


# Adding Network Flow Analysis to Your Security Architecture

**Sid Faber**  
Member of the Technical Staff  
CERT Network Situational Awareness  
*[sfaber@cert.org](mailto:sfaber@cert.org)*

 Software Engineering Institute | Carnegie Mellon

© 2011 Carnegie Mellon University

Welcome to this session on Network Flow Analysis.

# Outline

---

- What?
  - a flow primer
- Why?
  - whether it fits your needs
- How?
  - tools and analysis
- More?
  - more analysis, better tools, additional resources

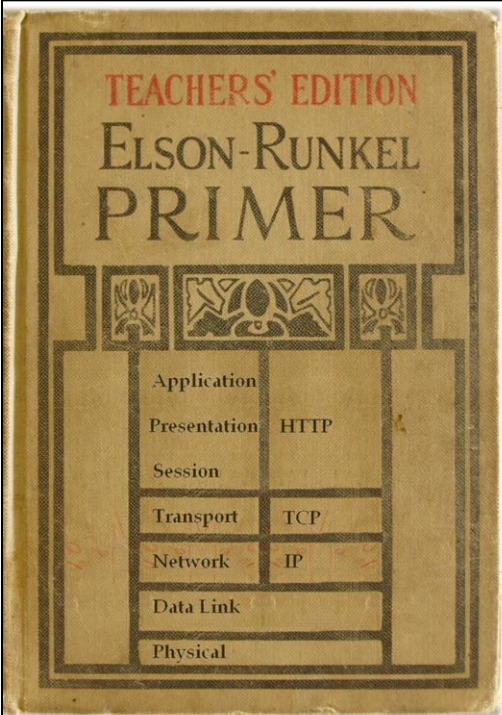


In this session we'll begin by covering the basics of network flow analysis. After a brief look at the history of flow, we'll dig into some examples of how common network traffic appears in flow.

Once the basics have been covered, we'll plug flow into a typical network security plan. There are some places gaps left by other security tools that can be filled by flow, so by the end of this discussion you should be able to decide whether flow might help solve some of your challenge problems.

Now that you've decided to add flow into your architecture, we need to find the right tools. In this section we'll talk about the architecture of the typical SiLK installation.

We'll close with a look at some of the more advanced tools and things that are in the works.



**TEACHERS' EDITION  
ELSON-RUNKEL  
PRIMER**

Application	
Presentation	HTTP
Session	
Transport	TCP
Network	IP
Data Link	
Physical	

**A Flow Primer**

*A is for Application,  
B is for Bulk Transport*

CERT | Software Engineering Institute | Carnegie Mellon

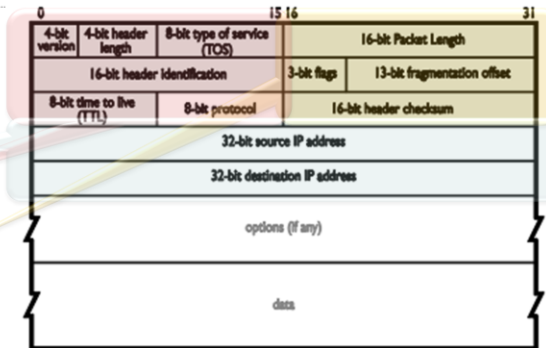
© 2011 Carnegie Mellon University 3

So let's get started looking at what flow is all about.


# In the Beginning...

First there were packets

- Source and destination IP
- Size & flag data
- Routing information (TTL, fragmentation, QoS, etc)



Then protocols added ports



Let's take a closer look at the relationship between data packets and flow records.

Here you see the IP and TCP packet header. As you already know, each packet contains the source and destination IP addresses. The IP packet also contains information about packet size, IP flags which help the client interpret the data contained in the packet.

Routing information, and other packet related data in the IP header help get the packet to its destination. TCP and UDP protocols encapsulated within IP packets contain source and destination ports which help the packet find the right service on the destination machine.

# Accounting for packets

The ISP asks, “who is using my bandwidth”

- Only routers know
- Very high volume
- Routers know packets, we need them to summarize

Netflow was developed by Cisco in 1996

- Proprietary
- Evolved into the primary network accounting method
- IETF standard on IPFIX (based on Cisco NetFlow v9), RFC5101



So then how did packets end up generating flow? In the early days of Internet routing, the service providers quickly realized that it was very important to know who was using network resources.

Some features of this problem quickly became apparent:

- Only routers tracked the information that was provided,
- The volume of traffic was very high so data collection had to match the volume, and
- Routers are very good at understanding and moving packets, so they can summarize packet data—but they can not summarize actual content data.

As a result, Cisco developed a proprietary standard for netflow data collection. In 1996 that was used for network accounting and network traffic analysis. This standard was quickly adopted by all major routers, and became formalized into an Internet Engineering Task Force Standard based on the most recent version of netflow. The proprietary “netflow” protocol has now become a standard wire-format network “flow” protocol known as IPFIX (Internet Protocol Flow Information Exchange).

# What is a Flow?

A flow is an aggregated record of packets

SiLK flows are defined by five unique attributes

- IP Protocol (TCP, UDP, ICMP, IPSec, etc)
- Source Address
- Destination Address
- Source Port
- Destination Port

SiLK flows are unidirectional



As we have already discussed, the flow record is an aggregated summary of packets observed by the flow sensor.

Each flow record combines all the packets that match the five unique packet attributes: IP protocol, source and destination address and source and destination port.

SiLK flows are unidirectional; that is, packets from a client to a server are **not** combined with the server to client packets because the source and destinations are swapped, so typical TCP sessions result in two flow records—one from client to server and another from server to client.



# What Does a Flow Know?

Each unique flow (tuple) has associated attributes

- Timing (start, stop)
- Volume (packets, bytes)
- TCP flags
- Collection location (sensor, traffic type)

Flows get flushed when they close

- Timeouts, TCP FIN/RST, router resources low



As individual packets are combined into flow records, specific attributes of the packets are recorded and saved with the flow record.


The flow records the time the first packet in the flow was seen, as well as the time the last packet was seen. Remember, individual packets are instantaneous and don't have time stamps or durations. A flow record counts the total number of packets combined into the record, and the total byte volume for all the packets that make up the flow.

A list of all the TCP flags seen in the flow is saved. The flow record also stores the location for where the flow was collected. Also, remember that flow records can not stay in memory forever, and they must be flushed to the collector. This happens when TCP sessions close and when the flows timeout.

## SiLK Flows are Half Duplex

For the TCP 3-way handshake, consider how flows are counted:

- Flow 1 is created when the sensor observes the first packet between hosts A and B.
- Flow 2 is created with the second packet. Swapped IPs means a new flow.
- The third ACK packet updates flow 1 since the source and destination addresses and ports match.



Let's take another look at the concept of unidirectional flow records.

Consider the TCP three-way handshake, where the client sends a SYN packet to the server, the server responds with a SYN-ACK packet, and the client finishes the handshake with an ACK packet.

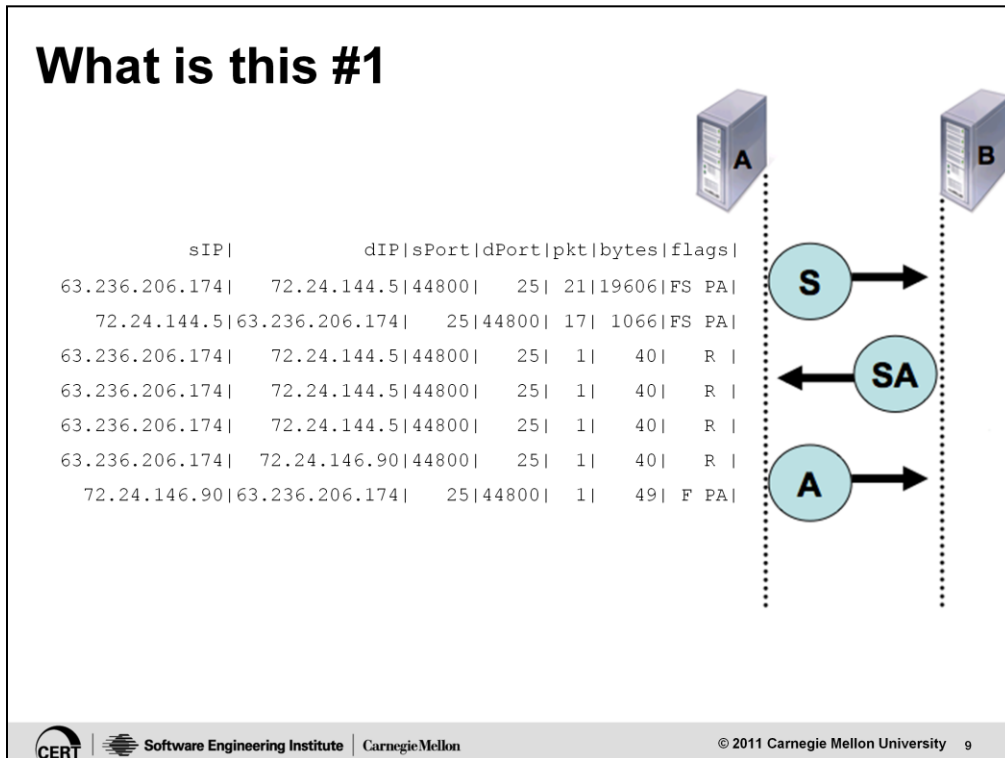
As the first SYN packet passes the sensor, our first flow record is created. The tuple includes the protocol (TCP), the source address A and the destination address B, and has the "flags" value set to SYN only.

B responds with a SYN-ACK packet. This creates a `_new_` flow record where the tuple has the source address B and the destination address A.

The third packet, the ACK packet, matches the first flow record, since it has source address A and destination B. This flow is updated with a packet count of two, the byte volume is increased, and the "flags" value is changed from "SYN only" to "SYN plus ACK". From here, you should be able to judge how additional packets in this TCP session will be added to these two existing flow records.



# What is this #1



At this point we will begin looking at some sample flow data. What is shown in this example? Can you tell which IP address appears to be the client and which is the server?

We see communications on port 25, which means this is probably email (SMTP) traffic. IP address 72.24.177.5 appears to be the server since it is hosting the SMTP port 25 service.

Since IP address 63.236.206.174 is using a high numbered ephemeral port, it is the SMTP client.

Take a look at the byte volume. There's a big difference between the client and the server—the client set about five times as much data as the server, a total of about 20k bytes. That makes sense if we consider 63.236.206.174 to be sending a message to 72.24.144.5.

What about all those RST packets? Well, they're all sent from the client to the server, and it seems like they're sent after the FIN connection teardown has completed. Although we can't tell exactly what is going on here, it's not uncommon to see spurious RST packets at the end of a TCP session, particularly with high-volume clients and servers.

Finally, did you notice that there are some extra records here that don't really belong? The last two records may have escaped your attention. They have the same client address and port, and they're SMTP traffic, but they actually point to a different SMTP server.

## What is this #2

sIP	dIP	pro	pkts	bytes	sTime
66.142.134.179	72.24.150.186	1	2	122	00:00:00.582
66.142.134.179	72.24.148.123	1	2	122	00:00:00.911
66.142.134.179	72.24.146.95	1	2	122	00:00:01.783
66.142.134.179	72.24.159.123	1	2	122	00:00:01.895
66.142.134.179	72.24.145.227	1	2	122	00:00:02.220
66.142.134.179	72.24.154.87	1	2	122	00:00:02.329
66.142.134.179	72.24.149.212	1	2	122	00:00:02.550
66.142.134.179	72.24.158.18	1	2	122	00:00:02.766
66.142.134.179	72.24.150.34	1	2	122	00:00:02.875
66.142.134.179	72.24.153.102	1	2	122	00:00:02.879
66.142.134.179	72.24.144.61	1	2	122	00:00:03.421
66.142.134.179	72.24.129.2	1	2	122	00:00:03.530
66.142.134.179	72.24.129.224	1	2	122	00:00:03.642
66.142.134.179	72.24.151.196	1	2	122	00:00:04.184

"WhatIsThis-2.txt" 15L, 871C 15,1 All

What do you notice in this set of flow records? Does this look like normal traffic to you?

*Taking a quick look at the IP addresses, you see that one address—the source address—remains the same, while the destination address changes. In fact, you'll notice that while the destination address changes, it only changes within a given range of addresses; the first two octets of the destination address remain the same while the second to octets change. It looks like the IP address 66.142.134.179 is scanning the 72.24 network. This is a scan.*

*Can you tell anything else about the scan? Let's take a closer look. The protocol is always 1. That means this is an ICMP scan. Each scan target receives two packets for a total of 122 bytes. That means each packet is probably half that, or 61 bytes. Each packet probably has a 40-byte IP header and four bytes of ICMP header, leaving room for 17 bytes of data.*

*Finally, look at the packet timing. We see about two to four packets per second. Is this a fast scan, or maybe even a denial-of-service attack? Not at that rate.*

*Considering everything we've observed, this just looks like routine scan activity.*

## What is this #3

sIP	dIP	sPort	dPort	pkt	flags	sTime
72.24.144.12	68.8.27.65	63126	80	7	FS PA	00:01:31.232
68.8.27.65	72.24.144.12	80	63126	5	FS PA	00:01:31.232
72.24.144.12	68.8.27.65	63277	80	8	FS PA	00:01:42.642
68.8.27.65	72.24.144.12	80	63277	8	FS PA	00:01:42.642
72.24.144.12	68.8.27.65	63330	80	7	FS PA	00:01:51.052
68.8.27.65	72.24.144.12	80	63330	5	FS PA	00:01:51.052
[pause]						
72.24.144.12	68.8.27.65	63707	80	8	FS PA	00:02:47.722
68.8.27.65	72.24.144.12	80	63707	8	FS PA	00:02:47.831
[pause]						
72.24.144.12	68.8.27.65	63957	80	8	FS PA	00:03:20.036
68.8.27.65	72.24.144.12	80	63957	8	FS PA	00:03:20.036
[pause]						
72.24.144.12	68.8.27.65	64504	80	8	FS PA	00:04:12.501
68.8.27.65	72.24.144.12	80	64504	8	FS PA	00:04:12.501
"WhatIsThis-3.txt" 16L, 909C						All

Next example. At first glance, this seems a bit more confusing than our last example. Is it normal traffic, or something to be concerned about?

*Our first clue in unraveling this traffic is to look at the ports. You should quickly recognize port 80, the port used for normal web or HTTP traffic. Every flow record on the list uses this port, so it's probably all HTTP traffic. Looking at the first record, since 80—the service port—is the destination port, that means that the destination IP address is the server. So we know that 68.8.27.65 is service HTTP traffic.*

*Similarly, looking at the high port, the ephemeral port, we identify the client as 72.24.144.12. Also notice that the ephemeral port increases with new connections, which is to be expected.*

*We see three separate connections from the client to the web server all within 20 seconds, and then three later connections. Each connection has normal flags—SYN, ACK and FIN—and a reasonable number of packets. All in all, the first set of six flows—three TCP connections—looks like a standard web page loading and then pulling down some included reference data like images or style sheets.*

*Finally, take a look at all the flows together. Everything's pretty similar, except for the timing. There's repetitive behavior here, but it's pretty slow—there's almost a minute between each connection—and the timing is not exact. This does not look like system behavior; instead, it looks like normal user interaction with a web site: click, read, click, read, click.*

*Overall, this looks like normal web browsing activity.*

## What is this #4

sIP	dIP	sPort	dPort	pkts	flags	sTime
72.24.129.20	82.80.30.150	80	1220	152	S PA	00:00:23.602
82.80.30.150	72.24.129.20	1220	80	90	SRPA	00:00:23.602
72.24.129.20	82.80.30.150	80	1221	1126	S PA	00:00:23.710
82.80.30.150	72.24.129.20	1221	80	413	SRPA	00:00:23.710
72.24.129.20	82.80.30.150	80	1223	63	S PA	00:00:26.341
82.80.30.150	72.24.129.20	1223	80	39	S PA	00:00:26.341
72.24.129.20	82.80.30.150	80	1224	8	S PA	00:00:26.883
82.80.30.150	72.24.129.20	1224	80	7	SRPA	00:00:26.883
82.80.30.150	72.24.129.20	1223	80	1	R A	00:01:33.068

"WhatIsThis-4.txt" 10L, 630C 10,1 All

CERT | Software Engineering Institute | Carnegie Mellon © 2011 Carnegie Mellon University 12

Here's another example. Once again, you see the port 80 traffic right away, and can assume that this is HTTP web browsing traffic. But there's something a bit different about this set of flows. What do you see?

*You probably noticed right away that these flows are larger than the ones in the previous example. Even though these sessions are large, this is still fairly normal for web traffic. This could come from large images, or another rich media type.*

*However, did you notice the RST packets in these flows? In this case, we see resets being sent from the client back to the server, which is not uncommon—the client probably no longer wants the content being delivered. In fact, it's also very common for high-volume web servers to send resets to clients as well, because it's much quicker and more efficient to shut down a TCP session with a RST packet than with a FIN exchange.*

*Finally, note how the last flow happens long after the actual data exchange in the TCP session. When the client timed out the TCP session it sent this reset packet. In cases like this you have to watch the ephemeral port to make sure you don't accidentally tie the last RST packet in with the wrong TCP session.*

# It's All a Matter of Timing


The flow buffer needs to be kept manageable

Inactivity timeout:

- If there is no activity within [30] thirty seconds, flush the flow

Active timeout:

- Flush all flows open for [30] thirty minutes



Software Engineering Institute | Carnegie Mellon

© 2011 Carnegie Mellon University 13

We briefly mentioned that flows have to time out to keep the flow buffer manageable. Let's take a closer look at what causes a flow to time out.

Keep in mind that a “flow” means the tuple—source and destination ports and addresses. An inactive timeout occurs when no packets are seen within the inactive window that match the tuple. Active timeouts occur when a flow has been open for a specific period; this ensures that long-lived flows are flushed off the sensor within a reasonable time.

Taking a look at this activity in a timeline, here you see individual packets with the same tuple, packets that are aggregated into a single flow. Since inter-packet arrival is relatively quick, the packets aggregate nicely. Then at some point, communications on the socket pauses for a period, which causes the flow to time out. This is an inactivity timeout.

After the pause, the conversation becomes active again, and data continues to flow for a long time. At a certain point—30 minutes in this case—the flow collector flushes the flow and starts a new flow. This is an active timeout.

Now let's take a look at flow records that have timeouts.

## What is this #5

sIP	dIP	sPort	dPort	pro	Pkt	bytes	sTime	dur
8.97.138.194	72.24.145.68	500	500	17	1	112	00:02:31	0.000
72.24.145.68	8.97.138.194	500	500	17	1	112	00:02:31	0.000
8.97.138.194	72.24.145.68	500	500	17	2	224	00:13:53	40.498
72.24.145.68	8.97.138.194	500	500	17	2	224	00:13:53	40.498
8.97.138.194	72.24.145.68	500	500	17	2	224	00:25:10	45.582
72.24.145.68	8.97.138.194	500	500	17	2	224	00:25:10	45.582
8.97.138.194	72.24.145.68	500	500	17	1	112	00:36:03	0.000
72.24.145.68	8.97.138.194	500	500	17	1	112	00:36:03	0.000
8.97.138.194	72.24.145.68	500	500	17	1	112	00:43:19	0.000
72.24.145.68	8.97.138.194	500	500	17	1	112	00:43:19	0.000
8.97.138.194	72.24.145.68	500	500	17	3	336	00:47:30	46.088
72.24.145.68	8.97.138.194	500	500	17	3	336	00:47:30	46.088
72.24.145.68	8.97.138.194	500	500	17	1	112	00:53:32	0.000
8.97.138.194	72.24.145.68	500	500	17	1	112	00:53:32	0.000
72.24.145.68	8.97.138.194	500	500	17	2	208	00:58:42	0.000
8.97.138.194	72.24.145.68	500	500	17	20	2232	00:58:49	90.095

"WhatIsThis-5.txt" 17L, 1207C 17,1 All

CERT | Software Engineering Institute | Carnegie Mellon © 2011 Carnegie Mellon University 14

Take a look at this flow data, see if you can tell how flow timeouts affect what you see. Do you recognize this traffic?

*It might seem odd that both the source and destination ports are the same for this protocol 17 (UDP) traffic. However, that's actually typical for VPN clients.*

*Look at the packet and byte counts, however. They're pretty small. Also look at the timing of these packets—there's a lot of time in between each one. What do you think this is?*

*To me this looks like a mostly quiet VPN, the only thing being sent on it are small keep-alive type packets. That would also explain why the last flow is larger than the rest—at some point, a small amount of data actually traveled across the VPN.*

## What is this #6

sIP	dIP	sPort	dPort	pkts	flg	sTime	dur
72.24.147.6	58.210.70.72	35282	22	29640	PA	00:00:11.361	1800.63
58.210.70.72	72.24.147.6	22	35282	29633	PA	00:00:11.911	1800.08
72.24.147.6	58.210.70.72	35282	22	30824	PA	00:26:23.092	1800.82
58.210.70.72	72.24.147.6	22	35282	30825	PA	00:26:23.092	1800.82
72.24.147.6	58.210.70.72	35282	22	29346	PA	00:56:24.020	1800.90
58.210.70.72	72.24.147.6	22	35282	29347	PA	00:56:24.020	1800.90
72.24.147.6	58.210.70.72	35282	22	31107	PA	01:00:10.783	1800.20
58.210.70.72	72.24.147.6	22	35282	31113	PA	01:00:11.301	1800.68
72.24.147.6	58.210.70.72	35282	22	29227	PA	01:26:25.036	1800.95
58.210.70.72	72.24.147.6	22	35282	29228	PA	01:26:25.036	1800.95
72.24.147.6	58.210.70.72	35282	22	30880	PA	01:56:26.096	1800.82
58.210.70.72	72.24.147.6	22	35282	30878	PA	01:56:26.096	1800.82
72.24.147.6	58.210.70.72	35282	22	30302	PA	02:00:11.301	1800.65
58.210.70.72	72.24.147.6	22	35282	30287	PA	02:00:11.843	1800.10
72.24.147.6	58.210.70.72	35282	22	31998	PA	02:26:27.028	1800.90
58.210.70.72	72.24.147.6	22	35282	31999	PA	02:26:27.028	1800.90

"WhatIsThis-6.txt" 17L, 1191C 17,1 All

What do we have here?

*The first thing you'll probably notice about this traffic is the service port, 22. This is the port for SSH so these all appear to be SSH client / server flows.*

*All the flags are PSH-ACK--no SYN, FIN or RST flags--so this definitely looks like one very long lived TCP session.*

*Now look closely at the times. Each flow is 1800 seconds long—30 minutes, which is a common setting for the active timeout. Then look at the start time for each flow—they're almost exactly 30 minutes apart. However, there are two sensors in play here: one rolls the flow over at :26 and :56 minutes after the hour; the other one is on the hour.*

*This is normal SSH traffic for a long-standing, high-volume traffic.*

# What is this #7

sIP	dIP	sPort	dPort	pkt	flags	sTime	dur
72.24.144.17	10.225.235.38	40395	80	45	S PA	01:59:34.81	1759.18
10.225.235.38	72.24.144.17	80	40395	44	S PA	01:59:34.81	1759.07
10.225.235.38	72.24.144.17	80	40395	40	PA	02:29:39.82	1797.62
72.24.144.17	10.225.235.38	40395	80	40	A	02:29:39.93	1797.51
10.225.235.38	72.24.144.17	80	40395	40	PA	03:00:23.46	1800.17
72.24.144.17	10.225.235.38	40395	80	40	A	03:00:23.57	1800.17
10.225.235.38	72.24.144.17	80	40395	40	PA	03:31:09.83	1797.52
72.24.144.17	10.225.235.38	40395	80	40	A	03:31:09.93	1797.52
10.225.235.38	72.24.144.17	80	40395	40	PA	04:01:53.42	1797.72
72.24.144.17	10.225.235.38	40395	80	40	A	04:01:53.51	1797.64
10.225.235.38	72.24.144.17	80	40395	35	RPA	04:32:37.18	1560.50
72.24.144.17	10.225.235.38	40395	80	34	A	04:32:37.29	1520.89
72.24.144.17	37.52.253.241	40395	80	13	FS PA	05:18:41.57	0.48
37.52.253.241	72.24.144.17	80	40395	18	FS PA	05:18:41.63	0.43
72.24.144.17	42.215.190.19	40395	80	9	FS PA	10:21:01.15	4.14
42.215.190.19	72.24.144.17	80	40395	6	FS PA	10:21:01.15	4.14
42.215.190.19	72.24.144.17	80	40395	1	A	10:21:05.29	0.00
72.24.144.17	10.46.227.72	40395	80	7	FS PA	12:21:24.36	0.22
10.46.227.72	72.24.144.17	80	40395	6	FS PA	12:21:24.47	0.22
72.24.144.17	18.113.57.14	40395	80	6	FS PA	13:39:43.67	0.11

"WhatIsThis-7.txt" 21L, 1492C 21,2 All

Here's the last example for you to take a look at. Don't worry that the IP addresses are in the 10.x network, this data has been anonymized.

*This is port 80 web traffic. The first thing that should stick out immediately is that this is a very long TCP session for web traffic. Normally HTTP sessions deliver a web page and then close; occasionally streaming media will be delivered over HTTP, but that still rarely lasts for more than a few minutes. In this case, we see a single session that's been open from 2:00 until almost 5:00.*

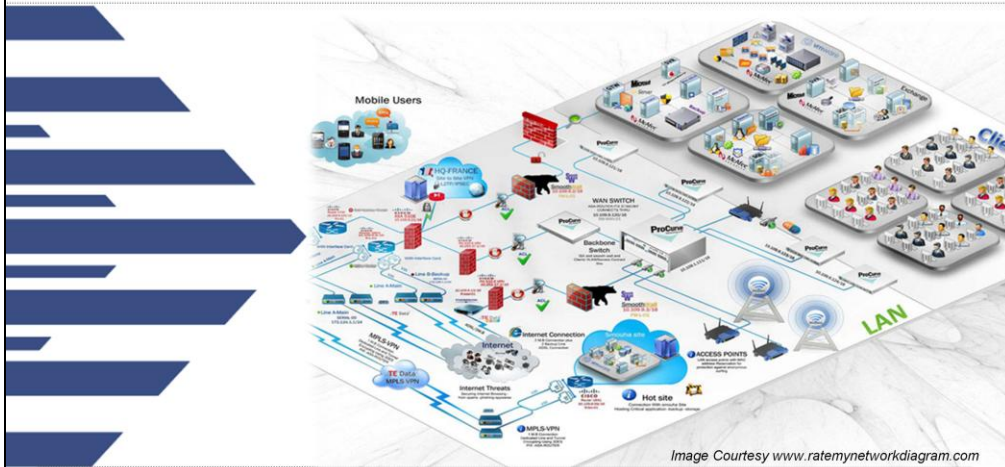
*Now take a look at the rate packets are being sent: 40 packets every 30 minutes. That's slow, certainly not fast enough for streaming media.*

*Finally, look at the last set of connections. They're all new TCP sessions that occur over the course of a few hours, yet they all use the same ephemeral port. This looks very suspicious.*

*Overall, this does not appear to be normal HTTP traffic. What is it? We can't tell for sure, just that it does not fit the typical profile for web traffic.*



# Why Flow?



That was interesting, but I can get all that from packet capture. Why bother with flow?

# Sensor Types

---

## Alerting

- Network IDS/IPS
- Host-based IDS

## Data Collection

- Firewall, router
- Metadata collection
- Full packet capture



Let's start by considering the different types of sensors that might already be in your environment.

In general, alerting sensors feed some type of workflow so you can take action when something “interesting” happens.

On the other hand, data collection sensors just pull information off the wire and store it until you need it to investigate an event.

Flow fits in the “Data Collection” sensor; it's similar to collecting firewall or router logs, metadata, or even full packet capture.

So if I have those data collection techniques already, why bother with flow?

## Why Flow? Reason #1

---

Complaint: There's no content!

Counter-complaint: That means privacy concerns are minimal. Many privacy-conscious organizations have legally approved flow collection.



The major complaint is that there's no content with flow, so there's no value.

That actually can be a significant strength if your collection efforts are hampered by privacy concerns. Even public service providers have been able to safely collect flow without violating privacy concerns.

## Why Flow? Reason #2

---

Complaint: There's still no content!

Counter-complaint: Because there's no content, you can store **lots** of data. Using SiLK:

- Fully saturated 100mb link = 50Gb / month
- With a good storage back-end, you can easily query 50Gb of data in 10-20 minutes.

*Details for SiLK storage and bandwidth usage can be found in the SiLK Provisioning Worksheet at <http://tools.netsa.cert.org/releases/SiLK-Provisioning-v204.xls>*



Software Engineering Institute | Carnegie Mellon

© 2011 Carnegie Mellon University 20

Another reason to use flow is to be able to store lots of historical data. Most flow compresses nicely and any traffic that doesn't (mostly DNS and scanning) can be easily trimmed from the long term data store if necessary.

## Why Flow? Reason #3

Complaint: No, really, there's *still* no content!!

Counter-complaint: Content is a distraction for many problems.

- For many problems you need an index to get to the right content in your huge data store
- For other problems you need to search gobs and gobs of history
- Encrypted? So what!

That's the niche where flow lives



Software Engineering Institute | Carnegie Mellon

© 2011 Carnegie Mellon University 21

Believe it or not, there are a lot of problems you face that really don't need content. Or problems that you can't ask to your entire full content data set; problems where you need a pointer to focus your query. That's where flow fits in nicely: you can do a broad sweep of your data set with flow, and then do a fine-grained query against your full content data store.

Since flow doesn't use content, it also has lots of advantages when analyzing encrypted traffic. Remember that VPN tunnel from earlier that wasn't being used?

## Got a Question? Flow Can Help

- What's on my network?
- What happened before the event?
- Where are policy violations occurring?
- What are the most popular web sites?
- How much volume would be reduced with a blacklist?
- Do my users browse to known infected web servers?
- Do I have a spammer on my network?
- When did my web server stop responding to queries?
- Who uses my public DNS server?



Software Engineering Institute | Carnegie Mellon

© 2011 Carnegie Mellon University 22

We've spent time looking at where flow records come from, what data they contain, and seen how network traffic appears in flow data. Hopefully you now can begin to consider many different ways that flow can be used as part of your overall security suite. It can act as a tool to support forensic and historical analysis, it can monitor your infrastructure, help you gain situational awareness about your network, and locate some classes of security events.



# The Big Picture

---

Q. Why Flow?

A. So you can get beyond signature-based detection

That's ***Network Situational Awareness***:

- Perceive
- Comprehend
- Predict




Software Engineering Institute | Carnegie Mellon

© 2011 Carnegie Mellon University 23

So at the end of the day, flow is your bridge between your data collection sensors and your alerting workflow. It's the place where you do things that are generally NOT signature-driven (although that's starting to change), the place where you really begin to understand what's on your network.

For more information about situational awareness, see Mica Endsley's paper "Toward a Theory of Situation Awareness in Dynamic Systems", *Human Factors*, 1995, 37(1), p32-64.



Argus is another popular open source flow collection and analysis package. Commercial products exist as well. There's even a module you can bind to a Linux Ethernet adapter to generate flow.

For the rest of this class we'll be looking specifically at the SiLK toolkit. It's very mature (over a decade of large installations), scalable (used by large organizations and service providers), actively maintained and it's open source.

Let's dig into some of the tools in the SiLK toolkit and some things you can do with them.



# The Flow Suite

---

## Generate

- Aggregates packets into flow records
- Sends flow records somewhere

## Store

- Gathers the flow records into a storage architecture

## Analyze

- Retrieves and processes stored flow record

*A Pluggable Architecture*  
*We'll focus on SiLK but others exist*



A complete flow implementation requires three things: something that generates flow records from packets on the wire, something that turns the wire-format flow (or netflow) records into a disk format and something that allows you to analyze the on-disk flow records.

Commercial solutions based on netflow often will combine two or more of these things together, but it's important to keep them all in mind as you plan for interoperability between sensors and security event management systems.

# Flow Generation



## Hardware Option

- Router generates and forwards Cisco Netflow
- Pro: easy
- Con: router's main job is routing, not accounting

## Dedicated Sensor Option

- Tap or span session to a dedicated appliance
- SiLK option is to use YAF (Yet Another Flowmeter) to generate IPFIX and forward to the store



Most commercial routers can generate flow, so it might be an easy matter to simply turn on flow accounting and have records generated for you. However, remember that a router's primary job is to route packets, not to count flow. As the router becomes overburdened—particularly in a denial-of-service condition—the router will stop generating flow.

One alternative to reduce the load of generating flow records is to use sampling. For sampled flows, the router will only count every 100<sup>th</sup> or 1000<sup>th</sup> packet. This works fine for determining large-scale network trends, but it tends to mask many of the network behavior details that make flow valuable.

Another alternative is to use a dedicated sensor that exists specifically to monitor traffic and generate flow. The SiLK suite includes "Yet Another Flowmeter" (YAF), a service designed specifically to create flow from a network tap.

## Flow Storage



### Flow record on-disk format

- Highly tuned, very compact
- Most analyses are I/O constrained
- SiLK uses binary flat file storage
  - Limited only by file system space, IO performance

### Alternatives:

- SIM/SEIM or an RDBMS
- Generally don't scale well



The flow records are sent to a collector in Cisco Netflow or IPFIX format, and the collector is responsible for storing the records on disk.

SiLK uses a very tightly compressed proprietary data structure for storing flow on disk. Since analysis tasks are often bound by the amount of time it takes to retrieve data from disk, it's important to keep the disk footprint small.

It's also possible to store flow records directly in your security information manager or in a relational database. That's a great way to correlate flow with other events, but it generally will not scale well in larger organizations.

# Flow Analysis



Generally an I/O problem

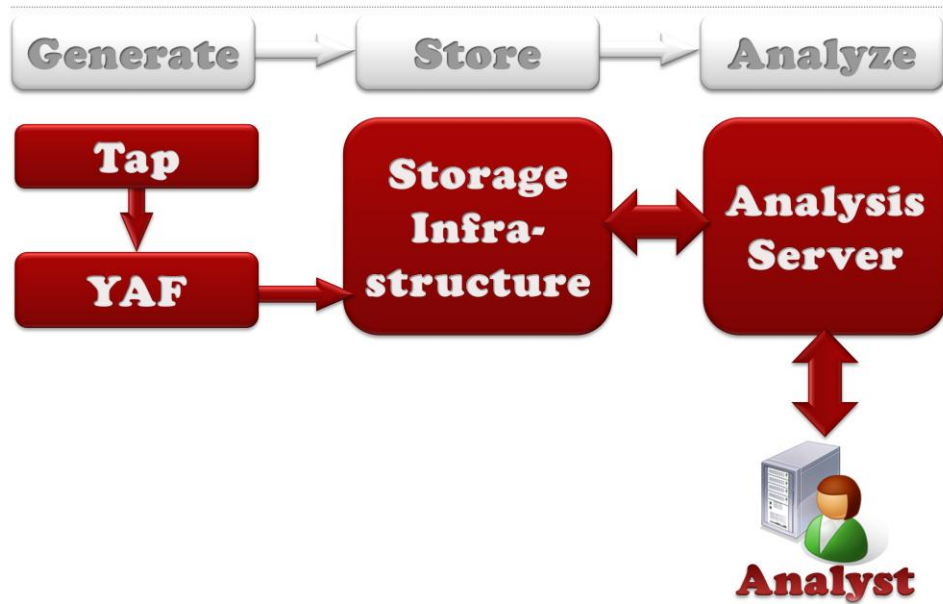
- SiLK's "rwfilter" pulls records from the store
- Many supporting tools manipulate binary flow records
  - `rwcut`, `rwuniq`, `rwsort`, `rwcat`, etc.
  - Keep it binary as long as possible
- Specialized tools
  - IP sets, arbitrary IP labeling, statistics, top-N analysis



That leads us to analysis of the stored flow data. If you've used SiLK as your storage back end, you can take advantage of a number of specialized analysis tools in the suite. The tools are designed to keep flow data in the compressed binary form for as much of the analysis as possible—although you may be tempted to print out flow records as text and bring them into Excel, you'll be much more successful in large scale analysis by keeping the data in binary.

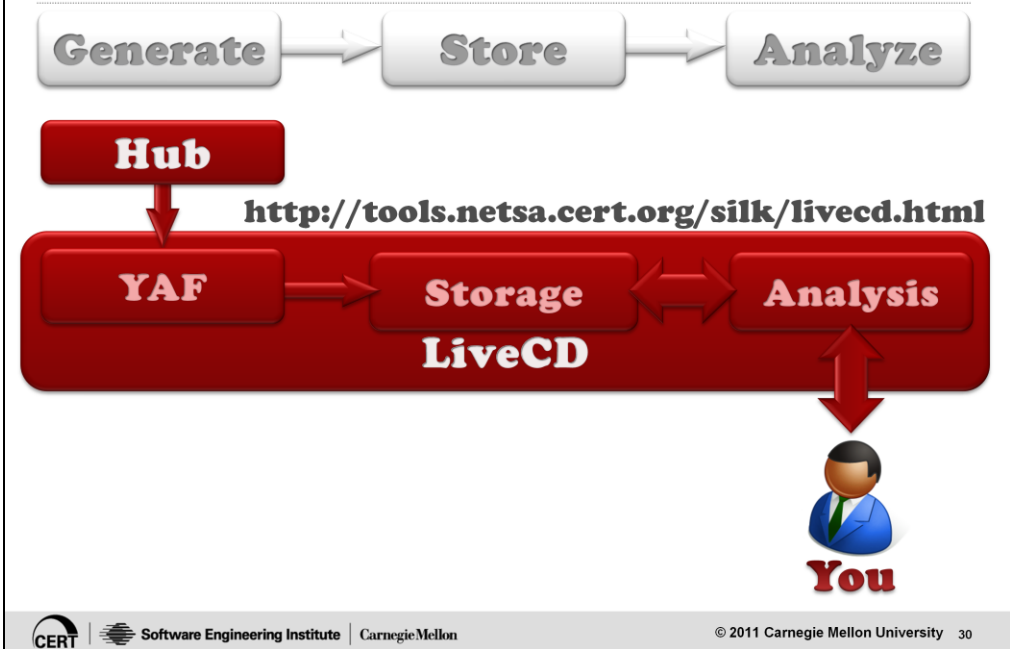
Although the analysis tools are mature, you should still think of them as middle-tier tools. They're generally all Unix command-line driven and produce text output.

## Typical Installation



The typical large organization will have three distinct tiers for its flow collection and analysis setup. A dedicated tap will feed a flow sensor, the sensor will send data to a distributed storage infrastructure, and the analysis server will access data in the storage cloud.

# Minimal Installation



If you're a tinkerer, you can bypass all that complexity with the SiLK live CD. This CD has all the moving parts loaded into a single virtual image, just about all you have to do is to define a network interface to begin sensing.



# Let's Get Practical

*Some results of flow analysis*

  Software Engineering Institute | Carnegie Mellon

© 2011 Carnegie Mellon University 31

Now let's get back into some of the analysis that can be done once the flow infrastructure is set up.

## Example #1: Poison Ivy

```

sIP|sPort|      dIP|dPort|pro|pkts|bytes|ini|sess | sTime+msec| dur+msec|
120.16.86.207| 1637|235.185.173.131| 3460| 6| 3| 144|S | S |12:08:50.487| 1.373|
235.185.173.131| 3460| 120.16.86.207| 1637| 6| 1| 40| | R A|12:08:50.722| 0.000|
235.185.173.131| 3460| 120.16.86.207| 1637| 6| 1| 40| | R A|12:08:51.316| 0.000|
235.185.173.131| 3460| 120.16.86.207| 1637| 6| 1| 40| | R A|12:08:52.105| 0.000|
120.16.86.207| 1638|235.185.173.131| 3465| 6| 3| 144|S | S |12:09:52.093| 1.444|
235.185.173.131| 3465| 120.16.86.207| 1638| 6| 1| 40| | R A|12:09:52.313| 0.000|
235.185.173.131| 3465| 120.16.86.207| 1638| 6| 1| 40| | R A|12:09:53.020| 0.000|
235.185.173.131| 3465| 120.16.86.207| 1638| 6| 1| 40| | R A|12:09:53.771| 0.000|
120.16.86.207| 1639|235.185.173.131| 3666| 6| 3| 144|S | S |12:10:53.774| 1.670|
235.185.173.131| 3666| 120.16.86.207| 1639| 6| 1| 40| | R A|12:10:54.025| 0.000|
235.185.173.131| 3666| 120.16.86.207| 1639| 6| 1| 40| | R A|12:10:54.787| 0.000|
235.185.173.131| 3666| 120.16.86.207| 1639| 6| 1| 40| | R A|12:10:55.661| 0.000|
120.16.86.207| 1642|235.185.173.131| 3460| 6| 3| 144|S | S |12:14:37.755| 9.004|
120.16.86.207| 1643|235.185.173.131| 3465| 6| 24| 2040|S |F PA|12:15:58.791| 513.393|
235.185.173.131| 3465| 120.16.86.207| 1643| 6| 40| 16809|S A| PA|12:15:59.026| 513.398|
120.16.86.207| 1649|235.185.173.131| 3460| 6| 26| 3688|S | PA|12:24:32.184| 95.440|
235.185.173.131| 3460| 120.16.86.207| 1649| 6| 38| 23113|S A| PA|12:24:32.424| 95.550|
235.185.173.131| 3460| 120.16.86.207| 1649| 6| 94| 6526| PA| PA|12:28:51.561| 1789.016|
120.16.86.207| 1649|235.185.173.131| 3460| 6| 55| 6664| PA| PA|12:28:51.674| 1788.529|
235.185.173.131| 3460| 120.16.86.207| 1649| 6| 81| 5207| PA| PA|12:59:25.227| 1800.622|
120.16.86.207| 1649|235.185.173.131| 3460| 6| 41| 3608| PA| PA|12:59:25.265| 1800.619|
...

```



This showed up one day while watching the default poison Ivy port (3460). Can you explain what's going on here based on what you know about poison Ivy? What characteristics might help you decide if this is botnet command/control traffic or just routine traffic on a non-standard port?

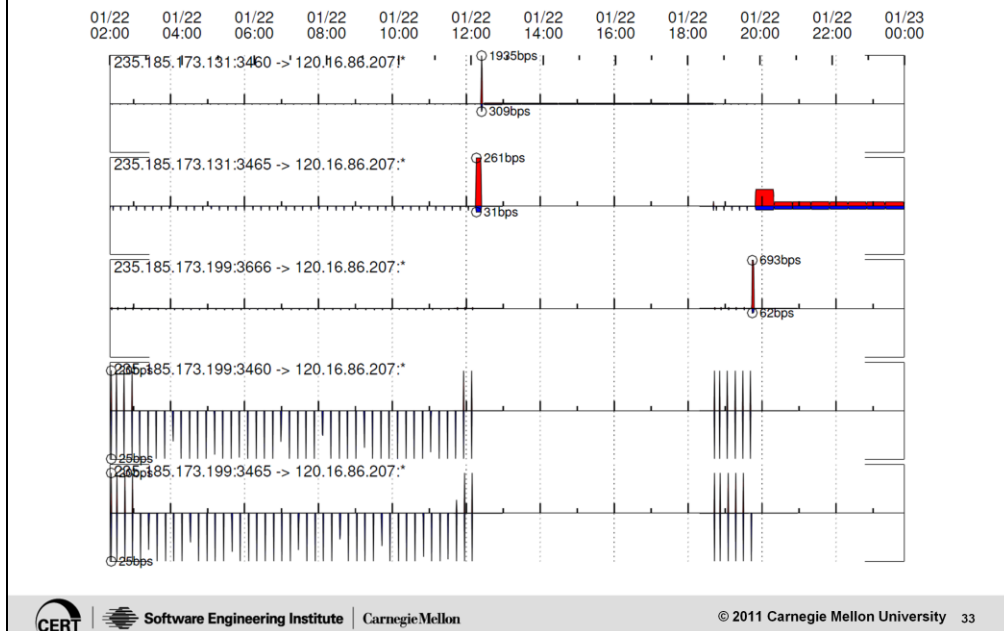
*120.16.86.207 might have an implant that's trying to connect to 235.185.173.131, but it looks like we're missing the first SYN packet in this screen shot. It seems to be retrying the connection attempt every 60 seconds or so.*

*There's a point where the connection succeeds, there's a lot of data sent from the server (182) back to the client (117), and then the connection stays open for over an hour.*

Even with those features, it's difficult to get a feel for what's going on here. Let's take a look at how the activity proceeded over a longer time period.



## Example #1: A Picture



Here's an example of essentially the same traffic plotted over a 24-hour time span. Now can you spot the beacon? And what do you think happened at 12:15? How about 18:30 and 19:45?

Originally we thought there was a beacon every 60 seconds; what's actually happening are three retries 60-seconds apart, and then a rest for 30 minutes.

This plot was generated using the "stripplot.py" python script with the command

```
./stripplot -v -count=6 -fields=* controller.rw
```


This script uses SiLK, gnuplot and ghostscript to generate simple time-series flow visualizations as a .pdf file. It's available here:

<https://tools.netsa.cert.org/confluence/display/tt/Strip+Plots>

And there's a presentation on how to use it here:

[https://tools.netsa.cert.org/confluence/download/attachments/10027010/Faber\\_StripPlots.pdf?version=1&modificationDate=1263239683000](https://tools.netsa.cert.org/confluence/download/attachments/10027010/Faber_StripPlots.pdf?version=1&modificationDate=1263239683000)

## Example #2: 30b UDP Packets

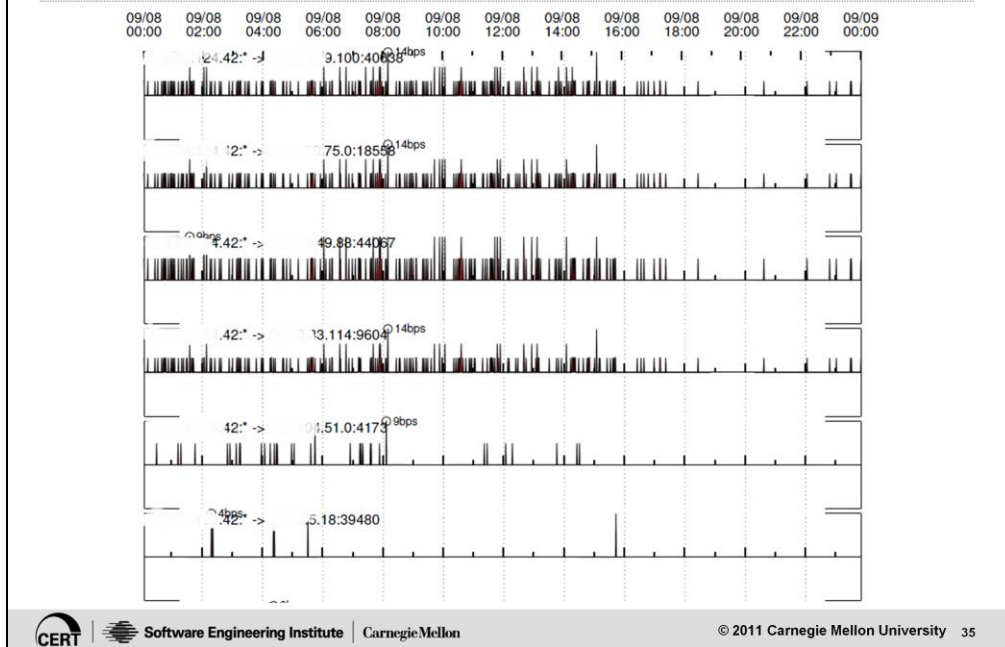


One day I started looking at 30 byte UDP packets transiting my network. I found tons of them coming from all over the place, destined for a small number of IP addresses in my network. Many were heading to addresses that ended in .0 or .4, but never for a live host. They were coming from hundreds of thousands of sources. Roughly 90% came from China, maybe another 8% from the Pacific Rim, and the remaining mostly from the United States.

Look at the third strip: that's what I found really confusing. All at once tens of thousands of clients almost simultaneously started generating the same type of traffic.

Any idea what this might be?


## Example #1: 30b UDP Packets (Single Host)



Here's a view of the packets generated by a single IP address. Note how the source will hit up four destinations at a time. Whenever we see traffic like this we're tempted to think the worst—that it's some type of attack—but it just doesn't have the hallmarks of a distributed denial-of-service attack or anything else particularly malicious.

So what is this? I'm still not sure, but I believe it's some type of broken peer-to-peer software; likely peer-to-peer IP Television which is widely used in China. But I never found out for sure.

## Example #3: IP Watchlist




Can you spot the beacon? Any idea what this is?


This is actually a view of traffic destined to known Conficker sinkholes. Review what you know about conficker—can you explain what's going on in these plots?

## Example #4: Beacons


### FTP Anti-Virus Updates



### SMTP Resends



### HTTPs strangeness



All beacons are not malicious. Here are some examples of two benign beacons. The first one is for anti-virus updates; in this case it looks like the update check is about every 10 minutes. This might represent updates for multiple clients behind a single NAT'd IP address. Can you tell when the clients are shut off for the evening?

The second strip shows some aggressive SMTP retries. There are a number of scenarios where email can get caught in a resend loop for a few days, can you think of any?


The final plot is just plain odd HTTPS traffic. It's relatively slow and flat, it looks kind of like a stale HTTPS channel...but it shouldn't stay active for a whole day!

## Routine Behavior: Streaming Media

Streaming video: browse, watch, browse watch, done.



Streaming Audio: Browse, browse, listen



Software Engineering Institute | Carnegie Mellon

© 2011 Carnegie Mellon University 38

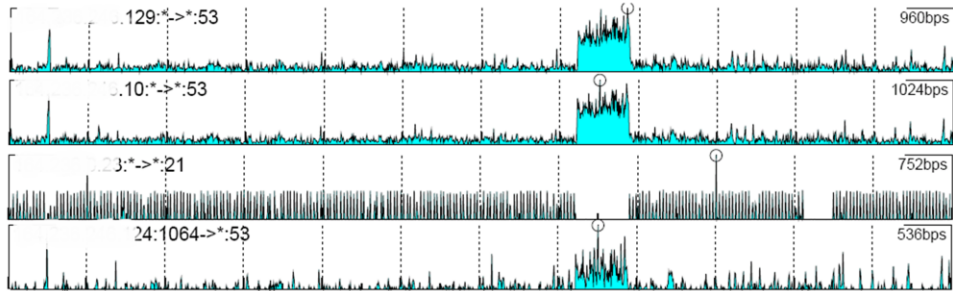
Here's a comparison of traffic for streaming media, typically the largest consumer of bandwidth on your network.

Note how streaming video—youtube in this case—shows a browse, watch, browse, watch pattern and the videos are all rather short.

Contrast that with the streaming audio example in the second strip. Bandwidth is smaller, but you see a steady utilization that covers a few hours.

# A Network Anomaly

Different hosts and protocols on the same network:



Software Engineering Institute | Carnegie Mellon

© 2011 Carnegie Mellon University 39

And what's going on here? I'm not quite sure, but there's definitely an anomaly. Probably there was a network outage that resulted in an increased volume of DNS retry attempts.



## Better Tools, More Resources

  Software Engineering Institute | Carnegie Mellon

© 2011 Carnegie Mellon University 40

Let's take a brief look at some of the things that are on-the-way for the SiLK tool suite.



# YAF Improvements

## Application labeling

- “app” field tags HTTP, SSH, Teredo, 20+ protocols

## Metadata Capture

- DNS (becomes your pDNS store)
- HTTP server, user-agent, referrer, etc
- X.509 certificates (in progress)



Recall that YAF is the tool that counts packets and creates flow records. YAF includes an “application labeling” feature that allows you to assign an integer to a flow based on deep packet inspection. This is very useful to identify protocols—for example, a relatively simple regular expression can identify HTTP traffic running on any port and assign it an application label of “80”.

Since YAF is inspecting packets, it has been re-engineered to pull deeper metadata out of the conversation. Since IPFIX allows for more generic passing of information about a flow, YAF can pull some metadata off the wire and add that to the flow record. Eventually YAF will become a passive DNS collector, HTTP monitor and pull metadata out of any number of protocols.

# YAF Inspector (almost live!)

The screenshot shows a web browser window titled "yInspector: YAFDPI - Windows Internet Explorer" with the URL "http://10.65.48.2/yInspector/". The page content includes the title "yInspector" and the tagline "DPI - you know you want to look". Below the title is a navigation menu with "Home", "Query", "Top 10", and "Yinzers", along with a search bar labeled "Search..." and a "Go" button. The main content area has an "Introduction" section and an "About YAF" section. Below the screenshot is a flow diagram with three buttons: "Generate", "Store", and "Analyze", connected by arrows. The "Store" button is highlighted in red.

**Generate** → **Store** → **Analyze**

CERT | Software Engineering Institute | Carnegie Mellon | © 2011 Carnegie Mellon University 42

Although IPFIX allows collection of metadata, a mature storage solution for the data does not yet exist. yInspector seeks to fill that gap with flexible storage and query of IPFIX metadata.

# YAF Inspector Reports and Queries

The screenshot displays two data tables from the YAF Inspector. The first table, titled 'Top 10 Referrers', lists the following referer URLs:

Referrer	Total
http://www.linux.com/archive/feature/120	
http://twitter.com/	
http://www.ustream.tv/socialstream/6951	
http://www.cnn.com/	
http://www.ustream.tv/channel/one-track	
http://www.wired.com/wiredscience/2010	
http://www.google.com/search	

The second table, titled 'Top 10 User Agent Strings', lists the following user agents and their counts:

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

Below the reports is a workflow diagram consisting of three buttons: 'Generate', 'Store', and 'Analyze'. The 'Store' button is highlighted in red, and arrows indicate a flow from 'Generate' to 'Store' to 'Analyze'.

[intentionally blank]

# Pipeline

Process flows as they are collected

- Moves some analysis into the storage phase
- Alert data: watchlists, beacon detection
- Situational Awareness data: inventories, IPv6 usage, current activity



Pipeline is a new and significant improvement that will allow flow analysis to feed directly into the alerting flow in real time—but it won't require you to throw all your flow records into your security correlation engine.

Analysis pipeline takes a streaming feed of events, maintains state, performs analytics and feeds an alerting library.

# iSiLK (when you need something clicky)

rwfilter

```
--type=out  
--start=2011/04/12:00  
--end=2011/04/12:12  
--dport=80,8080,443  
--pass=tmp.rw
```



**Generate**

**Store**

**Analyze**



Software Engineering Institute | Carnegie Mellon

© 2011 Carnegie Mellon University 45


As mentioned at the beginning, all of the SiLK analysis tools run from the Unix command-line. iSiLK is a graphical user interface that will generate all the commands for you, display the output in a spreadsheet and produce simple visualizations.



Software Engineering Institute | Carnegie Mellon

© 2008 Carnegie Mellon University

# Prism



Software Engineering Institute | Carnegie Mellon

© 2011 Carnegie Mellon University 46

Prism is a situational awareness tool used to carve traffic up into flexible bins and display long term trends.



Software Engineering Institute | Carnegie Mellon

© 2008 Carnegie Mellon University

## Other Analysis Tools

---

Rayon

pySiLK; the NetSA python toolkit

IPA: the IP Address Annotation System



Stay tuned, there's much more in the works!

## Where to go for more

---

The entire SiLK training class is on-line!

Also available at <http://tools.netsa.cert.org>:

- LiveCD
- Software
- Documentation
- Wiki / Tooltips
- Scripts

FloCon 2012 (mid-January)




Software Engineering Institute | Carnegie Mellon

© 2011 Carnegie Mellon University 48

For the SiLK training class go to <https://tools.netsa.cert.org/> and select the “Online Training” link in the left column. This brings up the public Virtual Training Environment course on “Using SiLK for Network Traffic Analysis”. This five-hour course covers most of the SiLK analysis tools and techniques and includes a hands-on lab.



You can find the SiLK LiveCD at <https://tools.netsa.cert.org/silk/livecd.html> or by clicking on “SiLK” under the Projects section and selecting LiveCD from the menu.





## Questions?

<http://tools.netsa.cert.org>  
[netsa-help@cert.org](mailto:netsa-help@cert.org)  
[sfaber@cert.org](mailto:sfaber@cert.org)

  Software Engineering Institute | Carnegie Mellon © 2011 Carnegie Mellon University 49

[intentionally blank]

---

#### NO WARRANTY

THIS MATERIAL OF CARNEGIE MELLON UNIVERSITY AND ITS SOFTWARE ENGINEERING INSTITUTE IS FURNISHED ON AN "AS-IS" BASIS. CARNEGIE MELLON UNIVERSITY MAKES NO WARRANTIES OF ANY KIND, EITHER EXPRESSED OR IMPLIED, AS TO ANY MATTER INCLUDING, BUT NOT LIMITED TO, WARRANTY OF FITNESS FOR PURPOSE OR MERCHANTABILITY, EXCLUSIVITY, OR RESULTS OBTAINED FROM USE OF THE MATERIAL. CARNEGIE MELLON UNIVERSITY DOES NOT MAKE ANY WARRANTY OF ANY KIND WITH RESPECT TO FREEDOM FROM PATENT, TRADEMARK, OR COPYRIGHT INFRINGEMENT.

Use of any trademarks in this presentation is not intended in any way to infringe on the rights of the trademark holder.

This Presentation may be reproduced in its entirety, without modification, and freely distributed in written or electronic form without requesting formal permission. Permission is required for any other use. Requests for permission should be directed to the Software Engineering Institute at [permission@sei.cmu.edu](mailto:permission@sei.cmu.edu).

This work was created in the performance of Federal Government Contract Number FA8721-05-C-0003 with Carnegie Mellon University for the operation of the Software Engineering Institute, a federally funded research and development center. The Government of the United States has a royalty-free government-purpose license to use, duplicate, or disclose the work, in whole or in part and in any manner, and to have or permit others to do so, for government purposes pursuant to the copyright license under the clause at 252.227-7013.



Software Engineering Institute | Carnegie Mellon

© 2011 Carnegie Mellon University 50