Network Modeling and Simulation (NMS) Project

Project Title: Macroscopic Internet Data Measurement and Analysis

Organization: University of California - San Diego

AO Number: L018/00

Grant Number: **N66001-01-1-8909**

Start Date: June 6, 2001

End Date: June 6, 2004

Principal Investigator:

Dr. Kimberly Claffy
9500 Gilman Dr.
CAIDA at San Diego Computer Center UCSD MS#0505
La Jolla, CA 92093-0505
Phone: (858) 534-8333
Fax: (858) 534-5113
Email: kc@caida.org

Funding Received to Date: **$1,032,514**

Unexpended Funds on Hand: **$ 101,485**

Project URL: **http://www.caida.org/projects/nms/**

Overall Objective: The UCSD/CAIDA project focuses on advancing the capacity to monitor, depict, and predict traffic behavior on current and advanced networks, through developing and deploying tools to better engineer and operate networks and to identify traffic anomalies in real time. CAIDA will concentrate efforts in the development of tools to automate the discovery and visualization of Internet topology and peering relationships, monitor and analyze Internet traffic behavior on high speed links, detect and control resource use (security), and provide for storage and analysis of data collected in aforementioned efforts.

NMS Task 1, Monitoring

Objective:

Develop and evolve measurement and monitoring tools. Expand the scope and breadth of existing monitoring measurement initiatives. Create techniques for correlation of these data sets.

NMS Task 2, Archiving and Serving Data Sets to the Community

Objective:

Determine how details about underlying networks, such as data on the topology and performance associated with specific workloads or events, can provide valuable insights to the development of data flow information for simulators. Develop data sets from real network measurements for use by the network modeling and simulation community.

Approach

Task 1: Topology Measurement Using Active Probes

skitter (http://www.caida.org/tools/measurement/skitter/) is a CAIDA tool that measures both the forward path and round trip time (RTT) to a set of destination hosts by sending probe packets through the network. It does not require any configuration or cooperation from the remote sites on its target list. In order to

reveal global IP topology, CAIDA's Macroscopic Topology Measurement and Mapping project builds software and infrastructure to: 1) collect path and round trip time (RTT) data; 2) acquire infrastructure-wide global connectivity information; 3) analyze the visibility and frequency of IP routing changes; and 4) visualize network-wide IP connectivity. An essential design goal of *skitter* is to execute its pervasive measurement while placing minimal load on the infrastructure and upon final destination hosts. To achieve this goal, skitter packets are small (52 bytes in length), and we restrict the frequency of probing to 1 packet every 2 minutes per destination and 300 packets per second to all destinations. To improve the accuracy of its round trip time calculations, CAIDA added a kernel module to the FreeBSD operating system platform used by its *skitter* monitors. Kernel timestamping does not solve the synchronization issue required for one-way measurements, but reduces variance caused by multitasking processing when taking round trip measurements. This feature helps to capture performance variations across the infrastructure more effectively. By comparing data from various sources, we can identify points of congestion and performance degradation or areas for potential improvements in the infrastructure.

Task 1: Workload / Performance Measurement Using Passive Monitors

CAIDA will collect workload data for real Internet traffic using tools such as OCxmon/ *CoralReef* (see http://www.caida.org/tools/measurement/coralreef/) and *NeTraMet* (see http://www.caida.org/tools/measurement/netramet/). CAIDA continues to enhance both *CoralReef* and *NeTraMet* under this project in order to maintain their usefulness for measurements involving changes in network transport and applications protocols as well as detection of denial-of-service attacks. In particular, we will evolve our new "backscatter analysis" technique for estimating denial-of-service attack activity in the Internet.

Task 1: Routing Monitoring and Analysis

Several BGP4 route mirror sites are available for analysis of routing data. The foremost among these is the industry-supported University of Oregon's RouteViews project (see http://www.routeviews.org). CAIDA plans to correlate these routing data with other data sets, as well as to expand monitoring of routers co-located with *skitter* and OCxmon/*CoralReef* monitors (for correlation with active and passive data sets). File formats for various data sets will be published on the web.

Task 2: Archiving and Serving Data to the Community

CAIDA believes that details about underlying networks, such as data on the topology and performance associated with specific workloads or events, can provide valuable insights to the development of data flow information for simulators. NMS project collaborators will be consulted to determine the best methodology for monitoring NMS efforts.

Recent Accomplishments:

Task 1: Monitoring

Current monitors support research projects concerning macroscopic topology, workload/performance characterization, routing dynamics, and virus or worm propagation. Twenty-one *skitter* active measurement monitors are currently deployed. Each monitor probes one of several different destination lists designed to target specific macroscopic topology studies. Daily summaries of *skitter* measurements are published at: http://www.caida.org/cgi-bin/skitter_summary/main.pl . CAIDA currently publishes results from three passive measurement monitors for workload/performance characterization of real Internet traffic: 1) CoralReef analysis of SDNAP exchange point traffic: http://www.caida.org/dynamic/analysis/workload/sdnap/ ; 2) CoralReef analysis of MFN OC48 traces http://www.caida.org/analysis/workload/oc48/ ; and 3) NeTraMet measurement of DNS root and gtld name server performance: http://www.caida.org/cgi-bin/dns_perf/main.pl . In addition, routing analysis based on a combination of *skitter* macroscopic topology data and U-Oregon Routeviews project routing data has resulted in a visualization of the Internet's AS core as well as analysis of Internet expansion, refinement, and churn. Other important analysis results include: 1) Use of walrus hyperbolic 3D visualization tool to show traffic anomalies such as high delays or the spread of the CodeRed worm through Internet routing infrastructure; 2) A methodology for ranking the relative importance for Autonomous Systems (AS) seen on a topology sample; 3) Internet routing geopolitical analysis and visualization; 4) Analysis results that

refute commonly held assumptions about Internet domain growth and routing dynamics; and 5) Use of NeTraMet passive monitors from several sites to track DNS root and gtld nameserver performance on a daily basis.

Task 2: Archiving and Serving Data to the Community
Various research groups throughout the U.S. and abroad are using raw *skitter* data from this effort. See: http://www.caida.org/projects/nms/reports/skitter_comuse.xml . In addition, CAIDA recently added an archive of the first of a potential series of Macroscopic IP Topology Data Kits, containing a combination of data and processing code of four major types: 1) skitter-related; 2) iffinder-related; 3) U Oregon's RouteViews BGP tables; and 4) DNS domain names.

Current Plan:

Additional monitors will be deployed with the remaining four DNS root name servers during 2002. CAIDA also hopes to place an OC48 passive monitor at the PAIX exchange point. CAIDA will support and assist NMS participants by analyzing traffic traces collected using *CoralReef* and/or *NeTraMet* during scheduled NMS Demos. CAIDA will continue to pursue analysis projects based on interactions with NMS PIs and the network research and operations community, and will post research questions and results in our quarterly reports.

Technology Transition:

CAIDA efforts to study Internet virus, worm, and denial-of-service attacks are likely to result in techniques of great interest to homeland security related products and services. CAIDA's *skitter* tool and routing analysis techniques are of potential interest to router manufacturers as well as the network operations community. CAIDA is working with both academic and industrial R&D groups to provide access to tools and data. Products of these efforts will directly benefit Defense Department agencies through their use in developing and evaluating new Internet protocols (e.g., protocols that are more resilient, secure, and scalable) and enhanced network planning, management, and control capabilities.