Implementing the ERSPAN Analytics Feature on Cisco Nexus 6000 Series and 5600 Platform Switches

White Paper

October 2014

Contents

What You Will Learn

This document describes how to enable the Cisco® Encapsulated Remote Switched Port Analyzer (ERSPAN) data analytics feature on Cisco Nexus® 6000 Series Switches and Cisco Nexus 5600 platform switches. It also describes the benefits of this feature for data analytics.

A common way to analyze the network is to send mirrored traffic to a traffic analyzer or to a server with a traffic analyzer application. Sometimes, however, that is not possible because all connected servers are in production. The only way to analyze traffic is to send it to a remote server or traffic analyzer. ERSPAN is a feature that helps forward replicated traffic to a remote site for traffic analysis. ERSPAN does this by replicating data frames on a source device, encapsulating the data in a generic route encapsulation (GRE) header, and sending traffic to a destination device. On the destination device, traffic is decapsulated and sent to a network analyzer or server with a traffic analyzer application.

Introduction

Traffic analysis is not always easy to perform. SPAN cannot always be used to replicate and forward packet to a local port because you may not have a server or traffic analyzer connected directly to the device. RSPAN also cannot always be used because you may lack of control of the whole network to make a clear path for RSPAN traffic to a destination. ERSPAN provides a solution in both cases: traffic can be transported to a remote site and can travel over an unknown network because all data is encapsulated in packets with GRE header.

With ERSPAN, network troubleshooting is easier, because ERSPAN transfers exact copies of frames to a remote site, allowing network administrators to discover the source of the problem. Another advantage of ERSPAN is that you can configure multiple ERSPAN source sessions to send traffic to the same destination and get a clear picture of your network at multiple points.

With ERSPAN, network analysis and troubleshooting of traffic latency in the network also is easier. An ERSPAN type-3 packet contains a Precision Time Protocol (PTP) time stamp, allowing you to see and track latency between two or multiple points in the network. ERSPAN forwards time-stamped traffic to a remote site, where a traffic analyzer can compare packet time stamps and calculate latency.

Concept of ERSPAN

An advantage that ERSPAN brings to data analytics is its capability to transport mirrored traffic to a remote destination. With SPAN support on Cisco Nexus 6000 Series and 5600 platform devices, the traffic analyzer needs to be attached directly to a switch. Sometimes that is not possible, because the source of the SPAN traffic is on the aggregation layer of the network and no servers are directly attached, or because a traffic analyzer cannot be attached directly to the aggregation layer of the network.

ERSPAN can send mirrored traffic to a remote site, over an IP network, and retain data integrity because traffic is encapsulated in the GRE header. On the ERSPAN source, traffic is replicated and encapsulated with the IP destination address of the ERSPAN destination device. When traffic is transferred over the IP network, traffic is encapsulated, protecting the original packet. At the destination, traffic is decapsulated and forwarded to the destination port on the destination switch. Figure 1 shows the ERSPAN topology and packet format during travel through the network.

Figure 1. ERSPAN Topology

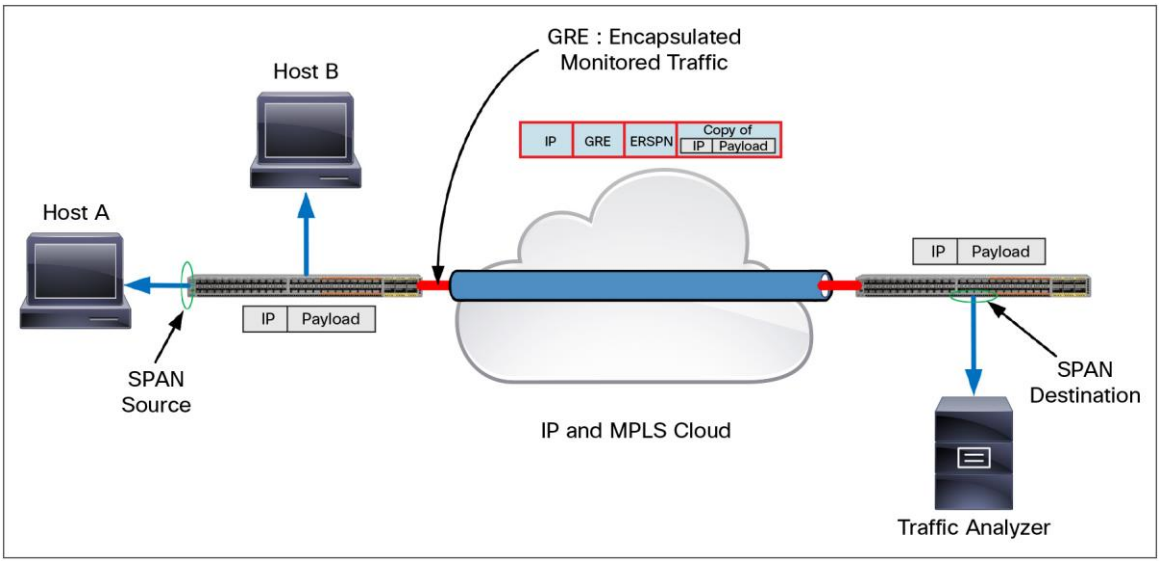

ERSPAN Type-3 Frame Format

Cisco Nexus 6000 Series and 5600 platform devices can encapsulate traffic in ERSPAN type-3 frames with GRE and ERSPAN headers. Parts of the frame have new MAC and IP addresses that refer to the destination of the session. In addition, part of the ERSPAN type-3 frame is a RFC 1588 PTP time stamp, contained in the ERSPAN header. The time stamp is 64 bits and has a time precision value of nanoseconds (Figure 2).

Figure 2. ERSPAN Type-3 Frame

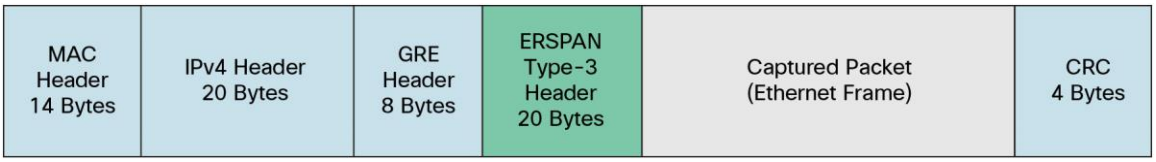

Cisco Nexus 6000 Series and 5600 platform devices can measure latency per device, by using the latency monitoring feature. To allow end-to-end network latency to be measured, packets need time stamps, and a way is needed to compare those time stamps on a single device. One of benefits of the ERSPAN type-3 header is the time stamp, which allows you to measure end-to-end latency using ERSPAN. ERSPAN can forward packets to a unique remote location from multiple points in the network. With RFC 1588 PTP time stamping enabled on the device, each replicated packet will be time stamped, and at the point at which all packets are collected, time stamps can be compared and end-to-end latency can be calculated (Figure 3).

Figure 3. End-to-End Latency Calculation

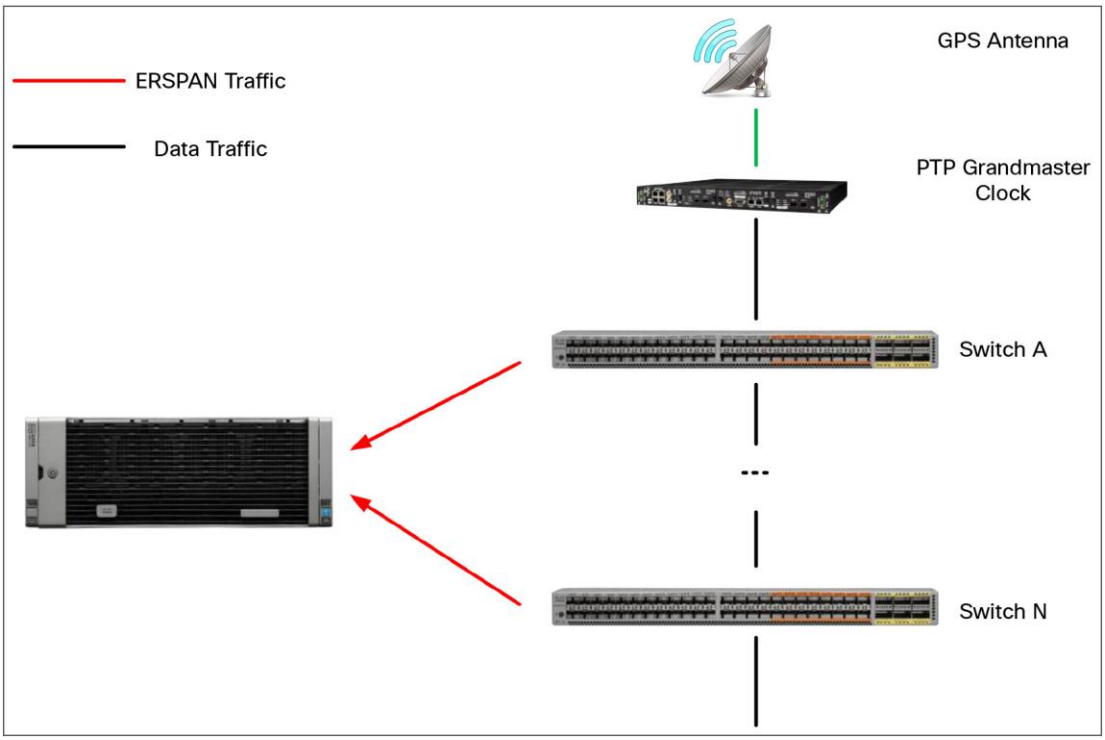

ERSPAN Source Sessions

An ERSPAN source session collects traffic from source ports, VLANs, or VSANs and replicates, encapsulates, and sends traffic to the ERSPAN destination. Each ERSPAN session has a unique session ID and an ERSPAN flow ID that must be unique on the source and destination for ERSPAN traffic. An ERSPAN session destination for the traffic has the IP address of a remote destination. Furthermore, the source ERSPAN session can be assigned to a Virtual Routing and Forwarding (VRF) table. As an optional parameter, the GRE envelope can be assigned IP type-of-service (TOS) and time-to-live (TTL) values.

An ERSPAN source session, by default, monitors all traffic, including multicast and Bridge Protocol Data Unit (BPDU) frames. On a source port, ERSPAN can monitor either ingress or egress traffic or traffic in both directions. For a source VLAN or source VSAN, ERSPAN can monitor only ingress traffic. Furthermore, PortChannels can be the source for an ERSPAN session. If the ERSPAN source is a VLAN, traffic will be mirrored from all interfaces that are members of that VLAN.

On Cisco Nexus 6000 Series and 5600 platform devices, up to 16 active sessions are supported per ERSPAN source device, but hardware can support up to 31 active sessions. An ERSPAN session can support 128 source ports per session and up to 32 source VLANs per ERSPAN source session. Furthermore, both ports and VLANs can be used as sources in the same ERSPAN source session.

To configure an ERSPAN source session, you must define the session ID and session type. These two parameters cannot be changed after they have been entered.

```
switch(config)# monitor session 1 type erspan-source
switch(config-erspan-src)#
```

You next need to define the source for the ERSPAN session: an interface, VLAN, or VSAN from which traffic will be mirrored. By default, traffic is monitored in both directions if the source is a port or PortChannel. To monitor traffic in only one direction, you must specify the transmit (tx) or receive (rx) parameter. To define a port and a PortChannel as the source of the ERSPAN session, use these commands:

```
switch(config-erspan-src)# source interface ethernet 1/1 [tx|rx|both]
switch(config-erspan-src)# source interface port-channel 101 [tx|rx|both]
```

To define a VLAN and a VSAN as the source of the ERSPAN session, use these commands:

```
switch(config-erspan-src)# source vlan 1
switch(config-erspan-src)# source vsan 1
```

The ERSPAN source session needs a defined destination IP address to which to forward traffic. The ERSPAN session can have only one destination IP address. Use this command to configure the destination IP address:

```
switch(config-erspan-src)# destination ip 1.1.1.1
```

Define the ERSPAN ID for the ERSPAN flow with this command:

```
switch(config-erspan-src)# erspan-id 1
```

To define a specific VRF to use instead of the global VRF, use either of these commands:

```
switch(config-erspan-src)# vrf default
```

or

```
switch(config-erspan-src)# vrf vrf-name
```

You can also define optional parameters: access control list (ACL), TTL, maximum transmission unit (MTU) for truncated ERSPAN, and differentiated services code point (DSCP) values.

In the source session, you can configure ACLs to filter packets in the ERSPAN session:

```
switch(config-erspan-src)# filter access-group erspan_acl_filter
```

To limit the number of jumps from the source to the destination and stop traffic from going to an unvented device, you can define the IP TTL. To configure the IP TTL for packets in ERSPAN traffic, use this command:

```
switch(config-erspan-src)# ip ttl 1
```

To prioritize ERSPAN traffic over some less important production traffic, you can change the DSCP value. By default, ERSPAN traffic is assigned a DSCP value of 0. To define a different DSCP value for packets in ERSPAN traffic, use this command:

```
switch(config-erspan-src)# ip dscp 1
```

If ERSPAN traffic is oversubscribing a network link, you can reduce the load by defining an MTU value. The MTU value for truncated ERSPAN packets can be between 64 and 1518 bytes. Use this command:

```
switch(config-erspan-src)# mtu 64
```

The ERSPAN session is by default in the shutdown state. To enable or disable the session, you need to use these commands:

```
switch(config-erspan-src)# no shutdown
switch(config-erspan-src)# shutdown
```

or

```
switch(config)# no monitor session 3 shutdown
switch(config)# monitor session 3 shutdown
```

Furthermore, to bring up an ERSPAN session, you need to specify the origin IP address. The origin IP address is specified in the GRE header as the source IP address. To specify the origin IP address, use this command:

```
Switch(config)# monitor erspan origin ip-address 1.1.1.1 global
```

ERSPAN Destination Sessions

An ERSPAN destination session collects traffic sent by the ERSPAN source session over an IP network and sends it to the destination ports. Before collected mirrored traffic is sent to the destination, the traffic is decapsulated from the GRE envelope. After the destination ports receive the copied traffic, they send it out of the switch to the traffic analyzer or the server with a sniffer application.

Each ERSPAN destination session is defined by a source IP address and ERSPAN ID. The IP address allows multiple source sessions to send traffic to the same destination. The ERSPAN ID allows the destination to identify multiple source sessions with the same destination IP address.

The ERSPAN destination port cannot be configured as a source and as a destination at the same time. The destination port does not participate in spanning-tree instances or Layer 3 protocols. The frames learning option is not supported on the destination port. Fabric extender interfaces are not supported as destination ports in an ERSPAN destination session.

All ports that participate in the ERSPAN destination session need to be configured as monitor ports. To configure a port, apply the **switchport monitor** command to a specific interface:

```
switch(config)# interface Ethernet 1/1
switch(config-if)# switchport monitor
```

To create an ERSPAN destination session, you need to define the session type and ERSPAN ID. The ERSPAN ID should be the same as for the corresponding ERSPAN source session. Use this command to configure the destination session with the session ID and session type:

```
switch(config-if)# monitor session 1 type erspan-destination
```

The next step in configuring the ERSPAN destination session is to assign the source IP address. The address is the same as for the source session, and it determines which traffic gets forwarded to the destination ports. Only one IP address can be assigned per destination session. Use this command:

```
switch(config-erspan-dst)# source ip 1.1.1.1
```

Only physical ports can be assigned as destinations for an ERSPAN destination session. Furthermore, ports in trunk mode can be assigned as destination of the ERSPAN session. A server or network analyzer will be connected to the destination ports. To assign ports to the session, use this command:

```
switch(config-erspan-dst)# destination interface ethernet 1/1
```

You need to configure the ERSPAN ID for the destination session. The ERSPAN ID must be the same as the ERSPAN ID for corresponding ERSPAN source session. To configure the ERSPAN ID, use this command:

```
switch(config-erspan-dst)# erspan-id 1
```

The only supported VRF instance for the ERSPAN destination session is the default:

```
switch(config-erspan-dst)# vrf default
```

To bring up the ERSPAN destination session, you need to run the **no shutdown** command:

```
switch(config-erspan-src)# no shutdown
switch(config-erspan-src)# shutdown
```

ERSPAN with ACLs and Truncated ERSPAN Source Session

Cisco Nexus 6000 Series and 5600 platform devices can have link bandwidth of 1, 10, and 40 Gbps. Therefore, if an ERSPAN source session is mirroring line-rate traffic on the source ports, during traffic forwarding over an unknown IP network some or all traffic may be dropped because of lower link bandwidth between the ERSPAN source and destination ports, or lower bandwidth on the destination ports.

To overcome this problem, you can apply ACLs to the source session. With an ACL applied, only desirable traffic will be forwarded to the destination. Overall bandwidth will be reduced, and lower link bandwidth within the IP network will affect traffic less or not at all.

ACLs can be applied only on an ERSPAN source session. ACLs are not supported on an ERSPAN destination session. If the source session is monitoring one or more ports in both directions, the size of the ACL is limited. If the session is monitoring in only one direction, this limitation does not apply to the ACL.

Note: For more information about ACL limits, see the ERSPAN configuration guide chapter in the Cisco NX-OS System Management Configuration Guide, section ERPSN.

A traffic filter is defined as an ACL and applied to the ERSPAN source session. First, you need to define the ACL with the specific parameter that will filter traffic and a specific ACL ID:

```
switch(config)# ip access-list ERSPAN_ACL
switch(config-acl)# permit ip 10.0.0.0 0.0.0.255 any
switch(config-acl)# deny ip any any
```

To apply a specific ACL to the ERSPAN source session, you need to go into the session and add the ACL as a filter:

```
switch(config)# monitor session 1 type erspan-source
switch(config-erspan-src)# filter access-group ERSPAN_ACL
```

Truncated ERSPAN limits the frame size to the defined MTU value. This feature helps reduce the overall bandwidth of mirrored traffic by taking only a defined number of bytes of the original packet and frame. Furthermore, if any of the devices in an unknown IP network has an MTU limitation, packets exceeding that limitation are dropped. The MTU value can be between 64 and 1518 bytes. By default, truncation is not enabled, and all traffic is forwarded in its original size. If mirrored traffic includes a jumbo frame of 2000 bytes and the defined MTU value for the source session is 1518, the device will forward only the first 1518 bytes of the original packet and frame.

To enable truncated ERSPAN, you need to define it in the source session:

```
switch(config)# monitor session 1 type erspan-source
switch(config-erspan-src)# mtu 1518
```

ERSPAN Session Verification

After you have configured all the parameters, the session is configured. You can use the command **show monitor** to see all the configured sessions on the device:

```
switch# show monitor
Session  State        Reason                   Description
-------  -----------  ----------------------  ------------------------------
1        up           The session is up
```

The output shows the state of the session (up or down) and the reason for this status.

To see more detailed session status information and destination IP address, ERSPAN ID, and session-type information, use the command **show monitor session** with the specific session ID:

```
switch(config)# show monitor session 1
```

The **show monitor session 1** command provides the following output:

```
switch# sh monitor session 1
   session 1
---------------
type              : erspan-source
state             : up
erspan-id         : 147
vrf-name          : default
destination-ip    : 10.23.21.8
ip-ttl            : 255
ip-dscp           : 0
acl-name          : ERSPAN_ACL
mtu               : 1518
origin-ip         : 10.254.254.30 (global)
source intf       :
    rx            : Eth1/47
    tx            : Eth1/47
    both          : Eth1/47
source VLANs      :
    rx            :
source VSANs      :
    rx            :
```

In the output, the first line shows the session ID that was specified by the **show** command. The session type is shown: in this case, a source session. The output also shows the state, and if the state is down, it shows the reason why, making troubleshooting easier. The unique ERSPAN ID and VRF table are also specified in the output. As previously described, if TTL, IP DSCP, and ACL values are configured for the session, they will be specified in the output. The origin IP address specified in the output represents the source address of the new packet. The destination IP address is the unique IP address that specifies where to send traffic, and the path to the destination IP address needs to be in the routing table. The source of the traffic to be mirrored is the interface Ethernet 1/47, and traffic from both directions will be mirrored.

The same command can be used to get output for the destination session. The output for the destination session is different, because fewer parameters are specified by configuration.

```
switch# show monitor session 1
   session 1
---------------
type              : erspan-destination
state             : up
erspan-id         : 147
source-ip         : 10.23.21.8
destination ports : Eth1/47
```

The output for the ERSPAN destination session shows the session ID that was specified by the command. The output also shows the session state; in this case, the session is up. If the session state is down, the output shows the reason why the session is down. The output shows the unique ERSPAN ID for the source and destination sessions. The destination session has IP address as source of traffic. The destination for this session is an interface, more interfaces can be specified as a destination, but in that case, each specified destination is counted as one destination session.

Conclusion

The ERSPAN feature brings the advantages of traffic mirroring to a remote device if no available local ports support a SPAN session. ERSPAN helps you discover the source of unwanted traffic in the network. In addition, each mirrored frame includes an ERSPAN header that contains a time stamp, and allowing you to measure end-to-end latency in the network.

The Cisco Nexus 6000 Series Switches and Cisco Nexus 5600 platform support source and destination ERSPAN sessions and can perform line-rate traffic mirroring. They can support multiple sessions to the same destination, and they can support up to 128 source ports per session. ERSPAN thus allows great flexibility for network traffic analysis.

For More Information

For more information about ERSPAN, see the ERSPAN configuration guide chapter in the Cisco NX-OS System Management Configuration Guide.

For more information about Cisco Nexus 5600 platform switches, see http://www.cisco.com/c/en/us/products/collateral/switches/nexus-5000-series-switches/datasheet-c78-730760.html.

For more information about Cisco Nexus 6000 Series Switches, see http://www.cisco.com/c/en/us/products/switches/nexus-6000-series-switches/datasheet-listing.html.

Americas Headquarters
Cisco Systems, Inc.
San Jose, CA

Asia Pacific Headquarters
Cisco Systems (USA) Pte. Ltd.
Singapore

Europe Headquarters
Cisco Systems International BV Amsterdam,
The Netherlands

Cisco has more than 200 offices worldwide. Addresses, phone numbers, and fax numbers are listed on the Cisco Website at www.cisco.com/go/offices.

Printed in USA C11-733110-00 10/14