H5 network

White Paper: NetFlow Collector Performance

Executive summary

NetFlow, like all other equivalent network performance solutions, such as sFlow or IPFIX, is widely supported by the network hardware and software providers. It enables comprehensive and distributed application visibility across the global corporate network with little limitation in terms of coverage and throughput.

Being able to discover and monitor applications traffic processed by each critical network link is a key to enforce SLA and QoS and is probably the must have management solution when deploying large, complex and advanced networks such as MPLS or SD-WAN.

NetFlow relies on two main components:

- NetFlow Exporters such as routers processing the traffic, extracting and sending applications flows data, known as NetFlow tickets, to third party NetFlow Collectors.
- NetFlow Collectors: powerful IT machines responsible for collecting NetFlow tickets and providing the network performance visibility.

Unlike SNMP based network monitoring servers whose performance scales linearly with the network size in terms of count of network devices and is not affected by high network usage conditions, NetFlow collectors must be dimensioned to support the NetFlow ticket throughput and activity peaks growing not only with the count of network devices, but also with the count of active users, applications and servers intensively using the network every business day.

The traditional NetFlow collector approach, inherited from syslog processing, consisting in storing the NetFlow tickets first and post processing the flow data afterwards, works fine for some IT security use cases such as intrusion detection but is not optimized for network performance monitoring use cases where a much higher and more versatile NetFlow tickets throughput is expected.

Relying on more than 10-year experience in network performance monitoring, including zero loss 40Gbps traffic processing, H5 NetFlow collector breakthrough approach enables real time network performance data **while** removing process bottlenecks **and** allowing the NetFlow peaks to be absorbed by the system without data loss.

Flow monitoring & NetFlow input data

As a contention[1] based IT Process, Flow monitoring is subjected to performance limitation. Flow monitoring data is extracted from the network traffic by the network devices[2] responsible for processing the network packets across the network and exported to Flow Monitoring Collectors for aggregation and consolidation.

[1] Receiving unsolicited flow of information without any means to request lost data retransmission

[2] Routers, VRFs, L3-Switches, etc.

Layer-4 sessions and NetFlow tickets vs. IP Conversations and Applications

Typically NetFlow tickets are exported by routers for each Layer-4[3] session detected within the network traffic. User client software like Outlook or Chrome use one or more UDP or TCP[4] Session to access servers like Exchange or Web applications . The exact number of TCP/UDP sessions for one single application ranging from one to hundreds.

Let's take an example of a 1-minute network journey of an average user connected for instance to Teams Visio Conference, reading some email from Exchange and browsing some pages one from the company Web site:

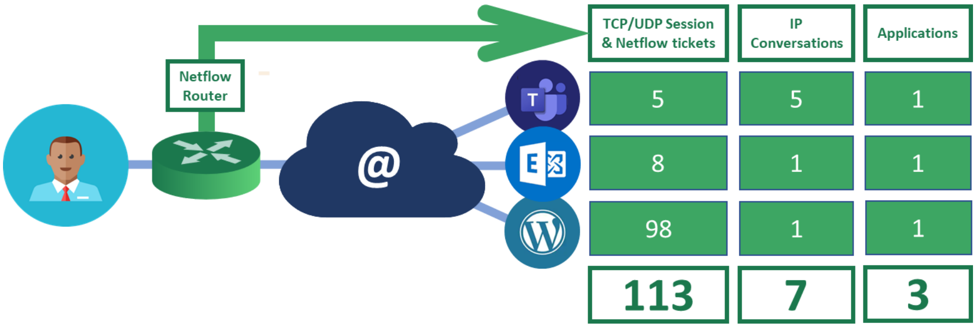

TCP/UDP Session & Netflow tickets	IP Conversations	Applications
5	5	1
8	1	1
98	1	1
113	7	3

Figure 1 : Example of NetFlow Tickets, conversations and applications for a single user during 1-minute interval

The user, yet connected to 3 applications only, still use 113 TCP & UDP connections. The least optimized application being the WordPress Web site using 98 TCP connections[5]. Teams application uses 5 separate IP conversations for two Voice, two for Video and Signaling[6].

This single – and finally not too chatty – user will make his local router export a total of 123 NetFlow tickets in one minute for just 7 IP conversations and 3 applications.

Flow Monitoring Point-of-View

From network monitoring prospective Layer-4 session, the lowest pertinent data granularity are the IP conversations mainly used for application performance trouble shooting. In the example of Figure 1 for instance an expert, while troubleshooting Teams, may need to separate Real Time Voice and Video from Signaling fand check individual IP conversations figures.

The Business user has even less requirements in terms of granularity as she/he will look to larger aggregates like applications. But her/his real time requirement is the same therefore any flow monitoring solutions shall process applications and all other business relevant aggregates with the same priority as IP Conversations.

Count of NetFlow tickets, IP Conversation and Application Cheat Sheet

NetFlow tickets ⬄ Layer-4 Sessions ⬄ UDP Sessions & TCP Connections
IP conversation ⬄ sequence of layer-4 sessions serving specific application purpose ⬄ smaller pertinent network flow monitoring granularity for expert users
Applications ⬄ smaller pertinent flow monitoring aggregate for the business users

[3] Layer-4 protocols also known as IP protocols include TCP, UDP and ICMP

[4] For TCP layer-4 sessions are better known as TCP Connections

[5] This is a simplified count: for the detailed flow analyses: Applications Flows

[6] For detailed teams flow analytics: MS Teams flows

Flow Monitoring Ticket processing

While Flow Monitoring reports and dashboard is based on larger aggregates, NetFlow tickets remains the raw data the Flow monitoring solution is fed with and since NetFlow packets are based on UDP datagrams the ability to capture and process all tickets without loss is critical.

Typical flow monitoring Approach

Many NetFlow monitoring collecting architecture was inherited from Intrusion detection and SIEMS where NetFlow is considered like any other source of network data like syslog for instance and are processed the same way: NetFlow ticket captured on the listening interface are saved on permanent storage drives and the analyses is done latter by a separate process.

This approach is efficient when the primary goal is to look for threats or intrusion signatures – on DMZ to WAN traffic for instance – but is **not optimized** and present significant processing bottlenecks for network performance monitoring where **all NetFlow Tickets** must be processed from **much more** network devices in the in **very large numbers**.

Even worse: for many solutions on the market, HW dimensioning and SW licensing is based on **raw Network tickets throughput** making them much more sensible to the frequent peak of activity inherent to corporate traffic of tens of Gigabit per second.

H5 Approach: On-the-fly NetFlow tickets processing

Thanks to H5 experience in application monitoring, H5-app appliances can handle 40Gbps at wire speed with zero loss, H5-Flow appliance NetFlow processing is optimized to handle large NetFlow ticket throughput:

H5-Flow Collects NetFlow tickets and process all aggregates calculation on-the-fly for one minute and stores the aggregate following data in real time databased every minute.

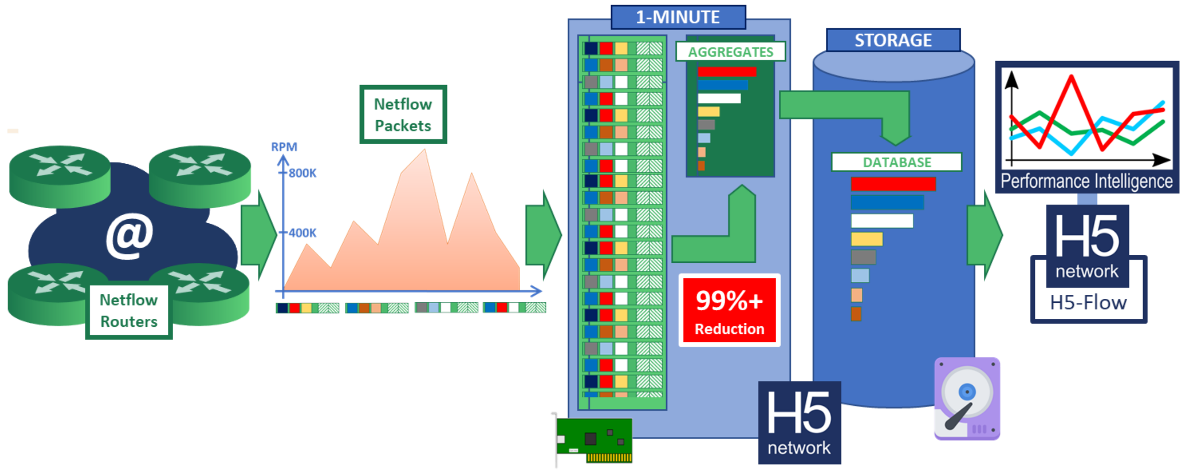

Figure 2: H5-Flow ticket Processing

The following 1st level aggregates are computed in real time every minute and do not need further post-processing before the flow data is stored in the Database:

- *Network Equipments*
- *Network Interfaces*
- *Top IP addresses*
- *Top IP pairs*
- *Monitored Objects : user defined IP@ ranges and lists.*
- *Business Flows : user defined Applications.*

In addition to the 1st level aggregates a significant number of sub aggregates are also computed to meet most data drill-down need, for instance:

Remote Site VRF ⇨ SD-WAN Voice QoS Virtual Interface ⇨ User IP@ ⇨ Teams Calls (udp-3480)

Flow monitoring data may be accessed from many different ways: flat reporting, deep analytics, north bound third-party integration for advanced Data Viz, Automation, Supervision, Ticketing, etc… It is therefore critical that all 1st level and sub-level aggregates are process in **before** data storage instead of being processed when the reporting requests are made. In H5-Flow data is stored in a form that is ready to be rendered and displayed.

H5-Flow approach also removes the most painful storage bottlenecks :

- Storage media stress is reduced by more than 99%
- Minute-data database insertion process is 100% separated from the NetFlow ticket listening process
- Sub minute NetFlow ticket Throughput peaks are absorbed flawlessly with Zero data loss

Finally, H5-Flow approach is scalable over the lifetime of the flow monitoring deployment: more aggregates or metrics can be introduced without jeopardizing the solution performance.

Flow monitoring performance cheat sheet

NetFlow Tickets collector ⇔ On-the-fly
Flow Monitoring Aggregates Processing ⇔ On-the-fly
Data Storage ⇔ Pre Aggregated data
Reporting, Analytics, API ⇔ render ready data with Zero preprocessing