REMARKS by Prof. Peter A. Freeman
for
G-Lab Kickoff Meeting
TU Kaiserslautern
24 July 2009

Good afternoon!

I have visited Kaiserslautern a number of times over the past ten years and it is my pleasure to be back. It is my distinct honor to be able to help "kickoff" the G-LAB experimental platform project. I was delighted to be asked to join the G-LAB Scientific Advisory Board, and am looking forward to attending EuroView2009 in Würzburg next week.

From my perspective of involvement with GENI (three days ago I was at the Fifth GENI Engineering Conference in Seattle) and other advanced networking projects (I attended and spoke at the New Generation Networking Forum in Tokyo in April) I want to do two simple things this afternoon: Review the case for extensive, experimental research such as you are undertaking and emphasize the importance of international collaboration.

Those of you here today that are engaged in G-LAB and similar projects don't need to be reminded of why you're doing this research, but if Germany is anything like the rest of the world in this regard, there are others that may not fully grasp the importance of advanced experimental research in networking.

Most of us can remember when there wasn't an Internet – and thus no users – but it is worth remembering that the growth of the Internet has been truly phenomenal:

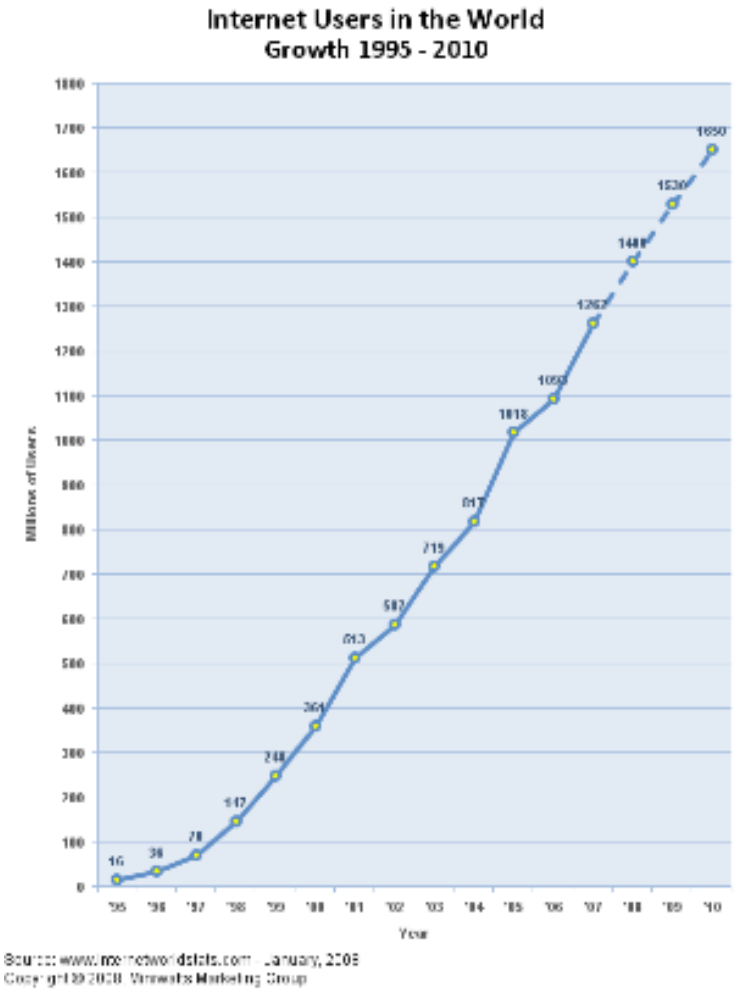

Just considering the number of users, there are easily over 1.5 billion users today – a growth by a factor of over 100 in just fifteen years. This isn't just a phenomenon in the advanced economies, but something that is worldwide. The penetration rate in North America is almost 75% and worldwide it is approaching 25% of the world's population.

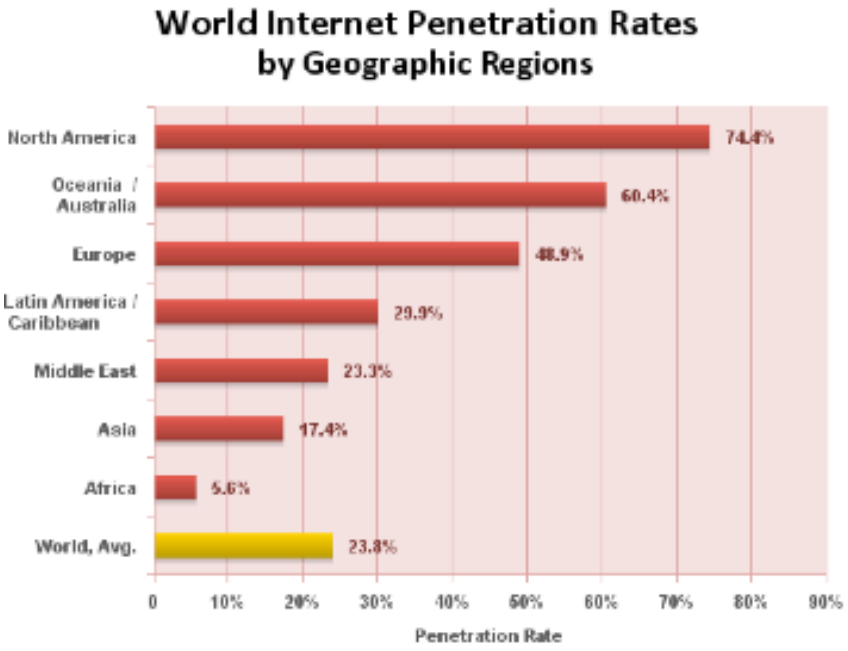

Source: Internet World Stats - www.internetworldstats.com/stats.htm
Penetration Rates are based on a world population of 6,710,029,370
and 1,596,270,108 estimated Internet users for March, 2009.
Copyright © 2009, Miniwatts Marketing Group

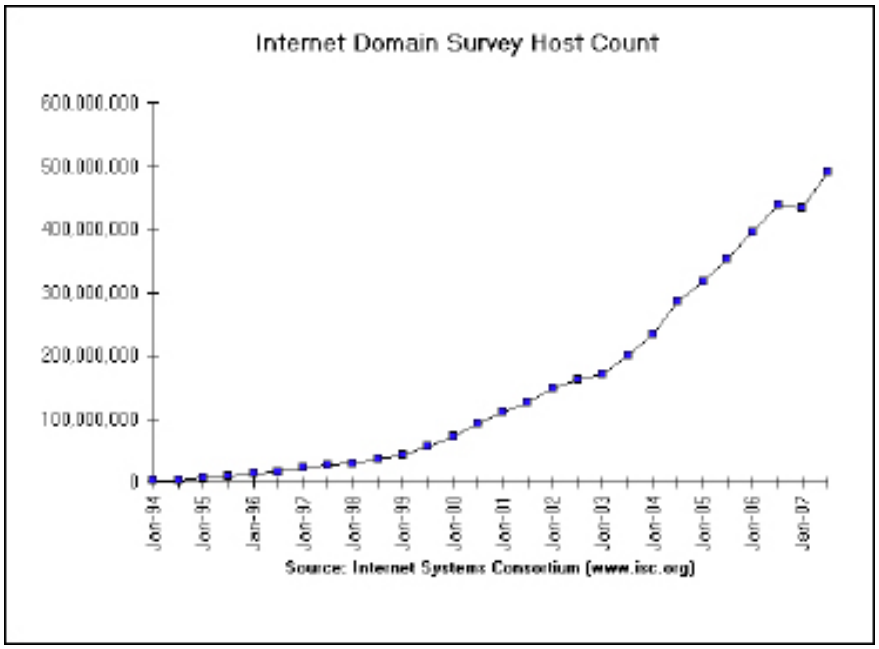

The sources of information on the Internet has also grown rapidly from zero to over 500 million hosts today – and a single host often provides information on multiple websites, so that there may be upwards of 30 billion websites in existence today – five times the world's population!

Sheer numbers, while impressive, don't begin to characterize the impact the Internet has already had:

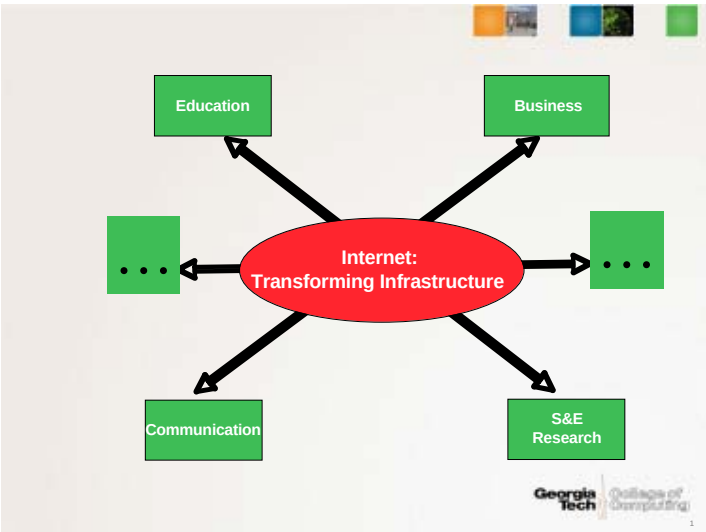

It has transformed communication almost beyond comprehension, profoundly changed scientific research and technological development, created entire new industries and completely changed others, has made significant inroads in the education world, has added whole new dimensions to individual behavior patterns, and is starting to revolutionize even government.

If the current networking structure has permitted such unprecedented growth, then one can rightfully ask “Why undertake research aimed at changing it?”

There are several answers that seem to have wide acceptance among those most knowledgeable about the technical structure of the Internet.

First, other than the basic TCP/IP routing suite, today’s Internet is NOT the same as twenty years ago. There have been over a thousand engineering changes and workarounds. While these have enabled much of what we experience today when using the Internet, they are becoming increasingly difficult to implement and are sometimes mutually exclusive.

Think of the classic VW Beetle. Even though it had all the basic features of a modern car (motor, steering wheel, seats, et cetera), today we expect much more. Designers might have extended the VW beetle with many new features (stronger motor, power steering, power seats, air bags, and so on), but it got to the point where to add a new feature they would have to rip out or radically change an existing feature. At some point it made more sense to design a new car from the ground up that would support these new features – and others to come - without ripping out and losing previous features.

Second, there are a number of problems and issues with the Internet today for which we believe sufficiently rigorous and efficient solutions cannot be devised without fundamental changes in the architecture and basic components of the Internet. These include technical (e.g. security from attack), social (e.g. protecting children without

restricting access by adults), policy (e.g. insuring open access), and legal (e.g. protecting intellectual property). While piecemeal solutions can and are being proposed, most are only partial solutions at best and may have undesirable side effects. The level of strategic concern is reflected in the fact that President Obama has made cybrersecurity one of his top priorities for the United States.

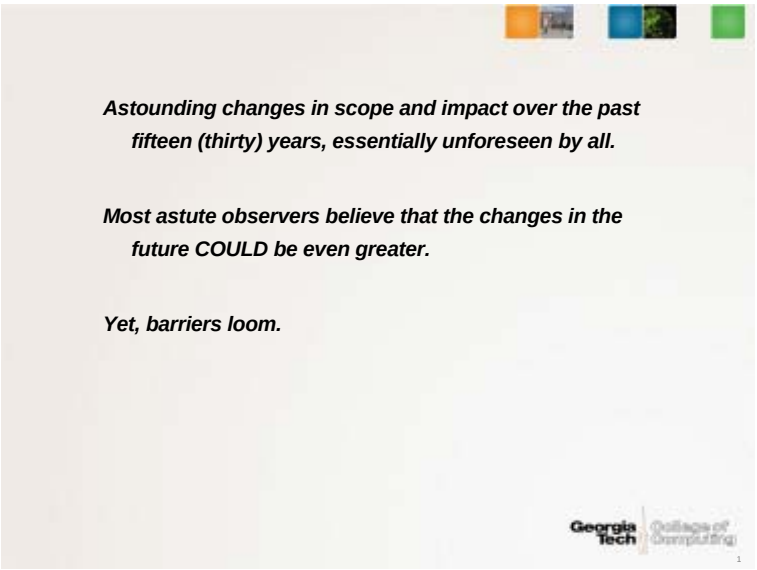

Finally, we can see that, based only on the kinds of services and applications that are already operating in prototype, the growth in traffic and use of the Internet is going to continue to expand far beyond today's usage – and, that the problems we are already seeing (e.g. security) will only multiply.

This has led many of the best networking researchers and experts, including most of those responsible for designing and building the current Internet, to come to the conclusion that we need a deeper understanding of the Internet and its parts, and a fresh start. A simple example is the cover story of an important engineering journal, featuring a picture of Larry Roberts, truly one of the creators of the current Internet, and the line "The Internet is broken – let's fix it."

As with any complex, engineered artifact - like an airplane - we don't just start designing and building a new version by the seat of our pants. Before designing a new model of plane, for example, much preliminary work on new subsystems, materials, and even basic structures is carried out. Once there is a design, full-scale manufacture and usage doesn't begin before many hours of experimental work – first on the pieces and then by building and flying prototype versions whose characteristics can be thoroughly studied.

In the U.S., the conclusion that significant, new research and experimentation was arrived at after several years of discussion. Similar conclusions have been reached independently in Japan, Western Europe in general, and – obviously – Germany.

Let me turn briefly to the issue of international cooperation, something that from the start we understood the value of and planned for in the current networking research activities in the U.S. I would note in passing, that some of the very earliest nodes on the original ARPAnet were in England and Norway.

The Internet is dramatically international today and thus any attempts to understand and/or develop new networking mechanisms clearly must consider international aspects of networking and networking research. Further, as opposed to the early days of the existing Internet (and, indeed, most of networking development in general), there are now strong cohorts of active network researchers in a number of countries.

Scientific research has always valued international collaboration because it is well understood that only in that way can we bring the best minds in the world to bear on a problem. Increasingly today we also understand that different cultural and social outlooks may enable contributors with different outlooks to make important contributions that would be overlooked by those with different perspectives.

Some argue, however, that the Internet, because it is a practical technology with extremely large strategic and economic value, is different and that thus a country must jealously protect any advances its researchers and developers can make. Clearly, there is truth in this as we get closer to the design and implementation of specific products such as routers or servers. On the other hand, understanding how traffic patterns on a network may oscillate under certain conditions, while having value in some aspects of design, is more of a scientific fact than it is a technological breakthrough. When we think about creating structures and technologies that will work for an improved Internet, they must be compatible worldwide and thus cooperation on insuring this needs to begin long before technologies become products.

The last reason for valuing international cooperation on networking research is more abstract but some would say is the most important – creating an open "creative commons" that will benefit all. In fact, no country controls the current Internet and it runs largely on technology that is not patented or copyrighted or trademarked by any individual or company or organization. The U.S., because it initially developed much of the technology, has dominated the Internet in the past as a practical matter not because of secrecy but because of moving quickly to exploit the open and free use of the basic ideas for products and services built on the Internet. This has been one of the strongest possible arguments for openness and sharing.

As illustration, let me simply list a few of the international cooperative efforts of which I'm aware: In Japan, the New Generation Networking Forum has an aggressive program of communication with efforts in Asia, Europe, and North America; in the U.S. the National Science Foundation has sponsored workshops with Europe (one is being held in Seattle this week) and Japan, and several of the GENI prototyping projects (PlanetLab, ORBIT, ORCA) have active international collaborations; and, in Europe invitations have been extended to a number of people from the U.S. to attend and/or speak at important meetings such as EuroView and the EC's Future Internet Assembly. There are also many valuable, international researcher-to-researcher connections.

While this is only a sample, almost all are only informational exchange at this point. Most important moving forward will be to follow these contacts up with substantial collaborations that directly connect some of the technical work – for example, a project that might involve the linkage of an experimental wireless network in one country with an experimental high-capacity real-time network in another country.

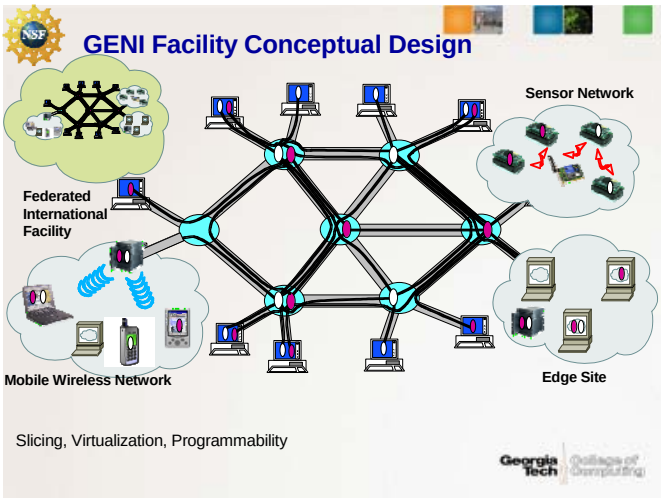

While we envisioned this in the initial conceptualization of GENI, it has not been made reality yet. The G-Lab experimental platform and associated research is one of the first realizations anywhere of the type of experimental research that is so badly needed. You have a great opportunity help lead in establishing international, cooperative networking research. I strongly encourage you to take this opportunity to establish strong, collaborative ties between G-Lab and other projects such as GENI.

If there are those that doubt the need for better networks and the research that is essential to creating them, they need only look at the attention that the U.S., Japan, and much of Europe are paying to creating more capable future networks. The best way to convince doubters, however, is to do innovative and excellent research with the resources you have and demonstrate the value of such research. But, it takes two to dance, so I hope that you and your sponsors will respond favorably to the invitations from abroad for meaningful collaboration. I wish you the best of luck in all your endeavors!

Thank you!