Vision for Interagency Network R&D Needs Over the Next Five Years

Background and Context

Over last five years, it has become clear that researchers and educators are demanding two research tools from research internetworks: virtual lans (VLANS) and optical switching and the tools to make connections easy to manage, end-to-end. However, the value of these internetworks is only as good as their interconnections – both within the US and internationally.

In order to support the research platforms of the US government, with its various open science and international research use of high performance networks, it is apparent that serious basic and applied research needs to be directed to the network of the future. Terabits per second networks will be essential components of distributed Petascale science and advanced analysis. These terabit networks will likely be based on advanced optical transport infrastructure, ultra high-speed protocols and dynamically reconfigurable services. Drafting a concrete and robust roadmap to develop, test and implement these advanced capabilities will be critical to agency missions in the next decade.

Key Federal requirements for Petascale analysis and science will emerge years before they are required by the commercial communities. Such requirements can be met only through agency network R&D directly targeted at addressing key needs for internetwork infrastructure and services. Agencies cannot compromise mission-critical programs by assuming commercial communities will address their needs in a timely manner. It is clear that no other R&D community is positioned to address these needs within the required time frame and only an organization like NLR can readily step up to the developing needs for infrastructure that can and will support the required research efforts.

Networking R&D Requirements for Agency Petascale Science and Analysis

The federal government has a broad spectrum of network requirements ranging from routine production IP services to very specialized high-throughput capabilities that are required to support emerging distributed Petascale and Exascale science and analysis applications. With DOD, DOE, NASA, and NOAA all moving towards Petascale and Exascale scientific discovery, it becomes imperative to interconnect these distributed facilities for the advancement of science. In terms of end-to-end networking, Petascale and Exascale applications will require both *capacities* and *capabilities* unprecedented in currently envisaged network infrastructures and associated support technologies*:*

Network Capacity utilizing Optics: To support Petabyte and Exabyte-scale data distribution and other applications, Terabit-capable networks will be needed. This implies protocols and services that can efficiently operate at ultra high speeds. The agencies' supercomputing, storage, visualization and experimental facilities will be required to have sufficient capacity to handle Petabyte or larger data sets. The one hundred Gbps/lambda circuit technologies, currently emerging from development laboratories, barely meet this requirement. It still takes 24 hours to

transfer a Petabyte of data at such rates. Nonetheless, this serves as a starting point in planning for Petascale- and Exascale analysis, since current experience shows that conducting large-scale data transfers using soft aggregation of circuits or data streams is difficult and support-intensive. While planning 100-1000 Gbps capacity requirements for agency cores, it is critical that the needed network capacity be provisioned end-to-end: including metro, campus, edge and host. Technologies capable of providing connections with such capacities are necessarily disruptive and therefore require theoretical and experimental research. Hence, capacity solutions based on such technologies must be fostered and developed through highly focused efforts on experimental networks and systems, since the advanced technologies may prove to be too disruptive for production network infrastructures.

End-to-End Capabilities*:* In addition to the transport network's path capacities, an extremely important and vital part of the solution consists of optimized systems of software and edge/host technologies that will enable users to achieve throughputs commensurate with the provisioned capacities. As we already see today, research networks will include multiple vendors' hardware and software, multiple protocols and encompass multiple service-providers.

The challenges of developing the needed capacities and capabilities for Petascale and Exascale applications are multi-fold, spanning the wide-area connections, edge and hosts systems, systems and application software and middleware tools. The development is very complex and often requires non-traditional solutions to achieve the needed quantum leaps in the capabilities. Examples might include special interconnects to supercomputers and direct wide-area InfiniBand interfaces to storage systems. Such developments would require specific combinations of specialized technologies in this area and would be extremely unlikely to become available as incidental byproducts of other projects.

Research Recommendations

We see the need for R&D in five core areas. We believe a strong, coordinated interagency approach will be needed to deliver robust working solutions in all five areas in a five-year time frame

- Transport protocols, technologies and data distribution services to enable networks to move the large datasets required
- End-to-end federated network measurement based on Petascale-derived performance characteristics
- Multi-layer federated network provisioning to enable regional, national and international network use on an end-to-end basis
- High-performance end-system middleware that augments a multiprotocol, multiprovider infrastructure
- Experimental network testbeds to support the breakable and configurable research of network scientists