perfSONAR Workshop Report

July 8-9, 2010
Arlington, VA

This material is based upon work supported, in part, by the National Science Foundation (NSF) under Grant No. CNS-1019008. Any opinions, findings, conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the NSF.

This material is based upon work supported, in part, by the Department of Energy. The views and opinions expressed in this material are those of the author(s) and do not necessarily reflect those of the United States government or any agency thereof.

Introduction

On July 8-9, 2010, Internet2 hosted the first perfSONAR Workshop, focusing on the perfSONAR Network Measurement Infrastructure. The workshop, held at the Crystal Gateway Marriott in Arlington, Virginia, was sponsored through a grant from the National Science Foundation and a supplemental grant from the Department of Energy. Information about the workshop is publicly posted at: http://www.internet2.edu/workshops/perfSONAR/

Background

perfSONAR is an extensible, standards-based network performance monitoring middleware infrastructure, that not only facilitates the ability to solve end-to-end performance problems on paths crossing several networks but also enables network-aware applications. This infrastructure supports performance data collection and exchange between multiple networks using well-known protocols and formats. perfSONAR was designed and developed internationally, by technologists within the research and education networking community. perfSONAR is based on schemas developed in an open standards body, the Open Grid Forum. More than 68 entities around the world, ranging from backbone networks to universities to government laboratories have deployed perfSONAR. For more information, see: psps.perfsonar.net (Perl-based effort in the U.S., including links to use-cases) and www.perfsonar.net (overview of the international project).

The Large Scale Networking (LSN) Coordinating Group of the Networking and Information Technology Research and Development Program (NITRD) worked with the National Science Foundation's (NSF) Directorate for Computer and Information Science and Engineering and the Department of Energy's (DoE) Office of Science, Advanced Scientific Computing Research programs to fund an international perfSONAR Workshop to engage researchers, applications developers, network managers, and tool developers in a discussion of the uses to which perfSONAR can be put, future goals for the infrastructure, and methods by which members of the community can work together to ensure perfSONAR is widely deployed and used.

The workshop, which involved 90 attendees (roster posted at: http://www.internet2.edu/workshops/perfSONAR/roster.cfm) representing 57 separate organizations around the world, included international participants from Brazil, Germany, Switzerland, Japan, Spain, England, Poland, Norway and Canada. R&E networks represented include: Dante/GÈANT, ESnet, Internet2, RNP, PSNC/PIONIER, SWITCH, CEFET-MG, UFES, APAN, DFN-Verein, and UNINETT.

Researchers benefitted by presenting relevant research projects, hearing operational and engineering needs, and learning about the protocols and system as a distribution channel for research results. Operators and engineers benefitted through exposure to research results and by articulating engineering challenges and research needs, and learning how perfSONAR might meet some of those needs.

A stated goal of the workshop was the creation of an ongoing community of interest coupled with a report summarizing the current community consensus on potential research focus areas, deployment strategies, open issues, and operational requirements.

The Process

In late January 2010, a number of stakeholders were identified and invited to participate on the Workshop Executive Committee. The first meeting of this group was held in conjunction with the Winter 2010 ESCC/Internet2 Joint Techs in Salt Lake City; a Birds of a Feather (BoF) session was held to include any interested parties at an early stage, to answer questions on process, and to solicit feedback from the community as to what should be covered, who should be invited, and what outcomes reached.

Shortly after Joint Techs concluded, members of the Executive Committee began meeting weekly to establish a timeline, identify potential speakers and invitees, frame the agenda, and craft a Call for Proposals (CFP) to ensure participants would be bringing a range of expertise to share. The CFP was issued in early April; members of the Executive Committee identified a wide range of mailing lists to use, thereby ensuring word of the event was widely disseminated. Internet2 hosted web pages for the event, linking information as it became available (i.e., agenda, rosters and participant requests for inclusion/proposals for presentations).

When many of the responses included a request to present material, the Executive Committee decided to allow each of those who requested to give a talk (21 individuals) a short presentation opportunity, which could be bolstered by posters. The posters were on display throughout the first day of the meeting and a pre-dinner poster event was held with great success. The Executive Committee identified five general topic areas for these presentations: Wide Scale Deployments, Implementation Barriers, Research Uses/Findings, End-to-End Performance Problems, and Deployment Experiences. Several members of the Executive Committee volunteered to chair these topic areas, working with the selected speakers to form a coherent topic message for the group.

The Executive Committee decided that the agenda should include both presentations from attendees and a significant block of time for breakout sessions on topics of interest to the perfSONAR community: Community Building, Research Uses & Findings, Operational Issues, and Expanding perfSONAR Technology. The group decided that every breakout session should provide information on some general questions as well as session-specific ones:

- What are the open questions?
- Which are missing definition?
- Which are missing software?
- What are the Best Common Practices?
- Are the existing options sufficient?

Carla Hunt (MCNC) led the Community Building session, aided by Roberto Sabatino (DANTE) and Grant Miller (NITRD). This session was held in the main meeting room because the Executive Committee felt this was a topic of interest to all but one that would not garner sufficient participation if held at the same time as the more participant role-specific groups.

Malathi Veeraraghavan (UVA) led the Research Uses & Findings session, aided by Mark Crovella (BU) and Jose Augusto Suruagy Monteiro (RNP). This was a very well-attended side session; the group covered a wide range of topics but the consensus was that they had merely

begun good discussions; very few reported that they had reached consensus on many issues. In any future perfSONAR Workshops, this topic would need to be included in the breakout group discussions.

Joe Metzger (ESnet) led the Operational Issues session, aided by Ann Harding (SWITCH), E. Paul Love (NOAA/NCO), and Yasuichi Kitamura (APAN). This group was unique in that participants largely reported that consensus had been reached on many of the key issues. In a future iteration of this workshop, participants may want to determine how to implement the resolutions the group made and consider developing some of the materials (Best Common Practices guide, etc.) recommended.

Martin Swany (UDel) led the Expanding perfSONAR Technology session, aided by Brian Tierney (ESnet) and Jeff Boote (Internet2). In this session, the smallest of the groups, many topics were raised and the group was in complete agreement that consensus had not been reached. This group voted overwhelmingly for continuing the discussions they started at this meeting.

Jason Zurawski (Internet2) offered a 4-hour pre-workshop tutorial to participants not familiar with perfSONAR. This session was both very well-attended (55 participants; roster included in Appendices) and very well-received by participants. Tutorial materials are posted at: http://www.internet2.edu/workshops/perfSONAR/materials.html

Acknowledgements

Internet2 would like to thank the Executive Committee for participating in all the calls/planning:

> Jeff Boote (Internet2), Aaron Brown (Internet2), Grover Browning (University of Indiana), Rich Carlson (DoE), Mark Crovella (Boston University), Susan Evett (Internet2), Dale Finkelson (Internet2), Victor Frost (NSF), Ann Harding (SWITCH), Carla Hunt (MCNC), Yasuichi Kitamura (APAN), Paul Love (NOAA/NCO), Joe Metzger (ESnet), Grant Miller (NITRD), José Augusto Suruagy Montiero (RNP), Roberto Sabatino (DANTE), Martin Swany (University of Delaware), Brian Tierney (ESnet), Malathi Veeraraghavan (University of Virginia), and Jason Zurawski (Internet2)

Internet2 would especially like to thank the co-authors of this report:

> Eric Boyd (Internet2), Joe Breen (University of Utah), Aaron Brown (Internet2), B. Ann Cox (Laboratory for Telecommunications Services), Susan Evett (Internet2), Tom Hacker (Purdue University), Ann Harding (SWITCH), Carla Hunt (MCNC), Yasuichi Kitamura (APAN), Joe Metzger (ESnet), Inder Monga (ESnet), José Augusto Suruagy Montiero (RNP), Malathi Veeraraghavan (University of Virginia), Matt Zekauskas (Internet2), Artur Ziviani (RNP), Lixia Zhang (University of California at Los Angeles), and Jason Zurawski (Internet2)

The community would like to thank the workshop sponsors – the Large Scale Networking (LSN) Coordinating Group of the Networking and Information Technology Research and Development

Program, the National Science Foundation's Directorate for Computer and Information Science and Engineering, and the Department of Energy-Office of Science's Advanced Scientific Computing Research Group – and their staff – Victor Frost (NSF), Grant Miller (NITRD), and Rich Carlson (DoE) – who worked closely with the Executive Committee throughout the planning process for their leadership and vision in this area.

Breakout Report: Community Building

Chair: Carla Hunt (MCNC); Committee: Roberto Sabatino (DANTE) and Grant Miller (NITRD)

Significance of Community Building

Community is an important component of perfSONAR. perfSONAR was designed and developed through international community collaboration, by technologists within the research and education networking community. perfSONAR is based on schemas developed in an open standards body, the Open Grid Forum. More than 68 entities around the world, ranging from backbone networks to universities to government laboratories have deployed perfSONAR. Continued growth of the community, in terms of both people and deployments, will: (1) Strengthen the framework, services and tools encompassed by perfSONAR, and (2) Enable better research by ensuring the needs of researchers and scientists are met.

In preparation for a discussion about community building at the workshop, we conducted an analysis of where the community is today. The data sources for our analysis included: (1) meta data from the list of active perfSONAR measurement endpoints that were registered with the global lookup service as of the first day of the workshop, and (2) meta data from the registered list of workshop attendees.

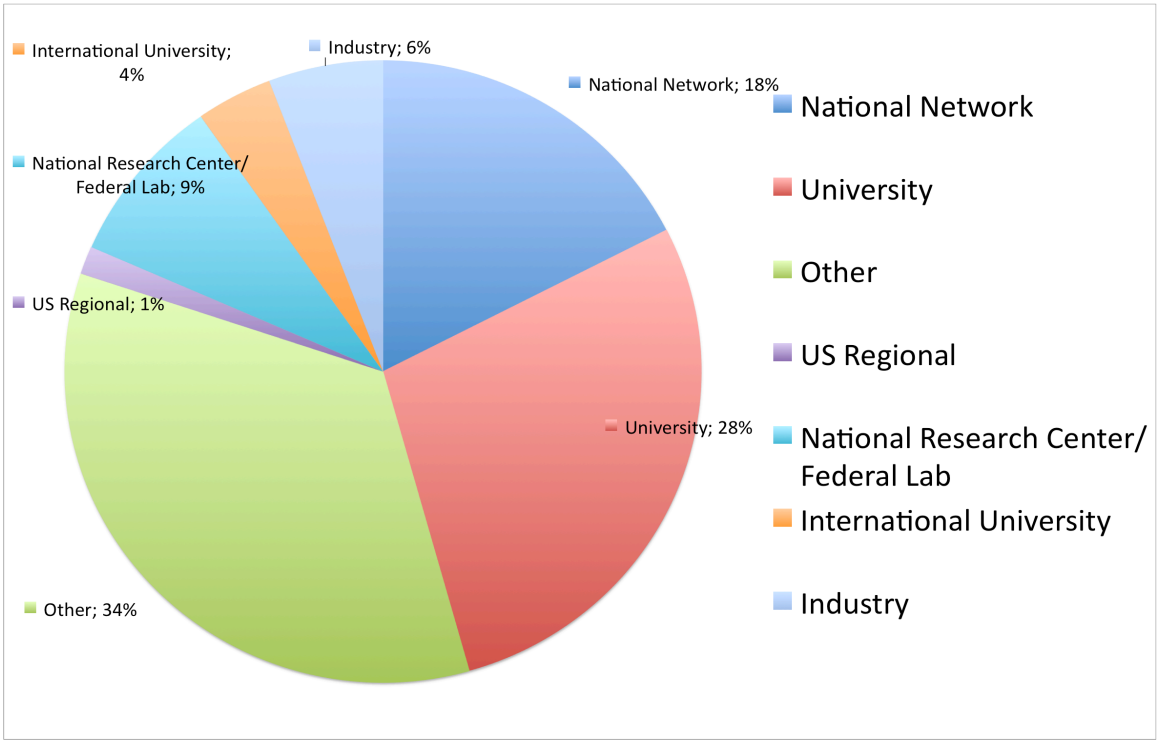

We discovered 208 measurement endpoints were registered with the lookup service. We reviewed the list of endpoints and classified the endpoints as belonging to one of the following categories of institutions: National Research Center or Federal Lab, National Research Network (US or International), U.S. Regional Network, International University, U.S. University, Industry, Other. Endpoints were classified in the "Other" category if the name of the endpoint was not intuitive enough to place it in one of the preceding categories.

We also reviewed the same list of endpoints and classified the endpoints according to geographic location. The following geographic locations were identified: Argentina, Australia, Brazil, Canada, Chile, China, Guatemala, Hong Kong, Europe, Japan and the United States. Approximately 15% of endpoints were located outside of the United States.

We conducted a similar geographic analysis of the list of attendees registered for the workshop: 19% of attendees were international or from locations outside the United States and 81% of attendees were from the United States.

Approach to Community Building

The community building breakout session was the only breakout session where everyone attending the conference was invited to participate together. The other three breakout sessions ran concurrently; therefore, participants were asked to choose one to attend. On the morning of the first day of the workshop, we introduced the community building breakout session. We shared the demographic analysis of measurement endpoints and workshop attendees described above. Then, we asked participants to begin thinking about this breakout session as an opportunity to engage in strategic planning for the perfSONAR community. Together, we would talk about where we wanted the perfSONAR community to go and how we (as a community) would be able to get there. We asked the participants to review and consider some specific questions in preparation for the community building breakout session we had scheduled for that afternoon.

To capture the community voice and to encourage honest feedback, we collected written feedback during the afternoon community building breakout session. We taped eight poster size pieces of paper with provocative straw man goals or questions around the room. Each participant was given a "sticky note" and was asked to provide "one good idea." Stickers were available to "vote for" or emphasize any idea with which the participant was in agreement. The questions were:

- *Where is the perfSONAR community today?*
- *Grow the deployment footprint by 20% in 2010. How?*
- *Grow the community by 20% in 2010. How?*
- *Should we have another workshop?*
- *Expand the availability of tools that support operation & growth of the community. How?*
- *Ensure Community Is Organized. How?*
- *Ensure the needs of researchers/scientists are met. How?*
- *Where (else) do we want the perfSONAR community to go?*

That night, we transcribed written feedback from the sticky notes into an online document shared with all participants. Some clear themes emerged during the process, which are described below. During the final session of the workshop we reported our findings in terms of these themes.

Recommendations

We collected over 70 distinct written ideas from the community building breakout session. Raw input from the notes is available online. The following themes emerged from our synthesis of the

raw data: (1) We should continue to grow the community. (2) We should continue to improve documentation and online resources for perfSONAR. This was requested by a number of different groups within the community including application developers, network operators, and researchers. (3) In general, we should make things easier; for example, providing software packages for easier installation. (4) We should provide more workshops and more tutorials with hands-on components. Specifically:

About Workshops

- Yes. Have more workshops!
- Include hands-on component to the tutorial.
- Keep short presentations.
- Increase International Participation.
- Pre-post slides to encourage more discussion.
- Consider separate tracks for implementation and policy.
- Expand the poster sessions.

About Community Growth

- Yes. Grow the Community!
- Develop innovative outreach activities to connect with those who may not have heard about perfSONAR.
- Update the web presence.
- Consider developing a portal to target different audiences.
- Consider developing a matchmaking service or clearinghouse/knowledge base listing of tools, deployments, projects, researchers, and success stories.
- Develop more documentation that is clear and centralized.
- Capture and disseminate more statistics on perfSONAR usage.

About Tools

- Offer more workshops and tutorials for developers/tool builders.
- Better documentation is needed on how tool developers can integrate their tools into perfSONAR.
- A centralized repository of documentation is needed.
- Make packages available for another Linux distribution.

About Tools – the Lookup Service

- Better lookup, search and discovery are needed.
- Some services may not be registering with the Lookup Service. How can we address this?

About Community Organization (lightest feedback)

- Organize mailing lists by community type (research, deployment, etc.).
- Continue organization through Internet2.
- Start a perfSONAR group, like a GLIF group.
- Encourage collaboration and communication between different groups.

About meeting the needs of Researchers/Scientists:

- Build a knowledge base.
- Build tools together.
- Engage the researcher in test-bed development and operation.
- Engage the researcher to get requirements.
- Highlight success stories.
- Convene user-oriented workshops to identify their requirements. Participation in these workshops would attract new users and user communities.
- Maintain a presence at major research conferences, offer perfSONAR tutorials, and publish success stories.
- Provide standardized data for researchers.

Where (else) do we want the community to go?

- Engage the commercial community.
- Add support for IPFIX/Netflow.
- Find out if there are things we need to measure that we are not currently measuring.
- Promote performance information sharing by users.

In addition to the results described above, some member of the community offered additional ideas verbally on Community Building. In terms of software packaging and testing new software contributions, a representative from the Open Science Grid (OSG) suggested it may be useful to consider leveraging the Virtual Data Toolkit (VDT) and OSG community's testing and packaging infrastructure. A couple of interesting ideas for increasing campus deployments included: tighter integration between perfSONAR and existing popular network management tools such as Cacti and Nagios, and increased focus on topology or path analysis capabilities.

Also, the workshop evaluations that were collected after the workshop provide a nice baseline for planning future workshops. Some highlights of respondent feedback include:

- 51 total evaluations (of 90) turned in; 8 of 12 Expanding pS Technology group, 15 of 20 Operational Issues group, and 26 of 31 Research Uses/Findings group
- 40 participants registered for the Tutorial; with room reconfiguration, we were able to include 55 participants
- Workshop attendees represented: 35 network engineers/operators, 30 researchers, 19 tool developers, and 6 (+2) government staff or meeting support personnel
- Travel support was provided to 7 International participants (2 from Germany and 5 from Brazil) and 12 U.S. participants (3 network researchers, 2 network operators, 1 tool developer, and 6 graduate students in network research)
- Areas of general consensus were: a) Workshop should be an annual event (95.6%), b) amount of time allowed for presentations was “about right” (82%) and c) amount of time allowed for breakouts was “about right” (75%)
- Areas of consensus among breakout group respondents were: a) my breakout group “started good discussion” (73.5% overall but only for the non-Operational Issues group – which had 66% “reached consensus” for same question and only 33% “started good discussion”), b) splitting breakouts onto 2 days was good (97.8%) and c) continuing discussions on these breakout topics in future years is important (76.1%)

- Respondents generally approved of the Tutorial; 48.6% for it “informative” and 42.9% found it “useful” (52.1% attended the tutorial and 73% of those who did NOT attend accessed the Tutorial online)

Conclusion

We collected some valuable feedback, not only during the community building session but also throughout the workshop and, at the end, using workshop evaluations. Activities the community is interested in pursuing include hosting additional workshops and tutorials, developing an information repository, and generally simplifying perfSONAR deployment and use.

To determine how to make using perfSONAR easier, it is instructive to review the specific findings from the other breakout sessions.

Breakout Report: Operational Issues

Chair: Joe Metzger (ESnet); Contributing authors: Joe Breen (University of Utah), Ann Harding (SWITCH), Yasuichi Kitamura (APAN), and Joe Metzger (ESnet)

Introduction

The goal of the Operational Issues breakout session was to characterize what the community should do to foster useful ubiquitous perfSONAR deployments. According to the Global Lookup Services available in the perfSONAR nodes, many Research and Education networks, Government networks, International networks, Regional Networks, University networks and even some Commercial networks have deployed perfSONAR services. But the question is, how to expand these deployments? How to increase the use and the different tools?

The current state of the art of perfSONAR deployments is a mish-mash of varying degrees based on time, skill-sets, security policies, and access to resources. Great strides have happened in making the software framework an easy bundle to deploy but work remains in several areas such as:

- security,
- ease of remote configuration/upgrades,
- usability

Outreach and PR activity is one obvious answer to increase the perfSONAR footprint but is not sufficient to grow deployments without a proper understanding of the positioning of perfSONAR within different communities. It is possible that current deployments based on interest in existing toolsets and technologies alone are reaching a plateau. Yet perfSONAR is a flexible and scalable infrastructure likely to be of value to further sites within the existing deployment groups and even beyond. It is, therefore, necessary to base further outreach on an understanding of the non-technological factors related to deployment, specifically making a business case for perfSONAR within different communities, developing Best Common Practices so that perfSONAR services as well as the technology can scale, addressing security and data management issues in a way that is compatible with each environment, and ensuring ongoing ease of operation and management so that perfSONAR services are sustainable in the long term.

The community is also in need of additional documentation and Use Case scenarios describing what groups can support and what they have permission to release to others.

Deployments

The initial European involvement in perfSONAR, via the GÈANT2 project, was focused on research into network monitoring. perfSONAR provided a means by which multiple organizations could share information in a standard format, increasing the understanding of an end-to-end path. European implemented perfSONAR components are primarily Java based but include some managed services based on dedicated hardware and some Perl-based tools. Towards the end of that project and continuing into the current iteration, GÈANT3, the primary perfSONAR priority is to combine existing toolsets with support and service definitions to deliver monitoring services in an R&E environment.

There are two approaches to this. For the needs of the LHC OPN a private, managed deployment of perfSONAR took place. Tier 0 sites were provided with a full deployment, including hardware (two devices per site plus GPS antenna) and operational support to actually run the deployment. Custom development on visualization was also provided in the initial deployment.

The second approach concentrates on federated delivery of monitoring service. Participating NRENs (National Research and Education Networks) will deploy and operate a more lightweight set of their own perfSONAR instances, export a core agreed set of metrics and allow access to the data to other operators as regulated by a policy encouraging open sharing. This facilitates a consistent service among multiple, independent organizations. An important strategic principle of this federated deployment is also to increase the value of the service by ensuring global interoperability with equivalent services. Key to a successful federated perfSONAR service over different regulatory environments is good authentication and authorization (AA), both in terms of policy agreed and underlying infrastructure to support this. AA to support these services requires similar principles to that required by a multi-domain monitoring infrastructure - to enable sharing while preserving domain independence and control over their own data and to scale over multiple sites. European perfSONAR deployments will leverage the eduGAIN service to deliver this.

In the case of the Asia Pacific area, there are two kinds of activities. One is the annual workshop of the human resource development (HRD) supported by DANTE. The workshop invites the operators or future operators of Trans-Eurasian Information Network (TEIN). These 1.5-day workshops involve a hands-on and, in these past three years, also a "getting started" style program. Some of the attendees have installed the perfSONAR infrastructure at their NRENs; ThaiREN (Thailand) and VINAREN (Vietnam) started their own perfSONAR infrastructures. India is very interested in joining this community.

A second type of activity is the documentation of the installation by one of the APAN Tokyo XP operators. The document is available at http://www.jp.apan.net/NOC/perfSONAR/. This is a password-protected document; registration helps the Tokyo XP operators to support the newcomers to perfSONAR use.

Security

Security is a top concern for many entities that wish to deploy the perfSONAR framework. For many entities, the integration of perfSONAR into the local security model is NOT a trivial matter. Security concerns can be divided into two main areas – security of the data gathered or accessed via perfSONAR and security on a systems- and application-level of the deployments themselves.

For data security, one key area is authentication and authorization. Questions arise regarding who has access to the data and what are the governing policies regarding the sharing of the data. In an End-to-End measurement situation, whose security applies to the data sets? Security teams want to see full Authentication, Authorization and Accounting (AAA) for those using the tools so that an audit trail exists in case of misuse of the tools. Privacy Compliance teams have concerns regarding the identification of traffic flows within the measurement archive. This issue

grows more complex as the network protocols migrate from IPv4 to IPv6 where the end-host MAC address may be imbedded in the Global Unique Address.

For Systems/Application security, policies of some entities wishing to install perfSONAR can require a security audit of the application, a demonstration or available documentation on what means and mechanisms are in place to ensure the application does not become a threat to the local network or become an attack point for other networks.

Questions arise regarding how secure the measurement applications are, what kind of patch management/update facility exists and what kind of data the measurement archive stores. Security teams want to ensure that the measurement devices do not turn into attack devices. These teams want to understand the risk and how network measurement groups will mitigate the risks. Security and Privacy Compliance teams want to verify that no data leaks out regarding traffic flow content.

Other entities may prefer a 'black box' approach, where it is not required to understand the means by which perfSONAR deployments are secured but it is important that a third party organization can be held accountable for this by means of a contract.

Security for many entities will govern how the respective entities will deploy the perfSONAR framework, if at all, and to what extent the entity will participate in the global perfSONAR community.

Operations and Maintenance

The Operations and Maintenance of the perfSONAR framework requires thought by those would potentially use it. These end users break into three large Use Case scenarios: regional network/backbone network operations, campus network operations, and Scientists. The last use case has two shades to it because network research scientists have slightly different requirements than a scientist who is just using the tool to verify/qualify a network.

Regional & Backbone Operators

Backbone networks typically have a profile whereby they are operating a wide range of services beyond monitoring. To do this, they frequently operate internal best practice in server deployment and systems administration so that they can achieve economies of scale over their full service delivery and can have dedicated, separate systems and network teams – both of whom can be involved in deployment decisions. Supported operating systems tend to be Linux variants – in Europe the dominant distributions are Debian/Ubuntu. Deployments must generally be managed remotely, manageable by the organization's own staff with standard system administration tools and best practices. These organizations may put their own effort into integrating perfSONAR with their own tools and workflows if they are also a target user of the perfSONAR service. Typical use of perfSONAR is in troubleshooting, service quality verification and network planning. Support requirements are predominantly software support (i.e., bug fixing, upgrades, documentation, etc.).

Regional Network Operation Centers and Backbone Network Operation Centers (NOCs) face similar issues, requirements and uses for the perfSONAR framework. One of the first hurdles in existing Network Operation Centers is the plethora of existing tools and how to integrate another

suite of tools into the existing menagerie. For this integration to be successful, the Network Operators must find the perfSONAR implementation to be a low requirement entry. The business case must be simple or pre-built so that Operators do not really have to argue their case of why another suite of tools. The installation requirements of the perfSONAR framework must be low, easy to maintain via scripts and must have configuration tools available to configure/update several simultaneous installations. These configuration tools should support mechanisms for both security and software updates.

A core strength of perfSONAR is flexibility. It empowers deploying organizations to implement their own policies regarding collection and sharing of information. This flexibility also has a flip-side whereas it is not possible to enforce a standard deployment by technical means which means delivery of a predictable end-to-end performance picture is a significant challenge. However, as backbones and large campuses share particular characteristics and requirements, the creation of and buy-in into a core set of practices can resolve this.

A detailed Best Common Practices document should be available for help in standardization among groups and for help for varying levels of skill. These Best Common Practices should identify the support requirements and how to best deliver for various constituents. Two examples of how to deliver to constituents include: (1) Delivery of Measurement as a Service where the Network Operation Center may deploy and manage remote measurement points, (2) End-user managed measurement points where the end entity takes ownership for resources and management. These Best Common Practices should also incorporate or tie to documented Use Case scenarios and deployment scenarios in the style of "cookbook" recipes. These recipes should provide detailed, useful information such as common parameters, etc. The recipes should also provide information regarding authorization and how to tackle some of the security issues identified by different groups.

For success in the Regional and Backbone NOCs, the framework must have a source code license that multiple organizations can support. The software should also support various features such as: multiple interfaces for support of different networks, different views for different users, fault management, and alarming for data and perfSONAR warnings.

Primary organizations providing developers of existing perfSONAR deployments come from the backbone community. As such their organizational core competencies mean they are well placed to continue development effort to improve the security and management mechanisms for perfSONAR toolkits.

Campus Network Operators

Larger campus networks can resemble backbone networks in having dedicated teams, existing policies and large-scale deployments to manage. However, campuses operate in an even more complex environment than backbones in terms of service provision e.g. they may also provide campus e-mail, ldap, calendaring services etc. and therefore may have less flexibility in deployment choices. Issues of privacy and security are also more difficult to manage due to the greater number and variability of users, especially in mixed-discipline campuses.

Smaller campus networks do not have the scale of services of larger ones, but they frequently have to provide the same range of services, albeit to a smaller footprint of users and they have to

do this with a smaller headcount of predominantly generalists. In this context, outsourced management of services becomes attractive as expertise can be bought without increasing headcount. Where a service is not outsourced, support requirements from a provider of perfSONAR software increase - more advanced installation support, configuration help, regular security notices supported by help in upgrading. Deployments are not necessarily geographically remote and where the campus operates their own infrastructure, use of CD or DVD distributions can be helpful. Some campuses may find the 'appliance' route desirable – where perfSONAR is a black box conforming to their policies.

Campus Network Operation Centers (NOCs) have many of the same issues as the Regional and Backbone Network Operation Centers. The Campus NOC has to face the hurdles of integrating the software, developing a business case, and supporting the software on an ongoing basis. However, the Campus NOC has the additional issue of how to fit a measurement framework of this nature into a campus security policy. Campus security policies tend to have more detailed application specific nuances than Regional or Backbone NOC security policies. Campus security policies also have to deal with the end hosts themselves and how the end hosts access or expose data on the Internet. Campus organizations with a particular scope e.g. those in governmental or semi-governmental organizations in particular may have restrictive data sharing policies for information sharing. perfSONAR needs to allow them to share what they wish with who they wish remotely and with a verifiable audit trail. The Best Common Practices for deploying a measurement framework on campus must also accept and reflect these nuances and additional requirements. Use Case scenarios and deployment scenarios should also reflect the parameters and information that campuses use. Specific cases for campuses are, therefore a distinct requirement.

Use of perfSONAR by the campus is similar to that of the backbone. However, smaller campuses may not have the direct resource to put into integration and development tasks and the backbone developers do not have the same understanding of the needs and concerns of this segment and it is not within their mandate to scale support requirements in depth to support all end users. This provides an opportunity for the growth and development of consultancy services for perfSONAR deployment, management and integration outside of existing partners. It his therefore necessary to choose software licensing options which do not restrict the growth of small (or large) business around perfSONAR but still preserve the open principles of the original concept.

Application Scientist/End Users

Application Scientists and End Users should be able to use the software to the extent they need or choose. Their uses may vary from deploying something simple for quick testing and verification of their location, to something as sophisticated as a Network Operations Center for continuous actively measuring a remote research lab. Application Scientists and End Users will typically emphasize the ease of use and ease of installation in their Use Case. They will often only have a small subset of the requirements that a campus NOC might need with almost no resources to accomplish the task. The ability to easily deploy in a virtual machine for some "quick and dirty" testing and simple measurements is a very useful variant for their implementation. The Best Common Practices and deployment case scenarios should reflect these simple cases and stress how to use the applications for end use testing to the desktop.

The profile of researchers regarding perfSONAR use depends on the position of perfSONAR within the research program. Some researchers will need to operate long term, reasonably complete and wide scale perfSONAR deployments similar to those operated by campus networks and where they control the information collected and shared.

Other researchers may need to have access to perfSONAR data from existing deployments for a particular ongoing purpose. The focus, then, is appropriate user interfaces and reliable delivery of data that is clearly understandable. (This doesn't mean 'simple'; rather 'complete'). Such interfaces must be: a) quick and easy to install and set up and b) are not necessarily remote.

Finally, researchers may have sporadic needs to 'dip in' to perfSONAR for short-term goals. Again, any tools or user interfaces need to be easy to find and both quick and straightforward to set up or the user will simply look for an alternative tool. Thirty minutes could be considered a generous timescale.

Data Management

Data Management is key to the community being able to utilize perfSONAR measurement data beyond simple real-time test scenarios. Long-term data management and archiving is necessary to discover trends for operations, mine data for research and allow the community as a whole to truly utilize and interpret the data. Data curation will be necessary as data sets grow to make sure that relevant data stays around for an appropriate time period for research. Proper architecture of the data and the data stores is necessary to allow ease of migration of the data to new boxes as hardware/software upgrades occur. The proper architecture will also allow rapid mining of the data for important clues in operational issues and research into changes over times and other possibilities. To share this data, the community requires sample data sharing policies and more discussion regarding data provenance issues. Legal retention policies differ from group to group and require some consideration and recommendations for basic policies.

Best Common Practices

The perfSONAR framework is very flexible. It is possible to deploy the tools in a manner that makes it easy to collect useful measurements, compare them with measurements from other domains, and develop meaningful analysis of the results. It is also possible to deploy it in a fashion where collecting useful measurements is very difficult, the data can't be compared across domains, and most of the artifacts in the results are unrelated to the underlying network performance.

The community needs to develop a set of Best Common Practices that describe how perfSONAR should be deployed in a way that meets the needs of Backbone Network Operators, Campus Network Operators, End Users/Scientists and Network Researchers. These Best Common Practices documents should be customized for each user community in a manner that makes it easy for the different types of users to select, install, configure, operate and use perfSONAR tools.

The Best Common Practices should identify appropriate sets of measurement capabilities for the different types of networks and the user classes they support.

The Best Common Practices should describe minimal common sets of measurement schedules and parameters, provides lots of value to the community, or:

- Break down by use cases of Backbone Network Operations, Campus Network Operations, Scientist, Network Researcher
- Minimal sets of functions, measurements, measurement schedules
    - Reasonable consumption of resources for measurements
- Deployment use cases
- Pointers for scientists on how to interpret results, tools and support
    - The case specification of the perfSONAR application will be helpful. APAN involves the telemedicine activities by the medical doctors. They want to use the high-resolution video transferring system.

Public Relations

- Documenting success stories - business drivers
- Way to cite perfSONAR data sets – what citations are appropriate?
    - Recommend work collecting case studies/success stories
        - Framework for collecting case studies (See Nagios case studies website.)
        - Summarizing them
        - Identify what data was collected for what use
        - Documentation of the value provided.

Recommendations and Conclusions

The breakout session identified a number of recommendations for the community that would facilitate deployment and enhance the usefulness of perfSONAR.

Security:

- The breakout session participants recommend that the community members continue to deploy perfSONAR in as open fashion as possible, while acknowledging that some data cannot be widely shared.
- We commend the perfSONAR-PS development community on recent progress that simplifies the process of integrating a perfSONAR deployment within a local security environment, and encourage more work in this area.
- We feel it is important that effort be made to integrate perfSONAR with existing widely deployed authorization and authentication frameworks.

O&M:

- The next major feature requested by the breakout session participants is the ability to integrate perfSONAR into network management systems, such as Nagios by supporting alarming capabilites.
- The session concluded that it is critical that funding be provided to build the tools and support the people that provide security & software updates, software packaging and simplifying installation and ongoing maintenance processes.
- The community should continue to work on refining measurement infrastructure business cases and support models. This effort should focus on providing examples that can be used by others to aid the adoption of perfSONAR, while being sure to identify innovative research opportunities for small or disadvantaged businesses.

Data Management:

- The community needs to do more research into data provenance & curation (especially as applies to research use) & legal retention policies.
- The community needs to do more research into analyzing the data for both research & operational purposes

Documenting Best Practices:

- Developing Best Common Practices that describe recommended measurement schedules, and hardware & software configurations will make the infrastructure more useful because it will lead towards a common set of data being available to network operators and researchers to analyze.

PR:

- Developing a framework for collecting and managing case studies that describe how perfSONAR protocols and tools have been deployed and used in the campus, regional, and backbone environments would facilitate additional deployments by demonstrating the value of perfSONAR deployments.

Acknowledgments

We would like to thank the members of the community who participated in the Operational Issues breakout sessions.

References

The Workshop Summary presentation is posted at:
http://docs.google.com/present/edit?id=0AWUgAuvCA9rlZGRtc2tndndfN2Rid3JtOWdy&hl=en

Breakout Report: Research Uses/Findings

Chair: Malathi Veeraraghavan (UVA); Contributing authors: Eric Boyd (Internet2), B. Ann Cox (LTS), Tom Hacker (Purdue University), Inder Monga (ESnet), José Augusto Suruagy Monteiro (RNP), Malathi Veeraraghavan (UVA), Matt Zekauskas (Internet2), Artur Ziviani (RNP), Lixia Zhang (UCLA), Jason Zurawski (Internet2)

Introduction

The current primary use case and initial driver for perfSONAR is to support network operations. Therefore, a key area identified for research is in the domain of diagnosing faults and identifying the source of performance problems. Data correlation and other techniques can be used to help develop autonomous systems that can pinpoint a problem as it happens without the intervention of engineers. Tools developed from these techniques would fit into the perfSONAR architecture at its top Analysis/Visualization layer.

Two additional use cases where the research community could benefit and contribute were identified for perfSONAR: (1) usage of the collected data to inform network research, e.g., future internetwork design, and (2) design and implementation of perfSONAR-aware applications, e.g., GridFTP clients request current status information from perfSONAR measurement points to choose the "best-available'' paths for high throughput.

All three Use Cases would benefit from research that adds new, and or better, data collection tools to the perfSONAR framework. Data collection tools form the bottom layer of the three-layer perfSONAR architecture, and the aforementioned Analysis/Visualization layer forms the top layer. This breakout group discussed tools that fit both these layers. However, research contributions to the middle – Measurement Infrastructure – layer of the perfSONAR architecture, which includes services such as lookup, topology, service configuration, authentication/authorization were not discussed at length in this breakout group, as this topic was handled by the Expanding perfSONAR Technology breakout group.

The rest of this section is organized as follows. Using a bottom-up approach, "Research on Data Collection Methods and Tools" describes open issues and lists major recommendations for research on data collection methods/tools (the bottom layer of perfSONAR architecture), and "Research on Analysis/Visualization Methods and Tools" similarly describes issues and recommendations for analysis/visualization methods and tools (the top layer of perfSONAR architecture). "General Recommendations" lists a set of general recommendations that apply to both sets of tools. "Research for R&E Operations" addresses research issues related to the primary use case listed above; "Research Use of Data" and "Application Use of Data" address research issues related to the two additional use cases listed above. "General Topics" groups three topics together as these are general issues that apply to perfSONAR as a whole. These include scaling and archiving, privacy, and tie-into GENI and other frameworks. Finally, the section concludes with major conclusions listed in "Conclusions".

Research on Data Collection Methods and Tools

There have been considerable research efforts to develop data collection techniques and tools. For example, over 400 tools are available through the MOnitoring and MEasurements (MOME) website [R1].

Internet topology inference and routing

Internet topology inference has been the subject of much research over the last ten years. perfSONAR can be used in two ways: (1) to provide additional vantage points for inferring topology, and (2) to provide ground truth.

Additional Vantage Points

Latency or traceroute data can be used as input to algorithms. With more measurement points, you have additional points to probe, and more data to infer topology. Where there is duplication with other infrastructure (for example, PlanetLab nodes), comparison of the utility of the two infrastructures can be made.

Provide Ground Truth

perfSONAR already has some knowledge of network topology, which is used to locate test points or relevant measurement archives. It also shares some topology with the dynamic circuit network services. The internal knowledge can be detailed, or it may be abstract. Even when the knowledge is detailed, it might prefer to export an abstracted topology, as is done for the dynamic circuit network service when sharing topologies with neighbor networks. Gathering these direct representations of topology would allow for direct comparison of ground truth with algorithmic results.

In addition to internal knowledge of topology, perfSONAR could be used to export IGP and BGP information kept by some networks (and shared with projects like RouteViews and the RIPE Routing Information Service). This information can be used to generate a complete current network topology, and in addition, can be used to understand when underlying IP networks (when there are underlying IP networks) recompute their topologies.

Any new perfSONAR documentation in support of research should describe how to extract perfSONAR's own view of network topology. perfSONAR should be extended to support exporting network routing data (BGP and IGP information) and, networks should be encouraged to export routing information via perfSONAR.

Physical-layer measurements

The GENI project is supporting research to define measurements at the physical layer for optical links. DHS is interested in such measurements for disaster recovery reasons. To make a consistent set of measurements available to network operators, perfSONAR schema and techniques for collecting and storing these measurements have to be designed. Data gathered from this set of tools could be useful input for cross-layer analysis/visualization tools that enable root cause analysis to isolate sources of errors.

Host/Application monitoring

It would be useful to obtain monitor host applications and instruments, and obtain corresponding measurements through the perfSONAR framework. Examples include GridFTP logs, Unidata logs, etc. These should provide an end-to-end view, which includes CPU, memory, and disk related performance metrics. Can cluster and grid monitoring systems such as Ganglia [R2] be integrated into perfSONAR? First-generation applications could self-report statistics into the perfSONAR framework. Second-generation applications could integrate perfSONAR clients for

dynamic adaptation to live conditions. The perfSONAR Event Demon (pSED) in the GENI LAMP project (see "Tie in to GENI and other frameworks" below) is contributing to this topic.

Dynamic circuit services/L2

An increasing amount of research is currently focused on multi-layer networking and automated provisioning of circuits/virtual circuits. DOE has funded a network research program on cross-layer provisioning. Hybrid networks, packet-optical convergence, and Carrier Ethernet are other trends in the marketplace. There are multiple flavors of dynamic circuits, e.g., wavelength-routed paths, SONET/SDH circuits, Ethernet VLAN based "L2 circuits." Measurements are required to support troubleshooting and network fault correlation across dynamic circuits in both intra- and inter-domain scenarios.

After a circuit is established, call backs can be used to trigger data collection for performance measures, such as throughput and latency, across the circuit. It may not be feasible to collect such measures for all circuits, in which case a sampled set can be used, from which average, minimum and maximum values can be archived. Parameters characterizing circuits, such as number of hops and circuit rate, should be archived along with collected measurements. Other measurements of interest include call arrival rates, rejection rates (with causes such as congestion), call durations, and lead times of requests (how much ahead of the required service time are book-ahead requests received). Frequency of failures and/or performance problems on dynamic circuits should also be archived.

At Layer-2, since a traceroute-like capability is not built into Ethernet switches, external tools can be implemented to provide operators with path information. This is invaluable for trouble-shooting.

The On-Demand Secure Circuits and Advance Reservation System (OSCARS), and other schedulers implemented for dynamic circuit services [R3] and perfSONAR, use common services, such as lookup service and topology service.

Security related tools

Network traffic anomalies denote significant and unusual changes in traffic patterns on one or more network links. Such anomalies do not necessarily imply malicious activities directly related to security issues, but rather any out-of-the-ordinary events. There are many causes of network traffic anomalies, including, for example, Distributed Denial of Service (DDoS) attacks, changes in IP forwarding tables due to mistakes in router configurations, device failures, and modifications in the Border Gateway Protocol (BGP) routing policies. Accurate and efficient anomaly detection and identification are crucial to tackling network anomalies correctly. Detection consists of determining when and where anomalies occur in network traffic. Identification involves the classification of detected anomalies, possibly indicating security risks. Diagnosing anomalies in network traffic thus presents great challenges due to the variety of possible anomaly types to be considered and the large volume of network information to be processed, usually demanding the adoption of sampling techniques. This also raises challenges on the feasibility of triggering data gathering in real-time as soon as an anomaly is detected.

A network anomaly of particular interest, due to its prevalence and potential damage to network usage, is that caused by a Distributed Denial of Service (DDoS) attack. Typical defense against DDoS attacks consists of three steps: (1) detection of an attack, which is usually done by intrusion detection and prevention systems such as network anomaly detection systems; (2) identification of the route(s) taken by attacking packets, which can be undertaken by IP traceback systems; and (3) mechanisms to block attacking packets at key points along route(s). It is a research challenge to achieve deployable and operational tools and services that thoroughly address these issues, in particular IP traceback to allow distributed attack traffic filtering.

A tool for multi-hop packet tracking with coordinated sampling, described in [R4], was recommended as a useful security tool.

Research on Analysis/Visualization Methods and Tools

Analysis tools

There are numerous papers that describe correlating network measurements on several axes, e.g. metrics of different types, path characteristics, timestamps and trouble ticket systems. Tools are being created, but this research area is of critical value to operators. An ideal solution would be an autonomous system that can pinpoint the source of problems. A prime example is a problem we term "find the closest measurement point." This would be useful in identifying performance problems on a particular path.

DANTE has used Performance Enhancement & Response Teams (PERT), which are groups of experts consisting of individuals selected from different provider organizations, who can jointly troubleshoot problems that require inter-domain cooperation to resolve. A presentation on PERT organization and approach [R5] notes "They work on an as-and-when basis," i.e., the teams solve problems as they arise, and work together until they are resolved. Such teams could be excellent sources for researchers interested in developing automated tools to meet this need.

Assessment tools

Assessment tools are a layer above analysis tools. They examine results from analysis and generate an assessment of the network comparable to a dashboard that indicates the health of the network so that a non-expert could quickly determine if the network is performing well or not.

These tools could greatly improve usability. For perfSONAR to have a measurable impact on practice, it must be usable by, and useful to, average IT staffers who are not necessarily networking experts or "wizards." An example case involving the Network for Earthquake Engineering Simulation (NEES) Program was cited. It required IT personnel across 14 U.S. universities to diagnose and fix poor end-to-end network performance. Many of the problems encountered were the standard host tuning problems, but some of them were not. If these IT personnel could install and use perfSONAR as a dashboard tool that could notify them when there is a performance problem, or they could use it to check performance (speedtest.net is a an example), more enterprise IT staff could and would use perfSONAR, which consequently would have a larger impact on practice. Some of the problems encountered in this recent example of diagnosing and fixing end-to-end networking problems in NEES were encountered 10 years ago in the NIH Visible Human Project. While perfSONAR has made significant advances toward

resolving these problems, addressing usability would significantly increase its usage base and correspondingly, its value.

General Recommendations

- Develop and publish a step-by-step procedure on the perfSONAR websites (perfSONAR-PS and perfSONAR-MDM) for researchers interested in contributing new data collection methodologies or tools to the perfSONAR framework. For new tools that propose to collect data of a type that is already collected, e.g. one-way latency, new schema definition is not required as it is for tools that propose to collect new types of data. Once this is done and the tools are developed using the perfSONAR schema and protocols, the next steps include submission of open-source software for posting on the perfSONAR websites, identification of a base set of servers on which to run the service for active measurement tools, registering the new service with appropriate lookup servers, and publishing archival policies (e.g., data deletion policies). Transformation tools that sanitize data for privacy and other reasons should be submitted along with the new data collection tools, and mechanisms put in place for an automatic execution of these tools on collected data at regular intervals.
- The community should discuss and decide on an organization to use as the "keeper" of perfSONAR schema and protocol standards. The OGF has been the vehicle for perfSONAR standardization up until now, but are there other candidates? For example, the Climate Studies community manages its own set of standards.
- Determine and publish a policy on tool validation procedures. Tool development is in the purview of the whole perfSONAR community, while service deployment is controlled by operators. First, tools have to be validated as providing accurate measurements. This can be done locally with an open-source effort, where researcher A could test a tool developed and posted on the perfSONAR website by researcher B and submit a "Thumbs-up" vote, or provide feedback to researcher B for further improvement. A website (perfsonar.net) is available for this purpose. Tool developers are encouraged to provide confidence intervals along with mean values of measured data. This open-source approach to tool validation is one of several approaches, which should be discussed, and the chosen policy posted on the perfSONAR websites. Once the accuracy of a tool has been validated, service providers can independently decide whether or not to deploy a new service using this tool.
- If perfSONAR is viewed as a framework into which others may contribute capabilities/tools for obtaining measurements, performing experiments (e.g., running network tests), and analyzing performance, each of these tools will either need to be modified to work with perfSONAR, or there could be a definition of a framework, schemas, and standards to which the tools must conform to allow them to "snap-in" to perfSONAR with few modifications. This conceptual model is similar to the approach used in the community climate system model (CCSM), which is described in [R6]. CCSM provides a framework and data standard that supports community development of simulation and analysis modules (e.g., land model, ocean model, cloud model) that operates within a complete earth climate simulation framework that can be used as a whole.
- A workflow model can be used to characterize the perfSONAR framework. As perfSONAR has data collection, analysis, and visualization components, the data and

control flows amongst these components must be well defined. One approach to do this is a haphazard definition-on-the-fly as components are written and integrated. Another approach is to understand that the framework is a workflow, in which there are processing steps, temporal ordering between processing steps, and well-defined data interfaces between the components that need to be defined.

Research for R&E Operations

The group identified a set of topics to pursue to benefit the operations of research and education (R&E) networks. These topics, for the most part, emphasize or give an operations view on the open research issues mentioned previously.

Tools for non-expert IT staff

As noted in "Assessment Tools" above, new assessment tools that can be used by non-expert IT staff are needed to solve problems and allow end users to utilize the full network potential. perfSONAR-based assessment tools would show a measurable impact on practice. Such a tool could provide a direct answer, or act as an expert system guiding IT staff (and users) along a path to identify the source of a problem.

Anomaly detection

Anomaly detection is useful to recognize security incidents and also to recognize performance problems. Anomaly detection finds deviation from "normal" or a baseline. Any given anomaly does not necessarily indicate a problem; correlation with other data sets (or training) is needed to understand what the anomaly means. We need more data analysis and interpretation, both for security incidents and performance problems.

Data correlation

Data correlation can be one-way for anomaly detection. Anomaly detection software could greatly benefit from correlating data from a number of metrics to understand what an anomaly means.

perfSONAR makes raw data (and analysis results) available for correlation. Research in what correlations, and algorithms, are useful for various problems is needed. These are likely "analysis tools" as mentioned above. In addition, publishing network events (topology changes, equipment changes, configuration changes, etc.) would allow anomalies (or problems) to be correlated, or traced back, to a network event.

Information on system logs, router configurations, and firewalls

The more situational awareness (interpreted as a network events above), the better analysis and assessment can be done. To discover the source of an end-to-end problem, it would be valuable to have access both to configurations settings and system logs. The goal is to query for events indicating something has changed in the network. This includes maintenance windows and trouble tickets.

There are commercial tools (e.g. NetView) that work with thresholds to tell when something seems to be wrong. The current systems don't work very well together – and there are lots of

them out there. There is no one common place or common schema for this type of information. perfSONAR could provide both.

One type of situational awareness that would be extremely useful is to be able to discern problems related to firewalls (or access control lists). Most operators are more concerned with security than with performance. These can be competing goals.

There will be a compromise between useful tools to locate performance problems and securing the enterprise. The group discussed the creation of two new measurement archives (MAs): a 'Firewall MA' and a 'Configuration MA'. The Firewall MA would be a way to learn what is the filter configuration for a given domain, while a 'Configuration MA' would be useful to expose other configuration details for this domain. It is likely that some additional authentication and authorization would be required to access this data.

To have effective problem and anomaly detection, we also need a model of what the ideal network is, so that we can say when something does not match this model. But, it was pointed out that the creation of a model is hard to do (e.g. establishing signatures, rules, etc).

Verification tools

The group identified that it would be helpful to have two kinds of verification tools. First, verify that data sent from this AS is received by the next AS. Second, end-to-end verification, including verification of end host configuration.

Research use of data

The perfSONAR community has developed a successful platform for data measurement, and collected large volumes of network performance data. Separately, in carrying out various research projects, many other efforts have also gone into tool development and measurement, e.g. PREDICT [R7] Netalyzr [R8] SecSpider [R9], Cyclops [R10]. There is a clear potential for community development/expansion through open dialog between the perfSONAR community and the networking research community at large.

Open issues

- How should the availability of perfSONAR tools and data sets be advertised more widely to stimulate both usage and community growth?
- How can the research community mine the performance data collected by the perfSONAR community from operational networks? This issue aims to discover new research questions and answers from the data.
- How can the perfSONAR community learn more about the types of data needed by the network research community at large to make the toolset more useful for research usage? This issue aims to enable the procurement of data required to meet specific research needs.

Recommendations

For better publicity to the general network research community, schemes should be developed for promoting perfSONAR in classrooms. There is a good potential for perfSONAR resources to be utilized in both graduate and undergraduate curricula. One example is the Stanford virtual

router project (http://yuba.stanford.edu/vns/, funded by the NSF). Students could be encouraged to design measurement experiments.

Researchers using perfSONAR data and tools (early adopters) should be required to include a "powered by perfSONAR" statement with a reference to the perfSONAR website. Such publications should also provide a description of precise access methods and reference pointers to the data so that others can repeat the experiments. Permanence and provenance of data are important for repeatability. The NIH has established requirements for data that can be used by multiple researchers for repeating experiments [R12]; a similar set of requirements should be developed for network research.

The research community's focus evolves over time finding new challenges to supplement existing ones, such as characterizing delay, throughput, etc., but there is a wealth of open questions that need to be addressed. Building a 2-way channel for fostering communications between the two communities will, no doubt, result in new research opportunities. The next workshop should include more people from the network research community at large (i.e., those outside the perfSONAR community).

While some research problems can be answered with data from R&E networks, others, such as understanding AS-level topologies, AS-based traffic characterization, or what percentage of Internet traffic comes from youtube.com, require commercial Internet data. Nevertheless, there is likely a significant set of problems that can utilize data collected by the perfSONAR framework from R&E networks. By advertising the availability of this data to the network community through some of the mechanisms outlined above, these open research problems can be exposed, which in turn can lead to new data collection and corresponding improvement in the perfSONAR framework.

Application use of data

GridFTP and other high throughput applications are examples in which the selected network path is critical to end-to-end performance. As noted in "Host/Application monitoring", in addition to self-reporting statistics about CPU usage, memory, and disk metrics into the perfSONAR framework, such applications could query perfSONAR measurement points to obtain current status on which path choices can be made. The perfSONAR measurement points and services should be designed to respond quickly to such queries.

General topics

Scaling and Archiving

- What happens if perfSONAR becomes widely deployed and the machines start collecting 100's of GB per day? Are the current measurement archive architectures up to this task?
- On the issue of data archiving, clearly as archiving "everything" and "forever" is not feasible, how should data archiving be organized? Easy access is a key requirement in the design of the archiving process. Also, the important question of "who pays for the resources required for archiving" needs to be discussed and answered.

Major recommendations

- Develop sampling mechanisms to store small amounts of data longer term. For example, average, minimum or maximum operations can be executed on data collected over

specified periods of time (e.g., one average point per hour), and these summarized values stored for longer periods of time. ESnet makes available network topology information from 1987 on its web site [R11]. Some detailed data may not be relevant in the long term as network link speeds or server configurations change. In such cases, storing just the average, minimum and maximum, along with the server and network configurations used to obtain this data, could be sufficient. Another example is to store anomaly related data since this may be useful to researchers or operators in the future as they examine historical data.

- Owners of archived data should use a well-defined data labeling process to make it easy for researchers and operations to locate data as per their interests.
- Owners of archived data should clearly state their deletion policies allowing any users of data to determine whether or not they need to save personal copies of the data.
- Just as NIH has a published data sharing policy [R12], if the perfSONAR framework is used to support networking research, research & education network providers can offer their collected data with some level of "permanence,'' subject to their clearly stated deletion policies. As repeatability is a key component of peer-reviewed research publications, and leading journal publications often take months to even years to publish an article, a researcher could download data used in their manuscripts if such data is scheduled for deletion and store a personal copy. However, for repeatability, personal copies of data could be subject to suspicions of tampering. As research usage of such operational network data grows, this issue needs further discussion.
- An issue of scalability of the perfSONAR deployment was raised. For example, if the current deployment has 200 instances and that number increases to 200K, will this cause a problem? The other two breakout sessions discussed this point. It appears that the structure is scalable for two reasons: (i) measurement archives are not centralized but rather distributed (e.g., an OWAMP server stores its own collected data), and (ii) the two level hierarchy of lookup services allows for an expansion in the number of lookup servers as the number of measurement points grows.

Tie in to GENI and other frameworks

There was a general consensus in the group that perfSONAR needs to work collaboratively with other databases and measurement frameworks to enable research. The NSF-funded Global Environment for Network Innovations (GENI) [R13] project is a virtual laboratory for exploring future Internets at scale. It has several clusters, including ORCA (Open Resource Control Architecture, Duke University) [R14] and ProtoGENI (University of Utah) [R15] among others. The ORCA cluster includes a project to define a measurement framework for physical-layer measurements. Several attendees of the research breakout session are working on this project, and defining perfSONAR schema for these measurements. Schema development is a challenging process, so this should be tasked to a working group or individuals with expertise in the specific area for which the schema is being defined. These schemas should then be available on a web site for ease of access to the research community. Documentation for the prospective new user should also be available as a web download. Another project, which is in the ProtoGENI cluster, is the Leveraging and Abstracting Measurements (LAMP) with perfSONAR project [R16], which directly links perfSONAR with GENI. The PlanetLab cluster also has support for measurement tools, which could be tied into perfSONAR.

Other databases of interest are PREDICT [R17] (Department of Homeland Security), CRAWDAD [R18], a wireless dataset (Dartmouth University), DATCAT [R19] (Cooperative Association for Internet Data Analysis (CAIDA)), and ONELAB [R20] in Europe. Finally, there was a discussion that the Resource Description Framework (RDF) describes topology using Network Description Language (NDL), while perfSONAR uses a different XML schema. These are currently being reconciled within the Network Markup Language (NML) working group in the Open GRID Forum (OGF). NDL and perfSONAR's schema use similar concepts and are thus translatable.

Privacy and sensitivity concerns with data – Anonymization

Some of the data that would be useful to researchers is sensitive. It might reveal individual identities; different jurisdictions have different laws about personally identifiable information (PII). Packet flow data, packet header data, and packet payload data are particularly sensitive (in order of increasing sensitivity). Data might reveal details about the network itself that operators try to hide or obscure, perhaps as another layer of security to prevent attacks on the infrastructure, perhaps for commercial reasons.

Open Issues

No one anonymization or data hiding technique works in all cases; furthermore, with enough data, or enough different data sets, many anonymization techniques can be defeated [R21]. Therefore, a combination of methods may be needed: varying anonymization techniques based on data types, privilege, or intended use; moving code to data (with the output being aggregates that are deemed to be non-sensitive); agreements between parties (employment, MOU, NDA, etc.), and no doubt others. In perfSONAR, anonymization or other data hiding might be performed by a Transformation Service. The PRISM project [R24] in Europe is looking at these issues.

A. Given the above, perhaps what we need is a framework for data privacy, a "Belmont Report" [R22] for PII in the network. CAIDA researchers have been working on this [R23]. In the area of network security research, the DHS-funded PREDICT project [R17] has been working on legal agreements and a mechanism to release data to U.S. researchers that is modeled on data used for medical research. How might ideas from these projects be applied to perfSONAR?
B. What anonymization techniques are appropriate for different data sets? Different uses? The PREDICT project has also worked on this for security research.
C. How support different levels of anonymization based on location (laws in various jurisdictions – for example, there are clear differences between US and EU law), privilege, usage.
D. What's the best way to bring code or algorithms to data? Existing projects include the CoMo project [R25] and more recently work in Europe to put CoMo in the ETOMIC testbed [R26,R27]. How does bringing code to live data interact with scientific method and repeatability?
E. Is there an appropriate way to express privilege? Work with federated authentication techniques? [R28]

Recommendations

- Research, and allow, different anonymization strategies based on location and privilege.
- Investigate bringing code to data along with anonmization strategies.
- NSF should facilitate the creation of community standards for both data sharing and privacy.
- Create templates that can be used by data providers and researchers.

Conclusions

This section looked at how researchers could use perfSONAR to further their own research, could research improvements to perfSONAR, and in particular, could research improvements in support of network operations and application performance. In addition, the section discusses improvements in perfSONAR and documentation that would help the research community: step-by-step guides for adding tools or analyses; creation of a repository for tools, schemas, and documents; creating tool validation procedures. Open areas with respect to safely sharing sensitive data, were cited, along with recommendations to work on both anonymiztion (better share data) and bringing code to the data, only exporting less-sensitive summaries. Open issues and recommendations are spread throughout the section.

Acknowledgements

Our thanks to the many members of the perfSONAR community who participated in the research uses/findings breakout sessions.

References

[R1] Cluster of European Projects aimed at MOnitoring and MEasurements (MOME), http://www.ist-mome.org/
[R2] http://ganglia.sourceforge.net/
[R3] DANTE, Internet2, Canarie and ESNet (DICE), "Inter-domain Controller (IDC) Protocol Specification." [Online]. http://www.controlplane.net/idcp-v1.1/idc-protocol-specification-v1.1-feb092010.doc.
[R4] Tanja Zseby, Javier Aracil, Sergio Lopez-Buedo, " Measurement Tools in PlanetLab Europe," http://www.internet2.edu/workshops/perfSONAR/presentations/2010-07-08-Zseby-Measurement-in-PLE.pdf
[R5] T. Rodwell, PERT, Performance Enhancement & Response Team, http://www.internet2.edu/presentations/jt2006feb/20060207-pert-rodwell.htm
[R6] http://www.cesm.ucar.edu/
[R7] www.predict.org
[R8] N. Weaver, Netalyzr, http://www.internet2.edu/workshops/perfSONAR/presentations/perfsonar.key-weaver.pdf
[R9] http://secspider.cs.ucla.edu/
[R10] http://cyclops.cs.ucla.edu/
[R11] ESnet Network Maps, http://www.es.net/pub/maps/topology.html
[R12] NIH Data Sharing Policy, http://grants.nih.gov/grants/policy/data_sharing/
[R13] http://www.geni.net/?page_id=2
[R14] https://geni-orca.renci.org/trac/
[R15] http://www.geni.net/?p=1246
[R16] http://groups.geni.net/geni/wiki/LAMP
[R17] https://www.predict.org/
[R18] http://crawdad.cs.dartmouth.edu/
[R19] http://imdc.datcat.org/Home
[R20] http://onelabinitiative.blogspot.com/
[R21] Allman, M., and Paxson, V. Issues and etiquette concerning use of shared measurement data. In IMC (2007).
[R22] The National Commission for the Protection of Human Subjects of Biomedical and Behavioral Research. "Ethical Principles and Guidelines for the Protection of Human Subjects of Research" or, "The Belmont Report". 18-Apr-1979. http://www.hhs.gov/ohrp/humansubjects/guidance/belmont.htm

[R23] Erin Kenneally and Kimberly Claffy. Dialing privacy and utility: a proposed data-sharing framework to advance Internet research. In IEEE Security and Privacy, vol. 8, no. 4, pp. 31-39, July/Aug 2010. Available from http://www.caida.org/publications/papers/2010/dialing_privacy_utility/.
[R24] The PRivacy Aware Secure Monitoring project (PRISM). http://www.fp7-prism.eu/
[R25] The CoMo Project. http://como.sourceforge.net/index.php
[R26] CoMo in OneLab. http://www.onelab.eu/index.php/services/testbed-components/como.html
[R27] European Traffic Observatory Measurement Infrastructure (ETOMIC). http://www.onelab.eu/index.php/testbeds/onelab-testbeds/etomic.html
[R28] Shibboleth. http://shibboleth.internet2.edu

Breakout Report: Expanding perfSONAR Technology

Chair: Martin Swany (University of Delaware); Committee: Brian Tierny (ESnet) and Jeff Boote (Internet2)

Introduction

perfSONAR is not language dependent; currently, implementations exist in Java, Perl, and Python. perfSONAR is dependent on web services. While it currently uses SOAP, it is not dependent upon it; however, perfSONAR tools are dependent on an XML database used for the Lookup Service. The group agreed it was important, when talking about expanding perfSONAR Technology, to separate protocol from architecture. A basic taxonomy would include:

- MA, or Measurement Archive – this is the "Workhorse;" it provides access to measurement data
- MP, or Measurement Point – this can be either an Active Measurement or a Passive Measurement actor
- IS, or Information Service – this is a Service containing service and topology information
- AS, or AuthN/AuthZ Service – this is a piece, not entirely implemented, that would allow deployers to regulate what data is available to users based on their identities or roles in their organizations
- RP, or Resource Protector – this is a service that mediates access to shared resources
- MC, or Measurement Coordinator – this is a proposed service that would allow the measurement infrastructure to recognize that it was in use and refuse or put on hold a subsequent request that might have an effect on outcome of the first test.

The group proposed a series of scenarios; how will perfSONAR meet these needs?

- A test between a host at FNAL to a host in Uruguay. If the BWCTL test to Uruguay fails, who do I contact?
- Real-time physical layer measurements. Folks have asked for a Skeleton service. We have ways of extracting information from elements with TL1. We want to put up a perfSONAR-like measurement point. We want shortest path between consumer and producer. What do we do?
- High-scalability. an MP with 10,000 consumers, and they don't want to poll. Relates to the pub-sub aspects.
- Ask a service, "give me data; if you don't have it, run a test".

The group focused discussion on the missing components in the architecture, components that should be reviewed or re-implemented, flaws in the architecture, and goals for the future.

Missing Architecture Components

The group discussed a variety of architectural components they felt were missing.

Participants thought that a Publish-Subscribe (pub-sub) interface for perfSONAR would be useful; they suggested considering WS-Notification and AMQP, both of which are available. The

latter is a framework, designed for low-latency, that has open-specification and open-implementations and is supported by the financial industry.

The non-active measurement actor that fills in the MA is never truly visible. In GENI, there is a Measurement Controller that describes the service behind the scenes filling in the MA. Another concept is a Broker Service that would take the data generated by the Measurement Controller, and push it to the appropriate MA. This would facilitate pub-sub.

The group was also concerned about the Transformation Service, or TS (e.g. the Nagios alert mechanism planned for the upcoming perfSONAR-MDM release but also anomaly detection services, find me 'close' MPs, etc.).

The group recommended developing a list of the open problems and how they align with the capabilities of the components, as well as a list of obstacles. This would improve the research options and would allow you to see what development is needed for the future.

Another need identified is the capability for self-organization of measurements. If a user is measuring over a backbone, then another user's measurement ant must notice the first user's measurement ant, and back off. This would require development of a "measurement coordinator."

Authentication and Authorization (AA)

Participants reported that, in several Big Science projects, the traditional attribute-driven model is being changed to an attribute-store method, which allows users to make assertions in the policy that can be signed by an identity provider. The assertions could be things like organization, user, etc. These groups also want to make their AA model something that can be plugged in to the control framework. perfSONAR also wants something that just plugs in.

The current 3.3 perfSONAR-MDM release contains an AA service but it was unclear whether this service merely identified the provider and user or if there was a matching of policy to resources. The group expressed concern about what is in place for users who are not Shib-enabled; they recommended that perfSONAR have a default of "open" so that anonymous users can access most data. Access to sensitive data and resources, of course, would continue to be limited. A Best Common Practices document is strongly recommended so that both users and providers know what is expected.

The group agreed that perfSONAR should make identity pervasive but transparent; however, there was strong concern that perfSONAR should not try to be ahead of AA deployment so as not to limit who can deploy perfSONAR. In general, they didn't want to splinter the old perfSONAR community and new adopters, but they do recommend making sensible choices to allow for seamless transitions; AA is an important component to the end-product perfSONAR deployments but deployment should not be delayed until an identity infrastructure is established.

Fault Tolerance

One thing missing from current architecture is the specific notion of redundancy; an MA can register with multiple LS', but there is no way to identify which is primary or secondary, and

users may get duplicate data. perfSONAR-MDM is considering using a self-documenting data model so users don't need to do a different schema for every new type of data. The schema includes a template to get information about what is responded.

Reimplementation Needs

The group argued the values of REST vs. Web Services and determined that it would be sensible to support both (and not cause inconsiderable effort), if only to be able easily to do transitions.

The group also discussed the need to reevaluate the original design decisions. SOAP and WSDL are over-designed, and they seem to be losing public support. Moving in the way the community is moving makes working with community easier.

Architecture Flaws

Scalability

The XML databases have some scaling issues, as do some of the MAs. The LS dumps data unless it is refreshed. The Topology Service keeps information about the makeup of the network for months or years. The group discussed what scalability limits are in the use of an XML DB and recommended compiling a list of what is needed with the XMLDB, since that is much less likely to scale.

Architecture Issues

The group discussed theoretical architecture issues (i.e., if 100,000 BWCTL servers were all talking to each other) and agreed that these types of problems are already manageable with the Resource Protector. Policy determination is hierarchical – each domain sets their own policies but is there anyone who can set or change policy across domains? If a backbone provider is running a full-mesh of tests, end-users may find only limited opportunities to run tests between end-points crossing that backbone. The ability to reserve a time for a test or to access backbone data is critical to the end-user value proposition.

Goals

This group agreed that the goal should be not “everyone will use perfSONAR-PS” or “everyone will use perfSONAR-MDM,” but “everyone will use perfSONAR (the architecture).” This allows everyone to implement a version that works for their network environment and end-user needs.

A second goal the group reached was every router on the planet having embedded perfSONAR capabilities. As more customers demand integrated measurement and monitoring, there will be forces demanding implementation of an international standard. In addition, perfSONAR should reside in the applications.

The group discussed application-level monitoring (within the application, within the middleware, within the operating system, and within the network, to produce the complete performance picture).

Community Building

The toolkit is really a community-building toolkit – it provides the mechanism for testing (can be deployed on many test platforms), it can be packaged, and users can create specialized deployments. It's a mechanism whereby the community contributes and then tests the end-to-end solution.

Problems might not show until you have 50 or more users; participants are given a list of components to download; all the virtual test sites will test the solution and ensure it works before the new version is released. This would allow the larger community to participate in end-to-end solutions.

Conclusion and Recommendations

The group came to recommendations in a variety of areas:

- *Missing components:*
  - Add a Publish-Subscribe (pub-sub) interface for perfSONAR and considering WS-Notification and AMQP.
  - Develop a list of the open problems and how they align with the capabilities of the components, as well as a list of obstacles.
  - Develop a "measurement coordinator" to facilitate self-organization of measurements
- *AA:*
  - Don't be ahead of AA deployment so as not to limit who can deploy perfSONAR. Also, that perfSONAR default to "open" sharing so that anonymous users can access most data.
  - Publish a Best Common Practices document so that both users and providers know what is expected.
- *Reimplementation Needs:* Support both REST vs. Web Services to be able easily to do transitions.
- *Architectural Flaws/Scalability:* Compile a list of what is needed with the XMLDB, since that is much less likely to scale.
- *Goals:* Develop a plan to ensure that every router on the planet has embedded perfSONAR capabilities by 2013. Also, perfSONAR in the network applications by 2015.

Conclusion

Workshop Outcomes

Overall information is available in the perfSONAR Workshop Survey included in the Appendices. In general:

- 51 total evaluations (of 90 or 57%) turned in; 8 of 12 (67%) Expanding perfSONAR Technology group, 15 of 20 (75%) Operational Issues group, and 26 of 31 (84%) Research Uses/Findings group
- 40 participants pre-registered for the Tutorial; with room reconfiguration, were able to include 55 participants
- Workshop attendees represented: 35 network engineers/operators, 30 researchers, 19 tool developers, and 6 (+2) government staff or meeting support personnel
- Travel support was provided to 7 International participants (2 from Germany and 5 from Brazil) and 12 U.S. participants (3 network researchers, 2 network operators, 1 tool developer, and 6 graduate students in network research)
- Areas of general consensus were: a) Workshop should be an annual event (95.6%), b) amount of time allowed for presentations was "about right" (82%) and c) amount of time allowed for breakouts was "about right" (75%)
- Areas of consensus among breakout group respondents were: a) my breakout group "started good discussion" (73.5% overall but only for the non-Operational Issues group – which had 66% "reached consensus" for same question and only 33% "started good discussion"), b) splitting breakouts onto 2 days was good (97.8%) and c) continuing discussions on these breakout topics in future years is important (76.1%)
- Respondents generally approved of the Tutorial; 48.6% for it "informative" and 42.9% found it "useful" (52.1% attended the tutorial and 73% of those who did NOT attend accessed the Tutorial online)

The three physically separate breakouts (in order of total participation) had varied interests; the **Research Findings/Uses** group (84% returned evaluations):

- Presentation topics of interest were overwhelmingly those on Research Uses/Findings (80%) with Deployment Experiences (42%) and E2E Problems (38%) as secondary interest
- 85% reported that the group "started good discussion"
- No real consensus on posters (7 found them informative, 9 saw them as supplements to presentations with 4 reporting them not useful) – only 5 suggested having more at subsequent workshops
- Significant consensus on adding new topics to future workshops – 58% suggested this; 35% also suggested more presentations

The **Operations Issues** group (75% returned evaluations):

- Presentation topics of interest were overwhelmingly those on Wide-Scale Deployments (60%) and Implementation Barriers (60%) with E2E Problems (47%) as secondary

interest; interestingly, both Research Uses and Deployment Experiences received 33% votes, showing this group to have very broad interests
- 66% reported that the group "reached consensus"
- Consensus on posters being informative & good supplements to presentations (7 vs. only 1 reporting them not useful) – but only 3 suggested having more at subsequent workshops
- No real consensus on adding new topics to future workshops (4 people suggested this) or including more presentations (5 people suggested that).

The **Expanding perfSONAR Technology** group (67% returned evaluations):

- Presentation topics of interest were overwhelmingly those on Wide-Scale Deployments (63%) and Implementation Barriers (63%) with Deployment Experiences (50%) as tertiary interest; interestingly, only 1 vote was cast for E2E Performance Problems!
- 100% reported that the group "started good discussion"
- Consensus on posters being informative (50% with no one reporting them not useful) – but only 2 suggested having more at subsequent workshops
- No votes for adding new topics to future workshops or including more presentations.

Recommendations

Evaluations

Based on feedback from the evaluations, the primary recommendation is that the perfSONAR Workshop should be an annual event. Because most respondents agreed that the agenda provided a proper balance between presentations and group discussions, we would propose continuing with this agenda model for any future workshops. The value of poster sessions was clearly identified; however, based on feedback, we would recommend more coordination with the presentation topics. Overwhelming support for the pre-workshop tutorial (or being able to access slides online) indicates that this should be a part of any future perfSONAR Workshops.

A significant number of participants recommended that new topics be considered for future workshops. Of the topics covered in this workshop, there was overwhelming interest in the Research Uses/Findings, Wide-Scale Deployments, and Implementation Barriers topics. The E2E Performance Problems talks garnered the least support; if new topics are included in future workshops, we would recommend first replacing the E2E Performance Problems area and then, if indicated by submissions or Executive Committee vision, the Deployment Experiences topic.

Specific recommendations from participants included several requests for a hands-on tutorial, as well as requests to offer separate tutorials for tool building and for deployers/end-users. One suggestion was having participants run the pS-Performance Toolkit (via ISOs on their laptops) and build a small demonstration measurement lab for attendees to work through.

The breakout sessions used shared online documents to collect comments; this was appreciated (especially by people who needed to attend more than one breakout) but several participants recommended using projectors in the breakout rooms but several of the breakout leaders, and a few participants, informed Internet2 staff that they would have preferred having flipcharts in the room to collect the data.

A final recommendation, culled from several evaluations, was that future workshops include demonstrations (of data retrieval, for example) of using perfSONAR and having several published use cases for participants to review.

Future perfSONAR Workshops are of interest to and requested by the R&E community. The overall approach should be similar to the 2010 perfSONAR Workshop with a more hands-on tutorial, several perfSONAR use demonstrations, and poster sessions that are fully coordinated with the scheduled presentations. A goal for future workshops should be to develop a list of action items that could be tackled by members of a working group

Breakout Sessions

Details of recommendations can be found at the end of the report for each of these sessions; in summary:

Community Building

- Expand deployment
- Grow the community
- Hold another Workshop
- Expand availability of tools and services
- Organize the community
- Meet the needs of researchers

Operational Issues

- Address security concerns
- Integrate perfSONAR into network management systems
- Provide funding to build the tools and support needed by people involved in security & software updates
- Data Management Research
- Document Best Practices
- Develop a framework for collecting and managing case studies

Research Uses/Findings

- Make it easy to add tools or analyses
- Develop assessment tools, especially for non-expert IT staff
- Export as much as possible to improve situational awareness
- Use of configuration and flow data will require improved authentication and authorization
- Address scalability and data provenance
- NSF should facilitate community standards – for data sharing and for privacy

Expanding perfSONAR Technology

- Address missing components
- Support both REST vs. Web Services
- Compile a list of what is needed with the XMLDB.
- Every router has embedded perfSONAR capabilities

APPENDICES

Appendix A: perfSONAR Workshop Executive Committee

Internet2 would like to thank members of the Executive Committee who met for several hours at the Winter 2010 ESCC/Internet2 Joint Techs to begin planning the event, and those who joined weekly calls to shape the agenda, identify speakers and participants to invite, and provided support for the breakout sessions during the meeting:

Jeff Boote (Internet2)
Aaron Brown (Internet2)
Grover Browning (University of Indiana)
Rich Carlson (DoE)
Mark Crovella (Boston University)
Susan Evett (Internet2)
Dale Finkelson (Internet2)
Victor Frost (NSF)
Ann Harding (SWITCH)
Carla Hunt (MCNC)
Yasuichi Kitamura (APAN)
E. Paul Love (NOAA/NCO)
Joe Metzger (ESnet)
Grant Miller (NITRD)
José Augusto Suruagy Montiero (RNP)
Roberto Sabatino (DANTE)
Martin Swany (University of Delaware)
Brian Tierney (ESnet)
Malathi Veeraraghavan (University of Virginia)
Jason Zurawski (Internet2)

Internet2 would especially like to thank the co-authors of this report:

Joe Breen (University of Utah)
Aaron Brown (Internet2)
B. Ann Cox (Laboratory for Telecommunications Services)
Susan Evett (Internet2)
Tom Hacker (Purdue University)
Ann Harding (SWITCH)
Carla Hunt (MCNC)
Yasuichi Kitamura (APAN)
Joe Metzger (ESnet)
Inder Monga (ESnet)
José Augusto Suruagy Montiero (RNP)
Malathi Veeraraghavan (University of Virginia)
Matt Zekauskas (Internet2)

Artur Ziviani (RNP)
Lixia Zhang (University of California at Los Angeles)
Jason Zurawski (Internet2)

Appendix B: perfSONAR Workshop Agenda

Wednesday, July 7		
Time	**Topic**	**Description**
Noon-5:00 pm	Registration	Pickup name tags, etc.
1:00-5:00 pm	Tutorial	Jason Zurawski, Internet2, will provide an overview of the perfSONAR project, with the purpose of bringing Workshop attendees to similar levels of understanding. Materials are located here.
Thursday, July 8		
Time	**Topic**	**Description**
7:00 am -5:00 pm	Registration	Pickup name tags, select breakout groups, etc.
8:00-8:10 am	Opening Remarks	Dr. Victor Frost (NSF) & Rich Carlson (DoE)
8:10-8:30 am	perfSONAR Overview	Eric Boyd
8:30-9:00 am	Motivation for Breakouts	Breakout Group Leaders
9:00-9:40 am	Wide Scale Deployments	Christos Papadopoulos Nick Weaver Wolfgang Fritz
9:40-10:00 am	BREAK	
10:00-10:40 am	Implementation Barriers	Joe Metzger Yee-Ting Lee Andy Germain
10:40 am-11:40 pm	Research Uses/Findings	Malathi Veeraraghavan Prasad Calyam Marcos Portnoi Martin Swany Tanja Zseby
11:40-12:30 pm	E2E Performance Problems	Guy Almes Chris Welti Chris Heerman Martin Swany
12:30-1:20 pm	LUNCH	
1:20-2:30 pm	Deployment Experiences	Takatoshi Ikeda Stephan Kraft Ramiro Voicu Katsushi Kobayashi Katsuichi Nakamura Suruagy Monteiro
2:30-3:10 pm	Breakout (general discussion)	Community Building [slides]
3:10-3:30 pm	BREAK	
3:30-5:30 pm	Breakout Sessions	Research Uses/Findings -- Roslynn Room Expanding pS Technology -- Lee Room Operational Issues -- Jefferson Room
6:00-6:30 pm	Poster Sessions	Tanja Zseby, Kenji Shimizu, Jeff Boote, Sergio Lopez-Buedo, Stephan Kraft, Artur Ziviani, and Prasad Calyam will be available to discuss their posters
6:30-8:30 pm	Dinner	
Friday, July 9		
Time	**Topic**	**Description**

8:00-9:30 am	Breakout Sessions	Research Uses/Findings -- Rosslyn Room Expanding pS Technology -- Lee Room Operational Issues -- Jefferson Room
9:30-9:50 am	BREAK	
9:50-11:00 am	Breakout Sessions	Research Uses/Findings -- Rosslyn Room Expanding pS Technology -- Lee Room Operational Issues -- Jefferson Room
11:00-Noon	Report Outs	Each break-out discussion group will present conclusions, followed by general discussion and wrap-up.
Noon-1:00 pm	Wrap-Up	Report planning & next steps
1:00-3:00 pm	Lunch & Continued Discussion	General Session and Breakout Rooms available

Appendix C: perfSONAR Workshop Roster

There were 90 attendees at the perfSONAR Workshop:

Last	First	Organization
Almes	Guy	Texas A&M
Angu	Pragatheeswaran	University of Nebraska-Lincoln
Baldine	Ilia	University of North Carolina at Chapel Hill
Bathula	Balagangadhar	Columbia University
Blatecky	Alan	NSF
Boote	Jeff	Internet2
Boyd	Eric	Internet2
Brauner	Daniela	RNP
Breen	Joseph	University of Utah
Brown	Aaron	Internet2
Calyam	Prasad	OARnet
Carlson	Richard	DoE Office of Science
Cox	Beverly Ann	Laboratory for Telecommunications Sciences (LTS)
Crovella	Mark	Boston University
DeMar	Phillip	Fermi National Laboratory
Downey	Traci	Internet2
Dykstra	Chonghui	WareOnEarth Communications
Dykstra	Phillip	DREN
Evett	Susan	Internet2
Fernades	Guilherme	University of Delaware
Fernandes	Stenio	RNP
Finkelson	Dale	Internet2
Fraser	Dan	Open Science Grid
Frey	Michael	Bucknell University
Fritz	Wolfgang	DFN-Verein
Frost	Victor	NSF
Germain	Andy	NASA GSFC
Greenfield	Richard	University of Alaska-Fairbanks
Gurkan	Deniz	University of Houston
Hacker	Thomas	Purdue University
Harding	Ann	SWITCH
Hartzell	David	NASA Advanced Supercomputing
Hausheer	David	Univeristy of California-Berkeley
Heerman	Chris	University of North Carolina at Chapel Hill
Huang	Shu	University of North Carolina at Chapel Hill
Hunt	Carla	MCNC/SCGPoP
Ibarra	Julio	Florida International University

Ikeda	Takatoshi	National Institute of Information and Communications Technology (NICT)
Jackson	Keith	Lawrence Berkeley National Laboratory
Kettimuthu	Rajkumar	Argonne National Laboratory
Kissel	Ezra	University of Delaware
Kitamura	Yasuichi	APAN
Kobayashi	Katsushi	National Institute of Advanced Industrial Science and Technology (AIST)
Kraft	Stephen	DFN-Verein
Lee	Yee-Ting	Stanford Linear Accelerator Center
Lewis	Joe	Texas A&M
Lopez-Buedo	Sergio	Universidad Autonoma De Madrid
Love	E. Paul	NOAA
Majumder	Debjyoti	University of Houston
Martinello	Magnus	Federal University of Espirito Santo - UFES)
Metzger	Joe	ESnet
Miller	Grant	National Coordination Office for NITRD
Monga	Indermohan	ESnet
Montiero	Jose Augusto Suruagy	RNP
Moraes	Priscilla	University of Delaware
Murta	Cristina	CEFET-MG
Nakamura	Katsuichi	Kyushu Institute of Technology
Ndousse-Fetter	Thomas	DoE Office of Science
Oslebo	Arne	UNINETT
Papadopoulos	Christos	Colorado State University
Portnoi	Marcos	University of Delaware
Pyun	Yul	University of Southern California
Radulovic	Predrag	University of Tennessee
Sabatino	Roberto	GEANT/DANTE
Sheldon	Paul	Vanderbilt University
Shimizu	Kenji	Nippon Telegraph & Telephone Corp (NTT)
Sim	Alexander	Lawrence Berkeley National Laboratory
Small	Chris	Indiana University at Bloomington
Song	Larry	George Mason University
Swany	Martin	University of Delaware
Throckmorton	Thomas	MCNC/SCGPoP
Tierney	Brian	ESnet
Trocha	Szymon	PSNC/PIONIER
Turnbull	Susan	DoE Office of Science
Uhl	George	NASA GSFC
Veeraraghavan	Malathi	University of Virginia
Verlo	Alan	University of Illinois at Chicago

Voicu	Ramiro	California Institute of Technology
Wac	Katarzyna	Carnegie Mellon University
Weaver	Nicholas	Univeristy of California-Berkeley
Welti	Chris	SWITCH
Wolff	Stephen	Cisco Systems
Xu	Qian	George Mason University
Yu	Dantong	Brookaven National Laboratory
Zekauskas	Matt	Internet2
Zhang	Lixia	University of California, Los Angeles
Ziviani	Artur	RNP
Zseby	Tanja	Fraunhofer FOKUS
Zurawski	Jason	Internet2

In addition, 45 attendees registered for the pre-workshop tutorial but a total of 55 (room capacity) actually attended. Those who registered included:

Last	**First**	**Organization**
Baldine	Ilia	University of North Carolina at Chapel Hill
Bathula	Balagangadhar	Columbia University
Breen	Joseph	University of Utah
Calyam	Prasad	OARnet
Cox	Beverly	Laboratory for Telecommunications Sciences (LTS)
Fernandes	Stenio	RNP (Rede Nacional De Ensino E Pesquisa)
Fraser	Dan	Open Science Grid
Frey	Michael	Bucknell University
Fritz	Wolfgang	DFN-Verein
Frost	Victor	National Science Foundation
Gurkan	Deniz	University of Houston
Hacker	Thomas	Purdue University Main Campus
Harding	Ann	SWITCH
Hartzell	David	NASA Advanced Supercomputing
Huang	Shu	University of North Carolina at Chapel Hill
Hunt	Carla	MCNC/SCGPoP
Ikeda	Takatoshi	National Institute of Information and Communications Technology (NICT)
Kettimuthu	Rajkumar	Argonne National Laboratory
Kitamura	Yasuichi	APAN
Kobayashi	Katsushi	National Institute of Advanced Industrial Science and Technology (AIST)
Li	Yee-Ting	Stanford Linear Accelerator Center

Love	E. Paul	NOAA (National Oceanic & Atmospheric Administration, Washington, D.C.)
Majumder	Debjyoti	University of Houston
Metzger	Joe	ESnet
Miller	Grant	National Coordination Office for NITRD
Monga	Indermohan	ESnet
Moraes	Pricilla	University of Delaware
Nakamura	Katsuichi	Kyushu Institute of Technology
Portnoi	Marcos	University of Delaware
Pyun	Yul	University of Southern California
Sim	Alexander	Lawrence Berkeley National Laboratory
Veeraraghavan	Malathi	University of Virginia
Voicu	Ramiro	California Institute of Technology
Wac	Katarzyna	Carnegie Mellon University
Wang	Michael	Columbia University
Weaver	Nicholas	University of California-Berkeley
Welti	Chris	SWITCH
Wolff	Stephen	Cisco Systems
Zhang	Lixia	University of California, Los Angeles
Ziviani	Artur	RNP (Rede Nacional De Ensino E Pesquisa)

Appendix D: perfSONAR Workshop Evaluation

The evaluation was created in SurveyMonkey and printed copies were distributed to all attendees. We received 51 evaluations (57%) and made the evaluation available online (http://www.surveymonkey.com/s/M58RW76) for those who did not turn in paper copies at the conclusion of the Workshop. A copy of the evaluation summary is now posted at http://www.internet2.edu/workshops/perfSONAR/

In general:

- 51 total evaluations (of 90 or 57%) turned in; 8 of 12 (67%) from the Expanding perfSONAR Technology group, 15 of 20 (75%) from the Operational Issues group, and 26 of 31 (84%) from the Research Uses/Findings group
- 40 participants pre-registered for the Tutorial; with room reconfiguration, were able to include 55 participants
- Workshop attendees represented: 35 network engineers/operators, 30 researchers, 19 tool developers, and 6 (+2) government staff or meeting support personnel
- Travel support was provided to 7 International participants (2 from Germany and 5 from Brazil) and 12 U.S. participants (3 network researchers, 2 network operators, 1 tool developer, and 6 graduate students in network research)
- Areas of general consensus were: a) Workshop should be an annual event (95.6%), b) amount of time allowed for presentations was "about right" (82%) and c) amount of time allowed for breakouts was "about right" (75%)
- Areas of consensus among breakout group respondents were: a) my breakout group "started good discussion" (73.5% overall but only for the non-Operational Issues group – which had 66% "reached consensus" for same question and only 33% "started good discussion"), b) splitting breakouts onto 2 days was good (97.8%) and c) continuing discussions on these breakout topics in future years is important (76.1%)
- Respondents generally approved of the Tutorial; 48.6% for it "informative" and 42.9% found it "useful" (52.1% attended the tutorial and 73% of those who did NOT attend accessed the Tutorial online)

The three physically separate breakouts (in order of total participation) had varied interests; the **Research Findings/Uses** group (84% returned evaluations):

- Presentation topics of interest were overwhelmingly those on Research Uses/Findings (80%) with Deployment Experiences (42%) and E2E Problems (38%) as secondary interest
- 85% reported that the group "started good discussion"

- No real consensus on posters (7 found them informative, 9 saw them as supplements to presentations with 4 reporting them not useful) – only 5 suggested having more at subsequent workshops
- Significant consensus on adding new topics to future workshops – 58% suggested this; 35% also suggested more presentations

The **Operations Issues** group (75% returned evaluations):

- Presentation topics of interest were overwhelmingly those on Wide-Scale Deployments (60%) and Implementation Barriers (60%) with E2E Problems (47%) as secondary interest; interestingly, both Research Uses and Deployment Experiences received 33% votes, showing this group to have very broad interests
- 66% reported that the group "reached consensus"
- Consensus on posters being informative & good supplements to presentations (7 vs. only 1 reporting them not useful) – but only 3 suggested having more at subsequent workshops
- No real consensus on adding new topics to future workshops (4 people suggested this) or including more presentations (5 people suggested that).

The **Expanding perfSONAR Technology** group (67% returned evaluations):

- Presentation topics of interest were overwhelmingly those on Wide-Scale Deployments (63%) and Implementation Barriers (63%) with Deployment Experiences (50%) as tertiary interest; interestingly, only 1 vote was cast for E2E Performance Problems!
- 100% reported that the group "started good discussion"
- Consensus on posters being informative (50% with no one reporting them not useful) – but only 2 suggested having more at subsequent workshops
- No votes for adding new topics to future workshops or including more presentations.

Access the evaluation online at: http://www.surveymonkey.com/s/M58RW76