



Network Monitoring and Management

NetFlow Overview



These materials are licensed under the Creative Commons *Attribution-Noncommercial 3.0 Unported* license (<http://creativecommons.org/licenses/by-nc/3.0/>)


Agenda

- Netflow
 - What it is and how it works
 - Uses and applications
- Generating and exporting flow records
- Nfdump and Nfsen
 - Architecture
 - Usage
- Lab

What is a Network Flow?

- A set of related packets
- In practice this means: packets which belong to the same transport connection.
e.g.
 - TCP, same src IP, src port, dst IP, dst port
 - UDP, same src IP, src port, dst IP, dst port
 - Some tools consider "bidirectional flows", i.e. A->B and B->A as part of the same flow
- [http://en.wikipedia.org/wiki/Traffic_flow_\(computer_networking\)](http://en.wikipedia.org/wiki/Traffic_flow_(computer_networking))

Simple flows



 = Packet belonging to flow X

 = Packet belonging to flow Y

Cisco IOS Definition of a Flow

Unidirectional sequence of packets sharing:

1. Source IP address
2. Destination IP address
3. Source port for UDP or TCP, 0 for other protocols
4. Destination port for UDP or TCP, type and code for ICMP, or 0 for other protocols
5. IP protocol
6. Ingress interface (SNMP ifIndex)
7. IP Type of Service

IOS: which of these six packets are in the same flows?

	<i>Src IP</i>	<i>Dst IP</i>	<i>Protocol</i>	<i>Src Port</i>	<i>Dst Port</i>
A	1.2.3.4	5.6.7.8	6 (TCP)	4001	80
B	5.6.7.8	1.2.3.4	6 (TCP)	80	4001
C	1.2.3.4	5.6.7.8	6 (TCP)	4002	80
D	1.2.3.4	5.6.7.8	6 (TCP)	4001	80
E	1.2.3.4	8.8.8.8	17 (UDP)	65432	53
F	8.8.8.8	1.2.3.4	17 (UDP)	53	65432

Flow Accounting

- A summary of all the packets seen in a flow (so far)
 - Flow identification: protocol, src/dst IP/port...
 - Packet count
 - Byte count
 - Start and end times
 - Maybe additional info, e.g. AS numbers, netmasks
- Records the volume of traffic, not the content

Uses and Applications

- You can answer questions like:
 - Which user / department has been uploading / downloading the most?
 - Which are the most commonly-used protocols on my network?
 - Which devices are sending the most SMTP traffic, and to where?
- Identification of anomalies and attacks
- More fine-grained visualisation (graphing) than can be done at the interface level

Working with flows


1. Configure device (e.g. router) to generate flow accounting records
2. Export the flows from the device to a collector
 - Configure protocol version and destination
3. Receive the flows, write them to disk
4. Analyse the flows

There are many tools available, both free and commercial

Where to generate flow records

1. On a router or other network device
 - If the device supports it
 - No additional hardware required
 - Might have some impact on performance
2. Passive collector (usually a Unix host)
 - Receives a copy of every packet and generates flows
 - Requires a mirror port
 - Resource intensive

Router Collection




Flow collector
stores exported flows from router.

Router Collection

- With this method, all flows through the router can be observed
- However, more work for the router in processing and exporting the flows
- Optionally, one can choose on which interfaces netflow collection is needed and not activate it on others
- If there is a router on each LAN, netflow can be activated on those routers to reduce the load on the core router

Passive Monitor Collection



Passive Collector

- Examples: `softflowd` (Linux/BSD), `pfflowd` (BSD), `ng_netflow` (BSD)
- The collector will only see flows passing through the network point it is connected on
- However, this does relieve the router from processing netflows and exporting them
- Useful on links with only one entry into the network, or where only flows from one section of the network are needed
- Can deploy in conjunction with an IDS

A thought:

- Your network probably already has a device which is keeping track of IP addresses and port numbers of traffic flowing through it.

What is it?

Flow Export Protocols

- Cisco Netflow, different versions
 - v5: widely deployed
 - v9: newer, extensible, includes IPv6 support
- IPFIX: IETF standard, based on Netflow v9
- sFlow: Sampling-based, commonly found on switches
- jFlow: Juniper
- We will concentrate on Netflow, but many tools support multiple protocols

Cisco Netflow

- Unidirectional flows
- IPv4 unicast and multicast
 - (IPv6 in Netflow v9)
- Flows exported via UDP
 - Choose a port. No particular standard, although 2055 and 9996 are commonly used
- Supported on IOS, ASA and CatOS platforms - but with different implementations

Cisco IOS Configuration

- Configured on each input interface
- Define the version
- Define the IP address and port of the collector (where to send the flows)
- Optionally enable aggregation tables
- Optionally configure flow timeout and main (v5) flow table size
- Optionally configure sample rate

Cisco Command Summary (1)

- Enable CEF (done by default)

- `ip cef`

- Enable flow on each interface

- `ip route cache flow`

- OR

- `ip flow ingress`

- `ip flow egress`

- View flows

- `show ip cache flow`

- `show ip flow top-talkers`

Cisco Command Summary (2)

```
ip flow-top-talkers
  top 10
  sort-by bytes
```

```
gw-169-223-2-0#sh ip flow top-talkers
```

SrcIf	SrcIPaddress	DstIf	DstIPaddress	Pr	SrcP	DstP	Bytes
Fa0/1	169.223.2.2	Fa0/0	169.223.11.33	06	0050	0B64	3444K
Fa0/1	169.223.2.2	Fa0/0	169.223.11.33	06	0050	0B12	3181K
Fa0/0	169.223.11.33	Fa0/1	169.223.2.2	06	0B12	0050	56K
Fa0/0	169.223.11.33	Fa0/1	169.223.2.2	06	0B64	0050	55K
Fa0/1	169.223.2.2	Local	169.223.2.1	01	0000	0303	18K
Fa0/1	169.223.2.130	Fa0/0	64.18.197.134	06	9C45	0050	15K
Fa0/1	169.223.2.130	Fa0/0	64.18.197.134	06	9C44	0050	12K
Fa0/0	213.144.138.195	Fa0/1	169.223.2.130	06	01BB	DC31	7167
Fa0/0	169.223.15.102	Fa0/1	169.223.2.2	06	C917	0016	2736
Fa0/1	169.223.2.2	Local	169.223.2.1	06	DB27	0016	2304

```
10 of 10 top talkers shown. 49 flows processed.
```

Cisco Command Summary (3)

- Exporting Flows to a collector

```
ip flow-export version 5 [origin-as|peer-as]  
ip flow-export destination <x.x.x.x> <udp-port>
```

- Origin-AS* will include the origin AS Number in the flow, while *Peer-AS* will only include the AS Number of the peering neighbor. Applies only if your router talks BGP.
- Exporting aggregated flows (if required)


```
ip flow-aggregation cache as|prefix|dest|source|proto  
enabled  
export destination x.x.x.x <udp-port>
```

Questions?

Collecting flows: nfdump

- free and open source
- *nfcapd* listens for incoming flow records and writes them to disk (flat files)
 - typically starts a new file every 5 minutes
- *nfdump* reads the files and turns them into human-readable output
- nfdump has command-line options to filter and aggregate the flows

nfdump architecture

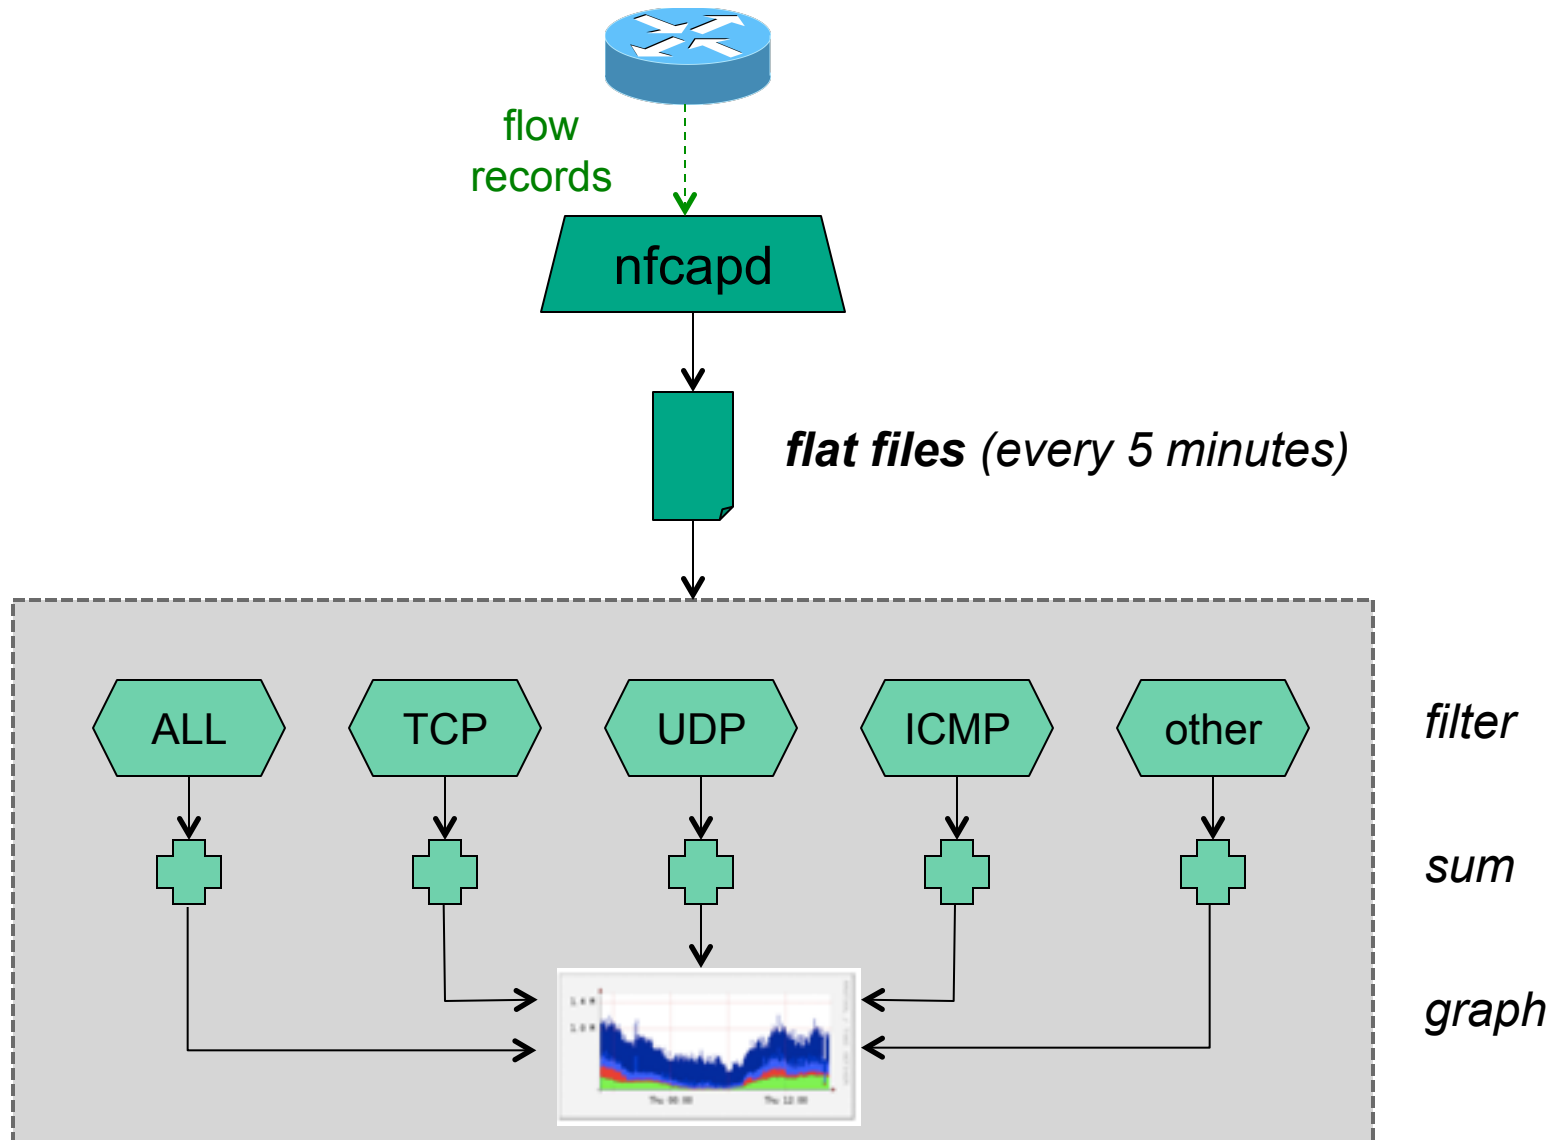


Date flow start	Duration	Proto	Src IP Addr:Port		Dst IP Addr:Port	Packets	Bytes	Flows
2013-04-18 13:35:23.353	1482.000	UDP	10.10.0.119:55555	->	190.83.150.177:54597	8683	445259	1
2013-04-18 13:35:23.353	1482.000	UDP	190.83.150.177:54597	->	10.10.0.119:55555	8012	11.1 M	1
2013-04-18 13:48:21.353	704.000	TCP	196.38.180.96:6112	->	10.10.0.119:62099	83	20326	1
2013-04-18 13:48:21.353	704.000	TCP	10.10.0.119:62099	->	196.38.180.96:6112	105	5085	1

Analysing flows: nfsen

- Companion to nfdump
- Web GUI
- Creates RRD graphs of traffic totals
- Lets you zoom in to a time of interest and do nfdump analysis
- Manages nfcapd instances for you
 - It can run multiple nfcapd instances for listening to flows from multiple routers
- Plugins available, e.g. port tracker, surfmap

nfsen architecture



nfusen: points to note

- Every 5 minutes nfcapd starts a new file, and nfusen processes the previous one
- Hence each graph point covers 5 minutes
- The graph shows you the *total* of selected traffic in that 5-minute period
- To get more detailed information on the individual flows in that period, the GUI lets you drill down using nfdump

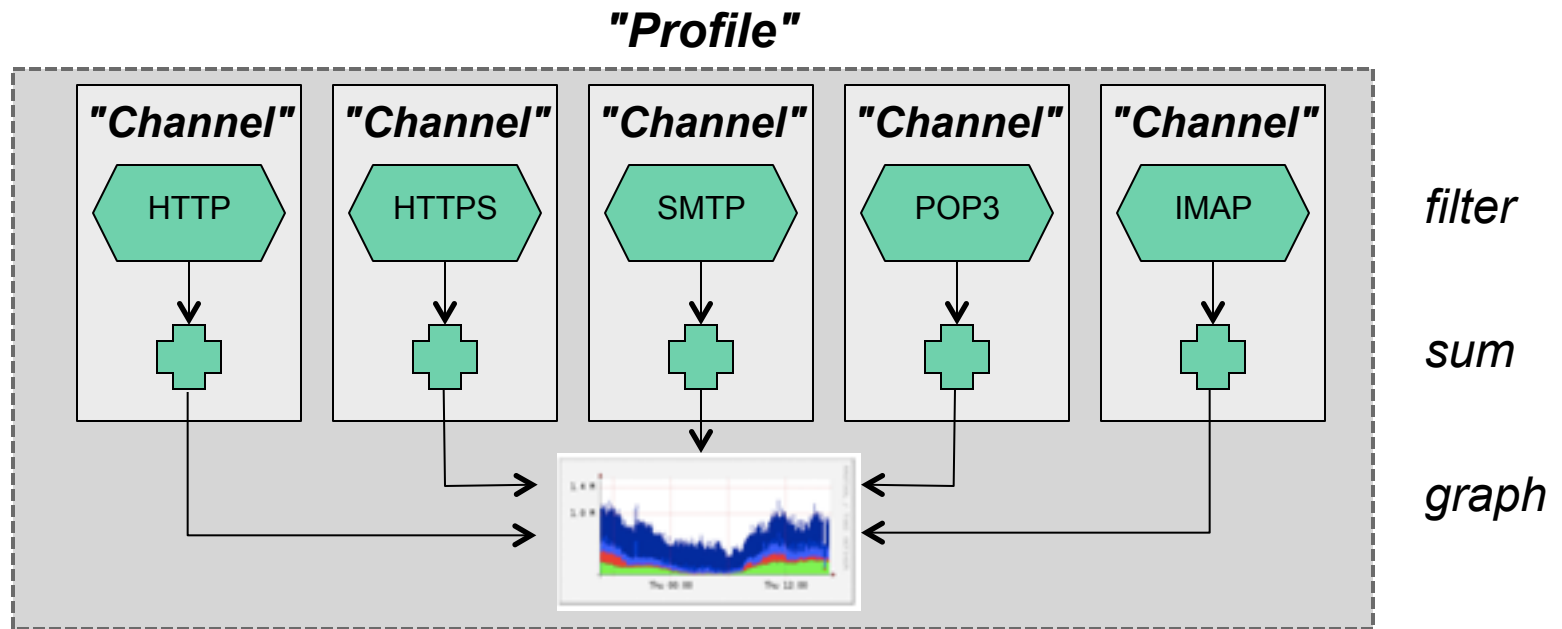
Demonstration

- Using nfsen to find biggest users of bandwidth

Profiles and Channels

- A "channel" identifies a type of traffic to graph, and a "profile" is a collection of channels which can be shown together
- You can create your own profiles and channels, and hence graphs. e.g.
 - Total HTTP, HTTPS, SMTP traffic (etc)
 - Traffic to and from the Science department
 - ...
- Use filters to define the traffic of interest

Profiles and Channels



References – Tools

- nfdump and nfsen:
<http://nfdump.sourceforge.net/>
<http://nfsen.sourceforge.net/>
<http://nfsen-plugins.sourceforge.net/>
- pmacct and pmgraph:
<http://www.pmacct.net/>
<http://www.aplivate.org/pmgraph/>
- flow-tools:
<http://www.splintered.net/sw/flow-tools>

References – Further Info

- Wikipedia:
<http://en.wikipedia.org/wiki/Netflow>
- IETF standards effort:
<http://www.ietf.org/html.charters/ipfix-charter.html>
- Abilene NetFlow page
<http://abilene-netflow.itec.oar.net/>
- Cisco Centric Open Source Community
<http://cosi-nms.sourceforge.net/related.html>
- Cisco NetFlow Collector User Guide
http://www.cisco.com/en/US/docs/net_mgmt/netflow_collection_engine/6.0/tier_one/user/guide/user.html

The end

(Additional reference materials follow)

Filter examples

any	<i>all traffic</i>
proto tcp	<i>only TCP traffic</i>
dst host 1.2.3.4	<i>only traffic to 1.2.3.4</i>
dst net 10.10.1.0/24	<i>only traffic to that range</i>
not dst net 10.10.1.0/24	<i>only traffic <u>not</u> to that range</i>
proto tcp and src port 80	<i>only TCP with source port 80</i>
dst net 10.10.1.0/24 or dst net 10.10.2.0/24	<i>only traffic to those nets</i>
dst net 10.10.1.0/24 and proto tcp and src port 80	<i>only HTTP response traffic to that net</i>
(dst net 10.10.1.0/24 or dst net 10.10.2.0/24) and proto tcp and src port 80	<i>...more complex combinations possible</i>

Flows and Applications: More Examples

Uses for NetFlow

- Problem identification / solving
 - Traffic classification
 - DoS Traceback (some slides by Danny McPherson)
- Traffic Analysis and Engineering
 - Inter-AS traffic analysis
 - Reporting on application proxies
- Accounting (or billing)
 - Cross verification from other sources
 - Can cross-check with SNMP data

Detect Anomalous Events: SQL “Slammer” Worm*

peakflow | DoS


Recent Anomalies: Anomaly 125772 : Detailed Statistics

11:51:49 EST 27 Jan 2003

Status Topology Ongoing Recent Dark IP Admin About

Anomaly 125772 Detailed Statistics

ID	Importance	Severity	Duration	Direction
125772	High	958.2% of 3.40 Kpps	09h 06m 47s	Outgoing



Affected Network Elements

Router net8 1.2.3.4

	Triggering	Expected	Difference	Max
Bitrate	71.69 Mbps	2.34 Mbps	69.35 Mbps	105.26 Mbp
Packet Rate	22.20 Kpps	712 pps	21.49 Kpps	32.58 Kpps

Summary | Source Addresses | Destination Addresses | Source Ports | Destination Ports | Protocols | Output Interfaces | Input Interfaces | Generate Filter

Summary of all Data Snapshots Collected:

	Bytes	Packets	Bytes/Pkt	bps
	308.01 GB	762,849,500	404 B	76.05 Mbps



Summary | Source Addresses | Destination Addresses | Source Ports | Destination Ports | Protocols | Output Interfaces | Input Interfaces | Generate Filter

Source Addresses

Network / Mask	Bytes	Packets	Bytes/Pkt	bps
192.168.20.217/32	168.22 GB	416,436,800	404 B	41.54 Mbps
192.168.18.187/32	139.53 GB	345,372,800	404 B	34.45 Mbps

Summary | Source Addresses | Destination Addresses | Source Ports | Destination Ports | Protocols | Output Interfaces | Input Interfaces | Generate Filter

Destination Addresses



Flow-based Detection (cont)*

Once baselines are built anomalous activity can be detected

- Pure **rate-based** (pps or bps) anomalies may be legitimate or malicious
- Many **misuse** attacks can be immediately recognized, even **without** baselines (e.g., TCP SYN or RST floods)
- **Signatures** can also be defined to identify “interesting” transactional data (e.g., proto udp and port 1434 and 404 octets(376 payload) == slammer!)
- Temporal compound signatures can be defined to detect with higher precision

Flow-based Commercial Tools...*


Anomaly 150228

Get Report: [PDF](#) [XML](#)

ID	Importance	Duration	Start Time	Direction	Type	Resource
150228	High 130.0% of 2 Kpps	17 mins	03:34, Aug 16	Incoming	Bandwidth (Profiled)	Microsoft 207.46.0.0/16 windowsupdate.com

Traffic Characterization

Sources	204.38.130.0/24
	204.38.130.192/26
	1024 - 1791
Destination	207.46.248.234/32
	80 (http)
Protocols	tcp (6)
TCP Flags	S (0x02)




Affected Network Elements

	Importance	Expected	Observed bps		Observed pps		
		pps	Max	Mean	Max	Mean	
Router nl-chi3 198.110.131.125	High						
Interface 67 at-1/1/0.14 pvc to WMU		26	832 K	563.1 K	2.6 K	1.7 K	Details

Anomaly Comments



Commercial Detection: A Large Scale DOS Attack



Accounting

Flow based accounting can be a good supplement to SNMP based accounting.

Thu Nov 17 00:50:00 2011 Bits/s any protocol



Cisco Netflow Versions

NetFlow Version 1

- Key fields: Source/Destination IP, Source/Destination Port, IP Protocol, ToS, Input interface.
- Accounting: Packets, Octets, Start/End time, Output interface
- Other: Bitwise OR of TCP flags.
- Does not have sequence numbers – no way to detect lost flows
- Obsolete

NetFlow Versions 2-4

- Cisco internal
- Were never released

NetFlow v5

- Key fields: Source/Destination IP, Source/Destination Port, IP Protocol, ToS, Input interface.
- Accounting: Packets, Octets, Start/End time, Output interface.
- Other: Bitwise OR of TCP flags, Source/Destination AS and IP Mask.
- Packet format adds sequence numbers for detecting lost exports.
- IPv4 only

NetFlow v8

- Aggregated v5 flows.
- Not all flow types available on all equipments
- Much less data to post process, but loses fine granularity of v5 – no IP addresses.

NetFlow v9

- IPv6 support
- Additional fields like MPLS labels
- Builds on earlier versions
- Periodically sends "template" packet, all flow data fields reference the template