InTraBase: Integrated Traffic Analysis Based on a DBMS

M. Siekkinen[1], E.W. Biersack[1], V. Goebel[2], T. Plagemann[2], G. Urvoy-Keller[1]

[1]Institut Eurecom, France, {siekkine, erbi, urvoy}@eurecom.fr

[2]Department of Informatics, University of Oslo, Norway, {goebel, plageman}@ifi.uio.no

Overview

Internet traffic analysis today consists of using handcrafted scripts and large number of software tools specialized for a single task. Data and results are archived into plain files. Traffic analysis also involves huge amounts of data. Due to problems of management and scalability, we started to devise an open platform for traffic analysis that would facilitate researchers' work. We base our solution on a database management system (DBMS) that provides the infrastructure for the analysis and management of data from measurements, related metadata, and obtained results.

Goals:
- conserve the semantics of data during the analysis process;
- enable the user to easily manage his own set of analysis tools and methods and share them with colleagues;
- allow the user to quickly retrieve small pieces of information from analysis data and simultaneously develop tools for heavier and more advanced processing.

Introduction

- Problems in analyzing Internet traffic:
  1. Management
     - Data, metadata, and tools
     - Getting lost with files containing data and scripts
  2. Analysis cycle
     - Data loses semantics
  3. Scalability
     - Cannot analyze very large data sets

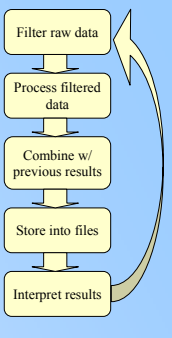

Approach		Data mgnt	Metadata mgnt	SW mgnt	Scalable	Publicly available	Integrated approach	On-line
ad-hoc scripts								
specialized tools (e.g. tcptrace)						X		(X)
toolkits (e.g. CoralReef)				X		X		(X)
ISP database projects	Sprint IPMon		X		X			
	Gigascope			X	X			X
	Internet Traffic Warehouse	X	X	X	X			
IntraBase		X	X	X	X	X	X	

Existing approaches do not fit our needs.

InTraBase Approach

Goal: Devise an open platform for traffic analysis that would facilitate researchers' work

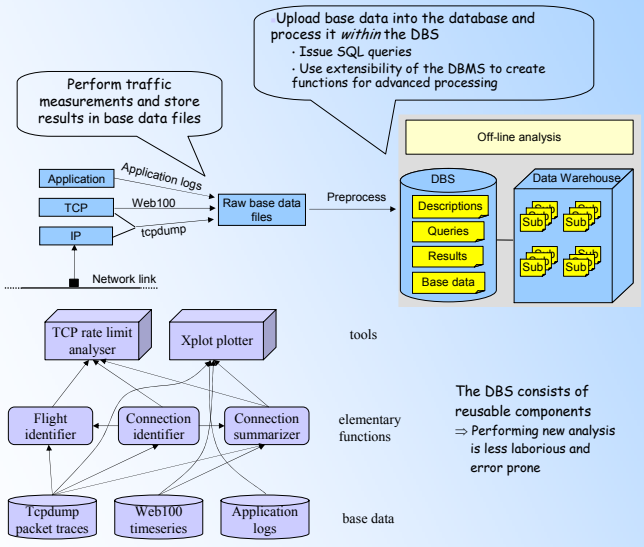

The DBS consists of reusable components
⇒ Performing new analysis is less laborious and error prone

Data conserves semantics
- Store reusable intermediate results
- It is possible to easily combine different data sources
  - E.g. application level events explain some of the phenomena in the traffic at TCP layer

Data becomes structured
- Processing and updating data is easier
- Searching is more efficient (indexes)

InTraBase Prototype

The prototype allows analysis of tcpdump packet traces with PostgreSQL.

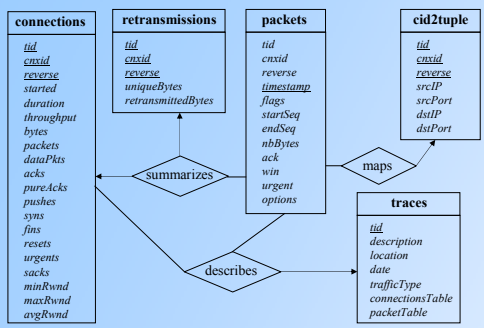

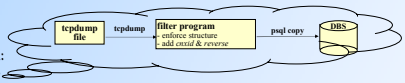

Five steps to process a tcpdump packet trace:
1. Copy packets from a file into the *packets*
2. Build an index for the *packets* based on *cnxid*
3. Create connection level statistics into *connections*
4. Insert unique 4-tuple to *cnxid* mapping data into *cid2tuple*
5. Count retransmitted data per connection into *retransmissions*

SQL queries | **Advanced queries w/ PL functions** | **Develop new PL functions**

A set of PL functions provided
- pl/pgSQL
  - Plot time-sequence diagram in xplot format
  - Produce timeseries (packet inter-arrival times, throughput…)
  - Perform analysis on a whole packet trace using timeseries
- pl/R
  - Produce graphs
  - Statistical calculations

Performance

Processing time grows approximately linearly with input data size

Also the more connections within the input data, the longer the processing time
- BitTorrent traffic vs. mixed Internet traffic

A total overhead of 50% in disk space consumption compared to flat file storage
- This is the price to pay for the structured data
- Nowadays disk space is cheap

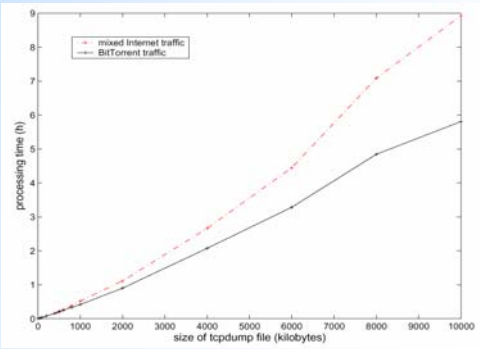

Indexing the *packet* table adds 15% of the disk space overhead

Tested on Linux 2.6.3. , 2x Intel Xeon 2.2GHz, SCSI RAID, 6GB RAM

Processing times of different steps with respect to trace file size

file size	copying (%)	indexing (%)	connections (%)	cid2tuple (%)	retransmissions (%)
10 MB	54	<1	23	8	15
100 MB	48	4	19	13	16
1 GB	44	4	23	13	16
10 GB	33	9	24	15	18

When the file size grows, the fraction of time spent for copying packets becomes smaller at the expense of the other operations, which are performed within the DBS.

The prototype is freely available at

http://metrojeu2.eurecom.fr:8080/intrabase.html