

Netflow

6/12/07

Overview


- Why use netflow?
- What is a flow?
- Deploying Netflow
- Performance Impact

Caveats


- Netflow is a brand name like Kleenex. It was developed by Cisco
- Juniper uses the term cflowd for flow export
- The term “netflow” will be used generically
- NETS, as of this presentation, only exports flow data from Junipers
- Application configuration is beyond the scope of this presentation

Why Use Netflow?

- Enterprise
 - protocol distribution
 - monitor users/applications
 - identify malicious traffic
- Service Provider
 - peering
 - Identify malicious traffic
 - planning
 - traffic engineering
 - accounting/billing



Daily graph



The solution: Netflow

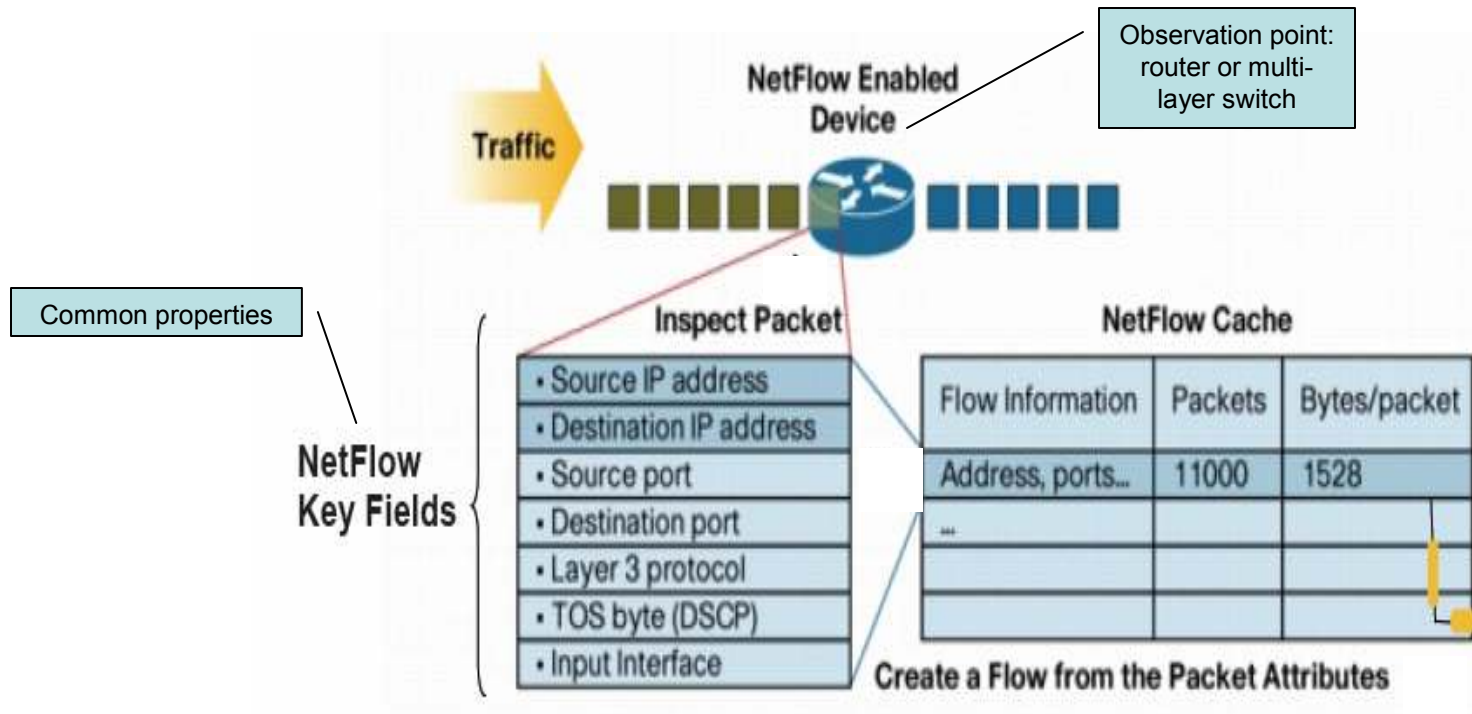
- Developed and patented by Cisco in 1996
- Classifies network traffic into “flows” by inspecting packets at layers 2 – 4.
- Currently on standards track - IPFIX
- “Flows” can be analyzed to provide network and security monitoring, network planning, traffic analysis and IP accounting.

What is a flow?

As defined by the IPFIX WG


- A flow is defined as a set of IP packets passing an observation point in the network during a certain time interval. All packets belonging to a particular flow have a set of common properties...A packet is defined to belong to a flow if it completely satisfies all the defined properties of the flow.

Flow Example



1. A flow is unidirectional
2. Defined by inspecting a packet's key fields (common properties) and identifying the values
3. If the set of key field values is unique create a flow record or cache entry

Flow example: part 2



Netflow Export Versions

- Multiple netflow export options (v1, v5, v7, v8, v9).
- Each version defines their own “common properties” and export packet format
- Most common is v5
- v7 specific to 6500s (now obsolete)
- v8 allows aggregation
- v9 (aka flexible netflow) used as basis for upcoming IPFIX (IP flow information export) standard. User defined.


Other Exported fields

- In addition to the key fields, the following non-key fields are exported in netflow v5:
 - *source and destination ASs*
 - *source and destination IP subnet masks*
 - *IP address of next-hop router*
 - *TCP flags*
 - *output interface*
- A list of all exported fields can be found here:
<http://www.cisco.com/univercd/cc/td/doc/cisintwk/intsolns/netflsol/nfwhite.htm#wp1095965>

Deploying Netflow

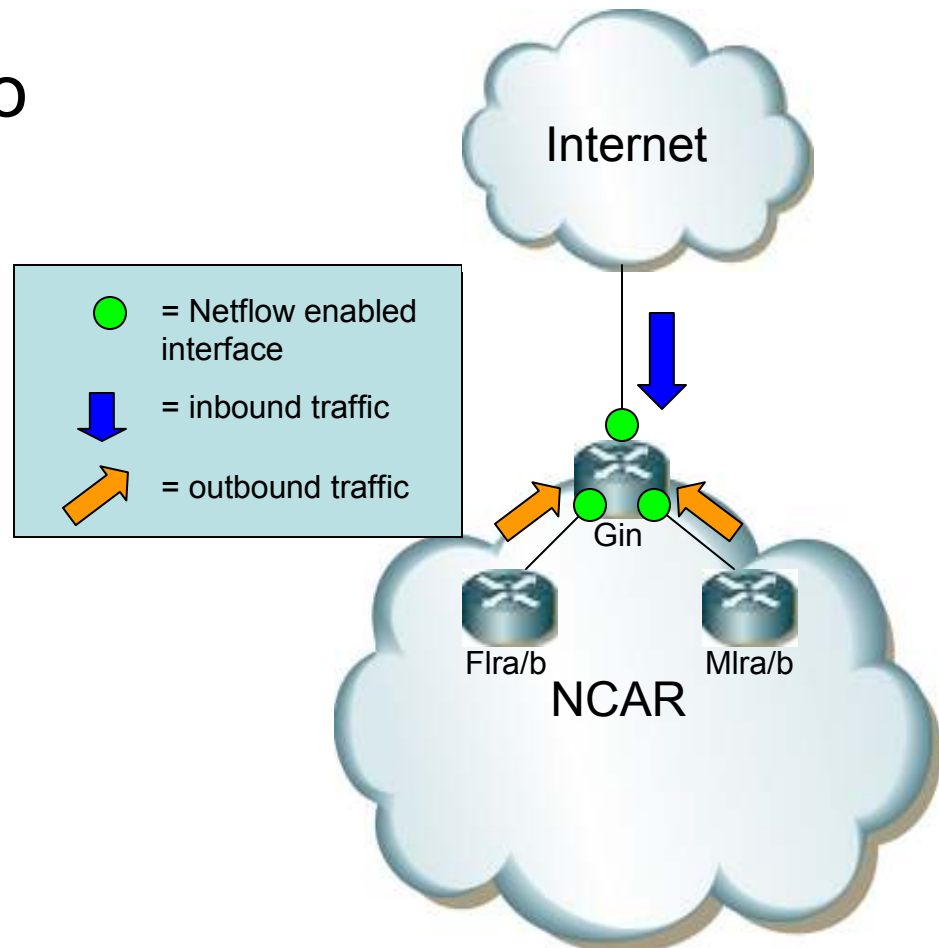
- Overview – Typical Deployment
- Basic steps to Deploy Netflow
 - Determine which routers/interfaces to enable netflow
 - Configure Routers
 - Juniper
 - Cisco
 - Setup netflow collectors
 - Choose and configure an application

Overview - Typical Deployment



Determine which routers/interfaces to enable netflow

- Enable netflow on selected interfaces to capture all inbound/outbound traffic
- Netflow only enabled inbound on an interface
- Avoid double counting!!

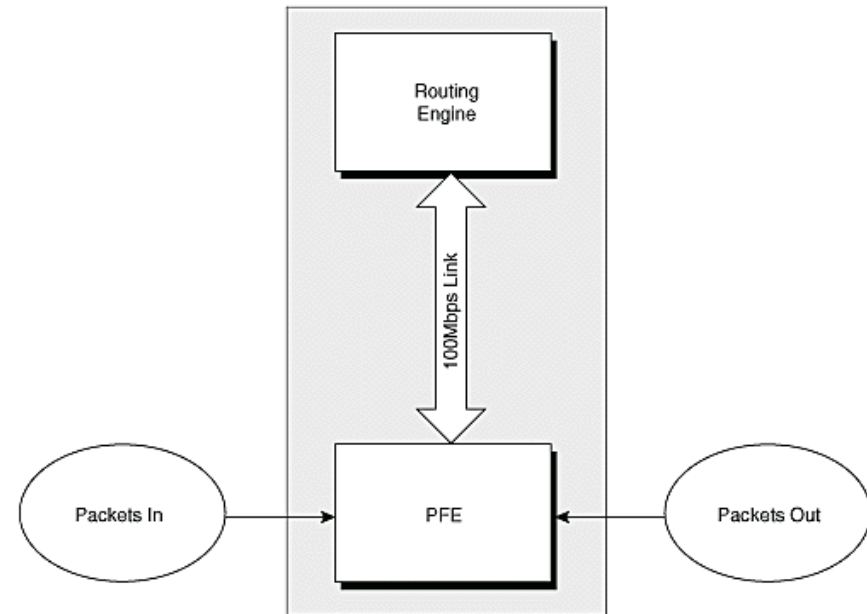


Configure Routers/Switches

- Juniper
 - Create a firewall filter
 - Apply the filter to an interface
 - Configure sampling
 - Configure netflow data export
- Cisco
 - Enable netflow/sampled netflow
 - Enable interfaces for sampled netflow (un-sampled netflow automatically applied to all interfaces)
 - Configure netflow data export
 - Configure netflow cache timers

Configure Router – Juniper M20 Architecture

- Two ways to configure netflow – basic or advanced.
Advanced would require an ASP2 PIC – list price \$35k.
NETS uses the basic config.
- Basic netflow config – netflow runs as a unix process on RE.
Limited by 8000pps across the RE to Forwarding Board 100 Mbps Ethernet link (not a factor if using ASP2)
- No support for v9 with basic netflow



Configure Routers – Juniper M20

Create & Apply Firewall

- Create Firewall using the key term “sample”.

- Apply Firewall to interface

```
firewall {
  filter NCARInput {
    term CatchAll {
      then {
        sample;
        accept;
      }
    }
  }
}

interfaces {
  ae0 {
    unit 303 {
      description "----- link with mlra and mlrb";
      vlan-id x;
      family inet {
        mtu 9000;
        filter {
          input NCARInput;
          output NCAROutput;
        }
        address x.x.x.x/29;
      }
    }
  }
}
```

Configure Routers – Juniper M20

Sampling and Export


- Configure Sampling:
 - Sampling Rate = $(\text{run-length} + 1) / \text{rate}$
 - max-packets-per-second - the maximum number of packets to be sampled
- Configure output
 - Cflowd – IP address of flow collector
 - Port – set the UDP port the collector is listening for netflow data.
 - Version – set the flow export version
 - No-local-dump – do NOT write flow files to disk before exporting
 - Autonomous-system-type [peer|orgin] – write the specified AS number in flow export file.

```
Sampling & Export config from Gin:

forwarding-options {
  sampling {
    input {
      family inet {
        rate 100;
        run-length 1;
        max-packets-per-second 1000;
      }
    }
  }
  output {
    cflowd x.x.x.x {
      port xxxx;
      version 5;
      no-local-dump;
      autonomous-system-type peer;
    }
  }
}
```

Configure Router/Switch – Cisco 6500 Architecture

- 1st packet in flow sent to RP, software switched.
- All other packets in flow switched in hardware.
- Must enable netflow in hardware and software.
- Flows stored in flow cache on router



Configure Router/Switch – Cisco 6500

Enable Netflow/Sampled Netflow

- Global config, enable sampled netflow in hardware, 1 out of 64 packets.
- Configure the flow mask, “interface full” required for sampled netflow in hardware
- Per interface, enable sampled netflow in hardware.
- Per interface, enable sampled netflow in software

```
C6500(config)# mls sampling time-based  
64
```

```
C6500(config)# mls flow ip interface-full
```

```
C6500(config)# interface Gx/x  
C6500(config-if)# mls netflow sampling
```

```
C6500(config-if)# ip flow ingress
```

Configure Router/Switch – Cisco 6500

Netflow Data Export (NDE)

- Set the NDE version
- Populate the following additional fields in the NDE packets
 - Egress interface SNMP index
 - Source-autonomous system number
 - Destination-autonomous system number
 - IP address of the next-hop router
- Configure NDE export destination (ip address) of collector

```
C6500(config)# mls nde sender
version 5
C6500(config)# mls nde interface

C6500(config)# ip flow-export
destination x.x.x.x
```


Netflow Collector

- NETS uses a freeware version called flow-tools written by Mark Fullmer
- Collects and aggregates data from multiple routers and writes it to a file for processing by a netflow application.
- Typical configuration looks like this:
 - `/usr/bin/flow-capture -w /var/netflow/flows 0/0/9996 -z0 -V5 -E1G -n 287 -N 0`
 - `-w` Store flows in `/var/netflow/flows`
 - `0/0/9996` Accept data from any source sending to port 9996
 - `-z0` Compression level, 0 = no compression
 - `-V5` PDU version
 - `-E1G` Retain max # of flow files up to 1 Gb
 - `-n 287` Number of times per day a flow file will be created (5min)
 - `-N 0` Nesting level for storing flow files

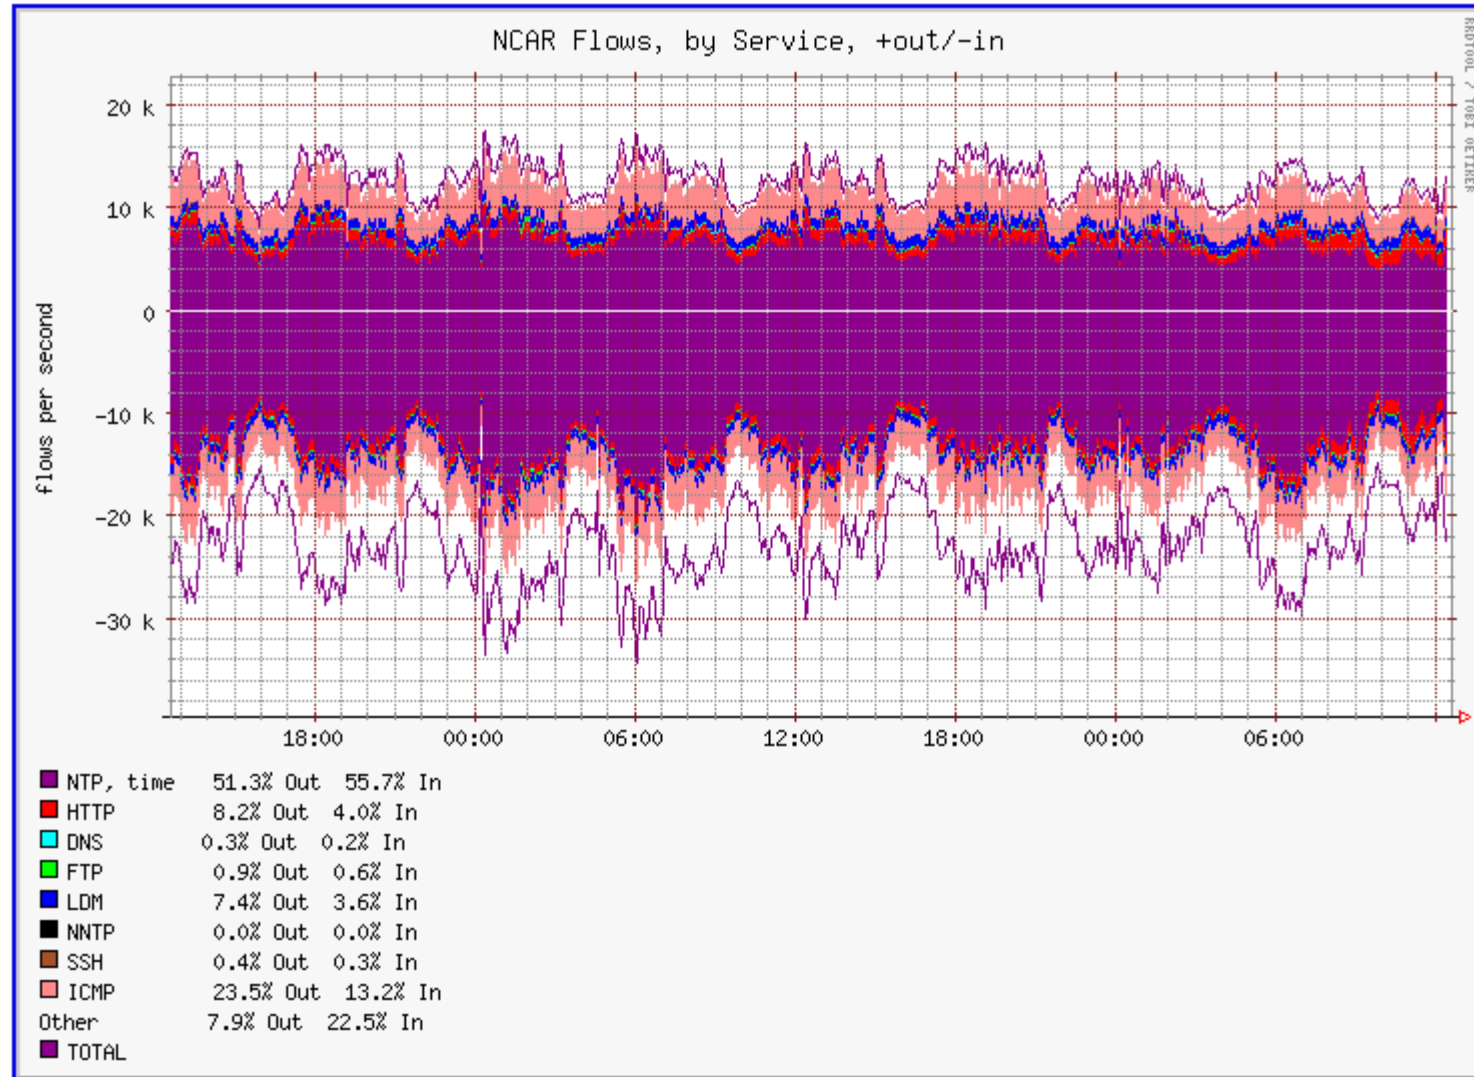
Netflow Applications

- NETS uses Flowscan developed by Dave Plonka at UW.
 - Report Modules
 - CampusIO – shows traffic in/out through a peering point or border router.
 - SubnetIO – shows traffic in/out per defined subnet
 - TopN – reports the top talkers
- Cisco CLI

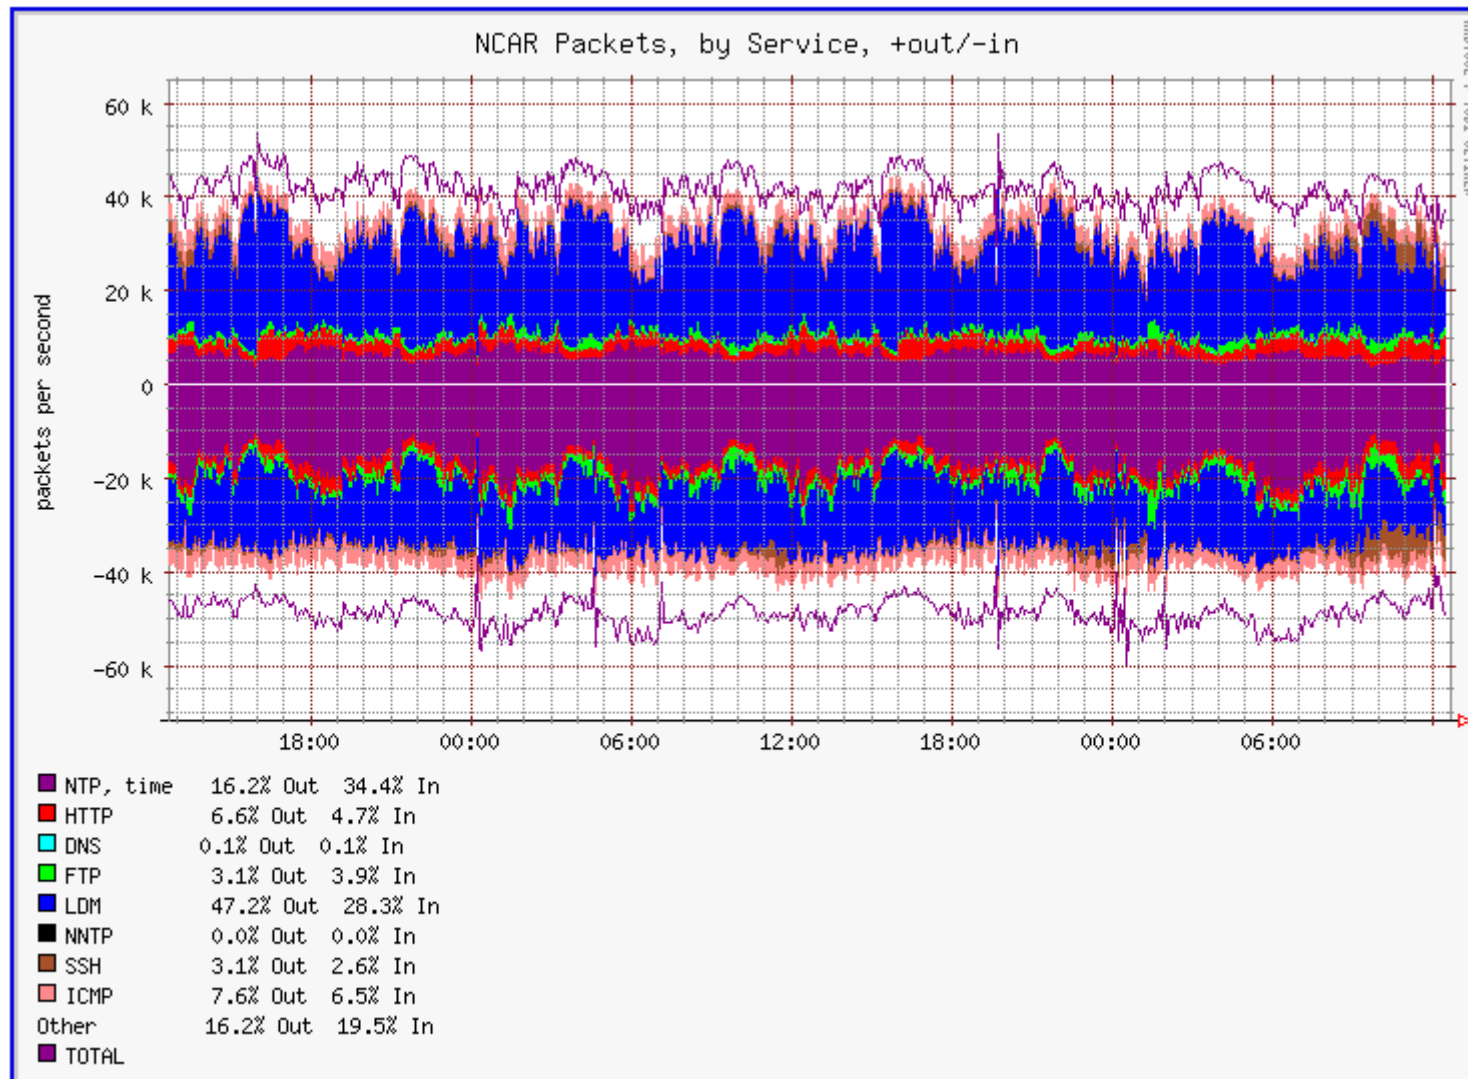
Flowscan – NCAR BPS



Flowscan – NCAR Flows



Flowscan – NCAR Packets



Flowscan – NCAR TopTalkers

Top 10 128.117.0.0/16 hosts by **bits out**
 based on five minute flow sample from gin
 dated Mon Jun 11 12:55:01 2007

rank	src Address	bits/sec in	bits/sec out	pkts/sec in	pkts/sec out	flows/sec in	flows/sec out
#1	idd.unidata.ucar.edu	3.3 M (3.1%)	142.0 M (69.5%)	7.2 k (35.4%)	14.0 k (55.7%)	543.7 (18.9%)	726.0 (24%)
#2	bv1103ge.ucar.edu	230.6 k (0.2%)	9.5 M (4.7%)	512.3 (2.5%)	803.0 (3.2%)	10.0 (0.3%)	6.0 (0%)
#3	huron.scd.ucar.edu	1.4 M (1.3%)	7.5 M (3.7%)	502.3 (2.5%)	791.3 (3.1%)	21.7 (0.8%)	24.7 (1%)
#4	cosmic-io.cosmic.ucar.edu	3.0 M (2.8%)	6.2 M (3.0%)	670.3 (3.3%)	672.3 (2.7%)	231.0 (8.0%)	271.7 (9%)
#5	ultrazone.ucar.edu	243.4 k (0.2%)	4.9 M (2.4%)	469.7 (2.3%)	577.3 (2.3%)	152.0 (5.3%)	157.3 (5%)
#6	yakov.unidata.ucar.edu	715.9 k (0.7%)	4.1 M (2.0%)	287.0 (1.4%)	451.3 (1.8%)	10.0 (0.3%)	16.3 (1%)
#7	tempest.ucar.edu	31.4 k (0.0%)	3.2 M (1.6%)	73.7 (0.4%)	292.7 (1.2%)	3.3 (0.1%)	3.3 (0%)
#8	adds.rap.ucar.edu	134.2 k (0.1%)	2.2 M (1.1%)	161.3 (0.8%)	241.3 (1.0%)	113.3 (3.9%)	147.7 (5%)
#9	bv1203ge.ucar.edu	40.5 k (0.0%)	2.2 M (1.1%)	92.3 (0.5%)	186.7 (0.7%)	3.3 (0.1%)	2.3 (0%)
#10	muffet.rap.ucar.edu	101.6 k (0.1%)	2.2 M (1.1%)	161.3 (0.8%)	234.0 (0.9%)	109.7 (3.8%)	143.3 (5%)

Netflow Applications – Cisco CLI

Show NetFlow Information 'show ip cache flow'

```
router_A#sh ip cache flow
IP packet size distribution (85435 total packets):
  1-32   64   96  128  160  192  224  256  288  320  352  384  416  448  480
  .000 .000 .000 .000 .000 .000 .000 .000 .000 .000 .000 .000 .000 .000 .000

  512  544  576 1024 1536 2048 2560 3072 3584 4096 4608
  .000 .000 .000 .000 1.00 .000 .000 .000 .000 .000 .000
```

Packet sizes

```
IP Flow Switching Cache, 278544 bytes
2728 active, 368 inactive, 85310 added
463824 age polls, 0 flow alloc failures
Active flows timeout in 30 minutes
Inactive flows timeout in 15 seconds
last clearing of statistics never
```

of active flows

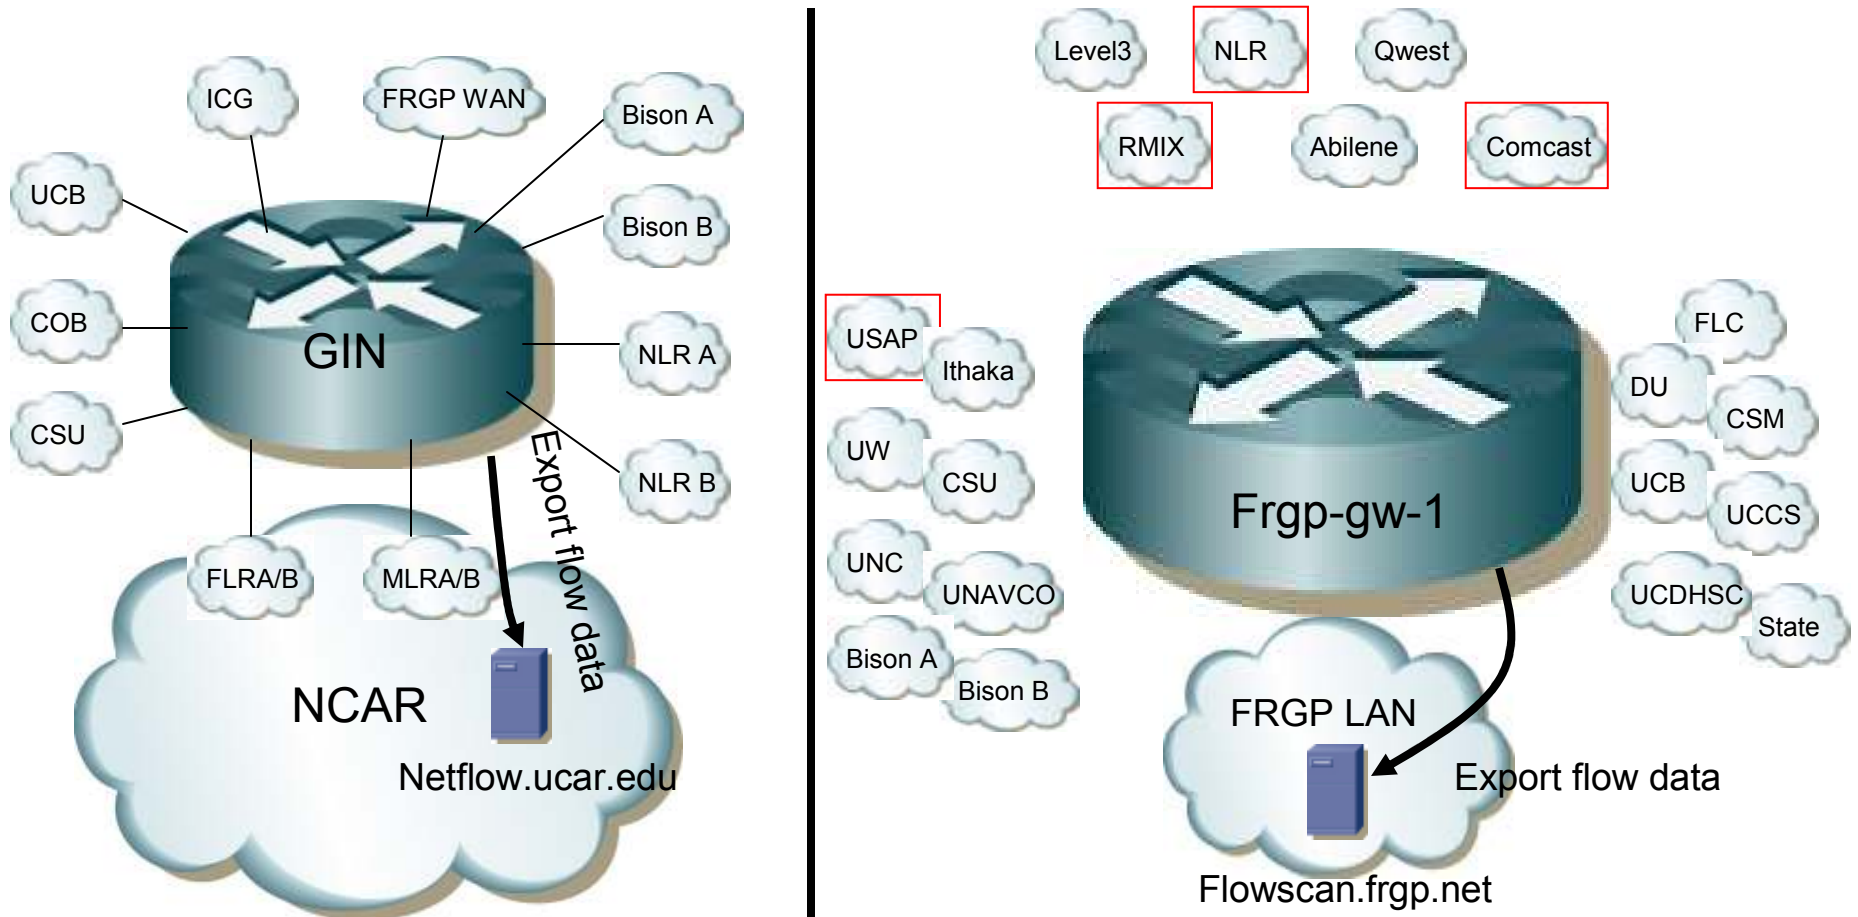
Rates and duration

Protocol	Total Flows	Flows /Sec	Packets /Flow	Bytes /Pkt	Packets /Sec	Active (Sec) /Flow	Idle (Sec) /Flow
TCP-X	2	0.0	1	1440	0.0	0.0	9.5
TCP-other	82580	11.2	1	1440	11.2	0.0	12.0
Total:	82582	11.2				0.0	12.0

Flow details cache

SrcIf	SrcIPAddress	DstIf	DstIPAddress	Pr	SrcP	DstP	Pkts
Et0/0	132.122.25.60	Se0/0	192.168.1.1	06	9AEE	0007	1
Et0/0	139.57.220.28	Se0/0	192.168.1.1	06	708D	0007	1
Et0/0	165.172.153.65	Se0/0	192.168.1.1	06	CB46	0007	1


Current NCAR Netflow Deployment



= Netflow not enabled


Performance Impact – Juniper

- CPU -Initial spike on RE process when new interfaces are enabled.
- Memory – none, does not keep state (cache)
- Limited by 8000pps across the RE to Forwarding Board 100 Mbps Ethernet link
 - ROT – 3 samples are bundled per packet; $3 * 8k = 24k$ max samples/sec. Look at all interfaces and total bps. Choose sampling rate $< 24k$ samples/sec
- Current BW used for export
 - Gin, 200kbps (1:100) Sampling
 - Frqp-gw-1, 10kbps (1:10000) Sampling
- Alleviate load problems by
 - using sampled netflow
 - Use firewall filters to include/exclude traffic
 - Enable on specific interfaces



Performance Impact – Cisco 6500

- CPU load - depends on number of flows and flow cache timer settings. Lower the timer setting, the higher the CPU because it is constantly looking through the cache for flows to export.
- Memory -Netflow Cache
 - Timers:
 - Inactive timer (Normal Aging) ; on 6k its default is 256 sec - should set it to 30 sec
 - Active timer (Long Aging) ; 32 minutes
 - PFC3B – can hold approx. 115k flow entries
 - if cache has too many flows then flows are dropped (lower Inactive timer)
- Alleviate load problems by
 - using sampled netflow
 - Use flow masks on 65k
 - Use “exclude” filters on 65k
 - Tweak timers
 - Enable on specific interfaces.



Future

- Move to v9/IPFIX
- Enable NDE on Cisco 65ks (tcom & L3 gw).
- Send all export data to collectors.

References

- Introduction to IP Accounting and Netflow, Cisco Networkers 2006.
- Juniper Networks Solutions for Network Accounting; Chuck Semeria, Marketing Engineer; Hannes Gredler, Professional Services Engineer; 07/01.
- <http://www.juniper.net/techpubs/software/junos/junos83/swconfig83-services/download/flow-monitoring-config.pdf>
- <http://www.juniper.net/techpubs/software/junos/junos80/feature-guide-80/html/feature-guide-80TOC.html>
- Catalyst 6500 Series Cisco IOS Command Reference, 12.2SX, http://www.cisco.com/en/US/customer/products/hw/switches/ps708/products_command_reference_chapter09186a00801ea88c.html#wp1179040
- Check the flow-tools April 2007 mailing list for reasons on to enable sampling on the 65k switches.
- Juniper Networks Routing Architecture; <http://www.awprofessional.com/articles/article.asp?p=30631&seqNum=2>

Backup Slides

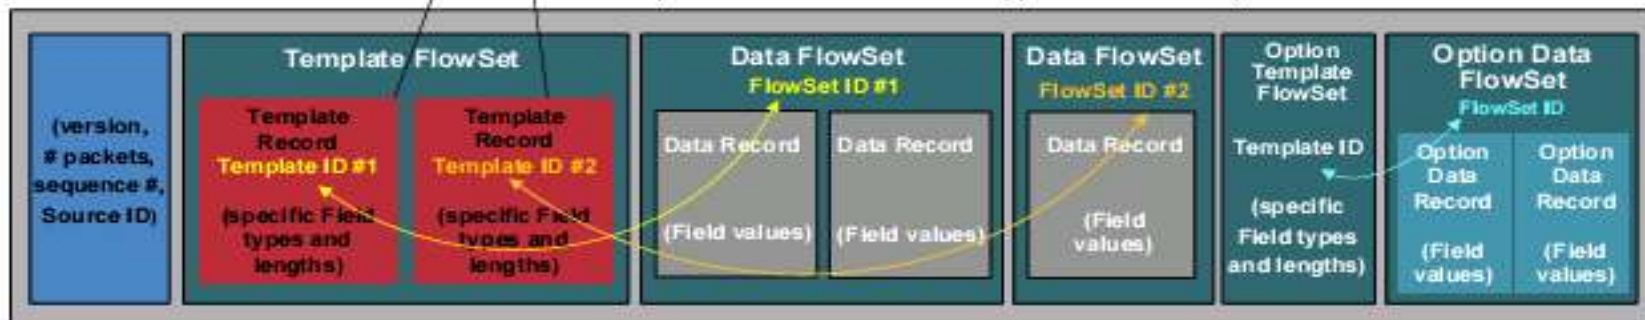
v9

NetFlow v9 Export Packet

To support technologies such as MPLS or Multicast, this export format can be leveraged to easily **insert new fields**

Flows from Interface A

Flows from Interface B



- Matching ID numbers are the way to associate template to the data records
- The header follows the same format as prior NetFlow versions so collectors will be backward compatible
- Each data record represents one flow
- If exported flows have different fields, they cannot be contained in the same template record (i.e.: BGP next hop cannot be combined with MPLS aware NetFlow records)

V5 Header

Table 23 NDE Version 5 Header Format

Bytes	Content	Description
0-1	version	Netflow export format version number
2-3	count	Number of flows exported in this packet (1-30)
4-7	SysUptime	Current time in milliseconds since router booted
8-11	unix_secs	Current seconds since 0000 UTC 1970
12-15	unix_nsecs	Residual nanoseconds since 0000 UTC 1970
16-19	flow_sequence	Sequence counter of total flows seen
20-21	engine_type	Type of flow switching engine
21-23	engine_id	Slot number of the flow switching engine

V5 Flow Record, Part 1

Table 24 NDE Version 5 Flow Record Format

Bytes	Content	Description	Flow masks: • X=Populated • A=Additional field (see the "Populating Additional NDE Fields" section)				
			Destination	Destination Source	Destination Source Interface ¹	Full	Full Interface ¹
0-3	srcaddr	Source IP address		X	X	X	X
4-7	dstaddr	Destination IP address	X	X	X	X	X
8-11	nexthop	Next hop router's IP address	A ²	A	A	A	A
12-13	input	Ingress interface SNMP ifIndex			X		X
14-15	output	Egress interface SNMP ifIndex	A ²	A	A	A	A
16-19	dPkts	Packets in the flow	X	X	X	X	X
20-23	dOctets	Octets (bytes) in the flow	X	X	X	X	X
24-27	first	SysUptime at start of the flow	X	X	X	X	X
28-31	last	SysUptime at the time the last packet of the flow was received	X	X	X	X	X
32-33	srcport	Layer 4 source port number or equivalent				X	X
34-35	dstport	Layer 4 destination port number or equivalent				X	X

V5 Flow Record, part 2

36	pad1	Unused (zero) byte					
37	tcp_flags	Cumulative OR of TCP flags					
38	prot	Layer 4 protocol (for example, 6=TCP, 17=UDP)				X	X
39	tos	IP type-of-service byte					
40-41	src_as	Autonomous system number of the source, either origin or peer		A	A	A	A
42-43	dst_as	Autonomous system number of the destination, either origin or peer	A	A	A	A	A
44-45	src_mask	Source address prefix mask bits					
46-47	dst_mask	Destination address prefix mask bits					
48	pad2	Pad 2 is unused (zero) bytes					

Competing Technologies

- SNMP – Simple Network Management Protocol
- NBAR – Network Based Application Recognition
- BGP PA – BGP Policy Accounting
- AAA – Authentication, Authorization, Accounting

Scenario	Technology
Network Monitoring	NetFlow, BGP PA
Network Planning and Traffic Engineering	NetFlow, BGP PA
Application Monitoring	NBAR, NetFlow
User Monitoring	AAA, NetFlow
QoS/CoS Monitoring	CB-QoS MIB, IP SLAs, NetFlow
Security Analysis	NetFlow, NBAR
Peering and Transit Agreements	SNMP, NetFlow, BGP PA
Time and Usage-Based Billing	AAA, NetFlow
Destination and Source-Sensitive Billing	BGP PA, NetFlow