Fun with Flow

Richard Friedberg
rf [at] cert.org

Objectives

- Flow Primer
- Why do I care?
- Tools
- Capabilities and examples
- Almost live demo
- Build it!
- Where to go for more

What is flow?

- The simple version: a very brief summarization of a network connection
  - The key values
    - IP address source & destination
    - Protocol
    - Transport source & destination port
  - Other stuff that can be useful
    - Time/date
    - Flags
    - Number of packets sent / received
    - Number of bytes sent / received

Why flow?

Pros

- Cheap
- Quick (relative)
- Plenty of tools
- Minimal Privacy concerns
- Scale
- Provides a large data set

Cons

- No payload
- Provides a large data set

Why do I care?

- Network Billing
- Utilization
- Network modeling
- Behavior-based analytics
- QoS monitoring
- Data Loss / Exfil Detection
- Malware Detection
- Watchlist / Intrusion set monitoring
- Botnet tracking
- Acceptable use monitoring
- Application troubleshooting
- Forensics / Incident Response

Visibility! Know your network

Signatures are not sufficient

Why flow? A perspective on storage

Fully saturated 100mb link

1 month of storage = 50GB

Fully saturated GigE link

1 month of storage = 500GB

Fully saturated 10GigE link

1 month of storage = 5TB

A little more reasonable than full packet capture

http://tools.netsa.cert.org/releases/SiLK-Provisioning-v204.xls

Tools

YAF

Yet Another Flowmeter

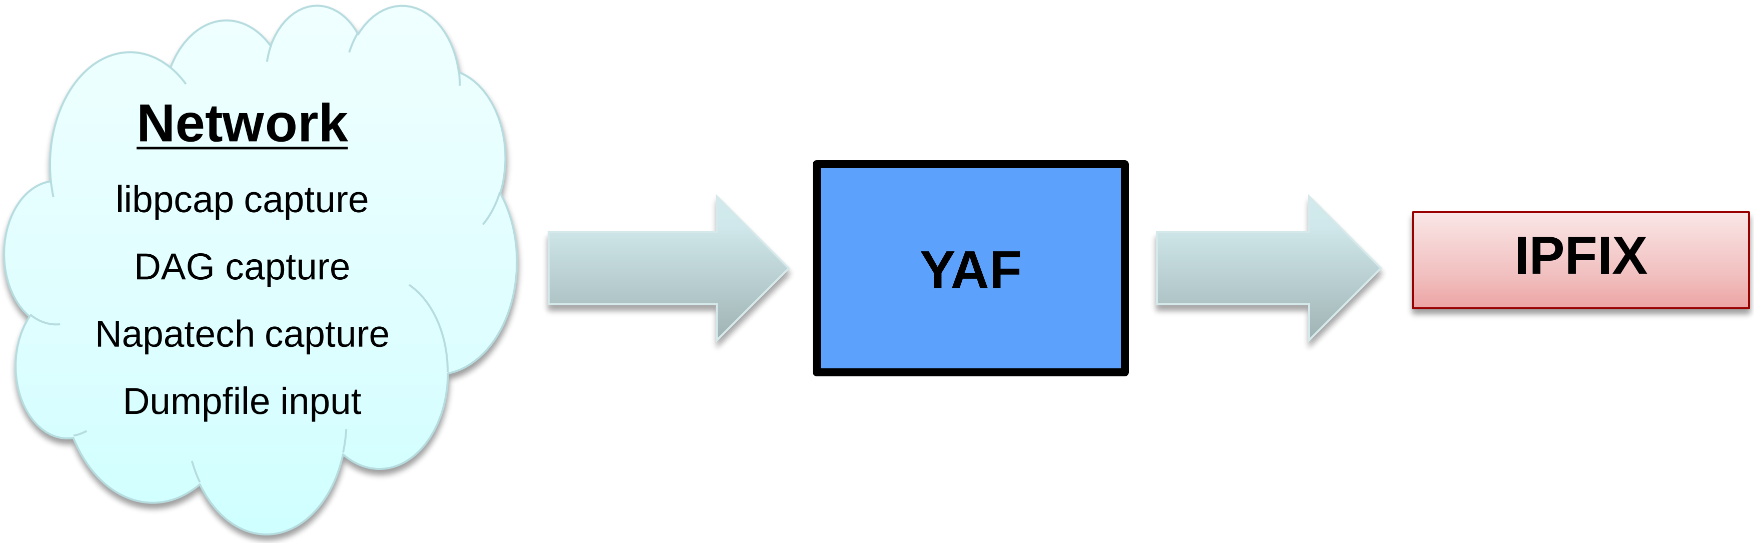

To be used with SiLK suite or other IPFIX compliant tools

YAF – Flow and then some…

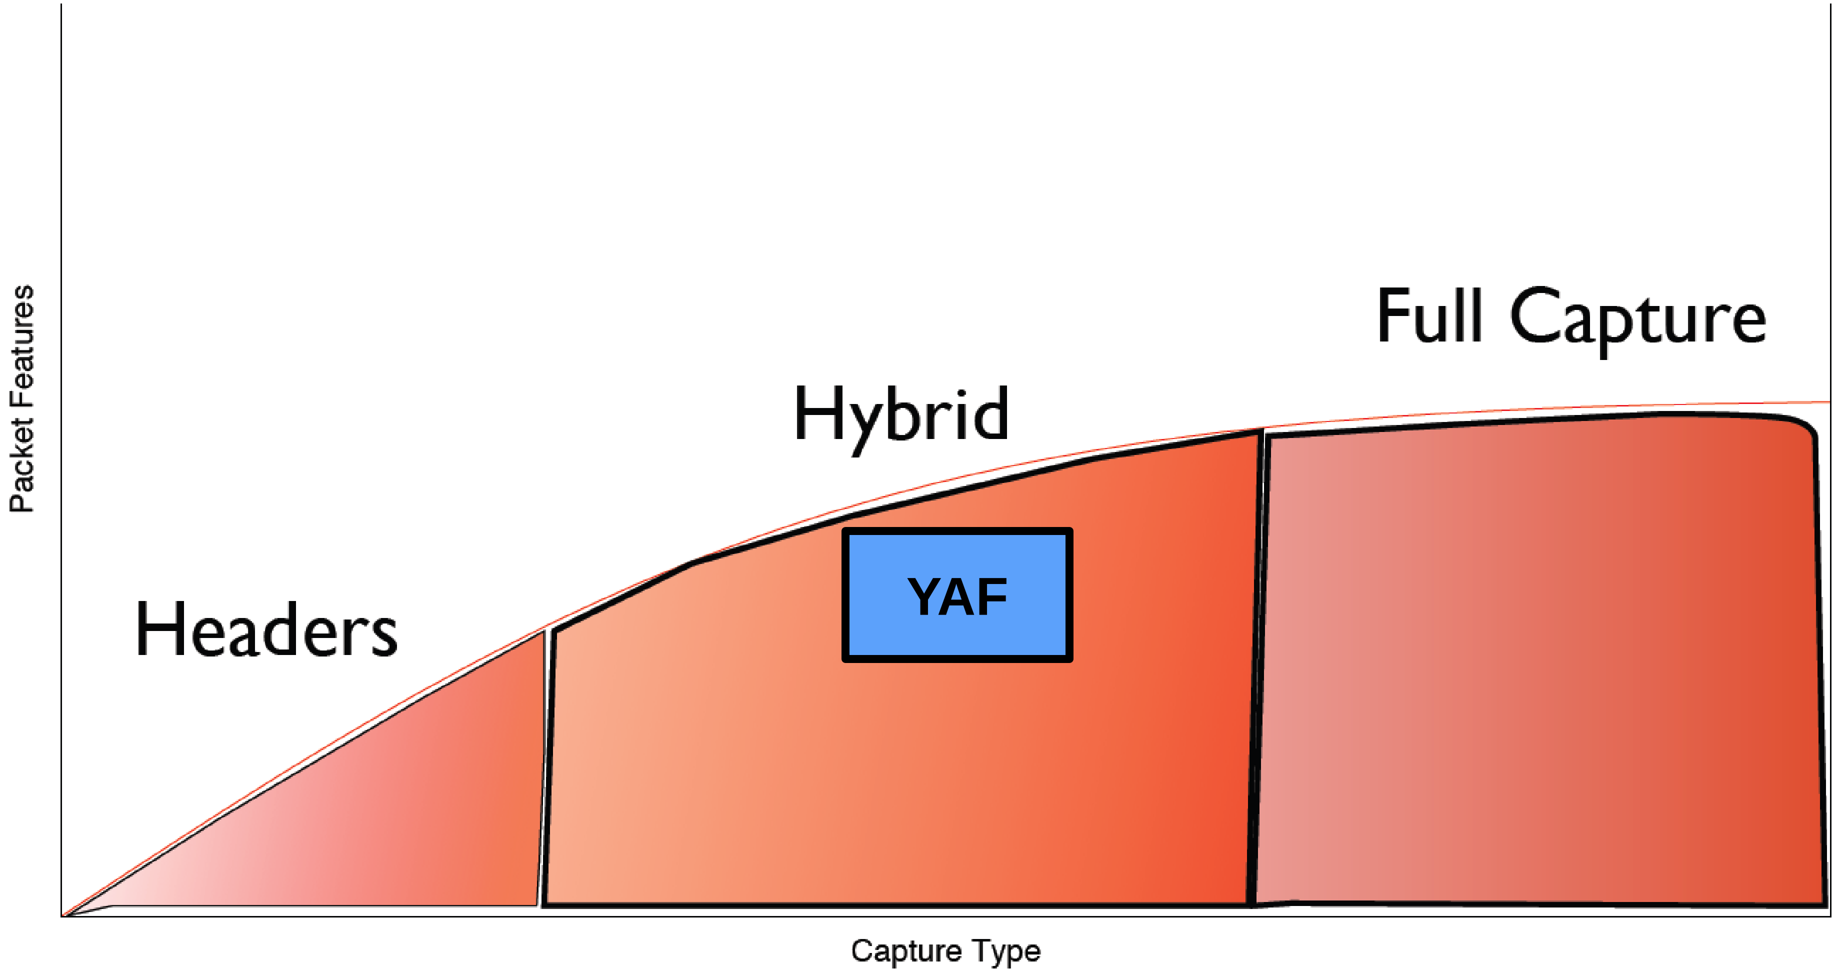

YAF

Application labeling can recognize

- HTTP
- SSH
- SMTP
- Gnutella
- Yahoo Messenger
- DNS
- FTP
- SSL/TLS
- SLP
- IMAP
- IRC
- RTSP
- SIP
- RSYNC
- PPTP
- NNTP
- TFTP,
- Teredo
- MySQL
- POP3

YAF Capture

DNS – All or just Authoritative and NXDomain responses

HTTP

- Server, User-Agent, GET, Connection
- HTTP, Referer, Location, Host
- Content-Length, Age, Content-Type
- Accept, Accept-Language,(Result Code)

FTP, IMAP, RTSP, SIP, SMTP, SSH

Soon to be added X.509 Certificates

Primarily from recognized SSL/TLS protocol negotiations

SiLK

What is SiLK?

Set of collection and analysis tools (30+)

Accepts IPFIX, Netflow v5, Netflow v9

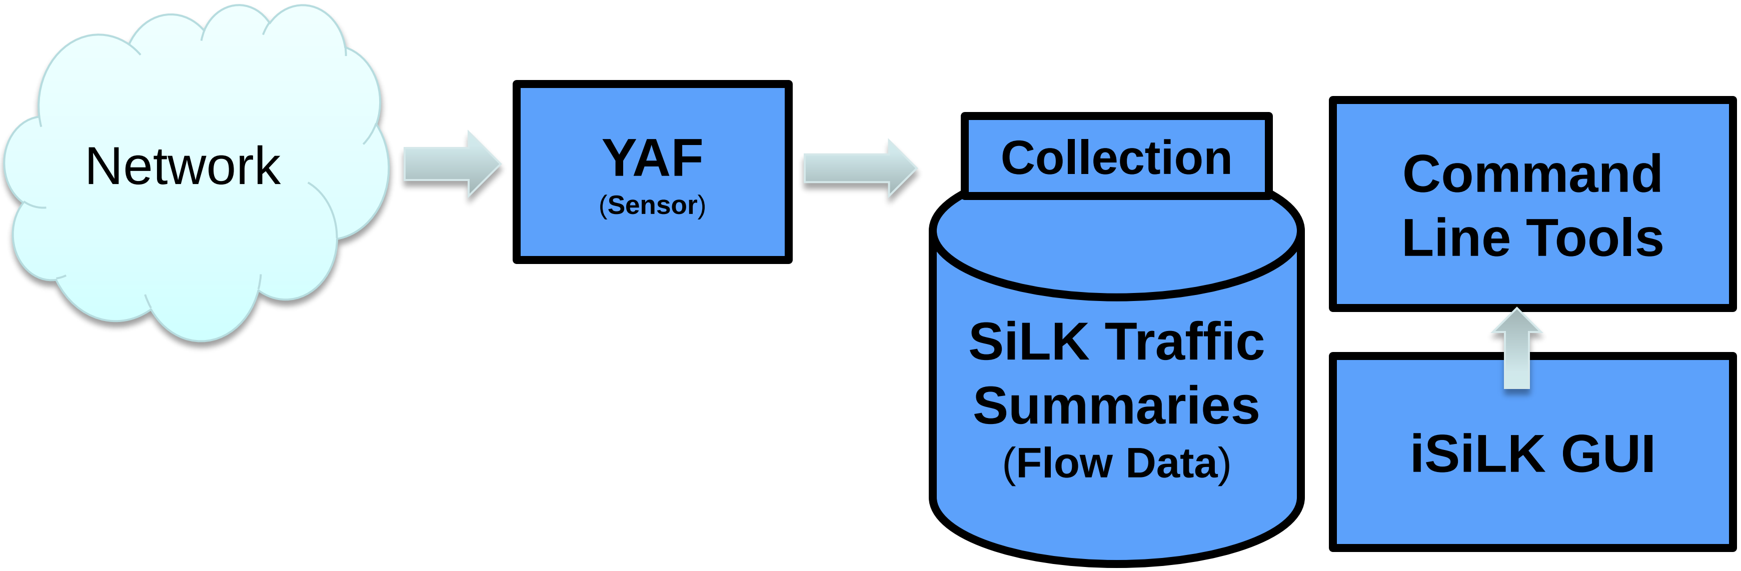

SiLK Repository

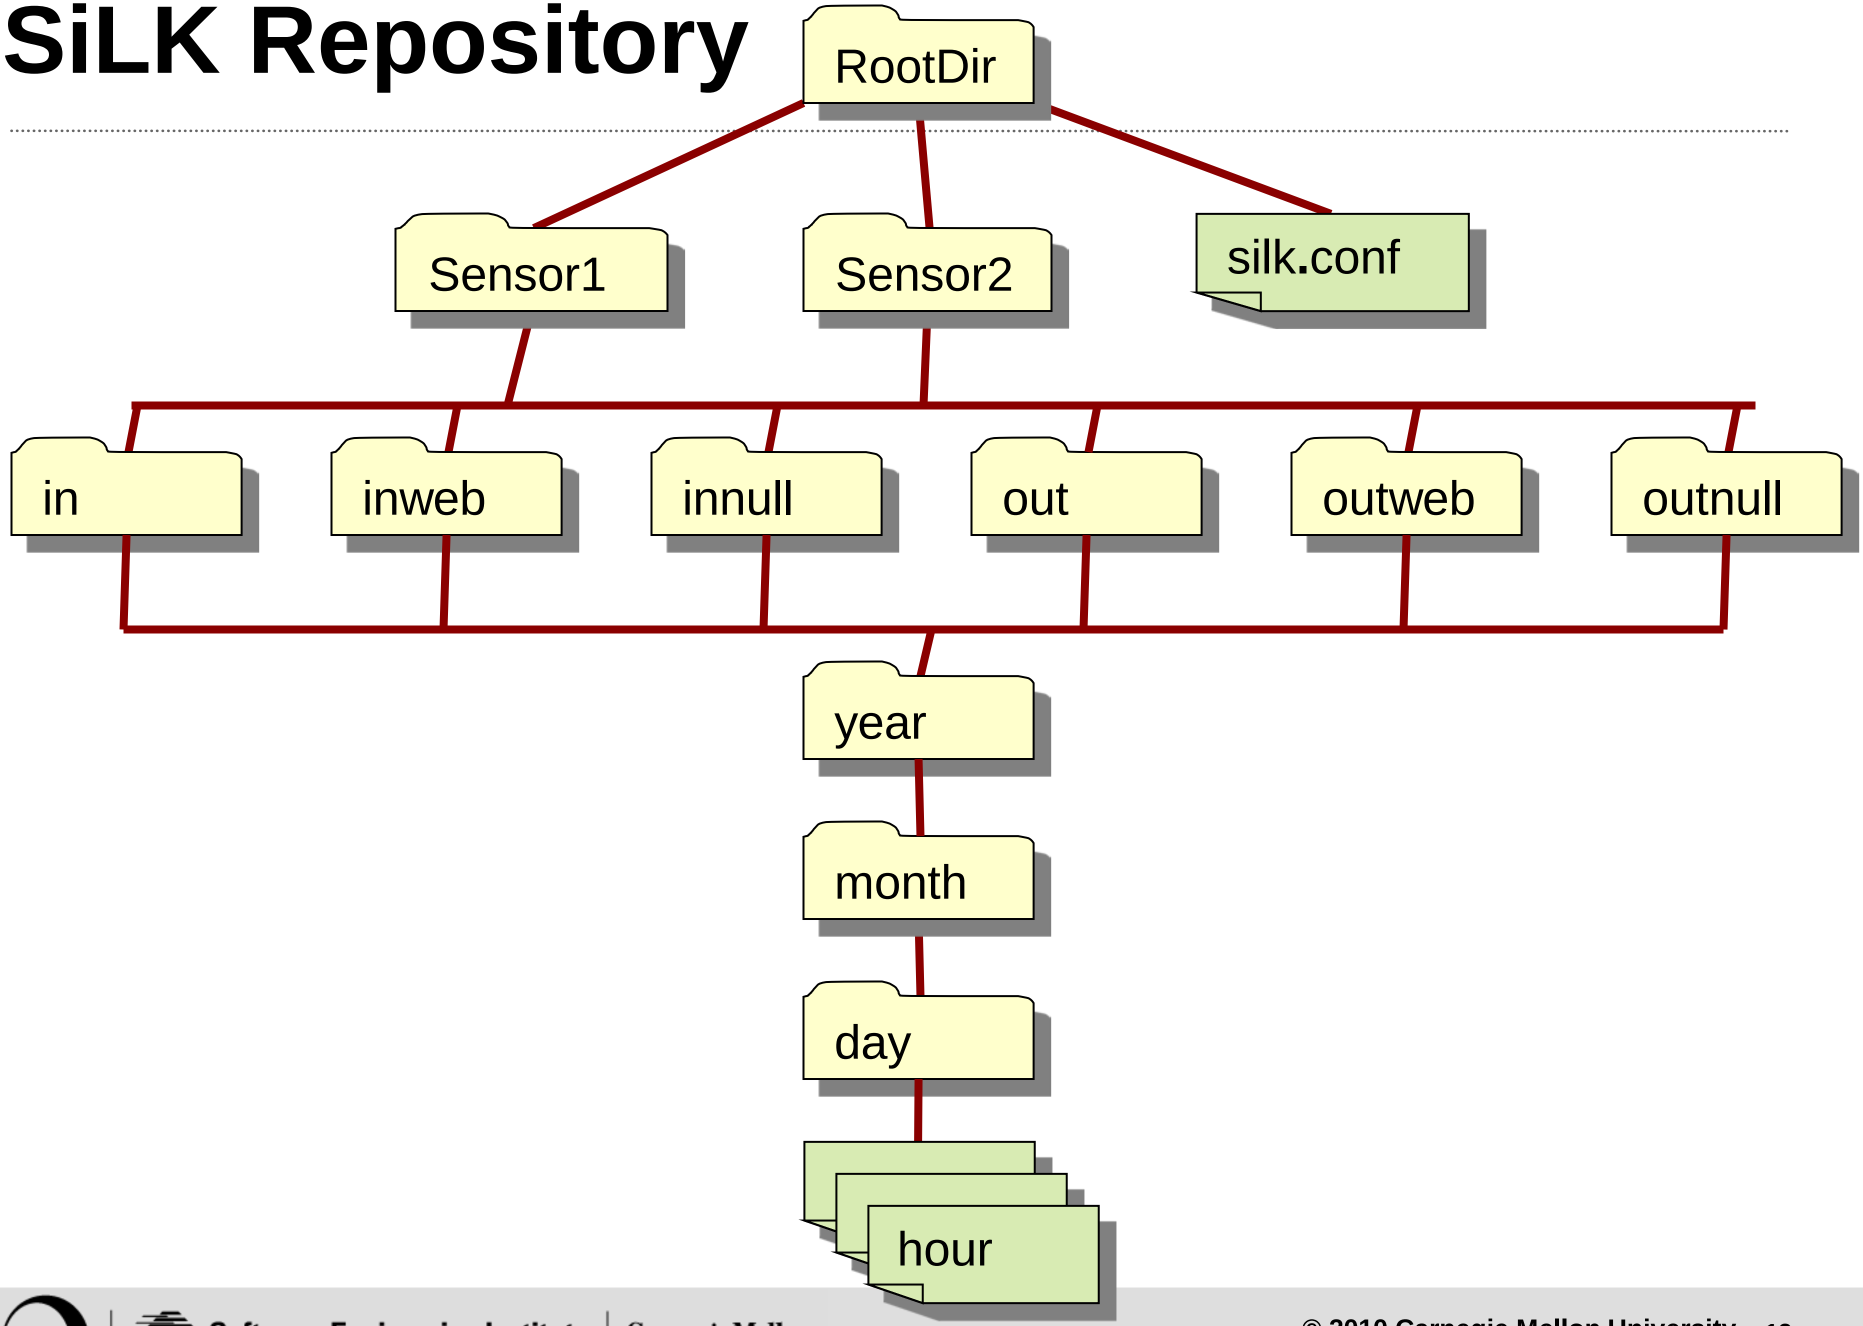

Basic commands

rwfilter

Partition SiLK Flow records into one or more 'pass' and/or 'fail' output streams. rwfilter is the primary tool for pulling flows from the data store.

rwsort

Sort SiLK Flow records using a user-specified key comprised of record attributes, and write the records to the named output path or to the standard output. Users can define new key fields using plug-ins written in C or PySiLK.

rwcut

Print the attributes of SiLK Flow records in a delimited, columnar, human-readable format. Users can define new printable attributes using plug-ins written in C or PySiLK.

rwuniq

Bin (group) SiLK Flow records by a user-specified key comprised of record attributes and print the total byte, packet, and/or flow counts for each bin. rwuniq can also print distinct source IP and destination IP counts. Users can define new key fields and value fields using plug-ins written in C or PySiLK.

rwcount

Summarize SiLK Flow records across time, producing textual output with counts of bytes, packets, and flow records for each time bin.

rwstats

Summarize SiLK Flow records by a user-specified key comprised of record attributes, compute values from the flow records that match each key, sort the results by the value to generate a Top-N or Bottom-N list, and print the results. Users can define new key fields and value fields using plug-ins written in C or PySiLK.

iSiLK – Why?

It helps me to choose SiLK tools

- Toolbar buttons allow quick perusal of tools

It lets me avoid SiLK tool syntax

- Menus & other GUI elements show my choices

It lets me avoid Linux command syntax and file names

- iSiLK organizes my data sets and results

It has an integrated graphing capability

GUI - iSiLK

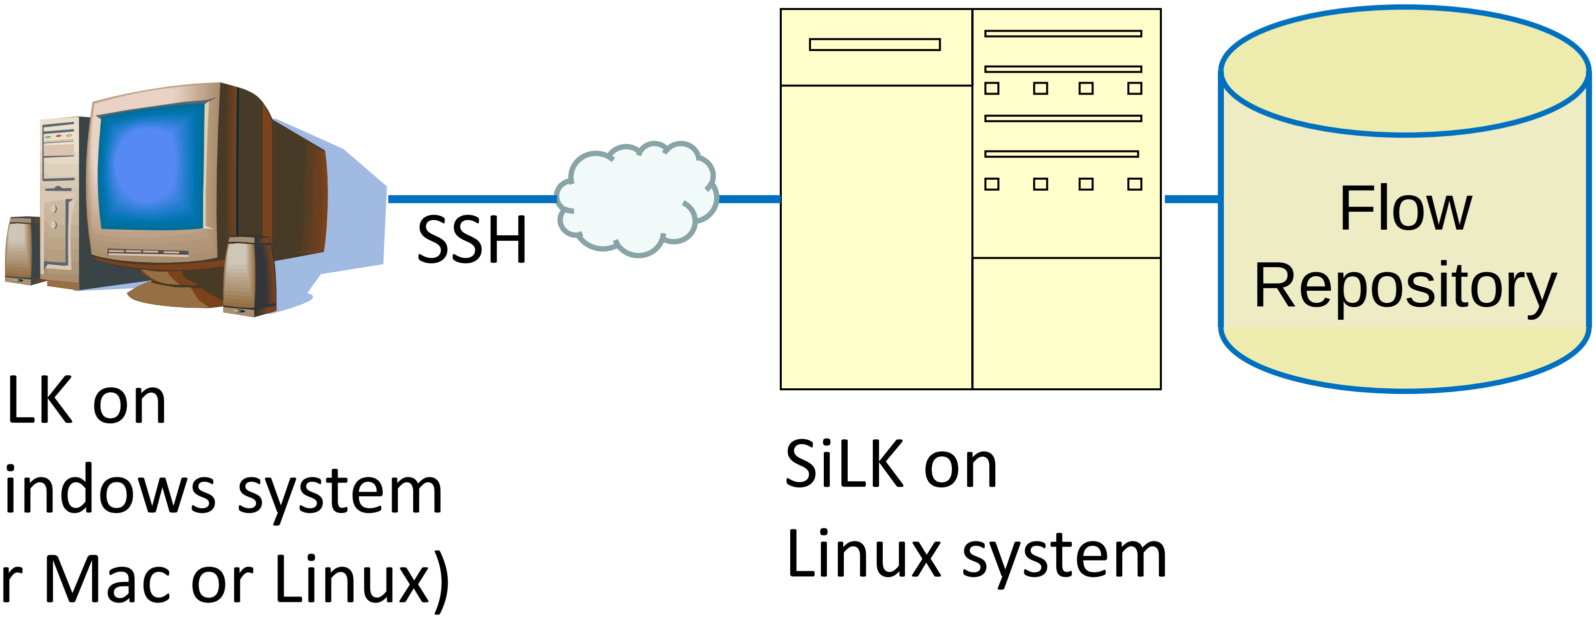
SSH
Flow Repository
iSiLK on
Windows system
(or Mac or Linux)
SiLK on
Linux system

iSiLK Query Builder

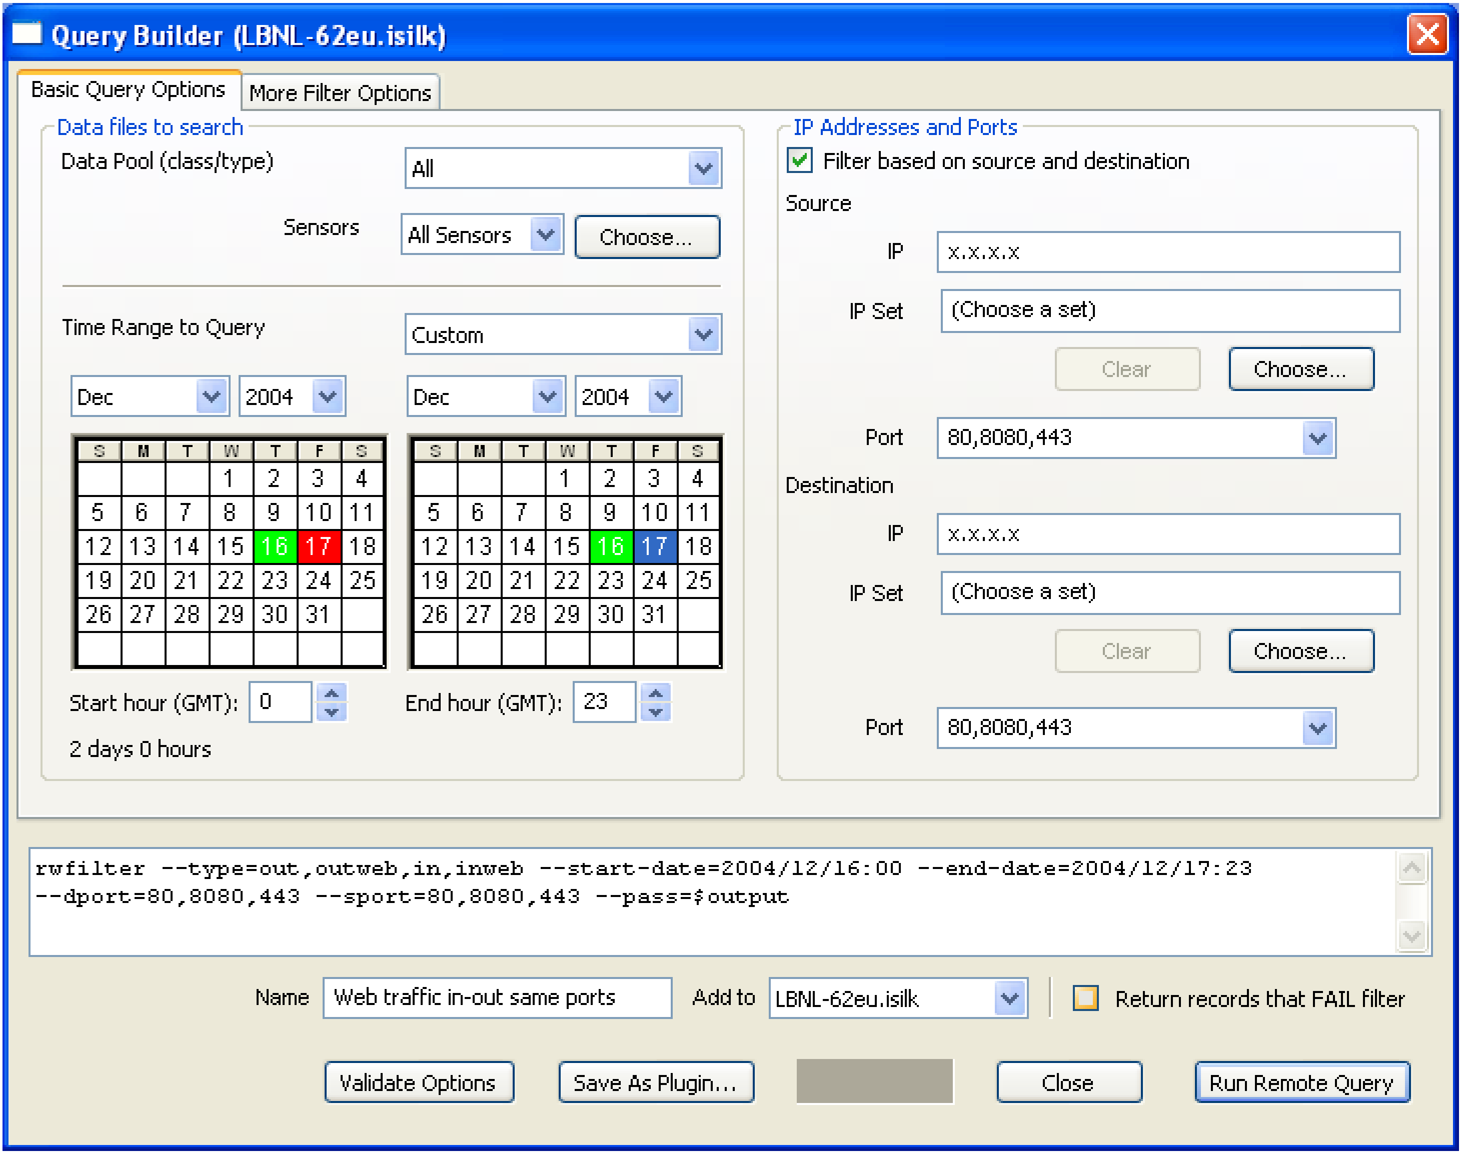

Data Pool (class/type): All

Sensors: All Sensors

Time Range to Query: Custom

Dec 2004

S	M	T	W	T	F	S
			1	2	3	4
5	6	7	8	9	10	11
12	13	14	15	16	17	18
19	20	21	22	23	24	25
26	27	28	29	30	31	

Dec 2004

S	M	T	W	T	F	S
			1	2	3	4
5	6	7	8	9	10	11
12	13	14	15	16	17	18
19	20	21	22	23	24	25
26	27	28	29	30	31	

Start hour (GMT): 0

End hour (GMT): 23

2 days 0 hours

☑ Filter based on source and destination

Source

IP: x.x.x.x

IP Set: (Choose a set)

Port: 80,8080,443

Destination

IP: x.x.x.x

IP Set: (Choose a set)

Port: 80,8080,443

```
rwfilter --type=out,outweb,in,inweb --start-date=2004/12/16:00 --end-date=2004/12/17:23
--dport=80,8080,443 --sport=80,8080,443 --pass=$output
```

Name: Web traffic in-out same ports

Add to: LBNL-62eu.isilk

□ Return records that FAIL filter

iSiLK Query Builder

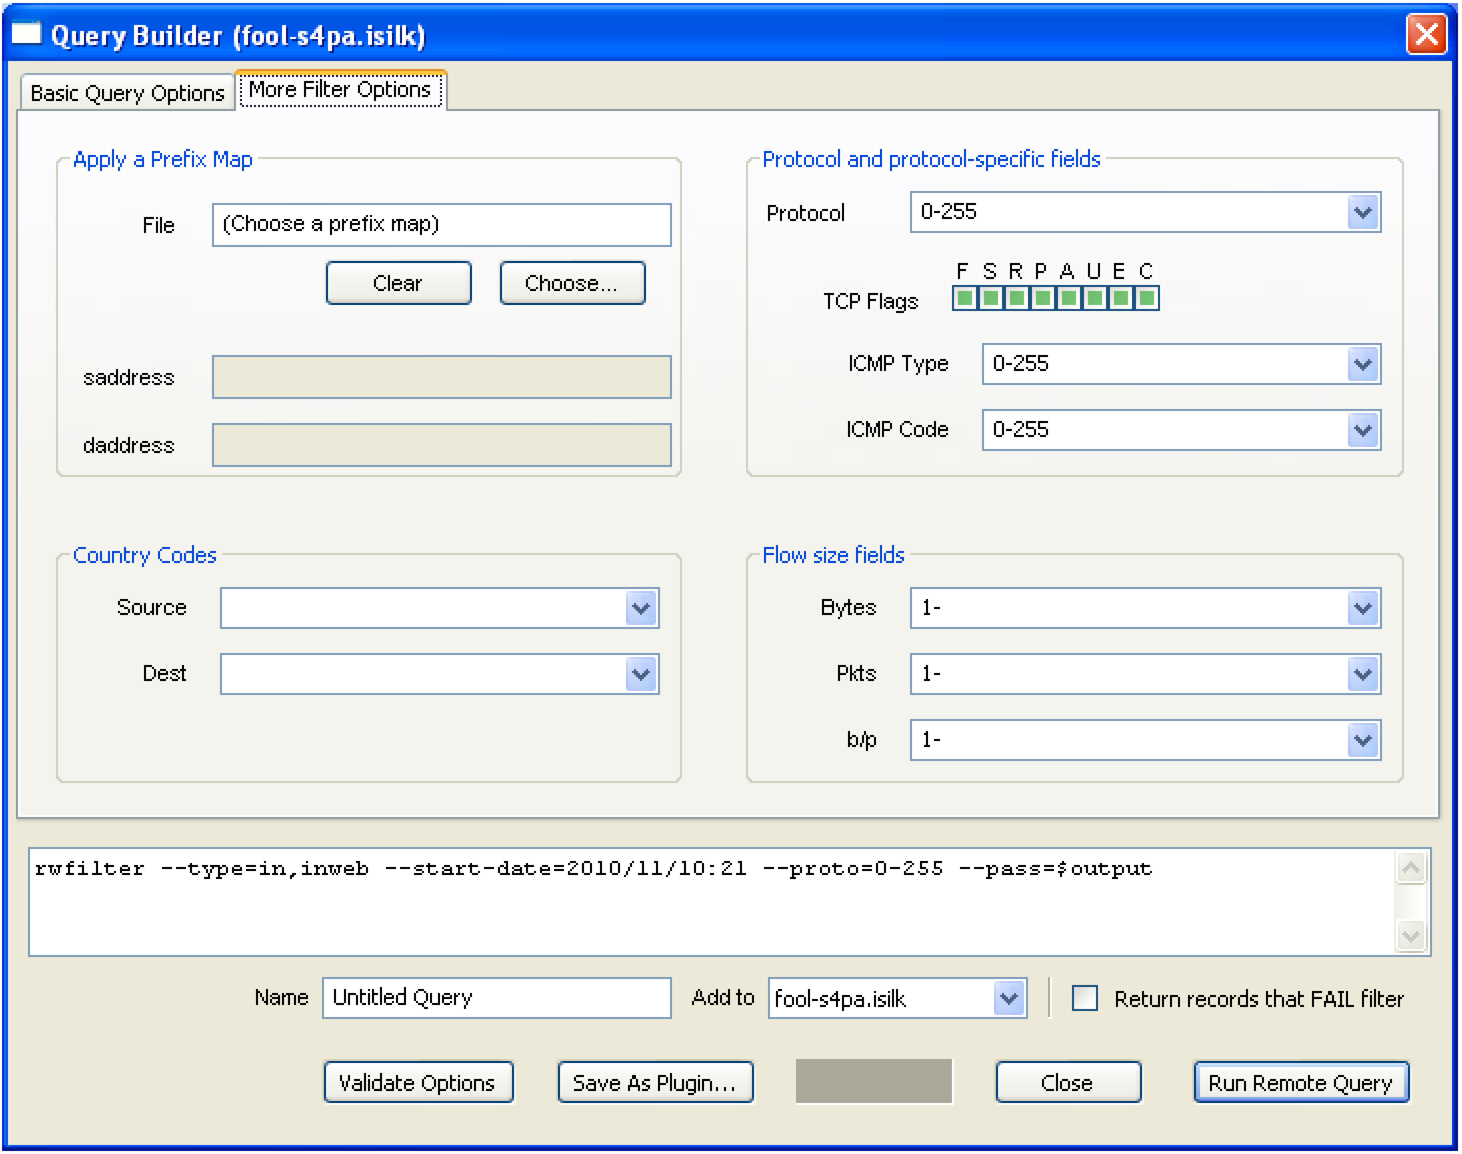

File: (Choose a prefix map)

saddress:

daddress:

Protocol: 0-255

TCP Flags: F S R P A U E C
🟩🟩🟩🟩🟩🟩🟩🟩

ICMP Type: 0-255

ICMP Code: 0-255

Source:

Dest:

Bytes: 1-

Pkts: 1-

b/p: 1-

```
rwfilter --type=in,inweb --start-date=2010/11/10:21 --proto=0-255 --pass=$output
```

Name: Untitled Query

Add to: fool-s4pa.isilk

□ Return records that FAIL filter

iSiLK Query Builder

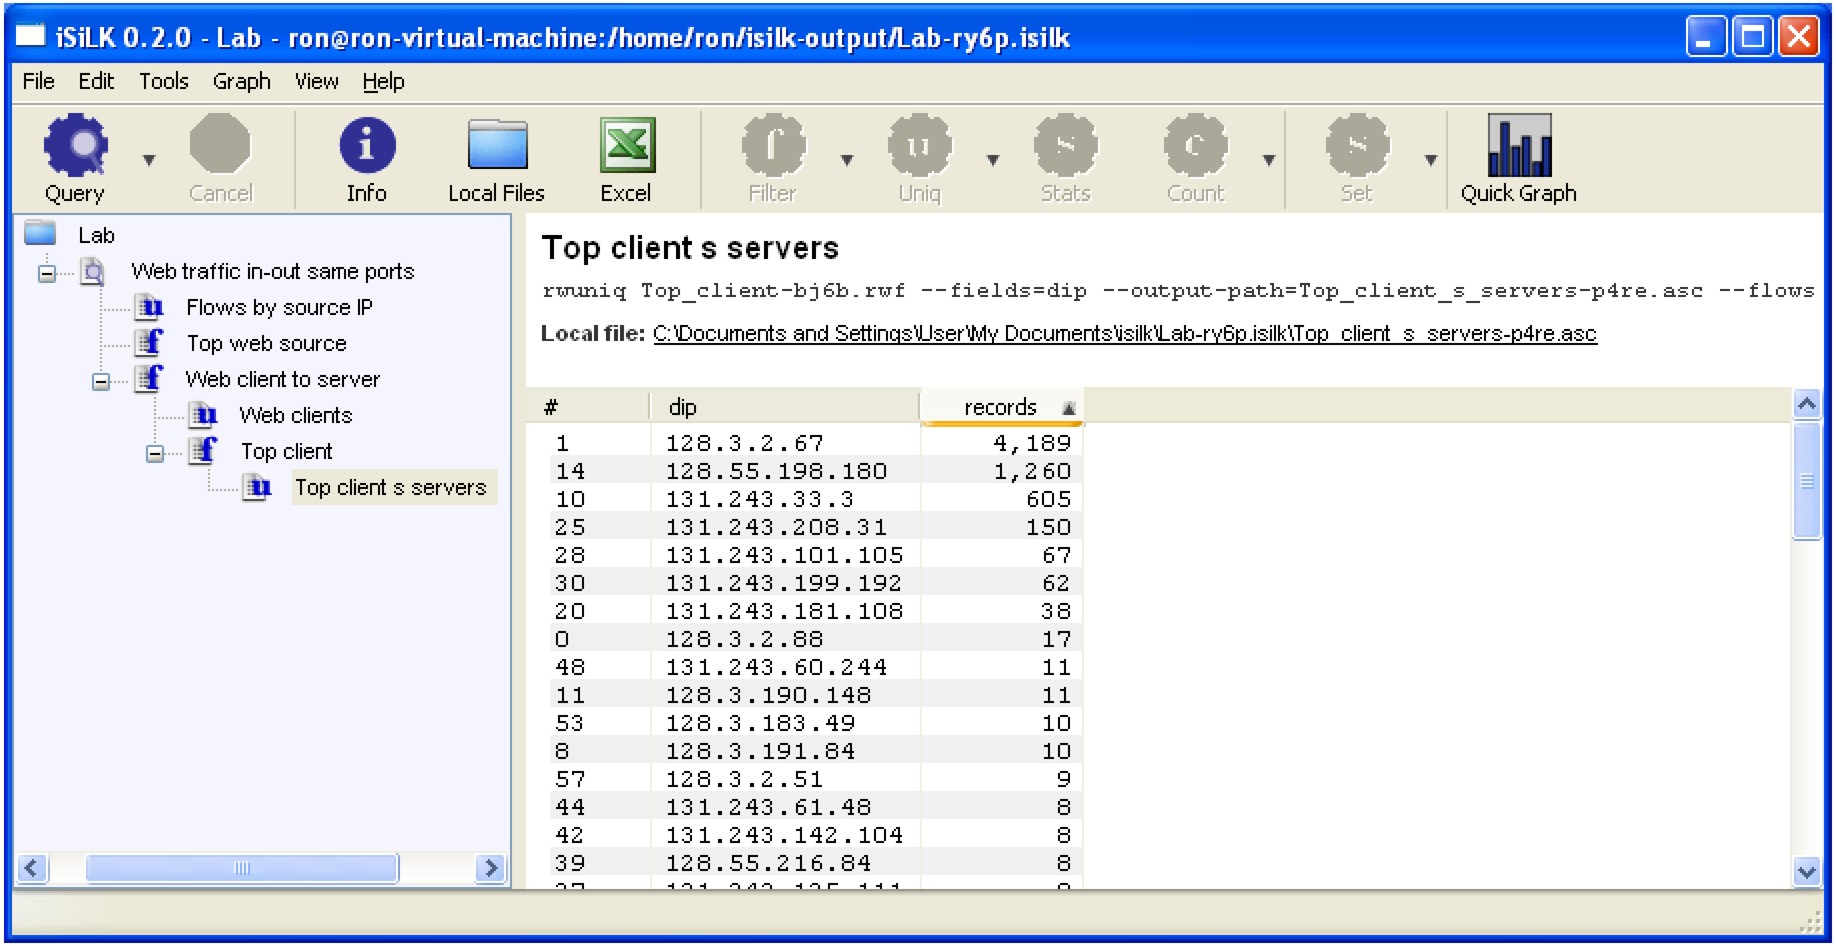

Top client s servers

```
rwuniq Top_client-bj6b.rwf --fields=dip --output-path=Top_client_s_servers-p4re.asc --flows
```

Local file: C:\Documents and Settings\User\My Documents\isilk\Lab-ry6p.isilk\Top_client_s_servers-p4re.asc

#	dip	records
1	128.3.2.67	4,189
14	128.55.198.180	1,260
10	131.243.33.3	605
25	131.243.208.31	150
28	131.243.101.105	67
30	131.243.199.192	62
20	131.243.181.108	38
0	128.3.2.88	17
48	131.243.60.244	11
11	128.3.190.148	11
53	128.3.183.49	10
8	128.3.191.84	10
57	128.3.2.51	9
44	131.243.61.48	8
42	131.243.142.104	8
39	128.55.216.84	8

SiLK Analysis Pipeline

SiLK was built to effectively query a repository

- Everything is retroactive

Closest to real time is batched jobs

Pipeline is a streaming analysis engine for SiLK flow files

- Versus retroactive SiLK tool analysis

“Near-time” alerting

SiLK Analysis Pipeline

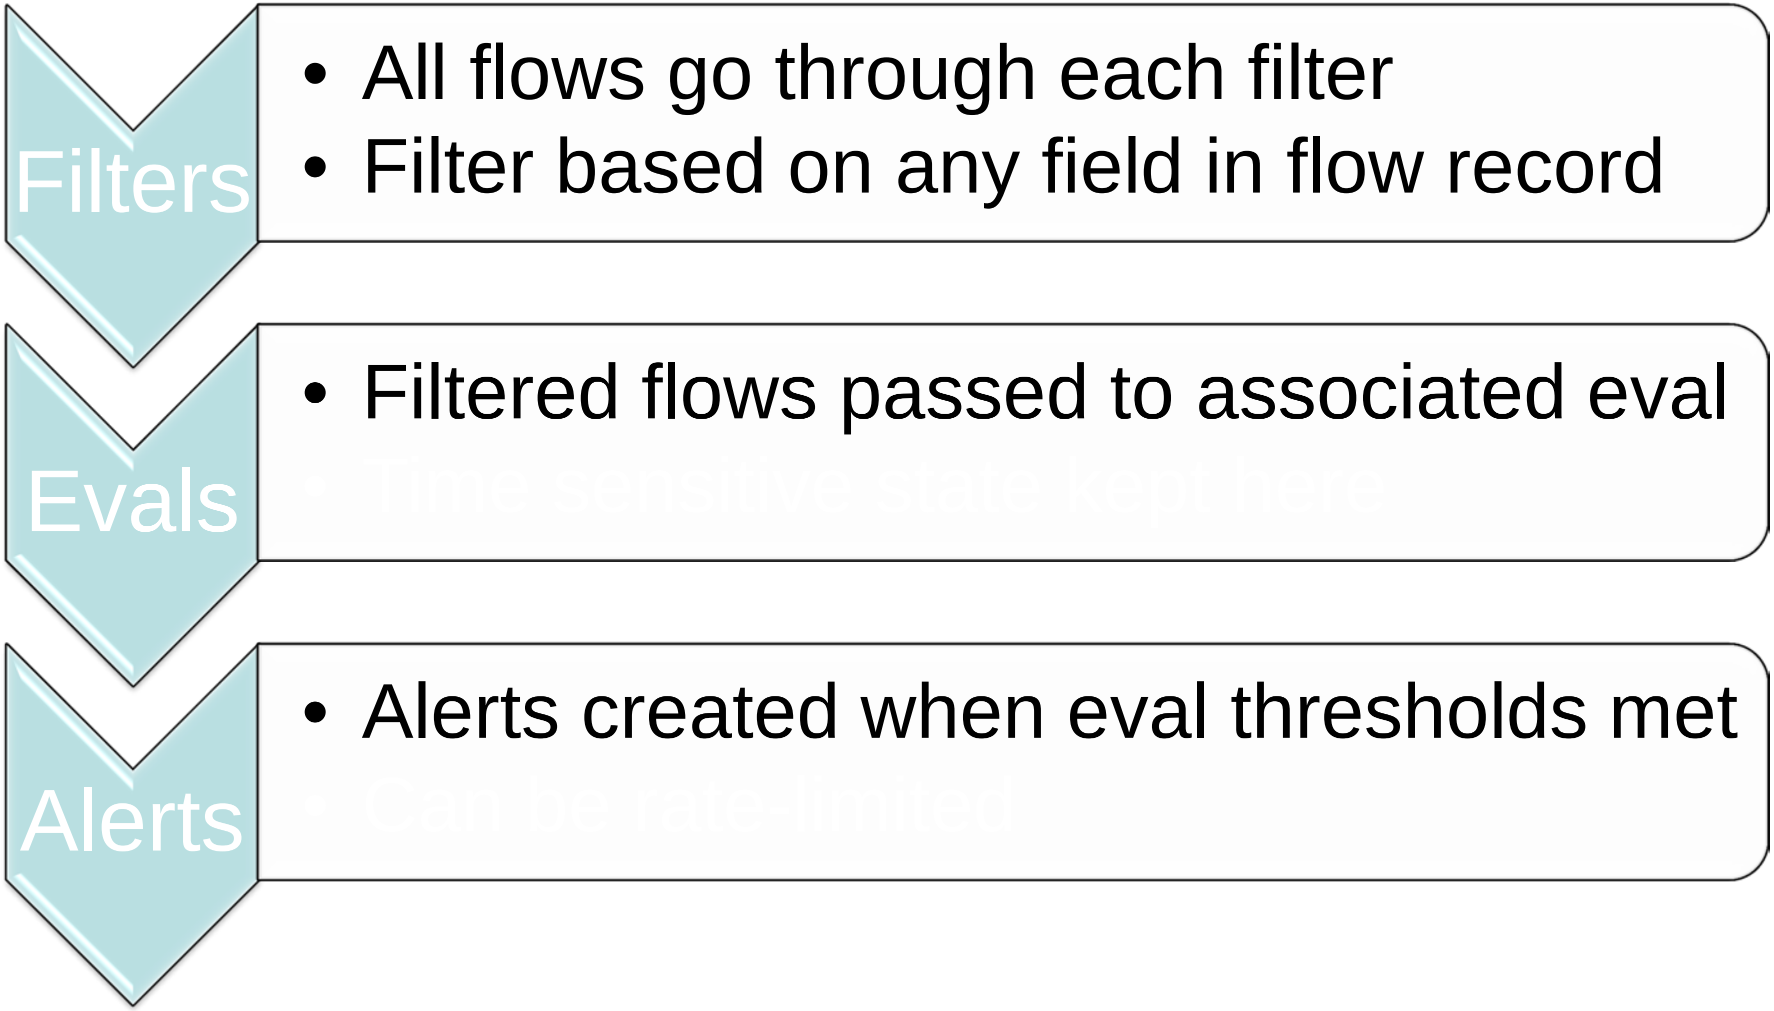

Filters

- All flows go through each filter
- Filter based on any field in flow record

Evals

- Filtered flows passed to associated eval
- Time sensitive state kept here

Alerts

- Alerts created when eval thresholds met
- Can be rate-limited

SiLK Pipeline Filters

Each evaluation gets its flows from **one** filter

A filter can provide for multiple evaluations

A single filter is specified in the configuration file for each evaluation.

Operators for any field in flow record

- <, <=, >, >=, ==, !=, IN_LIST, NOT_IN_LIST
- Each filter can have multiple “anded” comparisons

SiLK Pipeline Evaluations

Can have time restrictions:

- Alert if “this” happens in any 5 minute period

Made up of a number of independent checks

- E.g. Bytes > 1000 and packets > 500 in 5 minutes

Each check can be limited by its own time window

- Examples
  - Sum of Packets > 1000 in 10 minutes
  - Number of Unique Source IP Addresses > 10 in an hour
  - Total Flow Count > 10000 in 1 minute
- If all checks meet threshold, the evaluation alerts

SiLK Analysis Pipeline Capabilities

Finite State Beacon Detection

Sensor Outage Detection

IPv6 Tunnel Detection

Passive FTP Detection

Watchlists

Flow counts

Flow field based capabilities (Can be combined)

- Sum or Average of the field value (bytes, packets, durations, etc)
- Proportion of flows with a given field value (TCP, Web, etc)

SiLK Pipeline Alerts

When deemed able to alert, they contain:

- The flow record
- Evaluation name as identifier
- Metrics that triggered alert and its threshold
- Timestamp

Currently output to arcSight files

Can output to files and logs

Tools – Bringing it all together

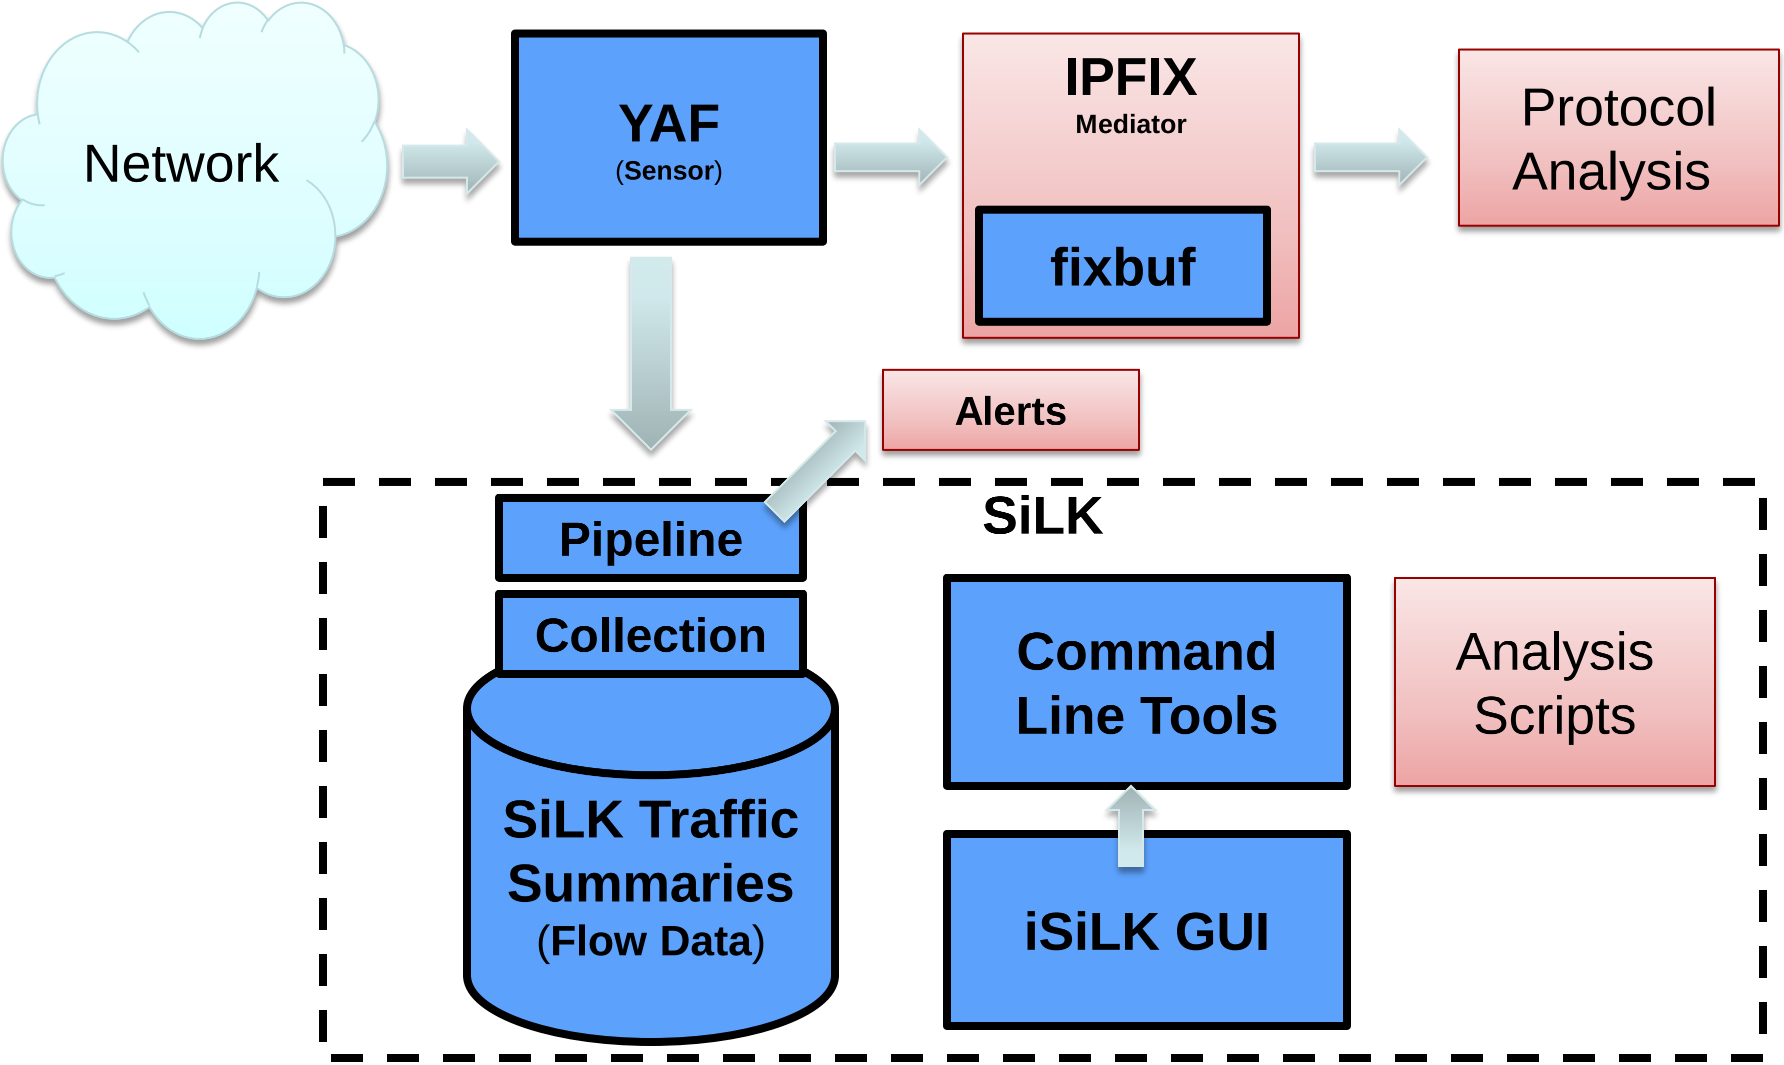

Network

YAF
(Sensor)

IPFIX
Mediator

fixbuf

Protocol Analysis

Alerts

SiLK

Pipeline

Collection

SiLK Traffic Summaries
(Flow Data)

Command Line Tools

iSiLK GUI

Analysis Scripts

Prism

Data examined on a schedule (e.g. once daily)

Data diverted by first filter hit (filters designed to match known traffic)

Partitions data by filter expressions, order matters!

Visualization of known traffic examined for anomalies

Remainder of data analyzed for anomalies, malicious activity

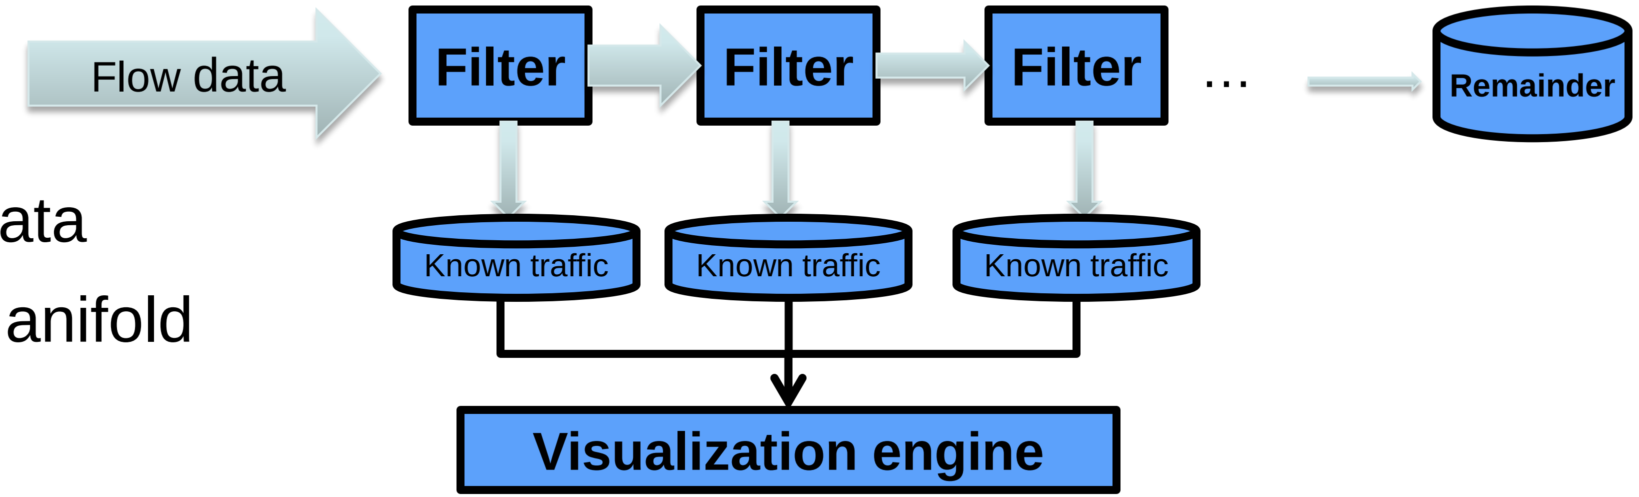

Data

Manifold

Prism Output example

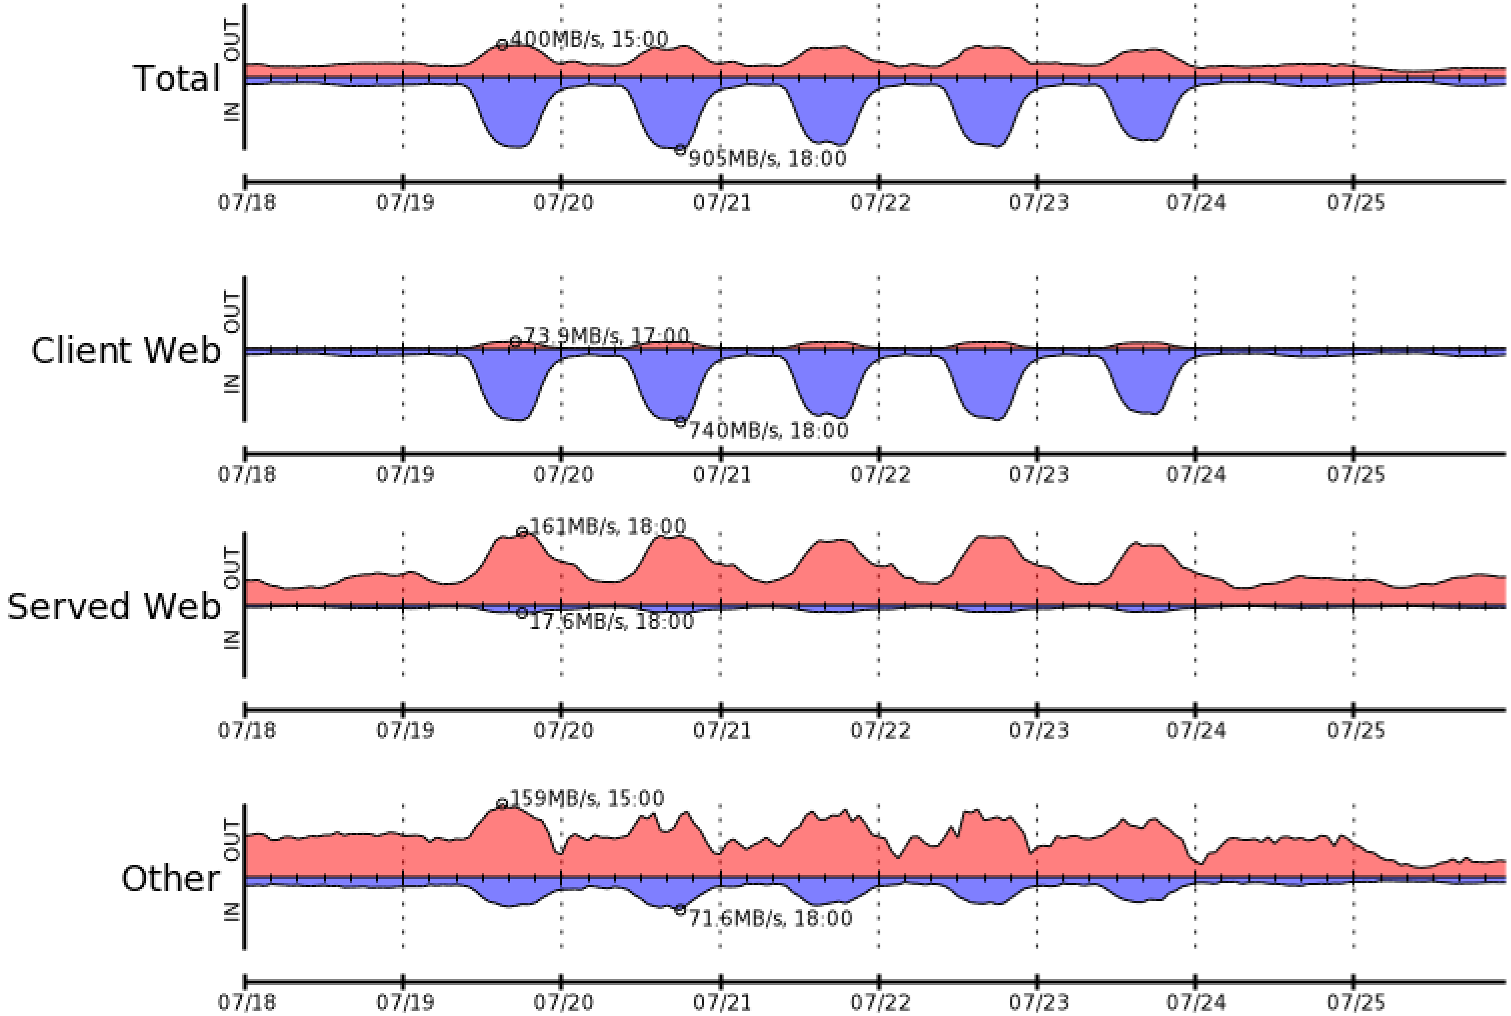

Beacons

SiLK scripts or watchlists to identify possible beaconing

Time series plots of possible beacons

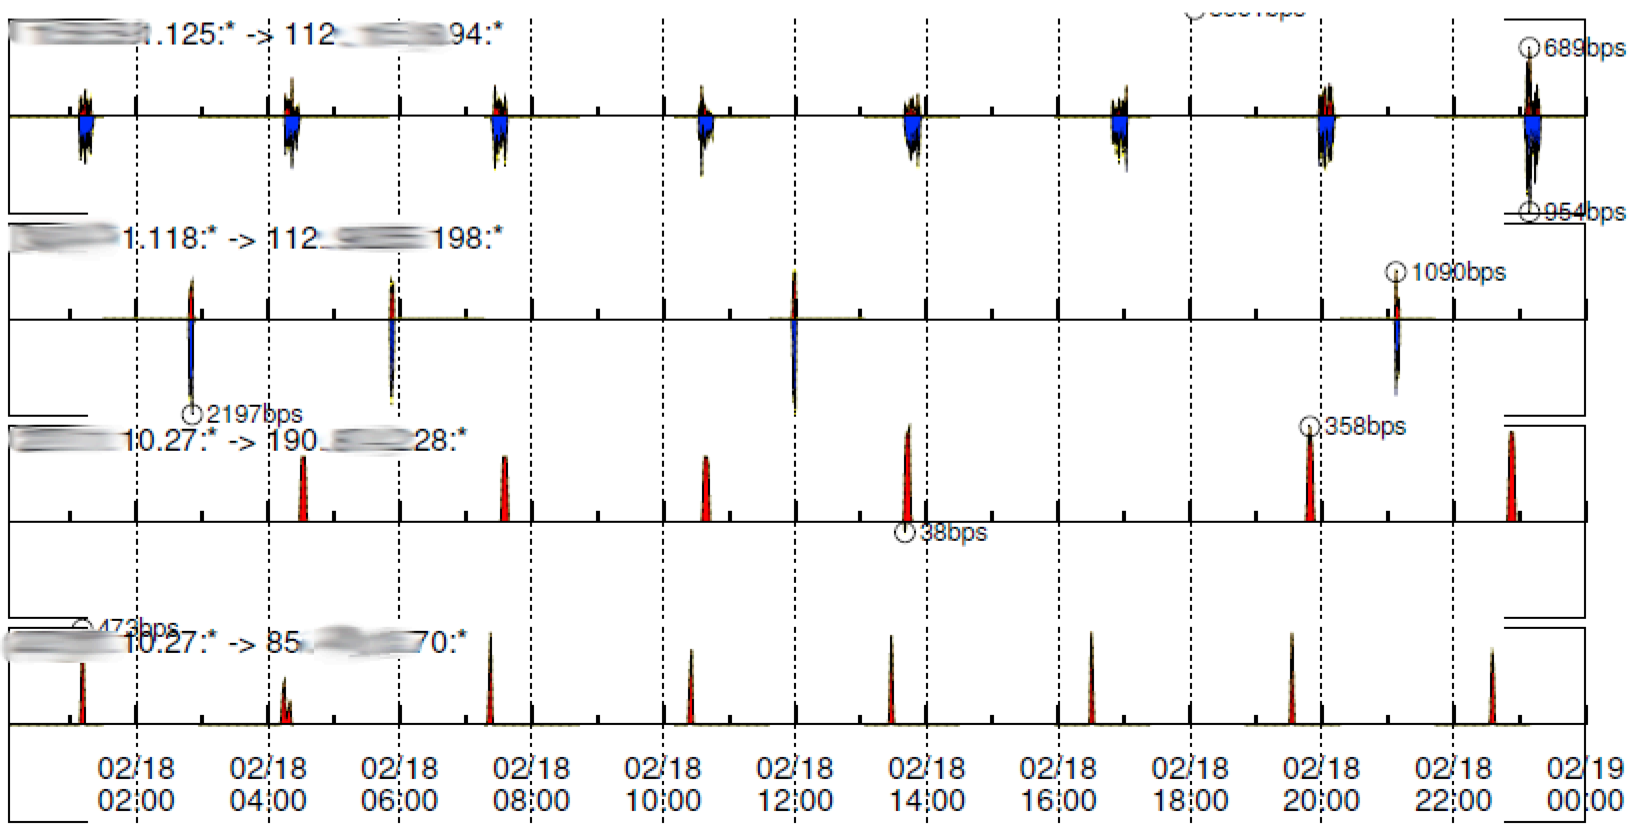

After the fact forensics/modeling

Who talked to who?

What else did this box / emplyoee do?

Who else had this malware?

Large outbound flows?

Who are the top talkers on the network?

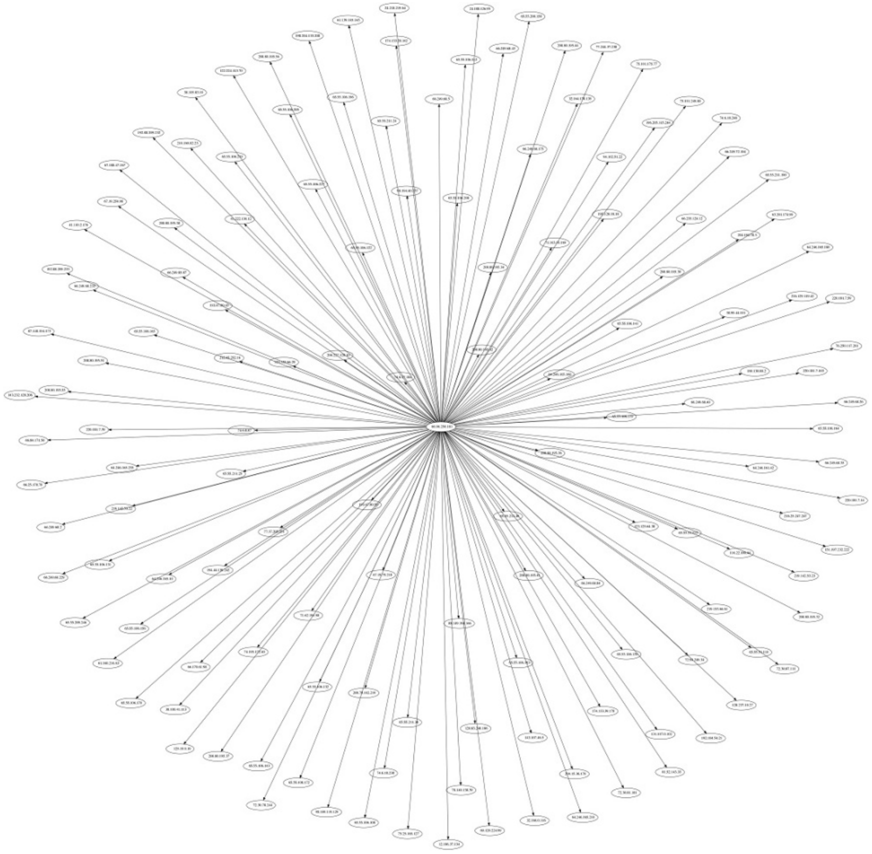

Rayon – Viz Library and Tools

Motivation

- Improve transition and uptake of analytics
- Provide basic visualization SOC analysts can use easily
- Integrate well into existing workflow (i.e. command-line)
- Inclusion in iSiLK

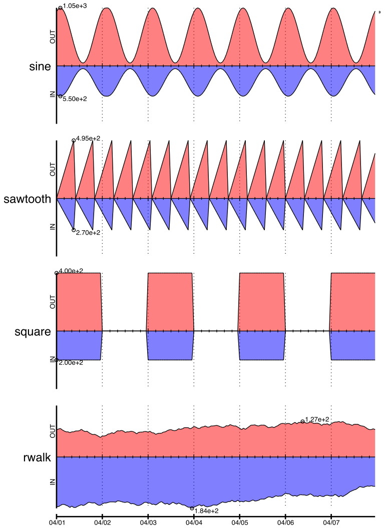

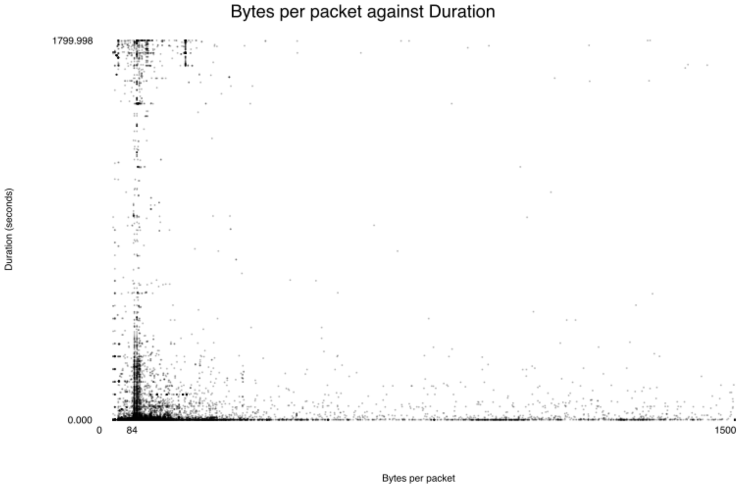

Data fusion

DNS Data - Fast Flux domains, Phishing URLs

Malware analysis - Network Touchpoints

Simple analytics from flow feeding SIEM for correlation

Leveraging DPI

YAF Inspector – Initial proof of concept for YAF's extended capabilities

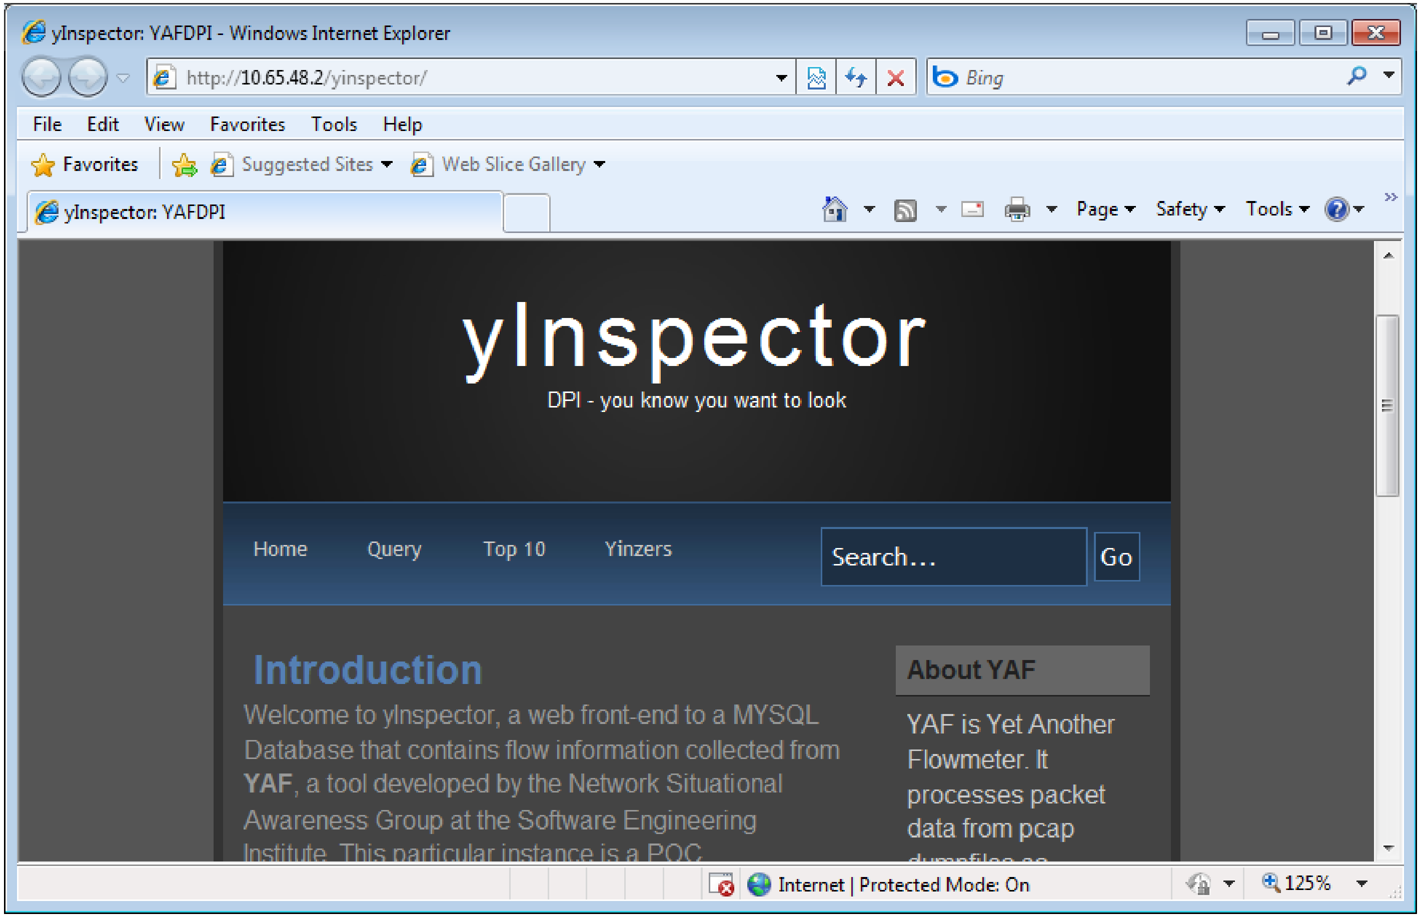
yInspector: YAFDPI - Windows Internet Explorer
http://10.65.48.2/yinspector/
Bing
File Edit View Favorites Tools Help
Favorites
Suggested Sites
Web Slice Gallery
yInspector: YAFDPI
Page Safety Tools
Home
Query
Top 10
Yinzers
Search…
Go
Internet | Protected Mode: On
125%

yInspector

DPI - you know you want to look

Introduction

Welcome to yInspector, a web front-end to a MYSQL Database that contains flow information collected from **YAF**, a tool developed by the Network Situational Awareness Group at the Software Engineering Institute. This particular instance is a POC

About YAF

YAF is Yet Another Flowmeter. It processes packet data from pcap dumpfiles as

Leveraging DPI

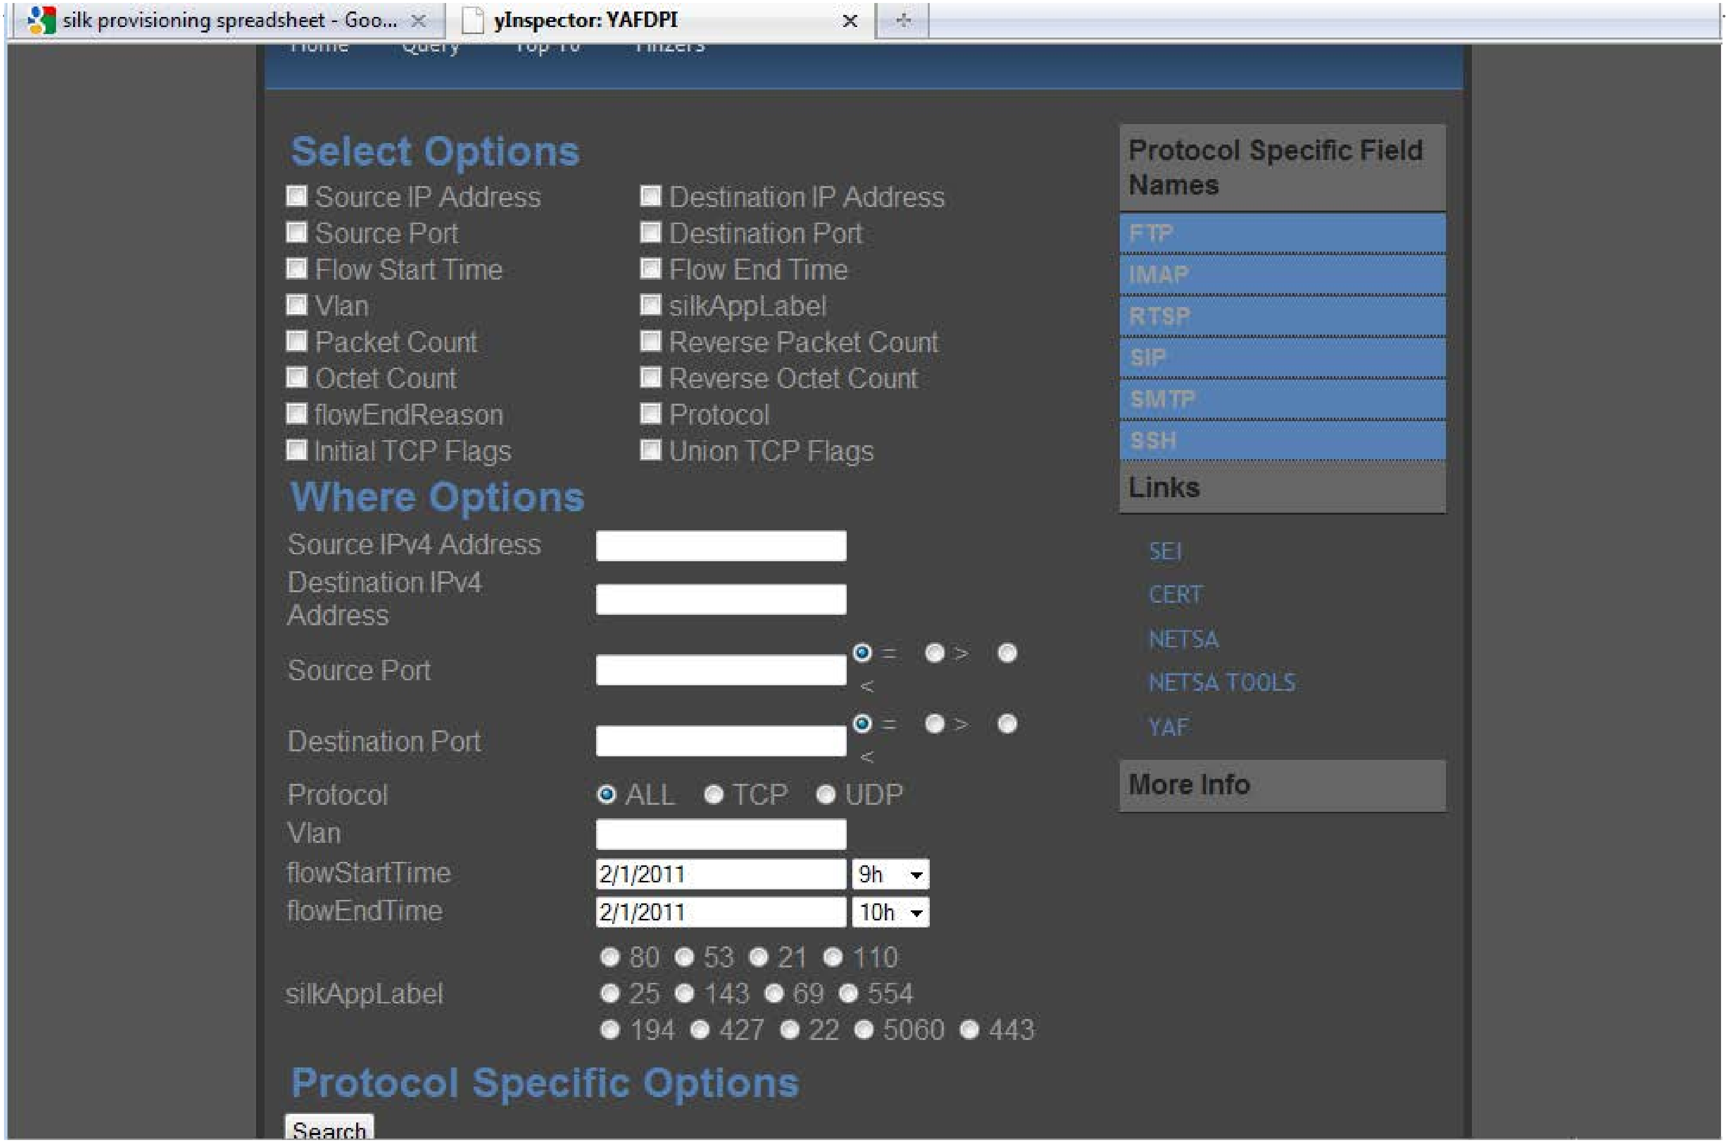

Select Options

☐ Source IP Address
☐ Source Port
☐ Flow Start Time
☐ Vlan
☐ Packet Count
☐ Octet Count
☐ flowEndReason
☐ Initial TCP Flags

☐ Destination IP Address
☐ Destination Port
☐ Flow End Time
☐ silkAppLabel
☐ Reverse Packet Count
☐ Reverse Octet Count
☐ Protocol
☐ Union TCP Flags

Where Options

Source IPv4 Address

Destination IPv4 Address

Source Port ● = ○ > ○ <

Destination Port ● = ○ > ○ <

Protocol ● ALL ○ TCP ○ UDP

Vlan

flowStartTime 2/1/2011 9h

flowEndTime 2/1/2011 10h

silkAppLabel ○ 80 ○ 53 ○ 21 ○ 110 ○ 25 ○ 143 ○ 69 ○ 554 ○ 194 ○ 427 ○ 22 ○ 5060 ○ 443

Protocol Specific Options

Yinspector – Almost Live!

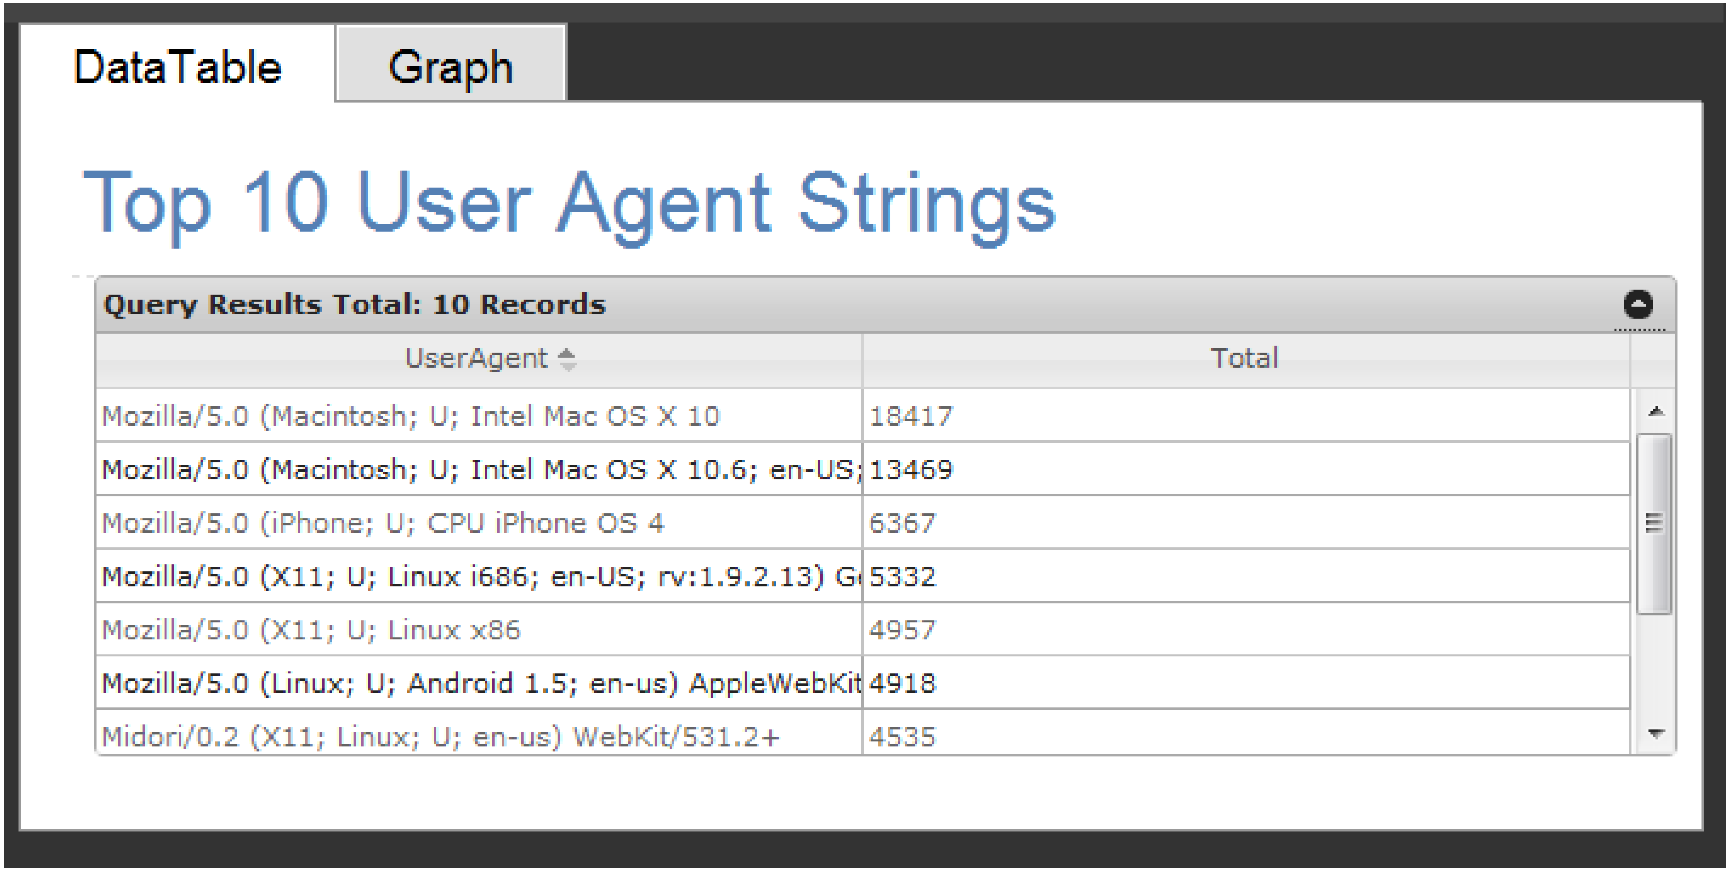

Top 10 User Agent Strings

Query Results Total: 10 Records

UserAgent	Total
Mozilla/5.0 (Macintosh; U; Intel Mac OS X 10	18417
Mozilla/5.0 (Macintosh; U; Intel Mac OS X 10.6; en-US;	13469
Mozilla/5.0 (iPhone; U; CPU iPhone OS 4	6367
Mozilla/5.0 (X11; U; Linux i686; en-US; rv:1.9.2.13) G	5332
Mozilla/5.0 (X11; U; Linux x86	4957
Mozilla/5.0 (Linux; U; Android 1.5; en-us) AppleWebKit	4918
Midori/0.2 (X11; Linux; U; en-us) WebKit/531.2+	4535

Yinspector – Almost Live!

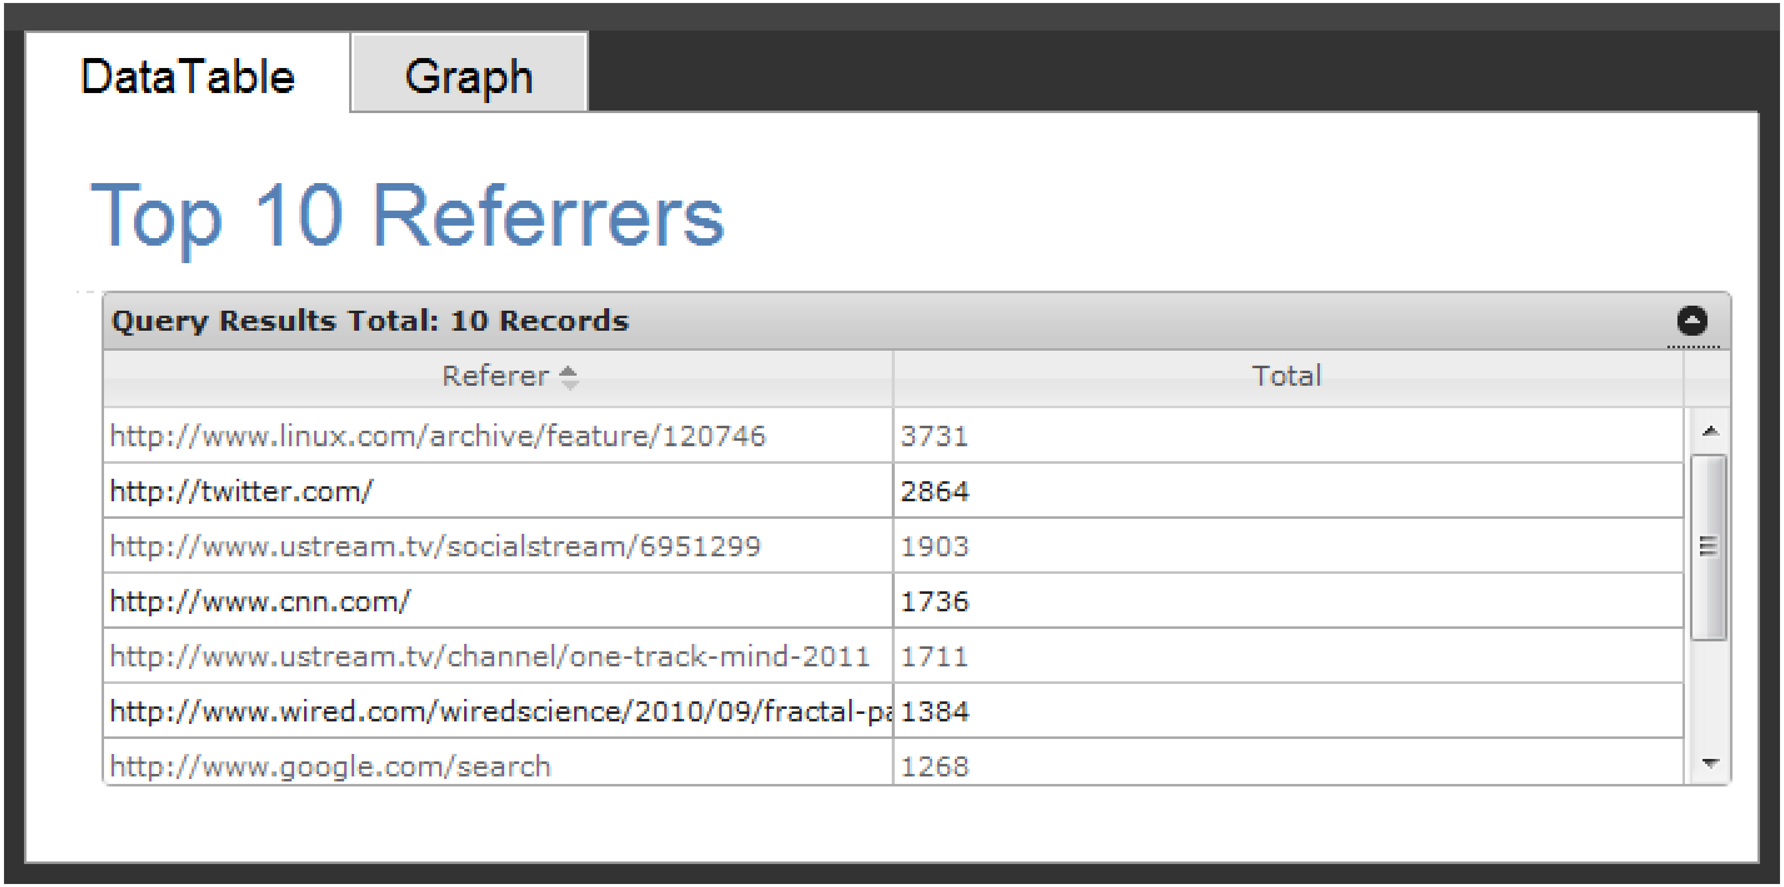

Top 10 Referrers

Query Results Total: 10 Records

Referer	Total
http://www.linux.com/archive/feature/120746	3731
http://twitter.com/	2864
http://www.ustream.tv/socialstream/6951299	1903
http://www.cnn.com/	1736
http://www.ustream.tv/channel/one-track-mind-2011	1711
http://www.wired.com/wiredscience/2010/09/fractal-pa	1384
http://www.google.com/search	1268

Yinspector – Almost Live!

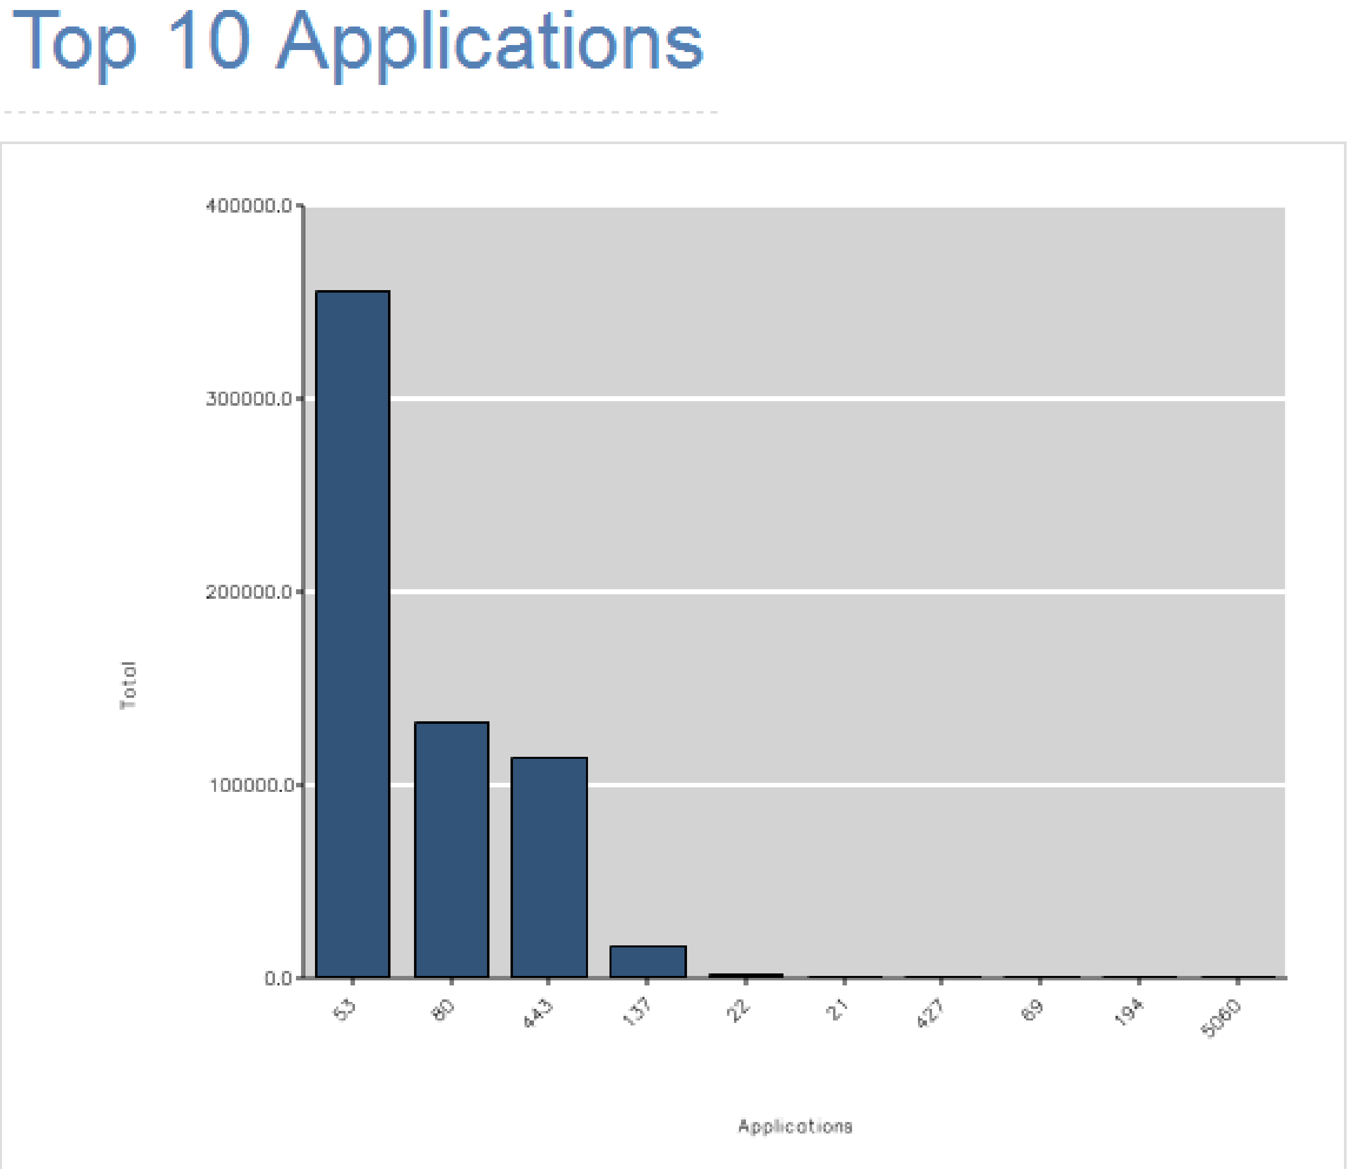

Top 10 Applications

Yinspector – Almost Live!

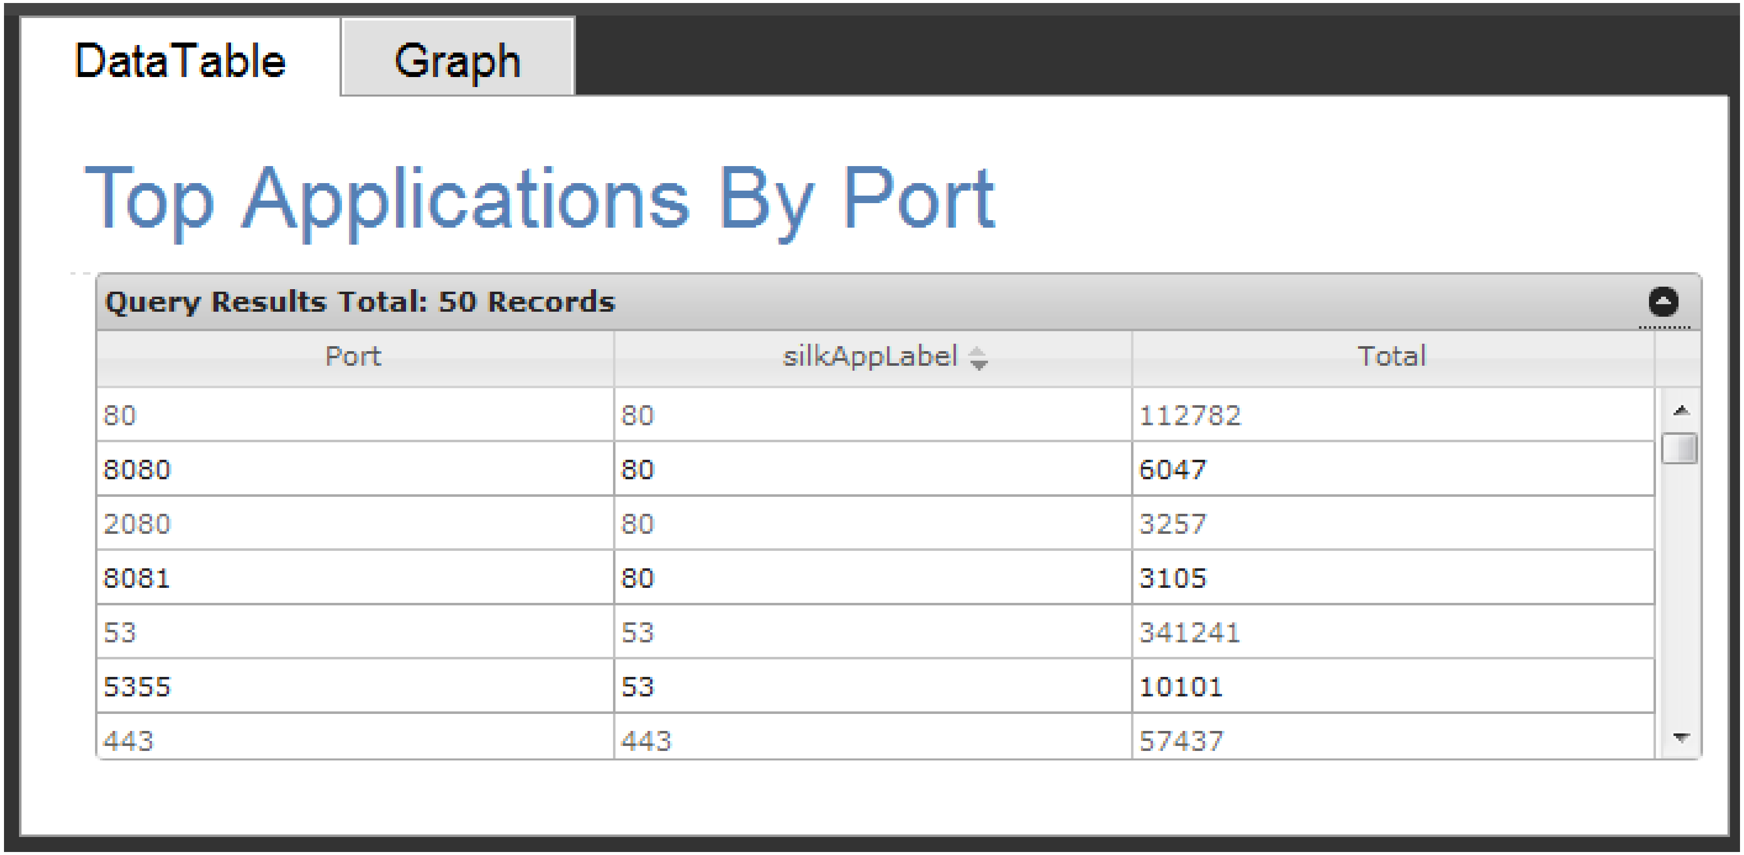

Top Applications By Port

Query Results Total: 50 Records

Port	silkAppLabel	Total
80	80	112782
8080	80	6047
2080	80	3257
8081	80	3105
53	53	341241
5355	53	10101
443	443	57437

Yinspector – Almost Live!

Top Talkers

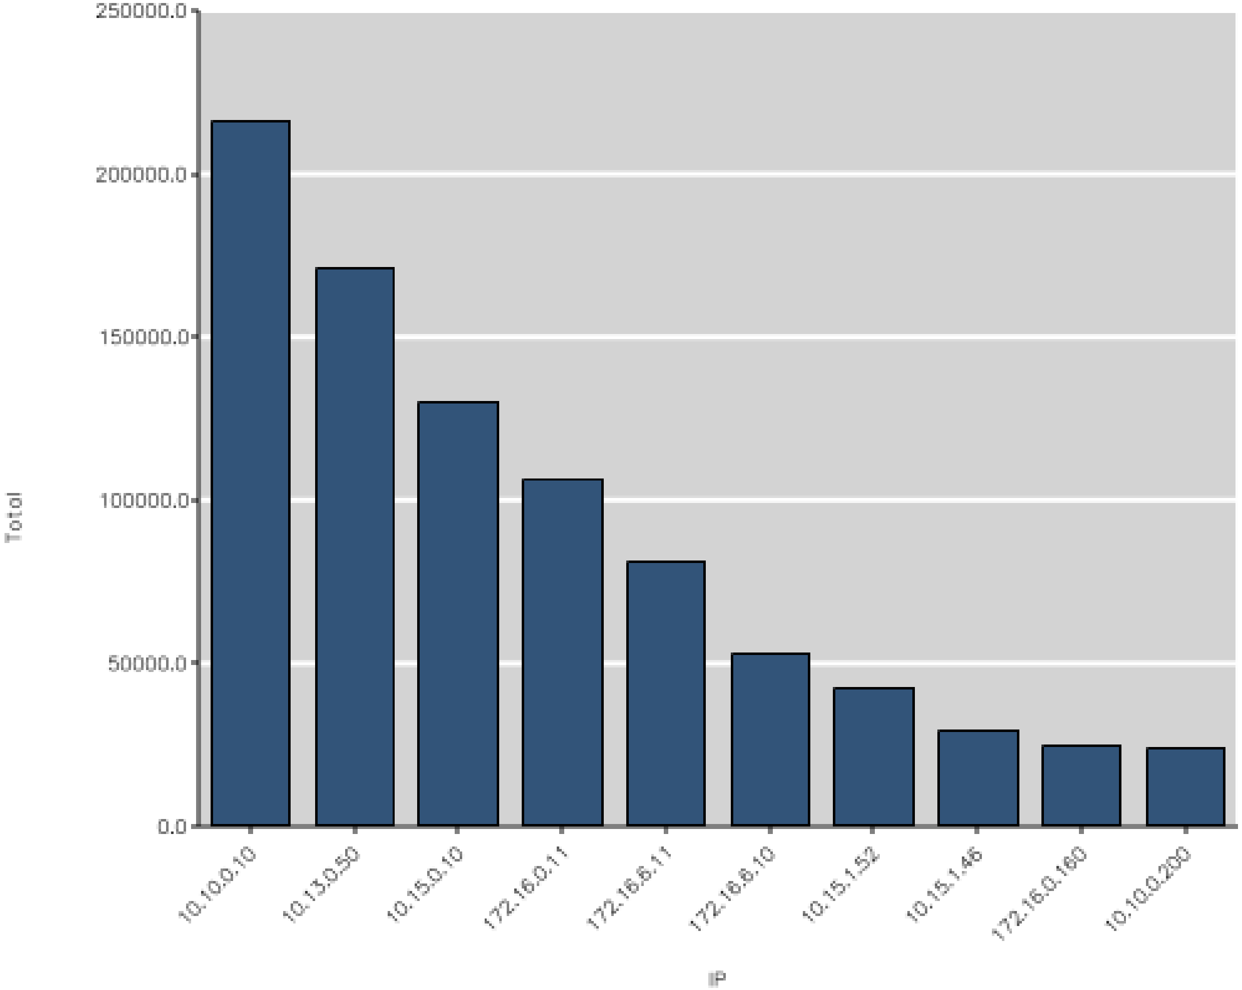

Build it!

Tools available: http://tools.netsa.cert.org

- Source
- RPM
- LiveCD
- Reference Data

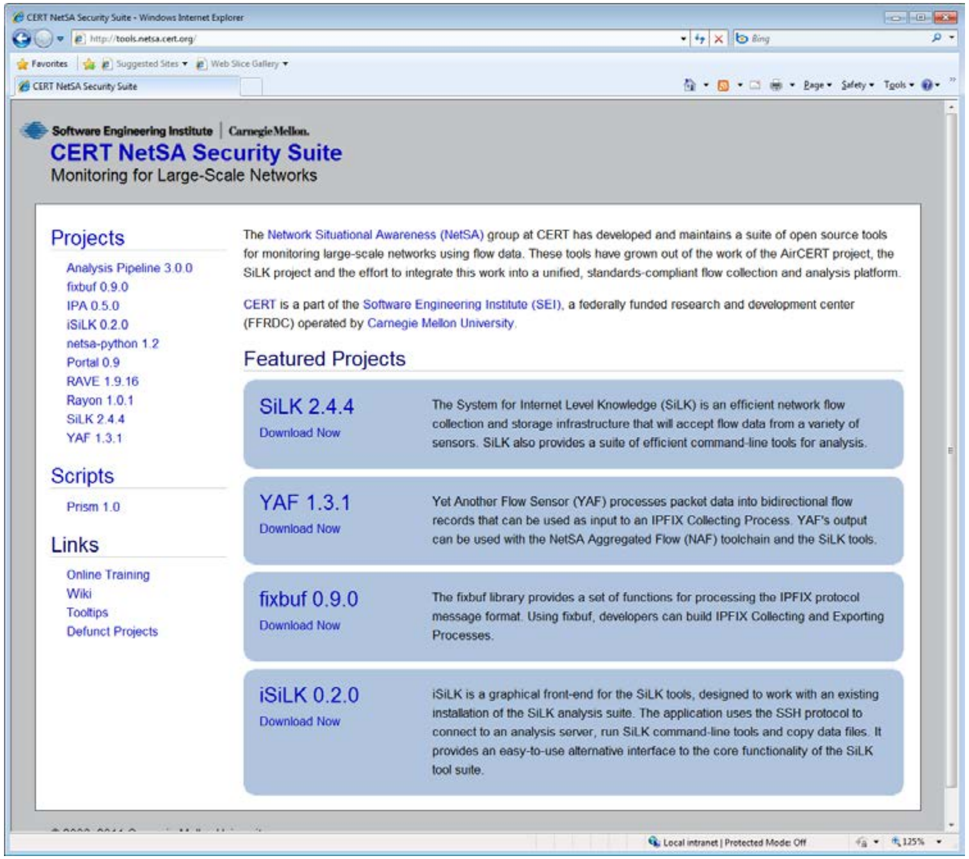
CERT NetSA Security Suite - Windows Internet Explorer
http://tools.netsa.cert.org/
Bing
Favorites
Suggested Sites
Web Slice Gallery
CERT NetSA Security Suite
Page
Safety
Tools
Local intranet | Protected Mode: Off
125%

Software Engineering Institute | Carnegie Mellon.

CERT NetSA Security Suite

Monitoring for Large-Scale Networks

Projects

Analysis Pipeline 3.0.0
fixbuf 0.9.0
IPA 0.5.0
iSiLK 0.2.0
netsa-python 1.2
Portal 0.9
RAVE 1.9.16
Rayon 1.0.1
SiLK 2.4.4
YAF 1.3.1

Scripts

Prism 1.0

Links

Online Training
Wiki
Tooltips
Defunct Projects

The Network Situational Awareness (NetSA) group at CERT has developed and maintains a suite of open source tools for monitoring large-scale networks using flow data. These tools have grown out of the work of the AirCERT project, the SiLK project and the effort to integrate this work into a unified, standards-compliant flow collection and analysis platform.

CERT is a part of the Software Engineering Institute (SEI), a federally funded research and development center (FFRDC) operated by Carnegie Mellon University.

Featured Projects

SiLK 2.4.4
Download Now

The System for Internet Level Knowledge (SiLK) is an efficient network flow collection and storage infrastructure that will accept flow data from a variety of sensors. SiLK also provides a suite of efficient command-line tools for analysis.

YAF 1.3.1
Download Now

Yet Another Flow Sensor (YAF) processes packet data into bidirectional flow records that can be used as input to an IPFIX Collecting Process. YAF's output can be used with the NetSA Aggregated Flow (NAF) toolchain and the SiLK tools.

fixbuf 0.9.0
Download Now

The fixbuf library provides a set of functions for processing the IPFIX protocol message format. Using fixbuf, developers can build IPFIX Collecting and Exporting Processes.

iSiLK 0.2.0
Download Now

iSiLK is a graphical front-end for the SiLK tools, designed to work with an existing installation of the SiLK analysis suite. The application uses the SSH protocol to connect to an analysis server, run SiLK command-line tools and copy data files. It provides an easy-to-use alternative interface to the core functionality of the SiLK tool suite.

Where to go for more

http://tools.netsa.cert.org

- Software
- Documentation
- Installation Guide, Analyst Handbook
- Online (CBT) training
- Wiki / Tooltips
- Scripts

FloCon® 2011

FloCon 2012 (mid-January) www.flocon.org

Specifics coming soon

We are hiring! Come see me. http://cert.org/jobs

Q&A

http://tools.netsa.cert.org

netsa-help [at] cert.org

rf [at] cert.org