

Site-to-Site VPN

- Site-to-site VPN, on page 1
- Types of site-to-site VPN topologies, on page 4
- License requirements for site-to-site VPN, on page 4
- Prerequisites for site-to-site VPN, on page 5
- Manage site-to-site VPNs, on page 6
- Configure a policy-based site-to-site VPN, on page 7
- Configure virtual tunnel interfaces, on page 21
- Deploy a SASE tunnel on Umbrella, on page 46
- Secure traffic using Cisco Secure Access and Firewall Threat Defense devices, on page 52
- Monitor site-to-site topologies, on page 57

Site-to-site VPN

A site-to-site VPN is a secure, encrypted VPN tunnel between two networks in different geographic locations. After the VPN connection is established, the hosts behind the local gateway can connect to the hosts behind the remote gateway through the secure VPN tunnel. You can create site-to-site IPsec connections between managed devices, and between managed devices and other Cisco or third-party peers.

Site-to-site VPN features

Secure Firewall Threat Defense site-to-site VPN supports these features:

- IPsec IKEv1 & IKEv2 protocols.
- Certificates and automatic or manual preshared keys for authentication.
- IPv4 & IPv6 addresses. All combinations of inside and outside are supported.
- Static and dynamic interfaces.
- HA environments for Firewall Threat Defense devices.
- VPN alerts when the tunnel goes down and tunnel statistics.
- IKEv1 and IKEv2 back-up peer configuration for point-to-point extranet and hub-and-spoke VPNs.
- Extranet device as hub in 'Hub and Spokes' deployments.

- Dynamic IP address for a managed endpoint pairing with extranet device in 'Point to Point' deployments and for extranet device as an endpoint.

VPN topologies

VPN topology configuration requires specifying a unique name, a topology type, IKE version selection, and authentication method. The Cloud-Delivered Firewall Management Center configures site-to-site VPNs on Firewall Threat Defense devices only.

You can select from three types of topologies, containing one or more VPN tunnels:

- Point-to-point (PTP) deployments establish a VPN tunnel between two endpoints.
- Hub and Spoke deployments establish a group of VPN tunnels connecting a hub endpoint to a group of spoke nodes.
- Full Mesh deployments establish a group of VPN tunnels among a set of endpoints.

In the Cloud-Delivered Firewall Management Center, site-to-site VPNs are configured based on IKE policies and IPsec proposals that are assigned to VPN topologies. Policies and proposals are sets of parameters that define the characteristics of a site-to-site VPN, such as the security protocols and algorithms that are used to secure traffic in an IPsec tunnel. Several policy types may be required to define a full configuration image that can be assigned to a VPN topology.

For authentication of VPN connections, configure a preshared key in the topology, or a trustpoint on each device. Preshared keys allow for a secret key, used during the IKE authentication phase, to be shared between two peers. A trustpoint includes the identity of the CA, CA-specific parameters, and an association with a single enrolled identity certificate.

IPsec and IKE

In the Cloud-Delivered Firewall Management Center, site-to-site VPNs are configured based on IKE policies and IPsec proposals that are assigned to VPN topologies. Policies and proposals are sets of parameters that define the characteristics of a site-to-site VPN, such as the security protocols and algorithms that are used to secure traffic in an IPsec tunnel. Several policy types may be required to define a full configuration image that can be assigned to a VPN topology.

Authentication

For authentication of VPN connections, configure a preshared key in the topology, or a trustpoint on each device. Preshared keys allow for a secret key, used during the IKE authentication phase, to be shared between two peers. A trustpoint includes the identity of the CA, CA-specific parameters, and an association with a single enrolled identity certificate.

Extranet Devices

Each topology type can include extranet devices, devices that you don't manage in Cloud-Delivered Firewall Management Center. These include:

- Cisco devices that Cloud-Delivered Firewall Management Center supports, but for which your organization isn't responsible. Such as spokes in networks managed by other organizations within your company, or a connection to a service provider or partner's network.
- Non-Cisco devices. You can't use Cloud-Delivered Firewall Management Center to create and deploy configurations to non-Cisco devices.

Add non-Cisco devices, or Cisco devices not managed by the Cloud-Delivered Firewall Management Center, to a VPN topology as "Extranet" devices. Also specify the IP address of each remote device.

Guidelines for configuring site-to-site VPN in Secure Firewall Threat Defense devices

General guidelines

- Site-to-site VPN supports ECMP zone interfaces.
- Configure all nodes in a topology with either crypto ACL or a protected network. You cannot configure a topology with crypto ACL on one node and protected network on another.
- Configure a VPN connection across domains by using an extranet peer for the endpoint not in the current domain.
- You can backup Firewall Threat Defense VPNs using the Cloud-Delivered Firewall Management Center backup.
- Configure unique local IKE identity for all tunnels across all your VPN topologies.
- Ensure that IKE ports 500 and 4500 are not already in use and that no active PAT translations exist on those ports before configuring a site-to-site VPN. Configuring a site-to-site VPN on ports that are already in use will cause the service to fail to start.
- When configuring a site-to-site VPN between two devices managed by the same Cloud-Delivered Firewall Management Center, do not configure the devices as backup peers. Instead, configure one of the peer devices in the topology as an extranet device.
- If a spoke is part of multiple topologies with the same hub, for the Local Tunnel (IKE) Identity, we recommend that you use Key ID instead of an IP address.

  If the spoke uses an IP address as its local identity, the hub creates only one tunnel group. The spoke is then mapped to a specific virtual template interface, so the connection works with only one topology.

Limitations for configuring site-to-site VPN in Threat Defense devices

General limitations

- IKEv1 does not support CC/UCAPL-compliant devices. IKEv2 is recommended for these devices.
- VPN does not support network objects with a 'range' option.
- Firewall Threat Defense VPNs do not support PDF export.
- Tunnel status is not updated realtime, but at an interval of five minutes in the Cloud-Delivered Firewall Management Center.
- You cannot use the double quote character (") in pre-shared keys, replace it if already in use.

Crypto ACL limitations

- Cloud-Delivered Firewall Management Center supports only point-to-point VPN with crypto ACL and does not support tunnel health events.

- Cloud-Delivered Firewall Management Center does not verify the device interface address verification for transport mode when you select a crypto ACL.
- There is no support for automatic mirror ACE generation. Mirror ACE generation for the peer is a manual process on either side.

Topology management limitations

- There is no per-tunnel or per-device edit option for Firewall Threat Defense VPNs, you can edit only the whole topology.

Types of site-to-site VPN topologies

This table provides an overview of the different site-to-site VPN topology options available for secure network connectivity configuration.

Site-to-Site VPN Topology	Description	More Information
SD-WAN Topology	Configure a secure branch network using the SD-WAN wizard. This wizard simplifies and automates VPN and routing configuration of your network using a hub and spoke topology, enabling SD-WAN capabilities.	SD-WAN wizard for secure branch network deployment
Route-Based VPN	Configure secure traffic dynamically between peers in a network based on routing over Virtual Tunnel Interfaces (VTIs).	Create a route-based site-to-site VPN, on page 30
Policy-Based VPN	Configure secure traffic between peers in a network based on a static policy using protected networks.	Configure a policy-based site-to-site VPN, on page 7
SASE Topology	Configure an IPsec IKEv2 tunnel from a Firewall Threat Defense device to Secure Access or an Umbrella Secure Internet Gateway (SIG). This tunnel forwards all internet-bound traffic to Secure Access or Umbrella SIG for inspection and filtering.	Secure traffic using Cisco Secure Access and Firewall Threat Defense devices, on page 52 Configure a SASE tunnel for Umbrella, on page 50

License requirements for site-to-site VPN

This topic lists the license requirements for setting up various types of site-to-site VPN connections.

License requirements for policy-based and route-based VPN

With the Essentials license, you can set up policy-based and route-based VPNs on your Firewall Threat Defense devices.

Depending on whether export-controlled functionality is enabled in your Smart License account, Cloud-Delivered Firewall Management Center determines whether to allow or block the usage of strong crypto on devices. To verify if export-controlled functionality is enabled for your Smart License account, choose **Administration** > **Licenses** > **Smart Licenses**.

Note If you use an evaluation license, or if you have not enabled the export-controlled functionality, you cannot use strong encryption for your VPN connections.

License requirements for deploying a SASE tunnel on Umbrella

To deploy SASE tunnels on Umbrella from Cloud-Delivered Firewall Management Center, you must enable your Smart License account with the export-controlled functionality. If this functionality is not enabled, you can only create a SASE topology, you cannot deploy tunnels on Umbrella.

Licenses for automatic tunnel between Secure Accessand Firewall Threat Defense devices

- Cloud-Delivered Firewall Management Center Essentials (formerly Base) license must allow export-controlled functionality. Choose **Administration** > **Licenses** > **Smart Licenses** to verify this in Cloud-Delivered Firewall Management Center.

Note Evaluation license does not support this feature.

- Cisco Secure Access Essentials subscription license must be available.

Prerequisites for site-to-site VPN

Supported interfaces

Topology Type	Interface Type

Policy-Based	• Physical interfaces
	  • Non-management
	  • Interface Mode must be either Routed or None
	• Subinterface interfaces
	• Redundant interfaces
	• Etherchannel interfaces
	• VLAN interfaces

Route-Based	Static Virtual Tunnel Interfaces



Note The maximum number of supported security associations (SA) depends on the maximum VPN peers for your platform. For more information about maximum VPN peers, refer the data sheet of the Secure Firewall Threat Defense device.

General prerequisites

- You must have administrator privileges.

Manage site-to-site VPNs

You can view, create, modify, and deploy site-to-site VPN configurations and manage their deployment status.

The **Site-to-Site VPN** page provides a snapshot of site-to-site VPN tunnels. You can view the status of the tunnels and filter the tunnels based on the device, topology, or tunnel type. The page lists 20 topologies per page and you can navigate between pages to view more topology details. You can click individual VPN topologies to expand and view details of the endpoints.

Before you begin

For certificate authentication of your site-to-site VPNs, you must prepare the devices by allocating trustpoints as described in Certificates.

Procedure

Select **Secure Connections** > **Site-to-Site VPN & SD-WAN** to manage your Firewall Threat Defense site-to-site VPN configurations and deployments.

The page lists the site-to-site VPNs topologies and indicates the status of tunnels using color codes:

- Active (Green)–There is an active IPsec Tunnel.

- Unknown (Amber)–No tunnel establishment event has been received from the device yet.
- Down (Red)–There are no active IPsec tunnels.
- Deployment Pending–Topology has not been deployed on the device yet.

Choose one of these actions:

- **Refresh**—View the updated status of the VPNs.
- **Add**—Create new policy based or route-based Site to Site VPNs.
- **Edit**—Modify the settings of an existing VPN topology.

Note You cannot edit the topology type after you initially save it. To change the topology type, delete the topology and create a new one.

Two users shouldn't edit the same topology simultaneously; however, the web interface doesn't prevent simultaneous editing.

- **Delete**—To delete a VPN deployment, click **Delete** ( ).
- Deploy—Choose **Deploy** > **Deploy**; see Deploy configuration changes.

Note Some VPN settings are validated only during deployment. Be sure to verify that your deployment was successful.

Configure a policy-based site-to-site VPN

Configure a policy-based site-to-site VPN to establish secure communication channels between network endpoints using crypto maps and specified network topologies.

Procedure

Step 1 Choose **Secure Connections** > **Site-to-Site VPN & SD-WAN**, click **Add**, and enter the topology name in the **Topology Name** field.

Step 2 Click the **Policy-Based VPN** radio button and select the VPN topology and click **Create**.

Step 3 Choose the IKE versions to use during IKE negotiations. Check the **IKEv1** or **IKEv2** check box.

Default is IKEv2. Select either or both options as appropriate; select IKEv1 if any device in the topology doesn't support IKEv2.

You can also configure a backup peer for point-to-point extranet VPNs. For more information, see Firewall Threat Defense VPN endpoint options, on page 8.

Step 4 Required: Add Endpoints for this VPN deployment by clicking **Add** (+) for each node in the topology.

Configure each endpoint field as described in Firewall Threat Defense VPN endpoint options, on page 8.

- For Point to point, configure **Node A** and **Node B**.
- For Hub and Spoke, configure a **Hub Node** and **Spoke Nodes**
- For Full Mesh, configure multiple **Nodes**

Step 5 (Optional) Specify non-default IKE options for this deployment as described in Firewall Threat Defense VPN IKE options, on page 13

Step 6 (Optional) Specify non-default IPsec options for this deployment as described in Firewall Threat Defense VPN IPsec options, on page 15

Step 7 (Optional) Specify non-default Advanced options for this deployment as described in Firewall Threat Defense advanced site-to-site VPN deployment options, on page 17 and click **Save**.

The policy-based site-to-site VPN is configured with the specified endpoints, IKE settings, and topology options.

What to do next

Deploy configuration changes.

Note Some VPN settings are validated only during deployment. Be sure to verify that your deployment was successful.

If you get an alert that your VPN tunnel is inactive even when the VPN session is up, follow the VPN troubleshooting instructions to verify and ensure that your VPN is active.

Firewall Threat Defense VPN endpoint options

Configure the basic parameters for a point-to-point topology in a policy-based VPN as described in Configure Endpoints for a Point to Point Topology and click the **Endpoints** tab.

Fields

Device

Choose an endpoint node for your deployment:

- A Firewall Threat Defense device managed by this Cloud-Delivered Firewall Management Center
- A Firewall Threat Defense high availability container managed by this Cloud-Delivered Firewall Management Center
- An **Extranet** device, any device (Cisco or third party) not managed by this Cloud-Delivered Firewall Management Center .

Device Name

For extranet devices only, provide a name for this device. We recommend naming it such that it is identifiable as an unmanaged device.

Interface

If you chose a managed device as your endpoint, choose an interface on that managed device.

For 'Point to Point' deployments, you can also configure an endpoint with dynamic interface. An endpoint with a dynamic interface can pair only with an extranet device and can't pair with an endpoint, which has a managed device.

IP Address

- If you choose an extranet device, a device not managed by the Cloud-Delivered Firewall Management Center, specify an IP address for the endpoint.

  For an extranet device, select **Static** and specify an IP address or select **Dynamic** to allow dynamic extranet devices.

- If you chose a managed device as an endpoint, choose a single IPv4 address or multiple IPv6 addresses from the drop-down list. These IP addresses are already assigned to this interface on the managed device.

- All endpoints in a topology must have the same IP addressing scheme. IPv4 tunnels can carry IPv6 traffic and vice versa. The **Protected Networks** define which addressing scheme the tunneled traffic uses.

- If the managed device is a high-availability container, choose from a list of interfaces.

This IP is Private

Check the check box if the endpoint resides behind a firewall with network address translation (NAT).



Note Use this option only when the peer is managed by the same Cloud-Delivered Firewall Management Center and don't use this option if the peer is an extranet device.

Public IP address

If you checked the **This IP is Private** check box, specify a public IP address for the firewall. If the endpoint is a responder, specify this value.

Connection Type

Specify the allowed negotiation as bidirectional, answer-only, or originate-only. Supported combinations for the connection type are:

Table 1: Connection Type Supported Combinations

Remote Node	Central Node
Originate-Only	Answer-Only
Bi-Directional	Answer-Only
Bi-Directional	Bi-Directional

Certificate Map

Choose a preconfigured certificate map object, or click **Add** (+) to add a certificate map object. The certificate map defines what information is necessary in the received client certificate to be valid for VPN connectivity. See Certificate map objects for details.

Protected Networks

Caution Hub and Spoke topology—To avoid traffic drop for a dynamic crypto map, ensure that you don't select the protected network *any* for both the endpoints.

If the protected network is configured as *any*, on both the endpoints, the crypto ACL that works upon the tunnel is not generated.

Defines the networks that are protected by this VPN endpoint. Select the networks by selecting the list of Subnet/IP Address that define the networks that are protected by this endpoint. Click **Add** (+) to select from available Network Objects or add new Network Objects. See Create a network object . Access control lists are generated from the choices made here.

- **Subnet/IP Address (Network)**—VPN endpoints can't have the same IP address and protected networks in a VPN endpoint pair cannot overlap. If protected networks for an endpoint contain IPv4 or IPv6 entries, the other endpoint's protected network must have at least one entry of the same type (IPv4 or IPv6). If it doesn't, the other endpoint's IP address must be of the same type and not overlap with the entries in the protected network. (Use /32 CIDR address blocks for IPv4 and /128 CIDR address blocks for IPv6.) If both of these checks fail, the endpoint pair is invalid.

  By default, **Reverse Route Injection is enabled** check box is enabled in Cloud-Delivered Firewall Management Center.

  Subnet/IP Address (Network) remains the default selection.

  When you've selected Protected Networks as *Any* and observe default route traffic being dropped, disable the Reverse Route Injection.

  1. Choose **Secure Connections** > **Site-to-Site VPN & SD-WAN**.
  2. Edit a VPN and then click **IPsec** > **Enable Reverse Route Injection** .
  3. Deploy the configuration changes to remove set reverse-route (Reverse Route Injection) from the crypto map configuration and remove the VPN-advertised reverse route that causes the reverse tunnel traffic to be dropped.

- **Access List (Extended)** —An extended access list provides the capability to control the type of traffic that will be accepted by this endpoint, like GRE or OSPF traffic. Traffic may be restricted either by address or port. Click **Add** (+) to add access control list objects.

Note Access Control List is supported only in the point to point topology.

Enable NAT Traversal

Check this check box to allow seamless communication between the peer Firewall Threat Defense devices when there are NAT devices between these devices. For hub and spoke topologies, this option is available only for the spoke. You can uncheck this check box to disable this feature for an endpoint. This parameter is a per peer setting for an endpoint within a topology.

To view or configure the global NAT Traversal (NAT-T) setting, in the **Add/Edit VPN Topology** dialog box:

1. Click the **Advanced** tab.
2. In the left navigation pane, click **Tunnel** .
3. Under **NAT Settings**, the **Keepalive Messages Traversal** is a global setting that enables NAT-T for all endpoints within a topology.

Exempt VPN traffic from network address translation

Check this check box to exempt the VPN traffic from the Network Address Translation (NAT) rules.

If you do not exempt the VPN traffic from the NAT rules, the traffic gets dropped or is not routed through the VPN tunnel to the remote device. After you enable this option, you can view the NAT exemptions for the device in the **NAT policy** page:

- If Unified Firewall Management is enabled in your organization, choose **Policies** > **FTD network policies** > **NAT**. For classic experience, perform the next step.
- Choose **Policies** > **Network policies** > **NAT**
- Click **NAT Exemptions**

Inside interfaces directly connected to the internal network

Specify the security zone or interface group for the inside interface(s) where the protected networks reside. By default, the inside interface is any.

Click + to configure one or more interfaces from a security zone or an interface group that can map to one or more inside interfaces. Ensure that the interface type of the security zone or an interface group is Routed.

Advanced Settings

Enable Dynamic Reverse Route Injection—Reverse Route Injection (RRI) enables routes to be automatically inserted into the routing process, for the networks and hosts protected by a remote tunnel endpoint. Dynamic RRI routes are created only upon the successful establishment of IPsec security associations (SA's).

Note

- Dynamic RRI is supported only on IKEv2, and not supported on IKEv1 or IKEv1 + IKEv2.
- Dynamic RRI isn't supported on originate-only peer, Full Mesh topology, and Extranet peer.
- In Point-to-Point, only one peer can have dynamic RRI enabled.
- Between Hub and Spoke, only one of the endpoints can have dynamic RRI enabled.
- Dynamic RRI cannot be combined with a dynamic crypto map.

Send Local Identity to Peers—Select this option to send local identity information to the peer device. Select one of the following **Local Identity Configuration** from the list and configure the local identity:

- **IP address**—Use the IP address of the interface for the identity.
- **Auto**—Use the IP address for pre-shared key and Cert DN for certificate-based connections.
- **Email ID**—Specify the email ID to use for the identity. The email address can be up to 127 characters.
- **Hostname**—Use the fully qualified hostname.
- **Key ID**—Specify the key-id to use for the identity. The key ID must be fewer than 65 characters.

The local identity is used to configure a unique identity per IKEv2 tunnel, instead of a global identity for all the tunnels. The unique identity allows Firewall Threat Defense to have multiple IPsec tunnels behind a NAT to connect to the Cisco Umbrella Secure Internet Gateway (SIG).

For information about configuring a unique tunnel ID on Umbrella, refer to Cisco Umbrella SIG User Guide.

VPN Filter—Select an extended access list from the list or click **Add** to create a new extended access list object to filter the site-to-site VPN traffic.

The VPN filter provides more security and filters site-to-site VPN data using an extended access list. The extended access list object selected for the VPN filter lets you filter pre-encrypted traffic before entering the VPN tunnel and decrypted traffic that exits a VPN tunnel. The **sysopt permit-vpn** option, when enabled, would bypass the access control policy rules for the traffic coming from the VPN tunnel. When the **sysopt permit-vpn** option is enabled, the VPN filter helps in identifying and filtering the site-to-site VPN traffic.

Note The VPN filter is supported only on point to point, and hub and spoke topologies. It isn't supported on mesh topology.

For hub and spoke topology, you can choose to override the hub VPN filter on the spoke endpoints in case a different VPN filter needs to enabled on a specific tunnel.

Select the **Override VPN Filter on the Hub** option to override the hub VPN filter on the spokes. Select the **Remote VPN Filter** extended access list object or create an access list to override.

Note For an extranet device as a spoke, only the **Override VPN filter on the Hub** option is available.

For more information about sysopt permit-VPN, see Firewall Threat Defense advanced site-to-site VPN tunnel options , on page 19 .

NAT Traversal (NAT-T)

NAT-T allow seamless communication between the peer Firewall Threat Defense devices when there are NAT devices between these devices.

Firewall Threat Defense VPN IKE options

Configure the basic parameters for a point-to-point topology in a route-based VPN as described in Configure Endpoints for a Point to Point Topology and click the **IKE** tab.

For the versions of IKE you have chosen for this topology, specify the **IKEv1/IKEv2 Settings**.

Note Settings in this dialog apply to the entire topology, all tunnels, and all managed devices.

Fields

Policy

Choose the required IKEv1 or IKEv2 policy objects from the predefined list or create new objects to use. You can choose multiple IKEv1 and IKEv2 policies. IKEv1 and IKEv2 support a maximum of 20 IKE policies, each with a different set of values. Assign a unique priority to each policy that you create. The lower the priority number, the higher the priority.
We recommend that you do NOT use values such as 10, 20 and so on because the default IKEv2 policy for remote access VPN can have these values as its priority value. Before deployment, verify that the priority values of the IKE policies (site-to-site and remote access VPN) do NOT conflict.
For details, see IKE policies

Authentication Type

Site-to-site VPN supports two authentication methods, pre-shared key and certificate. For an explanation of the two methods, see Decide VPN authentication methods.

Note In a VPN topology that supports IKEv1, the **Authentication Method** specified in the chosen IKEv1 Policy object becomes the default in the IKEv1 **Authentication Type** setting. These values must match, otherwise, your configuration will error.

- **Pre-shared Automatic Key**—The Cloud-Delivered Firewall Management Center automatically defines the pre-shared key for this VPN. Specify the **Pre-shared Key Length**, the number of characters in the key. The range is 1 to 27.

  The character " (double quote) isn't supported as part of pre-shared keys.

- **Pre-shared Manual Key**—Manually assign the pre-shared key for this VPN. Specify the **Key** and then reenter the same to **Confirm Key**.

  When you choose this option for IKEv2, the **Enforce hex-based pre-shared key only** check box appears, check if desired. If enforced, you must enter a valid hex value for the key, an even number of 2-256 characters, using numerals 0-9, or A-F.

- **Certificate**—When you use certificates as the authentication method for VPN connections, peers obtain digital certificates from a CA server in your PKI infrastructure, and trade them to authenticate each other.

  In the **Certificate** field, select a preconfigured certificate enrollment object. This enrollment object generates a trustpoint with the same name on the managed device. The certificate enrollment object should be associated with and installed on the device, post which the enrollment process is complete, and then a trustpoint is created.

A trustpoint is a representation of a CA or identity pair. A trustpoint includes the identity of the CA, CA-specific configuration parameters, and an association with one enrolled identity certificate.

Before you select this option, note these requirements:

- Ensure you've enrolled a certificate enrollment object on all the endpoints in the topology—A certificate enrollment object contains the Certification Authority (CA) server information and enrollment parameters that are required for creating Certificate Signing Requests (CSRs) and obtaining Identity Certificates from the specified CA. Certificate Enrollment Objects are used to enroll your managed devices into your PKI infrastructure, and create trustpoints (CA objects) on devices that support VPN connections. For instructions on creating a certificate enrollment object, see Add certificate enrollment objects, and for instructions on enrolling the object on the endpoints see one of these as applicable:
  - Install a certificate using self-signed enrollment
  - Install a certificate using EST enrollment
  - Install a certificate using SCEP enrollment
  - Install a certificate using manual enrollment
  - Install a certificate using a PKCS12 file

Note For a site-to-site VPN topology, ensure that the same certificate enrollment object is enrolled in all the endpoints in the topology. For further details, see the table below.

- Refer to this table to understand the enrollment requirement for different scenarios. Some of the scenarios require you to override the certificate enrollment object for specific devices. See Manage object overrides to understand how to override objects.

Certificate Enrollment Types	Device identity certificate for all endpoints is from the same CA		Device identity certificate for all endpoints is from different CAs
	Device-specific parameters are NOT specified in the certificate enrollment object	Device-specific parameters are specified in the certificate enrollment object	
Manual	No override required	Override required	Override required
EST	No override required	Override required	Override required
SCEP	No override required	Override required	Override required
PKCS	Override required	Override required	Override required
Self-signed	NOT applicable	NOT applicable	NOT applicable

- Understand the VPN certificate limitations mentioned in Secure Firewall Threat Defense VPN certificate guidelines and limitations.

Note If you use a Windows Certificate Authority (CA), the default application policies extension is **IP security IKE intermediate**. If you use this default setting, you must select the **Ignore IPsec Key Usage** option in the Advanced Settings section on the **Key** tab in the **PKI Certificate Enrollment** dialog box for the object you select. Otherwise, the endpoints can't complete the site-to-site VPN connection.

Firewall Threat Defense VPN IPsec options

Configure the basic parameters for a point-to-point topology in a route-based VPN as described in Configure Endpoints for a Point to Point Topology and click the **IPsec** tab.

Note Settings in this dialog apply to the entire topology, all tunnels, and all managed devices.

Crypto-Map Type

A crypto map combines all the components required to set up IPsec security associations (SA). When two peers try to establish an SA, they must each have at least one compatible crypto map entry. IPsec security negotiation uses the proposals defined in the crypto map entry to protect the data flows specified by that crypto map's IPsec rules. Choose static or dynamic for this deployment's crypto-map:

- **Static**—Use a static crypto map in a point-to-point or full mesh VPN topology.
- **Dynamic**—Dynamic crypto-maps essentially create a crypto map entry without all the parameters configured. The missing parameters are later dynamically configured (as the result of an IPsec negotiation) to match a remote peer's requirements.

  Dynamic crypto map policies are applicable to both hub-and-spoke and point-to-point VPN topologies. To apply these policies, specify a dynamic IP address for one of the peers in the topology and ensure that the dynamic crypto-map is enabled on this topology. In a full mesh VPN topology, you can apply only static crypto map policies.

IKEv2 Mode

For IPsec IKEv2 only, specify the encapsulation mode for applying ESP encryption and authentication to the tunnel. This determines what part of the original IP packet has ESP applied.

- **Tunnel mode**—(default) Encapsulation mode is set to tunnel mode. Tunnel mode applies ESP encryption and authentication to the entire original IP packet (IP header and data), hiding the ultimate source and destination addresses and becoming the payload in a new IP packet.

  The major advantage of tunnel mode is that you don't need to modify the end systems to receive the benefits of IPsec. This mode allows a network device, such as a router, to act as an IPsec proxy. That is, the router performs encryption on behalf of the hosts. The source router encrypts packets and forwards them along the IPsec tunnel. The destination router decrypts the original IP datagram and forwards it onto the destination system. Tunnel mode also protects against traffic analysis; with tunnel mode, an attacker can only determine the tunnel endpoints and not the true source and destination of the tunneled packets, even if they are the same as the tunnel endpoints.

- **Transport preferred**—Encapsulation mode is set to transport mode with an option to fall back to tunnel mode if the peer doesn't support it. In transport mode only the IP payload is encrypted, and the original IP headers are left intact. Therefore, the admin must select a protected network that matches the VPN interface IP address.

  This mode has the advantages of adding only a few bytes to each packet and allowing devices on the public network to see the final source and destination of the packet. With transport mode, you can enable special processing (for example, QoS) on the intermediate network based on the information in the IP header. However, the layer 4 header is encrypted, which limits examination of the packet.

- **Transport required**— Encapsulation mode is set to transport mode only, falling back to tunnel mode is allowed. If the endpoints can't successfully negotiate transport mode, due to one endpoint not supporting it, the VPN connection is not made.

Proposals

Click **Edit** (✎) to specify the proposals for your chosen IKEv1 or IKEv2 method. Select from the available **IKEv1 IPsec Proposals** or **IKEv2 IPsec Proposals** objects, or create and then select a new one. See Configure an IKEv1 IPsec proposal object and Configure an IKEv2 IPsec proposal object for details.

Enable Security Association (SA) Strength Enforcement

Enabling this option ensures that the encryption algorithm used by the child IPsec SA isn't stronger (in terms of the number of bits in the key) than the parent IKE SA.

Enable Reverse Route Injection

Reverse Route Injection (RRI) enables static routes to be automatically inserted into the routing process for those networks and hosts protected by a remote tunnel endpoint.

Enable Perfect Forward Secrecy

Whether to use Perfect Forward Secrecy (PFS) to generate and use a unique session key for each encrypted exchange. The unique session key protects the exchange from subsequent decryption, even if the entire exchange was recorded and the attacker has obtained the preshared or private keys used by the endpoint devices. If you select this option, also select the Diffie-Hellman key derivation algorithm to use when generating the PFS session key in the Modulus Group list.

Modulus Group

The Diffie-Hellman group to use for deriving a shared secret between the two IPsec peers without transmitting it to each other. A larger modulus provides higher security but requires more processing time. The two peers must have a matching modulus group. For a full explanation of the options, see Decide which Diffie-Hellman modulus group to use.

Lifetime Duration

The number of seconds a security association exists before expiring. The default is 28,800 seconds.

Lifetime Size

The volume of traffic (in kilobytes) that can pass between IPsec peers using a given security association before it expires. The default is 4,608,000 kilobytes. Infinite data isn't allowed.

ESPv3 Settings

Validate incoming ICMP error messages

Choose whether to validate ICMP error messages received through an IPsec tunnel and destined for an interior host on the private network.

Enable 'Do Not Fragment' Policy

Define how the IPsec subsystem handles large packets that have the do-not-fragment (DF) bit set in the IP header.

Policy

- Copy DF bit—Maintains the DF bit.
- Clear DF bit—Ignores the DF bit.
- Set DF bit—Sets and uses the DF bit.

Enable Traffic Flow Confidentiality (TFC) Packets

Enable dummy TFC packets that mask the traffic profile which traverses the tunnel. Use the **Burst**, **Payload Size**, and **Timeout** parameters to generate random length packets at random intervals across the specified SA.



Note You can enable dummy Traffic Flow Confidentiality (TFC) packets at random lengths and intervals on an IPsec security association (SA). You must have an IKEv2 IPsec proposal set before enabling TFC.

Enabling TFC packets prevents the VPN tunnel from being idle. Thus the VPN idle timeout configured in the group policy doesn't work as expected when you enable the TFC packets.

Firewall Threat Defense advanced site-to-site VPN deployment options

The sections describe the advanced options you can specify in your site-to-site VPN deployment. These settings apply to the entire topology, all tunnels, and all managed devices.

Firewall Threat Defense VPN advanced IKE options

Advanced > IKE > ISAKAMP settings

IKE Keepalive

Enable or disables IKE Keepalives. You can set this option to EnableInfinite so that the device never starts the keepalive monitoring itself.

Threshold

Specifies the IKE keep alive confidence interval. This interval is the number of seconds allowing a peer to idle before beginning keepalive monitoring. The minimum and default interval is 10 seconds; the maximum interval is 3600 seconds.

Retry Interval

Specifies number of seconds to wait between IKE keep alive retries. The default is 2 seconds, the maximum is 10 seconds.

Identity Sent to Peers:

Choose the identity that the peers will use to identify themselves during IKE negotiations:

- **autoOrDN**(default)—Determines IKE negotiation by connection type: IP address for preshared key, or Cert DN for certificate authentication (not supported).
- **ipAddress**—Uses the IP addresses of the hosts exchanging ISAKMP identity information.
- **hostname**—Uses the fully qualified domain name of the hosts exchanging ISAKMP identity information. This name comprises the hostname and the domain name.

Note Enable or disable this option for all your VPN connections.

Peer Identity Validation
During IKE tunnel establishment, the peer provides its identity: either an IP address, a Fully Qualified Domain Name (FQDN), or a Distinguished Name (DN). It also presents a certificate, which contains none, some, or all of these fields.

If IKE peer identity validation is enabled, the Firewall Threat Defense compares the peer's identity to the respective field in the certificate to see if the information matches. If the information matches, then the peer's identity is validated and the Firewall Threat Defense establishes the tunnel. If the information does not match, the tunnel is not established.

- **Do not check**—Firewall Threat Defense does not validate the peer identity.
- **Required**—Firewall Threat Defense validates the peer identity.
- **If supported by cert**—Firewall Threat Defense validates the peer identity only if the peer provides a certificate.

Enable Aggressive Mode
Select this negotiation method for exchanging key information if the IP address isn't known and DNS resolution might not be available on the devices. Negotiation is based on hostname and domain name.

Enable Notification on Tunnel Disconnect
Allows an administrator to enable or disable the sending of an IKE notification to the peer when an inbound packet that is received on an SA does not match the traffic selectors for that SA. This notification is disabled by default.

Advanced > IKE > IKEv2 security association (SA) settings

More session controls are available for IKE v2 that limit the number of open SAs. By default, there's no limit to the number of open SAs.

Cookie Challenge
Whether to send cookie challenges to peer devices in response to SA initiate packets, which can help thwart denial of service (DoS) attacks. The default is to use cookie challenges when 50% of the available SAs are in negotiation. Select one of these options:

- Custom
- Never (default)
- Always

Threshold to Challenge Incoming Cookies
The percentage of the total allowed SAs that are in-negotiation. This triggers cookie challenges for any future SA negotiations. The range is zero to 100%.

Number of SAs Allowed in Negotiation
Limits the maximum number of SAs that can be in negotiation at any time. If used with Cookie Challenge, configure the cookie challenge threshold lower than this limit for an effective cross-check.

Maximum number of SAs Allowed
Limits the number of allowed IKEv2 connections. Default is unlimited.

Enable Notification on Tunnel Disconnect
Allows an administrator to enable or disable the sending of an IKE notification to the peer when an inbound packet that is received on an SA doesn't match the traffic selectors for that SA. Sending this notification is disabled by default.

Firewall Threat Defense VPN advanced IPsec options

Advanced > IPsec > IPsec Settings

The advanced IPsec settings include these configuration options:

- **Enable Fragmentation Before Encryption**: This option lets traffic travel across NAT devices that don't support IP fragmentation. It doesn't impede the operation of NAT devices that do support IP fragmentation.
- **Path Maximum Transmission Unit Aging**: Check to enable Path Maximum Transmission Unit (PMTU) Aging, the interval to reset the PMTU of a Security Association (SA).
- **Value Reset Interval**: Enter the number of minutes at which the PMTU value of an SA is reset to its original value. The range is 10 to 30 minutes, default is unlimited.

Firewall Threat Defense advanced site-to-site VPN tunnel options

This topic lists the advanced site-to-site VPN tunnel options.

Tunnel options

Only available for Hub and Spoke, and Full Mesh topologies. This section doesn’t appear for Point to Point configurations.

- **Enable Spoke to Spoke Connectivity through Hub**—Disabled by default. Choosing this field enables the devices on each end of the spokes to extend their connection through the hub node to the other device.

NAT settings

- **Keepalive Messages Traversal**—Enabled by default. This parameter is a global setting that enables NAT-T for all endpoints within a topology. Check this check box to enable keepalive messages for NAT traversal. NAT traversal keepalive is used for the transmission of keepalive messages when there’s a device (middle device) located between a VPN-connected hub and spoke, and that device performs NAT on the IPsec flow.

  NAT traversal allows seamless communication between the peer Firewall Threat Defense devices when there are NAT devices between these devices. For hub and spoke topologies, this option is available only for the spoke.

  If you select this option, configure the **Interval**, in seconds, between the keepalive signals sent between the spoke and the middle device to indicate that the session is active. The value can be from 5 to 3600 seconds. The default is 20 seconds. You can disable this feature for an endpoint in a topology when you add it using the VPN wizard (**Enable NAT Traversal** check box in the **Add Endpoint** dialog box).

Session settings

- **Enable VPN Idle Timeout**—Enabled by default. This timeout is the amount of time that the VPN connection can be idle, without any activity in the tunnel, before it is disconnected. The default timeout is 30 minutes.

Bidirectional Forwarding Detection (BFD) settings

- **Enable Bidirectional Forwarding Detection Routing**—BFD is a failure detection protocol to detect forwarding path failures. This option enables BFD routing protocol on the SVTIs and DVTIs. When BFD detects a path failure, traffic is rerouted over the newly identified path.

  1. From the **Interval** drop-down list, choose the unit of interval, microseconds or milliseconds, at which BFD control packets are sent.
  2. In the **Multiplier** field, enter the number of consecutive missed BFD control packets allowed before declaring that the peer is unavailable. The range is 3 to 50 packets.
  3. In the **Minimum Transmit and Receive Interval** field, enter the minimum transmit and receive interval of the BFD control packets. The range is 50 to 999 milliseconds or 50000 to 999000 microseconds.



Note BFD over VTI is supported only with Firewall Threat Defense or ASA devices as it requires the peer to advertise its tunnel IP address using the IKEv2 CFG_SET payload.

TrustSec (SGT) settings

- **Enable SGT propagation over Virtual Tunnel Interfaces**—Cisco TrustSec uses Security Group Tags (SGTs) to control access and enforce traffic on a network. This option enables SGT propagation over SVTIs and DVTIs of the VPN topology. To enable SGT propagation on a specific SVTI or DVTI, configure it in individual devices. Note that the Firewall Threat Defense device must be Version 10.0.0 and later.

Access control for VPN traffic

Bypass Access Control policy for decrypted traffic (sysopt permit-vpn)—By default, the Firewall Threat Defense applies access control policy inspection on the decrypted traffic. Enable this option to bypass the ACL inspection. The Firewall Threat Defense still applies the VPN Filter ACL and authorization ACL downloaded from the AAA server to the VPN traffic.

Enable or disable the option for all your VPN connections. If you disable this option, ensure that the traffic is allowed by the access control policy or prefilter policy.



Note For route-based VPNs, **sysopt permit-vpn** does not work. You must always create access control rules to allow route-based VPN traffic.

Certificate map settings

- **Use the certificate map configured in the Endpoints to determine the tunnel**—If this option is enabled (checked), the tunnel is determined by matching the contents of the received certificate to the certificate map objects configured in the endpoint nodes.
- **Use the certificate OU field to determine the tunnel**—Indicates that if a node isn't determined based on the configured mapping (the above option) if selected, then use the value of the organizational unit (OU) in the subject distinguished name (DN) of the received certificate to determine the tunnel.
- **Use the IKE identity to determine the tunnel**—Indicates that if a node isn't determined based on a rule matching or taken from the OU (the above options) if selected, then the certificate-based IKE sessions are mapped to a tunnel based on the content of the phase1 IKE ID.
- **Use the peer IP address to determine the tunnel**—Indicates that if a tunnel isn't determined based on a rule matching or taken from the OU or IKE ID methods (the above options) if selected, then use the established peer IP address.

Configure virtual tunnel interfaces

Cloud-Delivered Firewall Management Center supports a routable logical interface called the Virtual Tunnel Interface (VTI).

Virtual tunnel interfaces

A virtual tunnel interface is a routable logical interface that

- does not require static mapping of IPsec sessions to a physical interface
- associates the IPsec tunnel endpoint with a virtual interface, and
- supports static and dynamic routing policies.

Virtual tunnel interface capabilities

Cloud-Delivered Firewall Management Center supports VTIs as an alternative to policy-based VPN. VTIs support route-based VPN with IPsec profiles attached to the end of each tunnel. VTIs use static or dynamic routes. The device encrypts or decrypts the traffic from or to the tunnel interface and forwards it according to the routing table. Deployments become easier, and having VTI which supports route-based VPN with dynamic routing protocol also satisfies many requirements of a virtual private cloud. Cloud-Delivered Firewall Management Center enables you to easily migrate from crypto-map based VPN configuration to VTI-based VPN.

You can configure route-based VPN with static or dynamic VTI using the site-to-site VPN wizard. Traffic is encrypted using static route, BGP, OSPFv2/v3, or EIGRP.

You can create a routed security zone, add VTI interfaces to it, and define access control rules for the decrypted traffic control over the VTI tunnel.

There are two types of VTI interfaces: static VTI and dynamic VTI.

You can create VTI-based VPNs between:

- Two Firewall Threat Defense devices.

- A Firewall Threat Defense and public cloud.
- One Firewall Threat Defense and another Firewall Threat Defense with service provider redundancy.
- A Firewall Threat Defense and any other device with VTI interfaces.
- A Firewall Threat Defense and another device with policy-based VPN configuration.

For more information, see Static VTI, on page 22 and Dynamic VTI, on page 23.

Static VTI

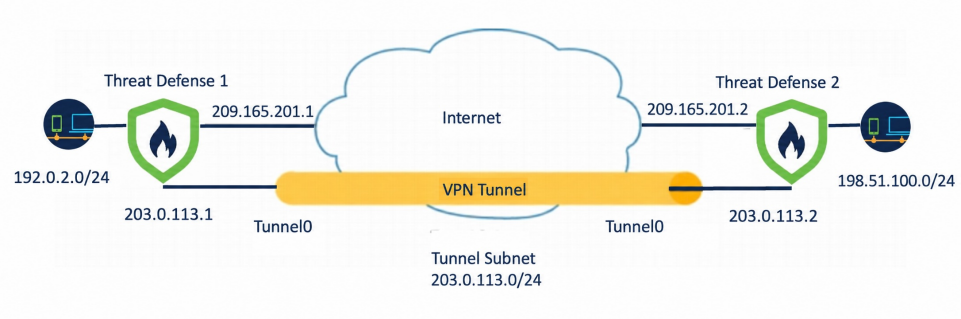
A static VTI is a virtual tunnel interface that

- uses tunnel interfaces to create a tunnel that is always-on between two sites
- requires defining a physical interface as a tunnel source, and
- supports a maximum of 1024 VTIs per device.

To create a static VTI interface in the Cloud-Delivered Firewall Management Center, refer to Add a VTI interface, on page 29.

The figure shows a VPN topology using static VTIs.

Figure 1: Static VTI VPN topology



On Threat Defense 1:

- Static VTI IP address is 203.0.113.1
- Tunnel source is 209.165.201.1
- Tunnel destination is 209.165.201.2

On Threat Defense 2:

- Static VTI IP address is 203.0.113.2
- Tunnel source is 209.165.201.2
- Tunnel destination is 209.165.201.1

Benefits of SVTI

Static VTI provides these benefits:

- Minimizes and simplifies configuration.

  You do not have to track all remote subnets for a crypto map access list, and configure complex access lists or crypto maps.

- Provides a routable interface.

  Supports IP routing protocols such as BGP, EIGRP, and OSPFv2/v3, and static routes.

- Supports backup VPN tunnels
- Supports load balancing using ECMP.
- Supports virtual routers.
- Provides differential access control for VPN traffic.

  You can configure a VTI with a security zone and use it in an AC policy. This configuration:

  - Allows you to classify and differentiate VPN traffic from clear-text traffic and permit VPN traffic selectively.
  - Provides differential access-control for VPN traffic across different VPN tunnels.

Dynamic VTI

A dynamic VTI is a VPN interface that

- uses a virtual template for dynamic instantiation and management of IPsec interfaces
- dynamically generates a unique virtual access interface for each VPN session
- supports multiple IPsec security associations and accepts multiple IPsec selectors proposed by the spoke.

Benefits of dynamic VTI

- Minimizes and simplifies configuration.

  You do not have to configure complex access lists or crypto maps.

- Simplifies management.
  - Easily manage peer configuration for large enterprise hub and spoke deployments.
  - Use only one dynamic VTI for multiple spokes, instead of configuring one static VTI per spoke.
- Provides a routable interface.

  Supports IP routing protocols such as BGP, EIGRP, and OSPFv2/v3.

- Simplifies scaling

  Addition of new spokes does not require any additional VPN configuration on the hub. You may need to update NAT and routing configurations depending upon the setup.

- Support for backup VPN tunnels.
- Supports dynamic spokes.

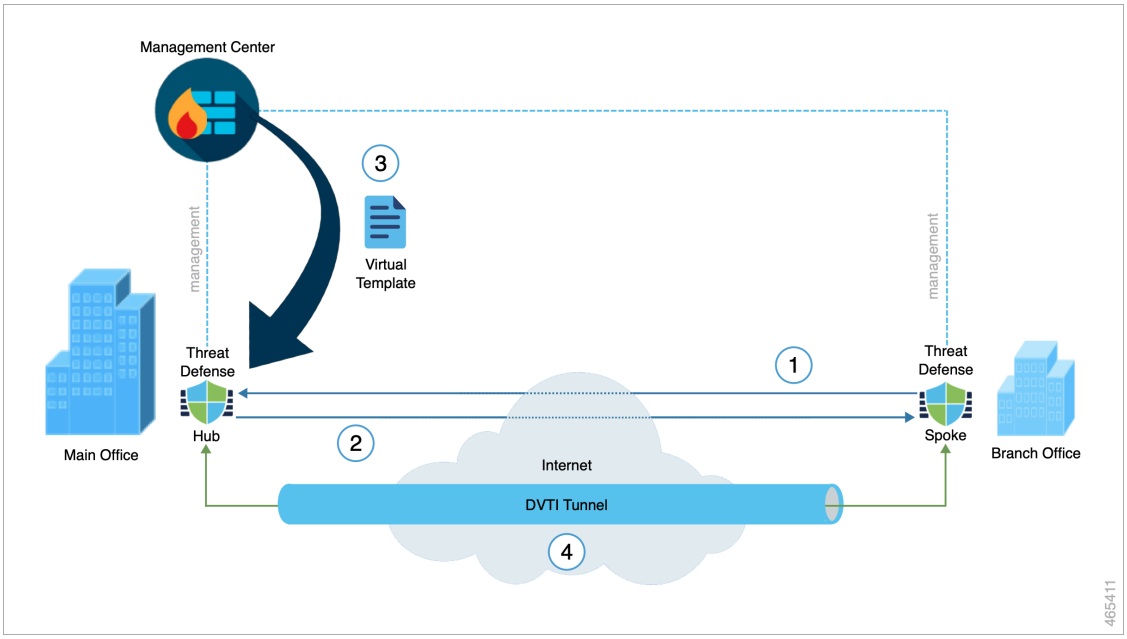
  You do not have to update the hub configuration for spoke's DHCP IP address changes.

- Conserves IP addresses.
  - Uses the IP unnumbered interface functionality to borrow the IP address from another physical or loopback interface.
  - All virtual access interfaces associated to a dynamic VTI use the same IP address.
- Supports virtual routers.
- Provides differential access control for VPN traffic.

  You can configure a VTI with a security zone and use it in an AC policy. This configuration:
  - Allows you to classify and differentiate VPN traffic from clear-text traffic and permit VPN traffic selectively.
  - Provides differential access-control for VPN traffic across different VPN tunnels.

How Does the Cloud-Delivered Firewall Management Center Create a Dynamic VTI Tunnel for a VPN Session

When a spoke initiates a tunnel request with the hub:



1. The spoke initiates an IKE exchange with the hub for a VPN connection.
2. The hub authenticates the spoke.
3. The Cloud-Delivered Firewall Management Center assigns a dynamic virtual template on the hub for the spoke.

   The virtual template dynamically generates a virtual access interface on the hub. This interface is unique for the VPN session with the spoke.
4. The hub establishes a dynamic VTI tunnel with the spoke using the virtual access interface.
   a. The hub and spoke exchange traffic over the tunnel using:

- Specific traffic proposed by the spokes over IKE exchanges.
- BGP/OSPF/EIRGP protocols over the IPsec tunnel.

b. After the VPN session ends, the tunnel disconnects and the hub deletes the corresponding virtual access interface.

To create a dynamic VTI interface in the Cloud-Delivered Firewall Management Center, see Add a VTI interface, on page 29.

To configure a route-based site-to-site VPN using dynamic VTI, see Configure dynamic VTI for a route-based site-to-site VPN, on page 42.

Virtual Routers and Dynamic VTI

You can create virtual routers, associate dynamic VTIs with these virtual routers, and extend the capabilities of dynamic VTIs in your network. You can associate dynamic VTIs either with global or user-defined virtual routers. You can assign a dynamic VTI to only one virtual router.

A virtual router associated with:

- A dynamic VTI is called an Indoor VRF (IVRF).
- A tunnel source interface is known as Front Door VRF (FVRF).

A dynamic VTI and its corresponding protected network interface must be part of the same virtual router. You must map the borrow IP interface and the dynamic VTI to the same virtual router. A tunnel source interface can be part of multiple virtual routers.

To configure virtual routers using dynamic VTI for a route-based site-to-site VPN, refer to Configure a virtual router with dynamic VTI, on page 42.

For more information about a configuration example, refer to Secure traffic from networks with multiple virtual routers over a site-to-site VPN with dynamic VTI

Guidelines for virtual tunnel interfaces

General configuration guidelines

- VTIs are only configurable in IPsec mode.
- Dynamic VTI supports only Firewall Threat Defense devices from Version 7.3.
- You can use static, BGP, EIGRP IPv4, OSPFv2/v3 routes for traffic using the tunnel interface.
- In an HA configuration with dynamic routing, the standby device cannot access the known subnets through the VTI tunnels as these tunnels are created with the active IP address.
- For static and dynamic VTI, ensure that you don't use the borrow IP interface as the tunnel source IP address for any VTI interface.
- When you configure a route-based site-to-site VPN using static or dynamic VTI interfaces, ensure that the value of the TTL hop is more than one if you use BGP.
- If NAT has to be applied, the IKE and ESP packets are encapsulated in the UDP header.

- Ensure that you remove unused transform sets from route-based IKE VPN configurations. If you have unused transform sets, the VPN tunnels may not come up and encapsulation and decapsulation will not happen. Use **show run crypto** and **show VPN-sessiondb detail l2l** commands to view the transform sets.

IPv6 address guidelines

- VTI supports IPv6 addresses.
- You can use an IPv6 address for the tunnel source interface and use the same address as the tunnel endpoint.
- The Cloud-Delivered Firewall Management Center supports the following combinations of VTI IP (or internal networks IP version) over public IP versions: IPv6 over IPv6, IPv4 over IPv6, IPv4 over IPv4, and IPv6 over IPv4.
- VTI supports static and dynamic IPv6 addresses as the tunnel source and destination.
- The tunnel source interface can have IPv6 addresses and you can specify the tunnel endpoint address. If you don't specify the address, by default, the device uses the first IPv6 global address in the list as the tunnel endpoint.

IKE support and guidelines

- Static VTI supports IKE v1 and v2.
- Dynamic VTI supports only IKE version v2, and uses IPsec for sending and receiving data between the tunnel's source and destination.
- IKE and IPsec security associations are re-keyed continuously regardless of data traffic in the tunnel. This ensures that VTI tunnels are always up.
- Tunnel group name must match what the peer sends as its IKEv1 or IKEv2 identity.
- For IKEv1 in LAN-to-LAN tunnel groups, you can use names which aren't IP addresses, if the tunnel authentication method is digital certificates and/or the peer is configured to use aggressive mode.

MTU behavior

- MTU for VTIs is automatically set based on the underlying physical interface.
- For dynamic VTI, the virtual access interface inherits MTU from the configured tunnel source interface. If you don't specify the tunnel source interface, the virtual access interface inherits the MTU from the source interface from which the device accepts the VPN session request.

VTI scaling and guidelines

- Configure a maximum of 1024 static and dynamic VTIs in a device.

  While calculating the VTI count, consider nameif subinterfaces, portchannel configurations, and platform VLAN limitations.
- If you configure more than 400 VTIs in a device in a high-availability setup, configure 45 seconds as the unit holdtime.

Routing protocol support

VTI supports multiple routing protocols for different IP versions.

- VTI supports IPv6 BGP.
- VTI supports IPv4 EIGRP.
- VTI supports IPv4 and IPv6 OSPF.

ECMP configuration guidelines

- Configure spokes' static VTIs in an ECMP zone to load balance application traffic. If you do not configure the ECMP zone, the remaining paths act as backup paths when the primary path goes down.
- The VPN wizard allows you to enable ECMP on the virtual access interfaces associated with the hub's dynamic VTI. Hub device must be Version 10.0 and later for ECMP.
- You can associate VTI interfaces with ECMP zones and configure ECMP static routes to achieve these benefits:
  - Load balancing (Active/Active VTIs)—Connection can flow over any of the parallel VTI tunnels.
  - Seamless connection migration—When a VTI tunnel becomes unreachable, the flows are seamlessly migrated to another VTI interface that is configured in the same zone.
  - Asymmetric routing—Forward traffic flow through one VTI interface and configure the reverse traffic flow through another VTI interface.

  For information on configuring ECMP, see ECMP routing methods.

Access control rule guidelines

- By default, all traffic sent through a VTI is encrypted.
- Access control rules can be applied on a VTI interface to control traffic through VTI.
- For route-based VPNs, Bypass Access Control policy for decrypted traffic (**sysopt connection permit-VPN**) does not work. You must always create access control rules to allow route-based VPN traffic.

Crypto map guidelines

- If you use dynamic crypto maps and dynamic VTIs in your site-to-site VPNs, only the dynamic VTI tunnels will come up. This behaviour occurs because both the crypto maps and dynamic VTIs try to use the default tunnel group. We recommend that you migrate your site-to-site VPNs to dynamic VTIs or use static crypto maps with their own tunnel-groups.
- VTI and crypto map configurations can coexist on the same physical interface if the peer address configured in the crypto map and the tunnel destination for the VTI are different.

Topology guidelines

- Dynamic VTI supports only the hub and spoke topology in the Cloud-Delivered Firewall Management Center.
- We recommend that you configure only a single hub for a route-based hub and spoke topology. To configure a topology with multiple hubs for a set of spokes, with one hub as the backup hub, configure multiple topologies with a single hub and the same set of spokes.

Dynamic VTI and virtual router guidelines

- A dynamic VTI and its corresponding protected network interface must be part of the same virtual router.
- You must map the borrow IP interface and the dynamic VTI to the same virtual router.
- User-defined virtual routers support only BGPv4/v6 and OSPFv2 routing protocols.
- A tunnel source interface can be in a different user-defined virtual router than that associated with the dynamic VTI.

Firewall mode guideline

Configure VTI only in routed mode.

Limitations for Virtual Tunnel Interfaces

This topic lists the limitations for configuring Virtual Tunnel Interfaces in Cloud-Delivered Firewall Management Center.

General limitations

- VTIs do not support clustering.
- VTI tunnel interfaces are not supported for SNMP or Syslog in Platform Settings.

Static VTI limitations

- Only 20 unique IPSec profiles are supported.
- In policy-based routing, you can configure VTI only as an egress interface.
- You cannot configure a VTI interface as a network interface for a remote access VPN policy.

Dynamic VTI limitations

- Dynamic VTI does not support ECMP , VRF in multi-instance, Clustering, IKEv1, or QoS.
- If a spoke has a dynamic IP address and a hub has a dynamic VTI behind a NAT, the tunnel status will be unknown.
- For a dynamic extranet, when multiple spokes establish a connection, the site-to-site monitoring dashboard does not show the individual tunnels.
- If you configure a hub with dynamic VTI behind NAT with dynamic spokes, the VPN monitoring data will not be accurate.
- You cannot configure a VTI interface as a network interface for a remote access VPN policy.

Backup VTI limitations

- Flow resiliency across tunnel failovers isn't supported. For example, the clear text TCP connection gets lost after a tunnel failover, and you need to reinitiate any FTP transfer that took place during the failover.
- Certificate authentication isn't supported in backup VTI.

Add a VTI interface

For configuring a route-based site-to-site VPN, you must create a VTI interface on the devices at both the nodes of the VTI tunnel.

When you specify the tunnel type as dynamic and configure the related parameters, the Cloud-Delivered Firewall Management Center generates a dynamic virtual template. The virtual template dynamically generates the virtual access interface that is unique for each VPN session.

Before you begin

Configure a loopback interface for redundancy of static and dynamic VTI VPN tunnels. For more information, see Configure a loopback interface.

For a Secure Firewall 1200/3100/4200/6100/200 device, IPsec flow offload is also used when the device's VTI loopback interface is enabled.

Procedure

Step 1 Choose **Devices** > **Device Management**.

Step 2 Click the **Edit** icon next to the device on which you want to create a VTI interface.

Step 3 Choose **Add Interfaces > Virtual Tunnel Interface**.

In the **Add Virtual Tunnel Interface** dialog box, configure these parameters:

a) Select the **Tunnel Type** as **Static** or **Dynamic**, and enter the name and description for the interface. By default, the interface is enabled.

b) (Optional) Choose a security zone from the **Security Zone** drop-down list to add the static VTI or dynamic VTI interface to that zone. In the **Priority** field, enter the priority to load balance the traffic across multiple VTIs.

If you want to perform traffic inspection based on a security zone, add the VTI interface to the security zone and configure an access control (AC) rule. To permit the VPN traffic over the tunnel, you need to add an AC rule with this security zone as the source zone.

The priority range is from 0 to 65535. The lowest number has the highest priority. This option is not applicable for dynamic VTI.

c) **Security Group Tag** configuration:

Check the **Propagate Security Group Tag** check box to enable Security Group Tag (SGT) propagation over the VTI.

TrustSec uses SGTs to control access and enforce traffic on a network. To enable SGT propagation over SVTIs and DVTIs of a VPN topology, check the **Enable SGT propagation over Virtual Tunnel Interfaces** check box in the **Advanced > Tunnel** tab of the VPN wizard.

Check the **Enforce Inline Security Group Tag** check box to enable inline SGT tagging over the static VTI interface. Use this option only when the tunnel peer cannot perform Cisco Vendor ID negotiation for SGT.

d) Depending on the tunnel type, do one of the following:

- For a dynamic VTI, enter a unique ID in the range of 1 to 10413 in the **Template ID** field.
- For a static VTI, enter a unique tunnel ID in the range of 0 to 10413 in the **Tunnel ID** field.

e) (Optional for dynamic VTI) Choose the tunnel source interface from the **Tunnel Source** drop-down list.

The VPN tunnel terminates at this interface, a physical or loopback interface. Choose the IP address of the interface from the drop-down list. You can select the IP address irrespective of the IPsec tunnel mode. In case of multiple IPv6 addresses, select the address that you want to use as the tunnel endpoint.

f) Under **IPSec Tunnel Mode**, click the **IPv4** or **IPv6** radio button to specify the traffic type over the IPsec tunnel.

g) Under **IP Address**:

- **Configure IP**: Enter the IPv4 or IPv6 address for the static VTI interface. You cannot configure an IP address for a dynamic VTI interface. Use the **Borrow IP** field for the dynamic VTI interface.
- **Borrow IP (IP unnumbered)**: Choose a physical or loopback interface from the drop-down list, the VTI interface inherits this IP address.

  Ensure that you use an IP address different from the tunnel source IP address. You can use this option for a static or dynamic VTI interface.

  Click Add + to configure a loopback interface. The loopback interface helps to overcome path failures. If an interface goes down, you can access all interfaces through the IP address assigned to the loopback interface.

Step 4 Click **OK**.

Step 5 Click **Save**.

Create a route-based site-to-site VPN

You can configure a route-based site-to-site VPN for the following two topologies:

- **Point to Point**: Configure VTIs on both nodes of the tunnel and use the wizard to configure the VPN.
- **Hub and Spoke**: Configure VTIs on the hub and the spokes. Configure the hub with a dynamic VTI and spokes with static VTIs.

You can configure an extranet device as the hub and managed devices as spokes. You can configure multiple hubs and spokes, and also configure backup hubs and spokes.

- For extranet hubs and spokes, you can configure multiple IPs as backup.
- For managed spokes, you can configure a backup static VTI interface along with the primary VTI interface.

For more information on VTI, see Virtual tunnel interfaces, on page 21.

Note All references to VTI stands for static VTI and dynamic VTI, unless mentioned.

Procedure

Step 1 Choose **Secure Connections** > **Site-to-Site VPN & SD-WAN**, and click **Add**.

Step 2 Enter a name for the VPN topology in the **Topology Name** field.

Step 3 Choose **Route Based (VTI)** and do one of the following:

- Select **Point to Point** as the network topology. To configure endpoints for a route-based **Point to Point** topology, refer to Configure endpoints for a point-to-point route-based VPN topology, on page 31.
- Select **Hub and Spoke** as the network topology. To configure endpoints for a route-based **Hub and Spoke** topology, refer to Configure endpoints for a hub and spoke route-based VPN topology, on page 34.

Step 4 Click **Create**.

Step 5 (Optional) Specify the **IKE**, **IPsec**, and **Advanced** options and click **Save**.

For more information, refer to these topics:

- Firewall Threat Defense VPN IKE options, on page 13
- Firewall Threat Defense VPN IPsec options, on page 15
- Firewall Threat Defense advanced site-to-site VPN deployment options, on page 17

What to do next

After you configure VTI interfaces and VTI tunnel on both the devices, you must configure:

- A routing policy to route the VTI traffic between the devices over the VTI tunnel. For more information, refer to Configure routing and AC policies for VTI, on page 43.
- An access control rule to allow encrypted traffic. Choose **Policies** > **Access Control**.

Configure endpoints for a point-to-point route-based VPN topology

Configure these parameters to configure endpoints for a route-based site-to-site VPN for the **Point to Point** topology nodes:

Before you begin

Configure the basic parameters for a point-to-point topology in a route-based VPN as described in Create a route-based site-to-site VPN, on page 30 and click the **Endpoints** tab.

Procedure

Step 1 Under **Node A**, in the **Device** drop-down menu, select the name of the registered device (Firewall Threat Defense) or extranet as the first endpoint of your VTI tunnel.

For an extranet peer, specify the following parameters:

a. Specify the name of the device.

b. Enter the primary IP address in the **Endpoint IP address** field. If you configure a backup VTI, add a comma and, specify the backup IP address.

c. Click **OK** .

After configuring the above parameters for the extranet hub, specify the pre-shared key for the extranet in the **IKE** tab.

Note The AWS VPC has **AES-GCM-NULL-SHA-LATEST** as the default policy. If the remote peer connects to AWS VPC, select **AES-GCM-NULL-SHA-LATEST** from the **Policy** drop-down list to establish the VPN connection without changing the default value in AWS.

Step 2 For a registered device, you can specify the VTI interface for Node A from the **Virtual Tunnel Interface** drop-down list.

The selected tunnel interface is the source interface for Node A and the tunnel destination for Node B.

If you want to create a new interface on Node A, click the Add + icon and configure the fields as described in Add a VTI interface, on page 29 .

If you want to edit the configuration of an existing VTI, select the VTI in the **Virtual Tunnel Interface** drop-down field and click **Edit VTI**.

Step 3 If your Node A device is behind a NAT device, check the **Tunnel Source IP is Private** check box. In the **Tunnel Source Public IP Address** field, enter the tunnel source public IP address.

Step 4 **Send Local Identity to Peers**—Select this option to send local identity information to the peer device. Select one of the following **Local Identity Configuration** from the list and configure the local identity:

- **IP address**—Use the IP address of the interface for the identity.
- **Auto**—Use the IP address for pre-shared key and Cert DN for certificate-based connections.
- **Email ID**—Specify the email ID to use for the identity. The email ID can be up to 127 characters.
- **Hostname**—Use the fully qualified hostname.
- **Key ID**—Specify the key-id to use for the identity. The key ID must be fewer than 65 characters.

The local identity is used to configure a unique identity per IKEv2 tunnel, instead of a global identity for all the tunnels. The unique identity allows Firewall Threat Defense to have multiple IPsec tunnels behind a NAT to connect to a Cisco Umbrella Secure Internet Gateway (SIG).

For information about configuring a unique tunnel ID on Umbrella, refer to Cisco Umbrella SIG User Guide.

Step 5 (Optional) Click **Add Backup VTI** to specify an extra VTI as the backup interface and configure the parameters.

Note Ensure that both peers of the topology do not have the same tunnel source for the backup VTI. A device cannot have two VTIs with the same tunnel source and tunnel destination; hence, configure a unique tunnel source and tunnel destination combination.

Though the virtual tunnel interface is specified under Backup VTI, the routing configuration determines which tunnel to be used as primary or backup.

Step 6 Expand **Advance Settings** to configure additional configurations for the device. For more information, see Configure advanced settings for a point-to-point topology in a route-based VPN, on page 33.

Step 7 Repeat the above procedure for Node B and click **OK** .

What to do next

- (Optional) Specify the **IKE** options for the deployment as described in Firewall Threat Defense VPN IKE options, on page 13.
- (Optional) Specify the **IPsec** options for the deployment as described in Firewall Threat Defense VPN IPsec options, on page 15.
- (Optional) Specify the **Advanced** options for the deployment as described in Firewall Threat Defense advanced site-to-site VPN deployment options, on page 17.
- Click **Save**.
- To route traffic to the VTI, choose **Devices** > **Device Management** , edit the threat defense device and click the **Routing** tab.

  You can configure the static routes, or use BGP, OSPF v2/v3, or EIGRP for routing the VPN traffic.
- To permit VPN traffic, choose **Policies** > **Security policies** > **Access Control**. Add a rule specifying the security zone of the VTI. For a backup VTI, ensure that you include the backup VTI in the same security zone as that of the primary VTI.

Configure advanced settings for a point-to-point topology in a route-based VPN

Configure advanced settings to enhance the functionality and security of a point-to-point topology in a route-based VPN beyond the basic configuration parameters.

Before you begin

Configure the basic parameters for a point-to-point topology in a route-based VPN as described in Configure endpoints for a point-to-point route-based VPN topology, on page 31 and expand **Advance Settings**.

Procedure

Step 1 Check the **Send Virtual Tunnel Interface IP to the peers** check box to send the VTI IP address to the peer device.

Step 2 Check the **Allow incoming IKEv2 routes from the peers** check box to allow incoming IKEv2 routes from the spokes and peers.

Step 3 Choose one of the following from the **Connection Type** drop-down list:

Answer Only: The device can only respond when a peer device initiates a connection, it can't initiate any connection.

Bidirectional: The device can initiate or respond to a connection. This is the default option.

Configure endpoints for a hub and spoke route-based VPN topology

You can create a route-based site-to-site VPN using dynamic VTI only for hub and spoke topologies. The hub can use only a dynamic VTI and the spokes can use only static VTI interfaces. You can also configure an extranet device as a hub.

Configure these parameters to configure endpoints for a route-based site-to-site VPN for the **Hub and Spoke** topology nodes:

Before you begin

Configure the basic parameters for a hub and spoke topology in a route-based VPN as described in Create a route-based site-to-site VPN, on page 30 and click the **Endpoints** tab.

Procedure

Step 1 Under **Hub Nodes**:

a) Click Add + to configure the hub node in the **Add Endpoint** dialog box.
b) Choose a hub from the **Device** drop-down list.

Note A Firewall Threat Defense device running on software version 7.2 can't be configured as a hub. It must be an extranet or a device running on software version 7.3 or later.

For an extranet hub, configure these parameters:

1. Enter the name of the device.
2. Enter the primary IP address. If you configure a backup VTI, add a comma, and then specify the backup IP address.
3. Click **OK**.

After configuring the above parameters for the extranet hub, specify the pre-shared key for the extranet in the **IKE** tab.

Note The AWS VPC has **AES-GCM-NULL-SHA-LATEST** as the default policy. If the remote peer connects to AWS VPC, select **AES-GCM-NULL-SHA-LATEST** from the **Policy** drop-down list to establish the VPN connection without changing the default value in AWS.

c) Choose a dynamic VTI from the **Dynamic Virtual Tunnel Interface** drop-down list.

Tunnel source configuration is mandatory for a dynamic VTI as the management center needs this information to determine the tunnel destination of the spoke.

Click Add + to add a new dynamic VTI. We recommend that you configure the Borrow IP for the dynamic interface from a loopback interface.

If you want to edit an existing dynamic VTI, select the interface, and click **Edit VTI**.

d) (Optional) If your endpoint device is behind a NAT device, check the **Tunnel Source IP is Private** check box and configure the tunnel source IP address in the **Tunnel Source Public IP Address** field.

e) Click **Routing Policy** to configure the routing policy for the hub.

f) Click **AC Policy** to configure the access control policy.

g) Expand **Advance Settings** to configure additional configurations on the hub. For more information, see Configure advanced settings for hub and spokes in a route-based VPN, on page 37 .

h) Click **OK**.

i) Check the **ECMP** check box to enable ECMP on the dynamic VTIs of hub devices with Version 10.0 or later. All virtual access interfaces on the hub connecting to the same spoke are grouped into an ECMP zone.

Note that this option does not enable ECMP on the spoke SVTIs. You must manually enable ECMP on the spoke SVTIs.

Step 2 Under **Spoke Nodes**:

a) Click Add + to configure the spoke in the **Add Endpoint** dialog box.

b) Choose a spoke from the **Device** drop-down list.

For an extranet spoke, configure these parameters:

1. Enter the name of the device.
2. In **Endpoint IP Address**, choose one of these options:
   - **Static**: Enter the IP address of the device, and the backup IP address, if required.
   - **Dynamic**: Choose this option to dynamically assign the IP addresses for the extranet spokes.
3. Click **OK**.

c) Choose a static VTI from the **Static Virtual Tunnel Interface** drop-down list.

Click Add + to add a new static VTI. Tunnel IP of the static VTI is auto populated, ensure that this IP address is unique for the spoke.

If you want to edit an existing static VTI, select the interface, and click **Edit VTI**.

d) (Optional) If your endpoint device is behind a NAT device, check the **Tunnel Source IP is Private** check box. The Cloud-Delivered Firewall Management Center needs the tunnel source interface address to

configure the tunnel destination IP address on the spokes. In the **Tunnel Source Public IP Address** field, enter the tunnel source public IP address.

e) (Optional) **Send Local Identity to Peers**—Check this check box to send the local identity information to the peer device. Choose one of the following parameters from the **Local Identity Configuration** drop-down list and configure the local identity:

- **IP address**—Use the IP address of the interface for the identity.
- **Auto**—Use the IP address for pre-shared key and Cert DN for certificate-based connections.
- **Email ID**—Specify the email ID to use for the identity. The email ID can be up to 127 characters.
- **Hostname**—Use the fully qualified hostname.
- **Key ID**—Specify the key-id to use for the identity. The key ID must be less than 65 characters.

The local identity is used to configure a unique identity per IKEv2 tunnel, instead of a global identity for all the tunnels. The unique identity allows Firewall Threat Defense to have multiple IPsec tunnels behind a NAT to connect to the Cisco Umbrella Secure internet Gateway (SIG).

For more information about configuring a unique tunnel ID on Umbrella, see Cisco Umbrella SIG User Guide.

f) (Optional) Click **Add Backup VTI** to specify an extra VTI interface as the backup interface.

Note Ensure that both peers of the topology do not have backup VTI configured on the same tunnel source. For instance, if Peer A is having two VTIs (primary and a backup) configured with a single tunnel source interface, say, 10.10.10.1/30, then Peer B also can't have its 2 VTIs with a single tunnel source IP, say 20.20.20.1/30.

Though the virtual tunnel interface is specified under backup VTI, the routing configuration determines which tunnel to be used as primary or backup.

g) Click **Routing Policy** to configure the routing policy for the spoke.
h) Click **AC Policy** to configure the access control policy.
i) Expand **Advance Settings** to configure additional configurations on the spoke and click **OK** . For more information, see Configure advanced settings for hub and spokes in a route-based VPN, on page 37.

What to do next

1. (Optional) Specify the **IKE** options for the deployment as described in Firewall Threat Defense VPN IKE options, on page 13.
2. (Optional) Specify the **IPsec** options for the deployment as described in Firewall Threat Defense VPN IPsec options, on page 15.
3. (Optional) Specify the **Advanced** options for the deployment as described in Firewall Threat Defense advanced site-to-site VPN deployment options, on page 17.
4. Click **Save**.

Configure advanced settings for hub and spokes in a route-based VPN

After configuring basic parameters for a hub and spoke in a route-based VPN, you can configure additional advanced settings to optimize the VPN topology's functionality and performance.

Before you begin

Configure the basic parameters for a hub and spoke in a route-based VPN as described in Configure endpoints for a hub and spoke route-based VPN topology, on page 34 and expand **Advance Settings**.

Note Only the **Connection Type** field is applicable to the device with Version 7.2. The remaining fields don't apply to this version of the device.

Procedure

Step 1 Check the **Send Virtual Tunnel Interface IP to the peers** check box to send the VTI IP address to the peer device.

For a hub, you must check this check box if you use BGP as the routing protocol. This configuration ensures that the loopback IP address is shared in the BGP routing table.

For a spoke, this option is enabled by default.

Step 2 Add the **Protected Networks** to define the networks protected by the VPN endpoint. Click Add + to select a protected network.

For a hub, configure the protected networks behind the hub. This information and the spoke's protected network generate the spoke access list.

You cannot create a static route for a virtual access interface on a hub with dynamic VTI. The hub creates and deletes these interfaces dynamically during tunnel establishment and termination.

For a spoke, configure the spoke's protected network.

To enable static routing for the spokes, after you configure the endpoints for your topology, click the **IPsec** tab and check the **Enable Reverse Route Injection** check box.

You do not need this option if you use BGP, OSPF, or EIGRP.

Step 3 Check the **Allow incoming IKEv2 routes from the peers** check box to allow incoming IKEv2 routes from the spokes and peers.

For a hub: During an IKE exchange, the hub advertises the dynamically created virtual access interfaces to the spokes, and the spokes advertise their VTI IP addresses to the hub.

For a spoke: By default, this option is enabled.

Step 4 In the **Connection Type** drop-down list, choose one of the following:

Answer Only: The device can only respond when a peer device initiates a connection, it can't initiate any connection.

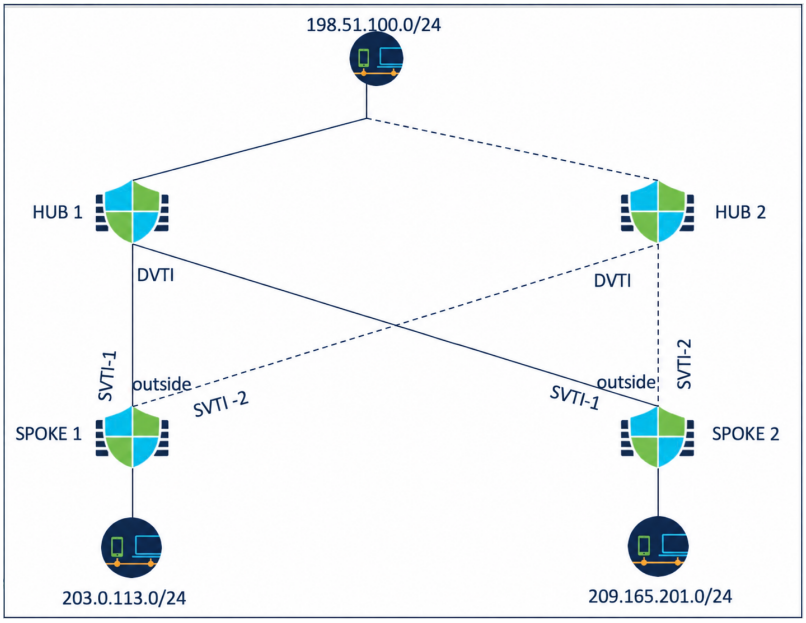
Bidirectional: The device can initiate or respond to a connection. This is the default option.

Configure multiple hubs in a route-based VPN

You can configure a topology with multiple hubs for a set of spokes. With one hub as the backup hub, you can configure multiple topologies with a single hub and the same set of spokes.

In this example, there are two hubs connected to the same set of spokes. Hub 1 is the primary hub and Hub 2 is the secondary hub. To configure this network in the Cloud-Delivered Firewall Management Center, you must configure two route-based hub and spoke topologies:

- Topology 1: Hub 1 connected to spoke 1 and spoke 2.
- Topology 2: Hub 2 connected to spoke 1 and spoke 2.



To configure topology 1:

Procedure

Step 1 Choose **Secure Connections** > **Site-to-Site VPN & SD-WAN**.

Step 2 Enter a name for the VPN topology in the **Topology Name** field.

Step 3 Choose **Route Based (VTI)** and do one of these topologies:

- Select **Point to Point** as the network topology. To configure endpoints for a route-based **Point to Point** topology, see Configure endpoints for a point-to-point route-based VPN topology, on page 31.
- Select **Hub and Spoke** as the network topology. To configure endpoints for a route-based **Hub and Spoke** topology, see Configure endpoints for a hub and spoke route-based VPN topology, on page 34.

Step 4 Configure the IKE version and click the **Endpoints** tab.

Step 5 Under **Hub Nodes**:

a) Click Add + to add the hub.
b) Choose Hub 1 from the **Device** drop-down list.
c) Choose a dynamic VTI from the **Dynamic Virtual Tunnel Interface** drop-down list or click Add + to add a new dynamic VTI.

We recommend that you configure the Borrow IP for the dynamic interface from a loopback interface.

d) (Optional) If your endpoint device is behind a NAT device, check the **Tunnel Source IP is Private** check box and configure the tunnel source IP address in the **Tunnel Source Public IP Address** field.
e) Click **Routing Policy** to configure the routing policy for the hub. You can configure dynamic routing using BGP.
f) Expand **Advance Settings**. You can configure these advanced settings for the hub to enable IKEv2 routing, which can be used if you do not use dynamic routing.

- (Optional) Check the **Send Virtual Tunnel Interface IP to the peers** check box.
- Check the **Allow incoming IKEv2 routes from the peers** check box for the hub to accept routes from the spokes and update the routing table.
- Choose **Connection Type** as Bidirectional from the drop-down list.

g) Click **OK**.

Step 6 Under **Spoke Nodes**:

a) Click Add + to add a spoke.
b) Choose Spoke 1 from the **Device** drop-down list.
c) Choose SVTI-1 as the static VTI for the spoke from the **Static Virtual Tunnel Interface** drop-down list or click Add + to add a new static VTI.

Choose the outside interface as the tunnel source of SVTI-1. Tunnel IP of SVTI-1 is autopopulated, ensure that this IP address is unique for spoke 1 across peers in both the topologies.

d) Expand **Advance Settings**. If you do not use dynamic routing, you can configure these settings to enable IKEv2 routing for the spoke.

- Check the **Send Virtual Tunnel Interface IP to the peers** check box to send the VTI IP address to the peer device.
- Check the **Allow incoming IKEv2 routes from the peers** check box to allow incoming IKEv2 routes from the peers.
- Choose **Connection Type** as Bidirectional from the drop-down list.

e) Click **OK**.
f) Repeat steps 5a to 5e to add spoke 2. Configure SVTI-1 as the static VTI of spoke 2.

Step 7 Configure the IKE and IPSec parameters as required or use the default values.

What to do next

1. Configure topology 2 with hub 2, spoke 1, and spoke 2.

Configure SVTI-2 as the static VTI of spoke 1 and SVTI-2 as the static VTI of spoke 2 (refer the above illustration). Tunnel source for SVTI-2 should be the same outside interface.

2. For each spoke, configure the routing policy. For more information, see Configure routing for multiple hubs in a route-based VPN, on page 40.
3. Verify the configuration and tunnel statuses. For more information, see Verify multiple hubs configuration in a route-based VPN, on page 41.

Configure routing for multiple hubs in a route-based VPN

Configure dynamic routing using BGP on both hub and spoke devices, and implement Policy Based Routing on spokes to enable proper traffic routing in a route-based VPN topology with multiple hubs.

Before you begin

Configure topology 1 and 2 as explained in Configure multiple hubs in a route-based VPN, on page 38.

Procedure

Step 1 Configure dynamic routing for the hub using BGP.

a) Choose **Devices** > **Device Management**, and click **Edit** (✎) then click the **Routing** tab.
b) In the left pane, choose **General Settings > BGP**.
c) Check the **Enable BGP** check box and enter the **AS number**.

You can configure the other fields AS per your requirement.

d) Click **Save**.
e) On the left pane, choose **BGP > IPv4**.
f) Check the **Enable IPv4** check box.
g) Click the **Neighbor** tab, click **Add** and configure these parameters.

1. **IP Address**: Enter the tunnel interface IP address of Spoke 1.
2. **Remote AS**: AS number of Spoke 1.
3. Check the **Enabled Address** check box.
4. Click **OK**.

Repeat the above steps to add Spoke 2 as a neighbor.

h) Click **Save**.
i) Click the **Networks** tab and click **Add** to advertise the network behind the hub to the peers.

Step 2 Configure dynamic routing for the spokes using BGP.

The BGP configuration for the spokes is similar to that of the hub except for these differences:

- Configure Hub 1 and Hub 2 AS the neighbors for both the spokes and use the tunnel interface IP address of the hubs.
- When you configure networks, use the network behind each spoke.

Step 3 Configure Policy Based Routing on the spokes.

a) In the left pane, choose **Policy Based Routing** and click **Add**.
b) Choose the **Ingress Interface** from the drop-down list.
c) Click **Add** to configure a Match ACL.

   For example, for spoke 1, source network is 192.168.20.0/24 and destination network is 192.168.10.0/24.

d) Choose Egress Interfaces from the **Send to** drop-down list.
e) Choose Order from the **Interface Ordering** drop-down list.
f) Select the SVTI-1 and SVTI-2 interfaces AS the egress interfaces.
g) Click **Save**.

If you want to use the hubs as a load-balancing pair, you must configure ECMP.

Step 4 Deploy the configurations on the hub and spokes.

What to do next

Verify the configurations and tunnel statuses. For more information, see Verify multiple hubs configuration in a route-based VPN, on page 41.

Verify multiple hubs configuration in a route-based VPN

To verify the multiple hub configurations and the tunnel statuses.

- After the deployment, check the tunnel status.
- Use these show commands for each endpoint to verify the configurations:
  - **show run route-map**
  - **show run access-list**
  - **show route-map**
  - **show route**

Route traffic through a backup VTI tunnel

A backup VTI tunnel ensures traffic continuity by providing an alternative path when the primary VTI tunnel fails to route traffic in route-based VPN configurations.

You can deploy the backup VTI tunnel in these scenarios:

- Both peers having service provider redundancy backup.

  In this case, there are two physical interfaces, acting as the tunnel sources for the two VTIs of the peers.

- Only one of the peers has service provider redundancy backup.

  In this case, there's an interface backup on only one side of the peer and on the other end, there is only one tunnel source interface.

Before you begin

Ensure that you review the guidelines and limitations. For more information, refer to Guidelines for virtual tunnel interfaces, on page 25.

Procedure

Step 1 Create the VTI interface.

For more information, refer to Add a VTI interface, on page 29

Step 2 In the **Add Endpoint** dialog box of the **Create New VPN Topology** wizard, click **Add Backup VTI** to configure the respective backup interface for each peer.

For more information, refer to

- Configure endpoints for a point-to-point route-based VPN topology, on page 31
- Configure endpoints for a hub and spoke route-based VPN topology, on page 34

Step 3 Choose **Devices** > **Device Management**, and edit the threat defense device and click **Routing** to configure the routing policy.

Step 4 Choose **Policies** > **Access Control** to configure the access control policy.

Configure dynamic VTI for a route-based site-to-site VPN

To configure dynamic VTI for a route-based site-to-site VPN in the Management Center:

Step	Do This	More Information
1	Create a dynamic VTI interface on the hub.	Add a VTI interface, on page 29
2	Create static VTI interfaces on the spokes.	Add a VTI interface, on page 29
3	Create a route-based site-to-site VPN.	Create a route-based site-to-site VPN, on page 30
4	Configure the routing policy and the access control policy.	Configure endpoints for a hub and spoke route-based VPN topology, on page 34

Configure a virtual router with dynamic VTI

To configure a virtual router with dynamic VTI for a route-based site-to-site VPN in the Management Center:

Step	Do This	More Information
1	Create a route-based site-to-site VPN with a dynamic VTI interface on the hub and static VTI(s) on the spoke(s).	Create a route-based site-to-site VPN, on page 30
2	Create a virtual router.	Create a virtual router

Step	Do This	More Information
3	Assign the interfaces to the virtual router.	Configure a virtual router
4	Configure routing policies for the hub and spoke.	Configure endpoints for a hub and spoke route-based VPN topology, on page 34
5	Configure access control policies for the hub and spoke.	Configure endpoints for a hub and spoke route-based VPN topology, on page 34

Configure routing and AC policies for VTI

After you configure VTI interfaces and the VTI tunnel on both the devices, you must configure these parameters:

- A routing policy to route VTI traffic between the devices over the VTI tunnel.
- An access control rule to allow encrypted traffic.

Procedure

Step 1 Configure routing for VTI interfaces.

For the VTI interfaces, you can configure static route or routing protocols such AS BGP, EIGRP, OSPF/OSPFv3.

a. Choose **Devices** > **Device Management**, and edit the Firewall Threat Defense device.

b. Click **Routing**.

c. Configure static route, or BGP, EIGRP, OSPF/OSPFv3.

Routing	Parameters	More Information
Static Route	• **Interface**—Select the VTI interface. For a backup tunnel, select the backup VTI interface. • **Selected Network**—Remote peer's protected network. • **Gateway**—Remote peer's tunnel interface IP address. For a backup tunnel, select the remote peer's backup tunnel interface IP address. • **Metric**—For a backup tunnel, configure a different metric to handle the failover of the traffic flow over the backup tunnel. **Note** DVTI does not support static routes.	Add a static route

Routing	Parameters	More Information

BGP	• Under **General Settings** > **BGP**, enable BGP, provide the AS number of the local device, and add Router ID (if you choose Manual).
	• Under **BGP**, enable IPv4/IPv6 and click the **Neighbor** tab to configure the neighbors.
	  • **IP Address**—Remote peer's VTI interface IP address. For a backup tunnel, add a neighbor with the remote peer's backup VTI interface IP address.
	  • **Remote AS**—Remote peer's AS number.
	• Click the **Redistribution** tab, select the **Source Protocol** AS Connected to enable connected route redistribution.
		Configure BGP

Routing	Parameters	More Information
EIGRP	• Enable EIGRP, provide the AS number of the local device, and select the networks or hosts that participate in the EIGRP routing process. • Click the **Neighbors** tab and define the static neighbors for the EIGRP process. • To advertise summary addresses from a VTI interface, click the **Summary Address** tab, choose the VTI interface from the **Interface** drop-down. From the **Network** drop-down, choose the network to be summarized. • Click the **Interfaces** tab to configure the interface-specific EIGRP routing properties for the VTI interface. To enable EIGRP split-horizon on the interface, check the **Split Horizon** check box. You can also configure the **Hold Time** that is advertised by the device in the EIGRP hello packets.	
OSPF	• Check the **Process 1** check box, and choose the OSPF role. • Click the **Interface** tab and choose a VTI interface.	OSPFv2 configuration
OSPFv3	• Check the **Process 1** and **Enable Process 1** check boxes, and choose the OSPFv3 role. • Click the **Interface** tab and choose a VTI interface.	OSPFv3 configuration

Step 2 Add an access control rule to the access control policy on the device to allow encrypted traffic between the VTI tunnels with these settings:

a. Create the rule with the Allow action.

b. Select the VTI security zone of the local device AS the source zone and the VTI security zone of the remote peer AS the destination zone.

c. Select the VTI security zone of the remote peer AS the source zone and the VTI security zone of the local device AS the destination zone.

For more information about configuring an access control rule, see Create and edit access control rules.

View virtual tunnel interface details

You can view details of dynamic and static VTIs of route-based VPNs on the device. For every VPN topology, you can also view details of all the dynamically generated virtual access interfaces associated with each dynamic VTI.

Before you begin

- Supported device for dynamic VTIs: Firewall Threat Defense Version 7.3 and later

Procedure

Step 1 On the Cloud-Delivered Firewall Management Center, choose **Devices** > **Device Management** and click **Edit** (✎) for your Firewall Threat Defense device. The **Interfaces** page is selected by default.

Step 2 Click the **Virtual Tunnels** tab.

For each VTI, you can view details such as name, IP address, IPsec mode, tunnel source interface details, topology, and remote peer IP. You can also view if path monitoring is enabled for each interface.

Deploy a SASE tunnel on Umbrella

Guidelines for configuring SASE tunnels on Umbrella

SASE topology supports:

- Only PSK-based authentication
- IKEv2
- High availability

General configuration guidelines

- Cloud-Delivered Firewall Management Center does not discover tunnels created directly on Umbrella or by other applications.
- You can add only devices managed by the Cloud-Delivered Firewall Management Center as endpoints for the SASE topology. You cannot add extranet devices.

  For high availability pairs, the HA pair names appear in the endpoint list.

- If you check the **Deploy configuration on threat defense nodes** check box in the wizard, the Umbrella SASE topology configuration is deployed on the Firewall Threat Defense only after the tunnels are deployed on Umbrella.

  The Cloud-Delivered Firewall Management Center requires the local tunnel ID to deploy the Umbrella configuration in Firewall Threat Defense. Umbrella generates the complete tunnel ID (*<prefix>@<umbrella generated ID>*-umbrella.com) only after the Cloud-Delivered Firewall Management Center deploys the tunnel on Umbrella.

Deployment status and monitoring guideline

You can track the tunnel deployment status in three places:

- Cisco Umbrella Configuration dialog box of the wizard
- Notifications page under the Deployments and Tasks tabs
- Site to Site VPN Monitoring dashboard

Tunnel deletion guidelines

- You cannot edit or delete a SASE topology while deployment to Umbrella is in progress.
- When you delete a tunnel from Cloud-Delivered Firewall Management Center and if it is unable to delete the tunnel from Umbrella, you must manually delete it by logging into Umbrella.

Limitations for configuring SASE tunnels on Umbrella

SASE topology does not support these features:

- Clustering
- Certificate-based authentication
- IKEv1

Deploy a SASE tunnel on Umbrella

This section provides instructions to deploy a SASE tunnel on Umbrella from a Firewall Threat Defense device using the Cloud-Delivered Firewall Management Center.

Step	Do This	More Info
1	Review the guidelines and limitations.	Guidelines for configuring SASE tunnels on Umbrella, on page 46
2	Ensure that you meet the prerequisites.	Prerequisites for configuring Umbrella SASE tunnels, on page 48
3	Configure Cisco Umbrella connection settings.	• Configure Cisco Umbrella Connection Settings • Map Management Center Umbrella parameters and Cisco Umbrella API keys, on page 48

Step	Do This	More Info
4	Configure a SASE tunnel on Umbrella.	Configure a SASE tunnel for Umbrella, on page 50
5	View the status of the SASE tunnel.	View SASE tunnel status, on page 51

Prerequisites for configuring Umbrella SASE tunnels

License prerequisites

- You must have a Cisco Umbrella Secure Internet Gateway (SIG) Essentials subscription.
- You must enable your Smart License account with the export-controlled features to deploy tunnels on Umbrella from the Cloud-Delivered Firewall Management Center. If this license is not enabled, you can only create a SASE topology. You cannot deploy tunnels on Umbrella.

General prerequisites

- You must establish an account with Cisco Umbrella at https://umbrella.cisco.com, log into Umbrella at http://login.umbrella.com, and obtain the required information to establish connection to Cisco Umbrella.
- You must register Cisco Umbrella with the Cloud-Delivered Firewall Management Center and configure the management key and the management secret in the Cisco Umbrella Connection Settings. The Cloud-Delivered Firewall Management Center requires the management key and management secret to fetch the datacenter details from the Cisco Umbrella cloud. You must also configure the Organization ID, Network Device Key, Network Device Secret, and the Legacy Network Device Token in the Cisco Umbrella Connection Settings.

  For more information, refer to:

  - Configure Cisco Umbrella Connection Settings
  - Map Management Center Umbrella parameters and Cisco Umbrella API keys, on page 48

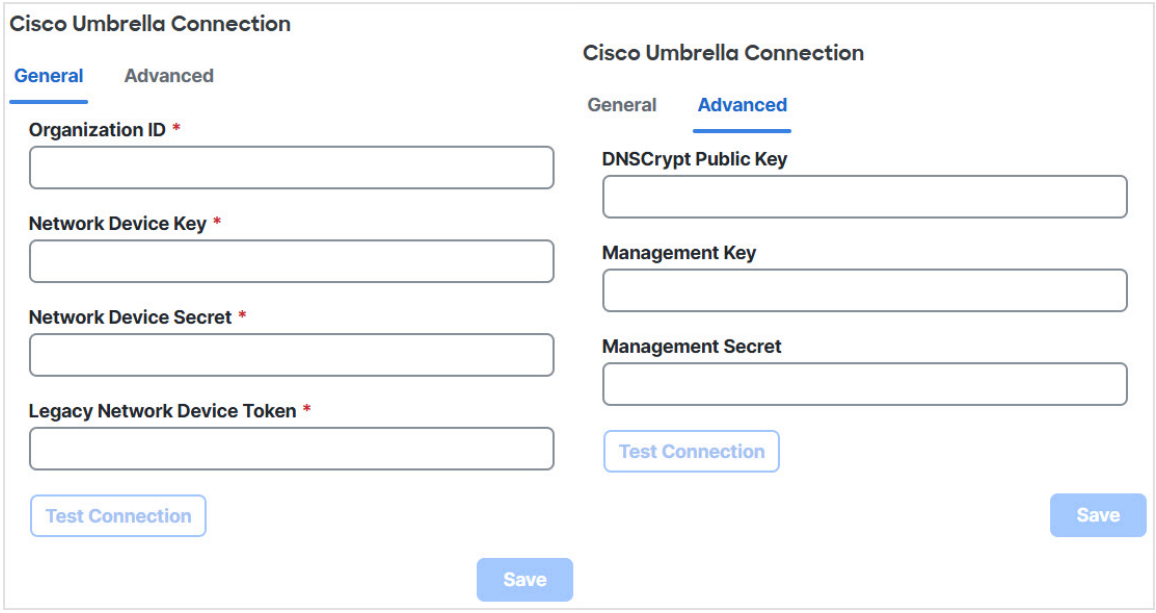
- Ensure that Umbrella data center is reachable from the Firewall Threat Defense.

Map Management Center Umbrella parameters and Cisco Umbrella API keys

To register Cisco Umbrella with the Cloud-Delivered Firewall Management Center and configure the Umbrella parameters in the Cloud-Delivered Firewall Management Center, you must complete these steps:

1. Log in to Cisco Umbrella.
2. Choose **Admin > API Keys > Legacy Keys**.
3. Generate and copy the required API keys.
4. Use the API keys to configure the Cisco Umbrella connection parameters in the Cloud-Delivered Firewall Management Center.

This figures show the parameters that you must configure in the Cisco Umbrella Connection in the Cloud-Delivered Firewall Management Center. DNScrypt Public Key is an optional parameter.



This figure shows the Cisco Umbrella API keys that you must use to register the Cisco Umbrella with the Cloud-Delivered Firewall Management Center.

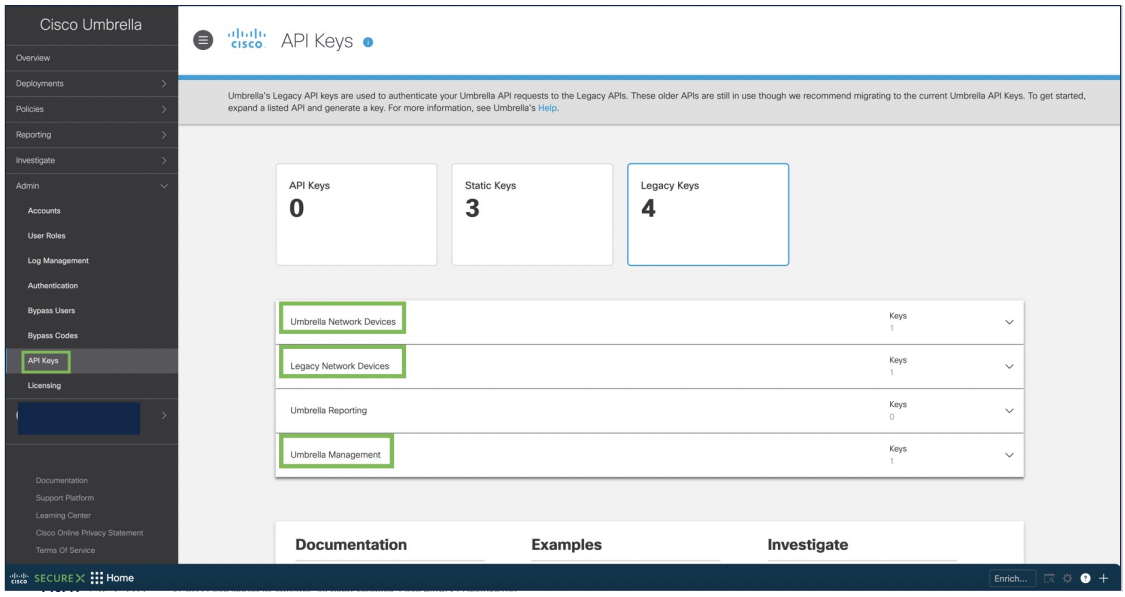


Table 2: Map Cloud-Delivered Firewall Management Center Umbrella Parameters and Cisco Umbrella API Keys

Management Center Parameter	Cisco Umbrella API Keys
Network Device Key Network Device Secret	Umbrella Network Devices
Legacy Network Device Token	Legacy Network Devices

Management Center Parameter	Cisco Umbrella API Keys
Management Key Management Secret	Umbrella Management

Configure a SASE tunnel for Umbrella

Configure a SASE tunnel to establish secure connectivity between threat defense devices and Cisco Umbrella data centers for cloud-delivered security services.

SASE (Secure Access Service Edge) tunnels provide secure connections to Umbrella data centers, enabling cloud-based security functions. These tunnels use IKEv2 with pre-shared keys for authentication and static VTI interfaces for routing traffic.

Before you begin

Ensure that you review the prerequisites and guidelines in Prerequisites for configuring Umbrella SASE tunnels, on page 48 and Guidelines for configuring SASE tunnels on Umbrella, on page 46.

Follow these steps to configure a SASE tunnel for Umbrella:

Procedure

Step 1 Choose **Secure Connections** > **Site-to-Site VPN & SD-WAN**, and enter a unique **Topology Name**.

Step 2 **Pre-shared Key**: This key is auto-generated according to the Umbrella PSK requirements. For a single topology, the pre-shared key is common for all threat defense spokes and Umbrella.

The device and Umbrella share this secret key, and IKEv2 uses it for authentication. If you want to configure this key, it must be between 16 and 64 characters in length, include at least one uppercase letter, one lowercase letter, one numeral, and have no special characters. Each topology must have a unique pre-shared key. If a topology has multiple tunnels, all the tunnels have the same pre-shared key.

Step 3 Choose a data center from the **Umbrella Data center** drop-down list. (Configure routing on the Firewall Threat Defense to ensure reachability of the umbrella data center from the Firewall Threat Defense.)

Step 4 Click **Add** to add a Firewall Threat Defense node.

a) Choose a Firewall Threat Defense from the **Device** drop-down list.

Only devices managed by the Cloud-Delivered Firewall Management Center appear in the list. For high availability pairs, the HA pair names appear in the endpoint list.

b) Choose a static VTI interface from the **VPN Interface** drop-down list.

To create a new static VTI interface, click Add +. The **Add Virtual Tunnel Interface** dialog box appears with the following pre-populated default configurations.

- Tunnel Type is static.
- Name is <*tunnel_source interface logical name*>+ static_vti +<*tunnel ID*>. For example, outside_static_vti_2.
- Tunnel ID is auto-populated with a unique ID.
- Tunnel Source Interface is auto-populated with an interface with an 'outside' prefix.

- IPsec tunnel mode is IPv4.
- IP address is from the 169.254.x.x/30 private IP address range.

c) Enter a prefix for the local tunnel ID in the **Local Tunnel ID** field.

The prefix can have a minimum of eight characters and a maximum of 100 characters. Umbrella generates the complete tunnel ID (<*prefix*>@<*umbrella generated ID*>-umbrella.com) after the Cloud-Delivered Firewall Management Center deploys the tunnel on Umbrella. The Cloud-Delivered Firewall Management Center then retrieves and updates the complete tunnel ID and deploys it on the threat defense device. Each tunnel has a unique local tunnel ID.

d) Click **Save** to add the endpoint device to the topology.

You can add multiple endpoints in a SASE topology.

Step 5 Click **Next** to view the summary of the Umbrella SASE tunnel configuration.

- **Endpoints** pane: Displays the summary of the configured endpoints.
- **Encryption Settings** pane: Displays the default IKEv2 policies and IKEv2 IPsec transform sets for the topology.

Step 6 Check the **Deploy configuration on threat defense nodes** check box to trigger deployment of the network tunnels to the threat defense. This deployment occurs after the tunnels are deployed on Umbrella. Local tunnel ID is required for the threat defense deployment.

Step 7 Click **Save**.

This action:

a. Saves the topology in the Cloud-Delivered Firewall Management Center.

b. Triggers deployment of the network tunnels to the Umbrella.

c. Triggers deployment of the network tunnels to the threat defense devices, if the option is enabled. This action commits and deploys all the updated configurations and policies, including non-VPN policies, since the last deployment on the device.

d. Opens the **Cisco Umbrella Configuration** window and displays the status of the tunnel deployment on Umbrella. For more details, see View SASE tunnel status, on page 51.

What to do next

For the interesting traffic intended to flow through the SASE tunnel, configure a PBR policy with a specific match criteria to send the traffic through the VTI interface.

Ensure that you configure a PBR policy for each endpoint of the SASE topology.

View SASE tunnel status

After configuring Umbrella SASE topology, you need to monitor the tunnel deployment status and verify that the configuration has been successfully pushed to Umbrella.

Procedure

Step 1 Choose **Secure Connections** > **Site-to-Site VPN & SD-WAN**.

Step 2 Enter a unique **Topology Name**, **Pre-shared Key**, choose a data center, add a device, and click **Next**.

Step 3 View the summary of the Umbrella SASE Tunnel configuration and click **Save**.

The **Cisco Umbrella Configuration** window appears.

You can view the topology details such as name, data center, data center IP address, start, and end time of the tunnel deployment.

You can view the deployment status of the tunnels on Umbrella. The different tunnel deployment statuses are:

- Pending: The Cloud-Delivered Firewall Management Center hasn't pushed the configuration to Umbrella.
- Success: The Cloud-Delivered Firewall Management Center successfully configured a tunnel on Umbrella.
- In Progress: The Cloud-Delivered Firewall Management Center is deploying the tunnel on Umbrella.
- Failure: The Cloud-Delivered Firewall Management Center couldn't configure a tunnel on Umbrella.

If the status appears as pending, or failure, use the transcript to troubleshoot the tunnel creation. Click the **Transcript** button to view the transcript details such as the APIs, request payload, and the response received from Umbrella.

Step 4 Click **Umbrella Dashboard** to view the network tunnels in Cisco Umbrella.

Step 5 View the Umbrella tunnel deployment status in:

- The **Notifications** page under the **Deployments** and **Tasks** tabs.
- The **Site to Site** VPN dashboard (**Insights & Reports** > **VPN dashboards** > **Site-to-Site VPN**).

  The dashboard provides a summary of the site to site VPN topologies including the Umbrella SASE topologies. You can view the tunnels between the peer devices and the status of each tunnel. You can also troubleshoot any problems with the tunnel deployment using CLI commands and packet tracer.

Secure traffic using Cisco Secure Access and Firewall Threat Defense devices

Cisco Secure Access is a cloud-based security service edge (SSE) solution that protects against internet threats and enables secure access to the internet, SaaS applications, and private resources from any location.

Benefits of using Secure Access

Secure Firewall seamlessly integrates with Secure Access for various use cases primarily for:

- **Enhanced visibility and control for remote users**: Secure Access extends security policies beyond the physical branch perimeter to protect users regardless of location, providing the same level of threat protection and access controls equivalent to on-premises, creating consistent security enforcement.
- **Simplified security architecture**: Secure Access integrates with a Firewall Threat Defense device to consolidate multiple security functions such as Secure Web Gateway, Cloud Access Security Broker (CASB), Data Loss Prevention (DLP), and cloud firewall. This consolidation provides a single-pane-of-glass experience for policy management across branch deployments and cloud-delivered security services.

Benefits of using SASE wizard to automate Secure Access tunnels from Firewall Threat Defense

You can automatically create secure tunnels between Firewall Threat Defense and Secure Access using Cloud-Delivered Firewall Management Center. The SASE wizard provides a guided workflow to handle the tunnel setup across both the products, helping you avoid switching between consoles or manually matching tunnel parameters.

- Reduce tunnel setup time and configuration errors.
- Secure internet, cloud, and private application access through Secure Access.
- Simplify operations for distributed enterprise environments.

Licenses for automatic tunnel between Secure Access and Firewall Threat Defense devices

This topic lists the required licenses for using the SASE wizard to establish automatic tunnels between Threat Defense devices and Secure Access.

- Cloud-Delivered Firewall Management Center Essentials (formerly Base) license must allow export-controlled functionality. Choose **Administration** > **Licenses** > **Smart Licenses** to verify this in Cloud-Delivered Firewall Management Center.

Note Evaluation license does not support this feature.

- Cisco Secure Access Essentials subscription license must be available.

Prerequisites for an automatic tunnel between Secure Access and Firewall Threat Defense devices

This topic provides the prerequisites required to establish an automatic tunnel between Secure Access and Firewall Threat Defense devices for secure traffic integration.

- Integrate a Cisco Security Cloud Control tenant with a Secure Access organization.
- Configure basic routing and access control policies to route traffic between the Secure Access data center and the Firewall Threat Defense device.

Guidelines for an automatic tunnel between Secure Access and Firewall Threat Defense devices

General guidelines

- Set up routing in the Firewall Threat Defense device to route traffic over the tunnel.
- Delete the Firewall Threat Defense device from the topology and add it again with the required routing configuration to update routing settings in a SASE topology.
- When change management is enabled, tunnels are not deployed automatically to Secure Access. Manually deploy tunnels to Secure Access from the **Site-to-Site VPN** page after the change management ticket is approved.

Multi-ISP SASE topology guidelines

- To create a multi-ISP SASE topology, you must establish multiple tunnels from a Firewall Threat Defense device to the same Secure Access region using different interfaces. Since each SASE topology supports only one WAN interface, use the SASE wizard to create multiple topologies.
- When you configure multi-ISP SASE topologies that use the same region and device but different VPN interfaces, ensure that all topologies have identical settings. The tunnel ID prefix, passphrase, and routing configurations must match across all topologies.
- To update routing settings in a multi-ISP SASE topology, delete the Firewall Threat Defense device from each topology and add it again with the required configuration. Ensure that routing settings are consistent across all topologies.

Limitation for automatic tunnel between Secure Access and Firewall Threat Defense devices

Clustering is not supported.

Configure automatic tunnel between Secure Access and Firewall Threat Defense devices using SASE wizard

The SASE wizard simplifies tunnel creation from Firewall Threat Defense devices to Secure Access by automating multiple steps.

Before you begin

Ensure that you review Prerequisites for an automatic tunnel between Secure Access and Firewall Threat Defense devices, on page 53 .

Procedure

Step 1 Choose **Secure Connections** > **Site-to-Site VPN & SD-WAN** , and click **Add** .

Step 2 In the **Topology Name** field, enter a name for the SASE topology.

Step 3 Click the **SASE Topology** radio button and click **Create** .

Step 4 Configure a Secure Access region by choosing a region from the **Region** drop-down list:

A Secure Access region is a cluster of data centers in a specific geographic area.

Step 5 Configure Firewall Threat Defense nodes:

a) Click **Add** and configure the parameters in the **Add Threat Defense Node** dialog box.
b) From the **Device** drop-down list, choose a Firewall Threat Defense device.

Note Only Firewall Threat Defense devices managed by Cloud-Delivered Firewall Management Center are displayed in this list. Extranet devices won't appear in the list.

c) From the **VPN Interface** drop-down list, choose a WAN-facing or internet-facing physical interface that establishes a VPN connection with Secure Access .

Step 6 Configure **Tunnel ID and passphrase** .

a) In the **Prefix for tunnel ID** field, enter a prefix containing 8 to 30 characters. The prefix can contain alphabets, numbers, period (.), underscore (_), and dash (-).

Secure Access uses this prefix to generate the complete tunnel ID after the Cloud-Delivered Firewall Management Center deploys the tunnel in Secure Access . The tunnel ID format is < *prefix* >@< *org* >< *hub* >.sse.cisco.com.

b) The passphrase is auto generated. However, you can regenerate or enter a different passphrase. In the **Passphrase** field, enter a passphrase containing 16 to 64 characters. The passphrase must contain at least one upper-case alphabet, one lower-case alphabet, and one number, and cannot include special characters.
c) Confirm the passphrase.
d) Click **Next** .

Step 7 Configure **NAT or routing** :

a) Check the **Enable NAT/outbound only** check box to enable Network Address Translation (NAT) for correct routing and identification.

Use this option if the IP address of the subnet behind this tunnel group overlaps with the IP address in your network. When this option is selected, the tunnel group supports only outbound traffic and cannot provide access to private applications hosted at this site.

Note Routing options are unavailable when the network contains overlapping subnets.

b) Configure one of these options for **Routing** :

- Click the **Static routing** radio button.

  Use the **Network IP addresses** field to add the IP addresses of networks for this tunnel group. Add all the public and private addresses used internally by your organization.

- Click the **Dynamic routing** radio button.

  Use this option to advertise your internal networks over BGP.

  In the **AS number** field, enter the AS number of the device. This value must match the device's BGP AS number in the routing settings of the device.

c) Click **Next** .

Step 8 Check the **Deploy to Threat Defense devices** check box to trigger deployment of all the configurations besides the Secure Access auto tunnel configurations that are yet to be deployed on the device.

Step 9 Click **Finish** to save and validate the configurations, and create the SASE topology.

The wizard performs these actions:

- Saves all Firewall Threat Defense device configurations in Cloud-Delivered Firewall Management Center .
- Pushes the configurations to Secure Access.
- Creates virtual tunnel interfaces (VTIs) on the device.
- Creates a SASE VPN topology.
- Triggers deployment of all the configurations besides the Secure Access auto tunnel configurations that are yet to be deployed in the device if the **Deploy to Threat Defense devices** option is enabled.
- Opens the **Configuration of Secure Access Tunnels** dialog box, which displays the status of the tunnel deployment in Secure Access .

  Click the **Transcript Details** ( ) button to view transcript details such as the APIs, request payload, and the responses received from Secure Access .

You can view the SASE topology in the **Site-to-Site VPN & SD-WAN** page (**Secure Connections** > **Site-to-Site VPN & SD-WAN**).

What to do next

1. Create an extended Access Control List (ACL).

   This ACL defines the specific DNS and web traffic intended for routing through the tunnel to Secure Access . For more information, refer to Configure an extended ACL object .
2. Configure the static route to the next hop of the VTI interface.
3. Create a policy-based routing (PBR) policy.

   Use the newly created extended ACL within a policy-based routing policy to direct the defined DNS and web traffic through the tunnel to Secure Access for security inspection. For more information, refer to Configure policy-based routing policy .
4. When you create multiple SASE topologies for a multi-ISP setup, configure ECMP zones with the VPN interfaces to balance the load of application traffic.
5. Perform validation. For more information, refer to Validate Secure Access integration with Firewall Threat Defense devices, on page 56 .

Validate Secure Access integration with Firewall Threat Defense devices

You can validate the integration between Secure Access and Threat Defense devices by checking tunnel statuses and viewing transcript details.

Verify tunnel statuses in Cloud-Delivered Firewall Management Center

To view the SASE topology tunnel details in the site-to-site VPN dashboard, choose **Insights & Reports** > **VPN dashboards** > **Site-to-Site VPN**.

To view more details about each tunnel:

1. For each tunnel, hover your cursor over a topology and click the **View** ( ) icon.
2. Click the **CLI Details** tab.
3. Click **Maximize View**. You can view the output of these commands:
   - **show crypto ipsec sa peer**: Shows the number of packets that are transmitted through the tunnel.
   - **show VPN-sessiondb detail l2l filter ipaddress**: Shows detailed data for the VPN connection.

Verify tunnel statuses in Secure Access

1. In Secure Access, choose **Connect > Network Connections > Network Tunnel Groups**.
2. Click a network tunnel group to view details.

   In the **Network Tunnels** area, the primary and secondary tunnels will be displayed:
   - **Primary1** tunnel originates from the branch's outside interface and connects to the primary Secure Access data center.
   - **Secondary1** tunnel also originates from the branch's outside interface but connects to the secondary data center, serving as a standby for continuous connectivity.

View transcript details

In the **Configuration of Secure Access Tunnels** dialog box, click the **Transcript Details** ( ) button to view the transcript details such as the APIs, request payload, and the responses received from Secure Access.

Transcript details are not available in these scenarios:

- When you delete a Network Tunnel Group (NTG) or an endpoint.
- When you configure any multi-ISP tunnel after configuring the primary tunnel.

Monitor site-to-site topologies

This section provides information about the different ways to monitor your site-to-site topologies.

Monitor site-to-site VPNs using Site-to-Site VPN summary page

You can view a summary of your site-to-site VPN topologies in the **Site-to-Site VPN Summary** page. For each topology, you can view details such as VPN type, network topology, VPN interfaces, VPN devices, and tunnel statuses.

You can edit or delete the topology using the edit and delete buttons. For SASE topology VPNs, you have options to deploy, edit and delete any topology.

You can use the Filter option to filter the connections and topologies by Devices, Topologies, Type and Tunnel Statutes.

Monitor site-to-site VPNs using Site-to-Site VPN dashboard

Choose **Insights & Reports** > **VPN dashboards** > **Site-to-Site VPN** to open the site to site dashboard.

The **Site-to-Site VPN** dashboard provides a snapshot and status of the site-to-site VPN tunnels. You can view the list of tunnels between peer devices and the status of each tunnel: Active, Inactive, or No Active Data. You can filter the data in the table according to the topology, device, and status. The dashboard shows live data and you can also configure the refresh interval.

Benefits of the dashboard

The dashboard provides these benefits:

- Identify problematic VPN tunnels and troubleshoot them.
- Verify connectivity between the site-to-site VPN peers devices.
- Monitor the health of the VPN tunnels to provide uninterrupted VPN connectivity between sites.
- Use Packet tracer to troubleshoot VPN tunnels.

Widgets of the dashboard

- **Tunnel Status**—A table listing the tunnel status of the site-to-site VPNs , including the SASE tunnels, configured using the Cloud-Delivered Firewall Management Center. For more information, refer to Tunnel status widget, on page 59.
- **Tunnel Summary**—Aggregated status of the tunnels in a donut graph.
- **Topology**—Status of tunnels summarized by topology.

Status of VPN tunnels

The different VPN tunnel statuses that you can view in the Site-to-Site dashboard are:

- **Inactive**—A policy-based (crypto map-based) VPN tunnel is inactive if all the IPSec tunnels are down. A VTI or and SASE topology VPN tunnel is down if the tunnel encounters any configuration or connectivity issues.
- **Active**—In the Cloud-Delivered Firewall Management Center, policy-based site-to-site VPNs are configured based on IKE policies and IPsec proposals that are assigned to VPN topologies. A policy-based VPN tunnel is in the Active state if the Cloud-Delivered Firewall Management Center identifies interesting traffic through the tunnel after the deployment. An IKE tunnel is up only if a minimum of one IPsec tunnel is up.

  Route-based VPN (VTI) and SASE topology VPN tunnels do not require interesting traffic to be in the Active state. They are in the Active state if they are configured and deployed without errors.
- **No Active Data**—Policy-based and SASE topology VPN tunnels remain in the No Active Data state until there is a traffic flow event through the tunnel for the first time. The No Active Data state also lists the policy-based and route-based VPNs that have been deployed with errors.

Important notes about tunnel statuses in Cloud-Delivered Firewall Management Center

- The VPN statuses in the Cloud-Delivered Firewall Management Center are event-based. The Cloud-Delivered Firewall Management Center does not initiate status updates. Hence, there might be mismatches between the tunnel statuses in the dashboard and the Firewall Threat Defense. You can view the correct status in the **CLI Details** tab of the **Tunnel Status** widget. Click **Refresh** to see the correct status of the tunnels in the main tunnels table.
- When a Firewall Threat Defense switches over to a secondary Firewall Threat Defense , there is a mismatch between the status of the VPN tunnels in the Cloud-Delivered Firewall Management Center and the Firewall Threat Defense. When the device switches back to the primary device, the correct tunnel status appears.

Limitations of the dashboard

- The table shows the list of site-to-site tunnels , including SASE topology, VPNs that are deployed. It does not show the tunnels that are created and not deployed.
- The table does not show the information about the backup tunnels of policy-based VPNs and backup VTIs.
- For cluster deployments, the table does not show director change in real-time data. It shows only the director information that existed when the VPN was deployed. The director change reflects in the table only after the tunnel is redeployed after the change.

Tunnel status widget

This widget lists the site-to-site VPNs, including SASE topology VPN, configured using the Cloud-Delivered Firewall Management Center and their tunnel statuses.

Features of the widget

Hover over a topology and click **View** ( ) to view these details about the topology:

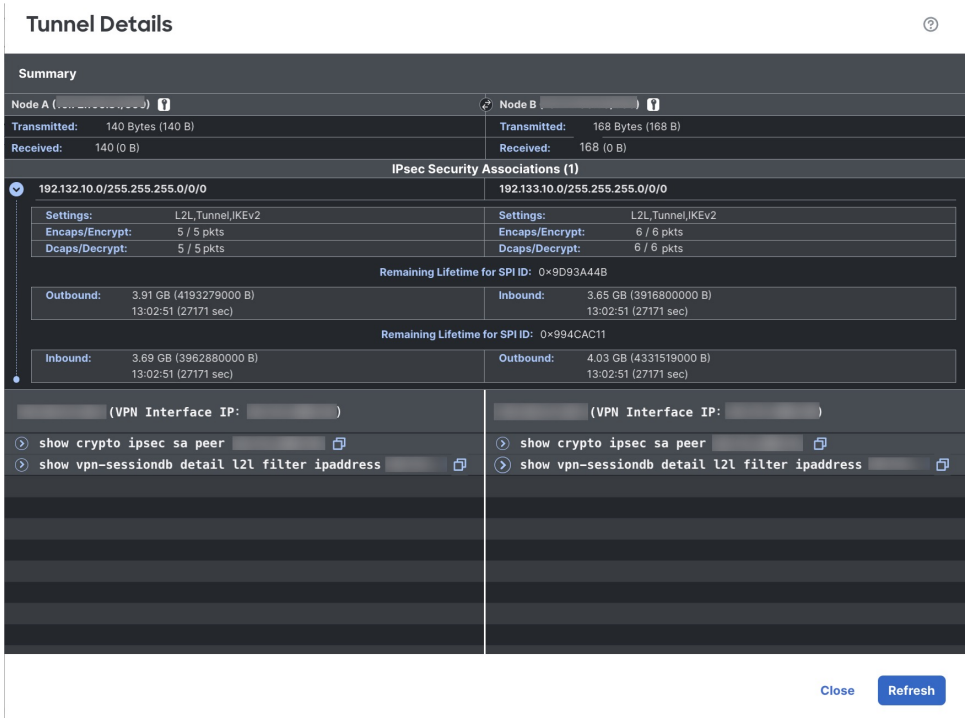
- **General**—Displays more information about the nodes such as IP address and interface name.
- **CLI Details**—Displays the CLI outputs for the following commands:
  - **show crypto ipsec sa peer**: Displays the IPsec SAs built between Node A and B.
  - **show vpn-sessiondb l2l filter ipaddress** : Displays information about VPN sessions.

Note If the tunnel status in not sync, click **Refresh** to see the correct status of the tunnels in the main tunnels table.

For an extranet device, no command output appears.

Important details about the IKE and IPsec sessions, derived from the above command outputs, appear in a summarized, and user-friendly format. You can view the details of both nodes at a time. The icon next to the node name specifies the authentication type: preshared key or client certificate. The details include IKE statistics per tunnel and IPsec SA statistics as shown below:

Figure 2: Tunnel Status > View > CLI Details



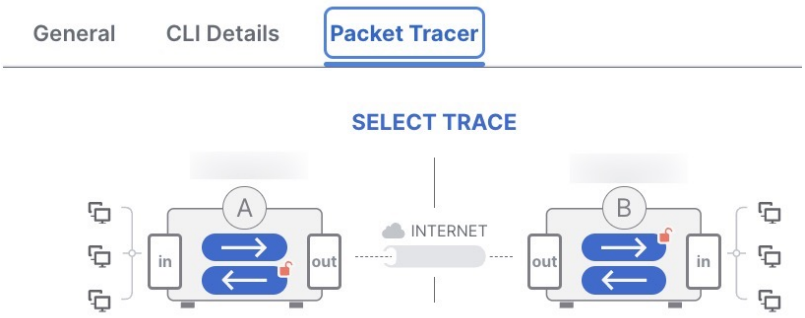
- **Packet Tracer**—Use packet tracer to troubleshoot the VPN tunnels.

Packet Tracer

Packet tracer allows you to troubleshoot VPN tunnels between two threat defense devices. You can check if the VPN connection between device A and device B is up. This tool performs these actions:

- Injects a packet into the device and tracks the packet flow from the ingress to the egress ports.
- Simulates traffic after you configure the ingress interfaces of the devices along with the protected networks.
- Evaluates the packet against modules such as flow and route lookups, ACLs, protocol inspection, NAT, and QoS.

Figure 3: Packet Tracer



For each device, the tool runs an encrypted trace and a decrypted trace (packet is treated as decrypted VPN traffic). You can run four different traces between the ingress and egress ports of the devices. Click the individual encrypt and decrypt options to enable or disable the trace.

When you run the trace, the tool executes the trace sequentially in the following order:

1. Encrypted trace of A.
2. Decrypted trace of B.
3. Encrypted trace of B.
4. Decrypted trace of A.

After the trace completes, you can view the output of the trace with the results of each module.



Note You cannot run a decrypt trace for route-based (VTI-based) VPNs.

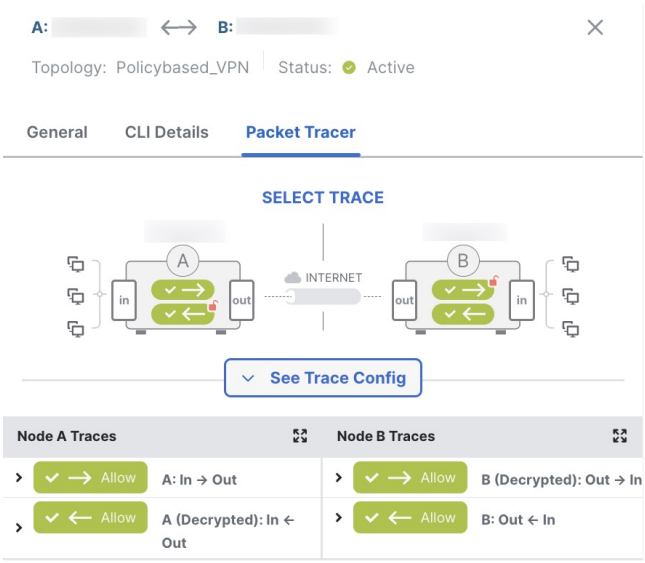
To run the Packet Tracer:

1. Click **See Detailed Config** to view the VPN interface name, VPN interface IP address, VTI interface name, and the VTI Interface IP address.
2. (Optional) Choose a protocol from the **Protocol** drop-down list. You can choose ICMP/8/0, TCP, or UDP.

   ICMP/8/0 is the default option. If you choose ICMP/8/0, 8 indicates the ICMP type as Echo Request and 0 indicates the ICMP code. If you choose TCP or UDP, choose the destination port from the **Destination Port** drop-down list. The range is from 0 to 65535.
3. Choose the ingress interface for both the devices on which to trace the packet from the **Ingress Interface** drop-down lists.
4. Enter an IP address from the same subnet as the ingress interface in the **Protected Network IP Address** fields.
5. Click **Trace Now**.

   After you initiate the trace, you can view if the trace is successful or not for each module. If the tunnel is down, the path appears in red. If the tunnel is up, the path appears in green. If a tunnel is down, click **Re-trace** to run the tool again. For a crypto-map based VPN, when the tunnel is inactive with no interesting traffic, the initial trace can be red. Click **Re-trace** to run the trace again.

Figure 4: Packet Tracer after a Successful Trace



Extranet Nodes: You can initiate a packet trace for VPN tunnels with one node as an extranet. For an extranet node, you cannot choose the ingress interface. The remaining steps for the packet trace are the same. You can't run trace on the extranet side.

For example, if Node A is a managed threat defense and Node B is an extranet:

- Configure the ingress interface for node A.
- Configure the protected network for Node A and B.
- Click **Trace Now**. The traces appear for Node A and not for Node B.

Site-to-site VPN connection event monitoring

The site-to-site VPN connection event allows you to know if the VPN encrypts or do not encrypts the connection and helps you to troubleshoot connectivity issues, especially in multi-hop VPN deployments. The event dashboard of the Cloud-Delivered Firewall Management Center displays the IP address of the VPN peer (peer's IKE address) which encrypts or decrypts the traffic and displays the VPN action.

The VPN actions displayed are:

- If the connection is decrypted by the VPN, the column **Decrypt Peer** displays the IP address of the peer's address from where the flow was received and displays **Decrypt** as the VPN action.
- If the connection is encrypted by the VPN, the column **Encrypt Peer** displays the IP address of the VPN peer to which the flow is sent and displays **Encrypt** as the VPN action.
- If the VPN server cascades the connection, it gets decrypted on one tunnel and gets re-encrypted on another tunnel. In this case, both **Encrypt Peer** and **Decrypt Peer** IP addresses get appears in the event. The column **VPN Action** displays **VPN Routing** as the action to indicate that the connection transit through the VPN server.

If you enable the **Bypass Access Control Policy for decrypted traffic (sysopt permit-VPN)** option, the system bypasses the access control policy and do not log events for decrypted traffic. This option is disabled by default and all decrypted traffic in the VPN tunnel undergoes ACL inspection.

You can integrate Cloud-Delivered Firewall Management Center with Security Analytics and Logging (SAL) and Splunk to view events.

View site-to-site VPN connection events

Access the connection event viewer of the Cloud-Delivered Firewall Management Center to know if the VPN encrypts or do not encrypts the connection traffic and to retrieve the VPN peer details.

Before you begin

Ensure that you enable logging of connection events at the beginning of connection and at the end of connection in the access control rule.

Procedure

Step 1 Choose **Events & Logs** > **Analysis** > **Unified Events**.

Step 2 Choose these columns:

- Decrypt Peer
- Encrypt Peer
- VPN Action

Step 3 Click **Apply**.