# *Net100*

***http://www.net100.org***
***PIs:*** *Gwendolyn Huntoon. PSC (huntoon@psc.edu)*
*Thomas Dunigan, ORNL, ( thd@ornl.gov)*
*Brian Tierney (BLTierney@lbl.gov)*

*Executive Summary*

*September 2001*

**Vision:**

The well-known difficulties in obtaining good network performance for TCP-based applications without expert tuning or case-by-case application optimization can be overcome by building expertise into the operating system where it will benefit all users and applications.

The Net100 Collaboration (PSC, NCAR, UT, LBL, and ORNL) will to develop a model for network-aware operating systems using Web100 as the means for incorporating network information and its analysis into host operating systems to improve performance. To investigate how effective network-aware operating systems can be, we will use a three-phase approach. First, we will use the network-aware; Web100 based operating system that we develop to create a simple, bulk-transport application and demonstrate its use over high performance network links. We will then extend this model to support more advanced and complex applications, moving from point-to-point optimization to optimizations for fully distributed environments. Finally, as proof that a network-aware operating system can tune and optimize performance on behalf of applications, we will also develop application-internal tools (based on NetLogger) to monitor the efficiency of application support, and provide an external monitoring methodology (based on the Network Weather Service) to gauge the impact this system has on the rest of the network.

In addition to serving the needs of high-performance computing and network users, this project will serve as an (open source) showcase for network-aware operating systems, beginning with a Web100-based O/S. The tools, sources, measured results, and methods will be showcased on a "Closing the Wizard Gap" web site and made available to the high performance networking community.

**Major Goals and Technical Challenges:**

The core component of the Net100 project is a network aware operating system (NAOS). The NAOS collectively refers to kernel level modifications as well as API's, libraries, and daemons developed as part of the project to support the network-aware functions. In many cases actual NAOS components may depend on the base operating system and/or the underlying network technologies.

Net100 will use the above technologies and expertise as mechanisms for testing and refining Web100. however, there are some critical missing components which must be developed to allow a thorough analysis. These will also be provided by Net100, and include:

- The Network Tools Analysis Framework (NTAF), a framework for triggering NWS network measurement tools as well as other network monitoring tools, such as pathrate, pipechar, and Iperf. The NTAF will be provided by LBNL.
- The Network Analysis Information Base (NAIB), a monitoring data archive database that provides sophisticated relational queries and analysis. The NWS team at UTK will provide the NAIB.
- Instrumented applications and tools. This involves using the Web100 TCP instrumentation interface to collect TCP data from a variety of applications and monitoring tools, and use NetLogger to format the data send it to the NAIB. LBNL, ORNL, and the NWS team will do this.

The results of this monitoring and analysis will then be fed back to the Web100 team for further refinement of the Web100 kernel.

The goal of Net100 is to eliminate what has been called the "*wizard gap*". Through the integration of end-to-end and application-level monitoring capabilities with the tuning and diagnostic capabilities provided by Web100, we

will develop a unique and general-purpose system for optimizing and understanding end-to-end network and application performance. This will allow us to create a network-aware operating system that will be able to maximize network utilization for a wide variety of applications and, without help from the wizards, eliminate the wizard gap.

**Major Milestones and Activities:**

**Year 1:**

- Analyze the instrumentation requirements from the bulk-tramport application and integrate existing Web100 kernel into the host operating system and application itself, and modify and enhance Web100 instrument sets, API and libraries as necessary.
- Begin development of mechanisms for using NAOS generated information with the Network Tool Analysis Framework.
- Develop and demonstrate Web100-aware bulk data transfer application for Probe/HPSS testing between NERSC and ORNL.
- Contribute to test and evaluation of existing end-to-end tools.
- Explore transport optimizations for single TCP flows.
- create the core NTAF framework
- instrument *iperf, pipechar, pathrate* and GridFTP using Web100 TCP-KIS and NetLogger, and integrate into NWS
- begin analysis of which monitoring tools will be required for network-aware operating systems
- interface NTAF with NWS
- design and begin to prototype NWS NAIB analysis modules
- analyze Web100 autotuning of GridFTP using NTAF and NWS

**Year 2:**

- Integrate NAOS into multi-site distributed application.
- Identify key operating systems NAOS should be expanded to. Work with Net100 collaborators to get vendors to adopt the NAOS and Web100 standards.
- Integrate Net100 tools into HPSS transfer protocols (HSI).
- Net100 integration with other vendor TCP stacks (e.g.: Solaris, Irix, Compaq).
- Instrument and integrate new network tools as they become available
- analyze Web100 autotuning with Grid applications
- analyze impact of tuned applications on non-tuned applications, and impact of parallel streams
- analyze alternative protocols
- integrate with other monitoring projects
- begin integration of NAIB within NTAF

**Year 3:**

- Continue to expand and enhance NAOS instrumentation and functions based feedback from applications, primarily focusing on integrating NAOS into a range of DoE specific distributed applications.
- Begin the process of unifying performance measurements across all layers in the network. Add instrumentation and functions to the NAOS as necessary.
- Document and standardize all enhancements and provide mechanism for distributing NAOS code to DoE users.
- Continue Net100 integration with other vendor TCP stacks
- continue to instrument and integrate new network tools as they become available
- write IETF documents on network-aware operating system techniques
- continue to analyze Web100 autotuning with Grid applications
- continue to analyze impact of tuned applications on non-tuned applications
- continue to analyze alternative protocols
- continue to design and test NWS NAB analysis modules
- complete NAIB prototype, and deploy integrated NTAF and NAIB

**Current Connections with Other SciDAC Projects:**

This project will work closely with the LBNL “Self Configuring Network Monitoring” project, which plans to utilize the monitoring data archive that net100 will deliver. We also expect to collaborate with the “Bandwidth Estimation: Measurement Methodologies and Application” project lead by K. Claffy. We plan to test any tools provided by this project using the NTAF.