Internet2 Network

Overview

20070213

Executive Summary

The U.S. research and education community formed Internet2 in 1996 to help meet its advanced networking needs. For the past ten years the same community has deployed high-performance networking infrastructure at the campus, regional, and national level, and has collaborated to develop new capabilities in areas such as middleware and network performance measurement. The result has been a leading-edge and highly-reliable nationwide networking environment that supports the most demanding high-performance applications and experimentation in research, discovery, and teaching

Today, leaders in disciplines as varied as high-energy physics, radio astronomy, bioinformatics, network research, and musical performance have begun to look beyond advanced networking capabilities available today. Based on these emerging needs, in 2005 R&E networking leaders developed a shared vision of next-generation networking technologies, including new hybrid optical and packet networking capabilities. In 2006, ten years after it was created by the institutions it serves, Internet2 began deploying a new nationwide infrastructure to fulfill this technical vision and to meet the community's needs. The first link of the new network was operational in December 2006 and deployment continues on track for completion by the beginning of the 2007 academic year.

Through corporate partnerships, the Internet2 Network leverages the very latest optical networking technology, offering a major step forward in cost-effectiveness, capability and flexibility to provide a community-controlled, highly scalable nationwide optical infrastructure. These partnerships also enable the Internet2 community to focus its resources on developing and deploying advanced capabilities, rather than maintaining the underlying optical infrastructure. Internet2 is working closely with regional and state networking organizations as it deploys the nationwide network.

The network will extend over more than 13,000 miles of dedicated fiber, with complete community control of the optical layer and highly granular "lightpath" capabilities that can be provisioned dynamically across multiple domains. In addition to high-performance IP, the network will provide short-term and long-term dedicated wavelengths, and eventually, on-demand or advance reservation lightpath scheduling at the campus level. The network will build on the community's efforts in middleware, control plane, and end-to-end network performance technologies, as well as experience from the highly successful Internet2 Hybrid Optical and Packet Infrastructure (HOPI) testbed.

These capabilities and its extensive reach make the network an ideal platform for large-scale science, grid computing, remote instrumentation, tele-immersion, and tele-medicine. In its inaugural use, the first complete link of the network flawlessly supported an uncompressed, 1.5 gigabit-per-second HD videoconference from New York City to the Fall 2006 Internet2 Member Meeting in Chicago. And, the Internet2 network already has been selected by the U.S. Department of Energy as the backbone for ESnet*4.

The network also extends Internet2's commitment to, and ability to support, network research. The Internet2 Observatory, which now supports more than a dozen projects, including PlanetLab and the NSF-supported WAIL: The Wisconsin Advanced Internet Laboratory, will be expanded. With more extensive network monitoring and data-collection capabilities built into the network as it is deployed, even more comprehensive datasets and performance measurements will be available to network researchers. Network research projects that require dedicated facilities, such as the NSF's GENI, can be supported easily.

Through technical workshops and outreach, Internet2 is working with regional and state networks as well as individual member institutions to deploy and extend new capabilities offered by the network. Internet2 also has worked with the community to develop a fee structure that recovers fixed costs in a way that ensures its continued sustainability, distributes costs equitably among connectors and participants, and sets a predictable and solid foundation for future investment in new technologies.

The Internet2 Board is taking a number of measures to support community involvement in the use and development of the network. First, it is establishing a new category of members for regional and state networks. Second, a committee of the Board of Trustees last year engaged the community in a comprehensive review of the Internet2 governance structure, and their recommendations are rapidly moving towards implementation. Third, the Board will begin considering later this year a further improved membership structure that will ensure any interested institution is able to participate in advancing the use of networking in research and education.

With these changes, Internet2 will be able to support and contribute to new collaborations and partnerships between the regional networks and individual institutions. The universities, companies and other organizations that are at the heart of the Internet2 community will have an even more direct role in the organization that provides advanced networking for their faculty, staff, and students. Internet2 is committed to this inclusive membership principle today, just as strongly as it was at its founding. And Internet2 will remain committed to it into the future.

TABLE OF CONTENTS

OVERVIEW

Internet2 is working with regional networks and university campuses as well as corporate and research partnerships to provide the U.S. research and education community with a dynamic, innovative and cost-effective hybrid optical and packet network. The new Internet2 Network is designed to provide next-generation production services as well as a platform for the development of new networking ideas and protocols. With community control of the fundamental networking infrastructure, the network will enable a wide variety of bandwidth-intensive applications under development at campuses and research labs.

Through agreements reached with Level 3 Communications, Infinera, Indiana University and Ciena, the Internet2 Network provides multiple services, across multiple domains, at multiple speeds, over one seamless infrastructure. Extremely scalable and flexible in bandwidth, the new network will enable advanced applications and collaboration among academia, research and industry across the nation and globally. Key attributes include:

- National footprint - 13,000-miles dedicated optical fiber backbone provides low latency and diverse optical paths for reliability
- Rich services - dynamic and static wavelength services along with existing enhanced IP capabilities
- Large capacity - 100 Gbps capacity along entire network footprint initially
- Scalability – unlimited availability of additional wavelengths, without large capital expenditures
- Significant economies of scale - leverages Level 3's optical expertise and capital investments for maintenance and sparing
- Dense topology - Infinera's unique photonic integrated circuit architecture supports cost-effective add/drop capabilities, provisioning and redundancy
- Network efficiency - Ciena CoreDirector, high capacity multiservice switching system, provides capability for sub channels using Ethernet or next-generation SONET services
- Carrier-class reliability - Level 3 has the organizational and procedural structure in place to ensure network reliability, offering production-level support (service level agreements) on its layer 1 wavelength services. The Infinera PIC platform and the Ciena multi-service switch have exemplary service reliability records.
- Reconfigurability - Dynamic circuit provisioning within seconds and, in the future, advance reservations of circuits to support an even broader range of applications and research
- Network flexibility - Wave provisioning in sub-wavelength increments of 50 megabits per second (Mbps) to multiple 10 gigabits per second (Gbps) wavelengths, to allow more efficient bandwidth allocation
- Support for research - DWDM capabilities allow the creation of logical networks over the same fiber facility for production or experimental uses
- Community control – complete control over the provisioning of the network

The new network continues to provide IPv6 connectivity and multicast capabilities with performance identical to that provided for IPv4 unicast – supporting new modes of data distribution and interoperability with increasingly IPv6-based international networks, as well as global addressability for very large numbers of nodes that may be needed for large sensor networks or similar applications. The network leverages work by the Internet2 community in the areas of control plane technologies, middleware, and network performance measurement to develop and deploy new circuit-based capabilities that can dynamically provision and use multiple dedicated wavelengths on the nationwide footprint.

The network offers a number of other capabilities intended to add significant value to members while bringing new technologies to the commercial Internet. These capabilities and services include higher networking speeds (40 Gbps and 100 Gbps support), and options for connection to the commercial Internet for both peering and low-cost transit with commercial internet services and content providers.

MOTIVATIONS

Science and research needs

Over the past few years, the networking demands of many research and education communities have expanded tremendously. Massive data streams generated by complex instruments, the requirement to integrate geographically disparate information storage, processing and visualization capabilities, real-time remote operation of scientific sensors, high-definition collaboration facilities, and network-accessible computational clusters are all driving exponential growth in the networking needs of academia. The constituencies supported by Internet2 members – researchers, faculty, students – are increasingly looking for new networking capabilities to support their most important and ambitious work. Consequently, networking demands have led to the proliferation of parallel purpose-built network infrastructures. Additionally, networking advancements overseas have pressed the U.S. research and education communities to develop similar capabilities for the sake of global interoperability and collaboration among researchers and scholars everywhere.

Fields such as high-energy physics, bioinformatics, astronomy and earth sciences are among the leaders, but the movement is growing to include nearly every discipline. Network researchers engaged in projects such as GENI require large-scale testbeds to experiment with fundamentally new network protocols. Even more broadly, high-fidelity real time collaboration tools used in teaching and learning among not only universities, but also K20 institutions that connect through state and regional education networks, are driving new networking requirements. The Internet2 Network will have the capabilities to meet current and future networking demands, as well as integrate duplicate network efforts into a shared, cost effective community platform.

The objective is to link computational clusters, storage arrays, visualization facilities, remote sensors, collaboration sites and other instruments into globally distributed and application specific topologies. These advanced network services are motivated by the realization that science benefits when expensive resources such as radio telescopes, powerful computational clusters, and large data repositories can be easily shared among researchers regardless of location. Other communities and applications whose interest in these type of services is anticipated include emergency response, mission/business critical services, and building (or traffic engineering) of a best effort IP network.

Technical vision and capabilities

Internet2's 1997 agreement with Qwest anticipated the need for a dedicated fiber pair in support of the growing needs of the community. Exploring solutions to that need, in the fall of 2002 Internet2 joined a consortium of U.S. research universities and private sector technology companies to form National LambdaRail, Inc. (NLR), a not-for-profit corporation dedicated to building a national scale infrastructure for research and experimentation in networking technologies and applications. Internet2 committed $10 million over five years to the NLR effort. Internet2 used one 10 Gbps wavelength on the national NLR footprint, which in conjunction with the Abilene Network, became the testbed for the Hybrid Optical and Packet Infrastructure (HOPI) Project to explore advanced networking architecture and systems. http://networks.internet2.edu/hopi/

The development of a fiber-based next-generation infrastructure for research and education was a critical component of Internet2's strategic plan for 2004. Given the availability of dedicated optical fiber at the national, regional, and campus geographic scales, and Dense Wave Division Multiplexing (DWDM) optical equipment that enables multiple channels of light on a single fiber optic cable, network designers have been looking at optical capabilities to supplement IP networks. The HOPI Project, launched in 2004, provided rich experience and insight within a hybrid network of shared IP packet switching and the aggressive use of dynamically provisioned optical lambdas.

Formative reports published in the spring of 2005 by community advisory groups set the vision of Internet2's future network requirements. Group A, formed by the NLR Board of Directors and the Internet2 Networking Policy and Planning Council, stated the following vision in its Group A Report http://www.internet2.edu/network/library/GroupAReport2005.pdf:

> 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.

The Abilene Technical Advisory Committee (TAC) produced a requirements document describing the essential features for the next generation network. http://www.internet2.edu/files/AbileneTAC-Next-Gen-Requirements.pdf.

With an impending deadline to either continue the existing Abilene network through 2008 or transition to a new network to meet the evolving needs of the community, in January of 2006 Internet2 began a process to examine the possibilities for a new network.

Working from the vision of the Group A Report, Internet2 focused on moving to a network infrastructure that would continue to provide the capabilities of the existing IP network and also deliver point-to-point optical circuits and dynamic wave capabilities. Internet2 discussed the possibilities of a hybrid optical and packet infrastructure with a variety of different national optical carriers and equipment vendors. Several viable alternatives were developed, but as the discussions took place one was clearly less expensive and more capable: a Dedicated Wave System with Level 3 Communications. Indeed, many Internet2 members and connectors had built their networks on Level 3 fiber and into Level 3 facilities.

Cyberinfrastructure and the Internet2 Network

Beyond Internet2, the networking vision of the broader academic community, including the National Science Foundation (NSF), has coalesced around a concept of cyberinfrastructure (CI). First proposed by the report produced in 2003 by a Blue Ribbon panel convened by the NSF, "cyberinfrastructure refers to infrastructure based upon distributed computer, information and communication technology… required for a knowledge economy." Further, the report recommended that NSF fund "a major Advanced Cyberinfrastructure Program (ACP) to create, provision, and apply advanced cyberinfrastructure to advance, and ultimately revolutionize, the conduct of scientific and engineering research and allied education. Success for this far-reaching ACP will require synergy among constituencies with varied expertise as well as incentives for participation."

The NSF **Cyberinfrastructure Vision for 21st Century Discovery**, published in July 2006, articulated, "Cyberinfrastructure integrates hardware for computing, data and networks, digitally-enabled sensors, observatories and experimental facilities, and an interoperable suite of software and middleware services and tools." (http://www.nsf.gov/od/oci/ci-v7.pdf) This vision has engaged the Internet2 community, and will be furthered with the capabilities and applications potentials of the new network.

Informed by its membership and working in conjunction with collaborators (e.g. Open Science Grid, TeraGrid, discipline groups such as the Large Hadron Collider [LHC] project, the GÉANT2 project, ESnet, etc.), Internet2 collaborates with the community to offer, and in some cases develop, services and technology components needed for a coherent CI suite. In order for such CI to work, it has to be a workable end-to-end system; Internet2 emphasizes a systems approach towards CI integration. Internet2 does both systems integration work, assembling open source communication tools into a common veneer, and develops new open source software components, as needed. The resultant cyberinfrastructure suite is being deployed on the new network and offered to our connectors (e.g. regional networks) and partners (e.g. international partner networks, federal networks).

The strategy is to start locally (on the new network) and then move remotely, focusing on campus, national, and international CI. Internet2 is specifically targeting application-community CI (quasi-national or international) such as LHC. Generally, we seek the ability to effectively use authorized resources regardless of where they exist and to integrate new resources as they become available. This assumes a landscape with multiple, autonomous CI providers that need to work interoperably (assuming the autonomy of regional networks and our international partners). The basic tactic is a toolkit approach that looks like a wall jack to the end user but allows independent administrative entities to set policies locally.

Internet2 also plays the role of community CI coordinator. Partnered with other community coordinators such as TeraGrid and EDUCAUSE, we play a convening function in order to facilitate the development, use, and dissemination of CI (e.g. Bridging the Gap). We will facilitate conversations among various federal agencies (e.g. DOE, NSF, NIH), each of which is developing its own CI, and present a consistent vision back to the campuses. The objective is to work with campuses to build valuable CI. The goal is for campus CI to fit seamlessly into national CI, so that campuses have a clearer path to determine and accomplish their technology and networking opportunities.

Internet2 also leads in international outreach efforts at several different layers of CI (e.g. control plane, federated trust, measurement). DANTE, Internet2, CANARIE, and ESnet meet biannually to coordinate cyberinfrastructure development efforts. The goal of these efforts is to engender and coordinate area-specific software development collaborations that pool engineering resources and ensure seamless communication between jointly developed and deployed cyberinfrastructure systems.

NETWORK CAPABILITIES

Network services for the Internet2 community result from joint efforts by the whole community – our vendor partners, the Internet2 Network Operating Center (NOC), regional networks, and the campuses. Internet2 provides many of the capabilities (e.g. backbone, middleware, performance management, and security) and the organizing effort required to enable the community to leverage these capabilities into services. A key advance of the new network is the wide range of capabilities over the same infrastructure. While in some cases, similar services already exist, the technology used in the Internet2 Network allows the community to create the services it needs from capabilities that are scalable, more flexible, and more cost-effective. What follows is an overview of the capabilities that will be available on the network.

IP Capabilities

The IP network (Layer 3) will be built on top of the optical systems and initially provide 10 Gbps of v4 and v6 IP services, supporting both unicast and multicast. The highly reliable Juniper T640 routers, which continue to be state of the art and have the ability to scale to 40 Gbps per interface, will be used to implement this service. The network architecture provides a variety of protection options at the optical layer, which further ensures a highly reliable network.

Current high-performance IP networks are very efficient and effective for most types of network traffic, specifically smaller flows and multipoint connectivity that characterizes most of the network usage today, typified by lower-rate data transfers (roughly 1 Gbps), Web browsing, email, instant messaging, and the like. The current IP network, however, is not as well suited to traffic that: consumes large amounts of bandwith for long periods of time, is very sensitive to network performance, or utilizes network protocols other than TCP/IP.

Therefore, while the new Internet2 IP network will have an initial capacity of 10 Gbps, it will have the potential near term option of 40 Gbps, and the possibility of 100 Gbps capacity. Furthermore, the underlying optical infrastructure provides a foundation for scaling the capacity IP services well beyond any existing nationwide infrastructure dedicated to R&E. The Internet2 network also will provide an additional capability for serving network traffic that is not easily accommodated within the IP infrastructure, direct access to optical circuits

Dedicated Circuit Capabilities

The next step in the evolution of networking that will satisfy the needs of the leading-edge researchers and the most demanding applications is the ability to provision dedicated circuits on both a temporary and a persistent basis. An example of the types of traffic that would benefit from circuit-based capabilities would be large e-science data transfers, high-performance high-definition video conferencing, and instances where a high degree of network security is required. The Internet2 community views these two types of network capabilities, IP and circuit based, as complementary and mutually beneficial to the users of the network resources. The new network makes the advantages of both available to the user as the circumstance requires. The key point is flexibility, the flexibility of the network to provide services to users that best fit their needs.

- The Dedicated Circuit capabilities of the Internet2 Network will provide:
    - Short term dynamically provisioned services – granular services (for example, 1 GigE connections) providing sub-channel capabilities. These would have durations no longer than a day, with restrictions on monopolizing connections. Scheduled services are expected to be available through these capabilities to support connections such as CERN's Large Hadron Collider (LHC) Tier 2 distribution centers and eVLBI (electronic Very Long Baseline Interferometry). Depending on demand, one or two 10 gigabit waves on the backbone are expected to be dedicated to this service. Call blocking can be expected on this service, except for scheduled services.

- Long term dynamically provisioned services – provisioned through the point-to-point connections, on a different wave from the dynamic services described above. An additional cost will be associated with these services, determined by bandwidth and duration, but there will be no call blocking associated with such services.
- Long term static, full wave services – provisioned directly through the Infinera gear at an additional cost and reserved for one year or more. This type of service would advantage a GENI-like network, and it is the method by which ESnet is implementing its next network along the Internet2 backbone

- Additional Services, provisioned on the Level 3 infrastructure (*off-net* services).
  - Dedicated circuits cost-effectively provided by “WaveCo” on a yearly basis for RONs and other members to support and extend connectivity that is not on the Internet2 footprint.

The Internet2 community is working together to develop and deploy control plane technology that allows these capabilities to be managed at the national, regional, campus, and user level so they can be delivered as a service to the end user of the network.

Commercial Peering

One of the value-added services planned for the network is commercial peering (CP). By providing direct connections at high speed with commercial service providers, commercial peering would advantage community members’ access to content and potential content exchange with commercial providers. Coexistent with the research and education network, Internet2 interface at peering locations would extend greater bandwidth and other advanced network technologies to commercial content providers.

Potential cost savings to Internet2 members – particularly campuses – is a clear motivation for this type of service. CP service is included in the base connection fees for the new network. Like IPv6 or IP multicast, there is no additional cost, and the network connector organizations choose whether to use the service to meet their individual community’s needs.

However, Internet2 believes that there are other equally compelling reasons and potential benefits for our community to establish a successful presence at the national network crossroads (i.e., the major peering facilities) alongside of the large content, service, and connectivity providers. These include:

- Preserving net-neutral access
- Promoting advanced services such as IP Multicast to content providers, and
- Facilitating our community’s position as an advanced content and service source

We are currently undertaking a peering trial using the Abilene network. We anticipate having four Abilene Connectors testing the service in the coming weeks. Indiana University continues to report a shift of up to 600 Mbps from its commodity

Internet connection to the peering service. The peering trial offers over 47,000 routes, or one third of the global routing table, and we're adding additional peer networks on a regular basis.

Internet2 is working with the Network Technical Advisory Committee (NTAC) to further develop the technical implementation of the service to better meet the requirements of its users. A relatively new feature of the beta service is the ability for a connector to control, via the use of BGP communities, with which networks to exchange routes. Participants in the trial have the ability to choose all or any set of networks to exchange traffic with via the peering service.

Commodity Internet Transit

Internet2 believes reliable and cost-effective connectivity to the commodity Internet is important. We have worked to address the community need by including commodity access in the capabilities offered by the new Internet2 network.

Based on on-going discussions with members and connectors, Internet2 is exploring a collaborative effort with regional networks that would enable commodity service from an ISP by providing transport for commodity service through a connector's Internet2 network connection.

To gather a comprehensive understanding of the community's position on this issue, Internet2 will convene a group that will include members of the NPPAC Fee Committee, the Quilt Executive Committee, and the StateNets Steering Committee to discuss and report their consensus. Internet2 will also need to investigate its contractual obligations with its vendors and potentially seek new terms that are consistent with the desired policy.

International Peering

Through partnerships with the world's major national and multi-national research and education networks, Internet2 network users have access to research, university, school, hospital and other institutions in 86 countries and growing. The hybrid characteristics and capabilities of the new network, compatible with advanced networks overseas, will facilitate the growing interest of the Internet2 community to participate in global research and education initiatives and projects.

NETWORK DESIGN

Architecture

The Internet2 Network is provisioned on a pair of fiber that is exclusive to the Internet2 community and maintained by Level 3. The dedicated fiber is not provided by an IRU but rather as part of the network system as a whole for the period of the agreement. The fact that the system lies on dedicated fiber has substantial advantages. Dedicated wave equipment currently used by the carrier will provide the waves for the system. That wave system is provided by Infinera, which supports advanced technologies with substantial add/drop capabilities and significant advantages in provisioning and redundancy. Moreover, since the system is completely dedicated to Internet2, it is possible to leap beyond the carrier's standard offerings to utilize the advanced technology provided by Infinera. For example, while the carrier may not wish to migrate to 40 Gbps waves on its production system, it is possible for the Internet2 wave system to incorporate such technologies early in the deployment. It is even likely that 100 Gbps capabilities on a single interface will be available in the next few years.

The system will also incorporate a grooming capability – the ability to provide sub channels through waves using either an Ethernet or advanced SONET infrastructure. To implement this, Internet2 has worked in partnership with Ciena to purchase and deploy CoreDirector multiservice switching systems. The CoreDirectors enable the sub channel flexibility of the new network, including the ability to provide lightpath capabilities provisioned within seconds that last for durations ranging from hours to days.

Existing connections to the Abilene IP network will be maintained and migrated to the new network. Over the next few years, the network will evolve to a simplified connectivity model, in which each fully participating RON will probably provision two 10 Gbps connections to the hybrid backbone, one for IP and one for point-to-point services. To ensure cost-effectiveness and flexibility as needed, other bandwidths will be available to support corporate connections and the network research community.

Topology

The topology consists of optical nodes connected by waves maintained by the carrier that connect to RONs and other participants. The carrier footprint primarily determines the topology of the network, but the use of the Infinera equipment allows for great variability in drop/add locations. The network design encourages aggregation at the RON layer in the hierarchy. Approximately 20 to 24 connectors are expected to support regional aggregations, but additional, lower bandwidth connections can be accommodated.

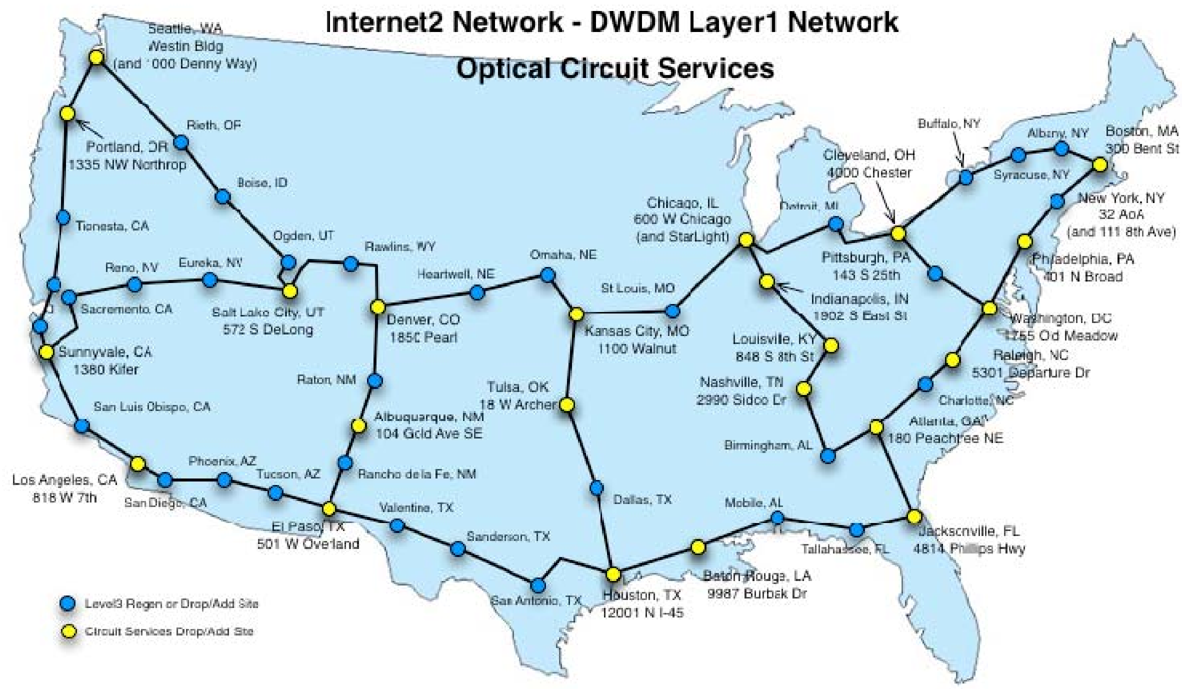

The IP network is built on top of the optical network. The IP backbone is provisioned across the waves in the system, and each RON connects to the IP network through the optical system. Since the carrier provides an SLA for the waves on the system, the IP network will be carrier-quality. Moreover, there are substantial redundancy options. The Infinera platform can provide control plane redundancy for IP connectors, and SONET restoration can be made available. These options will be determined by consulting the community as a whole. The IP network will initially use Internet2's existing Juniper routers. These routers remain state of the art and are capable of migrating to 40 Gbps interfaces.

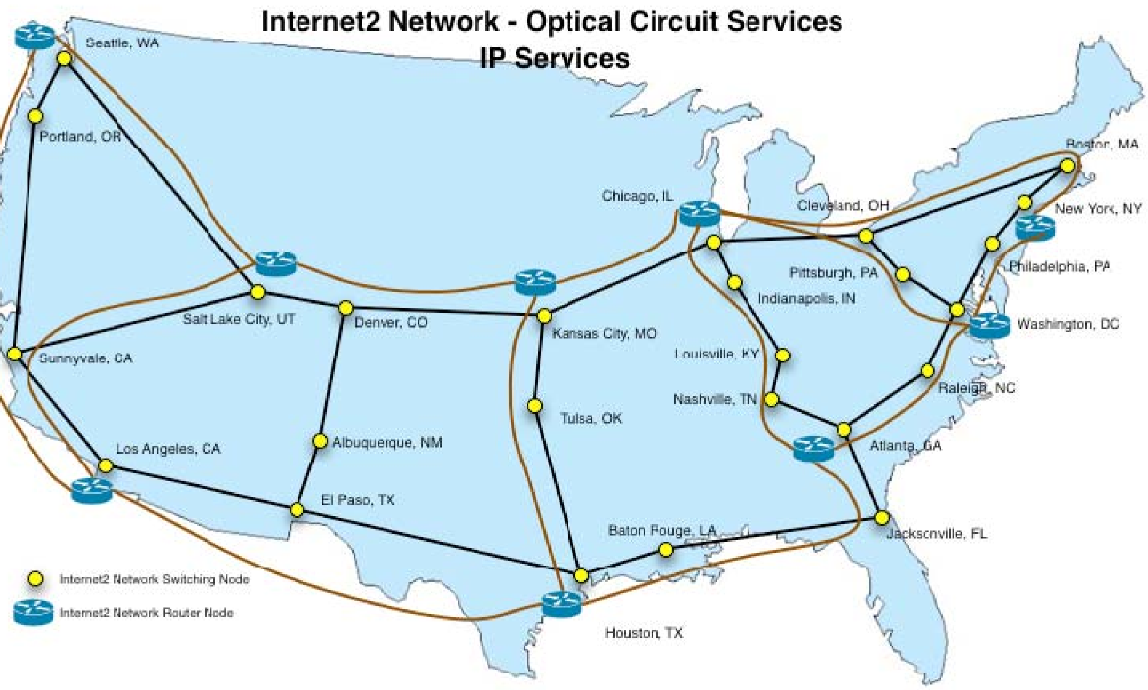

Node Architecture

The architecture of each of the optical nodes consists of the Infinera platform together with the grooming devices. The following diagram describes a sample optical node in Denver. As the network connections migrate to a standard connection model, one for point-to-point services and one for IP services, it will be possible to take full advantage of the hybrid nature of the new network. The point-to-point services connect to the grooming devices.

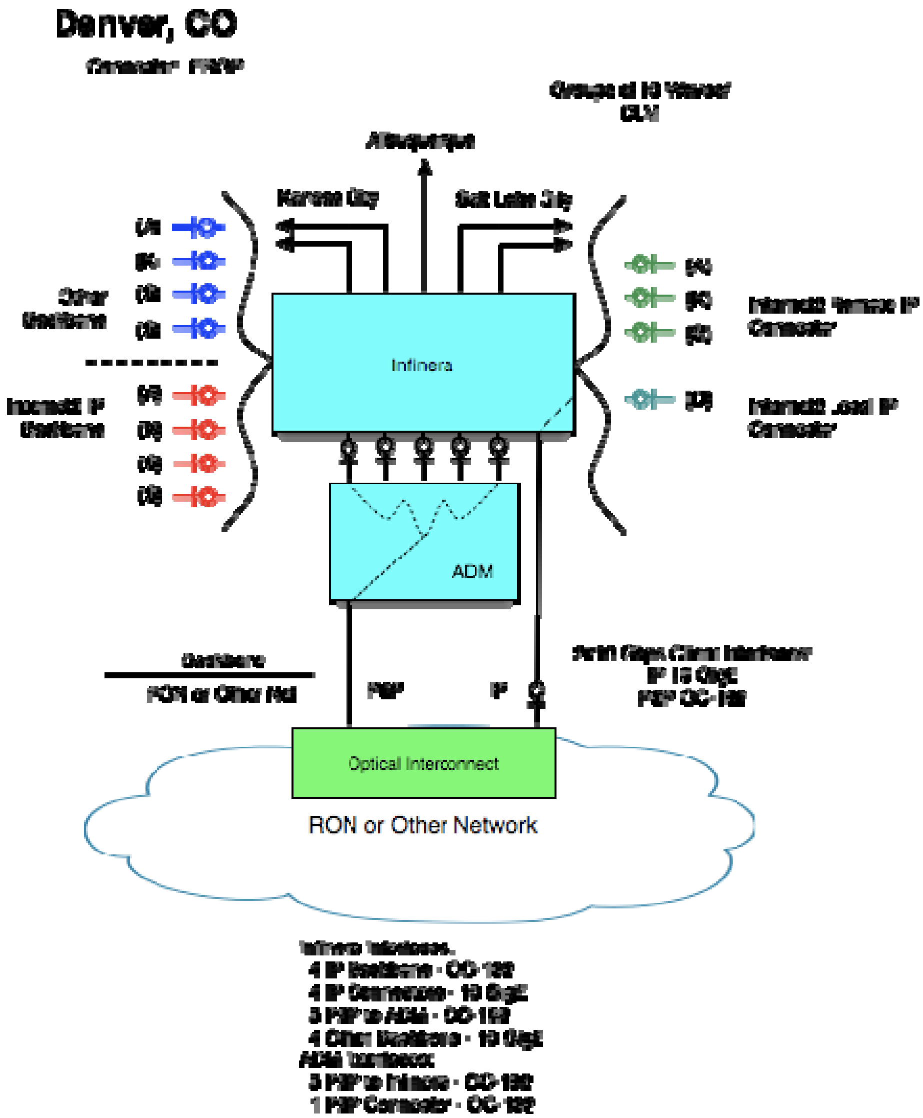

There may be special cases where it will take additional time to migrate to this model or where aggregation is difficult to achieve. However, since extensive backhaul is not necessary for the RONs, two 10 Gbps connections should be possible over dark fiber.

Control Plane

Key to the implementation of short-term dedicated circuits (DC) is the development of new control plane technology to provision lightpaths across networks. The control plane will enable Internet2 regional networks and network participants to use the circuit capabilities of the new infrastructure to control and deliver these circuit-based services. The Group A report and the Abilene Technical Advisory Committee requirements document both anticipate increasingly broad use of this service over the lifetime of the Internet2 Network.

To this end, Internet2 is supporting, along with the National Science Foundation, the Dynamic Resource Allocation via GMPLS Optical Networks (DRAGON) project. DRAGON, a partnership of the University of Maryland (UMD), Mid-Atlantic Crossroads (MAX), the University of Southern California (USC) Information Sciences Institute (ISI), and George Mason University (GMU), is developing technology that allows network applications to dynamically acquire dedicated and deterministic network resources. The DRAGON architecture utilizes Generalized MultiProtocol Label Switching (GMPLS) as the basic building block for network element control and provisioning. The key functional elements in the DRAGON control plane architecture include the Network Aware Resource Broker (NARB), the Virtual Label Switch Router (VLSR), the Client System Agent (CSA), and the Application Specific Topology Builder (ASTB).

The DRAGON control plane software is being deployed across the Internet2 Network. The DRAGON project is an active participant in the Global Lambda Integrated Facility (GLIF) and in the Open Grid Forum (OGF) Network Markup Language Working Group (NML-WG). Developers of the DRAGON software are working closely with partners in the GÉANT2 project, ESnet, and the University of Amsterdam through the DICE forum to establish a common protocol for provisioning circuits across multiple autonomous networks so that end users in the Internet2 community can eventually provision circuits from a US university to a national or international partner across five or more networks.

Federated Trust

An integrated network environment must include the technology to anticipate and manage security issues. Providing user-level access to dynamically provisioned network resources requires a federated trust system that enables distributed authentication and authorization. This simplifies management of identity and access permissions across a circuit-specific network, facilitating informed authorization decisions for individual access of protected online resources.

The Internet2 Middleware Initiative is fostering the development and deployment of the elements required to implement this capability across the Internet2 community. Specific projects include the development of Shibboleth, Signet, and Grouper. The new network will incorporate the work of the Internet2 Middleware Initiative into the dynamic Dedicated Circuit services.

Internet2 also operates InCommon, a federated authentication service built atop Shibboleth, eliminating the need for researchers, students, and educators to maintain multiple, password-protected accounts. Online service providers will no longer build and manage account provisioning systems. InCommon uses innovative Shibboleth authentication and authorization systems to enable cost-effective, privacy-preserving collaboration among its community of participants. Beyond the new network, InCommon provides the Internet2 community a federation upon which to build a common trust fabric.

Internet2 is collaborating with the GÉANT2 project through the DICE forum to ensure interoperability between the Shibboleth and EDUGain federated trust solutions and interoperability between other cyberinfrastructure systems (e.g. measurement, control plane) and the underlying deployed federated trust infrastructures.

Performance Measurement

Building on the experience and success of the Abilene Observatory, deployment of the Internet2 Network will be fully instrumented to enable network users, engineers, and network researchers to better understand network performance and to trouble shoot problems should they arise. Performance measurement capabilities may also provide the basis for implementing new capabilities that adjust the provisioning of network resources on a realtime basis depending on network load and capacity. This has proven useful to the community in the past as shown by numerous citations in research papers

Internet2 has built basic performance tools such as the Network Diagnostic Tool (NDT), performance measurement beacons that enable partial path analysis of performance and limited sharing of performance monitoring resources across administrative boundaries such as the One-Way Active Measurement Protocol tool (OWAMP) for measuring one-way latency, and Bandwidth Control (BWCTL, which schedules tools such as NLANR DAST's Iperf tool) for measuring throughput. All these tools have been deployed on the Abilene Observatory. These capabilities are available to network participants and connectors to assist them in their network diagnostic and problem-solving efforts.

These basic tools are being deployed across the new network as part of the Internet2 Observatory. Physically, the Observatory consists of rack space for servers and switches that connect between three and ten servers to each node's optical gear and/or router, with additional space for servers to support Internet2 member research projects. Installed at each node is software to support the dynamic circuit service control plane, to perform passive and active measurement gathering, and to try experimental advanced networking solutions. The measurement servers are being used to do regular and on-demand performance monitoring between network nodes and with our connectors. Universities, regional networks, and partners can also initiate tests to and from these measurement points.

The process of gathering and sharing network performance information across multiple administrative domains also leverages the work the Internet2 community has accomplished in middleware to simplify the process of debugging network issues and to provide performance data to network-sensitive applications without redundant measurements. Early prototypes of the perfSONAR framework are deployed in Europe and in North and South America and serve as a central component of the network monitoring functionality for the global network dedicated to serving the LHC project. The rapid adoption of the early prototype of the perfSONAR system by various National Research and Education Networks (NRENs) and the LHC project is a testament to the need that it fills. The perfSONAR framework is being deployed across the network, allowing access to backbone measurement data in the same via the same protocols as being deployed on other

R&E networks in the US and abroad. This capability will help network operators efficiently troubleshoot end-to-end performance issues along network paths that include the Internet2 network. Operators will be able to find and repair problems more quickly, enhancing their support of their user community (faculty and students). This is particularly important for Layer 1 and Layer 2 issues, where many traditional IP troubleshooting tools are ineffective for debugging network problems.

NETWORK PARTNERSHIPS, CONNECTORS AND COMMUNITY

The strength of the Internet2 membership's involvement is fundamental to driving the evolution of the next-generation optical network environment. The new network provides both IP and layers 1-2 connectivity, thus providing connection alternatives that have not existed before. A key strength of the deployment of the Internet2 network is the breadth of engagement by the Internet2 community and beyond in its deployment.

Community Design Workshop

In mid-June 2006, Internet2 hosted a Community Design Workshop (CDW) to present the proposed hybrid network for the community, which had been announced two months prior at the Internet2 Member Meeting. The workshop was attended by over 120 technical and executive staff representing both current and prospective connectors to the Internet2 national network infrastructure. The goal of the CDW was to provide more detailed information about the technical aspects of the new network plans and to gain input from the technical community on several elements of the network's design and implementation. The engagement of the attendees -- particularly those organizations providing member connections to the national network backbone – proved invaluable in further design development, deployment planning and policies. Subsequently, the roster of the CDW has served as the foundation for an ongoing program of connector communication and engagement, individually and collectively, by Internet2 staff.

Network Technical Advisory Committee

Created at the Community Design Workshop, the Network Advisory Group (NAG) was charged with providing recommendations on a technical advisory structure that would provide feedback and counsel on the design, deployment and operation of the new network. In a series of meetings that took place June through August, the NAG produced a report recommending a Framework for a new Internet2 Network Technical Advisory Committee (NTAC). Composed of representatives of regional networks and members, NTAC provides ideas, feedback and counsel to the Internet2 Network leadership on technical matters relating to the design, planning, implementation and operation of the network.

ESnet Partnership

ESnet is the high-speed network serving thousands of Department of Energy scientists and collaborators worldwide. Over the past few years, ESnet and Internet2 have shared a common technical vision for the evolution of dynamically delivered network capabilities for future scientific research. Formalizing their collaboration, at the end of August 2006 the two networking organizations announced a partnership to develop and deploy the next generation of ESnet with the Internet2 Network backbone. Thus, both the DOE and the Internet2 community are working together toward their mutual need for the development of a scalable, sustainable, high performance network backbone that connects researchers and scholars at Internet2 member institutions with scientists at the Department of Energy. The Department of Energy's ESnet build-out will follow that of the Internet2 Network.

Reach of the Network

The Internet2 community includes 208 universities, 61 companies, 22 regional networks, 52 affiliate organizations, 38 K-20 networks, virtual organizations of scientists scattered across the globe at many institutions, federal partners, and over 50 international partners. As one measure of network reach, there are over 9300 network routes in the current Abilene network, spanning 1385 Autonomous Systems (ASes)

INTERNET2 GOVERNANCE

Internet2 is a member-driven organization built on its community's expertise and leadership.

Our Board of Trustees consists predominantly of representatives from members, including university presidents and CIOs, as well as leaders from industry and research. The member-focused Board provides strategic direction, leadership and oversight. Leaders from the Internet2 community also serve on Advisory Councils and provide valuable input to the Internet2 Board of Trustees on matters related to advanced networking in higher education.

Governance and Nominations Committee

In response to community input that has encouraged revisions to Internet2's governance structure, the Internet2 Board of Trustees formed a new Governance and Nominations Committee (GNC) in May of 2006. The committee is a standing committee of the Board that replaces the former Nominations Committee, which was charged with the initial task of assessing the current Internet2 governance structure and processes, and recommending modifications to take effect in 2007.

In its report to the community, the GNC described its task as both descriptive and prescriptive. "The descriptive responsibility is to develop a clear picture of the current governance matrix, showing the extent of each set of interested institutions' and individuals' power and voice within each governance forum, and describing criticisms that have been presented of different elements of that picture. The prescriptive responsibility is to recommend how that governance matrix should change, in light of institutional experience and changes in our environment over the past decade."

A discussion draft of the GNC report was published on November 30, 2006, and posted on the Web at that time. It was the subject of discussion and focus groups at the Internet2 Member Meeting in December, as well as numerous communications to the GNC from individuals, members and other consortia. The final draft, reflecting the GNC's consideration of comments from the community, was accepted by the Board the following month. https://wiki.internet2.edu/confluence/download/attachments/3375/GNC+Final+Report+12-29-062.pdf

A key outcome of the GNC's work, supported by the Network Planning and Policy Advisory Council (NPPAC) was the establishment of a new membership category called Network Member. This category is available to any non-profit or not-for-profit organization that is sub-state, state or multi-state in scope, and whose principal mission is to provide network infrastructure and services primarily to the research and education community in the relevant geographic area, including, but not limited to, access to the Internet2 national network infrastructure and services. Members of Internet2 in this category will elect their own representatives to sit on the new councils, and will have representation on the new Board.

TRANSITION PLAN

The transition from Abilene to the new network began in the fall of 2006, anticipating complete build-out in July 2007. With the objective of minimizing disruption to existing connectivity and services, a deployment plan working in concert with RONs and connectors follows a geographical sequence, beginning in the Northeast. Internet2 will use the optical capabilities to transition connections in the near term and provide support for evolution to the longer-term, 2 x 10 Gbps connectivity, allowing connectors to upgrade at a measured pace. (See http://www.internet2.edu/network/deployment.html)

CONNECTING TO THE NETWORK

Researching future bandwidth requirements and developing usage models for the new network, Internet2 has created a sustainable business model for a network with approximately two dozen connectors.

Beyond simply transitioning to a new network with identical capabilities, the opportunity and challenge to the community is to take full advantage of the new capabilities the network offers, in advance of the demands that are likely to quickly emerge from a broad range of communities served by both individual institutions and regional networks. The transition plan for the new network sets targets around 2 x 10GigE connection, with lower speed connections from 1GigE supported during the four-year migration period. Institutions initially connecting to the network at lower speeds are expected to eventually migrate to a 10GigE interface, but their traffic will be rate-limited based on the terms of their connection agreement.

Based on conversations with regional networks, a 2 x 1GigE connection is now offered. This will enable dynamic services for connectors not able to go to 2 x 10GigE immediately. While OC48 connections can still be supported in most cases if necessary, migration to an OC192 or 10GigE interface wherever possible is encouraged to allow for a seamless future upgrade.

Considerations for connection sizing to the Internet2 Network

Bandwidth intensive applications targeted at multiple disciplines continue to grow in number and in use within the research and education community. The number of network devices and their ability to generate traffic continue to increase. It is crucial to consider these applications and their demands when determining the proper connection speed to the Internet2 Network.

Video applications continue to be some of the most demanding applications that currently run over the network, both in terms of bandwidth utilization and the quality of the underlying network. Uncompressed High Definition video currently consumes approximately 1.5 Gigabits per second (Gbps) of bandwidth and like many other video applications are very sensitive to loss, whether due to congestion delay or network errors. It is imperative that an underlying network that supports high performance video applications be both sufficiently high capacity and free of errors that these applications can run in concert with other normal traffic and not be degraded.

Large file transfers also represent a traditional challenge to networking. Advancements in e-science continue to grow the size of the available data sets and the importance of individual or groups of researchers working collaboratively across geographic divides is also increasing. The result of these two factors is an increasing emphasis on fast and clean networks as the transfer of large files at the highest data rates must traverse exceptionally error-free networks due to the throttling back and retransmit features of TCP when a packet loss event occurs.

The most common example of the need to support researchers depending on large file transfers is the global network of researchers analyzing data from experiments that utilize the aforementioned Large Hadron Collider at CERN. Health sciences have applications that have the same type of networking requirements, including visualizing massive amounts of biomedical data at disparate locations and moving images that cannot experience network jitter or delay in order to be of use by a clinician. Similarly, new techniques in radioastronomy increasingly rely on moving massive amounts of data from widely distributed locations.

In the past it was well beyond the reach of all but a few very specialized machines to generate enough traffic to saturate a 10Gbps capable network. Now standard desktop—and even laptop—hardware and operation systems can utilize significant portions of 1 Gbps of network capacity. Servers capable of saturating the backbone can be had for under $5000 and will increasingly allow broad deployment of the applications mentioned above along with new ones that are currently being developed. The history of networking has shown that leading-edge capabilities are often quickly adopted and diffused.

Furthermore, as noted above, early stages of testing the Commercial Peering service has increased flows over institutions' Internet2 connections by nearly 600 Mbps. Similar increases are expected as the Commodity Transit service is implemented. Therefore, current assumptions about requirements for Internet2 connections should be reconsidered in light of new services being offered, and the anticipation that expectations of network capabilities and performance will dramatically change in the near future.

Aggregation Considerations

The general architectural network hierarchy for the research and education community in the United States for the past 10 years has been a three-tiered system in which campuses connect to regional networks that in turn connect to a backbone. It has supported the users in the community with high-performance networking from the early days of the NSFNET network. The following diagram illustrates the typical architectural network hierarchy:

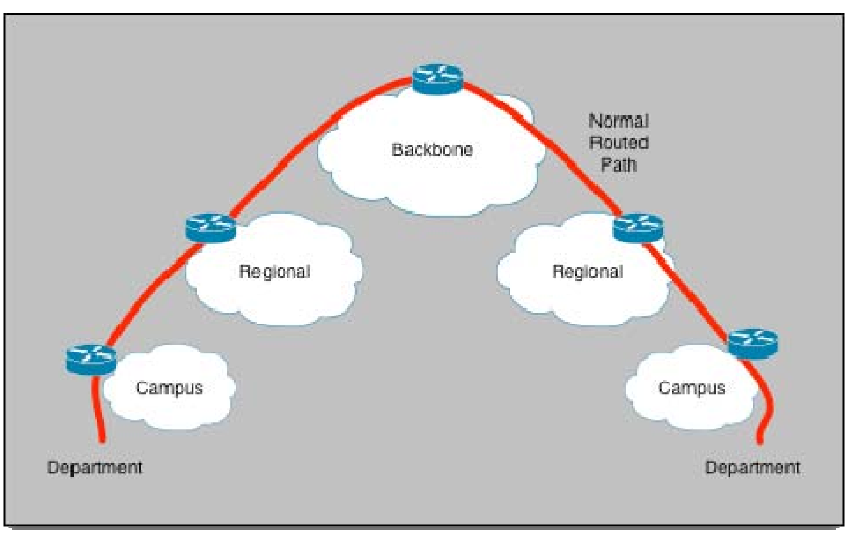

The Internet2 backbone was originally designed expecting approximately 15 to 25 regional connectors. The regional connectors provide connectivity for approximately 200 – 250 research institutions, with an average of approximately 10 research institutions per regional. The numbers vary depending on the regional network, but the pattern is roughly aggregations of 10 – 20 connections at each level in the hierarchy. Moreover, the connections between levels in the hierarchy are typically limited to 10 Gbps. This implies bandwidth averages of approximately one Gbps per institution down to the campus level.

As is currently under discussion in the Internet2 community, changing the hierarchy, either by aggregating a larger number of institutions for each regional network or by adding levels to the hierarchy, presents several potential challenges:

- Bandwidth per institution could be affected, at all layers in the network, because of funneling a much larger number of institutions through a single 10 Gbps connections. For example 50 institutions are connected through a single 10 Gbps connection, either at the IP layer or through the Dedicated Circuit (DC) services, the capacity per institution is limited and drastically reduced.
- Fair allocation schemes for institutions are more difficult. They also take on greater importance, and must be implemented by the aggregating regional network. In the past, this was typically handled by over-provisioning, but that may not be sufficient for large-scale regional networks. In addition, this applies to both IP and the dynamic services capabilities. Typical bandwidths for IP to a university are 1 Gbps, and typical dynamic service circuits are also expected to be on the order of 1 Gbps. Assuming 50 or more research institutions aggregated behind a 10 Gbps, adequate performance becomes problematic or in some cases impossible.
- If an additional layer in the hierarchy is added by creating mega-regional aggregation, it would add complexity, especially for the dynamic services capabilities, but also for IP connections. As a technical consideration, it adds an additional point of failure.

In either case—aggregating a larger number of institutions for each regional network or by adding levels to the hierarchy—the number of connections to the backbone would likely decrease and the number of institutions accessing the backbone through a given connection would likely increase.

It is clear that a drastic change in scaling of the architectural hierarchy would require a redesign of the backbone network. It would also likely result in a review of the fee structure that supports the recovery of costs related to implementing and operating the backbone network on a sustainable basis. Ultimately, this is an issue to resolve through community discussion, but the overall goal is a fair sharing of costs across the Internet2 membership for a network that meets the needs of that community.

Dedicated Circuit Services

To use the dedicated circuit service, a user must make a physical connection to the Ciena multiservice switching infrastructure using an Ethernet or SONET client interface. The Ethernet interfaces may be 1GE or 10GE. SONET interfaces may be OC-48 or OC-192, though OC-192 is encouraged where possible to allow for maximum flexibility.

Over the physical interface, the user may carry a single circuit or multiplex multiple sub circuits. An Ethernet interface multiplexes VLANs. A SONET interface multiplexes streams, where each stream is made up of multiples of OC-1 (51.84 Mbps) streams.

The bandwidth of the physical interface and the bandwidth of the circuit provided across the network do not need to be the same. Each incoming, possibly multiplexed, circuit is connected to a dynamic circuit across Internet2. The physical interface to the Internet2 circuit has a specific bandwidth, and this bandwidth sets the upper limit for the bandwidth of the circuit (or circuits) carried by this interface. The dynamic circuit may be configured in increments of OC-1.

For Ethernet connections, the circuit provided to a remote Ethernet device is a point to point Ethernet circuit. The circuit may be a tagged VLAN at one end of the circuit and untagged Ethernet at the other end, or be the same at both ends.

Ethernet circuits are carried on Internet2 dynamic circuit services as GFP encapsulated SONET. GFP encapsulation and decapsulation of Ethernet frames is done at network ingress and egress.

When the Internet2 segment of the dynamic circuit connects with a segment from another network using a SONET interface, the data carried between Internet2 and the other network is GFP encapsulated SONET. This means that the GFP encapsulation must be done at the termination of the SONET part of the circuit.

A good example occurs when the circuit includes a segment from Internet2 and a segment from GÉANT, where the users at each end connect using Ethernet, but Internet2 and GÉANT connect using SONET. In this case the circuit between Internet2 and GÉANT is GFP encapsulated SONET.

For SONET connections, the circuits provided are point to point SONET circuits at multiples of OC-1 bandwidth. Incoming physical SONET circuits may be channelized or unchannelized.

As noted, the Ciena infrastructure connection across Internet2 is always SONET. Dynamic circuits' SONET connections are always channelized. User SONET connections may be channelized or unchannelized. If unchannelized they are converted to channelized at the edges.

Following are ways that dynamic circuits may be connected to other users and networks. The methods of interconnecting are undergoing intense investigation.

Presented here are best guesses as to how this will happen, in order to provide insight for actual implementations.

Connecting Two Regional Networks and Their Users

The dynamic circuit service provides the ability for a RON to connect to Internet2 and provide dynamic circuit services across Intenet2 to a user on another regional network. In this case, the regional network makes a physical connection to Internet2.

In the most common case this will be an Ethernet connection that supports VLANs. The regional network on each end creates a VLAN circuit to its user and then makes a VLAN connection to Internet2. The VLAN segments are joined to create an end to end circuit between the users.

The following figure shows such a connection. The regional network provides its own circuit multiplexing capability that takes circuits from its users and multiplexes them over its connection to Internet2. Each multiplexed connection is sent by Internet2 on a separate SONET Virtual Container connection to a multiplexed connection on a different RON.

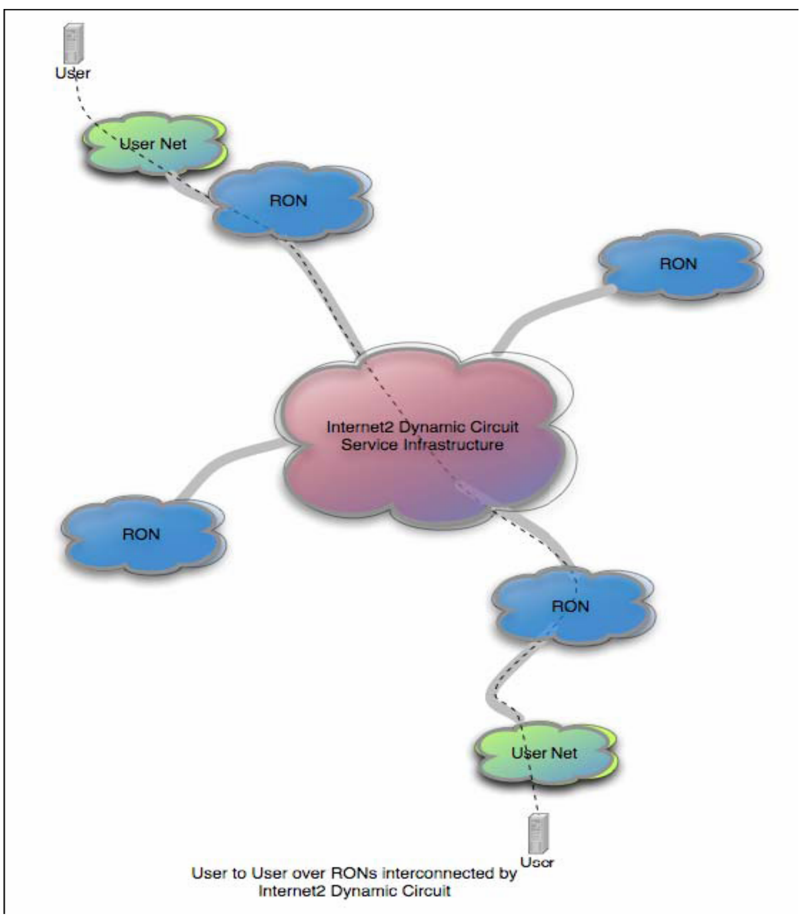

User to User over RONs interconnected by Internet2 Dynamic Circuit

Some notes about this setup. The physical connection between the RON and Internet2 is permanent. The dynamic circuits carried over the physical connection change over time, depending on the requirements of the users.

The dynamic circuits are set up and managed by control plane software, and if the RON and Internet2 have compatible control plane software, the circuits may be set up automatically across domains. Cross-domain control plane software is a heavily investigated research topic at this time.

Internet2 is working with others in the Internet2 community and through the DICE collaboration to develop software that will allow RONs to control Ethernet switches and interface with the Internet2 control plane software. Internet2 is also participating in research and standardization efforts to create control plane designs and specifications that will allow interoperability of control plane software developed by different organizations.

Connecting Internet2 with other National and International R&E Networks
Internet2 dynamic circuit services connect with similar services provided by other national and international R&E networks such as ESnet, GÉANT, and CANARIE. Each interconnection is a physically permanent connection over which dynamic circuits are multiplexed.

These interconnections take place either directly with other networks on a one-to-one basis, or — as proposed by the Global Lambda Integrated Facility (GLIF) — at exchange points where multiple networks come together and circuits may be switched among any of the connected networks. Following GLIF, we refer to these exchange points as Global Optical Lambda Exchanges (GOLEs).

Using these interconnections and appropriate manual and automatic control of circuit switching, Internet2 will be able to be a partner in creating circuits from users on regional networks in the United States to users on networks connected to other core networks.

The data interface between core networks may be done using either Ethernet or SONET/SONET interfaces. The control plane interface is logically very similar to the interface between Internet2 and regional networks. At the present time, the biggest differences in control plane software interfaces appear to be in defining how to authorize users requesting services, and how to share information about reachability between networks. Internet2 is involved with groups working on ways to provide and standardize these.

POLICIES AND PARTICIPATION

Acceptable Use Policy

As it is deployed, the Internet2 Network will offer the member community the advantages of greater bandwidth capabilities, dynamic circuit flexibility, and evolving value-added services including transit and peering with the commercial Internet. Considerable input from the community, including discussions with connectors, was integral in developing an open and simple policy to encourage the use of the advanced capabilities and services of the new network:

"The Internet2 network can be used for any legal purpose, so long as it does not interfere with or adversely affect the operation of the Internet2 network or any network user, as may be determined by Internet2.

In addition, when using the Internet2 network, users shall be subject to any posted guidelines or other rules, including Internet2's Privacy Policy, Spam Policy and Copyright Policy. All such guidelines and rules are hereby incorporated by reference into this AUP.

Internet2 reserves the right to modify this AUP (and its posted guidelines or other rules) from time to time and intends to provide advance notice of any such modifications."

Internet2 Service Operating Practices – effective January 2007 through 2008

I. Internet2 Network Participation – Interim Policy

The following outline, currently in draft form, represents the Interim Network Access Policy (NAP). Also included here are other network services for which usage policies are specified (e.g., commodity). It is considered Interim so that full consideration of a new network participation model can be carried out during 2007 as part of continued discussions about a new membership model. It is proposed that the final Interim NAP be developed by a committee that includes appropriate representation from Internet2 governance and members, particularly representatives from state and regional networking organizations.

The Network Access Policy serves as a companion to the Network Acceptable Use Policy and other relevant Internet2 policies governing who can connect to network services.

A. All institutions and organizations connected to Internet2's advanced IP Research & Education network service through a direct or indirect Connector are one of three kinds:
    1. Individual (primary) network participant
    2. Sponsored Educational Group participant
    3. Sponsored participant
    4. Other (see E below)

B. Individual network participants are:
    1. Individual organizations or institutions
    2. Members of Internet2
    3. Billed directly by Internet2 for the participation fee[1]

C. Sponsored Educational Group participants are:
    1. Serviced by an Internet2 Network Connector
    2. Billed by the Connector according to Connector's business model. Connector is billed by Internet2 for the group participation fee, which the Connector may or may not rebill.

D. Sponsored network participants are:
    1. Individual, non-member institutions and organizations that are sponsored by an Internet2 member.

E. Other
    1. Regional Health Network Participants
        - Proposed pilot implementation is under consideration.
    2. Working assumption is that these groups will be connected through the appropriate Connector and billed by the Connector for any associated fees.

F. All direct state and regional network organization Connectors[2] will be given gratis membership as Internet2 Network Members.

II. Network Capabilities Offered to Internet2 Network Connectors

A. IP Connectivity
B. Commodity Transit*[3]
C. Commercial Peering**[4]
D. Dynamic Wave Service
E. Static Wave Service

The Internet2 network Connector can in turn offer the use of these capabilities to build services for their members and connectors.

---

[1] Except by special arrangement with the Connector
[2] An Internet2 Network Connector is defined as an entity providing network services and aggregation for others within a region or state, and which has established a direct connection to the Internet2 network. These are most often state or regional networking organizations, and it is Internet2's standard practice to connect individual member institutions and organizations through these state or regional network connectors.
[3] Internet2 will facilitate our connectors' access to this service. Actual implementation will be advised by the group of regional network representatives described on page 11
[4] Actual implementation awaiting result of pilot experience

III. Participation in Other Internet2 Capabilities

The following Internet2 capabilities are offered to any Internet2 member or connector:

A. The Commons
B. WaveCo[5]

The following capabilities are restricted to non-profit research and education entities

A. FiberCo Fiber Services (metro and inter-city dark fiber, collocation)

Services offered to Members and Non-Members*

A. InCommon
B. USHER
C. FiberCo Professional Services

*Non-members may pay a higher fee.

SUSTAINABLE FINANCIAL PLAN

A key aspect of developing the new network has been ensuring its long-term financial stability and viability. Experience with both the NSFNET and the first round of high-performance networks used by the Internet2 community has demonstrated that networking infrastructure, especially infrastructure with new capabilities, require time to be widely adopted by the research and education community. Moreover, transition costs associated with moving from one nationwide infrastructure to another tend to divert resources and energy from developing or further deploying new capabilities or technologies.

Therefore, in developing the new network Internet2 has formulated a fee structure that recovers costs in a way that ensures its continued viability, and sets a predictable and solid foundation for future investment in new technologies. The fee structure formulation also ensures the new technologies embodied in the network will add significant capability and value while minimizing additional costs. To accomplish this, Internet2 has worked closely with regional networks and member organizations to ensure the smoothest possible transition.

Finally, Internet2 has been able to build key partnerships with corporate and university members in developing, implementing and operating the network, which not only help provide a beneficial cost structure, but also set the stage for working with corporate members to apply the lessons learned through the development and deployment of the Internet2 Network to the broader Internet.

---

[5] WaveCo (per the Level3 contract) is available to higher ed and research institutions that are members of Internet2 or not-for and for-profit institutions that are collaborating with Internet2 members.

Internet2 Financial Status and Plan

Planning Platform

Based on the agreement with Level 3 Communications announced in June 2006, the Internet2 Network is being deployed nationally over 13,000 miles of dedicated fiber. Level 3 will be responsible for meeting a Service Level Agreement (SLA). Subsequent to the Level 3 partnership, Internet2 announced agreements with Infinera and Ciena to provide the advanced optical platform and switching capabilities for the new network.

The advantages of the Level 3/Infinera/Ciena agreements are:

- Carrier-class reliability
- Economies of scale, using Level 3's infrastructure and expertise
- Community control of the infrastructure, including provisioning and switching of circuits
- Flexibility to use the infrastructure for layers 1-3 transport, as recommended in the Group A Report
- The Level 3 agreement, in the spirit of community interest, enabled incorporation of the NLR fiber IRU to achieve cost savings in the event of a merger of Internet2 and NLR

Basic Characteristics of Internet2's Financial Model

Internet2's overall financial model is based on cost recovery, with a planned annual amount used to build a reserve for major planned capital expenditures. This reserve is accumulated through depreciation charges built into the annual operating expense budget.

Internet2's budget comprises two primary components, Core and Network, supplemented by three much smaller components that collectively represent about 5% of the total (sponsored programs, projects, and special hosted projects).

The Core component consists of Internet2's membership programs and activities, such as member relations, member meetings, international relations, external relations, and non-network technical development projects such as middleware, security, and others. Its major revenue sources are membership dues, supplemented by contributions and other revenue such as Member Meetings and workshop fees.

The Network component consists of Internet2 network development, operations, and project expenses, including end-to-end performance, the Observatory, HOPI, the K20 initiative, and others. Its major revenue sources are connector fees, primary network participation fees, and SEGP fees, supplemented by partnerships and other service offerings. Allocated and administrative costs are accounted for in the Core budget, and the Network budget transfers a share of the revenues to the Core to cover an equitable share of these costs. Items included in the cost base for the transfer are financial and human resources support, space, insurance and other organizational infrastructure costs, plus a share of executive personnel expenses.

Business Model Going Forward

Internet2's business model for 2007 through 2010 is based on the assumptions below. As stated above, the overall goal is cost-recovery, coupled with the building of adequate reserves to enable future progress.

Revenue assumptions

- Internet2's current membership model will be in place through 2008
- A new membership model will be instituted in 2009 with the potential to increase membership by the addition of smaller institutions and organizations.
- Projections of numbers of members are conservative estimates and include potential loss in 2008, which is recovered going forward.
- Health-care organizations are projected to join Internet2 in increasing numbers due to the health network activities
- There will be some revenue from the FCC pilot program in 2007 and 2008 from health networks that connect to the Internet2 network.
- There will be a moderate increase in connector aggregation
- Connection fees are projected to remain constant
- SEGP fees are projected to remain constant

Expenses

- 4% added each year for inflation
- Increased personnel expenses are included to support the new network and other key priority areas.

The following tables summarize the financial position and plan for Internet2. Footnotes are included to explain assumptions behind many of the numbers.

Table I shows the **2006 versus 2007 Budgets**. Internet2's fiscal year is the calendar year. The financial results for 2006 are not yet final and are based on September 2006 projections. The annual audit is in process and the draft results will be presented to the audit committee of the board on March 12. Subsequent to the creation of the 2006 budget, the decision was made to pursue the development and deployment of the new Internet2 network. As a result, significant additional spending was required and supported by the board for the Level3 agreement, capital expenditures and other network build out requirements.

The 2007 budget was developed to support the completion of the new Internet2 network and support the other key priorities of the organization. Some specifics about the 2007 budget include:

Revenues

- Increase of approximately $7M primarily related to Network Services

Operational Expenses

- Increase in Personnel total of $2.9M related to additional staffing as well as $1M of expense previously captured in the Office Support category (Note offsetting decrease in Office Support)
- Increase in Network total reflects the impact of incurring costs for both Abilene and the Internet2 Network for what amounts to three quarters of the year
- Depreciation increase of $1.9M due to capital acquisitions for the new network being put into service

Capital Expenditures

- Increase in Network Equipment expenditures of $4.9M related to the build out of the network. Between the two years of 2006 & 2007 we expect the total of capital expenditures to be $13.4M.

Table II is an **Overview** showing five years of Internet2's financial plan. Internet2 has always maintained separate budgets and accounting for its network activities and core activities of the organization. During 2006 and 2007 substantial investments are being made to transition the Abilene network to the new Internet2 network causing negative cash flow for the organization each year but supported by accumulated cash reserves. In 2010, there is also a $10M capital expenditure planned to upgrade the network router platform.

Network Revenue (Table III) shows revenue anticipated from network participation, connections and SEGP participation during 2006 through 2010 with a set of assumptions for the forecast. It is important to note the connection fees are assumed to remain constant over time.

Table IV of **Ending Cash Comparison** summarizes the cash position of Internet2 on an annual basis. With the support of its members, Internet2 accumulated reserves to support upgrading the network. As of January 1, 2006 cash reserves were $27.4M. Investment in the new Internet2 Network has and will cause negative cash flow in 2006 and 2007, depleting the cash reserves to $15.4M and $5.6M respectively before rebuilding over time to support investments in subsequent years.

TABLE I: 2006 vs. 2007 BUDGET

	Calendar Year 2006		Calendar Year 2007	2006 to 2007 Percent Change	
	Budget	YE Projection*	Budget	Budget to Budget	YE Projection* to Budget
Revenues					
Member Dues	8,438,000	8,516,840	8,864,800	5.1%	4.1%
Meeting Fees	632,000	582,954	577,275	-8.7%	-1.0%
Network**	14,657,250	14,782,250	14,810,250	1.0%	0.2%
Network Services	250,000	125,000	6,629,811	2551.9%	5203.8%
Management Fees	95,000	150,000	185,000	94.7%	23.3%
Sponsored Program	1,523,607	1,523,607	1,798,000	18.0%	18.0%
Contributions	100,000	100,000	110,000	10.0%	10.0%
Miscellaneous	410,709	300,000	150,000	-63.5%	-50.0%
Investment Income	400,000	451,485	200,000	-50.0%	-55.7%
Project income	149,000	86,568	168,500	13.1%	94.6%
Interest	350,000	729,926	350,000	0.0%	-52.0%
Total Revenues	*27,005,566*	*27,348,630*	*33,843,636*	25.3%	23.7%
Operational Expenses					
Personnel	10,641,429	10,924,892	13,529,062	27.1%	23.8%
Meeting	1,526,577	1,472,128	1,509,336	-1.1%	2.5%
Travel	996,907	866,562	1,094,783	9.8%	26.3%
Print	89,750	29,468	90,575	0.9%	207.4%
MarComm	25,600	24,341	27,100	5.9%	11.3%
Office Support	3,456,921	4,208,993	2,737,327	-20.8%	-35.0%
Network	8,142,650	9,089,477	17,181,700	111.0%	89.0%
Depreciation***	1,940,000	2,159,621	3,861,846	99.1%	78.8%
Other	165,000	(6,344)	333,580	102.2%	-5358.2%
Total Operational Expenses	*26,984,834*	*28,769,138*	*40,365,309*	49.6%	40.3%
Net Operating Results	20,732	(1,420,508)	(6,521,673)		
Capital Expenditures					
Desktop / Laptop Hardware	25,000	5,761	35,000	40.0%	507.5%
Server Hardware	105,000	358,705	1,421,000	1253.3%	296.1%
Network Equipment	1,870,000	4,378,934	6,787,500	263.0%	55.0%
Other Equipment	7,500	-	8,000	6.7%	-
Leasehold Improvements	4,000	17,696	4,000	0.0%	-77.4%
A/V Hardware	60,000	56,735	100,000	66.7%	76.3%
Furniture & Fixtures	5,000	2,634	5,000	0.0%	89.8%
Software	32,500	7,104	28,000	-13.8%	294.1%
Total Capital Expenditures	*2,109,000*	*4,827,569*	*8,388,500*	297.7%	73.8%
Net Activity	(2,088,268)	(6,248,077)	(14,910,173)		
Level 3 Advanced Cash Payment		*(10,000,000)*			
Annual Net Cash Flow	(148,268)	(11,978,530)	(9,848,327)		

* Year End Projection by Executive Directors as of September, 2006

** Network consists of Participation, SEGP and Connection fees.

*** Depreciation represents the transfer to reserves set aside according to Internet2's capital replacement plan

TABLE II: OVERVIEW

		2006*	2007	2008#	2009#	2010#
Revenue						
	Core Membership	8,209,240	8,559,000	8,674,000	8,999,550	9,880,605
	Core Other	4,097,140	3,599,575	3,576,403	3,663,051	3,756,621
	Total Core Direct Revenue	*12,306,380*	*12,158,575*	*12,250,403*	*12,662,601*	*13,637,226*
	Core Admin Fee**	2,500,000	2,625,000	2,756,250	2,894,063	3,038,766
	Total Core Revenue	14,806,380	14,783,575	15,006,653	15,556,664	16,675,992
	Network Participation	4,935,000	5,368,000	5,664,000	6,183,950	6,829,778
	Network SEGP Participation	1,822,250	1,822,250	1,873,450	1,924,650	1,975,850
	Network Connection	7,900,000	7,620,000	10,600,000	11,940,000	12,430,000
	Network Services***	385,000	6,874,811	6,053,654	10,693,135	13,412,009
	Total Network Direct Revenue	*15,042,250*	*21,685,061*	*24,191,104*	*30,741,735*	*34,647,637*
	Network Admin Fee**	-2,500,000	-2,625,000	-2,756,250	-2,894,063	-3,038,766
	Total Network Revenue	12,542,250	19,060,061	21,434,854	27,847,672	31,608,871
Total Revenue		27,348,630	33,843,636	36,441,507	43,404,336	48,284,863
Expenses						
	Core Operating	14,422,210	14,483,877	14,111,650	14,817,233	15,558,095
	Core Capital	171,852	612,000	190,902	200,448	210,470
	Total Core Expenses	14,594,062	15,095,877	14,302,552	15,017,681	15,768,565
	Network Operating	24,346,926	25,719,858	22,278,002	24,091,640	25,812,745
	Network Capital	4,655,717	7,965,000	711,971	738,795	10,766,691
	Total Network Expenses	29,002,643	33,684,858	22,989,973	24,830,435	36,579,436
Total Expenses		43,596,705	48,780,735	37,292,525	39,848,116	52,348,001
Variance						
	Core Variance	212,318	-312,302	704,101	538,983	907,427
	Network Variance	-16,460,393	-14,624,797	-1,555,119	3,017,237	-4,970,565
Total Variance		-16,248,075	-14,937,099	-851,018	3,556,220	-4,063,138
Depreciation & Other Non-Cash						
	Core Depreciation	518,558	600,000	378,560	397,488	417,362
	Network Depreciation & Other Non-Cash	3,750,987	4,488,774	2,112,346	2,318,508	2,309,426
Total Depreciation & Other Non-Cash		4,269,545	5,088,774	2,490,906	2,715,996	2,726,788
Cash Variance##						
	Core Cash Variance	730,876	287,698	1,082,661	936,471	1,324,789
	Network Cash Variance	-12,709,406	-10,136,023	557,227	5,335,745	-2,661,139
Total Cash Variance		-11,978,530	-9,848,325	1,639,888	6,272,216	-1,336,350

* 2006 Year End Projection by Executive Directors as of September, 2006

** Administrative costs are accounted for in the Core budget, and the Network budget pays an annual allocation back to Core to cover an equitable share of these costs. Items included in the cost base for the administrative transfer are financial and human resources support, other organizational infrastructure costs, plus a share of executive personnel expenses. This fee will continue to be evaluated moving forward to determine the correct %.

*** Includes Anchor Tenants, MAN LAN, WaveCo, etc.

\# Fees for 2008 and beyond have not been officially established

\## Cash Variance is calculated as Revenue - Expenses + Depreciation & Other Non-Cash

TABLE III: NETWORK REVENUE

	2006	*	2007	*	2008	#	2009	#	2010	#
Revenue										
Network Participation	4,935,000	*a*	5,368,000	*b*	5,664,000	*c*	6,183,950	*d*	6,829,778	*e*
Network SEGP Participation	1,822,250	*f*	1,822,250	*g*	1,873,450	*h*	1,924,650	*i*	1,975,850	*j*
Network Connection	7,900,000	*k*	7,620,000	*l*	10,600,000	*m*	11,940,000	*n*	12,430,000	*o*
Total	**14,657,250**		**14,810,250**		**18,137,450**		**20,048,600**		**21,235,628**	
Year to Year Variance			*153,000*		*3,327,200*		*1,911,150*		*1,187,028*	

* *Budget*
\# *Business Model Projection*

a *Based on 235 Participants paying an annual fee of $21,000*
b *Based on 244 Participants paying an annual fee of $22,000*
c *Assumes 236 Participants paying an annual fee of $24,000 (Fees have not been established)*
d *Assumes new Membership Model implemented with tiered approach; Annual fee of $26,000 (Fees have not been established)*
e *Assumes new Membership Model continues with tiered approach; Annual fee of $27,300 (Fees have not been established)*

f *Based on projected Participant States paying a fee calculated as $30,000 base with a variable of $2,000 times the current number of Congressional Representatives for that state*
g *No change in participants or fees expected*
h *Expect the addition of one additional State Participant*
i *Expect the addition of one additional State Participant*
j *Expect the addition of one additional State Participant*

k *Based on 35 Connectors with various types of connections with no anticipated fee change (See schedule below)*
l *Based on 31 Connectors with various types of connections with no anticipated fee change (See schedule below)*
m *Based on 27 Connectors with various types of connections with no anticipated fee change (See schedule below)*
n *Based on 28 Connectors with various types of connections with no anticipated fee change (See schedule below)*
o *Based on 28 Connectors with various types of connections with no anticipated fee change (See schedule below)*

Connection Fees		
2, 10 GE	$550,000	
1, 10 Gbps	$480,000	
2, 1 GE	$340,000	
1, 2.5 Gbps	$340,000	
1, 1 GE	$250,000	
1, OC-12	$220,000	
1, OC-3	$110,000	

TABLE IV: ENDING CASH COMPARISON

	2006	2007	2008	2009	2010
Beginning Cash Balance	**27,391,263**	**15,412,733**	**5,564,408**	**7,204,296**	**13,476,512**
Cash Receipts	27,348,630	33,843,636	36,441,507	43,404,336	48,284,863
Operating Expense Budget (Less Depreciation)	24,499,591	35,114,961	33,898,746	36,192,877	38,644,052
Capital Expense Budget	4,827,569	8,577,000	902,873	939,243	10,977,161
Level 3 PrePayment	10,000,000	-	-	-	-
Annual Net Cash Flow	***(11,978,530)***	***(9,848,325)***	***1,639,888***	***6,272,216***	***(1,336,350)***
Ending Cash Balance	**15,412,733**	**5,564,408**	**7,204,296**	**13,476,512**	**12,140,162**

Network Fee Structure

Internet2 networking staff have spent considerable time researching future bandwidth requirements and developing usage models for the new network. This information was used to create a sustainable business model for a network with approximately 25 connectors (scaled down over time from Abilene's current 35 connectors as a result of aggregation and disconnections). Because national-scale networks require large capital outlays and have considerable start-up costs, deployment of the new network will be covered initially by the reserve that Internet2 has accumulated over the years. Based on current usage and membership models, cash flow will turn positive and Internet2 will begin to build back a reasonable reserve for future network upgrades.

Internet2 does not use Membership Dues to pay for the network – only participation, SEGPs, connection and service fees cover the costs of deploying, managing, and maintaining the network. Service fees cover the costs of providing those services that benefit an individual or small group of member institutions, such as dedicated waves or commodity Internet. Participation, SEGPs and connection fees cover costs associated with providing basic network services to the entire community, including: costs of the initial 100 Gbps network provided by Level 3 Communications, costs and maintenance of Internet2 owned network hardware (routers, switches, grooming boxes, etc.), fees associated with maintaining a presence at exchange points, network operations, technology evaluation and development expenses, and Network Services personnel costs. Equipment costs are amortized over time and depreciation is accrued as an expense.

Other value-added services are only minor investments in time and money relative to their value and are also covered by connection fees – such as an optional commercial peering service which aims to reduce member universities' commodity Internet costs and provide a hedge against unfavorable network neutrality legislation. It should be noted that, in response to member discussions at and since the CDW, Internet2 has unbundled the 500 Mbps/month of commodity Internet from the 2 x 10GigE connection fees. As noted on pages 8-9, Internet2 is exploring a collaborative effort with regional networks to provide transport for commodity service through a connector's Internet2 network connection.

Internet2 has chosen to standardize the new network connection model around a 2 x 10GigE connection, even though lower speed connections from 1GigE will be supported during the four-year migration period. Those institutions that connect to the network at lower speeds will be eventually required to connect using a 10GigE interface, but their traffic will be rate-limited based on the terms of their connection agreement. This achieves two objectives. First, this model provides a convenient and cost-effective way to manage separate IP and dynamic wave services. Second, it minimizes costs by increasing equipment density (minimizing collocation and power fees) and reducing the amount of sparing required to ensure a production quality network infrastructure. There is now a 2 x 1GigE offering that will allow dynamic services for connectors not able to go to 2 x 10GigE immediately. This offering was added as a result of our one-on-one conversations with the connectors following the CDW in June.

The fees for dedicated point-to-point static waves are in addition to the cost of a connection and are based on Internet2's costs including a charge, amortized over several years, for the hardware required to deploy the wave. The Internet2 Network Fees schedule breaks these costs down in 10 Gbps, 1 Gbps and STS-1 (50 Mbps) increments on a per-mile basis for specific time periods. Shorter duration waves have slightly higher fees based on the need to recover costs over shorter time periods. Of course, as was explained at the CDW, immediate reservation, dynamic waves (defined at this time as waves of less than 1 month duration) will be available at no cost and can be used to connect to any other site on the network. Further, advanced reservation, dynamic waves will be made available on this network. Users of advanced reservation waves may incur a fee for the guaranteed availability of this dynamic bandwidth, the fees to be determined by eventual demand and network costs. (See Attachment – Network Services Fees)

Beyond network fees, participation fees for 2008 have been reviewed by the NPPAC and the Internet2 Board will act on them in February. A fee committee convened as part of the implementation of the new governance structure will consider membership model and fees for 2009 and beyond.

The fee structure for Internet2 reflects a very cost-effective investment in next-generation networking for research and education. Further, it follows the Internet2 principles of stewardship of member resources, equitable sharing of costs, and sustaining a long-term, financially viable and technologically advanced network.

ATTACHMENT – Network Services Fees

Effective 2007 – 2010

Redundant IP Connections: Connectors can use the dynamic wave as an immediate reservation dynamic backup circuit to an alternative IP router. A dedicated backup wave to the second closest router can also be purchased so long as the aggregate usage across the two links remains below the connection fee allowance. Prices include the cost of the router interface. Other redundancy options will be considered upon request.

- Dynamic wave: $25,000/year. Pre-installed cross-connect would come out of grooming box and go directly into the router (or into Infinera to be backhauled to alternate router) to be immediately available in the event primary IP connection goes down.
- Dedicated 1 GigE: $75,000/year
- Dedicated 10 GigE: $175,000/year

Dynamic Waves: Only available to those connectors at 2 x 1GigE and 2 x 10GigE. These waves may be used reserved for hours, days or weeks.

- Immediate reservation, short-term (<1 month) point-to-point waves are free, subject to blocking.
- Advanced reservation, short-term point-to-point waves will have guaranteed availability for the period of the reservation. These waves may incur additional fees, to be determined when we can evaluate actual demand and costs.

Dedicated Waves: Available on the Internet2 Network footprint to Connectors. Members that aren't direct connectors need to go through a regional with a signed connection agreement with Internet2. Available at any core optical node or add/drop location as depicted on the network map. Includes 99.9% availability SLA.

WAVE FEES

	STS-1	**1 GigE**	**10 GigE**
< 1 year	$ 0.12	$ 2.34	$ 20.58
1 year	0.062	1.21	10.63
2 years	0.057	1.12	9.81
3 years	0.054	1.05	9.26
4 years	0.052	1.01	8.92
5 years +	0.05	0.98	8.58

$ based on fiber route miles per month

FiberCo Fees (through March 2008):

- Dark fiber: $850/strand mile for 20 year IRU
  - $40,000 transaction fee applied
- Colo: All racks have $3000 NRC and 5 year commitment if in Internet2 space, otherwise per FiberCo contract ($2,500 NRC).
  - Gateway: $800/month/rack includes 30 amps DC power
  - ILA: $600/month/rack includes 10 amps DC power
  - Additional power: $150/10 amps/month
- O&M: $200/route mile/year for up to 6 fibers
- Metro fiber: Individual case basis, dependant on market
  - Transaction fee: 10% up to $40,000 but not less than $5,000
- Professional Services: Individual case basis
  - Remote Hands: Discounted prices according to purchased plan.

WaveCo: In addition to dedicated waves on the Internet2 Network, members also will have access to dedicated circuits beyond the 13,000-mile Network footprint. Internet2 WaveCo, a new service in collaboration with the network carrier, Level 3 Communications, will provision circuits – optical carrier and digital signal – for regional networks and other Internet2 participants needing to cost-effectively extend their connectivity anywhere on the Level 3 optical network. WaveCo also provides the advantage of national-scale aggregated rates, which are especially beneficial for regional networks that wish to extend their reach to only a few sites with varying demands for circuit capacity or duration.

Fees determined on individual case basis.