Deploying distributed network monitoring mesh for LHC Tier-1 and Tier-2 sites

Jeff Boote[1], Eric Boyd[1], Aaron Brown[1], Maxim Grigoriev[2], Joe Metzger[3], Phil DeMar[2], Martin Swany[4], Brian Tierney[3], Matt Zekauskas[1] and Jason Zurawski[1]

[1] Internet2, 1000 Oakbrook Drive, Suite 300, Ann Arbor MI 48104, USA
[2] Fermilab, PO BOX 500, Batavia, IL 60510, USA
[3] ESnet, Lawrence Berkeley National Lab, 1 Cyclotron Rd, Berkeley, CA 94720, USA
[4] University of Delaware, Newark, DE 19716, USA

E-mail: [1] {aaron, boote, eboyd, matt, zurawski}@internet2.edu

E-mail: [2] {maxim, demar}@fnal.gov

E-mail: [3] {metzger, bltierney}@es.net

E-mail: [4] swany@cis.udel.edu

Abstract.
Fermilab hosts the US Tier-1 center for data storage and analysis of the Large Hadron Collider's (LHC) Compact Muon Solenoid (CMS) experiment. To satisfy operational requirements for the LHC networking model, the networking group at Fermilab, in collaboration with Internet2 and ESnet, is participating in the perfSONAR-PS project. This collaboration has created a collection of network monitoring services targeted at providing continuous network performance measurements across wide-area distributed computing environments. The perfSONAR-PS services are packaged as a bundle, and include a bootable disk capability. We have started on a deployment plan consisting of a decentralized mesh of these network monitoring services at US LHC Tier-1 and Tier-2 sites. The initial deployment will cover all Tier-1 and Tier2 sites of US ATLAS and US CMS. This paper will outline the basic architecture of each network monitoring service. Service discovery model, interoperability, and basic protocols will be presented. The principal deployment model and available packaging options will be detailed. The current state of deployment and availability of higher level user interfaces and analysis tools will be also be demonstrated.

1. Introduction

The computing facilities at Fermilab will be playing a critical role in the CMS [31] experiment at LHC. Fermilab is hosting a USCMS Tier-1 center for data storage and analysis. The same scale Tier-1 US ATLAS facility has been established at DOE managed Brookhaven National Lab and will be a centerpiece of physics data storage and analysis for ATLAS [1] experiment at the LHC. The LHC Tier-2 centers in the US are connected over DOE sponsored-ESnet and NSF-sponsored Internet2 networks and are moving terabytes of data every day from and to the Tier-1 center at Fermilab. At Fermilab alone there are currently two routed IP 10 gigabit/second channels and four ESnet Science Data Network (SDN) with a dynamic circuit reservation system based on the Lambda Station project. With present multi-gigabit rates and projected 100 gigabit physics data transfer rates, it's hard to overstate the significance of a robust high performance network.

To satisfy operational requirements for the LHC networking model, the networking group at Fermilab, in collaboration with Internet2 and ESnet, is participating in a multi-domain network monitoring project called *perfSONAR* . The deliverable for this project is a collection of the network monitoring services and protocols named perfSONAR-PS [9]. *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 and to facilitate the automatic processing of these metrics as much as possible.

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, 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.

We envision a number of future, higher-level services that will be able to aggregate the *perfSONAR* data from multiple sources and utilize it as part of the own workflow. For example, data transfer middleware could use *perfSONAR* to locate the best replica/copy of a file to request, or to help determine the optimal network protocol to use for a given link. Network engineers could use *perfSONAR* to help automate the detection of large bulk data flows that may require special handling, such as tagging the flow as high- or low-priority, depending on its source or destination. Finally, network researchers will find *perfSONAR*-enabled networks a convenient source of performance and topology information.

Another use case scenario may utilize network monitoring mesh built on *perfSONAR* services as first step in the network troubleshooting workflow with intent to separate network related problem from the end-host ones. There is a possibility of deploying on-demand monitoring end-to-end system on any network segment with an additional requirement on having virtualization layer installed on the end-host, for example [34].

The proliferation of dynamic virtual network circuits [7, 12, 24] have set new requirements for the close to real time multi-domain network monitoring and ad-hoc availability of the complete monitoring toolkit [22] deployed on any arbitrary segment of the network topology.

The topic of this paper is the large-scale deployment of *perfSONAR* for the LHC community, how it has been used thus far, and how it is intended be used in the future. The contribution of this paper is to demonstrate the practical value of the *perfSONAR* approach. We also hope to *inspire others to contribute to the effort* by building network-aware middleware and analysis applications on top of *perfSONAR*, and to help us find solutions to the security and privacy issues for this type of distributed decentralized system.

1.1. LHC Use of perfSONAR

Much of the current *perfSONAR* effort targets the immediate needs of the Large Hadron Collider (LHC) community. The LHC, located at CERN near Geneva Switzerland, will soon be generating about 100 Terabytes of data per day. The LHC data distribution model is a multi-tiered where data source is called "Tier-0" and the first level processing and storage is called "Tier-1." There are 11 Tier-1 sites; each site is expected to handle and store about one Petabyte of raw data per month. The 140 "Tier-2" sites are based at research universities and other scientific facilities and will play the major role in data analysis. There will be continuous exchange of high volumes of physics data between various Tier-1 and Tier-2 centers because Tier-1 centers are playing a "data-hub" role and data will be replicated among several Tier-1 sites. The expected wide area data rates into and out of the Tier-1 sites will be at least 20 Gbps, so this traffic will be segregated from the general Internet whenever possible, and the ability to collect both active and passive monitoring data is important. Although network circuits between Tier-0 and Tier-1 sites are built on a very well provisioned private optical network, called LHCOPN [14], the mesh of network connections between Tier-1 and Tier-2 sites might have frequent problems with connectivity and data transfer performance.

To make it easy for LHC sites to deploy, we have packaged perfSONAR-PS tools in a Knoppix-based [11] bootable CD, called the pS-NPToolkit. Sites only need to insert the pS-NPToolkit CD, boot up the host, and answer a few configuration questions to have an operational measurement point. The US-ATLAS part of LHC has deployed pS-NPToolkit hosts at 15 phisical locations. Each site will be running two monitoring hosts, one for latency services, and one for bandwidth services, as bandwidth testing adversely affects latency tests. These services are described in more detail below. US-ATLAS is planning to use *perfSONAR* to help monitor its circuits, and to provide realistic bandwidth expectations to its users. For US-CMS collaboration of the LHC, the plan is to deploy *perfSONAR*-based monitoring at Tier-1 sites and start deployment for the Tier-2 sites by the middle of year 2009.

2. perfSONAR Architecture

2.1. Overview

perfSONAR is an example of a Service Oriented Architecture (SOA), which offers the ability for specialized, autonomous services to join under a common access scheme. Thus, it is possible to separate the roles of monitoring, storage, processing, and visualization of data into specialized service instances.

The different *perfSONAR* components are implemented using Web Services (WS) technology. Therefore the interaction between services and between clients and services are performed using well defined language independent interfaces. All WS interfaces are defined using eXtensible Markup Language (XML). *perfSONAR* web services furthermore extend an XML schema defined by the Open Grid Forum (OGF) [23] Network Measurement Working Group NM-WG [21]. These schemas are used to provide a uniform encoding and exchange mechanism for performance information within and across administrative domains. This vastly simplifies system component interactions as well as storage requirements. *perfSONAR* also utilizes the network characteristic taxonomy defined by the OGF NW-WG [16].

We will be presenting an abbreviated version of the full *perfSONAR* architecture [9] since the full architecture is not needed for the LHC deployment. The core components of the *perfSONAR* architecture used in this case are the data producers (Measurement Point (MP), and Measurement Archive (MA) services), data consumers (Analysis clients), and discovery (Information Services (IS)). The MPs and MAs are responsible for exposing performance metrics, and in the MA case in potentially storing metrics for later retrieval. The IS is responsible for helping clients find available services and even finding relationships between specific network topology elements.

2.2. Topology

A key aspect of *perfSONAR* is the fact that the data model for measurements is based on the OGF-NMWG schema, which directly relates performance metrics to network topology. This relationship allows data produced by different services to be related to each other because the multiple metrics are related to the same network topology elements. This is particularly powerful when combining passive information about a path with active measurements as we show in section 5.

The fact that all data can be related to network topology also allows the cloud of IS services to work efficently as a globally distributed database of information concerning what services, performance metrics, and network topology components exist and how that information relates to the other information. Data distribution and locality of information is determined based on the relationship of the information to the underlying network topology. Additionally, this allows local policy controls to be put in place determining how much topology information to share from local domains.

A primary goal of *perfSONAR* is to allow "federations" among multiple disjoint monitoring deployments. Federated systems are designed to work together while tolerating different configurations and policies. Federated relationships between instances requires changes in the presentation and storage of information. Sharing data may require omitting items due to policy; thus the ability to efficiently transform data, yet still have it maintain meaning becomes critical.

3. perfSONAR-PS components

In this section, we briefly describe the applications and services that make up the *perfSONAR*-capable pS-NPToolkit, as deployed for monitoring LHC-related networks. The core components of the *perfSONAR* architecture used in this case are the data producers - Measurement Point (MP) and Measurement Archive (MA) services, data consumers (Analysis clients) and discovery - Information Services (IS).

3.1. Information Service

The *perfSONAR* Information Service (IS) is used for service registration, service discovery, data discovery, and network topology representation. These services were previously separated into a Lookup Service (LS) and a Topology Service (TS), but those systems overlap significantly in some cases. The query syntax of the two is essentially the same, and the infrastructure used to support local registration and global discovery is common as well, so these were merged into a single IS.

The discovery function of the IS involves accepting registration information from *perfSONAR* services. As each component updates its information, other components and clients may locate these deployed services via queries. All service descriptions and network metrics, (both actual data and descriptions of the types of data an MP may collect) are defined using XML schema and encoded in XML.

The topology service functionality within the IS stores a representation of the elements of the network topology. This is used for pathfinding, representing relationships between elements about which performance data has been gathered, and to make decisions about topologically-appropriate network services.

Local IS instances accept XML-based information and make it available via XQuery-based queries. These local instances must facilitate discovery of what information sets are contained, but at the same time must constrain the volume of information that is propagated. To address this, IS instances compute "summaries" and register these summaries with higher-level IS instances. Where a local IS instance would have complete information about the data in a given MA, the summarized information would contain information saying "I have metric X for some interfaces in network W.X.Y.Z/24." These summaries can be further summarized to the higher levels of the hierarchy.

When an entity is launching a query against the system, it can first engage in a "discovery phase" during which candidate IS instances are identified, then it can query the set of candidate IS instances for the desired information. Architecturally, there can be multiple levels in the hierarchy, but the currently-deployed software only supports 2 levels: a local and global scope. Additionally, services can be configured to register with multiple IS instances for redundancy.

3.2. *Diagnostic Tools*
A couple of high-level user network diagnostic tools, *NDT* and *NPAD*, are provided on the pS-NPToolkit. *NDT* [4] allows end users to test the network path for a limited number of common problems, such as inadequate TCP buffer sizes and duplex mismatches. *NDT* attempts to determine what kind of performance the user should expect, and what the current limiting factor is. *NPAD* [17] allows end-users to test limited portions of the network path and attempts to determine if there are issues that would adversely effect longer paths. The user provides a target data rate and round-trip-time (RTT) and NPAD attempts to determine if that should be possible, given the infrastructure on that limited portion of the path. Both *NDT* and *NPAD* are registered with the *perfSONAR* IS so that they can be easily located.

3.3. *Measurement Tools*
The pS-NPToolkit contains a collection of tools for collecting passive and active measurements. The specific tools were selected based on two criteria. One, they provide the specific metrics LHC Network administrators determined they needed for monitoring[13]; and, two, they have been extended, in some way, to integrate with the *perfSONAR* infrastructure.

3.3.1. SNMP Passive interface statistics delivered via SNMP [5], are a common non-intrusive health indication of the network. Metrics, such as utilization, errors, and discards at both the octet and packet level, can be especially important when detecting performance and related problems. The pS-NPToolkit incorporates a Cacti [3] instance that can be configured to collect these interface metrics using web-menus. The resulting Cacti round-robin database [27] of metrics is then published using a *perfSONAR* MA interface. There are numerous tools that exist to collect statistics via SNMP polling, many of which offering features that may apply to "power" users and beginners alike. Common tools to collect and mange SNMP data that have been integrated into the *perfSONAR* framework include Cacti [3], Cricket [6], and MRTG [20].

3.3.2. PingER *ping*-based monitoring is frequently used by many wide area network monitoring projects. *ping* monitoring is particularly useful because it is lightweight and only requires ICMP traffic to be allowed through a firewall. The *perfSONAR* PingER MA supports the same set of measured metrics as the PingER project [18], but is built on a completely new code base and integrates *perfSONAR* functionality. The *perfSONAR* PingER MA is configurable using a web-based GUI; it utilizes the *perfSONAR* IS to find other existing measurement nodes to which to run tests. PingER includes a *perfSONAR* MA interface for publishing the end-to-end connectivity metrics.

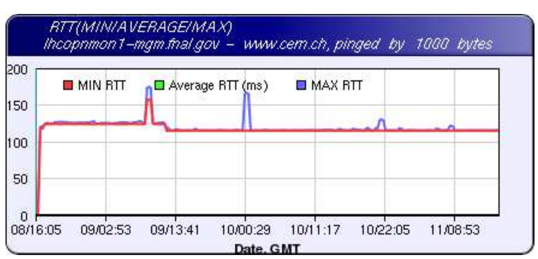

Figure 1. PingER data graph

3.3.3. OWAMP and perfSONAR-BUOY *owamp*[25] is an implementation of RFC 4656[29] and is used to run active tests to collect one-way latency and other related metrics such as loss and

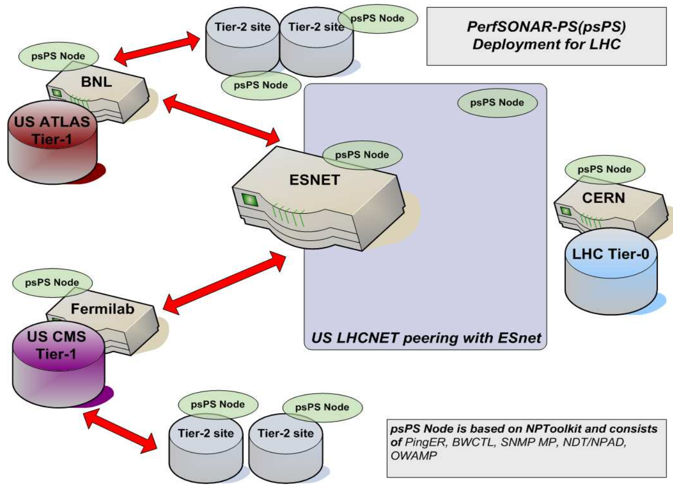

Figure 2. Deployment model for LHC wide area network monitoring

delay variation. One-way latencies are useful for isolating the direction of performance issues and can also be used to look for routing problems as well as interface queueing. perfSONAR-BUOY is a *perfSONAR* service that can be used to define sparse sets of active measurements to be performed and archived. The web-based configuration GUI utilizes the IS to find *owamp* test peers, again allowing user-specified affinities. perfSONAR-BUOY then exposes the *owamp* data using a *perfSONAR* MA interface.

3.3.4. BWCTL and perfSONAR-BUOY *bwctl*[2] is a tool that adds distributed scheduling and policy capabilities to the well known Iperf[10] throughput testing tool. This allows ad-hoc throughput tests to occur on the same host as regular measurements without worry of overlapping tests skewing the results.

For the LHC project, deployments will run regular TCP throughput tests. By default a 60 second test is run every 2 hours.The archived achievable throughput metrics are useful to the LHC participants as a way to set expectations. If the LHC data transfers are not performing similarly to the regular throughput tests, then further analysis is warranted.

As in the *owamp* case, perfSONAR-BUOY is used to configure the set of active throughput tests using *bwctl* in addition to making the archived metrics available through the *perfSONAR* MA interface.

4. Principal Deployment Model and Status

In order to satisfy the needs of LHC computing model we decided to deploy a mesh of the perfSONAR-PS (pS-PS) distributed service. There is no central facility or authority for the management of such system. Every authoritative domain will be deploying its own network monitoring services. For example on the Figure 2 one can see how pS-PS services will be installed at Fermilab and BNL and will be utilized with the same services at ESnet.

We anticipate every Tier-1 and Tier-2 site to be covered with pS-PS network monitoring services. It will allow the network administrator at the end-site to isolate any network related problem and schedule all necessary network tests on-demand for any static or dynamic network circuit. There are currently about 100 deployed pS-PS services in US. They interoperable with perfSONAR MDM services deployed in Europe. That means every trans-atlantic network path could be troubleshooted without any extra effort required for the deployment, configuration and collection of the extra network monitoring data. We anticipate by the end of Year 2009 the collection of pS-PS services will be installed on more than 150 sites and will cover about 30 research and production networks in US. We already achieved 100% coverage on ESnet network and Internet2.The ESnet network provides connectivity for every High Energy Physics research lab in US and Internet2 is the largest R&D network in US.

4.1. ESnet coverage
There are currently 26 monitoring boxes running OWAMP and PingER and 26 running BWCTL and perfSONAR BUOY. There are 3 centralized services where one is running SNMP based utilization Measurement Archive (MA), another one is running Layer2 End-to-End operational circuit monitoring service and last one is running Global Lookup Service. The plans are to add 3 more centralized servers at Layer2 sites and several servers for latency and throughput tests.

4.2. Internet2 coverage
The Internet2 backbone is a major research network in the US. Almost every LHC Tier-2 or Tier-3 site is connected to the respected Tier-1 over the circuits provided by Internet2. The Internet2 backbone contains 9 Points of Presence (PoPs) where perfSONAR tools are deployed for both continuous monitoring by the NOC staff and general availability to the user community. Each PoP location contains several nodes with specific purposes:

- 2 nodes for bandwidth testing (e.g. BWCTL and perfSONAR-BUOY)
- 1 node for latency testing (e.g. OWAMP and perfSONAR-BUOY)
- 1 node containing NPAD and NDT testers

In addition to the remote testing facilities, the Internet2 backbone is served by a single instance of an SNMP MA delivering router interface information for the entire network. This deployment is nearly complete, and will be fully deployed by the end of 2009.

4.3. LHC coverage
The LHC computing network is represented by US ATLAS Tier-1 and Tier-2 sites and US CMS Tier-1 and Tier-2 sites. Among US ATLAS sites we achieved almost 100% coverage with 17 servers been deployed and we are striving to achieve the same coverage on US CMS sites where at the present moment only Fermilab and MIT have installed required monitoring nodes.

4.4. Hardware requirements
There are very modest hardware requirements for the end-site to deploy LHC Network Monitoring Node. Each box is identical and based on 1U chassis, Intel Dual Core E2200 2.2Ghz systems with 2 GBytes RAM, dual 1GbE NICs, and single 160GB SATA disk. The dual NIC cards are required if there is a separate channel for the high impact physics data available at the end-site. It is recommend that each site install two boxes, one for the latency tests and another one for the throughput tests, as latency tests are easily perturbed by other activity on the host. The latency tests are based on PingER round trip time delay probes and one way delay probes provided by OWAMP. The throughput tests are implemented by BWCTL and perfSONAR BUOY services.

5. Experimental Results
As of this writing, the full-scale deployment of *perfSONAR* in the LHC community is underway. To see a list of currently active public *perfSONAR* tools and services, go to *http://www.perfsonar.net/activeServices.*

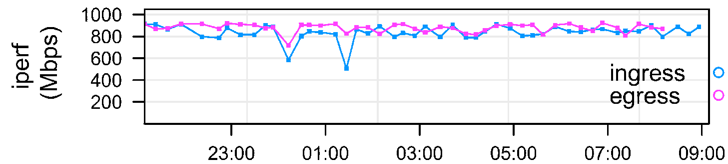

Figure 3. 8-hour history of achievable bandwidth

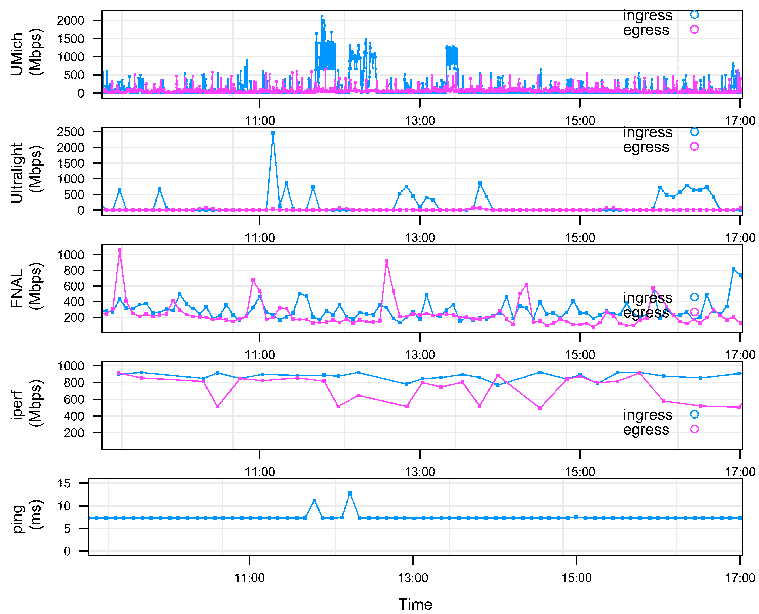

Figure 4. Example comparison of multiple metrics

A simple example of what is possible today is the ability to answer the question: "Give me all the network monitoring data along the path from Host A at Fermi National Lab (FNAL), a Tier-1 site, and Host B at the University of Michigan, a Tier-2 site." This network path crosses four network domains (FNAL, ESnet, Internet2, and U Mich), all of which are publishing SNMP data via a perfSONAR MA. There are perfSONAR MP's on every network segment collecting regular latency measurement, using PingER, and achievable bandwidth [16] measurements, using *iperf*.

Using *perfSONAR*'s Information Service, one can easily determine all available data related to the network path from Host A at FNAL to Host B at UMich. For example, if an LHC user wanted to know what the typical achievable bandwidth was from FNAL to UMich, they can query the perfSONAR-BUOY MA at FNAL for recent iperf results, as shown in Figure 3. This type of data helps greatly set performance expectations for users, allowing users to know what rates are possible on any given path.

If one wanted to look to see if cross traffic was affecting achievable throughput on a given path, they could query for all SNMP data along the path, and compare it with achievable bandwidth and latency data, as shown in Figure 4. This plot shows both ping and iperf results for an 8 hour window on the network path from FNAL to UMich. Note the latency spikes around 11:30 that are clearly related to the traffic spike on the UMich router during that same time.

This is a very simple example of the types of analysis that is enabled by wide deployment of *perfSONAR* services. A few prototype visualization and analysis tools have been written such as GMAPS (http://packrat.internet2.edu:8008/), which provides a Google Maps interface to locate *perfSONAR* data, and perfsonarUI [26], which provides a large number of ways to view various types of *perfSONAR* published data. There are also command line tools that allow one to query for raw data, as was used of the plots in this paper.

6. Future Work

The perfSONAR architecture enables a large number of opportunities for higher-level services and functionality. Current and planned uses for perfSONAR services for the LHC community include:

- monitoring link-by-link status of network circuits to provide general health and performance metrics
- using published topology to implement path-finding algorithms
- locating Inter-Domain Controllers for dynamic circuits
- notification services (e.g. generate an alarm whenever link utilization goes above 90%)
- publishing of middleware and application log data
- publishing of flow-related passive network data (e.g. note specific patterns which could indicate an event such as an intrusion)

As more perfSONAR hosts are deployed, we have quickly discovered the need for better scoping abilities in the IS user interfaces. For example, the query "show me all LHC-related *bwctl* services" currently returns a rather unwieldy list of URLs. Users will need to be given good ways to sort and group related services, perhaps based on administrative domains or geography. Scoping information can be represented in the IS schemas, but has not been used much yet. Growth in *perfSONAR* deployments will begin to require this use in practice. We are also working on the ability to query for perfSONAR services that are "close" to the network path of interest.

Additionally, there is the potential for client applications to utilize *perfSONAR* published performance data to modify application behavior. For the specific LHC use case, the performance data might allow a client application to determine which copy of a remote dataset can be most efficiently retrieved.

7. Security Considerations

Authentication and authorization will be critical for expanding perfSONAR usage. The US LHC sites will be using perfSONAR to make available data that their community policy has determined to be public. However, we are working with several groups that want to use perfSONAR to publish summaries of flow records, but only to a select group of network engineers. Other networks are reluctant to publish network utilization data, and network topology data is almost always deemed sensitive.

For the perfSONAR components to be generally useful, they must integrate with existing authentication and authorization deployments. The wide-variety of existing SAML[28] based mechanisms such as [19][15][32][30][33] used in the R&E community led the perfSONAR group to work with the eduGAIN[8] developers to define mechanisms for bridging authentication and authorization requests from *perfSONAR* to the SAML-based infrastructures. The perfSONAR architecture therefore includes an authentication and authorization-related service (AS), which is used by the other perfSONAR services. The AS enables domains to centralize their authentication and authorization interactions. Other *perfSONAR* services interact with the AS, which then is responsible for communicating with the specific authentication and authorization architectures in use by the domain. This solution requires domains to federate their authentication mechanisms to work. Because federated authentication and authorization architectures are still relatively immature, *perfSONAR* developers isolated these issues to the AS service, which can more easily be modified without causing excessive changes to the rest of the *perfSONAR* architecture.

Even without authentication there are a number of protections in place on the US-ATLAS deployment. The *owamp* and *bwctl* tools both give sites rudimentary control over who can request tests, what kinds of tests they can request, and how much network resources they can consume. Tools like TCP wrappers and firewalls can also be used to restrict access to the perfSONAR services.

8. Conclusion

We described a measurement framework for characterizing the behavior and usage of the network. Our approach for the implementation of the system is a scalable, distributed, service-oriented architecture. The framework combines information from different kinds of measurement tools that currently exist and is able to ea sily accommodate new ones. Full scale deployment of these services is currently underway, and early results show promise. Clearly we have barely begun to scratch the surface on the types of analysis that is enabled by wide deployment of *perfSONAR* services. We hope the network research community will take advantage of this wealth of publicly available information and develop additional interesting analysis tools and techniques that use the *perfSONAR* services.

9. Acknowledgements

This material is based in part on work supported by the National Science Foundation (NSF) under Grant No. OCI-0721902. Any opinions, findings and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the NSF. This work was also supported by the Director, Office of Science, Office of Basic Energy Sciences, of the U.S. Department of Energy under Contract No. DE-AC02-05CH11231 and Contract No. DE-AC02-07CH11359.

10. References

[1] US ATLAS collaboration. http://www.usatlas.bnl.gov/.

[2] Bandwidth Test Controller (BWCTL). http://e2epi.internet2.edu/bwctl/.

[3] Cacti Network Monitoring Tool. http://www.cacti.net/.

[4] R. Carlson. Developing the Web100 based network diagnostic tool (NDT). In *Passive and Active Measurement (PAM)*, 2003.

[5] J. Case, M. Fedor, M. Schoffstall, and J. Davin. A Simple Network Management Protocol (SNMP). RFC 1157, May 1990.

[6] Cricket Network Monitoring Tool. http://cricket.sourceforge.net.

[7] Internet2 Dynamic Circuit Network. http://www.internet2.edu/network/dc/.

[8] edugain.org. http://www.edugain.org/.

[9] 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.

[10] Iperf. http://dast.nlanr.net/Projects/Iperf/.

[11] Knoppix Linux Distribution. http://www.knoppix.net/.

[12] The Lambda Station Project. http://www.lambdastation.org/.

[13] Tier 2 Best Common Practices. http://code.google.com/p/perfsonar-ps/wiki/Tier2BCP.

[14] LHC Optical Private Network. http://lhcopn.cern.ch.

[15] D. R. López and R. Castro-Rojo. Ubiquitous Internet Access Control: The PAPI System. In *13th International Workshop on Database and Expert Systems Applications*, 2002.

[16] B. Lowekamp, B. Tierney, L. Cottrell, R. Hughes-Jones, T. Kielmann, and M. Swany. Enabling Network

Measurement Portability Through a Hierarchy of Characteristics. In *4th International Workshop on Grid Computing (Grid2003)*, 2003.

[17] M. Mathis, J. Heffner, P. O'Neil, and P. Siemsen. Pathdiag: Automated TCP Diagnosis. In *Passive and Active Measurement (PAM)*, April 2008.

[18] W. Matthews and L. Cottrell. The PingER Project: Active Internet Performance Monitoring for the HENP Community. IEEE Communications Magazine on Network Traffic Measurements and Experiments, May 2000.

[19] R. L. Morgan, S. Cantor, S. Carmody, W. Hoehn, and K. Klingenstein. Federated Security: The Shibboleth Approach. *EDUCAUSE Quarterly*, 27(4), 2004.

[20] MRTG - The Multi Router Traffic Grapher. http://oss.oetiker.ch/mrtg/.

[21] Network Measurements Working Group (NM-WG). http://nmwg.internet2.edu.

[22] Network Performance Toolkit. http://code.google.com/p/perfsonar-ps/wiki/NPToolkit.

[23] Open Gird Forum. http://www.ogf.org.

[24] ESnet On-demand Secure Circuits and Advance Reservation System (OSCARS). http://www.es.net/oscars/.

[25] One-way Ping (OWAMP). http://e2epi.internet2.edu/owamp/.

[26] perfsonarUI. http://wiki.perfsonar.net/jra1-wiki/index.php/PerfsonarUI.

[27] Round robin database. http://people.ee.ethz.ch/ oetiker/webtools/rrdtool/.

[28] SAML. http://docs.oasis-open.org/security/saml/v2.0/saml-2.0-os.zip, March 2005.

[29] S. Shalunov, B. Teitelbaum, A. Karp, J. Boote, and M. Zekauskas. A One-way Active Measurement Protocol (OWAMP). RFC 4656 (Proposed Standard), Sept. 2006.

[30] simplesamlphp. http://rnd.feide.no/simplesamlphp.

[31] US CMS collaboration. http://www.uscms.org/.

[32] T. Verschuren, M. Koopmans, and A. Odaci. Introducing A-select, a Web Initial Sign-On System. In *TERENA Networking Conference*, 2003.

[33] V. Welch, T. Barton, K. Keahey, and F. Siebenlist. Attributes, Anonymity, and Access: Shibboleth and Globus Integration to Facilitate Grid Collaboration. In *4th Annual PKI R&D Workshop*, 2005.

[34] Xen. http://www.xen.org/.