D-Lib Magazine
April 1998

ISSN 1082-9873

All Packets Should Not Be Created Equal

The Internet2 Project

William H. Graves[1]
COLLEGIS Research Institute
Chapel Hill, North Carolina

wgraves@collegis.com

Internet2: Why and From Whence

Today's packet-switched Internet had its genesis in the American research community at a technological moment when assigning equal status to all packets made sense and worked. That best-effort, one-level-of-service-fits-all model of data transmission, however, has developed some severe shortcomings as information technology has advanced dramatically beyond its capacities in the era of the seminal ARPAnet experiment. Consider, for example, the following two scenarios.

Data from a telescope might be transmitted in real time to a group of astronomers working at different sites around the world. They might collaborate in real time to analyze the data and to decide how to adjust the aim of the telescope to optimize the value of the data being collected during the session. The same data might be multicast to amateur astronomers who have "subscribed" to the appropriate "channel." These amateurs probably would not have the privilege of participating with the collaborating experts who are controlling the telescope, but they might have real-time (or delayed) access to the data generated by the telescope. The archiving of such massive visual and numerical data sets raises interesting indexing and retrieval issues beyond the scope of current research in textual indexing and retrieval.

A student studying the Spanish language might be watching and listening to a digitized video of a conversation between two Spanish-speaking young people that is streaming across the Internet from a distant archive of such conversations. How instructors might locate such materials and negotiate their use by students is one of the issues being addressed by the Instructional Management Systems(IMS) Cooperative, while the indexing of such materials is a concern of both the IMS Project and the more general research community interested in the information science of digital video information.

Both of these visions can almost be realized on today's Internet, but, in the final analysis, both demand a quality of service that cannot be extracted from today's Internet. That quality of service sometimes may depend solely on the high-probability availability of raw bandwidth, but it often depends on the availability of differentiated network services designed, for example, to insure against problems of latency in delivering 30-frames-per-second, synchronized video from distant server to personal computer.

The growing possibilities, such as the two described above, for bridging and expanding the capacities and capabilities of the multimedia computer and the Internet have attracted broad attention in the higher education community. Whether working from the inside or the outside, many with a stake in higher education envision a future when on-line tools for communication and collaboration and on-line learning resources are broadly deployed, not only to strengthen the traditional classroom, library, and laboratory experience, but also to extend higher education's reach with convenient and flexible anytime, anyplace modes of instruction, research, and public service. Today's Internet, however, cannot support the deepest expression of that vision - as illustrated by the two examples requiring differentiated network services and new applications.

The need for differentiated network services and related quality-of-service guarantees loomed large for the leaders from the higher education information technology services and information resources communities who gathered in 1995 at a conference in Monterey, California, with experts from the network industry and the federal and regional network communities. Those differentiated-service needs and their implications for higher education inspired a group of conference attendees to initiate a grass-roots movement that evolved into the Internet2 Project. That grass-roots initiative rapidly migrated into a more broadly representative "movement" supported by Educom, FARNET, the National Science Foundation, other organizations and agencies, and the information technology industry. The resulting series of meetings and workshops motivated over 40 university CIOs to announce the Internet2 Project on behalf of their universities in the fall of 1996. Those institutions committed enough project funding to hire a project director and a staff to operate under the aegis of a small Internet2 Steering Committee. Internet2 membership had grown to include more than 100 institutions by the time a year later when the Steering Committee won the membership's approval for creating the nonprofit University Corporation for Advanced Internet Development (UCAID) in the fall of 1997.

UCAID was created to advance and manage the Internet2 Project and to provide a permanent mechanism through which higher education could play a continuing role in the evolution of the Internet even beyond the Internet2 Project. The reader can gain both a more detailed historical perspective on the Internet2 project and an up-to-date perspective on Internet2 planning and progress by browsing www.internet2.edu and www.ucaid.edu. There are several principles at the heart of the UCAID mission.

Key UCAID Principles

UCAID is committed to advancing its agenda in an open-standards environment built on the leverage of a shared approach to resolving issues of advanced internetworking. The architecture for Internet2 emphasizes national (and global) approaches to interconnecting Internet2 "gigapops" -- the regional, high-capacity Internet2 interconnection "points of presence" that serve as aggregation points for traffic from participating institutions and that can be designed with considerable latitude to meet local or regional needs. Beyond such architectural considerations for continuously migrating to the next generation of internetworking technologies, however, lie some broad principles of policy and purpose.

A Strategic Role in Internetworking for Higher Education: The continuing development of the Internet is so strategic to higher education that those colleges and universities with both needs-driven interest and the resources to contribute to that development should have a well organized and supported mechanism -- UCAID -- for advancing their collective interests in a manner that over the long term benefits the entire higher education community and the broader Internet community.

An Open Membership Model: Membership in UCAID is open to any college or university that can direct resources and expertise to incorporating advanced technologies into its institutional network and to connecting its institutional network to the UCAID national internetworking fabric while also developing exemplary applications that require the advanced features of that fabric. UCAID also makes provision for the participation of the network industry and various nonprofit organizations with a stake in the goals of advanced networking in higher education.

A Representative Agenda and Governance Model: UCAID is managed and led by a chief executive and governed by an elected Board of Trustees comprised in the majority of CEOs from member institutions to ensure that UCAID serves long-term, strategic needs of higher education. Three expert advisory councils counsel the leadership. Two of these, the Network Planning and Policy Advisory Council and the Applications Strategy Council, help ensure that the UCAID production network fabric is planned, operated, and used in a manner that meets two goals. It is consistent with the advanced application needs of member institutions, and it meets the long-term need for open standards, sustainable affordability, and technology transfer to the commercial sector. The Network Research Liaison Council advises on the transfer of leading-edge research on computer and network systems into the UCAID networking fabric.

A Nimble Governance Model: The UCAID advisory, governance, and management structures are designed to ensure that UCAID is not only responsive to the strategic needs of member institutions, but also nimble enough to take leadership action on behalf of higher education at the forefront of the fast-paced transition from one life cycle of internetworking technologies to the next.

Internet2 in the Life Cycle of the Internet

Investing wisely in campus networks is a tall order, especially when the pace of technological change is accelerating and technological expertise is scarce and difficult for many institutions to hire, manage, and retain. An investment in networking provides a good institutional return on investment when the network and the application services it supports can be affordably infused into the fabric of the institution to increase institutional effectiveness. Affordability does not necessarily imply cost savings, although it might. But affordability, in combination with infusion, does say something about the nature of the technology and services in question and their support. The technology or service must scale to broad-based institutional use, and, in scaling beyond prototype use, must be supportable in an affordable manner. These are the characteristics of a *commodity* or *core* IT service, such as an institutional e-mail service or a set of Web-based services giving students access to institutional information about the academic and social services available to them, authenticated access to their transcripts, and authenticated access to instructional resources for their courses. UCAID institutions are trying to migrate to the next generation of core services. They are committed at any moment to providing ubiquitous access to today's core or commodity network services while at the same time investing in next-generation network technologies and application services. In other words, UCAID institutions are committed to riding the life cycles of internetworking and its applications. There is considerable risk in such a commitment, and one of UCAID's major roles is to reduce that risk by providing the leverage of a collective investment in the shared long-term goal of advancing the capacity and functionality of the Internet.

To discuss the relationship of Internet2 to the Internet, it will be useful to introduce a conceptual model based on the concept of the life cycles of innovative products and services - a model that guided the deliberations of the Internet2 Steering Committee. The following graphic depicts that model.

Life Cycle Model

Most innovations that depend on technology evolve through four stages: *experimentation*, *incubation*, *commercialization* (or, in the nonprofit sector, *implementation* or *roll out*), and "*commoditization*" - the delivery of the product or service to the largest possible audience at the lowest possible cost. In this model, experimentation and incubation costs are often high and at risk and, so, must be shared or in some way subsidized. In contrast, commercialization (or implementation) and commoditization must be self-sustaining (profitable in the for-profit context) - often with some portion of the profits defraying the costs of the next life cycle of experimentation and incubation.

These ideas apply very nicely to the Internet and Internet2. The Internet derived from the ARPAnet experiment and is evolving today into its next life cycle through the Internet2 Project and the federal Next Generation Internet initiative. A graphic and explanation follow.

ARPAnet → NSFnet → commercial Internet → commodity Internet → vBNS → Internet2

Experimentation: An innovative product or service often begins life as an experiment or a research study underwritten by a company or by some combination of commercial and nonprofit interests willing to share the pre-competitive costs of creating new products and markets. The Internet, for example, can trace it origins to ARPAnet, the experimental network designed to connect a few computers serving defense-related research projects. The government subsidized the original ARPAnet experiment.

Incubation: An experiment or research study that points the way to a new product or service is often followed by an incubation period in which a viable prototype is designed and tested. In the case of the Internet, the NSF recognized the potential of ARPAnet technologies to enhance research broadly in science and engineering and accordingly funded the NSFnet to support research and education among participating institutions. But NSFnet was also designed to demonstrate the scalability and affordability of open-standards, packet-switched TCP/IP inter-networks. As the national backbone of the network of networks, which came to be called the Internet, the NSFnet exemplified the leverage to be gained by cost sharing and the adherence to shared "standards" (among participating colleges and universities, companies, and government agencies).

Commercialization: A successful incubation effort typically leads to a host of market development activities and to commercialization (or production implementation, in the nonprofit case). To allow commercialization to extend the benefits of the Internet beyond the restricted domain of education and research, for example, the NSF stopped operating the NSFnet in 1995. This tactic succeeded. Commercial Internet Service Providers (ISPs) rushed in to provide a mesh of inter-connected wide-area TCP/IP networks and a multiplicity of means to connect private "corporate" networks, such as campus networks, to this new global Internet.

Commoditization: In today's competitive markets (profit and nonprofit), many products and services quickly become commodities available not just to the corporate market, but also to the consumer market -- VCRs, PCs, institutional home pages, and the Internet itself. Recognizing that the potential value of the Internet is increased by every new connection, the commercial sector extended the benefits of the Internet to the ends of the public telephone network through the dial-up modem pool. Internet service is today a $10-$20/month commodity available from a new breed of ISPs operating in the consumer market wherever there is universal telephone service with reliable telephone connections.

The Next Life Cycle: Commodities usually breed low-cost competition, which leads to a search for value-added services and eventually to a competitive race into the next life cycle of innovation. The next life cycle of the Internet is well underway. In 1995, the NSF initiated the next life cycle of experimentation by connecting a few supercomputer centers through its very high performance Backbone Network Service (vBNS). The announcement in the fall of 1996 of the Internet2 Project represented the need to move beyond experimentation and into incubation. Internet2 is designed to demonstrate the economic and technological viability of incorporating differentiated broadband services into wide-area networks on a national scale and with production values. The announcement shortly thereafter by President Clinton of the federal Next Generation Internet (NGI) initiative amplified the momentum of Internet2 by including the goals of Internet2 among the goals of the NGI. Internet2 thus became, in concept, higher education's agenda within the larger NGI effort. Or, to use a software metaphor, NGI/Internet2 will be a production-value, beta version of the next release of the commercial Internet and, eventually, the commodity Internet.

Common-Good Conclusion

Internet2 and the broader constructs it generates, such as UCAID, are attempts to serve the public interest in the most complex of environments -- the national and global internetworking environment. Patricia Battin captured the idea succinctly in "New Ways of Thinking about Financing Information Technology," (in Organizing and Managing Information Resources on Campus, ed. by Brian L. Hawkins, Educom Strategies Series on Information Technology, Educom, Washington, DC (1989), p. 369):

> [Information technology] makes possible an unprecedented decentralization of technical power to individual option while at the same time it requires a globally coordinated infrastructure to permit the effective individual exercise of that power.

What is optimal for a nationally (or globally) coordinated infrastructure may not always be optimal for individuals and their institutions, but UCAID is committed to an action agenda on behalf of higher education in advancing the Internet while leaving as much choice to individual option as possible.

About the Author

Dr. William H. Graves is President of the nonprofit COLLEGIS Research Institute, Senior Vice President of the higher education IT service company COLLEGIS, and a member of the COLLEGIS Board of Directors. He is responsible for the company's learning technology services that are based on the Institute's research and development efforts. These faculty professional development services and related instructional technology support services help colleges and universities incorporate on-line learning resources and communication tools in their instructional programs. Dr. Graves and his colleagues at the Institute came to their positions from the University of North Carolina at Chapel Hill's Institute for Academic Technology. Graves joined UNC's faculty upon completing his mathematics Ph.D. at Indiana University. He is on leave today after having served as senior information technology officer for eight years, having founded the IAT, and also having served in a variety of other academic administrative positions.

Dr. Graves is a member of the transition board steering the consolidation of CAUSE and Educom. He chairs the planning committee for Educom's National Learning Infrastructure Initiative and also serves on the steering committees for the Instructional Management Systems Project, the ARL/CAUSE/Educom Coalition for Networked Information, and the Educom Networking and Telecommunications Task Force. He served on the steering committee for the Internet2 Project and helped organize the University Corporation for Advanced Internet Development to carry forward the Internet2 agenda and subsequent efforts to advance the Internet to meet higher education's needs.

Dr. Graves has served in a consulting capacity for many institutions as they plan a role for information technology services in shaping their future. He is an accomplished public speaker on that theme and has delivered over 350 invited presentations on campuses and at conferences during the past 15 years. Dr. Graves has published over 40 articles on all aspects of information technology in higher education. He is a frequent contributor to CAUSE/EFFECT and the Educom Review. Some of his recent publications are accessible on-line from the Web Site of the COLLEGIS Research Institute.

hdl:cnri.dlib/april98-graves