NGI/CE/03

NOAA Ipv6 Based Grid

FY 2005 Proposal to the NOAA HPCC Program

September 20, 2004

| Title Page | Proposed Project | Budget Page |

Principal Investigator: **John Hernandez**

Line Organization: OAR
Routing Code: R/OM62
Address:

NOAA/OAR
325 Broadway St
MS R/OM62
Boulder, CO 80305

Phone: (303) 497-6392
Fax: (303) 497-6951
E-mail Address: John.Hernandez@noaa.gov

Daniel S. Schaffer
Daniel.S.Schaffer@noaa.gov

John V. Parker
John.V.Parker@noaa.gov

Mike Knezevich
Michael.T.Knezevich@noaa.gov

Christopher Moore
Christopher.Moore@noaa.gov

Jeff Flick
Jeff.Flick@noaa.gov

Michael Doney
Michael.A.Doney@noaa.gov

Paul Hyder
Paul.Hyder@noaa.gov

Proposal Theme: **NGI**

Funding Summary:

Signatures

John Hernandez	Jerry Janssen	Don Mock
Network Engineer	Network Manager	Executive Director
NOAA-Boulder NOC	NOAA-Boulder NOC	Boulder OED

NOAA Ipv6 Based Grid

Proposal for FY 2005 HPCC Funding

Prepared by: John Hernandez

Executive Summary:

Following on FY 04 grid development, we propose to upgrade the prototype NOAA grid currently under development to utilize Internet Protocol version 6 (IPv6) as the network communication protocol. This proposal builds on previous HPCC Grid and IPv6 proposals. If funded this project will investigate utilizing IPv6 to simplify the integration of NOAA-wide resources into a large Globus-enabled Grid. The first and primary task will be to expand the production quality IPv6 into the local area networks (LANs) at the participating sites. Hardware upgrades will be required at some sites. This will be necessary to avoid impacting production traffic. Traffic will be monitored on a dedicated host that will provide a platform for testing and evaluation of IPv6 specific security tools. Globus grid tools already support IPv6. Host and network components will need to be enabled and tested. Finally, IPv6 addressing will be investigated as a solution for a problem discovered in earlier grid development: inconsistent host addressing among the diverse NOAA sites.

Problem Statement:

High-speed networks such as the National Lambda Rail are rapidly becoming available to NOAA laboratories. With this high-speed access and grid software such as the Globus toolkit and other packages that build on it, NOAA will soon be in a position to more tightly integrate its computational and data resources. In fact, this effort is already underway. The next set of high performance computers will be obtained through a joint GFDL/NCEP/FSL procurement. A key assumption underlying this effort is that the resulting computational and data resources will be shared NOAA-wide. It is envisioned that, for example, an oceanographer at GFDL could run a model on resources located in Boulder, Colorado that utilizes initial conditions stored in a centralized data repository in Asheville, NC and then render graphics of the model output locally in Princeton, New Jersey. Doing so requires networks to seamlessly and quickly move the large data sets across the country. Ultimately it will be necessary to execute parallel data transfers between 10's or 100's of nodes in various locations. If you consider other NOAA laboratories and connections with resources outside of NOAA, then it becomes clear that the number of node addresses required will be in the 1000's.

Ipv4, the current network protocol used to connect NOAA resources, lacks the means to cleanly implement these kinds of parallel data transfers. RFC3330 private address space is commonly utilized throughout NOAA. Although policy has been established to standardize the assignment and use of these Internet Protocol version 4 (IPv4) addresses, actual compliance to this standardized addressing is difficult and hence implementation is proceeding slowly. In addition to address inconsistency, the total address space available in IPv4 is quite small. Much is already allocated and available blocks are often small. Utilizing existing addressing techniques leads to complex addressing schemes in an attempt to "fit" into available blocks.

Proposed Solution:

1. Implement full production IPv6 networking to grid components at FSL, PMEL, and GFDL. Most of the higher level NOAA routing architecture is already in place and will only need to be tested in this context. However, it will be necessary to complete implementation of IPv6 LAN connectivity at the participating sites. For example, native IPv6 links will need to be created from the NOAA outer-edge routers into each of the laboratories. This will be enabled by purchase of the supervisor modules listed in the budget and will avoid the necessity for tunneling of any kind. Individual nodes will need to be configured and tested, as will the supporting grid toolkits. Since the NOAA backbone is not expected to support IPv6, connections between the laboratories will go through Abilene or other backbones that do support it.

2. Enable existing IPv6 stacks on hosts and test connectivity.

3. Integrate IPv6 support into the existing DNS infrastructure or construct an experimental DNS.

4. Configure the Globus Grid toolkit to use IPv6 as the primary communication path.

5. As part of previous NOAA HPCC funded efforts, a coupled Weather Research and Forecasting (WRF)/Regional Ocean Modeling System (ROMS) model has been configured to execute across the prototype NOAA grid. While it is not envisioned that this mode of operation will become widespread in NOAA, it does provide a convenient means to analyze performance and security characteristics of the IPv6 implementation. Model runs and data transfers will be monitored to analyze and tune the IPv6 behavior. Tuning will be structured and repeated to evaluate the effect of specific changes and combinations.

6. Firewalls and monitoring tools currently in use at FSL do not support IPv6. It will be necessary to evaluate candidate IPv6 firewalls such as the Lawrence Berkeley Laboratory's BRO product for both security and performance. A separate host at FSL will be used to support security and traffic monitoring.

7. Increasing the connectivity between the various laboratories raises a number of security issues. For example, what happens if a security hole appears because GFDL installs a software update but FSL does not? NOAA-wide policies will need to be adopted to ensure software updates are made in a safe manner. Although no deliverables are promised in regard to these issues, our intention is to hold discussions of them in the course of implementing the IPv6 prototype.

Analysis

By implementing the Grid communication in IPv6 a number of problems are immediately resolved:

A. Private address space and its centralized assignment are avoided. Each communicating site will utilize a small part of its unique IPv6 address block.

B. Extremely large amounts of address space are available, eliminating the need for complex host addressing schemes and network address translation.

C. Routing is simplified and easily restricted.

D. Additional security can be implemented in IPv6.

Monitoring and security are critical and the experience gathered will enhance expanded use of IPv6. It will also protect the test resources during the project.

This project will leverage and/or facilitate several previous and on-going efforts. One is the construction of a prototype NOAA grid; funded largely by an FY '04 NOAA HPCC grant. Second is the Modeling and Assimilation Portability Project (MAPP); an FSL-led effort to enable significant NOAA weather and climate codes to execute seamlessly on various NOAA HPC resources. Third is the FY '03 NOAA HPCC funded work that enabled the NOAA headquarters, Norman, Boulder, Seattle, and Miami NOC's to construct a large-scale IPv6 deployment. Fourth is an FY '05 proposal (A Prototype for Earth Science Data on Demand) to construct a means for applications such as portals to seamlessly access distributed data from satellite, radar and other observations or from other model simulations. If successful, our IPv6 prototype could be extended to simplify the network connections needed for the Earth Science Data on Demand effort.

Performance Measures:

The project can be successfully accomplished according to the following timetable:

Milestones

Month 02 - Verify new individual site IPv6 address blocks have been allocated and configured across site LANs.
Month 03 - Verify IPv6 support in DNS at participating sites
Month 06 - Verify IPv6 connectivity on cluster hosts at participating sites via simple data transfers and other tests.
Month 07- Update Globus toolkit configurations and verify IPv6 functionality as the primary protocol at all participating sites.
Month 08- Analysis of traffic flows across the grid during model runs.

Month 12 - Documentation complete

Deliverables

- A prototype IPv6 based grid
- Documentation how to expand the grid