Monitoring the US ATLAS Network Infrastructure with perfSONAR-PS

For the ATLAS Collaboration, Shawn McKee[1], Andrew Lake[2], Philippe Laurens[3], Horst Severini[4], Tomasz Wlodek[5], Stephen Wolff[6] and Jason Zurawski[6]

[1] University of Michigan Physics Department
[2] Lawrence Berkeley National Laboratory
[3] Michigan State University Physics and Astronomy Department
[4] University of Oklahoma, Physics/IT
[5] Brookhaven National Laboratory
[6] Internet2

E-mail: `smckee@umich.edu`, `andy@es.net`, `laurens@pa.msu.edu`, `hs@nhn.ou.edu`, `tomw@bnl.gov`, `swolff@internet2.edu`, `zurawski@internet2.edu`

Abstract. Global scientific collaborations, such as ATLAS, continue to push the network requirements envelope. Data movement in this collaboration is routinely including the regular exchange of petabytes of datasets between the collection and analysis facilities in the coming years. These requirements place a high emphasis on networks functioning at peak efficiency and availability; the lack thereof could mean critical delays in the overall scientific progress of distributed data-intensive experiments like ATLAS.

Network operations staff routinely must deal with problems deep in the infrastructure; this may be as benign as replacing a failing piece of equipment, or as complex as dealing with a multi-domain path that is experiencing data loss. In either case, it is crucial that effective monitoring and performance analysis tools are available to ease the burden of management.

We will report on our experiences deploying and using the *perfSONAR-PS* Performance Toolkit at ATLAS sites in the United States. This software creates a dedicated monitoring server, capable of collecting and performing a wide range of passive and active network measurements. Each independent instance is managed locally, but able to federate on a global scale; enabling a full view of the network infrastructure that spans domain boundaries. This information, available through web service interfaces, can easily be retrieved to create customized applications. The US ATLAS collaboration has developed a centralized "dashboard" offering network administrators, users, and decision makers the ability to see the performance of the network at a glance. The dashboard framework includes the ability to notify users (alarm) when problems are found, thus allowing rapid response to potential problems and making *perfSONAR-PS* crucial to the operation of our distributed computing infrastructure.

1. Introduction

Scientific innovation continues to increase requirements for the computing and networking infrastructures of the world. Collaborative partners, instrumentation, storage, and processing facilities are often geographically and topologically separated, thus complicating the problems involved with data management. Global scientific collaborations, such as ATLAS[1], continue to push the network requirements envelope. Data movement in this collaboration is routinely

including the regular exchange of petabytes of datasets between the collection and analysis facilities in the coming years. This increased emphasis on the "network", now a vital resource on par with the actual scientific process, implies that it **must** be a highly capable and reliable resource to ensure success; the lack thereof could mean critical delays in the overall scientific progress of distributed data-intensive experiments.

Network operators routinely monitor the health and overall capabilities of their own networks. This can be done in an ad-hoc manner by simply responding to threats and complications when needed. Others may take a more structured approach by using monitoring tools to ensure certain success metrics and objectives are being met. Often the "end-to-end" nature of a network problem, e.g. the well known fact that scientific collaboration does not end at the network boundary, is ignored due to the overall complexity of the problem as well as the limitations on what a given organization is able to handle internally. Operations staff routinely must deal with problems deep in the infrastructure; this may be as benign as replacing a failing piece of equipment, or as complex as dealing with a multi-domain path that is experiencing data loss. In either case, it is crucial that effective monitoring and performance analysis tools are available to ease the burden of management, and that they are capable of extending usefulness well beyond the campus boundary.

The *pS Performance Toolkit*[2], an "all in one" network performance monitoring solution developed by the *perfSONAR-PS* project[3], aims to ease the burden of end-to-end performance monitoring. This software creates a dedicated monitoring server, capable of collecting and performing a wide range of passive and active network measurements as well as providing an archive of those measurements. The goal of this software is to "lower the bar" in terms of setting up performance testing between facilities through an easy to deploy framework. Each independent instance is managed locally, but able to federate on a global scale; enabling a full view of the network infrastructure that spans domain boundaries. Metrics including network utilization, available bandwidth, point to point latency, packet loss and connection stability can be gathered through this system. Deploying the software and establishing tests is the first step; interpreting the results through a user-focused GUI comes next. This information, available through web service interfaces, can easily be retrieved to create customized applications.

We will report on our experiences deploying and using the *pS Performance Toolkit* at ATLAS sites in the United States. US ATLAS[1] has developed a centralized "dashboard" offering network administrators, users, and decision makers the ability to see the performance of the network at a glance. The dashboard framework includes the ability to notify users (alarm) when problems are found, thus allowing rapid response to potential problems and making *perfSONAR-PS* crucial to the operation of our distributed computing infrastructure.

The remainder of the paper will proceed as follows. Section 2 will discuss the ATLAS collaboration, as well as data movement requirements and expectations. Section 3 will discuss the pS Performance Toolkit, a software package designed to ease the burden of multi-domain network monitoring. Section 4 will discuss efforts to create a single point within the collaboration that displays the status of the network and is capable of alerting partners when problems are observed. Section 5 will cover our experience deploying, configuring and using the perfSONAR-PS toolkit along with the dashboard. Section 6 will discuss future plans for this work and present ways it may be adapted to other Virtual Organizations (VOs) and projects. Section 7 and Section 8 conclude the paper.

2. ATLAS Collaboration

The ATLAS collaboration consists of over 3000 physicists and 1000 students from 38 countries and 178 Universities and Laboratories worldwide. This large group of scientists is working

---

[1] US ATLAS is the subset of the ATLAS collaboration at institutions within the United States and has its own management structure and facilities

together at the Large Hadron Collider[4] (LHC) to learn about the basic forces that have shaped our Universe since the beginning of time and that will determine its fate. ATLAS physicists are exploring the frontiers of high-energy physics, explaining the origin of mass, probing the existence of extra dimensions, searching for microscopic black holes and looking for evidence for dark matter candidates in the Universe, (just to name a few topics of interest).

To undertake these explorations, the collaboration has constructed the ATLAS detector[1] over a 15 year period and assembled it at Point 1 in the LHC ring. The detector is 45 meters long, 25 meters high and weighs about 7000 tons. The ATLAS detector employs a number of types of sub-detectors to measure attributes of the various particles resulting from the collision of counter-rotating beams of protons in the LHC. There are millions of electronic channels associated with the readout of the ATLAS detector. In effect, the set of all of these sub-detectors and associated readouts can be viewed as a very large 3-dimensional digital camera, capable of taking precise "pictures" 40 million times a second (proton beam bunches cross one another every 25 nanoseconds). The detailed information collected allows the ATLAS physicists to reconstruct the underlying event and search for new physics.

If all the data ATLAS produces could be stored, it would fill more than 100,000 CDs per second, a rate (and corresponding data-volume) which is not possible to support. Instead a set of hardware, firmware and software systems makes fast decisions about what data is interesting to keep and result in a data rate of 400 MBytes/sec into "offline" disk storage. Even so, this rate of data production results in many petabytes of data being produced by ATLAS each year. In addition, detailed simulations also produce Petabytes of data required to understand how the ATLAS detector responds to various types of events and validate that the ATLAS software works as expected. It is important to note that these large data volumes are common to all the LHC experiments and not just ATLAS.

Because of the data-intensive nature of the ATLAS scientific program, the ATLAS collaboration implicitly relies upon having a ubiquitous, high-performing, global network to enable its distributed grid-computing infrastructure. Providing effective access to petabytes of data for thousands of physicists all over the world just wouldn't be possible without the corresponding set of research and education networks that provide 1 to 10 to 100 Gigabits per second of bandwidth to enable ATLAS data to flow to where it is needed. Typical network paths that ATLAS data traverses consist of multiple administrative domains (local area networks at each end and possibly multiple campus, regional, national and international networks along the path). The ability of the Internet to allow these separate domains to transparently inter-operate is one of its greatest strengths. However, when a problem involving the network arises, that same transparency can make it very difficult to find the cause and location of the problem.

Because of both the criticality of the network for ATLAS normal operations and the difficulty in identifying and locating the source of network problems when they occur, the US ATLAS facility began deploying and configuring perfSONAR-PS in 2008. Our goal was to provide our sites with a set of tools and measurements that would allow them to differentiate network issues from end-site issues and to help localize and identify network specific problems to expedite their resolution.

3. Performance Measurement with the perfSONAR-PS Toolkit

perfSONAR is a framework that enables network performance information to be gathered and exchanged in a multi-domain, federated environment. The goal of *perfSONAR* is to enable ubiquitous gathering and sharing of this performance information to simplify management of advanced networks, facilitate cross-domain troubleshooting and to allow next-generation applications to tailor their execution to the state of the network. This system has been designed to accommodate easy extensibility for new network metrics, to archive performance information for later reference and to facilitate the automatic processing of these metrics as much as possible.

perfSONAR is a joint project started by several national R&E networks and other interested partners. The aim of this project is to create an interoperable framework which allows network metrics to be gathered and exchanged in a multi-domain, heterogeneous, federated manner. *perfSONAR* is targeting a wide range of use cases. For example current use cases include: collection and publication of latency data, collection and publication of achievable bandwidth results, collection and archiving of packet loss, publication of utilization data, publication of network topology data, diagnosing performance issues, and several others. While *perfSONAR* is currently focused on publication of network metrics, it is designed to be flexible enough to handle new metrics from technologies such as middleware or host monitoring. Previous papers on *perfSONAR* have described the original overall architecture[5], the data model and schemata[6], and the information and location services [7].

To ease the burden of deployment, the *perfSONAR-PS* project packaged select measurement tools in a CentOS Linux [8] bootable CD format, called the *pS Performance Toolkit*. Sites only need to insert the distribution CD, boot up the host, and answer a few configuration questions to have a completely operational measurement point. Additional feedback from the operations community, including the US ATLAS project, led to the construction of a version capable of persisting on the target machine's hard disk. The advantages of this disk-based install are that it allows easy maintenance via the *yum*[9] system, and eases the burden of patching or customizing for a given use case.

The US ATLAS project[10], an early adopter of this solution, has operated a dedicated monitoring infrastructure since the fall of 2008. The Tier 1 and Tier 2[2] members of the US ATLAS project (currently 10 distinct locations[3]) coordinate testing of the following regularly scheduled network measurements:

- **Available Bandwidth:** a measure of how much bandwidth is available to a well tuned host and application, delivered via the BWCTL[11] tool.
- **Round Trip Latency:** a measure of latency statistics between two points, collected via the PingER [12] tool.
- **One Way Latency:** a measure of latency statistics, separated into independent directions, and collected via the OWAMP [13] tool. A critical attribute of this measurement is the amount of packet loss along the path.
- **Traceroute:** a calculation of the Layer 3 (e.g. the "network" layer) path traveled, using the traceroute tool.

Best practice is for each site to maintain two monitoring hosts, one for latency services, and one for bandwidth services, as bandwidth testing adversely affects latency tests. Experience has shown this separation valuable, as each form of measurement can be run in a simultaneous fashion and provide accurate observations about the network, without having to be concerned with host specific behaviors or self-interference. Figure 1 demonstrates this deployment on a per-site basis.

Implemented as a web services-based infrastructure for collecting and publishing network performance monitoring, a primary goal of *perfSONAR* is making it easier to solve end-to-end performance problems on paths crossing several networks. The software contains a set of services delivering performance measurements in a federated environment. These services

[2] Worldwide, ATLAS has organized its computing facilities into a hierarchical "tiered" infrastructure, with national-scale facilities designated Tier 1 and regional-scale facilities designated Tier 2.

[3] The Tier 1 at Brookhaven National Laboratory and the five Tier 2 centers (some Tier 2s span more than one site) at Boston University, Indiana University, Michigan State University, Oklahoma University, SLAC National Accelerator Laboratory, University of Chicago, University of Illinois at Urbana-Champaign, University of Michigan, University of Texas at Arlington

Figure 1. An example pS performance deployment at 3 sites showing the inter-site testing and central collection of test metrics.

act as an intermediate layer, between the performance measurement tools and the diagnostic or visualization applications. This layer is aimed at making and exchanging performance measurements across multiple networks and multiple user communities, using well-defined protocols. Section 4 describes an ongoing project by members of the US ATLAS collaboration to collect these measurements and present them in an easy to read an interpret fashion.

4. Monitoring Dashboard

As US ATLAS began to deploy perfSONAR-PS instances, we found the federated design architecture to have some shortcomings for our intended use-case. Each site was independently installed, configured and controlled and it was difficult to see the status of the sites or the inter-site measurements without visiting each site and viewing multiple graphs. This was time-consuming and we often found some tests or services had hung or stopped responding.

In order to provide a high-level summary of the site test status and to visualize the results of our *perfSONAR* measurements, a dedicated monitoring system was proposed and developed. The architecture of the system consists of multiple functional components. The first is a set of *collectors* responsible for gathering monitoring information. The results from the collectors are stored in a *data store* component. Information from the data store is presented to users via a *data presentation web interface*. Finally, new monitoring jobs and alerting is defined through a *configuration web interface*. This section outlines the details of these components.

The collectors are a set of programs running on separate monitoring nodes and a central data repository. The collectors periodically query the data store and are issued instructions about tests (also referred to as *probes*) to be executed. After executing the required probes, the

collectors upload the results back to the central data store. Since the collectors are independent of each other, they can run on machines belonging to separate sub-networks, allowing them to probe different network domains from within firewall boundaries. Also, the fact that there can be multiple data collectors allows us to distribute the load associated with executing the monitoring probes to more than one system.

The central data store consists of a web application server (currently Apache Tomcat) and a back-end database (currently MySQL). The server interacts with collectors by accepting connections, posting instructions on probes to run, and storing uploaded probe results. It also accepts connections from a web browser and displays a summary of data in human readable format.

RACF
Grid Group

The Experimental Independent perfSONAR Dashboard

Status as of: Thu Jun 21 11:31:03 EDT 2012

US cloud throughput measurement

---	0	1	2	3	4	5	6	7	8	9
0:BNL (lhcmon.bnl.gov)	---	0.84	0.40	0.81	0.70	0.82	0.88	0.56	0.00	0.77
	---	0.68	0.84	0.87	0.00	0.57	0.86	0.70	0.21	0.48
1:AGLT2 (psmsu02.aglt2.org)	0.92	---	0.94	0.84	0.89	0.94	0.89	0.93	0.73	0.90
	0.88	---	0.94	0.92	0.93	0.94	0.92	0.78	0.92	0.65
2:AGLT2 (psum02.aglt2.org)	0.70	0.94	---	1.85	2.38	3.42	0.00	0.21	0.08	0.26
	0.30	0.94	---	1.75	2.56	0.00	0.16	0.17	0.13	0.27
3:MWT2 (iut2-net2.iu.edu)	0.67	0.41	0.00	---	3.09	5.17	0.15	0.12	0.08	0.29
	0.54	0.36	1.74	---	2.71	4.20	0.15	0.15	0.05	0.27
4:MWT2 (uct2-net2.uchicago.edu)	0.40	0.70	0.55	1.43	---	1.66	0.16	0.28	0.13	0.32
	0.33	0.71	0.57	1.55	---	1.92	0.12	0.19	0.10	0.26
5:MWT2 (mwt2-ps02.campuscluster.illinois.edu)	0.00	0.00	2.91	2.76	2.16	---	0.92	0.93	0.00	0.90
	0.89	0.94	2.37	0.00	2.53	---	0.00	0.85	0.00	0.00
6:NET2 (atlas-npt2.bu.edu)	0.84	0.80	0.78	0.79	0.60	0.00	---	0.00	0.79	0.75
	0.72	0.48	0.82	0.67	0.78	0.80	---	0.58	0.53	0.52
7:SWT2 (ps2.ochep.ou.edu)	0.88	0.92	0.92	0.91	0.81	0.93	0.91	---	0.91	0.91
	0.85	0.93	0.92	0.91	0.83	0.00	0.91	---	0.91	0.73
8:SWT2 (netmon2.atlas-swt2.org)	0.85	0.92	0.92	0.91	0.76	0.00	0.91	0.93	---	0.89
	0.88	0.91	0.92	0.91	0.92	0.93	0.91	0.93	---	0.62
9:WT2 (psnr-bw01.slac.stanford.edu)	0.54	0.77	0.75	0.71	0.82	0.00	0.60	0.79	0.60	---
	0.82	0.00	0.89	0.88	0.76	0.87	0.86	0.89	0.75	---

(c) 2010 Brookhaven National Laboratory - send suggestions and comments to tomw@bnl.gov

Figure 2. A snapshot of the perfSONAR-PS dashboard illustrating the throughput matrix between the ten Tier 1 and Tier 2 sites in US ATLAS.

A typical data presentation page is shown in figure 2[14]. The table shown represents the throughput measurements between sites in the US ATLAS cloud and should be interpreted as follows: each row in the table represents a site which serves as source of the measurement while each column represents its destination. Thus the fields of tables correspond to unique source-destination pair. In addition, each measurement between a source-destination pair is performed twice, i.e. both by the source node and by the destination node, with the corresponding results displayed in the upper and lower number in the table field, respectively. The status of the measurement is color coded: green means OK, yellow means warning, red means critical, brown means unknown, and gray means probe timeout. The definitions for "OK/warning/critical" are adjustable depending upon the needs of the users. Other pages allow the users to view plots with history of individual probes.

The configuration of the dashboard is performed via a dedicated graphical user interface (GUI) which allows a dashboard administrator to define hosts, sites, probes, measurement matrices and groups of sites (clouds). The authentication and authorization of users is performed using HTTPS client authentication. The client certificate must be signed by a certificate

authority which is known and trusted by the server. The list of trusted authorities is currently configurable by the server administrator. A dedicated GUI for user management is also provided. Through this interface it is possible to define alerts which will notify selected users by e-mail in the case that a site service fails to respond or returns errors.

5. US ATLAS Experiences with *perfSONAR-PS*

The combination of a centralized dashboard for summarizing and monitoring the US ATLAS perfSONAR-PS installations along with the perfSONAR-PS toolkit allowed us to make an effective system for monitoring our inter-site networks. In this section we describe how we configured our setup and some experiences we had in deploying and using it.

US ATLAS uses the following configuration guidelines for a selection of network measurements. Note that all tests described below are configured as a "full-mesh", e.g. each host performs tests to all other hosts in the US ATLAS "cloud".

- **Available Bandwidth:** The BWCTL tool will invoke an iperf test to all other hosts once every 4 hours. This test will run for 20 seconds, and use the TCP protocol. TCP window size is adjusted automatically (i.e. via "autotuning").
- **Round Trip Latency:** The PingER tool will run a test every 5 minutes, and send 10 total packets spaced a full second apart. These packets are 1000 bytes in size, and are ICMP based.
- **One Way Latency:** The OWAMP tool will send a continuous stream of UDP packets, each 20 bytes in size, at a rate of 10 packets per second. Each packet is time stamped by the source and destination to determine the latency and carefully accounted for to measure the rate of packet loss.
- **Traceroute:** The traceroute tool will be invoked every 10 minutes, and will send 40 byte UDP packets during the test.

For the US ATLAS case of 10 distinct sites in the cloud, the average bandwidth consumed by the latency tests is about 6.6 KBytes/sec and (assuming we achieve 925 Mbits/sec and test both ways from each end) about 46.2 Mbits/sec for the bandwidth tests.

This configuration has been in place since the initial setup of the measurement framework. The combination of lightweight latency tools with the heavier bandwidth based measurement (i.e. a simulation of a well tuned data movement application) has resulted in the discovery of several serious performance abnormalities between members of the US ATLAS collaboration. Figure 3 shows observed latency measurements between two US ATLAS sites (the University of Oklahoma and Indiana University). The observations can be broken into two categories:

- **Asymmetric Latency/Routing:** Different paths were being taken through the network infrastructure in each direction, depending on where the traffic originated.
- **Symmetric Latency/Routing:** A correction was made by network operations to prefer a specific route, thus returning latency to a symmetric pattern.

The observations of Figure 3 **do not** necessarily imply a serious network problem. Asymmetry can be present in routing protocols and will not adversely affect traffic by itself (other than adding additional latency to a path). Problems arise when one of the paths **does** contain a performance problem (e.g. packet loss, congestion, etc.). Figure 4 shows a snapshot of bandwidth observations between the same two facilities. The performance returns to a symmetric behavior after the routing change, thus implying that one of the paths contained a problem in need of further investigation.

Feedback from the US ATLAS collaboration has strengthened the overall operation of the *pS Performance Toolkit*. Suggestions to implement additional metrics (e.g. traceroute) as well as
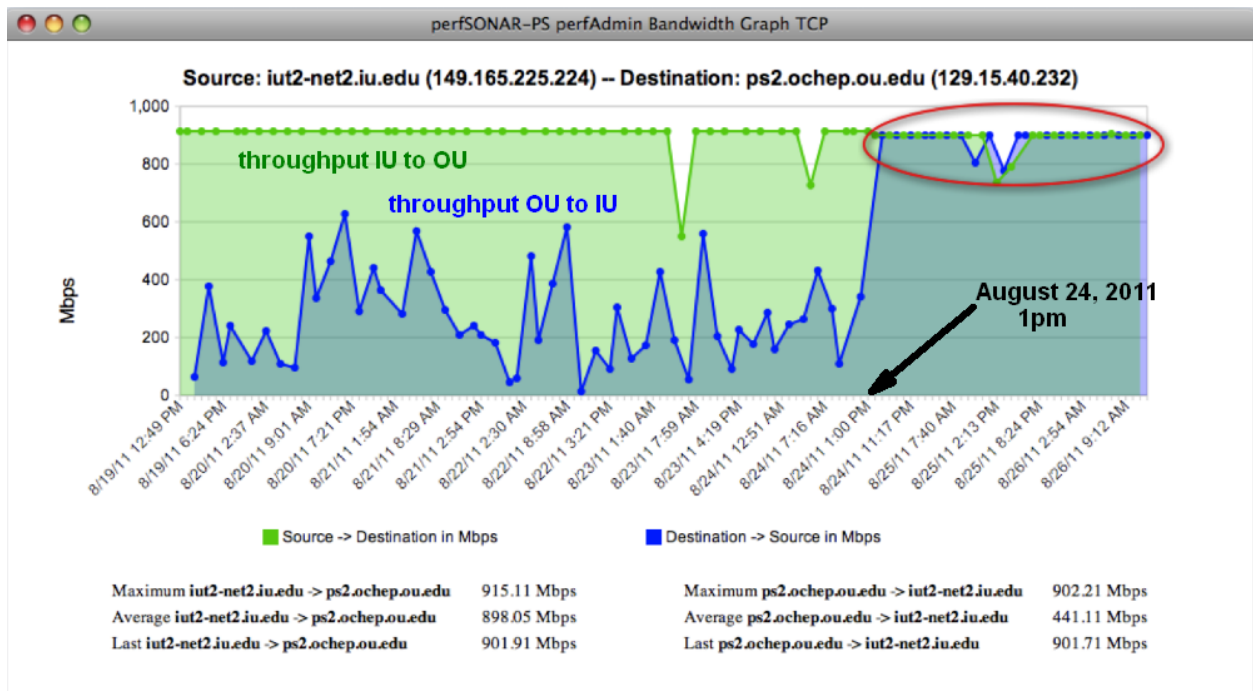
Figure 3. Shown is a plot of the one-way delay (in ms) between the Indiana and Oklahoma instances. The routing for the path from Oklahoma to Indiana was changed about the time indicated.

background improvements to facilitate easier maintenance over time (e.g. data rotation scripts) have been incorporated into the development life-cycle of the project. While not every new VO will have the time or patience to experiment with a software package in this manner, the experience has also forced change within the facilities management of ATLAS. Often performance measurements are observed in concert with the results from data transfer tools, and serve as the "evidence" that is provided to network operators around the world to investigate performance reports.

In our experience, both the dashboard and the perfSONAR-PS toolkit have worked well for US ATLAS. The dashboard has evolved based upon our needs and experiences and the perfSONAR-PS developers have incorporated our feedback as they move forward. However we still have some issues that will need addressing in future work:

- The dashboard is often slow to display results.
- The web interface on each perfSONAR-PS instance can be slow to render graphs.
- The dashboard collector (partially because of the large-scale ramp-up of other communities) is often overloaded and unable to query all needed metrics from all sites
- Cleanup and maintenance of the perfSONAR-PS installations is mostly automated as installed, but still requires additional work to ensure more robust operation at the scale used by US ATLAS.
- The LookupService (hLS) in perfSONAR-PS often becomes non-responsive
- The current alarm/email capability only checks "primitive" service status to make sure a particular service is responsive. There is no "higher-level" alarming on the measurement results themselves

Figure 4. Shown is a plot of the achieved bandwidth (in Mbps) between the Indiana and Oklahoma instances in each direction. The routing for the path from Oklahoma to Indiana was changed near the time indicated, resulting in significantly improved throughput.

- Bandwidth tests, while scheduled to run in every 4 hour window, may often fail to run. There is no easy way to see this or flag the reason for the test not being run.
- Some services experience reliability problems when used for long duration and with many sites under test. This is frequently related to insufficient clean-up or tuning for the particular use-case.
- Graphs from perfSONAR-PS require careful interpretation to ensure you understand what source and destination are being shown.

Even with the shortcomings noted above, the combination of dashboard and perfSONAR-PS toolkit has proven to be an indispensable tool. The success of using these in US ATLAS led to the quick adoption of these tools for the LHCOPN[15] community. After the system was demonstrated at the June 2011 LHCOPN meeting it was voted to become the network monitoring framework for LHCOPN and was substantially in place at all 11 Tier-1 sites by September 2011. It has also been chosen to provide the transitional monitoring for the LHCONE[16] project and details are available at `https://twiki.cern.ch/twiki/bin/view/LHCONE/SiteList`.

6. Future Work

Several changes to the monitoring infrastructure are planned based on experiences gathered while running the system and watching its expansion over time. Some of these changes will need to be made to the software developed specifically for the monitoring dashboard system; other changes will need to be integrated into related projects like the *pS Performance Toolkit* and will be leveraged by display tools such as the dashboard. These changes, and the projects they impact, are detailed in the following subsections.

6.1. Modular Design

There is an intention to split the monitoring dashboard into several inter-operating modules, in an effort to allow the system to be deployed on several physical machines. This will help mitigate the risk of increased system load from one component negatively, and impacting the performance of others. It also allows one component to be enhanced or upgraded without affecting the others. In this way, the modularization of the dashboard will allow for more efficient deployment and development.

6.2. Data Storage

In order to improve response time, changes will be made to rely on storage of network status data entirely in system memory, and not in the back-end database. The on disk database will be used solely for storing persistent configuration data, as well as results of historical measurements. This change is motivated by the observation that as the size of the monitoring systems grows, so does the overhead associated with accessing data from back end database, and therefore leading to noticeable degradation of performance.

6.3. Alarming Subsystem

As noted in the previous section there is additional work required to improve the functionality required for generating alarms in the current infrastructure. The alerting system will be expanded to allow for more sophisticated definitions of alarm conditions. For example, by comparing short and long term changes of status for selected probes, it is possible to determine the true network status in a more accurate fashion. The Adaptive Plateau Detection scheme is one possibility that we have been testing with the Ohio Supercomputer team[17]. The desire for this change is to increase the quality of alarms, and prevent things like scheduled maintenance from falsely reporting a problem with the infrastructure. Making alarms accurate and indicative of a real problem is critical in ensuring that human administrators do not ignore alarms.

6.4. Application Programming Interface

To facilitate data retrieval, there are plans to add a data access API to the system, which will give other computing systems access to the results collected by the probes. This allows the data presentation to be extended beyond the currently implemented display interface, and will also allow external sites to easily integrate *perfSONAR-PS* measurements into their internal alarming or monitoring systems.

6.5. Measurement Server Configuration Management

A centralized configuration management system is planned, where a site can determine what tests it should run. At present, an addition of a new site to the monitoring matrix requires changing the configuration of each site that belongs to the current "cloud". In the long-term, this approach will not scale, and the need for centralized configuration management is apparent. This is one area where future changes to the *pS Performance Toolkit* may be helpful. The *pS Performance Toolkit* provides the software that currently generates and manages performance testing configuration. This configuration currently resides in a file on disk that is not publicly accessible. The *perfSONAR-PS* project is currently exploring how configuration information can be dynamically shared and updated within communities such as ATLAS. Future changes such as this should make managing the dashboard probes much simpler.

6.6. Open Science Grid

In addition, there are plans to work with the Open Science Grid[18](OSG) on integrating the monitoring and measurement tools into the MyOSG portal. This step will ensure that multiple

VOs for a given site (e.g. a campus that houses physics, biology, and chemistry research) have access to accurate performance measurement statistics for common infrastructure.

7. Conclusion

To support the growing scale of global scientific projects, network operators and VOs alike must be cognizant of network performance considerations to assure proper operation. Frameworks, such as the *pS Performance Toolkit*, are capable of monitoring internal and external network performance metrics when properly deployed and managed. Presentation layers, such as the US ATLAS dashboard, can deliver the raw results of this performance assurance in an easy to use and interpret format. This holistic approach to network measurement has resulted in the correction of numerous performance abnormalities, and saved the time and resources of strained operations staff.

8. Acknowledgements

References

[1] The ATLAS Collaboration 2008 JINST 3 S08003 `http://iopscience.iop.org/1748-0221/3/08/S08003`
[2] The pS Performance Toolkit `http://psps.perfsonar.net/toolkit`
[3] The perfSONAR-PS Project `http://psps.perfsonar.net`
[4] L. Evans and P. Bryant (editors), JINST 3 (2008) S08001 `http://iopscience.iop.org/1748-0221/3/08/S08001/`
[5] A. Hanemann, J. Boote, E. Boyd, J. Durand, L. Kudarimoti, R. Lapacz, M. Swany, S. Trocha, and J. Zurawski. Perfsonar: A service oriented architecture for multi-domain network monitoring. In *Third International Conference on Service Oriented Computing - ICSOC 2005, LNCS 3826, Springer Verlag*, pages 241–254, Amsterdam, The Netherlands, December 2005.
[6] J. Zurawski, M. Swany, and D. Gunter. A scalable framework for representation and exchange of network measurements. In *IEEE/Create-Net Conference on Testbeds and Research Infrastructures for the Development of Networks and Communities*, Barcelona, Spain, March 2006.
[7] J. Zurawski, J. Boote, E. Boyd, M. Glowiak, A. Hanemann, M. Swany, and S. Trocha. Hierarchically federated registration and lookup within the perfSONAR framework (short paper, poster session). In *Tenth IFIP/IEEE International Symposium on Integrated Network Management (IM 2007)*, 2007.
[8] The Community ENTerprise Operating System `http://centos.org/`
[9] The Yellowdog Updater, Modified `http://en.wikipedia.org/wiki/Yellowdog_Updater,_Modified`
[10] The US ATLAS organizational page `http://www.usatlas.bnl.gov/USATLAS_TEST/Organization.shtml`
[11] Bandwidth Test Controller (BWCTL) `http://www.internet2.edu/performance/bwctl/`
[12] PingER (Ping End-to-end Reporting) `http://www-iepm.slac.stanford.edu/pinger/`
[13] One-Way Ping (OWAMP) `http://www.internet2.edu/performance/owamp/`
[14] The modular dashboard for US ATLAS is available at `https://perfsonar.usatlas.bnl.gov:8443/exda`
[15] The LHC Optical Private Network (LHCOPN) information is found at `http://lhcopn.web.cern.ch/lhcopn/`
[16] The LHC Open Network Environment is documented at `http://lhcone.net/`
[17] Prasad Calyam, Jialu Pu, Weiping Mandrawa, Ashok Krishnamurthy. OnTimeDetect: Dynamic Network Anomaly Notification in perfSONAR Deployments, `http://www.osc.edu/~pcalyam/ontimedetect_mascots10.pdf`
[18] The Open Science Grid `http://iopscience.iop.org/1742-6596/78/1/012057`