Networking Research Challenges Workshop Report

September 28-30, 2008
Seattle, Washington

Sponsored By

Department of Energy: Office of Science
National Science Foundation: Directorate for Computer and Information Science and Engineering

Coordinated By
Large Scale Networking Coordinating Group
Networking and Information Technology Research and Development Program (NITRD)

Workshop Breakout Session Chairs

This report was prepared by the workshop breakout session chairs based on the participation, discussion, and contributions of the workshop participants (See Appendix 1).

Next Generation Heterogeneous Networking:
Bill Wing
Oak Ridge National Laboratory

Network Security:
Joe St Sauver
University of Oregon and Internet2

Federated Optical Networking:
Joe Mambretti
Northwestern University

Network Science and Engineering:
Karen Bergman
Columbia University

The opinions, findings, and recommendations provided in this report do not necessarily reflect the views of the Department of Energy, the National Science Foundation, or the NITRD Program.

Table of Contents

Table of Contents (Continued)

Summary: Networking Research Challenges Workshop

Over the last four decades, U.S. government research and development (R&D) in advanced networking has fueled a technological, economic, and social transformation, not only in the US, but world-wide. Today, networks have become fundamental infrastructure for government services, operations, national defense and homeland security, commerce, communication, research, education, and leisure-time activities. However, the current generation of network technologies has limitations and vulnerabilities that can threaten our national security and economic competitiveness.. Basic capabilities have not kept pace with emerging critical requirements. The Networking Research Challenges (NRC) Workshop convened between September 28th and 30th in Seattle Washington to produce recommendations for the narrowing of this gap. Workshop sponsors included the Department of Energy and the National Science Foundation . It was coordinated by the Large Scale Networking Coordinating Group (LSN CG) of the Networking and Information Technology Research and Development (NITRD) Program.

Visionary Networking Applications

Continuing investments in networking research and development will foster new applications and capabilities in a vast range of commercial, societal, healthcare, collaboration, and other application areas. Some of the new capabilities will include:

- Transport of unprecedented amounts of data, particularly for science community collaborations
- Powerful new applications linking distributed people, data, resources, and sensornets with secure reliable, private, real time communications
- Flexible, rapid provisioning and restoration of resources to optimize use of networking resources by bringing required resources to bear on an application and by releasing those resources when they are no longer needed
- Lower cost communications through optical technologies
- Green communications based on optical networking which has much lower power and cooling costs

In January 2008, the Director of the Office of Science and Technology Policy (OSTP) recognized the need for a long-range plan for advanced networking research and. called for the Federal agencies to develop a *Federal Plan for Advanced Networking Research and Development*). The resulting plan was published in September, 2008 (see: http://nitrd.gov/Pubs/ITFAN-FINAL.pdf). The plan is centered on a vision for advanced networking to provide secure network services anytime, anywhere; make secure global federated networks possible; manage network complexity and heterogeneity; and foster innovation through coordination between the Federal and private sector research communities.

Traditionally, in the U.S., the federal sector has funded the basic research needed to discover and develop new capabilities and breakthroughs, supporting longer-term research by university, laboratory, Federal lab, and commercial researchers. The commercial sector has generally focused on the shorter term development of new

commercial products and services based on the discoveries and developments of the longer-term research community.

NRC Workshop Objectives

Workshop participants assembled to:

- Provide broad-based, in-depth discussion and review of networking research needs from all national and international networking communities
- Develop guidance and recommendations on networking research priorities
- Identify coordination and cooperation among networking research communities to accelerate progress on new networking capabilities
- Provide guidance on roadmaps and timing for networking research

The workshop was held in conjunction with the international Global Lambda Integrated Facility (GLIF) Conference. The goal of the workshop was to generate a report that summarizes the current state of global internetworking capabilities and highlights fundamental technical network research and engineering challenges. Workshop recommendations therefore provide guidance to industry, academia, and Federal agencies on formulating and prioritizing networking research and engineering roadmaps.

Workshop presentations on critical networking issues included key topics such as security, network convergence, integrated optical networking, and dynamic secure mobile wireless technology. Attendees (a list of the workshop participants appears in Appendix 1) participated in one or more of the following four breakout groups whose discussions and recommendations are described in this report:

- Technical challenges for future network security
- Federated optical networking
- Next generation heterogeneous and embedded networking
- Network science and engineering challenges

Joint Recommendations

The joint recommendations of the four breakout groups include the need for:

- Revisiting fundamental design principles through interdisciplinary basic research in modeling, architectures, and protocols for developing future networks that can address the accelerating demands for end to end functionality and provide a robust and scalable socio-economic infrastructure
- Improved and continuing coordination among all Large Scale Networking (LSN) constituencies to identify research priorities, promote cooperation, increase visibility of research results, and accelerate commercialization of new capabilities
- New standardization processes, programs, and organizations to enable advances in the design, development, and adoption of new innovative capabilities
- Lowering barriers to adoption of new capabilities (e.g., security and cross-domain optical networking) including methods for evaluating economic considerations, such as new forms of risk reward frameworks

- Designing and implementing multiple large scale national and international test beds with differing characteristics to enable multiple simultaneous cooperative and interdisciplinary experiments by different communities of researchers
- End-to-end performance measurement across network domains

Recommendations specific to the topical breakout groups included:

Security Recommendations

Research is needed to improve:

- Security in-depth: the identification of assets for protection tailored to their value, intrusion detection, automated response, automated recovery, automated early alerting on attacks, and related capabilities
- Identity management, trust, privacy, and policies for sharing trust
- Security for optical networks: authentication before resource allocation, sharing performance and incidence data across domains, distributed intrusion detection systems, end-to-end security for control planes, security when optical networks bypass firewalls, and leveraging the inherent security capabilities of optical and photonic technologies at the physical level
- Other security topics including metrics, human factors, deployment of two-factor authentication at scale, eliminating spam and unsecure practices, site-configurable NIC cards, sensornets, wireless networks, and physical security such as Electro-Magnetic Pulse (EMP)

Federated Optical Networking Recommendations

Research is needed to improve:

- Architectures, protocols, and prototypes for a large scale optical networking systems of systems, including innovative control and management, dynamic optical switching, and hybrid packet/circuit switched optical networking
- Interdisciplinary research on optical networking technology including interfaces, switching elements, service granularity, and specialty appliances
- New technologies and capabilities including:
  - Generalized Multi-Protocol Label Switching and related control plane architecture and protocols
  - Optical burst switching
  - Network virtualization at the optical layer
  - Optical wireless technology

  Time Division Multiplexing and related technologies
  - Lowering the cost of high capacity transport using optical channels
  - Dynamic provisioning at all levels
  - Scalability for many techniques and technologies
  - Research program structure and organization
  - Policy development
  - Facility design and implementations
  - Economic evaluations
  - Technology transfer
  - Applications-oriented capabilities
  - Standards organizations

- Terabit Local Area Networks

Heterogeneous Networking Recommendations

- Open the educational TV spectral allocation in the 2.4 - 2.6 Ghz band for R&D use
- New high voltage transmission line rights of way, should be made available for national-scale fiber infrastructure
- The National Coordination Office (working with Federal networking agencies) should coordinate a series of network workshops on:
  - Technology-bridging protocols
  - Self-adaptive protocols
  - Control and management in the context of partially hidden link-state information
  - End-to-end security in heterogeneous networks
  - Advances in the physics and engineering of new network devices

Network Science and Engineering (NetSE) Recommendations

- Develop a dual approach to creating new network architectures: (1) define and fund a research program to explore theory in network virtualization, modularity, and composability; (2) fund applied research and development of these concepts to create fully functional – but experimental – environments that can act as a proving ground for these new service models
- Create new and aware protocols at different levels that can address real time flows, mobility, dynamically changing environments, and different channel characteristics
- Develop methodologies for studying network-centric social applications and environments

Section 1: Introduction and Overview

Today, public- and private-sector enterprises throughout the United States and around the world depend upon a pervasive infrastructure made up of computing and storage systems, software, and devices interconnected across a vast web of wired and wireless networks. High-speed network connectivity links sensors, data, devices, and applications to users on the move, enabling near-instantaneous communication and global transmission, storage, and retrieval of enormous amounts of data (e.g., text, images, sound, multimedia, signals). In the U.S., networking capabilities have become indispensable, accelerating industrial and commercial innovation, advancing science and engineering, and supporting vital government missions and services.

Emerging classes of advanced networking applications play increasingly critical roles in national defense and homeland security, as well as in aviation and transportation; management of key physical infrastructures such as power and water supply; medicine and health care; emergency preparedness and response; environmental monitoring; and large-scale, data-intensive, and domain-specific scientific research.

Historically, the Federal government's long-term investments in networking research and development (R&D) have provided the core technical foundations for networking. Federal research led the way to the Internet, wireless mobile and optical networking, and a broad range of networking applications such as search engines and grid computing that continue to transform our society and economy. The Federal government's productive wellspring of networking R&D advances nourishes the development of the multi-billion-dollar IT industry and a vast array of new technological capabilities for both individuals and organizations.

In its August 2007 report entitled *Leadership Under Challenge: Information Technology R&D in a Competitive World*, the President's Council of Advisors on Science and Technology (PCAST) stated that "U.S. leadership in advanced networking is a strategic national priority" essential to sustaining the nation's military, scientific, economic, and technological preeminence. Noting the inherent limitations of the decades-old Internet architecture, the PCAST recommended the development of an R&D agenda to upgrade the Internet, strengthen critical-infrastructure networks, and meet federal needs for advanced networking capabilities such as secure wireless mobile networks.

The PCAST endorsed the January 2007 call by the President's Science Advisor for Federal R&D agencies to develop a long-range plan for advanced networking research to meet Federal agency mission requirements and commercial-sector networking needs as well as national and homeland security requirements. The resulting *Federal Plan for Advanced Networking Research and Development,* published in September 2008, presented a vision for advanced networking with four main goals: 1) Provide secure network services anytime, anywhere; 2) Make secure global federated networks possible; 3) Manage network complexity and heterogeneity; and 4) Foster innovation among the Federal, research, commercial, and other sectors through development of advanced network systems and technologies.

1.1 Networking Research Challenges Workshop

In this context of identified critical needs – including the rapidly changing environment of Federal networking mission requirements, national security needs,

commercial networking needs, and science and other application needs – the National Science Foundation (Directorate for Computing and Information Science and Engineering (CISE) and the Department of Energy, Office of Science, with the endorsement of the Large Scale Networking Coordinating Group (LSN CG) of the Networking and Information Technology Research and Development Program (NITRD), co-sponsored The Networking Research Challenges Workshop with broad participation from the networking research community to:

- Provide broad-based discussion and review of networking research needs from all networking communities including network researchers, developers, and users across the commercial, Federal, national laboratory and international communities
- Develop guidance and recommendations on networking research priorities
- Identify coordination and cooperation among networking research communities to accelerate progress on new networking capabilities
- Provide guidance on roadmaps and timing for networking research

The workshop sought inputs from domain-specific scientists, advanced networking researchers, program managers, the commercial sector and developers with interests in advanced networking capabilities and research. The participants were asked to focus on a vision for the network and networking needs for the 2015 time frame. The participants were also asked to focus on four goals:

- Technical challenges for future network security
- Federated optical networking
- Next-generation heterogeneous and embedded networking
- Network science and engineering challenges

Each of these four areas was the focus for a workshop breakout group in which participants identified key issues, discussed the needs and barriers associated with these issues, and provided recommendations on research needed to address the issues. For each of the four goal areas, the key issues and recommendations are summarized below.

1.2 The Strategic Vision for Networking

Workshop participants, under the leadership of an organizing committee consisting of federal agency, academic, laboratory, and international representatives, formulated the following strategic vision for networking.

Unprecedented data transport

Increasingly, scientific research requires the gathering, analysis, visualization, and transport of extraordinary volumes of data. Much scientific discovery is dependent on models and simulations based on extremely large data sets. Advanced optical networking is the only option available that can provide the ultra-high-capacity transport required for future science. The capabilities of optical networking for high-volume transport far exceed any other technology. This capability remains especially important for interdisciplinary scientific research requiring the integration and examination of multiple, large sets of data from many disciplines.

Powerful new applications

Advances in networking generate innovative, powerful capabilities that enable new services and applications for science. These capabilities, based on optical networking, provide for continuous, high quality, reliable streams of information, in part by transporting information directly on waves of light. This enables scientists to view and interact with very large sets of data in real time from any location as long as they have access to appropriate optical communication services.

Flexible, rapid provisioning and restoration

Previous communication architectures were based on static inflexible components that limited the expansion and enhancement of services. Advanced optical networking, by contrast, enables enhanced flexibility, fast deployment of new and expanded platforms, greater reliability, and extremely rapid service restoration. A mesh architecture in advanced optical networks eliminates single points of failure within the communications infrastructure.

Cost-effective communications

Advanced optical networks have proven the most cost effective platforms for transporting large amount of data across an organization, a metro area, a nation or the world. New optical core technologies are far more economical than traditional methods. In the future, devices, such as computers and scientific instruments will incorporate individual components using optically based techniques for the generation and transmission of data.

Green communications

Advanced optical networking has the potential to reduce power, cooling, and environmental requirements for communications significantly. Light-based technology requires substantially less electricity than electronic technology and the optical components generate much less heat than standard electron-based equipment per volume of information communicated.

New infrastructure

Enabling these capabilities will require a new, dynamic networking infrastructure using wavelength-routing optical switches with switching times on the order of a few nanoseconds. The infrastructure will span sub-wavelength circuits, wavelengths, and entire wavebands and fibers. Higher-layer nodes will provide interoperability among heterogeneous services (IP, MPLS, SONET, MSPPs, etc.). Distributed users of the network will have the ability to configure resources (networking, compute, storage, security, management, etc.) for the creation of dynamic virtual private networks. Connectivity to the infrastructure will be supported across network domains and heterogeneous technologies. Recognizing the growing importance of commercial mobile radio technologies and applications, we envision the integration of existing wired, wireless, and IP-based infrastructures into a Next Generation Network fabric to support secure, end-to-end, heterogeneous, multimedia networking.

Section 2 Network Security[1]

The core of the Internet is based on a simple architecture that provides universal connectivity, universal communications, and allows the creation of new applications and link technologies. The core architecture is built on many types of routers, domain name services (DNS), firewalls, Internet service providers (ISPs), network information centers (NICs) and other technologies and organizations. While each of the relevant technologies contains vulnerabilities, they also have capabilities to mitigate attacks, albeit at an economic and performance cost. The vulnerabilities include violation of confidentiality such as router password compromise; violation of integrity by erroneously modifying router tables or poisoning DNS caches; and impeding availability by flooding routers or spamming ISPs that cause denial of information. A wide range of attacks such as Man In The Middle (MITM), spoofing, spam, phishing (see: http://en.wikipedia.org/wiki/Phishing) , intrusion attacks, eavesdropping on network traffic, botnets, identity theft, insider compromise, malware and Trojan horses, worms, and viruses exploit these vulnerabilities. Thus, current Internet security challenges include tracing attackers to their source.

Network monitoring is the key to detection, analysis, and response to attacks. But the constant increases in network traffic rates and volumes, heterogeneity across network domains, and differing policies across those domains complicate the task of monitoring. Performance monitoring on larger amounts of data at higher speeds across network domains will require new capabilities.

Identity management is a critical component of security. Malicious actors can forge source addresses for UDP traffic making trace-back generally impossible. The networking community must develop new protocols and services to provide a binding between a packet's source address and the identity of the sender (a trusted third party holding the link of a user to a source address could protect the anonymity of the user).

A new Internet protocol could provide a new addressing scheme for networks and hosts that enables self-certifying addresses and provides anti-spoofing, secure routing, prevention of distributed denial of service (DDoS) attacks, and other capabilities.

Botnets are a persistent problem. They enable 90% of all spam, all denial of service attacks (DoS), and contribute significantly to phishing and pharming (see: http://en.wikipedia.org/wiki/Pharming) attacks, key logging and identity theft, and anonymized terrorist and criminal communication. Botnets are controlled by identifying and attacking their command and control channels which generally requires human resources. . Furthermore, false positives lead to complaints. A rich inter-site analysis for mitigating cooperative attacks might provide a clearinghouse architecture whereby cooperating sites would receive early warning of attacks on resources. The Bro Intrusion Detection System, (see http://www.bro-ids.org) and Snort (see www.snort.org) provide examples of some of the capabilities that could be deployed.

---

[1] **Source materials for the development of this section on Network Security may be found at:**
http://www.uoregon.edu/~joe/nitrd/
and
http://www.uoregon.edu/~joe/nitrd/november20th.pdf
http://www.uoregon.edu/~joe/nitrd/december10th.pdf

2.1 Security Breakout Session

Network security often is focused on near-term operational needs and responses to intrusions. This focus makes it difficult for network managers, planners, and researchers to think strategically over the longer term. Additionally, developing and deploying new security technologies (such as DNSSEC; see: http://www.dnssec.net/) can easily take a decade or more. If the networking community identified a new security technology today, it might easily take until 2018 or 2019 before deployment. The security breakout sessions subsequently discussed both near-term operational needs for network security research and the longer term research perspectives for network security. The discussion focused only on unclassified aspects of network security.

2.2 Vision of Network Security in 2015

Six years from now, network security will operate in a significantly more complex, faster, and heterogeneous environment than today's environment. Network security will be deployed at not only Layer 1 (physical layer), Layer 2 (data link layer) and Layer 3 (network layer) but also at upper layers of the protocol stack and most traffic likely will be encrypted. Other factors impacting future network security include:

- A huge legacy/production base where new security technology introduction and diffusion may be constrained by equipment replacement lifetimes resulting in a slow rollout of new technology
- Costs and benefits are often asymmetric and constitute barriers to adoption of new security technology (my expenditure on network security may help your security, but paradoxically may not necessarily do much for my own security)
- The commercial sector will build the equipment that we need, but commercial differentiation favors new features and increased complexity over simplicity, performance, and economy
- The supply of trained network engineers and security people remains insufficient.
- Deployed complexity (e.g., firewalls) currently exceeds the administrative ability of amateurs to operate
- Compliance-related activities (e.g., paperwork) may drain additional resources away from fighting the cyber "wars"
- One size cannot fit all; flexibility is important
- We will continue to overlook obvious solutions

2.3 Trust Models

Traditional trust models include hierarchical trust models rooted at a trusted origin (e.g., PKI and other certificate-based models), and less structured "web-of-trust" models, as used by PGP/Gnu Privacy Guard, whose trustworthiness of a credential is a function of attestation by multiple trusted peers. Federated trust models, such as those based on Shibboleth and InCommon or Kerberos, also are being actively developed and are experiencing widespread deployment in some communities.

Trust is sometimes tightly coupled to notions of identity and reputation, although a trusted party's ultimate "real life" identity may remain anonymous. Many practical

problems remain unsolved, e.g., revocation lists are still problematic, and the ad hoc nature of PGP/Gnu Privacy Guard can deter adoption in some business applications.

Recommendations:

Research is needed on:

- How to lower barriers to adoption and use of existing trust models (e.g., digitally signing mail streams with either Pretty Good Privacy (PGP), , Gnu Privacy Guard or Secure/Multipurpose Internet Mail Extensions (S/MIME)
- The linkages among trust, identity (or anonymity), and reputation

2.4 End-to-End Network Security

End-to-end secure, dynamic, seamless, transparent and heterogeneous network environments are possible today through the use of such protocols as Secure Shell (SSH). However, in practice, security issues are inherent in all layers of the protocol stack.

2.4.1 Down the OSI Stack

Lower protocol stack layers must be secure to secure the higher protocol layers. Yet in recent years, only limited research has taken place on insecurities at the Layer 1 (physical layer) or Layer 2 (data link layer), despite changes in network operational practice that include nation-wide Layer 2 networks, and national and regional optical networks. Currently known/familiar threats at lower levels of the Open System Interconnection (OSI) stack include Address Resolution Protocol (ARP) spoofing, MITM attacks at Layer 2, and Layer 1 attacks such as passive optical taps or the interception of wireless network signals by attackers. While these attacks are well known, networking researchers need to focus on detecting and addressing these threats in scalable ways.

Less familiar attacks affecting the lower levels of the OSI stack (such as the physical layer) include intentional attempts at kinetic (physical) destruction of key national network infrastructure by terrorists or hostile nation state actors; and electromagnetic destruction of network assets using high power microwave weapons, or high altitude electromagnetic pulse (EMP) effects.

Recommendation:

Research is needed to address known and anticipated security threats to lower layers of the protocol stack, including:

- Identifying key Internet assets (such as transoceanic cable landing points, major network traffic exchange points, locations where multiple long haul networks are channeled into common corridors due to a lack of alternatives, etc.)
- Developing methods to harden and improve the survivability of high value networking assets in a physically hostile environment
- Systematic testing to quantify the vulnerability of commercial network equipment to electromagnetic pulse (EMP) effects
- Identifying approaches to countering EMP risks, e.g., shielding at time of manufacturing, post hoc shielding, and the use of all optical (rather than OEO) interconnects

2.4.2 Up the OSI Stack

Networks worldwide have come under increasing attacks "up the OSI stack," particularly at the application layer. The SANS@RISK data from November 2006 to October 2007 indicate that over half of the 4,396 total vulnerabilities relate to web application vulnerabilities such as Structured Query Language (SQL) injection attacks, cross-site scripting, cross-site request forgeries, and Hypertext Preprocessor (PHP) remote file inclusions (see www.sans.org/top20/#s1). The increasing emphasis on application layer vulnerabilities reflects efforts to obtain sensitive financial information such as credit card numbers or other personally identifiable information in the government and commercial sector. In the context of this report, we presume a focus on information for both counterintelligence and the protection of proprietary competitive information.

Proper application of encryption to data in transit and data at rest, along with improved application development practices to eliminate complications like SQL injection attacks, would mitigate network security risks. Since these tools have not been ubiquitously operationally deployed, we need to understand and eliminate the barriers to their use.

Phishing, a social engineering attack on confidential data, can impair system integrity when users volunteer their passwords. Research is needed into human factors to better understand how to keep human participants in complex security systems from serving as the "weakest link."

SSH and Secure Sockets Layer/Transport Layer Security (SSL/TLS) encryption, along with two factor authentication (the use of both something you know, such as a password, and something you have, such as a hardware cryptographic token), could largely defeat technical credential capture. Although SSH and SSL/TLS are widely used, research is needed on economic and human factors issues to understand why they and two factor authentication approaches have not been more widely used.

We don't know how to deploy two-factor authentication at scale. Most users currently require one token for each service, e.g., routers and other network devices, and commercial services such as personal bank access and stock brokers, potentially acquiring a large number of tokens since services are not set up to use the tokens of other services, . Federated approaches based on Shibboleth have great potential in this area, but deployment/adoption has been slow.

Research is needed on how to eliminate continued reliance on simple passwords transmitted in plain text, an outdated and insecure foundation technology still widely used across the Internet.

Pretty Good Privacy (PGP), e.g., Gnu PrivacyGuard, has the potential to substantially improve the privacy and integrity of the email application. Its deployment, however, remains limited. We need to understand how to overcome barriers to the adoption of this capability.

Spam now constitutes 90% of all email; within 5 to 15 years it could constitute over 99% of all email unless we deploy effective countermeasures.

Recommendations

Security research is needed on:

- Identifying and lowering barriers to deploying currently available effective means of network security
- Addressing human factors in security to eliminate the user as the "weak link"

- How to deploy two-factor authentication at scale
- How to eliminate dependence on unsecure practices, e.g., simple passwords transmitted in plain text
- Eliminating spam

2.5 End to End Security in Diverse Network Environments

Network security should be independent of knowledge, or assumptions about underlying transport technologies. A user might take advantage of ethernet, wireless, optical lambdas, packet over SONET, ATM, FDDI, etc., and combinations of these technologies on a single connection. If the network is a passive transport media (rather than an active participant in the security process), security becomes an application layer problem on the host rather than a network layer challenge.

2.5.1 Closer Coordination

Siloed communities (security, networking, application users and developers, security and system administrators, and Federal networking and security program managers) currently implement security for networking using multiple entities. Cooperation and improved communication among these groups is needed to improve the effectiveness of security research, development, deployment, and management. Similarly, Federal networking and security program managers need to coordinate more closely with:

- Higher education operational security practitioners
- Higher education security researchers
- Commercial system and networking security entities
- Civil and criminal cyber law enforcement agencies

Recommendation

- Develop mechanisms and opportunities to foster planning, cooperation, and interaction among network and security users, application users and developers, security and system administrators, and Federal networking and security program managers

2.5.2 Moving the Security Perimeter Into the Host

Network speeds continuously increase such that today's Internet speeds of 10 Gbps will become 100 Gbps and 1000 Gbps over the next 15 years. Security chokepoints (e.g., firewalls) will present significant challenges to continually keeping up with this dramatic increase in speeds. To scale border protection, the security perimeter needs to be moved to "two inches into the host", for example, by putting network security policy onto a trusted network interface card/chip.

Recommendations

- Research is needed on making the Network Interface Card (NIC) site-configurable, rather than host-configurable, and auditable. It would report events as required by configured security policy

The workshop participants indicated that verifying host and operating system integrity was out of scope for purposes of this report.

2.5.3 Security Implications of Circuit-Oriented Architectures

Circuit-oriented, point-to-point, wide-area optical architectures have emerged as a focus of the government and academic advanced networking computing communities, particularly for high bandwidth science applications. The limitations of security technology in dealing with the high speed and performance requirements of the advanced networking/computing community may result in circuit-oriented architectures bypassing traditional perimeter security appliances such as firewalls or intrusion detection systems.

If one system, or a small subnet of systems, connects via a switched optical network connection to another small subnet, thereby forming a small closed collaborative enclave, the potential population of attackers would be limited to those who have access to the collaborative enclave. However, if the circuit-based architecture bridges sensitive networks to public networks, and the optical network explicitly avoids public firewalls, the sensitive networks can be exposed. Imagine a scenario with two sites interconnected by a point-to-point optical network:

Internet ==> host at site one ==> optical network
==> host at site two ==> sensitive internal network

If the optical network element explicitly avoids institutional firewalls, this architecture may enable asynchronous or phased undesirable access. For example, content from the Internet could be introduced at one time, allowing access to sensitive internal networks at a subsequent time.

A partitioned "red/black" network architecture can prevent cross-contamination, much as secure government networks are currently air-gapped from the Internet. But strict partitioning comes with substantial real and intangible costs. This approach, applied to an unclassified environment, needs careful study. These concerns apply beyond optical networks to other point-to-point environments, including tunnels and virtual private networks (VPNs).

Recommendation

- Fundamental research is needed to develop strategies to address security vulnerabilities introduced by optical networks that may both avoid firewalls and be attached to sensitive networks

2.6 Network Security Meets Secure Network Traffic

Future network traffic will likely use end-to-end encryption. Network traffic analysis (an important component of future network-based security systems) will be impeded by this encryption since traffic monitoring and filtering may have access only to source and destination addresses and traffic history. Even when users send traffic in the clear, the norm for open science data, the sheer volume of data flows guarantees that pattern-based detection will often misfire, shifting the burden once again to traffic analysis.

Rcommendation

- Research is needed on security systems using content-blind rules or heuristics for intrusion detection and prevention systems. Inputs to such rules can include source and destination addresses, security associated ID, times of observation, and some key negotiation traffic

2.7 Challenges of Distributed Security

Security attacks are increasingly distributed such that their detection and defense often require a distributed solution. Optical circuit-switched paths may cross several administrative domains, adding to the complexity of solutions. Traffic flows often take asymmetric paths, making monitoring and control from a single location impossible. No current intrusion prevention system works in the face of distributed asymmetric flows. Coordination is often very limited among incident response groups across domains.

Recommendations

Research is needed on:

- Performing authentication prior to establishing connections over optical-switched paths
- Distributed intrusion prevention systems
- Methods for more direct sharing of performance and incident detection data across domains

2.8 Control Plane Security

Dynamically switched optical networks employ control planes and signaling technologies. Systems at the endpoints of dynamically switched optical paths may make assumptions about the origin of traffic arriving on those paths. Compromise of the control plane – or accidental flaws in its design or operation – can invalidate those assumptions, with unpredictable effects. Control plane traffic is commonly carried in-band. Even when the traffic falls out-of-band, the possibility of its appearing in-band by error may exist.

Recommendation

- Research is needed on end-to-end security mechanisms for the control plane (in addition to the isolation of control plane traffic) that provide robust protection against both partial network failures and active attacks through the physical media

2.9 Network Security Test Beds

A security test bed is needed immediately to test and deploy capabilities and to determine how the community of users and network engineers responds to them. Researchers should also use test beds in verifying the usability of security designs. Currently available security test beds remain small in scale, have limits on acceptable testing, or are classified and thus unavailable for non-classified research.

Recommendations

- A network security test bed is needed with inputs from attack traffic datasets
- Applications should include security metrics and a discussion of security considerations
- New network architectures should include a security model

2.10 Additional Network Security Research Topics

Additional recommendations of the Security breakout group included the following.

Recommendations

Network security research is needed to address:

- Embedded systems, sensornets, and wireless networks

- Soliciting, analyzing and synthesizing existing network security R&D roadmaps and plans
- Control theory and Kirchoff-type laws (to capture normal behaviors for routers) to identify whether secure systems can be composed from insecure components (or even from secure components)
- Metrics for characterizing system security or privacy

Section 3: Federated Optical Networking

Next generation optical networking holds the promise of enabling substantial advances for communications technology. These advances will enable not only improved versions of current applications, but also many applications that do not exist today. New macro architectural concepts will lead to advances in capacity, robustness, security, real time capabilities, and economic models.

To achieve the potential of advanced optical communications, the networking community needs to define and systematically address major challenges, starting with architectures and prototypes for a "system-of-systems" or "meta system" for next generation optical networks. Macro architecture investigation should encourage high risk, high potential interdisciplinary research areas. A wide ranging research program is required for progress in these areas, which would address multiple key topics, within a macro context. This research is beyond the capability of individual laboratories. It will require major national and international facilities capable of supporting multiple large scale experimental network test beds.

3.1 Basic Research

Enabling optical networks to operate at 100 Gbps and beyond (e.g., 200, 300, and 1,000 Gbps, etc.) will prove essential for meeting future communication requirements. Emerging applications such as petascale science, high energy physics, and digital media will soon demand 100 Gbps and beyond. The IEEE has established a 100 Gbps working group. Optical capacity per se is not the major issue (capabilities are being demonstrated today that can transport 8 Tbps). The basic issues relate to optimizing the use of such large capacity resources for network architecture, technology, and infrastructure. A comprehensive approach is needed to optical networking research across a wide range of research areas, as opposed to addressing elements individually. Addressing switching requirements, for example, constitutes a key research need. Networking researchers can only address these requirements by providing solutions within the context of a wide range of other research areas. There are multiple architectural and technology considerations that need addressing, including such fundamental issues as developing serial or parallel solutions. A comprehensive research program must encompass high risk, high potential interdisciplinary research areas incorporating long term objectives.

Together, the design, construction and operation of cost-effective nodal switching units and elements are a major research area. A set of key issues relates to electronic, rather than optical, elements, including optical-electronic boundaries such as those within switching devices. Increasingly, the life-cycle energy cost of equipment can exceed capital equipment expenditures, both for power and for cooling. Such life-cycle costs are

a basic consideration in the research and development of new technology. All photonic switching provides significant power, cost, capacity, and flexibility advantages.

Other research priorities include:

- New types of transponders
- Customized components, especially those with simple interfaces
- Optimizing network interfaces
- High capacity network edge technologies, e.g., faster disk transfer technology
- Optical back planes using new types of materials
- Custom optical fibers providing enhanced capacity and flexibility

Recommendations

- A broad-based research program is needed in optical networking to address architectures and prototypes for a "system-of-systems" or "meta system" supported by multiple large-scale national and international experimental test beds
- An interdisciplinary basic research program is needed to support large scale, longer term research in optical networking to provide fundamental change
- Basic research on network architecture should include high risk, high potential interdisciplinary research areas incorporating long tem objectives

3.2 Integration of Optical Networking Technology into Existing Capabilities

Emerging capabilities for 100 Gbps (and higher) capacities impact all network elements: protocols, interfaces, switching elements, service granularity, specialty appliances, and others. Generally, research in these areas has focused on specific individual elements versus comprehensive perspectives. These research areas underscore the need for a new general architectural framework as a context for federated optical networking research.

Recommendation

- Basic research is needed for interdisciplinary conceptualization and experimentation on optical networking technology with all network elements, protocols, interfaces, switching elements, service granularity, specialty appliances, etc.; Experimental general architectural frameworks should guide these areas of research

3.3 Federated Optical Networking Architecture

New architectures are needed to enable the operation of large-scale federated optical systems. Network scalability requires a high degree of decentralization, inter-domain provisioning, and new capabilities, which collectively present many challenges. New multi-layer capabilities will require the communication of more information among individual devices and across multi-domains in a systematic way, including over-dedicated channels. Locating, isolating, and responding to multi-domain problems is difficult. New debugging capabilities adding attributes of self diagnostics and self repair to protocols could make management and debugging of multi-domain networks more automated. As increasing Layer 1 and Layer 2 resources become available across multiple domains, mechanisms will be needed to provide management and control without traditional centralized capabilities.

Currently, the majority of research frameworks for optical networking are oriented toward single investigator efforts or small groups. New processes, organizational

structures, support mechanisms, and facilities for optical research and development are needed to enable large scale research and development projects.

Recommendation

- Research is needed on designing and developing a new system-of-systems (meta system) architecture and prototypes; Multiple options for centralized, decentralized, and hybrid centralized/decentralized approaches for innovative control and management should be experimentally investigated

3.4 Dynamic Switched Optical Networking

High-end science pushes the envelope of networking capabilities, fostering the development of new needed services and capabilities such as high-rate data transfers and dynamic collaboration across network domains. To provide these services, current federated optical networking researchers are developing dynamically switched multi-domain optical networks in Layers 1-3 (L1, L2, L3, respectively) including control and signaling, data plane peering/circuit exchange, end-to-end (E2E) secure circuits, E2E circuit monitoring, E2E circuit protection and restoration, interplay of circuits, and best-effort Internet protocol (IP) traffic. We will still need point to point connection services, both those complementary to Internet services and ones that provide alternatives to traditional IP routed networks. .

This research would benefit from the development of a generalized architectural framework to provide admission control and cost considerations. Designs should include capabilities for non-traditional policy-driven access mechanisms and for better approaches to cost considerations. There are major challenges to deploying new architecture, protocols, and technology in today's networks. System level issues need to be addressed when large capacity E2E paths are introduced into networks that are based primarily on delay tolerant networks with many packet buffers. Rapid link changes can cause instabilities in L3 networks. Lower network layer changes can severely disrupt L3 protocols. Border Gateway Protocol (BGP), for example, is optimized for stability while route dampening is used.

Millisecond changes at the optical level cause problems with line cards because route convergence takes seconds to minutes. If topologies change at L2, routers will attempt to adjust. L2 changes will have to be hidden or protected from these types of topology changes. Timing differences among layers have been a long-recognized, fundamental problem. New protocols will be required to reroute traffic to lower network layers for the optimization of traffic flows to utilize capacity while preventing instability and disruption. To some degree, this issue relates to packet formats.

Several research projects focus on the development of all-optical routers. However, progress has been slow in this area. New dynamic optical path provisioning techniques, are currently under development for control planes. A fundamental reexamination of the common practice of overprovision as a method to achieve capacity should take place. E2E dynamic optical path provisioning can be used to implement capabilities such as segmentation, ultra high capacity, and specialized admission control to enhance security. State information is particularly important to the deployment of these capabilities.

Recommendation

- A comprehensive research program is needed to explore dynamic switched optical networking within a large scale architectural framework.

3.5 Hybrid Packet/Circuit Switched Optical Networking

Research projects currently address issues of mixed technology E2E paths, integration of IP-Quality of Service (QoS), Multi Protocol Label Switching (MPLS), and Generalized MPLS (GMPLS) as separate issues. Few R&D efforts address them comprehensively. In addition, as new capabilities become available, they will require new service and operational models and capabilities. New methods for unambiguous alarming may be useful, especially if combined with enhanced communication of operational information.

In this area, the Defense Advanced Research Projects Agency (DARPA) funds the Core Optical Networks (CORONET) research project, which presumes 100 Gbps, 50 Ghtz networks and expects capabilities for interactions among layers.

Recommendation:

- Develop a comprehensive architecture of systems ("meta-system") for hybrid packet/circuit-switched optical networking

3.6 Optical Networking Standards

Most standards organizations concentrate on short term, incremental improvements of existing technologies, not major challenges, high risk, and long term projects. The agenda of these organizations also focuses on narrowly defined areas such as the International Telecommunication Union/Automatically Switched Optical Network (ITU/ASON), Internet EngineeringTask Froce (IETF)/GMPLS and MPLS, and Optical Internetworking Forum (OIF). Standards organizations tend to react to research activities and conclusions and they usually do not set research agendas. The GLIF R&D community has formed a partnership with the Optical Grid Forum (OGF) specifically to address issues related to developing new architecture for network interfaces that would take into consideration emerging techniques such as those required by federated optical domains. Individual members of the GLIF R&D community participate in various standards organizations, including the IETF and IEEE. A comprehensive systems approach is required, for example, to enable enhanced communication between network layers. Applications that require these new capabilities have begun to emerge. An application project in the U.K., for example. focuses on the design of a very high speed serial optical channel to support an 8k format at 380 Gbps uncompressed, using L1 multicast techniques. Some emerging radio astronomy applications require 10 Gbps per interface and multiple channels (e.g., 32,000 channels) to central correlators. These are examples of applications that do not require high performance edge processors to generate multiple streams.

Recommendations

- Develop new standardization processes to expedite advances in the design, development, and adoption of new innovative capabilities
- Develop processes for migrating research results to wider communities through standards organizations
- New standards organizations may be needed to address optical networking topics outside the scope of existing organizations

3.7 Optical Network Technologies

All-optical-networks remain a potentially high risk, high reward research direction. Further investigation is required to evaluate the level of risk. Optical GMPLS is moving from a high-risk technology toward a lower-risk capability and it has become a useful tool for supporting large scale science applications. Optical Burst Switching (OBS) is a high risk, potentially high reward area. Network virtualization at the optical layer remains an essential objective. As both optical and wireless technologies evolve, it is important to consider the intersections of these technologies. Each of these areas has attracted motivated researchers that would like to investigate the full potential of the focal technology. OBS may be useful for delay tolerant networking (DTN). Low dispersion technology over long distances is needed; however, attenuation is a more critical issue. High capacity does not necessitate low dispersion fiber because new techniques, such as electronic dispersion removal, can assist in breaking up non-linearities. Nanotechnology research provides customized materials for functionality within fiber. The potential for using state information to improve networking services and technology support capabilities has not been sufficiently investigated; this topic deserves additional research efforts."

Recommendation

- A comprehensive long-term research program is needed for a wide range of potential new optical networking technologies including optical GMPLS and related control plane protocols and technologies, optical burst switching, network virtualization at the optical layer, optical wireless technologies, dynamic provisioning, mechanisms for utilizing state information and many other technologies

3.8 Test Beds in Optical Networking

Optical networking research requires large scale national and international test beds. While much solid research progress can be made in labs using modeling, simulations, and large spools of optical fiber, these limited lab activities cannot substitute for large scale test beds. Test beds at both the national and international level are needed to explore new methods for achieving economies of scale based on new technologies, for example, using approaches based on meta systems. Test beds enable researchers to experiment with new "clean slate" concepts that are not merely extensions of existing technologies. New processes and technologies provide test beds with greater flexibility than previous test beds to enable a wide range of experiments.

Researchers can employ optical test beds to experiment over several layers of the protocol stack using new components, including fiber, at large scale. Optical test beds can also play a role in the investigation of physical effects and the relationship of those effects to new types of component and equipment designs, including issues related to interoperability. Production and operational networks are difficult to use for experimental optical research and few large scale research optical test beds have been implemented recently. Therefore, there is a strong need for new test bed facilities designed and implemented to support optical networking research.

Research on international inter-domain issues (e.g., heterogeneity and policy issues) requires international optical test beds. Designing and using new capabilities would benefit from novel architectural and technical approaches that operate at a global scale. International test beds are needed to support experimental research in this area. The development of global test beds will depend on international partnerships.

Recommendations

Research is needed to:

- Design and implement multiple large scale national optical test beds with differing characteristics to enable multiple simultaneous cooperative and interdisciplinary experiments by different communities of researchers
- Design and implement international optical test beds for experiments that leverage existing test bed and fiber resources (e.g., the GLIF community and segments of undersea fiber)

3.9 Additional challenges in federating optical networks

Other fundamental research topics for federated optical networking that warrant investigation include the definition of granularity, Time Division Multiplexing (TDM), lowering the cost of transport using optical channels, and scalability for many techniques and technologies. Optical networking research will need increased programmatic support to address program structure and organization, policy development, facility design and implementations, and funding structure.

Recommendation

Optical networking research programs are needed to address a wide range of capabilities, including:

- Optimizing granularity capabilities
- TDM and related technologies
- Signaling and control plane architecture and technologies
- Lowering the cost of transport using optical channels
- Scalability for many techniques and technologies

Additional efforts are need in the areas of:

- Research program structure, organization, and processes
- Policy development
- Facility design, implementation, and operation

3.10 Coupling of Optical Network Research to Application Requirements

Application requirements often drive networking research and development of new capabilities. Models exist for determining a balance between pure R&D without reference to applications and those with application contexts (e.g., 70-30, 20-80 pure-applied research). A mix of both basic and applications-related research is needed.

Recommendation

- Research programs for federated optical networks need to support application oriented requirements as well as basic R&D needs

3.11 Funding Mechanisms for Optical Networking

The networking community should re-evaluate traditional models for funding research to provide a balance between very large scale projects (e.g., Manhattan Project scale) and single investigator scale projects, as well as mid tier projects. Funding mechanisms are needed that allow for rapid changes in research objectives.

Recommendations

- Design and implement new processes, organizational structures, support mechanisms, and facilities for optical R&D
- New research programs should provide a balance between large scale, mid-range, and single investigator programs

3.12 Economics of Networking

Various models exist for the analysis of the economic contexts for the development and deployment of new networking architectures and services, ranging from the macro scale (e.g., the requirements of the national economy and national security) to the micro scale (e.g., low cost manufacturing leading to low cost commodity optical components). Significant economic considerations also factor into barriers that impede the adoption of new technology and technology transfer from the research sector to the commercial sector. These economic considerations are generally not well understood.

Recommendation

- Methods for evaluating economic considerations in research and network deployment, including risk reward frameworks, merit consideration within the context of a general network architectural framework

3.13 Commercialization of Technology and Technology Transfer

The demise of the major commercial communications R&D labs has diminished opportunities for both R&D and technology transfer. The short term focus of standards organizations also reduces opportunities for the technology transfer of new research. Many technology transfer and commercialization processes concentrate on provider services. We must create a broader context to decrease the time from research results to commercial availability and user adoption.

Recommendations

- New research mechanisms and processes are needed to promote basic R&D processes that are no longer supported by the commercial sector
- Additional technology transfer processes are needed to address the requirements of non-traditional constituencies

3.14 Critical Research Areas

The research areas enumerated below are of critical importance for advancing optical networking capabilities.

Recommendation

Networking research is needed to address:

- Wavelength switching (e.g., cheaper, smaller, better, faster…)
- Architecture and protocols to maximize the use of photonic layer adaptive topologies
- Maximum use of capabilities, functions, components
- Wave length selective switches
- Enhanced access capabilities
- Terabit LANs
- Photonic RAM
- Optical packet switching

In support of these research goals, facilities are needed that are capable of supporting many test beds, many experiments, high density photonics and electronics (P&E), P&E integration, and optimized infrastructure.

Limited mechanisms exist for the coordination of inputs across the many constituencies of the networking community to establish priorities among these areas of potential opportunity.

Recommendation

- A process is needed to coordinate inputs across the many networking research community constituencies to determine priorities among research opportunity areas

Section 4: Heterogeneous Networking

Future networks will have more complexity and heterogeneity than the current Internet. They will link circuit-switched and packet-switched networks, high-speed optical paths, intermittent planetary-scale paths, sensor networks, and dynamic ad hoc networks. These varied network forms will involve millions or billions of interfaces that will change dynamically. Understanding the behavior of such systems remains, in itself, an enormous technical challenge. Researchers need to develop new tools, based on models and analysis of complexity in heterogeneous networks, to enable network administrators to manage and control these networks, diagnose their faults and failures, and recognize and respond to attacks. Emphasis is needed on technical approaches to attain simplicity and transparency of design.

Visionary Heterogeneous Application

In the future a new type of attack, capable of causing massive system failure and release of sensitive data, will be launched against a networked system. The attack hits the kernel of the operating system and hardware, making it resistant to re-booting. With new technology designed to automatically manage a response across a complex system, the attack is quickly detected; a signature to stop it is synthesized, distributed out of band, and applied throughout the network, slowing the spread of the attack; the attack code is reverse-engineered so that a patch can be synthesized and distributed; the patch is installed to eliminate the vulnerability and restore all systems to an operational state.

4.1 Information across Interfaces

Heterogeneous networks must transport (and interpret) services across the technology interfaces, requiring data transformation or adaptation. The range of cross-domain interpretation issues is extensive. Management information exchange across domain or administrative boundaries has generally proven inadequate to achieve end-to-end path, link, or application optimization. This circumstance makes Service Oriented Architecture services, such as mobility, security, and walled garden markets, difficult to implement.

Performance measurement for heterogeneous networks (either for control plane feedback, or application feedback) is inadequate and ill-defined. TCP, for example, interprets loss as congestion. Span-by-span link state information such as lost packets is needed across the boundaries of the heterogeneous network segments. Useful information is needed even in partially hidden (domain-referenced) environments.

Recommendations:

- Research is needed to develop capabilities for modeling and analyzing unprecedented levels of network complexity
- Research networks need to implement end-to-end performance measurement, and the ability to share performance information, across heterogeneous technologies and domains
- The NITRD Program (through its constituent agencies) should sponsor a series of workshops on next-generation heterogeneous networking to address:
    - Technology-bridging protocols
    - Self-adaptive protocols
    - Control and management in the context of partially hidden link-state

_ End-to-end security in heterogeneous networks
_ Advances in the physics and engineering of new network devices

4.2 RF Spectrum

Radio frequency (RF) spectrum for networking is limited. This spectrum is also highly fragmented inside and outside national boundaries, and, in the U.S. is constrained by military allocations. These limitations create difficult barriers to next generation RF R&D. Commercial providers exacerbate the problem by walling off their market. Current demand for spectral space (e.g., 802…) has far exceeded available spectrum space. RF congestion has become so pervasive, particularly in the U.S., that it places U.S. R&D efforts at a competitive disadvantage.

National and commercial "walls" guarantee a lack of true mobility in the large scale. Intelligent or cognitive radio can alleviate some, but not all of the problems.

Recommendation

- Open the educational TV spectral allocation in the 2.4 – 2.6 Ghz band for networking research and development

4.3 Dark Fiber

Dark fiber, available for sale at modest prices in the U.S. since the mid 1990s, has become increasingly scarce and expensive to acquire. Lack of access to all layers of the protocol stack (facilitated by use of dark fiber), seriously hinders much basic research in networking.

Recommendation

- As the U.S. mandates creation of new high voltage transmission line rights of way, it should make them available for national-scale fiber infrastructure

4.4 Test Beds

Recommendation

- National scale test beds should extend to both ends of the heterogeneous technology spectrum including RF and fiber media. These test beds need to support transparent development from the application to the infrastructure layer

4.5 Barriers to Commercialization

Research addressing barriers to commercialization is needed to facilitate uptake of emerging technologies and broad user adoption. The Federal government and its private-sector partners should enhance existing research programs to carry out comprehensive, complementary, and synchronized actions focused on attaining these high-priority goals. As networking visions, capabilities and research needs advance as a result of these actions, the Federal government and its private-sector partners should coordinate their efforts on changing research requirements. Federal R&D progress toward the goals, in conjunction with complementary private-sector efforts, is needed to accelerate the evolution of advanced networks.

Recommendation

- Research is needed on methods to reduce barriers for the commercialization and user adoption of new technologies

Section 5: Network Science and Engineering

The networking vision for a 20 year horizon is for users to experience a natural and richer interaction with the network that will encompass complex social, economic, and policy related interactions mediated by the physical communications infrastructure. The user experience – including the applications that interact with the network – will evolve seamlessly and in a manner responsive to user/stakeholder desired outcomes. The key challenge in this arena centers about providing seamless transparency to the wide diversity of users in a complex highly heterogeneous environment. Transparency through simplicity of design should be one of the goals of networking research. Additional goals include stakeholder access anytime-anywhere, providing:

- Minimal cost
- Reliability
- Security and privacy
- Limitations imposed only by fundamental physical constraints
- Responsiveness to users (stakeholders) needs

The complex network will have to satisfy dynamic, evolving requirements for conflicting interests at differing scales. A holistic cyber infrastructure of resources will provide the fabric to create and enable these capabilities. It will combine highly heterogeneous devices, availability, reliability, speed, performance, channel characteristics, and connectivity. Security will include not only protection of information but also privacy and usability. To achieve this security, the network will incorporate law, policies, social, and economic tools (for multiple, self-interest driven autonomous entities) to enable resource sharing and to foster collaboration.

5.1 Fundamental Research

Fundamentals research is needed on how to model, simulate, analyze, measure, design, build, deploy, manage, monitor and evolve envisioned future networks. We will need to understand complex networks at different levels of granularities using models, abstractions, and tools to understand the interplay among different components of the complex structure. Researchers need to integrate social, economic, and regulatory aspects of networks into their design, management, and operations.

The network influences the growth of new social, cultural, economic, and political norms and the emergence of new network-centric forms of social interaction have a profound influence on network development, deployment, and adoption. Understanding how network architecture affects these emergent social applications and vice versa, constitutes a significant challenge to present day network science and engineering.

In future networks, intelligent agents will increasingly act in the "user" role – requesting services, collecting and analyzing information, and generally acting as a representative of the human user within cyber-space. Even today, we see the beginning of network-mediated social evolution through the pervasive use of Google search, Facebook, Twitter, Amazon, Ebay, YouTube, 2nd Life, global gaming and entertainment,

cyber-warfare, environmental sensing and surveillance, and many other applications. Easily available information on a global basis has had a truly significant impact on society. On-demand access to global information is likely to continue to influence social, economic, scientific, cultural, and political evolution in the future.

Network scientists have not developed the ability to study the impacts of such cyber-mediated activities and relate them directly to aspects of network design. The correlation between high level applications activities and network performance is not well understood and, within the context of present day network architectures and commercial network services, is all but impossible to study. We have no effective way to study whether other network architectures could influence the development and evolution of these emerging classes of networked applications.

Future networks should enable social scientists and computer or network scientists to study usage patterns, postulate new architectures and/or designs, instantiate those concepts, and then evaluate the effect these innovations have on the relevant user community. Existing research techniques are unable to capture and study resulting networking data or to effectively characterize the effect these cyber-space social structures may have on the network or vice versa.

Non-technical issues such as privacy and policy prevent much detailed information from capture and study by even the most well-meaning and secure research teams. Technical aspects of commercial service providers – both the transport/network service providers as well as the content providers – often remain closely guarded competitive and proprietary secrets.

Recommendation

- A multi-disciplinary research program is needed to study key socio-technical questions arising from the emergence of new network environments. The program should develop effective and rigorous methodologies for studying network-centric social applications and environments

5.2 Architectural Frameworks and Design Principles

The current networking architectural framework and available design methodologies are inadequate for developing future networking architectures that can encompass the complex requirements of the future socio-technical infrastructure.

Network researchers need to explore the layered architectural concepts to assess whether they are fundamental and how cross-layer approaches might be incorporated. New architectural paradigms are needed that are hardware independent, evolvable, and extendible, to allow for flexibility in the designs across multiple platforms. Within these new architectural approaches, research is needed to explore the design of uniform (universal) frameworks for virtualization at scale across different levels of the architecture and for different resources. This research must reassess fundamental network design principles and develop new end-to-end arguments in light of evolving socio-economic infrastructure and environments. Within this context, scientists need to develop paradigms that can explore issues such as network neutrality. Specific goals for architectural research include:

- Rethinking the principles of naming and the tight coupling of paths addresses with the goal of achieving a richer and secure name management system

- Creating new and aware protocols at different levels that can address real time flows, mobility, dynamically changing environments, and different channel characteristics
- Addressing how the network handles data intensive collection, storage, and dissemination – in multi-point to multi-point environments – and provides support for dynamic resource allocation at scale to incorporate traffic engineering

Research should explore new control and management paradigms, frameworks, and tools that provide adequate interfaces and effectively integrate physical, logical, and human components and resources of the complex networking structure. This research area is closely coupled with the development of new computationally efficient, scalable, algorithms that allow planning, scheduling, and resource allocation, and enable and enforce diverse policies for diverse stakeholders.

Recommendation

- A comprehensive research program is needed for fundamentally new end-to-end networking architectures in the context of the emerging complex socio-economic infrastructure. The architectural approaches must provide for flexible design across multiple platforms using evolvable and extensible technology, enable virtualization at scale across multiple platforms and for different resources, and design for cross-layer functionalities.

5.3 Services and Applications

The current networking services model lacks scalability and the ability to cover service requirements of emerging global applications. Future applications will also require a high level of predictability from the network.

Our view of the network continues to change from that of a transparent cloud to a set of managed resources that can be used to address a problem. This new paradigm introduces several research areas that merit exploration:

- Virtualization – "Virtual machine" environments have proven critical to the efficient use and sharing of computer resources. Going forward, the ability to abstract basic network services and then to virtualize these services – i.e. separate the functional capability from physical hardware - will provide distributed applications with a novel new capability to create network environments designed to address the requirements of that specific application. This capability represents a significant departure from the basic principles of a single shared "best-effort" Internet of today. Virtualization will allow the network science and engineering community to develop a set of formal methodologies for dynamic management of network resources for users that open a wide range of capabilities to the broader networking community.
- Modularity – The design and development of future network technologies – through software and architectures – should incorporate open modularity. "Modularity" in this context means to provide publicly defined, openly accessible functional blocks within the data protocols and data handling that simplify implementation, configuration, testing, research, and experimentation. A network fundamentally architected and designed to function as a set of interoperating modules would provide an open network to allow experimentation, instrumentation, reconfiguration, and substitution of modules. Such object functionality could potentially create an "object market place" where vendors could sell specific implementations of particular objects or the

research community could develop an experimental object that has novel or unique features to facilitate networking.

- Composability – Modular network architecture should allow efficient composability, (grouping a set of modular functions together to create a new or customized set of modules incorporating the resource characteristics and constraints of the constituent modules). Composability also has implications for the mating of islands of similar functionality to create a new service environment with a broader scope. These principles should be intrinsic to the design of a modular, virtualized future network.

These future network characteristics – virtual services, modularity, and composability – create a much richer network service environment for the end user. They provide the user with a dramatic increase in predictability and manageability of the inter-process communication function among a set of cooperating cyber agents. Yet current networks deploy little technology that provides such capabilities.

Recommendation

- Develop a dual approach to creating a new network architecture: (1) define and fund a research program to explore theory in network virtualization, modularity, and composability; and (2) fund applied research and development of these concepts to create fully functional – but experimental – environments that can act as a proving ground for these new service models

5.4 Experimentation and Test Beds

The development of new networking capabilities depends on networking test beds. Such test beds could, for example, support collaborative research activities in network science and engineering that address the scalability challenges and complex multi-disciplinary issues in the required depth and breadth for future networks. Test beds provide a critical bridge between the theoretical architectural development and deployed production. This experimental stage – during which the results of detailed and rigorous research and prototypes can be reworked with robustness and scaling as core objectives, and undergo testing in real world conditions without the same hard reliability requirements of a production environment –is essential to moving technologies as quickly as possible from theory to widespread deployment.

Limited-scale lab test beds are insufficient to support our understanding of complex realistic future networks that include massive scaling of technologies and applications. They can contribute to the construction of operational prototypes of new concepts, but currently, they cannot operate effectively under real-world conditions.

Large experiments require significant amounts of time and money in addition to careful planning and construction. They can be performed on test beds, as well as operational networks, or by simulation. Rigorous methods are needed to design experiments, conduct them, and evaluate the results. Users must be included as key participants in these experiments.

Human factors are critical for certain types of experiments, and test bed environments should tap the user environment. Instrumentation of the test beds should adequately support the evaluation of experiment metrics. Existing capabilities (e.g., performance monitoring), should be leveraged in the design and use of the test beds. The relationship between simulations and test beds warrants exploration.

Recommendation

- Objective metrics are needed to evaluate the efficacy of test beds in emulating real environments

Appendix 1

Workshop Registrants

Tomonori Aoyama,	Keio University
Suman Banerjee	University of Wisconsin-Madison
Karen Bergman	Columbia University
Joseph Berthold	Ciena
Bos Erik-Jan	SURFnet
John Bowers	UCSB
Misha Brodsky	AT&T Labs
Maxine Brown	University of Illinois at Chicago
Natalia Bulashova	Russian Institute for Public Networks (RIPN)
Ian Chakeres	Motorola
Gee-Kung Chang	Georgia Institute of Technology
Amy Clark	United States Department of Energy
Steve Cotter	ESnet/Lawrence Berkeley National Laboratory
Matt Crawford	Fermi National Accelerator Laboratory
Cees de Laat	University of Amsterdam
Thomas DeFanti	University of California, San Diego
Freek Dijkstra	University of Amsterdam
Phillip Dykstra	DREN
Chip Elliott	BBN Technologies- GENI
Franz Fidler	Columbia University New York
David Foster	CERN
Nasir Ghani	University of New Mexico
Terry Gibbons	MIT Lincoln Laboratory
Chris Greer	NCO/NITRD
Robert Grossman	University of Illinois at Chicago
Jan Gruntorad	CESNET
Chin Guok	ESnet
Robert Hartman	DREN
David Hartzell	NASA Ames Research Center
Xiangyang Huang	Computer Networking Information Center,Chinese Academy of Sciences
Wendy Huntoon	Pittsburgh Supercomputing Center / NLR
Suzanne Iacono	National Science Foundation
Julio Ibarra	Florida International University
William Johnston	ESnet
Admela Jukan	Technische Universität Carolo-Wilhelmina zu Braunschweig
Osamu Kamatani	NTT Network Innovation Laboratories
Gigi Karmous-Edwards	MCNC
Dongkyun Kim	KISTI
Franko Kueppers	University of Arizona

Minsun Lee	Korea Institute of Science & Technology Information
Tom Lehman	USC/ISI
E Paul Love	NCO
Bryan Lyles	Telcordia Technologies
Joe Mambretti	Intrntl Ctr Adv. Internet Research, Northwestern Univ
Martha Matzke	NCO/NITRD
McLaughlin, George	Asia Pacific Advanced Network (APAN)
Grant Miller	NCO/NITRD
Debasis Mitra	Bell Labs, Alcatel-Lucent
Gabriel Montenegro	Microsoft
Ruth Ann Mullen	Photon Futures
Makoto Naruse	National Institute of Information and Communications Technology (NICT)
Thomas Ndousse	Department of Energy
Kees Neggers	SURFnet
Peter O'Neil	Mid-Atlantic Crossroads
Drew Perkins	Infinera Corporation
Ivan Philips	AARNet PTY LTD
Dave Pokorney	Florida LambdaRail, LLC
Nageswara Rao	Oak Ridge National Laboratory
Kristin Rauschenbach	BBN Technologies
David Reese	CENIC
Kim Roberts	Nortel
Sumit Roy	U Washington
David Salmon	JANET(U.K.)
Michel Savoie	Communications Research Centre Canada
Srinivasan Seshan	Carnegie Mellon University
Afrodite Sevasti	GRNET SA
Fay Sheu,	National Center for High-performance Computing (Taiwan)
John Silvester	University of Southern California
Dimitra Simeonidou	University of Essex
Jerry Sobieski,	NORDUnet
Joseph St. Sauver	Internet2 and the University of Oregon
Bill St. Arnaud	CANARIE Inc
Michael Stanton	RNP
Su David	National Institute of Standards & Technology (NIST)
Rick Summerhill	Internet2
Toshiaki Suzuki	National Institute of Information and Communications Technology (NICT)
Atsushi Takahara	NTT Network Innovation Labs
Tieniu Tan	Computer Network Information Center
Brian Tierney	ESnet
William Turnbull	NOAA
Malathi Veeraraghavan	University of Virginia

Naoya Wada	National Institute of Information and Communications Technology (NICT)
Alan Welday	DREN (DoD)
Kenneth White	NASA Integrated Services Network (NISN)
Alan Willner	Univ. of Southern California
William R. (Bill) Wing	Oak Ridge National Lab
Linda Winkler	Argonne National Lab
Jie Wu	National Science Foundation
Baoping Yan	Computer Networking Information Center,Chinese Academy of Sciences
Eugene Yeh	NCHC
S. J. Ben Yoo	University of California, Davis
Dantong Yu	Brookhaven National Lab
Taib Znati	National Science Foundation

Appendix 2
Workshop Organizers and Breakout Session Chairs

Organizing Committee

Karen Bergman: NSF Principle Investigator, Columbia University
Bill Wing: DOE Principle Investigator, Oak Ridge

Hank Dardy, Department of Defense
Cees de Laat, University of Amsterdam
Suzanne Iacono, NSF
Paul Love, NCO
Joe Mambretti, Northwestern University
Alison Mankin, NSF
Grant Miller, NCO
Thomas Ndousse, DOE
Joe St Sauver, University of Oregon
Taib Znati, NSF

Breakout Session Chairs

Next Generation Heterogeneous Networking: Bill Wing
Network Security: Joe St Sauver
Federated Optical Networking: Joe Mambretti
Network Science and Engineering: Karen Bergman