Protocol Stripping/De-Encapsulation using NETSCOUT nGenius Packet Flow Switches

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. Competitive offerings do not typically offer de-encapsulation or stripping which makes nGenius® Packet Flow Switch series superior to any competitive offerings.

INTRODUCTION

Typical protocols that present challenges that limit handling of network traffic include Generic Routing Encapsulation (GRE), GPRS Tunneling Protocol (GTP), MPLS (Multi-Protocol Label Switching), VLAN (Virtual Local Area Network), VN (Cisco's Virtual Network)-tag, and many more. To address each of these challenges, NETSCOUT has features for de-encapsulating or stripping protocols from traffic.

Note: GPRS stands for General Packet Radio System, used in data communication across mobile networks.

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.

FabricPath

Cisco® FabricPath is a tunneling protocol that was developed for the Cisco Unified Fabric, as a part of Cisco NX-OS, with the aim of providing a virtual layer 2 switching network for scalability. A tool not specifically designed for handling FabricPath may not be able to either recognize FabricPath, or account for the added tunnel headers, resulting in the tool being unable to analyze the traffic encapsulated inside FabricPath.

GRE

GRE is a tunneling protocol that was developed for various generic applications, including CDMA 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.

GTP

GTP is a specific cellular mobile communications protocol. A tool not specifically designed for cellular mobile may not be able to either recognize GTP, or account for the added GTP tunnel headers, resulting in the tool being unable to analyze the traffic encapsulated inside GTP.

MAC-in-MAC

MAC-in-MAC is a layer 2 tunneling method used for interconnecting multiple service provider networks. A tool not specifically designed for handling MAC encapsulation may not be able to either recognize MAC-in-MAC, or account for the added MAC header, resulting in the tool being unable to analyze the traffic encapsulated inside the double headers.

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.

TRILL

Transparent Interconnection of Lots of Links (TRILL) is a tunneling protocol that was developed for Routing Bridges, providing combination of link-state routing and VLANs. A tool not specifically designed for handling TRILL may not be able to either recognize TRILL, or account for the added tunnel headers, resulting in the tool being unable to analyze the traffic encapsulated inside TRILL.

VLAN

In the case of VLAN, a tool may have been designed for handling traffic that contains only one VLAN tag. 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 needed to analyze the traffic encapsulated inside VLANs. Also, although VLAN tags can have up to 4094 unique VLAN IDs, the number of unique VLAN IDs that can be handled by a tool may be limited to much less than this number.

VN-tag

In the case of VN-tagging, almost no tools have been designed for handling traffic that contains a VN-tag. In virtual environments that utilize Cisco's Nexus virtual distributed switching fabric, VN-tags are used to indicate the source and destination virtual machines that

the packet is flowing between, When this traffic is forwarded out of the Nexus from a virtual SPAN or TAP port, the traffic will still have this VN-tag in each packet and the tools are unlikely to be able to recognize or account for this tag and there are unable to analyze the traffic from this virtual environment.

VXLAN

Virtual Extensible LAN (VXLAN) is a tunneling protocol that was developed for virtual environments which encapsulates the full layer 2 and up packets. A tool not specifically designed for handling VXLAN may not be able to either recognize VXLAN, or account for the added tunnel headers, resulting in the tool being unable to analyze the traffic encapsulated inside VXLAN.

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, GTP, 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.

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, GTP, 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.

Removing Protocol Headers

To remove various protocol headers from traffic, the protocol de-encapsulation and stripping hardware options have support for specific protocols (GRE, GTP, MPLS, VLAN, VN-tag) and a generic capability for other protocols. The current implementation of the generic capability has four pre-canned protocols defined (FabricPath, MAC-in-MAC, TRILL, VXLAN).

FabricPath De-encapsulation

Encapsulation in FabricPath means that a packet's content, including the layer 2 header, is encapsulated inside an outer MAC header and FabricPath header. These headers are used to differentiate this traffic flow from other flows for routing purposes, and do not bear any direct relationship to the encapsulated flows themselves.

Generic de-encapsulation identifies the packet by EType 0x8903 and then removes the outer MAC and FabricPath headers 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. It also provides the ability to distinguish the flows from one FabricPath to another.

Figure 1: De-encapsulating a FabricPath tunneled packet.

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 mobile 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 at line rate; 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.

Figure 2: De-encapsulating a GRE tunneled packet.

Figure 3: De-encapsulating a NVGRE Tunneled Packet.

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 seen in the layer 3 and layer 4 headers inside the GTP encapsulation.

GTP 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.

Figure 4: De-encapsulating a GTP tunneled packet.

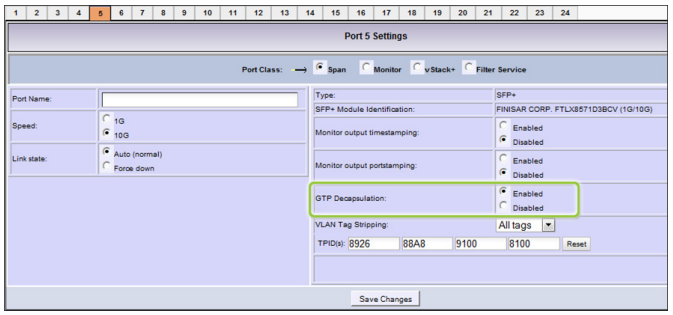

Figure 5: GTP De-encapsulation Selection in the GUI.

MAC-in-MAC Stripping

Encapsulation in MAC or layer 2 means that a packet's content, including the layer 2 header, is encapsulated inside an additional outer MAC header. The headers are used to differentiate this traffic flow from other flows for routing purposes, and do not bear any direct relationship to the encapsulated flows themselves.

Generic de-encapsulation identifies the MAC-in-MAC packet by EType 0x88E7 and then removes the outer MAC header 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. It also provides the ability to distinguish the flows from one outer MAC to another.

Figure 6: De-encapsulating a MAC-in-MAC Tunneled Packet.

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 (which 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. There are two types of MPLS:

- MPLS for Layer 3, i.e. MPLS-L3
- MPLS for Layer 2, i.e. MPLS-L2

MPLS-L3 Labeling

For MPLS-L3, up to nine MPLS inner most labels can be specified to have specific EtherType and/or MAC address values inserted into the MAC header (0x0800 for IPv4 is the default EtherType), which helps to provide some reference information regarding the original MPLS labeling as well as facilitate label-based filtering and MAC address filtering (the inner most label is usually the first encapsulation). MPLS labeling does not contain the original layer 3 EType, so it cannot easily be automatically deduced once all the MPLS labels have been removed.

Now filtering and load balancing can be performed on the user session's layer 3 and layer 4 headers, and beyond without difficulty. It also provides the ability to distinguish the flows from one MPLS label to another.

Figure 7: Stripping a MPLS-L3 labeled Packet.

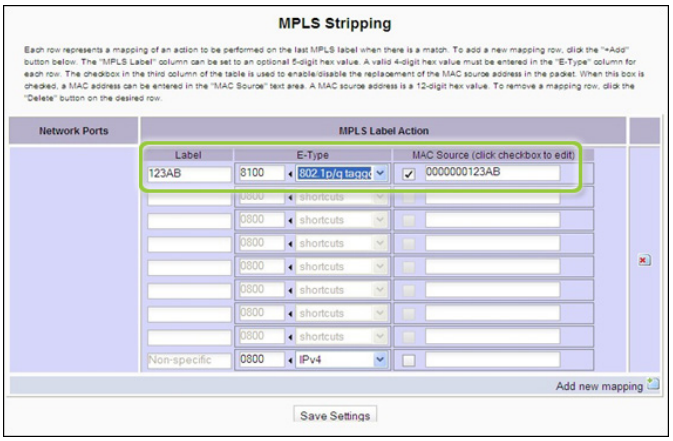

Figure 8: MPLS Stripping/De-encapsulation Configuration in the GUI.

MPLS-L2 Encapsulation

For MPLS-L2, also known as Psuedowire, the outer MAC header and all MPLS labels are removed from each packet, including the optional Control Word which may occur in legacy versions of pseudowire. This then leaves the original packet as it was prior to encapsulation.

Figure 9: De-encapsulating a MPLS-L2 Encapsulated Packet.

TRILL De-encapsulation

Encapsulation in TRILL means that a packet's content, including the layer 2 header, is encapsulated inside an outer MAC header and TRILL header. These headers are used to differentiate this traffic flow from other flows for routing purposes, and do not bear any direct relationship to the encapsulated flows themselves.

Generic de-encapsulation identifies the TRILL packet by EType 0x22F3 and then removes the outer MAC and TRILL headers from each packet, returning the packet to its original form prior to encapsulation.

Now filtering and load balancing can be performed on the user session's layer 3 and layer 4 headers, and beyond without difficulty. It also provides the ability to distinguish the flows from one TRILL to another.

Figure 10: De-encapsulating a TRILL Tunneled Packet.

VLAN and VN Tag Stripping

VLAN Tags

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 four 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. Typical TPID values are 0x8100 for single VLAN tags and 0x88A8 for Q-in-Q or multiple VLAN tags.

Figure 11: Stripping a VLAN-tagged Packet.

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

VN-Tags

Similarly, VN-tags can also be removed by defining a TPID or EType value of 0x8926.

Figure 12: Stripping VN_Tag packet.

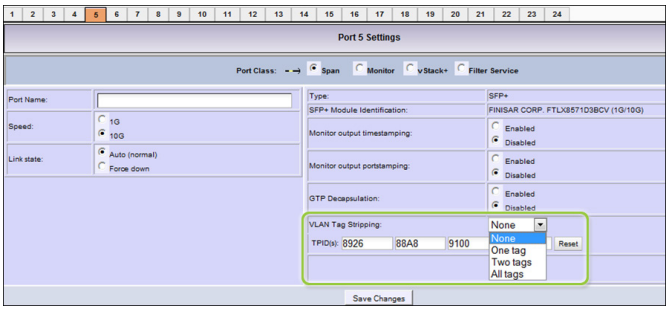

Figure 13: VLAN & VN-Tag Stripping Configuration in the GUI.

VXLAN De-encapsulation

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

Generic de-encapsulation identifies the packet by detecting a MAC header, an IP header (v4 or v6), a UDP header, and a UDP port number of 4789 (0x12B5). It then removes the outer Ethernt, IP, and UDP headers as well as the VXLAN header at line rate; thereby restoring the packet to what it was prior to VXLAN encapsulation.

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

Figure 14: De-encapsulating a VXLAN Tunneled Packet.

PLATFORM SUPPORT

These de-encapsulation and stripping capabilities are supported on the advanced packet-processing chassis modules and line cards of the PFS 2204, PFS 4204, PFS 6010, VB220, VB420, and VB6000 platforms.

NETSCOUT

Americas East
310 Littleton Road
Westford, MA 01886-4105
Phone: 978-614-4000
Toll Free: 800-357-7666

Americas West
178 E. Tasman Drive
San Jose, CA 95134
Phone: 408-571-5000

Asia Pacific
17F/B
No. 167 Tun Hwa N. Road
Taipei 105, Taiwan
Phone: +886 2 2717 1999

Europe
One Canada Square
29th floor, Canary Wharf
London E14 5DY, United Kingdom
Phone: +44 207 712 1672

NETSCOUT offers sales, support, and services in over 32 countries.

For more information, please visit www.netscout.com or contact NETSCOUT at 800-309-4804 or +1 978-614-4000