Header Stripping in the Packet Flow Switching Layer

Network monitoring, analysis, and security tools are typically either unable to handle or have limitations handling traffic that has certain tunneling or encapsulation protocols present in the packets. Furthermore, the presence of such protocols in the packets can restrict or limit the ability to apply filtering and flow-based load balancing to the traffic as it is forwarded to specific tools. To address each of these challenges, NETSCOUT® provides features for de-encapsulating or stripping protocols from traffic.

Challenges

Monitoring Tool Limitations

Monitoring, analysis, and security tools are often developed for targeted specific applications and tend to be implemented on general computing platforms with general purpose network interface cards (NIC). Even though the tools function as intended, it can mean that the software and/or hardware implementations are not designed to handle certain protocols.

GRE

GRE is a tunneling protocol that was developed for various generic applications, including mobile core networks and virtual mirror port forwarding (e.g. NVGRE, ERSPAN). A tool not specifically designed for handling GRE may not be able to either recognize GRE, or account for the added tunnel headers, resulting in the tool being unable to analyze the traffic encapsulated inside GRE.

MPLS

In the case of MPLS, a tool may have been designed for handling traffic that contains only one MPLS label. In bridged or cross-provider networking where multiple MPLS labels can be present, the tool may not be able to recognize or account for the additional labels in order to analyze the traffic encapsulated inside MPLS.

GTP

In the case of GTP, this is a specific cellular mobile communications protocol which a tool that was not specifically designed for cellular mobile may not be able to either recognize GTP or account for the added GTP tunnel headers to be able to analyze the traffic encapsulated inside GTP.

VLAN

In the case of VLAN, a tool may have been designed for handling traffic that contains only one VLAN tag and so, in bridged or cross-provider networking where multiple VLAN tags can be present, the tool may not be able to recognize or account for the additional tags to be able to analyze the traffic encapsulated inside VLANs. Also, although VLAN tags can have up to 4094 unique VLAN IDs, in some cases, the number of unique VLAN IDs that can be handled by a tool may be limited to much less than this.

Filtering Limitations

Normal pre-canned filtering, for elements beyond layer 2, is not always able to account for the presence of additional protocol headers between layers 2 and 3, particularly when there are multiple labels or tags.

GRE and MPLS will cause these limitations when there are multiple labels or tags. However, with VLAN tagging, limitations only come into play when there are more than one VLAN tag present or the VLAN tag uses TPID values that are not 0x8100, which implies more than one tag present.

There are two ways to address this:

- Use custom offset filtering which is a little more intricate to use and may have some inherent limitations such as the 128-byte depth limit within the packet that often times renders filtering on L4 and beyond useless for encapsulated packets, or
- Strip off the offending protocol headers so that normal filtering mechanisms can be applied to the un-encapsulated packet

Load Balancing Limitations

Normal layer 3 and layer 4 flow-aware load balancing is unable to account for presence of additional protocol headers between layers 2 and 3, particularly when there are multiple labels or tags.

GRE and MPLS will cause these limitations when there are multiple labels or tags. However, with VLAN tagging, limitations only come into play when there are more than one VLAN tag present or the VLAN tag uses TPID values that are not 0x8100, which implies more than one tag present. Other tunneling protocols will also highlight these limitations as well.

One effective way to address these issues is to strip off the offending protocol headers so that normal balancing mechanisms can be applied to the un-encapsulated packet.

NETSCOUT Solution

Architecture

There are two approaches that users can select from for their monitoring fabric infrastructure, being the traditional integrated software and hardware versus the disaggregated hardware and software options.

The advantage of the nGenius Packet Flow Switch (PFS) solutions is that the header stripping and de-encapsulation itself can be performed by any of the products, regardless of the architecture. It is only if and when a user needs to de-encapsulate the monitored traffic, for one or more monitoring applications, that the decision of which architecture to adopt may become important. 1 and Figure 2 show examples of the two architectural approaches using the nGenius PFS family.

Integrated Architecture Deployments

The nGenius 2200, 4200, and 6000 series packet flow switches offer fully integrated custom hardware architecture where the header stripping and de-encapsulation for the traffic is conducted in the integrated advanced hardware.

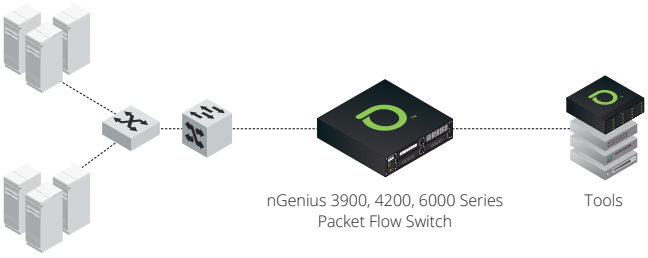

Figure 1: Integrated Hardware & Software with Packet Flow Switch.

Disaggregated Architecture Deployments

The nGenius Packet Flow eXtender (PFX) software is designed to offer header stripping and de-encapsulation when needed. In visibility network deployments with the nGenius 3900 or 5000 series packet flow switches, PFX can be added to perform the task.

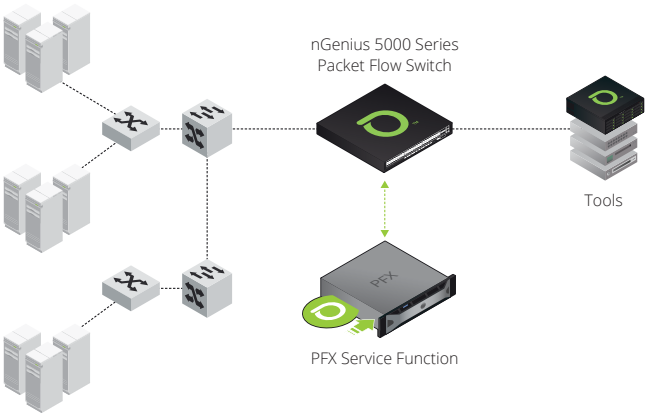

Figure 2: Disaggregated Architecture with PFX for header stripping.

GRE De-encapsulation

Encapsulation in GRE means that a packet's content, inside the layer 2 header, is encapsulated inside new layer 2 (MAC), layer 3 (IP), and optionally layer 4 (usually UDP) headers. These new headers represent the two main network nodes that the GRE tunnels have been established between, and do not bear any direct relation to the actual user as seen in the layer 3 and layer 4 headers inside the GRE encapsulation.

GRE de-encapsulation removes the outer IP and optional UDP headers as well as the GRE header, thereby restoring the packet to what it was prior to GRE encapsulation, except that it retains the same MAC header as the encapsulated packet.

Now filtering and load balancing can be performed on the user session's layer 3 and layer 4 headers and beyond without difficulty.

MPLS Label Stripping & De-encapsulation

MPLS labeling or encapsulation in MPLS (as it is sometimes known) means that a packet's content, inside the layer 2 header, is encapsulated inside one or more MPLS labels (i.e. headers). These labels are used to differentiate this traffic flow from other flows for quality of service (QoS) control, VPN, and other routing purposes, and do not bear any direct relationship to the encapsulated flows themselves. The reason for the presence of multiple MPLS labels is that when traffic from one network that uses MPLS labeling traverses another network,which also uses MPLS labeling, it needs the nested labels for the traversing of more than one network. Stripping, or de-encapsulation, removes all MPLS labels from each packet, including single labels, double-stacked, and n-stacked labels.

GTP De-encapsulation

Encapsulation in GTP means that a packet's content, inside the layer 2 header, is encapsulated inside new layer 2 (MAC), layer 3 (IP), and layer 4 (usually UDP) headers. These new headers represent the two main network nodes (e.g. GGSN and SGSN) that the GTP tunnels have been established between, and do not bear any direct relation to the actual mobile user as the layer 3 and layer 4 headers inside the GTP encapsulation do.

De-encapsulation removes the outer IP and UDP headers as well as the GTP header at line rate, thereby restoring the packet to what it was prior to GTP encapsulation, except that it retains the same MAC header as the encapsulated packet.

Now filtering and load balancing can be performed on the user session's layer 3 and layer 4 headers, and beyond, without difficulty.

VLAN Tag Stripping

VLAN tagging means that a packet's content, inside the layer 2 header, has one or more VLAN tags. These tags are supposed to represent the virtual private networks (VPNs), and do not bear any direct relationship to the tagged flows themselves.

Stripping removes one, two, or all VLAN tags from each packet, depending on VLAN Tagged Packet ID (TPID), including Q-in-Q or bridging VLAN tags. Up to three different TPID values can be specified in order to identify the VLAN tags to be stripped. Thus, a combination of the number of tags to be stripped and the TPID values will determine which tags will be removed from each packet.

Now filtering and load balancing can be performed on the user session's layer 3 and layer 4 headers, and beyond, without difficulty.

NETSCOUT

Corporate Headquarters
NETSCOUT Systems, Inc.
Westford, MA 01886-4105
Phone: +1 978-614-4000
www.netscout.com

Sales Information
Toll Free US: 800-309-4804
(International numbers below)

Product Support
Toll Free US: 888-357-7667
(International numbers below)

NETSCOUT offers sales, support, and services in over 32 countries. Global addresses, and international numbers are listed on the NETSCOUT website at: www.netscout.com/company/contact-us

PFSAN_010-1700 09/2017