Title Page

Report to the National Science Foundation
Directorate for Computer and Information Science and Engineering (CISE)
Advanced Networking Infrastructure & Research Division (ANIR)

NSF ANIR Workshop on Experimental Infostructure Networks

Submitted to NSF November 5, 2002

FINAL REPORT

The NSF ANIR Workshop on Experimental Infostructure Networks was supported by NSF under grant ANI-0227640.

November 2002. Report prepared by Larry Smarr, California Institute for Telecommunications and Information Technology [Cal-(IT)2], University of California, San Diego; George Clapp, Telcordia Technologies, Inc.; and, Thomas A. DeFanti and Maxine D. Brown, Electronic Visualization Laboratory, University of Illinois at Chicago.

Authors and Contributors

The following individuals attended the Workshop and contributed to the writing of this report.

Attendees/Participants/Committee		
Name	**Institution**	**email**
Aiken, Robert J. (Bob)	Cisco	raiken@cisco.com
Bajcsy, Ruzena	UC-Berkeley/CITRIS	bajcsy@eecs.berkeley.edu
Benner, Alan	IBM	bennera@us.ibm.com
Birkner, Franz (Cal-(IT)^2 staff)	San Diego Telecom Council	franz@fxbventures.com
Bortolini, Ed	Network Photonics	ed.bortolini@networkphotonics.com
Brown, Maxine (program co-chair)	UIC and UCSD/Cal-(IT)^2	maxine@uic.edu
Clapp, George	Telcordia	clapp@research.telcordia.com
Corbató, Steve	UCAID	corbato@internet2.edu
DeFanti, Tom (program co-chair)	UIC	tom@uic.edu
Dennison, Larry	Avici	larry@avici.com
Edwards, Hal	Nortel	hal.edwards@nortelnetworks.com
Hidley, Greg (Cal-(IT)^2 staff)	UCSD/Cal-(IT)^2	ghidley@soe.ucsd.edu
Huntoon, Wendy	PSC/The Quilt	huntoon@psc.edu
Hutchins, Ron	Georgia Tech	ron.hutchins@oit.gatech.edu
Jamison, John	Juniper	jjamison@juniper.net
Johnson, Ron	University of Washington	ronj@cac.washington.edu
Kamrath, Anke	UCSD/SDSC	kamratha@sdsc.edu
Lehman, Tom	USC	tlehman@isi.edu
Mambretti, Joe	Northwestern University	j-mambretti@nwu.edu
Myers, Tom	TeraBurst	tmyers@teraburst.com
Obsitnik, Paul	ONI	pobsitnik@oni.com
Papadopoulos, Phil (Cal-(IT)^2 staff)	UCSD/SDSC/Cal-(IT)^2	phil@sdsc.edu
Shalf, John	LBNL	jshalf@lbl.gov
Sharma, Rohit	ONI	rsharma@oni.com
Smarr, Larry (general workshop chair)	UCSD/Cal-(IT)^2	lsmarr@ucsd.edu
Sobieski, Jerry	MAX	jerrys@maxgigapop.net
Verrant, Jeff	ONI	JVerrant@oni.com
Wallach, Steven J.	Chiaro	swallach@chiaro.com
West, Tom	CENIC	twest@cenic.org
Wing, Bill	ORNL	wingwr@ornl.gov
Winkler, Linda	ANL	winkler@mcs.anl.gov
Wolff, Rich	Telcordia	rsw@research.telcordia.com

Ed Lazowska of the University of Washington, <lazowska@cs.washington.edu>, did not attend the Workshop, but helped to organize it and reviewed drafts of the report.

Acknowledgments

We wish to express our appreciation to the NSF Directorate for Computer and Information Science and Engineering (CISE) and its Divisions of Advanced Networking Infrastructure & Research (ANIR) and Advanced Computational Infrastructure & Research (ACIR) for their support and encouragement. In particular, we wish to acknowledge the assistance and participation of the following individuals.

Name	**Institution**	**email**
Blatecky, Alan	NSF	ablateck@nsf.gov
Bush, Aubrey	NSF	abush@nsf.gov
Greene, Tom	NSF	tgreene@nsf.gov
Hirsh, Rich	NSF	rhirsh@nsf.gov
Maeda, Mari	NSF	mmaeda@nsf.gov

Table of Contents

A. Executive Summary

A.1. Creating a National Program for Experimental Infostructure Networks

On May 20-21, 2002, a NSF-funded Workshop was held at the Arnold & Mabel Beckman Center of the National Academies of Sciences & Engineering in Irvine, California, to gather the opinions and recommendations of a dozen key corporate technology leaders regarding the future of advanced research and education networks. In particular, the focus was on Experimental Networks as differentiated from Operational Research (i.e., Production) and Research Networks. Also attending were NSF personnel who fund advanced networking and computing, and academics who help build computing and networking information infrastructures (“infostructures” for short).

The goal of the May Workshop was to answer these questions: *Does the computer and telecommunications industry agree on a need for Experimental Infostructure Networks? If yes, for what reasons, and what should they look like in the near and long term? What roles should industry, government and academia play?*

A.2. Workshop Conclusions

Workshop attendees concluded that an *applications-driven focus* for Experimental Networks is of utmost importance. The key concept is not so much the technology research but rather the vertical integration effort (from network to middleware to application to user interface) over multiple application domains. Definition and bounding of the applications and technical communities are also important. The various relationships and purposes of Experimental, Research, and Operational/Production Networks to each other were discussed in detail.

The benefits of Experimental Networks to industry were concluded to be several. Foremost is “lessons learned” from exploring and testing advanced networking concepts through collaborations among researchers from one or more companies and academic organizations who would not otherwise work together; academic organizations with major distributed applications then become company-neutral “meet-me” places. NSF’s essential role is seen by industry to be a supporter of very high-risk experiments staged in more realistic environments than could possibly be set up in a laboratory. Industry benefits from NSF’s investment by reducing their costs and associated risks while garnering early customer validation of ideas, without fear of market embarrassment.

Industry researchers should ideally have opportunities for initial and continuous participation in the program, as well as opportunities for funding, rather than solely have a supplier relationship with academic grantees. For sustainable partnering relationships, companies need to conduct research on topics that are on the critical path to new developments, as well as gain realistic expectations of the benefits that result from their involvement.

Workshop participants concluded that private companies and venture capitalists are not generally able to fund Experimental Network infrastructure and research. If the Federal government does not fund Experimental Networks, and does not leverage the installed and emerging base of regional/campus infrastructures, Experimental Infostructure Networks will not be constructed.

A.3. Workshop Recommendations

The Workshop recommended a budget for a NSF Experimental Infostructure Networks Program: $50M-$75M spread over 5-7 projects with 5-year awards (administered as a $10M-$15M/year program), with the last two years of funds subject to mid-award reviews. Projects need 5-year funding to reasonably attract allied faculty, enable applications, and engage industry.

Multi-institution awards and company sub-awards are recommended, as is support for academic professional staff as peers with industry researchers.

Another recommendation is that an NSF Experimental Infostructure Networks Program should especially leverage the installed (and emerging) base of local and regional infrastructure.

The Program should fund delivered *end-to-end* connection of all resources involved in each experiment, *and* a demonstration of controls of network capabilities that facilitate applications.

The recommended highest-level review criterion is *innovation* rather than *geographic scope*.

The Experimental Infostructure Networks Program should call an annual meeting to encourage interchange between the grantees, vendors and national Operational Research/Production Network providers. In addition, NSF should encourage strong interactions between its funded Research Network and Experimental Network recipients, as well as encourage ties to other Federal Agency and higher-education networks.

B. Workshop Motivation

On May 20-21, 2002, the California Institute for Telecommunications and Information Technology [Cal-(IT)^2][1] hosted a NSF-funded Workshop at the Arnold & Mabel Beckman Center of the National Academies of Sciences & Engineering in Irvine.

The goal of this Workshop was to assemble the opinions and recommendations of a dozen key corporate visionaries regarding future development of advanced research and education networks. Also invited were NSF and other Federal Agency personnel who fund advanced networking, and several academics who help build these networks.

This Workshop was organized at the request of the NSF CISE/ANIR Division as follow-on to a December 2001 workshop called "Grand Challenges in e-Science." The report from the December Workshop (available at <http://www.evl.uic.edu/activity/NSF/index.html>) contains recommendations from the academic community on what "delivered services" future networks must provide in order to best enable solutions to complex e-Science problems.

Industry participants were asked to read this report's recommendations and present a 15-minute talk reacting to and/or amplifying the concept of Experimental Infostructure Networks, as a way to create a common ground for the main workshop sessions and discussions. Most of these talks are available on <http://www.calit2.net/events/2002/nsf/index.html>.

The stated purpose of the May Workshop was to answer the questions: *Does the computer and telecommunications industry agree on a need for Experimental Infostructure Networks for e-Science? If yes, for what reasons, and what should they look like in the near and long term? What roles should industry, government and academia play?*

This Report is the outcome of the May Workshop, primarily offering industry-academic-government infrastructure-building concepts for Experimental Infostructure Networks to help NSF define substantial new funding programs. Section C offers definitions, conclusions and recommendations that arose from three breakout sessions, which met independently for several hours each (see Section E for the precise schedule). Section D documents each breakout group's discussions, which are retained for reference; any overlaps were subsequently combined into the Report's conclusions and recommendations, which were discussed during the plenary sessions and writing activities.

---

[1] Cal-(IT)^2 is part of the University of California San Diego (UCSD) and University of California Irvine (UCI).

C. Experimental Infostructure Networks

C.1. Definitions

The following definitions are from the final report of the "Grand Challenges in e-Science" Workshop.[2]

Cyber-infrastructure encompasses the emerging computational, visualization, data storage, instrumentation and networking technologies that support our Nation's major science and engineering research facilities, enabling *e-Science*, or large-scale science that studies very complex micro to macro-scale problems over time and space. Cyber-infrastructure includes computers and networks, and the Grid middleware that enables coordinated resource sharing and problem solving among distributed facilities.

Networks are the key enabling technology for transforming cyber-infrastructure from geographically separated computational facilities and instruments into a National Information Infrastructure. NSF ANIR director Aubrey Bush describes three classes of Research & Education (R&E) networks beyond the commodity Internet:

> ***Production, or Operational Research, Networks*****:** High-performance networks, which are always available and dependable. Examples are the *FedNets* (ESnet, DREN, NREN, etc.), UCAID/ Internet2's Abilene and WorldCom's vBNS+. These networks reach all US researchers who need them and must be 24/7 reliable.
>
> ***Experimental Networks*****:** High-performance trials of cutting-edge networks that are based on advanced application needs unsupported by existing production networks' services. Experimental Networks must be robust enough to support *application-dictated* development of application software toolkits, middleware, computing and networking. Experimental Networks must provide *delivered experimental services* on a persistent basis and yet encourage experimentation with innovative and novel concepts. Experimental Networks are seen as the (missing) link between Research and Production Networks.
>
> ***Research Networks*****:** Smaller-scale network prototypes, which enable basic scientific and engineering network research and the testing of component technologies, protocols, network architectures, and so on. Research Networks are not expected to be persistent, nor are they expected to support applications beyond their own testing and evaluation.

C.2. Workshop Conclusions

C.2.a. Differentiating Experimental, Research and Production Networks

Workshop attendees concluded that an *applications-driven focus* for Experimental Networks is of utmost importance. The key concept is not so much the technology research but rather the vertical integration effort (from network to middleware to application to user interface) over multiple application domains. Definition and bounding of the applications and technical

---

[2] The workshop report can be found at: <http://www.evl.uic.edu/activity/NSF/index.html>

communities are also important. The various relationships, interactions and purposes of Production, Experimental and Research Networks are summarized below.

C.2.b. Relationship of Experimental to Research Networks

- Experimental Networks are encouraged to make use of novel ideas emerging from Research Networks.
- Research Networks are encouraged to port new components to Experimental Networks when ready for persistent use.
- Experimental Networks should specify new architectures and capabilities, based on application requirements, and identify critical bottlenecks, in order to motivate future Research Network activities.
- An Experimental Networks Program should fund the development and testing of systems-level integration – of leading-edge telecommunications hardware with network management software, operating systems, Open Source middleware and applications – to investigate pre-market strategies.
- A Research Networks Program should fund access to dark fiber, wavelengths, lightpaths, and so on, so academic and corporate research partners can investigate strategies pertaining to both open and proprietary (potentially patentable) hardware and middleware.

Moreover, Workshop participants concluded that NSF should develop strong *interaction* between its funded Research and Experimental Network programs. Such interaction could facilitate the successful migration of application requirements from Experimental to Research and of prototype technologies from Research to Experimental. While Research and Experimental Networks have distinct purposes, it is desirable, when possible, to interconnect them via a common infrastructure, enabling teams to work together to solve problems and to expand objectives from specific applications to a more varied traffic mix. For example, collaboration on new applications – such as those requiring massive data transfer – might lead to solutions for mass-market delivery of services, analogous to the success of Napster. Of course, Research and Experimental Networks should leverage the installed base of regional infrastructures whenever possible.

C.2.c. Relationship of Experimental to Production Networks

- Experimental Networking teams should explore multiple paths for technology transfer and knowledge transfer to Production Networking engineers, and are encouraged to make their innovative successes available.
- An Experimental Network Program should have an annual meeting to encourage interchange between the grantees, vendors and national Operational Research/Production Network providers.

C.2.d. Relationship of Experimental Networks to Industry

- Industry researchers should have opportunities for initial and continuous participation in the program (and process), as well as opportunities for funding, rather than solely have a supplier relationship with academic grantees.
- Companies need to be offered long-term opportunities; however, they should be able to opt-out in the short term if circumstances warrant.

- Joint academic/industry teams should be encouraged to adopt Open Source models as outcomes; otherwise, an Intellectual Property Rights plan for sharing results is advised, in which case a consortium agreement should a part of any proposal.
- Industry sees NSF's role as a supporter of very-high-risk, yet realistic, research and development infrastructure environments, than can possibly be created in a laboratory. By providing a venue for demonstrating the impact of new or proposed infrastructure features before they reach production, Experimental Networks benefit industry by reducing their costs and associated risks while garnering early customer validation of ideas.
- The ability to manage change is difficult in Production Networks, but possible in Experimental Network environments.
- Experimental Networks should re-examine the impedance match between applications and technologies, and provide the necessary "hooks" to redefine the network based on classes of applications, and provide non-production-ready services.
- Experimental Networks open up opportunities for industry to experiment with new, novel and undeveloped uses of network technologies.
- Experimental Networks should bring together researchers from one or more companies and academic organizations that would not otherwise work together; academic organizations with major distributed applications then become company-neutral "meet-me" places. Thus, Experimental Networking programs can encourage academic/industry/government partnerships that push the envelope on developing technologies and demonstrating network management structures within accelerated timeframes.

C.2.e. Summary

To summarize, the important benefits of Experimental Networks to industry were concluded to be several. Foremost is bringing together researchers from one or more companies and academic organizations that would not otherwise work together; academic organizations with major distributed applications then become company-neutral "meet-me" places. NSF's essential role is seen by industry to be a supporter of very-high-risk experiments staged in more realistic environments than could possibly be set up in a laboratory. Industry benefits from NSF's investment by reducing their costs and associated risks while garnering early customer validation of ideas without fear of market embarrassment.

Industry researchers should ideally have opportunities for initial and continuous participation in the program, as well as opportunities for funding, rather than solely have a supplier relationship with academic grantees. Companies need to be offered long-term opportunities; however, they should be able to opt-out in the short term if circumstances warrant.

Workshop participants concluded that private companies and venture capitalists are not generally able to fund Experimental Network infrastructure and research. If the Federal government does not fund Experimental Networks, and does not leverage the installed and emerging base of regional/campus infrastructures, Experimental Infostructure Networks will not be constructed.

C.3. Workshop Recommendations

The Workshop recommended a budget for a NSF Experimental Infostructure Networks Program: $50M-$75M spread over 5-7 projects with 5-year awards (administered as a $10M-$15M/year program), with the last two years of funds subject to mid-award reviews. Projects need 5-year funding to reasonably attract allied faculty, enable e-Science applications, and engage industry.

Multi-institution awards and company sub-awards are strongly recommended, as is support for academic professional staff as peers with industry researchers.

Another clear recommendation is that this Experimental Infostructure Networks Program should especially leverage the installed (and emerging) base of local and regional infrastructure. Several of the larger metro/regional networks (the *GigaPoPs/MetaPoPs*) are demonstrating advanced operational capabilities, meeting the requirements of cutting-edge applications and their users, thus proving that local- to regional-sized Experimental Networks are manageable today. Leveraging off this operations base, national and even international Experimental Network building should be considered if affordable. Phased funding should be based on measurable criteria specified in the proposals, keeping in mind that a substantial initial investment is needed to build the infrastructure. An important strategy is to maintain research openness by encouraging the deployment of technology coming from the researchers, rather than limiting technology to what is provided by the Production Networks to the researchers.

The Program should grant funds for delivered *end-to-end* connection of all resources involved in each experiment, *and* a demonstration of controls of network capabilities that facilitate applications. The term “controls of network capabilities” is somewhat ambiguous, and can refer to *strong* controls (allowing researchers to access every part of the network, including routers and switches, or to install diagnostic equipment) or *weak* controls (allowing researchers to access parts of the network, so they can manipulate the infrastructure in particular ways, but not all ways). The former definition explicitly differentiates Production and Experimental Networks, regardless of the applications being tested; the latter definition differentiates the degree of control and access to network parameters available to applications. Regardless of the interpretation, a Program should encourage the development of networks whose capabilities are well differentiated from, and not simply redefinitions of, existing Production Networks.

The recommended highest-level review criterion is *innovation* rather than *geographic scope*. NSF should take care in selecting reviewers to achieve balance between application and network researchers.

The Experimental Infostructure Networks Program should call an annual meeting to encourage interchange between the grantees, vendors and national Operational Research/Production Network providers. In addition, NSF should encourage strong interactions between its funded Research Network and Experimental Network recipients, as well as encourage ties to other Federal Agency and higher-education networks.

D. Breakout Session Discussions

Three parallel breakout sessions were held during the Workshop:

- *Technologies and Early Adopters of Experimental Networks*
- *Network Infrastructure Requirements*
- *Creating Industry/Government/Academic Partnerships*

This section summarizes the questions and discussions held in each Breakout. Session leaders summarized discussions in PowerPoint slides, which they presented to the entire group during plenary sessions. Subsequent group discussions led to the conclusions and recommendations presented in Section C.

D.1. "Technologies and Early Adopters of Experimental Networks" Breakout

D.1.a. Technologies and Early Adopters

The technologies needed to support Experimental Networks are dark fiber, wavelengths, switching, routing, and network control and management.

Possible early adopters of these technologies are metro/regional networks, state networks, national networks, international networks, carriers, RBOCs (Regional Bell Operating Companies), computer equipment vendors, network equipment vendors, and university labs and institutes.

Workshop industry participants suggested the following technology drivers for Experimental Networks, noting that *the key effort needed was not in the technologies themselves but in the integration of these technologies*:

- Next-generation packet formats and protocols
- Network processors
- Optical packet switching and label switching
- Lambda conversion/agility
- Wavelength services: multiple wavelengths for multiple purposes
- Wireless access
- Free space optics
- Integration of electronic packet switching and optical circuit and packet switching
- Common infrastructure sufficiently flexible and re-configurable to support multiple experiments
- Routing
- New paradigms in network control and management

D.1.b. Recommended Goals of the Experimental Networks Program

Given an anticipated budget of $10M-$15M/yr for 5 years, funding for 5-7 campus or regional-sized Experimental Networks is most likely, but national-scale connectivity between some of them may also be achievable through leveraging existing Research & Education Networks. The relevant criterion for funding should be innovation rather than geographic scope. *End-to-end* connectivity, from laboratory-to-laboratory, is key, as is crisp definition and bounding of the applications and the technical communities. NSF should take care in selecting reviewers to achieve balance between application and network researchers.

D.1.c. Experimental Networks Should Interconnect with Research Networks

Both Experimental Networks and Research Networks serve bounded communities of researchers who want to advance the state of the art; these research communities are independent of those communities that rely upon Production Networks in order to do their jobs, though production traffic may well be carried as part of an experiment. Industry sees benefit in collaborations between application and network (physical layer) researchers.

Experimental Networks provide infrastructure appropriate to support the proposed applications, with potential for technology transfer to Production Networks and production-grade products. If an application works perfectly well over a Production Network, it is probably not a good candidate for stress testing an Experimental Network. Therefore, it is essential to encourage environments that allow novel applications to scale and it is important to articulate the impact of the application requirements on the network's architecture and capabilities.

Research and Experimental Networks are distinct, yet there can be a fruitful collaboration between the two groups of researchers. While Research and Experimental Networks have distinct purposes, it is desirable to interconnect them via a common infrastructure, enabling teams to work together to solve problems and to expand objectives from specific applications to a more varied traffic mix. Experimental Networks can "glue" Research Networks together, where appropriate. Experimental Network requirements can migrate to Research Networks, and new technologies from Research Networks can be incorporated into Experimental Networks. To assure interconnectivity, NSF should consider funding fiber to and equipment in University-controlled spaces if public carrier facilities are not supportive of researcher access and control.

It is important that those who operate Experimental Networks be supported to assist networking researchers in the often "transitory" deployment of Research Networks for specific purposes.

D.1.d. Themes of Experimental Network Research and Measures of Success

There is a continuum between application and network research that must now address the mismatch among Moore's Law for computing, the increase in transmission speeds, the rate of growth of storage technologies, and the massive increase in the numbers and speeds of access technologies (e.g., broadband wireless and free-space optics).

A key objective is to support very demanding applications, such as scientific applications utilizing massive datasets, scalable real-time collaboration, interactive remote visualization or high-definition video streaming. Developing capabilities and capacity to do what cannot currently be done may lead to solutions for mass-market delivery of large-scale data transfer services. Expanded objectives should lead from specific scientific applications to a range of data types, creating greater use of the networks and a more varied traffic mix. For example, Experimental Networks and applications on campuses might support research and education peer-to-peer services (e.g., analogous to Napster).

Wherever possible, Experimental Networks should leverage the installed base of regional infrastructure and State and local funding sources. **Industry participants stressed that vendors and venture capitalists are not able to fund Experimental Networks.** If national and local government initiatives do not fund Experimental Network infrastructure and research, it will not be undertaken. An important strategy is to maintain research openness by encouraging the deployment of technology coming from the researchers, rather than from the backbone to the researchers. Industry also recommended phased funding based on measurable criteria specified in proposals, keeping in mind that a large initial investment is needed to build the infrastructure.

D.2. “Network Infrastructure Requirements” Breakout

D.2.a. Services, Not Just Bandwidth, Differentiate Applications

The “Grand Challenges” Workshop determined that *delivery of services*, not backbone aggregate bandwidth, was most important to advanced e-Science applications. While many services are built into today’s hardware, there is no way to access them given current network management issues, and lack of Application Programmer Interfaces (APIs) and usable feedback about network capabilities. Resources are not schedulable or negotiable; if multiple scientists want access, who gets priority and who loses? We have neither the policies in place nor the technologies needed to manage infostructure access and services.

Experimental Networks must be firmly grounded in improvements that directly impact the effectiveness of e-Science applications. Typically, e-Science application developers are asked to distill their network requirements down to bandwidth and latency constraints; however, application-level services like multicast audio/video distribution, shared data spaces, and remote control of instruments are common requirements of many e-Science applications, and they often fail for reasons that have no direct relationship to bandwidth or latency restrictions. Usually failure is due to a combination of factors, such as a lack of any level of service guarantee (not just fixed bandwidth, but failure notification and adaptation), lack of a robust monitoring infrastructure, lack of transparency of the network infrastructure, and the sheer complexity of attempts to implement such capabilities entirely at the application level.

These failures involve elements that fall outside the “application scope” as defined by the ISO/OSI model, and cannot be controlled by the application or the scientists running the application for lack of robust application-accessible monitoring, brokering, and control infrastructures. The interior of the network remains largely opaque to applications and network users. Remedies require personal interactions with people who manage the network, or attempts at application-level workarounds for problems far below the application layer. This is not a sustainable approach to providing infostructure, as it enables no more than a handful of simultaneous e-Science applications at best.

To manage the complexity of a more advanced, transparent, application-oriented infostructure, the network architecture must be organized around services that are of common interest to more than one e-Science application scenario. One could implement middleware to provide these services, but that software infrastructure would still be hamstrung by the same lack of transparency and reliability that hampers application-layer implementations. Ultimately, the network architecture must expand in the following ways:

- It must offer more transparency for both control and monitoring of shared resources so that the middleware/application can manipulate network behavior to deliver these services effectively.

- It must provide higher-level services, or elements thereof, as a feature of the hardware/firmware infrastructure.

Either approach will raise alarm in the industry because it will increase the complexity of the infrastructure. Such increases in complexity must be justified by rigorous case studies. This is a chicken-and-egg problem, in that usage scenarios cannot be tried without a realistic network testbed and vice versa. Experimental Networks can break this cycle by providing such a resource and helping qualify the cost/benefits of such deep and potentially disruptive infrastructure modifications.

D.2.b. Addressing Network Problems as a Systems Issue

Should an Experimental Infostructure Networks Program be largely computer-centric? Applications, whether they are computationally-based, sensor-based, or collaboration-based, all have computers at the edges. Or, should the focus be on network technology that enables applications? How does the fact that this is an infrastructure initiative, not a pure research initiative, affect the balance between computer and network R&D efforts?

The ISO/OSI and TCP/IP networking models manage network complexity by dividing capabilities into layers. Current realities of optimizing network performance have pushed many low-level deficiencies up into the application layer; thereby breaking this paradigm of managing complexity. We now recognize that current deficiencies in service delivery to e-Science applications are a systems-level problem rather than residing in any particular layer, so these problems **must** be addressed by considering the architectural decisions that cut across all layers. Whereas much current research funding addresses “layers” of the network, Experimental Network funding is more of a systems problem in which all layers need to be addressed and funded together.

Participants cautioned to avoid *stove-pipe* vertical approaches that support one specific problem, since the resulting techniques often do not transfer well to other disciplines. As a generalization, point-to-point connections rarely solve general infostructure problems.

Experimental Networks must be addressed as a systems problem, but the social/policy/economic problems need to be considered as well. Networks interconnect people, organizations, and nations with radically different needs and value systems. Any Experimental Network that considers issues of Quality of Service (QoS) or other mechanisms to arbitrate resource-sharing policies must take social implications into account. QoS alone poses as many social quandaries as it does technical challenges. While sound on technical grounds, QoS designs that were developed in isolated research environments may be rejected as socio-politically impractical. The social implications of the resource arbitration methods can **only** be explored adequately in-vivo using Experimental Networks. Otherwise, QoS designs will not be subjected to this dimension of evaluation until they are already solidified into production designs. Social policy issues include the following questions:

- Who resolves resource conflicts on Experimental Networks? How do we decide importance of applications? For instance, medical data may save one life at a time; earthquake sensor data may save thousands of lives at a time, but less frequently.
- How rapidly can one establish an urgent connection, or must it be up and ready to use all the time?

- Regarding policy and control of the network, how does one express, implement and enforce policy?
- How do inter-domain and international networks, even experimental ones, complicate policy resolution?
- What is an appropriate vocabulary for expressing network sharing and access policies. Currently, there is no common language for expressing complex parameters in a manner that supports automated service arbitration.
- Is arbitration to be done solely on the basis of who has the best ability to pay?

Clearly these questions must be addressed hand-in-hand with the technological implementation. Such questions can only be addressed with real implementations, and any such implementations must demonstrate an ability to properly evaluate these issues.

D.2.c. Applications Needing Experimental Networks

Participants listed the following application categories and example applications that need Experimental Networks:

Science	• Remote sensing data combined with supercomputer simulations • Interactive visualization of large-scale earth sciences data • Large-scale shared astronomical instruments combined with the National Virtual Observatory, including real-time remote control of instruments and data mining and data storage
Medical/ Health-Care	• Health-monitoring of elderly and children in home-care situations • Guaranteed level of service for critical interactive applications (e.g., surgery) • Blood flow with human sensor data • Mutli-site/collaborative biomedical volume visualization
Arts & Entertainment	• Museums and cultural heritage virtual reality • Shoah Foundation-type video, image, sound and text databases • Video to/from cell phones; rich communication to devices • Immersive 3D gaming environment in the home
National Security, Governance, Policy and Emergency Response	• Wildfire response incorporating multiple agency data input • Sensor data for energy control, bio-terrorism, water quality, etc. • Enhanced airport and airplane scanning
Computer Science and InfoTech (Grid and Infrastructure building efforts)	• Network-distributed storage and caching • Standards for integrating grid computing • Peer-to-peer participation on networks with security to enable peer-to-peer management and end-user management • Scavenger service that finds unused resources • Encapsulated private networks that give the illusion of one's own network running over a shared network

iGrid 2002 Applications. For iGrid 2002 (September 23-26, 2002, in Amsterdam) <www.igrid2002.org>, 16 countries proposed 28 demonstrations: Australia, Canada, CERN/Switzerland, Finland, France, Germany, Greece, Italy, Japan, Netherlands, Singapore, Spain, Sweden, Taiwan, United Kingdom, and the United States. These applications include: art, bioinformatics, chemistry, cosmology, cultural heritage, education, high-definition media streaming, manufacturing, medicine, neuroscience, physics and tele-science. Major emphasis is on grid technologies: grid middleware, data management grids, data replication grids, visualization grids, data/visualization grids, computational grids, access grids, and grid portals.

D.2.d. Addressing Elements in the Infrastructure by Layers

Each application has a mix of measurable services involving low jitter, low latency, reliability, high throughput (fault-resilient reliable vs. unreliable transport) and low reconfiguration/ connection time. It is unclear whether Service Level Agreements (SLAs) address these needs. For that matter, low-level network characteristics (such as jitter and latency) may be ill posed, as it is unclear if IP should be the basis for all of these services. From the standpoint of base-level application requirements, how can we provide for:

- Remote control of resources
- Reliable transport
- Fault-resilient unreliable methods (graceful failure)
- Software standards
- Middleware capabilities
- Network management
- Security (management and end user)
- And, effective end-user feedback on the network (peer participation/grid)

It is critically important that we have robust instrumentation and monitoring information structures with both static and dynamic capabilities in order to effectively evaluate solutions to the above-posed application needs. In addition to evaluating such technical solutions, we must also evaluate the social/political/financial interactions and impacts of resolving disputes over resources.

D.2.e. Political/Financial/Social Issues

How will resource conflicts be resolved in the following areas?

- Production vs. Experimental Networking
- Human resources
- Computing resources
- Social values behind connectivity
- Vertical integration of effort; improving communication among layer experts
- Adoption of technical standards
- Intellectual property
- Privacy and security

D.2.f. Technical Issues

Experimental Networks need to facilitate research on ways to provide:

- Access to existing capabilities in routers to build a common security model with scalable methodologies

- Access to lambdas
- Network introspection and feedback
- Fault-resilient unreliable methods allowing graceful failure
- Exploitation of the symbiosis of wired and wireless networks
- Computer/network interfaces hardware
- End-to-end connections for both the *last mile* and the *last yard* of connectivity

D.2.g. Infostructure Management Issues

Management of Experimental Networks involves system integration efforts with:

- Cooperation among Experimental Networking grantees
- Resource allocation
- Conflict resolution
- Managing multiple domain interfaces (WAN to campus; campus to lab; testbed to testbed)

D.3. "Creating Industry/Government/Academic Partnerships" Breakout

D.3.a. Motivating Partnerships

Workshop participants want to see truly scalable distributed computing; that is, more than single grid, linked networks or Linux clusters. Networking researchers need the opportunity to address system-level problems – bringing together companies, universities and government laboratories to work together on Experimental Networks.

Workshop industry representatives noted that to motivate partnerships, Experimental Networks must be tangible; that is, engineered and built, with useful lessons that can be learned. For example, for companies to participate in field trials and in the development of novel technologies, they want more than simply learning about research results after the fact, which is common among university-defined industrial affiliation programs.

When companies are asked to participate in a typical NSF Center proposal, it is usually just an opportunity to provide cash, get reports and recruit graduating students. Companies would instead prefer a partnership to develop networks that neither party could do on its own, with rewards for both parties. Corporate research labs have a hard time justifying long-term investments now. If multi-institution partnerships and awards (including company sub-awards) are developed with a *specific set of aims* that can be evaluated concordant with corporate goals and with tangible evidence of markets and windows of opportunity, broad corporate support is also possible.

Networks and applications *working together* will also help break down some of the traditional barriers between academia and companies, specifically by providing network measurement, monitoring, and routing information that would typically be carrier-proprietary, and providing realistic deployment, ahead of what leading commercial customers would purchase as services. It is recommended that NSF fund academic researchers with major distributed applications whose organizations can serve as company-neutral "meet-me" places.

D.3.b. Benefits of Experimental Networks

Industry noted the clear benefits of the availability of Federally-funded Experimental Networks:

- Experimental Networks have real data and management needs
- Testing is possible at “intermediate” scale: larger than lab-scale but before production use
- A focus is possible on longer-term return on investment than companies can justify
- Access to performance data that might otherwise be proprietary is normal
- Change is possible without the stigma of market “embarrassment”
- In short, ***experiments*** are possible!

D.3.c. Award Length

The group strongly encourages 5-year awards because it takes a year or more to build a useful Experimental Network, and a guarantee of multi-year persistence is needed to attract faculty, enable e-Science applications to progress to completion, and secure industry involvement.

D.3.d. Full-Time Staff Support

Companies often prefer dealing with their workforce peers in academia, who are typically academic professionals (full-time staff researchers), rather than students or faculty. Thus, adequate funding for full-time academic staff is deemed essential to attracting and retaining industry participation in any Experimental Network effort.

D.3.e. Experimental/Research Networks Scope

Experimental Network initiatives should fund the development and testing of systems-level integration for leading-edge hardware to investigate pre-market strategies, providing:

- Management (“Network with the Right Knobs”)
- System software
- Applications development and testing
- Open Source middleware

Research Network initiatives may best focus on access to dark fiber, wavelengths, lightpaths and technology that are not ready for applications. Research Networks are seen as ideal for investigating strategies pertaining to open and proprietary (potentially patentable) hardware and middleware.

D.3.f. Getting Applications to Use Experimental Networks

Across the NSF Directorates, future funding must be made available to pay for risky and novel applications on already funded Experimental Networks. One avenue may be to work closely with the DTF/ETF networks, with Experimental Network initiatives providing metro/regional on-ramps to experimental infrastructures that have applications already planned and/or running.

Applications can be encouraged to focus on advanced prototyping of Grid activities, like high-speed visualization of huge datasets, creating a *grid of grids* by providing, for example, scheduling and management of Layer2/Layer3 Virtual Local-Area Networks (VLANs). Peer review of applications, and support for technology professionals and travel (in support of iGrid and I-WAY large-scale demonstrations) can help assure a yearly influx of novel applications.

D.3.g. Industry Incentives

Industry offered advice on ways to help ensure private sector participation:

- Companies need to clearly see the business opportunities
- Companies want a founding seat at the table, not just a supplier relationship with grantees
- Companies want a long-term opportunity with a short-term opt-out available

Industry sees NSF's role as a supporter of very-high-risk research and development, providing rewards for successful Open Source models and Intellectual Property Rights sharing plans. A consortium agreement should be in place, where needed, as part of proposals.

While the supercomputing market was pleased to have lots of machines with serial #1, the networking market is the opposite. Network managers want to wait three years before touching new technology. A major benefit of Experimental Networks would be to bring together one or more companies as well as academics, who otherwise would not work together on experiments, to experiment with the newest purchasable ("beta") equipment; thus, NSF-funded academic organizations become intellectual "meet-me" places.

It was noted that business models for profit-making data-only services remain elusive at this time. Most telecommunications revenue is still in voice; carrying data is considered a loss leader. Yet, many users of data services think they should be free. The private sector cannot afford to fund research in data networks since there is no acceptable revenue model. This makes industry more interested in teaming with academia.

NSF can help build infostructure into more realistic environments than individual companies and/or researchers could possibly set up in corporate or academic laboratories with spools of fiber. The ability to manage change is difficult in Production Networks, but possible in Experimental Network environments. Companies want NSF to reduce their cost and risk by funding Experimental Networking, providing testbeds for early customer validation of ideas.

E. Appendix: Final Workshop Agenda

Sponsored by the National Science Foundation (NSF)
Computer and Information Science and Engineering Directorate (CISE),
Advanced Networking Infrastructure and Research Division (ANIR)

Hosted by California Institute for Telecommunications and Information Technology [Cal-(IT)^2],
University of California, Irvine (UCI)
Irvine, California

May 20-21, 2002

General Chairs
Larry Smarr and William Parker
California Institute for Telecommunications and Information Technology [Cal-(IT)^2]
University of California, San Diego (UCSD)
University of California, Irvine (UCI)

Program Committee Chairs
Tom DeFanti and Maxine Brown
Electronic Visualization Laboratory
University of Illinois at Chicago

Monday May 20, 2002

8:30 am **Shuttle** available outside Hyatt Regency Irvine to transport Workshop participants to the Arnold & Mabel Beckman Center

9:00 **Continental breakfast;** Informal introductions

10:00-10:15 **Welcome and Charge to Workshop Attendees –** Aubrey Bush, NSF ANIR Division Director

- Workshop goals
- NSF’s taxonomy for R&E research, experimental and production networks

10:15-12:00 **15-minute presentations from carriers and equipment manufacturers**

- Steve Wallach, Chiaro
- Alan Benner, IBM
- Larry Dennison, Avici Systems
- Bob Aiken, Cisco
- John Jamison, Juniper
- Hal Edwards, Nortel

11:45-1:00 **Lunch and informal discussions at Beckman Center**

1:00-2:30 **More presentations from carriers and equipment manufacturers**

- Rohit Sharma, ONI
- Ed Bortolini, Network Photonics
- Tom Myers, Teraburst

- Rich Wolff, Telcordia

2:30-3:00 **Coffee Break**

3:00-4:45 **Breakout Discussions**

Three rooms will be provided. Scribes will take notes and record major opinion points using PowerPoint slides created in real-time on projectors/large plasma panels.
Goal: “Grand Challenges in e-Science” Recommendations to NSF

Breakout #1: Technologies and Early Adopters

- What are the technologies for Experimental Networks?
  a. Dark fiber
  b. Wavelengths
  c. Switching
  d. Routing
- What does industry see as the drivers for Experimental Networks?
- Who are the likely early adopters?
  a. Metro/regional networks
  b. State Networks
  c. National Networks
  d. International Networks
  e. Carriers
  f. RBOCs
  g. Computer Equipment Vendors
  h. Network Equipment Vendors
  i. University Labs/Institutes

Breakout #2: Network Infrastructure Requirements

- What are the network infrastructure elements needed to:
  a. Address current and future applications and data storage needs?
  b. Work for applications that need more than 100Mb, 1Gb and/or 10Gb end-to-end connectivity?
  c. Incrementally build 1Gb to 1Tb national networks?
  d. Lead to progress in wavelength provisioning middleware to applications users?
  e. Provide reasonable security to high-bandwidth applications?
  f. Allow experimental networks to interoperate with current R&E and commercial infrastructures (i.e., what routing and policy issues need to be addressed in addition to security issues)?

Breakout #3: CISE’s Role in Creating Industry/Government/Academic Partnerships

- Given that e-Science needs specific services not available in today’s production networks, are experimental networks the only answer, or are there other suggested modes of networking infrastructure support?
- What percentage of the funding should industry, government and academia each contribute and why?

- What level of Federal funding should this program optimally have?
- What percentage of funding should address security needs in the coming decade?

4:45-6:00	**Breakout session reports (PPTs)**
6:00-6:15	**Discussion**
6:15	**Shuttle** returns participants to the Hyatt Regency Irvine hotel
6:30-8:30	**Dinner at the Hyatt Regency Irvine (6:30 Cash bar; 7:30 Dinner)**

Tuesday May 21, 2002

8:30 am	**Shuttle** available outside Hyatt Regency Irvine to transport Workshop participants to the Arnold & Mabel Beckman Center
9:00am	**Continental breakfast**
9:15-10:45	**Writing assignments** Small teams form to create recommendations for future NSF solicitation (advocacy groups evolving from each breakout, assuming 2-3 positions from each breakout emerge).
10:45-11:00	**Coffee Break**
11:00-12:30	**10-minute summaries (PPTs) from each team**
12:30-1:15	**Lunch and informal discussions at Beckman Center**
1:15-3:00	**Continued refinement of recommendations.** People catching planes can leave anytime.
3:00-3:30	**Coffee Break**
3:30-5:00	**Final design of structure of recommendation document;** text and diagrams refined as necessary
5:00	**Shuttle** returns participants to the Hyatt Regency Irvine to pick up luggage and take hotel shuttle to John Wayne airport

F. Appendix: Breakout Participants

The following individuals participated in the Breakout Sessions.

Breakout #1: Technologies and Early Adopters

Scribe: George Clapp
Bob Aiken
Franz Birkner
Ed Bortolini
Steve Corbató
Hal Edwards
Greg Hidley
Wendy Huntoon
Ron Johnson
Anke Kamrath
Tom Lehman
Joe Mambretti
Tom Myers
Jerry Sobieski
Jeff Verrant
Steve Wallach
Observers: Aubrey Bush

Breakout #2: Network Infrastructure Requirements

Scribe: Maxine Brown
Phil Papadopoulos
Paul Obsitnik
Tom West
John Shalf
Linda Winkler
Larry Dennison
Ruzena Bajcsy
Observers: Alan Blatecky, Aubrey Bush

Breakout #3: CISE's Role in Creating Industry/Government/Academic Partnerships

Scribe: Tom DeFanti
Alan Benner
Ron Hutchins
John Jamison
Larry Smarr
Bill Wing
Rich Wolff
Observers: Tom Greene, Rich Hirsh

G. Appendix: Workshop Presentations

Aubrey Bush opened the Workshop with a Welcome address, after which there were eight 15-minute presentations from industry representatives from equipment vendors, computer experts and start-up companies, to set the tone for subsequent Workshop discussions. All PowerPoint presentations are archived on the web: <http://www.calit2.net/events/2002/nsf/index.html>.