_o.png _o.bb _o.png height

Contribution ID: **147** Type: **not specified**

ERSPAN Support for Linux

Tuesday, 13 November 2018 14:35 (35 minutes)

Port mirroring is one of the most common network troubleshooting techniques. SPAN (Switch Port Analyzer) allows a user to send a copy of the monitored traffic to a local or remote device using a sniffer or packet analyzer. RSPAN is similar, but sends and received traffic on a VLAN. ERSPAN extends the port mirroring capability from Layer 2 to Layer 3, allowing the mirrored traffic to be encapsulated in an extension of the GRE (Generic Routing Encapsulation) protocol and sent through an IP network. In addition, ERSPAN carries configurable metadatas (e.g., session ID, timestamps), so that the packet analyzer has better understanding of the packets.

ERSPAN for IPv4 was added into Linux kernel in 4.14, and for IPv6 in 4.16. The implementation includes both transmission and reception and is based on the existing ip_gre and ip6_gre kernel module. As a result, Linux today can act as an ERSPAN traffic source sending the ERSPAN mirrored traffic to the remote host, or an ERSPAN destination which receives and parses the ERSPAN packets generated from Cisco or other ERSPAN-capable switches.

We've added both the native tunnel support and metadata-mode tunnel support. In this paper, we demonstrate three ways to use the ERSPAN protocol. First, for Linux users, using iproute2 to create native tunnel net device. Traffic sent to the net device will be encapsulated with the protocol header accordingly and traffic matching the protocol configuration will be received from the net device. Second, for eBPF users, using iproute2 to create metadata-mode ERSPAN tunnel. With eBPF TC hook and eBPF tunnel helper functions, users can read/write ERSPAN protocol's fields in finer granularity. Finally, for Open vSwitch users, using the netlink interface to create a switch and programmatically parse, lookup, and forward the ERSPAN packets based on flows installed from the userspace.

I agree to abide by the anti-harassment policy

Presenters: TU, William (VMware); ROSE, Greg (VMware)

Session Classification: Networking Track