NGI/CE/04

Network Flow Profiling

FY 2004 Proposal to the NOAA HPCC Program

August 11, 2003

| Title Page | Proposed Project | Budget Page |

Principal Investigator: **Alex Hsia**

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

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

Phone: (303) 497-6351
Fax: (303) 497-6951
E-mail Address: Alex.Hsia@noaa.gov

Mike Knezevich
Michael.T.Knezevich@noaa.gov

John Kyler
John.C.Kyler@noaa.gov

Robert Kohler
Robert.E.Kohler@noaa.gov

Bruce Marshak
Bruce.Marshak@noaa.gov

Gary Skaggs
Gary.Skaggs@noaa.gov

Proposal Theme: **NGI**

Signature 1 (required)Signature 2 (required) Signature 3 (optional)

Alex Hsia	Jerry Janssen	Don Mock
Network Engineer	Network Manager	Executive Director
NOAA-Boulder NOC	NOAA-Boulder NOC	Boulder OED

Network Flow Profiling

Proposal for FY 2004 HPCC Funding

Prepared by: Alex Hsia

Executive Summary:

One of the common questions asked of a network administrator is "Howmuch of our network is being used?" Software packages that monitor network utilization by polling interface counters are common and in use at most of the major NOAA campuses. The next question that is often asked is "What are our networks being used for?" This question can be answered through the use of flow profiling which provides information on the who and what is using our networks.

Problem Statement:

As network bandwidth requirements expand at unprecedented rates, it is increasingly important to have good information on network usage, patterns and characteristics.

When attempting to meet the challenges of managing a heavily utilized IP network, near real-time traffic analysis and visualization quickly becomes an essential technology. One way to provide these capabilities is by utilizing Internet traffic flow profiling based on technology available in most networking equipment.

Network administrators who collect measurement data often find that they either have collected too little data or too much of it. In a sense, flow profiling is a "sweet spot" between those extremes. Flows strike a balance between detail and summary. They are neither captured packets, nor are they merely aggregate totals tallied as packets travel across a given port or interface. Flows are an expressive abbreviation in which each flow represents a series of packets traveling between "interesting" end points. While flow features within the network infrastructure are a convenience, the presence of this feature alone is not sufficient for reliable continuous use in production networks. We need software tools to extract, record, and help us understand the flows.

This proposal aims to provide a common view for the detailed flow analysis of IP traffic that crosses the major NOAA campus borders. Flow profiling provides information on network conversations termed flows. The type of information provided includes such details as source and destination IP address, source and destination port, destination and path Autonomous Systems (AS), and IP protocol. Thus flow profiling can be used to determine the who and what of network traffic. This information is invaluable for the capacity planning purposes.

Proposed Solution:

The flow analysis will be obtained through the use of a Monitoring Physical Interface Card (PIC) on the Juniper routers, obtained from the 2003 IPv6 HPCC project, and public domain software to analyze and report on the information exported from the Monitoring PIC.

The Juniper routers are capable of performing high rate flow profiling through the use of their Monitoring PIC. Relevant interfaces are configured to be monitored and the data is exported to an external workstation for processing.

A Unix workstation with dual processors will provide the analysis and reporting functions of the Flow Profiling system. Public domain software, FlowScan, will be used for the processing of the flow data which will be stored in Round Robin Database (RRD) files for efficient storage of archival information. Data presentation will be via the web through the use of CGI programs that will dynamically build the requested graphs on demand.

Installation and configuration information will be documented as appropriate so the installation can be replicated at each of the major NOAA campuses with tweaks to accomodate the specific campus topology.

Analysis:

The Monitoring PIC family offers a scalable and high-performance solution for monitoring high-volume traffic with flow export on the M-series routers for both accounting and security applications. The Juniper Monitoring PIC family delivers the only solution in the industry capable of monitoring multiple OC-12c/STM-4 and OC-48c/STM-16 interfaces.

One of the disadvantages to the FlowScan system is that FlowScan's near real-time processing can lag behind when processing flow files containing mostly "pathological" flows, or when monitoring heavily utilized high bandwidth links. In part, for tools other than FlowScan, these backlog issues have been avoided by aggregating totals on the router itself and then exporting those totals less frequently as summary PDUs. Another alternative is to adopt SNMP polling-based gathering of flow statistics rather than collecting unsolicited, exported flow data. However, both of these methods prevent nearly all the interesting post-processing that FlowScan currently performs, and also eliminate the archivable record of network traffic that detailed flow PDUs provide, which has proven invaluable during investigations of security compromises and network abuse.

Performance Measures:

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

Milestones

Month 02 - Obtain Monitoring PICs
Month 04 - Deploy requisite software on monitoring workstation
Month 06 - Configure FlowScan software
Month 08 - Develop graphing CGI scripts
Month 10 - Analyze Flow Profiling data for accuracy

Month 12 - Document installation

Deliverables

Provide a list of the final products from this project

- Flow profiling monitoring hardware and software system
- Flow profiling analysis data available via the web

- Documented installation and configuration procedure