

The Catalyst SDWAN v20.15.1 (Updated) Master Class

Platform Changes in version 20.15.1

This document describes device configuration templates, the course focuses on Configuration Groups (as shown in the video) as our primary mode of device management, but Templates are still commonly seen in the majority of SDWAN deployments. We discussed Templates in the Lab Build class, so I am presenting this as a refresher to describe and explain the components of device and feature templates, the variable options, template configuration sections, and system settings. It also describes how to apply the device template and device provisioning process, CLI templates, and Cisco vSmart device templates.

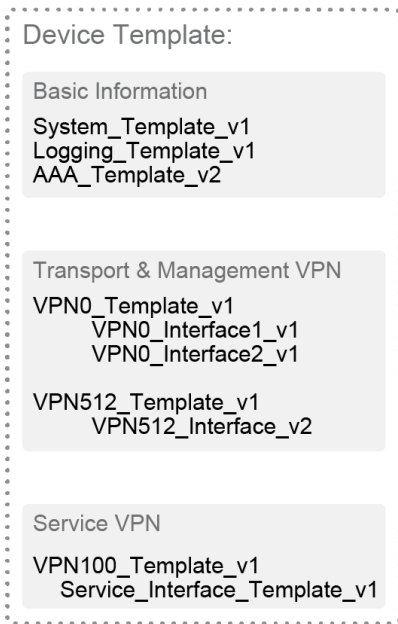
Device templates are similar to configuration groups, and they contain a complete operational configuration for a device. Device templates are supported both on Viptela OS and IOS XE based devices. You can create device templates by consolidating individual feature templates. With feature templates, each individual feature is configured within its own template and each feature template can be used by multiple device templates. This differs from feature profiles used in configuration groups, where the individual features are grouped into a feature profile and each feature profile, along with all individual features, can be used in multiple configuration groups.

Each device template is specific for a type of device. For each device type, if multiple devices have the same configuration, you can use the same device template for them. For example, many network routers might have the same basic configuration, so you can configure them with the same templates. You specify the differences in the templates using configuration variables. If the configurations for the same type of devices are different, you create separate device templates.

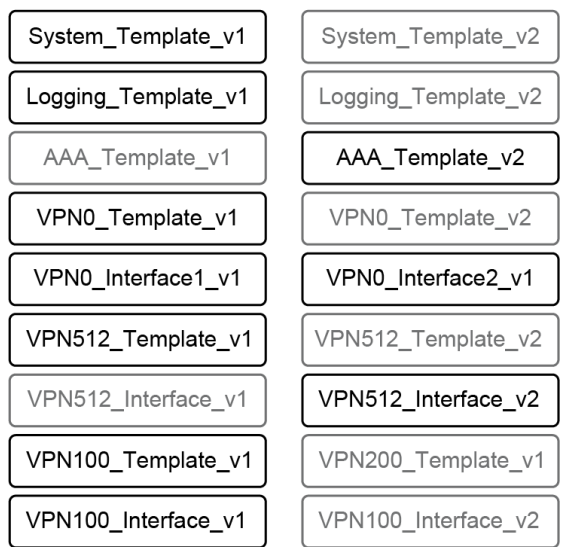
You can also create a device template by entering a CLI text-style configuration directly in Cisco vManage. Typically, you upload a text file containing the configuration text (or cut the configuration text from a text file and paste it into Cisco vManage). You can also directly type the configuration text into Cisco vManage.

Templates exist on two basic levels:

- Device templates are specific to a particular device type. Device templates describe the device overall, with all its high-level components, such as VPN0 and VPN512. Within each component, device templates point to specific feature templates.
- Feature templates describe one particular feature in detail, with all the settings and parameters required for that feature to function. For example, the VPN0 feature template contains all the settings for the VPN, the default routes, and so on. You can reuse feature templates for various device types.



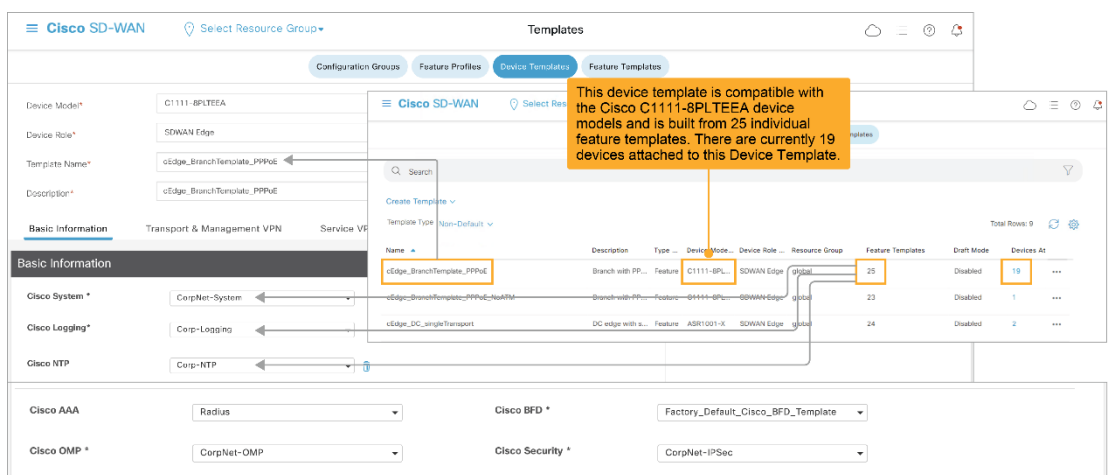
Feature Templates:



When working with templates, it is important to remember whether you are looking at a device or a feature template and how they fit together. A simple analogy would be that of a puzzle. The feature templates are the individual puzzle pieces that come together to form the complete puzzle (that is, the device template).

Centralized Device Configuration Through Templates

A device template is valid for a specific device model, it has a name and description, and has all the sections of the configuration.



- Centralized Feature Templates
- Configuration with Variables
- Self-Recover on Misconfiguration

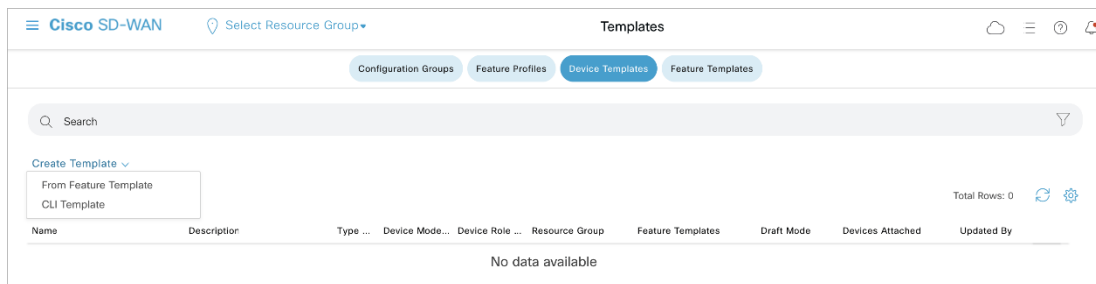
On the device template overview page in the Cisco vManage GUI, you can view the name of the defined templates, the description, which device model the template applies to, and how many feature templates it uses. You can also verify how many devices are currently attached to this template (that is, devices using thi

s template). In the example, the device template applies to the WAN Edge router Cisco C1111-8PLTEEA and has 25 feature templates that are associated with it.

Device Templates

Device templates are described as follows:

- Device templates define the complete operational configuration of a device.
- They can reference several feature templates.
- They are model-specific (device type).
- No device templates are defined in a new installation, by default.



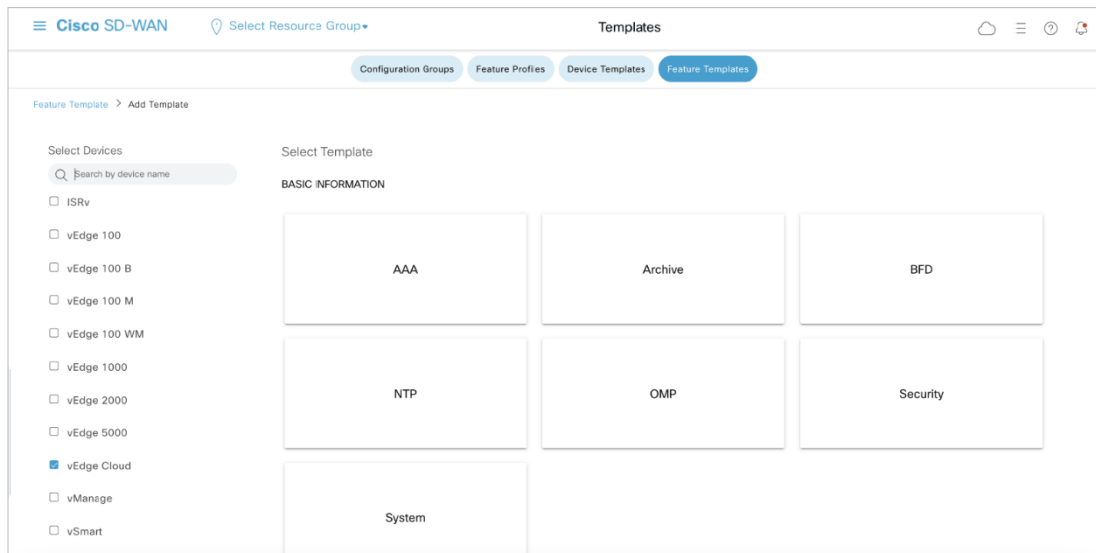
A device template consists of several feature templates. Each feature template defines the configuration for a particular Cisco SD-WAN software feature. Some feature templates are mandatory, indicated with an asterisk (*), and some are optional. Each mandatory feature template, and some of the optional ones, have a factory default template. For software features that have a factory default template, you can use either the factory default template (named `Factory_Default_<feature-name>_Template`) or you can create a custom feature template. You can define templates in Cisco vManage GUI under Configuration > Templates. No device templates are defined in a new installation by default.

Feature Templates

Within a feature template, some configuration commands and command options are identical across all device types. Others—such as a device system IP address, its geographic latitude and longitude, the time zone, and the overlay network site identifier—are variables that change from device to device. When you attach the device template to a device, you are prompted to enter actual values for these command variables. You can either enter the values manually, by typing the values for each variable and for each device, or you can upload an Excel file in CSV format that contains the values for each device.

The roles of feature templates and how they are used:

- Feature templates represent individual building blocks of the configuration.
- Feature templates are device model-specific.
- Feature templates are attached to devices by using device templates.



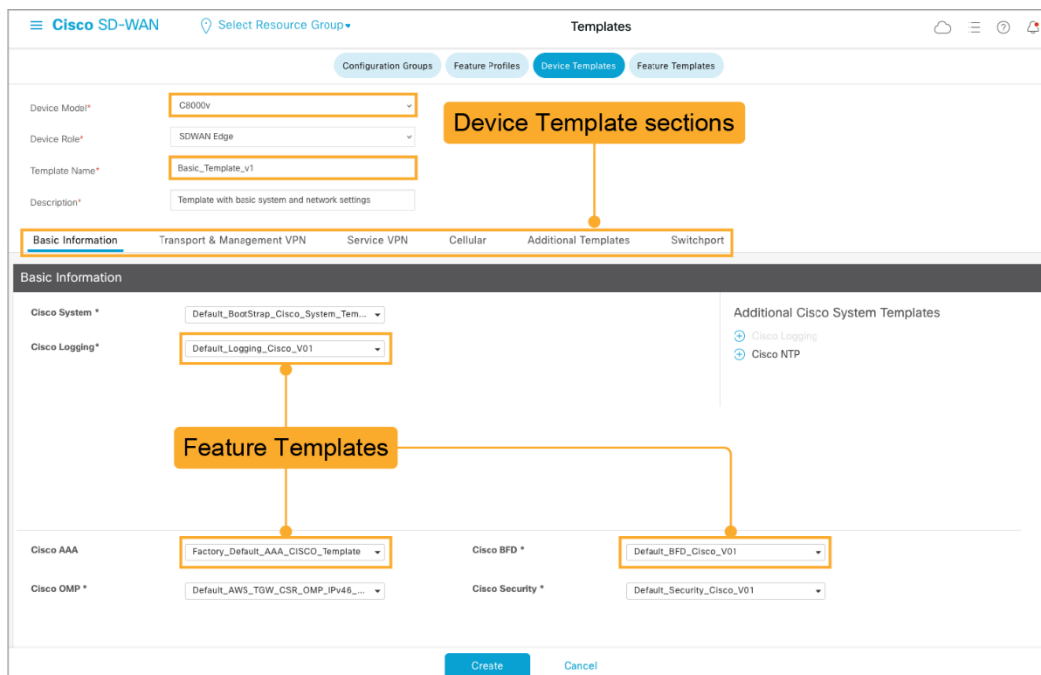
A feature template is used for each feature of a device configuration, which explains the high number of templates that are displayed in

Configuration > Templates > Feature Templates. For each feature template, you can choose which device types to use the template with. Recall the analogy of the puzzle: each feature template is a piece of the puzzle; in this case, the device template is that complete puzzle. The figure shows different types of feature templates: AAA, Archive, BFD, NTP, OMP, Security, and System.

Device Template Configuration Sections

Besides name, description, and device model, a device template has four main sections of the device configuration:

- Basic Information
- Transport and Management VPN
- Service VPN
- Additional Templates



You can click a section heading to go immediately to that section. Each section of the template then points to

individual feature templates. Before you define your own templates, the device template always points to factory default templates.

Creating a Device Template—System Settings

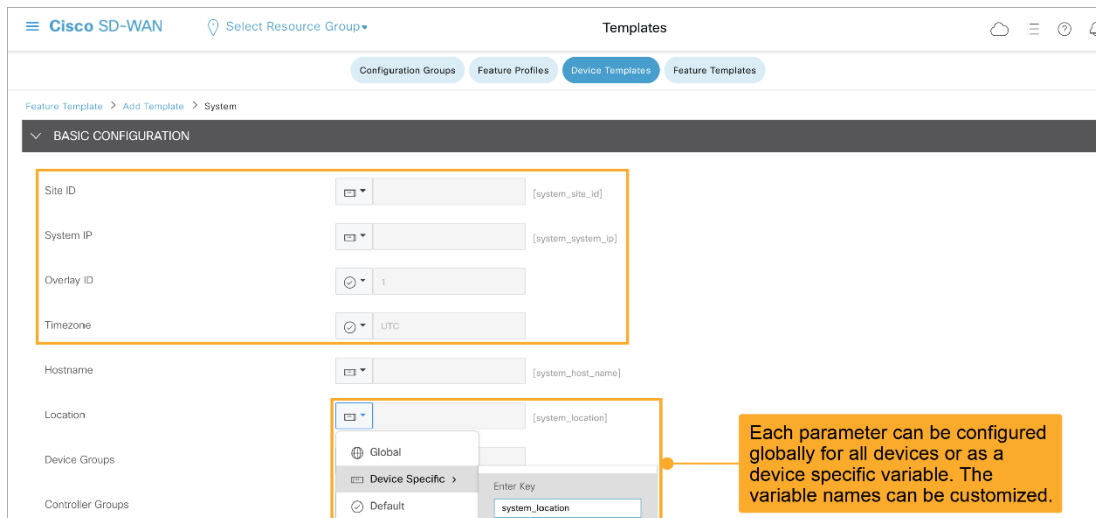
The first feature template in a device template is the system settings. Here, you will rediscover many parameters that were mentioned previously, such as System IP and Site ID.

Each variable in a template has the following three potential settings:

- **Global:** The value entered for this variable is a global one. All devices using this template use the same value, such as the name of the transport network. Examples of such parameters are DNS server, syslog server, and interface MTUs.
- **Device Specific:** Each device has its own value, such as interface IP addresses.
- **Default:** The template does not touch that variable; it remains with its default value.

The first feature template in a device template is the system settings template, and it includes:

- Site ID
- System IP
- Overlay ID
- Timezone

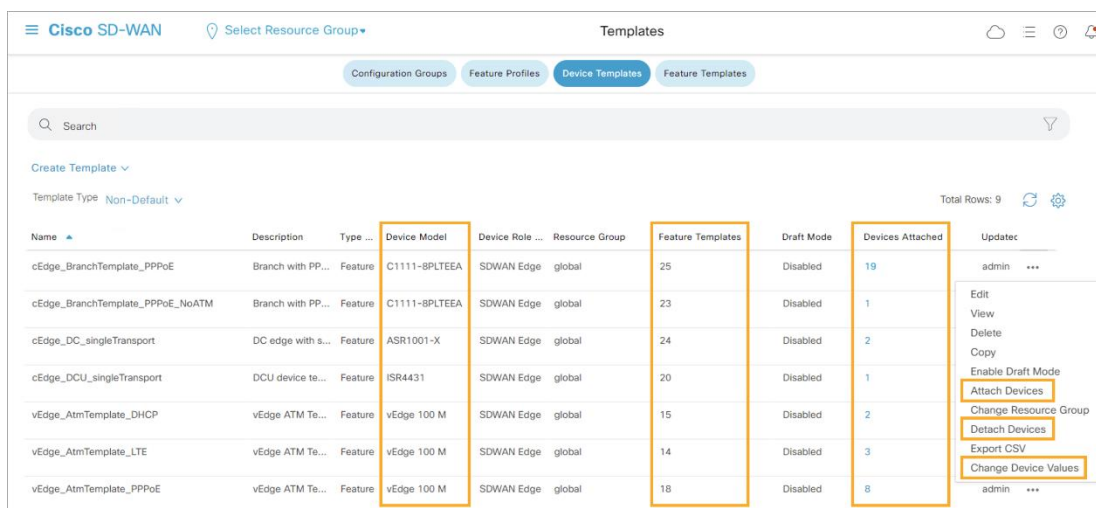


If you select the device-specific option, you can adapt the variable names to potentially already existing variables in other components in a software-defined environment. Unless there is a good reason, the recommendation is to not change those variable names.

Applying the Device Template

After you finalize the device template and associate its feature templates, you must attach the device template to devices. Initially no devices are attached to the template.

In the template overview screen, click the More Options icon (...) on the right of a template, and choose **Attach Devices**.



After applying the template to devices, a pop-up screen displays all available devices of this device type. You will not see devices of other device types here.

Choose the devices on the left to which the template applies and add them by using the arrows to the right. When all devices are selected, click **Attach** to attach the template to those devices.

You can perform the attach and detach device template operations on different devices, from one or more Cisco vManage servers, at the same time. However, if any one of these operations is in progress on one Cisco vManage server, you cannot edit any feature templates on any of the servers until the attach or detach operation completes.

You can also change any variable values for any device using the Change Device Values option. This action opens a form where you define the values for the variables in the template.

Defining Variables for Device Templates

Templates often include parameters that are device-specific. The system must provide those device-specific values.

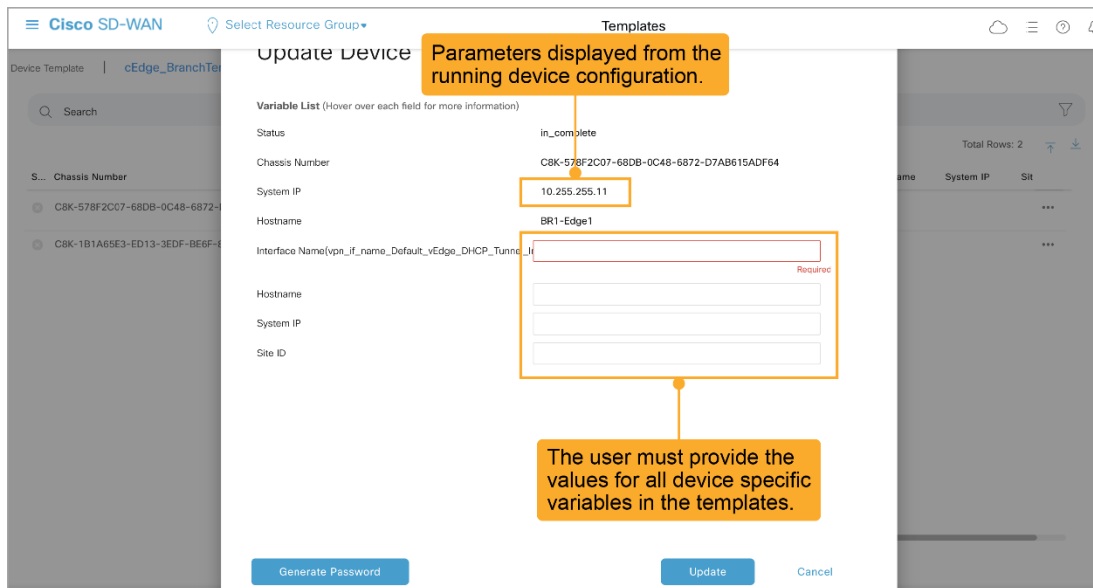
The screenshot shows the Cisco vManage interface for managing device templates. The page title is "Templates" and the current template is "cEdge_BranchTemplate_PPPE". A table lists two devices with columns for Chassis Number, System IP, Hostname, Interface Name, and Site. Two callout boxes are present: one pointing to the download/upload CSV buttons and another pointing to the "Edit Device Template" button.

| S... | Chassis Number | System IP | Hostname | Interface Name(vpn_if_name_Default_vEdge_DHCP_Tunnel_Interface) | Hostname | System IP | Site |
|------|-------------------------------------|---------------|-----------|---|----------|-----------|------|
| ● | CBK-578F2C07-68DB-0C48-6872-D7AB... | 10.255.255.11 | BR1-Edge1 | | | | |
| ● | CBK-1B1A65E3-ED13-3EDF-BE6F-862C... | 10.255.255.12 | BR1-Edge2 | | | | |

There are two methods for providing those values:

- You can download a CSV file that contains all the parameter fields (with current values, that were already configured) for all the associated devices. You can edit this file offline and upload it again to the tool by using the up arrow button.
- You can edit the values of a specific device by clicking the More Options icon (...) next to the device and clicking **Edit Device Template**.

When editing a specific device template to provide all parameters, a popup window opens, displaying all parameters. The parameters that are already configured are provided in the template; the parameters that are still required are empty. The figure shows the "Update Device Template" popup window. Notice that some values are already populated, while others are blank.



In an overlay network, you might have multiple devices of the same type that have identical or effectively identical configurations. For example, in a network with redundant Cisco vSmart controllers, each controller must be configured with identical policies. Another example is a network with Cisco WAN Edge routers at multiple sites, where each Cisco WAN Edge router provides identical services at each site.

Because the configurations for these devices are essentially identical, you can create one set of feature templates, which you then consolidate into one device template that you use to configure all the devices. You can create an Excel file in CSV format that lists the variables and defines each device specific variable value for each device.

Then you can load the file when you attach a device template to a device.

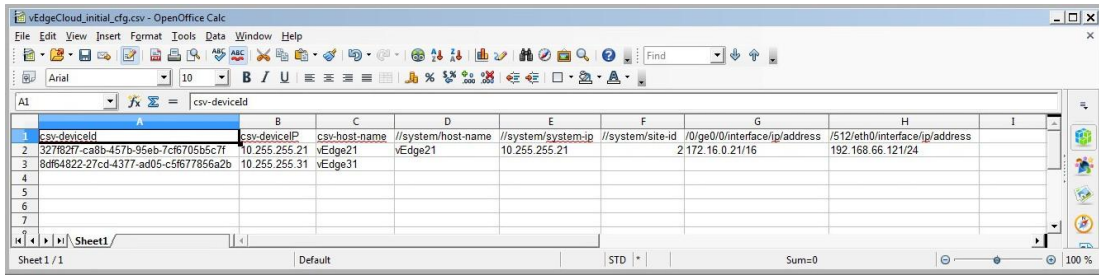
In the CSV file, the header row contains the variable name, and each row after that corresponds to a device, defining the values of the variables.

The first three columns in the spreadsheet must be in this order:

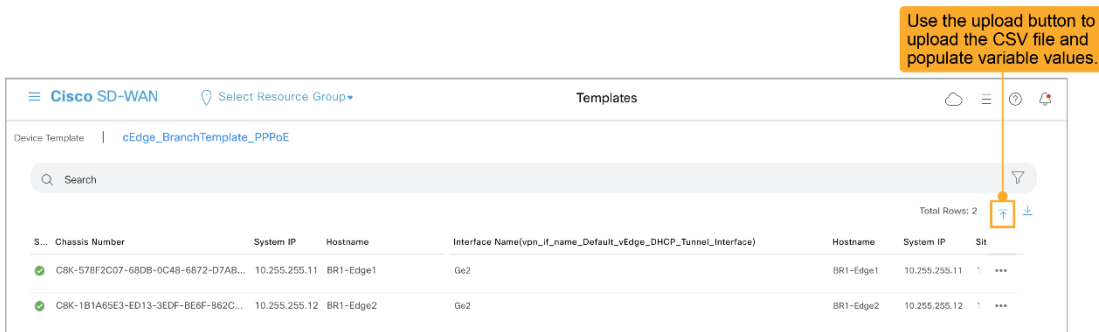
- **csv-deviceId:** Serial number of the device (used to uniquely identify the device). For WAN Edge routers, you receive the serial numbers in the authorized serial number file that is sent to you from Cisco. For other devices, the serial number is included in the signed certificate that you receive from Symantec or from your root CA.
- **csv-deviceIP:** System IP address of the device (used to populate the system IP address command).
- **csv-host-name:** Hostname of the device (used to populate the system hostname command).

You can create a single spreadsheet for all devices in the overlay network—Cisco vSmart controllers, Cisco vBond orchestrators, and WAN Edge routers. You do not need to specify values for all variables for all devices.

You can open the CSV file with a favorite spreadsheet tool (Excel, Calc, and so on). The values are separated by commas, and all previously defined variable values are preserved.

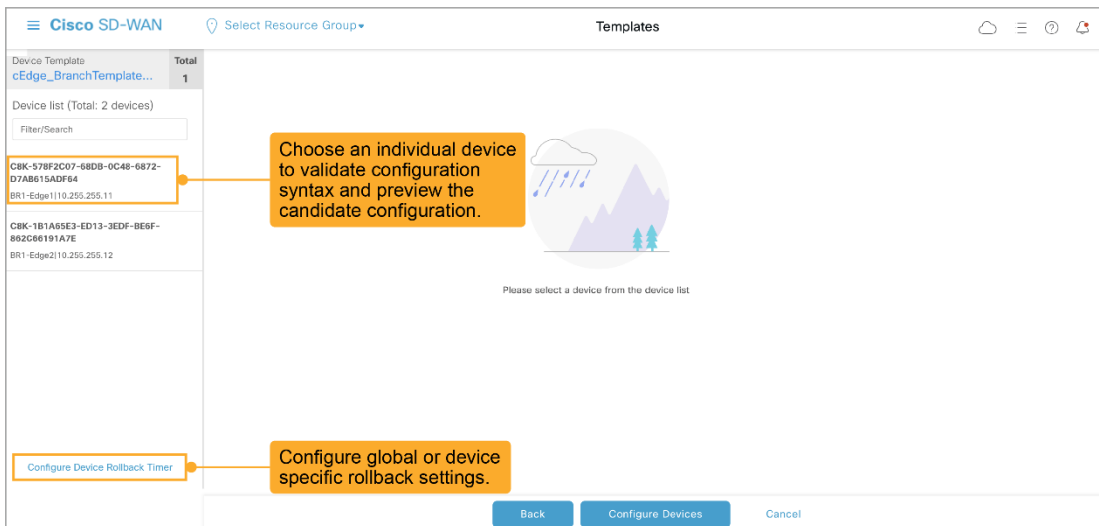


When finished, you save the CSV file (in the CSV format) and upload it to Cisco vManage.



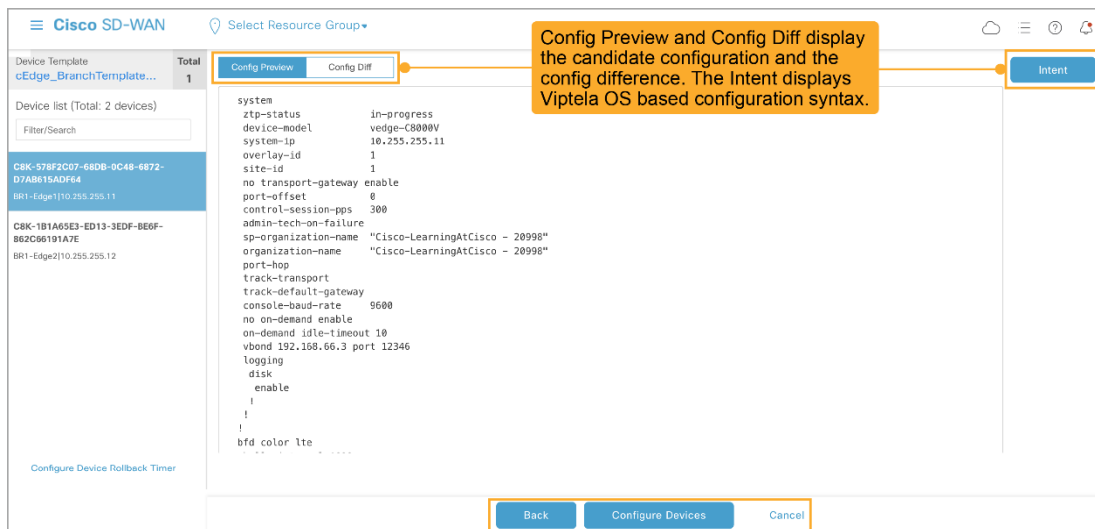
Device Provisioning

Before the templates with the new values can be applied to the network, a control window is displayed. On the left, the list of devices to be configured are listed. You can click on a device to view the configuration that Cisco vManage will push to the device.



In the case of a serious misconfiguration, such as a wrong interface address or a missing default route, you may lose the control connection to the WAN Edge router. All Cisco WAN Edge devices are configured with a rollback function. If a device loses control connections after a configuration, it can automatically roll back to the previous configuration. The rollback timer can be configured at the bottom left of the screen. The default rollback timer value is 5 minutes.

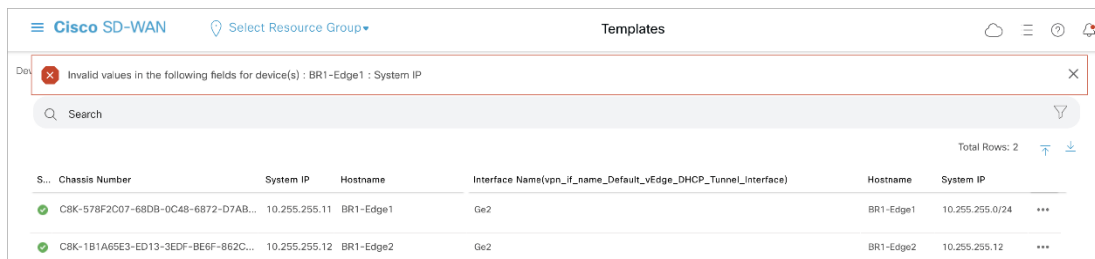
You can examine and verify the configuration that Cisco vManage will push to the device. You can also view the configuration differences (using the Config Diff option) to see what has changed. This way, you can catch mistakes before they take effect on the network.



If all configurations look correct, click the **Configure Devices** button at the bottom.

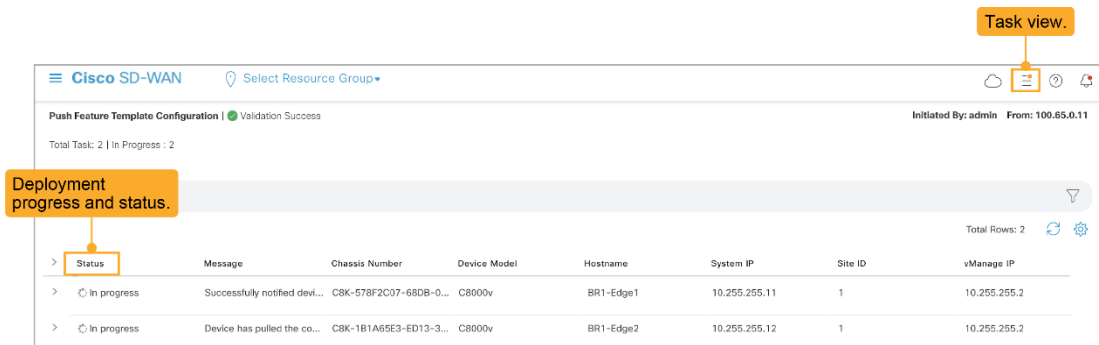
Error in Variable Value Example

If you made an error in the parameters, the error might be detected before deployment. In the example in the figure, the system-ip of one of the devices is misconfigured. Cisco vManage validates the configuration parameters and prevents you from deploying a misconfigured template. In such a case, update the misconfigured value and try again.



Template Deployment Status

Depending on the number of devices and the complexity of the template, the deployment may take a while. In the Task View, you can verify the status of the deployment.

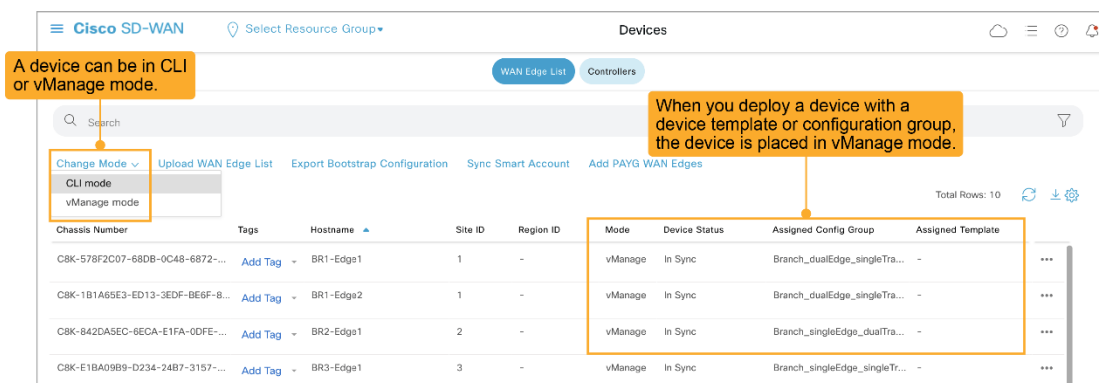


At any time, you can switch from other tabs and return to the task view by clicking the **Task View** icon on the top right.

Configuration Mode

A device can be in either CLI mode or vManage mode, as follows:

- CLI mode means that the device is being managed directly through the device CLI.
- vManage mode means that the device is being managed from Cisco vManage.



```
BR1-Edge1#config-transaction
admin connected from 127.0.0.1 using console on BR1-Edge1
BR1-Edge1(config)#hostname Branch1-Router1
BR1-Edge1(config)#commit
Aborted: 'system is-vmanaged': This device is being managed by vManage. Configuration through CLI is not allowed.
```

Devices in vManage mode can not be locally managed through the CLI.

When you apply a template for the first time, a device is automatically put into Cisco vManage mode. Notice that in the example, once a device is in Cisco vManage mode, you cannot configure it through the CLI.

Editing Device Templates

If you modify a template in a running network, or the individual device variable values, the next step is always to deploy those changes. If you cancel the deployment, the template changes are discarded. You can enable Draft Mode for a device template which enables you to modify a device template and save the changes but delay the template deployment for a later time. This makes it possible to implement significant changes to a template over a period of time before deploying the changes.

When you develop a new template, it is recommended that you copy an existing template and edit the copy, instead of editing the original—so you can go back to the last working status if an error occurs. After you edit a template or create a new one, the device values might require updating again. Note that if you add a new field to a template, the format of the downloaded CSV file also changes.

The screenshot shows the Cisco SD-WAN Templates interface. At the top, there are tabs for Configuration Groups, Feature Profiles, Device Templates (selected), and Feature Templates. Below the tabs is a search bar and a 'Create Template' dropdown. The main area contains a table with columns: Name, Description, Type, Device Mode, Device Role, Resource Group, Feature Templates, Draft Mode, Devices Attached, Updated By, and Last Update. A context menu is open over the second row, showing options: Edit, View, Delete, Copy, Enable Draft Mode, Attach Devices, Change Resource Group, Detach Devices, Export CSV, and Change Device Values.

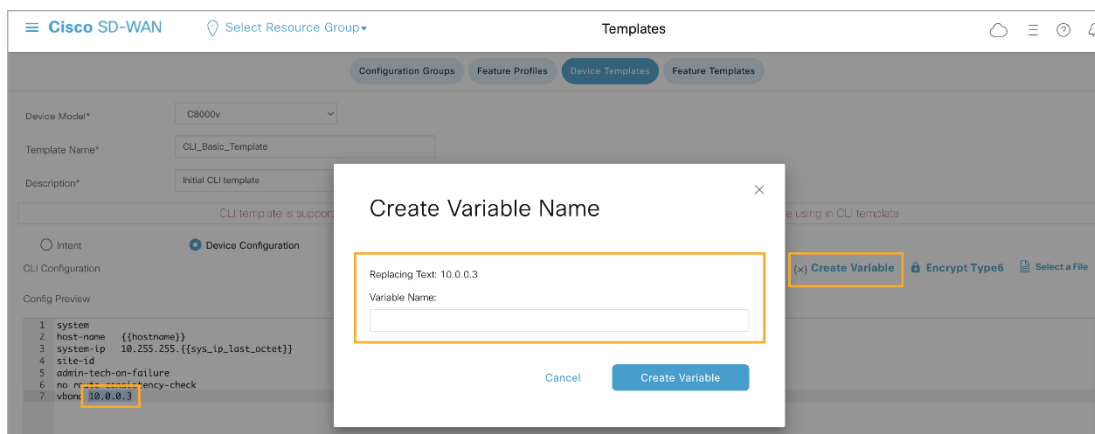
| Name | Description | Type | Device Mode | Device Role | Resource Group | Feature Templates | Draft Mode | Devices Attached | Updated By | Last Update |
|------------------------------|-------------------|---------|--------------|-------------|----------------|-------------------|------------|------------------|------------|-------------|
| cEdge_BranchTemplate_PPPE | Branch with PP... | Feature | C1111-BPL... | SDWAN Edge | global | 25 | Disabled | 19 | admin | 19 Oct 202 |
| cEdge_BranchTemplate_PPPE... | Branch with PP... | Feature | C1111-BPL... | SDWAN Edge | global | 23 | Disabled | 1 | admin | |
| cEdge_DC_singleTransport | DC edge with s... | Feature | ASR1001-X | SDWAN Edge | global | 24 | Disabled | 2 | admin | |
| cEdge_DCU_singleTransport | DCU device te... | Feature | ISR4431 | SDWAN Edge | global | 20 | Disabled | 1 | admin | |
| vEdge_AtmTemplate_DHCP | vEdge ATM Te... | Feature | vEdge 100 M | SDWAN Edge | global | 15 | Disabled | 2 | admin | |
| vEdge_AtmTemplate_LTE | vEdge ATM Te... | Feature | vEdge 100 M | SDWAN Edge | global | 14 | Disabled | 3 | admin | |
| vEdge_AtmTemplate_PPPE | vEdge ATM Te... | Feature | vEdge 100 M | SDWAN Edge | global | 18 | Disabled | 8 | admin | 03 Nov 202 |

CLI Templates

Cisco vManage also allows you to define and use CLI templates. A CLI template is a method that service providers and large enterprises used in the past: scripts that take a base CLI configuration with variables, apply to it a set of values for those variables, and push the resulting CLI.

Characteristics of CLI templates:

- Any text can be replaced with a variable.
- CLI templates are deployed the same way as device templates.
- Heavily used in the past before the introduction of device templates.




CLI templates are powerful because they are very generic. Cisco vManage does not need to “understand” a feature to push configuration this way. However, CLI templates are significantly more error prone and more difficult to develop. Therefore, the recommendation is to always use feature templates. In some specific scenarios, though, it is possible that Cisco

vManage does not support all feature options in a feature template. In such case, you need to use CLI templates.

Cisco vSmart Device Templates

In a production network, the recommendation is to have at least two Cisco vSmart controllers. However, there are significantly more controllers in larger production networks. To make sure that the network is in a consistent state, all Cisco vSmart controllers must have the same view of the network. Conflicting Cisco vSmart configurations can lead to unpredictable network behavior.



The screenshot shows the 'Templates' configuration page in Cisco SD-WAN vManage. The page is titled 'Cisco SD-WAN' and 'Templates'. It has a navigation bar with 'Configuration Groups', 'Feature Profiles', 'Device Templates', and 'Feature Templates'. The 'Device Templates' tab is active. The 'Device Model*' dropdown is set to 'vSmart'. The 'Template Name*' is 'vSmart_GoldenTemplate_v1' and the 'Description*' is 'The vSmart golden template'. Below this, there are tabs for 'Basic Information', 'Transport & Management VPN', and 'Additional Templates'. The 'Basic Information' tab is selected, showing fields for 'System *' (vSmart_System_v1), 'Logging*' (Factory_Default_Logging_Template_V01), 'AAA *' (Factory_Default_AAA_Template), 'OMP *' (Factory_Default_vSmart_OMP_Template), and 'Security *' (Factory_Default_vSmart_vManage_Sec...). There are 'Update' and 'Cancel' buttons at the bottom.

Therefore, you should use templates to configure Cisco vSmart controllers so that all vSmart controllers are automatically in sync. For centralized policy deployment, Cisco vManage mode is required. When applying templates to Cisco vSmart, they convert to Cisco vManage mode (just like the WAN Edge routers). You can use CLI templates to configure Cisco vSmart controllers, but it is not recommended.

Introducing SDWAN User Experience 2.0

The Cisco SD-WAN vManage interface provides an intuitive and centralized approach for managing SD-WAN deployments. It simplifies configuration, monitoring, and policy management across a distributed WAN environment, significantly improving network agility, security, and operational efficiency compared to traditional device-centric management methods. The introduction of **Configuration Groups**, **Feature Profiles**, and **Policy Groups**

in the vManage User Experience (UX) marks a significant evolution in how administrators interact with and control their SD-WAN environments.

This detailed overview focuses on the capabilities and features of **Configuration Groups**, **Feature Profiles**, and **Policy Groups** in vManage and how they enhance network management over traditional device-centric methods.

Configuration Groups in Cisco SD-WAN vManage

Capabilities:

Configuration Groups in Cisco SD-WAN vManage allow administrators to logically group multiple devices that share similar configurations, simplifying the management and deployment of consistent configurations across an SD-WAN environment. This approach significantly reduces manual configuration efforts and minimizes the risk of configuration errors.

- **Centralized Group Management:** Configuration Groups provide a centralized method for defining and managing the configuration of multiple devices simultaneously. This helps to maintain uniform configurations across devices that perform similar roles or are located in the same geographical region.
- **Efficient Onboarding:** When onboarding new devices, they can be automatically added to a predefined Configuration Group. The devices then inherit the group's configuration, reducing the time and complexity involved in deploying new devices.
- **Hierarchical Configuration:** Supports hierarchical configuration management, allowing base configurations to be defined at a higher level and then inherited by subordinate groups. This simplifies configuration management and enables scalability.

Features:

- **Configuration-Group Based Configuration:** Configuration Groups leverage Feature Profiles and Feature Profiles are made up of individual features and can be applied to multiple devices. These templates include system settings, VPN configurations, and security parameters, ensuring consistent policy enforcement across the network.

- **Inheritance and Overriding Capabilities:** Supports inheritance of configuration settings from parent groups to child groups, allowing for flexibility in defining group-specific configurations while maintaining overall consistency.
- **Bulk Device Updates:** Allows administrators to apply updates or configuration changes to all devices within a group in a single operation, significantly reducing the time required for routine updates.
- **Dynamic Device Membership:** Devices can be dynamically added to or removed from groups based on predefined criteria, ensuring that Configuration Groups remain relevant and up-to-date as the network evolves.

Advantages Over Old CLI Device Management Methods:

- **Reduced Configuration Complexity:** Traditional device management required individual configuration of each device, often leading to inconsistencies and increased chances of human error. Configuration Groups provide a more streamlined and error-free approach.
- **Scalability:** Managing a large number of devices becomes more scalable with Configuration Groups, as opposed to manual device-by-device configurations.
- **Consistency Across the Network:** Ensures uniform policy and configuration enforcement across all devices within a group, reducing configuration drift that is common in traditional management methods.

Feature Profiles in Cisco SD-WAN vManage

Capabilities:

Feature Profiles in vManage allow for the granular configuration of specific network functions and services, such as routing protocols, security features, quality of service (QoS), and interface settings. They provide a modular approach to defining and applying network features, enhancing flexibility and control over device configurations.

- **Granular Control of Device Features:** Feature Profiles allow administrators to configure specific features or functions independently, which can then be combined to form a complete device configuration. This modularity enhances the flexibility and control of SD-WAN deployments.

- **Reusable and Modular Profiles:** Once created, Feature Profiles can be reused across multiple Configuration Groups or individual devices, allowing for scalable and efficient network management.
- **Dynamic Configuration Adjustments:** Administrators can quickly adjust specific network functions without altering the entire device configuration, reducing the risk of widespread network disruption.

Features:

- **Predefined and Custom Profiles:** Provides both predefined profiles for common network functions and the ability to create custom profiles tailored to specific needs.
- **Support for Multiple Protocols and Services:** Feature Profiles cover a wide range of networking functions, including OSPF, BGP, VPN, DHCP, DNS, NAT, ACLs, QoS, and more, allowing comprehensive network feature management.
- **Role-Based Access Control (RBAC):** Integrates with vManage's RBAC system, ensuring that only authorized personnel can create, modify, or apply Feature Profiles to devices or groups.
- **Simplified Profile Management:** Profiles can be managed centrally, providing a clear overview of which profiles are in use, where they are applied, and their current status.

Advantages Over Old CLI Device Management Methods:

- **Enhanced Modularity and Reusability:** Traditional methods often required repetitive configurations for each device. Feature Profiles allow for modular configurations that can be reused, simplifying management.
- **Reduced Errors and Configuration Time:** By breaking down configurations into smaller, manageable components, Feature Profiles reduce the likelihood of configuration errors and the time needed to deploy changes.
- **Streamlined Change Management:** Changes to specific network functions can be made without impacting the entire device configuration, providing more control and reducing the risk of network issues.

Policy Groups in Cisco SD-WAN vManage

Capabilities:

Policy Groups in vManage provide a streamlined way to manage traffic policies, security rules, and routing behavior across the entire SD-WAN fabric. They allow for centralized policy definition and distribution, ensuring that traffic management and security policies are consistently enforced.

- **Centralized Policy Management:** Policy Groups enable administrators to centrally define and manage policies that govern traffic behavior, security rules, and application performance across the network.
- **Granular Policy Control:** Provides fine-grained control over how traffic is routed, prioritized, or blocked, based on criteria such as application type, source/destination IP, and user-defined conditions.
- **Dynamic Policy Adaptation:** Supports dynamic policy adjustments based on network conditions, allowing for real-time traffic optimization and enhanced application performance.

Features:

- **Application-Aware Routing and Traffic Engineering:** Policies can be created to optimize traffic flows for specific applications, ensuring high performance and low latency for critical business applications.
- **Security Policies:** Supports defining security policies that include firewall rules, segmentation, and service chaining to external security services, such as firewalls or intrusion detection/prevention systems (IDS/IPS).
- **Policy-Based Traffic Steering:** Allows for dynamic routing of traffic based on policy criteria, such as link quality, latency, jitter, or bandwidth availability, optimizing the use of available WAN links.
- **Integration with vAnalytics:** Policies can be monitored and adjusted based on real-time analytics and insights, ensuring optimal performance and security.

Advantages Over Old CLI Device Management Methods:

- **Centralized and Consistent Policy Enforcement:** Traditional methods often required policy configuration on each device, which was time-consuming and prone to inconsistencies. Policy Groups ensure uniform policy enforcement across all devices.

- **Simplified Troubleshooting and Optimization:** Centralized management of policies makes it easier to troubleshoot issues, optimize performance, and adjust policies without affecting device configurations directly.
- **Adaptability to Network Changes:** Policies can be dynamically adjusted based on real-time network conditions, something that is difficult to achieve with static, device-centric configurations.

Conclusion

The Cisco SD-WAN vManage User Experience, with its use of **Configuration Groups**, **Feature Profiles**, and **Policy Groups**, provides a modern and efficient approach to managing SD-WAN environments. These capabilities offer significant improvements over traditional device management methods, enabling centralized, scalable, and flexible network management. They reduce configuration complexity, improve consistency and reliability, and provide granular control over network functions and policies.

- **Configuration Groups** simplify device management by allowing bulk updates and hierarchical configurations.
- **Feature Profiles** offer modular and reusable configurations for specific network features, enhancing flexibility.
- **Policy Groups** provide centralized control over traffic behavior and security, enabling dynamic and consistent policy enforcement.

These features, combined with a user-friendly interface and centralized management capabilities, make Cisco SD-WAN vManage a powerful tool for modern network administrators, significantly enhancing operational efficiency, agility, and network security.

Lab 1: Create Common Feature Profiles

Design Goals:

Create a System Profile for use in all sites:

Name: "Common-System-Profile"

- Global Profile: SDU-GLOBAL-PROFILE
 - Accept Defaults

- AAA Profile: SDU-AAA-PROFILE
 - Local User: admin | 1234QWer
- BFD Profile: SDU-BFD-PROFILE
 - Accept Defaults
- OMP Profile: SDU-OMP-PROFILE
 - Accept Defaults
- Logging Profile: SDU-OMP-PROFILE
 - Accept Defaults
- Basic Profile: SDU-BASIC-PROFILE
 - Accept Defaults

Create a Feature Profile that defines VPN0 (Transport VRF) for use in all sites:

Name: “Common-Transport-Profile”

- VPN 0 should include 2 Ethernet Interfaces
 - Interface GigabitEthernet 1
 - Static IP Address
 - no shut
 - Allow All Services for Testing
 - Connected to Private1 in all sites
 - We will not be using GRE anywhere in this POC Topology
 - Interface GigabitEthernet 2
 - Static IP Address
 - no shut
 - Allow All Services for Testing
 - Connected to Private2 in all sites
 - We will not be using GRE anywhere in this POC Topology
- VPN 0 will support 2 Static Gateways:
 - 39.1.1.x
 - 40.1.1.x
- VPN 0 will use 10.255.255.1 as DNS

Create a CLI Add-on Profile that defines the following:

Name: “Common-CLI-Profile”

- platform console serial

Underlay vs Overlay Routing

Cisco Catalyst SD-WAN operates through a dual-layer architecture comprising the **underlay** and the **overlay** networks. Each layer serves a distinct role in delivering secure, resilient, and flexible wide-area network (WAN) services across distributed locations. The **underlay** provides the foundational physical connectivity, while the **overlay** creates a virtualized, policy-driven network on top of this foundation, leveraging secure tunnels to connect distributed sites. A core component within the overlay is the **Overlay Management Protocol (OMP)**, which manages route, policy, and control information across the SD-WAN fabric.

This document provides a comprehensive overview of the behaviors, roles, and operations of both the underlay and overlay networks in Cisco Catalyst SD-WAN, explaining how routing works in each and detailing how OMP facilitates the operation of the overlay network.

The Underlay Network in Cisco Catalyst SD-WAN

Definition and Role of the Underlay

The **underlay network** is the foundational physical or virtual infrastructure that connects all SD-WAN devices, such as branch routers, data centers, and hubs. It consists of traditional network transport mechanisms, including MPLS, LTE, and internet broadband, and provides the basic IP connectivity essential for the SD-WAN overlay network to operate.

The underlay effectively serves as the "pipe" through which all overlay traffic flows, enabling SD-WAN routers to communicate across diverse geographic locations.

Key Functions of the Underlay

1. **Transporting IP Packets:** The underlay network is responsible for the raw transportation of IP packets between SD-WAN nodes. Overlay traffic, encapsulated within IPsec tunnels, relies on the underlay for the actual data transfer.
2. **Supporting Diverse Transports:** The underlay leverages various transports (MPLS, broadband, LTE, etc.) to provide path diversity, redundancy, and resilience.
3. **Providing Reachability for Overlay Establishment:** The underlay allows SD-WAN routers to establish IP reachability and initiate connections necessary for the overlay network's formation.

Routing in the Underlay

Routing within the underlay is independent of the SD-WAN overlay and adheres to conventional IP routing principles:

- **Static and Dynamic Routing Protocols:** The underlay uses traditional routing protocols like OSPF, BGP, and static routes to establish IP reachability across different transports.
- **Path Selection Based on Native Metrics:** Routing decisions are based on native metrics (e.g., cost, bandwidth, latency), independent of SD-WAN policies.
- **Independence from Overlay Policies:** Underlay routing is configured independently and strictly follows network administrator settings, unaffected by overlay policies or dynamic SD-WAN requirements.

Underlay Network Behavior in Catalyst SD-WAN

The underlay in Cisco Catalyst SD-WAN operates transparently with respect to the overlay, ensuring all SD-WAN devices have IP reachability. Its main responsibility is providing a reliable foundation that remains resilient against link failures and transport degradation.

The Overlay Network in Cisco Catalyst SD-WAN

Definition and Role of the Overlay

The **overlay network** is a virtualized, policy-driven network created over the underlay infrastructure, enabling secure, encrypted communications across the SD-WAN fabric. It consists of IPsec tunnels established between SD-WAN edge devices (such as vEdge and cEdge routers) and forms either a full-mesh or hub-and-spoke topology as required. The overlay is where Cisco SD-WAN applies advanced policies for application-aware routing,

security, and traffic engineering to meet specific business and application performance requirements.

Key Functions of the Overlay

1. **Establishing Secure Tunnels:** The overlay establishes secure IPsec tunnels between SD-WAN edge devices, ensuring data confidentiality and integrity over potentially untrusted underlay paths.
2. **Applying Policies for Traffic Control:** The overlay enables the enforcement of SD-WAN policies, such as application-aware routing, QoS, and security policies.
3. **Dynamic Path Selection and Traffic Engineering:** The overlay uses real-time monitoring to make routing decisions based on performance metrics (e.g., latency, jitter, and packet loss).
4. **Centralized Control and Management:** The overlay provides centralized control through vManage and vSmart controllers, which distribute routing, security, and policy information to SD-WAN edge devices.

Overlay Management Protocol (OMP)

The **Overlay Management Protocol (OMP)** is a control protocol integral to the operation of the Cisco SD-WAN overlay network. OMP operates between SD-WAN edge devices and the vSmart controller, facilitating the exchange of routing, policy, and security information across the SD-WAN fabric.

Key Roles and Functions of OMP

1. **Route Distribution:**
 - OMP distributes routes within the SD-WAN overlay, allowing each SD-WAN edge device to advertise its local routes (e.g., LAN subnets) to the vSmart controller.
 - The vSmart controller aggregates these routes and redistributes them to other SD-WAN edge devices, enabling end-to-end reachability across the SD-WAN fabric.
2. **Service and TLOC Routes:**
 - **Service Routes** represent specific subnets or services located behind an SD-WAN edge device.
 - **Transport Locator (TLOC) Routes** define the physical location, transport encapsulation (e.g., IPsec), color (labeling for different transport types like MPLS or internet), and preference for each SD-WAN edge. TLOC routes are essential for establishing IPsec tunnels between edge devices.
3. **Policy Distribution:**
 - OMP distributes centralized policies defined in the vSmart controller, such as traffic engineering, service chaining, and security segmentation policies. These policies govern how traffic flows through the SD-WAN overlay, enforcing path selection and access control.
4. **Path Optimization:**
 - With BFD (Bidirectional Forwarding Detection) continuously monitoring path performance, OMP dynamically selects optimal paths based on real-time conditions, ensuring efficient routing.
5. **Security and Encryption:**
 - OMP manages IPsec tunnels between SD-WAN edge devices by distributing key and encryption information, ensuring data security across the overlay.

Routing in the Overlay with OMP

Routing within the overlay network is driven by OMP, which provides an adaptive and application-aware approach:

- **BFD Monitoring and Real-Time Adjustments:** OMP relies on BFD to monitor path metrics like latency, jitter, and packet loss for each TLOC, enabling real-time path adjustments.
- **Centralized Policy and Control:** The vSmart controller, through OMP, distributes policy-driven routing information across the network.
- **Application-Aware Routing:** OMP enables path selection based on application requirements, ensuring latency-sensitive applications can follow low-latency paths.
- **Service Chaining and Traffic Steering:** OMP supports service chaining, directing traffic through specific network services (e.g., firewalls) as needed.

Overlay Network Behavior in Catalyst SD-WAN

The overlay, controlled by OMP, is highly responsive to network conditions:

- **Self-Healing and Resiliency:** Through BFD and dynamic path adjustments via OMP, the overlay can self-heal, rerouting traffic automatically in case of link degradation.
- **Application-Aware Routing:** OMP allows the overlay to route traffic based on specific application requirements, enhancing performance and user experience.
- **vSmart and vManage Integration:** The vSmart controller handles OMP route and policy distribution, while vManage provides an intuitive interface for monitoring and policy configuration.

Interaction Between Underlay and Overlay

While the underlay provides the physical connectivity and IP reachability for SD-WAN devices, the overlay operates as an independent layer that dynamically optimizes traffic flow:

- **Tunnel Establishment:** The underlay provides the connectivity for initial tunnel establishment between SD-WAN edges.
- **Path Diversity and Redundancy:** The underlay's transport diversity allows for physical redundancy, while the overlay utilizes this redundancy for path selection and failover.
- **Resilience and Failover:** If an underlay path experiences issues, the overlay can detect the problem through BFD and use OMP to reroute traffic across other available paths.

TL/DR

In Cisco Catalyst SD-WAN, the **underlay** and **overlay** networks provide complementary functions:

- **Underlay:** Supplies the physical transport, IP connectivity, and path diversity needed for basic connectivity, using traditional routing protocols that operate independently of SD-WAN policies.
- **Overlay:** Manages secure, policy-driven communication across the SD-WAN fabric, with OMP serving as the key control protocol for dynamic routing, application-aware policies, and security enforcement.

The Overlay Management Protocol (OMP) enables the overlay to be adaptive, flexible, and responsive to both network conditions and business requirements, delivering the performance, security, and scalability that modern WAN architectures demand. By separating the control and data planes and leveraging the strengths of both layers, Cisco Catalyst SD-WAN achieves an optimal balance between performance, security, and manageability.

Control Plane Connections

Introduction

In Cisco SD-WAN, secure control plane tunnels are established between SD-WAN edge devices (such as cEdge routers) and Cisco SD-WAN controllers, including the vBond Orchestrator, vManage (Network Management System), and vSmart (SD-WAN Controller). These control plane tunnels use either Datagram Transport Layer Security (DTLS) or Transport Layer Security (TLS) protocols to ensure secure communication. The secure DTLS/TLS tunnels facilitate control and management functions, including device onboarding, configuration management, and policy distribution across the SD-WAN fabric.

This document provides a comprehensive overview of how DTLS/TLS tunnels are formed, utilized, and maintained in Cisco SD-WAN. It details the tunnel establishment process between cEdge routers and the vBond Orchestrator, vManage, and vSmart controllers, focusing on the control plane's security, resiliency, and operational integrity.

Overview of DTLS/TLS in Cisco SD-WAN Control Plane

Cisco SD-WAN uses **DTLS** or **TLS** protocols to establish secure, encrypted control plane tunnels between cEdge routers and SD-WAN controllers. These protocols enable authentication, encryption, and data integrity across control plane connections.

- **DTLS** (Datagram Transport Layer Security): Used in most cases for control plane tunnels. DTLS operates over UDP, which can handle packet loss more efficiently in networks with higher latency or loss rates.
- **TLS** (Transport Layer Security): Often used as a fallback for DTLS, TLS operates over TCP. It provides reliable delivery at the cost of slightly increased overhead, suitable for networks that demand strict reliability.

Both DTLS and TLS ensure that control plane communications between SD-WAN components are secure, authenticated, and resistant to tampering.

DTLS/TLS Tunnel Formation

The process of DTLS/TLS tunnel formation in Cisco SD-WAN is a stepwise operation where each cEdge router establishes secure tunnels with three main controllers: **vBond (Orchestrator)**, **vManage (Network Manager)**, and **vSmart (Controller)**.

1. Tunnel Formation with the vBond Orchestrator

The **vBond Orchestrator** serves as the initial contact point for cEdge routers joining the SD-WAN network. It performs several critical functions in tunnel formation:

1. **Device Authentication:**
 - When a new cEdge router attempts to join the SD-WAN fabric, it first contacts the vBond orchestrator.
 - The vBond and the cEdge mutually authenticate each other using a pre-configured certificate-based authentication process.
 - Both devices use Organization Validated Certificates (OVCs) or Trusted Root CA Certificates issued by Cisco or a trusted Certificate Authority (CA).
2. **Tunnel Establishment:**
 - After successful authentication, the vBond orchestrator and the cEdge router establish a DTLS/TLS tunnel.
 - The vBond uses this tunnel to securely exchange information about the SD-WAN fabric, including IP addresses of the vManage and vSmart controllers.
3. **Allocation of Controllers:**
 - Through the DTLS/TLS tunnel, the vBond orchestrator provides the cEdge router with IP addresses of its designated vManage and vSmart controllers.
 - This dynamic allocation allows for efficient scaling and load balancing by ensuring that each cEdge is assigned an optimal set of controllers.
4. **Firewall Traversal and NAT Detection:**
 - The vBond orchestrator performs NAT detection and assists with firewall traversal, ensuring that the cEdge can reach other controllers even if it's behind NAT.

After successfully establishing this initial tunnel with the vBond, the cEdge can proceed to set up tunnels with the vManage and vSmart controllers.

2. Tunnel Formation with the vManage Controller

The **vManage** controller is the centralized network management system that provides configuration, monitoring, and policy deployment capabilities for the SD-WAN fabric.

1. **Tunnel Authentication and Establishment:**
 - Using the information provided by vBond, the cEdge initiates a DTLS/TLS tunnel with vManage.
 - Mutual authentication occurs, typically using certificates, similar to the process with vBond.
2. **Configuration Management:**
 - The vManage uses this secure tunnel to push initial configurations and updates to the cEdge device.
 - Any changes in configurations or policies are sent over the DTLS/TLS tunnel, ensuring secure and consistent configuration management across the network.
3. **Telemetry and Monitoring:**
 - The vManage collects operational data and telemetry from the cEdge router over this tunnel.
 - This data includes statistics on performance, application usage, and real-time metrics, which vManage uses for monitoring and troubleshooting.

3. Tunnel Formation with the vSmart Controller

The **vSmart Controller** is responsible for policy and route distribution across the SD-WAN fabric. It plays a crucial role in enforcing centralized policies and managing the SD-WAN overlay routing.

1. Tunnel Authentication and Establishment:

- Using details provided by vBond, the cEdge router initiates a DTLS/TLS tunnel with the vSmart controller.
- Like the other controllers, mutual authentication is performed to ensure secure access.

2. Policy Distribution:

- The vSmart controller uses the DTLS/TLS tunnel to distribute routing, security, and application-aware policies to the cEdge router.
- These policies dictate the routing behavior, traffic segmentation, and access control rules that the cEdge router applies to its traffic.

3. Route Exchange:

- Through this secure tunnel, the vSmart controller and cEdge router exchange routing information.
- The vSmart controller uses the Overlay Management Protocol (OMP) to advertise overlay routes, Transport Locator (TLOC) routes, and service routes to the cEdge, enabling it to make dynamic, policy-based routing decisions.

Utilization of DTLS/TLS Control Plane Tunnels

Once established, the DTLS/TLS control plane tunnels between the cEdge router and each SD-WAN controller serve specific ongoing operational purposes:

1. **Secure Communication:** DTLS/TLS tunnels ensure that control plane traffic, including configuration, routing, and policy information, is encrypted and authenticated, preventing unauthorized access and data tampering.
2. **Configuration Synchronization and Updates:**
 - The vManage controller uses these tunnels to synchronize configurations across the SD-WAN fabric. Changes made in the vManage interface are securely distributed to each cEdge router, maintaining network consistency.
3. **Real-Time Monitoring and Telemetry:**
 - The cEdge router continuously sends performance data to vManage over the secure tunnel, including health metrics, traffic statistics, and application usage data. This information aids in real-time monitoring and troubleshooting.
4. **Policy and Routing Updates:**
 - The vSmart controller uses the DTLS/TLS tunnel to send policy and routing updates to the cEdge router. This enables dynamic, centralized control over routing decisions and traffic flow across the SD-WAN overlay.
5. **Fault Tolerance and Recovery:**
 - Cisco SD-WAN has mechanisms to re-establish DTLS/TLS tunnels in the event of a failure, maintaining connectivity between cEdge routers and controllers. If one controller becomes unavailable, the cEdge router can use the vBond orchestrator to reallocate resources and re-establish secure tunnels with other controllers.

Operations and Maintenance of DTLS/TLS Control Plane Tunnels

1. Tunnel Resilience and Re-establishment:

- Cisco SD-WAN monitors DTLS/TLS tunnels for connectivity and performance issues. In the event of a tunnel failure, the cEdge will attempt to reconnect automatically.
- The vBond orchestrator assists with re-authenticating the cEdge and re-establishing control plane tunnels with available vManage and vSmart controllers.

2. Certificate Management:

- Certificates are essential for authenticating devices and ensuring secure communication. Cisco SD-WAN supports centralized certificate management, allowing administrators to update or renew certificates without disrupting tunnel operations.

3. Troubleshooting and Diagnostics:

- Cisco SD-WAN provides tools to diagnose tunnel issues, such as packet captures and real-time logging. vManage includes dashboards for monitoring the status of each DTLS/TLS tunnel, allowing for quick troubleshooting.

4. Protocol Selection:

- The SD-WAN system automatically selects DTLS or TLS based on network conditions. If UDP traffic is blocked or degraded, TLS over TCP can be used as an alternative for reliable communication.

TL/DR

In Cisco Catalyst SD-WAN, DTLS/TLS tunnels provide a secure, resilient control plane foundation, enabling robust and centralized management of the SD-WAN fabric. Through distinct tunnel formations with the vBond, vManage, and vSmart controllers, cEdge routers can:

1. **Authenticate and Onboard Securely:** The initial DTLS/TLS tunnel with vBond ensures secure device onboarding, authentication, and dynamic controller allocation.
2. **Receive Configurations and Report Telemetry:** The DTLS/TLS tunnel with vManage allows for continuous configuration management, telemetry, and monitoring.
3. **Implement Policies and Exchange Routes:** The DTLS/TLS tunnel with vSmart facilitates the dynamic distribution of routing and security policies.

This tunnel-based architecture, using DTLS/TLS protocols, is fundamental to Cisco SD-WAN's secure, scalable, and policy-driven approach, making it highly adaptable for modern WAN environments.

Overlay Management Protocol (OMP)

The **Overlay Management Protocol (OMP)** is the core protocol in Cisco Catalyst SD-WAN that manages routing, policy distribution, and overall control plane operations across the SD-WAN fabric. Operating between SD-WAN edge devices (such as cEdge and vEdge routers) and the vSmart controllers, OMP enables a dynamic, policy-driven overlay network by managing routes, security parameters, and service advertisements. This document provides a

detailed explanation of how OMP functions, focusing on its adjacency mechanism, message types, and the information it learns and advertises across the SD-WAN network.

Overview of OMP

OMP operates as a control plane protocol, managing the communication between edge devices and vSmart controllers. Through OMP, SD-WAN routers can exchange routing information, service routes, security policies, and other critical data necessary to establish and maintain the SD-WAN overlay. OMP is designed to operate in a scalable and resilient manner, providing dynamic, application-aware routing across a distributed network.

Key Functions of OMP in Cisco SD-WAN

1. **Route Distribution:** OMP shares routes among SD-WAN edge devices via the vSmart controller, enabling end-to-end reachability across the SD-WAN overlay.
2. **Policy Distribution:** OMP distributes SD-WAN policies (such as traffic engineering and security policies) from vSmart to edge routers, ensuring consistent policy application.
3. **Service Advertisement:** OMP advertises services, such as data center applications or firewall resources, enabling service chaining and advanced routing configurations.
4. **TLOC Management:** OMP manages Transport Locator (TLOC) routes, which define the physical connections of each SD-WAN device, allowing for multi-path and diverse transport management.

OMP Adjacency Mechanism

To establish a stable control plane, OMP uses an **adjacency mechanism** that enables edge devices and vSmart controllers to form OMP adjacencies and exchange critical network information.

OMP Adjacency Formation

1. **Initial Contact with vBond:** Before an OMP adjacency can form, a new SD-WAN device must first authenticate with the **vBond Orchestrator**. The vBond provides the device with information about its designated vSmart controller(s).
2. **Connection with vSmart:** Once vBond has completed the authentication process, the SD-WAN device initiates a secure connection to the assigned vSmart controller(s). OMP adjacency forms between the edge device and the vSmart controller using a DTLS or TLS tunnel.
3. **Bidirectional Communication:** The OMP adjacency enables bidirectional communication between the vSmart controller and SD-WAN edge devices, allowing continuous exchange of routing, policy, and reachability information.
4. **Keepalive Messages:** OMP uses keepalive messages to maintain adjacency status. These keepalives help verify the health of the OMP session and ensure connectivity between the edge and the vSmart controller.
5. **Resilient Reconnection:** If an OMP adjacency is lost (e.g., due to network failure), the SD-WAN device will attempt to reconnect and re-establish adjacency, ensuring robust control plane resilience.

OMP Convergence

The OMP adjacency mechanism enables quick convergence across the SD-WAN fabric. When changes occur (e.g., a route update or policy change), OMP rapidly propagates these updates to all relevant edge devices via the vSmart controller, ensuring network-wide consistency and minimizing downtime.

OMP Message Types

OMP uses various message types to exchange routing, policy, and reachability information. Each message type serves a distinct purpose in maintaining and optimizing the SD-WAN overlay.

1. OMP Route Advertisements:

- **Purpose:** Distribute routes between SD-WAN edge devices and the vSmart controller, enabling end-to-end reachability.
- **Content:** Contains details about the prefixes reachable within the SD-WAN fabric, such as local LAN subnets or connected services.
- **Example:** An edge router advertises its local LAN prefix to the vSmart controller, which then redistributes it to other edge routers needing to reach that subnet.

2. TLOC Advertisements:

- **Purpose:** Share Transport Locator (TLOC) information, which defines the physical location, encapsulation type, color (transport label), and preference for each SD-WAN device.
- **Content:** TLOC advertisements include information such as IP address, encapsulation method (e.g., IPsec), and color (representing the transport, such as MPLS, internet, or LTE).
- **Example:** An edge device advertises its available TLOCs, allowing other devices to establish IPsec tunnels across different transport types.

3. Service Advertisements:

- **Purpose:** Advertise specific services, such as data center applications, firewall services, or other network resources available at certain locations.
- **Content:** Information about the available services, including the prefix and service type.
- **Example:** A data center router advertises a firewall service, enabling branch sites to route traffic through this firewall for inspection or security processing.

4. Policy Advertisements:

- **Purpose:** Distribute centralized policies configured on the vSmart controller to edge devices, ensuring consistent policy application.
- **Content:** Policy rules, including traffic engineering, access control, and security policies.
- **Example:** A traffic steering policy is sent to an edge device, instructing it to prioritize certain application traffic over a low-latency link.

5. Key and Security Parameters:

- **Purpose:** Exchange encryption keys and security parameters for IPsec tunnel establishment.
- **Content:** Information about encryption algorithms, keys, and security settings.
- **Example:** The vSmart controller advertises IPsec encryption keys to an edge device, enabling secure tunnel formation with other SD-WAN devices.

6. Keepalive Messages:

- **Purpose:** Maintain the OMP adjacency and verify the health of the control plane connection.
- **Content:** Minimal data indicating the active status of the connection.
- **Example:** An edge device sends periodic keepalives to the vSmart controller to maintain the OMP adjacency.

Information Learned and Advertised by OMP

OMP plays a central role in collecting, learning, and advertising critical information that enables SD-WAN devices to make intelligent routing and policy decisions. Below are the key types of information OMP learns and advertises across the SD-WAN fabric:

1. Routing Information:

- **Learned:** Each SD-WAN edge device learns about available routes from the vSmart controller. These routes include LAN prefixes from remote sites and service routes from specialized devices.
- **Advertised:** Edge devices advertise their local LAN prefixes to the vSmart controller, which aggregates and redistributes these routes to other edge devices.
- **Impact:** This enables end-to-end reachability and connectivity across the SD-WAN overlay.

2. TLOC Information:

- **Learned:** Each SD-WAN device learns about available TLOCs (Transport Locators) from the vSmart controller, representing possible transport paths to reach other devices.
- **Advertised:** Edge devices advertise their own TLOCs, including information about transport type (e.g., MPLS, internet), color, and encapsulation, to the vSmart controller.
- **Impact:** TLOC information is used to establish IPsec tunnels, enabling multi-path routing and transport redundancy.

3. Service Information:

- **Learned:** Edge devices learn about available services (e.g., firewalls, application servers) that other devices advertise, enabling service chaining and traffic optimization.
- **Advertised:** Devices advertise their local services to the vSmart controller, which redistributes them to other devices that may need access to these services.
- **Impact:** Enables intelligent service-based routing and application-aware traffic management.

4. Policy Information:

- **Learned:** Edge devices receive policy information from the vSmart controller, which is configured centrally on vManage.
- **Advertised:** Policies are generally not advertised by edge devices but are distributed from the vSmart controller, ensuring centralized control.
- **Impact:** Policies such as traffic engineering, QoS, and security rules are applied across the SD-WAN, enabling consistent application-aware routing and access control.

5. Security Information:

- **Learned:** Edge devices learn about security parameters for IPsec tunnel formation, such as encryption keys and algorithms, from the vSmart controller.
- **Advertised:** Security information, including keying material, is primarily distributed by the vSmart controller.

- **Impact:** Enables secure, encrypted communication between SD-WAN devices, ensuring data confidentiality and integrity.

TL/DR

The **Overlay Management Protocol (OMP)** is fundamental to the operation of Cisco Catalyst SD-WAN, enabling a highly dynamic, resilient, and policy-driven overlay network. By establishing OMP adjacencies, distributing a variety of message types, and continuously learning and advertising critical information, OMP provides:

1. **Reliable Control Plane Connectivity:** Through OMP adjacencies and keepalives, devices maintain consistent communication with the vSmart controller, ensuring the stability of the control plane.
2. **Dynamic Routing and Service Advertisements:** OMP facilitates the exchange of routing and service information, enabling intelligent, application-aware routing decisions.
3. **Comprehensive Policy Distribution:** OMP distributes centralized policies that ensure consistent behavior across the SD-WAN fabric, supporting advanced features like traffic engineering, QoS, and security.

VPNs & Labels

Cisco Catalyst SD-WAN leverages Virtual Private Networks (VPNs) and label-based forwarding to segment traffic and provide secure, isolated transport across the SD-WAN fabric. VPNs in Cisco SD-WAN function similarly to VRFs (Virtual Routing and Forwarding) in traditional networking, creating isolated, logical routing domains on a single physical network. Labels, meanwhile, are used to identify and route traffic between these VPNs across the SD-WAN overlay.

This document provides a detailed explanation of how VPNs and labels operate in Cisco Catalyst SD-WAN, exploring their roles in traffic segmentation, policy enforcement, and label-based forwarding across the overlay network.

Overview of VPNs in Cisco Catalyst SD-WAN

In Cisco SD-WAN, **VPNs** are used to create logical, isolated networks for different types of traffic or organizational needs. Each VPN operates independently, with its own routing table, forwarding policies, and segmentation, enabling secure multi-tenancy within a single physical infrastructure.

- **VPN Segmentation:** VPNs enable traffic segmentation, allowing organizations to isolate different types of traffic (e.g., employee traffic, guest traffic, IoT devices, voice/video) and apply unique policies to each.
- **Logical Routing Domains:** Each VPN functions as a separate routing domain, maintaining an independent routing table and IP address space.

- **Enhanced Security and Policy Control:** VPN segmentation allows administrators to enforce unique security and traffic policies for each VPN, ensuring controlled access and compliance with organizational requirements.

Key VPNs in Cisco Catalyst SD-WAN

Cisco SD-WAN uses VPNs extensively to separate different types of control and data traffic, including:

1. **VPN 0: Transport VPN** – Manages the underlay transport connections (e.g., internet, MPLS) and connects SD-WAN devices to the physical network. VPN 0 is where WAN-facing interfaces are placed, facilitating connectivity to other SD-WAN routers and controllers.
2. **VPN 512: Management VPN** – Handles device management and out-of-band management traffic. VPN 512 connects the SD-WAN device to external management networks or the vManage console, enabling centralized configuration and monitoring.
3. **User-Defined VPNs: Data VPNs** – User-defined VPNs, typically numbered from 1 to 511, and from 513 upwards, are used to create separate logical networks for specific user traffic, applications, or services. Each user-defined VPN represents an isolated routing domain with its own policies and access controls.

Functionality of VPNs in Catalyst SD-WAN

VPNs in Cisco Catalyst SD-WAN function as isolated, logical networks within the SD-WAN fabric. Here's how they operate:

1. **Independent Routing Tables:** Each VPN maintains its own routing table, which is isolated from other VPNs. Routes in one VPN are not visible to another VPN, ensuring strict traffic segmentation and security.
2. **Traffic Isolation and Policy Enforcement:** VPNs allow organizations to apply custom routing, security, and application-aware policies to each logical network. For example, one VPN might be prioritized for critical applications, while another VPN could have stricter security controls for sensitive data.
3. **Inter-VPN Communication Control:** By default, VPNs in Cisco SD-WAN do not communicate with each other. Any necessary inter-VPN traffic must be explicitly allowed through policies, ensuring that data flows are strictly controlled.
4. **Service and Resource Allocation:** VPNs allow the allocation of specific resources, such as IP subnets, QoS settings, and security policies, to different traffic segments, making it easy to allocate bandwidth, prioritize traffic, and maintain service quality for critical applications.
5. **Multi-Tenancy Support:** VPNs support multi-tenancy by enabling multiple isolated environments on a shared SD-WAN infrastructure, useful for organizations that serve different departments, clients, or locations with distinct networking requirements.

Overview of Labels in Cisco Catalyst SD-WAN

In Cisco SD-WAN, **labels** function as identifiers that allow for traffic forwarding across the overlay based on VPN membership. Labels provide a way to tag and route traffic, ensuring it remains within the appropriate VPN as it traverses the SD-WAN fabric. Labels are applied to

packets when they enter the SD-WAN fabric and are used to maintain VPN association across the overlay.

Labels work alongside the Overlay Management Protocol (OMP), enabling dynamic routing and efficient forwarding across multiple transport networks.

Key Functions of Labels

1. **VPN Identification:** Labels are used to identify which VPN a packet belongs to, ensuring that it is forwarded only within the corresponding VPN.
2. **Efficient Forwarding:** Labels enable efficient forwarding across the SD-WAN overlay by marking packets for their intended VPN, allowing for rapid, policy-based routing decisions.
3. **Integration with OMP:** Labels are exchanged via OMP between SD-WAN edge devices and vSmart controllers, allowing each device to understand the VPN association of incoming traffic.
4. **Interoperability with TLOCs:** Labels are used in conjunction with Transport Locators (TLOCs), which represent the physical paths over which VPN-labeled traffic is sent. TLOCs handle the actual transport selection, while labels ensure the traffic remains within the correct VPN context.

Label-Based Forwarding Process in Catalyst SD-WAN

When a packet enters the SD-WAN fabric, it is assigned a label that corresponds to its VPN. This label is used to ensure that the packet remains within its assigned VPN as it traverses the SD-WAN network. Here's how the label-based forwarding process works:

1. **Label Assignment:**
 - When traffic enters a cEdge router, the router assigns a label based on the VPN to which the traffic belongs.
 - This label is used to tag the packet, identifying it as belonging to a specific VPN.
2. **Label Propagation through OMP:**
 - The cEdge router advertises labeled routes to the vSmart controller using OMP.
 - These advertisements include the label, VPN identifier, and other necessary routing information, allowing vSmart to distribute the labeled route to other edge routers in the SD-WAN fabric.
3. **Label-Based Forwarding Decision:**
 - When a packet arrives at an SD-WAN edge router, the router checks the label to determine the VPN membership and forwarding policy for that packet.
 - The router then makes a forwarding decision based on the label and policy information, ensuring that the packet is routed according to its VPN's specific requirements.
4. **TLOC Selection and Transport Encapsulation:**
 - In addition to the label, the router selects a **TLOC** (Transport Locator) that represents the physical transport path for the packet. The TLOC selection considers factors like path latency, jitter, and policy preferences.
 - The packet is then encapsulated within an IPsec tunnel and forwarded over the chosen TLOC, while maintaining its VPN association via the label.
5. **End-to-End VPN Integrity:**
 - The packet retains its VPN label as it travels across the SD-WAN fabric, ensuring that it remains isolated from other VPNs.

- When the packet reaches its destination edge router, the router removes the label and forwards the packet according to the local VPN routing table.

Example of VPN and Label Functionality

Consider an organization that has segmented its network traffic into multiple VPNs: one for corporate traffic (VPN 1), one for guest traffic (VPN 2), and another for IoT devices (VPN 3).

1. **Traffic Segmentation:**
 - Corporate traffic is assigned to VPN 1, guest traffic to VPN 2, and IoT traffic to VPN 3. Each VPN has unique routing and security policies based on its traffic type.
2. **Label Assignment and Forwarding:**
 - When corporate traffic from VPN 1 enters a branch cEdge router, the router assigns a label corresponding to VPN 1.
 - This label accompanies the packet as it traverses the SD-WAN network, ensuring it is handled according to VPN 1's routing and policy settings.
3. **OMP-Based Route Distribution:**
 - The cEdge router advertises VPN 1's routes, along with the corresponding label, to the vSmart controller via OMP.
 - The vSmart controller propagates these routes to other SD-WAN routers, allowing them to recognize and forward VPN 1 traffic appropriately.
4. **Transport Locator and Policy Enforcement:**
 - The packet is forwarded across the SD-WAN fabric using the most optimal TLOC path, which is selected based on network conditions and configured policies.
 - When the packet arrives at the destination edge router, it is delivered within the VPN 1 context, preserving end-to-end segmentation and policy compliance.

TL/DR

In Cisco Catalyst SD-WAN, **VPNs** and **labels** work together to achieve secure, segmented, and policy-driven traffic management:

- **VPNs:** Provide logical separation of network traffic, enabling isolated routing domains with unique security and access policies. This segmentation supports multi-tenancy, traffic isolation, and efficient resource allocation.
- **Labels:** Identify and maintain VPN association as traffic traverses the SD-WAN fabric. Labels, in conjunction with TLOCs, facilitate efficient label-based forwarding, ensuring that each packet adheres to its VPN's routing policies and remains within its designated logical domain.

By integrating VPNs and labels, Cisco Catalyst SD-WAN delivers a highly flexible, secure, and scalable solution for managing complex, multi-tenant WAN environments. This approach allows organizations to maintain strict control over network segmentation, enabling advanced traffic policies, service chaining, and secure multi-path forwarding across distributed locations.

TLOC OMP Route Type

Cisco Catalyst SD-WAN: OMP TLOC Route Type

Cisco Catalyst SD-WAN architecture is highly dependent on the Overlay Management Protocol (OMP) for routing, policy management, and ensuring proper communication between different components of the SD-WAN fabric. OMP functions as the control plane protocol that advertises routes, policies, and tunnel information (TLOCs) between Cisco vSmart controllers and vEdge routers (or Catalyst routers). One of the key elements in SD-WAN routing is the **Transport Locator (TLOC)**, which represents the attachment point of a WAN edge device to the physical transport networks, such as MPLS, Internet, or LTE.

This write-up provides an in-depth look at the **OMP TLOC route type** and its critical operations, including its role in Network Address Translation (NAT) and how it interacts with the **Session Traversal Utilities for NAT (STUN)** protocol as defined by **RFC 5389** to traverse NAT boundaries.

Understanding OMP and TLOCs

OMP is the core protocol within the Cisco SD-WAN overlay and carries essential control plane information, including the following:

- **OMP Routes:** These are prefixes (IP subnets) that represent the data-plane routing destinations across the SD-WAN.
- **TLOC Routes:** These routes identify the transport locations of WAN edge devices. A TLOC is a combination of an IP address, color (identifying the transport type), and an encapsulation type (GRE or IPsec).
- **Service Routes:** Service-related routing information for firewall, IPS/IDS, load balancers, etc.

Components of a TLOC

A TLOC is defined by three primary attributes:

1. **System IP:** This is the unique identifier of the WAN edge device (similar to a router ID in traditional routing).
2. **Color:** Defines the transport link type (e.g., MPLS, public-internet, or LTE).
3. **Encapsulation:** Specifies whether the communication between edge devices happens over **GRE** or **IPsec** tunnels.

In SD-WAN networks, TLOCs are used to form the transport infrastructure between sites. Edge devices advertise their TLOCs via OMP to the vSmart controllers, which in turn advertise these TLOCs to other edge devices. TLOCs form the foundation of how SD-WAN edge routers reach each other over multiple transport networks.

OMP TLOC Route Type: Key Operations

TLOC routes play a central role in building the overlay network and determining the paths through which data plane traffic flows. OMP uses TLOCs to determine the next-hop paths for forwarding traffic. When an edge router advertises an OMP route, it includes TLOC

information, indicating how the destination can be reached via specific transport locators. These advertisements contain critical metadata:

- **TLOC IP Address:** The actual IP address of the WAN interface on the edge router connected to a transport (Internet/MPLS).
- **TLOC Color:** The identifier for the transport link (e.g., "mpls," "internet," "lte").
- **TLOC Encapsulation:** Indicates whether the transport link uses GRE or IPsec for tunneling.

NAT Considerations for OMP TLOCs

In environments where devices sit behind NAT (Network Address Translation), special considerations are required for TLOC operations. NAT plays a significant role when WAN edge devices are deployed behind firewalls or NAT-enabled routers, which often happens when edge routers connect to public internet transport.

When a TLOC is behind a NAT device, the IP address that the router advertises as part of the TLOC might be translated by the NAT device before being forwarded to the vSmart controller. The OMP must account for this NAT behavior and ensure that devices behind NAT can still establish IPsec or GRE tunnels to other SD-WAN nodes. This is where **STUN (Session Traversal Utilities for NAT)** becomes relevant.

STUN (RFC 5389) and NAT Traversal

In the context of Cisco SD-WAN, the **STUN protocol** as defined in **RFC 5389** is used to traverse NAT and establish connectivity between devices. STUN helps WAN edge devices discover their **public IP addresses** when they are behind a NAT. This process is essential for ensuring that TLOCs advertised from NATed devices contain accurate information about how they can be reached by other SD-WAN nodes.

STUN operates by sending special requests from the SD-WAN edge device to a STUN server (typically the vBond orchestrator in Cisco SD-WAN). The response from the STUN server helps the edge device learn its public-facing IP address and port as seen by external peers. This public IP is then used in TLOC advertisements so that other devices can initiate tunnels to the edge device, even when it is behind a NAT.

Key Steps in STUN Operations for NAT Traversal:

1. **Discovery of Public IP and Port:** When a WAN edge device behind NAT needs to establish a connection with another device (e.g., a vSmart controller or another edge router), it initiates a STUN request to the vBond orchestrator. The vBond replies with the public IP address and port from which the request originated. This public IP and port are what other devices will use to reach this TLOC.
2. **TLOC Advertisement with Public IP:** Once the WAN edge device behind NAT learns its public IP and port via STUN, it advertises its TLOC route to the vSmart controller, including the public IP as the TLOC IP. Other edge devices will use this public IP to reach the NATed device.
3. **Encapsulation with NAT Detection:** Cisco SD-WAN detects the NAT state of the connection and adjusts the encapsulation accordingly. If NAT is detected, IPsec encapsulation will include necessary NAT traversal settings, such as NAT-T (NAT

Traversal in IPsec), which allows IPsec traffic to pass through NAT devices by encapsulating IPsec in UDP.

4. **Connectivity Establishment:** Once the correct TLOC information (including the public IP) has been advertised, other WAN edge devices can establish secure tunnels (GRE or IPsec) to the NATed device using the public IP obtained via STUN.

How OMP and TLOC Routes Handle NAT

In Cisco SD-WAN, OMP and TLOC routes must handle two primary aspects of NAT traversal:

- **Public IP Mapping:** OMP relies on STUN to ensure that the correct public IP of NATed devices is advertised. If this mapping is not properly handled, devices will be unable to form tunnels, leading to communication failures.
- **Port Preservation and Firewall Traversal:** For successful tunnel establishment, the correct combination of public IP and port must be known. STUN helps in determining both, enabling seamless communication across firewalls and NAT devices.

When a WAN edge device behind NAT sends traffic to an external site, the NAT device replaces the private source IP address with the public IP address. STUN ensures that this public IP is used in OMP TLOC advertisements. This way, other edge devices can correctly form tunnels to the NATed device using the public IP.

Use Cases of OMP TLOC with NAT and STUN

1. **Edge Device Behind a Home Router:** In scenarios where branch devices are connected to the internet through a home router performing NAT, the edge device must rely on STUN to discover its public IP and advertise this IP as part of its TLOC. This ensures that external edge routers can establish tunnels with the device using its public-facing address.
2. **MPLS and Internet Transport Separation:** Cisco SD-WAN deployments often involve hybrid transports, such as MPLS and public internet. When edge devices use MPLS for private communication and internet transport for backup or public communication, the public IP used on the internet transport may be NATed. STUN helps ensure that the correct IP is advertised to facilitate communication over the public internet.
3. **Multi-Transport with NATed Public Internet:** For multi-transport SD-WAN architectures where edge devices use both MPLS and internet, devices behind NATed public internet must rely on STUN to establish reliable connectivity via public IPs. OMP ensures that both MPLS and internet TLOCs are properly advertised, even when NAT is involved.

Conclusion

The Cisco Catalyst SD-WAN solution relies heavily on the operations of OMP and TLOC routes to build the overlay network across various transport mechanisms like MPLS, internet, and LTE. When edge devices are placed behind NAT, STUN (RFC 5389) is essential for discovering public IPs and ensuring that these are accurately reflected in TLOC route advertisements. The NAT traversal processes, aided by STUN, allow secure and reliable

tunnel formation between edge devices, even when NAT devices are in the communication path. This capability is crucial for enabling seamless, scalable, and resilient SD-WAN deployments across diverse network environments.

Color Restrict via Configuration Group

The **Color Restrict** feature in Cisco Catalyst SD-WAN is a critical functionality that allows network administrators to control how traffic is routed across different transport circuits in a multi-transport network. In a typical SD-WAN deployment, traffic is routed over multiple WAN transports such as MPLS, broadband, LTE, and other WAN links. These transports are identified using a tag called **Color**, which categorizes the type of transport. By using **Color Restrict**, administrators can granularly restrict certain traffic types from using particular transport paths, providing control over how and where traffic flows across the network.

This detailed technical write-up explores the **Color Restrict** feature in depth, explaining its mechanics, use cases, and how it enhances traffic management in an SD-WAN environment.

Overview of SD-WAN Colors

In Cisco SD-WAN, each WAN transport link is assigned a **Color** that defines the type of transport network being used. These colors are essential in identifying paths across the SD-WAN fabric and are central to how data traffic is distributed over multiple paths. For example, a color can represent different types of WAN connections like **MPLS, public-internet, private1, biz-internet, lte**, etc.

Cisco has predefined certain colors to distinguish between public and private networks:

- **Public Colors:** These are colors associated with internet-based links, such as:
 - public-internet
 - lte
- **Private Colors:** These are colors assigned to private WAN circuits like MPLS or dedicated leased lines:
 - mpls
 - private1, private2

The colors associated with each transport path are advertised between SD-WAN routers using the Overlay Management Protocol (OMP) and are part of the **Transport Locator (TLOC)**, which is essential for establishing tunnels between edge devices over these paths.

Color Restrict Feature: Overview and Purpose

The **Color Restrict** feature allows administrators to prevent the formation of data plane tunnels over specific WAN transports. Normally, SD-WAN devices establish secure tunnels between sites based on available transport paths, which are identified using TLOCs. Without Color Restrict, SD-WAN attempts to establish tunnels across all available TLOCs for

redundancy and high availability. However, certain situations may call for preventing tunnels from being established over a particular transport color.

The **Color Restrict** feature is applied at the transport level. It prevents OMP from using certain colors for creating tunnels between edge devices, effectively limiting or restricting traffic from traversing those paths.

Reasons for Using Color Restrict

There are several scenarios where **Color Restrict** is valuable in Cisco SD-WAN deployments:

1. Transport Circuit Cost Management

In some environments, certain WAN circuits are significantly more expensive than others. For instance, MPLS links often incur higher costs compared to broadband internet circuits. In such cases, you may want to avoid routing less critical or non-business traffic over more expensive MPLS circuits. By using Color Restrict, you can ensure that specific traffic types avoid MPLS and instead use cheaper circuits like public internet or broadband.

2. Latency-Sensitive Traffic Routing

Real-time applications like VoIP, video conferencing, and other latency-sensitive traffic must be routed over low-latency links. For example, you might want to restrict traffic from using LTE or satellite links, which tend to have higher latency compared to MPLS or private circuits. Color Restrict allows administrators to block these latency-sensitive applications from establishing tunnels over high-latency circuits, ensuring they use faster, more reliable paths.

3. Link Reliability and Stability

Certain WAN links, such as those using cellular LTE or satellite, may be less reliable due to higher variability in bandwidth and link stability. In environments where reliability is crucial, administrators may choose to restrict traffic from using these less stable paths, except for failover scenarios. **Color Restrict** helps ensure that primary and backup transport policies are enforced according to the performance and reliability requirements of the organization.

4. Security Considerations

In environments where security is paramount, administrators may wish to prevent sensitive traffic from traversing public networks like the public internet. Public internet links are more prone to attacks and can present security risks, even though SD-WAN employs strong encryption techniques such as IPsec. By using Color Restrict, network admins can ensure that sensitive data only flows through more trusted private links like MPLS or dedicated leased lines.

5. Selective Tunneling for Failover and Redundancy

Color Restrict allows for better control over failover policies. For example, a company may want to use a primary MPLS link for its main transport and only fail over to a broadband or LTE link in case of failure. By restricting the broadband or LTE color for normal operations, the company can ensure that these links are only used as backup transport options, avoiding unnecessary tunnel formation on backup links unless a failure occurs.

6. Compliance and Regulatory Reasons

Certain industries have strict compliance and regulatory requirements regarding data transmission, such as healthcare and finance. These industries may have policies that prohibit the use of public internet links for transmitting certain types of data. By using **Color Restrict**, administrators can prevent certain traffic from traversing public networks, ensuring compliance with regulations and avoiding penalties for data transmission violations.

Technical Operation of Color Restrict

The **Color Restrict** feature works by preventing the formation of data plane tunnels over specified colors between SD-WAN edge devices. Here's how it operates technically:

1. **TLOC Advertisement with Restricted Color:** Each SD-WAN edge device advertises its TLOC information to the **Controller** (formerly vSmart) via OMP. The TLOC advertisement includes the transport color, encapsulation type (GRE or IPsec), and other attributes such as bandwidth and latency.

When Color Restrict is applied to a particular color, the TLOC advertisement for that color is flagged in such a way that OMP will not establish tunnels over it. The restriction is enforced on both sides of the tunnel, preventing traffic from using the restricted color.

2. **Control Plane (OMP) Behavior:** OMP running on the **Controller** is responsible for receiving and processing all TLOC routes from the SD-WAN edge devices. When a device advertises a TLOC with a restricted color, the **Controller** marks that route as unusable for establishing tunnels. The restricted TLOC is not advertised to other edge devices, ensuring that no edge devices attempt to use it for data traffic.

If two devices try to communicate and one of them is advertising a restricted color, the other edge device will not form a tunnel using that TLOC, even if the transport is operational.

3. **Data Plane (IPsec Tunnel) Behavior:** Once the color is restricted, no IPsec (or GRE) tunnels are established over the corresponding transport path. The restriction is enforced on the data plane, meaning that even though the interface might be active and have network connectivity, no encrypted tunnels are formed over that link.
4. **Policy-Based Control:** Color Restrict is typically configured as part of an overall SD-WAN policy. These policies are defined using **Manager** (formerly vManage), where an administrator can configure a device-specific or network-wide policy to restrict certain colors. The policy can be applied based on application type, SLA class, or other criteria, allowing for granular control over traffic flow.

For example, you might create a policy that restricts **public-internet** color for real-time traffic, while allowing best-effort traffic to use it as a backup. Similarly, MPLS circuits might be restricted for non-business applications.

Configuration of Color Restrict in Cisco SD-WAN

Configuring the **Color Restrict** feature involves the following steps in **Manager** (formerly vManage):

1. **Access Manager and Navigate to Configuration:** Log into **Manager**, navigate to the **Configuration** tab, and select **Policies**.
2. **Define Traffic Policy:** Create a traffic policy that specifies which colors should be restricted. This policy can be defined based on application type, device, or transport preference.
3. **Apply Color Restrict to TLOCs:** In the TLOC configuration section, apply the **Color Restrict** option to the colors you want to restrict. For example, you might apply this restriction to **Ite** or **public-internet** colors depending on your use case.
4. **Policy Application:** After defining the policy, apply it to the appropriate SD-WAN edges or device groups within your topology. Once deployed, the policy will take effect, preventing the formation of tunnels on restricted colors.

Use Cases and Benefits of Color Restrict

Color Restrict provides several advantages in multi-transport SD-WAN deployments:

- **Cost Efficiency:** Prevents less critical traffic from using expensive MPLS links, ensuring that lower-cost links such as broadband or public internet are utilized appropriately.
- **Enhanced Performance:** Ensures that latency-sensitive applications like voice and video only use the best-performing links, leading to improved user experience and network performance.
- **Security:** Helps prevent sensitive traffic from using less secure public networks, ensuring that data is routed over trusted private links in compliance with organizational policies.
- **Operational Simplicity:** Allows for straightforward implementation of failover strategies, ensuring that backup links are only used when necessary, reducing overall operational complexity.

Conclusion

The **Color Restrict** feature in Cisco Catalyst SD-WAN is a powerful tool that provides precise control over traffic flow in multi-transport networks. By allowing network administrators to restrict certain colors from being used for tunnel establishment, it enhances cost management, improves traffic performance, and adds an extra layer of security. This feature is particularly beneficial in environments with varied WAN links, enabling optimal routing based on business and technical requirements.

By using **Manager** (formerly vManage) to configure policies and employing **Controller** (formerly vSmart) for control plane operations, the **Color Restrict** feature integrates

seamlessly into the overall SD-WAN fabric. The ability to restrict certain colors allows enterprises to prevent traffic from traversing costly or insecure links, ensuring that business-critical traffic is routed over preferred and optimal paths.

From improved performance of latency-sensitive applications, to reducing WAN costs, to ensuring compliance with industry-specific regulations, **Color Restrict** provides an essential level of granularity in managing SD-WAN deployments. By controlling which paths traffic can use, organizations can deliver more efficient, reliable, and secure network operations, leveraging Cisco SD-WAN's flexible architecture to meet modern network demands.

Overall, **Color Restrict** is an essential feature that contributes to an optimized SD-WAN architecture, allowing enterprises to better manage and fine-tune the utilization of their available network resources.

The Impact of Public and Private TLOC Colors on IPsec Tunnel formation

In Cisco Catalyst SD-WAN, **IPsec tunnels** are used to secure communication between edge devices (such as routers) across various transport networks, including public networks where **Network Address Translation (NAT)** is often deployed. NAT presents a challenge when forming IPsec tunnels because it alters the IP address and possibly the port numbers involved in communication, which can break traditional IPsec tunneling mechanisms. However, Cisco SD-WAN uses a combination of NAT traversal techniques, including **NAT-T (NAT Traversal)** and **STUN (Session Traversal Utilities for NAT)**, to allow IPsec tunnels to function correctly across NAT environments.

Mechanism of IPsec Tunnels Across NAT in Cisco Catalyst SD-WAN

1. NAT Detection via STUN (RFC 5389)

Before IPsec tunnels can be formed, Cisco Catalyst SD-WAN uses the **STUN (Session Traversal Utilities for NAT)** protocol to detect whether a WAN edge device is behind a NAT. STUN operates by sending a request from the edge router to a remote server (often the **Validator** or a vBond orchestrator), which in turn replies with the public IP address and port number that the request came from. This allows the edge router to discover its public-facing IP address as seen from outside the NAT.

Here's how the NAT detection process works:

- The WAN edge device sends STUN messages to the **Validator** (formerly vBond) to determine if NAT is being used on the public-facing interface.
- The **Validator** replies with the observed public IP and port. If the observed IP differs from the device's local IP, it confirms that NAT is present.

Once the device detects NAT, it stores this information and adjusts its IPsec configuration accordingly, ensuring that the correct public IP is used during tunnel setup.

2. NAT Traversal (NAT-T) in IPsec

After detecting NAT, Cisco Catalyst SD-WAN automatically enables **NAT Traversal (NAT-T)** for IPsec tunnels. NAT-T allows IPsec traffic to traverse NAT by encapsulating the IPsec traffic in UDP packets. This encapsulation ensures that the NAT device treats the IPsec traffic like regular UDP traffic, preventing the NAT device from breaking the IPsec encapsulation.

The NAT-T process works as follows:

- **UDP Encapsulation:** IPsec ESP (Encapsulating Security Payload) traffic is encapsulated inside UDP packets. By using UDP port 4500, NAT devices can handle IPsec traffic without breaking the encryption or causing communication issues.
- **Keepalives:** NAT-T often uses periodic **keepalive messages** to keep the NAT mapping active. NAT devices typically timeout and remove entries from their NAT table if traffic is inactive for a certain period. The IPsec endpoints send these keepalives to prevent the NAT device from prematurely closing the session.

3. TLOC Advertisement and NAT Information

In Cisco SD-WAN, WAN edge devices advertise **Transport Locators (TLOCs)** to inform other devices about how to reach them across different WAN transports (e.g., MPLS, internet, LTE). When a WAN edge device is behind NAT, it advertises the **public IP address** and associated **UDP port** that was detected via STUN in its TLOC information. This allows other devices in the SD-WAN fabric to form tunnels using the correct public IP and port, even if the device is behind a NAT.

This process ensures that all edge devices in the SD-WAN fabric can correctly form IPsec tunnels to edge devices behind NAT.

4. IPsec Tunnel Establishment

Once NAT-T and STUN have been used to determine the correct public IP and port, IPsec tunnels are established over the public network as follows:

1. **TLOC Advertisement:** The edge device behind NAT advertises its TLOC with the public IP and port number learned from STUN. This TLOC includes additional information, such as the **Color** (e.g., `public-internet`, `mpls`) and **Encapsulation** (IPsec in this case).
2. **Tunnel Formation:** Other WAN edge devices use this TLOC to establish IPsec tunnels to the device behind NAT. The IPsec tunnel uses UDP encapsulation (NAT-T) to ensure that traffic can traverse the NAT.
3. **Data Encryption:** Once the tunnel is established, all data traffic between the edge devices is encrypted using IPsec, ensuring security over public networks.

5. Tunnel Maintenance and Resiliency

- **Keepalive and Tunnel Monitoring:** To keep the tunnel alive and ensure that NAT mappings remain active, Cisco SD-WAN regularly sends keepalive messages. This

also ensures that in case of any network issues or changes in NAT behavior, the IPsec tunnel can be reestablished without significant downtime.

- **Automatic Failover:** If the NAT or public IP address changes, the Cisco SD-WAN fabric is designed to automatically re-establish tunnels using the updated NAT information.

Key Benefits of Cisco SD-WAN's NAT Handling for IPsec

1. **Seamless Tunnel Creation Across NAT:** By leveraging STUN and NAT-T, Cisco SD-WAN allows secure IPsec tunnels to be created seamlessly across NAT devices. This is crucial in deployments where WAN edge devices may be connected to public networks behind NAT routers.
2. **Automatic Detection and Configuration:** The system automatically detects when NAT is present and adjusts configurations accordingly, reducing the need for manual intervention by network administrators.
3. **End-to-End Encryption:** Even when traversing NAT devices, IPsec ensures that traffic remains encrypted end-to-end, providing high levels of security even over public internet connections.
4. **Resiliency:** By using keepalives and automatic failover mechanisms, Cisco SD-WAN ensures that IPsec tunnels remain stable and reliable even when NAT configurations change.

Summary of Steps for IPsec Tunnel Formation Across NAT

1. **STUN Discovery:** The WAN edge device uses STUN to discover its public IP address and port.
2. **TLOC Advertisement:** The edge device advertises its public IP and port in the TLOC via OMP.
3. **NAT-T (UDP Encapsulation):** NAT Traversal is enabled, and IPsec traffic is encapsulated in UDP packets.
4. **Tunnel Establishment:** IPsec tunnels are established using the public IP and UDP port, allowing traffic to traverse NAT.
5. **Tunnel Maintenance:** Keepalive messages and automatic NAT handling ensure the tunnel remains operational.

This combination of STUN and NAT-T allows Cisco Catalyst SD-WAN to handle the challenges of NAT in public WAN environments, ensuring that secure IPsec tunnels can be established and maintained even when edge devices are behind NAT.

TLOC Carrier

In Cisco Catalyst SD-WAN, **Transport Locators (TLOCs)** are one of the foundational elements for creating overlay tunnels between SD-WAN edge devices across various transport networks (MPLS, broadband, LTE, etc.). **TLOC Carrier designations** are an advanced feature within the TLOC framework that allow network administrators to control how traffic traverses different transport networks, particularly in complex environments with multiple service providers or carriers. This feature provides finer granularity in traffic management by differentiating TLOCs based on **carrier networks**.

In this detailed explanation, we'll explore the application and operation of TLOC Carrier designations in Cisco Catalyst SD-WAN, how they are configured, and their impact on network performance, traffic engineering, and service provider interaction.

Key Concepts: What Are TLOCs in Catalyst SD-WAN?

Before diving into TLOC Carrier designations, let's quickly review the concept of **TLOCs** in Cisco SD-WAN.

- A **TLOC (Transport Locator)** is a unique identifier that represents the attachment point of a WAN edge device to a transport network. It identifies how a particular device is connected to a specific WAN link, whether it's MPLS, broadband, or LTE.
- A TLOC is defined by three key attributes:
 1. **System IP Address:** The unique identifier of the device (similar to a router ID).
 2. **Color:** Defines the type of transport (e.g., MPLS, public-internet, private1, etc.).
 3. **Encapsulation:** The tunneling method used between devices, usually **GRE** or **IPsec**.

These TLOCs form the backbone of SD-WAN tunnel creation, with the **Overlay Management Protocol (OMP)** exchanging information about available TLOCs between the WAN edge devices and the **vSmart Controller** (or **Controller** in newer terminology).

What is a TLOC Carrier Designation?

A **TLOC Carrier** is a label that is used to group TLOCs that are associated with specific carrier networks. In environments where multiple WAN transports are used—especially in global networks where multiple ISPs or carriers are involved—TLOC Carrier designations help define which service provider or WAN transport a particular TLOC belongs to.

Why Use TLOC Carrier Designations?

There are several reasons for employing TLOC Carrier designations in an SD-WAN environment:

1. **Traffic Engineering and Path Selection:**
 - In large-scale deployments with multiple WAN providers or when multiple circuits are available, TLOC Carrier designations allow for more granular control of traffic flow. Network administrators can specify how traffic should be routed across different carrier networks by associating TLOCs with specific carriers.
2. **Service Provider Segmentation:**
 - By using TLOC Carrier designations, different service providers (or WAN circuits) can be segmented. This segmentation can help manage traffic between different carriers and avoid unintended routing or packet loss that might occur due to carrier-level issues.
3. **Policy-Based Routing:**
 - TLOC Carrier designations enable the creation of policies that direct traffic based on the carrier network associated with the TLOC. This is useful for situations where traffic over a certain carrier may need to be preferred or avoided based on performance, cost, or contractual obligations.
4. **Optimized Redundancy:**

- TLOC Carriers help enhance redundancy. Instead of treating all WAN circuits equally, administrators can assign primary traffic paths to specific carriers and backup paths to secondary carriers. This ensures more intelligent use of resources and provides better control over failover scenarios.
5. **Control Over Regulatory or Compliance Requirements:**
- In some regions or industries, there may be restrictions on which service providers can be used for certain types of traffic due to compliance or regulatory requirements. TLOC Carrier designations allow administrators to enforce policies that ensure traffic is routed only through approved carriers.

Application of TLOC Carrier Designations in a Cisco Catalyst SD-WAN Environment

1. **Classifying WAN Transports by Carrier:**
 - When a WAN edge device is connected to multiple WAN circuits (e.g., one MPLS link and one internet link), each WAN circuit can be associated with a **TLOC**. By using TLOC Carrier designations, these TLOCs can be further classified by the WAN carrier. For instance, the MPLS link from **Carrier1** might be given one designation, while the internet link from **Carrier2** might be given a different designation.
2. **Policy Creation Based on TLOC Carrier:**
 - Once the TLOC Carrier designation is set, Cisco Catalyst SD-WAN allows the administrator to create traffic routing policies that use these designations as criteria. For example, high-priority traffic (such as VoIP) might be routed exclusively through **Carrier1** (MPLS), while best-effort traffic might be routed through **Carrier2** (internet).
 - Similarly, failover policies can be created to ensure that if **Carrier1** becomes unavailable, traffic automatically shifts to **Carrier2** without any manual intervention.
3. **Service-Level Agreements (SLA) Enforcement:**
 - In situations where different carriers provide different service levels (latency, bandwidth, reliability), TLOC Carrier designations can help enforce service-level agreements (SLAs). Cisco SD-WAN can monitor performance metrics (such as latency, jitter, and loss) and adjust traffic flows accordingly. Traffic can be moved from one carrier to another based on SLA violations or degradation in service quality.
4. **Simplified Multi-Cloud Connectivity:**
 - With businesses increasingly adopting multi-cloud architectures, WAN edge devices connected to different cloud providers can use TLOC Carrier designations to manage cloud-bound traffic. For example, traffic to **AWS** might be designated to use **Carrier1**, while traffic to **Azure** might be routed through **Carrier2**.

Operation of TLOC Carrier Designations

TLOC Carrier designations are part of the **Overlay Management Protocol (OMP)**, which manages the distribution of route and TLOC information between WAN edge devices and the central SD-WAN controller (vSmart or Controller). The operation can be broken down into the following steps:

1. **TLOC Advertisement:**

- Each WAN edge device advertises its TLOCs to the Controller, including not only the System IP, Color, and Encapsulation information, but also the **Carrier** designation, if configured.
 - This information is shared across all edge devices, allowing other devices to make informed decisions about how to route traffic to or from that particular device based on the carrier information.
2. **Policy Enforcement:**
 - Once TLOC Carrier designations are in place, policy enforcement becomes easier. Policies can be written in **vManage** (or **Manager**) to dictate how traffic should be routed based on the carrier of the underlying TLOC.
 - For example, a policy might state that only critical traffic should use TLOCs associated with **Carrier1**, and all other traffic should use TLOCs associated with **Carrier2**. Similarly, you can specify backup routes to kick in when the primary carrier's TLOCs are down.
 3. **Traffic Steering Based on Carrier:**
 - During the operation of the SD-WAN fabric, the **vSmart Controller** evaluates traffic flow policies based on TLOC information, including the TLOC Carrier designation.
 - When routing decisions are made, the Controller selects the appropriate TLOC based on the routing policy and the TLOC's Carrier designation, ensuring that traffic is forwarded using the preferred carrier whenever possible.
 4. **Performance and Redundancy Monitoring:**
 - Cisco SD-WAN monitors the performance of each transport circuit and its associated TLOCs. If a TLOC tied to a specific carrier experiences degradation (e.g., higher latency or packet loss), the SD-WAN fabric can automatically shift traffic to a different carrier based on the policies defined.
 - The system ensures continuous connectivity by using alternative TLOCs associated with different carriers in case of failure or suboptimal performance on the primary carrier.

Configuring TLOC Carrier Designations

TLOC Carrier designations can be configured via **Cisco vManage** (or **Manager**) as part of the WAN Edge configuration. Below is a simplified outline of the steps:

1. **Login to Cisco vManage (Manager):**
 - Access the vManage portal and navigate to the device template for the WAN edge router that will be configured.
2. **Configure the Transport Interface:**
 - Under the WAN transport interface configuration (e.g., an interface connected to MPLS or the internet), assign the appropriate **Color** (e.g., `mpls`, `biz-internet`) and specify the **TLOC Carrier**. This ties the TLOC to a specific carrier for later use in policies.
3. **Define Traffic Policies:**
 - Create policies based on the TLOC Carrier designation. For example, you can create a centralized control policy to route voice traffic over MPLS (Carrier1) and best-effort traffic over the internet (Carrier2).
4. **Apply and Monitor:**
 - After applying the policy, monitor the traffic flow and performance in **vManage** to ensure that the desired routing behavior is occurring based on the TLOC Carrier designations.

Benefits of Using TLOC Carrier Designations

- **Enhanced Traffic Control:** More precise control over which WAN carrier is used for different types of traffic, improving traffic management across large networks.
- **Optimized Use of WAN Links:** Ensures that traffic flows through the best available WAN carrier for its needs, reducing costs or improving performance.
- **Redundancy and Resiliency:** Allows for intelligent failover between different carriers, ensuring continuous network availability even if one WAN provider fails.
- **Simplified Operations:** Provides an easy way to manage multi-provider environments by abstracting away the complexity of dealing with different carriers.

TL/DR

TLOC Carrier designations provide a powerful mechanism in Cisco Catalyst SD-WAN to manage and optimize traffic flow in complex WAN environments where multiple transport providers are used. By associating TLOCs with specific carriers, administrators can fine-tune traffic engineering, improve redundancy, enforce SLAs, and meet compliance requirements. TLOC Carrier designations integrate seamlessly with the SD-WAN fabric, allowing traffic to be routed efficiently, ensuring continuous service, and providing high levels of operational control.

Secure Data Plane Operations

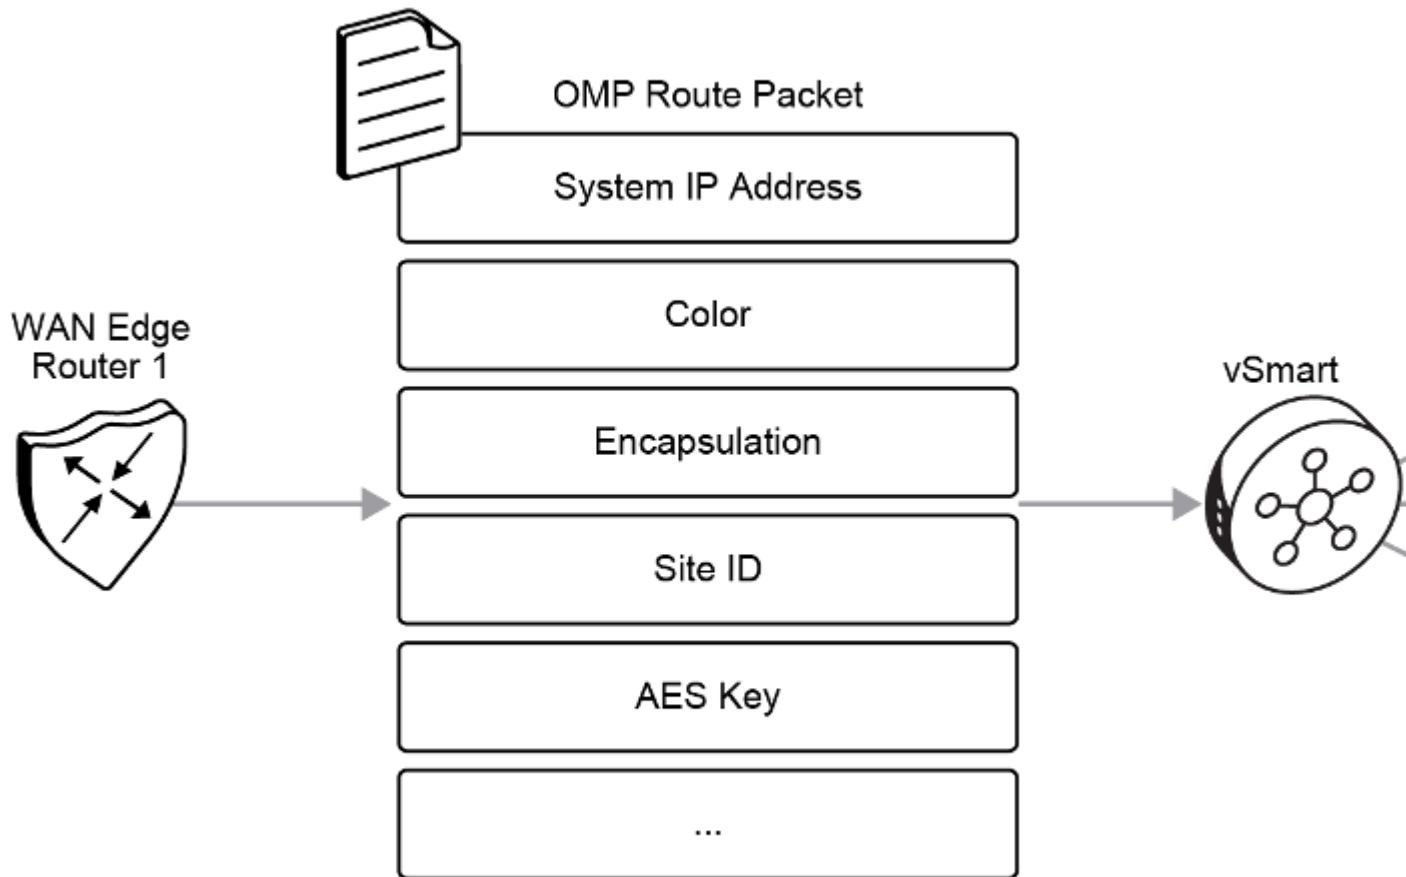
Catalyst SD-WAN Data Plane Security Overview

Cisco Catalyst SD-WAN implements robust security features to ensure secure transmission of data across wide-area networks. These key components include **authentication**, **encryption**, and **integrity checks** that protect data in transit, even across public networks.

1. Control Plane: The Foundation of Data Plane Security

In Cisco SD-WAN, the **control plane** forms the backbone of the overall security framework. The control plane validates and authenticates all devices in the network, establishing secure, tamper-proof communication channels. Once this is done, the **data plane** can use these validated connections to transmit data securely without requiring additional authentication steps for the data plane itself.

All control plane communication is encrypted using **DTLS (Datagram Transport Layer Security)**, ensuring secure exchanges of routes, policies, and device validation. This secure control plane facilitates the setup of **IPsec tunnels** between WAN Edge devices in the data plane, where data is transmitted.



2. Data Plane Authentication and Encryption

Cisco SD-WAN eliminates the need for traditional IPsec's **IKE (Internet Key Exchange)**, which can become a bottleneck in large-scale networks. Instead, SD-WAN leverages the secure control plane for efficient key exchange and data plane authentication.

Authentication

- **Traditional Key Model:** The **vSmart controller** centrally manages and distributes IPsec encryption keys to each WAN Edge device. While this method is effective, it lacks the granularity needed for more advanced security deployments.
- **Pairwise Key Model (ECDH):** A more secure model where the **vSmart controller** sends **Diffie-Hellman public values** to each device. Each WAN Edge router then generates pairwise IPsec keys using **Elliptic Curve Diffie-Hellman (ECDH)**, creating unique keys for each connection, ensuring more granular security.

Encryption

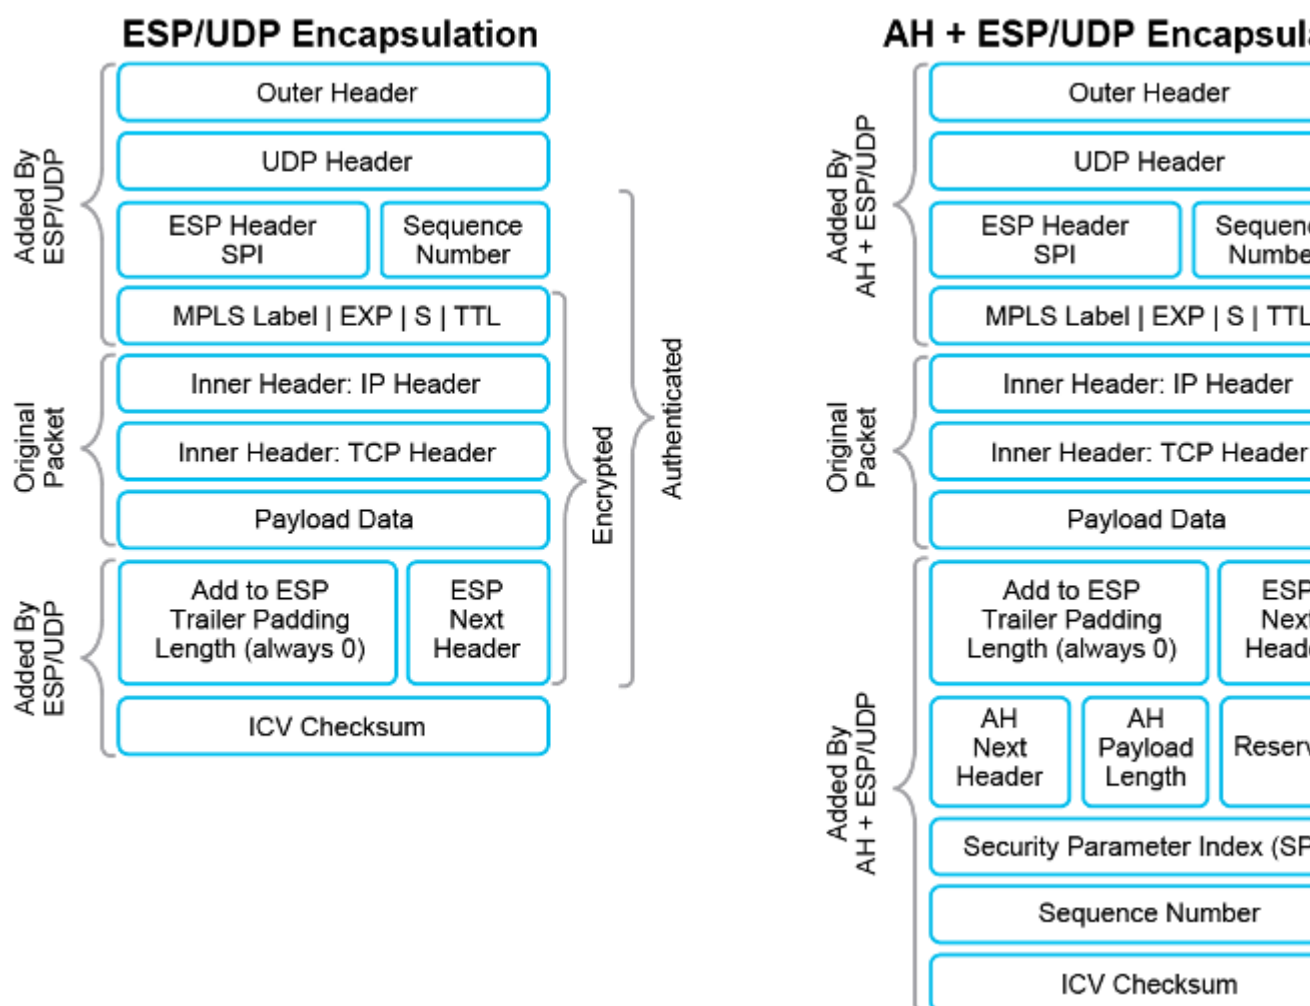
- **AES-256-GCM Encryption:** All data transmitted across the SD-WAN overlay is encrypted using **AES-256-GCM** (Advanced Encryption Standard - Galois/Counter

Mode). This ensures data confidentiality and integrity as it moves through public or private networks.

- **Modified ESP Protocol:** Cisco SD-WAN uses a modified version of **ESP (Encapsulating Security Payload)**, which protects the data packet's payload and verifies the integrity of both inner and outer packet headers. This modification mimics the **Authentication Header (AH)** protocol, ensuring that both encryption and integrity checks are applied.

Integrity

- **ESP for Data Integrity:** The **ESP protocol** ensures that data packets are protected from tampering. ESP encrypts the data payload and uses **AES-GCM** for generating hash values that verify data authenticity. If the hashes don't match, the packet is dropped.



- **Modified ESP for Outer Headers:** The modified ESP also protects the **outer IP and UDP headers**, extending data integrity to more packet components.
- **Anti-replay Protection:** Each data packet is assigned a unique sequence number, and the receiving WAN Edge router checks that the packet has a valid sequence number. Any duplicated packets are rejected, ensuring that attackers cannot replay captured packets to disrupt the network.

3. Key Exchange Without IKE

Traditional IPsec environments use **IKE** to manage secure key exchanges, but in large networks, this can result in scalability issues due to the exponential growth in key exchanges. Cisco SD-WAN replaces IKE with a more scalable, control-plane-based solution.

- **Key Distribution via vSmart:** Each WAN Edge device generates an **AES key** for its data path. This key is distributed via **OMP (Overlay Management Protocol)** route packets, along with **TLOC** information (system IP and color). The **vSmart controller** then distributes these keys to all other WAN Edge devices, avoiding the need for numerous individual handshakes.
- **Pairwise Key Generation:** When needed, **Diffie-Hellman public values** are exchanged between devices, and unique encryption keys are generated for each IPsec tunnel. This ensures that each pair of devices has its own secure encryption key, enhancing overall network security.

4. Aggressive Key Regeneration

Cisco SD-WAN implements an aggressive key regeneration mechanism to further strengthen security. By default, each WAN Edge router regenerates its AES encryption keys every **24 hours**. This ensures that keys are always fresh, limiting the potential exposure window in case of key compromise. Importantly, key regeneration occurs without dropping any traffic.

5. Monitoring Data Plane Connections with BFD

To ensure that IPsec connections between WAN Edge routers remain active, **Bidirectional Forwarding Detection (BFD)** packets are periodically exchanged. If BFD packets stop being received, the connection is marked as lost, and the **vSmart controller** is notified. This enables the SD-WAN overlay to maintain awareness of connection states between routers and take action if a connection fails.

Cisco SD-WAN does not require an explicit **SA idle timeout**. As long as BFD packets confirm the connection is alive, the associated **Security Associations (SAs)** remain active. Frequent key regeneration also eliminates the need for such timeouts, further ensuring continuous secure communication.

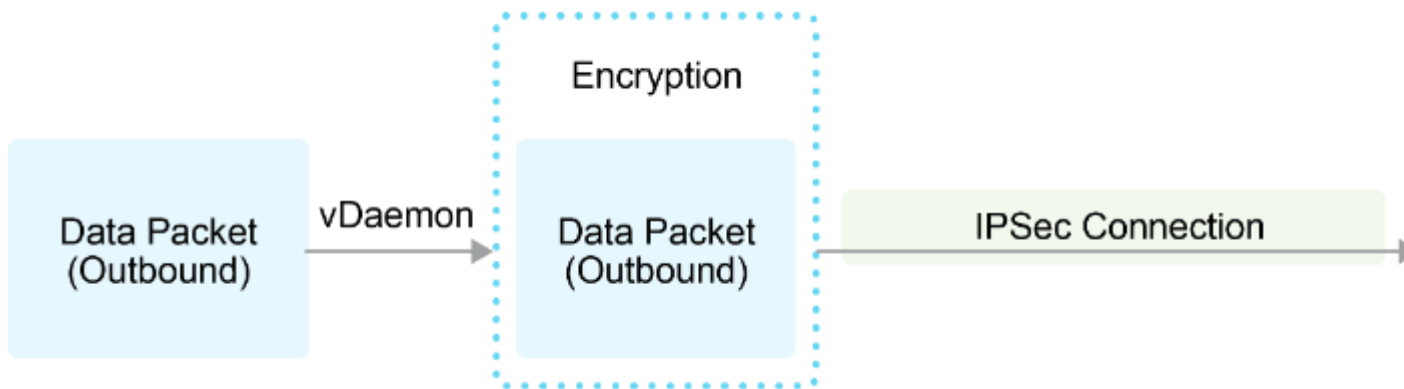
6. Data Integrity Mechanisms in SD-WAN

Cisco SD-WAN employs several mechanisms to ensure that data remains intact and unaltered during transit, which is critical for maintaining data authenticity across public and private networks.

Encapsulating Security Payload (ESP)

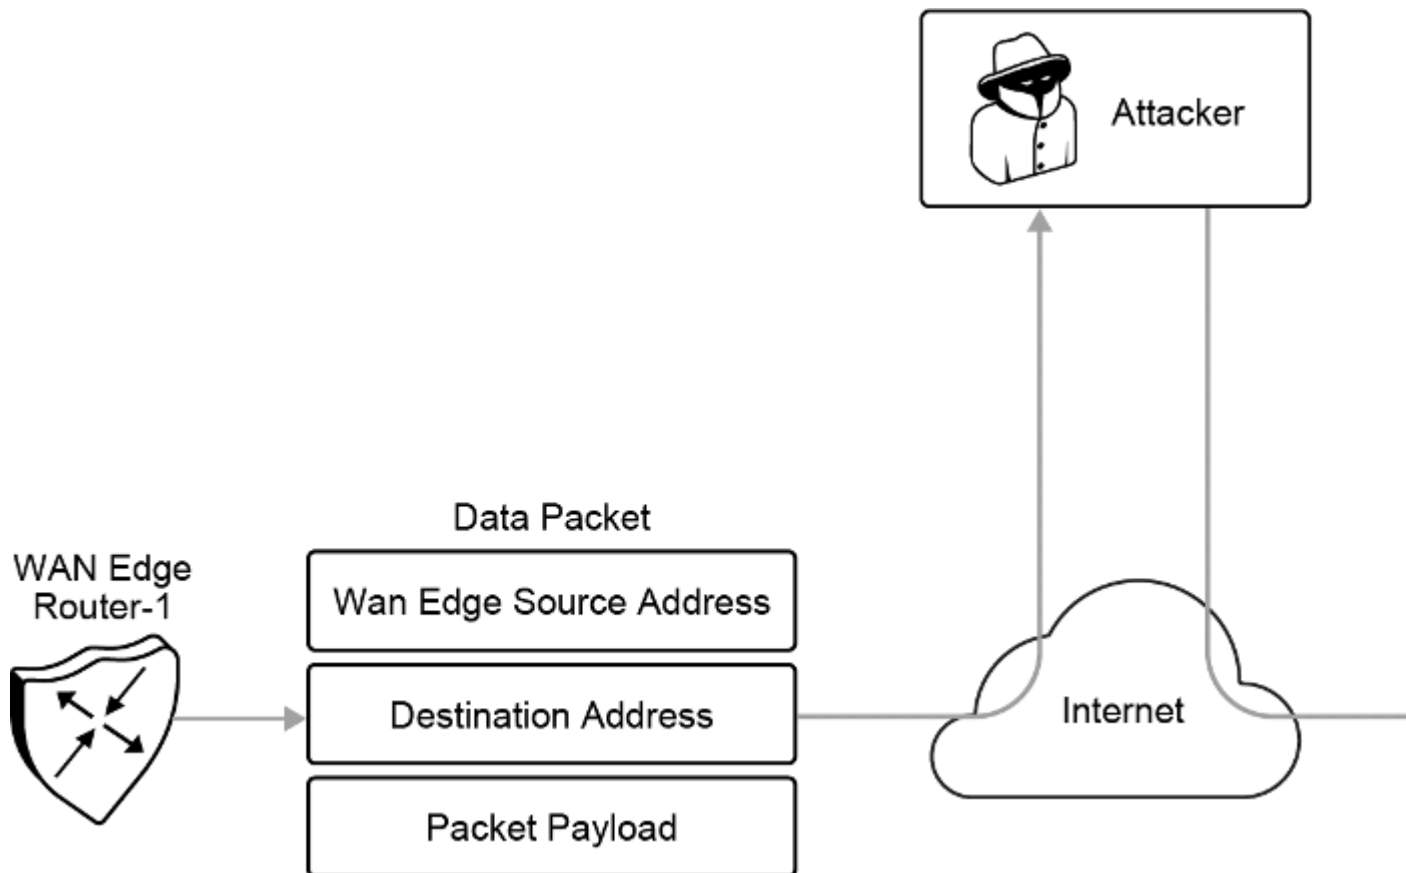
- **ESP:** The standard **IPsec ESP** protocol encrypts the data payload and inner IP header while verifying the integrity of each packet through **AES-GCM**-based hashes. This ensures that only untampered packets are accepted.

- **Modified ESP for Enhanced Protection:** Beyond standard ESP, the modified version also checks the integrity of outer IP and UDP headers, providing broader protection similar to **Authentication Header (AH)** functionality.



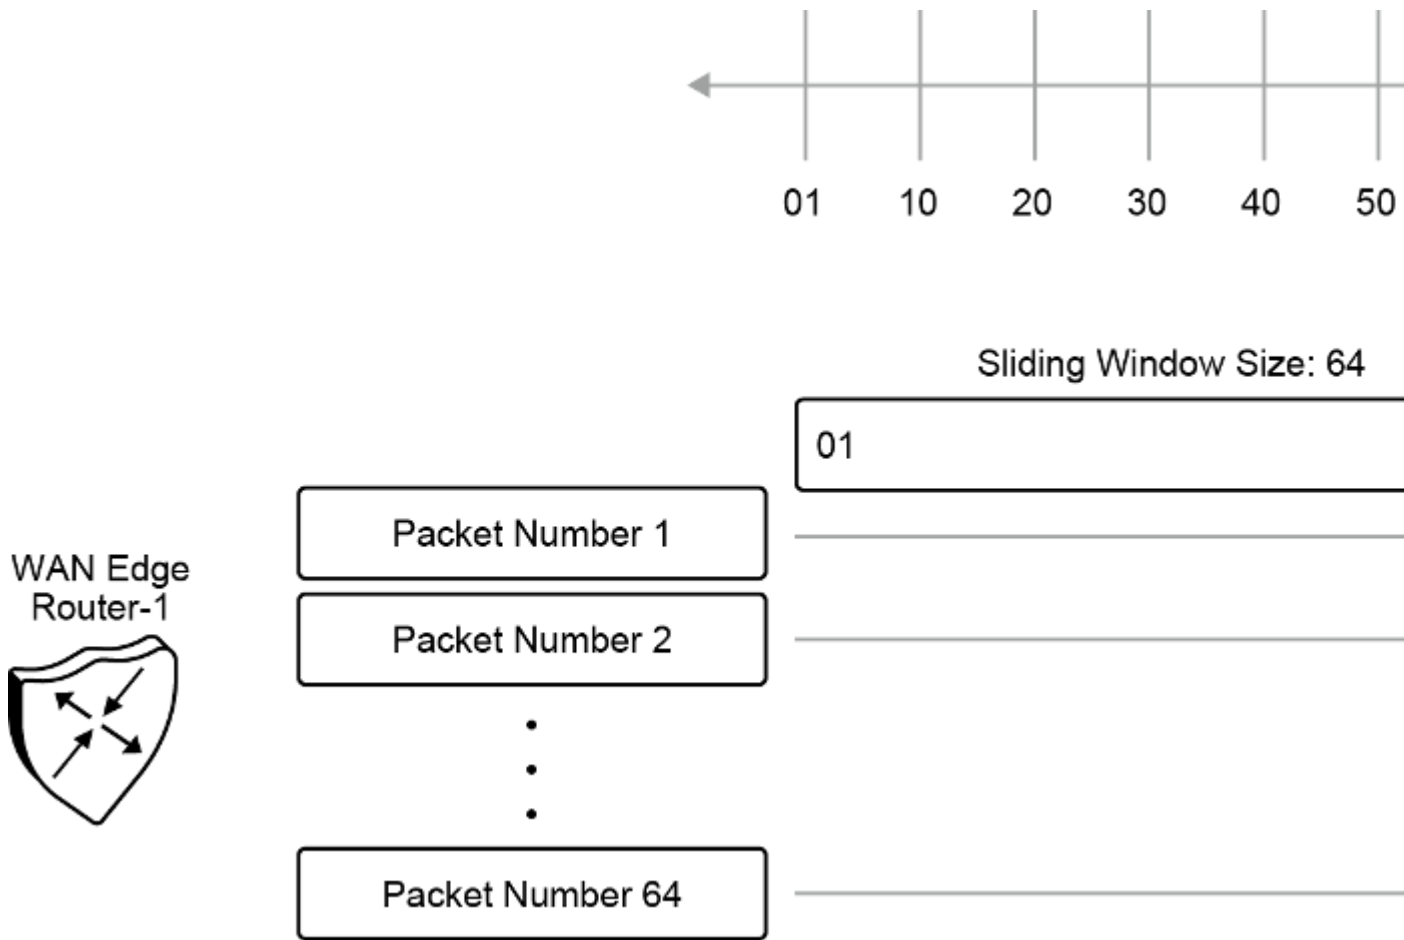
Anti-replay Protection

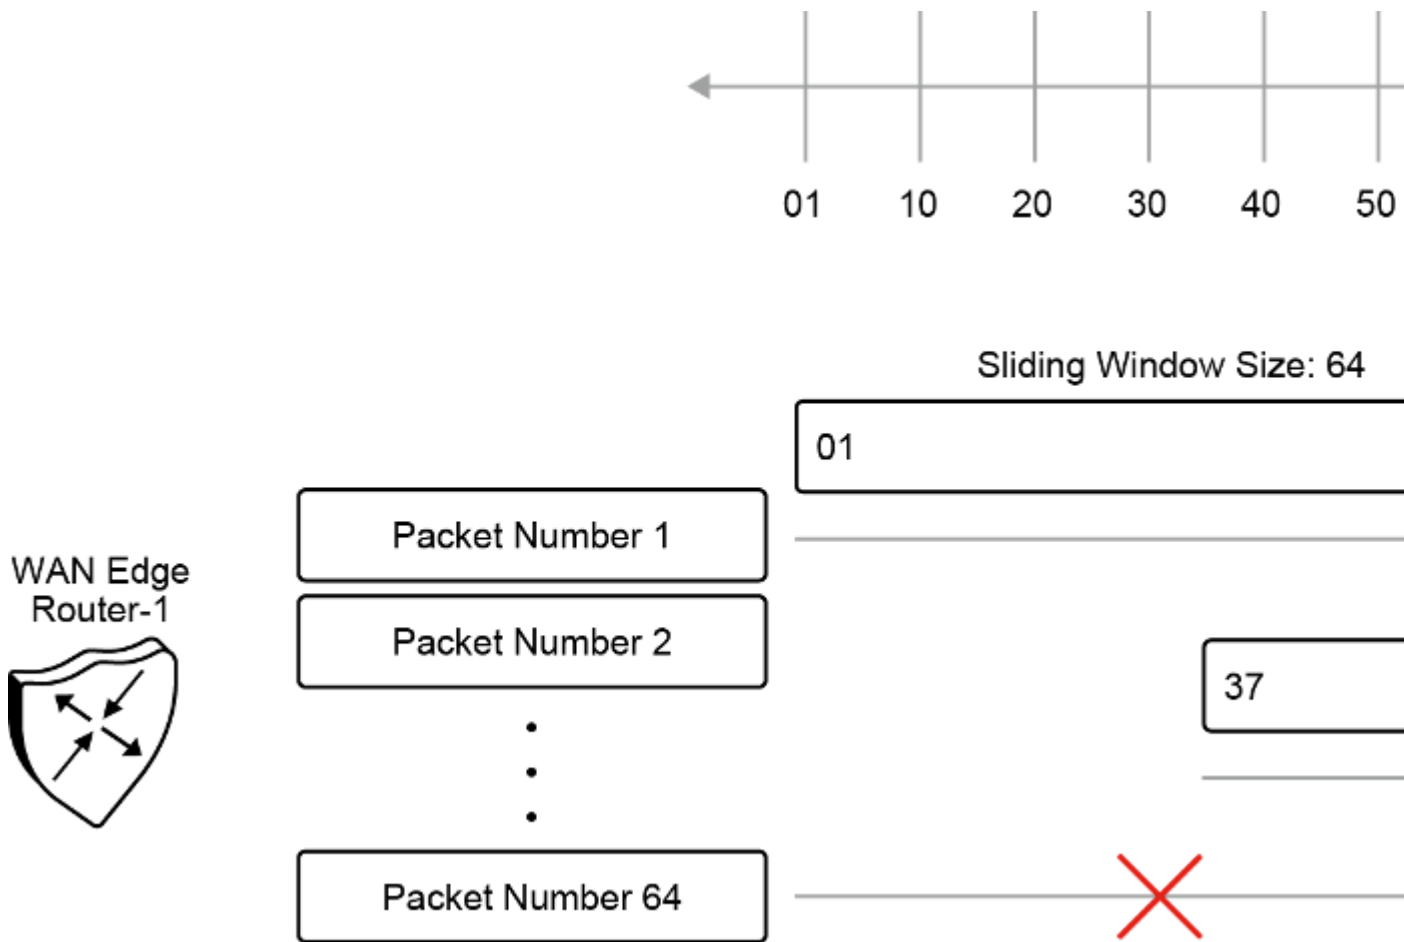
To protect against **replay attacks**, each data packet is assigned a unique sequence number. If a packet's sequence number is a duplicate, the receiving router rejects the packet, preventing unauthorized message replay.



- **Sliding Window Mechanism:** Cisco SD-WAN uses a **sliding window** to define the range of acceptable sequence numbers. If a packet's sequence number is outside this range or a

duplicate, it is discarded. This mechanism protects against **man-in-the-middle (MITM)** attacks, ensuring that only valid, in-order packets are processed.






7. Support for VPNs and MPLS in Data Packets

Cisco SD-WAN supports **MPLS extensions** for data packets transported within **IPsec connections**, ensuring security and network segmentation for enterprise-wide VPNs.

- **MPLS and VPNs:** MPLS information is embedded in the packet header, enabling secure support for multitenancy across branches or campuses. This segmentation is critical for organizations requiring separate traffic paths for different applications or business units.



Summary: Key Advantages of Cisco SD-WAN Over Traditional IKE

Cisco SD-WAN improves significantly over traditional **IKE-based IPsec** environments in several ways:

- **Scalability:** SD-WAN reduces the number of required key exchanges from n^2 in IKE to $n + 1$, making it highly scalable for large networks.
- **Frequent Key Regeneration:** AES keys are refreshed every 24 hours, minimizing security risks.
- **Efficient Key Management:** Keys are managed and exchanged via the **vSmart controller**, ensuring streamlined and secure key distribution.

- **BFD-Based Monitoring:** **BFD packets** provide real-time monitoring of the liveness of IPsec tunnels, ensuring continuous secure communication.

By integrating these mechanisms, Cisco SD-WAN ensures scalable, secure, and reliable data transmission across both public and private networks, offering a superior solution to traditional IPsec approaches.

Bonus Content:

In Cisco Catalyst SDWAN, **vSmart Key Distribution** and **Pairwise Key Generation** are two distinct mechanisms used for managing encryption keys that secure data transmission across the SDWAN fabric. Each method has its unique characteristics, use cases, and roles in ensuring the security of IPsec tunnels between SDWAN edge routers. Below is a detailed explanation of both mechanisms and the scenarios in which they are applied.

vSmart Key Distribution

vSmart Key Distribution is a centralized approach to key management within the SDWAN environment, where the **vSmart controller** is responsible for generating and distributing encryption keys to the SDWAN edge routers (cEdge or vEdge devices). This method streamlines the management of encryption keys, especially in large networks, by using a centralized controller to securely distribute keys to each edge device, eliminating the need for direct peer-to-peer exchanges of encryption keys between routers.

How vSmart Key Distribution Works:

1. **Key Generation by vSmart:** The vSmart controller generates the IPsec encryption keys and securely distributes them to all WAN Edge devices.
2. **Key Distribution with OMP:** The vSmart controller uses **OMP (Overlay Management Protocol)** to deliver the keys, along with **TLOC (Transport Locator)** information, to the edge devices. TLOC information includes the system IP, transport type (color), and encapsulation type, which are used to establish secure communication between routers.
3. **Shared Key Model:** Each edge router receives the same encryption key from the vSmart controller. This method is straightforward but provides less granularity compared to Pairwise Key Generation. All routers share the same key for encrypting and decrypting traffic between them.

When vSmart Key Distribution is Used:

- **Simple Key Management:** vSmart Key Distribution is often used in deployments where simplicity and scalability are key factors. The centralized management model reduces the complexity of distributing and managing keys across a large number of routers.
- **Small-to-Medium Scale Networks:** This method is well-suited for networks where security requirements are not extremely granular, and the network can operate with a shared encryption key for secure communication between all devices.

Pairwise Key Generation

Pairwise Key Generation, on the other hand, is a more advanced and secure key management method that uses **Elliptic Curve Diffie-Hellman (ECDH)** to generate unique encryption keys for each IPsec tunnel between pairs of SDWAN edge routers. This method provides more granular control over encryption, enhancing security by ensuring that each tunnel between two routers has its own unique encryption key, rather than sharing a single key for multiple connections.

How Pairwise Key Generation Works:

1. **Diffie-Hellman Public Values:** The vSmart controller sends **Diffie-Hellman public values** to each WAN Edge device. These public values are used in the key exchange process between routers.
2. **Local Key Generation:** Each WAN Edge device uses the received Diffie-Hellman values, along with its own private keys, to independently generate **pairwise IPsec encryption keys** for each connection it establishes with other edge devices.
3. **Unique Keys for Each Tunnel:** Unlike vSmart Key Distribution, each IPsec tunnel between two routers has its own unique encryption key, which provides greater security. Even if a key for one tunnel is compromised, it does not affect other tunnels.

When Pairwise Key Generation is Used:

- **Higher Security Requirements:** This method is typically used in environments where security is a high priority, and there is a need for **granular control** over encryption keys. It ensures that each pair of routers has a distinct, unique encryption key, making it harder for attackers to compromise multiple connections.
- **Large, Distributed Networks:** Pairwise Key Generation is often used in large, distributed networks where sensitive data may be traversing different transport types (MPLS, Internet, LTE), and stronger security is required for individual tunnels between sites.

Key Differences

1. **Key Management Model:**
 - **vSmart Key Distribution:** A **centralized** key management model where a single encryption key is distributed to all routers by the vSmart controller.
 - **Pairwise Key Generation:** A **decentralized** key management model where each router generates its own unique encryption key for each IPsec tunnel, using Diffie-Hellman public values sent by the vSmart controller.
2. **Security Level:**
 - **vSmart Key Distribution:** Provides a basic level of security with a shared encryption key across the SDWAN fabric. It is simpler to manage but less secure compared to pairwise key generation.
 - **Pairwise Key Generation:** Offers higher security by generating unique encryption keys for each tunnel, preventing the compromise of one key from affecting other tunnels.
3. **Use Cases:**
 - **vSmart Key Distribution:** Suitable for **small-to-medium networks** or environments where simplicity and ease of management are prioritized over granular security.

- **Pairwise Key Generation:** Ideal for **large, complex networks** where **granular security** is critical and where different sites or tunnels require individual encryption keys for enhanced protection.

TL/DR: When to Use Each Method

- **vSmart Key Distribution** is the best choice when you need **simplified key management** in smaller networks where high granularity in security isn't critical. It is suitable for scenarios where centralized control and ease of key distribution are prioritized.
- **Pairwise Key Generation** should be used when **high-level security** and **granularity** are needed. This method is especially effective in larger, distributed networks where each connection requires its own distinct encryption key, providing a more secure SDWAN environment.

Tunnel Groups

Cisco Catalyst SDWAN offers a robust and scalable solution for securely connecting distributed networks over various transport types, such as MPLS, broadband, and LTE. One of the key mechanisms to ensure the efficient and secure operation of the SDWAN fabric is the concept of **Tunnel Group Numbers**, which control how **IPsec tunnels** form between SDWAN edge devices. This document explains how Tunnel Group Numbers are leveraged within the SDWAN fabric to manage, organize, and optimize IPsec tunnels, enhancing traffic flow and ensuring secure communication.

Introduction to Tunnel Group Numbers in Catalyst SDWAN

In a Catalyst SDWAN environment, IPsec tunnels serve as the foundation for interconnecting edge devices securely over different transport networks. These tunnels form the secure data plane, ensuring that traffic between branch sites, data centers, and the cloud is encrypted and protected. The use of **Tunnel Group Numbers** helps to categorize and organize these tunnels, enabling fine-grained control over traffic forwarding, path selection, quality of service (QoS), and redundancy.

Tunnel Group Numbers offer a logical grouping mechanism, allowing SDWAN policies to prioritize traffic, ensure redundancy, and apply specific security or performance settings to different tunnels based on their classification.

Role of Tunnel Group Numbers in IPsec Tunnel Formation

Transport Locators (TLOCs) and IPsec Tunnels

Before diving into Tunnel Group Numbers, it is essential to understand the concept of **Transport Locators (TLOCs)** in Catalyst SDWAN. TLOCs represent the specific transport connections used by edge devices to communicate with each other. A TLOC consists of three main components:

- **System IP address:** Identifies the SDWAN device.
- **Transport Color:** Represents the type of transport network (e.g., MPLS, Internet, or LTE).
- **Encapsulation Type:** The type of encapsulation used (e.g., IPsec or GRE).

Each SDWAN edge device uses TLOCs to establish IPsec tunnels with remote sites, providing secure communication across multiple transport types.

Tunnel Group Numbers and IPsec Tunnel Classification

Tunnel Group Numbers are assigned to categorize IPsec tunnels based on their transport characteristics, security policies, and traffic requirements. The key purpose of Tunnel Group Numbers is to:

- Organize IPsec tunnels into logical groups based on their transport characteristics (e.g., MPLS, Internet).
- Enable SDWAN edge devices to make more intelligent path selection decisions based on tunnel characteristics.
- Apply different QoS or SLA policies based on the group classification.

For example, tunnels formed over high-priority MPLS circuits could be assigned to Tunnel Group 0, while tunnels over lower-priority broadband links might belong to Tunnel Group 1. This separation allows the SDWAN fabric to differentiate between the performance and security requirements of different types of traffic.

Tunnel Group Numbers and Control Policies

Control policies configured within the **vSmart** controller use Tunnel Group Numbers to determine how tunnels are established and how traffic flows between sites. These policies can dictate which tunnels should be prioritized based on the group's performance and security characteristics.

Examples of control policies influenced by Tunnel Group Numbers include:

- **Path Selection Policies:** These policies ensure that high-priority traffic is routed through tunnels in specific groups (e.g., over MPLS), while less critical traffic may use lower-priority groups (e.g., broadband Internet).
- **Failover Policies:** If a tunnel in one group fails, the control policy can automatically redirect traffic to tunnels in another group, ensuring uninterrupted communication.
- **Load Balancing:** Based on tunnel group classification, the SDWAN fabric can load balance traffic across multiple tunnels in different groups, optimizing resource usage and preventing congestion.

Tunnel Group Numbers in Data Plane Operations

Data Traffic Forwarding and Path Selection

Tunnel Group Numbers play a crucial role in how the SDWAN data plane forwards traffic across the fabric. Based on the assigned group, traffic can be forwarded over different paths based on performance metrics such as latency, jitter, and packet loss. For example:

- Traffic classified as mission-critical (e.g., voice, video conferencing) could be routed through tunnels in Tunnel Group 0, which have lower latency and jitter requirements.
- Less critical traffic (e.g., general web traffic or email) may be forwarded through tunnels in Tunnel Group 1, which use cost-effective Internet or broadband links with less stringent performance metrics.

This path selection is managed through a combination of control policies and the performance metrics collected from each tunnel in the SDWAN fabric.

QoS and SLA Enforcement

Tunnel Group Numbers enable the enforcement of **QoS** (Quality of Service) and **SLA** (Service Level Agreement) policies in the data plane. By associating IPsec tunnels with specific group numbers, administrators can enforce traffic policies that match the business requirements for different applications. For instance:

- Tunnels in a higher-priority group might be allocated more bandwidth, lower latency, and higher QoS parameters to support real-time applications like VoIP or video.
- Tunnels in a lower-priority group may have less stringent SLA requirements, allowing for more flexible traffic handling for non-critical applications.

Redundancy and Failover with Tunnel Group Numbers

Tunnel Group Numbers also facilitate redundancy and failover strategies. In the event that a tunnel in one group experiences issues (e.g., link failure, increased latency), the SDWAN fabric can automatically redirect traffic to another available tunnel in a different group. This process ensures that the traffic is still securely delivered, even if the primary path becomes unavailable.

For example, if a primary IPsec tunnel in Tunnel Group 0 (MPLS) fails, the system can switch to an IPsec tunnel in Tunnel Group 1 (Internet) as a backup. This failover occurs seamlessly, ensuring minimal impact on end-user experience.

Security Implications of Tunnel Group Numbers

Tunnel Group Numbers can also be used to define different security policies for IPsec tunnels. For instance, different encryption algorithms, key management techniques, or authentication protocols can be applied based on the tunnel's group classification. This allows network administrators to tailor the level of security for different tunnels depending on the sensitivity of the traffic they are carrying.

For example:

- Tunnels in Tunnel Group 0 (carrying sensitive financial data) might use stronger encryption (e.g., AES-256) and more stringent authentication mechanisms.
- Tunnels in Tunnel Group 1 (carrying general Internet traffic) might use a more flexible security configuration, with a lower overhead encryption algorithm like AES-128.

Configuring and Managing Tunnel Group Numbers in Catalyst SDWAN

Configuring Tunnel Group Numbers in vManage

Tunnel Group Numbers are typically configured and managed through the Cisco **vManage** interface, which provides a centralized point for configuring SDWAN policies. When configuring control and data policies, administrators can assign Tunnel Group Numbers to different transport types and apply policies based on their classification.

Key steps for configuring Tunnel Group Numbers include:

- Defining tunnel groups and associating them with different TLOCs (based on transport type).
- Applying control policies in **vSmart** to enforce routing and forwarding decisions based on the tunnel group classification.
- Monitoring tunnel performance and ensuring that SLA policies are being enforced for each group.

TL/DR

Tunnel Group Numbers are a critical feature in Cisco Catalyst SDWAN that offer a flexible and powerful way to organize and manage IPsec tunnels across different transport networks. By grouping tunnels based on transport type, security level, and performance requirements, administrators can create more efficient, secure, and resilient SDWAN fabrics. Through the use of Tunnel Group Numbers, SDWAN edge devices can make intelligent path selection decisions, prioritize traffic, and ensure redundancy and failover across the SDWAN fabric.

The control and data plane policies in the SDWAN fabric are heavily influenced by Tunnel Group Numbers, making them an essential element of a well-architected SDWAN deployment.

SDWAN Fabric and NAT

Working with NAT

This topic describes the deployment challenges in a Network Address Translation (NAT) environment. The topic explains different Session Traversal Utilities for NAT/Simple Traversal of User Datagram Protocol (STUN) methods, NAT traversal, and its variations.

The NAT environment includes these features:


- NAT allows hosts with private IP addresses in a LAN to communicate with devices in public address spaces, such as the internet.
- Network devices configured with NAT can function as hardware firewalls to prevent unwanted data traffic from passing through a WAN Edge (and to LAN networks in the service side networks connected to the WAN Edge device).

- To enhance the security at branch sites, you can place the WAN Edge router behind a NAT network device or firewall.
- The WAN Edge router can interact with NAT devices configured with the following Session Traversal Utilities for NAT (STUN) methods.
 - Full-cone NAT
 - Address-Restricted Cone NAT
 - Port-Restricted Cone NAT
 - Symmetric-NAT

NAT functions that are used at branch sites must be carefully considered in your Cisco SD-WAN design, because they can affect whether sites can form connections and communicate directly with each other. All NAT functions can create mappings for source IP address, source port, destination IP, and destination port in an IP network packet.

Recall that NAT allows the use of private IP addresses in a LAN to communicate with hosts in the public address spaces, such as the internet. NAT network devices also function as hardware firewalls to prevent unwanted data traffic from reaching a WAN Edge router (and then reaching the service-side LAN networks that are connected to the WAN Edge router).

To enhance the security at branch sites, you can place the WAN Edge router behind a NAT network device. These four STUN methods are described in more detail in the figures that follow.



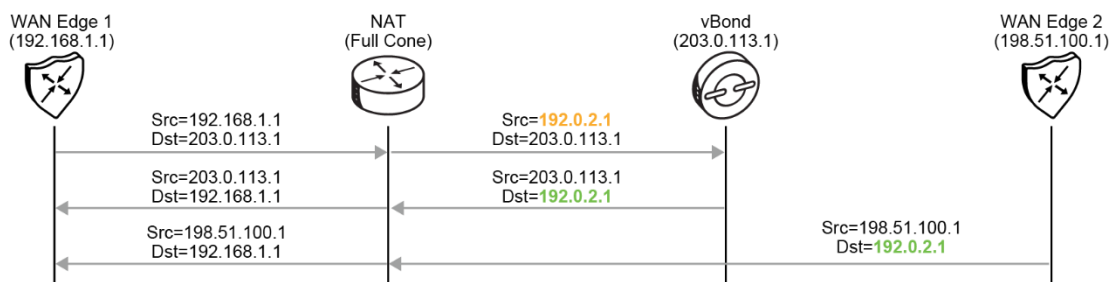
By default, all Cisco SD-WAN devices use base port 12346 for establishing the connections that handle control and data traffic in the overlay network. Each device uses this port when establishing connections with other Cisco SD-WAN devices.

Recall that when multiple Cisco SD-WAN devices are installed behind a single NAT device, you can configure different port numbers for each device. That way, the NAT can properly identify each individual device. You do this by configuring a port offset from the base port 12346. The default port offset is 0.

Also, in the context of Cisco SD-WAN overlay network, port hopping is the process by which devices try various ports when attempting to establish connections with each other if a connection attempt on the first port fails. After such a failure, the port value is incremented, and the connection attempt is retried. The software rotates through five base ports, waiting longer between each connection attempt. The default base port is 12346. Port hopping is done sequentially among ports 12346, 12366, 12386, 12406, and 12426, and then returns to port 12346.

Full-Cone NAT

With full cone NAT (1:1 NAT), you map an internal address and port pair with an external address and port. Any external host can send packets to LAN devices behind the WAN Edge router by addressing them to the external address and port.



Full-cone NAT is also known as one-to-one NAT. You can use full-cone NAT to map an internal address and port pair with an external address and port.


Any external host can send packets to LAN devices behind the WAN Edge router by addressing them to the external address and port.

WAN Edge routers always reach out to the vBond controller first to learn about the rest of the components in the fabric. During this process, they also learn whether they are behind a NAT device. When the WAN Edge initially connects to the vBond, it inserts its real IP address into the exchange. When this packet passes through the NAT device, the source IP and possibly the source port are translated. Because the message still contains the WAN Edge real IP and port, the vBond can send a message back to the WAN Edge. The message notifies the WAN Edge that it is behind a NAT (because the real IP differs from the NAT-translated IP that was received in the exchange). The WAN Edge then inserts this information into its OMP TLOC route and sends it to the vSmart controller. If these values are different, the WAN Edge is behind a NAT device. This information is then reflected to all WAN Edges in the overlay, and the routers use this information to build its data plane. The way to achieve this

NAT detection is by using STUN (RFC 5389). In the example, WAN Edge 2 has received an OMP TLOC route from the vSmart route to reach WAN Edge 1 through its public address (192.0.2.1).

Address-Restricted Cone NAT

Any external host can send a packet to an internal IP address if the internal device initiates a connection to the external host by using the previously created address mappings.




The address-restricted cone NAT method also maps an internal address and port to an external address and port. However, an external host can send packets to the internal device only if the external address (and any port at that address) has first received a packet from the internal device address and port.

In other words, an external host can send packets to the internal address and port by sending packets to the

external address and port, but only if the internal address and port have previously sent a packet to the external host.

Port-Restricted Cone NAT

Similar to address restricted cone NAT, with ports added to the mapping.



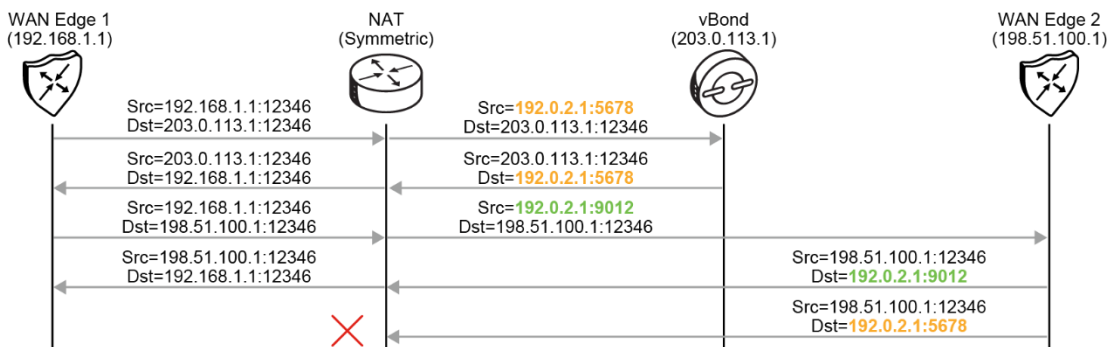
The port-restricted cone NAT method is a stricter version of the address-restricted cone NAT. In the address-restricted cone NAT, an external host can send packets to the internal address and port only if the external address and port pair have received a packet from that internal address and port. The external device must send packets from the specific port to the specific internal port.

This method is similar to an address-restricted cone NAT, but the restriction includes port numbers. After an internal address and port (socket) are mapped to an external address and port, any packets from the internal socket are sent through an external socket.

An external host can send packets to the internal socket by sending packets to an external socket only if the internal socket has previously sent a packet to an external host and port.

Symmetric NAT

Request from the same internal socket to a specific destination IP address and port is mapped to a unique external source socket. Only an external host that receives a packet from an internal host can send a packet back.



With the symmetric NAT method, each request from the same internal socket to a specific destination socket is

mapped to a unique external source socket. If the same internal host sends a packet with the same source socket but to a different destination, the NAT device creates a different mapping.

Only an external host that receives a packet from an internal host can send a packet back.

WAN Edge routers support symmetric NAT only on one side of the WAN tunnel. That is, when a WAN Edge router operates behind a NAT device that is running symmetric NAT, only one NAT device at either end of the tunnel can use symmetric NAT.

The WAN Edge router that is behind a symmetric NAT cannot establish a BFD tunnel with a remote WAN Edge router that is behind symmetric NAT, address-restricted NAT, or port-restricted NAT.

To allow a WAN Edge router to function behind a symmetric NAT, you must configure the vManage and vSmart control connections to use TLS. DTLS control connections do not work through a symmetric NAT.

NAT Traversal Combinations

The following table shows all possible combinations of NAT for two sites of Cisco SD-WAN network.

| cEdge-1 | cEdge-2 | IPsec tunnel can form | GRE tunnel can form |
|--------------------------|------------------------|-----------------------|---------------------|
| No-NAT (Public IP) * | No-NAT (Public IP) | YES | YES |
| No-NAT (Public IP) | Symmetric | YES | NO |
| Full Cone (One-to-one) * | Full Cone (One-to-one) | YES | YES |
| Full Cone (One-to-one) | Restricted-Cone | YES | NO |
| Full Cone (One-to-one) * | Symmetric | YES | NO |
| Restricted-Cone | Restricted-Cone | YES | NO |
| Symmetric | Restricted-Cone | NO | NO |
| Symmetric * | Symmetric | NO | NO |

(*) indicates the most commonly encountered NAT Combinations

A yes means that you can establish a direct connection, whereas no signifies that you cannot establish a direct IPsec tunnel.

As a general rule when there is no NAT or if there is full-cone NAT, there are no restrictions. Also, when both sides require first an outbound connection and restrict inbound to the correct IP and port, no direct connection can be established.

When you cannot establish a direct IPsec tunnel, you can still establish a tunnel to the hub site, which typically has no NAT or uses static NAT relying on a Hub and Spoke Architecture using various means.

In real deployments where hub sites have no NAT and spoke sites might have PAT, a hub and spoke can always be established.

There are several types of NAT that are supported for use with WAN Edge routers. For full-mesh connectivity, at least one side of the WAN Edge tunnel must initiate a connection inbound to a second WAN Edge, regardless of the presence of a firewall in the path. It is recommended that you configure full-cone or 1-to-1 NAT at the data center or hub site, which ensures that WAN Edge routers can always successfully establish an IPsec tunnel with the hub site. Two sites with firewalls running symmetric NAT will have issues forming a tunnel connection.

Symmetric NAT translates the source port of each side to a random port number, and traffic cannot be initiated from the outside. Symmetric NAT requires full-cone NAT or a public IP with no NAT at the destination to establish a direct IPsec tunnel. Sites that cannot connect directly should be set up to reach each other through the data center or other centralized site.

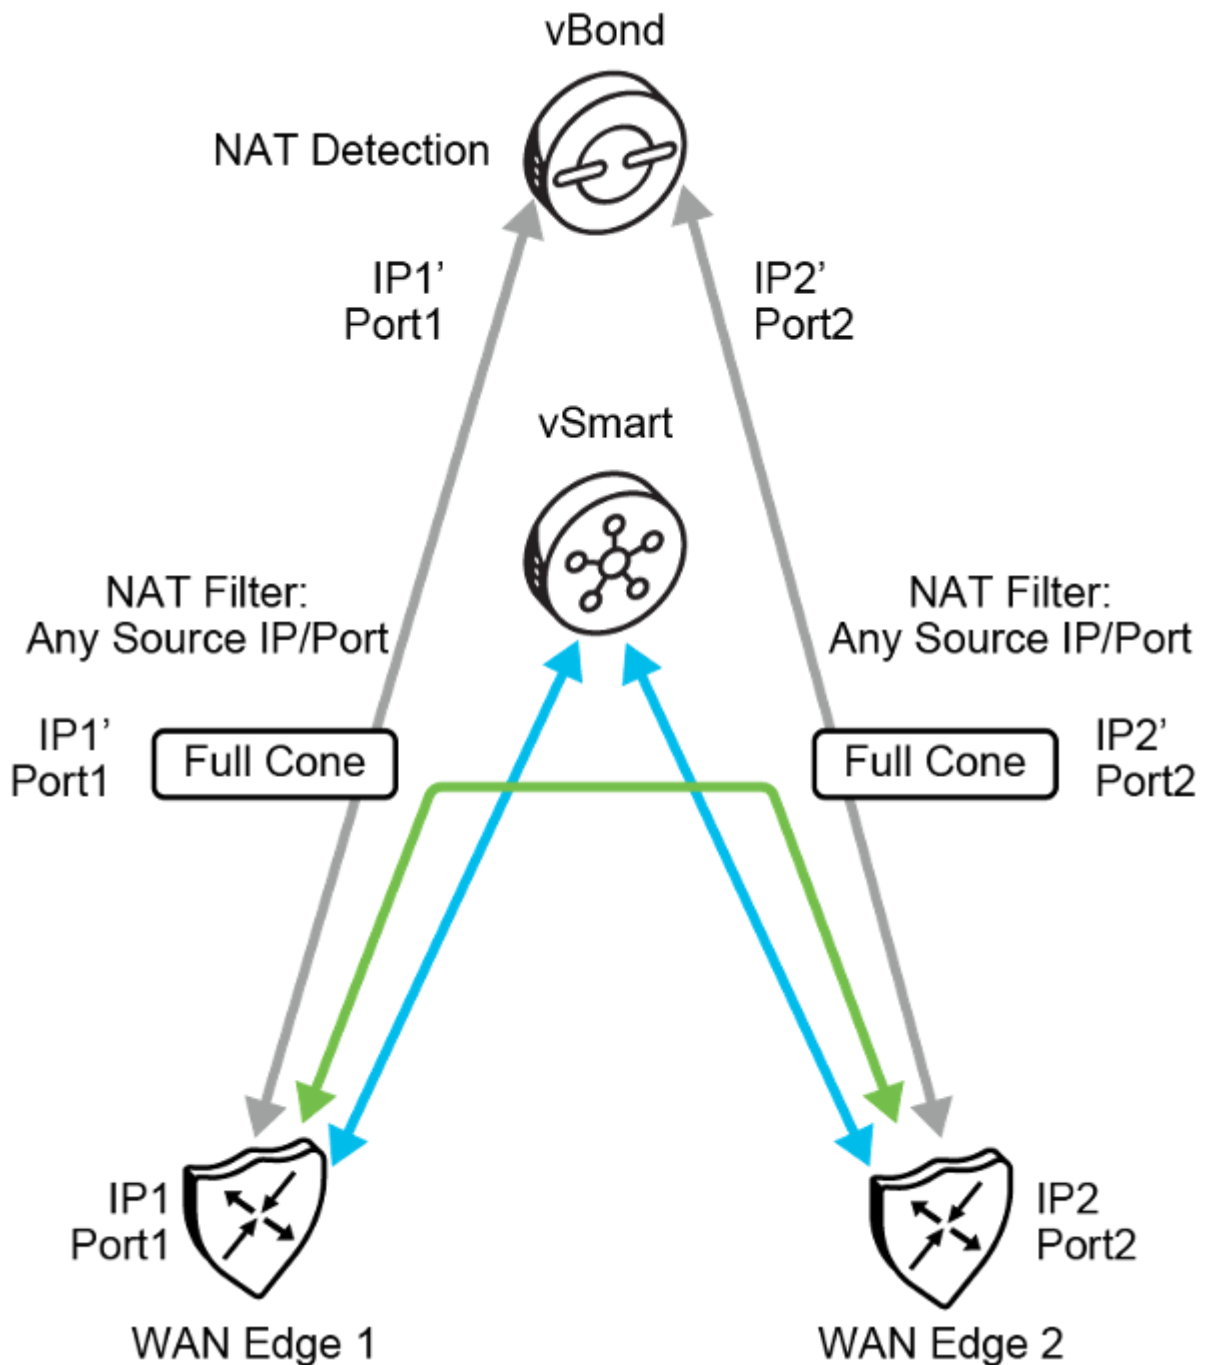
NAT Traversal-Dual-Sided Full Cone

NAT Traversal-Dual-Sided Full Cone workflow:

- vBond discovers post-NAT public IP and communicates back to WAN Edge routers.
 - STUN server

WAN Edge routers notify the vSmart controllers of their post-NAT public IP address. NAT network devices enforce no filter.

Full-cone NAT



WAN Edge router establishes initial connection with the vBond orchestrator at bootup.

For NAT detection to work in Cisco SD-WAN, WAN Edge routers always reach out first to the vBond orchestrator at startup. The vBond orchestrator automatically detects whether there is a NAT function in the path. The vBond orchestrator acts as a STUN server. The vBond orchestrator receives a packet in which the WAN Edge router puts its own address and source port into the payload. vBond then compares the external IP and port to the IP and port inside the payload to determine a NAT function.

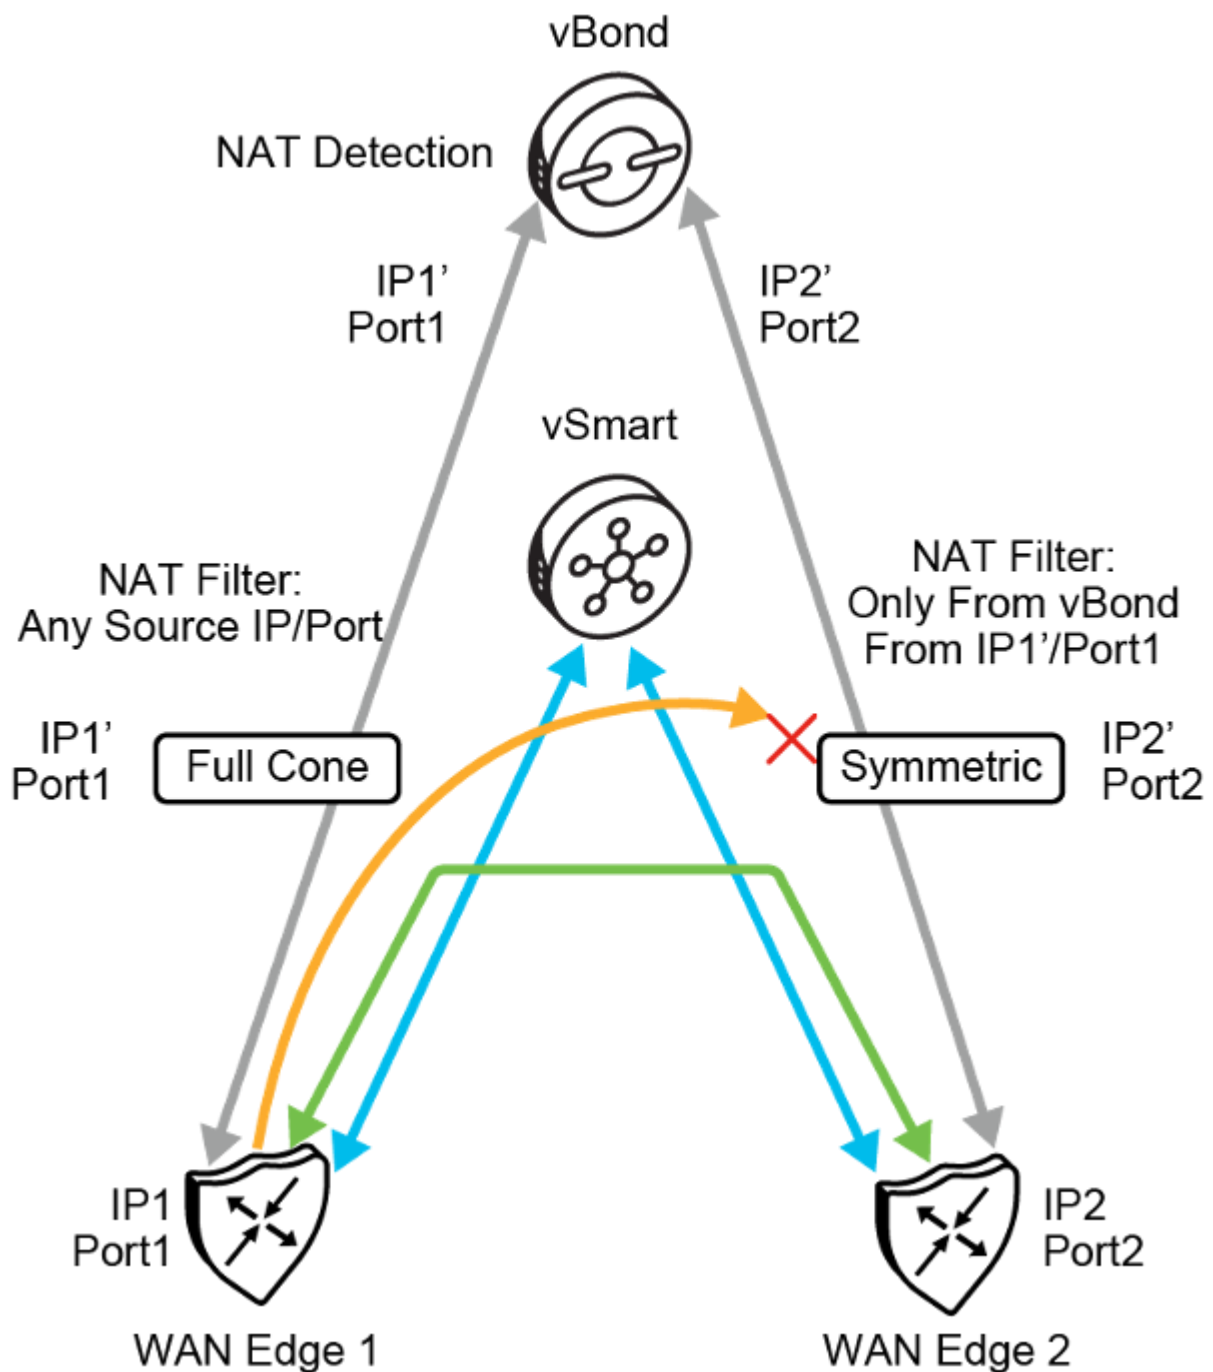
The vBond orchestrator informs the WAN Edge router about how NAT is deployed, and the WAN Edge router communicates the translated address to the vSmart controller. The vSmart controller announces the public addresses to other WAN Edge routers, so that they can communicate.

In the example, WAN Edge 1 uses IP1, which gets translated to IP1'. The vBond orchestrator detects the translated address and informs the vSmart controller, and vSmart sends IP1' to the peer WAN Edge router. That way, both WAN Edge routers know which address to use, and NAT can be traversed without a problem.

NAT Traversal-Full Cone and Symmetric

NAT Traversal-Full Cone workflow:

- vBond discovers post-NAT public IP and communicates back to WAN Edge.
 - STUN server
- WAN Edge routers notify vSmart of their post-NAT public IP address.
 - Symmetric NAT devices enforce filtering:
 - Allows traffic only from vBond.



Bidirectional Forwarding Detection (BFD)

In Cisco Catalyst SDWAN, **Bidirectional Forwarding Detection (BFD)** is a critical protocol used for **rapid link failure detection**. BFD operates at the data plane layer and provides a fast mechanism for detecting faults in the forwarding path between two SDWAN edge routers. In the context of SDWAN, where multiple transport links such as MPLS, broadband, and LTE are often used concurrently, BFD ensures that the SDWAN fabric can

maintain reliable and efficient connectivity by quickly identifying issues in the underlying transport paths.

This document provides an in-depth explanation of how BFD is used in Cisco Catalyst SDWAN, its operational mechanisms, and its impact on path selection, failover, and overall network reliability.

Purpose and Importance of BFD in SDWAN

The main objective of BFD in an SDWAN environment is to **quickly detect failures in the path between SDWAN routers** (vEdge or cEdge) across various transport networks. Since traditional routing protocols like OSPF or BGP rely on timers that may take several seconds to detect a failure, BFD improves failure detection times by reducing the time interval to milliseconds. This rapid detection allows for quick failover to alternate paths, ensuring minimal disruption to network services.

In a typical SDWAN deployment, edge routers use multiple transport types (e.g., MPLS, Internet, LTE) to connect to remote sites, data centers, or cloud services. BFD constantly monitors these links and provides real-time information about their status to the SDWAN controllers (vSmart). This feedback is essential for making decisions regarding path selection, traffic steering, and failover.

How BFD Works in Cisco Catalyst SDWAN

BFD operates by establishing a session between two SDWAN edge routers that are connected via IPsec tunnels across one or more transport networks. It continuously exchanges **BFD control packets** between these routers over each tunnel, providing a mechanism for monitoring the health of the transport links. If BFD detects a failure in the forwarding path, it notifies the SDWAN control plane (vSmart), which then takes appropriate actions to reroute traffic over an alternative path.

BFD Session Establishment

In Catalyst SDWAN, BFD sessions are established between SDWAN edge routers (cEdge or vEdge devices) over each available transport link. For example, if two edge routers are connected via three transport links (MPLS, broadband, and LTE), BFD sessions will be established over each of these tunnels.

The BFD session is established as follows:

- BFD sends periodic **control packets** (known as BFD Hello packets) between the two SDWAN routers over the IPsec tunnels.
- The routers expect to receive a response from the remote router within a specified interval (known as the **detection interval**).
- If the BFD session is established, it indicates that the forwarding path between the routers is operational and healthy.

BFD Failure Detection

The key benefit of BFD is its ability to rapidly detect failures in the forwarding path. This detection occurs as follows:

- If one router does not receive a response to its BFD Hello packets from the remote router within the detection interval, it considers the path to be down.
- BFD then informs the SDWAN fabric that the path is no longer available.
- The SDWAN controllers (vSmart) are notified of the failure and can trigger failover mechanisms to redirect traffic over alternate tunnels.

BFD operates at a much faster rate than traditional routing protocols. For example, typical failure detection times for BFD are in the range of **50 milliseconds to 300 milliseconds**, depending on the configuration. This speed is significantly faster than routing protocols that might take several seconds to detect a failure.

BFD Timers

Two key timers are used in BFD sessions:

- **Hello Interval:** This is the time between the periodic BFD Hello packets sent from one router to another. The default interval in Cisco SDWAN is often set to **1000 milliseconds** (1 second), but it can be reduced for faster detection.
- **Detection Interval:** This is the time within which a response (BFD Hello packet) must be received from the remote router. If no response is received within the detection interval, the path is considered down.

The detection interval is typically configured to be a multiple of the Hello Interval, allowing BFD to quickly detect failures.

BFD in Action: Path Monitoring and Failover in SDWAN

BFD plays a vital role in Catalyst SDWAN's ability to monitor transport links, detect failures, and enable **rapid failover**. Below is an overview of how BFD is used to optimize the network's performance and ensure reliability:

Real-Time Path Monitoring

BFD continuously monitors the health and performance of each transport link by exchanging control packets. If any degradation in the performance of a link is detected (e.g., increased packet loss, delay, or jitter), BFD can inform the SDWAN control plane. This information is essential for making informed decisions about path selection based on the **application-aware routing** policies configured in the SDWAN fabric.

For instance, if a link begins to experience high levels of jitter or latency that could impact real-time traffic such as voice or video, the vSmart controller can use this BFD data to reroute traffic over a more stable path, ensuring that application SLAs (Service Level Agreements) are met.

Rapid Link Failure Detection and Failover

In the event of a transport link failure (e.g., a link goes down or becomes severely degraded), BFD detects the failure within milliseconds and notifies the SDWAN control plane. The vSmart controller can then trigger an immediate failover to an alternative link.

For example, if a branch router has active BFD sessions over both an MPLS link and an Internet link, and the MPLS link fails, BFD will quickly detect the failure. The SDWAN fabric will then automatically route traffic over the Internet link, ensuring that connectivity is maintained without noticeable downtime to users or applications.

Load Balancing and Path Optimization

In addition to failover, BFD can also help the SDWAN fabric optimize traffic flows by providing real-time information on link performance. The SDWAN control plane can use this information to intelligently balance traffic across multiple links, ensuring that network resources are used efficiently.

For example, if BFD detects that the primary link (e.g., MPLS) is experiencing congestion or higher delay, the SDWAN fabric can dynamically shift some of the traffic to secondary links (e.g., Internet or LTE) to prevent congestion and improve performance.

Tunnel Health Monitoring

BFD is not only used to monitor physical links but also the **tunnels** that are established between SDWAN routers. These tunnels, which can be IPsec-encrypted for security, are the transport mechanisms used to carry traffic between remote SDWAN sites. BFD continuously monitors the health of these tunnels to ensure that they are operational.

If a tunnel becomes unavailable (even if the physical link is still up), BFD will detect this and initiate the failover process to a backup tunnel.

Configuration of BFD in Cisco Catalyst SDWAN

BFD is typically configured using the Cisco **vManage** interface, which provides a centralized point for managing SDWAN devices and policies. The configuration process involves defining the BFD parameters, such as the Hello Interval and Detection Interval, to control how quickly failures are detected.

In most cases, BFD is enabled by default on the tunnels between SDWAN routers, but network administrators can fine-tune the timers and behavior to meet specific network requirements.

Key Benefits of BFD in Cisco Catalyst SDWAN

- **Fast Failure Detection:** BFD allows for the detection of path failures in milliseconds, reducing downtime and improving the overall availability of the SDWAN network.

- **Improved Application Performance:** By quickly detecting path degradation (such as increased latency or jitter), BFD ensures that critical applications are rerouted over optimal paths, maintaining performance and meeting SLAs.
- **Seamless Failover:** In the event of a transport link or tunnel failure, BFD enables seamless failover to backup links, ensuring continuous connectivity without manual intervention.
- **Enhanced Resilience:** BFD provides real-time monitoring of both physical links and IPsec tunnels, ensuring that the SDWAN fabric is always aware of the network's health and can react to changes quickly.
- **Better Path Selection:** By feeding real-time path information to the vSmart controller, BFD helps the SDWAN fabric make intelligent path selection decisions, optimizing network performance and resource utilization.

TL/DR

BFD is a vital component in Cisco Catalyst SDWAN, providing rapid failure detection, real-time path monitoring, and seamless failover for both physical transport links and IPsec tunnels. Its ability to detect forwarding path failures within milliseconds allows the SDWAN fabric to maintain high availability and performance across multiple transport networks. With BFD, Catalyst SDWAN ensures that critical applications are always routed over optimal paths, improving network resilience, reducing downtime, and maintaining consistent performance levels.

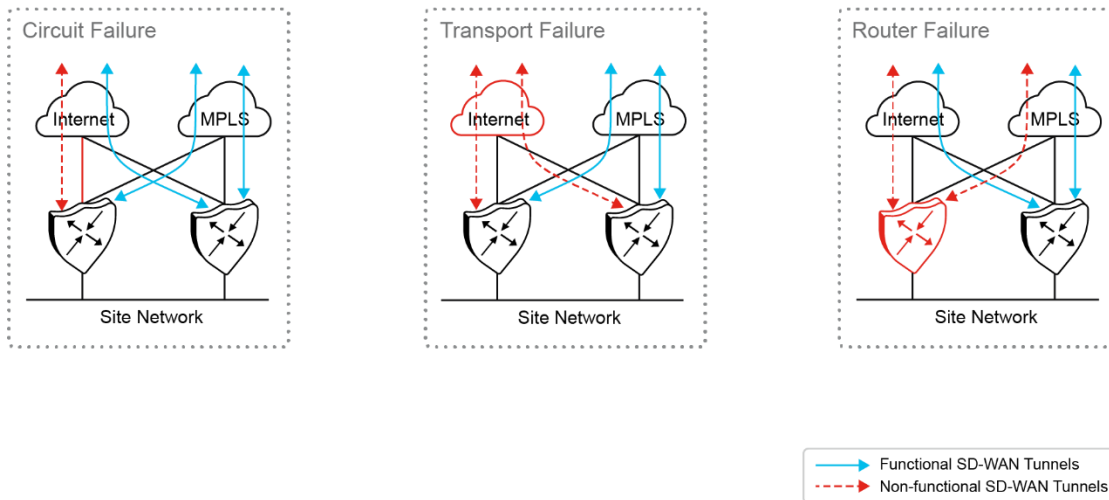
This fast and lightweight protocol significantly enhances the overall reliability and agility of the SDWAN environment, making it a key enabler for modern, distributed network architectures.

TLOC Extension

Implement TLOC Extensions

This topic describes the benefits, types, transport choices, and configuration options for transport locator (TLOC) extensions.


The ideal scenario is that every WAN Edge router is multi-homed and directly connected to every transport provider. Cisco SD-WAN overlay tunnels are built through all directly connected transports.



During a circuit failure, both edge routers maintain WAN connectivity to at least one transport provider. During a transport failure, both edge routers maintain WAN connectivity to the remaining transport provider. In case of a router failure, the remaining edge router maintains connectivity to both transport providers.

At times WAN Edge routers cannot be connected to each transport directly, and only one WAN Edge router can be connected to a single transport. Alternatively, a switch can be connected to each transport, and the Cisco SD-WAN routers can connect to each transport through the connected switches. Such topology is not usually recommended at a branch because it adds cost to the solution and results in having another device to manage.

TLOC extensions allow each WAN Edge router to access the opposite transport through a TLOC-extension interface on the neighboring WAN Edge router.

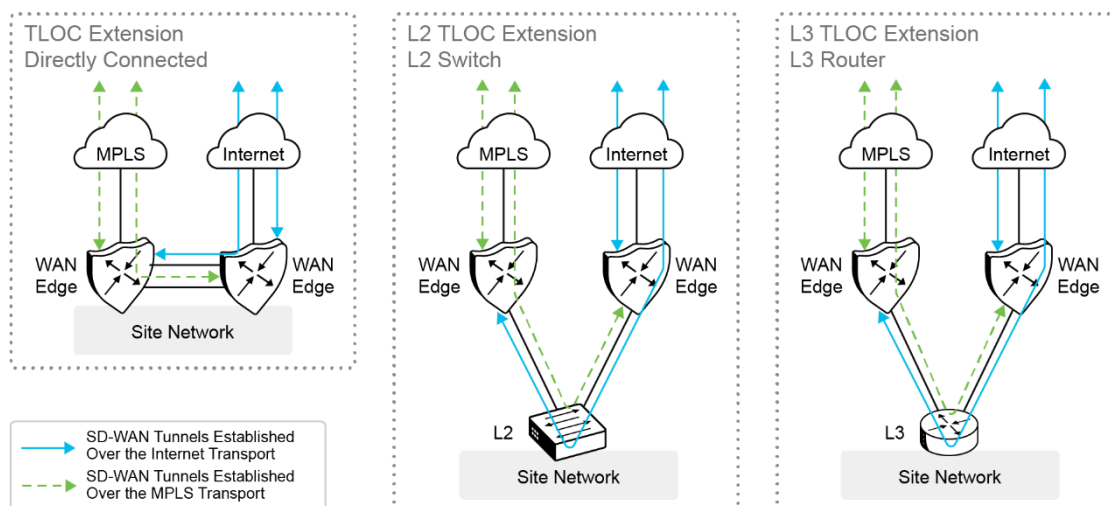


In the figure, WAN Edge 1 connects directly to the MPLS transport and uses the TLOC extension interface on WAN Edge 2 to connect to the internet transport. In turn, WAN

Edge 2 connects directly to the internet transport and uses the TLOC extension interface on WAN Edge 1 to connect to the MPLS transport. The connection from a TLOC extension interface through to transport is transparent. WAN Edge 1 router in the diagram still has two physical interfaces with tunnels configured – one to the MPLS and one to the internet and is unaware that the tunnel to the internet passes through another Cisco SD-WAN router.

TLOC Extension Types

TLOC extensions on Cisco SD-WAN routers can be connected in multiple ways. The Cisco SD-WAN routers can be directly connected, connected through a Layer 2 switch, or connected through a Layer 3 switch/router. Layer 2 TLOC extensions describe TLOC extensions between two routers that are L2-adjacent to each other, and the links are in the same subnet. Layer 3 TLOC extensions describe TLOC extensions between two routers separated by a Layer 3 switch or router where the links are in different subnets. Layer 3 TLOC extensions are implemented using GRE tunnels. Note that TLOC extensions can be separate physical interfaces or sub-interfaces (if bandwidth allows). The figure illustrates different Layer 2 and Layer 3 TLOC extension deployments.

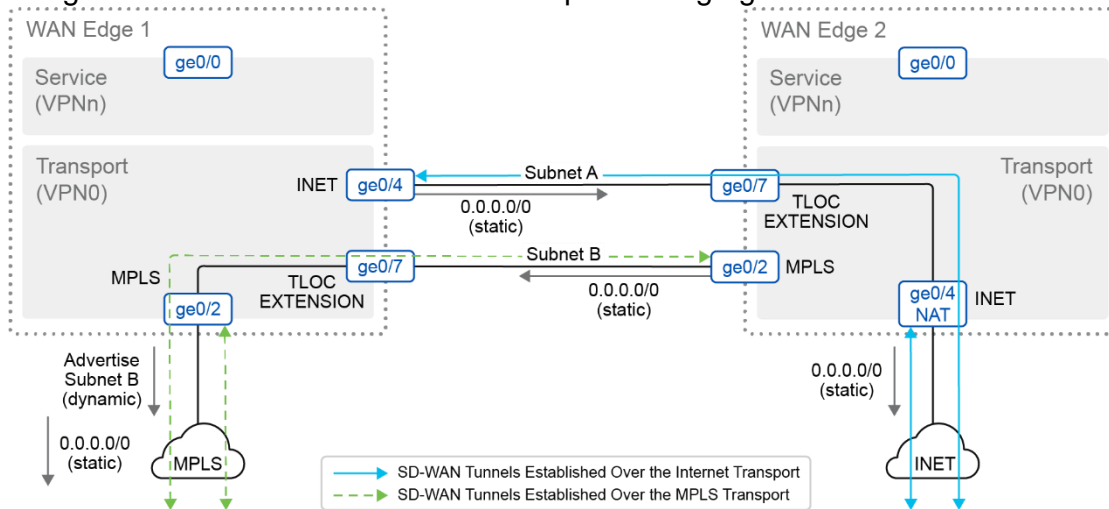


There are some limitations to the use of TLOC extensions:

- TLOC and TLOC extension interfaces are supported only on Layer 3 routed interfaces. Layer 2 switch ports and SVIs cannot be used as WAN or Tunnel interfaces and can only be used on the service side. LTE also cannot be used as a TLOC extension interface between WAN Edge routers.
- TLOC extension does not work on transport interfaces which are bound to loopback tunnel interfaces.

TLOC Extension Routing

When you configure the TLOC extension interface, configure it in VPN 0, assign it an IP address, and then specify the WAN interface to which it is bound. In the figure, WAN Edge 1's TLOC extension interface is ge0/7 and is bound to the MPLS transport through ge0/2. WAN Edge 2's TLOC extension interface is ge0/7 and is bound to the INET transport through ge0/4.



Some routing considerations need to take place for controller reachability to occur and for IPsec tunnels and BFD sessions to come up with other sites over the TLOC extension interfaces. Static default routes should be configured in the underlay (transport VPN 0) on each WAN Edge router, pointing to the service provider router as the next hop.

To reach the INET transport, WAN Edge 1's INET interface (ge0/4) should be configured with a default route pointing to WAN Edge 2's ge0/7 IP address. If subnet A is in a private address space, then NAT should be configured on WAN Edge 2's ge0/4 transport interface to ensure that traffic can be routed back from the internet to WAN Edge 1 over the TLOC Extension.

To reach the MPLS transport, WAN Edge 2's MPLS interface should be configured with a default route pointing to WAN Edge 1's ge0/7 IP address. To ensure that traffic can be routed back to the TLOC extension interface, a routing protocol (typically BGP, or OSPF) can be run in the transport VPN (VPN 0) of WAN Edge 1 to advertise subnet B so that the MPLS provider has a route to subnet B through WAN Edge 1. Typically, a route map is also applied inbound to deny all incoming dynamic routes from the service provider since the static default route is used in the transport VPN for control plane and IPsec tunnel establishment. As an alternative to a routing protocol, the MPLS PE router can implement a static route to subnet B through WAN Edge 1, which can then be redistributed through the service provider network. Static routes are not recommended because the method is not as manageable or scalable as using a dynamic routing protocol when you have many sites.

Transport Choices

There are a many different transport choices and different combinations of transports that can be used. Transports are deployed in an active-active state, and how you use them is extremely flexible.



A very common transport combination is MPLS and internet (Figure A). MPLS can be used for business-critical traffic, while the internet can be used for bulk traffic and other data. When one transport is down, the other transport can be used to route traffic to and from the site. Internet is reliable in most places and able to meet the SLAs of most applications, so often, sites will deploy two internet transports instead. LTE is used frequently as a transport choice and can be deployed in active mode or as a circuit of last resort, which does not become active unless all other transports become unavailable (Figure B).

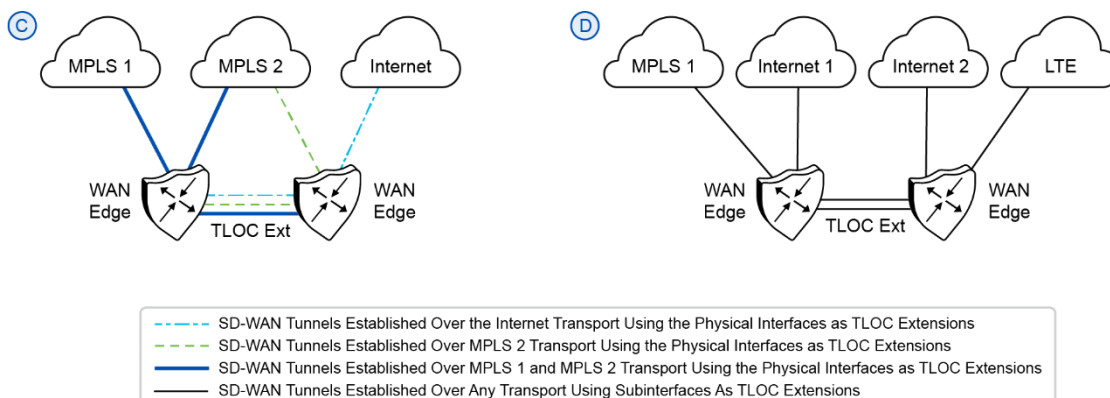


Figure C shows TLOC Extensions on separate physical interfaces while Figure D shows multiple TLOC extensions using subinterfaces across two physical interfaces.