UltraNet – *enabling scientific insight*

Background

The Department of Energy is about to embark on the construction of a computational capability that will return leadership in scientific computing to the United States. The problems to be addressed in climate modeling, bioinformatics, astrophysics, and other disciplines, will all produce data sets that regularly approach or exceed a petabyte (peta = 1,000 trillion) in size. In a very real sense, the large, high-value, data sets produced by computational simulation are analogous to the data produced by large-scale instruments. Experience shows that these data sets are saved and poured over in the search for insight, by teams of scientists around the world. Therefore, part of the Ultrascale Capability (USSCC) includes planning for multiple tightly integrated, grid-based, data storage and visualization centers where subject matter experts can (virtually) meet to analyze, visualize, explore, and discuss these results. An ultrascale network, "UltraNet" is needed to support this capability.

Why UltraNet is Required

Manipulating a petabyte data set is not a project to be undertaken lightly. Petabyte storage systems are at the leading edge of what can be supported on-line today. Moving a petabyte data file through today's fastest TCP/IP networks would require weeks, if it could be done at all. (Today's network protocols are statistically unable to perform an error-free file transfer this large.) In addition, no commercial or research network today has the backbone bandwidth to support this volume of traffic.

Nonetheless, it is well recognized, that scientific insight results from the ability to visualize, render, and display computational results – frequently from playing back results like a stop-action sports replay - fast-forward, rewind, jog-forward, stop, there! These capabilities will only be made possible by integrating Ultrascale computational, storage, and data analysis systems into a coherent Grid. Thus, a critical part of delivering Ultrascale capabilities will be the ability to move results from computational, to storage and visualization resources residing at different laboratories in a timely, accurate, and efficient manner.

What UltraNet will do for the Department of Energy

Simply put, UltraNet will be the fastest research network in the world. Figure 2, below, illustrates the comparative bandwidths of agency and research networks over the last several years. Note that today's most advanced networks run at 40 Gb/s (40,000 million bits/second). As can be seen, UltraNet represents a major step forward. Quantitatively, moving a petabyte file between, for example, a storage resource and a visualization center in less than 24 hours will require a network throughput of 160 Gbs or better. This will only be possible with fundamental improvements in the underlying network protocols. As a result of the fact that Internet protocols have evolved in the context of very large, shared networks (millions of users), the core data transport protocol, TCP, is optimized for "fairness" in a widely shared environment. While good for general purpose Internet services, this optimization actually presents serious limitations to mission-critical capabilities even when delivered via dedicated networks because TCP prevents any individual application from using more than half the available bandwidth.

Adding parallel streams helps this situation, but doesn't solve it, and does not scale to a 160 or 200 Gbs network. Further, today's protocols are statistically unable to transfer a petabyte file in an error-free form (because the existing error detection algorithms did not envision petabyte transfers, just as some software in the last century did not envision the need for more than 2 bytes to represent the year portion of a date). Protocol extensions or modifications necessary to support high-performance and high-accuracy transfers are straightforward, but this work is not being tested and validated elsewhere.

Finally, both today's commercial and agency research networks are fundamentally based on routing technology as developed for the general purpose Internet. However, it is becoming increasingly evident that an optically switched network might offer more flexibility and improvements in cost-effectiveness. Although a number of vendors have optical switching projects in development, many have been seriously delayed as part of the current economic downturn. UltraNet represents the perfect opportunity for DOE to partner with one or more optical equipment vendors (e.g. Lucent, or Nortel) to serve as a catalyst to pull these projects into the mainstream.

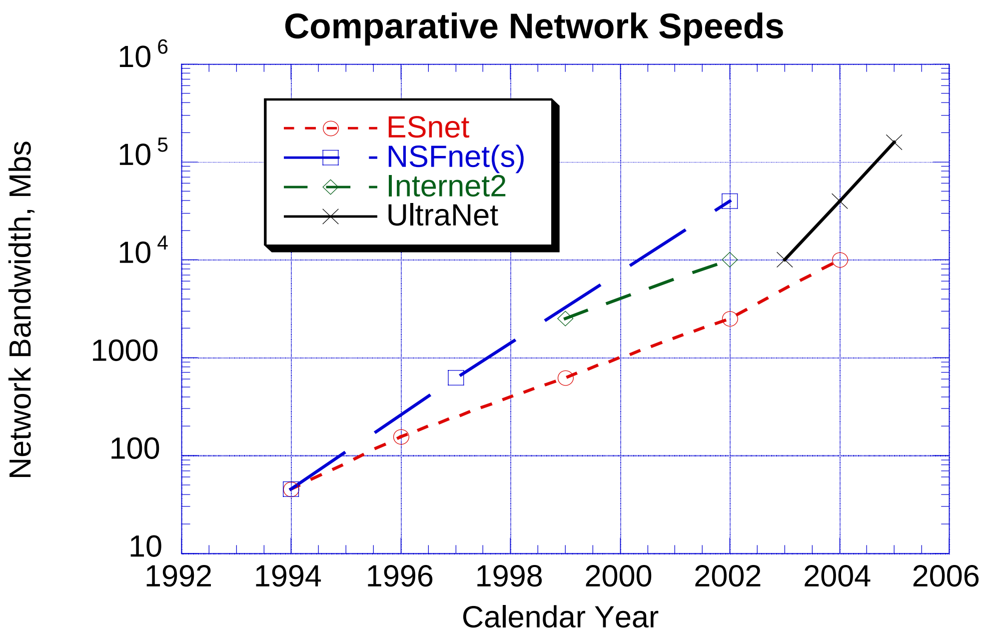

Figure 1: A comparison of research and education network speeds over time.

What UltraNet will do for the country

DOE is widely acknowledged to have provided the impetus that moved ATM networking technology from the laboratory and into the commercial mainstream in the early 1990s, effectively preparing the foundation upon which the Internet could grow exponentially. The timing is exactly right for this to happen again, and DOE could become the catalyst for rolling out the next generation of high performance commercial networks – precisely the type of network that will be required to support such new commercial capabilities as 3rd generation cell phones and Ethernet-in-the-last-mile. These are the capabilities being looked to as drivers of economic recovery (in the high-tech economy). Furthermore they are the technologies that will require far more backbone bandwidth than current networks can provide. In addition, the high-accuracy protocols that will be required for UltraNet, will provide armor for both commercial, and homeland defense networks.