# SCOPING THE FUTURE MULTI-SERVICE ADVANCED NETWORK INFRASTRUCTURE FOR U.S. HIGHER EDUCATION

## A Preliminary Assessment

## – Group A Report –

May 2005


– GROUP A –

Erv Blythe (Virginia Tech), Javad Boroumand (Cisco Systems), Steven Corbató (Internet2),
Wendy Huntoon (committee chair; Pittsburgh Supercomputing Center),
Ron Johnson (University of Washington), Michael Krugman (Boston University),
Richard Summerhill (Internet2), and Douglas Van Houweling (Internet2)

# Contents

[ Linked – Click to Follow ]

# Scoping the Future Multi-Service Advanced Network Infrastructure for U.S. Higher Education

A Preliminary Assessment

**– Group A Report –**

May 2005

## Executive Summary

This is the report of Group A, one of two teams formed in February 2005 by the National LambdaRail (NLR) Board and the Internet2 Networking Policy and Planning Council (NPPAC) as a first step in a two-stage, strategic planning process. Group A was charged with analyzing NLR and Internet2 service plans, identifying technical and operational requirements, and recommending any changes needed to accommodate service requirements and facilitate convergence.[1] Group A's efforts are directed at supporting the activities and analysis of Group B, another team formed as part of the NLR/Internet2 planning process charged with reviewing broader issues and offering subsequent recommendations.

Group A met via weekly conference calls and once in person from February through April 2005. It drew on the perspectives of its eight representatives without formal input from others, although the NPPAC was kept apprised and a draft report was issued on April 4, 2005, to Group B and the NLR Board for preliminary review and reaction. Refinements followed, culminating in this final report issued in May 2005.

This report opens in Section 1 with a brief, broad, and idealized vision of the global network that might one day serve research and education and the importance of regional networks in materializing this vision. This section also provides a preliminary needs assessment, an overview of services required, and the role and relationship of the regional and backbone networks in meeting these needs and delivering services. Section 2 outlines the context in which this report is framed, the attributes of our evolving national network infrastructure and the need for a common interconnect and service model. In Section 3, current and anticipated services are described, and Section 4 offers a brief perspective on financial considerations. Appendix B elaborates on our longer-term vision. A centerpiece of this report is contained in Appendix C, which provides a succinct table of current and anticipated NLR/Internet2 services and capabilities.

It is noteworthy that the Group A charge specifically focuses on the relationship between NLR and Internet2 resources and services. This focus is reflected in the scope of this report. While Group A recognizes that regional networks are critical to the long-term evolution of research and education (R&E) networks overall, given the focus of our charge and the timeframe allotted for our discussions and the development of this report, it was impractical to fully or fairly incorporate a regional perspective. Nonetheless, *regional input is essential to a thorough and comprehensive analysis, and Group A recommends that the overall NLR/Internet2 strategic planning process solicit and integrate input from R&E regional network service providers*.

Over the course of our deliberations and the compilation of this report, Group A reaffirmed its conviction that evolving requirements, advancing technology, and the expansion of facility-based network infrastructure promise vastly greater options and opportunities, but also carry

[1] The complete Group A charge is included in Appendix A.

substantially increased complexity and costs. At the same time, financial pressures continue to mount, underscoring the urgency of avoiding the promulgation of competing and non-complementary organizational structures, resources, and services. *It is therefore more important than ever for NLR, Internet2, and the larger regional, national, and international networking communities to redouble our collective commitment to planning, outreach, coordination, aggregation, and collaboration*. Facility-based optical networks demand new and significantly expanded models for interconnection, support, and service, spanning international, national, regional, and campus domains, as well as government entities and commercial service and equipment providers. Our progress to date building and supporting networks to serve research and education represents an excellent foundation, but even greater challenges lie ahead.

This report surveys this landscape, provides reconnaissance, offers perspective, and elaborates on the issues outlined above. However, it is important to note that this report is by no means exhaustive or conclusive. It is offered as a preliminary step in an ongoing process that should increasingly engage others—regionally, nationally, and internationally—in a comprehensive and inclusive effort to establish the architectural, structural, organizational, and financial foundations necessary to support both the anticipated and unanticipated requirements of our community in the years to come.

# Scoping the Future Multi-Service Advanced Network Infrastructure for U.S. Higher Education

A Preliminary Assessment

**– Group A Report –**

May 2005

## 1. Global Network for Research and Education

A brief, broad, and idealized vision of the global network that one day might serve the research and education community is provided below. The evolution of the architecture and the stakeholders associated with the vision are described as well. The drivers for this vision are the networking needs of the broad research and education community, specifically the aggregation of the diverse and changing end-to-end connectivity and services needs of the community's individuals, organizations, projects, and end-users. The needs of this community, as well as how these needs are currently supported through regional aggregation points, are also described below.

### 1.1. Vision

Our evolving global research and education network environment should constitute a general purpose, flexible, interoperable, cohesive, and secure communications platform that accommodates current, anticipated, and unforeseen requirements gracefully, reliably, and cost-effectively in ways which maximize options and accommodate change. Its architecture should minimize complexity and maximize flexibility. As always, the overarching goal is to evolve a pervasive communications fabric that fosters innovation and serves as the foundation for the development of applications that encourage collaboration, the sharing and integration of human and material resources, technology transfer and commercialization, and substantial improvements in research and education.

### 1.2. Evolution of Architecture

Over the past decade, the unprecedented scalability and favorable economics of the Internet Protocol (IP) have created an environment where most advanced infrastructure providers view IP, both in its current and next generation versions (IPv4 and IPv6, respectively), as the common bearer network service. However, the recent availability of cost-effective dark fiber, especially regionally, enables a new networking paradigm. We expect a new network architecture to emerge that supports a broad range of network services, from dark fiber and co-location at the most fundamental level, through more deterministic services such as Multiprotocol Label Switching (MPLS) tunnels via the IP substrate, to shared or dedicated Gigabit Ethernet channels and dedicated 10 Gigabit per second (Gbps) wavelengths. In addition to the increased diversity of service offerings, the community anticipates services of varying class and duration. The challenge will be to provide the combination of services and infrastructure that are most popular, while maintaining the flexibility to accommodate special needs.

## 1.3. Stakeholders

The Group A charge specifically identifies the core stakeholders as Internet2, NLR and the regional entities, often referred to as gigapops, but now evolving into Regional Optical Networks (RONs). However, the broader array of stakeholders includes, among others, the various national research and education network organizations, federal research agency networks, and international research network organizations. Examples include the U.S. network and systems research communities, the International Equal Educational Access Foundation (IEEAF), EDUCAUSE, The Quilt, corporate partners, and research and education broadly construed.

## 1.4. Needs of End-Users and Sites

The networking needs of the research and education community are not specifically tied to any particular backbone network or backbone network facility or facilities, nor to any particular regional network or gigapop. Meeting the ever-changing networking needs of the research community requires a set of interconnected fabrics. These fabrics comprise a continuum from departmental and campus network infrastructure through regional fabric to gigapops or similar interconnection, aggregation, and exchange points; through those gigapops, or aggregation points, either directly to other gigapops via regional exchange or else on to backbone networks, to best reach other gigapops or other types of transit, exchange, interconnection points, and related local and regional fabric; and ultimately on to endpoints.

Of the huge array of network services facilities, capabilities, and even backbones of crucial importance to many or most in the higher education R&E community, this document is very narrowly focused on issues relating to one of the major shared backbone networks (Internet2's Abilene Network) and one of the backbone facilities (NLR). It is important to remember that they are only two of a profusion of elements of importance to research and education in the overall ecology of networking.

It is not yet clear what the right number of Higher Education networking organizations is at the backbone level. As that issue is discussed, it is important to remember that a large number of other entities provide services at least as important as national backbone capabilities, and that the future will need to be an auspicious mix of services and capabilities from a wide range of national, regional, and local levels and a set of complex relationships among them that cannot be determined in this treatment of NLR and Internet2.

### 1.4.1. Needs of the Community

R&E networking needs are an aggregation of diverse and changing end-to-end connectivity and service needs of our community's individuals, organizations, projects and end-users. This group includes a vast population of faculty, clinicians, staff, students, researchers and external partners in the US and around the globe. It includes the US and international research agencies such as NASA, NIH, DoE, DoD, NOAA, NSF, their research facilities, instruments and labs, and that, in some cases, have research networks themselves. It includes myriad global communities of corporate entities, suppliers, partners and collaborators with information technology (IT), data, and transactional resources that interact with those individuals and their devices, projects, and organizations. R&E networking requirements also include sustaining a diverse array of access by and interaction with the larger public, including alumni, prospective students, faculty and staff, patients, prospective patients, donors and potential donors, and an expansive domain of other external governmental, international, educational, project, discipline, corporate and small business organizations, and associations.

In sum, this broad community needs a wide array of networking services and facilities including:

- Commodity Internet;
- Local and regional interconnection peering and exchange facilities and services;
- Local, regional, national, and international private point-to-point circuits and private layer 2 and 3 networks[2];
- Direct access to raw fiber at metro- and inter-regional levels;
- Access to each of the US R&E backbones, including ESnet, DREN, NLR, UltraScience Net and Abilene, and;
- Access to on-demand, switchable and/or dedicated wavelengths and/or light paths that are part of the emerging fabric of waves and switching points represented by the Global Lambda Integration Facility (GLIF).

The combination of local and regional circuits, including waves over owned and lit fiber, and the gigapops or other similar forms of inter-connection, aggregation, exchange, and peering are the principal means for providing our community, both individuals and institutions, with access to these commodity, commercial, and R&E backbone services and facilities (e.g., Abilene, ESnet, NLR, etc.) and hence to each other and the larger community with whom they need to connect. Gigapops and the evolving RONs provide Abilene access to over 90% of the roughly 200 Internet2 university members, as well as to many thousands of colleges, community colleges and K-12 sites, which are aggregated as Sponsored Educational Group Participants (SEGPs) that can transit Abilene.

### 1.4.2. Services Required

The end users, sites, gigapops, and other local and regional entities that house most of R&E networks and network users provide a wide range of network services and capabilities. These include:

1. FULL END-TO-END NETWORK CONNECTIVITY, of the kind needed, and with the performance needed, and at the cost needed, by any individual or project between any endpoints needed anywhere they may be on the globe.
2. COMMODITY INTERNET SERVICES AND/OR COMMODITY PEERING that is reliable, high-performance, and as inexpensive as possible. Our community spends considerably more money on commodity Internet access than on R&E networking, and commodity service is a higher priority for universities, colleges, and K12 institutions. Commodity services are essential to interconnecting us with the world. The regional networks or gigapops, as both the primary level of aggregation and critical-mass market influence, have been and are likely to remain the primary Internet connectors for members of Internet2 and NLR (and state K20 networks, etc.), as well as the rest of the R&E community nationwide.

---

[2] See section 2.1.3. for a description of network layers

3. LOCAL, INTRA-REGIONAL, AND TYPICALLY ACCEPTABLE USE POLICY (AUP)-FREE CIRCUITS, BANDWIDTH AND FIBER INFRASTRUCTURE to enable [a] cost-effective, 'high-enough' performance, and appropriately reliable layer 1 capacity to/from (and, in some cases, among) campuses, labs, and hospitals for carrying both commodity and R&E traffic (and, increasingly, dynamic light paths/lambdas) between campuses and other entities, and [b] cost-effective, nearby, logical regional access and interconnection points providing optimal price/performance commodity services and the necessary range of R&E network access for layer 2 and 3 networks, as well as any necessary lambda and light path switching fabrics.
4. TYPICALLY AUP-FREE INTRA-REGIONAL SHARED OR AGGREGATED LAYER 3 BANDWIDTH AND SERVICES to aggregate, transit, and route both local and regional commodity and R&E traffic. Regional layer 3 services (e.g., CENIC's) are often the best means for providing cost-effective and reliable access to and aggregation of the various kinds of layer 3 traffic that users and sites need. In many places and for many sites, it is an economic and performance imperative to share bandwidth and to seek greater aggregations of traffic at regional levels.
5. LOCAL PEERING, EXCHANGE POINTS, AND DIRECT INTERCONNECTION CAPABILITIES AT LOCAL, REGIONAL, AND MEGA-REGIONAL LEVELS to provide both the performance improvements and cost savings associated with direct, fewest-hop, and shortest-distance traffic exchange. It is almost never useful or cost-effective to send traffic beyond the first place it can be successfully exchanged.
6. COLLOCATION SPACE available from a local campus, regional gigapop, or regional mega-pop that is conveniently accessible to researchers, experimenters, and others.
7. SUPPORT AND INFRASTRUCTURE for specialized community, local, and regional networks, including K12, K20, health care, and other affinity networking that often includes necessarily localized programmatic support and services.
8. POLICY AND LOBBYING for information technologies and telecommunications at local, state, regional, and national levels that reflects the challenges and opportunities of specific geo-political regions and leverages common regional interests.
9. DIRECT ACCESS TO ALL NORTH AMERICAN R&E BACKBONES AND FACILITIES, such as ESnet, DREN, and Abilene shared layer 3 services, and multi-wave and owned optical facilities like UltraScience Net and NLR.
10. INTERNATIONAL CONNECTIVITY via densely interconnected common connection and exchange points and transit mechanisms that are largely AUP-free and which provide for allocated or dedicated lambda and/or light path capabilities in addition to high performance layer 2 and 3 exchange and transit.

### 1.4.3. Relationship to the Backbone

As the aggregating entities between national and campus networks, the regional networks represent an increasingly critical component of the core network infrastructure for US research and education. In recent years, many of these regional efforts have been transformed fundamentally through the acquisition of dark fiber and the subsequent development of facilities-based networks known as regional optical networks or RONs. Interestingly, with a few notable exceptions (e.g., the NIH-funded Lariat project, NEREN), the sources of capital funding for the RONs have limited their scale to that of individual states.

In the evolving network environment, the regional gigapops and RONs will play a critical role in both delivering and contributing to the evolution of the new services and capabilities. Each entity may be expected to maintain common interconnection and service interfaces with their member campuses and end users, downstream advanced national networks, and any other regional networks or RONs with which it may peer optically or otherwise. The degree of communication and overall collaboration required to maintain service interconnectivity and interoperability with many other networks (e.g., deploying new services identified within the GLIF) is anticipated to drive new expectations for technical staffing and expertise within the RONs.

The potential for significant management efficiencies and other economies of scale may lead to the emergence of multi-state overlay structures in support of RON organization and operation. Currently, The Quilt serves as the national locus for collaboration among the regional networks, but the anticipated increase in overall system complexity likely motivates the investigation of a new organizational model that would be capable of providing some core and ancillary supporting services and possibly supplementary financial backing to the RONs. However, as indicated above, the regional aggregation points, whether gigapops, RONs, or something in between, are a very diverse set of organizations providing a broad range of services and capabilities. It is conceivable that one organizational model may not suffice. This matter is also left for Group B to consider.

## 2. Network and Technical Environment

The network and technical environment can be described fundamentally as the common way that we interconnect a variety of network components to provide connectivity between two or more nodes. The initial environment may have a set of specific characteristics, but these characteristics will evolve over time, influenced by both new technologies and new requirements. This environment supports both *capabilities* and *services*. Capabilities are things that the network can do, while services are specific parts of the infrastructure that users have access to, whether specifically provisioned or simply just available. Capabilities and services available within the network environment can be described by the specific combination of characteristics or attributes required.

The network and technical environment must be built to accommodate as many services and capabilities as possible, then frequently tested to probe any limitations. When capabilities and options have been narrowed, it is essential to preserve the ability to rebuild or broaden the environment at a later date. The availability of a broad range of capabilities, while not critical for the majority of R&E entities, is essential for the experimental and leading edge segments of the community. Retaining the latitude to accommodate longer-term requirements, specifically the longer-term development of capabilities and services, is critical to the evolution of the network. Capabilities must be preserved, not because they are specifically or immediately required, but in order to foster the development of new applications and services. Capabilities can be prioritized based on strategic needs and decisions, with the highest priority capabilities often becoming the actual services. A common interconnection, support, and business model is needed in order to articulate and to sustain the range of services and capabilities as an end-to-end infrastructure through all its dimensions and attributes.

The sections below describe the range of attributes, or multiple dimensions, associated with the network and technical environment. A common interconnection and service development model that will promote delivering the services and capabilities to the end user is also described.

## 2.1. Attributes

Below is a brief outline of the primary network attributes identified by Group A and considered critical to the new networking environment.

### 2.1.1. Capacity

Capacity is the amount of data, or information, which can be transferred over a specific time period, typically referred to as the available data rate or bandwidth.

### 2.1.2. Duration (Short-Term to Long-Term)

Duration represents the length of time that a service or capability must be provided. The duration of service delivery can refer to a specific component the infrastructure itself, such as a point-to-point connection, or for a specific period of general access, such as with a shared IP network. In general, duration can range from essentially zero time to indefinite (permanent). Two wide ranging examples of this attribute include a single application requiring just ten seconds of connectivity over a 40-Gbps, shared IP network and a 20-year contractual commitment (usually via an IRU) for a dark fiber segment.

### 2.1.3. Service Layer (0 – 10)

Capabilities and services can be described within the framework of the traditional OSI (or layered) service model. This concept includes all of the original OSI layers (1-7) as well as subsequent extensions (layers 8 through 10) that take financial, political, and human issues into consideration. Most recently, R&E network services typically have been defined within layer 3 – the network or IP layer. However, it is increasingly common to find network requirements within the following layers:

- Layer 0 (dark fiber or wireless spectrum)
- Layer 1 (wavelengths)
- Layer 2 (Gigabit Ethernet lightpaths, SONET circuits)
- Layer 2.5 (MPLS tunnels)
- Layer 3 (IPv4/v6)

In the future, we anticipate that service requirements at layers 4 and higher may become increasingly common (e.g., middleware federations).

### 2.1.4. Availability and Robustness

Availability and robustness are based on the expectations of subscribers and the stability requirements of their applications. While the broad terms "production" and "experimental" typically have been used to characterize the expected level of availability for the network infrastructure, the evolving practical requirements for this attribute are now more complex. For example, while the term production is typically characterized by the long-term availability of core network components (such as “four 9’s”, representing 99.99% average service availability over time), the mechanisms for achieving this objective vary. Robustness, the ability of the networking environment to withstand disruption, plays a large role in the overall system availability. For example, a very robust service could respond seamlessly to component failures and could fail over quickly to alternate paths in the event of a circuit outage. The factors influencing availability and robustness also depend upon the particular network layer of the intended service. For a layer 3, IP-based service, latency and packet loss are among the critical factors influencing this attribute.

### 2.1.5. Cost, Price, and Business Model

Cost is the level of resources that are required (i.e., manpower, money, equipment) to support a specific capability or service. Cost considerations can involve tradeoffs such as risk or the exclusion of other capabilities or services. Price is the actual monetary cost of the service to the consumer and ranges from essentially free services to significant levels of long-term financial commitment.

### 2.1.6. Reach

Since networking infrastructure provides connectivity between at least two end points and often times among many more, the ability of the network to reach the set of desired destinations is a crucial component. The emergence of new networking architectures and models has changed the networking concept of reach. Components of the reach attribute include availability of desired network components in specific geographic areas; availability of end point infrastructure (i.e., the traditional last mile or, in some cases, ‘last inch’ problem); and support of specific protocols, services, or capabilities across network boundaries, particularly as private networks emerge.

### 2.1.7. Coordination and Collaboration

As the networking environment grows in complexity, the successful delivery of the desired capabilities and services is critically dependent on the ability of the groups managing and overseeing the various network components to collaborate and to coordinate with each other. In the simplest example (e.g., services provided across an individual campus), a single organization might be responsible for all the network components and characteristics. However, for essentially all other services and capabilities, this attribute will need to be considered. Coordination and collaboration will include the ability to match end-user requirements with the actual and planned infrastructure characteristics across network boundaries. With multiple organizations influencing what needs to be done, efficient mechanisms for developing standards and best practices and timely sharing of this information are necessary.

### 2.1.8. Support

Broadly defined, support is the set of actions that are necessary in order to make a specific capability or service work as desired. Support can range from essentially non-intervention, when the service is available ubiquitously to all users, to a service that must be supported with intensive manual engagement and multi-party communication for the duration. As the complexity of the network environment increases, so will the levels of support required for the end users to utilize the full range of available services effectively.

### 2.1.9. Security/Privacy

The requirements for security and privacy can range from a completely open environment to one with specific security or privacy regulations (e.g., encryption) to one that is fully private . This pervasive attribute spans all of the network layers and includes any requirements encompassing the physical environment of the end users. Increasingly, the need to specify and then to assure desired levels for this attribute is a significant concern across the entire community.

### 2.1.10. Organizational Structure/Strength

R&E network environments – both current and anticipated – are complex in practice as they are supported by multiple distinct, but cooperating, organizations. The suite of end-to-end capabilities and services can require different levels of organizational support, ranging from essentially none for a transparent, readily accessible service to intensive support for the most complex service requirements. Similarly, from an end user perspective, the service experience is critically dependent on both the individual strengths and the overall coordination of all the organizations along the service path. Organizational structures and strengths vary greatly – from *ad hoc,* volunteer-driven initiatives to organizations that are tightly managed and operated. We observe that a fairly mature organizational infrastructure will be required in order to maximize the benefits from the fuller set of capabilities and services potentially available in the new networking environment.

### 2.1.11. Product Space

Specific products and services depend critically on both the capabilities of the underlying infrastructure and the community's ability to provide support. The product space attribute ranges from specific telecommunications services, such as dark fiber availability and leased services through network equipment, such as optical gear, switches, and routers, all the way to specific software and applications. The relationship between the community's defined product space and the offerings of the commercial marketplace is critical as well. Typically, a mature product can be procured directly from commercial entities. Alternatively, experimental products that are under development within the community often are available to a limited set of users.

### 2.1.12. Automation

Automation describes the minimization of intervention actually needed to support the capabilities and services within the advanced network environment. This attribute is closely tied to the maturity of a particular service, the amount of relevant support available, and the level of coordination and collaboration within the environment. The degree of automation achievable depends greatly on the aggregate capabilities of the organizations involved in provisioning and maintaining the service.

### 2.2. Common Interconnection and Service Development Model

This richer network architecture and environment will drive a new model for network interconnection, support, and service within the community. A common interconnection and service development model must be created to provide standards on which to provide a more flexible interface for enabling new capabilities and services, which is much different than what has been done in the past. A key objective is the rapid and efficient matching of network capabilities and services to user requirements, as well as the rapid deployment and ongoing performance assurance of these capabilities. At the same time, a greater degree of support and coordination may be required not only for the discipline communities, but also between cooperating network operators at all scales. The network support function must have the capability to work effectively to coordinate campus, regional, and national network requirements in a collaborative fashion and to leverage existing relationships with commercial carriers and product vendors when the requisite optical facilities and services are not available within the community. This environment must also promote and support persistent application development along with the evolution of new and critical applications into the commercial community. It must be able to identify new ideas proactively at early stages, particularly ones that may have a dramatic impact on the network, and then disseminate these concepts to others.

Most recently, while the concept of the IP common bearer service has served the community well, this common model has become increasingly straightforward – either Ethernet or SONET at layer 2, IPv4 at layer 3, BGP-4 for routing, and usually TCP or UDP at layer 4. In fact, several of these standards stem from effective practices put in place during the NSFNET era, or even earlier. During the upcoming period, we expect the technical options at each layer to grow. Nevertheless, the need for establishing a common service model that supports experimentation in new technologies and services while fostering interoperability and maximizing connectivity is seen as fundamental to the success of this endeavor. In addition, with the growth in fiber optic based interconnectivity, the common model must expand to provide standards for optical interconnectivity (e.g., ITU grid specifications), signaling, and switching. A new organizational framework is anticipated for defining and disseminating the standards within the common model. In particular, the community must ensure a process that motivates all participating networks, including regional and campus networks, to adopt the common service model and to stay current as it evolves. The further definition of this structure is left as an exercise for Group B.

## 3. Expected Services and Capabilities

It is expected that the evolving network and technical environment will provide a continuum of services and capabilities that grow and change over the next five years. It is expected that these services will span the service and capability attributes described above. Appendix C describes the current and near term (i.e. 18 months from now) services and capabilities expected to be supported by NLR and Internet2. More detailed examples from the table are described below in order to provide a better understanding of what the services are and how applications and end users might utilize them. Examples are given for both current and near-term services. A longer-term view is necessary in order to understand the continuum and to put the current and near-term services and capabilities into proper perspective.

## 3.1. Current and Near Term Examples

Various types of wave-based services are currently in demand within the networking community. Core waves provide the basis for a range of other services such as IP-based networks and application-specific Grids. Similarly, the availability of flexible waves has the potential for providing cost-effective infrastructure for a number of network and scientific research projects.

### 3.1.1. Core Waves

The primary attributes of high-reliability core waves are carrier class or similar levels of service with optimized availability. These waves would support a wide range of applications that include IP network backbone and backhaul, grid applications, such as weather forecasting grids, radiological image transfer, etc. While the core infrastructure, capability, and services will range across the technical attribute space, there are some consistencies between them. Below is an attribute description of two possible core services. These services tend to be supported from standard telecommunications providers or typical research and education networks.

#### a. IP Network

An IP network typically has the following attributes: Long duration, service layer 3, high reliability and robustness, and a fairly standard economic model (with the cost consistent with the type of IP service ranging from commodity to typical high performance IP network). Reach is fairly ubiquitous, reaching anyone who wants to use the service. The service requires a level of support and collaboration that is consistent with what is currently available from Internet service providers. The service is well automated so that changes are typically transparent to the end user.

#### b. Application Specific Grid

An Application Specific Grid is network infrastructure configured specifically to meet the network needs of an application or application community. This service will have a number of the same attributes that supports an IP network, specifically high reliability and robustness, reaching all possible users, support and collaboration consistent with what is currently available, and well automated. Variations from the typical IP network are tied to the specific needs and requirements associated with the specific application. For example, duration could be short-term or long-term depending upon the lifetime of the application; service layer will typically range from layer 1 to layer 3.

### 3.1.2. Flexible Waves

Bandwidth requirements for flexible waves are similar to that for core waves, typically ranging from 2.5Gbps to 10 Gbps, and possibly to even 40 Gbps. However, they are differentiated by less stringent reliability requirements. For example, the end users can withstand some level of outages – either unexpected or planned. Cost becomes a factor as flexible waves are typically less expensive to maintain than core services. A wide range of durations will be seen, although typically shorter than the durations of core waves. Additional requirements for quick provisioning and teardown will be necessary as well. Flexible waves could support a wide range of applications, from Grid computing clusters, IP network overflow, network research projects, highly redundant IP networks, etc. NLR's network infrastructure has been designed as a possible source for these light paths.

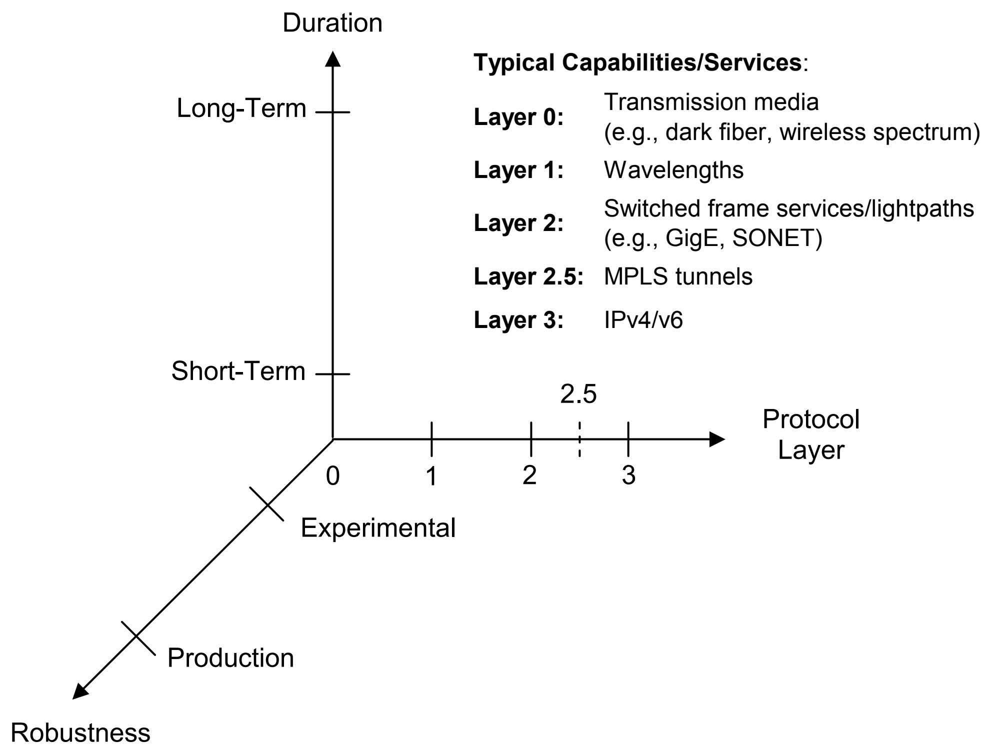


**Figure 1**: A POTENTIAL MODEL FOR SERVICES AND CAPABILITIES TO BE OFFERED IN NEXT GENERATION HIGHER EDUCATION NETWORK INFRASTRUCTURE. This diagram presents examples from three attribute types of the range of services and capabilities to be offered in the next generation higher education network infrastructure.

## 3.2. Long Term Perspective (3-5 years)

As specified in its charge, the primary focus of Group A's analysis is focused on the near term, through June 2006. However, the charge invites a longer-term view, and Group A believes the horizon holds strategically important cues critical to setting direction and helping us respond to uncertainty and challenges ahead.

We have before us an unprecedented opportunity. For the first time, our community is securing its own national fiber optronic infrastructure, enabling us to exploit the raw materials of advanced communications technology with a degree of control and reach never before available to the research and education community. This extraordinary resource enables a fundamental shift in the way we think about networking and provisioning networks. No longer confined by narrowly defined market priced services from commercial providers, we can move beyond the notion of "a network" or "the network" and into the domain of "the virtual network," akin, in the context of computers, to the virtual machine. The virtual network isn't really a network *per se*, but rather a programmable research, development, and provisioning platform comprising a dynamically configurable collection of communications resources that can be systematically cast and recast, interconnected and partitioned, in response to need.

Appendix B elaborates on this principle and provides illustrative examples of longer-term hypothetical architectures and applications.

## 4. Organizational Environment – Financial Considerations

Consistent with the charge to Group A by the NLR Board, this document focuses on the range of technical services and infrastructure requirements which may either facilitate or hinder convergence in the development and management of advanced communications capabilities serving the interests of R&E. However, Group A is extraordinarily sensitive to the probability that competing and non-complementary financial strategies among NLR, Internet2, and both the RONs and gigapops pose the most immediate and long-term threat to the viability of these organizations critical to the advanced network interests of the R&E community

Just one example of a potential high-tension financial issue is the use, and proposed use, of membership fees as the primary revenue source underpinning Internet2, NLR, and the RONs and Gigapops. An uncomfortably high percentage of NLR participants viewed the annual commitment of one million dollars over a period of five years as an investment which would lead quickly to direct and measurable cost savings, such as the elimination of Internet2 related membership and services fees. Likewise, there are Internet2 members who are not supportive of the NLR vision, who believe Internet2 can meet any advanced communication network need, and who will pressure Internet2 to aggressively compete with NLR in providing access to equivalent services and infrastructure. In the context of such pressures, it is a credit to the leaders of NLR and Internet2 that they recognize the interdependencies of the two organizations, and the extraordinary shared opportunity in merging their respective visions and coordinating development plans. Internet2 is organized, and operates, like a membership organization. NLR, by design, does not operate like a membership driven organization. The R&E community, over the long run will not support two national organizations with membership fees at the level required to develop and support the array of services and capabilities identified in the spreadsheet accompanying this report. However, the R&E community has significant capacity to fund a vast array of services, where those services are differentiated to meet the variable demands of this community, and where the pricing of specific services reflects the full direct, indirect, and capital costs of providing the services.

We have an unprecedented opportunity to develop a combination of revenue and financial strategies that leverage the base infrastructure being developed by NLR, Internet2, and the RONs, strengthen all of these organizations, and advantage the R&E community. However, first, Internet2 and NLR must implement financial strategies and mechanisms that are complementary and coordinated. Ultimately, the development of sound financial plans reflective of the services and requirements array identified in this report will drive the future organization structure and relationship of NLR, Internet2, and the RONs.

## APPENDIX A: Charge to Group A

Below is the specific charge given to Group A by the NLR board and endorsed by the NPPAC.

Over the next 18 months,

- Internet2 will renew and enhance the Abilene high-performance network service, and begin expanding beyond it;
- NLR will expand beyond its current services focused on specific research connections; and
- Regional networking entities will expand their services and membership in ways that benefit from and influence NLR and Internet2.

We must plan the evolution of the high-performance networking environment over the next few years. During this period, NLR will move from being a startup aimed solely at providing experimental and breakable network infrastructure to researchers. Internet2 will move beyond its current focus on production layer 3 services. Regional entities will become ever more complex and important links in the overall networking fabric. The whole must deliver robust, reliable networking for production efforts such as Abilene, ESnet, and other national services, and accessible, flexible, and affordable networking for research.

We need technical and infrastructure convergence rather than divergence, particularly between Internet2 and NLR. To satisfy this need, it is essential that design of the entities' networks, infrastructure, and operating practices proceed in conscious recognition of the need for convergence. Internet2 and NLR both prefer that expanded Internet2 and NLR services use each other's infrastructure and other capabilities, and seek to avoid unnecessary duplication of capacity, technology, staff, and other resources between them or with regional entities.

To this end, Group A, as knowledgeable individuals rather than representatives of specific organizations, will do following:

1. Jointly identify and document current Internet2 and NLR service plans through June 30, 2006, and beyond if possible.

2. Where the two organizations' service plans might rely on common infrastructure or operating resources, identify the specific technical and operational requirements that would enable the services to achieve this end.

3. Where current infrastructure and operations practices fall short of service requirements, propose changes or enhancements that would close the gap and enable convergence.

Group A members: Erv Blythe (Virginia Tech), Javad Boroumand (Cisco Systems), Steven Corbató (Internet2), Wendy Huntoon (committee chair; Pittsburgh Supercomputing Center), Ron Johnson (University of Washington), Michael Krugman (Boston University), Richard Summerhill (Internet2), and Douglas Van Houweling (Internet2)

## APPENDIX B: Longer Term Perspective

As specified in its charge, the primary focus of Group A's analysis is near-term — through June 2006. However, the charge invites a longer-term view, and Group A believes the horizon holds strategically important cues critical to setting direction and helping us respond to uncertainty and challenges ahead. This section elaborates on this principle and provides illustrative examples of longer-term hypothetical architectures and applications.

Fundamentally, networks occupy three physical dimensions:

1. *Space*—two or more parallel or tangential wires, cables, and/or fibers;
2. *Time*—fixed slots (e.g., ATM) or dynamically interspersed packets (e.g., IP);
3. *Frequency*—radio or optical.

Communications technology exploits these dimensions in order to maximize two essential parameters or goals:

1. *Connectivity*—pervasively, among people and resources;
2. *Speed*—in establishing connectivity and in the exchange of information (subsuming high-bandwidth and low-latency).

The idealized upper bounds on these parameters are:

1. *Connectivity*—everyone/thing
2. *Speed*—instantaneous/light-speed

Network evolution toward these ideals is driven by incentives (real or perceived value), as well as reach and availability (pervasiveness and cost), speed (bit rate and bandwidth), and sophistication (logic and intelligence) of components available to process signals and distribute, exchange, and deliver information.

With that distillation of the physics of our universe as backdrop, we have before us an unprecedented opportunity. For the first time, our community is securing its own national fiber optronic infrastructure, enabling us to exploit the physical dimensions of this universe—space, time, and frequency—with a degree of control and reach never before available to research and education. High-speed time division multiplexing (e.g., 10Gbps/40Gbps layer 3 packet switching and routing) promises to vastly expand the quality, quantity, and integration of services and the interconnectedness of our community. Dense wave division multiplexing will multiply these channels by orders of magnitude, allowing us to expand capacity and trial and transit to experimental and advanced services efficiently and gracefully via resources and technologies never before directly available to our communities. Space division multiplexing (e.g., scores, hundreds, and even thousands of fibers bundled together and distributed pervasively throughout the nation and the globe) offer even greater flexibility, multiplying capabilities yet again, and promising essentially unlimited capacity in the foreseeable future.

This extraordinary expansion in technologies and facilities suggests a fundamental shift in the way we think about networking and provisioning networks. No longer confined by narrowly defined market priced services from commercial providers, we can move beyond the notion of "a network" or "the network" and into the domain of "the virtual network," akin, in the context of computers, to the virtual machine. The virtual network isn't really a network per se, but rather a programmable research, development, and provisioning platform comprising a dynamically configurable collection of communications resources that can be systematically cast and recast, interconnected and partitioned, in response to need.

The comparison between the virtual network and the virtual computer is a useful one. A computer is a collection of facilities and resources (e.g., computational elements, memory systems, storage hierarchies, communications devices, and applications) closely integrated and made manageable by an abstraction layer called an operating system. A virtual machine imposes another abstraction layer, creating the illusion of multiple machines. Our evolving national fiber optronic infrastructure offers an opportunity to impose just these sorts of abstraction layers on top of this new communications platform, providing multiple network partitions and operating environments, complete with APIs, end-user interfaces, authentication and authorization mechanisms, access methods, schedulers and priority schemes, and a highly integrated yet vastly distributed network operating system supporting research, development, and production services.

What sort of applications might such architecture serve? We're all familiar with discipline specific, high-bandwidth requirements associated with the high energy physics and astrophysics communities, calling for the distribution of vast datasets associated with collider experiments and telescope arrays. Research into grid computing and other advanced computing architectures (e.g., OptIPuter and TeraGrid) require dedicated optical backplanes spanning geographic expanses. Other unprecedented demands for network resources are anticipated in biosciences, climatology, and even entertainment. Instead of revisiting and recounting these, let's consider instead an application that might serve them all: Telepresence.

Imagine a future generation of Access Grid-like facilities in which any number of participants can meet remotely using high definition desktop video conferencing technology. However, instead of participants appearing in discrete windows representing their specific context and locations, image processors extract them from their backgrounds in real time and place them in a shared, virtualized environment. All participants are materialized in this virtual environment, allowing them to occupy a common space more conducive to interaction and exchange than a shared audio space alone. Using tracking technologies (head and eye movement), each participant's view of this space is from his or her perspective. Because the meeting space is virtualized, an array of capabilities unavailable in the real world are now possible, including, through linkages to discipline specific applications, virtual objects such as molecular models, atomic structures, and research prototypes, as well as visualizations drawn from databases associated with the aforementioned big-science research initiatives, including immersive representations of collider, astrophysical, or genetics datasets from network-attached instruments and storage arrays.

Clearly the network resources, computational requirements, and the degree of integration and resource management associated with such an application are staggering. Here's where the notion of a virtual network operating system comes to the fore. Vast amounts of bandwidth must be allocated and bound to computational elements dynamically and under the systematic control of an array of interrelated applications. Participants may come and go, demanding and releasing vast amounts of bandwidth on demand. High-definition audio and video feeds from and to each participant and among visualization, virtualization, and storage elements require network resources and capacity on a per event basis that dwarf current requirements. An application of this nature requires comprehensive end-to-end performance, conformance, and interoperability across the entire fabric of the network.

How do we move effectively and efficiently toward such a horizon? Is consistency and conformity in the evolution of our global optical infrastructure an immediate imperative?

Progress along our networks' evolutionary timeline is marked by periods of scattering and gathering, experimentation and convergence. Experimentation is healthy and necessary and allows new approaches to be developed, tested, and proven or abandoned concurrently. Convergence represents seminal milestones, at which point specific approaches crystallize and achieve critical mass, and then standards are established, conformity ensues, and uniformity, consistency, and interoperability prevail. Examples of seminal milestones include the development of packet switching and routing; wire, wireless, and fiber technologies and standards; Ethernet; the Internet Protocol; and wave division multiplexing. Counterexamples of approaches pursued, deployed in considerable scale, and largely set aside include coaxial cable, DECnet, SNA, and Token Ring.

As we consider where we are today and where we're headed in this scheme of things, we should recognize the practical realties inherent in this evolutionary process and the need to strike a practical and effective balance between experimentation and convergence. It may, for example, be unnecessary and even unfortunate for us to deploy inconsistent core technologies across our national backbone(s) (e.g., DWDM), but we'll have less influence and certainly less control over what happens as optical networks extend internationally, regionally (RONs), into metropolitan areas (MANs), and onto campuses and into communities (LANs). This may place limitations on end-to-end capabilities, and, at the outset, will certainly increase complexity and costs, but it will also provide diversity in the evolution of next generation technologies and techniques, the subsequent survival of the fittest, and, ultimately, consensus, consistency, and continuity founded in proven techniques and best practices.

# APPENDIX C: NLR and Internet2 Aggregate Service Plans

| Capability or Service | Attributes | NLR and Internet2 Services<br>Including Layer 1, 2, 3 Capabilities + Abilene, HOPI, FiberCo Services | |
|---|---|---|---|
| | | **Currently Committed**<br>(Now to 6 Month) | **Potential Future Plans**<br>(12 to 18 Month Projection) |
| Fiber & Colo Facility | **Implementation** | Contracts and cooperative relationships with specific vendors, including FiberCo and NLR specific arrangements. | Expansion to additional carrier contracts as well as into the aggregate metro fiber relationship. |
| | **Reach**<br>Geographical Coverage | All participating carriers' footprints, including NLR-specific footprint. | Expansion based on end-user requirements |
| | **Cost** | Typically per contract, but NLR foot print, first 6 months is free for experimental use. | On major carrier footprint, price standardized through March 2006. |
| Wave Service | **Implementation** | Facilities based - fiber and long haul optronics ownership. Carrier managed wave services available as needed. | Complementary relationship with carriers for managed services as needed. |
| | **Production** | Production with 2-3 nines availability. | Improved availability to 3-4 nines; experimentation with GMPLS and inter-domain provisioning on a separate testbed or in labs. |
| | **Experimental** | Experimental Services available. | Options for lower cost waves with reduced availability commitments. |
| | **Architecture & Network Design** | DWDM linear systems. | Addition of OXCs; early experimentation with L1-L3 integration and virtualization and associated network management. Carrier operated DWDM linear systems with potential for overlap of OXCs and other switching equipment by higher ed. |
| | **Service Features** | Point-to-point unprotected waves, for durations of 1 year or longer. | Addition of multi-levels of services (different nines for different price); diverse-path protection and short duration (weeks, months) waves; GMPLS dynamic provisioning and "lambda switching"; layer-1 multicast through OXC. Point-to-point unprotected & protected variable duration waves. |
| | **Reach**<br>Geographical Coverage / # of PoPs | Tied to NLR backbone, "nationwide" at minimum 28 PoPs in 20 states. | Service dependent. Few new PoPs for sparse backbone configurations, expansion is left to the regional aggregation points. Other services, nationwide with typically 50-60 PoPs. |
| | **Speed & Capacity** | 10G, 32-40 waves per segment. | 40G, 10G, and 2.5G as appropriate. |
| | **Cost** | $85M over 5 years plus per-wave cost. | Dependent on market conditions. |
| | **Interconnection & Peering** | NA | Campus, regional, national and international GMPLS-based interconnect and peering (overlay the carrier network). |
| | **Instrumentation** | None | Layer-1 monitoring tools. |
| | **AUP** | No AUP | No change |
| Ethernet Switched Frame Service | **Implementation** | Ethernet switches ownership, using NLR waves. Includes Internet2 HOPI testbed and NLR Layer2 GigE service. | Evolving architecture for testbed. |
| | **Production** | Experimental at first, moving towards production. | Production baseline service available. |
| | **Experimental** | Available | Separate experimental testbed. |
| | **Architecture & Network Design** | Partial mesh Ethernet switching. | Early experimentation with L1-L3 integration and virtualization. |
| | **Service features** | Point-to-point "circuit-like" Ethernet channels and over-subscribed channels, dynamic provisioning. | Point-to-multipoint channels; Ethernet over MPLS; dynamic provisioning (e.g., UCLP). |
| | **Reach**<br>Geographical Coverage / # of PoPs | Nationwide, 16+ PoPs. | Footprint expansion with addition of new waves. |
| | **Speed and Capacity** | 10G trunks and 10G/1G user ports | No Change |
| | **Cost** | Varies depending on service level and backbone infrastructure. | TBD |
| | **Interconnection & peering** | Regional, national and international MPLS/GMPLS-based interconnect and peering (e.g, GLIF). | Regional, national and international MPLS/GMPLS-based interconnect and peering (e.g, GLIF). |
| | **Instrumentation** | Network dependant, layer-2 monitoring available. | Layer-2 monitoring. |
| | **AUP** | No AUP | No change |
| IP Routed Packet Service | **Implementation** | IP routers ownership; using carrier or NLR waves depending on the specific service. | Addition of carrier managed waves if needed to achieve desired availability. Upgrade of routers as desired. |
| | **Production** | Advanced production; general-purpose R&E backbone- roughly 4.5 nines (support experiments - e.g., reduced buffer size - and deployment of advanced services). | Production general-general purpose, production special-purpose/limited scope production at 4 nines. |
| | **Experimental** | Starting experimental and moving into both experimental and limited-scope production (e.g., subset of regional/campuses/subnets to have real applications traffic to drive experiments). | Experimental on same platform as production. |
| | **Architecture & Network Design** | Partial mesh IP routing with IGP. | Early experimentation with L1-L3 integration and virtualization; separation of edge/core routers. |
| | **Service Features** | IPv4 unicast/multicast, IPv6 unicast/multicast, MPLS tunnels, experimentation with logical/virtual routers; Research facilitation - Observatory (data access and experiment co-location). | Addition of logical/virtual router to baseline. |
| | **Reach**<br>Geographical Coverage / # of PoPs | Nationwide with Qwest or NLR backhauls dependent on specific backbone and service. | Addition of new PoPs with edge routers. |
| | **Speed & Capacity** | 10G trunks and 10G/2.5G/1G/lower speed user ports available. | Addition of 40G trunks and 1G user ports. |
| | **Cost** | Dependent on backbone and level of service. | No change |
| | **Interconnection & Peering** | Regional, national & international BGP interconnect & peering for experimentation. | Addition of production peering; addition of commodity Internet peering. |
| | **Instrumentation** | Layer-3 monitoring; Abilene Observatory. | No change. |
| | **AUP** | Infrastructure dependent. Ranges from AUP-free to Abilene COU. | No change |
| Operations & Support | **Traditional NOC service Layer1** | Provided | Carrier provided available. |
| | **Traditional NOC service Layer 2/3** | Provided | No change |
| | **Experiments & Testbed Support** | Provided | No change |
| | **Coordination & Collaboration** | NLR and Internet2 organizations through various initiatives (e.g., middleware, e2e performance) councils and WGs. | Expansion of organization to scale as new production services and new experiments are added. |