

# An Overview of High-Performance Networking & Computing at LANL

Research at Los Alamos National Laboratory (LANL) in high-performance networking and computing addresses the communication needs of Grand Challenge applications over a wide range of environments—wide-area network (WAN) in support of grids and local-area network (LAN) and system-area network (SAN) in support of networks of workstations and clusters.

While the high-performance computing (HPC) community generally groups clusters and grids together as commodity supercomputing infrastructures, the networking aspects of clusters and grids are fundamentally different.

In networks of workstations and clusters, the primary communication bottleneck is the host-interface bottleneck, whereas in grids, the bottlenecks are adaptation bottlenecks; in particular, flow control and congestion control. To address these problems, we offer a set of general solutions for each of the aforementioned environments.

Specifically, our research in high-performance networking and computing encompasses the following projects, listed to the right.

Primary Contact: Wu Feng  
feng@lanl.gov

## **NETWORK PERFORMANCE**

- **10-Gigabit Ethernet** (see SC 2003 technical paper)
- **DRS: Dynamic Right-Sizing**
- **CHEETAH: Circuit-switched High-speed End-to-End Transport Architecture**
- **Transport Protocol Compatibility and Interoperability**

## **NETWORK MONITORING, MEASUREMENT, AND TRAFFIC CHARACTERIZATION**

- **MAGNeT: Monitor for Application-Generated Network Traffic**
- **TICKET: Traffic Information-Collecting Kernel with Exact Timing** (a.k.a. “tcpdump on steroids”)

*In turn, the above projects have generally supported our research in high-performance computing, including:*

## **SOFTWARE OSCILLOSCOPE FOR CLUSTERS AND GRIDS**

- **MAGNET: Monitoring Apparatus for General kernel-Event Tracing + MUSE: MAGNET User-Space Environment**, integrated with Autopilot (UIUC/NCSA) and Globus (ANL) as well as NetLogger (LBNL) and TAU (U. Oregon)

## **EFFICIENT SUPERCOMPUTING**

- **Green Destiny: A 240-Node Supercomputer in a Telephone Booth**

## **BIOINFORMATICS**

- **mpiBLAST: An Open-Source Parallelization of BLAST Delivering Super-Linear Speed-Up**



*For a complete list of publications, please visit:*  
<http://www.lanl.gov/radiant/pubs.html>

*This research is supported primarily by the DOE Office of Science through LANL contract W-7405-ENG-36.*