RZ 3561 (# 99571) 10/18/04
Computer Science 4 pages

Research Report

A Light-Weight and Scalable Network Profiling System

Andreas Kind, Paul Hurley and Jeroen Massar

IBM Research GmbH
Zurich Research Laboratory
8803 Rüschlikon
Switzerland
{ank,pah,jma}@zurich.ibm.com

A Light-Weight and Scalable Network Profiling System

by Andreas Kind, Paul Hurley and Jeroen Massar

Long-term network profiling in high-speed networks with high flow rates requires new ways for collecting, storing and analyzing flow-based network traffic information. Our project at the IBM Zurich Research Laboratory looked at alternatives to the conventional flow-based network profiling approach with the objective of improving scalability for high flow rates. The result is a light-weight and scalable network profiling system for NetFlow and IPFIX.

The continuing trend toward distribution of computing resources increases the need to tightly control the networks enabling remote access to resources such as servers, storage and databases. An important means for controlling networks is network profiling. A typical system collects and analyzes information about the traffic flows passing an observation point in the network, e.g. a router or traffic meter. A flow is a sequence of packets with common properties (i.e., protocol and source/destination addresses/ports).

In the past, flow-based network profiling has proven to be useful for a number of applications, including network monitoring, billing, security and planning. Profiling will be required even more in the future for smooth operation of service access in distributed computing architectures (e.g., SANs, computational Grids).

Unfortunately, most network profiling systems have a critical scalability problem regarding the storage, analysis and access of collected profiling information. In high-speed networks with average flow rates of 1,000 flows/s and peak flow rates of as much as 10,000 flows/s, a storage capability for 180 MB/h (1.8 GB/h at peak times) must be provided and maintained. Over longer time periods (i.e., months, years), the data accumulates, resulting in an over-loaded system with high capacity requirements and slow report generation times. In fact, generation of monthly and yearly reports cannot be done within a couple of seconds because the number of flow records to be considered becomes too large.

Aggregation Database

The network profiling system developed at the IBM Zurich Research Laboratory addresses the scalability problem by using a novel aggregation database (ADB) for time-series information (see Figure 1). ADB stores information in circular buffers of degrading resolution and is, in this aspect, similar to the *Round Robin Database* (RRD). But ADB is designed specifically for storing and accessing large time-series datasets and, therefore, takes in other aspects a different approach:

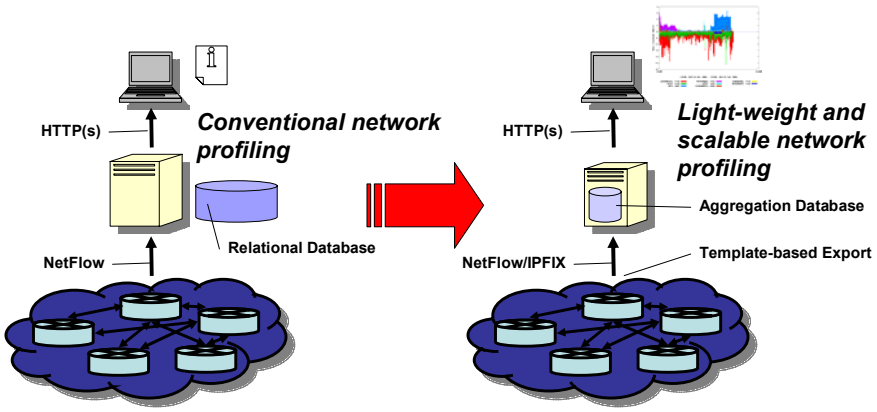

Figure 1: Using aggregation database for storing network profiling information.

- ADB groups sets of data arrays and stores these sets in single files. This representation reduces the number of file system operations and allows maintaining a sorted view onto the set of arrays regarding the sum of all array values to be maintained. This feature is important, for instance, for determining top talkers, top protocols, top flows, etc. in the network profiling context.

- ADB is designed for volume-based time-series streams, i.e., each data point in a time-series stream has a start and end stamp, and an associated volume. Typically, stamps represent time, but other representations (e.g., distance) follow naturally.

- ADB offers an array allocation optimization: When the database is updated with values whose timestamps are nondecreasing, no preceding array space is allocated because it will never be needed. Furthermore, array space is allocated in fixed chunks. These allocation optimizations reduce the storage requirements considerably for only sparsely filled time-series data streams. If ADB is used for network profiling, this optimization is very useful when, for instance, a dynamically assigned IP address is only observed during a certain time period (e.g., a week). In this case, array space in a monthly ADB set is only allocated around the actual observation period and not for the entire month.

- ADB is used with an external plotting package that is able to generate graphs in vector graphics format for inclusion in PDF reports.

- ADB accounts for concurrent array access with a locking mechanism.

- ADB has an interface to *R* for statistical analysis (e.g., trend detection and prediction).

In addition to assuming an upper bound for storage consumption — thereby reducing the resulting database maintenance as well as the hardware installation requirements — fast access to the profiling information is enabled.

Handling High Flow Rates

Bursts of high flow rates can be caused by port or host scans. In these cases, a single packet may be considered a flow since no other proceeding packet will have the same properties with regard to the source/destination addresses and ports. The rate of the flow export can in fact exceed the data rate. Unfortunately, bursts of high flow rates can not only provoke flow table overflows at the observation points but could also cause the analysis and storage to be no longer able to keep up with the incoming flow information.

For these cases, our profiling system has an automatic mechanism to aggregate all flow records of a port or host scan into a single record. The information regarding the duration and characteristics of the scans is kept and presented in a security report.

Benefits

The developed system shares many features with other flow-based network profiling systems. These include identification of network congestion causes, profiling of actual protocol and application usage as well as of actual traffic volume between network hosts, servers and domains. The system has the following additional specific characteristics:

Light-weight and scalable: Reduced storage requirements by using ADB and scan detection. Furthermore, highly configurable and modular installation.

Future-proof: Collection and analysis for NetFlow v5/v9 and the emerging IETF IPFIX standard as well as IPv6 support at data and control plane.

High level of details: The reports show traffic information in graphs and tables on domains, protocols, QoS tags, hosts/servers, individual flows, packet and flow statistics, port/host scans and other aspects. The reporting period can be chosen from hours to years. Custom zoom reports regarding specific traffic aspects can be generated on demand using a filter mask.

The described profiling system, including the aggregation database (ADB), has been developed at the IBM Zurich Research Laboratory for the past two years. The system has been installed at a number of IBM locations and is currently being tested at two European ISPs. In some installations it is combined with a newly developed network meter, which, likewise, supports NetFlow v5/v9 and IPFIX (as currently defined) in IPv4 and IPv6 environments. Snapshots of sample reports are shown in Figure 2.

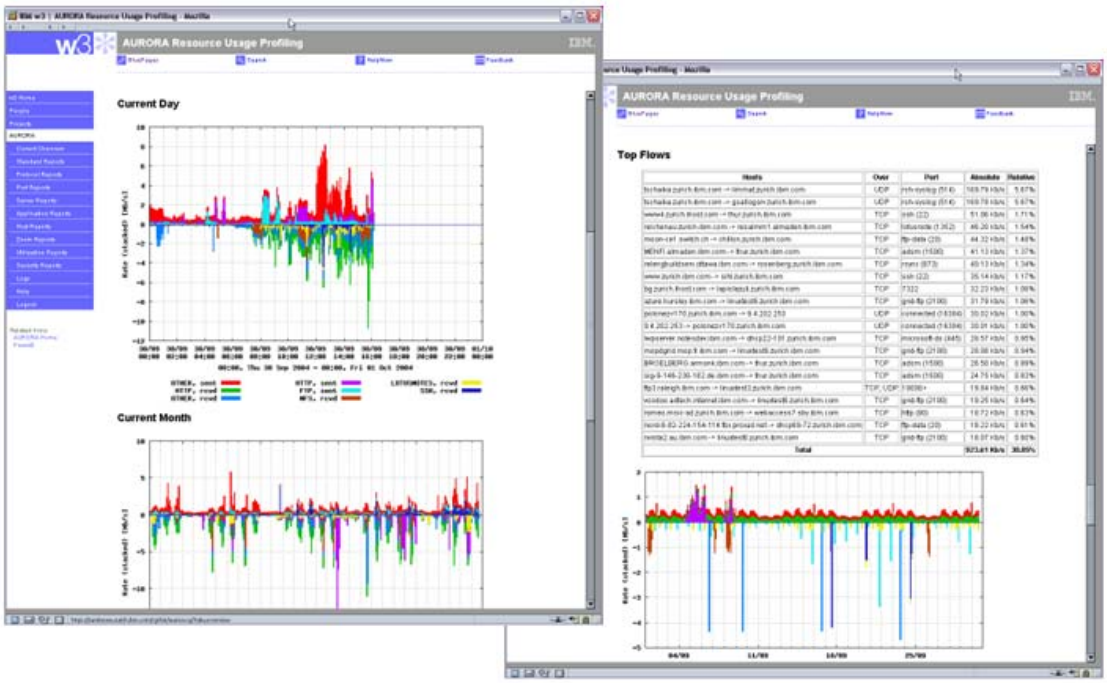

Figure 2: Sample reports of the traffic profiling system.

Links:
Project page:
http://www.zurich.ibm.com/sys/storage/resource.html

Please contact:
Andreas Kind, Paul Hurley and Jeroen Massar
IBM Zurich Research Laboratory, Switzerland
E-mail: {ank,pah,jma}@zurich.ibm.com