Workshop on New Visions for Large-Scale Networks: Research and Applications

Large Scale Networking (LSN) Coordinating Group
Of the
Interagency Working Group (IWG) for Information Technology
Research and Development (IT R&D)

March 12-14, 2001
Vienna, Virginia

Table of Contents

Executive Summary

This paper documents the findings of the March 12-14, 2001 **Workshop on New Visions for Large-Scale Networks: Research and Applications**. The workshop's objectives were to develop a vision for the future of networking (10 to 20 years out) and to identify needed Federal networking research to enable that vision. The meeting was sponsored by the Large Scale Networking Coordinating Group (LSN CG) of the Interagency Working Group in Information Technology Research and Development (IWG/IT R&D). The IWG, which functions under the White House National Science and Technology Council, coordinates Federal multiagency IT R&D efforts. The LSN CG agencies are DARPA, DOE, NASA, NIH, NIST, and NSF. The Network Research Team (NRT) of the LSN CG coordinated the workshop. It was attended by more than 160 leading networking researchers from universities, industry, government, and laboratories.

The participants concluded that industry is not prepared to do the long-term research needed to enable the workshop visions for future networking. Industry is oriented toward near-term development and is currently scaling back the corporate ability to provide networking research. This places increased responsibility on Federal agencies to fund and conduct the research needed to support the continuing growth of the Internet.

The workshop organizers developed scenarios of visionary uses of the Internet, including:

- **Zero-Casualty War**: The intelligent, automated, densely sensored battlefield
- **Smart World**: Intelligent, aware, secure embedded sensors for maintaining battlefield systems
- **Crisis Management**: On-line emergency resources supported by distributed sensors, dynamic networking, and distributed high-performance modeling
- **Collaboratories**: Proactive, intelligent, dynamic, "natural" interactions
- **Networked Medical Care**: Distributed medical services through collaboration with high security, high assurance, and guaranteed Quality of Service (QoS)
- **High-Energy Physics**: Collaboration with high-end, on-line resources

Discussion in workshop breakout sessions identified networking research needed over the next five years to enable these visions. These research needs include:

- **Adaptive, dynamic, and smart networking**: Networks in the future will need to be intelligent and able to adapt to dynamic and evolving situations. They will need to be self-organizing, dynamic, and responsive to applications, to support application responsiveness to networks, and to provide automated network management and QoS.
- **Measurement, simulation, modeling, and scalability**: Scalable networking technologies are critically needed to support projected massive increases in nodes, traffic, users, and complexity of the Internet. Simulation, modeling, and standardized end-to-end performance measurement are needed to support development of new technologies, standards, and network management. A long-term archive of continuous performance measurements is needed to support retrospective studies and trends analysis.

- **Trust (security, privacy, and reliability)**: Security, privacy, and reliability are ubiquitously needed to provide user trust in using the Internet and in the information provided, and to protect sensitive information transported over the network.
- **Networking applications**: The workshop participants found great merit in the visionary scenarios developed for the workshop and in additional applications identified during their discussions. Significant networking research is needed to vertically integrate applications and to provide the tools and services needed to support the applications.
- **Middleware**: Middleware assures that distributed resources and applications work in a transparent and synchronized manner to provide end user services such as seamless transport of information in a trustworthy framework and functionality across heterogeneous network elements to meet user requirements. The Grid (a high-end scientific, distributed, modeling infrastructure) developed with IT R&D funding, is an example of a developing middleware capability that needs to be extended through network research.
- **Testbeds**: Testbeds are needed to support networking research in performance measurement, security, privacy, reliability, active networking, adaptive mobile networks, intelligent networking, applications, and middleware. They are also needed to bridge the transition from the research stage to successful commercialization of technologies. Industrial participation is critical in developing and refining standards to provide technologies capable of near-term development by industry.
- **Collaboration environments**: Collaborative environments need to support transparent, intuitive human interactions with such services as automatic configuration, ubiquitous access, security, multimedia capabilities, immersive environments, expert consultation, and side conversations.
- **Revolutionary research**: With many orders of magnitude increases expected in the scale of the Internet, basic research is needed to understand network behavior and to study the complexity of networked systems. Research is needed to adapt relevant science from other fields such as chaos theory, stochastic processes, economics, catastrophe theory, and generalized control theory to understand networked systems' complexity and to address network modeling for scalability.
- **Revisit networking fundamentals**: In the future, network services and their associated protocols will need to extend across heterogeneous access technologies with seamless functionality. Fundamental changes may be needed in addressing, routing, forwarding, and transport modes to support the increased scale and functionality of the network. Research is needed to address basic issues in complexity, performance, and technology evolution/revolution.

Workshop participants emphasized the importance of networking research as a critical foundation for continuing the information technology revolution we are experiencing and for realizing their compelling visions of future networking. The workshop findings provide the LSN CG and LSN agencies with expert guidance from a wide range of private sector and government experts on Federal networking research needed in FY 2002-2006 and beyond to enable those visions of future networking.

1. Introduction

The Federal agencies of the Large Scale Networking Coordinating Group (LSN CG) held a March 12-14 **Workshop on New Visions for Large-Scale Networks: Research and Applications** to stimulate bold thinking and to explore new directions that could revolutionize networks and applications. The objectives of the workshop were to develop a vision of the networking needed to support societal transformations over the next 10 to 20 years, discuss networking capabilities needed to enable those visions, and identify networking research needs to provide those networking capabilities. Participants with expertise and interest in large-scale and advanced networking research from the academic, commercial, research laboratory, and Federal networking communities were invited. They provided their recommendations on networking capabilities and research needed to achieve their vision for networking and applications. This serves as guidance to the LSN agencies in developing their Federal agency IT R&D networking research programs for FY 2002-2006 and beyond.

1.1 Background

Under the Federal IT R&D Program, the LSN CG coordinates large-scale networking research programs among the six LSN agencies: DARPA, NSF, DOE, NIH, NASA, and NIST. In FY 2001 the LSN CG initiated a new phase of planning and coordination for agency advanced networking research programs. These research programs address LSN agency networking mission requirements and provide networking technologies for the future growth of the Internet. They also respond to the recommendations of the President's Information Technology Advisory Committee (PITAC) and other private sector inputs.

The PITAC was established February 1997 to provide expert independent guidance to the Federal government on maintaining America's preeminence in high performance computing and communications, information technology, and the Next Generation Internet. In its influential February 1999 report, "Information Technology Research: Investing in Our Future," the PITAC called for a visionary expansion of the Federal investment in information technology R&D. The Committee cited the need for fundamental research in networking to support societal transformations such as providing scaling, reliability, and growth so that at least one billion people can simultaneously access the Internet.

1.2 Workshop Scope and Purpose

In its FY2001 planning and coordinating of advanced networking research programs the LSN CG wanted to take advantage of the existing networking knowledge and expertise available from researchers in universities, industry, and government. The LSN CG decided to hold the Workshop on New Visions for Large Scale Networking: Research and Applications to enable leading researchers to provide input into the LSN planning activity on needed advanced networking research. The LSN CG issued an announcement of the workshop and a call for white papers (Appendix 1). In the white papers, researchers, applications developers, and network users applied their vision and creativity to help define broad research needs for networking and distributed applications. The white papers formed

the basis for presentations and panel discussions at the workshop. Most of these white papers are available on the Web at: http://www.ngi-supernet.org/conferences.html.

The workshop participants were asked to envision network technology possibilities that could revolutionize the way we live and work in the decades ahead, but that are beyond the scope of today's commercial, profit-driven R&D programs, and to identify network research needs to help realize those possibilities. Emphasis was placed on identifying exciting new areas of research that are radically forward-looking and that hold the potential to yield unexpected results. The workshop asked the participants to identify critical research barriers and the limitations of existing approaches. They were encouraged to bridge the gap between the broad-scale vision and the specific technologies, however difficult to achieve, needed to realize the vision.

1.3 Workshop Format

The LSN workshop was held over a three-day period from March 12 to 14, 2001. On the first day, three panel sessions provided discussion of the white papers in the areas of:

- **Adaptive networking**: Network-aware distributed applications, proactive self-tuning systems for ubiquitous computing, and custom channel building for large-scale network systems
- **Infrastructureless networking**: Ad-hoc disposable networks; dynamically forming, self-organizing hierarchy; and precision geo-location and ultra-wideband radios to support sensornets
- **Heterogeneous networking**: Heterogeneity of future network services; hierarchical addressing to simplify switching; and improving TCP/IP with features developed for User-Level Network Interfaces (ULNI)

The first day also included 10-minute talks on papers submitted by selected participants and the presentation of the six scenarios described below.

On the second day of the workshop, six breakout sessions were each tasked to consider one networking scenario in depth to identify the networking elements of that scenario and the networking research and applications needs to enable that scenario. The breakout sessions reported on their results on the third day.

2.0. Summaries of Workshop Scenarios

The workshop organizers developed six visionary scenarios of future capabilities that could transform major aspects of society and that depend on networking. The scenarios were designed to stimulate thinking about a wide range of potential uses of the Internet, the networking and applications capabilities that will be needed to support those visions, and the near-term to far-term research needed to realize those visions. The scenarios, which were developed and refined by the workshop participants in breakout sessions, were:

- Zero-Casualty War
- Deeply Networked World/SWARMS
- Crisis Management
- Collaboration
- Networked Medical Care
- Networked High Energy Physics

The following summaries focus on the networking capabilities needed to support the scenarios. The scenarios are presented in their entirety in the appendices.

2.1 Zero-Casualty War

A highly dynamic and intelligent battlefield surveillance and command and control capability has been developed, supported by a wide variety of on-line sensors and analysis, processing, and communications resources. This capability allows rapid deployment of forces armed with real-time intelligence to incisively respond to battlefield conditions with a minimum of casualties.

2.1.1 Scenario Description

A massive conflict has broken out between two Asian countries. Chemical and biological agents may be used. In preparation for a United Nations response, the U.S. has deployed networks of sensors in and around the conflict region. They include chemical sensors (vehicle exhaust fumes, urine, chemical agents, etc.), broad-spectrum acoustic sensors, seismic sensors, video sensors, and imaging sensors. Some are mobile.

A forward operating base is established in a neighboring neutral country to perform tactical reconnaissance. Terrain, street, and building information are updated based on visual and acoustic information from Unmanned Aerial Vehicles (UAVs). Signals Intelligence data are gathered from the UAVs and relayed to the forward operating base for analysis and correlation. Sensors provide acoustic, seismic, and visual signatures of each of hundreds to thousands of motorized vehicles that are cataloged and characterized. Air defense artillery and surface-to-air missile sites are identified. Automated analysis of visual information provides data about approximate numbers and locations of dismounted troops, enemy command posts, and command vehicles. Visual, chemical, and acoustical indications of weapons fire are all enunciated within the forward operating base, and video of that region is either initiated or tagged. Forward operating base data are relayed to the continental U.S.

(CONUS). CONUS or forward operating base personnel can initiate video streams and live sensor reports.

While en route aboard transport planes, Future Future Combat Systems (FFCS) and Air Cavalry units monitor vehicle tracking information and dismounted troop movements and develop their unit plans. Up-to-date terrain, street, building, and weather information is loaded into FFCS and Air Cavalry onboard databases via satellite from CONUS and the forward operating base.

The mission succeeds with minimal casualties as a result of:

- Fusion of data from large numbers of sensors
- Large-scale target identification and tracking
- Large-scale video acquisition, transmission, analysis, and directing of this information to appropriate command and control entities
- Remote command and control of robotic surface and UAV resources
- Rapid insertion of overwhelming force

A key component of the mission is the transmission of critical, sensitive information over reliable, secure networks. The networks need to be rapidly deployable and configurable to support command and control as well as tactical operations.

2.1.2 Zero-Casualty War Networking Research Needs

The discussion of this scenario identified the need for research in:

Scalable networking for large numbers of low-data-rate nodes

Future combat systems will have thousands to millions of nodes with very low throughput, high delays, and high redundancy. Networking needs to support scalability to large numbers of nodes with very low data rates by allowing redundant nodes, highly sub-optimal routes, high tolerance to losses and errors, and adaptation to changes. Research is needed on programming techniques for large-scale redundancy-based computing paradigms.

Network self-organization, automated configuration, reconfiguration, and management

Dynamic configuration and reconfiguration of networks are needed to support rapid initial deployment of sensors and their networks, changing conditions and locations, and mobile elements with asset tracking and handoffs. To support these capabilities a wide range of information is needed such as topography, sensor location, and user requirements. It also requires network capabilities such as:

- Tool sets for network design and deployment
- Performance measurement throughout the network
- Network discovery of applications and their requirements

- Self-diagnosing, self-healing capacity

Research is needed to automatically generate, propagate, and maintain the optimal communications, network, and application configurations required to rapidly establish and maintain mobile ad hoc tactical networks. To support crisis or conflict situations, network resources should be deployable and configurable in an operational state in the time required to physically transport those resources to their destination. These situations also would benefit from an ability to establish *virtual* configurations of networking assets that may include mobile field nodes and fixed end user sites. This capability must support decisions about frequency assignment (optimizing spatial reuse), application location, and network addressing.

Self-organizing networks have the potential to reduce the large manpower requirements to set up and configure networks. Research is needed to reduce the number of networking infrastructure personnel required to establish, operate, and maintain networks from 20 percent of a rapid insertion force to no more than 1 percent of the force. This will require networking capabilities such as:

- Automated diagnosis and fault isolation of mobile ad hoc networks
- Non-destructive automated network reconfiguration mechanisms to maintain system integrity and performance
- Mechanisms for network evolution, including interface definition and standardization
- Automated mechanisms for the diagnosis and correction of problems in mobile ad hoc networks

Hierarchical networking

CONUS, forward operating bases, Task Force commanders, and FFCS cell team leaders have access to common views of the tactical situation, but typically need different networking and aggregation to operate at different levels of the hierarchy.

Seamless, transparent service across heterogeneous elements

In a dynamic ad-hoc environment, networking will rely on heterogeneous technologies (wireless, satellite, land line) that must seamlessly and transparently work together to support the end users. Network-to-network interfaces must be interoperable. Standards are needed to support seamless and transparent service.

End-to-end performance

Applications, networking, and services must cooperate to satisfy the end-to-end needs of the user in a seamless, transparent, cost-effective, trustworthy, and timely manner. The network must be able to adapt to mobile and ad-hoc sensors and nodes, accommodate in-situ sensors and nodes, and provide access to widely distributed computational resources (for example processing, modeling, and data resources). A knowledge-based, rule-driven tool is

needed to tailor sensor performance to specific mission requirements and to tune the sensor array for deployment patterns, transmission frequencies, and power levels.

The networks need to support fusion of sensor data and to provide information tailored to end users requiring information at different levels of granularity – e.g., data covering a corps level or a battalion level. Sensor data may be aggregated in the field to minimize data transmission if the results will still meet end user requirements.

End-to-end performance measurement is critical to tuning end-to-end performance and trust.

Power management

Mobile networks will rely on finite power sources, usually batteries. It is critical that mobile network elements accurately measure and effectively use the power needed for sensing, processing, transmitting, and receiving information.

Trust: security, assurance, and reliability

Functions such as telemedicine, weapon's fire control, voice, and image transfers require high levels of end user trust. This trust will depend on end-to-end system reliability, security, responsiveness, and predictable performance. System responsiveness relies on channel access methods and end-to-end route generation. Predictable performance will require system and network redundancy and fault tolerance. Information must meet user and application requirements for trust by providing throughput, timeliness, fidelity, assurance, reliability, latency, location (e.g., for weapon's fire control), error, and other factors. QoS may address some of these factors.

Security must be provided throughout the system since each sensor and network node is subject to compromise. Differing levels of end user trust may accrue to different network paths, data aggregation from different sensors, cross-correlation of sensor data, and other system characteristics. Research is needed on decreasing information uncertainty through configuration and management of sensor and networking resources.

End user trust is dependent on establishing a common operational view and QoS. To achieve this, research is needed on:

- Techniques for capturing minimum mission requirements
- Adaptive middleware to map application-level requirements into network-level QoS mechanisms
- Network-level mechanisms (QoS and techniques) for resolving conflicting needs

Multimedia

Multimedia technologies will accommodate voice, data, video, and still images.

Revisiting networking fundamentals

Future systems will have to seamlessly integrate components with a dynamic range many orders of magnitude larger than today's networks (with speeds commonly ranging from hundreds of gigabits per second to a few kilobits per second) in a changing environment. Research is needed to revisit the network protocol stack to determine what types of control information are needed at each layer (including the application layer) to allow the other layers to effectively adapt to rapidly changing network conditions. Intermediate steps in this research include identifying characteristics of the potential links, interfaces, and component networks, and developing a control plane application/platform interface.

2.2 Deeply Networked World/SWARMS (Smart World Airforce Repair and Maintenance System)

In the future, networking and networked devices will be broadly and deeply deployed to make possible the truly *smart world* in which intelligent agents query, collate, and manage systems. An agent model will be developed and deployed, in which agents function correctly as individuals and collaborate with each other effectively in order to make higher-level decisions than any individual agent might make. Both individually and in groups, agents will provide higher-level, composite functions responsive to societal policy constraints that may change or evolve over time. Humans will not be responsible for managing the large numbers and heterogeneity of devices in such a deeply networked world. The network will be self-organizing and self-healing. This will require the ability to measure and evaluate its behavior, and either mask or correct problems when they arise.

2.2.1 Scenario Description

In the future, the Airforce has implemented an architecture called the Smart World Airforce Repair and Maintenance System, or SWARMS, where every repair component and parts depot is "smart." SWARMS predicts when and where specific repairs will be needed and, at the level of the whole aircraft, understands flight schedules and uses this information to plan where and when work should be done. SWARMS informs the global inventory system, which in turn makes sure that by the time a plane arrives at a destination the appropriate parts are there, with enough information that the repair can be made.

In SWARMS, a part knows where it is. A depot knows how many of which kinds of parts it has and has a model of what is needed based on reports about the schedule of arriving planes. Every part, when installed in an aircraft, is also introspective. Each one knows how well it is functioning and can predict when it will need to be repaired or replaced. The composite systems not only integrate over all their parts, but also have a higher level understanding of the emergent system to support mission planning.

In this scenario, there are critical issues relating to trust, assurance, and security, including privacy, authenticity, authorization, and denial of service. The system must maintain security to prevent exploitation by the enemy. For example, SWARMS data are of great use for espionage and sabotage by the opposing force, and maintenance schedules and procedures are critical to the safety of the aircraft.

2.2.2 SWARMS Networking and Networking Research Needs

The SWARMS discussion identified key networking elements of the scenario, including:

- Smart devices that monitor their location, query function and status, and identify situations requiring attention
- Sensors and systems considered at different levels of aggregation – e.g., as individual devices or as components of larger systems
- Automated functions implemented to evaluate conditions and to control behavior

- Multiple simultaneous agents acting collectively

The SWARMS discussion identified the research needed to support these elements, including:

Trustworthiness of complex self-organized networks

The end user must know that the end-to-end system can be trusted to meet requirements. Characteristics of a complex system contributing to end user trust in the system include reliability, robustness, and security. Technical capabilities for implementing these characteristics include digital signatures, authentication, authorization, path quality, information source, and quality of the information. The trustworthiness of a complex system is a function of the trustworthiness of its components and how they are integrated. It may change as sensors, networks, and other resources change over time. Since some network paths may be more trustworthy than others, the algorithms chosen to organize, select, and establish network paths contribute to the trustworthiness of the system. Research is needed to develop these algorithms.

Sensor data may have different "value" to an end user depending on the trust associated with the specific sensors that produced the data. Distributed sensor design characteristics, such as reliability and communications mechanisms, contribute to the end user trust in a sensor and its data. These characteristics are a consideration in the design and cost of producing the sensors. Trust will also be affected by the sensors chosen and the data paths used to obtain data from these sensors.

Adaptive distributed systems

Adaptivity may enable a greater functional range for a distributed system. In a bandwidth- and sensor-limited environment, the system can adapt the sensors chosen and the data they transmit to produce information tailored to specific levels of the decision and operational hierarchy. Several alternatives for adaptation exist:

- For simple network, the application may adapt
- For a simple smart network, the applications may adapt based on network-provided information, including initial information and operational feedback
- For a complex highly controllable network, the network may adapt based on information provided by applications
- For a complex implicitly adaptive network, the applications run and the network adapts

Adaptive networking depends on performance measurement and evaluation. Tools are needed for development, implementation, evaluation, and use of adaptation algorithms.

Scalability and self-organizing communications algorithms

We expect orders of magnitude increases in both the number of networked devices and network traffic on the future Internet. It is critical that the network scale to accommodate those increases. Research is needed to understand network behavior with these increases and to study networked systems' complexity. Research is needed to adapt relevant science from other fields such as chaos theory, economics, catastrophe theory, stochastic processes, and generalized control theory to promote breakthroughs and revolutionary solutions to scalability. Network performance measurement is critical to providing information on the functioning of the network to guide real-time network management, and to provide an understanding of network behavior to support design of the future Internet. Network performance measurement tools need to be developed, standardized, and ubiquitously deployed to provide performance data. A performance data archive is needed to provide an historical record for understanding operational network behavior, complexity, and trends and to support network simulation and design.

Currently, network scalability is implemented using hierarchical and cluster network organization. However, such organization is difficult to implement for mobile network elements, for responding to dynamic conditions, and for responding to administrative constraints. These require that self-organizing networks be able to continually change the network organization.

Research is needed to identify core networking functions and parameters, to develop algorithms that will enable highly flexible multimodal routing to support scaling and QoS, and to implement more flexible addressing schemes to accommodate emerging optical technologies. Routers need broader semantics for topology, name, attributes, and coordinates (grid location, hierarchy, etc.), and scaling for orders of magnitude expansion in numbers of networked devices and network traffic. Network simulation tools are needed to determine performance limits of a network to anticipate problems before they occur.

2.3 Crisis Management

A crisis management system has been developed that enables crisis response teams to respond with critical resources where needed in real time. It uses deployable and in-situ sensors that configure themselves to provide real-time monitoring of the environment. Networking supports reliable, dynamically configurable, and highly secure communications to enable real-time delivery of information to distributed decision-makers and real-time information development using remote, on-line resources.

2.3.1 Scenario Description

In 2015, “perfect” conditions exist for multiple fires in the U.S. DoD has deployed a "staring" missile launch detection satellite system approved for military and civilian use. NASA has orbited Firesat, capable of providing twice-a-day high-resolution data over multiple spectral bands (for higher informational content), and multi-instrument views of forest fire activity around the world.

A large collection of wireless devices with embedded chips has been deployed in cellular phones. They support “spotcasting,” ad hoc communications, a sensor mode, and general purpose programming and processing capabilities. Optical fiber is extensively deployed at the core of the network. A rich assortment of sensors is deployed in homes, commercial buildings, public infrastructure, and the natural environment. Computation, communications, and sensor resources are widely available over the Internet.

On the day of crisis, DoD reports dozens of new fires from a single highly charged lightning storm. By mid-day, the early morning Firesat images and data have been processed and disseminated to hundreds of state and Federal agencies. Hot-spot data are combined with vegetation cover and dryness models to produce detailed next-24-hour maps for the worst-hit regions. These maps and digital models are disseminated instantly to government command and control centers and are spotcast to the individual homes and businesses most in danger.

Department of Interior and other Federal agency supercomputing systems are organized on-line to model the existing situation and begin “nowcasting” the predicted tracks of the worst fires. The models and nowcasts are transmitted by satellite communications to the forest fire field units, which return validation and update information. This field information, along with real-time atmospheric, chemical, and other environmental data from sensors deployed throughout the area – both in-situ microsensor platforms deployed in advance as well as self-contained, self-powered sensors dropped from aircraft that same morning – are continuously integrated into the nowcast models. Customized warning and evacuation messages are automatically provided to all the homes and businesses in the area.

Emergency mobilization forces are directed by computer into the affected area. They establish a high performance field network instantly capable of local area and remote communications, using truck-based wireless technologies tied into regional networks via high performance satellite communications.

Telemedical facilities are established to attend to fire victims in the worst-hit areas. Mobile whole-body scanners, sophisticated medical instruments, and mini chemical analysis labs are plugged into the network. This allows deep resources of medical specialists, data and information resources, and analysis facilities to support on-site paramedics in real time. Command and control units have real-time high performance network access to all needed statewide and Federal resources.

2.3.2 Disaster Scenario Networking and Networking Research Needs

The disaster scenario discussion identified networking challenges including:

- Sensornet: An ad hoc network of sensors configured for and attached to the existing infrastructure. High bandwidth connections, e.g. gigabyte satellite to reach rural areas.
- Heterogeneous environment of sensors, networking capabilities, and administrative structures
- Dynamic environments and changing user requirements providing a need for new network management and visualization tools and automatic reconfiguration, management, and control
- Technology reuse: Using surviving resources for purposes other than the primary purpose they were designed for
- Data resources: Satellite sensors and deployed video sensors that produce data at the rate of hundreds of megabytes per second. These data are used in modeling and by command centers. Rapidly changing loads place emphasis on QoS based on media type (sensor data, voice, video) and user.
- Real-time modeling: Significant distributed computational and communications resources to support nowcasting

The disaster scenario discussion identified research needed to meet these challenges, including**:**

Interoperability

Organized sensors and networks will have to operate seamlessly with the existing infrastructure and with each other to overcome existing incompatibilities, routing mismatches, and security mismatches between different providers.

Robustness and dynamic reconfiguration

The infrastructure must be designed to cope with a wide variety of faults and dynamically changing resources by providing redundant resources and paths and the ability to actively reconfigure. Redundant technologies should be used so that their failure modes are as distinct as possible to decrease the probability of system failure.

Reuse of technologies

Reuse of wireless devices (including routing, spotcasting, ad hoc communication, sensing, and application software download) could help ensure that local resources are available during a disaster response. Reuse could also support functions needed to transform from short-term crisis management to longer-term emergency response.

Self-organizing, self-healing networks

Self-organizing, self-healing networks will expedite the organization of remaining and newly deployed sensors and technologies to establish routes and to connect to the existing infrastructure with minimal human intervention. The involvement and coordination of government agencies, companies, and individuals may require establishing a temporary administrative domain including components from the different organizations.

Dynamic, adaptive, time-varying QoS

In a crisis response, bandwidth resources may not match the workload and workloads may vary significantly over time and space. For example, time criticality and video quality requirements may vary depending on whether it supports telemedicine or media reporting. Thus, mechanisms are needed to deliver QoS within an ad hoc network that are appropriate to the application and network technology.

Discovering resources and their location

Establishing an ad hoc infrastructure for disaster response requires resource discovery such as identifying and locating available links and their capacities; information, computational, and other resources; and QoS capabilities to support priority information distribution and delivery of telemedical resources.

Trust: security, privacy, and reliability

Issues of trust, encompassing security, privacy, and reliability, pervade the disaster scenario. The disaster response resources must provide differing levels of security, assurance, and reliability based on the needs of the end users and their applications such as medical data transmission and patient privacy over heterogeneous, ad hoc networks and devices. Research needs to address:

- Heterogeneity of parties involved: A major disaster will involve many government agencies (local, state, and Federal), companies, and individuals. Disaster response networks must be responsive to their diverse security and trust policies that may contain incompatibilities and hinder sharing data and other resources. This issue can be further complicated if other sovereign nations are involved.
- Flexibility: Disaster responses may require temporary flexible modification or violation of security and trust policies. For example, an emergency medical team may need to access patient records for which it ordinarily would not have authorization.

- Reuse of technologies: Technologies may be designed so they can perform actions in crises that are not their primary functions. They also need to be designed so they are not then susceptible to third party invasion during normal times of operation using their crisis response capabilities.
- False alarms: Research should be conducted on detecting a false alarm by an intruder and being able to identify that intruder.

Network visualization and network management

Current network visualization and management tools are not able to handle the ad hoc heterogeneous networks needed for disaster response. New network monitoring and measurement tools are needed to support visualization and management.

Spectrum conflicts

Spectrum conflicts that arise whenever different technologies (for example, Medium Access Control (MAC) protocols and cellular standards) share the same portion of the spectrum will need to be overcome.

Metrics and performance

Metrics are needed to measure the time to set up a network and the amounts of traffic supported at different levels of QoS. Simulation and analysis tools are needed to deal with time dependent response problems and networks with many orders of magnitude difference in speeds from one part of the network to another. Solutions to the time dependent response problems should be evaluated in multiple ways including simulation using benchmarks. In addition, training exercises are needed to stress and test different solutions.

2.4 Collaboration

Our society increasingly relies on geographically distributed collaborations for human interactions in business, science, the arts, and other areas, both nationally and internationally. These collaborations improve communication among individuals with a common purpose; promote sharing, development, and dissemination of information; and foster interdisciplinary interactions.

The Internet supports distributed collaboration teams in which collaborators at multiple sites can interact visually and verbally, augmented by additional tools and services such as virtual reality and immersive environments. Distributed collaborations increasingly require realistic, "natural" interactions supported not only by high bandwidth but also by a wide array of enhancing services to provide ease of use, completeness of information, and appropriate levels of trust and assurance. Collaboration tools assist these groups in performing complex tasks, such as providing multimodal access to remote sites of scientific interest and supporting coordination to overcome problems and failures as they occur.

2.4.1 Collaboration Scenario

Current-generation network collaboration tools have not been widely used because they support limited exchange of information, provide limited "visibility" of remote collaborators, and are often difficult to set up and maintain. To be more widely accepted and used, collaboration environments (collaboratories) will need to provide automated setup and enhanced exchange of information and visualization. In addition they will need to be active, adapt to individual work patterns, and provide interoperable services over heterogeneous applications. Collaboration services, provided over a network, will need to adapt to end user personal preferences and work patterns, provide interoperable services over heterogeneous applications, and provide multimedia interactions, whiteboarding, and access to data and computational resources.

It is envisioned that future collaboration environments will make remote collaboration natural by being very easy to use, proactive, and engaging. They proactively will offload tasks from collaborators to enhance any given collaboration. For example, an automated "assistant" will retrieve information relevant to the collaboration. The system will configure the network and network services to support the particular individuals in the collaboration.

Future collaboration environments will need to deal with unpredictable and emergent changes, meet hard real-time constraints, and handle asynchronous events as they occur. Scientific collaborations will be particularly effective if scientific experiments can be conducted from local laboratories with all phases of the experiment coordinated so that it appears to be a local endeavor to each participant. The right level of coordination must be selected dynamically, depending on the task at hand and the feedback from the participants. Each participant should be able to augment the physical world with virtual worlds to consider "what-if" scenarios.

A collaborative scientific environment might include a large number of small sensors and robots with varying capabilities, capable of being embedded into the natural environment with minimum disturbance. Low-power nodes, with limited communication bandwidth, need to understand local conditions and together collaborate to identify and monitor global environmental conditions. Network traffic loads may be reduced if data and information can be aggregated and correlated at a local site to the level of granularity required by the collaborators.

2.4.2 Collaboration Scenario Networking Needs That Require Networking Research

To support collaboration environments in the future, networks will have to provide:

New middleware services

- Transparency among the collaborators, by accommodating heterogeneous technologies and interfaces and asynchronous events among end users
- User trust including security and reliability
- A virtual whiteboard
- The ability to convey body language (e.g., eye contact), visual cues, haptic, and olfactory information
- Ability to accommodate cultural differences among collaborators
- An automated "scribe" to record the collaboration including intrinsic information such as voice inflections and body language

Ubiquitous access

End users will have a wide range of technologies for accessing the network, depending on the technologies available to them (such as wireless and optical access) where they are located. The system must automatically implement the system interface for the different access technologies.

Intelligent collaborations

The collaboration system must be able to support the varying capabilities of different collaborators' systems. This will require automated mechanisms to:

- Correctly identify every collaborator
- Retrieve collaborator preferences and permissions
- Detect each collaborator's network speed, system protocols, and system capabilities – e.g., different modes of operation such as a Personal Digital Assistant (PDA) versus a dedicated multimedia facility
- Configure the system to support the above capabilities

This system will need to meet the needs of widely varying collaborative groups from small to large, across diverse disciplines, and operating in differing environments that could

range from a laboratory supported by an array of technologies and high bandwidth connectivity to a remote field site with more limited technologies and wireless connectivity.

In addition, the system needs to be able to extract real-time data from the scientific instruments, computing and data resources, and human collaboration as it happens. It also needs to extract contextual data (e.g. voice pitch and intensity) and bring additional relevant information into the collaboration for use, reference, and/or citation. Using advanced pattern recognition techniques and artificial intelligence agents to mine this complementary information as the collaboration happens will help enhance the collaboration. These capabilities contribute to ease of use, thereby encouraging system adoption, and they provide additional information on the collaboration for establishing an historic record.

Virtual environments and coupling issues

A scientific collaboratory needs to support scientists performing both virtual simulations and physical experiments. It should provide seamless support for interaction with science discipline models, virtual reality environments, and on-line databases. It must also address issues of coupling that can take several different forms. For example, two large-scale simulations being used by different groups of collaborators may operate in the same physical space but use different physical units such as meters versus millimeters. Other collaborations may use multiple virtual environments or virtual environments interfacing with physical experiments. Multiple databases accessed by different groups of collaborators may need to present data in different physical units or differing formats are a third example.

Multiple modality expert consultation

Collaborations need to be able to include ad hoc consultation with experts wherever they are located using interaction technologies ranging from PDAs to sophisticated collaboratory environments.

Network measurement

Instrumentation for performance measurement should be provided throughout the network to measure network performance and enable isolation of system faults. It needs to be implemented for multiple link types (e.g., optical, electronic, wireless) and measure end to end performance. Network measurements should be standardized across network providers to provide this end-to-end capability. Performance measurement should also support network reconfiguration for active networking.

Automated supervisory oversight

The system must support supervisory oversight for monitoring performance and identifying problems. It must ensure that standards are adhered to when carrying out experiments that have safety and/or environmental implications and/or when performing experiment-critical functions. The system needs to be able to warn participants when requirements for critical functions are not met.

Virtual meeting maker

The system must be able to schedule, establish, and record virtual meetings. It must support schedule conflict resolution, scribing, attendance authentication, and archiving. In addition, the system must support asynchronous access and coordination for meeting absentees who access archived meeting materials.

Security and privacy

Security and privacy tools must be able to handle a wide range of requirements such as authorization, end-to-end key management, and revocation of authorization. In addition, the system must be able to support advanced security features such as allowing selective anonymous collaborators to participate and retrospectively access archived collaboration materials.

Other features

The system will also need to support a variety of additional capabilities such as shared and private workspaces, an "electronic whisper" capability that allows two collaborators to hold a private conversation during a collaboration session, and language translation.

2.5 Networked Medical Care

In the future, networks will support expert medical care, including surgery, delivered to patients in remote and mobile locations on line, in real time, and collaboratively in a highly secure, intelligent, dynamic, and reliable environment. Additionally, doctors will access distributed medical records and medical expertise wherever it is located.

2.5.1 Medical Scenario Description

A middle-aged man at home begins to suffer chest pains. He uses a medical sensor to take automated medical readings that are relayed to a medical center that determines he is having a heart attack. In an ambulance dispatched to take him to the hospital, sensors monitor his vital signs and cardiac function. A remote cardiologist monitors these data and accesses the patient's medical records. She orders an angiogram to be taken when the patient reaches the hospital. The angiogram shows a possible anomaly and a remote consultant is shown the angiogram over the Internet. The angiogram displays a warning that the resolution of the image, as delivered by the Internet, does not meet the standard required for angiogram interpretation.

Heart surgery is performed on the patient at the hospital. An anomalous cardiac vasculature found in the patient leads the surgeon to consult an on-line 3-D anatomical library in real time. The library finds a consultant surgeon who, also in real time, assists in the operation, occasionally taking control of the haptic surgical robot.

2.5.2 Networked Medical Care Research Needs

Ubiquity

Multimegabit-per-second effective wireless bandwidth from multiple sources is needed. Bandwidth available during the ambulance ride may occasionally be degraded so the network should be able to identify networking alternatives, choose the best alternative for the application, and reconfigure the network.

Trustworthiness

Trustworthiness has many components that collectively assure the end users of the quality, timeliness, security, and reliability of the services provided by the network:

- Security and data integrity: First and foremost, the network must meet legal standards for medical data privacy and security as currently documented in the Health Insurance Portability and Accountability Act (HIPAA). This requires the networks to support authentication of the patient and the end user, authorization for end users, encryption to support privacy requirements, and traffic diversity to prevent identification of restricted information through traffic analysis. Authorization and access should be logged to provide an historical security record. Data security is required for restricted data and to assure data accuracy and integrity.

- Quality of Service: QoS is required to support the strict demands of distributed medical care delivery and collaboration. To support the cardiologist at a remote site, the wireless channel must provide real-time video and real-time data indicating the quality of the video display. Although this scenario may tolerate a fair amount of latency, it will not tolerate jitter and the video, audio, and data channels must be synchronized. Medical service must be provided across network service boundaries in a dynamic and sometimes mobile environment. All devices need to support QoS, and the system must be able to adapt in real time to networking or data content changes.

Sensors and end user devices

Networks must support dynamic sensors and end user devices that must be identifiable and locatable. Sessions may migrate from one device to another – for example, migrating from a fixed end station in a patient's home to PDAs in an ambulance requires networking services that support significantly different access interfaces.

Collaboration environments

The networks must support ad hoc establishment of collaboration sessions for specific access modes, locations, service needs, and networking capabilities. For example, one participant may need voice-only capability while others may need varying degrees of video, voice, and whiteboarding. The networks also need to support access to on-line resources such as distributed computing and database access to support the collaboration. Security, discussed above, is critical to collaboration environments.

Intelligent networking, end-to-end performance

The medical scenario angiogram procedure illustrates the need for end-to-end knowledge of the network data path including the end user display devices to assure that angiogram interpretation standards are met. Thus, an intelligent, scalable network needs to be able to reconfigure itself and automatically resolve any QoS problem to meet medical standards. The network must report any unresolvable problem to the participants.

Assured real-time service

For the bypass surgery scenario, the surgeon needs to retrieve 3-D image data sets, each of which may be several gigabytes in size. The consultant must be able to view the surgery in real time and accurately guide the surgical robot using its haptic controls. This requires a network operating at high bandwidth with minimal latency and minimal jitter while maintaining the security and integrity of the transmissions and the privacy of the patient data.

2.6 High-Energy Physics

High-energy physics (HEP) has pushed against the limits of networking and computing technologies for decades. Twenty years ago, the largest HEP experiment involved 100 physicists from many nations and acquired tens of thousands of magnetic tapes of data per year; graduate students spent months reading those tapes to perform data queries. Life is not so different for today's physicists. The new BaBar detector at the Stanford Linear Accelerator (SLAC) was designed by a large international collaboration of physicists at 72 institutions. The BaBar collaboration enables hundreds of physicists worldwide to query its 300-terabyte and rapidly growing database in hours or days rather than months. In the next 10 years, the Large Hadron Collider (LHC) experiments at CERN, the European Physics Laboratory, where some 600 U.S. physicists form the largest national group, will face the challenge of distributed analysis of hundreds of petabytes of data.

2.6.1 High-Energy Physics Scenario

The physics community greatly values being able to distribute digitized data electronically at the rate at which it is produced from the site of an experiment to collaborators worldwide who can analyze them. The HEP community has the goal of using affordable network and computational resources to provide physicists with transparent access to a distributed data-analysis system that uses all available resources as efficiently as possible. By 2005 to 2010, HEP computing will involve queries on databases containing exabytes (10^{18} bytes) of data structured as up to10^{16} individually addressable objects. These massive amounts of data will require the distribution of terabits per second of real-time data to major HEP data analysis centers.

Challenging networking and other information technology research needed to enable distribution of data, analysis, and collaboration includes:

- Multicast service delivered to multiple remote centers with diverse firewall filters
- Network error rate and robustness control *without* impacting the experiment's data-acquisition system
- Massive applications software – e.g., 3 million lines of BaBar C++ code
- Commercial object database management software
- Interfaces of the database with the network and storage
- Technology improvements including:
  - Computing technologies
  - Computer science
  - Networking
  - Computing system-to-network interfaces
  - Fiber technologies
  - Data storage

Improvements in HEP applications must be accomplished at minimal incremental costs. To help contain costs, network engineering labor, required to configure, optimize, and

maintain networks, should be minimized by developing automated network engineering and management.

HEP collaborations are increasingly international in composition. It is difficult to adopt standards across the resulting international boundaries, so that the implementation of uniform, collaboration-wide middleware, security, or hardware technologies is almost always unrealistic. The best that can be achieved is the adoption of a set of protocols and interfaces to link components that will almost certainly be implemented in different ways.

The international HEP research community is increasingly using Grid technologies, an integrated suite of services developed with Federal IT R&D funding. The Grid is a set of middleware tools and capabilities that enable seamless end user access to applications, data storage, and compute resources to support high-end modeling. The Globus project (http://www.globus.org/) is one state-of-the-art example of Grid development. Grid middleware faces many hard computer science problems. Vertical integration of existing components to provide Grid services to demanding, well-defined communities is essential to progress on Grid architecture and technologies.

2.6.2 High-Energy Physics Networking and Networking Research Needs

Networking underlies many of the services and applications being developed to support HEP. Progress in networking is expected to be evolutionary over the next five years, with revolutionary capabilities being developed over the longer term. The following table presents the current state of the art in various networking areas supporting HEP, what could evolve by around 2006, and the requirements to approach meeting the HEP goals. The current HEP capabilities are what is affordable, not what could be obtained with unlimited funding.

Current HEP Capabilities	Evolution to 2006	HEP Goals
Links Between Major Centers		
• 1 or 2 x 155 Mbps	• 10 Gbps	• 1 Tbps
Bulk Transfer Protocol		
• TCP/IP + fixes	• TCP/IP + more fixes	• New, *widely adopted,* transport protocol
Differentiated Services (CoS, QoS, Mixture of Packet and Circuit Switching, etc.)		
• Provide ~1.1 differentiated services (best effort + some Voice over IP (VOIP))	• Provide ~2 differentiated services	• Provide ~6 differentiated services that are application-negotiated, on-demand, and responsive to cost and policy
Network Measurement, Analysis, Interpretation, and Action and Network Modeling		
• Limited measurement, analysis, and modeling	• More/better measurement and analysis, and some interpretation • Models begin to predict non-obvious failure modes	• Automated measurement, analysis, and interpretation • Automated action based on measured and modeled information
Support for Collaboration		
• Some proof-of-concept (PoC) prototypes • Some commercial tools	• New PoC prototypes • Some mature components • Still incomplete	• Collaborations form via the Internet • Real sense of working together
Data-Grid: Authentication and Authorization		
• Local and manual	• Cross-authentication via proxies	• New approaches to regulating access to resources
Data-Grid: Information Infrastructure (Replica Catalog, Resource Catalog, Software Catalog, Operation/Task Catalog, etc.)		
• Manual and local • Limited *ad hoc* automation	• Evolution of Globus by 2+ generations	• Efficient distributed information management for more than 10^{16} virtual objects using millions of operations each using millions of lines of code (MLoC)
Data-Grid: Data Payload Infrastructure (Exabyte Databases, Reliable Replication, Storage Management, etc.)		
• Few x 100-Tbyte databases • PoC replication prototypes • PoC storage management	• Bleeding-edge 10^{19} Byte databases • Grid replica management • Grid storage management	• Industry-standard exabyte databases, replication, and storage management
Data-Grid: Resource Discovery		
• Telephone, e-mail	• Telephone, e-mail, partial automation	• Automated discovery • Standardized information models
Data-Grid: Distributed Resource Management, Distributed Job (Task, Operation) Management		
• Local batch systems • Prototype systems	• Grid job management • Early distributed resource management	• New approaches to regulating access to resources
Data-Grid: Virtual Data		
• Conceptual phase	• Starting to work for cutting-edge HEP experiments	• A generally accepted and implemented paradigm
The Grid an Integrated "Network" Service		
• Manually integrated services have been in use for more than 10 years	• Vertical integration of fabric and data payload services • Incomplete information services • Incomplete resource management services	• Easy creation of vertically integrated, worldwide information management and processing systems from standard industry components

Notes:

1. HEP technologies that work well locally but do not become widely adopted and supported may inhibit collaboration and prove costly. Qualifiers like "widely adopted, industry-standard," and "generally accepted" are vitally important.
2. Elegant approaches to authentication and authorization appear to be available for organizations that are part of a single administrative structure. Worldwide collaborations seem unlikely ever to fit this model. Discussion identified that a totally new approach to regulating access to resources might foster more open scientific research.

The Role of Industry in HEP Networking R&D

Wherever possible, high-end science takes advantage of capabilities that are developed and commercialized by industry. For example, the HEP community has benefited from cost reductions and reliability increases provided by industrial commercialization of individual middleware components, such as databases and well-defined information systems. Also, the HEP community has benefited from the availability of commercial high-end computing systems, high-bandwidth networks, and extensive middleware. It is likely that higher bandwidth will be more affordable in the future due to economies of scale, greater supply, and competition among providers. Carriers are beginning to make individual wavelengths available to major customers. Affordable links between major HEP computer centers should exceed 10 Gbps within five years and may approach 1 Tbps in less than a decade. However, it is likely to be difficult to exploit the available bandwidth using industry-standard transport protocols. TCP/IP requires fixes such as multiple streams to use today's affordable bandwidth. Additional fixes will be needed to accommodate the expected increases in numbers of users, number of nodes, and network traffic. It is possible to develop a new protocol or to extend TCP to work over dedicated links, but the extensive investment of industry and users in the current protocols would likely hinder acceptance of alternatives.

Workshop participants identified a need for a vertically integrated HEP solution for managing and processing the massive amounts of data expected from HEP experiments. Networking research, development of faster computing systems and more capable computational algorithms, and commercial development and marketing (productization) together deliver components that provide part of this vertically integrated HEP solution. New component technologies emerging from networking research and computer science are funded normally only to the proof-of-concept stage and fall short of the level of product hardening and support needed to provide technologies that can be reliably integrated into a complex operational system. Collaboration by network researchers, computer scientists, and application scientists required to provide vertical integration of the component capabilities are also research and development and, in the view of the workshop participants, should be funded by the Federal IT R&D funding agencies.

The HEP community is rapidly taking advantage of the Grid infrastructure to enable transparent, distributed, and international collaborations, resulting in improvements in the ability to cooperatively carry out science and to analyze increasingly large volumes of HEP data. However, the Grid primarily has been developed in universities and industry is currently largely decoupled from development of an integrated Grid capability. Thus, Grid software and infrastructure have not benefited from the standardization, cost reductions, and increased reliability often provided by commercial productization. This productization will take place only if industry perceives the potential for profitably marketing the technologies. Federal funding could help bridge the gap between the proof-of-concept prototype and the point at which successful vertical integration has demonstrated commercial viability.

Section 3.0 Summary of Networking Research Needs

This section summarizes the networking research needs identified in the Section 2.0 scenarios organized by research categories.

3.1 Adaptive, Dynamic, and Smart Networking

Most of the breakout sessions identified a need for research into elements of adaptive and dynamic networking to support ad hoc and mobile wireless access. The discussions of zero-casualty war and the medical applications scenarios identified the need to dynamically respond to developing situations with ad hoc, high-assurance networks supporting secure multimedia capabilities. In these scenarios, not only are the situations changing but also the participants in the networking sessions are changing with resulting changes in service requirements such as networking services, security levels, and end user devices to be supported.

Ad hoc networking to support deployable sensors for on-site chemical or temperature monitoring will require knowing the locations of the sensors and organizing them into a network capable of meeting requirements such as cost, location precision, measurement capability, power, and networking capabilities. Research is needed on self-configuration, connectivity to existing infrastructure, organization, and adaptation. Tradeoffs among functionality, performance, and cost will need to be understood and managed. For example, data aggregation and compression within the sensors may reduce communication requirements but increase sensor costs. Aggregation and compression may also affect the quality of the information sent to the monitors since such processing often changes the informational content and precision.

Future networks will be orders of magnitude more complex than current networks and must be able to respond to changing environments and dynamic networking as sensor elements are added and deleted. For example, large sensor arrays will be subject to sensor attrition that will require adaptive, dynamic, and smart networking to maximize the effectiveness of the remaining sensors. Engineering and managing these networks increasingly will require incorporating smart elements to automatically respond to the changing elements and environments. Research should also address the dynamic trustworthiness of the system and the information it is producing as the sensors and network change. Network measurement is fundamental to determining the status of networking elements to provide a basis for smart networking.

Smart networking research is needed for:

- Enabling sensors, networks, and applications to work together to increase the range of data granularity the system responds to and reports; for example, applications may vary significantly in the precision of a specific data parameter they require, thereby allowing data, system, or cost tradeoffs
- Automatically managing networks of increasing complexity including self-organizing, self-diagnosing, and self-healing networks

- Anticipating and automatically responding to network instabilities
- Network-aware applications that automatically respond to available networking resources
- Application-aware networks that automatically reconfigure networks to improve applications support
- Adaptive distributed systems: Applications may adapt based on network-provided information and system feedback, or the network may adapt based on information provided by the applications, as in implicitly adaptive networks.

3.2 Measurement, Modeling, Simulation, and Scalability

The Internet has expanded at a phenomenal rate, often driven by the need for increased capacity and capabilities to support new "killer applications" that in the past have included TCP/IP, e-mail, the Web, and Web browsers. With the continued evolution of Internet-based applications, types of media transmitted (for example large images and video), increasing connectivity of embedded devices, and increased support for arrays of sensors, the Internet over the next 15 years is expected to grow by many orders of magnitude in the number of nodes connected, amounts of information passed, and the number of users and their usage. Revolutionary new applications barely foreseen today are expected to lead to even faster expansion of the Internet and demand for Internet services. Instabilities may appear because existing Internet technologies and their evolutionary extensions could be severely strained to cope with this growth. Several of the breakout sessions discussed the need for research to address the growth, scaling, and stability of the Internet.

Network measurement

Each of these breakout sessions discussed the need for metrics and measurement of network performance. We do not have standardized technologies for measuring end-to-end Internet performance, let alone standardized reporting of measurements for most Internet nodes. This is a fundamental requirement for identifying current and developing bottlenecks and instabilities and for measuring improvements in performance as new capabilities are incorporated into the Internet. In the recent past, researchers have consistently observed that increases in network link bandwidth do not translate into proportional increases in end-to-end throughput for their applications. Measurement is imperative to study the causes of such behavior and to support engineering and management of the network links to improve performance for the end user.

Measurement research needs include:

- Intrinsic instrumentation: Make measurement a fundamental part of all systems on the network
- Extensible Application Platform Interfaces (APIs): APIs must provide measurement details to support network engineering to improve the end-to-end performance of the application
- Data reduction, formatting, and storage: Network measurement data collected in this environment must be reduced and stored in a format useful to end users. This

requires not only efficient data storage but also data synthesis, analysis, and formatting capabilities.

- National networking measurement archives: Provide a national archive to store network measurement data on a permanent basis. (Individual companies do not have the incentive to record and archive these data. Individual researchers do not have the resources required to provide the long-term archival storage.) A wide range of information should be stored, since we cannot now know what will later prove to be important.
- Ubiquitous inter-domain cooperation: Separate administrative authorities must agree upon a common set of measurement data that will be made available outside of their specific domains
- Correlation of measurements across levels: Data collected at the network connectivity, routing, end-to-end, and application levels must be correlated to provide a complete understanding of the system, to support network modeling and to provide a basis for network management.
- Synchronization of measurements: Measurements made at different levels and in different logical areas of the network must be time-synchronized to provide an instantaneous snapshot of network status and performance. The times when measurements are initiated must be synchronized and the ways timestamps are applied must be standardized.
- Modeling support: The measurement technologies must support network modeling to predict network failure modes, carry out network design and development, and enable network management.
- Security: The measurement system must support threat-evaluation models used to configure the network to withstand a wide range of possible attacks.
- Privacy: Measurement data must be securely transmitted to assure the privacy of individuals and administrative entities.
- New link types: The measurement and monitoring mechanisms will need to be adaptable to new link technologies such as optical networking.

Network modeling and simulation

Network modeling is needed to support research on network behavior and network management as the Internet grows in magnitude and complexity. Modeling is also needed to understand current network behavior and predict future behavior for assessing how new technologies will affect the stability of the Internet as they are introduced.

Network scalability

The network is expected to grow potentially by orders of magnitude, in the numbers of nodes, the amount of information, and the management complexity of the Internet. Current network architectures do not scale to handle these increases. A “science” of networked systems modeling is needed to understand how the Internet is likely to fail under increased loads, to fix potential problems before they occur, and to develop scalable architectures.

3.3 Trust: Security, Privacy, and Reliability

The Internet will be used for commercial, medical, scientific, and other uses only if users trust its security, privacy, and reliability. All of the workshop's scenarios inherently relied on this user trust. With the projected expansion of the Internet and the applications and media it carries, current issues in developing trust relationships will become more pervasive. The medical scenario relied on data access, consultations, and real-time collaborations with high assurance, security, and privacy. The disaster scenario identified the need to use distributed data and computing resources in near real time to support modeling and prediction and to support field units. The SWARMS scenario required security to protect against intrusion or espionage.

In some scenarios, it may be possible to quantify elements of security, privacy, and reliability associated with network elements or network links. Under these circumstances, "chains of trust" may be developed such that a user can choose networking paths based on highest overall trust or use information from specific network nodes that provide the highest confidence in the end product.

Trust may vary over time. Corroborating information may increase our confidence in information from some nodes or the networking architecture may change to a more reliable or secure pathway.

Security, privacy, and reliability research needs include:

- Quality of Service for critical applications in a complex environment that includes multiple providers, mobile and distributed access, and multimedia service (for example, for collaborations)
- Security, privacy, and reliability in dynamic, complex, and heterogeneous systems
- Scalability to accommodate heterogeneous environments and changing needs and hierarchies
- Trust modeling, configuring for trust, and responding to changing trust over time
- Trust retractability

3.4 Networking Applications

Each of the workshop scenarios relies on multiple networking applications. The workshop participants stressed the importance of developing these applications. Some of the applications identified in the scenarios that requiring networking research include:

- Telemedical remote collaboration with high assurance and security
- Sensornet: Self-organizing, dynamic, heterogeneous networks of sensors with network connectivity to remote resources
- Collaboratories: Support for interactions that are natural, intelligent, and secure, with multimedia capabilities and automated configuration
- Grid

- Hierarchical data delivery: Automatically develop and deliver data tailored to differing levels of an hierarchy

The workshop participants indicated that some of the potentially largest uses of the Internet will be for revolutionary applications not yet developed.

3.5 Networking Middleware

Networking provides connectivity among sensors, applications, end users, and distributed resources such as data repositories and computing facilities. Middleware assures that these elements work within a coordinated, transparent, and synchronized framework to provide end user services. Middleware can, for example, provide transparency among network service providers to seamlessly and securely transport information. Most of the workshop breakout sessions addressed the need to develop new middleware capabilities for networking.

The needs for enabling the Grid application illustrate many of the middleware networking needs, including:

- Vertically integrated, transparent, worldwide infrastructure for managing data and information, distributed storage, and access to computational resources
- Automated discovery of resources

Additional middleware needs identified in the scenarios include:

- Automated collaboratory setup, services, and toolsets
- Seamless, transparent service across heterogeneous network elements
- Control and management of dynamic networks
- Management of security, privacy, and reliability in a dynamic environment
- Automated measurement
- Productization to harden and standardize software by commercial developers

3.6 Testbeds

Workshop breakout session participants cited the need for testbeds to support networking research in performance measurement, security, privacy, reliability, active networking, adaptive mobile networks, intelligent networking, applications, and middleware. They also discussed the need for testbeds to bridge the transition from the research stage to successful commercialization of the technologies. An example is a Grid for high-energy physics research. Industrial participation in testbeds is often needed to develop and refine standards and to promote technology transfer.

3.7 Collaboration Environments

Networks support human interactions including human-to-human interactions such as collaborations and human-to-machine interactions such as access to distributed resources. Most of the breakout groups discussed the need for collaborations to be as good as face-to-

face meetings or to have enhanced capabilities such as immersive environments or automatic translations, for example for international collaborations.

Some collaboration environment capabilities that the Internet should support include:

- Ubiquitous access with a plug-and-play capability
- Automatic configuration to accommodate the personal preferences and characteristics of the participants and the heterogeneity of their environments (including extreme differences such as PDA versus CAVE environments)
- Authentication, authorization, security, privacy, and access control
- Resource-sharing with remote collaborators
- Natural and intuitive interactions supported by virtual, immersive, and integrated environments that provide body language, visual, audio, textual, haptic, and olfactory capabilities
- Language translation
- Large-scale on-line virtual and physical models
- Expert consultation
- Whisper mode (support for side conversations)
- Automated supervisory oversight

3.8 Revolutionary Research

This report has identified many research areas that are important in assuring the future growth, functionality, robustness, and usability of the Internet. Evolutionary networking research is expected to result in improvements in these areas. However, high-risk revolutionary research may provide unexpected dramatic improvements that accelerate the capacity to meet the growing networking needs. Revolutionary research comes from revolutionary visions of research groups or individuals, adaptation of research from widely different disciplines, interdisciplinary collaborations, and other research initiatives.

Some areas of networking are in need of revolutionary research. For example, revolutionary research is needed to address the scalability issues that will be increasingly critical with the projected orders of magnitude increases in the number of network nodes, network users, and network traffic. Revolutionary research is needed to understand network behavior with these order of magnitude increases and to study networked systems' complexity. Disciplines such as chaos theory, economics, catastrophe theory, stochastic processes, and generalized control theory may contribute to these complexity studies.

3.9 Revisit Networking Fundamentals

The Internet is based on fundamental concepts, technologies, and standards such as the TCP/IP protocol that were developed and implemented decades ago. These standards and technologies have provided a robust infrastructure for the phenomenal Internet growth we have experienced since then, and which was not foreseen when they were developed. They may not be able to meet the still growing demands. For example, the Internet growth may

exceed their ability to scale. Or some new technologies or a new protocol may provide greater efficiency, robustness, or ability to scale.

In 10 to 15 years, the core backbone network could well be Dense Wave Division Multiplexing (DWDM) optical with thousands of wavelengths per fiber. Broadly deployed network access will certainly be heterogeneous, incorporating broadband wireless, satellite, broadband wireline, and optical fiber with multiple wavelengths. Protocols and their associated services need to extend across these heterogeneous technologies with end-to-end transparent functionality. Fundamental changes may be required in addressing, routing, forwarding, and transport to support this increased scale and functionality. Revolutionary research, such as revisiting TCP/IP, can address basic issues of protocols, performance, complexity, and scalability.

4.0 Federal Networking Research is Needed: Industry Isn't Going to Do It

Federal networking research is needed to enable the scenarios described in this report and to provide for the long-term growth and viability of the Internet. Industry is focused on developing the commercial technologies required for the near-term (one to three years) growth of the Internet. Industry is not focused on the research needed for scalability, management, and improved services in networking required for the expected long-term Internet growth. The trend is for industry to conduct even less networking research. Indeed, many industrial research facilities, including Lucent Laboratories and Xerox Parc, are slated to be shut down by their parent organizations.

Current commercial networking technologies are largely based on technologies developed under long-term research funded by Federal agencies. Without continuing Federal funding of basic research in new networking technologies, such as optical networking, scalable protocols, active networking, dynamic networking, and intelligent networking, the pipeline of basic networking technologies needed to support further expansion of the Internet economy and the scenarios described in this report will not be available.

The Federal agencies engaged in networking research have an impressive record of supporting basic networking research and transitioning the results of this research to the industrial sector. Many networking technologies and applications are developed, refined, and tested in Federally funded testbeds with the active participation of industrial partners, who often contribute equipment, labor, and other resources. This Federal/industry partnership leverages Federal research funding and provides direct commercial experience in the developing technologies, thereby hastening technology transfer. This partnership is expected to continue for the research identified in this report.

Appendices

1. Invitation to submit white papers, January 5, 2001
2. List of LSN workshop attendees

Scenarios

3. Zero Casualty War Scenario
4. Deeply Networked World/SWARMS scenario (Smart World Airforce Repair and Maintenance System) Scenario
5. Crisis Management Scenario
6. Collaboration Scenario
7. Networked Medical Care Scenario
8. High Energy Physics Scenario

Appendix 1: Invitation to Submit White Papers, January 5, 2001

Interagency Working Group for Information Technology Research and Development (ITRD)

Large Scale Networking (LSN) Coordinating Group

Call for White Papers

Workshop on New Visions for Large-Scale Networks: Research and Applications

March 12 - 14, 2001
Vienna, VA
Paper Submission Deadline February 4, 2001

The phenomenal worldwide explosion in global networks and widespread dissemination of advanced technology is rooted in thirty years of patient investment by federal R&D agencies. As commercial investment has poured into this area, the research focus has naturally responded by seeking solutions to near-term problems. The goal of this workshop is to stimulate bold thinking that will take us off this evolutionary path and to explore new directions that could revolutionize future networks and applications.

Researchers from related disciplines are invited to share their perspectives in helping to define a broad research agenda for the future of networking and distributed applications. The goal will be to envision and identify networking technology needs and possibilities that would revolutionize the way we live and work in the decades ahead, but that are out of scope for today's profit-driven R&D programs. The academic, industrial and governmental research communities are invited to submit white papers that describe radically new visions for the future, as well as possible steps to realization. Accepted papers will form the basis for panel discussions and presentations, and will also be used to inform development of long-term research programs within the sponsoring organizations.

Of particular interest are submissions that elucidate an exciting new area of research that is radically forward-looking and that holds the potential to yield unexpected results. For any given focus area, quantitative explication of critical research barriers and limitations of existing approaches are required. Authors are encouraged to bridge the gap between the broad scale vision and the specific technologies, however difficult to achieve, needed for ultimate realization of the vision. For example, one may include specific research findings (original or cited) that may indicate expansive future successes or one may delineate assumptions that are made with regard to advances in supporting or enabling technologies. Alternatively, authors may also choose to structure the paper in the form of a proposal abstract to be submitted in the year 2005 or beyond. Such alternative submission should include a description of deliverables (hardware, software, system prototypes, and algorithms) that would be produced three to four years from the submission date.

The total length of the paper should not exceed six pages, including any figures, with minimum font size of 10 points. The first page (the cover sheet) must show the submission title, names and contact information for the author(s) and/or a contact person. The cover sheet should also include an abstract that succinctly describes the main idea, innovative claims and the critical technical barriers. Submissions must be formatted in Microsoft Word or Adobe PDF format.

Attendance will be by invitation; some limited support for travel and expenses will be available for invitees. Papers must be submitted electronically to lsn_workshop@snap.org by 5pm Sunday, February 4, 2001. Those selected to participate will be notified by February 16, 2001. Information related to this workshop will be posted at: http://www.eventmakeronline.com/sta/view/index.asp?meetingid=5. For further information, please send email to lsn_workshop@snap.org.

The workshop is sponsored by the Federal Large Scale Networking Coordinating Group, DARPA, DOE, NASA, NIST, NLM and NSF.

Appendix 2: List of LSN Workshop Attendees

Last Name	First Name	Organization	Email
Ackerman	Michael	NIH/NLM	ackerman@nlm.nih.gov
Agarwal	Deb	Lawrence Berkeley Laboratory	daagarwal@lbl.gov
Agarwal	Sharad	University of California - Berkeley	sagarwal@eecs.berkeley.edu
Agrawala	Ashok	University of Maryland	agrawala@cs.umd.edu
Ahmed	Mohin	HRL Laboratories	mohin@hrl.com
Almes	Guy	Internet2	almes@advanced.org
Aronson	Jules	NIH/NLM	aronson@nlm.nih.gov
Banerjee	Suman	University of Maryland	suman@cs.umd.edu
Barford	Paul	University of Wisconsin	pb@cs.wisc.edu
Bauer	Steve	Massachusetts Institute of Technology	bauer@lcs.mit.edu
Bernholz	David	NCO/ITR&D	bernholz@itrd.gov
Bhasin	Kul	NASA Glenn Research Center	kbhasin@grc.nasa.gov
Blumenthal	Marjory	Computer Science & Telecommunications Board	mblument@nas.edu
Boroumand	Javad	Cisco Systems	jborouma@cisco.com
Bradaric	Ivan	Drexel University	ivan.bradaric@drexel.edu
Braden	Bob	University of Southern California – Information Sciences Institute	braden@isi.edu
Brett	George	NLANR / Web100	ghb@ncsa.uiuc.edu
Brown	Bruce	Institute for Defense Analyses	bkbrown@ida.org
Burns	Catherine	University of Waterloo	c4burns@engmail.uwaterloo.ca
Bush	Aubrey	NSF	abush@nsf.gov
Bush	Stephen	GE Corporate R & D	bushsf@crd.ge.com
Calvin	Jim	MIT – Lincoln Laboratory	jcalvin@ll.mit.edu
Carlson	Rich	Argonne National Laboratory	racarlson@anl.gov
Carter	Bob	Honeywell Laboratories	carter_robert@htc.honeywell.com
Catlett	Charlie	Argonne National Laboratory	catlett@mcs.anl.gov
Claffy	KC	CAIDA	kc@caida.org
Clark	Dave	MIT	ddc@lcs.mit.edu
Cohen	Danny	CNRI	dannycohen@ieee.org
Corbato	Steve	UCAID/Internet2	corbato@internet2.edu
Cox	Chip	NSF/CISE/ANIR	chip@cox.net
Dao	Son	HRL Laboratories	skdao@hrl.com
Das	Sajal K.	University of Texas at Arlington	das@cse.uta.edu

desJardins	Dick	NASA Research & Education Network	rdesjardins@arc.nasa.gov
Dev	Parvati	Stanford University	parvati.dev@stanford.edu
Diersen	Dave	Chief of Naval Operation's Strategic Studies Group	diersend@nwc.navy.mil
Domich	Paul	Office of Science & Technology Policy	pdomich@ostp.eop.gov
Durst	Robert	The MITRE Corporation	durst@mitre.org
Echiverri	Kathy	Institute for Defense Analysis	kechiver@ida.org
Edwards	Hal	Nortel Networks	edwardsh@nortel.com
Eisenberg	Jon	Computer Science & Telecommunications Board	jeisenbe@nas.edu
ElBatt	Tamer	HRL Laboratories	telbatt@wins.hrl.com
Ephremides	Anthony	University of Maryland	tony@eng.umd.edu
Evans	Joe	University of Kansas	evans@ittc.ku.edu
Feng	Wu-chang	Oregon Graduate Institute	wuchang@cse.ogi.edu
Feng	Wu-chi	Ohio State University	wuchi@cis.ohio-state.edu
Feng	Wu-chun	Los Alamos National Laboratory & Ohio State University	feng@lanl.gov
Fisher	Darleen	NSF	dlfisher@eecs.berkeley.edu
Fleming	Robert	Aether Wire & Location, Inc.	bob@aetherwire.com
Foster	Ian	Argonne National Laboratory	foster@mcs.anl.gov
Foster	Mark	NASA/NREN	mafoster@arc.nasa.gov
Freeman	Ken	NASA Ames Research Center	kfreeman@arc.nasa.gov
Frost	Victor	University of Kansas	frost@eecs.ukans.edu
Furlani	Cita	NCO/ITR&D	furlani@itrd.gov
Gary	Pat	NASA Goddard Space Flight Center	pat.gary@gsfc.nasa.gov
Gilliam	David	NASA/Jet Propulsion Laboratory	david.gilliam@jpl.nasa.gov
Golubchik	Leana	University of Maryland	leana@cs.umd.edu
Govindan	Ramesh	University of Southern California – Information Sciences Institute	govindan@isi.edu
Greene	Tom	NSF/CISE/ANIR	tgreene@nsf.gov
Griggs	Kathleen	Puritan Research Corp.	kgriggs@puritanresearch.com
Griggs	Steve	Multi Spectral	sgriggs@multispectral.com
Gritter	Mark	Stanford University	mgritter@dsg.stanford.edu
Hayward	Gary	Telcordia Technologies	gah@lts.ncsc.mil
Hollebeek	Bob	University of Pennsylvania	bobh@nscp.upenn.edu
Howe	Sally	NCO/ITR&D	howe@itrd.gov
Hughes	Larry	Dalhousie University	lhughes2@is.dal.ca
Ingle	Jeff	Intelligence Community CIO Staff	jeffeti@odci.gov

Irwin	Basil	National Center for Atmospheric Research	irwin@ncar.ucar.edu
Izadpanah	Hossein	HRL Laboratories	hizad@hrl.com
Jackson	Deborah	NASA/Jet Propulsion Laboratory	deborah.j.jackson@jpl.nasa.gov
Jannotti	John	MIT	jj@lcs.mit.edu
Joa-ng	Mario	Telcordia Technologies	mjoang@research.telcordia.com
Johnson	Marjory	RIACS/NASA Ames Research Center	mjj@riacs.edu
Jones	Kevin	NASA Ames Research Center	kjones@arc.nasa.gov
Kandlur	Dilip	IBM TJ Watson Research Center	kandlur@us.ibm.com
Khan	Javed	Kent State University	javed@kent.edu
Kind	Pete	Institute for Defense Analysis	pkind@ida.org
Kittka	Kevin	Science & Technology Associates	kkittka@snap.org
Konishi	Kazunori	Asia-Pacific Advanced Networking	konishi@jp.apan.net
Koob	Gary	DARPA/ITO	gkoob@darpa.mil
Kovar	David	Western Disaster Center	kovar@webnexus.com
Kulik	Joanna	MIT	jokulik@lcs.mit.edu
Kulkarni	Amit	GE Corporate R & D	kulkarni@crd.ge.com
Kumar	Mohan	University of Texas at Arlington	kumar@cse.uta.edu
Kumar	Rakesh Teddy	Sarnoff Corporation	rkumar@sarnoff.com
Kumar	Sri	DARPA/ITO	skumar@darpa.mil
Kushner	Cherie	Aether Wire & Location, Inc.	cherie@aetherwire.com
Larsen	Ron	MAITI	rlarsen@deans.umd.edu
Lehman	Thomas	University of Southern California, Information Sciences Institute	tlehman@isi.edu
Lennon	Bill	Lawrence Livermore National Laboratory	wjlennon@llnl.gov
Liebeherr	Jorg	University of Virginia	jorg@cs.virginia.edu
Lockwood	John	Washington University	lockwood@arl.wustl.edu
Loyall	Joe	BBN Technologies	jloyall@bbn.com
MacKenzie	Robert	Solipsys	robert.mackenzie@solipsys.com
Maeda	Mari	DARPA/ITO	mmaeda@darpa.mil
Mandrekar	Ishan	Drexel University	ishan@io.ece.drexel.edu
Mankin	Allison	University of Southern California – Information Sciences Institute	mankin@isi.edu
Marquis	Jeff	University of Texas at Arlington	marquis@prismpti.com
Mathis	Matt	Pittsburgh Supercomputer Center	mathis@psc.edu
Maughan	Doug	DARPA/ATO	dmaughan@darpa.mil
McFarland Jr.	Ray	Laboratory for Telecommunication Science	rimcfar@afterlife.ncsc.mil
Miller	Grant	NCO/ITR&D	miller@itrd.gov

Mills	Dave	University of Delaware	mills@eecis.udel.edu
Mills	Kevin	NIST	kmills@nist.gov
Minden	Gary	University of Kansas	gminden@ittc.ukans.edu
Mishra	Amitabh	Virginia Tech	mishra@vt.edu
Monaco	Gregory	NSF	greg@greatplains.net
Montgomery	Doug	NIST	dougm@nist.gov
Mosse	Daniel	University of Pittsburgh	mosse@cs.pitt.edu
Mouchtaris	Petros	Telcordia Technologies	pmouchta@telcordia.com
Mount	Richard	Stanford Linear Accelerator Center	richard.mount@stanford.edu
Mundy	Russ	Network Associates Institute Labs	russ_mundy@nai.com
Murphy	Sandy	Network Associates Institute Labs	sandy@tislabs.com
Muthukrishnan	S.	AT&T	muthu@research.att.com
Nadeem	Tamer	University of Maryland	nadeem@cs.umd.edu
Ndousse	Thomas	DOE Office of Science	tndousse@er.doe.gov
Nelson	Kyle	Honeywell	kyle.nelson@honeywell.com
Newman	Harvey	California Institute of Technology	newman@hep.caltech.edu
Perkins	Colin	University of Southern California, Information Sciences Institute	csp@isi.edu
Personick	Stu	Drexel University	personick@ece.drexel.edu
Rajagopalan	Raj	Telcordia Technologies	sraj@research.telcordia.com
Rao	Nagi	Oak Ridge National Laboratory	raons@ornl.gov
Richeson	Anne	Qwest Communications	anne.richeson@qwest.com
Riedi	Rolf	Rice University	riedi@rice.edu
Roberson	Mark	MCNC	mwr@mcnc.org
Samad	Tariq	Honeywell Laboratories	tariq.samad@honeywell.com
Sawhney	Harpreet	Sarnoff Corporation	hsawhney@sarnoff.com
Sayyah	Keyvan	HRL Laboratories	krsayyah@hrl.com
Schantz	Rick	BBN Technologies	schantz@bbn.com
Schmidt	Doug	DARPA/ITO	dschmidt@darpa.mil
Scholtz	Jean	DARPA/ITO	jscholtz@darpa.mil
Scott	Keith	MITRE	scott@mitre.org
Semancik	Bill	National Security Agency	wjseman@lts.ncsc.mil
Seshan	Srini	Carnegie Mellon University	srini@cmu.edu
Seweryniak	George	DOE Office of Science	seweryni@er.doe.gov
Shankar	A. Udaya	University of Maryland	shankar@cs.umd.edu
Shapiro	Jonathan	University of Massachusetts	jshapiro@cs.umass.edu
Shirazi	Behrooz	University of Texas at Arlington	shirazi@cse.uta.edu
Shrikumar	H.	MIT - Media Lab	shri@mit.edu
Sieworek	Dan	Carnegie Mellon University	dps@cs.cmu.edu
Silverstein	Jonathan	University of Illinois at Chicago	jsilver@uic.edu
Sincoskie	Dave	Telcordia Technologies	sincos@research.telcordia.com

Singer	Ed	SES	ed.singer@erols.com
Smailagic	Asim	Carnegie Mellon University	asim@cs.cmu.edu
So	H. Wilson	University of California - Berkeley	so@cs.berkeley.edu
Sollins	Karen	MIT	ksollins@nsf.gov
St. Arnaud	Bill	CANARIE	bill.st.arnaud@canarie.ca
Stadler	Scott	MIT – Lincoln Laboratory	stadler@ll.mit.edu
Starobinski	David	Boston University	staro@bu.edu
Steenkiste	Peter	Carnegie Mellon University	prs@cs.cmu.edu
Sterbenz	James	BBN Technologies	jpgs@bbn.com
Stevenson	Dan	MCNC	stevenso@mcnc.org
Strawn	George	NSF	gstrawn@nsf.gov
Subrahmanian	V.S.	University of Maryland	vs@cs.umd.edu
Suda	Tatsuya	University of California - Irvine	suda@ics.uci.edu
Towsley	Don	University of Massachusetts	towsley@cs.umass.edu
Trachtenberg	Ari	Boston University	trachten@bu.edu
Turnbull	Bill	NOAA	wturnbull@hpcc.noaa.gov
Vogels	Werner	Cornell University	vogels@cs.cornell.edu
Walpole	Jonathan	Oregon Graduate Institute	Walpole@cse.ogi.edu
Watson	Robert	NAI Labs	robert_watson@nai.com
Willinger	Walter	AT & T Labs - Research	walter@research.att.com
Wroclawski	John	MIT	jtw@lcs.mit.edu
Yoon	Barbara	Consultant	byoon@erols.com
Zegura	Ellen	Georgia Tech	ewz@cc.gatech.edu
Zhang	Lixia	University of California, Los Angeles	lixia@cs.ucla.edu
Zirngibl	Martin	Lucent Technologies	mz@lucent.com
Znati	Ty	NSF	tznati@nsf.gov

Appendix 3: Zero-Casualty War Scenario

A massive conflict has broken out between countries A and B in Asia. Country B has invaded city "a," whose population is historically 40 percent of Country B origin. The attack includes possible use of chemical and biological agents. Country B has occupied the border valley leading up to city "a." The inhabitants of city "a" are locked in, but refugees from the surrounding region are pouring out into neutral Country C. Reports of atrocities and suffering within and around city "a" abound.

The U.S. has sent in humanitarian aid. The United Nations has imposed economic sanctions, which have not stopped the attacks, and has subsequently issued an ultimatum to Country B to withdraw its forces. The UN has now approved the use of force.

In preparation for conflict, the U.S. has deployed networks of sensors in and about the conflict region. These include chemical sensors (vehicle exhaust fumes, urine, chemical agents, etc.), broad-spectrum acoustic sensors, seismic sensors, video sensors, and imaging sensors. Some are mobile. The U.S. has coordinated with other regional powers for cooperative actions, and is preparing to engage in conflict promising — to the U.S. Congress — minimal casualty war.

The U.S. tasks the Army's FFCS (Future-Future Combat Systems) to:

- Contain or stop attacks on the civilian population and the refugees in the border regions
- Remove Country B's forces from Country A with minimal civilian casualty
- Establish a Zone of Separation between Country B and Country A
- Ensure that Country B's atrocities and aggression are displayed across the Internet and the media
- Conduct hand-over to follow-on forces when so ordered

This mission is to be achieved through a set of precision actions whose duration should not exceed 5 to 10 days.

To achieve these goals, the U.S. performs the following actions (this is an incomplete list):

- Establishes a forward operating base in Country C with reach-back to the continental U.S. (CONUS) over satellite (or whatever other communication links are reliably available)
- Deploys Unpiloted Aeronautical Vehicle (UAV) Networks (swarms of coordinated and cooperating UAVs) controlled via satellite and other relay links, and collaborating with each other via cross-link
- Establishes identification and tracking of each military air and ground vehicle in the area of operation
- Dispatches combined Air Cavalry and FFCS units from U.S. bases in Hawaii

The sensors are deployed via cruise missiles from ships in the Indian Ocean. The forward operating base is dispatched from Europe via transport planes, as are the UAV swarms. The UAVs consist of a mix of fixed- and rotor-wing aircraft, with altitudes of operation varying from tree-top level to 65,000 feet.

The forward operating base begins to perform tactical reconnaissance. Terrain, street, and building information are updated based on visual and acoustic information from the UAVs. Signals Intelligence data are gathered from the UAVs and relayed to the forward operating base for analysis and correlation. The acoustic, seismic, and visual signatures of each of the hundreds to thousands of motorized vehicles are cataloged with each vehicle being identified as military or civilian, and military vehicles are identified as hostile or friendly. Air Defense Artillery and Surface-to-Air missile sights are identified. Automated analysis of visual information provides data about the approximate number and locations of dismounted troops, enemy command posts, and command vehicles. Weapons and supplies stores are cataloged as identified. Visual, chemical, and acoustical indications of weapons fire are all enunciated within the forward operating base and video of that region is either initiated or tagged. Personnel within the forward operating base monitor tagged video for "newsworthy" clips and forward as appropriate to news media. Forward operating base data are relayed to CONUS. Video streams and live sensor reports may be relayed to CONUS at the initiation of CONUS or forward operating base personnel.

While en-route aboard transport planes, FFCS and Air Cavalry units begin to monitor vehicle-tracking information and dismounted troop movements. Although FFCS vehicles are disbursed among different transport vehicles, inter-vehicle communication is accomplished via air-to-air cross-links and each transport vehicle has high-rate connectivity to the satellite resources. Up-to-date terrain, street, building, and weather information is loaded into FFCS and Air Cav onboard databases via satellite from CONUS and the forward operating base. As the tactical picture emerges, officers in the forward operating base, CONUS, and aboard the transport planes collaborate to finalize operational plans. One or more FFCS units ("Cells") are assigned to each of the task forces that are to execute the following assignments:

- Task Force Alpha: Prevent the Country B troops occupying the border valley from moving further into Country A
- Task Force Bravo Cut off the logistics supply chain to Country B's forces
- Task Forces Charlie and Delta: Flank the Country B troops in the valley
- Task Forces Echo, Foxtrot, Golf, and Hotel: Encircle city "a" and use progressively lethal robotic means to disable or destroy Country B assets within the city. Coordinate with Special Operations Forces tasked to covertly evacuate Country A leaders from city "a" during the initial hours of the FFCS operation.
- Task Force India: Establish and occupy battle position to provide Air Cavalry support to other task forces
- Task Force Juliet: Reinforce forward operating base and establish refugee and enemy prisoner of war facilities

While still en route, task forces develop unit plans and assume responsibility for specific roads, intersections, and vehicles. Using intervisibility calculations derived from the terrain

data, FFCS units determine where to place robotic direct fire vehicles for optimal range fans and where to place robotic indirect fire vehicles, sensor vehicles, infantry carriers, and command and control vehicles.

Upon arrival at Country A, task forces deploy under cover of darkness, with some task forces being parachuted into position while others are dispatched from the forward operating base. Task forces deploy their organic UAV support, and the command and control vehicles establish processed sensor data and live video support as required from the UAVs under the control of the forward operating base.

Three hours before dawn, the Special Operations Forces enter the city and rendezvous with Country A's leaders who are to be extracted.

One hour before dawn, task forces launch indirect fire missiles to destroy all known Country B supply and weapons dumps in Country A. The Special Operations Forces depart the city during the ensuing confusion. At the same time, additional indirect fire missiles are targeted at Country B command posts and vehicles closest to the city. These missiles are instructed to loiter above their targets, above the range of enemy fire. The UN ultimatum is repeated to Country B and the surrender of Country B forces is solicited.

In response to the ultimatum, Country B forces within the city begin to fire on civilians. Muzzle-flash sensors identify the source of the firing, and loitering missiles are tasked to eliminate these vehicles. The call for surrender is repeated and Country B requests to withdraw. The task forces oversee the retreat of Country B troops and establishes the Zone of Separation along the border between Country A and Country B.

This mission succeeds with minimal casualties as a result of:

- Large scale sensor fusion
- Large scale target identification and tracking
- Large scale video acquisition, transmission, analysis, and routing
- Robotic command and control of surface and aerial forces
- Rapid insertion of overwhelming force

A key component of this mission is the transmission of critical, sensitive information over reliable, secure networks. The networks must support command and control as well as tactical operations. Functional networking requirements include the:

- Ability to disseminate information appropriately over high to low bandwidth links, assigning appropriate flows to appropriate links when multiple paths exist simultaneously
- Accommodation of voice, data, video, audio, still images, etc.
- Ability to provide high assurance, real-time closed-loop distributed control in a tactical environment while on the move in harsh terrain
- Seamless integration of heterogeneous component networks, including FFCS nets, other sensor nets, and command structure communications

- Provision of hierarchical views of the network structure, depending on the "user": forward operating base, CONUS, Task Force commanders, and FFCS Cell team leaders all have access to common views of the tactical situation, but typically operate at different levels of the hierarchy to enhance their effectiveness

Appendix 4: Deeply Networked World/SWARMS Scenario (Smart World Airforce Repair and Maintenance System)

We imagine in this future scenario the truly modernized Airforce, highly efficiently managing its plane maintenance and parts scheduling. The Airforce has recognized that they have several problems. First, in order to save money, the Airforce has learned from industry that they want to minimize their stock of parts and provide just-in-time delivery of repair parts. Second, the places at which the Airforce will need parts vary because the planes travel around the world and it is more efficient to have the parts at the locations where the planes will need them rather than requiring that the planes return to central repair facilities. The reasoning behind this is that the planes may be needed in action, and moving them far from their military arena for repairs or worse for routine maintenance is undesirable.

In that future time, we find a world where networking and networked devices have become more broadly and deeply deployed. In this world it is difficult to find places that are off "the net" because locations have available some form of wired, wireless, or other networking technology. Furthermore, the network is able to provide service to an extremely large number of devices essentially co-located. As a society, we have bought into a truly pervasive computing environment.

It has been recognized that humans cannot and should not be responsible for the detailed management of such a deeply networked world, nor for managing the quantity and heterogeneity of devices. The network must be not only self-organizing, but self-healing. To do this, there must be capabilities for measuring behavior, evaluating that behavior, and then either masking or correcting it, when problems arise. An "agent" model must be deployed, in which agents not only function correctly as individuals, but also understand how to collaborate with each other effectively in order to make higher level decisions than any individual agent might make. Both individuals and groups providing higher level, composite functions must be responsive to societal policy constraints that may change or evolve with time. The developers of this technology call it the *smart world*.

The truly modernized Airforce has bought into the *smart world* model and has determined that every repair component and parts depot will be "smart." A part will know where it is. A depot will know how many of each kind of part it has and have a model of what parts are needed based on reports about the schedule of arriving planes. The Airforce has gone further than that. Every part, when installed in a plane, is also introspective. Each knows how well it is functioning, and can predict when it will need to be replaced. Beyond that the composite systems not only integrate over all their parts, but also have a higher level understanding of the emergent system. Agents will query, collate, and manage the system. The architecture is called the Smart World Airforce Repair and Maintenance System or SWARMS. SWARMS can predict when and where specific repairs will be needed. At the level of the whole plane, it understands flight schedules, using this information to plan where and when work should be done. SWARMS will inform the global inventory system, which ensures that by the time a plane arrives at a destination the appropriate parts are there, with enough information that the repair can be made.

The Air Force is one of many branches of the government that have bought into the new "smart world" model. As a result, although many *smart worlds* are evolving, they will share some common network and services infrastructure, such as backbones. Furthermore, in some situations resources need to be shared among *smart worlds* (for example, requisitioning parts from outside the Airforce and interoperation under circumstances such as civilian crises in which the military helps provide support). These often reflect situations where coordination is needed or it is necessary to collaborate on the use of higher level resources.

In a military operation flexibility through over-provisioning is needed because predictive and planning capabilities will not be perfect and equipment reliability is critical.

Two objectives of the opposing force in this situation are espionage and sabotage. They would like to know what sorts of planes will be deployed at what locations and when. In addition, if possible they would like to cause those planes to be disabled.

The opposing force believes that SWARMS may be able to help them. Many of the individual parts and perhaps even collections of them, as well as the agents, were designed specifically to report on their state and make predictions about themselves, indicating when and under what conditions repairs or replacements will be needed. The agent-based network infrastructure may also help. The network is populated by network management agents that are collecting information, making decisions about network resources and behavior, and reporting as needed. Since these agents are "in the core of the network" they have to be trusted. However, different communities and individuals may have quite contradictory constraints on their operation and behavior, as well as differing policy controls. The opposing force believes that the combination of the smart network management world and SWARMS will allow for both the espionage and sabotage activities giving them the upper hand in the engagement.

In the context of SWARMS and this scenario there is a complex set of issues such as trust, assurance, and security, including privacy, authenticity, authorization, and denial of service. For example, one can consider that the "Unified Smart World" will consist not only of the Airforce's planes, but large numbers of additional planes and systems from the military and civilian communities. The motivation for ensuring smart behavior may range from economic forces to military security to life and death issues. If one of the Airforce's planes does not make it to a destination because of a part failure, this may cause not only loss of the plane, but also loss of life. To focus in further, one can ask what needs to be in the network and what can be pushed to the edges, recognizing that the true edges may be "dumb" devices with limited functionality and adaptability. Finally, one can consider the problem of providing assurance and trust in emergent systems.

Scenario networking elements:

- Deeply networked systems
- Security including authentication, authorization, and privacy
- Smart network management
- Network performance measurement

Appendix 5: Crisis Management Scenario

In 2010, NOAA began a five-year deployment of its Ultra Doppler/SAR radars across "Tornado Alley" in the central U.S. In the following year, DoD received "dual use" approval to supply early fire warning bulletins from its "staring" missile launch detection satellite systems. And in 2012, NASA orbited Firesat, capable of providing twice-a-day high resolution multi-spectral, multi-instrument views of forest fire activity around the world.

The year is now 2015, and a "perfect" fire and tornado season has descended on the U.S. Hot, dry Santa Ana winds have come to the West Coast of the U.S. with a vengeance, from San Diego all the way up to the Pacific Northwest. In the Central States, destructive twisters are beginning to form along nearly every low-pressure system that sets up across Oklahoma and Texas.

On the day of crisis, DoD reports dozens of new fires from a single highly charged lightning storm. By mid-day, the early morning Firesat images and data have been processed and disseminated to hundreds of state and Federal agencies. Hot-spot data are combined with vegetation cover and dryness sounder models to produce detailed next-24-hour maps for the worst hit regions. These maps and digital models are disseminated instantly to government command and control centers and are spotcast to individual homes and businesses most in danger.

Department of Interior and other Federal agency supercomputers begin "nowcasting" the present and predicted tracks of the worst fires. Their sophisticated models take into account detailed wind and temperature forecasts transferred in near real time from NOAA's National Weather Service computers as well as vegetation burn models built from satellite data and from digital elevation model data in U.S. Geodetic Survey computers. These models and data are transferred directly to Federal and state fire management command and control centers.

The models and nowcasts are also transferred by satellite communications to the forest fire field units, which return validation and update information. This field information, along with real-time atmospheric, chemical, and other environmental data from sensors deployed throughout the area – both in-situ microsensor platforms deployed in advance as well as self-contained, self-powered sensors dropped from aircraft that same morning – are continuously integrated into the nowcast models. Customized warning and evacuation messages are automatically provided to all the homes and businesses in the areas.

On the same day, tornadoes are touching down across a broad path in the central U.S. Data from the Ultra radars are continuously fed to NOAA supercomputers, which nowcast the locations and severity of the funnel clouds and touchdowns with half-mile accuracy. Within two minutes, every home that has a 60 percent or better chance of being in the path of a tornado receives a computer-generated call on its telephones and cell phones to take cover. Houses that have been automated go into protective mode, shutting off natural gas feeds, closing drapes, and taking other actions that they have been designed to do. Non-automated houses are shut down remotely by their owners from work, using secure Web browsers to access sensors and actuators plugged into their home networks. Police and fire units as well

as all individual homes in each affected area receive detailed nowcast maps of the tornado tracks and touchdowns in their neighborhoods. This allows the authorities to mobilize resources safely to provide emergency medical support to each affected site within two minutes after touchdown.

Immediately after each tornado hits, with power lines down and communication infrastructures demolished, emergency mobilization forces are directed by computer into the target zero area. The emergency units are set up quickly, establishing a high-performance field network instantly capable of local area and remote communications, using truck-based wireless technologies tied into regional networks via high performance satellite communications. Mobile whole-body scanners, sophisticated medical instruments, and mini chemical analysis labs are plugged into the network. This allows a formidable collection of medical specialists, data and information resources, and analysis facilities to be called on instantly to support the onsite paramedics. Command and control units have instant high-performance network access to all needed statewide and Federal resources. Within five minutes after touchdown, the President appears on television talking via two-way videophone to the area commander and a survivor at the worst touchdown site, declaring the region a disaster area and promising that the full resources of the Federal Emergency Management Agency will be made available to help the victims and promote a speedy recovery.

Scenario networking elements:

- Input from heterogeneous sensors
- Large-scale on-line modeling, QoS, and bandwidth requirements
- Automated delivery to diverse sites
- Self-contained small sensors
- Self-organization of networks
- On-line medical consultation capabilities

Appendix 6: Collaboration Scenario

Collaborative problem solving and decision making is a fundamental aspect of Dr. Clotho's research, which includes a large number of national and international partners and collaborators. Furthermore, much of the research is focused on domains in which complex tasks must be performed in environments that are largely inaccessible to human beings.

Dr. Clotho indicated that, "Collaboration can be much more fruitful if we can conduct, from the comfort of our labs, scientific experiments where all resources are located remotely, and yet all phases of the experiments are so well orchestrated that it appears a local endeavor to all participants. This should be possible, not only for immediate collaborators but also for everyone else that cares to join and participate in the research."

Dr. Clotho stated, "Current systems do not allow true collaboration and thwart natural, intuitive styles of interaction. They are passive in design and are difficult to adapt to particular work patterns. The services they provide are primarily fragmented and non-interoperable, which quickly makes them obsolete. How about security and the need to incorporate new applications in the collaboration space quickly and effectively?

"The scientific environment will include a large number of sensors and robots with varying capabilities, small enough to be embedded into the natural environment with minimum disturbances. The simple ones may be able to sense, compute, and act. The more sophisticated ones may be able to carry out more complex tasks related to detailed physical monitoring and manipulation. However, only through deployment of dense spatial sensing and the coordinated effort of a large number of these nodes can I explore my physical world and carry out my research goals.

"These low-power nodes, with limited communication bandwidth, need to understand local conditions and together collaborate to identify and monitor global environment conditions. They need to be able to form collaborative teams to accomplish complex tasks, such as surveying sites of scientific interests and coordinating to overcome potential problems and failures as they occur. The right level of coordination must be selected dynamically depending on the task at hand and the feedback from the scientists. For tasks where limited coordination is sufficient, the nodes can form teams and negotiate roles in a local yet distributed fashion. For tasks in which tighter coordination is required, the nodes can take a more global approach, potentially seeking the help of the scientists, in assigning roles."

"The problem," continued Dr. Clotho, "is that no single node has global knowledge and their capabilities preclude any centralized coordinated global sharing of state. So do not talk to me about centralized knowledge or control. It is simply not achievable for the large number of devices that I will be deploying at possibly fine granularities. It will be much easier for my colleagues and myself if coordination schemes that are inherently distributed and based on localized inputs, algorithms, and outputs were made available to us. Take the case of data collection, for example. This task is almost insurmountable in my small-scale simulated testbed. I cannot even begin to imagine the inevitable implosion of data that I will

be facing from the need to continuously monitor, at high resolution, the physical world I am interested in exploring. My task would be much easier if it were possible to execute local correlation and possible aggregation of data inside the network before I collect and process the data at the desired level of granularity."

Then in reference to her collaboration with her direct national and international peers, Dr. Clotho added, "Facilitating the interaction and collaboration among the large number of limited devices addresses only one aspect of my problem. We also need collaborative support at the scientist's level. We need tools for pro-active and dynamic support so that the system can behave as a problem-solving partner, adapting to that scientist's work pattern, potentially providing advice to advance collaboration. Each scientist should be able to augment the reality of the physical world with virtual worlds to allow the automated inclusion of 'what if' scenarios and conduct studies tailored to specific interests."

These requirements pose unprecedented challenges for future scalable, robust distributed system design. Add to that the unpredictability brought about by the micromobility of the devices and you will find that requirements and constraints cannot be addressed by current networking and distributed coordination technologies. Not only are traditional protocols too heavy-weight, but the building blocks of our current distributed systems and networking, namely layering, abstraction, and modularity, become questionable. The framework within which future collaborative research will be carried out must be flexible enough to deal with unpredictable and emergent changes, meet hard real-time constraints, and handle asynchronous events as they occur.

Appendix 7: Networked Medical Care Scenario

After eating a heavy Tex-Mex dinner, a middle-aged man is at home watching television one evening when he begins to suffer chest pains. Is he having a heart attack or is it just heartburn? He goes to his medicine cabinet where he finds a rod-like device that he presses against his bare chest. A yellow light comes on. The light soon turns green. The man presses the big red button. A few minutes later the telephone rings. A voice informs the man that he is having a heart problem and that an ambulance has been dispatched to his home.

The ambulance arrives to take the man to the hospital. On the way an EMT attaches the man to the vital signs and cardiac monitoring equipment. A cardiologist, who is at home but on duty, can be seen on the monitor in the ambulance. She has done a network search and has access to the patient's complete multi-media medical record. The cardiologist can see the patient and guides the EMT through a cardiac evaluation that includes heart sounds, a complete EKG and an echocardiogram. Based on her evaluation of the data received from the evaluation and her view of the patient, the cardiologist orders an angiogram to be performed as soon as the patient arrives at the hospital. The hospital angiogram team is assembled, scrubbed, and ready when the patient arrives.

The patient is prepared for the angiogram as soon as he arrives at the hospital. The procedure is performed. The hospital angiographer thinks that he has detected a cardiac anomaly but is unsure so he calls in a consultant. The consultant is reached at home and agrees to look at the angiogram on his home computer. When the consultant accesses the angiographic data over the network, the hospital angiographer receives a warning from the network that the video screen being used by the consultant does not meet the standard deemed necessary for angiogram interpretation. The consultant does not see the anomaly so, despite the warning from the network, it is decided that bypass surgery is necessary. Preparations are made to perform the surgery in the morning.

Bypass surgery begins. It appears to be a routine case until the surgeon realizes that the patient’s cardiac vasculature follows a rare anatomical anomaly for which the surgeon has no experience. From the operating room the surgeon is able to search a 3-D anatomical library for a similar case. He contacts the colleague responsible for the case he has found. The consultant is able to see the patient’s cardiac vasculature and advise on the proper way to successfully complete the operation. He might even occasionally take control of the haptic enabled surgical robot in order to help the resident surgeon through some of the more difficult or unfamiliar procedures.

Scenario networking implications:

Security

First and foremost, the network must be secure and must meet legal standards for medical data privacy and security, i.e., currently the proposed Health Insurance Portability and Accountability Act rules. The network must be able to securely carry stored data in a wireless environment.

Scalability and High Assurance Networking

To support the cardiologist at home, the wireless channel must provide real-time video and real-time data signal display. Although this application may tolerate a fair amount of latency, it will not tolerate jitter and the video, audio, and data channels must be synchronized.

End-to-End Performance, Intelligent Networking

The angiogram procedure identifies the need for end-to-end knowledge of the network data path including the display devices at the ends. An intelligent, scalable network should automatically find the work-around for the quality of service problem, but it must report the problem if a work-around is not possible.

Assured Real-time Service

For the bypass surgery scenario, the surgeon needs to retrieve 3-D image data sets, each of which may be several gigabytes in size. The consultant must be able to view the live procedure in real time and accurately guide the surgical robot through its haptic controls. This requires a network that is able to operate at high bandwidth with minimal latency and minimal jitter while maintaining the security and integrity of the transmissions and the privacy of the patient data.

Appendix 8: High Energy Physics Scenario

Never before has the scientific mission of particle physics research been so dependent on state-of-the-art information technology. Collaborations of hundreds to thousands of physicists and engineers are formed to create accelerators, detectors, and analysis systems with a productive life of tens of years. These analysis systems form a complex and widely distributed "fabric" of computing and storage resources.

The non-deterministic nature of quantum physics, uneasily understood during the last century, inevitably requires the measurement and analysis of billions of particle interactions to observe and understand fundamental processes. Particle physics experiments have pushed against the limits of technology, electronics, computing, and networking for decades. Detectors with millions of channels, each recording precise amplitudes with a resolution of picoseconds, have in the course of 40 years succeeded detectors with a few single-bit "yes/no" measuring devices. Information flows from such a detector at up to a terabit per second and must be drastically filtered in real time because of limited storage, analysis, and networking facilities.

The Large Hadron Collider (LHC) experiments at the European Organization for Nuclear Research (CERN) will rapidly reach tens of petabytes of stored data under intense analysis. The design, construction, and data analysis for an experiment require the combined intellect and dedicated work of international collaborations. However, technological limitations on the storage, transmission, and analysis of data impose difficult, even dangerous choices. For example, the LHC experiments expect to be able to record and share over networks less than one millionth of the collisions they observe. This draconian real-time selection will necessarily have to be optimized for "somewhat expected" new discoveries rather than the "totally unexpected" ones that are the dream of every scientist.

Even after the draconian selection, LHC collaborations will face the challenge of empowering thousands of geographically distributed physicists to use their intellect and wisdom to derive physics insight from tens of petabytes of data. Although the raw cost of bandwidth is no longer a crippling impediment, the end-to-end performance of applications is often unacceptable. Success will rely on middleware research to support data-intensive, worldwide collaborative science that is only beginning. A minimum requirement is the location-independent ability to analyze data to empower all of an experiment's physicists to work collaboratively on databases, growing to tens of petabytes in 2005-2010, using all computing resources to which they have access.

However a qualitative change in the way research is performed would be enabled if we could free the real-time selection of data from the constraint of a single filter system with selection algorithms decided by committees. The availability of networks with end-to-end terabit performance could make this possible, but speed alone is not enough.

The data acquisition and filtering systems might profitably become geographically distributed and operate as highly parallel, largely asynchronous data flow systems. The terabit systems that will become operational in 2005 for the LHC will include a multi-terabit

capacity switching fabric, but individual data acquisition nodes and filtering nodes will communicate at gigabit speeds. In addition to the substantial bandwidth requirement, challenges include:

- The multicast service required when more than one remote filtering center is available
- Achieving adequate error rate and robustness without *ever* allowing the implementation of the "wild idea" to impact the detector-site data acquisition system

Acknowledgements

The LSN workshop and this report were generated through the creativity, work, participation, and review of many individuals. Every participant listed in the attendee section of the report contributed his or her expertise, professional insights, and recommendations to the discussion.

Daniel Hitchcock (DOE) and George Strawn (NSF), are Co-Chairs of the Large Scale Networking Coordinating Group that sponsored the workshop. Mari Maeda (DARPA) and Karen Sollins (NSF), were the workshop co-chairs, and Robert Durst (MITRE), Kevin Mills (NIST), and Tatsuya Suda (University of California-Irvine) were the workshop panel chairs.

The draft workshop report was prepared by Grant Miller (NCO/Noesis) with the writing and editing assistance of David Bernholz (NCO), Sally Howe (NCO), Martha Matzke (NCO/Noesis), Richard Mount (Stanford Linear Accelerator Center), Karen Sollins (NSF), Richard Schantz (BBN Technologies), and Barbara Yoon (Consultant).

Workshop Organizing Committee
Michael Ackerman (NIH/NLM)
David Bernholz (NCO)
Richard desJardins (NASA Ames Research Center)
Daniel Hitchcock (DOE)
Marjory Johnson (NASA Ames Research Center)
Mari Maeda (DARPA)
Grant Miller (NCO/Noesis)
Douglas Montgomery (NIST)
Thomas Ndousse (DOE)
William Semancik (NSA)
George Seweryniak (DOE)
Karen Sollins (NSF)
George Strawn (NSF)
William Turnbull (NOAA)
Taieb Znati (NSF)

Scenario Authors

Robert Durst (MITRE)	Zero Casualty War
Karen Sollins (NSF)	SWARMS
Richard desJardins (NASA) and William Turnbull (NOAA)	Crisis Management
Michael Ackerman (NIH)	Networked Medical Care
Daniel Hitchcock (DOE) and Richard Mount (Stanford Linear Accelerator Center)	High Energy Physics
Taieb Znati (NSF)	Collaboratories

Breakout Session Chairs

Richard Shantz (BBN Technologies) and Barbara Yoon (Consultant)	Zero Casualty War
John Wroclawski (MIT) and Danny Cohen (CNRI)	SWARMS
Donald Towsley (University of Massachusetts)	Crisis Management
Charles Catlett (Argonne National Laboratory)	Networked Medical Care
Richard Mount (Stanford Linear Accelerator Center)	High Energy Physics
Darleen Fisher (NSF)	Collaboratories