Integrating network and transfer metrics to optimize transfer efficiency and experiment workflows

S. McKee[1], M. Babik[2], S. Campana[3], A. Di Girolamo[3], T. Wildish[3], J. Closier[3], S. Roiser[3], C. Grigoras[3], I. Vukotic[3], M. Salichos[3], Kaushik De[3], V. Garonne[3], J.A.D. Cruz[3], A. Forti[3], C.J. Walker[3], D. Rand[3], A. de Salvo[3], E. Mazzoni[3], I. Gable[3], F. Chollet[3], L. Caillat[3], F. Schaer[3], Hsin-Yen Chen[3], U. Tigerstedt[3], G. Duckeck[3], B. Hoeft[3], A. Petzold[3], F. Lopez[3], J. Flix[3], S. Stancu[3], J. Shade[3], M. O'Connor[3], V. Kotlyar[3] and J. Zurawski[3]

[1] Randall Laboratory, Physics Department, University of Michigan, 450 Church Street, Ann Arbor, Michigan, 48109-1040, USA

[2] European Organisation for Nuclear Research (CERN), 1211 Geneva 23, Switzerland

[3] Members of the WLCG Network and Transfer Metrics Working Group

E-mail: `Shawn.McKee@cern.ch, Marian.Babik@cern.ch, wlcg-ops-coord-metrics@cern.ch`

Abstract. The Worldwide LHC Computing Grid relies on the network as a critical part of its infrastructure and therefore needs to guarantee effective network usage and prompt detection and resolution of any network issues, including connection failures, congestion, traffic routing, etc. The WLCG Network and Transfer Metrics project aims to integrate and combine all network-related monitoring data collected by the WLCG infrastructure. This includes FTS monitoring information, monitoring data from the XRootD federation, as well as results of the perfSONAR tests. The main challenge consists of further integrating and analyzing this information in order to allow the optimizing of data transfers and workload management systems of the LHC experiments. In this contribution, we present our activity in commissioning WLCG perfSONAR network and integrating network and transfer metrics: We motivate the need for the network performance monitoring, describe the main use cases of the LHC experiments as well as status and evolution in the areas of configuration and capacity management, datastore and analytics, including integration of transfer and network metrics and operations and support.

1. Introduction

The Worldwide LHC Computing Grid (WLCG) relies on the network as a critical part of its infrastructure to inter-connect sites resources. Recently, networks have been evolving, both in terms of their speed and corresponding network traffic growth, as well as in the emergence of new technologies and paradigms. This evolution has caused rapid growth of the traffic between sites, both in terms of achievable peak transfer rates as well as in total amount of data transferred. Some of the major Research and Education Networks (R&E), such as ESnet, have seen traffic grow a factor of 10 every four years [1]. LHC experiments have adapted their computing models to benefit from this trend by introducing a more interconnected system, moving away from a strict tier-based hierarchies. As the scale and complexity of the current LHC network grows rapidly, network virtualisation technologies such as Software Defined Networking (SDN) are paving the way for performance and cost optimised networks. This has been recognized by

the LHC experiments and new projects, such as NSF funded ANSE [2], are trying to exploit the possibility to create on demand topologies (via advance network bandwidth allocations) to increase efficiency of the transfers between LHC data centres.

In this paper, we present one such project in the scope of the WLCG Operations Coordination initiative: Network and Transfer Metrics, which aims to integrate and combine all the network-related monitoring data collected by the WLCG infrastructure. This includes FTS monitoring information, monitoring data from the XRootD federations as well as results from the perfSONAR tests. We describe the main use cases driving the project as well current status of the perfSONAR network and status of its integration with other systems. Well also highlight advances made in the respective areas, such as configuration and capacity management, common datastore and analytics, integration with transfer systems, visualization as well as operations and support. We conclude with a brief description of the advanced networking systems currently being researched and developed.

2. Network Performance Monitoring

Network performance monitoring has been introduced in WLCG by a dedicated task force [3], which has established a pervasive network monitoring infrastructure based on the perfSONAR Toolkit [4]. The main goals that motivated a large scale deployment at all sites were the ability to find and isolate network problems, characterize network usage, and provide source of the network metrics to the higher level services. The choice of the open source perfSONAR toolkit was mainly driven by the consensus and prior experience of the R&E network providers that have formed the perfSONAR consortium to help develop the tools and establish a network that would allow to identify and better debug the Wide Area Network (WAN) issues. Currently, perfSONAR is deployed on over 300 domains around the world [5].

Network performance monitoring has been also deployed in parallel by several data management systems, which were focusing mainly on measuring transfer rates and data link quality from the perspective of the data transfers. It was therefore logical to join the two areas and combine metrics from both network and transfer systems. This has lead to establishing Network and Transfer Metrics working group in the fall 2014 [7]. The working group aims to enable use of network-aware tools to improve transfer efficiency and optimize experiment workflows. Among the initial objectives of the work group was the continuation of the commissioning and maintenance process of the WLCG network monitoring based on perfSONAR, but also to establish pilot projects that would facilitate the integration of the network and transfer metrics.

3. Experiments Use Cases

As part of the working group effort we have solicited desired use-cases for network and transfer metrics from the LHC experiments as well as feedback from the network related middle-ware and applications teams on what they are able to provide. This section summarizes the information we were able to gather.

In general, the core use case is to have the capability to define and understand slow transfers that are observed by the experiments. This involves identifying the weak links with a combination of different tools that provide network measurements and be able to narrow down to a particular source of the problem that can be then addressed by sites, experiment operations or network providers. The main focus so far has been on the ability to combine existing perfSONAR network measurements together with monitoring information from transfer systems such as FTS and XRootD and use them to distinguish between pure network and storage related issues. For network related problems, we benefit from the existing perfSONAR infrastructure to further debug the Wide Area Network (WAN) issues and localize the problem.

Another important aspect that has been requested is to introduce a coordinated response to the network performance problems. Since network problems can often become quite complex, defining procedure that would involve all the relevant parties and also inform all the experiments would be very beneficial.

Enabling network-aware tools is another common use case, mainly driven by the need to optimize transfers and/or experiment workflows. This involves providing a uniform way to access and integrate existing measurements and ability to define a so called distance metric between storage elements (and/or sites) that would integrate a range of different metrics such as link status, utilization, functional tests, occupancy, etc. and provide a cost matrix that can be used to decide on the job placement, finding closest replicas, determine closest storage where data can be uploaded, etc.

Finally, the experiments desire the ability to understand both new and existing network connections between sites. They want to easily commission new links by running on-demand throughput and latency tests. By also providing regular testing between all sites (full mesh) we can provide a baseline (expectation setting) for each network path and potentially replace some of the existing regular (synthetic) transfer tests that all experiments run to fill in the blank areas that have no production traffic. This also provides an opportunity to consolidate the number of existing regular testing activities that are currently run over the network.

4. perfSONAR in WLCG

One of the important sets of metrics our working groups is tasked with providing concerns measurements of the network along specific paths of interest. This is to be contrasted with end-to-end measurements (typically data transfers or data access over the WAN) that include the effects of the end-hosts and the applications involved in the process. Having network-only metrics is critical for identifying when there are issues in the network itself, which requires a very different resolution process, versus problems in the end-hosts, applications and/or their interaction with the network.

To gather network-specific metrics, we rely upon the WLCG/OSG perfSONAR deployment, that was mandated for all WLCG Tier-2 and above sites in Fall of 2013. Our working group is responsible for ensuring those metrics are consistently and correctly collected and made available for the experiments use. We rely upon the perfSONAR Toolkit to instrument our end-sites with the capability to make a standardized set of network related measurements.

Each WLCG/OSG Tier-[0/1/2] needs to provide two types of perfSONAR services: 1) latency and 2) bandwidth. Originally with perfSONAR version 3.3 and earlier, this was achieved by deploying two distinct systems, each dedicated to a specific set of latency or bandwidth measurements. With the release of version 3.4 this can also be achieved using a single host with two network interfaces, reducing the cost of providing these services. Our working group maintains a WLCG and OSG specific documentation on perfSONAR, its motivation, installation, configuration, use and troubleshooting [17].

As part the configuration instructions, each instance is registered with OIM (for OSG sites) or GOCDB (for non-OSG WLCG sites). This information can then be used to organize and manage the instances as will be described in the next section. We organize sites into groups based upon a combination of geography and virtual organization membership. Testing between instances is centrally controlled and perfSONAR provides us with 4 primary network metrics: 1) one-way delay, 2) packet-loss, 3) achievable bandwidth and 4) network path. These metrics are centrally gathered and made available for use.

5. Status and Areas of Evolution

Since inception, the working group main focus was to continue with the commissioning of the perfSONAR network and bring it to its full potential, but as well to initiate pilot projects that

would attempt to integrate existing network and transfer metrics, determine the main challenges and develop prototypes that would lead the way towards future production deployment. In this section we will describe the current status of our efforts and how we are intending to evolve moving forward.

5.1. Configuration and Testing Management

perfSONAR was designed to be a loosely federated infrastructure with no central control or management. For LHC deployments, this presented a significant problem in how to organize and manage many sites, each with typically two systems. The loosely organized model would have required us to coordinate the efforts of at least 130 WLCG system administrators: adding or removing perfSONAR instances, adding, changing or removing tests or altering groupings of instances would require communication-with and action-by many administrators before they would be effective. This just wouldnt work at the LHC scale or even for our initial deployments for USATLAS at 10 sites in 2008.
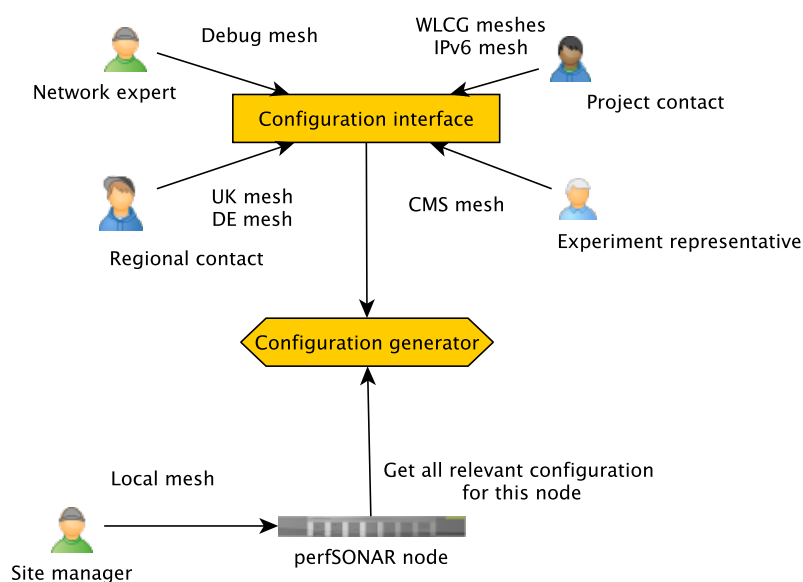
Figure 1. The WLCG/OSG perfSONAR config-management system.

An initial solution was provided by the perfSONAR developers who developed an agent based system that could reference a central configuration file stored at a specific URL. This became a part of the perfSONAR toolkit and included scripts to build the configuration JSON file from a text file. The agent on the toolkit checks the URL every day and implements any configuration changes since the last time it referred to the URL. This became referred to as the *mesh-configuration* because we typically organized tests amongst a set of sites as a full mesh of sources (rows) to destinations (columns).

While this was a significant improvement and allowed us to organize sites, tests and test parameters, it did have some drawbacks. Sites needed to configure the right set of URLs corresponding to all meshes they wanted to participate in. Central managers needed to regularly update the configuration files to include any changes: adding new sites, removing sites, renaming sites or altering test parameters. Creation of a new mesh still required that all participating sites be contacted to add a new mesh-URL to their agent configuration.

The next improvement was provided by OSG who added a mesh-configuration GUI based upon OIM and MyOSG. This system tracked the perfSONAR registration data in OIM and GOCDB and stored additional meta-data about meshes, tests and test parameters such that the JSON configuration URLs could be created and maintained by the system. Mesh administrators could change tests or groupings and site admins could change registration information and the mesh-URLs would be automatically updated to reflect any changes.

Once this system was in place it was simple to provide one more improvement that addressed the remaining issue of having sites update their agent configurations when they were added or removed from meshes: the *auto-mesh URL*. The idea was to provide a single URL, unique for each perfSONAR toolkit instance, that would provide customized JSON information needed by its agent. The URL is the same for each host except for the ending which is the fully-qualified domain name of that host. For example, the psum01.aglt2.org instance would configure its auto-mesh URL as https://myosg.grid.iu.edu/pfmesh/mine/hostname/psum01.aglt2.org

Once this is configured, the local site admin doesnt need to make any other changes to his agents configuration. Now meshes and tests can be centrally changed as needed and end-sites automatically update. The current system is shown schematically in Figure 1 above. This allows the working group to experiment with changes as required and then put them into production for everyone within a day.

5.2. Datastore

One of the responsibilities of the working group is to gather, organize and provide all the needed metrics to the experiments. Currently underway are a number of efforts to gather the various metrics and in this section we will describe those projects and their status.

For perfSONAR, the data is already being centrally collected by OSG as part of their network service and is scheduled to be production ready by the end of July, 2015. The OSG network service is designed around a few components: an *RSV master process* which parses the mesh-configs and creates *RSV probes* to query each perfSONAR toolkit instance and a *datastore*, built from the perfSONAR *Esmond* database system. OSG intends to become the network data provider for the various OSG and WLCG virtual organizations and to do this they will be gathering, storing and providing perfSONAR metrics for the foreseeable future. For FTS, each individual transfer publishes an event message that is available generally via dedicated messaging system. A common store, interface and dashboard is provided for the FTS data by the FTS dashboard [6, 10], which also performs different aggregations, such calculating transfer rates, aggregating FTS states, etc.

For XRootD, the monitoring capability is provided directly by the XRD monitoring [10], which can send the monitoring information over UDP to a set of pre-configured endpoints. The current collection and aggregation is performed by GLED [8], which supports various back-ends, including messaging, thus providing a uniform way to access the stream of XRD monitoring data. FAX and AAA dashboards provide a common store, programmatic interface and dashboard that also performs aggregations such as calculating the transfer rates [10].

5.3. Integration with Transfer Systems

Given the complexity of the integration effort, the working group is taking a pragmatic approach and decided to focus on two pilot projects, that would perform an initial research and development to understand the main challenges, develop prototypes and propose solutions.

The first project is the experiments interface to perfSONAR, which was proposed and started in collaboration with LHCb. It aims to establish a uniform way to access the perfSONAR measurements by the LHC experiments. Since both FTS and XRootD are providing their measurements via messaging, the initial idea is to establish a message bridge that would stream the data available in the centralized OSG data store, thus become key enabler for the integration

effort. An initial prototype has been developed and is currently under evaluation. One of the main challenges that become apparent during the design phase of the bridge was a need to establish a mapping between storages and sonars in order to understand how we could integrate throughput, latencies, but also link a given traceroute to a set of existing storages at a site. A separate service called proximity service has been prototyped in order to experiment on different ways how the mapping could be achieved. There are currently three different ways that are being investigated, using the existing topological services (GOCDB, OIM, AGIS) to map sonars to storages based on their assignment to a site, establishing a geo-location service to exploit geographical distance between services and finally exploit the existing traceroute data as a source for the linkage.

The second project is focused on a common FTS performance study for ATLAS, CMS and LHCb, that was initially started in ATLAS, and has shown very promising use cases for the integration with perfSONAR. The main motivation for choosing FTS was the fact that as a low level transfer service its currently used for majority of the LHC transfers and in addition its current coverage and granularity is a very good match to the available perfSONAR network. The initial goal for the project is to integrate FTS monitoring and perfSONAR traceroutes to determine points of congestion in the network and at the same time understand how we could optimize tcp streams to improve FTS perfromance. The initial results already provide interesting insights and show there is a great potential in integrating FTS and perfSONAR monitoring [14].

5.4. Visualization and Dashboards

Gathering and organizing the raw metrics related to the network and data transfers is just part of what we need to provide. A very important component is the ability to succinctly visualize this data to make it useful for the experiments and consumers of the metrics. In the case of the perfSONAR data we choose to follow the lead of the previous work done in OSG and USATLAS and incorporate MaDDash into our infrastructure. MaDDash [9] is a project sponsored by ESnet to collect, summarize and display perfSONAR metrics and we will provide some details below. For other metrics from FTS and Xrootd we are relying on the existing information systems and dashboards within WLCG and the experiments.

MaDDash has a number of nice features for displaying perfSONAR results because it was developed for the express purpose of gathering and visualizing such data. MaDDash understands mesh-configurations and can use them to determine which hosts and tests are to be displayed and then regularly query those remote host measurement archives and show the results on a web interface. Results are compared with customizable thresholds to define what is OK (green), WARNING (yellow) or CRITICAL (red). Without MaDDash it would be very difficult to understand the state of our networks without having to visit every perfSONAR toolkit web page and manually summarize the results of a particular set of measurements. Using MaDDash we can quickly understand if there are problems and in some cases where the problems likely are coming from.

5.5. Operations and Support

Because of the globally distributed and federated nature of our metric gathering infrastructure, it is very important that we have the right set of tools to monitor and manage that infrastructure. While MaDDash and our other dashboards are very useful for understanding the metrics being gathered, they are not that effective at finding problems in the systems generating those metrics. Since our working group is responsible for ensuring we are reliably gathering all the metrics continuously we needed to incorporate additional tools to help us do that.

Nagios [12] and related software has proven to be very powerful and monitoring services on a broad range of hosts. The challenge we faced was not only monitoring services on specific hosts but the results of tests that include two hosts with a multi-domain network in between. We

choose to use the Open Monitoring Distribution [11] (OMD) which packages Nagios and various other Nagios add-ins and Nagios-like components into a single RPM. The primary component we targeted to help with our monitoring task was Check_mk [13] which includes a number of nice features including service discovery, automated check configuration, automated graph creation and a highly customizable set of configuration options. OSG agreed to host both production and testbed instances of OMD for use in monitoring our global perfSONAR deployment. In addition the ATLAS Great Lakes Tier-2 (AGLT2) provided the initial prototype version of the service and continues to maintain it for WLCG use.

Because OMD uses Nagios, we are able to leverage existing Nagios plugins for our use. ESnet provided a number of perfSONAR Nagios checks that we could use to test individual services on perfSONAR toolkit nodes. In addition we wrote additional checks using either information scraped from the toolkit web interfaces or from JSON details exposed in v3.4+ of the toolkit. We also provided some scripts which parse the mesh-configs to update the OMD monitoring. As hosts are added or removed from meshes, OMD can automatically be updated this way. The last tweak we provided for WLCG use was the ability to use x509 credentials from any IGTF CA to allow guest access to the OMD/Check_mk monitoring pages.

The information in OMD is tagged with various meta-data to allow us to group hosts together by mesh (region or VO) or by host type (latency or bandwidth). Likewise we can view services by host or see the set of all OWAMP service results. Our working group has also setup various support mechanisms. We already noted the Wiki pages we created to document installation, configuration and troubleshooting of perfSONAR for WLCG. We have additionally setup a Support Unit (SU) within GGUS to address perfSONAR issues [1]. For network performance problems we have setup a separate mailing list (wlcg-network-throughput) and set of procedures documented on the working group web-site [7]. These procedures and support mechanisms are intended to evolve as we gain experience in addressing issues from the users and experiments.

5.6. Advanced Network Monitoring

One of the most requested features we have received from users and site administrators is the ability to be notified when there are network problems that those users/administrators need to do something about. This is a very challenging request in that network problems arise from network paths that may involve many entities (end-hosts, applications, local, regional and backbone networks) and are often complicated to diagnose and localize. Alerting every potential source of a problem only results in everyone learning to ignore those alerts, since, in most cases, the alert is not relevant for most of them. To try to address this, we are working closely with an OSG satellite project called PuNDIT [15], a 2-year NSF funded effort to use perfSONAR data to identify network problems and localize them. By collaborating with this effort we hope to provide the ability to accurately notify the right people, when there is good evidence of specific problems in their part of the global WLCG infrastructure.

6. Future Work

In summary, the working group has established and made progress in several areas of the WLCG network monitoring and plans to continue to evolve in the same areas also in the near term. As this paper presents a work in progress, we plan to continue our work within each area.

Configuration and capacity management has been the most mature area and there are no major developments foreseen. In order to improve on the current capacity management, we have plan to request and integrate utilization functions to be reported by the perfSONAR in the mesh configuration interface.

Storage and Analytics is currently being validated and is on track to enter production in July this year. This will be a major milestone for the working group as it will enable a central place

---

[1] https://wiki.egi.eu/wiki/GGUS:WLCG_perfSONAR_FAQ

to query, but also stream, perfSONAR measurements, which is a key enabling technology for the integration of network and transfer metrics. In addition, we plan to work with the exiting data analytics efforts to help integrate the available information to the key platforms.

Integration with transfer systems is foreseen to become the main driver of the evolution for the working group. The near-term plan is to finalize the work on the existing projects and evolve the existing prototypes to a production level. A particular interesting area of work that came out of the proximity service is to use graph databases (such as Neo4j [16]) to model the current state of the WLCG network, layer the experiments topologies on top of it as well as add network metrics that could help us better understand what is currently happening in the network.

In *visualization and dashboards* we foresee an increased activity in integrating the network measurements in the existing dashboards including, both general purpose ones such as FTS and XrootD dashboards as well as experiments specific ones such as Site Status Board [10].

In *operations and support*, we plan to finalize our work in the commissioning of the perfSONAR network, providing feedback while following up on the perfSONAR roadmap as well as improve the ways we currently test the release candidates. In addition, we foresee the need to tune the existing support channels, both perfSONAR support and recently established WLCG Network Throughput based on the feedback received from experiments and sites. We also plan to evolve the current infrastructure monitoring by adding a site alerting as well as integrate it with the current operational availability policies of the experiments. Finally, in *Advanced Network Monitoring*, work is planned on Automated APD (Adaptive Plateau Detection) on bandwidth data to identify when significant changes occur involving sites or specific paths.

Acknowledgements

References

[1] Bird I et al 2014 *Update of the Computing Models of the WLCG and the LHC Experiments* (CERN-LHCC-2014-014. LCG-TDR-002 http://cds.cern.ch/record/1695401)
[2] ANSE Advanced Network Services for Experiments, http://cern.ch/go/cSj6 2015
[3] Campana S et al 2014 *Deployment of a WLCG network monitoring infrastructure based on the perfSONAR-PS technology J. Phys.: Conf. Ser.* 513 062008
[4] Tierney B, Metzger J, Boote J, Boyd E, Brown A, Carlson R, Zekauskas M, Zurawski J, Swany M and Grigoriev M 2009 *perfSONAR: Instantiating a Global Network Measurement Framework* (4th Workshop on Real Overlays and Distributed Systems Co-located with the 22nd ACM Symposium on Operating Systems Principles)
[5] perfSONAR Deployment, http://www.perfsonar.net/about/who-is-uisng/ 2015
[6] Ayllon A A, Salichos M, Simon M K and Keeble O 2014 *FTS3: New Data Movement Service For WLCG J. Phys.: Conf. Ser.* 513 032081
[7] WLCG Network and Transfer Metrics Working Group, http://cern.ch/go/dDV72015
[8] Andreeva J et al 2014 *Monitoring of large-scale federated data storage: XRootD and beyond J. Phys.: Conf. Ser.* 513 032004
[9] MaDDash Monitoring and Debugging Dashboard, http://software.es.net/maddash/ 2015
[10] Andreeva J et al 2012 *Experiment Dashboard - a generic, scalable solution for monitoring of the LHC computing activities, distributed sites and services J. Phys.: Conf. Ser.* 396 032093
[11] OMD Open Monitoring Distribution, http://omdistro.org/ 2015
[12] Nagios, The Industry Standard in IT Infrastructure Monitoring, http://www.nagios.org/ 2012
[13] Check_MK IT monitoring solution, http://mathias-kettner.de/index.html 2015
[14] FTS Status and Performance, http://egg.bu.edu/LHC/fts-construction/index.htm 2015
[15] Batista J, Dovrolis C, Lee D and McKee S 2015 *Identifying and localizing network problems using the PuNDIT project* Journal of Physics: Conference Series CHEP2015
[16] Robinson I, Webber J and Eifrem E 2013 Graph Databases (O'Reilly Media, isbn: 1449356265, 9781449356262)
[17] WLCG and OSG perfSONAR Documentation, http://cern.ch/go/vn7d 2015