The Internet2
Prepared by: Siew Fan Wong
under the supervision of Dr. Blake Ives

What is Internet2?

The Internet first gained popularity in universities as a tool for research and education. But its widespread popularity and exponential growth came only after commercial adoption. Today, the Internet has grown into a multi-billion dollar commercial phenomenon with hundreds of million of users throughout the world. But its widespread usage has led to bandwidth bottlenecks as well as a shortage of Internet Protocol (IP) addresses – the Internet's equivalent of street addresses. Universities and research institutions that require priority access to a high performance network for collaborating research efforts are now addressing these shortcomings.

In October 1996, representatives from 34 American universities met in Chicago. They created a non-profit consortium project called the Internet2. The Internet2 was intended as neither a separate physical network nor a replacement for the existing Internet. Instead, the Internet2 was conceived as an application-driven project[1] designed to be a complementary network to support advanced scientific applications.

Currently, University Corporation for Advanced Development (UCAID) operates Internet2. It receives financial support from more than 180 universities, over 80 companies and the U.S. government[2]. Today, the Internet2 project is dedicated exclusively to research. The goals of the project are to:[3]

- Create and sustain a leading edge network capability for the research community,
- Direct network development efforts to enable and support advanced technology development,
- Ensure the rapid transfer of new network services and applications to the broader Internet community, and
- Accommodate anticipated explosive growth.

What are the benefits of the Internet2?

The benefits on Internet2 are best described by comparing it with the Internet we are all familiar with — let's call it Internet1.

Internet1 was also initially intended for research. Today, however, the bulk of its uses are in the commercial sector. In addition, Internet1 is accessible to anybody who has a connection to the network, including through commercial Internet service providers. Internet2 use, on the other hand, remains limited to the research community, with connections available only through participating research institutions.

Most home users of Internet1 still connect to the network using dial-up modems. These typically have a top speed of 56kbps, though increasingly home users are connecting through cable or DSL modems offering speeds an order of magnitude faster. Or, if you are in your office, the fairly standard 10BaseT Ethernet card moves data at 10 megabits per second while the increasingly popular 100BaseT Ethernet cards transfer data at 100 megabits per second. The Internet2 provides much greater speed and bandwidth. The network is currently operating at 2.5 Gbps (gigabits per second), and is expected to grow to 10 Gbps in the near future. That's two orders of magnitude greater than today's LAN speeds. Let's look at a practical example. To download a DVD version of the motion picture "The Matrix" from the Internet requires about 171 hours with a 56k modem. Downloading with an ISDN connection requires a 74-hour wait while users with a DSL or cable connection must still wait 25 hours. Even those fortunate to have a T1-connection must be patient for over 6 hours. But, should you be lucky enough to have an Internet2 connection you can download the movie in the time it takes you to walk to your refrigerator and retrieve a beverage.

Obtaining that difference in speed requires a completely new design of the network. Internet1 network infrastructure components, including subnetworks, operate at varying speeds, with slower components creating bottlenecks for the rest of the network. Internet2, in contrast, is faster and more robust because all the networks it connects have the same high-speed infrastructure.

	Internet1	**Internet2**
Purpose	Commercial and research usage	Research usage
Users	Everyone with a connection to the network	Currently available exclusively for research communities
Typical Speed	56kbps	Initially, the network will operate at 2.5 Gbps, but the capacity will be boosted to 10 Gbps in the near future[1]
Protocol	Ipv4 with 32-bit address	Ipv6 with 128-bit address
Number of IP Addresses	More than 1 billion (in actuality, it is about 2 million sets)	3.4×10^{38}
Prioritization of packets	No[4]	Yes[5]
Ways of Communication	Unicast (one-to-one)	Multicast (one-to-many)

Another major innovation in Internet2 is the addressing scheme. Internet1 employs Internet Protocol version 4 (Ipv4). This protocol uses a 32-bit IP address system that allows more than one billion IP addresses. One billion addresses may seem like a lot, however, the way Ipv4 assigns class systems in sets to users significantly limits the number of usable addresses. For example, an organization might be assigned a Class B address, which contains a total of 65,000 possible addresses. However, this organization may use only a few thousand of its available address space. The others will therefore go unused. Such class addressing systems significantly limit the number of possible IP addresses to about two million sets. Better management or IP addressing schemes could meet some of the future need. But even a billion addresses is a small number when we consider an emerging world where microprocessors, connected via wireless technology, are embedded in humans and animals, attached to packages, or integrated into consumer devices of all sizes. The employment of Ipv6, a new version of Ipv4, provides a solution to this problem. Ipv6, which is backward compatible with Ipv4, uses a 128-bit IP address system that allows a maximum of 3.4×10^{38} IP addresses[10]. Dividing this number by the earth's surface (about 511,263,971,197,990 square meters), Ipv6 allows approximately 665,570,793,348,866,943,898,599 addresses per square meter of the earth surface[10]. Such a dramatic increase in the number of IP addresses facilitates the handling of the fast growing number of users and devices connected to the network.

Another innovation in Internet2 is the capability to prioritize the data packets, thus improving transmission of live data feeds. In Internet1, all types of data packets, from email to audio to video files, have the same level of priority. In cases where too many emails flow into the network, the quality of video files might be affected. At times, some packages might be lost or delayed during the data transmission processes. Internet2 designers tried to resolve this issue by introducing the concept of Quality of Service (QoS). QoS prioritizes data transmissions and assigns the necessary bandwidth to the associated transmissions. Under QoS, packages are marked as higher (e.g., time sensitive packets such as video or audio) or lower priority (e.g., emails or web downloads). Higher transmission packages will be assigned the necessary bandwidth and be sent to the receiver faster than lower priority packages that will be queued until necessary bandwidth is available. With such prioritization, the performance of data transmission through the network will be improved and data loss and delay will be reduced or eliminated.

Finally, the multicasting feature of Interenet2 will enhance information distribution. Internet1 only allows unicast communication. This means if a sender wants to send information to multiple recipients, the originating server must transmit a separate data stream to each recipient. As the number of recipients increases, the transmission generally increases the congestion in the network. Multicasting technology in the Internet2 solves this problem by allowing the distribution of information to multiple recipients simultaneously on a real-time basis. The sender posts one copy of the information to the network. The routers in the network then duplicate the information and route it to the hosts that require the information. This will significantly reduce network traffic loads. Such technology is especially useful in supporting multimedia information distribution that involves interactions among a large number of people. For instance, astronomers located at various universities might share the data feed from a single large telescope, perhaps taking turns controlling what objects it will be focused upon. That same data feed might also be made available to high school astronomy classes.

How does the Intenet2 work?

The Internet2 projects are supported by two types of fiber-optic backbones, (1) The vBNS (very high-speed Backbone Network Service), and (2) Abilene.

The vBNS is a noncommercial network launched in April 1995 to interconnect a number of supercomputer centers in the U.S. for research purposes by providing connections to the four national network access points (NAPs). This network is funded by the National Science Foundation (NSF) and is currently run by MCI Worldcom based on a contract with the NSF. However, the NSF maintains the rights to determine the accessibility of this network for institutions and individuals (the time and use of the vBNS and supercomputers) based on the description of individual projects submitted to the NSF committee. The vBNS has the ability of handling data at up to 2.5 Gbps[6] using MCI's advanced switching and fiber optic transmission technologies. Currently this network serves as the initial interconnection for the Internet2 members.

Abilene is a 10,000 mile advanced backbone network launched in the fall of 1998 to support advanced Internet2 applications. This network, developed by UCAID, and currently operated by Qwest Communications, has the ability to operate at up to 10 Gbps[7]. Abilene connects regional network aggregation points, called gigaPoPs (Gigabit Points of Presence). GigaPoPs are high-speed circuits, distributed geographically across the U.S. and allow institutions and local or state networks to have a high-speed connection. Two types of gigaPoP exist; Type I gigaPoP provides access to the network shared with other gigaPoPs while Type II gigaPoP supplies connections between gigaPoPs and other networks such as Internet1. Hence, every institution just needs to build a connection to a local gigaPoP instead of building its own advanced network.

What are some of the applications developed with the Internet2?

The Internet2 application development is focusing on tele-immersion, digital libraries, and virtual laboratories.

Tele-immersion allows people from different locations to meet in a virtual environment with the perception and feeling that they are sharing the same physical space. In this type of environment, every presence and movement of physical objects and individuals are tracked and then projected through ‘tele-cubicle’ to other parties who shared the same virtual environment. Tele-cubicle is a canvas, stereo-immersive wall that appears as a window into another physical space to facilitate a transparent passage to other physical spaces. Using these passages, two parties working from two different locations will feel as if they are at the same location. This technology is useful in distance learning, telemedicine, and data mining environments. For example, in the future, doctors would be able to see 3D images of their patients, do tele-diagnosis, and even perform remote operations. Already a violinist and his teacher have used Internet2 to play in harmony from several hundred miles apart[12].

Digital libraries will incorporate texts, high-resolution graphics, continuous audio and video, animations, and three-dimensional maps into the information library. With digital libraries, information will flow more smoothly for network users by, for example, organizing user profiles that tracked each individual user’s interests. When new books or journal articles are released, the system would notify users who might be interested in them. If the materials are in digital form, systems could automatically e-mail these journal articles or even entire books to interested parties[7].

A virtual laboratory is a distributed problem-solving environment where scientists from all around the world can work on collaborative projects as if they are co-located. Virtual laboratories are particularly useful in areas such as astronomy, weather prediction, and design and manufacturing. For example, the process of designing an airplane might require designers to access many technical and manufacturing specifications databases that are provided by contractors at many different locations. Virtual laboratories facilitate simulations and provide members of the airplane design team with simultaneous access to these different databases[3].

When will these benefits be commercially available?

It is anticipated that advanced applications and technologies developed under the Internet2 project will eventually be transferred to the commercial network. Consourse Group predicts that the commercial use of some of these advanced applications and technologies will begin by 2005[8]. They predict that some capabilities such as multicasting and QoS, and the upgrading of routers to handle greater network speed may come as early as 2002. Some firms are also partnering with Universities to get an early introduction to the benefits Internet2 will provide. Ford Motor Company, for instance, has entered into a partnership with the University of Michigan Business School to explore how Internet2-enabled network applications, such as multi-site TV-quality videoconferencing, can improve work among teams at locations around the world[11].

Who is involved in the Internet2 at local Universities?

As mentioned before, more than 180 university members and over 80 industrial members are participating in the Internet2 project. The following list provides contact information for participants at some local universities.

Universities	Contact Name	Title
University of Houston	Charles Shomper	Vice Chancellor/Vice President Information Technology
Baylor College of Medicine	J. Robert Beck	Vice President for Information Technology
Texas A&M University	Richard Ewing	Dean, College of Science
Texas Tech University	Sam Segran	Asst. Vice President for Technology Services
University of Texas at Austin	Tom Edgar	Director, Academic Computing and Instructional Technology Services

What are the problems of the implementation of Internet2?

Many applications that are currently included among the Internet2 projects require greater infrastructure than are available through the existing Internet. Thus, implementing Internet2 applications needs a new generation of hardware such as routers and switches, and carrier services. It also demands new protocols, new network management tools, and a deeper understanding of the network needs of high-performance applications. All these may mean costly and time-consuming processes that some users might resist. Upgrading the public infrastructure to achieve the requirements of Internet2, therefore, might not be accomplished in the near future.

In addition, joining the Internet2 project requires considerable financial commitment. Participating universities have to spend about $500,000 to participate in the project and in the development of new applications. They must also pay $110,000 or more annually to connect to the Abilene network[9]. Such funding may not be feasible for every university and research institution and therefore would deprive some universities and research institutions of the chance to participate in the project and create obstacles to their future research endeavors.

What is the difference between Internet2 and NGI?

Some confusion exists on the difference between Internet2 and the Next Generation Internet (NGI). Such confusion is not surprising as both the Internet2 and NGI were launched in 1996, and share many common objectives. Both Internet2 and NGI intend to develop advanced technology that will benefit the research and practitioner communities but the Internet2 focuses more on application development while NGI focuses more on network infrastructure development. However, the main difference between Internet2 and NGI is that the Internet2 is a project led by university and jointly funded by universities, corporations and the U.S. government while NGI is a project led and funded only by the U.S. government. With the complementary nature of the Internet2 and NGI, both parties have already collaborated in developmental efforts. For example, a joint engineering team of NGI and the Internet2 members was formed to coordinate joint projects between the two parties. Furthermore, many federal agency representatives are currently participating in the Internet2's specialized workgroups to contribute their expertise to projects. Such collaborative efforts between the Internet2 and NGI will provide synergy in their effort to improve the future Internet technology.

What is going on in the rest of the world?

In a project similar to that of the Internet2, European academic and research communities launched a fast pan-European network called Geant in November 2000. The goals of Geant are to provide interconnection for each country's national research and education networks (NREN) in Europe, and to facilitate advanced technology development. Geant currently operates in 32 countries and serves more than 3,000 academic and research institutions. For each member country to tap onto Geant, it has to have its own NREN such as Janet in UK and RCTS in Portugal. With Geant, researchers across Europe can collaborate easily. For example, researchers from UK and researchers from Portugal can connect to each other through Geant. Similarly, researchers from Europe and the US continent can work together in a fast network through connections between Abilene and Geant[12]. Such network connection can enhance effective and efficient global research effort.

For More Information

Online resources

[3]To learn more about the Internet2 project, its developments and updates: http://www.internet2.edu

To know more about the Next Generation Internet (NGI) initiative: http://www.ngi.gov/

[10]For information on Ipv6: http://playground.sun.com/pub/ipng/html/INET-IPng-Paper.html

To read more about vBNS: http://www.vbns.net/

To find out more about Abilene: http://www.internet2.edu/abilene/

[12]To learn more about Geant: http://www.guardian.co.uk/internetnews/story/0,7369,589431,00.html and http://www.dante.net/geant/

[11]For more information on Ford Motor's involvement in the Internet2 project: http://archives.internet2.edu/guest/archives/I2-NEWS/log200112/msg00008.html

Articles

Anonymous. "Here Comes Internet2." *Popular Electronics*, August 1998, page 26.

[6]Cawkell, Tony. "Internet Futures: Information Capacity And Information Retrieval." *Information Services & Use*, 1999, page 107.

[2]Guerrero, Yazmin. "Internet2: The Widely Touted Sequel Looks Set to Bridge At Least Two Borders." *America's Network*, 2001, page 107.

[8]Lipschultz, David, "Internet2 Puts Broadband To The Test," Internetweek, 2001.

[4]McGarvey, Joe. "Internet2 Arrives, But Is It Still Relevant?" *Interactive Week*, 1999, pages 1-5.

[5]Smith, Veronica. "Internet2 Building The Next Backbone," *InternetWeek*, 1998, page 48.

[1]Van Houweling, Douglas E. "Building The Internet's Future: Internet2, ICAID and NGI." *Educom Review*, May/Jun98, page 16.

[9]Vincent Kiernan. "Will the Next Generation Internet Create a New Class of 'Have-Not Universities?" *The Chronicle of Higher Education*, October 23, 1998, page A23.

[7]Young, Jeffrey R. "Searching for 'Killer Applications." *The Chronicle of Higher Education*, August 8, 1997, page A22.

Important Books To Read

Quality of Service: Delivering QoS on the Internet and in Corporate Networks by Paul Ferguson and Geoff Huston.

Bauer College of Business Administration
University of Houston, Houston, Texas, 77204-6283
Phone: 713-743-4691; Fax: 713-743-4693; e-mail: isrc@uh.edu

ISRC