Bridging the Gap Workshop: Draft Report

Table of Contents

Executive Summary

Maximum performance for networked applications does not happen without special efforts but, unfortunately, resources are limited for these efforts. Wizards are too few in number and too costly in effort to assure maximum performance on their own. Therefore, end-user researchers, applications developers, and network operators must come closer to what experts (wizards) can achieve. The gap separating wizards from these other groups, however, has been increasing exponentially each year rather than decreasing. In an attempt to reduce it, the National Science Foundation (NSF) funded the "Bridging the Gap" workshop. Its aims were to: 1) identify and share each community's respective perspectives, expectations, and knowledge about network performance; 2) raise awareness of common ground and differences across communities; 3) identify appropriate tools, techniques, and potential obstacles to using them; and 4) determine communications required for collective efforts to improve network performance.

Thirty-two people participated for a day and a half, including almost equal numbers of scientific researchers, scientific application developers, network engineers, and network wizards. An emphasis was given to seeing problems from the end-user's point of view, with each community responding accordingly. Participants engaged in large group discussions and small team simulation exercises. In small teams, a mix of participants from each community role-played a problem-solving case involving network problems at multiple levels of the protocol stack. Large group discussions highlighted and extended insights from these small group simulations.

Findings from the workshop reveal that communities have greater common grounding than is assumed, thus making it more possible to know where to focus to bridge gaps that still exist. The workshop showed that communities' concerns fall into the following four areas, all of which need further attention, both individually and in relation to each other:

- What tools and information are available and how to use them
- How to provide resources in ways that assure their uptake
- How to make organizations responsive to collective efforts
- How to build and sustain a culture of partnership

Based on findings from the workshop, we present six recommendations, as follows:

1. Conduct studies on the best ways to deliver information and tools so that they truly will take hold with end-users.

2. Conduct expert workshops for assessing the completeness of documentation and tools available to users.

3. Create more repositories of information and tools and keep them current.

4. Exploit and create new opportunities for partnering across communities and professional networking.

5. Generate, from decision-makers, a commitment of resources for the goal of bridging the gap.

6. Use the NSF as a conduit to the scientific communities

Introduction

Background

The numbers and types of networking problems experienced by the scientific application community are growing, partly because of a 100-fold increase in the Wizard Gap[1]. Various projects and activities are working to solve specific aspects of the gap, including the Internet2 End-to-End Performance Initiative (http://e2epi.Internet2.edu), the NSF-funded Web100 (http://www.web100.org) and The Performance Advisor projects (http://dast.nlanr.net/Projects/Advisor/), the DOE-funded Net100 project (http://www.net100.org), and MonALISA, (http://monalisa.cacr.caltech.edu/). Thus far, however, network performance solutions have not been explored across scientific application communities. To do so involves bringing together researcher end-users, application developers, campus network engineers, and wizards. By working together they can develop a common language and understanding to discover, examine and resolve problems in end-to-end network performance that adversely affect scientists' work. To foster this collaboration and communication and to arrive at strategies for pressing problems that adequately account for the perspectives of all four groups, Internet 2 and the University of Michigan convened a workshop. Almost equal numbers of specialists from each group participated to achieve the following purposes:

- Share from all communities' perspectives the needs, expectations, and related challenges of using networks for scientific research purposes;
- Identify commonalities and distinct points of view of end-to-end network problems and issues across various communities and collectively prioritize problem solving, needs
- Raise awareness and inform the use of various new and existing diagnostic tools according to roles that each community plays a non-routine troubleshooting situation.
- Determine the communications, coordination, and organizational dynamics that cross-functional communities can work on together to improve joint efforts to enhance end-to-end performance

Highlights of Preparation

The one and one-half day workshop, which was organized jointly by Internet2 and the University of Michigan's School of Information, took place on August 29-30, postponed from its initially-planned June date. The site of the Workshop, the newly completed Michigan Information Technology Center (MITC), allowed the Workshop coordinators have ample space for group sessions and break-out rooms and refreshments.

A web site, http://e2epi.internet2.edu/btg/, was developed and maintained to describe the Wizard Gap and its relationship to the scientific application communities; to register participants; and to display an agenda, workshop logistics, details about attendees and presenters, and other shared resources. Presentations made during the workshop and links to the workshop white paper are available through this web site, as are the following problem-oriented scenarios contributed by specialists from various groups:

- Yoonjoo Kwon (KISTI): http://people.internet2.edu/~ghb/pmwiki/pmwiki.php/BridgingTheGap/YoonjooKwon
- Jason Gerk (NEEsit): http://people.internet2.edu/~ghb/pmwiki/pmwiki.php/BridgingTheGap/JasonGerk
- Ruszczyk & Whitney (MIT Haystack): http://people.internet2.edu/~ghb/pmwiki/pmwiki.php/BridgingTheGap/ChesterRuszczyk
- Matt Crawford (Fermi Lab): http://people.internet2.edu/~ghb/pmwiki/pmwiki.php/BridgingTheGap/MattCrawford
- Wilson Dillaway (Tufts): http://people.internet2.edu/~ghb/pmwiki/pmwiki.php/BridgingTheGap/WilsonDillaway

---

1 The Wizard Gap - the difference between the bandwidth a user (or application) can get over a testbed network where the application and network stack has been tuned by a network "wizard" and what a normal user will see.. the ratio between the gap was about 3:1 in 1988 and is now more than 300:1. This wizard gap is due to two very large effects: end-systems (hosts) are optimized for relatively low bandwidth public markets; and the network itself does extremely well at hiding its own bugs. – "A Traffic Dynamics Testbed," M. Mathis (PSC), G.L. Huntoon (PSC), K. Claffy (CAIDA) 2001 [www.ngi-supernet.org/lsn2000/PSC.pdf]

Invitation Process

A Call for Papers was announced to prompt individuals in research and networking communities linked to Internet2 to apply to the workshop. Network wizards were invited, and, to ensure an appropriate mix of participants, some campus network engineers, application experts, and researchers were also invited.

The organizers looked at four functional communities:

1. Network wizards from PSC, Internet2, LBL, SLAC, the University of Washington, Caltech, and CERN.
2. Researchers from BIRN, NEES Grid, HENP, Oceanography, Environmental Science (such as the NEON, NEPTUNE, and Earth Scope projects), Astronomy, and VLBI.
3. Network Application Developers from the researcher groups.
4. Network Engineers, from Network Operation Centers, GigaPoPs, and Campuses.

To identify individuals to invite, they scoured research communities for representatives of each of the four functional communities. The research communities researched and the individuals contacted are:

- High Energy Physics (e.g., HENP) : Harvey Newman (Researcher), Les Cottrell (Network wizard), Shawn McKee (Researcher)
- BIRN : Mark Ellsman, Mark James, Tom Hutton (Network wizard), Mike McGill
- Astronomers: Alan Whitney (Researcher), Paul Etzel (Researcher), Hans Wener-Braun (Network Engineer), Chester Ruczyczk (Network Engineer)
- NEES Earthquake Engineering : IT staff, Anke Kamrath (Researcher), Cliff Roblee (Researcher)
- Oceanography : Dempsey E. Lott III (Researcher) Art Gaylord (Network Engineer)
- NEON -- ecology/biology : Deborah Enstren (Researcher)
- GEON : Krishna Sinha (Researcher)
- Teragrid : Charlie Catlett (Researcher) and Linda Winkler (Network Engineer)
- National Center for Data Mining (NCDM) : Bob Grossman (Network Engineer)
- LOCI Logistical : Micah Beck (Application Developer)
- CENIC : SDSC staff (e.g. Kevin Walsh and K. C. Claffy) and CALIT^2 staff
- Backbone engineers : Brent Sweeny (Abilene), Brian Cort (CENIC), and Jerry Sobieski (MAX)

The organizers conferred by phone regularly. They included wizards in several of these calls, one of which resulted in the creation of a problem-solving scenario for break-out sessions. The break-out session groups were preset to ensure an optimal mix of members from each of the four functional groups. Many participants, in practice, span more than one group so the organizers asked registrants to specify their preferred role for the workshop.

The Workshop

Workshop Participation

Below is a list of Workshop participants with the functional and research communities they represented:

Wizards:

1. Jason Lee, LBL
2. John Estabrook, NCSA
3. Les Cottrell, SLAC
4. Matt Mathis, PSC
5. Matt Zekauskas, Internet2
6. Rich Carlson, Internet2
7. Terry Gray, UW
8. Laurence Kirchmeier, Merit

Researchers:

1. Alan Whitney (MIT), eLVBI
2. Scott Koranda (UWM), LIGO
3. David C. Bader (LLNL), Climate Model Diagnosis and Inter-comparison
4. Frank Vernon (UCSD), Seismology
5. John Heffner (PSC), Network Engineering
6. Juan Carlos Aragon (Stanford)
7. Tom Bray (UMich)

Applications Developers:

1. Steven Senger (UWI, LaCrosse), HAVnet and VHP
2. Erik Hoffer (UMich), Internet2 Commons
3. Matt Crawford (FNAL), Fermi National Accesserator
4. Bill Allcock (ANL), GridFTP Creator/Maintainer
5. Sean McKee (UMich), MonALISA
6. Ramaswamy Aditya (UMich)

Network Personnel:

1. Chester Ruszczyk (MIT), eVLBI
2. John Bigrow (BNL)
3. Peter O'Neil (UCAR)
4. David Mitchell (UCAR)
5. Nathan Lee (Caltech)
6. Brian Court (CENIC)
7. Scott Reed (PSU)
8. James Ward (PSI)
9. Roy Hockett (UMich)
10. Jason Gerk (SDSC)
11. Dale Fay (Merit)

Level Setting Activity

The first day began with presentations by representatives from each of these communities; the plan was to establish the degree of network knowledge, range of experiences, and special needs of these communities in their interactions with others. The Workshop organizers also gave presentations on their perspectives of the end-to-end problem.

Functional Community Presentations

Wizards: Matt Mathis (PSC), representing the Wizards, showed data on the Wizard Gap and discussed the problems inherent in focusing on helping a small group of researchers achieve wizard-like levels of performance. After years of work in the area, he stressed his conviction that more needs to be done to provide better functionality, 'out-of-the-box,' to all end-users.

He described why TCP/IP works so well (new and old applications work without knowing how the network works) and what is wrong with it (it hides the problems). TCP tuning seems to be more about debugging; he gave six classes of persistent, difficult problems that limit performance - too small TCP retransmission or reassembly buffers, packet losses/congestion, packets arriving out of order or duplicated, "scenic" IP routing or excessive round trip time, improper packet sizes (MTU/MSS). He also discussed inefficient or inappropriate application designs and explained why this means you'll always get worse performance in a crowded area of users (say a Starbucks). He noted that the debugging often tries to find the weakest link in an invisible chain – but if people bang on a link they think is broken this strategy could a) fix it, b) break it, c) break a different weak link, or d) other. For several years, Mathis worked on the Web100 project; it tried to better instrument TCP performance. The problem is that the current Web100 tools still require too much expertise (you still need a wizard in the loop).

Mathis also discussed why nearly all of the symptoms scale with round-trip time (RTT). The local client appears to be 'working' while the remote client sees the application 'fail'. He gave several examples of flaws that scale, such as chatty applications and fixed TCP socket buffer space. Mathis feels that this is 'the' problem. For nearly all network flaws, reduced performance (RTT) is the only evident symptom and, on short paths, it is virtually undetectable. The problem solver always gets a false 'pass' on tests for local equipment, even when the same test fails on long-distance paths.

Questions posed to Mathis ranged from specifics on Web100 to prognostications regarding how experimental TCP stacks will affect the 'fairness' balance. Questions for each of the presenters are documented in Appendix A.

Researchers: Alan Whitney (MIT Haystack), representing the Researchers, gave an overview of his work. Unlike many other scientists and often against his own preferences, he is drawn into network issues and troubleshooting activities because of the type of research he is trying to do. He described the problems faced by his field:

Research is inherently international (telescopes around the world with varying levels of connectivity)
Relatively small community of researchers (most of whom are not very expert in network tools)
Applications can tolerate some loss (they collect Gaussian 'noise'!)

His research community has had several good experiences and has had several successful real-time data transfers. The challenges they've faced have included 1) worrying too much about the details of network connectivity, 2) locating the expertise required for optimal performance in equipment, transport protocol, and parameter tuning , 3) affording to their own dedicated network experts (outside of grant funding for short periods), 4) dealing with competing technologies and uncoordinated efforts within the international community (jumbo frame support for example), and 5) anticipating instabilities in the network (you test on Tuesday for a demo on Friday, but it has changed on Thursday!).

Whitney provided researcher perspectives on needed improvements, including: more education for the science community, more 'simple' tools, more international coordination and standardization, and a catalog of available expertise (where to turn for specific problem types, advice on equipment, etc.).

Application Developers: As a representative Application Developer, Steve Senger (UWisc) talked about this group's diverse challenges, using his own situation as an example. For the Visible Human Project, scientists use Immersive Segmentation (ImSeg) to access anatomical data; Remote Stereo Viewer to make multicast images available and transport them 'on-demand;' and Nomadic Anatomy Viewer, similarly using multicast images that are available on-demand. The testbed includes four Internet2 nodes and three

different networks. He described the problems they've encountered with examples of each type.

One on-going problem is startup delays from LaCrosse into Stanford, sometimes up to 10 minutes. Senger and his team have made several attempts to solve this problem and have narrowed it to the path between Los Angeles and Stanford, via the CENIC network. After much discussion with stakeholders, Senger was told that "that's how it is supposed to work" via Cisco routers. He learned that the transition from Juniper to Cisco routers in Los Angeles started the problem. Cisco changed the 'automatic delay' at each router in CENIC such that the delay disappeared, essentially. Once this was 'solved', however, he began seeing this same problem on another path!

A second problem was a WiscNet Multicast 'weirdness' involving a router problem requiring a router reboot, and it occurred twice in the last six months. Cisco and Juniper routers both had this problem. At this point, the 'solution' is to reboot the router every couple of months but with no explanatory 'rationale.' The theme of better communication with vendors continued throughout the workshop, revealing yet another community that could be included in collaborative efforts.

A third type of problem was tied to conference demos: one required WiscNet to turn off peering with Illinois so that Senger's team could run a demo at McCormick Place in Chicago.

Network Engineers: Roy Hockett, as a network engineer, commented that the problems he encounters are taking more time to figure out, there are fewer engineers available to solve the problems, and there are more researchers using networks. He feels that appropriate terminology in information exchange is important so that diverse groups know each is trying to say. For example, Iperf reports a problem with 'window size' but it might mean the 'send window' or the 'receive window' size is a problem. The end-user needs to work hard to convince network engineers that there is a problem – the engineers don't even have time to solve every problem, but taking time to solve 'false positives' tends to make them less open to problem presentation by end-users. The community needs to change the mind set of network engineers but it also needs to increase the number of engineer hours or the utility of the solution paths.

Coordination, Communication, and Collaboration

Barbara Mirel (UMich) remarked on the end-to-end problem from a communications perspective. She has been looking at what communication breakdowns are occurring. Her specialization is the challenges folks face in complex troubleshooting. Problem-solving steps are ill-defined, there is inadequate information, problems are dynamic, and solving them requires improvisation. Mirel looks at the methods of communication and collaboration to assess how they affect technical outcomes. Awareness of one's own group dynamics and effective levels of communication are as important to solving problems with end-to-end performance as the tools available for this troubleshooting.

She gave an overview of a well-known problem with the EDUCAUSE conference in 2002. The outcome was that the video application both misread the problem and then exacerbated the problem in an attempt to solve it. A 'known' equipment failure was not well-communicated to the community of end-users. Worse, in the post-mortem that followed this failure, the group developed two mutually exclusive 'value sets' – zero-packet loss and 'acceptable packet loss' – based on the application.

Mirel emphasized that, in collaborative work, participants need to be attuned to each other's perspective. In some cases, information may not resolve misunderstandings; remember that we use language to construct identities and roles when in debate. It is important to formulate agreements (and disagreements) before 'consensus' is reached. It is important to pay attention to questions – what are the main things the group is talking about? How do we agree/disagree – who talks most to whom? About what? What kind of influence is being exerted? How much common understanding is shared within group? How readily does the group jump to agreement (and avoid exploring disagreement)? Some of the group dynamics ensure harmony, but this doesn't necessarily solve problems.

Cross-Community Insights: Scenario Collaborations

The central activity that engaged this mixed group of network stakeholders in collaborative discussion and

problem-solving was deliberating about a scenario of a complex network problem. Participants divided into five teams, each composed of members of each community, and worked on a scenario-based simulation of slow response time with the goal of identifying and resolving the causes. Each team member assumed his role and engaged in joint problem solving, starting with the researcher who experiences the slow response time and initiates the collaborative problem solving process. The problem in the scenario was complex, involving multiple problems in different layers of the protocol stack.

After teams reached a resolution to the problem, they reconvened as a large group and watched a wizard simulate the actual problem and demo problem solving processes, tool uses, and collaborative efforts. A group discussion followed. Teams then met in small groups once more to reflect on the implications of their problem solving processes and outcomes for similar everyday problems in network performance affecting the diverse communities. They reported in a large group discussion. (The scenario they were given and the reports presented by each group can be found in Appendix C. One group indicated the steps they took to reach a solution to the problem; those steps are listed in Appendix D.)

Community-Specific Insights: Wrap up Session

The second day of the workshop was devoted to wrap up and next steps. In the first wrap-up session, each functional community reported, as follows, insights its members derived based on their community-specific perspective, needs, goals, and priorities. (Comments collected from each Functional Community are summarized; the full list of comments is found in Appendix E.)

Researchers

Scientific researchers recognize that their cyber-infrastructure-enabled projects require them to be actively involved in end-to-end network problem solving when the technology impedes their work. But given that they are scientists and not engineering experts they expect the following:

- Technologists will provide expertise, as needed – ideally in timely ways and with procedures that do not continue to require researchers to be the main shepherd of problems by default.
- Shared sets of reasonable expectations disseminate across groups.
- Technological advances directed toward user experiences, such as cakeboxes[2]. Ideally, users want to experience that "the technology just happens."
- Education and training; also, an Internet2-sponsored repository of cases for learning geared to purpose. For example, researchers say it would be good to understand what a highly-tuned, high performance box should look like so that the designers of boxes will not just "top off" the valuable segments.

Network Engineers

From a network-centric, troubleshooting vantage point, network engineers stress that they strive to find the limiting factor when bandwidth is there but they're still not getting performance (is it the network? Disk? Query network and storage system?). As important as technical knowledge, network engineers discussed the need to get tools out to everyone and the necessity of getting people together and having open channels of communication through community, as follows:

- Developing a systemic method for getting problem-solvers together on a regular schedule and publishing the availability of this resource to end-users.
- Increasing the deployment of useful tools and coordinating with application developers to include variables to provide problem-solvers with the type of information needed.
- Developing the community of network engineers; this both provides mentoring for new engineers and quicker access to the 'right person'.
- Developing a systemic method for getting network engineers and their campus researchers

2 A 'cakebox' is a mini hand held CPU. It has known characteristics, has been tested with Iperf, and gets Gig performance. If you plug it into your network, it can find its way back to Internet2. Internet2 runs tests to it, etc. When you start off, it helps you validate the quality of your networks. It has been 'morphed' into other projects.

together to discuss performance expectations and realities on a regular basis.

Application Developers

Application developers give high priority to using the network efficiently, finding optimal ways to move data and choose between storing and caching. Awareness of the system is vital in order to know how to best access data. One main strategy they apply involves the distance between the data and jobs, with the goal of running jobs near the data so that latency will not be a worry. They believe they have some awareness – "not infinite awareness but enough so that it's not a matter of chance." An open question is what more needs to be known about the network, especially given that different scientific groups have different priorities and available resources must be balanced among them. They also do not always need "as fast as possible." For example, in scheduling grid computing, the important thing is moving data and knowing the bandwidth so they know/can project how long it with take.

A constant source of difficulty for application developers is that an application may work fine in a laboratory and they expect it to work across the country, but it doesn't. Additional comments include:

- Optimizing networks, which should be systemic.
- Education for developers on changes to the way applications are written that would make more efficient use of the way the network currently works.
- Cost-benefit analysis of increasing bandwidth over application changes.
- Availability of 'true bandwidth' so that applications can take advantage of that data; often, targeting an application at a known available bandwidth is easier and more efficient than attempting to go as fast as possible.

Wizards

Wizards had several key comments, including:

- The problem is bigger than they thought – there is a lot more to consider if researcher needs are considered.
- There is a continuing need for supporting/educational materials (such as cookbooks, war stories, success stories, best practices, etc.).
- Routers must provide more information; pressure on the developers must be widespread to ensure this occurs.
- Process and community-building at a student-level is very important. Learning to work together to solve or avoid problems takes time for trust relationships to be formed and information sharing habits to become routine.
- From a network wizard point of view, it is easy to write papers for peers but determining how much detail to include for less expert readers is a major problem.

Findings: Overriding Concerns

The workshop shows the success of gathering scores of domain experts from diverse communities for 1.5 days of extended discussion about actions, tools, responsibilities, and protocols for collectively solving hard network problems. It resulted in a great deal of newly shared understandings, resolutions of simulated and actual (outside) problems, and clarifications of misconceptions and misconstrued uses of terms and concepts. Unfortunately, it is costly and cannot scale to the level necessary to solve the problem directly.

What is clear from this success is that bridging the wizard gap is feasible – but, beyond interpersonal exchanges like the ones at this workshop, a more global approach must be established. On a widespread scale, we must ensure that the right information and tools get to the right people in the appropriate forms and content for the right times, places, and circumstances.

With that end in mind, the workshop showed that communities' concerns fall into four areas, all of which need further attention.

What tools/information are available and how to use them

Many challenges exist in disseminating and managing the vast and growing repertoire of knowledge and inventory of tools relevant to different communities' respective parts of troubleshooting and otherwise improving network performance.

Participants noted some information as being particularly valuable to have available. It includes:

- Glossaries of terms to use for optimal search results when seeking help and shared meanings of teems and concepts
- Discussions of problem solving by type of problem and tools relevant to each and why
- Explanations of why certain problem-solving strategies are effective and implications of various solutions
- Data pushed to desktops about what is going on in the network
- Exceptions (and the tools and investigations that may be helpful)
- Realistic expectations in each community
- Uses of tools in isolation and in combination
- Explanations of terms and concepts that are not always commonly understood

Focusing on knowledge and tool building, and their appropriate uses, is necessary but not sufficient. Approaches and solutions must be shaped and informed by insights from the other areas. For example, as expressed by one wizard: "We have ideas for disseminating knowledge but will it actually be used? How? What do they need? I can write for my peers but writing for people down the pipe is harder. I don't know what they know." In the same vein, a participant reflected, "I think we need to get at something deeper – more than just learning the knowledge; there's the process that is missing. The community is needed – everyone is doing their own thing in a vacuum."

How to provide tools/information in ways to assure their uptake

Many tools (and published instructional materials) now have advanced to the point of providing end-users with many means for working on performance problems within their designated scope – alone and consultatively. But we do not know if these resources are effectively taken up, understood, and applied by users. If knowledge and tools are to be transferred and used effectively, their presentation must be designed to fit recipients' needs and purposes.

Broad approaches for sharing and building knowledge and tools that participants proposed ranged from virtual office hours with knowledgeable technologists to tutorial road-shows to documentation and "smart" "Googling." Yet participants recognized that these general approaches needed to be designed in ways that

would assure their actual uptake by targeted audiences.

How to devise approaches that "speak" to recipients, are taken up by them, and have lasting effect is an open question. For example, these are some open issues:

- Is it more effective to take a top-down approach to learning, with experts training users and others "up" (e.g. through documentation, tutorials, road-shows); a bottom-up approach with grass-roots initiated blogs and WiKis; or both concurrently?
- More needs to be learned about how to address different kinds of problems and priorities in each community for people with different levels of technical expertise.
- What needs to be included in 'high-speed networking for dummies' material (that we know fairly well and generically) and what, by contrast, needs to be included for problems that are more intermittent and more demanding?
- What are the major confusions that people have about terms and mental models of the system and their interactions with it? What shared language, vocabulary, explanations, and knowledge will help to counter misconceptions and how can they be constructed in ways that "stick?"

How to make organizations responsive to collective efforts

Participants stressed numerous organizational processes that determine the success of joint efforts within and across communities succeed to achieve improved network performance. These include:

- Identifying who should be in charge of collecting war stories and other information for a knowledge base and who should be in charge of maintaining repositories.
- Working across organizations, including international organizations, and getting campus system administrators and vendors more involved in cross-community interactions.
- Getting the resources needed to fill key functions and roles – e.g. liaisons between campus network staff, researchers, and application developers, as well as single points-of-contact at each campus for training.
- Getting funding from appropriate sources for different levels of solution.
- Identifying and getting management buy-in for cost-effective ways to get everyone to work together concurrently – i.e. valuing and building time into people's schedules so that they can work together on a regular basis to share knowledge and tools and solve problems.
- Identifying organizational and disciplinary champions who will implement cross-community exchanges at professional organizations.

How to build and sustain a culture of partnership

Building trust, synergy and partnerships across communities involves group's coordinated uses of tools and collective problem-solving efforts and their other interactions. The small team role-playing scenario revealed positive methods that these cross-functional groups used to move troubleshooting processes from a user's experience through the other communities and their experts (see Appendix D).

Highlights for making these processes run smoothly include the following:

- A researcher should shepherd the problem and get the support he or she needs.
- A partnership should exist between the researcher and a network engineer or wizard.
- A system administrator must be brought into loop.
- Network engineers should have diagnostic scripts for problems that frequently occur to reduce the time invested in 'common' problems.
- Everyone should have a clear sense of the protocol of the help process: whom to ask, who will point the way to more assistance; a consistent information contact – e.g. NOC email address; and an understanding of the "formal" escalation process and "informal" channels.

The last point about protocols underscores the strong workings of social communities. For example, network engineers must know each other. As one network operator emphasized, it can be fatal for people to

spend a day finding the one person they need to have help them with a ‘simple’ fix.” Another noted: “You work out a problem in concentric circles. There’s a trust relationship that you have to build up. Every group is very busy with what they are doing and don’t want to be bothered. It is reciprocal – you have to be willing to give time out to get time back from someone.”

It is equally important to pay attention to what is not apparent in requests for help and what is never asked for at all as it is to attend to the obvious. A start at getting at these unknowns may be to pursue some of the following strategies for building a sense of partnership across communities:

- Strengthen liaison/boundary-crosser functions so that individuals take on a role of bringing the correct people to bear on specific problems and working on “big fixes.” Enacting this function requires being skilled and knowledgeable about engineering and scientific disciplines, able to move smoothly across those boundaries to effectively communicate the needs and expectations of one group to the others.
- Establish mentoring processes.
- Establish virtual office hours or other regular meeting times for cross-community interactions and communications.
- Identify a single point-of-contact at each campus for training, information sharing, and mentoring, preferably involving several layers of connections (Internet2, region, campus, scientific discipline).
- Orchestrate interactions between engineering specialists and scientists at scientists’ regular disciplinary conferences.

Recommendations and Next Steps

Based on findings, we make the following six recommendations:

1. Conduct studies on the best ways to deliver information and tools so that they truly will take hold with end-users.
Findings from these studies should enable us to publish materials and create tool interfaces that will enable users and their support staff to begin solving a problem without having to bring it, unnecessarily, to experts. These usefulness and usability studies should be the tip of a deeper research iceberg. That deeper research should focus on developing a better understanding of the knowledge and learning required to move novices – whose response to slow performance is "something doesn't seem right with the system" – toward becoming savvy users, who know the questions to ask, the basic problem-definition paths to pursue, and the information to provide when escalating a problem.

Suggested areas of investigation should include field studies of scientists to identify paths and patterns of problem solving differentiating novices from network-savvy users as they use published materials and tools; comparative evaluations of differently designed and presented materials sensitive to users' needs and practices; and distinctions across communities in user needs and internalized knowledge from instructive materials.

2. Conduct expert workshops for assessing the completeness of documentation and tools available to users.
These workshops should build on the success of the first workshop and should work in tandem with the studies suggested in (1). The workshops should engage experts in assessing published resources to assure that, from a technical perspective, they are comprehensive and coherent. These efforts should help attain the goals of users being able to easily start problem-solving processes on their own and of wizards gaining a better understanding of what others may not understand. To foster ongoing dialogue between the communities, these workshops also should help experts frame many of the language issues that lead to miscommunications and misconceptions.

3. Create more repositories of information and tools and keep them current.
Continued efforts should be made to enhance and maintain repositories of case studies, war stories, best practices, and problem-solving protocols for different roles and levels of expertise. Also necessary are documentation and glossaries of terms and concepts that often cause confusion and need to be defined for a shared understanding.

Moreover, actions should be taken to develop a repository of keywords so that, in publishing instructive web-based information, writers will be able to apply these terms to their pages and have their pages come up at the top of the list of Google search outcomes. Finally, tool development should continue, and experts should integrate their knowledge and efforts better to identify gaps that still exist and obstruct optimal network performance.

4. Exploit and create new opportunities for partnering across communities and professional networking.
An immediate and feasible action should be to hold "network office hours' - online or through conference calls. In them, people concerned about network performance issues should meet regularly, such as once a month. Other recommendations include establishing mentoring relationships, identifying a single point-of-contact at each campus, and cross-fertilizing ideas and relationships between engineers and scientists through attendance and presentations at each other's professional conferences. Efforts should be made, as well, to bring Sys Admins and vendors into the loop of cross-community interactions.

5. Generate, from decision-makers, a commitment of resources for the goal of bridging the gap.
Many immediate actions require buy-in from upper-level decision-makers so that various professionals have the time to commit to such activities as network office hours, mentoring, and the like. While most of the immediate actions recommended here should not be very resource-intensive, they will still require some

sanctioned support. Moreover, generating commitments of resources for the goal of "bridging the gap" will be essential for progressively implementing findings from the ongoing studies recommended previously.

6. Use the National Science Foundation (NSF) as a conduit to the scientific communities
The NSF should increase the leverage of the existing investments it and other agencies have made in end-to-end network diagnostic technologies by encouraging domain scientists in all divisions to adopt and coordinate these techniques. An important finding from this workshop was the group's overwhelming consensus that NSF should use its contacts at the research community level to spur involvement in cross-community efforts.

Appendix A: Questions from Presentations

Mathis Questions

Q: how many OS's have Web100 on them?
A: We only did Linux; IBM is doing one and there is a draft implementation in "LongHorn" for Microsoft. None of the BSD's have done this.

Q: Is there something about the state of the network that is 'fixed' for me that would cause someone to break it because I am causing problems for commodity internet apps?
A: the only time that happens 'auto-magically' is for window-fairness. If I take a hyper-clean machine going to a close server vs. one going to a far server, the near server starves the distance. There is bandwidth fairness and window fairness and they have different values for networks; it has to be fixed 'in the closets' based on your network's goals.

Q: do you have a sense of how RTTs are changing year by year?
A: they are approaching the speed of light for glass. I don't believe there are lots of new trenches. The current speed maximum is likely to be the maximum for some time to come.

Q: what is changing is where the data is coming from? It depends on where you're trying to find the data from – trying to get data from farther away.
A: Yes easier to imagine cross-country collaborations happening.

Q: how about experimental TCP stacks for Linux – will they be welcome additions? Is this something that we, at the research app developer level, should be worrying about?
A: the problem is that the congestion control algorithm that is in use for the standard TCP has a noise-immunity problem. It isn't (clearly) realistic so the control function needs to be different. The flip side to that problem is that the 'fairness' requires all the end-systems need to be roughly the same. There are so many bottlenecks outside that a few users in a non-standard aggressive protocol is not likely to mess up the network as a whole.

Whitney Questions

Q: any suggestions on good forums for science education?
A: Internet2 meetings are good forums but focused on the experts (that's the perception) and less on the novices. And most researchers have enough meetings of their own to go to. Some simple white papers that a researcher could download (provided by Internet2 or other folks) – "I want to use Internet2, how do I go about it?" Over the web is probably the best way to do that. Simple, non-technical introductions that would lead them to more information as they needed it.

Q: At Fermilab, we've instituted a weekly call – anyone with problems can bring them to the call and, though the problem solving is not 'promised'

Q: One of the problems is that you need to continually buy equipment that conforms to best performance. Need to see if there is a way to tap into another performance-driven group (gamers).
A: Great if there was a common place (aided by Internet2 in the US) to get that sort of information and/or links to places that could solve problems.

Q: Are there meetings in your community that could be used to disseminate information about other communities?
A: Lots of meetings in each community – but they don't 'invite' network experts to come present but having Internet2 folks go to each meeting doesn't scale ($$).

Senger Questions

Q: How many man hours were spent solving the first 2 problems?
A: don't even want to think about it. It took a long time to convince the folks at Stanford to understand that we weren't going to go away and weren't going to stop using multicast anytime soon before they would begin trying to solve the problem!

Hockett Questions

C: Going over multiple backbones owned by multiple groups. One segment goes down and, when you finally determine which segment had the problem it turns out there was a 'regularly scheduled update' that was 'posted' on their calendar. How do you find this out? Do you have to check each calendar (and there won't be a standard place for each place to post this on their website!) to see when the expected updates occur so you can determine how they will affect you.

(Q: couldn't a 'standard' be for the 'networks' to 'push' the data about scheduled updates to each 'user' – i.e., a standard holding station that end-users could 'pull')

Appendix B: Small Group Participants

Group 1

Alan Whitney, Researcher
Bill Allcock, Apps Developer
Edward Vielmetti, Apps Developer
Chester Ruszczyk, Network Engineer
Peter O'Neil, Network Engineer
Matt Mathis, Wizard

Group 2

David C. Bader, Researcher
Tom Bray, Researcher
Sean McKee, Apps Developer
Ramaswamy Aditya, Application Developer
Roy Hockett, Network Engineer
John Estabrook, Wizard
Laurie Kirchmeier, Wizard

Group 3

Frank Vernon, Researcher
Scott Koranda, Researcher
Erik Hoffer, Apps Developer
Brian Court, Network Engineer
Dale Fay, Network Engineer
Terry Gray, Wizard

Group 4

John Heffner, Researcher
Matt Crawford, Apps Developer
Jason Gerk, Network Engineer
David Mitchell, Network Engineer
Scott Reed, Network Engineer
Rich Carlson, Wizard

Group 5

Juan Carlos Aragon, Researcher
Steven Senger, Apps Developer
John Bigrow, Network Engineer
Nathan Lee, Network Engineer
Jason Lee, Wizard
Matt Zekauskas, Wizard

Appendix C: Small Group Reports

Below is the scenario the small groups were given, instructions for the Small Group activity, and a summary of the report each group presented to the whole, upon reassembly.

Scenario

As part of your research, you are trying to transfer a 50 MB file to a remote host. Your colleague must receive the file this morning to prepare for a crucial presentation today. You use SCP to transfer the file because FTP uses clear text passwords. You have every reason to expect the transfer be completed in a matter of seconds. In fact, if you calculated, it should take about 6 seconds. Each host has a 100 Mbps Fast Ethernet connection to the network; the path is over the Abilene backbone; and each campus network has access links should that support this transfer at 100Mbps. Therefore, 50 MB * 8b/B = 400MB/100Mbps = 4 seconds.

Yet the transfer takes much longer. You've run out of patience and you decide to complain.

Instructions

As a group, first decide on a recorder and other roles. Then walk through the problem, elaborating on real life conditions that are likely to occur.

Wizards in the group will observe -- though the group may call them in as part of the problem solving when and if team members feel it is appropriate to do so. Wizards are also there to provide results of any tests you decide to run as part of the problem solving process.

As you work through the problem case, detail your goals, sequences of steps, resources used (other people, tools, information sources), rationales, criteria for judging resources reliable, outcomes of each step, and their effects on choices for next steps.

Start by assuming the researcher role. What do you as a researcher do next?

After fleshing out the researcher's initial experiences -- but still keeping the researcher in the loop -- walk through next steps, shifting roles as needed.

Any time you bring new people and/or resources (e.g. tools, techniques) into the problem-solving, say why you turned to them, what you expect from them, and how you found them. Tie their expertise to evolving conditions in the problem situation. Elaborate on the scenario by bringing in real-life issues that may be likely to happen in problems.

As you proceed through the problem solving – and especially when you hit impasses – spend time talking in the group about what to do and why.

Respond to prompts from the facilitator asking you to elaborate on various issues.

Small Group Reports

Each of the five groups reported back on their progress solving the problem and the insights they garnered from this activity.

Group 1

Group 1 reported that, if "Joe researcher" attempted to solve a problem, what would he do? First: he would go to the local Sys Admin. The group felt it was crucial to focus at the level of the end-user; help the end-user solve problems in a manner that doesn't require the end-user to be a 'network expert.' The group envisioned two levels of support:

1. Web education, or more formal workshops, for the Sys Admin with the expectations being higher than with the end-user.
2. Provide the end-user with the appropriate vocabulary so they can communicate well with the Sys Admin but education at a lower level of problem-solving efforts.

The group recognized that there were many problems to address – hardware range, network problems, etc. It is not possible to test all these combinations. It would be useful if there were some recommendations for 'equipment that works well together' combinations, along with recommendations for default tuning. This would eliminate a number of equipment problems. Debugging difficulties the group identified include:

- A lot of the knowledge in the community is 'folk-lore' based because the technology changes so quickly ('it's the network', 'it's duplex mismatch', etc.).
- Problems in securing accounts on remote machines for testing – this can take a long time (the group discussed the possibility of using Internet2's Bandwidth Test Controller as a possible replacement for using Iperf Servers to reduce the difficulties in securing accounts on machines) and security issues associated with running tests.
- Lack of access to switches (which exists for perfectly good reasons, but this causes major difficulties for debuggers).
- Vocabulary at all levels, so that folks are talking about the same thing (buffer size vs. window size was one example of terms used interchangeably by novices).

Conclusions: Education: at all levels. Education of communities can be scaled but not for individual users. The community needs to develop a vocabulary that can be shared with end-users. "Google has to be the final answer" – end-users need to be able to search for material using 'known terms'; posters must hide related terms in headers to aid folks in finding the information. Overall, need more web pages with better links.

Group 2

Group 2 abandoned the 'scenario' when one of the members of the group (a researcher) reported on a real problem that had been plaguing his research for several months. The group discovered that, when you have all the concerned people at the table to solve a problem, it can take a few minutes to solve a problem; when these same folks are distributed, weeks and months can be required. The afflicted researcher is motivated to help solve the problem and is happy to do a large range of tasks to help the problem-solvers in their task. The group discovered that it is a good idea to have the researcher be the follow-through person because of this motivation to keep the ball rolling; however, having a network-guru to help them (perhaps have that person assigned to the project, for example, at the local level to help researchers find the people they need to solve their problem) greatly aids the problem-solving.

The group concluded that local Sys Admins need to know about NDT – it helps identify problems with the process and is less expensive than getting everyone working on a single problem. They felt that Iperf is for network engineers, not end-user. Until NDT, there was no way to get problem information from the end-user to the Sys Admin in a useful way.

Complex high-performance network problems for applications are lost in the fray of 'common problems' needing attention. The group talked about ways to discuss problems encountered routinely – auto-tuning stacks is not scalable. The community needs to communicate common problems across the research community. Is there an expert system? A problem debugging flow-chart? A Google search for that problem?

The biggest issue is finding ways to know *who* to ask. First- and second-level tech support staff needs to have links to someone who can point the way (described by the group as a 'Mythical Creature' or Network wizard); this person would be able to suggest next-steps and identify resources. The group also suggested that consistent contact information would be useful (i.e., information emailed to 'noc@domain' should go to the same place within each domain).

The group suggested collecting more ‘war stories’ and case studies; Internet2’s End-to-End Performance Initiative (E2Epi) has been doing that for several years but the group wanted both more and better publicized data. There was some talk about how solutions scale (How do you provide incentives for people to add to the knowledge base – FAQs? More meetings like this?) The group realized that over 100K students/faculty are connected to Internet2; they don’t all press for solutions as hard as we do, but they have legitimate problems that we need to solve, now, before they begin clamoring!

The group strongly recommended finding methods to fund the ‘Mythical Creature’ so that he/she would not be pulled into production work, but be free to work on specific problems and existing network problems that haven’t yet caused a ‘crisis.’

Conclusions:

- Local Sys Admins need to know about NDT
- Whomever has control over the host needs to patch scp
- Network engineers need to know what is "good enough" as far as performance goes
- All the layers of abstraction in organizations cause latency
- The researcher is singularly motivated to fix the problem, so they are willing to do a large class of tasks for the problem-solver
- Can knowledgeable and demanding users escalate directly to a "performance guru" (perhaps after a filter/vetting process)?
- How do we educate network engineers?
- Identify problems with the process
- Expensive to get all the participants to work together concurrently
- Iperf-like tools are for network engineers, not for the end-user; tools are only recently getting to be user-friendly (like NDT)
- Tools need flexibility to deal with firewalls, packet-shapers etc. so ability to run tests on specific interfaces/ports/etc. – get as close to the application as possible and need all the devices to be visible
- Complex, high-performance problems are currently lost in the noise of general, mundane (duplex-like) problems – this means that as the bar is pushed, the tolerance is tighter and so solving this now is important
- Discuss ways to deal with identifying problems with the process
- Auto-tuning stacks (not currently feasible because distributing resources on a host, esp. server, is not perfect)
- Setting expectations
- Communicating common problems (tcp buffer/window sizes/ssh buffer) and possibly tools to isolate problems – NDT/Iperf/pathdiag
- “Expert system” or network debugging flowchart
- Discuss ways to improve communications
- Researcher wants the first- or second-level of local tech support to know and figure out (perhaps by contacting a mailing list of similarly-minded people)
- Consistent contact information noc@domain or http://status.domain
- Express/publicize case studies/problems/solutions – NANOG?
- Talk about how solutions will scale
- A web-accessible anecdote archive/FAQ – how to provide incentives for people to add to it?
- Hold more meetings with interested parties when problems crop up or periodically
- 100K U. students and faculty connected to Internet2 who don't press for the problem solutions; they need to get the problems solved soon
- Have a network performance engineer on call who isn’t pulled into routine "production" level work

Group 3

A. Need for knowledge-based tools and techniques for diagnostics.

- Where is the training guide for the neophyte?
- Having a repository of the tools and well explained/understood set of problem protocol taxonomy (a checklist of things to do) –> prescriptive
- Useful: put up a list of expected performance metrics in the speak of end user, e.g. you should be able to transfer X over this segment using this tool

B. Need a 'Mythical Creature' that can bridge the gap between disciplines and engineering

- Can bridge knowledge-base and tool set or knows where to get the tools and techniques
- Not just building technical skills but acquiring/having the systems view...
- Dedicated middle person who would have network engineering experience and interest and some science application background
- Strong ability to facilitate/project manager – with strong project management skills

C. Where is the 'Mythical Creature'?

- Where do you find them, how do you get them, how do you keep them, local or national?
- Making the commitment to train these people – they must have training in networking
- NSF requiring a cyber-infrastructure management plan before funding projects
- Fractional donation of member institutions does not do it completely – accountability – nor does having one central organization be the one that provides (issues can be mainly local)
- Some problems will not be resolved at national level; they have to do with local/campus effort
- Institution-level compliance/campus expectations –> focus less on getting everyone to enlist cat5 but focus should be on upgrading technical skill set, have/recruit the right people (get local expertise locally)
- Can NSF help fund identifying, recruiting, and training a new class of expert (where does this person come from and how do they get funded)?

D. Bottling it – doing it…

- Internet2 has an obligation and an opportunity to increase its outreach activities : telling the discipline scientist "here is how you do it": Add it to the portfolio, enable bridges to technical support staff working for the disciplinary scientist – how do you reach them?, and come up with a tutorial-style road show, take it to the discipline conferences, have a demo but pack it up (open channel)
- More interaction/exploration needed on ways that Interet2 must help

Group 4

End-user education would benefit many people – from end-user researcher to basic administrator – and be scalable. There is also a need for auto-tuning tools and a tool to detect network problems. The group suggested several 'reactive' possibilities: 1) educating those who do NOC support and giving them access to information to help end-users tune their systems (removing buffer size problems and duplex mismatches) and 2) educating folks that there is an overall NOC. The group also strongly suggested that NDT be recommended as a tool you can give end-users access to that would help them walk through their problems. They were interested in funding the development of a global tool but realized that there was no 'ideal way' to address every issue an end-user confronts. This group emphasized developing a method to assist users to find their own help. **Active**:

- How to avoid the problems? [Auto-tuning can reduce buffer problems. Is buffer tuning really a job for application developers? Or should it be below the neck in the hourglass so to speak?]
- How to communicate to users that solutions exist? [Not all users realize that small changes may realize drastic performance improvements. Use your NOC / Help Desk to help get pointers out.]

Reactive: How to diagnose problems more quickly? [Tools such as NDT, etc.]

Group 5

1. Improve environment in which end-users and wizards work to solve network performance problems.

- Need culture of cooperative problem-solving
- Involves a kind of trust by wizards in abilities of end-users and responsibility incurred by providing information or tools
- Framework of common goals
- Process for developing documentation of problems and solutions that facilitates searching by users
- Analogy with medical profession's use of diagnostics scripts

2. Diagnostic knowledge-enabled searchable database – Combine semantic net of ideas with searchable database of problem/solutions (a.k.a., 'smart' Google)

3. Ideas for active efforts, such as: Watch traffic to look for common problems, report to users and work to improve default performance tuning of standard OS/App distributions

Note: Specific steps in the problem-solving process were collected from members of this group and are included in Appendix D.

Appendix D: Problem-Solving Steps

Group 5 reported on problems encountered in the morning session: firewall configuration can both prevent applications from working as well as degrade performance in a way opaque to the end-user, e.g. early PIX firewall behavior dropped tcp scaling options. Specific steps were collected from members of the group:

Researcher:

- Suggests performing non-tcp measurement of throughput, e.g. Chariot
- Stop and start scp to see if problem goes away
- Try using different hosts, isolate hosts versus network

Wizard:

- Reports problem(s) repeatable on same machine
- Reports Chariot would be just fine
- Possibly might get better with different hosts but if end-user switchs machines with same network ports, problem would remain

Researcher: Contact local network support, inform of test results

Wizard:

- Clarify that researcher would only do UDP in direction of transfer
- Clarify that researcher would do test from another source location to check problem but use same network connection
- Asks how long researcher would spend before contacting tech support?

Researcher: Approximately half an hour or until ideas exhausted

Wizard: If research doesn't want to spend time, what quick things to check or find a different way

Network Engineer: Suggests firewall. Look at where attached to infrastructure; mentions auto-negotiation.

Wizard: Asks what would you tell researcher to do after his results?

Network Engineer: Check for duplex mismatch

Wizard:

- Reports lots of errors on interface, switch interface would be nailed at 100 full
- Tells researcher to reboot
- Says with switch nailed at 100 full, reboot does not change results

Network Engineer: Try putting both switch and hosts in auto

Wizard: Would Network Engineer ask JCA to check interface config; if above does not show anything, what next?

Network Engineer:

- Check transfers at different distances
- Try different transfer mechanisms
- Look at endpoint IPs

Wizard: Would Network Engineer go watch transfer? Would put a machine on Researcher's subnet to check?

Application Developer: Asks Network Engineer how firewall would enter problem

Network Engineer: Is flow but slow, not suspect firewall (Interesting problems introduced by PIX firewall)

Application Developer: Reply to question about experience with Lab, have I involved others?

Wizard: Solving duplex improves speed but still not theoretical speed. What next?

Network Engineer: Look at tuning particular hosts

Researcher: Would enlist more support beyond campus network

Application Developer: Google for "slow scp transfer"

Network Engineer: Demonstrate best case connecting near machines

Application Developer: Compare iperf tcp against scp

Wizard: Asks Network Engineer again about tuning

Application Developer: Asks how much tuning Windows improves tcp flows

Network Engineer: Says substantial

Wizard: Asks Network Engineer about information available for WAN providers

Researcher: Look at reverse flow

Wizard: Asks if both directions have problem what next? (Back to tuning, can't install new kernels because of desire for consistent security situation)

Network Engineer: Under some circumstances, would point person to information on tuning, cannot involve help desk, would inform of dangers

Wizard: Back to TCP-based test works better then SCP.

Application Developer: Where does evidence come from that SCP should be better?

Appendix E: Functional Community Comments

Researcher comments included:

- The important thing is to have the expertise available to the researcher – not the expectation that the researcher WAS the expert
- People don't know what they need. Don't know what to expect. Don't know that they should contact someone about their 'performance'.
- If you want to move a gig file on Internet2 from pt A to B, a list of metrics so you know what you should see. Such as with GridFTP with x number of channels.
- That could end up being a list with about 30 variables – a whole matrix.
- Even if you could get people to the general order of magnitude of what to expect. General OS information.
- The netflows that Internet2 captures weekly can show the peak/curve – a lot of people are getting slow performance and there are specific groups that are getting good perf – they've 'figured it out.' Campuses could do that on their campus.
- Internet2 has cakeboxes – reference systems that could be made available so that people could run tests to check their network. Internet2 could provide. It might help problem-solvers if you could ship them a box that had been 'qualified'.
- Could Internet2 be a collection point for high-speed experiences – have the researchers say what they're working. Here's what we've used, here's what we've been getting with it, etc.
- You want to post it on the web but… that could be as simple as sending mail to a mail list or it could go on a web page, database, wiki, what? R: database – you could refer researchers to that dbase so they could see what people are using.
- One of the problems with this is the network is 'invisible' in this – unless you have Web100 and can grab the stats on that, the performance is based on the quality of the network you were using at the given moment of the test.
- You might not even be able to put up such a system that duplicates (parallel to the ones you've read about) another persons experience because technology changes so fast.
- People be come outraged that they cannot transfer data at 'x' rate, because they don't have reasonable expectations. The anecdotal evidence – set expectations.
- Why are we worrying about 'moving files'? the technology should just happen.
- Ultimately you're still moving files and ultimately people still care about the speed at which they can do it.
- It may be part of what we need to diffuse the benefit from this meeting is to have meetings at JTs, Internet2 meetings, NANOG to have sub-meetings that focus on performance tuning. It would be good to understand what a highly-tuned, high-perf box should look like. That way, the designers of boxes will not just 'lop off' the valuable segments.
- We've talked about that at Internet2, having a 'certified performance package' discussion so that people could just order it from x provider.
- We've also talked about having random tests to locations at universities in our membership to 'rank' the quality of the networks of our members.
- Somewhere in Internet2, there are tables on whether there are jumbo frames etc. there are things that should be added to those tables. You'd like to indicate whether there are hosts USING the jumbo frames, etc.
- This 'rating' is fraught with political problems – people don't want to 'look bad'.
- You could publicly list the 10 best and privately provide each campus with a report card that they could take to their provost.
- There's a project underway to get campuses and GigaPoPs to deploy some of our tools (NDT< BWCTL, OWAMP) at the GigaPoP and campus-edge level so that there are places to which people could run tests.
- CENIC, for example, have agreed to run these so they have testing points deployed.

Applications Developer comments:

- Researchers will run tests at one center, the data will be sent in small packets to each sub-center and then, when that packets has been reviewed, it gets heaved and another section of the whole (another small packet) gets sent along to the researchers.
- Meta-data is an important part of this – network data storage requires meta-data to make it easier to identify what pieces are where.
- How much thought goes into the decision to have storage in one place and the compute-cluster of machines in another?
- You need to meta-schedule that – how do you arrange to have your schedule of computers match up with the schedule of data transfer time and match up with specific storage areas. You can get latency if the storage for the data is far away so you want to meta-schedule it such that the transfers happen at the best time, etc.
- do you think most researchers understand the difficulties with this?
- I don't think most of them do – with physics, they understand that idea that they need to move large data sets long distances on a routine basis.
- More of it is determining the major problems that might be occurring rather than the twiddling of dials to fine tune the transfer itself.
- The system and the whole needs to worry about how it optimizes the jobs and get the most work done. Complicated calculation – need to understand what you can do with the network (any changes I can make to make the network more efficient?).
- I'm not sure, except in some cases (like Sean's), that instrumenting the application will take longer than increasing the bw so I don't need to.
- What I often see is that I ask something to go 'at x speed' (vs. 'as fast as possible') – in many cases, knowing the bw would be more useful than going as fast as possible. I can schedule transfer time, etc. if I know when it would start and how long it will take. We're going to add scheduling into GridFTP based on the BW/Disk capability.

Network Engineers comments:

- One thing I got out of this conference was 'availability' – the biggest hurdle is getting people together. Setting it up (weekly conference call, etc.) where there's a set time so people would know that 'these resources will be there'.
- For VDT, they have 'office hours' where they go to a virtual space at a specific time and people can join to ask questions.
- Getting the tools out there that the apps developers can run so the network engineers can use the data to solve problems.
- How do you bootstrap this – there is very little overlap between network engineers who are busy running their networks and researchers who are doing their research.
- Getting everyone into one place seems to be the answer. Shaving off the separation layers.
- NEES has done that pretty well. Could you describe.
- Everything comes in via one point-of-contact. We have user training, where they go around and hold workshops to show users how to run tests, etc. We are more interested in how we work with the sys admins at specific research sites. We need to figure that out.
- We work a lot with Europe and Asia – we can't just say 'we have tools, we're going to deploy them, etc.". We have to work with their 'system' of dealing with problems which can take too long for the data logs lifecycle. It comes out to be a bigger problem sometimes because of the
- We've run into that – system people have a single cube ticket; someone goes into that ticket from the network end and, if they don't get access into the system quickly, it can be a long and tedious process to solve a simple problem.
- Network engineers like to know each other
- It is required for folks to know each other – it can be fatal for people to spend a day finding the one person they need to have help them with a 'simple' fix.
- I'm just getting into this field and it is very hard to get up to speed to meet the right people, etc.
- Do the folks who know each other in Europe, do they bypass the 'system' to get things done?

- I haven't seen that, so far.
- it is important to have a culture where all the parties involved get TOGETHER to solve the problem vs. having the tools or handing the problem off to someone else to 'solve it' or solve that segment.
- The reason that Internet2 can quickly solve problems is that 1) we know lots of people and 2) we can call on a lot more people. It is very hard if you don't know people to get
- How do you grow that cadre of connections?
- I go to JTs to talk with people.
- You develop it over time – you work out a problem in concentric circles – I don't have a rolodex of names I could hand over to someone new. Besides, there's a trust-relationship that you have to build up.
- Problems DO get prioritized – even if you have a contact, they have other priorities.
- It is reciprocal – you have to be willing to give time out to get time back from someone.
- Is JTs international?
- We just changed the name – we've been getting lots of international registrations. Primarily US (co sponsored by Internet2 and ESnet) but, because we deal with global research projects we're going more global. The last one was with CANARIE and a previous one was with APAN.
- very informal networking.
- how do you socialize people in the network engineering area into the discipline issue problems? Maybe need internships – work in a community.
- every group is very busy with what they are doing and don't want to be bothered
- How many network engineers were invited, how were they selected and how many declined to come?
- as many of you know, this was rescheduled – it was done so because we had no researchers and apps developers who were available to attend. We had to go out and actively recruit people.
- the biggest response was from network engineers – they wanted to come.
- we arranged for the wizards ahead of time but the apps developers and researchers we had to solicit from people we know. So, even for this workshop, it was hard to create the interest in this intersection.

Wizards comments:

- I learned the problem is bigger than I thought it was. There are a lot more things I hadn't been considering that researchers need to keep in mind. Problems I've never had to worry about because, in my org, there's a 'line' on 'network performance debugging' that is far beyond what other orgs would consider necessary.
- Bit of déjà vu from the organizing meting on forming the E2Epi – need for cookbooks, war stories, success stories, etc. that we' heard about again here. There is a latent need for more mundane things on how to get the information out – such as best practices, etc. I'm also interested in the intersection between the perf diagnostic problem and the network failure diagnostic problem (introduced by the 'security' gizmos). Wondering what we need to do to put pressure on router developers, etc. to deal with why they don't tell us why they've dropped packets, etc. fundamental issues that get extremely difficult (info out of the black box). I'm not sure how to proceed to make progress on either of those but I think they're important.
- I'd reiterate the 'déjà vu' thing – in many areas, we've tried to do 'war stories' and put info on the web; still need to do that more.
- We've tried to do that and had a difficult time collecting them!
- Need Internet2 to sponsor the LAN diagnostic record!
- do we need to get at this when they're closer to their student years? When folks are learning to code. I don't know if we're discussing the root problem.
- if you were learning to code in 1992, you were concerned with congestion control, I think that if you're teaching them the problem will change underneath them.
- I think we're getting at something deeper – more than just learning the knowledge; there's the process that is missing. The community is needed – everyone is doing their own thing in a

vacuum. To get that going at a student-level would be very important.

- If we'd started training folks to work together 10 years ago, we'd have a different view of the process.
- I hope that, as we're leaving, folks understand that there's problems across all these areas and that they're interconnected. HENP has known that they are interconnected with their problems and have been working on it for years; VLBI is just learning it. How about the other discipline communities? Do we leave them to learn it on their own? Do we try to communicate with them about this? Do we get the funding agencies to understand that this is an important process for them.
- Is there a knowledge transfer in Internet2 communities to researchers in various communities?
- The business I'm in is, inherently, international. One of the things that's so useful to us is the International reception so that we can make f2f contacts with people we've been talking with for some period. Knowing who to talk to in various countries really helps out.
- one thing I'm curious about – do you think that the bidding on the new array is going to a) increase collaboration or b) increase competition?
- both. It is a multi-billion dollar project – lots of jockeying at this point. The radio astronomy community is pretty small so we're always working together anyway. It forces the collaboration.
- It also forces the competition, a lot with funding, depending on how much it is being pushed in specific countries.
- That's healthy – there needs to be competition for ideas and methods, etc. It is, however, a playing field that is tilted, hugely, by the amount of money that is available for the research…
- google is the answer because everyone uses it.
- google directs you to methods you can use to find ways to solve a problem.
- What are the keywords to put into web docs to get it googled?
- Will it be used? How? What do they need?
- How much time should be spent on this? How much time has been spent?
- The hard problem is the trickle-down aspect – I can write for my peers but writing for people down the pipe is harder. I don't know what they know?
- If you go down that path, people expect you to finish it.
- Who is in charge of this?
- Whoever touches it last. Sometimes I don't touch something because I don't want to 'own' it.
- How do you approach this?
- You feed a wiki with false information to anger one of the Matt's!
- You start a blog or wiki or message boards so you have a complete loop.
- Do you think this is a good idea? What is the reliability of the source?
- Post, post, post. They can try out whatever they want or appears to meet their needs. If it doesn't work, they'll provide feedback.
- That was one of the things we were trying to establish this – what you're missing is the 'why' – why should I search for this? If I'm searching for A, should I also be searching for B, C, and D? The Matt's know this but how do I know it? Only through talking to them!
- Need to get down to the next level – keep going over the same ground because you haven't trickled the info out and down or up to funding agents. What things are people needing? Need to identify the kinds of real problems and priorities you need to address – otherwise, it is a jump from a problem to a solution.
- Documentation is good and dissemination is not impossible but it takes a lot of time!
- Very important point – it is 'the last thing we do' but you keep talking about how it is 'built into the process' (as has been mentioned throughout) – but it keeps getting done at the end on people's personal time. It has been clear that you have a wide span of users – some are near expert and some need things 'idiot-proofed'. Knowing what your range of users are will address the range of documentation. Funding/focusing on the importance of documentation.
- Some orgs identify that they need to reach out to their users but the sticking point is funding – we need to get the funding agencies involved in this at a larger level.
- Would it scale? Every community needs a different Point of Contact.
- A central point of contact – point people to the right place.

- Another model is to ask a central organization, such as Internet2, to fund it out of subscriber fees.
- Work with your community but there are thousands of them.
- Next step up is the funding agency that is funding all the research projects.
- Incentive for the program officers needs to come from the research community and the network people.
- There's been research on this topic – knowledge management; it is widespread across many fields – people leaving an org (retirement) that have valuable info that needs to be collected and maintained. Very hard problem to solve based on the breadth of it.
- There's confusion between many terms; some look like the same thing but are very different. Sometimes someone uses one term 'correctly' but someone else will use the same term 'incorrectly' – how do you differentiate between what the user wants to use.
- Discussion of the different meanings of 'lag' – different causes, different symptoms.
- Note: this is clear that this is related to the dissemination of information. Akin to a patient telling a physician that they have a stomach ache – the physician knows of a wide range of causes for this and a wide range of smaller symptoms to suggest.
- Who are the right people to bring to the table: novice researchers, folks outside the U.s., resource allocation folks (those with control of the $), vendors (Cisco, etc.), community of people doing domain-science research (not 'network wizards' but 'network researchers') – they have a different view of what 'performance' means.
- If this wasn't in A2, would you (a researcher) have come to this workshop?
- If it was convenient and I didn't have anything of a higher priority, yes. I'm glad I cam and I learned a lot (besides having my problem solved – if I'd KNOWN it would be solved, I would have come, despite any difficulties!) that made it valuable but I didn't expect it to be as valuable as it was.
- We keep talking about tts and talking to the NOC… need representatives from the NOC/TT workers.
- We've dealt with the 'high-speed networking for dummies' material – some problems are more intermittent, more demanding. Some groups have a regularly-scheduled 'hack-fest' to work on problems. We'd need to id who has major problems that would respond well to group collaboration.
- Communication – shared vocabulary, create a 'Rosetta stone' of terms, and keep an open line of communication – regular phone calls/meetings? Intentional conferences. Dinner with 12 strangers – setup meeting with 12 people who don't really know each other to expand trust relationships. Gets more new people into the grouping.
- Re: setting expectations re: performance, would it be reasonable to setout specifications of what each out-of-the-box component is capable of performing?
- We're looking at documenting what we've achieved with various hardware but, at this point, the data is unreliable and incomplete. We have default recommendations for vendors on things like TCP stacks, etc. Vendors are concerned about changing defaults because that opens them up to additional potential problems.
- what is the time/cost tradeoff is for 'good enough' – everyone has a different view of what you need and what that's worth.
- the MonaLISA EMMA client gets local information when running tests like NDT that provides details re: NIC, etc. Over time, you could collect information on various pieces of equipment that would identify what the average user is actually 'getting' with specific hardware. But who hosts the central repository and how do you protect that from DoS attacks…

General Comments

- The one that crosses the boundaries and makes/maintains the lines of communication. Even if you talk about having that Liaison role, who funds it? Still need to understand that there is that jump between the problem and the solution – need to identify how to improve communication so that this breaks down the problems to more correctly address the problem range.
- What kind of action items should we take out of this to keep this effort going? Getting groups together – mixture of folks (discipline basis? Regional basis?)

- interact with user disciplines where they are already attending – get the network folks to the disciplines.
- A very hard to make a presentation there – is there a way around that? A poster session?
- Or via NSF? They fund much of this research – if the funding agency recognizes it it will help.
- who should come – the networker? The researcher? The Internet2 connection?
- Internet2 rep and/or someone from a similar basis. Develop a more-or-less canned presentation that opens the eyes of the researchers on what they can do, what they should expect, and serve as a conduit for more efforts.
- Network 'swat' team that could go to network-specific conferences.
- this is exactly the stuff that NLARN has done and , for various reasons, it hasn't continued.
- I have had a positive experience with this, via Internet2, but several other communities have NOT had a similar experience. Internet2 was there from the beginning and challenged/nurtured us.
- How much effort was involved on Charles' part – is that scalable?
- By going to some of their conferences, I managed to convince them to hold a 1day workshop on network performance. After badgering them, they grudgingly convinced them to do so but after holding the conference, they realized they could have done several more days worth – they will continue doing this. I was fairly successful with VLBI because they were already going pretty well, knew what they were doing, but the VLBI folks really needed a guide.
- We were directed at an early stage of our work with the physics community and wrote several proposals with them; but the problem that we ran into was that NSF funding was available for the physics portion but not for the astronomy portion.
- Ways to engage projects seems successful – another approach might be campus-level.
- Having done an Internet2 day at a campus, it is very difficult to engage with someone unless they are involved in a project and they have encountered problems. Campuses have so many diffuse needs but we haven't been as successful
- You might engage a single point of contact on campuses to provide them with information with so they can reach out to the researchers on campuses to offer help with problems.
- The work with a research group really requires someone listening to what the group is doing, seeing their needs and problems, and having discussions with them efore bringing in experts.
- Would campuses be willing to pay to bring 3-4 people to a campus to help a range of folks?
- Might offend network folks on campus.
- We can have some trainings that are done by video for a little less but folks would be willing to pay to fly experts out sometimes.
- Doesn't really scale.
- We need to find some method to 'hide' the network from the end-user, eventually, so they don't need to have any expertise; just run their app and not have to worry about the pipelines and speed etc.
- The biggest barrier to having a transparent network for the end-user is that, ironically, there is too much hidden from the network from the network engineers. Re: swat team coming out, how much could they do without the time/effort of the campus engineers?
- SWAT teams are stymied by things they cannot control – "the part will be in next week," or "the person you need to talk to isn't here," etc.
- We talked about 'virtual office hours' – if people knew what the hours were – or an email address for 'interesting network problems?
- What the Quilt was trying to do was getting adequate resources together to solve problems
- Many pronged attack – virtual office hours, email list, regular meetings, etc.
- Who isn't complaining loud enough?
- There are many out there who have either given up or don't know who to complain to, yet.