NGI/CE/04 (do not delete this)**Network Flow Profiling**

FY 2004 Final Report to the NOAA HPCC

May 24, 2005

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**

Funding Summary: FY 2004 $ 118,000

Network Flow Profiling

HPCC Final Report for FY 2004

Prepared by: Alex Hsia

Performance:

The main deliverable of a flow profiling monitoring hardware and software system was made available at the NOAA-Boulder campus. The system included flow profiling analysis data which is available in both graphical and tabular format via the web <http://flowmon.boulder.noaa.gov>.

The system includes a Juniper monitoring PIC to effectively process the Internet traffic into flow summaries, and an external collector server which processes the flow data.

On the external collector server the following open source software packages were used:

Flow-Tools: to ingest the cflowd, aka netflow, data exported by Cisco and Juniper routers.

FlowScan: to process, analyze and generate graphs for the cflowd data

CUFlow module: to account for the sampling required with the Juniper routers

The recommended server hardware configuration and example software installation and configuration is documented and available via the web <http://boulder.noaa.gov/noc/hpcc/2004/ntwk-stats>.

Unfortunately full utilization of the Adaptive Services (AS) PIC for monitoring was not accomplished in Boulder. As a result, the purchase of additional AS PICs and/or Monitoring PICs for other NOAA campuses has been delayed until the accounting mode for the AS PIC in our environment can be verified.

Additionally, sFlow data capture and analysis was explored through the use of Foundry switches which are capable of producing sFlow data, and commercial Inmon software for flow analysis and display.

Project Summary:

One of the common questions asked of a network administrator is "How much 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.

The flow profiling system originally investigated was through the use of the accounting mode of the Juniper Adaptive Services (AS) PIC and FlowScan software to process, analyze and display the data. The system was initially deployed without the use of the AS PIC which required a high sampling ratio so the routing engine would not be adversely effected by the additional load incurred from monitoring and producing netflow data. Later the AS PIC was purchased which provided the capability for much lower sampling ratios and thus less likelihood of losing interesting flows.

The FlowScan system provides a quick view of Internet traffic that is broken down by subnet definitions, protocol summaries, and Autonomous System (AS) peers broken down by either packet rate, flow rate or bit rate. The data is efficiently in Round Robin Database (RRD) format which provides a compact format for archiving and displaying historical data. FlowScan also provides the capability to provide Top Ten reports by host for packet rate, flow rate and bit rate. FlowScan is very good at providing quick predefined views, with historical trends, of Internet flow profiling. However the drawback to FlowScan is that data can not be easily queried for arbitrary parameters. While you can use the command line tools provided with flow-tools to view the data in as much detail as required, it still does not provide the ease of use that is sometimes required for interactive data analysis.

The Inmon software does provide the complete visibility of flow profiling data, but at the expense of providing historical trend data for service or subnet breakdown of flows. The Inmon software coupled with the deployment of Foundry switches was initially deployed to help answer the question of how much data was flowing between the various NOAA campuses <http://sflow.boulder.noaa.gov>. It also became very useful in answering specific questions such as "who is this particular server talking to?" or "what services are being provided by this particular server?" or "how much bandwidth is this set of servers consuming?"

Since sFlow data is sampled at a much higher ratio, data from the various campuses can be sent to a single centralized collector. This centralized view is useful for analyzing overall NOAA traffic, but the netflow data is being sampled at a much lower ratio and transmitting that netflow data to a single centralized collector would be not be practical.

All of the guidance and lessons learned are documented in the online documentation <http://boulder.noaa.gov/noc/hpcc/2004/ntwk-stats/>.

Expenditure Summary:

Category	Detailed Description	Amount FY2004	Matching FY2004
Personnel	Contract Personnel		
Compensenation	NOAA-Boulder NOC	$ 10,000	$ 4,500
	NOAA-Seattle NOC	$ 3,500	$ 4,500
Capital Expenses	Monitoring Hardware		
	NOAA-Boulder NOC	$ 13,000	$ 15,000
	NOAA-Seattle NOC	$ 13,000	$ 12,000
Totals		**$ 39,500**	**$ 36,000**

Future Direction:

Future work will involve deployment at the following NOCs: Seattle, SSMC, Miami, Norman & NWS-SRH. Since the Seattle NOC already has an AS PIC, initial deployment will take place there where the installation documentation will be tested out and corrected. The deployment will then take place at the other campuses which will involve the purchase of the remaining Monitoring or AS PICs.

The InMon software also has new packet signature analysis and traffic pattern analysis modules which could help supplement the IDS that is in place at the various major NOAA campuses.