


UNIVERSITY OF TWENTE.




Internet Management and Measurements Measurements

Ramin Sadre, Aiko Pras

Design and Analysis of Communication Systems Group

University of Twente, 2010



Measurements

- § What is being measured?
- § Why do you measure?
- § How do you do it?

Measurements

§ What is being measured?

§ Why do you measure?

§ How do you do it?

What is being measured?

- § Delay (one-way, round-trip)
- § Delay variation (jitter)
- § Throughput (average, peak,...)
- § Packet loss
- § Protocol/application usage
- § Nature of data exchanged between hosts
- § ...

Measurements

§ What is being measured?

§ Why do you measure?

§ How do you do it?

Why do you measure?

- § Traffic engineering
- § Intrusion detection
- § Accounting
- § Lawful interception
- § ...

Why do you measure?

§ Traffic engineering

§ Intrusion detection

§ Accounting

§ Lawful interception

§ ...

Traffic Engineering


- § Predict, plan, understand the behavior of the network
 - § Link usage (bottlenecks?)
 - § Packet losses
 - § Delays
 - § ...
- § Goal: optimize
 - § Quality of service provided to customers
 - § Costs

Traffic Engineering

- § In general, you can not just “try” with the real network
- § Needed:
 - § Formal models that describe the (future) behavior of the network
 - § Tools to evaluate the models by analysis or simulation
 - § Measurement data!
 - § Create and parametrize the models
 - § Validate the results

Example: Dimensioning Router Buffers


- § What is the packet loss L for a buffer size S ?
- § Router port can be modeled as a queueing station
- § Traffic input process?



Modeling Traffic Processes

- § What kind of modeling techniques do we need?
- § Traditional formal approaches based on “Markovian” behavior (Poisson distributions,...)
- § Measurements since the 1980s have shown that network traffic has “non-Markovian” properties:
 - § Self-similarity
 - § Long-range dependence
 - § Heavy-tailed distributions
- § Requires new models and tools


Self-similarity



(from: Traffic Characterisation for Telecommunication Networks, 1999)

Heavy-tailedness

WWW	min	Max	mean	median
2000	17 B	0.23 GB	12 KB	2.4 KB
2007	85 B	2.15 GB	68 KB	2.7 KB




Why do you measure?


- § Traffic engineering
- § Intrusion detection
- § Accounting
- § Lawful interception
- § ...

Intrusion Detection

Number of reported incidents (CERT):



Intrusion Detection



SOURCE: D1.4 SCAMPI PROJECT

Intrusion Detection: Classification

Main classification criteria:

§ Location of observation

§ Host-based: only observe traffic from/to a particular host

§ Network-based: observe the traffic in a whole (sub)network

§ Detection method:

§ Misuse-based: look for known patterns of misuse

§ Anomaly-based: everything that deviates from normality is suspicious

Two reasons to measure:

1. Learn from traffic measurements how good/bad traffic looks like
2. Protect the network

Intrusion Detection: Other Criteria

§ Analyzed data:

- § Log files

- § Packet payload

- § Only packet-headers

- § ...

§ Data collection

- § From multiple locations

- § Only one location

§ Adaptation

- § Static: has to be configured by user

- § Self-adapting: adapts automatically to changes in the network

- § ...

Internet Background Radiation

§ Idea: measure all traffic destined for unused IP addresses in a network

§ Study at Lawrence Berkeley National Laboratory

Characteristics of Internet Background Radiation. R. Pang, V. Yegneswaran, P. Barford, V. Paxson, L. Peterson. Proc. of the ACM Sigcomm Internet Measurement Conference, Taormina, Sicily, Italy 2004


§ “Background radiation” of the Internet consists of:

§ Non-productive traffic (misconfigurations)

§ Malicious traffic (scans, worms,...)

Measurement Results at LBL

TCP Port	# Source IP (%)	# Packets (%)
445	43.40%	19.70%
80	28.70%	7.30%
135	19.10%	30.40%
1025	4.30%	5.80%
2745	3.20%	3.60%



Why do you measure?

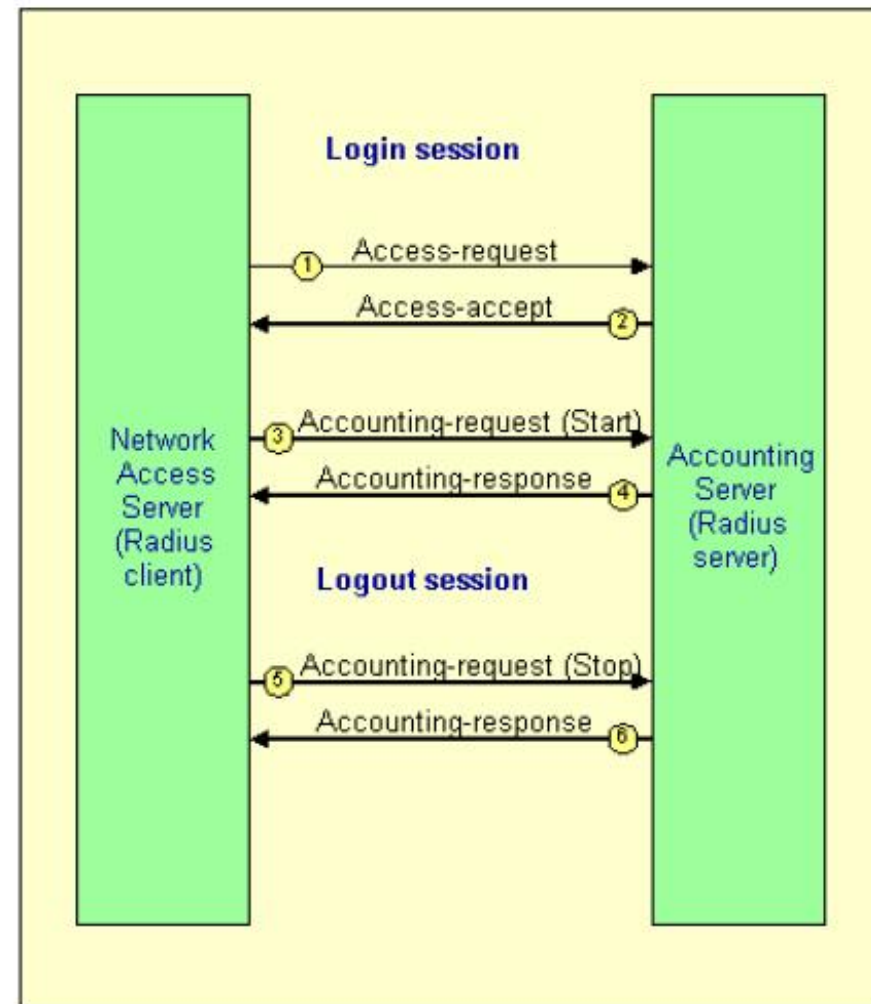
- § Traffic engineering
- § Intrusion detection
- § Accounting
- § Lawful interception
- § ...

Accounting

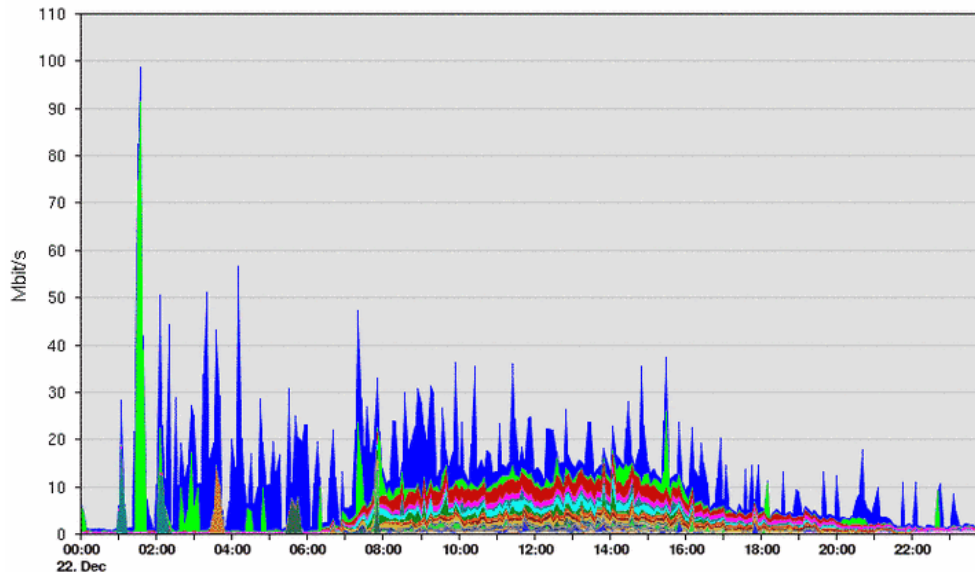
- § Who uses what, when, how often?
- § Reasons:
 - § Billing
 - § Limit resource usage
 - § Security
- § Can be done at different levels:
 - § Authentication services
 - § Applications
 - § Observe the traffic generated by the user
 - § ...

RADIUS

- § RFC 2866
- § Logout: NAS sends stop record
- § Session time
- § Packets transferred
- § Octets transferred
- § Disconnect reason
- § ...



Traffic-based Accounting: Example IsarFlow (IsarNet, Germany)



Nr.	Protokoll	Byte	min	avg	max	%
1	NETBIOS-SSN (tcp: 139)	70.12 GB	173.30 Kbps	6.97 Mbps	56.27 Mbps	45.94
2	MICROSOFT-DS (tcp: 445)	18.58 GB	26.69 Kbps	1.85 Mbps	91.01 Mbps	12.18
3	HTTP (tcp: 80)	12.13 GB	12.45 Kbps	1.21 Mbps	5.03 Mbps	7.95
4	AUTO: TCP/8000 (tcp: 8000)	6.86 GB	356.84 Kbps	682.43 Kbps	1.25 Mbps	4.50
5	TCP-Other	5.19 GB	7.66 Kbps	516.27 Kbps	18.07 Mbps	3.40
6	HTTP-PROXY (tcp: 8080)	4.71 GB	60.16 bps	468.19 Kbps	2.74 Mbps	3.09

Accounting PerCustomer_RX_TX

Es wurden keine Filter ausgewählt

29.06.2008 00:00 - 23:59

Nr.	Farmname	gesendete Bytes	empfangene Bytes	Bytes gesamt
1	Architektur	18.29 MB	56.32 MB	74.61 MB
2	Betriebswirtschaft	2.17 MB	1.74 MB	3.91 MB
3	Bibliothek	3.13 MB	36.43 MB	39.56 MB
4	Dial-In	69.12 MB	1.04 MB	70.16 MB
5	FH Augsburg (remaining)	4.28 GB	4.96 GB	9.25 GB
6	Gestaltung	9.75 KB	7.76 MB	7.77 MB
7	INTERNET	22.33 GB	21.90 GB	44.23 GB
8	Informatik	101.47 MB	310.85 MB	412.32 MB
9	Maschinenbau	887.32 KB	5.07 MB	5.93 MB
10	Multimedia	110.74 MB	582.51 MB	693.26 MB
11	PC-Pool Test	77.62 KB	3.88 MB	3.96 MB
12	Rechenzentrum	16.94 GB	8.52 GB	25.46 GB
13	Spider-Server	175.17 MB	7.64 MB	182.81 MB
14	VPN	871.02 MB	8.55 GB	9.40 GB
15	Verwaltung	76.20 MB	16.10 MB	92.30 MB
Tabellensumme		44.95 GB	44.95 GB	89.89 GB
CSV Export				

(from: isarflow.de)

Why do you measure?

- § Traffic engineering
- § Intrusion detection
- § Accounting
- § Lawful interception
- § ...

European Directive

Directive 2006/24/EC of the European Parliament and of the Council of 15 March 2006:

§ 1. *Member States shall ensure that the following categories of data are retained under this Directive:*

§ (a) *data necessary to trace and identify the source of a communication:*

§ (1) *concerning fixed network telephony and mobile telephony:*

§ (i) *the calling telephone number;*

§ (ii) *the name and address of the subscriber or registered user;*

§ (2) *concerning Internet access, Internet e-mail and Internet telephony:*

§ (i) *the user ID(s) allocated;*

§ (ii) *the user ID and telephone number allocated to any communication entering the public telephone network;*

§ (iii) *the name and address of the subscriber or registered user to whom an Internet Protocol (IP) address, user ID or telephone number was allocated at the time of the communication;*

§ (b) *data necessary to identify the destination of a communication:*

§ ...

Measurements

§ What is being measured?

§ Why do you measure?

§ How do you do it?

How to measure?

§ Passive measurement:

- § Log files

- § Observe traffic

- § ...

§ Active measurement:

- § Send probing packets

- § Example: ping

How to measure?

§ Passive measurement:

§ Log files

§ **Observe traffic**

§ ...

§ Active measurement:

§ Send probing packets

§ Example: ping

Record packets with tcpdump

- § Capture network traffic (not only TCP) at specified interface
- § Collected data stored as “pcap” files
- § Can be used as library by other applications
- § Many capturing options, filters,...
- § Example:

```
tcpdump 'tcp[tcpflags] & (tcp-syn|tcp-fin)!=0'
```

Wireshark

The screenshot displays the Wireshark interface with a packet capture of an HTTP GET request. The main pane shows a list of packets, with packet 16 selected. The packet list table is as follows:

No.	Time	Delta	Source	Destination	Protocol	Info
13	14.817570	14.817570	192.168.0.10	192.168.0.2	TCP	1242 > 80 [SYN] Seq=1404510823 Ack=0 Win=65536
14	14.817689	0.000119	192.168.0.2	192.168.0.10	TCP	80 > 1242 [SYN, ACK] Seq=3661615104 Ack=1404510824
15	14.818178	0.000489	192.168.0.10	192.168.0.2	TCP	1242 > 80 [ACK] Seq=1404510824 Ack=3661615105
16	14.819035	0.000857	192.168.0.10	192.168.0.2	HTTP	GET / HTTP/1.1
17	14.975815	0.156780	192.168.0.2	192.168.0.10	TCP	80 > 1242 [ACK] Seq=3661615105 Ack=1404511235
23	19.382555	4.406740	192.168.0.10	192.168.0.2	TCP	1242 > 80 [FIN, ACK] Seq=1404511234 Ack=3661615105
24	19.382634	0.000079	192.168.0.2	192.168.0.10	TCP	80 > 1242 [ACK] Seq=3661615105 Ack=1404511235
52	54.234482	34.851848	192.168.0.2	192.168.0.10	HTTP	HTTP/1.1 403 Forbidden (text/html)
53	54.235272	0.000790	192.168.0.10	192.168.0.2	TCP	1242 > 80 [RST] Seq=1404511235 Ack=366044707
54	58.137063	3.901791	192.168.0.10	192.168.0.2	TCP	1244 > 135 [SYN] Seq=1414452237 Ack=0 Win=65536
55	58.137176	0.000113	192.168.0.2	192.168.0.10	TCP	135 > 1244 [SYN, ACK] Seq=3672465192 Ack=1414452238
56	58.137527	0.000351	192.168.0.10	192.168.0.2	TCP	1244 > 135 [ACK] Seq=1414452238 Ack=3672465192
57	58.137992	0.000465	192.168.0.10	192.168.0.2	DCERPC	bind: call_id: 57 UUID: IOXIDResolver
58	58.188933	0.050941	192.168.0.2	192.168.0.10	DCERPC	bind_ack: call_id: 57 accept_max_xmit: 5840
59	58.189601	0.000668	192.168.0.10	192.168.0.2	IOXIDR	complexPing request AddToSet=0 DelFromSet=0
60	58.202631	0.013030	192.168.0.2	192.168.0.10	IOXIDR	complexPing response -> unknown (0x00000778)
61	58.203457	0.000826	192.168.0.10	192.168.0.2	IOXIDR	complexPing request Address=0

The packet details pane for packet 16 shows the following structure:

- Frame 16 (464 bytes on wire, 464 bytes captured)
- Ethernet II, Src: 00:04:61:4a:1e:95, Dst: 00:0b:5d:20:cd:02
- Internet Protocol, Src Addr: 192.168.0.10 (192.168.0.10), Dst Addr: 192.168.0.2 (192.168.0.2)
- Transmission Control Protocol, Src Port: 1242 (1242), Dst Port: 80 (80), Seq: 1404510824, Ack: 3661615105, Len: 410
- Hypertext Transfer Protocol
 - GET / HTTP/1.1\r\n
 - Host: 192.168.0.2\r\n
 - User-Agent: Mozilla/5.0 (Windows; U; Windows NT 5.0; en-us; rv:1.5) Gecko/20031007\r\n
 - Accept: text/xml,application/xml,application/xhtml+xml,text/html;q=0.9,text/plain;q=0.8,image/png,image/jpeg,image/gif;q=0.7\r\n
 - Accept-Language: en-us,en;q=0.5\r\n
 - Accept-Encoding: gzip,deflate\r\n
 - Accept-Charset: ISO-8859-1,utf-8;q=0.7,*;q=0.7\r\n
 - Keep-Alive: 300\r\n
 - Connection: keep-alive\r\n

The packet bytes pane shows the raw data of the captured packet, with a hex dump and ASCII representation.

Filter: tcp

Location of Measurement

§ Single hosts:

§ Only traffic directed to/generated by the host

§ All traffic seen on the network segment

§ Switches, routers:

1. Create mirror port (Cisco: SPAN-Port) for one or more source ports

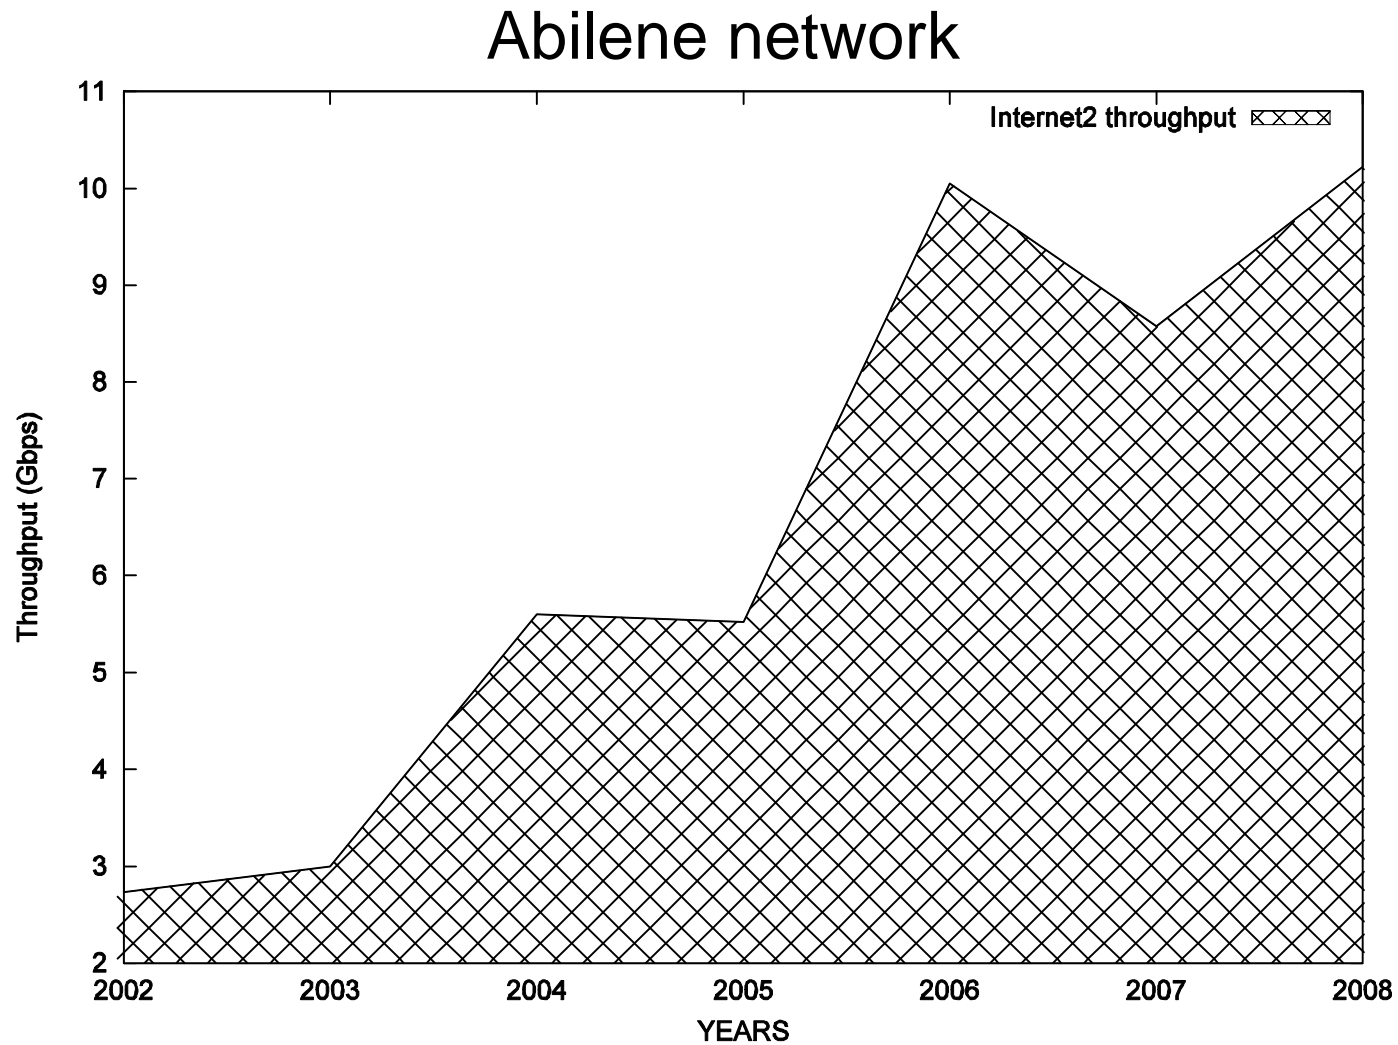
2. Send traffic from mirror port to a measurement host

§ Packet loss if traffic throughput is...

§ > mirror port bandwidth

§ > storage/processing rate at measurement host

Limitations of Packet Measurement



Limitations of Packet Measurement

Collect and process less information:

- § Only collect packet headers, not payload
- § Ignore single packets (aggregate)
- § Ignore some packets (sampling)

Make collection and processing faster:

- § Move to kernel space
- § Distributed collection & processing
- § Dedicated hardware

Limitations of packet measurement

Collect and process less information:

§ Only collect packet headers, not payload

§ Ignore single packets (aggregate)

§ Ignore some packets (sampling)

Make collection and processing faster:

§ Move to kernel space

§ Distributed collection & processing


§ Dedicated hardware

Sflow, Netflow, IPFIX

A flow “summarizes” a sequence of packets:

- § Source & destination IP address
- § Source & destination port number
- § Source & destination AS
- § Layer 3 protocol type
- § Size (aggregated number of bytes)
- § Timestamps of first and last packet
- § ...

NetFlow infrastructure



Protocols:


- Netflow v5, 9: quasi-standard. by Cisco
- IPFIX: from Netflow v9, by IETF

(from: Cisco)

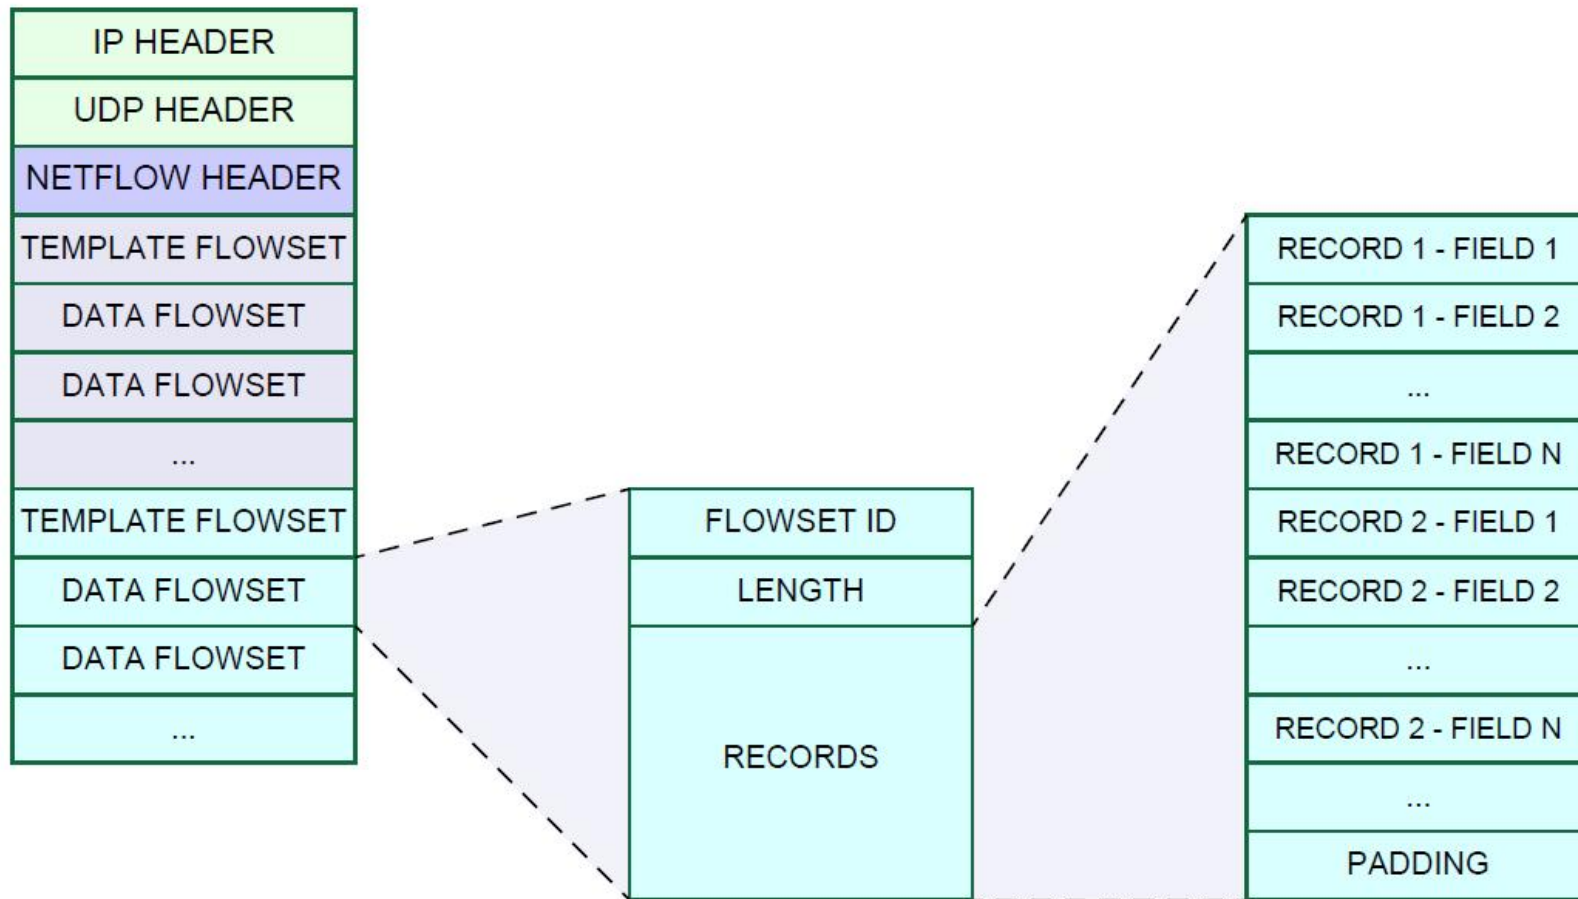
Flexible Netflow: flexible monitoring, by Cisco

Netflow v9

Exported data specified by templates (\neq v5)

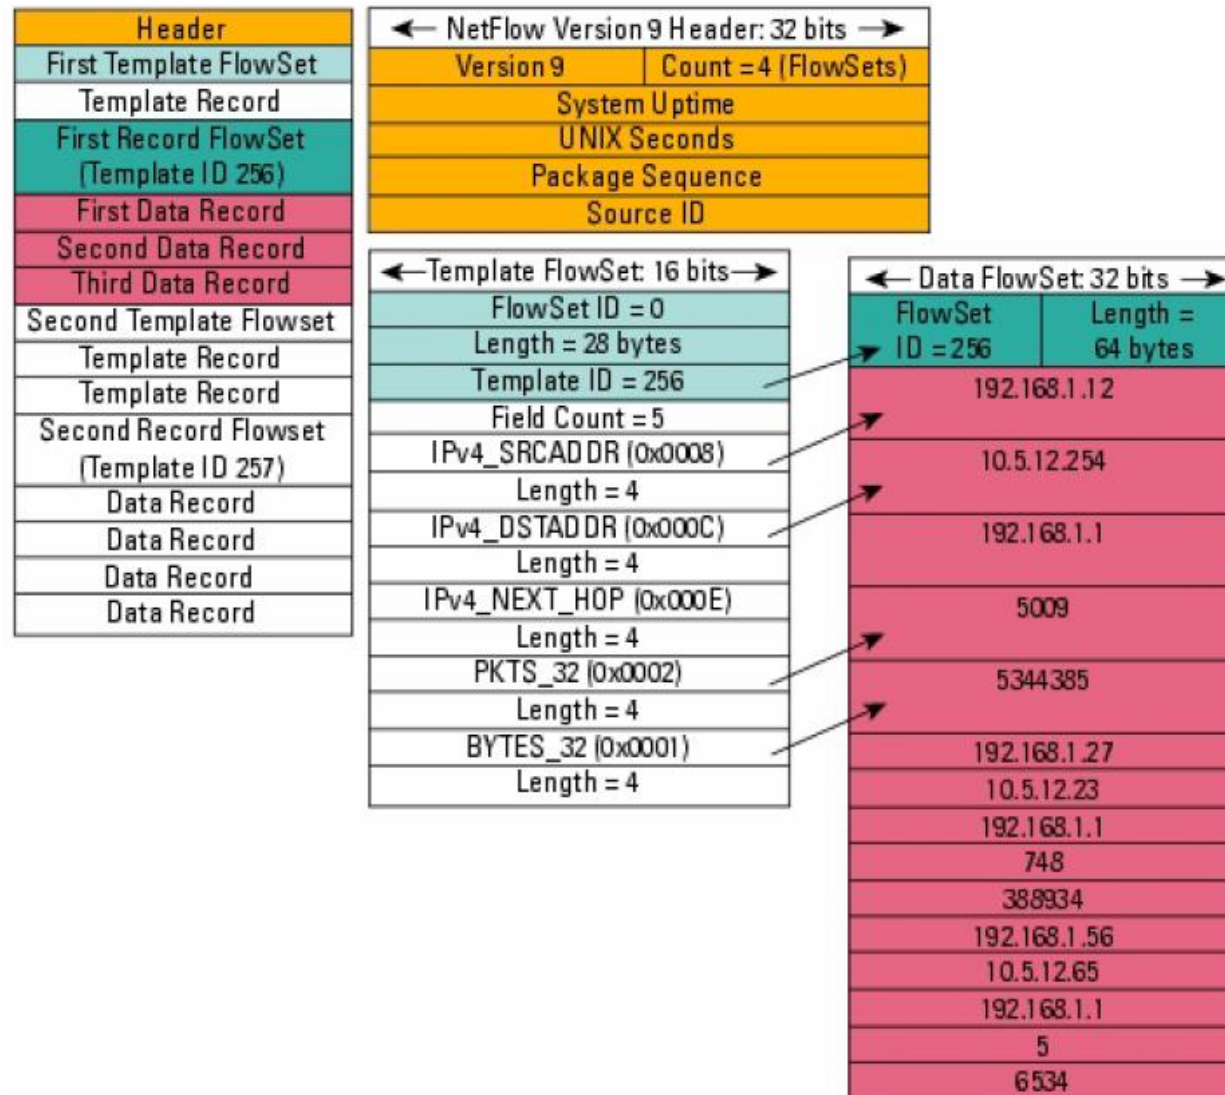


Netflow v9



(from: Cisco)

Netflow v9



Flows vs. Packets

- § Of course, loss of information (no payload, no details of packet headers)
- § But sometimes the only option
- § Still useful for:
 - § Accounting
 - § Interception
 - § Even intrusion detection:
 - § Scans
 - § Some kinds of DoS
 - § ...