Using NetFlow to Analyze Your Network

June 15th, 2011

Christopher J. White
Manager Applications and Analytics, Cascade
Riverbed Technology
cwhite@riverbed.com

SHARKFEST **'11**
Stanford University
June 13-16, 2011

NetFlow Data Collection

Collecting flow reports from devices throughout the network to provide flow-level visibility into network behavior and how traffic is delivered end-to-end across the network.

NetFlow Data Collection

- Overview
  - Flow Definition
  - Observation Points
  - Flow Records and Export
  - Flow Collection
  - Deduplication and Time Slices

Flows

- ***What is a Flow?***
  - A **Flow** is a set of IP packets in the network that all share a common set of key attributes
  - A typical flow is based on the 5-tuple:

    <SrcIP, DstIP, Protocol, SrcPort, DstPort>
  - In general, a flow is *unidirectional, e.g.* describing only half of a TCP connection
  - A flow may be defined only a subset of available attributes, such as just

    <SrcIP, DstIP>

Flows

- Flows are tracked at **Observation Points**
- The same logical flow may be observed at multiple different observation points
  - Different devices as the flow traverses the network
  - Within the same device, ingress and egress
- The same flow may (and probably *will*) yield different counts at each observation point

Flows

- **Ingress Flow**
  - Flow is incoming to the device at the observation point
  - Bytes / Packets received at the ingress interface prior to processing
- **Egress Flow**
  - Flow is leaving the device at the observation point
  - Bytes / Packets transmitted at the egress interface after processing

Flows

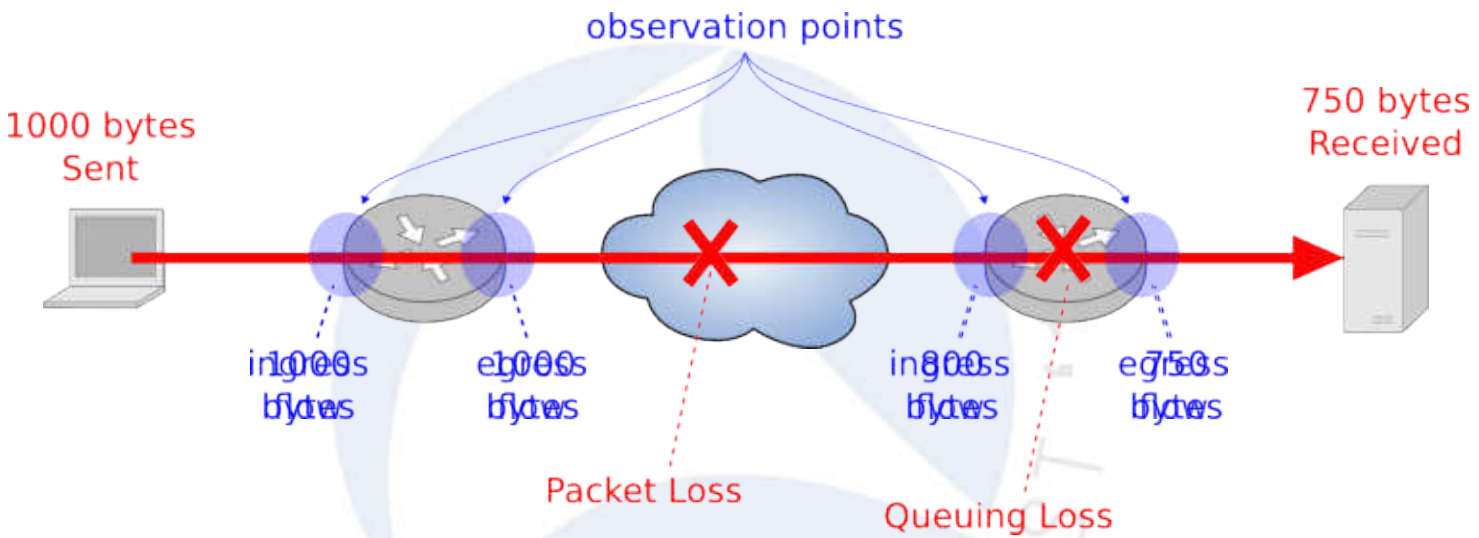
observation points
1000 bytes
Sent
750 bytes
Received
Packet Loss
Queuing Loss

Flow Records and Export

- **Flow Key** - set of attributes that distinguish one flow from another
  - Most common key is the 5-tuple:

    <SrcIP, DstIP, Protocol, SrcPort, DstPort>
  - Key is usually fully configurable based on standard IP Packet Fields

- **Flow Record** – collection of information gathered about a flow at an observation point
  - Flow Key values
  - Bytes / Packets observed
  - Ingress / Egress interface
  - QoS tags

Flow Records and Export

- **Metering** – the process of generating flow records based on packets at an observation point
- **Exporting** – sending flow records to one or more collectors for further processing

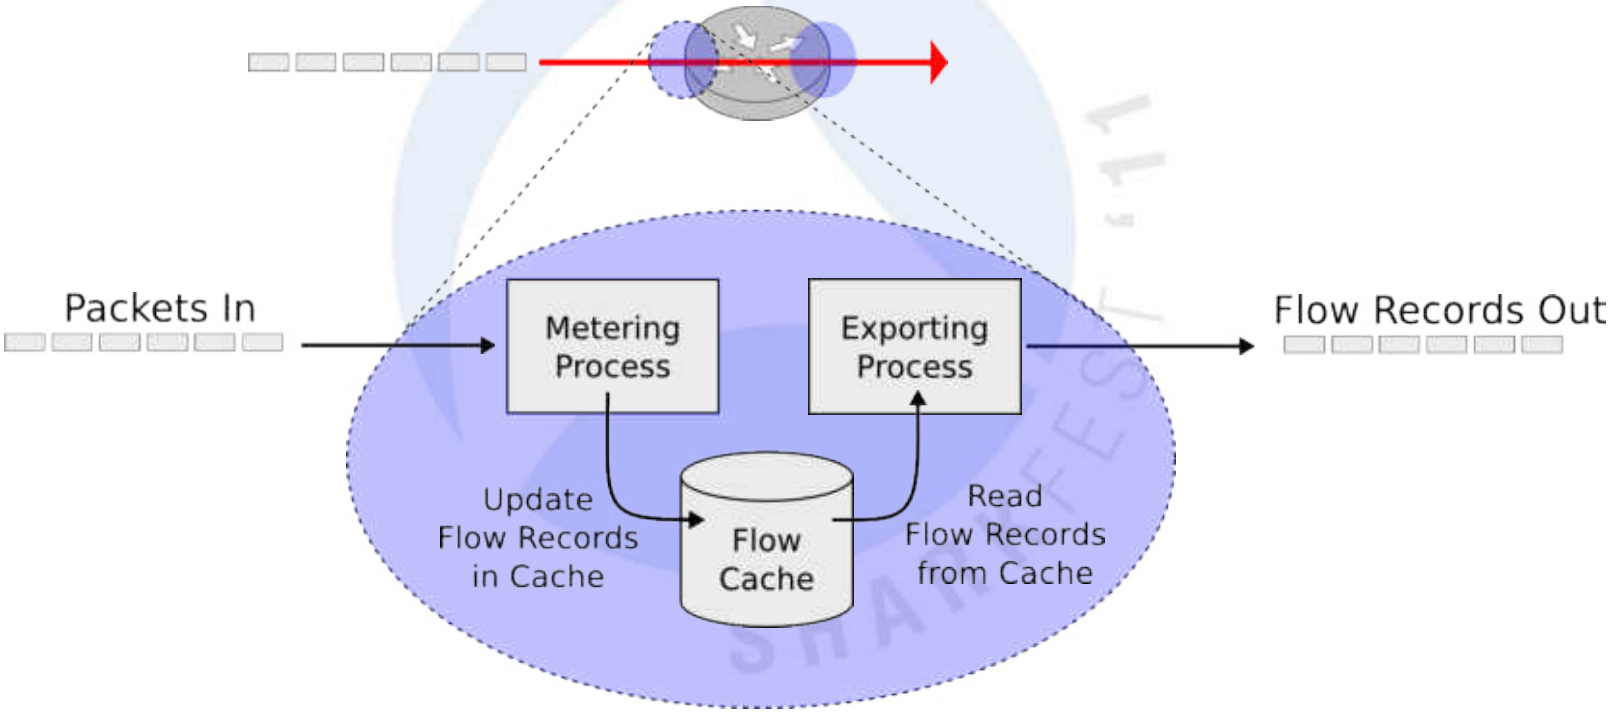

Flow Records and Export

Typical NetFlow v5 Flow Record

Exporter Address	10.99.1.1
IP Source	10.1.1.121
IP Dest	74.125.67.100
Protocol	TCP
TCP Source Port	16341
TCP Dest Port	80
Bytes	34,778
Packets	29
TCP Flags	SAPF
Ingress Interface	1
Egress Interface	2
QoS	Default
First Packet Timestamp	Oct 21, 8:24.12.321
Last Packet Timestamp	Oct 21, 8:24 14.929

Flow Records and Export

- ***When are Flow Records Exported?***
  - **Flow Terminated**
    - TCP connection terminated due to FIN / RST
  - **Inactive Flow Timeout**
    - the flow has been inactive for 15s
  - **Active Flow Timeout**
    - the flow has been active for 60s
  - **Forced Flow Discard**
    - flow cache is full and a new flow must be allocated
    - LRU / Random

Flow Records and Export

- ***Is it possible to have multiple flow record exports for the same flow?***

  (from the same observation point)

  - **Yes** – long lived flows, forced discard, etc.

  - Each exported Flow Record is an *update* to the flow
    - **Flow Key** is *identical* in each update
    - **Flow Record** *changes,* eg. bytes / packets counters are deltas seen since the last export for this flow

Flow Records and Export

- ***What is the mechanism for exporting flow?***
  - Flow Records are exported via IP to a configured **Flow Collector** IP Address / Port
    - UDP for NetFlow v5/v9
    - UDP/TCP/SCTP for IPFIX
  - Flow Export is a continual process for as long as the device is up and receiving packets

Flow Records and Export

- Flow Export is a continual process…

Start	End	SrcAddr	DstAddr	Proto	SrcPort	DstPort	Packets	Bytes	Flags
8:00:15.234	8:01:14.782	10.1.1.1	192.168.1.2	6	80	28456	40	72,554	SAPF
8:01:03.119	8:01:14.790	192.168.44.1	74.125.224.48	6	33152	80	52	14,032	SA
8:00:15.345	8:01:14.921	192.168.1.2	10.1.1.1	6	28456	80	220	112,005	SAR
8:01:15.448	8:01:15.448	192.168.14.1	10.4.1.2	17	58440	53	1	76	
8:01:15.552	8:01:15.552	10.4.1.2	192.168.14.1	17	53	58440	1	520	
8:00:19.445	8:01:15.634	10.1.1.1	10.1.2.4	1	0	8	50	6,104	
…	…	…	…	…	…	…	…	…	
8:01:19.021	8:02:16.789	192.168.1.2	10.1.1.1	6	28456	80	220	112,005	AF
8:02:15.313	8:01:16.899	10.1.122.4	204.14.234.33	6	47125	80	10	620	S

Flow Records and Export

- Flow Export is a continual process…

Start	End	SrcAddr	DstAddr	Proto	SrcPort	DstPort	Packets	Bytes	Flags
8:00:15.234	8:01:14.782	10.1.1.1	192.168.1.2	6	80	28456	40	72,554	SAPF
8:01:03.119	8:01:14.790	192.168.44.1	74.125.224.48	6	33152	80	52	14,032	SA
8:00:15.345	8:01:14.921	192.168.1.2	10.1.1.1	6	28456	80	220	112,005	SAR
8:01:15.448	8:01:15.448	192.168.14.1	10.4.1.2	17	58440	53	1	76	
8:01:15.552	8:01:15.552	10.4.1.2	192.168.14.1	17	53	58440	1	520	
8:00:19.445	8:01:15.634	10.1.1.1	10.1.2.4	1	0	8	50	6,104	
…	…	…	…	…	…	…	…	…	
8:01:19.021	8:02:16.789	192.168.1.2	10.1.1.1	6	28456	80	220	112,005	AF
8:02:15.313	8:01:16.899	10.1.122.4	204.14.234.33	6	47125	80	10	620	S

Flow Records and Export

- Flow Export is a continual process…

Start	End	SrcAddr	DstAddr	Proto	SrcPort	DstPort	Packets	Bytes	Flags
8:00:15.234	8:01:14.782	10.1.1.1	192.168.1.2	6	80	28456	40	72,554	SAPF
8:01:03.119	8:01:14.790	192.168.44.1	74.125.224.48	6	33152	80	52	14,032	SA
8:00:15.345	8:01:14.921	192.168.1.2	10.1.1.1	6	28456	80	220	112,005	SAR
8:01:15.448	8:01:15.448	192.168.14.1	10.4.1.2	17	58440	53	1	76	
8:01:15.552	8:01:15.552	10.4.1.2	192.168.14.1	17	53	58440	1	520	
8:00:19.445	8:01:15.634	10.1.1.1	10.1.2.4	1	0	8	50	6,104	
…	…	…	…	…	…	…	…	…	
8:01:19.021	8:02:16.789	192.168.1.2	10.1.1.1	6	28456	80	220	112,005	AF
8:02:15.313	8:01:16.899	10.1.122.4	204.14.234.33	6	47125	80	10	620	S

Flow Records and Export

- ***What devices export flow?***
  - ***Most devices already in your network:***
    - Routers
    - Switches
    - Riverbed Steelheads
    - Probes (Cascade Sensor, Packeteer..)

Types of Flow

- **NetFlow v5**
  - Widely in use, supported by multiple vendors
  - Fixed content flow record with basic counters and flow information
- **NetFlow v9**
  - Drastic increase in available fields
  - Templates allow customization of data collected
  - Official support for ingress and egress flows
- **IPFIX (NetFlow v10)**
  - IETF Working Group standard based on v9

Types of Flow

- **J-Flow / Packeteer**
  - NetFlow like variants
- **S-Flow**
  - Sampled
- **Cascade Sensor Flow**
  - L7 Application tag
  - Performance Metrics
    - Network RTT / Response Time
    - TCP Retransmissions
  - packets

Flow Collection

- **Flow Collector** – a device receiving flow records from one or more flow exporters

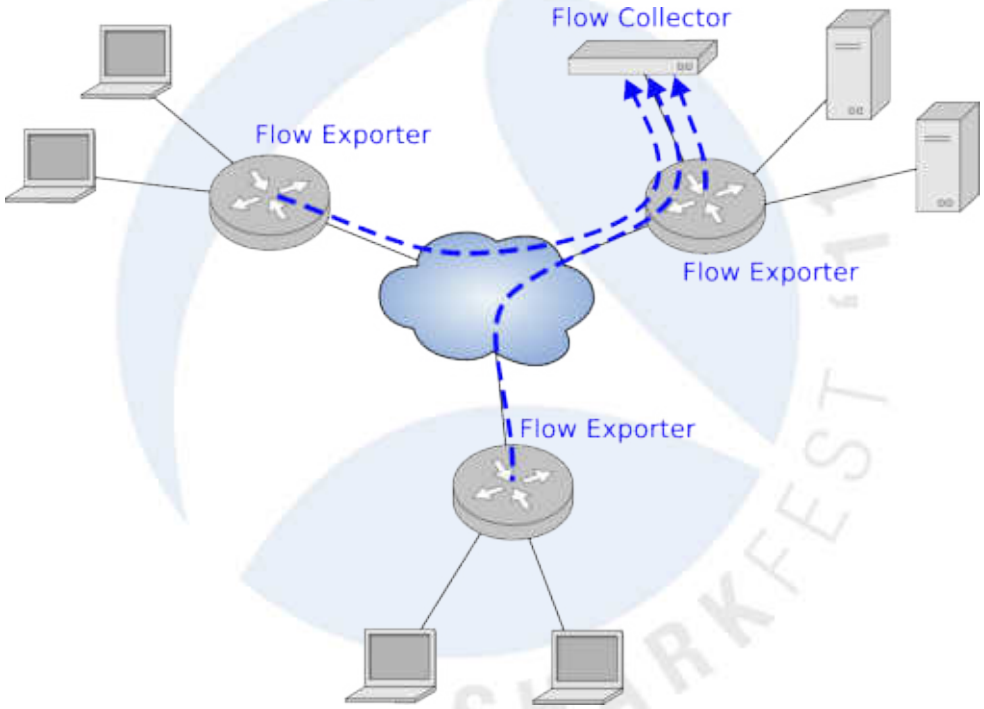

Flow Collection

- The Flow Collector typically serves three primary functions:
  - **Logging** – storing flow records on disk, with post-processing, may yield months to years of archival storage
  - **Reporting** – processing logged flow records to generate reports on network behavior in multiple dimensions
  - **Alerting** – real-time analysis of network behavior and alerting based on comparisons to thresholds and baselines

Flow Collection

- ***What about collecting multiple flow records for the same flow… but from different sources?***
  - Yes… each device that is exporting netflow will send a unique record for the flow
  - Collating records allows end-to-end analysis:
    - *What routers are involved when a client in San Francisco connects directly to the database server?*
    - *How did the byte counts change?*
    - *Was the QoS consistent across the network?*

Flows vs. Packets

- Three Primary Comparisons
  - Types of queries supported
  - Data storage
  - Visibility

Flows vs. Packets

- Supported queries

Query Type	Packets	Flow
TCP inter-packet timing analysis	✓	
TCP window size problems	✓	
SQL transaction times	✓	
Top host-pairs	✓	✓
Clients connecting to external server	✓	✓
Number of connections on Router1:En1/0	✓	✓
Link congestion investigation (hosts involved, ports in use, when the congestion started)	✓	✓

Flows vs. Packets

- Filter Criteria

Hosts

Server

Client

Client Port

Server Port

Application

Interface

Device

QoS

- Output Key

Hosts

Host Pairs

Host Pairs with Server Port

TCP Connection

Interface

Device

Server Port

Interface

Interface / QoS

Device

Application

Server / Application

Server / Server Port

- Data Columns

Bytes

Packets

Connections

TCP Flags

Inbound Bytes (Inteface)

Outbound Bytes (Interface)

Tx Bytes (Cli → Srv)

Rx Bytes (Srv → Cli)

Connection Duration

% Utilization (Interface)

TCP Retransmissions

Network RTT

Server Delay

Flows vs. Packets

- Data rates and storage

	Packets	Flow
Network load for 100Mb link	100 Mb/s	2Mb/s
Data Retention for 1TB Disk	1 day	2 months
Query over last hour	Scan 50GB data	Scan 1GB data

Flows vs. Packets

- Visibility
  - Packets
    - Install one or more probes at each physical location
    - Configure SPAN or install a Tap at each observation point
    - Wire up probes for each observation point
  - NetFlow
    - Install one collector in the data center
    - Configure NetFlow export at each observation point
      - Point export a the single collector

Flows *to* Packets

- Analysis with Flows
  - Historical reporting
  - Trending / Alerting
  - Network Problems
    - Faster queries
    - Broader visibility
    - Identify the actors involved (hosts, ports, devices)
  - Narrow the focus before diving in with Packets
    - Many problems can be fully diagnosed with flow-only

Thank you

Questions