



# CHAPTER 12

## Configuring Network Access

---

This chapter describes how to configure network access using Cisco Application Networking Manager (ANM).



### Note

When naming ACE objects (such as a real server, virtual server, parameter map, class map, health probe, and so on), enter an alphanumeric string of 1 to 64 characters, which can include the following special characters: underscore (\_), hyphen (-), and dot (.). Spaces are not allowed.

If you are using ANM with an ACE module or ACE appliance and you configure a named object at the ACE CLI, keep in mind that ANM does not support all of the special characters that the ACE CLI allows you to use when configuring a named object. If you use special characters that ANM does not support, you may not be able to import or manage the ACE using ANM.

---

This chapter includes the following sections:

- [Information About VLANs, page 12-2](#)
- [Configuring VLANs Using Cisco IOS Software \(ACE Module\), page 12-3](#)
- [Configuring Virtual Context VLAN Interfaces, page 12-6](#)
- [Configuring Virtual Context BVI Interfaces, page 12-19](#)
- [Configuring VLAN Interface NAT Pools and Displaying NAT Utilization, page 12-26](#)
- [Configuring Virtual Context Static Routes, page 12-29](#)
- [Configuring Global IP DHCP, page 12-30](#)
- [Configuring Static VLANs for Over 8000 Static NAT Configurations, page 12-32](#)
- [Configuring Gigabit Ethernet Interfaces on the ACE Appliance, page 12-33](#)
- [Configuring Port-Channel Interfaces for the ACE Appliance, page 12-36](#)

# Information About VLANs

This section provides an overview of how the ACE module and appliance use VLANs.

This section includes the following topics:

- [ACE Module VLANs, page 12-2](#)
- [ACE Appliance VLANs, page 12-2](#)

## ACE Module VLANs

The ACE module does not include any external physical interfaces to receive traffic from clients and servers. Instead, it uses internal VLAN interfaces. You assign VLANs from the supervisor engine to the ACE. After the VLANs are assigned to the ACE, you can configure the corresponding VLAN interfaces on the ACE as either routed or bridged for use. When you configure an IP address on an interface, the ACE automatically makes it a routed mode interface.

Similarly, when you configure a bridge group on an interface VLAN, the ACE automatically makes it a bridged interface. Then, you associate a bridge-group virtual interface (BVI) with the bridge group. For more information on bridged groups and BVIs, see the [“Configuring Virtual Context BVI Interfaces” section on page 12-19](#).

The ACE also supports shared VLANs, which are multiple interfaces in different contexts on the same VLAN within the same subnet. Only routed interfaces can share VLANs. Note that there is no routing across contexts even when shared VLANs are configured.

### Related Topics

- [Configuring VLANs Using Cisco IOS Software \(ACE Module\), page 12-3](#)
- [Configuring Virtual Context VLAN Interfaces, page 12-6](#)
- [Configuring Virtual Context BVI Interfaces, page 12-19](#)
- [Configuring Virtual Context Static Routes, page 12-29](#)
- [Configuring Global IP DHCP, page 12-30](#)

## ACE Appliance VLANs

The ACE appliance has four physical Ethernet interface ports. All VLANs are allocated to the physical ports. After the VLANs are assigned, you can configure the corresponding VLAN interfaces as either routed or bridged for use. When you configure an IP address on an interface, the ACE appliance automatically makes it a routed mode interface.

Similarly, when you configure a bridge group on an interface VLAN, the ACE appliance automatically makes it a bridged interface. Then, you associate a BVI with the bridge group.

The ACE appliance also supports shared VLANs, which are multiple interfaces in different contexts on the same VLAN within the same subnet. Only routed interfaces can share VLANs. Note that there is no routing across contexts even when shared VLANs are configured.

In routed mode, the ACE is considered a router hop in the network. In the Admin or user contexts, the ACE supports static routes only. The ACE supports up to eight equal cost routes for load balancing.

**Related Topics**

- [Configuring Virtual Context VLAN Interfaces, page 12-6](#)
- [Configuring Virtual Context BVI Interfaces, page 12-19](#)
- [Configuring Gigabit Ethernet Interfaces on the ACE Appliance, page 12-33](#)
- [Configuring Port-Channel Interfaces for the ACE Appliance, page 12-36](#)

## Configuring VLANs Using Cisco IOS Software (ACE Module)

To allow the ACE module to receive traffic from the supervisor engine in the Catalyst 6500 series switch or Cisco 7600 series router, you must create VLAN groups on the supervisor engine and then assign the groups to the ACE module. After the VLAN groups are assigned to the ACE module, you can configure the VLAN interfaces on the ACE module. By default, all VLANs are allocated to the Admin context on the ACE module.

This section includes the following topics:

- [Creating VLAN Groups Using Cisco IOS Software](#)
- [Assigning VLAN Groups to the ACE Module Through Cisco IOS Software](#)
- [Adding Switched Virtual Interfaces to the MSFC](#)

## Creating VLAN Groups Using Cisco IOS Software

In Cisco IOS software, you can create one or more VLAN groups and then assign the groups to the ACE module. For example, you can assign all the VLANs to one group, create an inside group and an outside group, or create a group for each customer.

You cannot assign the same VLAN to multiple groups; however, you can assign up to a maximum of 16 groups to an ACE. VLANs that you want to assign to multiple ACEs, for example, can reside in a separate group from VLANs that are unique to each ACE.

To assign VLANs to a group using Cisco IOS software on the supervisor engine, use the **svclc vlan-group** command. The syntax of this command is as follows:

```
svclc vlan-group group_number vlan_range
```

The arguments are as follows:

- *group\_number*—Number of the VLAN group.
- *vlan\_range*—One or more VLANs (2 to 1000 and 1025 to 4094) identified in one of the following ways:
  - A single number (*n*)
  - A range (*n-x*)

Separate numbers or ranges by commas, as shown in this example:

```
5,7-10,13,45-100
```

For example, to create three VLAN groups, 50 with a VLAN range of 55 to 57, 51 with a VLAN range of 75 to 86, and 52 with VLAN 100, enter:

```
Router(config)# svclc vlan-group 50 55-57  
Router(config)# svclc vlan-group 51 70-86
```

```
Router(config)# svclc vlan-group 52 100
```

### Related Topics

- [Assigning VLAN Groups to the ACE Module Through Cisco IOS Software, page 12-4](#)
- [Adding Switched Virtual Interfaces to the MSFC, page 12-5](#)

## Assigning VLAN Groups to the ACE Module Through Cisco IOS Software

The ACE module cannot receive traffic from the supervisor engine unless you assign VLAN groups to it. To assign the VLAN groups to the ACE module using Cisco IOS software on the supervisor engine, use the **svc module** command in configuration mode. The syntax of this command is as follows:

```
svc module slot_number vlan-group group_number_range
```

The arguments are as follows:

- *slot\_number*—Slot number where the ACE module resides. To display slot numbers and the devices in the chassis, use the **show module** command in Exec mode. The ACE module appears as the Application Control Engine Module in the Card Type field.
- *group\_number\_range*—One or more group numbers that are identified in one of the following ways:
  - A single number (*n*)
  - A range (*n-x*)

Separate numbers or ranges by commas, as shown in this example:

```
5,7-10
```

For example, to assign VLAN groups 50 and 52 to the ACE module in slot 5, and VLAN groups 51 and 52 to the ACE module in slot 8, enter the following commands:

```
Router(config)# svc module 5 vlan-group 50,52  
Router(config)# svc module 8 vlan-group 51,52
```

To view the group configuration for the ACE module and the associated VLANs, use the **show svclc vlan-group** command. For example, enter the following commands:

```
Router(config)# exit  
Router# show svclc vlan-group
```

To view VLAN group numbers for all devices, use the **show svc module** command. For example, enter the following command:

```
Router# show svc module
```



### Note

Enter the **show vlans** command in Exec mode from the Admin context to display the ACE module VLANs that are downloaded from the supervisor engine.

### Related Topics

- [Creating VLAN Groups Using Cisco IOS Software, page 12-3](#)
- [Adding Switched Virtual Interfaces to the MSFC, page 12-5](#)

## Adding Switched Virtual Interfaces to the MSFC

A VLAN defined on the Multilayer Switch Feature Card (MSFC) is called a switched virtual interface (SVI). If you assign the VLAN used for the SVI to the ACE module, then the MSFC routes between the ACE module and other Layer 3 VLANs. By default, only one SVI can exist between the MSFC and the ACE. However, for multiple contexts, you may configure multiple SVIs for unique VLANs on each context.

### Procedure:

- 
- Step 1** (Optional) If you need to add more than one SVI to the ACE module, enter the following command:
- ```
Router(config)# svclc multiple-vlan-interfaces
```
- Step 2** Add a VLAN interface to the MSFC. For example, to add VLAN 55, enter the following command:
- ```
Router(config)# interface vlan 55
```
- Step 3** Set the IP address for this interface on the MSFC. For example, to set the address 10.1.1.1 255.255.255.0, enter the following command:
- ```
Router(config-if)# ip address 10.1.1.1 255.255.255.0
```
- Step 4** Enable the interface. For example, enter the following command:
- ```
Router(config-if)# no shut
```
- 



### Note

To monitor any VLAN that is associated with more than two trunk ports, physical ports, or trunk-physical ports on the supervisor engine, enable the autostate feature by using the **svclc autostate** command. When you associate a VLAN to these ports, autostate declares that the VLAN is up. When a VLAN state change occurs on the supervisor engine, autostate sends a notification to the ACE module to bring the interface up or down.

---

To view this SVI configuration, use the **show interface vlan** command. For example, enter the following command:

```
Router# show int vlan 55
```

### Related Topics

- [Creating VLAN Groups Using Cisco IOS Software, page 12-3](#)
- [Assigning VLAN Groups to the ACE Module Through Cisco IOS Software, page 12-4](#)

# Configuring Virtual Context VLAN Interfaces

You can configure VLAN interfaces for virtual contexts on the ACE.



## Note

The options that appear when you choose **Config > Devices > context** depend on the device associated with the virtual context and the role associated with your account.

## Assumptions

This topic assumes the following:

- A Layer 3/Layer 4 or Management policy map has been configured for this virtual context. For more information, see the [“Configuring Traffic Policies” section on page 14-1](#).
- An access control list has been configured for this virtual context. Entering an ACL name does not configure the ACL; you must configure the ACL on the ACE appliance. For more information, see the [“Configuring Security with ACLs” section on page 6-80](#).

## Procedure

**Step 1** Choose **Config > Devices > context > Network > VLAN Interfaces**.

The VLAN Interface table appears.

**Step 2** In the VLAN Interface table, click **Poll Now** to instruct ANM to poll the devices and display the current values and click **OK** when prompted if you want to poll the devices for data now.



## Note

Even if the periodic polling is enabled, ANM polls all the devices thus ignoring the statistics defined during the periodic polling.

**Step 3** Click **Add** to add a new VLAN interface, or choose an existing VLAN interface and click **Edit** to modify it.



## Note

If you click **Edit**, not all of the fields can be modified.

**Step 4** Enter the VLAN interface attributes (see [Table 12-1](#)). Click **More Settings** to access the additional VLAN interface attributes.

By default, ANM hides the default VLAN interface attributes and the VLAN interface attributes that are not commonly used.



## Note

Fields and information related to IPv6 require ACE module and ACE appliance software Version A5(1.0) or later.



## Note

If you create a fault-tolerant VLAN, do not use it for any other network traffic.

Table 12-1 VLAN Interface Attributes


Field	Description
VLAN	VLAN identifier. Either accept the automatically incremented entry or enter a different value. Valid entries are from 2 to 4094.
Description	Brief description for this interface.
Interface Type	<p>Role of the virtual context in the network topology of the VLAN interface:</p> <ul style="list-style-type: none"> <li>• <b>Routed</b>—In a routed topology, the ACE virtual context acts as a router between the client-side network and the server-side network. In this topology, every real server for the application must be routed through the ACE virtual context, either by setting the default gateway on each real server to the virtual contexts server-side VLAN interface address, or by using a separate router with appropriate routes configured between the ACE virtual context and the real servers.</li> </ul> <p> <b>Note</b> A routed VLAN interface can support both IPv4 and IPv6 addresses at the same time. IPv6 requires ACE module and ACE appliance software Version A5(1.0) or later.</p> <ul style="list-style-type: none"> <li>• <b>Bridged</b>—In a bridged topology, the ACE virtual context bridges two VLANs, a client-side VLAN and a real-server VLAN, on the same subnet using a bridged virtual interface (BVI). In this case, the real server routing does not change to accommodate the ACE virtual context. Instead, the ACE virtual context becomes a “bump in the wire” that transparently handles traffic to and from the real servers.</li> <li>• <b>Unknown</b>—Choose Unknown if you are unsure of the network topology of the VLAN interface.</li> </ul>
IP Address	<p>Field that appears for the Routed Interface Type. Enter the IPv4 address assigned to this interface. This address must be a unique IP address that is not used in another context. Duplicate IP addresses in different contexts are not supported.</p> <p>If this interface is only used for IPv6 traffic, entering an IPv4 address is optional. IPv6 requires ACE module and ACE appliance software Version A5(1.0) or later.</p>
Alias IP Address	Field that appears for the Routed interface type. Enter the IPv4 address of the alias that this interface is associated with.
Peer IP Address	Field that appears for the Routed interface type. Enter the IPv4 address of the remote peer.
Netmask	Field that appears for the Routed interface type. Choose the subnet mask to be used.
BVI	Field that appears for the Bridged interface type. Enter the number of the bridge group to be configured on this VLAN. When you configure a bridge group on a VLAN, the ACE automatically makes it bridged. Valid entries are from 1 to 4094.
Admin Status	Administrative state of the interface. Specify whether you want the interface to be Up or Down.
Enable MAC Sticky	<p>Check box that instructs the ACE to convert dynamic MAC addresses to sticky secure MAC addresses and to add this information to the running configuration.</p> <p>Uncheck the check box to indicate that the ACE is not to convert dynamic MAC addresses to sticky secure MAC addresses.</p>

Table 12-1 VLAN Interface Attributes (continued)



Field	Description
Enable Normalization	<p>Check box that specifies that normalization is to be enabled on this interface. Uncheck the check box to indicate that normalization is to be disabled on this interface for IPv4, IPv6, or both. The IPv6 option requires ACE module and ACE appliance software Version A5(1.0) or later.</p> <p> <b>Caution</b> Disabling normalization may expose your ACE and network to potential security risks. Normalization protects your networking environment from attackers by enforcing strict security policies that are designed to examine traffic for malformed or malicious segments.</p>
Enable IPv6	<p>Field that appears only for ACE module and ACE appliance software Version A5(1.0) or later, and for the Routed interface type. Check the check box to enable IPv6 on this interface. By default, IPv6 is disabled. The interface cannot be in bridged mode. When you enable IPv6, the ACE automatically does the following:</p> <ul style="list-style-type: none"> <li>• Configures a link-local address (if not previously configured)</li> <li>• Performs duplicate address detection (DAD)</li> </ul> <p>Clear the check box to indicate that IPv6 is disabled on this interface.</p>
IPv6 Global Address	<p>Field that appears only for ACE module and ACE appliance software Version A5(1.0) or later, and for the Routed interface type. A global address is an IPv6 unicast address that is used for general IPv6 communication. Each global address is unique across the entire Internet. Therefore, its scope is global. The low order 64 bits can be assigned in several ways, including autoconfiguration using the EUI-64 format. You can configure only one globally unique IPv6 address on an interface.</p> <p>When you configure a global IPv6 address on an interface, the ACE automatically does the following:</p> <ul style="list-style-type: none"> <li>• Configures a link-local address (if not previously configured)</li> <li>• Performs duplicate address detection (DAD) on both addresses</li> </ul>
IPv6 Address	<p>To configure an IPv6 global address on an interface, enter a complete IPv6 address with a prefix of 2000::/3 to 3fff::/3. For example, enter 2001:DB8:1::0.</p> <p>Check the EUI-64 box to specify that the low order 64 bits are automatically generated in the IEEE 64-bit Extended Unique Identifier (EUI-64) format specified in RFC 2373. To use EUI-64, the Prefix Length field must be less than or equal to 64 and the host segment must be all zeros.</p>
Alias IPv6 Address	<p>When you configure redundancy with active and standby ACEs, you can configure a VLAN interface that has an alias global IPv6 address that is shared between the active and standby ACEs. The alias IPv6 address serves as a shared gateway for the two ACEs in a redundant configuration. You can configure only one alias global IPv6 address on an interface.</p> <p>To configure an IPv6 alias global address, enter a complete IPv6 address with a prefix of 2000::/3 to 3fff::/3. For example, enter 2001:DB8:1::0.</p> <p> <b>Note</b> You must configure redundancy (fault tolerance) on the ACE for the alias global IPv6 address to work.</p>



Table 12-1 VLAN Interface Attributes (continued)



Field	Description
Peer IPv6 Address	<p>To configure an IPv6 peer global address, enter a complete IPv6 address with a prefix of 2000::/3 to 3fff::/3. For example, enter 2001:DB8:1::0.</p> <p>Check the EUI-64 box to specify that the low order 64 bits are automatically generated in the IEEE 64-bit Extended Unique Identifier (EUI-64) format specified in RFC 2373. To use EUI-64, the Prefix Length field must be less than or equal to 64 and the host segment must be all zeros.</p> <p> <b>Note</b> The IPv6 peer global address must be unique across multiple contexts on a shared VLAN.</p> <p>Check the EUI-64 box to specify that the low order 64 bits are automatically generated in the IEEE 64-bit Extended Unique Identifier (EUI-64) format specified in RFC 2373. To use EUI-64, the Prefix Length field must be less than or equal to 64 and the host segment must be all zeros.</p>
Prefix Length	Enter the prefix length for all global addresses to specify how many of the most significant bits (MSBs) are used for the network identifier. Enter an integer from 3 to 127. If you use the optional EUI-64 check box for the global and peer addresses, the prefix must be less than or equal to 64.
IPv6 Unique-Local Address	Field that appears only for ACE module and ACE appliance software Version A5(1.0) or later, and for the Routed interface type. A unique local address is an optional IPv6 unicast address that is used for local communication within an organization and it is similar to a private IPv4 address (for example, 10.10.2.1). Unique local addresses have a global scope, but they are not routable on the internet, and they are assigned by a central authority. All unique local addresses have a predefined prefix of FC00::/7. You can configure only one IPv6 unique local address on an interface.
IPv6 Address	<p>To configure a unique local address, enter a complete IPv6 address with an FC00::/7 prefix in the first field. In the second field after the /, enter the prefix length to specify how many of the most significant bits (MSBs) are used for the network identifier.</p> <p>Check the EUI-64 box to specify that the low order 64 bits are automatically generated in the IEEE 64-bit Extended Unique Identifier (EUI-64) format specified in RFC 2373. To use EUI-64, the Prefix Length field must be less than or equal to 64 and the host segment must be all zeros.</p>
Peer IPv6 Address	<p>In a redundant configuration, you can configure an IPv6 peer unique local address on the active that is synchronized to the standby ACE. You can configure only one peer unique local IPv6 address on an interface.</p> <p>To configure a peer unique local address, enter a complete IPv6 address with an FC00::/7 prefix in the first field. In the second field after the /, enter the prefix length to specify how many of the most significant bits (MSBs) are used for the network identifier.</p> <p> <b>Note</b> The IPv6 peer unique local address must be unique across multiple contexts on a shared VLAN.</p> <p>Check the EUI-64 box to specify that the low order 64 bits are automatically generated in the IEEE 64-bit Extended Unique Identifier (EUI-64) format specified in RFC 2373. To use EUI-64, the Prefix Length field must be less than or equal to 64 and the host segment must be all zeros.</p>
Prefix Length	Enter the prefix length for all unique-local addresses to specify how many of the most significant bits (MSBs) are used for the network identifier. Enter an integer from 7 to 127. If you use the optional EUI-64 check box for the global and peer addresses, the prefix must be less than or equal to 64.

Table 12-1 VLAN Interface Attributes (continued)


Field	Description
IPv6 Link-Local Address	<p>Field that appears only for ACE module and ACE appliance software Version A5(1.0) or later, and for the Routed interface type. By default, when you enable IPv6 or configure a global IPv6 address on an interface, the ACE automatically creates a link local address for it. Every link local address must have a predefined prefix of FE80::/10. You can configure only one IPv6 link local address on an interface. This address always has the prefix of 64.</p> <p>To manually configure the link local address, enter a complete IPv6 address with an FE80::/10 prefix in this field. For example, enter FE80:DB8:1::1.</p>
IPv6 Peer Link-Local Address	<p>Field that appears only for ACE module and ACE appliance software Version A5(1.0) or later, and for the Routed interface type. In a redundant configuration, you can configure an IPv6 peer link local address for the standby ACE. You can configure only one peer link local address on an interface.</p> <p>To configure the peer link local address, enter a complete IPv6 address with an FE80::/10 prefix in this field.</p> <p> <b>Note</b> The IPv6 peer link local address must be unique across multiple contexts on a shared VLAN.</p>

Table 12-1 VLAN Interface Attributes (continued)


Field	Description
<b>More Settings</b>	
Enable ICMP Guard	<p>For ACE module and ACE appliance software versions earlier than A5(1.0), this field does not include the IP version number check boxes and is for enabling the IPv4 version only. Check the <b>IPv4</b>, <b>IPv6</b>, or both check boxes to indicate that ICMP Guard is to be enabled on the ACE.</p> <p>Clear the check boxes to indicate that ICMP Guard is not to be enabled on ACE.</p> <p> <b>Caution</b> Disabling ICMP security checks may expose your ACE and network to potential security risks. When you disable ICMP Guard, the ACE appliance no longer performs NAT translations on the ICMP header and payload in error packets, which can potentially reveal real host IP addresses to attackers.</p>
Enable DHCP Relay	<p>For ACE module and ACE appliance software versions earlier than A5(1.0), this field does not include the IP version number check boxes and is for enabling the IPv4 version only. Check the <b>IPv4</b>, <b>IPv6</b>, or both check boxes to indicate that the ACE is to accept DHCP requests from clients on this interface and to enable the DHCP relay agent. For IPv6, link local address for the</p> <p>Clear the check boxes to indicate that the ACE is not to accept DHCP requests or enable the DHCP relay agent.</p>
Reverse Path Forwarding (RPF)	<p>For ACE module and ACE appliance software versions earlier than A5(1.0), this field does not include the IP version number check boxes and is for enabling the IPv4 version only. Check the <b>IPv4</b>, <b>IPv6</b>, or both check boxes to indicate that the ACE is to discard IP packets if no reverse route is found or if the route does not match the interface on which the packets arrived.</p> <p>Clear the check boxes to indicate that the ACE is not to filter or discard packets based on the ability to verify the source IP address.</p>
Reassembly Timeout (Seconds)	<p>Enter the number of seconds that the ACE appliance is to wait before it abandons the fragment reassembly process if it doesn't receive any outstanding fragments for the current fragment chain (that is, fragments belonging to the same packet).</p> <ul style="list-style-type: none"> <li>• For IPv4, valid entries are 1 to 30 seconds. The default is 5.</li> <li>• For IPv6, valid entries are 1 to 60 seconds. The default is 60. IPv6 requires ACE module and ACE appliance software Version A5(1.0) or later.</li> </ul>
Max. Fragment Chains Allowed	<p>Enter the maximum number of fragments belonging to the same packet that the ACE appliance is to accept for reassembly. For IPv4 and IPv6, valid entries are integers from 1 to 256. The default is 24.</p>
Min. Fragment MTU Value	<p>Enter the minimum fragment size that the ACE appliance accepts for reassembly for a VLAN interface.</p> <ul style="list-style-type: none"> <li>• For IPv4, valid entries are 28 to 9216 bytes. The default is 576.</li> <li>• For IPv6, valid entries are 56 to 9216 bytes. The default is 1280. IPv6 requires ACE module and ACE appliance software Version A5(1.0) or later.</li> </ul>

Table 12-1 VLAN Interface Attributes (continued)

Field	Description
Action For IP Header Options	<p>For ACE module and ACE appliance software versions earlier than A5(1.0), this field does not include the IP version number and is for IPv4 only. Choose the IPv4, IPv6, or both action the ACE appliance is to take when an IP option is set in a packet:</p> <ul style="list-style-type: none"> <li>• <b>Allow</b>—Indicates that the ACE appliance is to allow the IP packet with the IP options set.</li> <li>• <b>Clear</b>—Indicates that the ACE appliance is to clear all IP options from the packet and to allow the packet.</li> <li>• <b>Clear-Invalid</b>—Indicates that the ACE appliance is to clear the invalid IP options from the packet and then allow the packet. This action is the default for IPv4.</li> <li>• <b>Drop</b>—Indicates that the ACE appliance is to discard the packet regardless of any options that are set. This action is the default for IPv6.</li> </ul>
Enable MAC Address Autogenerate	MAC address autogenerate option, which allows you to configure a different MAC address for the VLAN interface.
Min. TTL IP Header Value	<p>Minimum number of hops that a packet is allowed to reach its destination. Valid entries are from 1 to 255. This field is applicable for IPv4 and IPv6 traffic. IPv6 requires ACE module and ACE appliance software Version A5(1.0) or later.</p> <p>Each router along the path decrements the TTL by one. If the packet TTL reaches zero before the packet reaches its destination, the packet is discarded.</p>
MTU Value	Number of bytes for Maximum Transmission Units (MTUs). Valid entries are from 68 to 9216. The default is 1500.
Enable Syn Cookie Threshold Value	<p>Field that is applicable for ACE module software Version A2(1.0) and later, and ACE appliance software Version A3(1.0) and later. Embryonic connection threshold above which the ACE applies SYN-cookie DoS protection.</p> <p>Valid entries are as follows:</p> <ul style="list-style-type: none"> <li>• 2 to 65535 for ACE module software versions earlier than A4(1.0).</li> <li>• 1 to 65535 for ACE module software Version A4(1.0) and later, and ACE appliance software Version A3(1.0) and later.</li> </ul>
Action For DF Bit	<p>Action that the ACE takes when a packet has its DF (Don't Fragment) bit set in the IP header. Choose one of the following settings:</p> <ul style="list-style-type: none"> <li>• <b>Allow</b>—The ACE permits the packet with the DF bit set. If the packet is larger than the next-hop MTU, ACE discards the packet and sends an ICMP unreachable message to the source host. This is the default.</li> <li>• <b>Clear</b>—The ACE clears the DF bit and permit the packet. If the packet is larger than the next-hop MTU, the ACE fragments the packet.</li> </ul>

Table 12-1 VLAN Interface Attributes (continued)


Field	Description
ARP Inspection Type	<p>Type of ARP inspection, which prevents malicious users from impersonating other hosts or routers, known as ARP spoofing. ARP spoofing can enable a “man-in-the-middle” attack. For example, a host sends an ARP request to the gateway router. The gateway router responds with the gateway router MAC address.</p> <p>By default, ARP inspection is disabled on all interfaces, allowing all ARP packets through the ACE. When you enable ARP inspection, the ACE appliance uses the IP address and interface ID (ifID) of an incoming ARP packet as an index into the ARP table. ARP inspection operates only on ingress bridged interfaces.</p> <p> <b>Note</b> If ARP inspection fails, then the ACE does not perform source MAC validation.</p> <p>Choices are as follows:</p> <ul style="list-style-type: none"> <li>• <b>N/A</b>—ARP inspection is disabled.</li> <li>• <b>Flood</b>—Enables ARP forwarding of nonmatching ARP packets. The ACE appliance forwards all ARP packets to all interfaces in the bridge group. This setting is the default. In the absence of a static ARP entry, this option bridges all packets.</li> <li>• <b>No Flood</b>—Disables ARP forwarding for the interface and drops nonmatching ARP packets. In the absence of a static ARP entry, this option does not bridge any packets.</li> </ul>
UDP Config Commands	<p>UDP boost command options:</p> <ul style="list-style-type: none"> <li>• <b>N/A</b>—Not applicable.</li> <li>• <b>IP Destination Hash</b>—Performs destination IP hash during connection.</li> <li>• <b>IP Source Hash</b>—Performs source IP hash during connection lookup.</li> </ul>

Table 12-1 VLAN Interface Attributes (continued)


Field	Description
Secondary IP Groups	<p>Option that is available only for the ACE module A2(3.0), ACE appliance A4(1.0), and later releases of both device types. This option displays only when Interface Type is set to Routed.</p> <p>The number of secondary IP groups that you can enter for a VLAN depends on the ACE release as follows:</p> <ul style="list-style-type: none"> <li>• ACE module A2(3.0) and ACE appliance A4(1.0)—Up to 4 secondary IP groups.</li> <li>• ACE module A2(3.1) and later—Up to 15 secondary IP groups.</li> </ul> <p>The IP, alias IP, and peer IP addresses of each Secondary IP group should be in the same subnet.</p>  <p><b>Note</b> You cannot configure secondary IP addresses on FT VLANs.</p> <p>To create secondary IP groups for the VLAN, do the following:</p> <ol style="list-style-type: none"> <li>Define one or more of the following secondary IP address types: <ul style="list-style-type: none"> <li>– IP—Secondary IP address assigned to this interface. The primary address must be active for the secondary address to be active.</li> <li>– AliasIP—Secondary IP address of the alias associated with this interface.</li> <li>– PeerIP—Secondary IP address of the remote peer.</li> <li>– Netmask—Secondary subnet mask to be used.</li> </ul> <p>The ACE has a system limit of 1,024 for each secondary IP address type.</p> </li> <li>Click <b>Add to selection</b> (right arrow) to add the group to the group display area.</li> <li>Repeat the first two steps for each additional group.</li> <li>(Optional) Rearrange the order in which the groups are listed by selecting one of the group listings in the group display area and click either <b>Move item up in list</b> (up arrow) or <b>Move item down in list</b> (down arrow). Note that the ACE does not care what order the groups are in.</li> <li>(Optional) Edit a group or remove it from the list by selecting the desired group in the group display area and click <b>Remove from selection</b> (left arrow).</li> </ol>
Input Policies	<p>Policy map that is associated with this VLAN interface. From the Available list, double-click a policy map name or use the right arrow to move it to the Selected list. This policy map is to be applied to the inbound direction of the interface; that is, all traffic received by this interface.</p> <p>If you choose more than one policy map, use the Up and Down arrows to choose the priority of the policy map in the Selected list. These arrows modify the order of the policy maps for new VLANs only; they do not modify the policy map order when editing an existing policy map.</p>
Input Access Group	<p>ACL input access group to be associated with this VLAN interface. From the Available list, double-click an ACL name or use the right arrow to move it to the Selected list. Any ACL group listed in the Selected list specifies that this access group is to be applied to the inbound direction of the interface.</p>
Output Access Group	<p>ACL output access group that is associated with this VLAN interface. From the Available list, double-click an ACL name or use the right arrow to move it to the Selected list. Any ACL group listed in the Selected list specifies that this access group is to be applied to the outbound direction of the interface; that is, all traffic sent by this interface.</p>

Table 12-1 VLAN Interface Attributes (continued)


Field	Description
Static ARP Entry (IP/MAC Address)	<p>Static ARP entry.</p> <p>Do the following:</p> <ol style="list-style-type: none"> <li>In the ARP IP Address field, enter the IP address. This field accepts IPv4 addresses only.</li> <li>In the ARP MAC Address field, enter the hardware MAC address for the ARP table entry (for example, 00.02.9a.3b.94.d9).</li> <li>When completed, use the right arrow to move the static ARP entry to the list box. Use the Up and Down arrows to choose the priority of the static ARP entry in the list box. These arrows modify the order of the static ARPs for new VLANs only; they do not modify the static ARP order when editing an existing policy map.</li> </ol>
DHCP Relay Configuration	Enter the IPv4 address of the DHCP server to which the DHCP relay agent is to forward client requests. Enter the IP address in dotted-decimal notation, such as 192.168.11.2.
IPv6 DHCP Forward Interface VLAN	Field that appears only for ACE module and ACE appliance software Version A5(1.0) or later. Enter the VLAN to forward all received client requests with destination being the IPv6 DHCP address configured in the IPv6 DHCP Relay Configuration field.
IPv6 DHCP Relay Configuration	<p>Field that appears only for ACE module and ACE appliance software Version A5(1.0) or later. Enter the IPv6 address for the DHCP server where the DHCP relay agent forwards client requests.</p> <p>Select the VLAN when the server address is a link local address.</p> <p> <b>Note</b> When you enter a DHCPv6 server global IPv6 address, a VLAN is not required.</p>
Managed-Config	<p>Field that appears only for ACE module and ACE appliance software Version A5(1.0) or later. Check box to indicate that the interface use the stateful autoconfiguration mechanism to configure IPv6 addresses.</p> <p>Uncheck the check box to indicate that the interface does not use the stateful autoconfiguration mechanism to configure IPv6 addresses.</p>
Other-Config	<p>Field that appears only for ACE module and ACE appliance software Version A5(1.0) or later. Check box to indicate that the interface use the stateful autoconfiguration mechanism to configure parameters other than IPv6 addresses.</p> <p>Uncheck the check box to indicate that the interface does not use the stateful autoconfiguration mechanism to configure parameters other than IPv6 addresses.</p>
NS Interval	<p>Field that appears only for ACE module and ACE appliance software Version A5(1.0) or later. The ACE sends neighbor solicitation messages through ICMPv6 on the local link to determine the IPv6 addresses of nearby nodes (hosts or routers). You can configure the rate at which the ACE sends these neighbor solicitation messages.</p> <p>By default, the interval at which the ACE sends NS messages for DAD default is 1000 milliseconds (msecs). To configure the interval, enter an integer from 1000 to 2147483647.</p>
NS Reachable Time	<p>Field that appears only for ACE module and ACE appliance software Version A5(1.0) or later. The neighbor solicitation reachable time is the time period in milliseconds during which a host considers the peer is reachable after a reachability confirmation from the peer. A reachability confirmation can include neighbor solicitation or advertisement, or any upper protocol traffic.</p> <p>By default, this time period is 0 milliseconds. To configure this time, enter an integer from 0 to 3600000.</p>

Table 12-1 VLAN Interface Attributes (continued)

Field	Description
Retransmission time	Field that appears only for ACE module and ACE appliance software Version A5(1.0) or later. By default, the advertised retransmission time is 0 milliseconds. To configure the retransmission time, enter an integer from 0 to 3600000.
DAD Attempts	Field that appears only for ACE module and ACE appliance software Version A5(1.0) or later. By default, the number of attempts for sending duplicate address detection (DAD) is 1. To configure the DAD attempts, enter an integer from 0 to 255.
RA Hop Limit	Field that appears only for ACE module and ACE appliance software Version A5(1.0) or later. By default, the hop limit that neighbors should use when originating IPv6 packets is 64. To configure the hop limit in the IPv6 header, enter an integer from 0 to 255.
RA Lifetime	Field that appears only for ACE module and ACE appliance software Version A5(1.0) or later. The router advertisement lifetime is the length of time that neighboring nodes should consider the ACE as the default router before they send RS messages again. By default, this length of time is 1800 seconds (30 minutes). To configure the RA lifetime, enter an integer from 0 to 9000.
RA Interval	Field that appears only for ACE module and ACE appliance software Version A5(1.0) or later. By default, the rate at which the ACE sends RA messages is 600 seconds. To configure the rate, enter an integer from 4 to 1800.
Suppress RA	Field that appears only for ACE module and ACE appliance software Version A5(1.0) or later. Check box to instruct the ACE to not respond to RS messages. The ACE also stops periodic unsolicited RAs that it sends at the RA interval. By default, the ACE automatically responds to RS messages that it receives from neighbors with RA messages that include, for example, the network prefix. You can instruct the ACE to not respond to RS messages. Uncheck the check box to reset the default behavior of automatically responding to RS messages.
IPv6 Router Prefix Advertisement	Field that appears only for ACE module and ACE appliance software Version A5(1.0) or later. Click the Add button to configure the IPv6 prefixes that the ACE advertises in RA messages on the local link.
IPv6 Address/Prefix Length	To configure IPv6 address advertised in the RA messages, enter a complete IPv6 address in the first field. In the second field after the /, enter the prefix length to specify how many of the most significant bits (MSBs) are used for the network identifier.
No Advertisements	Check the check box to indicate that the route prefix is not advertised. Clear the check box to indicate that the route prefix is advertised.



Table 12-1 VLAN Interface Attributes (continued)

Field	Description
Lifetime	<p>Configure the prefix lifetime attributes as follows:</p> <ul style="list-style-type: none"> <li>• Lifetime Duration: <ul style="list-style-type: none"> <li>– Valid Lifetime—By default, the prefix lifetime is 2592000 seconds (30 days). To configure the prefix lifetime in seconds, enter an integer from 0 to 2147183647. Select Infinite to indicate that the prefix never expires.</li> <li>– Preferred Lifetime—By default, the prefix lifetime is 604800 seconds (10 days). To configure how long an IPv6 address remains preferred in seconds, enter an integer from 0 to 2147183647. This lifetime must not exceed the Valid Lifetime. Select Infinite to indicate that the preferred lifetime never expires.</li> </ul> </li> <li>• Lifetime Expiration Date: <ul style="list-style-type: none"> <li>– Valid Month/Day/Year/Time—Valid lifetime expiration date and time.</li> <li>– Preferred Month/Day/Year/Time—Preferred lifetime expiration date and time.</li> </ul> </li> </ul> <p>Use the drop-down lists to select a day, month, and year. To specify the time, use the hh:mm format.</p>
Off-link	<p>This option appears when you enter a Preferred Lifetime field.</p> <p>Check this check box to indicate that the route prefix is on a different subnet for a router to route to it.</p> <p>Clear the check box to indicate that the route prefix is on the same subnet for a router to route to it.</p>
No-autoconfig	<p>This option appears when you enter a Preferred Lifetime field.</p> <p>Check this check box to indicate to the host that it cannot use this prefix when creating an stateless IPv6 address.</p> <p>Clear the check box to indicate to the host that it can use this prefix when creating an stateless IPv6 address.</p>

**Step 5** Do one of the following:

- Click **Deploy Now** to immediately deploy this configuration on the ACE and save your entries to the running-configuration and startup-configuration files. This option appears for virtual contexts.
- Click **Cancel** to exit this procedure without saving your entries and to return to the previous window.

**Step 6** (Optional) To display statistics and status information for a VLAN interface, choose the VLAN interface from the VLAN Interface table, then click **Details**.

The **show interface vlan** CLI command output appears. See the “[Displaying VLAN Interface Statistics and Status Information](#)” section on page 12-18 for details.

#### Related Topics

- [Configuring VLAN Interface NAT Pools and Displaying NAT Utilization](#), page 12-26
- [Displaying All VLAN Interfaces](#), page 12-18
- [Displaying VLAN Interface Statistics and Status Information](#), page 12-18

## Displaying All VLAN Interfaces

You can display all of the VLAN interfaces associated with a specific virtual context by choosing **Config > Devices > context > Network > VLAN Interfaces**.

The VLAN Interface table appears with the information shown in [Table 12-2](#).

**Table 12-2** VLAN Interface Table Fields

Field	Description
VLAN	VLAN number.
Description	Description for this interface.
Interface Type	Role of the virtual context in the network topology of the VLAN interface.
IP Address	IP address assigned to this interface including the netmask for an IPv4 address or a prefix length for an IPv6 address. IPv6 requires ACE module and ACE appliance software Version A5(1.0) or later.  This table does not display the IPv6 link-local, unique-local, and multicast addresses for the interface. To display these addresses, click Details to display the output for the <b>show ipv6 vlan</b> command.
IPv6 Config Status	The status whether IPv6 is enabled or disabled on the interface. IPv6 requires ACE module and ACE appliance software Version A5(1.0) or later.
Admin Status	Status of the interface, which can be Up or Down.
Operational Status	Operational state of the device (Up or Down).
Last Polled	Date and time of the last time that ANM polled the device to display the current values.

### Related Topics

- [Configuring Virtual Context VLAN Interfaces, page 12-6](#)
- [Configuring Virtual Context BVI Interfaces, page 12-19](#)
- [Displaying VLAN Interface Statistics and Status Information, page 12-18](#)

## Displaying VLAN Interface Statistics and Status Information

You can display statistics and status information for a particular VLAN interface.

### Procedure

**Step 1** Choose **Config > Devices > context > Network > VLAN Interfaces**.

The VLAN Interfaces table appears.

**Step 2** Choose a VLAN interface from the VLAN Interfaces table, and click **Details**.

The **show interface vlan**, **show ipv6 interface vlan**, and **show ipv6 neighbors** CLI commands appears. The IPv6 commands require ACE module and ACE appliance software Version A5(1.0) or later. Click on the command to display its output. For details on the displayed output fields, see either the *Cisco ACE Module Routing and Bridging Configuration Guide* or the *Cisco ACE 4700 Series Appliance Routing and Bridging Configuration Guide*.

- Step 3** Click **Update Details** to refresh the output for the **show interface vlan** CLI command.
- Step 4** Click **Close** to return to the VLAN Interfaces table.

#### Related Topics

- [Configuring Virtual Context VLAN Interfaces, page 12-6](#)
- [Displaying All VLAN Interfaces, page 12-18](#)

## Configuring Virtual Context BVI Interfaces

You can configure Bridge-Group Virtual Interfaces (BVI) for virtual contexts. The ACE supports virtual contexts containing BVI interfaces. You can configure two interface VLANs into a group and bridge packets between them. All interfaces are in one broadcast domain and packets from one VLAN are switched to the other VLAN. The ACE bridge mode supports only two Layer 2 VLANs per bridge group.



#### Note

The options that appear when you choose **Config > Devices > context** depend on the device associated with the virtual context and the role associated with your account.

This section includes the following topics:

- [Configuring BVI Interfaces for a Virtual Context, page 12-19](#)
- [Displaying All BVI Interfaces by Context, page 12-25](#)
- [Displaying BVI Interface Statistics and Status Information, page 12-26](#)

## Configuring BVI Interfaces for a Virtual Context

You can configure BVI interfaces for a virtual context.

#### Procedure

- Step 1** Choose **Config > Devices > context > Network > BVI Interfaces**.
- The BVI Interface configuration table appears.
- Step 2** Click **Poll Now** to instruct ANM to poll the devices and display the current values, and click **OK** when prompted if you want to poll the devices for data now.
- Step 3** Click **Add** to add a new BVI interface.
- Step 4** Enter the interface attributes (see [Table 12-3](#)).



#### Note

When you create or edit a virtual context BVI, if either of the two VLANs do not exist, ANM creates the VLAN and populates the BVI with the description specified in the BVI Interface window.

If you delete the BVI and there are values specified in either of the two VLAN fields, ANM removes the BVI value from the VLAN.

**Note**

Fields and information related to IPv6 require ACE module and ACE appliance software Version A5(1.0) or later.

**Table 12-3 BVI Interface Attributes**


Field	Description
BVI	BVI identifier. Either accept the automatically incremented entry or enter a different, unique value for the BVI. Valid entries are from 1 to 4094.
Description	Brief description for this interface.
IP Address	IPv4 address assigned to this interface. This address must be a unique IP address that is not used in another context. Duplicate IP addresses in different contexts are not supported.   <b>Note</b> If this interface is only used for IPv6 traffic, entering an IPv4 address is optional. IPv6 requires ACE module and ACE appliance software Version A5(1.0) or later.
Alias IP Address	IPv4 address of the alias that this interface is associated with.
Peer IP Address	IPv4 address of the remote peer.
Netmask	Subnet mask to be used.
Admin Status	Administrative state of the interface: <b>Up</b> or <b>Down</b> .
Secondary IP Groups	Option that is available only for the ACE module A2(3.0), ACE appliance A4(1.0), and later releases of either device type. The number of secondary IP groups that you can enter for a BVI depends on the ACE release as follows: <ul style="list-style-type: none"> <li>• ACE module A2(3.0) and ACE appliance A4(1.0)—Up to 4 secondary IP groups.</li> <li>• ACE module A2(3.1) and later—Up to 15 secondary IP groups.</li> </ul> <p>To create secondary IP groups for this BVI, do the following:</p> <ol style="list-style-type: none"> <li>Define one or more of the following secondary IP address types: <ul style="list-style-type: none"> <li>– IP—Secondary IP address assigned to this interface. The primary address must be active for the secondary address to be active.</li> <li>– AliasIP—Secondary IP address of the alias associated with this interface.</li> <li>– PeerIP—Secondary IP address of the remote peer.</li> <li>– Netmask—Secondary subnet mask to be used.</li> </ul> <p>The ACE has a system limit of 1,024 for each secondary IP address type.</p> </li> <li>Click Add to selection (right arrow) to add the group to the group display area.</li> <li>Repeat the first two steps for each additional group.</li> <li>(Optional) Rearrange the order in which the groups are listed by selecting one of the group listings in the group display area and click either Move item up in list (up arrow) or Move item down in list (down arrow). Note that the ACE does not care what order the groups are in.</li> <li>(Optional) Edit a group or remove it from the list by selecting the desired group in the group display area and click Remove from selection (left arrow).</li> </ol>
First VLAN	First VLAN whose bridge group is to be configured with this BVI. This VLAN can be the server or client VLAN. Valid entries are from 2 to 4094.

Table 12-3 BVI Interface Attributes (continued)


Field	Description
First VLAN Description	Brief description for the first VLAN.
Second VLAN	Second VLAN whose bridge group is to be configured with this BVI. This VLAN can be the server or client VLAN. Valid entries are from 2 to 4094.
Second VLAN Description	Brief description for the second VLAN.
Enable IPv6	Field that appears only for ACE module and ACE appliance software Version A5(1.0) or later. Check box to enable IPv6 on this interface. By default, IPv6 is disabled. The interface cannot be in bridged mode. When you enable IPv6, the ACE automatically does the following: <ul style="list-style-type: none"> <li>• Configures a link-local address (if not previously configured)</li> <li>• Performs duplicate address detection (DAD) on both addresses</li> </ul> Uncheck the check box to indicate that IPv6 is disabled on this interface.
IPv6 Global Address	Field that appears only for ACE module and ACE appliance software Version A5(1.0) or later. A global address is an IPv6 unicast address that is used for general IPv6 communication. Each global address is unique across the entire Internet. Therefore, its scope is global. The low order 64 bits can be assigned in several ways, including autoconfiguration using the EUI-64 format. You can configure only one globally unique IPv6 address on an interface. <p>When you configure a global address, the ACE automatically does the following:</p> <ul style="list-style-type: none"> <li>• Configures a link-local address (if not previously configured)</li> <li>• Performs duplicate address detection (DAD) on both addresses</li> </ul>
IPv6 Address	To configure an IPv6 global address on an interface, enter a complete IPv6 address with a prefix of 2000::/3 to 3fff::/3. For example, enter 2001:DB8:1::0. <p>Check the EUI-64 check box to specify that the low order 64 bits are automatically generated in the IEEE 64-bit Extended Unique Identifier (EUI-64) format specified in RFC 2373. To use EUI-64, the Prefix Length field must be less than or equal to 64 and the host segment must be all zeros.</p>
Alias IPv6 Address	When you configure redundancy with active and standby devices, you can configure a VLAN interface that has an alias global IPv6 address that is shared between the active and standby devices. The alias IPv6 address serves as a shared gateway for the two ACEs in a redundant configuration. You can configure only one alias global IPv6 address on an interface. <p>To configure an IPv6 alias global address, enter a complete IPv6 address with a prefix of 2000::/3 to 3fff::/3. For example, enter 2001:DB8:1::0.</p>
	 <p><b>Note</b> You must configure redundancy (fault tolerance) on the ACE for the alias global IPv6 address to work.</p>

Table 12-3 BVI Interface Attributes (continued)



Field	Description
Peer IPv6 Address	<p>To configure an IPv6 peer global address, enter a complete IPv6 address with a prefix of 2000::/3 to 3fff::/3. For example, enter 2001:DB8:1::0.</p> <p>Check the EUI-64 box to specify that the low order 64 bits are automatically generated in the IEEE 64-bit Extended Unique Identifier (EUI-64) format specified in RFC 2373. To use EUI-64, the Prefix Length field must be less than or equal to 64 and the host segment must be all zeros.</p> <p></p> <p><b>Note</b> The IPv6 peer global address must be unique across multiple contexts on a shared VLAN.</p>
Prefix Length	<p>Enter the prefix length for all global addresses to specify how many of the most significant bits (MSBs) are used for the network identifier. Enter an integer from 1 to 128. If you use the optional EUI-64 check box for the global and peer addresses, the prefix must be less than or equal to 64.</p>
IPv6 Unique-Local Address	<p>Field that appears only for ACE module and ACE appliance software Version A5(1.0) or later. A unique local address is an optional IPv6 unicast address that is used for local communication within an organization and it is similar to a private IPv4 address (for example, 10.10.2.1). Unique local addresses have a global scope, but they are not routable on the internet, and they are assigned by a central authority. All unique local addresses have a predefined prefix of FC00::/7. You can configure only one IPv6 unique local address on an interface.</p>
IPv6 Address	<p>To configure a unique local address, enter a complete IPv6 address with an FC00::/7 prefix in the first field. In the second field after the /, enter the prefix length to specify how many of the most significant bits (MSBs) are used for the network identifier.</p> <p>Check the EUI-64 box to specify that the low order 64 bits are automatically generated in the IEEE 64-bit Extended Unique Identifier (EUI-64) format specified in RFC 2373. To use EUI-64, the Prefix Length field must be less than or equal to 64 and the host segment must be all zeros.</p>
Peer IPv6 Address	<p>In a redundant configuration, you can configure an IPv6 peer unique local address on the active that is synchronized to the standby ACE. You can configure only one peer unique local IPv6 address on an interface.</p> <p>To configure a peer unique local address, enter a complete IPv6 address with an FC00::/7 prefix in the first field. In the second field after the /, enter the prefix length to specify how many of the most significant bits (MSBs) are used for the network identifier.</p> <p></p> <p><b>Note</b> The IPv6 peer unique local address must be unique across multiple contexts on a shared VLAN.</p> <p>Check the EUI-64 box to specify that the low order 64 bits are automatically generated in the IEEE 64-bit Extended Unique Identifier (EUI-64) format specified in RFC 2373. To use EUI-64, the Prefix Length field must be less than or equal to 64 and the host segment must be all zeros.</p>
Prefix Length	<p>Enter the prefix length for all global addresses to specify how many of the most significant bits (MSBs) are used for the network identifier. Enter an integer from 1 to 128. If you use the optional EUI-64 check box for the global and peer addresses, the prefix must be less than or equal to 64.</p>

Table 12-3 BVI Interface Attributes (continued)


Field	Description
IPv6 Link-Local Address	<p>Field that appears only for ACE module and ACE appliance software Version A5(1.0) or later. By default, when you enable IPv6 or configure any other valid IPv6 address on an interface, the ACE automatically creates a link local address for it. Every link local address must have a predefined prefix of FE80::/10. You can configure only one IPv6 link local address on an interface. This address always has the prefix of 64.</p> <p>To manually configure the link local address, enter a complete IPv6 address with an FE80::/10 prefix in this field. For example, enter FE80:DB8:1::1</p>
IPv6 Peer Link-Local Address	<p>Field that appears only for ACE module and ACE appliance software Version A5(1.0) or later. In a redundant configuration, you can configure an IPv6 peer link local address for the standby ACE. You can configure only one peer link local address on an interface.</p> <p>To configure the peer link local address, enter a complete IPv6 address with an FE80::/10 prefix in this field.</p> <p> <b>Note</b> The IPv6 peer link local address must be unique across multiple contexts on a shared VLAN.</p>
<b>More Settings</b> (The More Settings option appears only for ACE module and ACE appliance software Version A5(1.0) or later.)	
Managed-Config	<p>Check box to indicate that the interface use the stateful autoconfiguration mechanism to configure IPv6 addresses.</p> <p>Uncheck the check box to indicate that the interface does not use the stateful autoconfiguration mechanism to configure IPv6 addresses.</p>
Other-Config	<p>Check box to indicate that the interface use the stateful autoconfiguration mechanism to configure parameters other than IPv6 addresses.</p> <p>Clear the check box to indicate that the interface does not use the stateful autoconfiguration mechanism to configure parameters other than IPv6 addresses.</p>
NS Interval	<p>The ACE sends neighbor solicitation messages through ICMPv6 on the local link to determine the IPv6 addresses of nearby nodes (hosts or routers). You can configure the rate at which the ACE sends these neighbor solicitation messages.</p> <p>By default, the interval at which the ACE sends NS messages for DAD default is 1000 milliseconds (msecs). To configure the interval, enter an integer from 1000 to 2147483647.</p>
NS Reachable Time	<p>The neighbor solicitation reachable time is the time period in milliseconds during which a host considers the peer is reachable after a reachability confirmation from the peer. A reachability confirmation can include neighbor solicitation or advertisement, or any upper protocol traffic.</p> <p>By default, this time period is 0 milliseconds. To configure this time, enter an integer from 0 to 3600000.</p>
Retransmission time	<p>By default, the advertised retransmission time is 0 milliseconds.</p> <p>To configure the retransmission time, enter an integer from 0 to 3600000.</p>
DAD Attempts	<p>By default, the number of attempts for sending duplicate address detection (DAD) is 1.</p> <p>To configure the DAD attempts, enter an integer from 0 to 255.</p>
RA Hop Limit	<p>By default, the hop limit that neighbors should use when originating IPv6 packets is 64. To configure the hop limit in the IPv6 header, enter an integer from 0 to 255.</p>

Table 12-3 BVI Interface Attributes (continued)

Field	Description
RA Lifetime	<p>The RA lifetime is the length of time that neighboring nodes should consider the ACE as the default router before they send RS messages again.</p> <p>By default, this length of time is 1800 seconds (30 minutes). To configure the RA lifetime, enter an integer from 0 to 9000.</p>
RA Interval	<p>By default, the rate at which the ACE sends RA messages is 600 seconds. To configure the rate, enter an integer from 4 to 1800.</p>
Suppress RA	<p>By default, the ACE automatically responds to RS messages that it receives from neighbors with RA messages that include, for example, the network prefix. You can instruct the ACE to not respond to RS messages.</p> <p>Check the check box to instruct the ACE to not respond to RS messages.</p> <p>Clear the check box to reset the default behavior of automatically responding to RS messages.</p>
IPv6 Router Advertisement Settings	<p>Click the <b>Add</b> button to configure the IPv6 prefixes that the ACE advertises in RA messages on the local link.</p>
IPv6 Address/Prefix Length	<p>To configure IPv6 address advertised in the RA messages, enter a complete IPv6 address in the first field. In the second field after the /, enter the prefix length to specify how many of the most significant bits (MSBs) are used for the network identifier.</p>
No Advertisements	<p>Check the check box to indicate that the route prefix is not advertised.</p> <p>Clear the check box to indicate that the route prefix is advertised.</p>
Lifetime	<p>Configure the prefix lifetime attributes as follows:</p> <ul style="list-style-type: none"> <li>• Lifetime Duration: <ul style="list-style-type: none"> <li>– Valid Lifetime—By default, the prefix lifetime is 2592000 seconds (30 days). To configure the prefix lifetime in seconds, enter an integer from 0 to 2147183647.</li> <li>Select Infinite to indicate that the prefix never expires.</li> <li>– Preferred Lifetime—By default, the prefix lifetime is 604800 seconds (10 days). To configure how long an IPv6 address remains preferred in seconds, enter an integer from 0 to 2147183647. This lifetime must not exceed the Valid Lifetime.</li> <li>Select Infinite to indicate that the preferred lifetime never expires.</li> </ul> </li> <li>• Lifetime Expiration Date: <ul style="list-style-type: none"> <li>– Valid Month/Day/Year/Time—Valid lifetime expiration date and time.</li> <li>– Preferred Month/Day/Year/Time—Preferred lifetime expiration date and time.</li> </ul> </li> </ul> <p>Use the drop-down lists to select a day, month, and year. To specify the time, use the hh:mm format.</p>



**Table 12-3** BVI Interface Attributes (continued)

Field	Description
Off-link	This option appears when you enter a Preferred Lifetime field. Check this check box to indicate that the route prefix is on a different subnet for a router to route to it. Clear the check box to indicate that the route prefix is on the same subnet for a router to route to it.
No-autoconfig	This option appears when you enter a Preferred Lifetime field. Check this check box to indicate to the host that it cannot use this prefix when creating an stateless IPv6 address. Clear the check box to indicate to the host that it can use this prefix when creating an stateless IPv6 address.

**Step 5** Do one of the following:

- Click **Deploy Now** to immediately deploy this configuration on the ACE and save your entries to the running-configuration and startup-configuration files. This option appears for virtual contexts.
- Click **Cancel** to exit this procedure without saving your entries and to return to the previous table.

**Step 6** To display statistics and status information for a BVI interface, choose the BVI interface from the BVI Interface table, and click **Details**.

The **show interface bvi**, **show ipv6 interface bvi**, and **show ipv6 neighbors** CLI command outputs appears. IPv6 commands requires ACE module and ACE appliance software Version A5(1.0) or later. See the “[Displaying BVI Interface Statistics and Status Information](#)” section on page 12-26 for details.

#### Related Topics

- [Configuring Network Access, page 12-1](#)
- [Configuring Virtual Context Primary Attributes, page 6-14](#)

## Displaying All BVI Interfaces by Context

You can display all of the BVI interfaces associated with a specific context by choosing **Config > Devices > context > Network > BVI Interfaces**.

The BVI Interface table appears with the information shown in [Table 12-4](#).

**Table 12-4** BVI Interface Fields

Field	Description
BVI	Name of the BVI interface.
Description	Description for the BVI interface.
IP Address	IP address assigned to this interface including the netmask for an IPv4 address or a prefix length for an IPv6 address. IPv6 requires ACE module and ACE appliance software Version A5(1.0) or later.
IPv6 Config Status	The status whether IPv6 is enabled or disabled on the interface. IPv6 requires ACE module and ACE appliance software Version A5(1.0) or later.

Table 12-4 BVI Interface Fields (continued)

Field	Description
Admin Status	Status of the interface, which can be Up or Down.
Operational Status	Operational state of the device (Up or Down).
Last Polled	Date and time of the last time that ANM polled the device to display the current values.

**Related Topics**

- [Displaying BVI Interface Statistics and Status Information, page 12-26](#)

## Displaying BVI Interface Statistics and Status Information

You can display statistics and status information for a particular BVI interface by using the **Details** button.

**Procedure**

- 
- Step 1** Choose **Config > Devices > context > Network > BVI Interfaces**.
- The BVI Interface table appears.
- Step 2** In the BVI Interface table, choose a BVI interface from the BVI Interface table, and click **Details**.
- The **show interface bvi**, **show ipv6 interface bvi**, and **show ipv6 neighbors** CLI command outputs appear. The IPv6 commands require ACE module and ACE appliance software Version A5(1.0) or later. For details about the displayed output fields, see either the *Cisco ACE Module Routing and Bridging Configuration Guide* or the *Cisco ACE 4700 Series Appliance Routing and Bridging Configuration Guide*.
- Step 3** Click **Update Details** to refresh the command output.
- Step 4** Click **Close** to return to the BVI Interface table.
- 

**Related Topics**

- [Displaying All BVI Interfaces by Context, page 12-25](#)

## Configuring VLAN Interface NAT Pools and Displaying NAT Utilization

You can configure Network Address Translation (NAT) pools for a VLAN interface. NAT is designed to simplify and conserve IP addresses. It allows private IP networks that use unregistered IP addresses to connect to the Internet. NAT operates on a router, usually connecting two networks, and translates the private (not globally unique) addresses in the internal network into legal addresses before the packets are forwarded to another network.

In addition to creating a NAT pool, you can display the utilization information associated with it.

This section includes the following topics:

- [Configuring VLAN Interface NAT Pools, page 12-27](#)
- [Displaying NAT Pool Utilization, page 12-28](#)

## Configuring VLAN Interface NAT Pools

This procedure shows how to configure NAT pools for a VLAN interface.

### Guidelines and Restrictions

- The ACE allows you to configure NAT so that it advertises only one address for the entire network to the outside world. This feature, which effectively hides the entire internal network behind that address, offers both security and address conservation.
- Several internal addresses can be translated to only one or a few external addresses by using Port Address Translation (PAT) in conjunction with NAT. With PAT, you can configure static address translations at the port level and use the remainder of the IP address for other translations. PAT effectively extends NAT from one-to-one to many-to-one by associating the source port with each flow.
- The options that appear when you choose **Config > Devices > context** depend on the device associated with the virtual context and the role associated with your account.
- Fields and information related to IPv6 require ACE module and ACE appliance software Version A5(1.0) or later.
- When server load balancing is IPv6 to IPv4 or IPv4 to IPv6, you must configure source NAT.

### Prerequisites

You have configured at least one VLAN interface (see the [“Configuring Virtual Context VLAN Interfaces”](#) section on page 12-6).

### Procedure

---

**Step 1** Choose **Config > Devices > context > Network > NAT Pools**.

The NAT Pools table appears.

**Step 2** In the NAT Pools table, click **Add** to add a new NAT pool, or choose an existing NAT pool and click **Edit** to modify it.



---

**Note** If you click **Edit**, not all of the fields can be modified.

---

**Step 3** Choose the VLAN interface that you want to configure a NAT pool for and click the **NAT Pool** tab. The NAT Pool configuration table appears.

**Step 4** In the NAT Pool configuration table, click **Add** to add a new entry.

**Step 5** In the VLAN ID field, from the drop-down list, choose a VLAN entry.

**Step 6** In the NAT Pool ID field, either accept the automatically incremented entry or enter a new number to uniquely identify this pool.

Valid entries are from 1 to 2147483647.

**Step 7** In the IP Address Type field, choose either **IPv4** or **IPv6**.

This field appears only for ACE module and ACE appliance software Version A5(1.0) or later, which supports IPv4 and IPv6.

- Step 8** In the Start IP Address field, enter an IP address for the selected IP Address Type.
- This entry identifies either a single IP address or, if using a range of IP addresses, the first IP address in a range of global addresses for this NAT pool.
- Step 9** In the End IP Address field, enter the highest IP address in a range of global IP addresses for this NAT pool.
- Enter the IP address for the selected IP Address Type. Leave this field blank if you want to identify only the single IP address in the Start IP Address field.
- Step 10** Depending on the IP address type that you chose, do one of the following:
- For IPv4, in the Netmask field, choose the subnet mask for the global IP addresses in the NAT pool.
  - For IPv6, in the Prefix Length field, enter the prefix length for the global IP addresses in the NAT pool.
- Step 11** Check the PAT Enabled check box to instruct the ACE to perform port address translation (PAT) in addition to NAT.
- Uncheck the check box to indicate that the ACE is not to perform port address translation (PAT) in addition to NAT.
- Step 12** Do one of the following:
- Click **Deploy Now** to immediately deploy this configuration on the ACE and save your entries to the running-configuration and startup-configuration files. This option appears for virtual contexts.
  - Click **Cancel** to exit this procedure without saving your entries and to return to the NAT Pools table.
  - Click **Next** to deploy your entries and to add another NAT Pool entry.

---

#### Related Topics

- [Configuring VLAN Interface NAT Pools and Displaying NAT Utilization, page 12-26](#)
- [Displaying NAT Pool Utilization, page 12-28](#)
- [Configuring Virtual Context VLAN Interfaces, page 12-6](#)
- [Configuring Virtual Context BVI Interfaces, page 12-19](#)

## Displaying NAT Pool Utilization

This procedure shows how to display the utilization of all configured NAT pools on all VLANs.

#### Procedure

- 
- Step 1** Select **Config > Virtual Contexts > *virtual\_context* > Network > NAT Pools**.
- The NAT Pools table appears.
- Step 2** Click **Show NAT Pool Utilization**.
- The **show nat-fabric nat-pool-utilization** command pop up window appears, displaying the following information:
- Pool ID—Unique NAT pool identifier.

- NP—ACE network processor to which the NAT is bound.
- Total/Usage/Utilization (%):
  - Total—Number of IP addresses configured in the NAT pool.
  - Usage—Number of IP addresses being used.
  - Utilization (%)—Percentage of configured IP addresses be used.
- LowerIP/UpperIP—Lower and upper IP addresses configured in the NAT pool IP address range.
- Context—Context to which the NAT pool belongs.

**Step 3** From the pop up window, do one of the following:

- Click **Update Details** to refresh the information displayed.
- Click **Close** to close the pop up window.

#### Related Topics

- [Configuring VLAN Interface NAT Pools, page 12-27](#)

## Configuring Virtual Context Static Routes

You can configure context static routes. Admin and user context modes do not support dynamic routing, therefore you must use static routes for any networks to which the ACE is not directly connected, such as when there is a router between a network and the ACE.



#### Note

Fields and information related to IPv6 require ACE module and ACE appliance software Version A5(1.0) or later.

#### Procedure

**Step 1** Choose **Config > Devices > context > Network > Static Routes**.

The Static Routes configuration table appears and displays the following information:

- Destination prefix
- Destination prefix mask
- Next hop IP address

**Step 2** In the Static Routes configuration table, click **Add** to add a new static route.



#### Note

You cannot modify an existing static route. To make changes to an existing static route, you must delete the static route and then add it back.

**Step 3** In the IP Address Type, choose either **IPv4** or **IPv6** for the route.

This field appears only for ACE module and ACE appliance software Version A5(1.0) or later, which supports IPv4 or IPv6.

**Step 4** In the Destination Prefix field, enter the IP address based on the address type (IPv4 or IPv6) for the route.

The address that you specify for the static route is the address that is in the packet before entering the ACE and performing network address translation.

- Step 5** Depending on the IP address type that you chose, do one of the following:
- For IPv4, in the Destination Prefix Mask field, choose the subnet to use for this route.
  - For IPv6, in the Destination Prefix-length field, enter the prefix length from 0 to 128 to use for this route.

- Step 6** (IPv6 IP address type only) For the Forward Interface Type, choose one of the following:
- N/A (Not applicable)
  - VLAN
  - BVI

If you select VLAN or BVI, select its number from the drop down menu. To configure an interface, click Plus. After configuring it, select its number from the drop down menu.

- Step 7** In the Next Hop field, enter the IP address of the gateway router based on the address type (IPv4 or IPv6) for this route.

The gateway address must be in the same network as a VLAN interface for this context.

- Step 8** Do one of the following:
- Click **Deploy Now** to immediately deploy this configuration on the ACE and save your entries to the running-configuration and startup-configuration files. This option appears for virtual contexts.
  - Click **Cancel** to exit this procedure without saving your entries and to return to the previous table.
  - Click **Next** to deploy your entries and to add another static route.

#### Related Topics

- [Configuring Virtual Contexts, page 6-8](#)
- [Configuring Virtual Context Primary Attributes, page 6-14](#)

## Configuring Global IP DHCP

You can configure the Dynamic Host Configuration (DHCP) relay agent at the context level so the configuration applies to all interfaces associated with the context. When you configure the ACE as a DHCP relay agent, it is responsible for forwarding the requests and responses that are negotiated between the DHCP clients and the server. By default, the DHCP relay agent is disabled. You must configure a DHCP server when you enable the DHCP relay agent.



#### Note

The options that appear when you choose **Config > Devices > context** depend on the device associated with the virtual context and the role associated with your account.



#### Note

Fields and information related to IPv6 require ACE module and ACE appliance software Version A5(1.0) or later.

### Procedure

---

- Step 1** Choose **Config > Devices > context > Network > Global IP DHCP**.  
The Global IP DHCP configuration table appears.
- Step 2** From the Global IP DHCP configuration table, in the **Enable DHCP Relay For The Context** field, click **IPv4**, **IPv6**, or both to enable DHCP relay for the context and all interfaces associated with this context.  
For ACE module and ACE appliance software versions earlier than A5(1.0), this field does not include the IP version number and is for IPv4 only.
- Step 3** In the Relay Agent Information Reforwarding Policy field, choose a relay agent information forwarding policy:
- **N/A**—Specifies to not configure the DHCP relay to identify what is to be performed if a forwarded message already contains relay information.
  - **Keep**—Specifies that existing information is left unchanged on the DHCP relay agent.
  - **Replace**—Specifies that existing information is overwritten on the DHCP relay agent.
- Step 4** In the IP DHCP Server field, choose the IP DHCP server to which the DHCP relay agent is to forward client requests.
- Step 5** In the IPv6 Forward Interface VLAN field, you can optionally enter the VLAN interface number that you configured in the [IPv6 DHCP Forward Interface VLAN](#) field on the interface where the multicast DHCP relay message is sent.  
This field appears only for ACE module and ACE appliance software Version A5(1.0) or later.
- Step 6** In the IPv6 DHCP server, specify one or more IP DHCP servers and IPv6 addresses to which the DHCP relay agent is to forward client requests.  
This field appears only for ACE module and ACE appliance software Version A5(1.0) or later.
- Step 7** Do one of the following:
- Click **Deploy Now** to immediately deploy this configuration on the ACE and save your entries to the running-configuration and startup-configuration files. This option appears for virtual contexts.
  - Click **Cancel** to exit this procedure without saving your entries and to return to the previous table.
  - Click **Next** to deploy your entries and to add another DHCP relay entry.
-

# Configuring Static VLANs for Over 8000 Static NAT Configurations



## Note

This feature applies to ACE modules only and was deprecated beginning with ACE software Version A5(1.0).

You can create more than 8,000 static NAT configurations (one static NAT configuration with a netmask is counted as one configuration). In addition, follow these restrictions and guidelines when using this feature:

- This feature is supported in routed mode only.
- Only one mapped interface is allowed per virtual context. However, each static NAT configuration must have a different mapped IP address.
- At any point, you can configure no more than one next-hop on the mapped interface.
- Bidirectional NAT, or in other words, source-address as well as destination-address translation, for the same flow is not supported.
- You must have fewer than 1,000 real IP addresses on the same subnet as the real interface. In addition, you must have fewer than 1,000 mapped IP address on the same subnet as the mapped interface.
- If you use this feature, we recommend that you do not use MP-based NAT for the same virtual context.

## Procedure

- 
- Step 1** Choose **Config > Devices > context > Network > Static NAT Overwrite**.  
The Static NAT Overwrite configuration table appears.
- Step 2** In the Static NAT Overwrite configuration table, click **Add** to add a new static NAT.
- Step 3** In the Mapped IP Address field, enter the IP address to which the real IP address is translated.  
In a context, the mapped IP address must be different in each static NAT configuration.
- Step 4** In the Real VLAN Number field, choose the VLAN number of the interface connected to the real IP address network.  
The list of available real VLANs includes routed mode VLANs only (for more information, see [Interface Type](#)).
- Step 5** In the Mapped VLAN Number field, choose the VLAN number of the interface connected to the mapped IP address network.  
The list of available mapped VLANs includes routed mode VLANs only (for more information, see [Interface Type](#)). In a context, the mapped interface must be the same in each static NAT configuration.
- Step 6** In the Real IP Address field, enter the real server IP address to be translated.  
In a context, you must configure a different address for configurations that have the same real server interface.
- Step 7** In the Real IP Netmask field, choose the subnet mask for the real server address.
- Step 8** Do one of the following:



- Click **Deploy Now** to immediately deploy this configuration on the ACE and save your entries to the running-configuration and startup-configuration files. This option appears for virtual contexts.
- Click **Cancel** to exit this procedure without saving your entries and to return to the previous table.
- Click **Next** to deploy your entries and to add another DHCP relay entry.

## Configuring Gigabit Ethernet Interfaces on the ACE Appliance



### Note

This feature is for ACE appliances only.

You can configure a Gigabit Ethernet interface on the ACE appliance, which provides physical Ethernet ports to connect servers, PCs, routers, and other devices to the ACE appliance. The ACE appliance supports four Layer 2 Ethernet ports for performing Layer 2 switching. You can configure the four Ethernet ports to provide an interface for connecting to 10-Mbps, 100-Mbps, or 1000-Mbps networks. Each Layer 2 Ethernet port supports autonegotiate, full-duplex, or half-duplex operation on an Ethernet LAN, and can carry traffic within a designated VLAN.

A Layer 2 Ethernet port can be configured as follows:

- **Member of Port-Channel Group**—The port is configured as a member of a port-channel group, which associates a physical port on the ACE appliance to a logical port to create a port-channel logical interface. The VLAN association is derived from port-channel configuration. The port is configured as a Layer 2 EtherChannel, where each EtherChannel bundles the individual physical Ethernet data ports into a single logical link that provides the aggregate bandwidth of up to four physical links on the ACE.
- **Access VLAN**—The port is assigned to a single VLAN. This port is referred to as an access port and provides a connection for end users or node devices, such as a router or server.
- **Trunk port**—The port is associated with IEEE 802.1Q encapsulation-based VLAN trunking to allocate VLANs to ports and to pass VLAN information (including VLAN identification) between switches for all Ethernet channels defined in a Layer 2 Ethernet data port or a Layer 2 EtherChannel (port-channel) group on the ACE appliance.

This section includes the following topics:

- [Configuring Gigabit Ethernet Interfaces, page 12-33](#)
- [Displaying Gigabit Ethernet Interface Statistics and Status Information, page 12-36](#)

## Configuring Gigabit Ethernet Interfaces

This section describes how to configure Gigabit Interfaces on the ACE.

### Procedure

- Step 1** Choose **Config > Devices > context > Network > GigabitEthernet Interfaces**.  
The GigabitEthernet Interfaces table appears.
- Step 2** In the GigabitEthernet Interfaces table, click **Poll Now** to instruct ANM to poll the devices and display the current values, and click **OK** when prompted to poll the devices for data.



**Step 3** Choose an existing gigabit Ethernet interface, and click **Edit** to modify it.

**Step 4** Enter the gigabit Ethernet physical interface attributes (see [Table 12-5](#)).

**Table 12-5** *Physical Interface Attributes*

Field	Description
Interface Name	Name of the Gigabit Ethernet interface, which is in the format <i>slot_number/port_number</i> where <i>slot_number</i> is the physical slot on the ACE for the specified port, and <i>port_number</i> is the physical Ethernet data port on the ACE for the specified port.
Description	Brief description for this interface.
Admin Status	Administrative state of the interface: <b>Up</b> or <b>Down</b> .
Speed	Port speed: <ul style="list-style-type: none"> <li>• <b>Auto</b>—Autonegotiate with other devices</li> <li>• <b>10 Mbps</b></li> <li>• <b>100 Mbps</b></li> <li>• <b>1000 Mbps</b></li> </ul>
Duplex	Interface duplex mode: <ul style="list-style-type: none"> <li>• <b>Auto</b>—Resets the specified Ethernet port to automatically negotiate port speed and duplex of incoming signals. This is the default setting.</li> <li>• <b>Full</b>—Configures the specified Ethernet port for full-duplex operation, which allows data to travel in both directions at the same time.</li> <li>• <b>Half</b>—Configures the specified Ethernet port for half-duplex operation. A half-duplex setting ensures that data only travels in one direction at any given time.</li> </ul>
Port Operation Mode	Port operation mode: <ul style="list-style-type: none"> <li>• <b>N/A</b>—Specifies that this option is not to be used.</li> <li>• <b>Channel Group</b>—Specifies to map the port to a port channel. You must specify: <ul style="list-style-type: none"> <li>• Port Channel Group Number—Specifies the port channel group number.</li> <li>• HA VLAN—Specifies the high availability (HA) VLAN used for communication between the members of the FT group.</li> </ul> </li> <li>• <b>Switch Port</b>—Specifies the interface switch port type: <ul style="list-style-type: none"> <li>• Access—Specifies that the port interface is an access port. You must specify a VLAN as an access port in the Access VLAN field.</li> <li>• Trunk—Specifies that the port interface is a trunk port. When you choose Trunk, you must complete one or both of the following fields: <ul style="list-style-type: none"> <li>- Trunk Native VLAN—Identifies the 802.1Q native VLAN for a trunk.</li> <li>- Trunk Allowed VLANs—Selectively allocates individual VLANs to a trunk link.</li> </ul> </li> </ul> </li> </ul>
HA LAN	High availability (HA) VLAN used for communication between the members of the FT group.

Table 12-5 Physical Interface Attributes (continued)

Field	Description
Carrier Delay	<p>Configurable delay at the physical port level to address any issues with transition time, based on the variety of peers. Valid values are from 0 to 120 seconds. The default is 0 (no carrier delay).</p> <p> <b>Note</b> If you connect an ACE to a Catalyst 6500 series switch, your configuration on the switch may include the Spanning-Tree Protocol (STP). However, the ACE does not support STP. In this case, you may find that the Layer 2 convergence time is much longer than the physical port up time. For example, the physical port would normally be up within 3 seconds, but STP moving to the forward state may need approximately 30 seconds. During this transitional time, although the ACE declares the port to be up, the traffic does not pass. In this case, you should specify a carrier delay.</p>
QoS Trust COS	<p>Quality of Service (QoS) for the physical Ethernet port. By default, QoS is disabled for each physical Ethernet port on the ACE.</p> <p>QoS for a configured physical Ethernet port is based on VLAN Class of Service (CoS) bits (priority bits that segment the traffic in eight different classes of service). When you enable QoS on a port (a trusted port), traffic is mapped into different ingress queues based on their VLAN CoS bits. If there are no VLAN CoS bits, or QoS is not enabled on the port (untrusted port), the traffic is then mapped into the lowest priority queue.</p> <p>You can enable QoS for an Ethernet port configured for fault tolerance. In this case, heartbeat packets are always tagged with CoS bits set to 7 (a weight of High).</p> <p> <b>Note</b> We recommend that you enable QoS on the FT VLAN port to provide higher priority for FT traffic.</p>

**Step 5** Do one of the following:

- Click **Deploy Now** to immediately deploy this configuration on the ACE and save your entries to the running-configuration and startup-configuration files.
- Click **Cancel** to exit the procedure without saving your changes and to return to the Physical Interface table.
- Click **Next** or **Previous** to go to the next or previous physical channel.
- Click **Delete** to remove this entry from the Physical Interface table and to return to the table.

**Step 6** (Optional) To display statistics and status information for a particular Gigabit Ethernet interface, choose the interface from the GigabitEthernet Interfaces table, and click **Details**.

The **show interface gigabitEthernet** CLI command output appears. See the “[Displaying Gigabit Ethernet Interface Statistics and Status Information](#)” section on page 12-36 for details.

**Related Topics**

- [Configuring Virtual Context VLAN Interfaces](#), page 12-6
- [Configuring Virtual Context BVI Interfaces](#), page 12-19
- [Configuring Port-Channel Interfaces for the ACE Appliance](#), page 12-36

## Displaying Gigabit Ethernet Interface Statistics and Status Information

You can display statistics and status information for a particular Gigabit Ethernet interface.

### Procedure

- 
- Step 1** Choose **Config > Devices > context > Network > GigabitEthernet Interfaces**.
- The GigabitEthernet Interfaces table appears.
- Step 2** In the GigabitEthernet Interfaces table, choose a Gigabit Ethernet interface from the GigabitEthernet Interfaces table, and click **Details**.
- The **show interface gigabitEthernet** CLI command output appears. For details on the displayed output fields, see the *Cisco ACE 4700 Series Appliance Routing and Bridging Configuration Guide*.
- Step 3** (Optional) Click **Update Details** to refresh the display.
- Step 4** Click **Close** to return to the GigabitEthernet Interfaces table.
- 

### Related Topics

[Configuring Gigabit Ethernet Interfaces on the ACE Appliance, page 12-33](#)

## Configuring Port-Channel Interfaces for the ACE Appliance

This section discusses how to configure port channel interfaces for the ACE appliance. It consists of the following topics:

- [Why Use Port Channels?, page 12-36](#)
- [Configuring a Port-Channel Interface, page 12-37](#)
- [Configuring a Catalyst 6500 Series Switch for an ACE Appliance Port-Channel Interface Connection, page 12-39](#)
- [Displaying Port Channel Interface Statistics and Status Information, page 12-41](#)

## Why Use Port Channels?

A port channel groups multiple physical ports into a single logical port. This is also called port aggregation or channel aggregation. A port channel containing multiple physical ports has several advantages:


- Improves link reliability through physical redundancy.
- Allows greater total throughput to the ACE appliance. For example, four 1-Gigabit Ethernet interfaces can be aggregated into a single 4-Gigabit channel.
- Allows traffic capacity to be scaled up in the future, without network disruption at that time. A port channel can do everything a switched port can do, but a switched port cannot do everything a port channel can do. We recommend that you use a port channel.
- Provides maximum flexibility of network configuration and focuses network configuration on VLANs rather than physical cabling.

The disadvantage of a port channel is that it requires additional configuration on the switch the ACE is connected to, as well as the ACE itself. There are many methods of port aggregation implemented by different switches, and not every method works with ACE. For an example of how to configure a Cisco Catalyst 6500 switch to enable a port channel connection to ACE, see the “[Configuring a Catalyst 6500 Series Switch for an ACE Appliance Port-Channel Interface Connection](#)” section on page 12-39.

Using a port channel also requires more detailed knowledge of your network's VLANs, because all “cabling” to and from the ACE will be handled over VLANs rather than using physical cables. Nonetheless, use of port channels is highly recommended, especially in a production deployment of ACE.

Figure 12-1 illustrates a port channel interface.

**Figure 12-1 Example of a Port Channel Interface**



#### Related Topics

[Configuring a Port-Channel Interface, page 12-37](#)

[Displaying Port Channel Interface Statistics and Status Information, page 12-41](#)

## Configuring a Port-Channel Interface



#### Note

This feature is for ACE appliances only.

You can group physical ports together on the ACE appliance to form a logical Layer 2 interface called the port channel. All the ports belonging to the same port channel must be configured with same values; for example, port parameters, VLAN membership, and trunk configuration. Only one port channel in a channel group is allowed, and a physical port can belong to a single port-channel interface only.

- 
- Step 1** Choose **Config > Devices > context > Network > Port Channel Interfaces**.  
The Port Channel Interface table appears.
- Step 2** In the Port Channel Interface table, click **Poll Now** to instruct ANM to poll the devices and display the current values, and click **OK** when prompted to poll the devices for data.
- Step 3** Click **Add** to add a port channel interface, or choose an existing port channel interface and click **Edit** to modify it.



**Note** If you click **Edit**, not all of the fields can be modified.

**Step 4** Enter the port channel interface attributes (see [Table 12-6](#)).

**Table 12-6** Port Channel Interface Attributes

Field	Description
Interface Number	Channel number for the port-channel interface, which can be from 1 to 255.
Description	Brief description for this interface.
Fault Tolerant VLAN	Fault tolerant (FT) VLAN used for communication between the members of the FT group.
Admin Status	Administrative state of the interface: <b>Up</b> or <b>Down</b> .
Load Balancing Method	Load balancing method: <ul style="list-style-type: none"> <li>• <b>Dst-IP</b>—Loads distribution on the destination IP address.</li> <li>• <b>Dst-MAC</b>—Loads distribution on the destination MAC address.</li> <li>• <b>Dst-Port</b>—Loads distribution on the destination TCP or UDP port.</li> <li>• <b>Src-Dst-IP</b>—Loads distribution on the source or destination IP address.</li> <li>• <b>Src-Dst-MAC</b>—Loads distribution on the source or destination MAC address.</li> <li>• <b>Src-Dst-Port</b>—Loads distribution on the source or destination port.</li> <li>• <b>Src-IP</b>—Loads distribution on the source IP address.</li> <li>• <b>Src-MAC</b>—Loads distribution on the source MAC address.</li> <li>• <b>Src-Port</b>—Loads distribution on the TCP or UDP source port.</li> </ul>
Switch Port Type	Interface switchport type: <ul style="list-style-type: none"> <li>• <b>N/A</b>—Indicates that the switchport type is not specified.</li> <li>• <b>Access</b>—Specifies that the port interface is an access port. You must specify a VLAN as an access port in the Access VLAN field.</li> <li>• <b>Trunk</b>—Specifies that the port interface is a trunk port. When you choose Trunk, you must complete the following fields: <ul style="list-style-type: none"> <li>– Trunk Native VLAN—Identifies the 802.1Q native VLAN for a trunk.</li> <li>– Trunk Allowed VLANs—Selectively allocate individual VLANs to a trunk link.</li> </ul> </li> </ul>

**Step 5** Do one of the following:

- Click **Deploy Now** to immediately deploy this configuration on the ACE and save your entries to the running-configuration and startup-configuration files.
- Click **Cancel** to exit the procedure without saving your changes and to return to the Port Channel Interface table.
- Click **Next** to deploy your entries and to add another port-channel interface.

**Step 6** (Optional) To display statistics and status information for a particular port-channel interface, choose the interface from the Port Channel Interfaces table, and click **Details**.

The **show interface port-channel** CLI command output appears. See the “[Displaying Port Channel Interface Statistics and Status Information](#)” section on page 12-41 for details.

#### Related Topics

- [Configuring Port-Channel Interfaces for the ACE Appliance, page 12-36](#)
- [Configuring Port-Channel Interfaces for the ACE Appliance, page 12-36](#)
- [Displaying Port Channel Interface Statistics and Status Information, page 12-41](#)
- [Configuring Virtual Context VLAN Interfaces, page 12-6](#)

## Configuring a Catalyst 6500 Series Switch for an ACE Appliance Port-Channel Interface Connection

This section provides information for you to configure a port-channel interface on a network device such as the Catalyst 6500 Series switch. After you configure the port channels for the ACE appliance through ANM and you physically connect the Gigabit Ethernet physical interfaces on the ACE appliance to the Catalyst 6500 Series switch ports, configure the port channels on the switch. The information outlined in this topic is intended as an example of configuring port channels on a switch. You can adapt this information for whatever switch the ACE appliance is connected to in your network.

For specific details on configuring the Catalyst 6500 Series switch, see the documentation set on [www.Cisco.com](http://www.Cisco.com).

This section includes the following topics:

- [Creating the Port Channel Interface on the Catalyst 6500](#)
- [Adding Interfaces to the Port Channel](#)

### Creating the Port Channel Interface on the Catalyst 6500

This section contains an example in which a Catalyst 6500 Series switch is configured with a port channel using an 802.1q trunk that allows the associated VLANs. The native VLAN of the trunk is VLAN 10.



#### Note

Default VLAN 1 should not be used for the native VLAN because this VLAN is used internally on the ACE appliance.

Port-channel load balancing is used to distribute the traffic load across each of the links in the port channel to ensure efficient utilization of each link. Port-channel load balancing on the Catalyst 6500 Series switch can use MAC addresses or IP addresses, Layer 4 port numbers, source addresses, destination addresses, or both source and destination addresses. By default, the ACE appliance uses Src-Dst-MAC to make a load balancing decision (see [Table 12-6](#)). We recommend that you use the source and destination Layer 4 port for the load-balancing decision.

The following example illustrates the CLI commands used to configure a port channel interface for the Catalyst 6500 Series switch:

```
Switch(config)# port-channel load-balance src-dst-port
Switch(config)# interface port-channel 1
Switch(config-if)# description For Connection with ACE Appliance
Switch(config-if)# switchport
Switch(config-if)# switchport mode trunk
Switch(config-if)# switchport trunk encapsulation dot1q
Switch(config-if)# switchport trunk native vlan 10
Switch(config-if)# switchport trunk allowed vlan 10,20,30,31, 40,50
Switch(config-if)# switchport nonegotiate
Switch(config-if)# mls qos trust cos
```

After you configure the port channel on the Catalyst 6500 Series switch, you can then add it to the configuration of the four interfaces as described in the [“Adding Interfaces to the Port Channel”](#) section on page 12-40.



**Note**

---

The ACE appliance does not support Port Aggregation Protocol (PAgP) or Link Aggregate Control Protocol (LACP) so the port-channel interface is configured using **mode on**.

---

## Adding Interfaces to the Port Channel

The following example illustrates the CLI commands used to configure the four switch ports 3/9 through 3/12 as members of the port channel on the Catalyst 6500 Series switch:

```
Switch(config-if)# int range Gig 3/9 - 12
Switch(config-if-range)# channel-group 1 mode on
Switch(config-if-range)# speed 1000
Switch(config-if-range)# spanning-tree portfast trunk
Switch(config-if-range)# no shut
```

On the ACE appliance, you can configure the Ethernet port speed for a setting of 10, 100, or 1000 Mbps by configuring the Speed field for a Gigabit Ethernet physical interface attributes (see [Table 12-5](#)). The default for the ACE appliance is the auto-negotiate interface speed. We recommend that you configure the speed to 1000 on both the Catalyst 6500 Series switch and the ACE appliance to avoid relying on auto negotiation of the interface speed. A speed setting of 1000 helps to avoid the possibility of the interface operating below the expected Gigabit speed and ensures that the port-channel interface reaches the maximum 4 Gbps throughput.

The ACE appliance does not implement Spanning-Tree protocol and does not take part in Spanning-Tree root bridge election process. PortFast is configured on the Catalyst 6500 Series switch to reduce the time required for spanning tree to allow traffic on the port connected to the ACE interface by immediately moving to the forwarding state, bypassing the block, listening, and learning states. The average time for switch port moving into a forward state is approximately 30 seconds. Using PortFast reduces this time to approximately 5 seconds.



**Note**

---

In virtual partitions operating in bridge mode, the ACE offers an option to bridge Spanning-Tree BPDUs between two VLANs to prevent the possibility of a loop. Such a loop may occur when two partitions actively forward traffic. This should not happen during normal operation; however, the option to bridge BPDUs provides a safeguard against this condition. Upon detecting BPDUs, the switch connected to the ACE appliance immediately blocks the port/VLAN from which the loop originated from. We recommend that you configure an ethertype ACL that includes the BPDU protocol and apply the ACL to Layer 2 interfaces in bridge mode.

---



## Displaying Port Channel Interface Statistics and Status Information

You can display statistics and status information for a particular port-channel interface.

### Procedure

- 
- Step 1** Choose **Config > Devices > context > Network > Port Channel Interfaces**.
- The Port Channel Interfaces table appears.
- Step 2** In the Port Channel Interfaces table, choose a port-channel interface from the Port Channel Interfaces table, and click **Details**.
- The **show interface port-channel** CLI command output appears. For details about the displayed output fields, see the *Cisco ACE 4700 Series Appliance Routing and Bridging Configuration Guide*.
- Step 3** (Optional) Click **Update Details** to refresh the display.
- Step 4** Click **Close** to return to the Port Channel Interfaces table.
- 

### Related Topics

[Configuring Port-Channel Interfaces for the ACE Appliance, page 12-36](#)

