International Committee for Future Accelerators (ICFA)
Standing Committee on Inter-Regional Connectivity (SCIC)
Chairperson: Professor Harvey Newman, Caltech

ICFA SCIC Network Monitoring Report

Prepared by the ICFA SCIC Monitoring Working Group
On behalf of the Working Group:
Shawn McKee smckee@umich.edu
Les Cottrell cottrell@slac.stanford.edu
Marian Babik Marian.Babik@cern.ch
Ilija Vukotic ivukotic@uchicago.edu

2016 - 2017 Report of the ICFA-SCIC Monitoring Working Group

Edited by Shawn McKee, Les Cottrell, Ilija Vukotic and Marian Babik On behalf of the ICFA-SCIC Monitoring WG

Contributions from Brian Tierney, Soichi Hayashi, Mike O'Connor

Created January 26, 2017, Finished February 1, 2017

ICFA-SCIC Home Page | Monitoring WG Home Page

This report is available from:

https://docs.google.com/a/umich.edu/document/d/17odQd2C3CLKt7ZkOtLohP_MVY6A-jnnUZuOFB3st0r0/edit?usp=sharing

Table of Contents

Figures

Please note that for all yearly graphs the data point is for the end of the year.

Executive Overview

This year's report has a truncated PingER contribution because of insufficient support for that effort. Much of the focus will be on monitoring and network management activities in High-Energy Physics (HEP).

Internet performance is improving each year with throughputs typically improving by 15% (for developed regions), 20% (for some developing regions such as Russia, South and Central Asia) per year and observed packet-loss rates improving by up to 25% per year. Most countries have converted from using Geostationary Satellite (GEOS) connections to terrestrial links. This has improved performance in particular for Round Trip Time (RTT) and throughput. GEOS links are still important to countries with poor telecommunications infrastructure, landlocked developing countries, remote islands, and for outlying areas. In some cases they are also used as backup links. In future, developing techniques such as weather balloons, solar powered drones, and low and medium earth satellites may assist in providing much reduced latencies and hence performance to remote areas[1].

In general, throughput measured from within a region is much higher than when measured from outside. Links between the more developed regions including N. America[2], E. Asia (in particular Japan, South Korea and Taiwan) and Europe are much better than elsewhere (2 - 10 times more throughput achievable). Regions such as S.E. Asia, S.E. Europe and Latin America are 5-9 years behind. However, in 2009, Africa was ~18 years behind Europe, also Africa's throughput was 12-14 times worse than Europe and extrapolating the data indicated that it would further degrade to almost 60 times worse by 2040. Since 2009, due in large part to the installation of multiple submarine fibre optic cables to sub-Saharan Africa, there has been a significant improvement in Africa's performance. It now appears to be catching up, such that if the present improvements are maintained, it could catch Europe by around 2040. However, since the initial bump in performance in 2010 and 2011, the growth rate has dropped. Hopefully future cable deployments[3] will assist in accelerating the performance once again, especially for East Africa.

Africa and South Asia are two regions where the Internet has seen phenomenal growth, especially in terms of usage. However, it appears that network capacity is not keeping up with demand in these regions. In fact many sites in Africa and India appear to have throughputs less than that of a well-connected (cable, DSL, etc.) home in Europe, North America, Japan or Australia. Further the end-to-end networking is often very fragile both due to last mile effects and poor infrastructure (e.g. power) at the end sites, and also due to lack of adequate network backup routes. Africa is a big target of opportunity with over a billion people of which in 2012

---

[1] Providing Internet access for hard to reach places, see https://confluence.slac.stanford.edu/display/IEPM/Providing+Internet+Access+for+hard+to+reach+places.

[2] Since North America officially includes Mexico, the Encyclopedia Britannica recommendation is to use the terminology Anglo America (US + Canada). However, in this document North America is taken to mean the U.S. and Canada.

[3] A Giant Leap 2016? Africa is narrowing its Techno-gap, see http://www.huffingtonpost.com/david-tereshchuk/a-giant-leap-in-2016-africa_b_8901556.html

only 15.6% were Internet users. This grew to 28.7% in 2016[4]. It also had a 7,448% (compared to 918% for the world) growth in number of Internet users from 2000-2016. However, there are many challenges including lack of power, import duties, lack of skills, disease, corruption, and protectionist policies. In almost all measurements, Africa stands out as having the poorest performance. Further Africa is a vast region and there are great differences in performance between different countries and regions within Africa.

There is a moderate to strong positive correlation between the Internet performance metrics and economic and development indices available from the UN and International Telecommunications Union (ITU)[5]. Given the difficulty of developing the human and technical indicators (at best they are updated once a year and usually much less frequently); having non-subjective indicators such as PingER that are constantly and automatically updated is a very valuable complement.

Between the Fertility Rate[6] and PingER derived throughput, there is a negative correlation[7] (i.e. countries with higher Fertility Rates have lower Internet performance). This is very concerning since Fertility Rates drive population growth and predictions indicate the world population will exceed 11 billion by 2100. This is driven by Africa for which the population could exceed 6 billion[8] by 2100. Such extreme growth threatens Africa's development and stability. Add to this that achieving significant fertility declines requires education and easy access to information, and this in turn is enabled by good internet access. Thus countries such as Niger, Burkina Faso and Zambia with high Fertility Rates and low Internet performance are particularly at risk.

For modern HEP collaborations and Grids there is an increasing need for high-performance monitoring to set expectations, provide planning and troubleshooting information, and to provide steering for applications. As link performance continues to improve, the losses between developed regions are decreasing to levels that are not measureable by PingER. Though the measurements for RTT, jitter, and unreachability[9] are still correct, as the measured losses go to zero this also makes the throughput derivation unreliable. Alternative solutions to measuring the throughput are available, however they can be harder to install and absorb more network bandwidth. Examples of other measurement projects using the more intense methods are the MonALISA[10] project that uses the pathload[11] packet pair technique as well as file transfers, and

---

[4] See http://www.internetworldstats.com/stats.htm

[5] Development Indices and PingER Correlations at https://confluence.slac.stanford.edu/display/IEPM/Development+Indices+and+PingER+Correlations

[6] List of sovereign states and dependent territories by fertility rate at https://en.wikipedia.org/wiki/List_of_sovereign_states_and_dependent_territories_by_fertility_rate.

[7] Correlation of PingER throughput and Fertility at https://confluence.slac.stanford.edu/display/IEPM/Correlation+of+PingER+throughput+and+Fertility

[8] Six Billion in Africa, Robert Engleman, Scientific American February 2016.

[9] A host is considered unreachable when none of the pings sent to it there is no response to any of the pings sent to it.

[10] MonALISA, see http:// monalisa.caltech.edu

[11] Pathload, see http://www.cc.gatech.edu/fac/Constantinos.Dovrolis/bw-est/pathload.html

perfSONAR[12] that uses the iperf[13] (and more recently iperf3[14]) TCP transport mechanism.

Introduction

This report may be regarded as a follow up to the previous ICFA Standing Committee on Inter-regional Connectivity (SCIC) Monitoring working group's Network reports[15] dating back to 1997.

The current report updates the January 2016 report. **As noted, PingER activities will be covered in the same depth as earlier reports because of a lack of funding for this effort**. We will be including some new areas related to network monitoring in HEP including updates and status on the perfSONAR efforts globally as well as the WLCG Network and Transfer Metrics Working Group activities.

Methodology

There are two complementary types of Internet monitoring reported on in this report.

1. In the first we use PingER[16] which uses the ubiquitous "ping" utility available standard on most modern hosts. Details of the PingER methodology can be found in the Tutorial on Internet Monitoring & PingER at SLAC[17]. PingER provides low intrusiveness (~ 100 bits/s per host pair monitored[18]) RTT, loss, jitter, and reachability (if a host does not respond to a set of up to 30 pings it is presumed to be unreachable). The low intrusiveness enables the method to be very effective for measuring regions and hosts with poor connectivity. Since the ping server is pre-installed on all remote hosts of interest, minimal support is needed for the remote host (no software to install, no account needed etc.)
2. The second method (perfSONAR[19] etc.) is for measuring high network and application throughput between hosts with excellent connections. Examples of such hosts are to be found at HEP accelerator sites and the Large Hadron Collider (LHC) tier 1 and 2 sites, major Grid sites, and major academic and research sites in N. America, Japan and Europe. The method can be quite intrusive (for each remote host being monitored from a monitoring host, it can utilize hundreds of Mbits/s or more for ten seconds to a minute, each hour). To minimize intrusion, the WLCG scheduling utilizes 30 second tests every

---

[12] What is perfSONAR available at http://www.perfsonar.net/

[13] Iperf home page is available at http://dast.nlanr.net/Projects/Iperf/

[14] Iperf3 at ESnet is available at http://software.es.net/iperf/

[15] ICFA/SCIC Monitoring Working Group's Annual Reports, see http://www.slac.stanford.edu/xorg/icfa/scic-netmon/#annual

[16] "PingER". Available http://www-iepm.slac.stanford.edu/pinger/; W. Matthews and R. L. Cottrell, "The PingER Project: Active Internet Performance Monitoring for the HEP Community", IEEE Communications Magazine Vol. 38 No. 5 pp 130-136, May 2002.

[17] R. L. Cottrell, "Tutorial on Internet Monitoring & PingER at SLAC". See http://www.slac.stanford.edu/comp/net/wan-mon/tutorial.html

[18] In special cases, there is an option to reduce the network impact to ~ 10bits/s per monitor-remote host pair.

[19] PERFormance Service Oriented Network monitoring Architecture , see http://www.perfsonar.net/

6-24 hours (depending upon host-pair groupings) rather than every hour. It also requires more support from the remote host. In particular either various services must be installed and run by the local administrator or an account is required, software (servers) must be installed, disk space, compute cycles etc. are consumed, and there are potential security issues[20]. The method provides expectations of throughput achievable at the network and application levels, as well as information on how to achieve it, and troubleshooting information.

PingER Status

Deployment

The current January 2017 deployment of PingER can be seen in Figure 1 below.

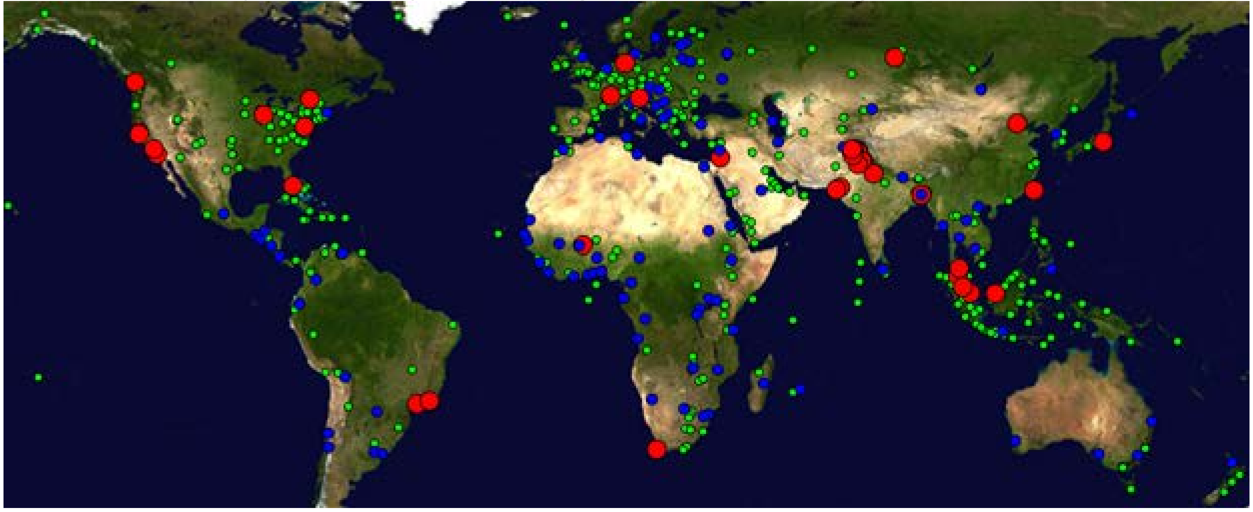

Figure 1: Locations of PingER monitoring and remote sites as of January 2017. Red sites are monitoring sites, blue sites are beacons that are monitored by most monitoring sites, and green sites are remote sites that are monitored by one or more monitoring sites.

Historical Growth of PingER Coverage Since 1998

Figure 2 shows the growth in the number of PingER monitoring sites, countries monitored, active remote sites monitored by PingER from SLAC and monitor-remote host pairs since 1998. Initially the main regions monitored were North America, Europe, East Asia, and Russia. These were the regions with the main HEP interest. Starting in 2003-2004, we increased the number of hosts monitored in developing regions such as Africa, Latin America, Middle East and South Asia. Starting in 2007 the number of countries monitored plateaued at 160-170, and in between 2011 and 2015 the number of monitor-remote site pairs plateaued fluctuating between 11,000 and 13,000. Also less apparent but more important the number of monitoring sites (Monitors) dropped from 97 in 2011 to just over 50 in 2016. This was driven by the reduced support from the Higher Education Commission in Pakistan and somewhat masked by the addition of support from Malaysian Universities. We expect a further drop in number of monitors in 2017 as we disable more non–responding monitors in Pakistan.

---

[20] WLCG/OSG perfSONAR details: https://twiki.opensciencegrid.org/bin/view/Documentation/DeployperfSONAR

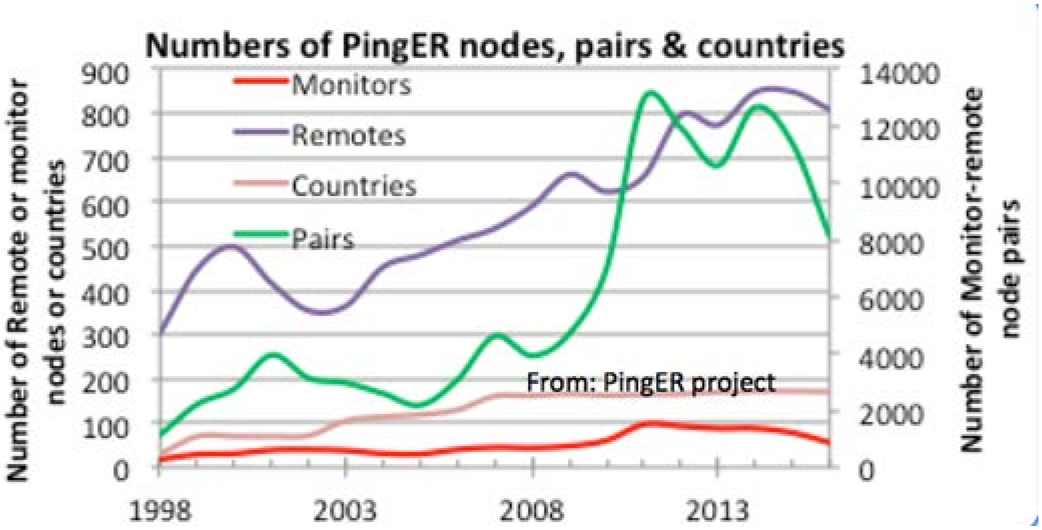

Figure 2: The growth in PingER monitoring hosts, remote hosts monitored, countries monitored & monitor-remote site pairs

Metrics

See the 2015 ICFA/SCIC monitoring report[21] for details on the use of the various PingER metrics and the measurements between 1998 and 2014. Here we only report on what is probably the most used metric, i.e. the throughput.

Yearly Throughput Trends

Figure 3 shows the annual throughput measured from SLAC using the Mathis formula[22] to derive the throughput from PingER measured RTT and loss. Since the throughput derivation depends critically on the inverse RTT we have not shown SLAC to N. America since the small RTTs would distort the results.

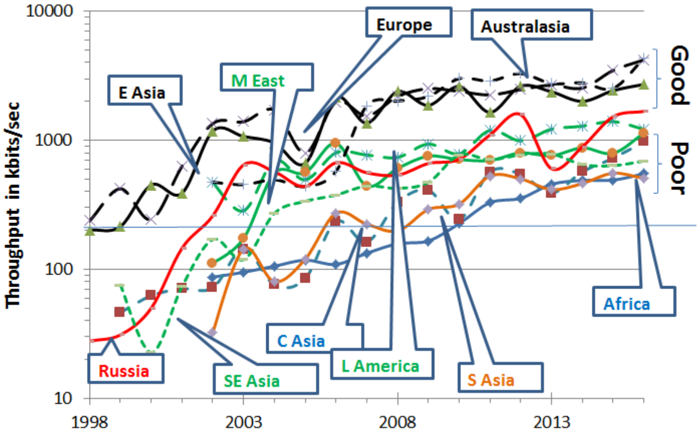

Figure 3: Derived Throughput kbits/sec from SLAC to the World (since the throughputs in this graph are not

---

[21] International Committee for Future Accelerators - Standing Committee on Inter-Regional Connectivity (ICFA-SCIC) 2015 Report on Networking, compiled by Les Cottrell and Shawn McKee on behalf of the working group, January 2015.

[22] M. Mathis, J. Semke, J. Mahdavi, T. Ott, "The Macroscopic Behavior of the TCP Congestion Avoidance Algorithm", *Computer Communication Review*, volume 27, number 3, pp. 67-82, July 1997

normalized we have not shown N. America) until Dec 2016

It is seen that throughput performance for the world divides roughly into two with Europe, Australasia, East Asia and North America leading, and the Middle East, Central Asia, South Asia and Africa bringing up the rear. Latin America and Russia are hovering on joining the leaders.

Parts of Latin America moved from satellite to fibre in 2000, and E. Asia in 1999. Also note the impact of moving the ESnet routing from E. Asia (in particular Japanese academic and research networks) to the US via New York in 2001 to a more direct route via the West Coast of the US. Also note that there is almost a 10 times difference in throughput between Africa and N. America, Europe and Oceania. Africa is the worst off region and had the slowest rate of improvement until 2008.

The improved performance for Russia 2001 is since the two Russian hosts being monitored at the time (one in Moscow and one in Novosibirsk) switched from using satellite links to land-line links with lower RTTs. The slow increase in performance starting in 2007 for Europe, Australia and East Asia in Figure 3 is partially an artifact of the difficulty of accurately measuring loss with a relatively small number of pings (14,400 pings/month at 10 pings/30 minute interval, i.e. a loss of one packet ~ 1/10,000 loss rate). We looked at using a method[23] that allows for zero packet loss, however it requires one to know the maximum congestion window size. Unfortunately this varies from host to host and can easily be changed, so we did not pursue it.

Looking at the data points one can see:

- East Asia and Oceania are catching Europe;
- Russia is 6 years behind Europe and catching up;
- Latin America and the Middle East are 8 years behind and falling further behind;
- S. E. Asia is also 8 years behind but is catching up;
- S. Asia and Central Asia are 13 years behind and keeping up;
- We deal with Africa in the next section.

Africa

Africa is 14 years behind Europe. See Figure 5 where it is seen:

- In 2008-2009 Africa was 12-14 years behind Europe and even worse was falling further behind such that in 2030 it would have been 60 times worse off or almost 28 years behind.
- Prior to 2008 the rate of improvement was a factor of 2 in 7 years
- Since 2008 the improvement is a factor of 3 in 5years and at the current rate it could catch up with Europe by around 2040.
- This remarkable improvement is largely a reflection of the impact of the multiple terrestrial links installed since 2008[24], initially driven by the soccer world cup.

---

[23] "Modeling TCP throughput: A simple model and its empirical validation" by J. Padhye, V. Firoiu, D. Townsley and J. Kurose, in *Proc. SIGCOMM Symp. Communications Architectures and Protocols* Aug. 1998, pp. 304-314.

[24] African Undersea Cables, see http://manypossibilities.net/african-undersea-cables/

- However, there is evidence that the rate of catch up has fallen off in 2013, 2014, 2015 and 2016. Shown in the figure is the estimate from the 2015 data of catch up by 2040 (solid redish-brown line), the 2016 data extends that to about 2046 (dashed red line).

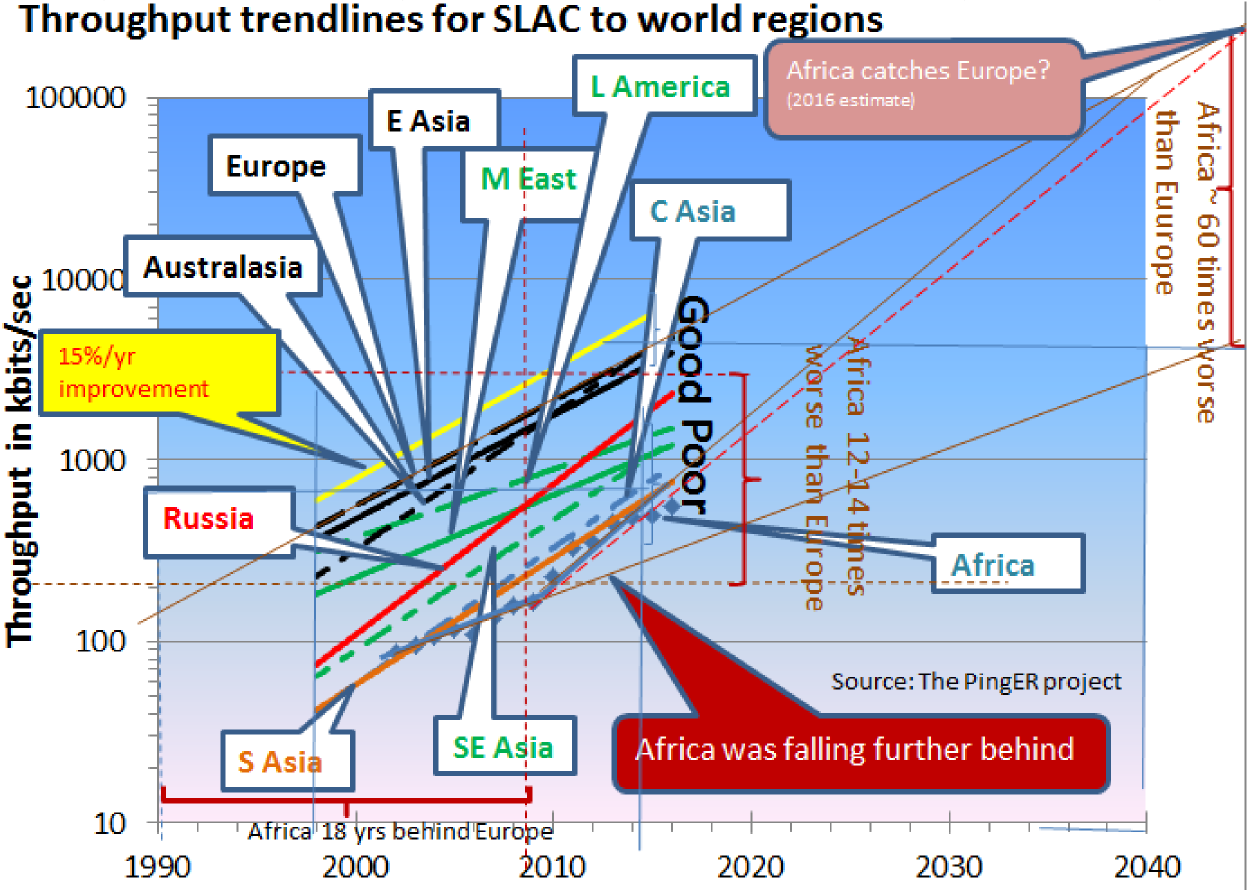

Source: The PingER project

Figure 4: Extrapolations on the throughput data with a focus on Africa.

More detail on the derived throughput seen from SLAC for the various African sub-regions is seen in Figure 5, where it is seen that:

- The instability in the All Africa data in the early years (up through 2003) is related to only monitoring <=3 sites in Africa,
- North Africa, for long the leader is being caught up to by South and West African countries,
- The instability and lack of growth 2009 on may be partially due to the “Arab Spring”.
- Sub-Sahara is tracking all Africa but about 10% lower.
- South Africa and East Africa appear to be catching up with North Africa
- East Africa and West Africa saw big improvement in 2010. It is still improving but much more slowly, possibly more linearly rather than exponentially.
- The East African region appears to have the worst performance.

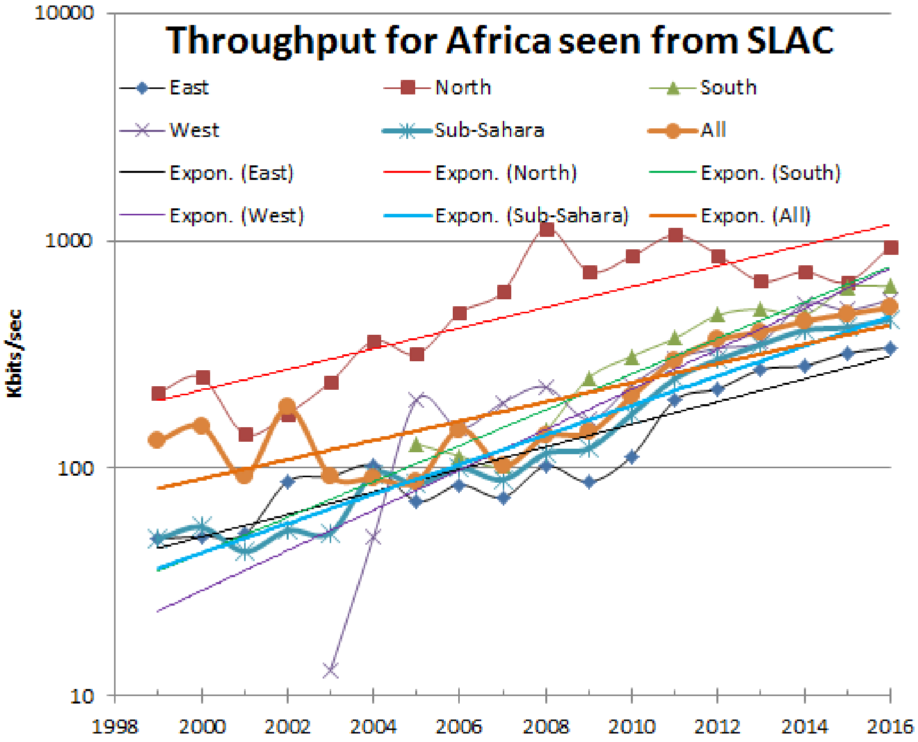

Figure 5: Derived annual throughputs seen from SLAC to African regions together with exponential trendlines

There are reports[3] that a new subsea cable, is to be laid off Africa's eastern coastline by Liquid Sea, a literally offshore subsidiary of Liquid Telecom, which already runs a pan-African network based on terrestrial fiber-optic cable supplemented by satellite links for rural and remote areas. “It will offer speeds of 20-30Tbps, up to 10 times the capacity of existing submarine cables in the region, enabling a reliable and affordable international connectivity service to both coastal and landlocked countries in Eastern, Central and Southern Africa.” Hopefully this will measurably assist in dramatically improving access in East Africa.

PingER Progress in 2016

There is a complete history of the PingER project this millennium[25]. Here we only report on 2016.

- Joao Rulff a student from Brazil spent 3 months at SLAC and developed an application to provide a PingER data warehouse, a system where all the PingER data would be stored in a efficient way. This task was divided in three main phases: the creation of a script that

---

[25] History of PingER this Millenium, see https://confluence.slac.stanford.edu/display/IEPM/History+of+growth+in+PingER+hosts+this+millenium

automates the update of the PingER public FTP, the creation of a four node cluster where the Hadoop file system (HDFS) is running and the creation of an ETL (Extraction, Transformation, Loading) process to populate the data warehouse everyday. This process allows the users to perform queries involving millions of rows in a much faster way.real time flexible queries of the PingER database.

- The automatic updating of the FTP site with the PingER data enables a collaborator to download the current archive. In addition there is a process to provide the last 6 months hourly data on a daily basis[26]. This enables a collaborator to update their archive with the latest data without loading the complete archive.
  - Since september of 2016 the PingERVis project is under development in Brazil. This project aims to develop an interactive visualization system to PingER Data where the user can select some attributes, generates different kinds of visualization and relate PingER data with a bunch of other metrics (economic and social index, for example). This project is focused on reducing the query response time, giving the user a very efficient system to explore PingER data and provide different approaches to data exploration. This project was divided in two main phases: data normalization and data visualization.

The data normalization pipeline was developed using the Python programming language and consists in three main phases: daily aggregation, monthly aggregation and country aggregation. The goal of this pipeline is keep a smaller, but representative, set of data which any kind of relational database can consume and perform queries in a low response time. This pipeline also includes a algorithm to remove outliers from the PingER dataset. This pipeline can be configured to run everyday using SciCumulus (tool to manage scientific workflows).

The data visualization phase allows the user to perform spatial data exploration and timely data exploration. The spatial data exploration consists in letting the user choose some set of countries and then visualize the data only for the selected countries. The timely data exploration allows the user to select a time range to visualize the data. Besides the capabilities mentioned above, the system can also consume any set of data, following a specific standard, to be related with PingER data (economic and social index, for example). During the development, the main tool used was Flask Framework to generate the web application and Javascript libraries like D3.js and Jquery. An example of the output for an RTT heat map for various countries is seen below:

---

[26] PingER Public FTP Archive daily update, see https://confluence.slac.stanford.edu/display/IEPM/PingER+Public+FTP+Archive+daily+update

Figure 6: RTT heat map derived from PingER

- The production PingER MA at SLAC was successfully migrated to a Virtual Machine after verifying that the Virtual Machine did not impact the results[27].
- Two projects were put together to explore migrating PingER to an Android smartphone. Two methods were addressed: porting the original Perl code (pinger2.pl); and writing a completely new version using native Android tools.
    - Sara Massood of the University of Agriculture Faisalabad (UAF) in Pakistan ported Perl to the Android and then installed pinger2.pl. There is a tutorial[28]. Currently, no data is available due to unavailability of the live IP address.
    - Amity University in India decided to use the native Android tools and rewrote pinger2 in java. They are also implementing a proxy which will enable the SLAC archiver to get the data on a daily basis. This has been submitted as a paper[29].
- Aqsa Hameed of UAF created a Data warehouse of PingER data. First they transformed the PingER text files containing the hourly data into CSV files. These were then uploaded into the HaDoop File System and populated Impala Tables and queries[30]. Line and Bar charts are created on a webpage running or executed by the localhost server. UAF are also working on detecting missing and anomalous values in the PingER archive. Missing values are then handled by replacing with average values.

---

[27] PingER on a virtual machine at SLC, see https://confluence.slac.stanford.edu/display/IEPM/PingER+on+a+Virtual+Machine+at+SLAC

[28] Tutorial on installing PingEr MA on Android, Sara Masood, see https://www.dropbox.com/s/v61as3h7wax940y/pinger%20Configuration%20tutorial.docx?dl=0

[29] Implementation of PingER on an Android, R. Sampson et. al 7th International Conference on Cloud Computing, Data Science and Engineering

[30] Applying big data warehousing and visualization techniques on pingER data, Aqsa Hameed, Saqib Ali, Roger Les Cottrell, Bebo White, submitted to BDCAT '16: Proceedings of the 3rd IEEE/ACM International Conference on Big Data Computing, Applications and Technologies. See http://dl.acm.org/citation.cfm?doid=3006299.3006337.

Collaboration

The PingER collaboration meets monthly by Skype. It consists of members from:

- SLAC
- National University of Sciences and Technology (NUST[31]), Islamabad, Pakistan
- University of Agriculture (Faisalabad) (UAF[32]), Faisalabad, Pakistan
- University of Malaysia in Sarawak (UNIMAS[33]),Kuching Malaysia
- University Utara Malaysia (UUM[34])), Sintok, Kedah, Malaysia
- Amity University, Noida, Uttar Pradesh, India[35]

In addition there are collaborators at the Federal University of Rio de Janeiro, and the Federal Rural University of Rio de Janeiro, with whom we meet on a less regular basis.

Papers:

- *Applying Big Data Warehousing and Visualization Techniques on PingER Data*, A Hameed, S. Ali, R. Cottrell, B. White, IEEE/ACM 3rd International Conference on Big Data Computing, Applications and Technologies, 2016
- *Implementation of PingER on Android*, R Sampson, S Rajappa, A Sabitha, A Bansal, B White, R Cottrell, 7th International Conference on Cloud Computing, Data Science and Engineering Jan 2017

Presentations

- *Applying Big Data Warehousing and Visualization Techniques on PingER Data*, A Hameed, S. Ali, R. Cottrell, B. White, IEEE/ACM 3rd International Conference on Big Data Computing, Applications and Technologies, 2016
- *Implementation of PingER on Android*, R Sampson, S Rajappa, A Sabitha, A Bansal, B White, R Cottrell, 7th International Conference on Cloud Computing, Data Science and Engineering Jan 2017

Publicity:

- *A Giant Leap in 2016? Africa is Narrowing its Techno gap*, Huffington Post with several mentions of PingER contributions[36]

---

[31] NUST, see https://en.wikipedia.org/wiki/National_University_of_Sciences_and_Technology_(Pakistan)

[32] UAF, see https://en.wikipedia.org/wiki/University_of_Agriculture_(Faisalabad)

[33] UNIMAS, see https://en.wikipedia.org/wiki/Universiti_Malaysia_Sarawak

[34] UUM, see https://en.wikipedia.org/wiki/Universiti_Utara_Malaysia

[35] Amity University, see https://en.wikipedia.org/wiki/Amity_University

[36] The Huffington Post, see http://www.huffingtonpost.com/david-tereshchuk/a-giant-leap-in-2016-africa_b_8901556.html

High Performance Network Monitoring

Introduction

PingER is an excellent light-weight way to measure global network trends as was shown earlier in this report, but doesn't provide enough detail for high-performance network monitoring requirements in the Worldwide LHC Computing Grid (WLCG) or in the Open Science Grid (OSG). Grid sites with significant amounts of storage and/or computing power critically rely upon the network to enable them to function effectively with their peer sites globally. Problems in the network can both be severely disruptive and hard to identify and locate. For example, very low levels of packet-loss can significantly degrade the throughput between sites with large round-trip times (RTT). Routes between sites can unexpectedly change pushing traffic onto less capable or improperly configured network paths. Bandwidth between sites may be significantly less than what is provisioned (wire-speed) because of misconfigurations or competing traffic. We need tools that can track and identify how our networks are performing in detail.

To gather more detailed metrics we have chosen to deploy the perfSONAR toolkit[37] at our WLCG and OSG sites worldwide. Over the last 5 years we have crafted a consortium amongst the perfSONAR developers, the WLCG experiments and the Open Science Grid to gather detailed network metrics amongst our largest grid sites, persistently store those metrics and provide visualization and analytics tools for users to understand how the networks are behaving. In the following sections we will discuss the components and collaborations we have in place to provide the needed network monitoring with HEP.

Network Performance Monitoring for HEP

High-energy physics grid sites rely upon the network to provide access to their computing and storage. The network provides the basis for users to access those resources and for virtual organizations to organize the sharing and use of their member's resources. When there are problems in the network, it can significantly degrade or even disable users and VOs ability to do their science.

Networking problems can be difficult to identify and isolate for numerous reasons:

- Network paths typically span multiple administrative domains with no single entity having complete access to the end-to-end infrastructure components
- Applications that work well on a Local Area Network (LAN) may behave significantly differently when run on a Wide Area Network (WAN) due to the impact of latency on the network communication involved
- End-host or LAN issues may be the actual source of problems and differentiating HOST vs LAN vs WAN problems can be difficult without sufficient expertise
- Problems that actually exist on a host or in the LAN may not be "visible" in local use and only show up when the application is used in the WAN. The tendency is to believe the

---

[37] The perfSONAR toolkit, see deployment information http://www.perfsonar.net/deploy/.

WAN is the problem even though it many cases it isn't

- Real WAN problems are hard to localize and it is not practical or effective to contact every entity managing a portion of the network path on which you see a problem.

For high-performance, data intensive sites, PingER is insufficient to address these issues. Our goal is to help users, Virtual Organizations (VOs) and site administrators better understand their network infrastructure and enable them to more effectively find problems and isolate their root cause. To do this WLCG is mandating its Tier-0/Tier-1/Tier-2 sites deploy perfSONAR Toolkit instances in their infrastructure. The following sections will highlight their management and use for HEP.

The perfSONAR Project

Because our working group relies upon perfSONAR for our high performance network monitoring it is appropriate to review the status of the perfSONAR effort. perfSONAR is an open source software project that enables seamless deployment of a network monitoring infrastructure. perfSONAR is currently a successful medium size open source project with ~2000 known public deployed instances, and likely an equal number of private deployments.

The global Research & Education (R&E) network ecosystem is comprised of hundreds of international, national, regional and local-scale networks. While these networks all interconnect, each network is owned and operated by separate organizations (called “domains”) with different policies, customers, funding models, hardware, bandwidth and configurations. This complex, heterogeneous set of networks must operate seamlessly from “end to end” to support science and research collaborations that are distributed globally.

The perfSONAR collaboration is an Open Source project lead by ESnet, Internet2, Indiana University, and GEANT. Each organization has committed 1.5 FTE effort to the project. The project also gets additional help from many others in the community, such as OSG, RNP, SLAC, and others.

Recent Changes to perfSONAR

The perfSONAR Roadmap is influenced by the following: requests on the project issue tracker; annual user surveys sent to everyone on the user list; regular meetings with VOs using perfSONAR such as the WLCG and OSG; and discussions at various perfSONAR related workshops.

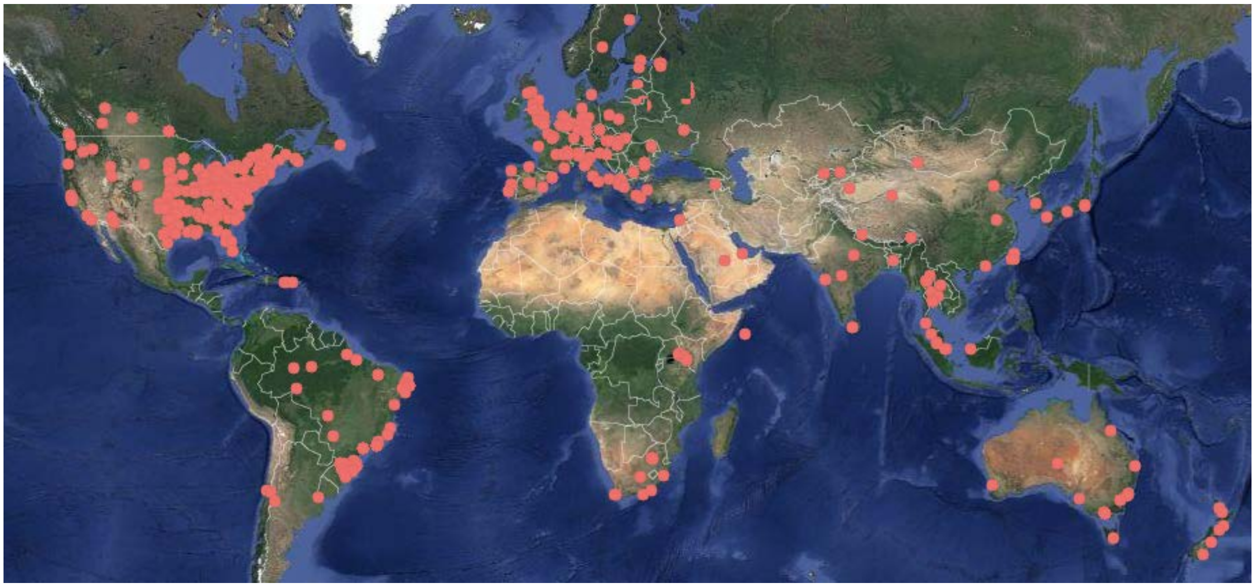

Figure 7: perfSONAR public network as of January 2017, there are currently around 2000 known deployed instances with likely an equal number of private deployments

During 2016 the total number of deployed perfSONARs increased to more than 2000 within more than 1000 domains, most of them running the latest production version (3.5.1.7). Almost half of all the deployed test points offers 10 Gbps connectivity, but there are more than 60 public test nodes offering 40 Gbps and around 17 offering 20 Gbps. Most of the toolkits are running on Linux/RHEL/CentOS6, but there is also a considerable portion of Debian deployments with a fraction of test points running on the virtual machines (8%). This year has been very active in terms of developments with major development release 4.0 scheduled to become official within next couple of weeks (February 2017). The main features of the perfSONAR 4.0 include:

- New web-based configuration interface focused on large deployments providing a centralized interface to manage hundreds of toolkits. It offers capability to search for test points in the global lookup service and configure regular testing between any set of endpoints.
- Brand new test scheduler (pscheduler), introducing new protocol and replacing the previous bandwidth controller (bwctl) and regular testing while adding pluggable support for adding new test tools (such as twamp) and archive backends (such as RabbitMQ). Also offering REST API and persistent store to maintain the test schedules during reboots, outages, etc.
- Improved graphing support and dashboards providing drill-down capabilities as well as basic diagnostics for most common problems.

perfSONAR next steps

Looking forward, perfSONAR collaboration aims to play an increasingly vital role in network operation and debugging with focused investments in automation. The key areas that could benefit from automation: configuration, execution, and analysis.

- While significant progress has been made already, there is still significant potential for improvement in the area of reconfiguring test meshes after an additional node is inserted into the mesh. Whereas today the system is designed to build the "start-up" mesh very well, the system is not well designed to add "one more node" in an evolutionary way. As such, automation to support a living and evolving web of nodes and tests could provide significant improvement to both WAN operations and data intensive science collaborations.
- There is potential for significant progress in the execution of tests built around machine learning. Specifically, perfSONAR does not currently have the ability to auto-detect problems along an end-to-end path or automatically run additional tests. We believe that many of the steps a Tier 1 network engineer might do to diagnose a network problem could be automated by running through the initial script of steps.
- There is potential for significant progress in the analysis of performance monitoring data. Specifically, through machine learning techniques, we believe that perfSONAR can be augmented with anomaly detection tools that effectively given intelligent alarms. Coupled with feedback about the rate of false positives or false negatives, this has the potential to create a living, breathing system for WAN operators and data intensive science collaborations to understand the performance of the network.

The rapid rise of cloud-based services and cloud-based workflows is transforming the science of data intensive scientific collaborations. Increasingly, networks, compute devices, and storage devices are assembled on the fly for a fee to solve a scientific problem. As such, performance monitoring cannot be a cumbersome "build it once infrastructure" but must become ephemeral, instantiated only when the underlying compute, storage, and network infrastructure is created.

As this is a very new area for the perfSONAR collaboration, we aim to focus on a specific, targeted use case to understand what can be accomplished.

Network Monitoring Platform

The Open Science Grid (OSG) facilitates access to distributed high throughput computing for research in the US. Since 2012, part of OSGs focus has included the network because of its centrality in connecting the components of science computing grids. OSG is providing networking information for its constituents and its partners (like the Worldwide LHC Computing Grid (WLCG)).

OSG is continuing to be the network information provider for its users and its partners. In collaboration with WLCG, OSG has been guiding and supporting the deployment of perfSONAR toolkit instances at its member and partner sites. More importantly OSG has developed and supported a comprehensive Network Monitoring Platform, which it operates to configure, gather, store and expose measurements from the WLCG/OSG perfSONAR network and makes the data available for anyone to use. In addition, it provides tools to register perfSONAR instances and organize testing between the sites. The components of the platform are depicted in Fig.8 and described in detail in the following chapters.

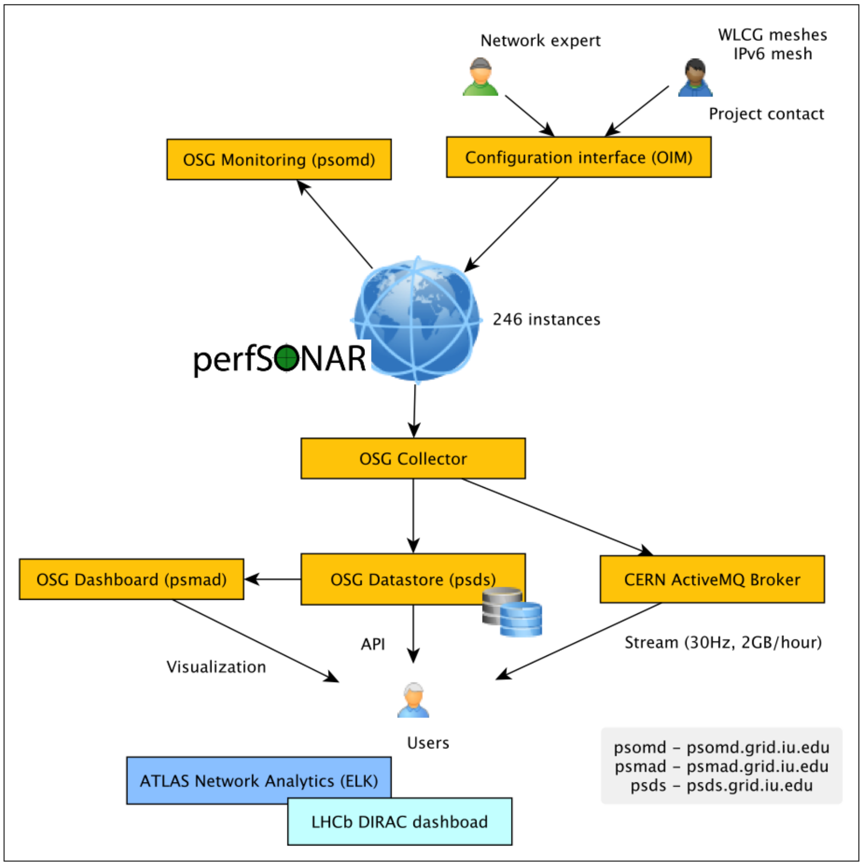

Figure 8: OSG Network Monitoring Platform

WLCG/OSG perfSONAR network

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, which is a subset of the global perfSONAR deployment dedicated to the needs of WLCG and OSG (see Figure 9). It has been established and is operated since 2013 and involves all WLCG T1 and T2 sites, all OSG sites as well as testing endpoints at the major R&E network hubs (ESNet, GEANT, Internet2) and associated HEP projects such as Belle II sites. OSG's Network Monitoring Platform ensures that 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 network measurement site (typically WLCG/OSG Tier-[0/1/2]) needs to provide two types of perfSONAR services: 1) latency and 2) bandwidth. The latency instances are measuring end-to-end latency, packet loss, packet reordering and number of TTLs hops by implementing the One-way latency measurement protocol (OWAMP). Bandwidth instances are measuring throughput (typically via iperf3) as well as registering the network paths using either tracepath or traceroute. All measured metrics are centrally collected and are available both as a (near real-time) stream as well as on request via OSG network datastore API.

WLCG and OSG specific documentation on perfSONAR is available covering motivation, deployment options, installation guide, configuration, use and troubleshooting.

Network Datastore

OSG and WLCG have worked closely together on perfSONAR for high-energy physics (and others). OSG, as a member of WLCG, has agreed to become the control hub for the global perfSONAR deployment and has developed a Network Datastore, based upon the Esmond[38] datastore in perfSONAR v3.5, to host all the perfSONAR metrics. This datastore collects in near real time all perfSONAR test results from OSG/WLCG perfSONAR instances and allows users to query the test results from a single Esmond instance. The datastore also publishes the test results through ActiveMQ server hosted by WLCG that users can subscribe to. Figure 10 provides an architecture diagram of the OSG network datastore, showing the use of a back-end

---

[38] ESnet Monitoring Daemon, http://software.es.net/esmond/

Cassandra[39] database easily scaled by adding additional instances. This datastore is the source of network metrics for OSG and WLCG and went to production in November 2015.

Figure 9: WLCG perfSONAR network as of January 2017, there are 250 registered and active instances measuring latency/packet loss, network path and achievable throughput. Each red dot represents a WLCG recommended deployment of the latency and bandwidth instance, where each instance contains measurement tools, configuration tools (regular testing), local datastore and visualization.

---

[39] See the Apache Cassandra project, http://cassandra.apache.org/

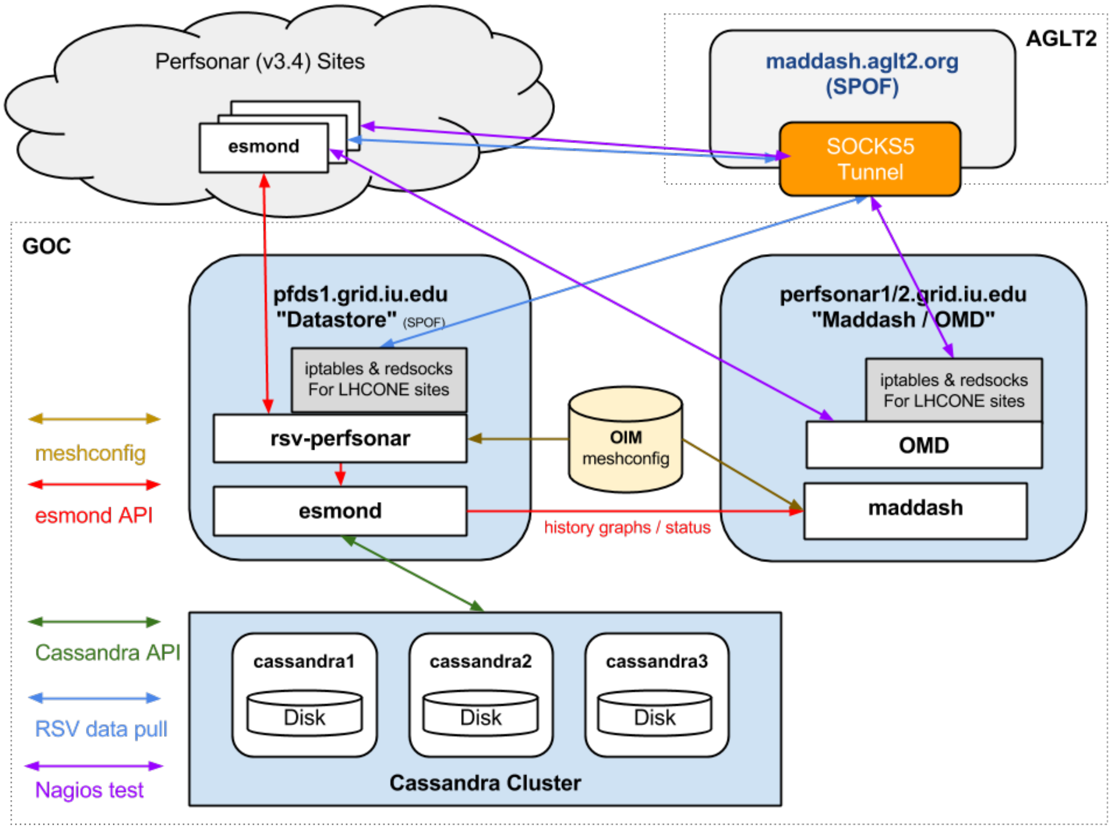

Figure 10: The OSG network datastore architecture used to gather, organize and archive network metrics from the global OSG and WLCG perfSONAR deployment and make them available for visualization or higher-level services.

Central Configuration Interface

One of the challenges for a large scale deployment of perfSONAR is managing the tests amongst the participating sites. When USATLAS began deploying perfSONAR in 2008, all configuration for each site was controlled by “emails” to the perfSONAR administrators. Every change (addition or deletion) required every administrator to update their configuration. The perfSONAR developers provided a solution with the so called “mesh-configuration” (see Figure 11). The perfSONAR toolkit was updated to provide a mesh agent that could get its configuration from URL. A web server could provide the JSON configuration for a whole mesh and changes could be made centrally. The perfSONAR administrators just needed to configure their agent to read from the specified URL for each mesh they participate in.

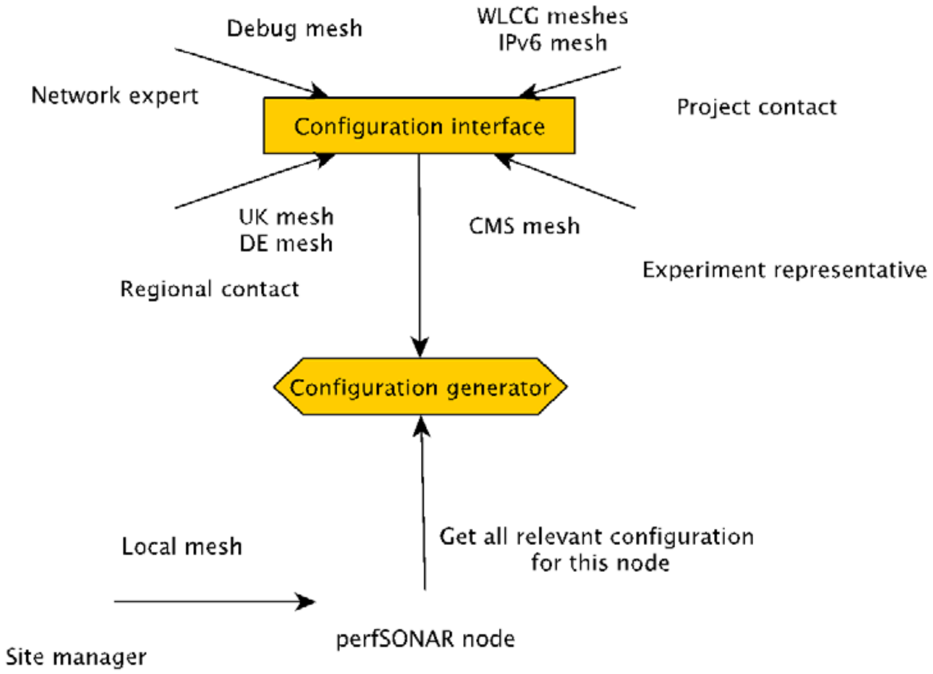

Figure 11: OSG centralized perfSONAR configuration system

In 2014, OSG improved upon this system by providing a secured GUI in MyOSG (see Figure 12) that could construct meshes based upon the perfSONAR registration information required by OSG and WLCG. All perfSONAR instances in OSG are required to be registered in OIM while all such instances in WLCG (not in OSG) are required to be registered in GOCDB. This allows OSG to centrally gather all needed information to create meshes for use by perfSONAR instances all over the world. Once created the meshes automatically update as registration information is updated.

One further interesting capability was enabled by OSG because of their mesh-management system: since OSG knows which perfSONAR hosts are participating in which meshes it is possible to have each perfSONAR instance configure a single URL (even if they participate in multiple meshes). We call this new URL the “auto-mesh” URL and is identical for all perfSONAR hosts except for the last part, which is the perfSONAR host’s fully qualified domain name, e.g., https://myosg.grid.iu.edu/pfmesh/mine/hostname/<FQDN>. This is very powerful in that now perfSONAR admins no longer need to update their configuration if meshes are added or changed.

In Fall of 2016 work on a new standalone version of the mesh configuration GUI was restarted, having languished for about a year. The OSG and perfSONAR developers agreed to fund the additional effort required to complete the work started back in 2015. The goal is to be able to replace the current mesh configuration with an improved version better integrated with the new features coming in the perfSONAR v4.0 release. OSG plans to migrate to using this version in production once it is ready. Longer term this software will be maintained by the perfSONAR team. An initial testing instance is viewable at https://meshconfig-itb.grid.iu.edu/

Using the mesh management system from OSG, we can now easily manage how, where and when tests are run. We need to walk a careful line of testing only enough to meet our needs. We have a tension between those who would like better test coverage vs those concerned about using available bandwidth to test.

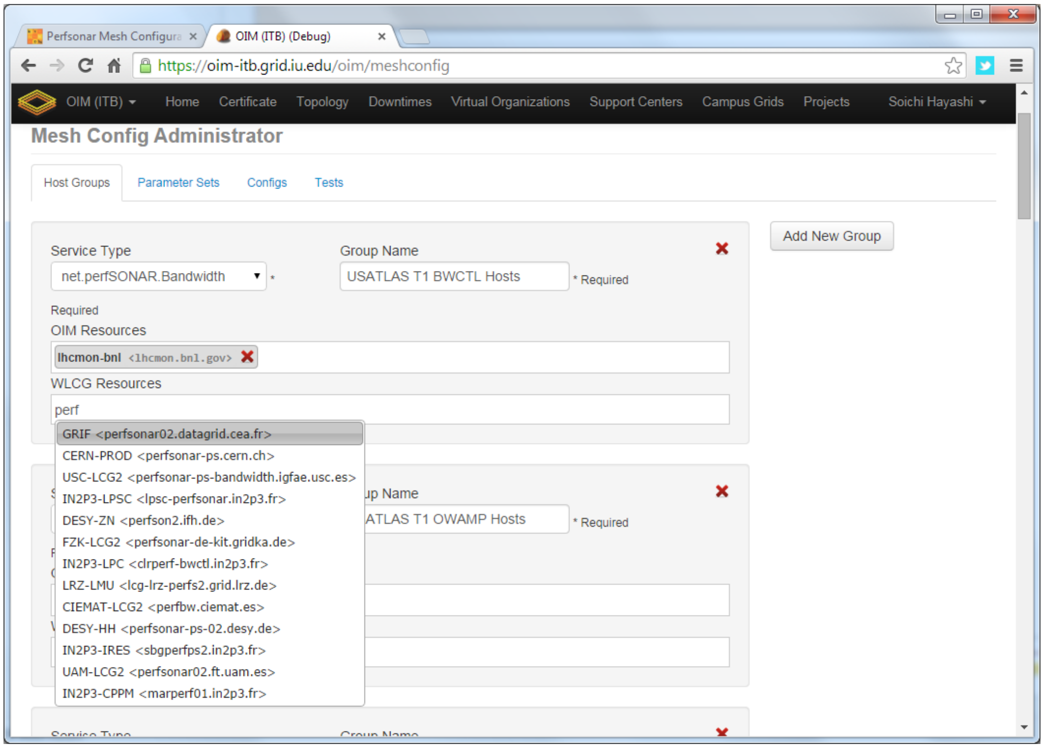

Mesh Config Administrator

Service Type: net.perfSONAR.Bandwidth *

Group Name: USATLAS T1 BWCTL Hosts * Required

Required

OIM Resources: lhcmon-bnl `<lhcmon.bnl.gov>`

WLCG Resources: perf

GRIF <perfsonar02.datagrid.cea.fr>
CERN-PROD <perfsonar-ps.cern.ch>
USC-LCG2 <perfsonar-ps-bandwidth.igfae.usc.es>
IN2P3-LPSC <lpsc-perfsonar.in2p3.fr>
DESY-ZN <perfson2.ifh.de>
FZK-LCG2 <perfsonar-de-kit.gridka.de>
IN2P3-LPC <clrperf-bwctl.in2p3.fr>
LRZ-LMU <lcg-lrz-perfs2.grid.lrz.de>
CIEMAT-LCG2 <perfbw.ciemat.es>
DESY-HH <perfsonar-ps-02.desy.de>
IN2P3-IRES <sbgperfps2.in2p3.fr>
UAM-LCG2 <perfsonar02.ft.uam.es>
IN2P3-CPPM <marperf01.in2p3.fr>

up Name: ATLAS T1 OWAMP Hosts * Required

Figure 12: Example of OSG's mesh-management tool interface

The latency test we have configured are managed by OWAMP and measure one-way delays between the latency node and its test partner at another site. We send 10Hz of small UDP packets to each test partner continuously (600 packets/minute). Since absolute time accuracy is critical for this test, part of a latency node configuration includes setting up a reliable time service (ntpd) configuration to ensure the node keeps accurate time. From this measurement we get the one-way delay to the partner site as well as information on any packet losses for each 1 minute interval (how many of the 600 packets were lost?). It is the packet loss measurement that is very sensitive to problems along the network path.

The second type of test measures throughput using Iperf3. Within a Tier-1 cloud mesh we schedule a 30 second throughput test in each direction (source to destination and destination to source) every 6 hours. In addition each end schedules it (both ways) so we end up with two 30 second tests each direction each 6 hours. We additionally are trying to sample ALL network paths but at a much lower cadence. To do this we have set up a WLCG-wide mesh which tries to run a 30 second throughput test each direction, once per week. Until we can determine the impact of this, we have limited the WLCG mesh to be the largest 50 sites (according to their published disk storage numbers). Starting with perfSONAR v3.5.0 we have been using BWCTL to schedule slots for bandwidth tests in order to avoid conflicts.

The last type of test is a critical one: traceroute. The traceroute test tracks the network path between the host and any destination and is run every 20 minutes to each destination which is being tested. If the route changes, we record the new route in the perfSONAR measurement archive. This is required to understand the topology we are measuring and can alert us to routing changes that may be correlated with observed network problems.

Metric Visualization

While perfSONAR provides a convenient way to gather standardized network metrics via deploying a toolkit instance, it can still be cumbersome to try to gather, check and interpret this data. As US ATLAS began deploying perfSONAR instances we realized that a critical missing component was a means of centrally monitoring and displaying the information we were collecting. ESnet had some initial efforts in this direction by creating Nagios[40] "plugins" that could query individual perfSONAR instances and check to see if the data returned was within bounds.

In 2013 new dashboard has been introduced called MaDDash, a project created and supported by ESnet. We now use MaDDash to monitor all our WLCG and OSG measurements while having it connected to the network datastore. The production instance[41] is shown in Figure 13. Colors indicate whether the metrics tracked are OK (green), WARNING (yellow), CRITICAL (red) or UNAVAILABLE (orange). MaDDash also supports "drilling-down" by clicking on the cell which will take the user to an interface with historical data, graphs and details of the test results.

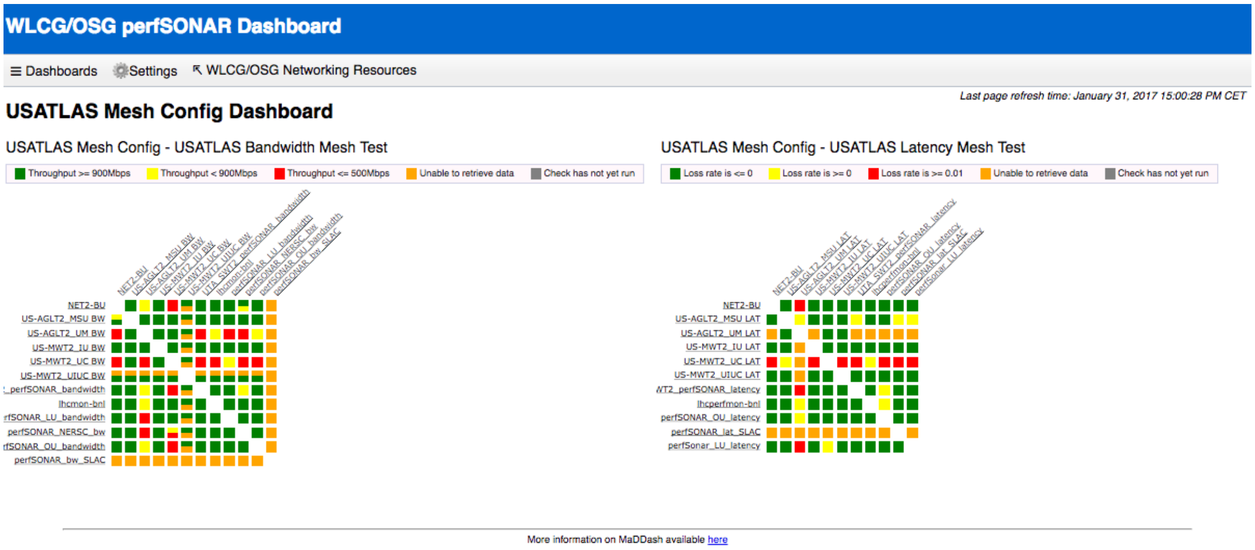

WLCG/OSG perfSONAR Dashboard

Last page refresh time: January 31, 2017 15:00:28 PM CET

USATLAS Mesh Config Dashboard

USATLAS Mesh Config - USATLAS Bandwidth Mesh Test

🟩 Throughput >= 900Mbps 🟨 Throughput < 900Mbps 🟥 Throughput <= 500Mbps 🟧 Unable to retrieve data ⬛ Check has not yet run

USATLAS Mesh Config - USATLAS Latency Mesh Test

🟩 Loss rate is <= 0 🟨 Loss rate is >= 0 🟥 Loss rate is >= 0.01 🟧 Unable to retrieve data ⬛ Check has not yet run

More information on MaDDash available here

Figure 13: MaDDash dashboard in the production instance showing the US ATLAS meshes for bandwidth and latency.

---

[40] http://www.nagios.org/

[41] http://psmad.grid.iu.edu/maddash-webui/

Infrastructure and Service Monitoring

In addition to metric visualization, the global perfSONAR deployment has another challenge. We need to be able to quickly find problems in the measurement infrastructure itself. While perfSONAR has evolved to be more robust over time, there are still cases where it has problems and fails to gather the needed metrics. To address this we have created a simple-to-deploy infrastructure monitoring system based upon OMD (Open Monitoring Distribution; http://omdistro.org/) which is a single RPM install of Nagios and many integrated applications. This has allowed us to quickly find infrastructure issues and better support end-sites when they have problems. We have created service checks for all perfSONAR services (by host type) as well as checks on registration information like admin name and email, latitude and longitude and program version. See Figure 14 for a view of OMD summarizing the WLCG perfSONAR host status.

Figure 14: OMD Check_MK screenshot showing some of the monitored perfSONAR hosts in WLCG.

OSG hosts the production version@ https://perfsonar2.grid.iu.edu/WLCGperfSONAR/check_mk

Changes in 2016 and Plans

During the year 2016 the platform has evolved to a stable production system that has been continuously available for experiments and end users. The following changes were introduced during the year:

- Following successful commissioning of the full mesh testing between all sites performed in 2015, new experiment specific meshes were introduced that helped us to increase the test frequency while decreasing the overall workload on the test endpoints. Currently, we have dedicated latency and bandwidth meshes for ATLAS, CMS, LHCb, Belle II as well as several regional meshes such as UK, US ATLAS, US CMS, etc. Each mesh is configured to perform continuous latency testing between all sites at 10 Hz (600 packets/minutes) as well as traceroute/tracepath tests every 30 minutes and throughput tests every 6 hours.

- We have established IPv6 testing at all T1 sites and started a campaign to add as many T2 sites as possible. A dedicated dual-stack mesh has been introduced to enable comparison of throughputs, latencies and network paths between IPv6 and IPv4.
- The total number of test endpoints increased slightly to 250 from 246 in 2015, there were new toolkits deployed in Ukraine, China and Turkey.
- New dashboard has been deployed in testing that can detect obvious issues with the infrastructure (endpoints not working correctly) as well as provide a summary report on the actual network performance issues (high packet loss, insufficient throughput, etc.)
- During 2016, OSG has open sourced a standalone instance of the mesh-configuration interface. This version has become part of the perfSONAR 4.0 and can be easily installed outside of OSG’s OIM infrastructure as it doesn’t require the OIM or GOCDB registration data to find perfSONAR toolkits to manage. Instead this version leverages the perfSONAR lookup service used by all installations, making it much more lightweight to deploy and use. The goal for this is year is to eventually replace the OSG OIM version with this one.
- One challenge for WLCG and OSG has been the data-lifecycle management tasks associated with stewardship of the ever increasing amount of network metrics we are collecting and making available. The dedicated OSG hardware storing the data will fill up by mid-2017. To address this issue we are pursuing two solutions: 1) purchasing new, much larger disks for the existing storage infrastructure, and 2) working with the perfSONAR developers to develop, test and maintain robust tools to handle data migration to other locations while maintaining access to that data. This second part is challenging because the perfSONAR development team is exploring replacing Esmond with another technology for the measurement archive. Thus we don’t expect to have such tools for at least another year.

Network Analytics and Diagnostics in HEP

Establishing the Network Monitoring Platform and making the data available for experiments and network researchers has triggered great interest from different communities that have started to look at the existing measurements and performed analysis with various different goals. At the same time, the platform has made it possible to diagnose and debug existing network issues, identify the problematic links or equipment and help fix the underlying problems, which has been the primary goal of the WLCG Network Throughput Working Group. Recently, the biggest efforts were focused on the network analytics covering the following topics:

- Real-time analytics to detect "obvious" issues with the networks as they arise, using both empirical methods as well as machine learning.
- Detecting how well network paths perform. This is based on computing a network cost-matrix for all existing links and using it to optimize job and data placement, job scheduling, etc., when there is a choice.
- Automated debugging of network issues and helping to find root causes in real-time (PuNDIT).

This chapter provides a high level summary of the most relevant activities and results in these areas.

WLCG Network Throughput Working Group

In order to ensure that sites and experiments can better understand and fix networking issues a working group has been established in late 2014 within WLCG, as part of the WLCG Operations Coordination, called Network and Transfer Metrics Working Group (recently renamed to Network Throughput Working Group). The working group meets on a monthly basis and follows up on the three areas of interest:

- Oversight of the perfSONAR network infrastructure - helping sites deploy and configure the endpoints as well as actively monitor the infrastructure and follow up on any operational incidents. It also provides WLCG/OSG specific support for perfSONAR deployments and coordinates the test parameters (frequencies, types of tests, etc.) to ensure efficient use of the infrastructure. The main goal of this activity is to gain visibility into how our networks operate and perform.

- Measure end-to-end network performance and use the perfSONAR measurements to single out on complex data transfer issues. At the end of 2015, the working group has established a support channel to work with sites and experiments on resolving the actual network performance issues. The main motivation was to use the Network Monitoring Platform and specifically WLCG/OSG perfSONAR network to debug complex network performance issues. This involved working with the sites and R&E partners to narrow down to a particular network link/segment and thus help identify and fix the root causes.

- Improve the overall transfer efficiency and thus our ability to fully utilize the existing network capacity by analysing the existing measurements and become pro-active in helping sites fix the existing network bottlenecks as well as promote network-aware design of the data management systems.

The working group has been also involved in organizing and contributing to the events and discussions on the future of the WLCG networking where significant challenges are foreseen. In the near-term, the 100 Gbit networks are midway in its (7-8 years) generation cycle while becoming more affordable, thus many sites plan to reach the 100Gbit uplink capacity by end of this year. This is likely going to increase the pressure on the R&E capacities in the near-term. In the longer-term, new technologies, such as those seen in DTNs, are paving the way to transfer large amounts of data and both HEP and non-HEP applications are preparing for transfer capacities at the LHC scale (or beyond). Sharing the future capacity will thus require greater interaction with the networks and could potentially lead to broader adoption of some form of the software defined networking. Finally, new infrastructure provisioning technologies, such as containers, have shown great potential in disrupting how the networks should be provisioned, configured and operated, thus impacting skills and effort required to manage the data centre networking.

Recent Progress and Plans

In 2016 the working group has directly contributed to the evolution of the Network Monitoring Platform by providing support and assistance in operating test endpoints to the sites, contributing to the test re-organization activities and registering new test endpoints as was summarized in Changes in 2016 and Plans. In addition, the working has been also active in the following areas:

- Network Throughput support channel that was started in late 2015 has seen its first year of operations. It was involved in several major cases, while most of them were due to issues at sites (configuration, insufficient or malfunctioning hardware, etc.), some cases have also involved R&E network providers. This was to some extend due to significant increase of the LHCOPN/LHCONE traffic, ESNet has seen LHCONE traffic growth of 113% YoY and GEANT reported 65% YoY [LHCOPN/LHCONE workshop], which was in turn caused by the exceptional performance and availability of LHC. During the year some of the LHCOPN links have become saturated and had to be upgraded. Finally, some cases also pointed at the sub-optimal connectivity for developed regions where performance over higher latencies (200ms) was very low due to poor selection of the network equipment.
- WG also contributed to the on-going discussions on the future of the LHC networks by presenting and contributing at various occasions including LHCOPN/LHCONE workshop, HEPiX, CHEP 2016 and WLCG workshop. It has also co-organized a dedicated one day workshop on the WLCG networking, where all four LHC experiments have presented their network requirements for both short and longer term.
- Innovative models of the simulated network throughput based on the perfSONAR measurements were investigated as part of the network analytics activities (for more details see Network Throughput Cost Matrix).
- Finally, WG contributed to running perfSONAR testbed to evaluate the next major perfSONAR release and prepared for global update campaign that will be one of the main activities of this year.

In 2017, we plan to further tune the network throughput support channel based on the feedback received from experiments and sites. We also aim to evolve the current infrastructure monitoring by adding site alerting/notifications (see Alarm and Alert Prototype) as well as its integration with the current operational availability policies of the experiments. Finally, we will be looking into adoption of the perfSONAR 4.1 to be released later this year, which will drop support for SL6 and will require all sites to re-install their test endpoints.

New and Ongoing Monitoring and Diagnostic Efforts in HEP

Most HEP users are not "network wizards" and don't wish to become one. In fact as pointed out by Mathis, the gap in throughput between what a network wizard and a typical user can achieve was growing significantly from the late 1980's to the late 1990's.

Within the last 15 years, because of improvements in default OS TCP stack settings, new protocols, hardware, firmware and software, this gap has significantly decreased but still remains in 2016. Because of HEP's critical dependence upon networks to enable their global collaborations and grid computing environments, it is extremely important that more end-user focused tools be developed to support these physicists and continue to decrease the gap between what an expert can achieve and what a typical user can get "out of the box".

In this report we have documented the effort in the HEP community to develop and deploy a network measurement and diagnostic infrastructure which includes end hosts as test points along end-to-end paths in the network. This is critical for isolating problems, identifying bottlenecks and understanding infrastructure limitations that may be impacting HEP's ability to fully utilize their existing networks. The history of the effort was documented in previous year's report and we won't repeat it here, instead highlighting updates during the last year in the following sections.

Network Analytics Platform

The ATLAS experiment, especially because of the involvement of the PANDA team with the ANSE project, has been actively updating their software to take advantage of network information in making higher-level decisions about workflows, data-access modes and data movement. Underlying this is getting all the relevant data easily accessible within the ATLAS information systems. During the last half of 2015, significant progress was made in incorporating the network metrics, gathered by OSG, into an analytics platform based upon ElasticSearch. This effort was led by Ilija Vukotic.

The idea was to create a network analytics service that indexes historical network related data while providing predictive capabilities for near term network throughput performance. Its primary functions are to:

- Aggregate, and index, network related data associated with WLCG "links"
- Combine perfSONAR data with information obtained from other sources e.g. file transfer service (FTS), derive near future throughput expectations for all the links.

- Serve derived network analytics to ATLAS production, DDM & analysis clients
- Provide a generalized network analytics platform for other communities in the OSG

Currently the system indexes data from four network performance/utilization measuring systems:

perfSONAR: Open Science Grid publishes all the measurements to Active Message Queue (AMQ) at CERN. Python collectors read AMQ messages and indexes part of the data in ElasticSearch.

FTS: File Transfer Service - the lowest-level data movement service doing point-to-point file transfers on behalf of Rucio. All the rucio events are sent to Elasticsearch using Logstash. While this provides all the details for each individual file transferred in this way, it does so only for ATLAS file transfers as only these are reported to Rucio.

FAX: Federated ATLAS storage system using XRootD protocol. Provides a global namespace, direct access to data from anywhere. We index two complementary types of data:

a) FAX cost matrix - obtained by continuously measuring file transfer rates using xrdcp from 64 storage elements to 20 largest ATLAS user analysis queues.

b) all the xrootd servers provide two monitoring streams (summary and detailed), each reporting every 60 seconds all the transfers, sources, destinations, rates, etc.

ESNET: We index rates measured by ESNET directly from their backbone devices, and provided via REST interface.

Shown in Figure 15 is the logical architecture that was developed to get all the relevant data into the envisioned analytics platform based upon ElasticSearch and Kibana.

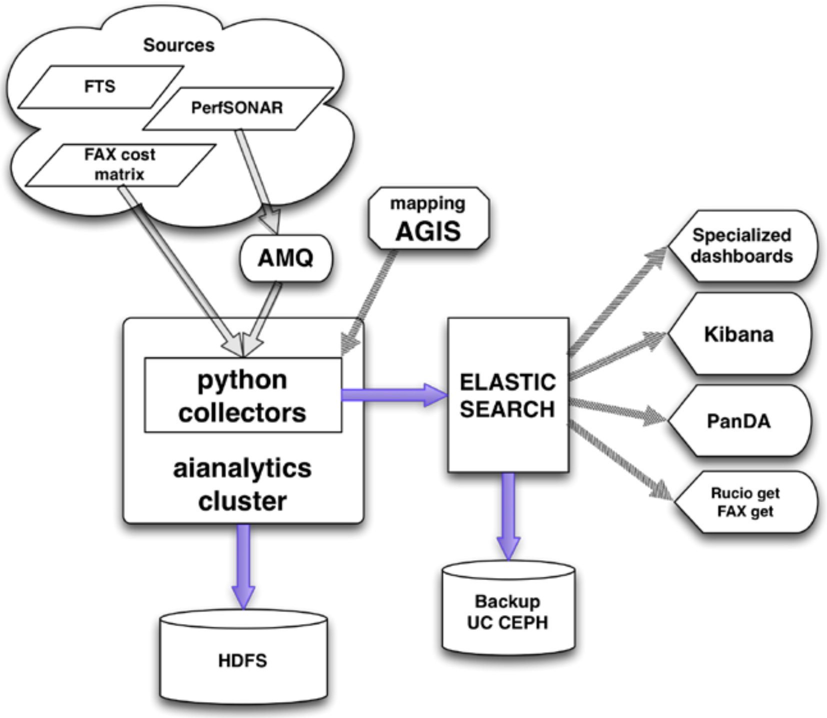

Figure 15: Network data analytics logical architecture

ATLAS has had this functioning since December 2015. The python collectors not only collect, and parse source information but also add mapping information loaded from AGIS. Collectors, (currently run on a single VM) were tested up to levels much higher than actually needed. The collectors can handle up to 24 h of ElasticSearch outage without dropping data. ElasticSearch index and search performance was measured to be higher than predicted needs of all our clients, both in terms of number of parallel requests and latency. The relevant data rates for the types of data gathered:

Data source	Data rate [MB/day]	Documents/day
perfSONAR	880	5M
FAX cost	15	120k
FTS transfers	50-150	200k – 600k

The ElasticSearch instance used by ATLAS is hosted by University of Chicago. Data are backed up on a weekly basis. As the same ElasticSearch instance hosts Panda and Rucio data, this presents an opportunity to examine and tune interaction of these system with the network. One more ElasticSearch cluster is being commissioned at CERN to be used exclusively by the production level ATLAS services. When ready, we will use that instance for alarm and alert services and University of Chicago instance for analytics, analysis and data mining tasks.

Users can access all the data and construct their own graphs and dashboards using Kibana[42], Elasticsearch frontend[43]. Timelion, a Kibana plugin for visualization and operation on time series can be used to combine information from disparate data source, either in the local Elasticsearch instance or available remotely (eg. graphite[44]), as shown on Figure 16.

Work is ongoing to test the impact of using this data as part of ATLAS operations. Various data enhancements (transforms and calculated values) are also being developed to augment the collected data. Additional sources of network information will be added starting with LHCONE flow data. This long term effort can be split in following subtasks:

- Data cleanup: validation of measurements validation, masking bad measurements, fixes.
- Set up of crude alarms/alerts: looking at local info (eg. single link), large time averages, simple cuts. Low sensitivity but high specificity test are sufficient at this step.
- Annotations of (ir)regular periods for supervised ML methods. Understanding of general correlations. Testing different models of anomaly detection. Testing performance/ free capacity of prediction models.
- Understanding of corner cases. Quick response, high sensitivity, high specificity, autonomous alarms/alerts.
- Integration: feedback loops of client services. Identification of problematic devices. Integration of all the sources.

Most of this work is performed on a Jupyter[45] server at University of Chicago. The server is a powerful machine with two NVidia GPUs, that make it very useful for all kinds of machine learning applications (e.g. training RNNs). Having all the modern ML libraries installed and being close (same local network) to the Elasticsearch cluster, make it easy to use and quite performant.

---

[42] https://www.elastic.co/products/kibana

[43] http://atlas-kibana.mwt2.org:5601/

[44] https://graphiteapp.org/

[45] http://jupyter.org/

Figure 16: A timelion visualization showing combined information from three different sources. Top: bandwidth seen by the sites border switch. Middle: bandwidth measurements reported by perfSONAR. Bottom: File Transfer Service (FTS) utilization.

Alarm and Alert Prototype

In the future we will have a number of algorithms (both in production and testing) that will generate alarms when a network performance anomaly is detected. We also need to store all the alarms for estimation of specificities and sensitivities, comparisons between different alarm generation models, and later reviews. Based on the alarms generated we need to send alerts to people that can do something about it.

A scheme of the prototype of the alarm and alert system is shown on Figure 17. It relies on the Jupyter instance described above, to periodically run both alarm generation algorithms and alert sending jobs. Currently people can (un)subscribe to alerts using a google form, while Jupiter jobs are scheduled using local cron.

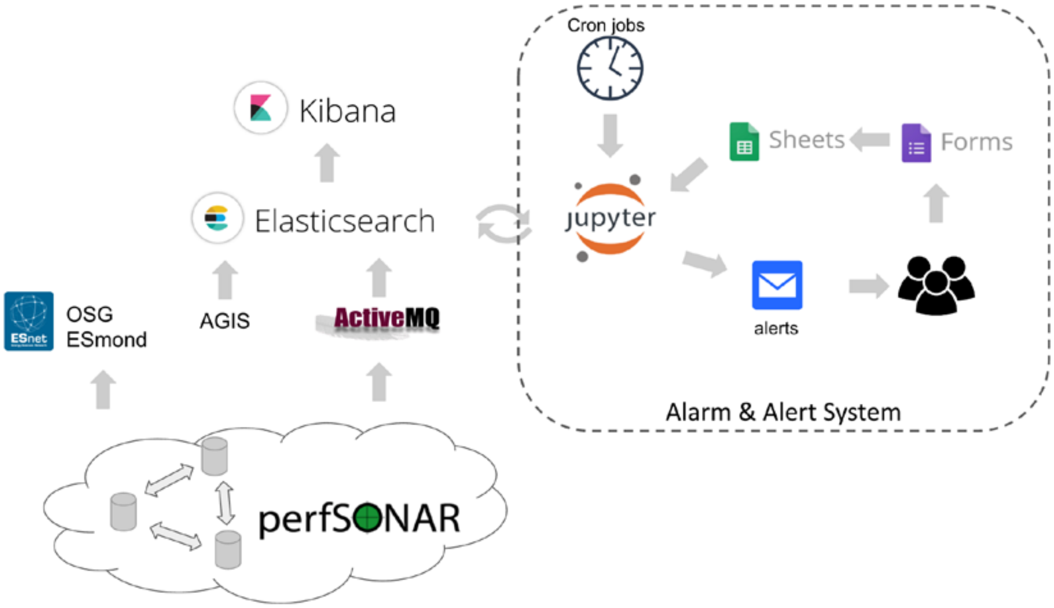

Figure 17: Schema of alarm and alert generation system.

In the future we will test running A&A jobs using Jenkins[46] framework and will need to establish reliable and up to date mapping of email addresses to parts of networking infrastructure they are responsible for.

Network Throughput Cost Matrix

WLCGs High Throughput Computing (HTC) model relies on network performance to transfer significant amounts of data between sites in a short period of time. Thus the primary focus is on the achievable network throughput. In order to obtain an overview of the performance for all existing links, a *network cost matrix* has been introduced to determine the status of all links between all sites. Computing the *network cost matrix* is grounded on the ability to determine the performance of a given network path, which is in turn determined by many different factors such as link capacity, latency, message transport unit (MTU), available buffers on the sending/receiving host, operating system, efficiency and available memory of the equipment along the path, etc. It therefore poses a substantial challenge.

From the network operations point of view, data on the link utilization can be gathered directly from the routers via certain protocols such as SNMP, NetFlow, SFlow. While this approach provides the best estimate of the actual traffic seen by the routers, it rarely gives an estimate of the actual end-to-end performance which involves also passive network equipment (switches) and well as both end hosts. In addition, particular the readout is dependent on the actual implementation and configuration at the site, which usually includes significant layer of virtualization and is therefore very challenging to obtain in a federated WAN environment, such

---

[46] https://jenkins.io/

as the one used by WLCG. The other option is to obtain an estimate by running a synthetic end-to-end network transfer between two hosts, with tools such as iperf, iperf3, nuttcp, etc. While this provides an accurate view on the available bandwidth at hand, the actual measurements, if done on a regular basis, will not exceed few per day given the number of links in LHCONE. Another popular approach among the experiments is to use the production data management systems directly and obtain an estimate of the achievable bandwidth from the production transfers[47], however since this usually involves disk-to-disk transfers and data management system itself (with protocol overheads, etc.), it's unclear what fraction of the measurement can be attributed to the actual network performance. In addition, there are usually several data management systems operated on the same link, competing for the same resources without any scheduling, which in practice leads to significant fluctuations in the measurements.

The investigated approach leverages from the perfSONAR measurements that can be done at high frequencies, such as one-way delay (OWAMP) measurement, which is measured by default at 10 Hz (600 packets/s) and was specifically designed to detect abnormal network behaviors such as packet loss (up to 0.0016%), packet reordering, jitter or significant path changes by detecting number of hops the packet has travelled. Since those are also the main attributes that determine the achievable throughput on the link, we have explored if we could combine them in a way that would allow us to see events when link is approaching its saturation. Sufficiently precise one-way-delay measurements would provide enough information on the devices along the path that are close to saturation and would either start filling their buffers and thus holding the packets for a little big longer or if out of buffers would start to drop the packets. Our primary focus was thus on detecting higher than usual latencies as well as any possible packet loss, which even at very low numbers (0.001%) is very detrimental to the achievable network throughput.

---

[47] As production transfers usually don't have regular pattern, coverage is improved by running additional "synthetic" transfers.

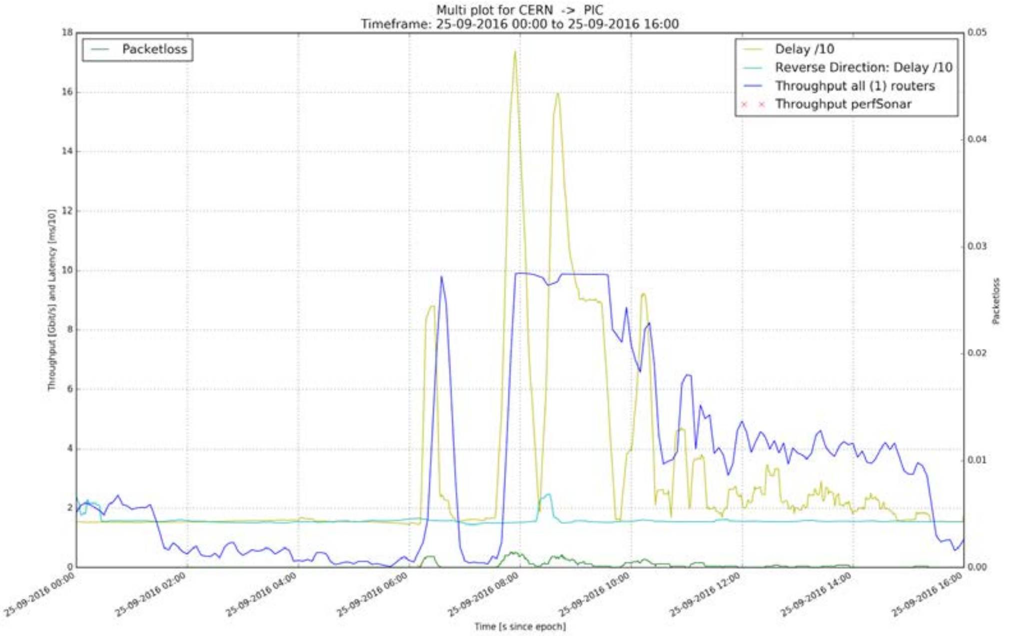

Figure 18: Plot showing combined perfSONAR latency/packet loss measurements with router traffic data gathered for LHCOPN link from CERN to PIC

The plot above shows part of the recorded data for the connection from CERN to the Tier 1 center "Port d'Informació Científica" (PIC). CERN and PIC are connected via a 10 Gbit/s dedicated optical link. In the plot the router traffic data is shown in five minute steps (blue) and data from perfSONAR (yellow, cyan, green) in one minute steps, both data sets have been smoothed via a moving average over 15 minutes in order to show tendencies clearer and filter high frequency noise.

As expected, delay (yellow/green) and packet loss (green) react strongly once the link utilization (blue) comes close to it's design maximum. Even though this reaction is strong, there is close to no warning in advance, that a congestion is about to happen. Similar patterns but with different strengths can be seen on all connections within LHCOPN.

Based on those insights, we have designed and developed a machine learning model trained using the existing LHCOPN router traffic and perfSONAR measurements. We have also developed and put into production a real-time publisher of the simulated throughput for all LHCONE links, which is available as part of the network stream together with other measurements from the Network Monitoring Platform.

We plan to continue tuning the current model to better fit with the complexity of the LHCONE and start looking into integration of the simulated throughput within the job scheduling and the data management systems.

PuNDIT

The PuNDIT project (NSF grants 1440571 and 1440585) targets analyzing perfSONAR latency measurements in near real-time to identify common network issues. The project intends to address two capabilities that are not possible using the perfSONAR toolkit but often requested:

1. **Identification of network problems in near real-time**
2. **Localization of the source of the problem within the network**

PuNDIT will build upon the de-facto standard perfSONAR network measurement infrastructure to gather and analyze complex real-world network topologies coupled with their corresponding network metrics to identify possible signatures of network problems from a set of symptoms. For example if the symptoms suggest a router along the path has buffers configured too small for high performance, Pythia will return a diagnosis of "Small Buffer". If symptoms indicate non-congestive packet-loss for a particular network segment, the user can be notified of a possible "Bad Network Segment". A primary goal for PuNDIT is to convert complex network metrics into easily understood diagnoses in an automated way.

To provide these capabilities the project will augment perfSONAR toolkit latency measurement nodes to analyze the raw packets used in the latency measurement for signatures of typical network problems. Each 5 second interval is studied for problem signatures for each network path being measured for latency. Each minute, all 5 second windows that indicated a symptom are reported via RabbitMQ to a central PuNDIT server. The central server analyzes all the reports it receives from all participating toolkit instances and correlates them to help localize the possible source of each problem.

The PuNDIT project is finishing up in 2017 and intends to provide two RPMs: one to install on any participating perfSONAR toolkit instance measuring latency and one to create the central PuNDIT server (including GUI). RPMs are intended to be available about a month or two after the release of perfSONAR v4.0. The assumption is the major high-energy physics collaborations already using perfSONAR would install PuNDIT to provide near real-time analysis of their latency data to expose problems in their network infrastructure.

Summary: Progress in HEP Network Monitoring for 2016

As noted above, efforts to ensure commonality in both monitoring and provisioning of networks have seen a significant increase in activity in HEP during 2015-2016. Originally, the GLIF[48] and DICE[49] communities were both working toward implementing “managed” network services and the corresponding monitoring that will be needed to support their efforts. A significant amount of this kind of exploratory work is now taking place in the LHCOPN/LHCONE communities.

To summarize, 2016 has seen new efforts to build-upon and benefit-from our standard perfSONAR network monitoring infrastructure. The goal is to have a single infrastructure providing network related metrics for HEP and the tools and services to make use of those metrics in beneficial ways. This is feasible because HEP’s use of the network is qualitatively the same between HEP collaborations. It is desirable because network providers and users don’t

---

[48] “Global Lambda Integrated Facility”, see http://www.glif.is/

[49] “DANTE-Internet2-CANARIE-ESnet collaboration, see http://www.geant2.net/server/show/conWebDoc.1308

want to have multiple “network measurement infrastructures” making redundant (and possibly interfering) measurements nor do they wish to have to develop, deploy and support many such instances when one will do.

Comparison with HEP Needs

Previous studies of HEP needs, for example from the Trans-Atlantic Networking Report (http://gate.hep.anl.gov/lprice/TAN/Report/TAN-report-final.doc) focused on communications between developed regions such as Europe and North America. In such reports packet loss less than 1%, vital for unimpeded interactive log-in, is assumed and attention is focused on bandwidth needs and the impact of low, but non-zero, packet loss on the ability to exploit high-bandwidth links. The PingER results show clearly that much of the world suffers packet loss impeding even very basic participation in HEP experiments and points to the need for urgent action. In addition the more detailed monitoring via perfSONAR shows that even many of our high-performance network links in nominally “good” areas of the world suffer from packet loss at a level that can significantly impact HEP dataflows.

The PingER throughput predictions based on the Mathis formula assume that throughput is mainly limited by packet loss. The 15% per year growth curve in Figure 10 is somewhat lower than the 79% per year growth in future needs that can be inferred from the tables in the TAN Report. True throughput measurements have not been in place for long enough to measure a growth trend. Nevertheless, the throughput measurements, and the trends in predicted throughput, indicate that current attention to HEP needs between developed regions could result in needs being met. In contrast, the measurements indicate that the throughput to less developed regions is likely to continue to be well below that needed for full participation in future experiments.

Recommendations

There is interest from ICFA, ICTP, IHY and others to extend the monitoring further to countries with no formal HEP programs, but where there are needs to understand the Internet connectivity performance in order to aid the development of science. Africa is a region with many such countries. The idea is to provide performance within developing regions, between developing regions and between developing regions and developed regions.

We should strive for >=2 remote sites monitored in each major Developing Country. All results should continue to be made available publicly via the web, and publicized to the HEP community and others. Typically HEP leads other sciences in its needs and developing an understanding and solutions. The outreach from HEP to other sciences is to be encouraged. The results should continue to be publicized widely.

We need assistance from ICFA and others to find sites to monitor and contacts in the developing and the rest of the world, especially where we have <= 1 site/country. A current list of countries with active nodes can be found at http://www-iepm.slac.stanford.edu/pinger/sites-per-country.html.

Work on high performance monitoring of the network using perfSONAR must continue. It is critical that we provide a robust toolkit deployment that requires minimal local administrator maintenance. The value of these deployments depends upon the broad-scale deployment so we have the metrics to identify and localize network problems. Beyond just having a resilient infrastructure operating worldwide, we need to make additional progress in alerting and alarming when specific network problems are found. This will maximize the value of our deployments for all HEP users.

Acknowledgements

We gratefully acknowledge the following: the assistance from NUST SEECS in improving the PingER toolkit and management has been critical to keeping the project running, with respect to this we particularly acknowledge the support of their leader Arshad Ali; and the students and lecturers who have assisted including recently: Umar Kalim, Anjum Navid, Raja Asad Khan of NUST SEECS who helped in updating some of the graphs, the case studies on Africa and Pakistan and implementation of PingER tools such as TULIP. Mike Jensen provided much useful information on the status of networking in Africa[50]. Alberto Santoro of UERJ provided very useful information on Latin America. Sergio Novaes of UNESP and Julio Ibarra of Florida International University provided useful contacts in Latin America. We received much encouragement from Marco Zennaro and Enrique Canessa of ICTP and from the ICFA/SCIC in particular from Harvey Newman the chairman. Also the leaders of the Malaysian collaboration: Johari Abdullah of UNIMAS, Adib Monzer of UUM and Saqib Khan of UAF and their teams have made major contributions. More recently over the last three summers, we have received help from three SLAC summer interns from Brazil. Unfortunately with the changes in government in Brazil last year, we do not expect this to continue. Also Les Cottrell is reaching retirement.

We must also not forget the help and support from the administrators of the PingER monitoring sites worldwide as well as all the perfSONAR site managers, the Open Science Grid and the WLCG working group members. Last, but not least, we would like to acknowledge the significant support the National Science Foundation has provided to perfSONAR and to various related networking projects.

Appendices

Appendix A: ICFA/SCIC Network Monitoring Working Group

The formation of this working group was requested at the ICFA/SCIC meeting at CERN in

---

[50] Mike Jensen, "Connectivity Mapping in Africa", presentation at the ICTP Round Table on Developing Country Access to On-Line Scientific Publishing: Sustainable Alternatives at ITCP, Trieste, October 2002. Available http://www.ictp.trieste.it/~ejds/seminars2002/Mike_Jensen/jensen-full.ppt

March 2002[51]. The mission is to: *Provide a quantitative/technical view of inter-regional network performance to enable understanding the current situation and making recommendations for improved inter-regional connectivity.*

When the ICFA/SCIC Network Monitoring team was assembled in 2002, the lead person for the monitoring working group was identified as Les Cottrell. Les assembled a team of 5 others to assist in preparing the report. By the end of 2015 only Les and Shawn McKee were still active and Les was unable to continue as team leader, in part because of the lack of funding for PingER.

For 2016 a new team needed to be put into place. Shawn McKee and Marian Babik agreed to take over assembling the report and a new team was recruited:

Table 7: Members of the ICFA/SCIC Network Monitoring team

Shawn McKee	Michigan	OSG, USATLAS, WLCG	smckee@umich.edu
Marian Babik	CERN	WLCG, CERN	Marian.Babik@cern.ch
Les Cottrell	SLAC	US, PingER	cottrell@slac.stanford.edu
Soichi Hayashi	Indiana	Pervasive Tech. Institute	hayashis@iu.edu
Brian Tierney	LBNL	perfSONAR	bltierney@es.net
Henryk Giemza	NCBJ	LHCb	henryk.giemza@ncbj.gov.pl
Ilija Vukotic	Chicago	ATLAS	ivukotic@uchicago.edu
Mike O'Connor	BNL	ESnet, LHCONE	moc@es.net

The working group wants to explicitly acknowledge Les's excellent leadership and the exceptional work that he has put into 15 years organizing and writing these annual reports. We hope he will be able to continue his participation for future reports!

Goals of the Working Group

- Obtain as uniform picture as possible of the present performance of the connectivity used by the ICFA community.
- Prepare reports on the performance of HEP connectivity, including, where possible, the identification of any key bottlenecks or problem areas.

---

[51] "ICFA/SCIC meeting at CERN in March 2002". Available http://www.slac.stanford.edu/grp/scs/trip/cottrell-icfa-mar02.html