



Configuring EVPN Microsegmentation

A secure BGP EVPN VXLAN fabric integrates Cisco TrustSec to provide microsegmentation and group-based policy enforcement.

- [Restrictions for EVPN Microsegmentation, on page 1](#)
- [Information About EVPN Microsegmentation, on page 1](#)
- [How to Configure EVPN Microsegmentation, on page 4](#)
- [Deployment of EVPN Microsegmentation, on page 4](#)
- [Configuration Example of EVPN Microsegmentation, on page 8](#)

Restrictions for EVPN Microsegmentation

Restrictions for EVPN Microsegmentation

- VXLAN-GPO encapsulation is not supported for VXLANv6. It is supported only with IPv4 underlay.
- If the host-facing interface is a trunk interface, you must configure device tracking on that interface.
- Overlay multicast traffic is not subject to SGACL enforcement.

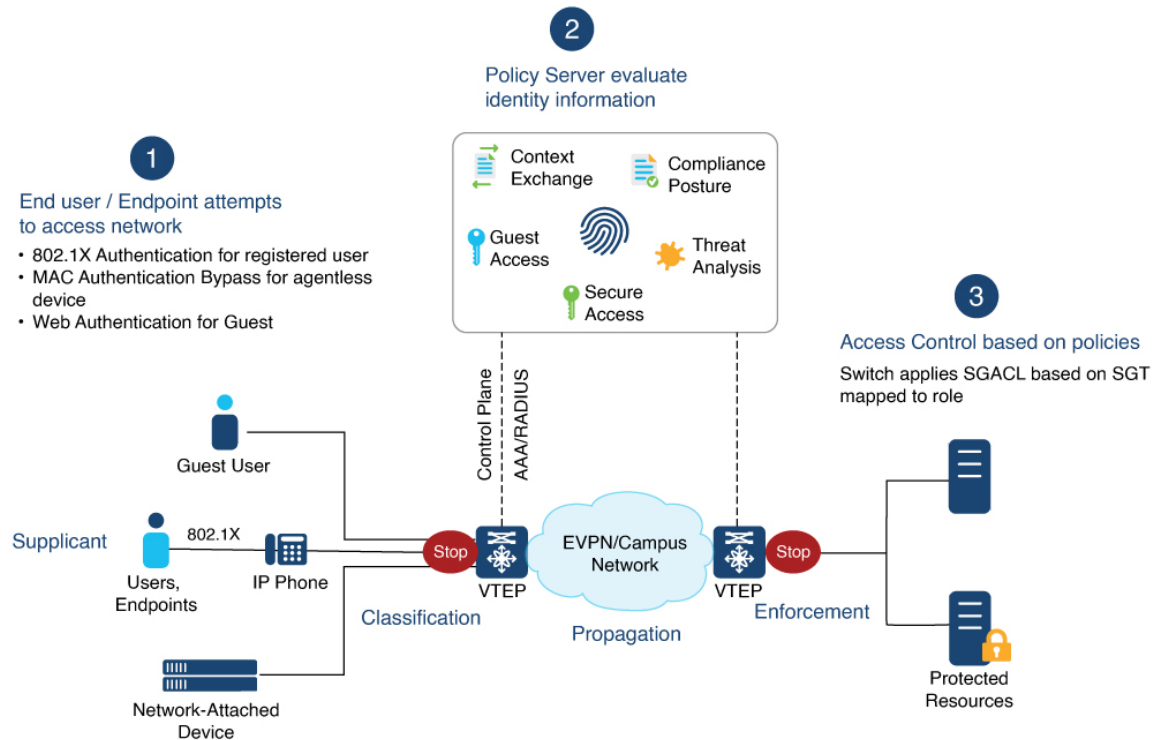
Information About EVPN Microsegmentation

In an enterprise network, users, devices, and applications, all utilize the network to access resources. But the access needs of a user may be different from that of a device in a network. Microsegmentation addresses the need for this isolation of network access. Using Security Group-Based Access Control, the endpoints within the overlay network can be permitted access to specific resources and denied access to others, based on their group membership.

BGP EVPN VXLAN integrates Cisco TrustSec to provide microsegmentation and end-to-end access control with propagation of the security group tag (SGT). Using security group access control lists (SGACLs), a network administrator can control the operations that users can perform based on their security group assignments and destination resources in a VXLAN campus fabric.

For more information on Cisco TrustSec, refer to “Cisco TrustSec Overview” in the [Cisco TrustSec Configuration Guide](#).

Figure 1: Secure Fabric: BGP EVPN VXLAN Fabric Microsegmentation



Cisco TrustSec solution encompasses many aspects as discussed in the Cisco TrustSec Overview module. The three fundamental components of CiscoTrustSec are: Classification, Propagation, and Enforcement.

Classification

When users or devices (endpoints) connect to a network, the network assigns it a specific security group tag (SGT). This process is called Classification and can be either static or dynamic assignment. Static assignment is done by associating the SGT with an IP, VLAN, or port-profile. Dynamic assignment of the SGT tag is based on the results of authentication of the endpoints and downloaded as an authorization option from Cisco Identity Services Engine (ISE).

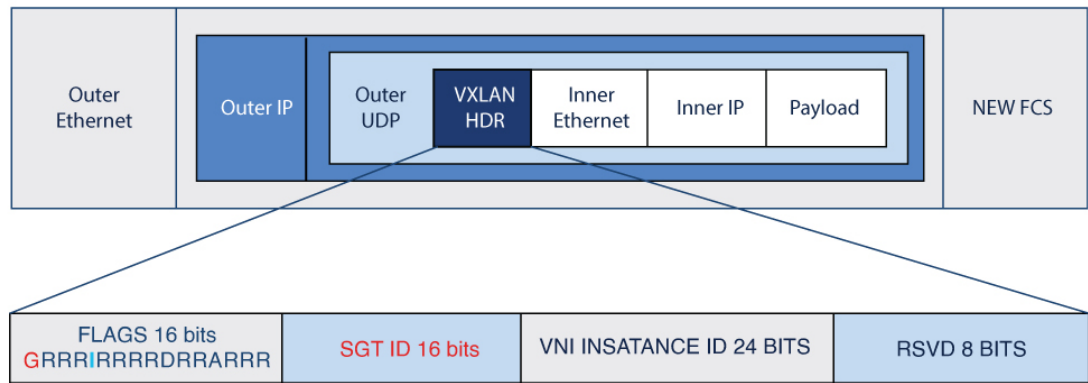
For information about ISE, refer to [Cisco Identity Services Engine](#).

Propagation

An SGT is a metadata value that is transmitted in the header of the VXLAN-encapsulated packets. It is a unique 16-bit tag that represents the privilege of the source endpoint.

After the user traffic is classified, the SGT is propagated from the node at which classification took place, to where the enforcement action is invoked. This process is called propagation. Propagation in a BGP EVPN VXLAN network occurs through the VXLAN Group-based Policy Option (VXLAN-GPO), as defined in [VXLAN Group Policy Option - draft-smith-vxlan-group-policy-0](#). SGT propagation through the EVPN VXLAN fabric is described in the later sections of this document.

A typical VXLAN header that contains an SGT tag is as shown:



Flag	Description
G Bit (bit 0)	Group-Based Policy Extension bit: Set to 1 when SGT is carried in the Group Policy ID field.
D bit (bit 9)	Don't Learn bit: Set to 1 to indicate that the egress VTEP must not learn the source address of the encapsulated frame. D bit is set to 0.
A bit (bit 12)	Policy Applied bit: A bit can be set to 0 (treat as reserved).
Group Policy ID	The Group Policy ID is a 16-bit identifier that identifies the source group of the packet encapsulated within VXLAN. Set to SGT (also called EPG Class).
VXLAN Network Identifier (VNI)	VNI is a 24-bit identifier that identifies the virtual network of the communicating hosts. It is set to the egress VNI associated with the destination subnet as determined by routing or bridging.

Enforcement

The enforcement of the policy based on the SGT occurs at an egress device like a firewall, router, or a switch. The enforcement device takes the source SGT and looks it up against the destination SGT to determine if the traffic should be allowed or denied. Security group access control list (SGACL) is one of the policy enforcement methods that provides state-less access control mechanism based on the security association or SGT value.

Security Group Access Control Lists

SGACLs define access control policies based on device identities instead of IP addresses as in traditional ACLs. Hence the network devices are free to move throughout the network and change IP addresses. If the roles and the permissions remain the same, changes to the network topology do not change the security policy. When an endpoint is connected to the device, you simply assign the endpoint to an appropriate security group and the endpoint immediately receives the permissions of that group. Using role-based permissions greatly reduces the size of ACLs and simplifies their maintenance. With Cisco TrustSec, the number of access control entries (ACEs) configured is determined by the number of permissions specified, resulting in a much smaller

number of ACEs than in a traditional IP network. The use of SGACLs typically results in a more efficient use of TCAM resources compared with traditional ACLs. For more information on the SGACL policies and enforcement, refer to “Configuring Security Group ACL Policies” in the *Cisco TrustSec Configuration Guide* for the relevant platform.

How to Configure EVPN Microsegmentation

Procedure

- Step 1** **Classify:** Authenticate an endpoint and authorize its network access by assigning an SGT tag.
- a) For configuration procedures to authenticate an endpoint at ingress, refer to the “Configuring Endpoint Admission Control” chapter in the *Cisco TrustSec Configuration Guide*.
 - b) For configuration procedures to assign an SGT to a device or a user at ingress, refer to the “Configuring SGT Exchange Protocol” chapter in the *Cisco TrustSec Configuration Guide*.

- Step 2** **Propagate:** Enable the propagation of the SGT information using the **group-based-policy** command on the NVE interface.

Example:

```
interface nve10
 no ip address
 source-interface Loopback0
 host-reachability protocol bgp
 group-based-policy
```

After the VXLAN-GPO encapsulation is enabled, it is applicable to all network overlay segments.

- Step 3** **Enforce:** Define the SGACL policies and enforce them through egress filtering of the packets.
- For the procedure to configure SGACL policy, refer to “Configure Security Group ACL Policies” chapter of the *Cisco TrustSec Configuration Guide*.
-

Deployment of EVPN Microsegmentation

The EVPN VXLAN fabric propagates the SGT tag (and thus security policy) in the following deployment scenarios:

- **EVPN Layer 2 Virtual Network Instance (L2VNI)**

An EVPN VXLAN Layer 2 overlay network allows host devices in the same subnet to send bridged traffic to each other. The network forwards the bridged traffic using a Layer 2 virtual network instance (VNI). Layer 2 overlay network traffic supports the SGT propagation.

To configure a Layer 2 overlay, refer to [How to Configure EVPN VXLAN Layer 2 Overlay Network](#).

- **EVPN Layer 3 Virtual Network Instance (L3VNI)**

An EVPN VXLAN Layer 3 overlay network allows host devices in different Layer 2 networks to send Layer 3 or routed traffic to each other. The network forwards the routed traffic using a Layer 3 virtual network instance (VNI) and an IP VRF.

For information about how to configure a Layer 3 overlay network, refer to [How to Configure EVPN VXLAN Layer 3 Overlay Network](#).

- **EVPN VXLAN Distributed Anycast Gateway**

EVPN VXLAN integrated routing and bridging (IRB) allows the VTEPs or leaf switches in an EVPN VXLAN network to perform both bridging and routing. The VTEPs in the network forward traffic to each other through the VXLAN gateways.

Distributed anycast gateway enables the use of the same gateway IP and MAC address across all the VTEPs in an EVPN VXLAN network. This ensures that every VTEP functions as the default gateway for the workloads directly connected to it.

To configure a Distributed Anycast Gateway, refer to [How to Configure EVPN VXLAN IRB using Anycast Distributed Gateway](#)

- **EVPN VXLAN Centralized Anycast Gateway**

A Centralized Anycast Gateway (CGW) VTEP performs the Layer 3 gateway function for all the Layer 2 VNIs. All the other VTEPs in the network perform only bridging. The CGW VTEP acts as the Layer 3 gateway and performs routing for the intersubnet VXLAN traffic.

To configure a CGW, refer to [How to Configure EVPN VXLAN IRB using CGW](#).

- **Multi-Homing**

Multi-homing allows you to connect a host or Layer 2 switch to more than one VTEP in the EVPN VXLAN network. This connection provides redundancy and allows network optimization. Redundancy in the connection with the VTEPs ensures that there's no traffic disruption when there's a network failure.

To configure Multi-homing, refer to [How to Configure Multi-Homing in a BGP EVPN VXLAN Fabric](#).

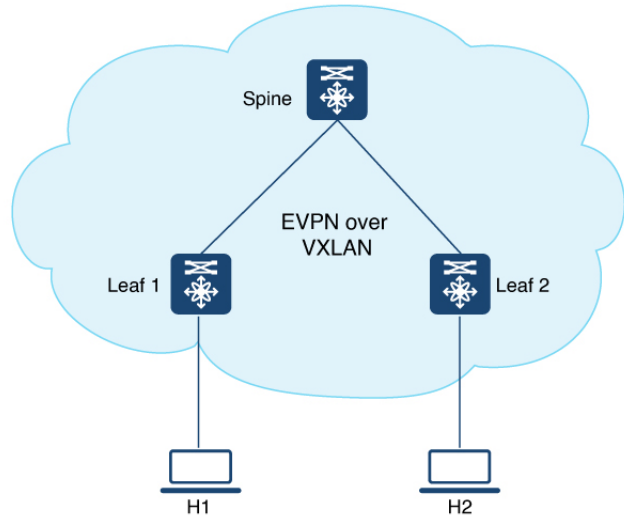
In each of the deployments, the leaf node is located either at the access layer or the distribution layer.

Leaf Node at Access or Distribution Layer

Leaf Node at Access Layer

The CTS solution is enabled on the leaf node to provide classification of packets and secure propagation and enforcement of policies. When a host/endpoint connects to a Leaf node that is located at the access layer:

- The leaf node first authenticates the endpoint using 802.1X protocol.
- On successful authentication, the endpoint is authorized to send and receive traffic. Traffic from the endpoint is classified and tagged with an SGT value.



Configuration of the Leaf Node


The leaf node is configured to provide the functionalities:

- VXLAN-GPO encapsulation
- CTS classification, propagation, and enforcement

Leaf Node at the Distribution Layer

The CTS solution is enabled on the access switch. When an endpoint connects to the switch at the access layer:

- Endpoint is authenticated with the IEEE 802.1X method.
- On successful authentication, the traffic from the endpoint is classified and tagged with an SGT value.
- The packets originating from the endpoints are encapsulated in a Cisco Meta Data (CMD) format at the access layer and forwarded to the leaf node at the distribution layer (L2VNI). The CMD header is inserted in the frame header before being sent out of the access switch to the VTEP (leaf node) in the fabric.
- The Leaf node encapsulates the packet with a VXLAN header, preserving the CMD data and sends the packet to the destination.



Configuration of Leaf Node and Access Switch

The leaf node at the distribution layer is configured to provide VXLAN-GPO encapsulation of the packets at egress, and propagation of SGT tags.

The access switch is configured to provide CTS classification.

Device Tracking for a Trunk Interface

When you configure the group-based policy, you must explicitly attach a device-tracking policy to the trunk interface of the Layer 2 VTEP. This explicit policy attachment is required to keep the IP-to-SGT binding active.

The following sample configuration shows how to define a policy and attach it to a specific target, which is the trunk interface.

1. Define a new policy.

```
VTEP(config)# device-tracking policy IPDT_POLICY
VTEP(config-device-tracking)# tracking enable
VTEP(config-device-tracking)# security-level glean
VTEP(config-device-tracking)# exit
```

2. Verify the policy that is created.

```
VTEP# show device-tracking policy IPDT_POLICY
Device-tracking policy IPDT_POLICY configuration:
security-level glean
device-role node
gleaning from Neighbor Discovery
gleaning from DHCP6
gleaning from ARP
```

```
gleaning from DHCP4
NOT gleaning from protocol unknown
tracking enable
```

3. Configure the policy on a specific target (trunk interface).

```
VTEP(config)# interface gil/0/1
VTEP(config-if)# switchport trunk allowed vlan 101,102,201,202
VTEP(config-if)# switchport mode trunk
VTEP(config-if)# device-tracking attach-policy IPDT_POLICY
VTEP(config-if)# exit
```

Configuration Example of EVPN Microsegmentation

Configuration of SGACL policies should be done primarily through the Policy Management function of the Cisco Secure Access Control Server (ACS) or the Cisco Identity Services Engine (ISE).

If you are not using AAA on a Cisco Secure ACS or a Cisco ISE to download the SGACL policy configuration, you can manually configure the SGACL mapping and policies.

The following is a snippet of a sample configuration of static SGT map on a Leaf node.

```
<snip: only relevant configuration shown>
cts role-based sgt-map 192.168.0.0/16 sgt 192
cts role-based sgt-map 192:168::/64 sgt 192
cts role-based sgt-map vrf green 10.2.201.10 sgt 101
cts role-based sgt-map vlan-list 101 sgt 101
cts role-based sgt-map vlan-list 102 sgt 102
<snip: only relevant configuration shown>
```

The following is a snippet of a sample configuration on the VTEPs for CTS enforcement.

```
<snip: only relevant configuration shown>
cts role-based enforcement
cts role-based enforcement vlan-list 101
cts role-based enforcement logging-interval 5
cts role-based permissions default d_permit
cts role-based permissions from 0 to 0 m_icmp
cts role-based permissions from 31 to 31 m_permit
cts role-based permissions from 101 to 192 m_icmp m_permit
<snip: only relevant configuration shown>
```

The following is a snippet of a sample configuration on the NVE interface that enables propagation of SGT information.

```
<snip: only relevant configuration shown>
interface nve1
  group-based-policy
<snip: only relevant configuration shown>
```

The following is a snippet of a sample configuration to accept SGTs carried via CMD header.

```
<snip: only relevant configuration shown>
interface HundredGigE2/0/29
  switchport trunk allowed vlan 101,102,201,202
  switchport mode trunk
  cts manual
  policy static sgt 2 trusted
<snip: only relevant configuration shown>
```