



# 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

## Flow-tools

- Architectural issues
- Software, tools etc

## Lab

# Prelude

Whether your organization is big or small, centralized or decentralized, you should have a network monitoring solution up and running. If you already have a solution in place, you should be thinking about using it for business critical applications, such as regulatory/HR compliance, transactional analysis, Industry specific billing, etc., as well as for network and security troubleshooting. And the time to do it is now, before you need digital evidence for any investigation or forensic analysis.

# Network Qs ...

REQUIREMENT	TECHNOLOGY	SOLUTION
<ul style="list-style-type: none"><li>✓ Who are the top talkers?</li><li>✓ How much<ul style="list-style-type: none"><li>✓ Traffic per user/group?</li><li>✓ Traffic per application?</li><li>✓ Traffic “on-net” / “off-net”</li></ul></li><li>✓ How many users are active on the network at any given time?</li><li>✓ Where<ul style="list-style-type: none"><li>✓ Does the traffic</li><li>✓ come from?</li><li>✓ Does it go to?</li></ul></li><li>✓ When was it transmitted?</li><li>✓ Security attacks?</li></ul>	<ul style="list-style-type: none"><li>✓ Source Address</li><li>✓ Destination Address</li><li>✓ Source Port</li><li>✓ Destination Port</li><li>✓ Layer 3 Protocol Type</li><li>✓ DSCP (DiffServ)</li><li>✓ Input Logical Interface</li><li>✓ BGP Next Hop TOS</li><li>✓ MPLS Label</li><li>✓ MPLS Label Type (LDP, BGP, VPN, ATOM, TE Tunnel MID-PT)</li></ul>	<ul style="list-style-type: none"><li>✓ Network Monitoring</li><li>✓ Network Planning</li><li>✓ Security Analysis</li><li>✓ Application Monitoring</li><li>✓ User Monitoring</li><li>✓ Traffic Engineering</li><li>✓ Peering Agreement</li><li>✓ Usage-Based Billing</li><li>✓ Destination-Sensitive Billing</li></ul>



# Answer is Historic

Originally developed by Darren Kerr and Barry Bruins in 1996 as part of the Cisco IOS, NetFlow™ was first used as a packet switching path selector designed to make the forwarding process within the router more efficient. Only later was NetFlow's true value as a network accounting technology fully realized, providing numerous benefits not available with other network or security technologies:

- network application and user monitoring
- accounting and billing
- network analysis and planning
- traffic engineering
- security analysis
- NetFlow data warehousing and data mining

# What is NetFlow ?

Nearly all Cisco powered networks contain routers and high end switches which support a technology called NetFlow or a competitive technology such as sFlow, IPFIX or Netstream. Enabling flow technology on each router and switch will cause the device to send a steady network flow to a collector.

# What Is NetFlow?

NetFlow is an open protocol developed by Cisco Systems for collecting IP packet information without the cost and complexity of hardware-based network probes. Built into routers and switches from Cisco and other vendors, NetFlow captures detailed information on streams of data that share a common source, destination and protocol — so-called network traffic flows.

*Using Network Flows correctly is the single most important step remaining for enterprises to secure their networks.*

*- Richard Stiennon*

# What Is NetFlow?

- A technology to gather information on forwarded packets
  - In one or several caches
  - Whose content can be queried real time for troubleshooting
  - And exported to collectors
- A protocol to export those records to collectors
  - To use sophisticated tools for analysis
  - Near real-time analysis
  - For long-term analysis, trending, whatever
  - To store and retrieve if needed (forensics) / when asked
- Comes from one vendor, but industry + IETF standard (not quite, but ...)
- Other concurrent technologies: CRANE, Diameter, LFAP, IPDR, sFlow, ... Yann

# What Is NetFlow?

## A loose definition

A set of packets having common characteristics


## Definition

A flow is a unidirectional set of packets that arrive at the router on the same sub interface, have the same source and destination IP addresses, Layer 4 protocol, TCP/UDP source and destination ports, and the same ToS (type of service) byte in the IP headers

# Cisco's 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



# Anatomy of a flow (v5)

- ❑ 7-tuple key fields

saddr, daddr, sport, dport, L3 proto, ToS,  
input ifIndex

- ❑ Additional fields

Byte count, packet count, start time,  
end time, output ifIndex, TCP flags,  
next hop, src AS, dst AS

# Some characteristics

- ❖ L3-L4
- ❖ Flows are unidirectionals (ingress / egress)
- ❖ Flow-cache comes before ACL lookup
- ❖ Comes at a cost (memory, CPU) for the router
- ❖ Not the perfect solution - but not a lot of other  
Candidates either



# NetFlow flavors

- ❑ Several versions of the export protocol
- ❑ Metering and export keep being worked on
- ❑ Depends on vendor, hardware, software, and combinations of those
- ❑ Even when and where supported, particularities to keep in mind

# Versions

- ✓ v5 - De facto standard
  - (supported by non-C vendors implementing NetFlow)
- ✓ v9 - template-based
  - Flexibility (for now means that the user has to be very flexible)
  - Complexity
  - May have some compelling features - SCTP, IPv6, some L2 fields, additional L3 (ttl, ipid), BGP next-hop, ...
    - ❑ But probably not processed by your collector of choice anyway
  - (Fast) moving target Some features are backported to v5 (e.g. transport independency)
- ✓ v10 aka IPFIX - IETF-blessed standard (in its way to, at least)
  - Based on NetFlow v9 ("forked" from)
- ✓ Other - v1 (obsolete), v7 (switches), v8 (router-based aggregation)

# 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.
- 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

# Hardware peculiarities

- ✓ Support is not equal on all devices
  - No support on the switches below the cat 45xx
- ✓ Peculiarities where supported
  - On L3 switches, no TCP Flags, many specificities depending upon type of sup engine, PFC, config, software, ...
  - Big CAUTION on every Cisco doco concerning NetFlow on the 12k

Entering NetFlow enabling command on a Cisco 12000 Series Internet Router causes packet forwarding to stop for a few seconds while NetFlow reloads the route processor and line card CEF tables. To avoid interruption of service to a live network, apply this command during a change window, or include it in the startup-config file to be executed during a router reboot



# Network Flows

- Unidirectional or bidirectional.
- Bidirectional flows can contain other information such as round trip time, TCP behavior.
- Application flows look past the headers to classify packets by their contents.
- Aggregated flows – flows of flows.

# Cisco NetFlow

- Unidirectional flows.
- IPv4 unicast and multicast.
- Aggregated and unaggregated.
- Flows exported via UDP.
- Supported on IOS and CatOS platforms.
- Catalyst NetFlow is different implementation.

# Working with Flows

- Generate the flows from device (usually a router).
- Export flows from the device to collector
  - Configure version of flows
  - Sampling rates
- Collect the flows
  - Tools to Collect Flows - Flow-tools, NFDump
- Analyze them
  - More tools available, can write your own

# Flow Descriptors

- A Key with more elements will generate more flows.
- Greater number of flows equals:
  - More post processing time to generate reports
  - more memory and CPU requirements for device generating flows
  - More storage needed on the flow processing server
- Depends on application. Traffic engineering vs. intrusion detection.

# Flow Accounting

- Accounting information accumulated with flows.
- Packets, Bytes, Start Time, End Time.
- Network routing information – masks and autonomous system number.

# Flow Generation/Collection


## **Passive monitor**

- A passive monitor (usually a Unix host) receives all data and generates flows.
- Resource intensive

## **Router or other existing network device**

- Router or other existing devices like switch, generate flows.
- Sampling is possible
- Nothing new needed

# Passive Monitor Collection




# Passive Collector

- Using passive collection, not all flows in the network will be seen as opposed to collection from the router
- The collector will only see flows from the network point it is connected on
- However this method 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



# Router Collection



# Router Collection

- With this method, all flows in the network 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
- Also, if there is a router on each LAN, netflow can be activated on those routers to reduce the load on the core router

# Cisco IOS Configuration

- Configured on each input interface.
- Define the version.
- Define the IP address 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 IOS Configuration

```
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

- 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

- 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
- Exporting aggregated flows

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



# Flows and Applications

# Uses for NetFlow

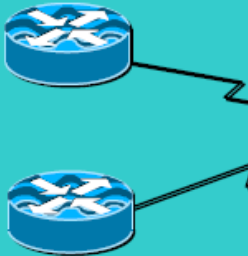




- Problem identification / solving
  - Traffic classification
  - DoS Traceback (some slides by Danny McPherson)
- Traffic Analysis and Engineering
  - Inter-AS traffic analysis
  - Reporting on application proxies
- Accounting (or billing)
  - Cross verification from other sources
  - Can cross-check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



# Uses for NetFlow

<b>Service Provider</b>	<b>Enterprise</b>
<b>Peering arrangements</b>	<b>Internet access monitoring (protocol distribution, where traffic is going/coming)</b>
<b>Network Planning</b>	<b>User Monitoring</b>
<b>Traffic Engineering</b>	<b>Application Monitoring</b>
<b>Accounting and billing</b>	<b>Charge Back billing for departments</b>
<b>Security Monitoring</b>	<b>Security Monitoring</b>

# Uses for NetFlow

Network Layer	Access	Distribution	Core	Distribution	Access
					
Applications	<ul style="list-style-type: none"> <li>• Attack Mitigation</li> <li>• User (IP) monitoring</li> <li>• Application monitoring</li> </ul>	<ul style="list-style-type: none"> <li>• Billing</li> <li>• Chargeback</li> <li>• AS Peer Monitoring</li> </ul>	<ul style="list-style-type: none"> <li>• Traffic Engineering</li> <li>• Traffic Analysis</li> </ul>	<ul style="list-style-type: none"> <li>• Billing</li> <li>• Chargeback</li> <li>• AS Peer Monitoring</li> </ul>	<ul style="list-style-type: none"> <li>• Attack Mitigation</li> <li>• User (IP) monitoring</li> <li>• Application monitoring</li> </ul>
NetFlow Features	<ul style="list-style-type: none"> <li>• Aggregation Schemes (v8)</li> <li>• “show ip cache flow” command</li> <li>• Arbor Networks</li> </ul>	<ul style="list-style-type: none"> <li>• NetFlow MPLS Egress Accounting</li> <li>• BGP Next-hop (v9)</li> <li>• Multicast NetFlow (v9)</li> </ul>	<ul style="list-style-type: none"> <li>• MPLS Aware NetFlow (v9)</li> <li>• BGP Next-hop (v9)</li> <li>• Sampled NetFlow</li> </ul>	<ul style="list-style-type: none"> <li>• NetFlow MPLS Egress Accounting</li> <li>• BGP Next-hop (v9)</li> <li>• Multicast NetFlow (v9)</li> </ul>	<ul style="list-style-type: none"> <li>• Aggregation Schemes (v8)</li> <li>• “show ip cache flow” command</li> <li>• Arbor Networks</li> </ul>

# Tracking Users

**Who are the top users?  
How long are the users on the network?**

**What Internet sites do they use?  
Where do the users go on the network?**

**What percentage of traffic do they use?  
What applications do they use?  
What are the user usage patterns?**

# **NetFlow for Security: Flow Information Helps Mitigate Attacks**

- ✓ **Identify the attack**

  - Count the Flows**

  - Inactive flows signal a worm attack**

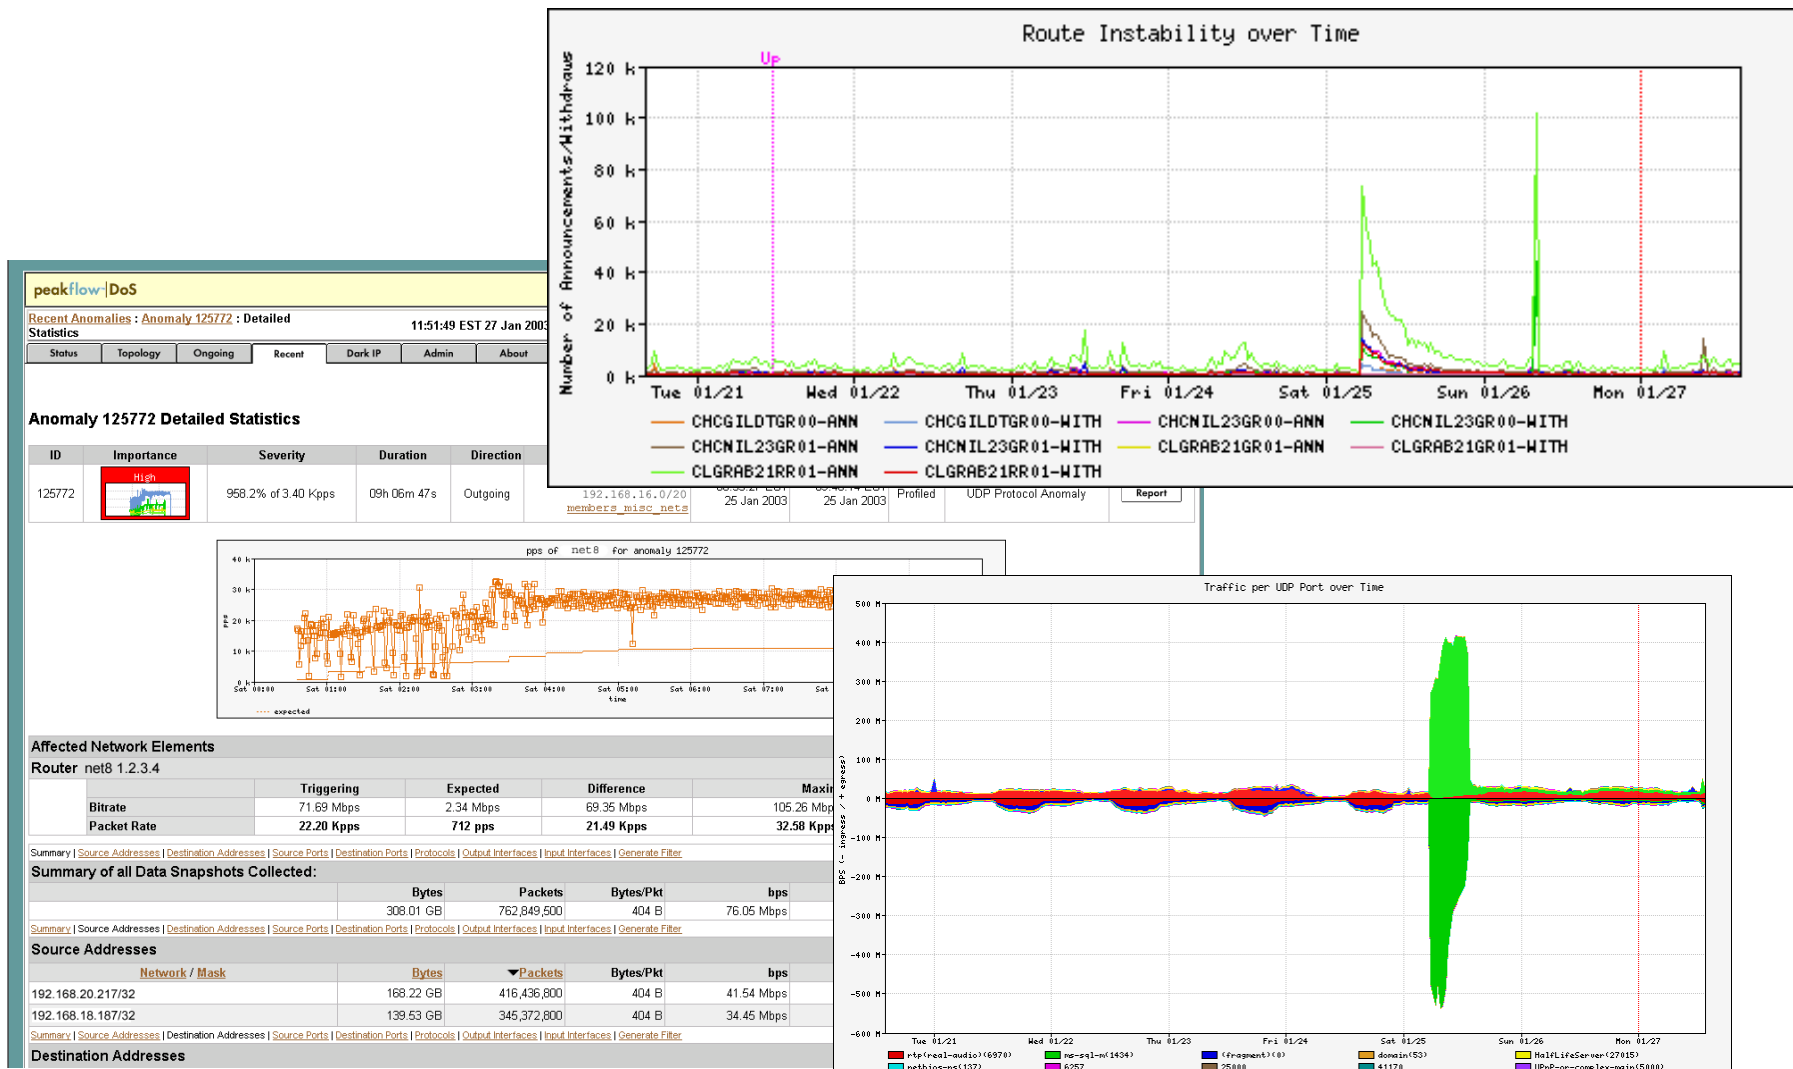
- ✓ **Classify the attack**

  - Small size flows to same destination**

  - What is being attacked and origination of attack**

- ✓ **Cisco IT prevented SQL slammer at Cisco by watching flows per port**

# Detect Anomalous Events: SQL “Slammer” Worm\*




# Billing

- ✓ **Flat-rate billing does not necessarily scale**  
Competitive pricing models can be created with usage-based billing
- ✓ **Usage-based billing considerations**
  - Time of day
  - Within or outside of the network
  - Application
  - Distance-based
  - Quality of Service (QoS) / Class of Service (CoS)
  - Bandwidth usage
  - Transit or peer
  - Data transferred
  - Traffic class

# NetFlow – Peering Agreement

Public Routers 1, 2, 3  
Month of September  
Outbound Traffic



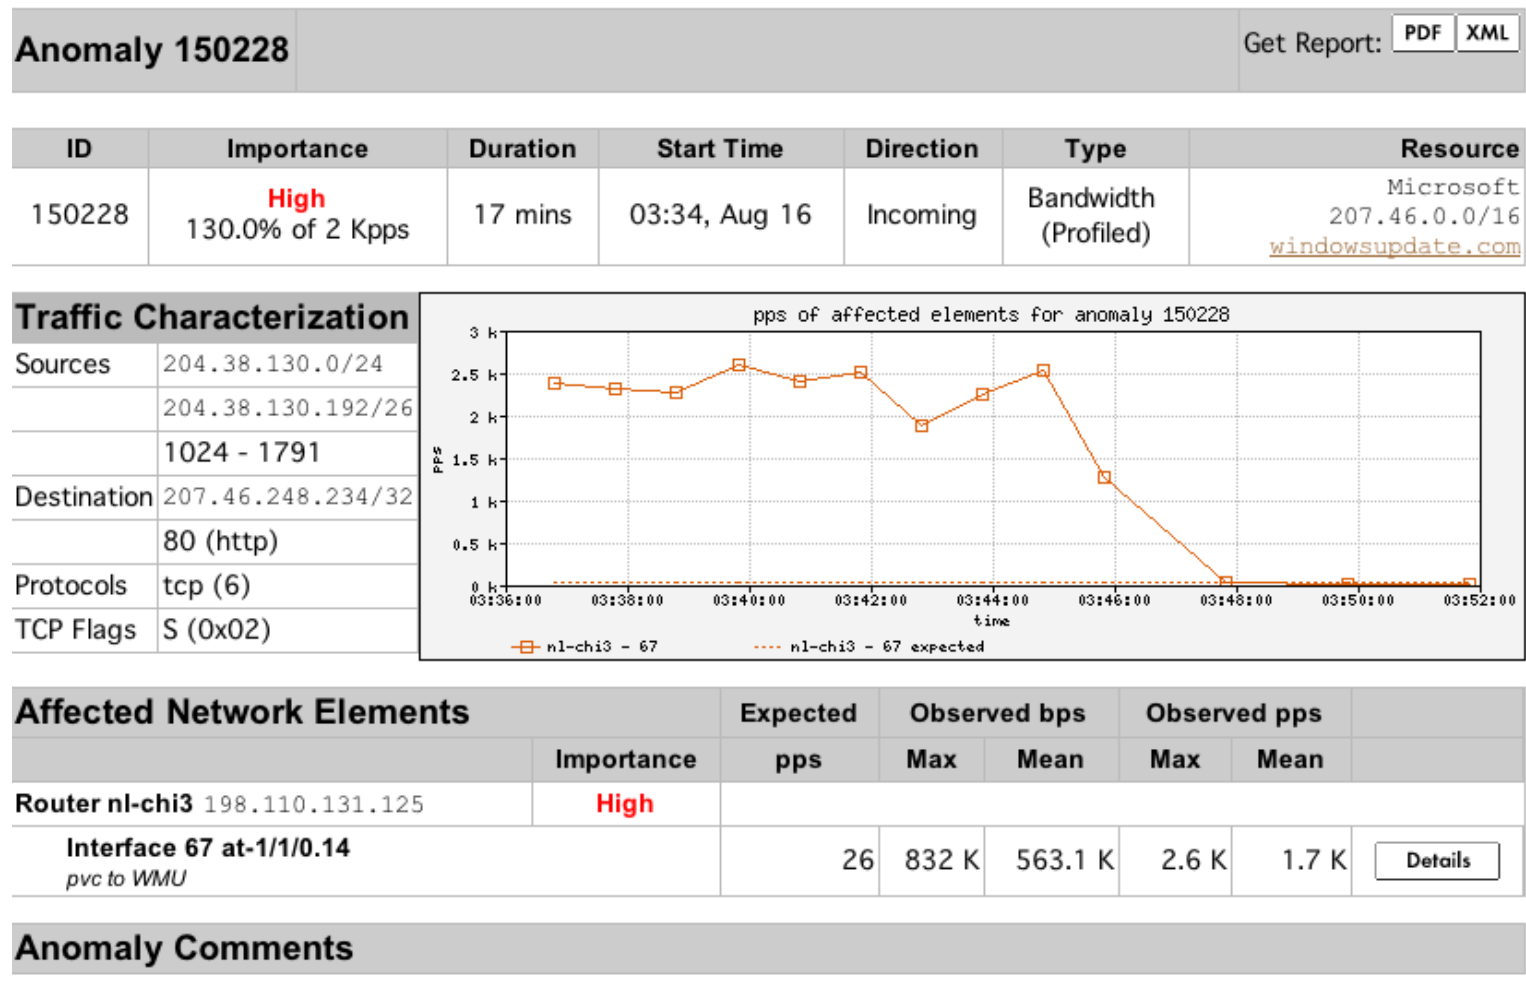
# Flow-based Detection (cont)\*

Once baselines are built anomalous activity can be detected

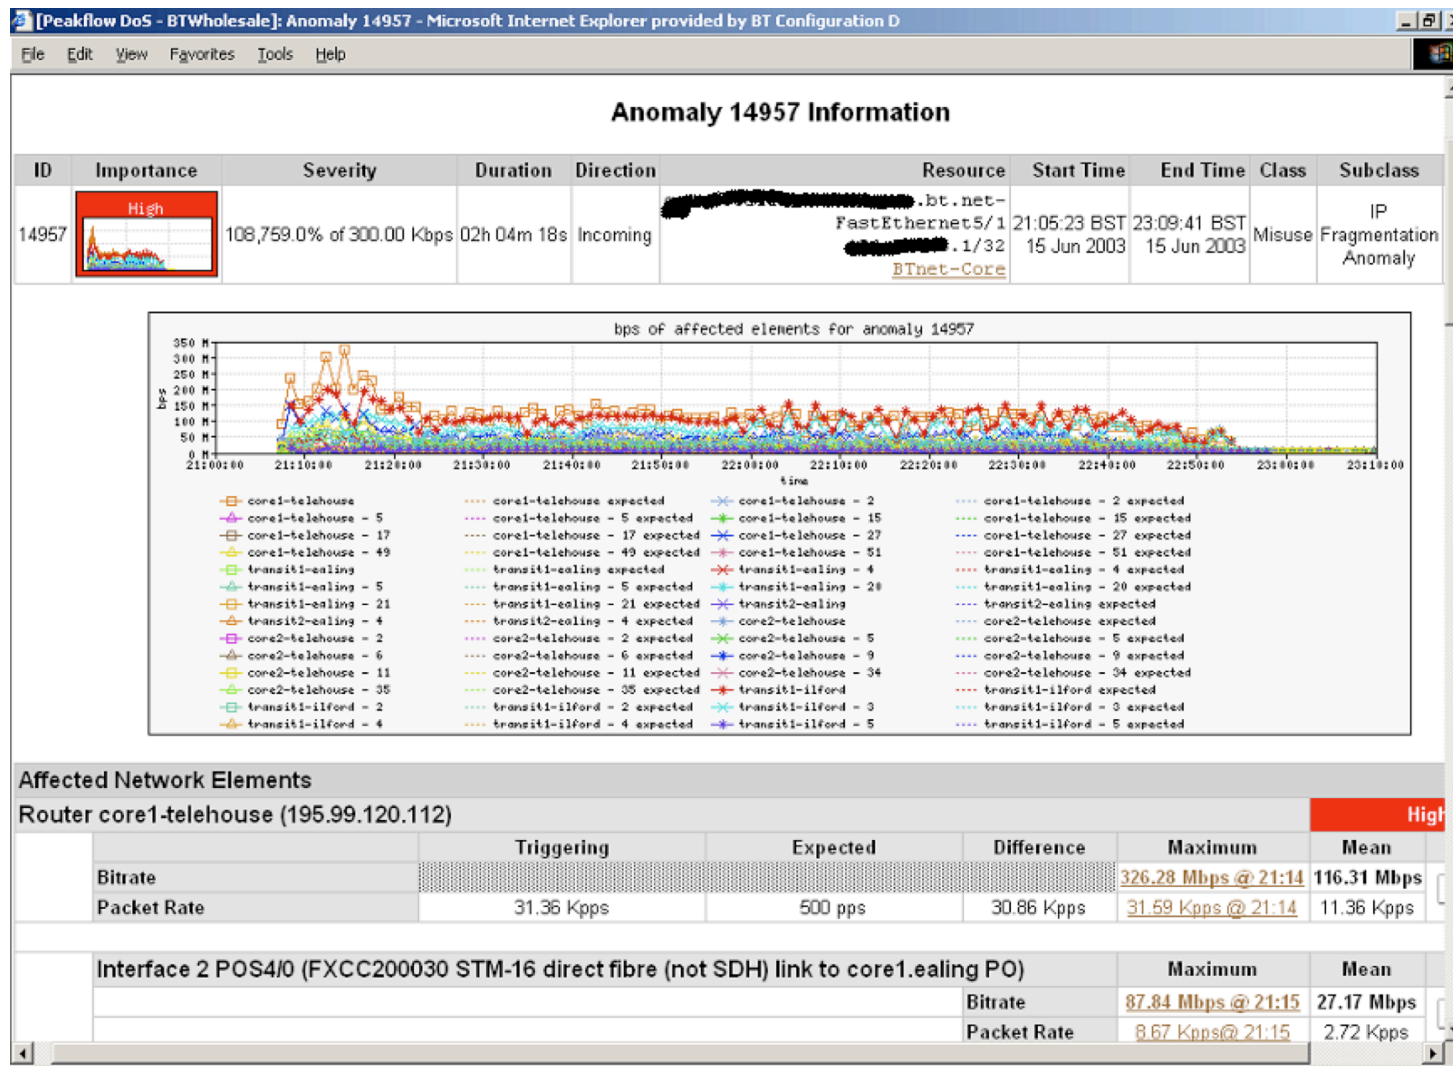
- Pure **rate-based** (pps or bps) anomalies may be legitimate or malicious
- Many **misuse** attacks can be immediately recognized, even **without** baselines (e.g., TCP SYN or RST floods)
- **Signatures** can also be defined to identify “interesting” transactional data (e.g., proto udp and port 1434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



# Flow-based Commercial Tools...\*





# Commercial Detection: A Large Scale DOS Attack



# Accounting

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



# Capacity Planning

**“Capacity planning is the process of determining the network resources required to prevent a performance or availability impact on business-critical applications.”**

# Capacity Planning

- ✓ **Key areas to monitor**

- Application usage

- Identify which applications consume bandwidth

- Who are the top ten nodes that consume bandwidth

- ✓ **Output data circuit forecasts**

- ✓ **Current network utilization and capacity being used**

# References

- flow-tools:  
<http://www.splintered.net/sw/flow-tools>
- Wikipedia:  
<http://en.wikipedia.org/wiki/Netflow>
- NetFlow Applications  
<http://www.inmon.com/technology/netflowapps.php>
- Netflow HOW-TO  
<http://www.linuxgeek.org/netflow-howto.php>
- IETF standards effort:  
<http://www.ietf.org/html.charters/ipfix-charter.html>

# References

- Abilene NetFlow page  
<http://abilene-netflow.itec.oar.net/>
- Flow-tools mailing list:  
[flow-tools@splintered.net](mailto:flow-tools@splintered.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](http://www.cisco.com/en/US/docs/net_mgmt/netflow_collection_engine/6.0/tier_one/user/guide/user.html)