H5 network performance intelligence

White Paper:
H5 NetFlow Zabbix GO Plugin

Executive summary

Netflow, like most other equivalent network performance technologies, such as sFlow or IPFIX, is widely supported by network devices such as routers, firewalls or packet gateways. It enables comprehensive and distributed application traffic 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 network segment is a key to enforce SLA and QoS and is a must have management solution when deploying large, complex and advanced networks such as MPLS or SD-WAN.

Zabbix, leader in SNMP based network monitoring solutions, provides Network managers with a wide and exhaustive ana extensive visibility and can monitor networks ranges tens of thousands of nodes, links and interfaces. With Zabbix Server you can, for instance, easily check whose of your WAN links are overloaded or dropping traffic. With H5 you get, in addition of native Zabbix network interface level charts, flow level charts that will not only tell you **if** your links are loaded but also **who** – network users – and **how** – network applications – is actually loading the link.

To achieve this, NetFlow collectors must be dimensioned to support the Netflow ticket throughput during activity peaks, the activity peaks growing exponentially as the network devices, active users and applications and servers grow.

NetFlow based Network monitoring is slightly different from SNMP based Network monitoring in large networks: In order to provide visibility for networks that can range to tens of thousands of nodes processing millions flows and hundred gigabit of traffic, H5 has a two steps approach: collecting NetFlow data on a separate virtual or bare metal appliance: H5-Flow, and make pre-aggregated data available to Zabbix Server though a dynamic, distributed, high performance and reliable GO plugin.

H5-Flow and H5 NetFlow Zabbix GO Plugin can be deployed in minutes and immediately feed your Zabbix Server with the NetFlow data collected from your L3/L4 network devices[1]. Please contact your H5 sales representative or partner to get a free unlimited trial.

NetFlow data collection

NetFlow data collection performance

On the receiving end of a high and volatile throughput of NetFlow tickets, data collection is a typical streaming contention[2] based IT Process, therefore it is subjected to specific constraints and performance requirements that usual shall require specific and dedicated processing. Please check H5 White Paper “NetFlow Collector Performance” for more details.

[1] Routers, VRFs, L3-Switches, Packet Gateways, etc.

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

Zabbix, on the other hand, has been designed to collect data from a (very) large number of IT elements either through SNMP pooling or through passive or active agents capable of retaining and retransmitting data and relies on robust protocols like TCP sockets or full REST APIs to meet the performance requirement. Zabbix data collecting process is usually referred as a pooling[3] based Process.

It is notoriously difficult to run concurrent contention and pooling processes in a single machine, therefore a two-stage approach appears much more efficient to meet the performance-reliability-scalability triple processing constraint.

Flow Monitoring data structure

What is true on data collection performance data-in-motion point of view is also true on data storage model data-at-rest point of view. NetFlow is about collecting flow data between network end points. The data model to store flow information is therefore driven by the 5-tupple identifying a flow:

- Sender identifier: IP address, IP ranges, etc.
- Recipient identifier: IP address, IP ranges, etc.
- Application: server TCP (or UDP) Ports
- Session: client TCP (or UDP) Ports
- Flow timestamp

This data model is quite different from element based monitoring solutions such as Zabbix that structure its data model to represent devices – Zabbix hosts – and a pre-defined number counters and metrics – Zabbix items. Even if Zabbix supports and performs well with a very large number of hosts and items: the numbers of items per host is fixed at a given point of time.

The two stage approach outlined above is again pertinent here: a dedicated NetFlow collector instance processing, buffering and aggregating hundreds of thousands records – known as NetFlow tickets – from Network devices and feeding Zabbix Server with much smaller data at host-item aggregation level.

Zabbix & NetFlow monitoring cheat sheet

NetFlow Monitoring	Full Contention
Zabbix Network Monitoring	**SNMP, passive & active Agent Pooling**
NetFlow collection optimal data model	**5-tuple IP flow**
Zabbix data model	**Host and Items**

H5 Zabbix integration approach

Further to performance and data storage considerations above, H5 NetFlow integration has been designed as a two stage process:

- H5-Flow appliance, Bare Metal or Virtual, first collect and aggregates NetFlow tickets from Network L3 and L4 devices
- Aggregated NetFlow data is then uploaded to feed Zabbix Server

The later process is carried out by the H5 Zabbix GO Plugin as shown in Figure 1 below:

[3] Blocking or Non-Blocking Process actively requesting data from an external source.

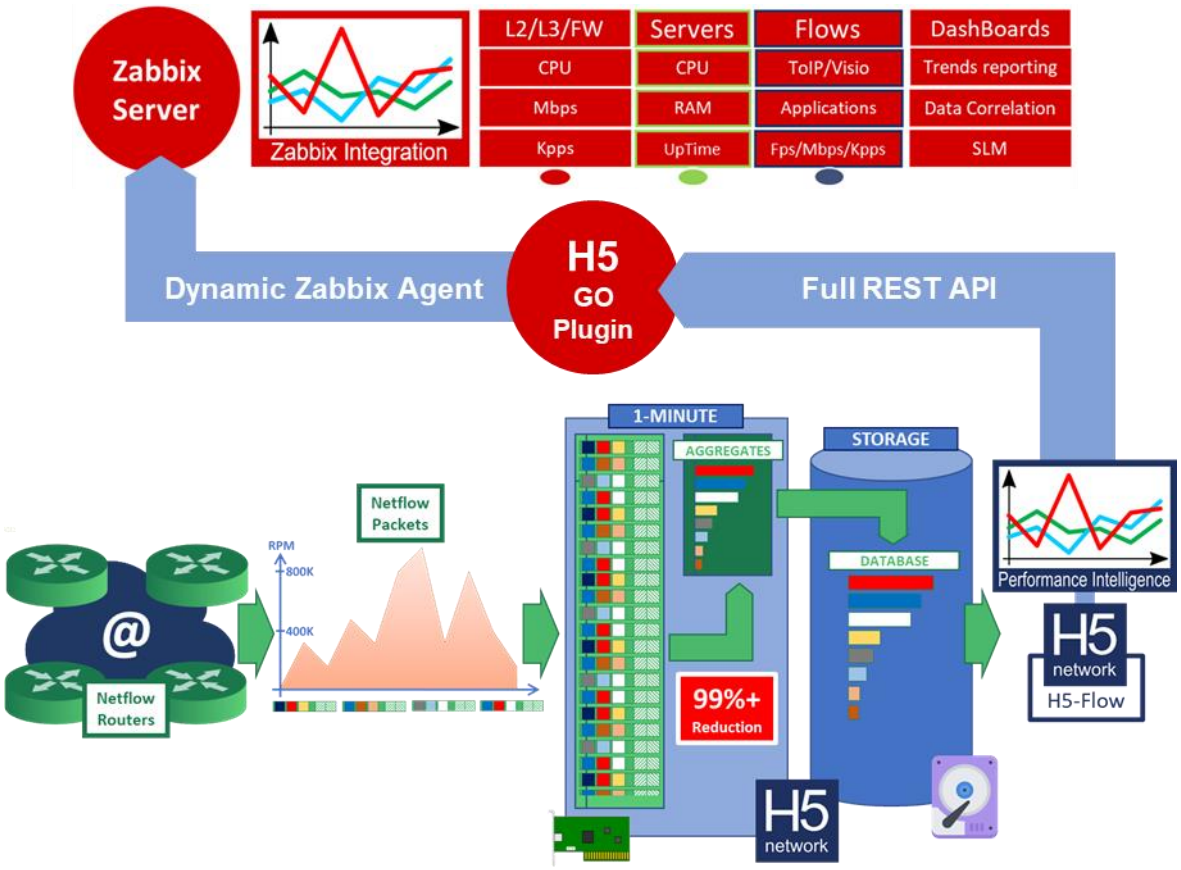

Figure 1: H5-Flow & Zabbix server integration

H5 Zabbix GO Plugin

H5 Zabbix GO Plugin has been developed in Golang and ships as a standard Zabbix Agent V2 plugin. It can be deployed in multiple and distributed locations, the most common being Zabbix Server itself.

Active Plugin

H5 Zabbix GO Plugin is an active Plugin. Active plugin execution (Plugin start, Plugin stop, Item readiness notifications, etc.) are still controlled by Zabbix Server but the actual aggregated NetFlow data is asynchronously pushed to Zabbix Server by the plugin without consuming unnecessary Server resources. This implementation has multiple advantages:

- NetFlow processing performance footprint (RAM, CPU, etc.) on Zabbix Server is very low
- Accurate NetFlow data timestamping is maintained throughout the process
- Peeks of NetFlow data can be absorbed flawlessly with up to 99% NetFlow ticket throughput reduction. NetFlow monitoring capability can be added to existing Zabbix Server/Proxies without upgrading CPU, RAM or Performance nor slowing down other Zabbix processes.
- Support of existing Zabbix Proxies
- Massively distributed Netflow collection architectures are natively supported without additional need for Zabbix Proxies
- Netflow collection is HA by design: natively supporting multiple N+1 redundancy on H5-Flows and Plugins
- Distributed Netflow collection is safe from Network downtimes or slowdowns: Remote H5-Flow being able to buffer up to 30-day Netflow data before pushing it to Central Zabbix Server.

Please check Zabbix blog for more details on active and passive plugins: Zabbix Agent: Active vs. Passive.

Zabbix V6.0 LTS Support

H5 Zabbix GO Plugin has been designed to support Zabbix Agent V6.0 LTS new features including “Support of loading of Agent2 plugins at startup (ZBXNEXT-6688)” but is also fully compatible with Zabbix Agent V5.0 LTS

and Zabbix Agent V5.4. Please contact your H5 support to get more information to deploy H5 Zabbix GO Plugin with Zabbix Agent V5.x.

H5-flow Zabbix items

H5-flow Zabbix GO Plugin dynamically creates, manage and populates a comprehensive Item structure for multiple hosts within the Zabbix Server.

H5 Zabbix integration cheat sheet

H5-Flow GO Plugin Time	Active
Zabbix V6.0 Dynamic loadable Plugin	Yes
Zabbix V5.x support	Yes
Zabbix Proxy support	Yes
Distributed Netflow collection	Yes
Flawless NetFlow Peak Absorption	Yes
High Availability	N+1 Collector redundancy
Network fail safe	30-day