SC2004 Keynote Presentation
9 November 2004
Pittsburgh, Pennsylvania

Remarks by Tom West, CEO of National LambdaRail

(as prepared accompanying slides can be found at:
http://www.nlr.net/presentations/SC2004-TWW-Slides.pdf)

[Slide 1]

INTRODUCTION

Pittsburgh represents something of a homecoming for me. I was born 30 miles north of Pittsburgh and spent the first 40 percent of my life growing up, being educated, and working in western PA. And, I still root for the Pirates and Steelers.

So I consider Pittsburgh my personal Mecca. In fact, diehards like myself feel this area of western PA and eastern Ohio, defined by the three rivers that flow in and out of Pittsburgh, has been the epicenter of significant events and people that have impact on our lives. To paraphrase David Letterman here are my Top 5 ways Pittsburgh has impacted our lives:

#5 If the Pittsburgh Supercomputing Center had not signed up a year ago to use three NLR 10 G waves between Pittsburgh and Chicago for the ETF, I would not be your guest speaker.

#4 If the British had not captured the Fort Duquense on the point down the street where the Monogehala and Allegheny rivers meet to form the Ohio we might in Duquense, Quebec Canada and all be speaking French today.

#3 If E. L Drake had not discover oil 60 miles north of here in 1859 and commercialized it, we might still be driving the horse and buggy.

#2 If some entrepreneurs and the federal government had not joined forces in 1938 to build and open the PA Turnpike 30 miles east of here in 1940 this country might not have an Interstate system.

#1 If H. J. Heinz had not created ketchup just across the river, what would we put on our sandwiches and French fries today?

As you can tell I think Pittsburgh has had an impact on our lives

[Slide 2]
The major theme of my remarks today is that a global network infrastructure is the key to building bridges of collaboration among the world's scientists and network researchers, and that it can best be achieved by **owning** and **managing** that infrastructure ourselves. Ownership gives control, and both the efficiency and effectiveness of research are enhanced.

Leadership from the R&E community and **partnership** with industry and the federal government agencies are the keys to success.

I have some supporting cast today via videos. Joining me will be David Farber, NLR's Chief Scientist and Distinguished Professor at Carnegie Mellon University; Professors Larry Smarr, University of California, San Diego and Tom DeFanti, University of Illinois at Chicago; and Doug Van Houweling, President and CEO of Internet2.

35 YEARS OF NETWORK DEVELOPMENT

[Slide 3]

Let's start by looking at the first 35 years of network development in the United States. For example, here is what the NSFNET backbone looked like in 1991. Note that T-1 speeds were the gold standard just 13 years ago. Today, most DSL and cable modems to the home far exceed that, and 10 gigabits per second (Gbps) is considered inadequate for many forms of network research and scientific applications.

[Slide 4]

If we go back to 1969, we can see that it took 25 years for the early Internet, and later the World Wide Web, to become commoditized.

[Slide 5]

When the research and education community realized that it had lost control of this critical resource to the traditional telecom industry, Internet2 was formed in 1996 and launched Abilene network with an OC-48 backbone in 1998, the GigaPops and regional networks followed, and finally the Abilene network moved to 10 Gbps speeds in 2003.

It is worth noting about that the research and education (R&E) community has always been in leadership positions in the nation's network development. On September 2, 1969, then graduate students Stephen Crocker, Vinton Cerf, and others joined UCLA professor Len Kleinrock in passing data between two computers. Three months later, there was a three-node network among UCLA, Stanford, and UC Santa Barbara. Visionary pioneers like these from the university community, with the support of government and industry, have always provided the leadership for network development. Why stop now?

CURRENT STATUS

Today, we have multiple networks for research and education running over multiple global, national, regional, and institutional infrastructures.

[Slide 6]

Some of these networks, such as the Internet2-Abilene OC-192c backbone that you see here are high performance networks that are shared by researchers from numerous disciplines.

[Slide 7]

Other networks are more specialized like ESnet that are designed to serve the needs of a subset of the entire research community.

An important fact is that most of the underlying infrastructure of these networks has not been owned by the R&E community. Rather, we have leased circuits from the traditional carriers. Often we need to cobble together circuits from multiple providers to meet our needs. As a consequence, research groups spend significant amounts of time and energy developing the network capability they need before they actually can do the scientific research.

SEA CHANGE

But “the times, they are a-changin’,” for at least two major reasons.

[Slide 8]

First, there is a growing urgency to develop new network technologies that “scale” to the growing needs of worldwide R&E community and downstream to the commodity internet users.

[Slide 9]

These projects, some of which you see here, require real-time collaboration among scientists and manipulation of enormous data sets. This, in turn, feeds the need for multiple dedicated, private research networks.
[Slide 10]

Applications in the arts, and for teaching and learning in almost every discipline, are now emerging with the same need for big, broadband networks.
The second big development is the fortuitous availability of dark fiber in the United States and elsewhere as a result of the downturn in the telecommunications industry. The last four years have provided a historic opportunity to migrate from leasing circuits from traditional telecommunications carriers to outright ownership of fiber.

[Slide 11]

This dark fiber, combined with optical dense wave division multiplexing (DWDM) equipment, enables multiple R&E networks to be built and run over the same fiber pair.
[Slide 12]

This changes everything in the development and management of dedicated research networks to support collaborative Big Science applications and large-scale network research. In short, we now have the components for owning and controlling a robust optical network infrastructure.

OPTICAL NETWORK INFRASTRUCTURE

[Slide 13]

I thought we needed a framework as a way to view what needs to be done to put into place a comprehensive network infrastructure to serve the R&E community. I offer to you this framework of rings that I believe make up the optical network infrastructure: institutional, regional, national, and international.

I also wanted to come up with a name to describe the groups and individuals leading efforts to deploy the optical network infrastructure, and NIPPERS seemed as good as any—Network Infrastructure Plumbing Persons for Education and Research. That's who we are—plumbers.

[Slides 14-17]

So the good news is that the United States is a player in global R&E optical infrastructure development. The bad news is that we are late to this game. You can see here that several other countries and continents already have operational optical fiber-based network infrastructures, notably Canada through CANARIE with its CA*Net 4, SURFnet6 in the Netherlands, and AARNet in Australia.

One reason these and other countries have leapfrogged the United States is the massive financial investments of their governments in optical infrastructure. Another is the very strong leadership of people like Bill St. Arnaud in Canada, Kees Neggers in the Netherlands, and George McLaughlin and Chris Hancock in Australia.

U.S. OPTICAL NETWORK DEVELOPMENT

[Slide 18]

Here is the relatively brief timeline of R&E optical network development in the United States. Notice that it began around 1995 with individual institutions deploying fiber campus-wide, moved to regional optical networks a couple of years later, and in 2003 and 2004, the nation's first R&E community owned optical network infrastructure—National LambdaRail or NLR—was formed and just recently completed Phase 1 of operational deployment.

[Slide 19]

Again, this pattern follows the framework of the infrastructure rings and players we saw earlier.

[Slide 20]

A large number of institutions, especially research institutions, are establishing on-campus, fiber-based network infrastructures that serve multiple networks, as illustrated here by the network architecture at the University of California Berkeley.

[Slide 21]

However, in many instances the institution goes far beyond the physical dimensions of the main campus and reaches out to university facilities in the surrounding community. Increasingly, it is important that these facilities be developed as part of the institutional infrastructure as illustrated here by the architecture at the University of California San Diego.

REGIONAL NETWORKS

[Slide 22]

Although regional optical network infrastructure development emerged a few years ago, the formation of National LambdaRail or NLR has spurred a virtual explosion in the number of regional efforts.

[Slide 23]

Here is where most of the action is. It is fair to say that research universities in states like Indiana, Illinois, Connecticut, and Ohio were early adopters.

[Slides 24-26]

In the last two years, a consortium of institutions within states like Texas and Florida have formed a not-for-profit corporation to undertake regional infrastructure development. Here is what the LEARN backbone of Texas looks like, and Florida's LambdaRail. California pioneered this model beginning in 1997 with the formation of the Corporation for Education Network Initiatives in California or CENIC, that brought together the public and private universities, and today the state's community colleges. CENIC's CalREN optical, fiber-based infrastructure provides multiple networks to serve this wide range of constituencies.

[Slide 27]

But this is just the tip of the iceberg. This chart shows the aggregate dark fiber assets acquired by U.S. R&E optical initiatives in terms of segment miles. Notice that of the roughly 27,000 total miles, about 55 percent is held by regional network organizations and the remainder is held by NLR. In the last two years, the NLR has purchased nearly 11,500 miles of dark fiber. FiberCo, a limited liability corporation created by Internet2, has been very instrumental in facilitating the acquisition of much of this fiber.

NATIONAL LAMBDARAIL

[Slides 28-29]

So now let's move to the national level of the optical network infrastructure. Less than three years ago the NLR was just a glimmer in the eyes of very few people in the United States. It started as a grassroots effort on the west coast to link from Seattle to San Diego. It then evolved to have a redundant path via Seattle to Denver and Denver to Los Angeles. The rest is history. In essence, NLR evolved from a regional effort but recently NLR has stimulated new regional developments.

On May 5, 2003 the NLR was established as a 501©3 not-for-profit organization, and this is what the national LambdaRail architecture looks like today. All of the backbone is owned fiber and optics.

[Slide 30]

Notice that the mission of the NLR is not to build and maintain an advanced network but an advanced infrastructure that will support many types and levels of networks for those working in the research, clinical, and educational fields.

[Slide 31]

The NLR supports both experimental and production networks, fosters networking research, promotes next-generation applications, and facilitates interconnectivity among high-performance research and education networks, whether regional or international. If there is one underlying theme to the NLR mission, it is collaboration.

[Slide 32]

The NLR members and associates span a wide geographic and organizational range including individual universities, boards of regents, consortiums of institutions, not for profit corporations, a supercomputing Center and the University of Pittsburgh; a limited liability corporation; and Internet2, a nationwide not for profit representing over 200 institutions, including some that belong to other NLR members. Associates include the Oak Ridge National Laboratory, Case Western Reserve University, and the Southeastern University Research Association.

[Slide 33]

The major strategic corporate participant has been Cisco Systems. The NLR would not have happened without the commitment of Cisco Systems to provide major resources---optics, routers and switches and a very strong focus on the goal of advancing the network research agenda of the nation. Level 3 Communications and WilTel Communications, as our predominate providers of fiber and related services gave special NLR consideration in the acquisition of fiber and in providing the related services.

[Slide 34]

All of this has yielded the five-year funding plan for NLR that you see here. The corporate commitment is in the form of discounts, donations, and human resources. The interesting thing about the NLR funding to this point is that direct support from the federal government has been entirely absent. This is in stark contrast to the history of R&E network funding in the U.S. generally where the federal government, educational institutions, and corporate partners shared in its development, a three-leg stool that has served the nation well.

So why haven't the federal agencies joined in? The short answer is they have not been asked—yet. Those of us involved in NLR felt that it was important for the R&E community to invest its own dollars first. That has been done. We are hopeful that in the next few months the agencies such as the Department of Energy, NASA, the National Science Foundation, the National Institutes of Health, and others will be convinced of the value that the NLR infrastructure brings to the nation and will invest in it.

[Slides 35-36]

The NLR infrastructure consists of owned fiber and equipment for Layers 1, 2, and 3. These wavelengths will support both experimental and production networks running over the NLR backbone. If you are interested in more details, please come to our booth 1153 here at SC2004.

COMMUNITIES SERVED BY NLR

[Slide 37]

There are two main audiences for the NLR—network researchers and researchers involved in big science applications, including supercomputing. Let me turn to NLR's Chief Scientist, Dave Farber to give you more insight.

[Slide 38: Video of David Farber, Chief Scientist, National LambdaRail]

[Slide 39]

As Dave stated, up to 50 percent of NLR capacity is being devoted to support network research projects under the auspices of an advisory council.

[Slides 40-41]

Here are the members of the NLR Network Research Council led by Dave.

CURRENT USES OF NLR

[Slide 42]

Next I would like to mention some of the uses of NLR. In fact, several applications being demonstrated at this conference are being supported by 10 Gbps waves of NLR. I encourage you to visit those booths and see how these applications are running on networks with their own lightwave or lambda.

[Slides 43-44]

The Extensible Terascale Facility or ETF supported by the National Science Foundation, is a multi-million dollar, multi-year effort that has built and deployed a world-class networking, computing and storage infrastructure—the TeraGrid --- designed to engage the science and engineering community to utilize to catalyze new discoveries. One of the original TeraGrid participants—the Pittsburgh Supercomputing Center was the first to use NLR to connect its facilities to the overall TeraGrid. Recently, the Texas Advanced Computing Center acquired a 10 Gbps wave from NLR to connect Austin to Chicago. Oak Ridge National Laboratory also is using the NLR for back-up waves between Atlanta and Chicago as part of ETF.

[Slides 45-46]

The HOPI project of Internet2 is exploring the evolution of the Internet's core. This project is engaging industry, regional, and international partners to examine a hybrid of packet switching and dynamically provisioned lambdas.

[Slides 47-48]

The CENIC organization in California and the Pacific Northwest GigaPOP have agreed to cooperate in a joint project to create, deploy, and operate an advanced, extensible peering facility along the entire Pacific Coast of the U.S. Known as Pacific Wave, this project will create a new peering paradigm by removing the geographical barriers of traditional peering facilities. Pacific Wave will enable any U.S. or international network to connect at any location along the U.S. Pacific Coast facility, as well as the option to peer with any other Pacific Wave participant regardless of their physical connection.

[Slides 49-50]

This is the overall Department of Energy (DoE) UltraScience project map between the DoE national lab and university partners, and the NLR portion linking Sunnyvale, Seattle, and Chicago. The UltraScience Net is an experimental research test bed funded by DOE Office of Science to develop networks with unprecedented capabilities to support distributed large-scale science applications.

[Slides 51-52]

The "OptIPuter" is a five-year research program led by the University of California, San Diego and the University of Illinois at Chicago with several partners. This is the map of NLR lambdas supporting the OptIPuter from University of California, San Diego and San Diego State University in the southwest to the University of Washington in the northwest to the University of Illinois at Chicago in the Midwest.

The OptIPuter is powerful distributed cyberinfrastructure supporting two major data-intensive scientific research and collaboration efforts in the Earth sciences and biosciences.

Here are Tom DeFanti of the University of Illinois at Chicago and Larry Smarr of the California Institute for Telecommunications and Information Technology at the University of California San Diego and their colleagues.

[Slide 53: OptIPuter Video]

FUTURE USES OF NLR

[Slide 54]

At its meeting last week the NLR Board of Directors endorsed the recommendation of the NLR Network Research Council to collaborate with the Virtualized Testbed for Networking Research Project led by professors Tom Anderson, University of Washington; Larry Peterson, Princeton; Scott Shenker, Berkeley and Jon Turner, Washington University, St. Louis. The purpose of this project is to bridge the gap between cutting-edge research and production usage. This testbed would allow multiple virtual networks to co-exist on a common substrate.

[Slide 55]

Looking ahead to 2007, ESnet is planning to use NLR waves to connect several DoE labs across the country to the existing ESnet at high speed, and to provide backup paths and additional capacity between the west and east coasts as one means for serving the Department of Energy's needs.

[Slide 56]

The decision is still pending on how the NLR backbone might be used by Internet2 for its next generation high performance shared research network. I would like you to hear from Doug Van Houweling about the relationship between Internet2 and NLR.

[Slide 57: Video of Douglas Van Houweling, President and CEO, Internet2]

NLR AND COLLABORATION WITH OTHERS

[Slide 58]

These principles suggest that cooperation and collaboration on common goals are the hallmark of the NLR. The relationship with Internet2 is one of the most important, and Internet2 has a major investment in NLR. However all of this requires cooperation and coordination between NLR and a number of organizations involved in this endeavor.

[Slide 59]

The NLR focus is on the full range of infrastructure needs of the R&E community. This includes stable and reliable production networks at the regional, national, and international levels, and "breakable" experimental networks in support of network research. The bottom line is that the NIPPERS I mentioned earlier cannot do the infrastructure development to facilitate applications research without the help of network researchers—it is a symbiotic relationship between both research communities and network organizations.

WORLDWIDE CONNECTIVITY IS ON THE WAY

[Slide 60]

The entire optical networking effort cuts across both nations and continents—North and South America, Europe, Asia, Africa, and Australia. The NIPPERS are busy worldwide.

[Slide 61]

For example, Ed Fantegrossi and Don Riley, representing the International Education Equal Access Foundation (IEEAF), have successfully secured donations from Tyco Global Networks that provide connectivity from the west coast of the United States to Asia, and from the east coast to Europe.

[Slides 62-65]

As you see here there are a number of activities going on in Europe to connect with the US. This slide illustrates the view from the Netherlands.

On the other side of the globe similar efforts are under way in Southeast Asia.

Our colleagues in South and Central America (CLARA) have recently launched the ALICE Network. This network will connect to the North America in Tijuana.

As you can see here there are efforts to ring Africa and link this ring to Europe.

[Slide 66]

Over the past four years there has been an effort to bring together the key NIPPERS from around the globe to share ideas and to collaborate on making interconnections happen. Spearheaded by Tom DeFanti, University of Illinois Chicago and Starlight and Kees Neggers, SURFnet 60 leading plumbers (NIPPERS) met in Nottingham England on September 2 for the fourth time and officially formed the Global Lambda Integrated Facility (World).

This event occurred 35 years to the day from the historic event when ULCA Professor Len Kleinrock and graduate students Vint Cerf, Stephen Crocker and others passed the first data from one computer to another in the UCLA lab.

SUMMARY

If it is not clear already, let's summarize *why* we NIPPERS are doing this, and why NLR exists as part of this effort.

[Slide 67]

Large scale scientific applications, many driven by supercomputing, are becoming more routine. However, there is a looming collision between applications requirements and network capacity. We have an historic opportunity for the R&E community to leverage technology and achieve control over advanced network resources, and lay the foundation for a new round of innovation.

Think about the 5 C's—collaboration, control, connectivity, comprehensiveness, and cost. In the final analysis, ownership and control of the basic infrastructure should provide the most cost-effective way to meet the full range of networking needs. We should spend zero time on how to connect the participants in large-scale research efforts. Instead, the network should be as ubiquitous and easy to use as a light switch or telephone. The R&E community has always led the way in network advances—it can continue to do so.

[Slide 68]

Bottom line—think globally and act locally. To that end, as a proud plumber in the trenches, I offer you the pledge of the NIPPERS. We, the NIPPERs, pledge to make sure that the light paths are many and that they are always clear to pass the vital data that advances research and education, which, in turn, contributes to a better worldwide community. NIPPERS unite for a bigger and better network infrastructure for everyone!

Thank you.