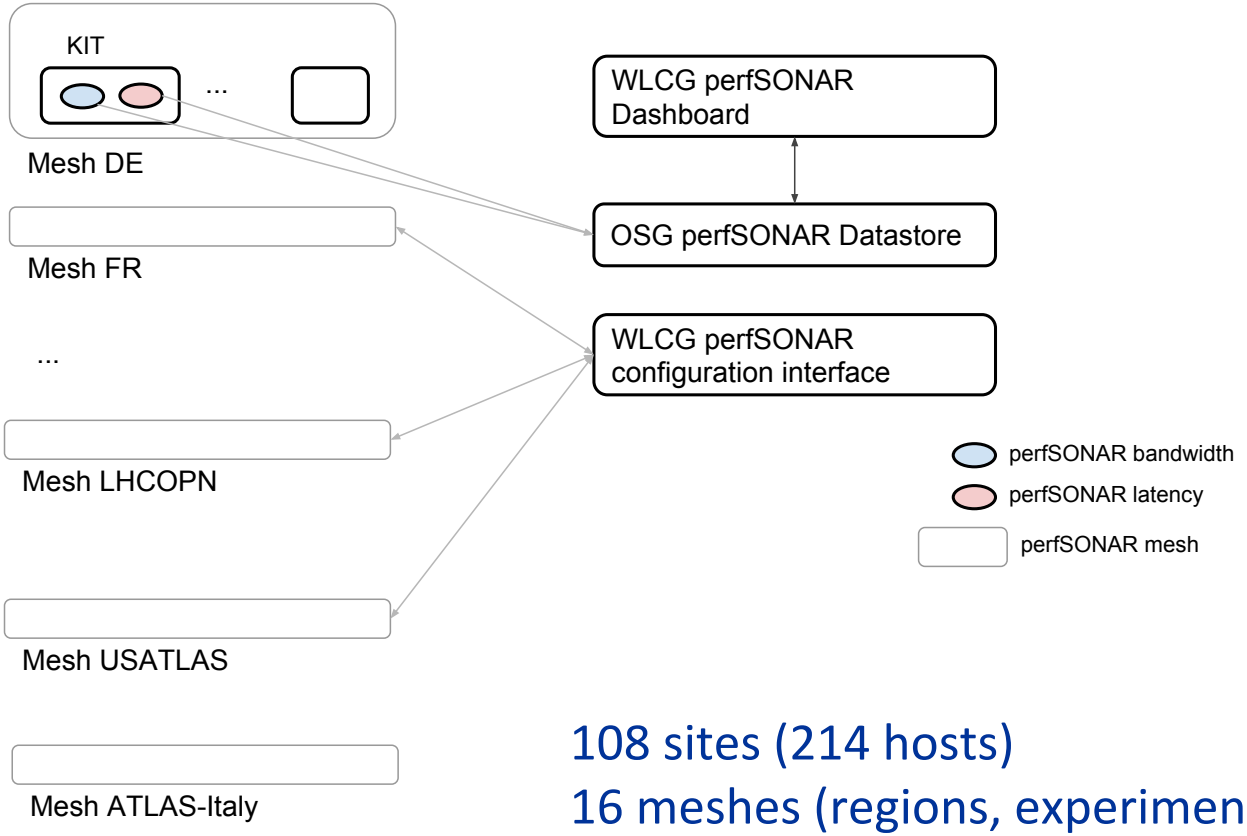

108 sites (214 hosts)
16 meshes (regions, experiments)
8 bandwidth tests each direction/day
Continuous latency tests (all-pairs)

Network Monitoring Topics

- perfSONAR is a “special” service
  - It tests a multi-domain network path, involving a service at the source and at the destination
  - It requires dedicated hardware and comes in a bundle with the OS.
- Installing a service at every site is one thing, but commissioning a NxN system of links is squared the effort.
  - This is why we have perfSONAR-PS installed but not all links are monitored.
- Tracking versions and updates is very challenging
- Firewalls issues. There are 2 kinds of firewalls rules to be considered:
  - For the hosts to be able to run the tests among themselves
  - For the hosts to be able to expose information to the monitoring tools.
- Test coverage and parameters
  - We need to ensure we’re consistently getting the needed metrics
  - Should we have more VO-specific meshes, i.e. WLCG-CMS, WLCG-ATLAS ?
  - What frequency of testing for traceroute (currently 1/hour)?
- Documentation (as usual) needs to be improved
  - How to use the information to diagnose network issues…

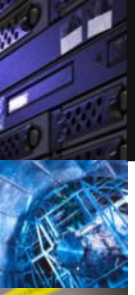


Transfer Monitoring Status

Motivation

- Many existing transfer tools/technologies used by experiments and in production
- Use cases:
  - Planning, design phase – network bandwidth reservations, algorithm design and optimization
  - Transfer routing phase – routing based on correlated network and transfer metrics
  - File-transfer phase – consulting network bandwidth for long transfer queues
- Current high-level overview focused on
  - FAX/AAA, FTS, PhEDEx
- Additional systems can be added to the mix

Metrics

- Transfer rates – bytes transferred between two endpoints at file-close event (volume/ throughput)
  - Other events published (file open, state)
  - Experiments metadata can be included
- Link status – error-rate for file transfers
- Errors – detailed error messages, mainly used for debugging purposes

Transfer Status

	FTS	FAX/AAA	PhEDEx
VOs	ATLAS, CMS, LHCb	ATLAS, CMS	CMS
Sources	Event messages from FTS service Command line client FTS3 monitoring service FTS dashboard service	XrdMon (GLED) FAX/AAA dashboards HammerCloud/SSB/FSB - synthetic tests	File download agents CMS PhEDEx dashboard (throughput and link)
APIs	FTS dashboard API Event messages in ActiveMQ	FAX/AAA dashboard APIs, SSB/AGIS APIs GLED publishes to ActiveMQ	Oracle DB with API
Status	Production	Production	Production
Detailed summary for each system in the backup slides			

Transfer Monitoring Topics

- Review of the transfer systems
  - Metrics published, deployment status, overlaps/gaps
- Transfer vs. Network monitoring methodology
  - Understand what exactly is measured in transfers and what with sonars – identify overlaps
- Coverage – transfer systems rely on SEs, perfSONAR relies on sonars
  - Proximity functions might be needed to correlate metrics
  - perfSONAR coverage tuning (meshes, test parameters)
- Topology
  - Granularity issues – measurements taken by hosts at different locations, but then aggregated to sites
  - Sites are specific to each experiment (sometime non-trivial mapping to physical locations)

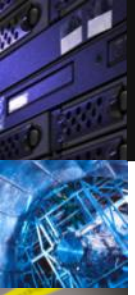


Topics/Tasks Review

Topics/Tasks Review

- T1.1: Gather requirements and use cases from experiments, data and workload management systems
- T1.2: Review existing transfer and network metrics and determine how they overlap (or can be extended wrt semantics), identify gaps (missing metrics)
  - Determine current operational status of underlying systems, ensure we can consistently make available all needed metrics
- T1.3: Determine current test coverage and propose how to tune parameters and mesh configurations
- T1.4: Propose how to map perfSONAR topology to transfer systems topologies (and vice versa)

Topics/Tasks Review

- T2.1 perfSONAR Commissioning/Operations
  - Tracking status of the network and developing procedures to follow up incidents/requests
  - Improve docs, create troubleshooting guide with a clear mapping to status tracking
  - Plan deployment campaigns
- T2.2 perfSONAR Storage
  - WLCG/OSG data store testing and deployment
- T2.3 perfSONAR Configuration
  - WLCG configuration interface – testing and deployment
  - Parameter tuning, host allocations tuning

Meetings

- Transfer Metrics Meeting
  - 13th – 17th October
  - http://doodle.com/xvwdvysdrdzap8wh
  - Gather requirements and use cases and review current status in transfer area
    - T1.1, T1.2
- WLCG perfSONAR meetings
  - 29th Sept – 3rd October http://doodle.com/e6epkkqmdx6ka3r7
  - 20th Oct – 24th October
  - http://doodle.com/qydib32fkv48er2r
  - Coordinate commissioning and maintenance of WLCG network monitoring
    - T2.1, T2.2, T2.3

AOB

- CHEP2015 abstract




Backup

Motivation

- Optimized experiment workflows
  - Network resource allocation along with CPU and storage in data and job placement
- Integration of network-awareness at different levels of experiment workflows
  - Planning, design phase – network bandwidth reservations, algorithm design and optimization
  - Transfer routing phase – routing based on correlated network and transfer metrics
  - File-transfer phase – consulting network bandwidth for long transfer queues
- Preserve investments already made in network and transfer monitoring

Areas

- Monitoring
  - Allows re-active use – react to events or situations in the networks based on network and transfer throughputs
    - Raise alarms and/or continue
    - Abort re-start transfers
    - Choose different source
  - Topology monitoring
    - Detect site isolation cases
    - Influence source selection
- Network control
  - Allow pro-active use
    - Network bandwidth allocations
    - Resource allocations based on CPU, storage and network
    - Optimize network topologies based on operational conditions

Use case: Find/Localize Problems

- Identifying and localizing a network problem
  - Often this is very difficult and time-consuming for Wide-Area Network (WAN) problems
- Scheduled perfSONAR bandwidth and latency metrics monitor WLCG network paths
  - Significant **packet-loss** or **consistent large deviation from baseline bandwidth** indicate a potential network problem (see in GUI or via alarms).
  - On-demand tests to perfSONAR instances can verify the problem exists. Different test points along the path can help pin-point the location.
  - Correlation with other paths sharing common segments can be used to localize the issue.
  - **The time** things change is also very useful to find the root causes. Scheduled tests provide this.

Use case 2: Optimize Net Use

- Both ATLAS and CMS are working together on an NSF project: **A**dvanced **N**etwork **S**ervices for **E**xperiments (ANSE), adding network awareness into PANDA and PhEDEx
  - In both cases they can benefit from network info
- Use network information for:
  - AAA/FAX brokerage
  - Job/task assignment
    - Improve flow of activated jobs
    - Improved accounting and diagnosis for transferring jobs
  - Site selection (data sources vs available CPU)
- The plan is to utilize network metrics to make intelligent decisions regarding our use of the network

File Transfer Service

- Low-level data movement service, moves data sets from one site to another (SE to SE)
  - Used for majority of LHC 3rd party transfers
- Transfer Metrics
  - FTS publishes event messages to ActiveMQ [MSG]
    - For each individual transfer (following event types: start, complete, state)
    - In addition, command line client reports aggregated snapshots
  - FTS3 monitoring service [CERN, RAL, BNL, FNAL]
    - Captures throughput and success rates that are used in optimizer [ref]
  - FTS dashboard service provides in-depth details [link]
    - Consumes raw data from message bus and calculates statistics
    - Supports ATLAS, CMS and LHCb
    - Calculates transfer rates (throughput, volume) per site, vo, host, country, token, job
    - Aggregates and reports on common errors
- Deployment
  - All services mentioned are deployed and operated in production

Federated XRootD services for ATLAS

- XRootD monitoring
  - Provides two types of monitoring: summary and detail [ref]
  - Monitoring information is sent over UDP (near real-time with low impact, medium latency - 5 minutes or full buffer 1-8kB)
- XrdMon
  - Based on GLED, collects and aggregates information
  - Maintains in-memory representation of all ongoing user sessions and file transfers, highly configurable
  - Supports multiple backends (ActiveMQ, Gratia, TTree, Http)
- FAX/AAA dashboards
  - Raw data consumed from ActiveMQ (published by XrdMon)
  - Events generated for every file close operation [ref]
  - Dashboard calculates statistics on transfer rates (throughput/volume) and active/finished transfer counts per site, country
- HammerCloud/SSB/FSB
  - HammerCloud continuously submits jobs to ~40 ATLAS analysis queues [ref]
  - These jobs copy special files from each of the FAX endpoints (every 15min, including traces)
  - Results reported to SSB and FSB [ref, ref] and via AGIS to JEDI for brokering

Physics Experiment Data Export

- Metrics
  - PhEDEX file download agents (at the sites) report on success/failed transfers, actual transfer are performed by FTS
  - This is aggregated centrally into throughput and link status (per CMS site names)
  - Throughput - counts the rate at which files arrive at a destination (into one hour bins) RatePlots
  - Link status - counts the error rates for transfers QualityPlots

- Throughput and link status are used to determine the best source-site for a given dataset/file
- PhEDEx doesn't need physical host awareness - it's "link" is a logical connection between two nodes (nodes don't have any association to physical hosts)
- Deployment Status
  - All metrics are stored in PhEDEx Oracle DB at CERN (including history), accessible via API
  - PhEDEx has Production instance, Development instance and Debug instance (used for heartbeat transfers/full mesh)

Rucio

- ATLAS Distributed Management Service
  - Manages accounts, files, datasets and distributed storage systems
  - Relies on FTS for 3rd party transfers
- Rucio Transfer/Network Monitoring
  - For each individual transfer, Rucio publishes event messages to the message bus
  - Raw data contain the following event types: transfer_done/fail, deletion_planned/done/fail, staging_done/fail [ref]
- ATLAS DDM dashboard service provides in-depth statistics [ref]
  - Consumes raw data and calculates statistics (with max 10 mins latency)
  - Transfer rates (throughput/volume), efficiency
  - Summaries per activity and event types
  - In addition, provides an overview of common errors
- Deployment status
  - There is an ongoing transition from DQ2 to Rucio, ATLAS DDM dashboard has production version (still on DQ2) and pre-production version (on Rucio, planned production 15th of Sept.)