CHAPTER 10

Configuring Network Access

This chapter describes how to configure network access. 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 bridge-group virtual interface (BVI) with the bridge group.

The ACE appliance also supports shared VLANs; 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.

Note When you use the ACE CLI to configure named objects (such as a real server, virtual server, parameter map, class map, health probe, and so on), consider that the Device Manager (DM) supports object names with an alphanumeric string of 1 to 64 characters, which can include the following special characters: underscore (_), hyphen (-), dot (.), and asterisk (*). Spaces are not allowed.

If you use the ACE CLI to configure a named object with special characters that the DM does not support, you may not be able to configure the ACE using DM.

This chapter contains the following sections:

Configuring Port Channel Interfaces

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

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.

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 10-1 illustrates a port channel interface.

Figure 10-1 Example of a Port Channel Interface

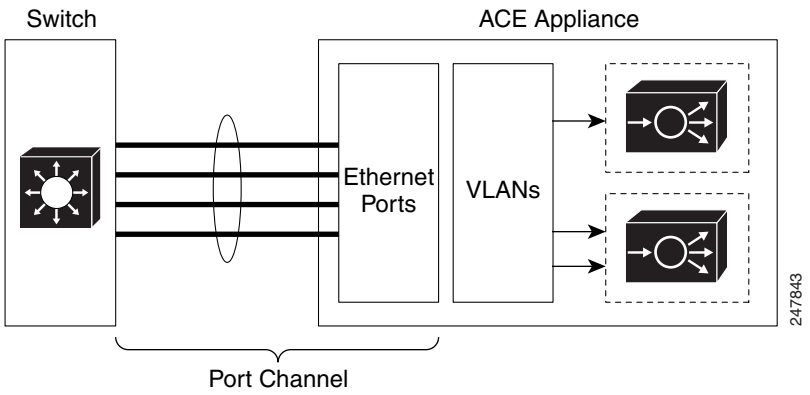

Related Topic

Configuring a Port-Channel Interface, page 10-3

Configuring a Port-Channel Interface

You can group physical ports together on the ACE 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 only to a single port-channel interface.

Step 1 Choose **Config > Virtual Contexts >** *context* **> Network > Port Channel Interfaces**. The Port Channel Interfaces table appears.

Step 2 Click **Add** to add a port channel interface, or select an existing port channel interface, and then click **Edit** to modify it.

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

Step 3 Enter the port channel interface attributes (see Table 10-1).

Table 10-1 Port Channel Interface Attributes

Field	Description
Interface Number	Specify a channel number for the port-channel interface, which can be from 1 to 255.
Description	Enter a brief description for this interface.
Fault Tolerance VLAN	Specify the fault tolerant (FT) VLAN used for communication between the members of the FT group
Admin Status	Indicate whether you want the interface to be Up or Down.
Load Balancing Method	Specify one of the following load balancing methods: • Dst-IP—Loads distribution on the destination IP address. • Dst-MAC—Loads distribution on the destination MAC address. • Dst-Port—Loads distribution on the destination TCP or UDP port. • Src-Dst-IP—Loads distribution on the source or destination IP address. • Src-Dst-MAC—Loads distribution on the source or destination MAC address. • Src-Dst-Port—Loads distribution on the source or destination port. • Src-IP—Loads distribution on the source IP address. • Src-MAC—Loads distribution on the source MAC address. • Src-Port—Loads distribution on the TCP or UDP source port.

Table 10-1 Port Channel Interface Attributes (continued)

Field	Description
Switch Port Type	Specify the interface switchport type: • N/A—Indicates that the switchport type is not specified. • Access—Specifies that the port interface is an access port. You must specify a VLAN as an access port in the Access VLAN field. • Trunk—Specifies that the port interface is a trunk port. When you select Trunk, you must complete one of the following fields: – Trunk Native VLAN—Identifies the 802.1Q native VLAN for a trunk. – Trunk Allowed VLANs—Selectively allocate individual VLANs to a trunk link.

Step 4 Do the following:

- Click **Deploy Now** to save your entries and to return to the Port Channel Interface table.
- Click **Cancel** to exit the procedure without saving your changes and to return to the Port Channel Interface table.
- Click **Next** to save your entries and to add another port-channel interface.

Step 5 (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 10-5 for details.

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 > Virtual Contexts >** *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 *Routing and Bridging Guide, Cisco ACE Application Control Engine*.

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 a Port-Channel Interface, page 10-3

Configuring Gigabit Ethernet Interfaces

The ACE appliance provides physical Ethernet ports to connect servers, PCs, routers, and other devices to the ACE. The ACE 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 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.

The following procedure describes how to configure a Gigabit Ethernet interface.

Procedure

Step 1 Choose **Config > Virtual Contexts >** *context* **> Network > Gigabit Ethernet Interfaces**. The GigabitEthernet Interfaces table appears.

Step 2 Select an existing Gigabit Ethernet interface, and then click **Edit** to modify it.

Step 3 Enter the Gigabit Ethernet physical interface attributes (see Table 10-2).

Table 10-2 Gigabit Ethernet Physical Interface Attributes

Field	Description
Interface Name	Name of the Gigabit interface, which is the *slot_number*/*port_number* where *slot_number* is the physical slot on the ACE for the specified port, and *port_number* is the physical Ethernet data port on the ACE for the specified port.
Description	Enter a brief description for this interface.
Admin Status	Indicate whether you want the interface to be Up or Down.
Speed	Specifies the port speed, which can be • Auto—Autonegotiate with other devices • 10 Mbps • 100 Mbps • 1000 Mbps
Duplex	Specifies an interface duplex mode, which can be: • Auto—Resets the specified Ethernet port to automatically negotiate port speed and duplex of incoming signals. This is the default setting. • Half—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. • Full—Configures the specified Ethernet port for full-duplex operation, which allows data to travel in both directions at the same time.

Table 10-2 Gigabit Ethernet Physical Interface Attributes (continued)

Field	Description
Port Operation Mode	Specifies the port operation mode, which can be: • N/A—Indicates that this option is not to be used. • Channel Group—Specifies to map the port to a port channel. You must specify – Port Channel Group Number—Specify the port channel group number – Fault Tolerant VLAN—Specify the fault tolerant (FT) VLAN used for communication between the members of the FT group. • Switch Port—Specifies the interface switchport type: – Access —Specifies that the port interface is an access port. You must specify a VLAN as an access port in the Access VLAN field. – Trunk—Specifies that the port interface is a trunk port. When you select Trunk, you must complete only one of the following fields: Trunk Native VLAN—Identifies the 802.1Q native VLAN for a trunk. Trunk Allowed VLANs—Selectively allocate individual VLANs to a trunk link.
Fault Tolerant VLAN	Specifies the fault tolerant (FT) VLAN used for communication between the members of the FT group.

Table 10-2 Gigabit Ethernet Physical Interface Attributes (continued)

Field	Description
Carrier Delay	Adds a configurable delay at the physical port level to address any issues with transition time, based on the variety of peers. Valid values are 0 to 120 seconds. The default is 0 (no carrier delay). **Note** If you connect an ACE to a Catalyst 6500 series switch, your configuration on the Catalyst 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 will not pass. In this case, specify a carrier delay.
QoS Trust COS	Enables Quality of Service (QoS) for the physical Ethernet port. By default, QoS is disabled for each physical Ethernet port on the ACE. QoS for a configured physical Ethernet port based on VLAN Classes 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. 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). **Note** We recommend that you enable QoS on the FT VLAN port to provide higher priority for FT traffic.

Step 4 Do the following:

- Click **Deploy Now** to save your entries and to return to the Physical Interface table.
- 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 5 (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 10-9 for details.

Related Topics

- Configuring Virtual Context VLAN Interfaces, page 10-10
- Configuring Virtual Context BVI Interfaces, page 10-23
- Configuring Virtual Context Static Routes, page 10-34

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 > Virtual Contexts >** *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 *Routing and Bridging Guide, Cisco ACE Application Control Engine*.

Step 3 (Optional) Click **Update Details** to refresh the display.

Step 4 Click **Close** to return to the GigabitEthernet Interfaces table.

Related Topic

- Configuring Gigabit Ethernet Interfaces, page 10-5

Configuring Virtual Context VLAN Interfaces

The ACE Appliance Device Manager uses class maps and policy maps to classify (filter) traffic and to direct it to different contexts. A virtual context uses VLANs to receive packets classified for that context.

Note When you create a new VLAN interface for a virtual context, you can configure one or more VLAN interfaces in any user context before you assign those VLAN interfaces to the associated user contexts in a virtual context through the Allocate-Interface VLANs field (see the "Creating Virtual Contexts" section on page 4-2).

Use this procedure to configure VLAN interfaces for virtual contexts.

Procedure

Step 1 To configure a virtual context, select **Config > Virtual Contexts >** *context* **> Network > VLAN Interfaces**. The VLAN Interface table appears.

Step 2 Click **Add** to add a new VLAN interface, or select an existing VLAN interface, and then click **Edit** to modify it.

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

Step 3 Enter the VLAN interface attributes (see Table 10-3). Click **More Settings** to access the additional VLAN interface attributes. By default, ACE appliance Device Manager hides the default VLAN interface attributes and the VLAN interface attributes which are not commonly used.

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

Table 10-3 VLAN Interface Attributes

Field	Description
VLAN	Either accept the automatically incremented entry or enter a different value. Valid entries are integers from 2 to 4094.
Description	Enter a brief description for this interface.

Table 10-3 VLAN Interface Attributes (continued)

Field	Description
Interface Type	Select the role of the virtual context in the network topology of the VLAN interface: • Routed—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. **Note** A routed VLAN interface can support both IPv4 and IPv6 addresses at the same time. • Bridged—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. • Unknown—Choose Unknown if you are unsure of the network topology of the VLAN interface.
IP Address	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. If this interface is only used for IPv6 traffic, entering an IPv4 address is optional.
Alias IP Address	Enter the IPv4 address of the alias this interface is associated with.
Peer IP Address	Enter the IPv4 address of the remote peer.
Netmask	Select the subnet mask to be used.
Admin Status	Indicate whether you want the interface to be Up or Down.
Enable MAC Sticky	Check the check box to indicate that the ACE appliance is to convert dynamic MAC addresses to sticky secure MAC addresses and add this information to the running configuration. Clear the check box to indicate that the ACE appliance is not to convert dynamic MAC addresses to sticky secure MAC addresses.

Table 10-3 VLAN Interface Attributes (continued)

Field	Description
Enable Normalization	Check the check boxes to indicate that normalization is to be enabled on this interface for IPv4, IPv6, or both. Clear the check box to indicate that normalization is to be disabled on this interface. ⚠ **Caution** Disabling normalization may expose your ACE appliance 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.
Enable IPv6	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: • Configures a link-local address (if not previously configured) • Performs duplicate address detection (DAD) Clear the check box to indicate that IPv6 is disabled on this interface.
IPv6 Global Address	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. When you configure a global IPv6 address on an interface, the ACE automatically does the following: • Configures a link-local address (if not previously configured) • Performs duplicate address detection (DAD) on both addresses
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. 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.

Table 10-3 VLAN Interface Attributes (continued)

Field	Description
Alias IPv6 Address	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. 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. **Note** You must configure redundancy (fault tolerance) on the ACE for the alias global IPv6 address to work.
Peer IPv6 Address	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. 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. **Note** The IPv6 peer global address must be unique across multiple contexts on a shared VLAN. 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.
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	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.

Table 10-3 VLAN Interface Attributes (continued)

Field	Description
IPv6 Address	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. 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.
IPv6 Peer Address	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. 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. **Note** The IPv6 peer unique local address must be unique across multiple contexts on a shared VLAN. 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.
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.
IPv6 Link-Local Address	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. 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.

Table 10-3 VLAN Interface Attributes (continued)

Field	Description
IPv6 Peer Link-Local Address	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. To configure the peer link local address, enter a complete IPv6 address with an FE80::/10 prefix in this field. **Note** The IPv6 peer link local address must be unique across multiple contexts on a shared VLAN.
More Settings	
Enable ICMP Guard	Check the IPv4, IPv6 or both check boxes to indicate that ICMP Guard is to be enabled on the ACE appliance. Clear the check boxes to indicate that ICMP Guard is not to be enabled on ACE appliance. **Caution** Disabling ICMP security checks may expose your ACE appliance 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.
Enable DHCP Relay	Check the IPv4, IPv6 or both check boxes to indicate that the ACE appliance is to accept DHCP requests from clients on this interface and to enable the DHCP relay agent. Clear the check boxes to indicate that the ACE appliance is not to accept DHCP requests or enable the DHCP relay agent.
Reverse Path Forwarding (RPF)	Check the IPv4, IPv6 or both check boxes to indicate that the ACE appliance 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. Clear the check boxes to indicate that the ACE appliance is not to filter or discard packets based on the ability to verify the source IP address.
Reassembly Timeout (Seconds)	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). • For IPv4, valid entries are 1 to 30 seconds. The default is 5. • For IPv6, valid entries are 1 to 60 seconds. The default is 60.
Max. Fragment Chains Allowed	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 1 to 256. The default is 24.

Table 10-3 VLAN Interface Attributes (continued)

Field	Description
Min. Fragment MTU Value	Enter the minimum fragment size that the ACE appliance accepts for reassembly for a VLAN interface. • For IPv4, valid entries are 28 to 9216 bytes. The default is 576. • For IPv6, valid entries are 56 to 9216 bytes. The default is 1280.
Action For IP Header Options	Select the IPv4, IPv6 or both action the ACE appliance is to take when an IP option is set in a packet: • Allow—Indicates that the ACE appliance is to allow the IP packet with the IP options set. • Clear—Indicates that the ACE appliance is to clear all IP options from the packet and to allow the packet. • Clear-Invalid—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. • Drop—Indicates that the ACE appliance is to discard the packet regardless of any options that are set. This action is the default for IPv6.
Enable MAC Address Autogenerate	Allows you to configure a different MAC address for the VLAN interface.
Min. TTL IP Header Value	Enter the minimum number of hops a packet is allowed to reach its destination. Valid entries are integers from 1 to 255. This field is applicable for IPv4 and IPv6 traffic. Each router along the packet's path decrements the TTL by one. If the packet's TTL reaches zero before the packet reaches its destination, the packet is discarded.
MTU Value	Enter number of bytes for Maximum Transmission Units (MTUs). Valid entries are integers from 68 to 9216, and the default is 1500.
Enable Syn Cookie Threshold Value	Embryonic connection threshold above which the ACE applies SYN-cookie DoS protection. Valid entries are integers from 1 to 65535.
Action For DF Bit	Indicate how the ACE appliance is to handle a packet that has it DF (Don't Fragment) bit set in the IP header: • Allow—Indicates that the ACE appliance is to permit the packet with the DF bit set. If the packet is larger than the next-hop MTU, ACE appliance discards the packet and sends an ICMP unreachable message to the source host. • Clear—Indicates that the ACE appliance is to clear the DF bit and permit the packet. If the packet is larger than the next-hop MTU, the ACE appliance fragments the packet. The default is Allow.

Table 10-3 VLAN Interface Attributes (continued)

Field	Description
ARP Inspection Type	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 IPv4 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. ARP inspection 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. **Note** If ARP inspection fails, then the ACE does not perform source MAC validation. The options are as follows: • N/A—ARP inspection is disabled. • Flood—Enables ARP forwarding of nonmatching ARP packets. The ACE appliance forwards all ARP packets to all interfaces in the bridge group. This is the default setting. In the absence of a static ARP entry, this option bridges all packets. • No-flood—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.
UDP Config Commands	Select the UDP boost command: • N/A—not applicable • IP Destination Hash—Performs destination IP hash during connection. • IP Source Hash—Performs source IP hash during connection lookup.

Table 10-3 VLAN Interface Attributes (continued)

Field	Description
Secondary IP Groups	This option appears only when Interface Type is set to Routed. Enter a maximum of four secondary IP groups for the VLAN. The IP, alias IP, and peer IP addresses of each Secondary IP Group should be in the same subnet. **Note** You cannot configure secondary IP addresses on FT VLANs. To create up to four secondary IP groups for the VLAN, do the following: **a.** Define one or more of the following secondary IP address types: – IP—Secondary IP address assigned to this interface.The primary address must be active for the secondary address to be active. – AliasIP—Secondary IP address of the alias associated with this interface. – PeerIP—Secondary IP address of the remote peer. – Netmask—Secondary subnet mask to be used. The ACE has a system limit of 1,024 for each secondary IP address type. **b.** Click **Add to selection** (right arrow) to add the group to the group display area. **c.** Repeat Steps 1 and 2 for each additional group. **d.** (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. **e.** (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).
Input Policies	From the Available list, double-click the policy map name that is associated with this VLAN interface 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. 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.
Input Access Group	From the Available list, double-click an ACL name for the ACL input access group to be associated with this VLAN interface 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.

Table 10-3 VLAN Interface Attributes (continued)

Field	Description
Output Access Group	From the Available list, double-click an ACL name for the ACL output access group that is associated with this VLAN interface 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.
Static ARP Entry (IP/MAC Address)	For the Static ARP entry, do the following: **a.** In the ARP IP Address field, enter the IP address. This field accepts IPv4 addresses only. **b.** In the ARP MAC Address field, enter the hardware MAC address for the ARP table entry (for example, 00.02.9a.3b.94.d9). **c.** 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.
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 Forward Interface VLAN	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	Enter the IPv6 address for the DHCP server where the DHCP relay agent forwards client requests. Select the VLAN when the server address is a link local address. **Note** When you enter a DHCPv6 server global IPv6 address, a VLAN is not required.
Managed-Config	Check the check box to indicate that the interface use the stateful autoconfiguration mechanism to configure IPv6 addresses. Clear the check box to indicate that the interface does not use the stateful autoconfiguration mechanism to configure IPv6 addresses.
Other-Config	Check the check box to indicate that the interface use the stateful autoconfiguration mechanism to configure parameters other than IPv6 addresses. Clear the check box to indicate that the interface does not use the stateful autoconfiguration mechanism to configure parameters other than IPv6 addresses.

Table 10-3 VLAN Interface Attributes (continued)

Field	Description
NS Interval	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. 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.
NS Reachable Time	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. By default, this time period is 0 milliseconds. To configure this time, enter an integer from 0 to 3600000.
Retransmission time	By default, the advertised retransmission time is 0 milliseconds. To configure the retransmission time, enter an integer from 0 to 3600000.
DAD Attempts	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	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	The router advertisement (RA) 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	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. This interval must not exceed the RA lifetime.
Suppress RA	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. Check the 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. Clear the check box to reset the default behavior of automatically responding to RS messages.

Table 10-3 VLAN Interface Attributes (continued)

Field	Description
IPv6 Routing Prefix Advertisement	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.
Lifetime	Configure the prefix lifetime attributes as follows: • Lifetime Duration: – 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. – 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. • Lifetime Expiration Date: – Valid Month/Day/Year/Time—Valid lifetime expiration date and time. – Preferred Month/Day/Year/Time—Preferred lifetime expiration date and time. Use the drop-down lists to select a day, month, and year. To specify the time, use the hh:mm format.
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 4 Do the following:

- Click **Deploy Now** to save your entries and to return to the VLAN Interface table.
- Click **Cancel** to exit the procedure without saving your changes and to return to the VLAN Interface table.

Step 5 (Optional) To display statistics and status information for a VLAN interface, choose the VLAN interface from the VLAN Interface table, and then click **Details**.

The **show interface vlan** CLI command output appears. See the "Displaying VLAN Interface Statistics and Status Information" section on page 10-23 for details.

Related Topic

- Viewing All VLAN Interfaces, page 10-22

Viewing All VLAN Interfaces

Use this procedure to view all VLAN interfaces.

Procedure

Step 1 Choose **Config > Virtual Contexts >** *context* **> Network > VLAN Interfaces**.

The VLAN Interface table appears listing all VLAN interfaces for the selected virtual context with the information shown in Table 10-4.

Table 10-4 VLAN Interface Fields

Field	Description
VLAN	Name of the interface.
Description	Description for this interface.
Interface Type	Role of the virtual context in the network topology of the VLAN interface: Routed, Bridged, or Unknown.
IP Address	IP address assigned to this interface including the netmask for an IPv4 address or a prefix length for an IPv6 address. 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 **show ipv6 vlan** command.
IPv6 Config Status	The status whether IPv6 is enabled or disabled on the interface.
Admin Status	The status of the interface, which can be Up or Down.
Operational Status	Operational state of the ACE (Up or Down).
Last Polled	Date and time of the last time that DM polled the ACE to display the current values.

Related Topic

- Configuring Virtual Context VLAN Interfaces, page 10-10

Displaying VLAN Interface Statistics and Status Information

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

Procedure

Step 1 Choose **Config > Virtual Contexts >** *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 vlan**, and **show ipv6 neighbors** CLI commands appears. Click on the command to display its output. For details on the displayed output fields, see the *Routing and Bridging Guide, Cisco ACE Application Control Engine*.

Step 3 Click **Close** to return to the VLAN Interfaces table.

Related Topics

- Configuring Virtual Context VLAN Interfaces, page 10-10

Configuring Virtual Context BVI Interfaces

The ACE Appliance Device Manager supports virtual contexts containing Bridge-Group Virtual Interfaces (BVI). Use this procedure to configure BVI interfaces for virtual contexts.

Procedure

Step 1 Choose **Config > Virtual Contexts >** *context* **> Network > BVI Interfaces**.

The BVI Interface tables appears.

Step 2 Click **Add** to add a new BVI interface, or select an existing BVI interface, and then click **Edit** to modify it.

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

Step 3 Enter the interface attributes (see Table 10-5).

Table 10-5 BVI Interface Attributes

Field	Description
BVI	Either accept the automatically incremented entry or enter a different, unique value. Valid entries are integers from 1 to 4094.
Description	Enter a brief description for this interface.

Table 10-5 BVI Interface Attributes (continued)

Field	Description
IP Address	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. **Note** If this interface is only used for IPv6 traffic, entering an IPv4 address is optional.
Alias IP Address	Enter the IPv4 address of the alias this interface is associated with.
Peer IP Address	Enter the IPv4 address of the remote peer.
Netmask	Select the subnet mask to be used.
Enable MAC Address Autogenerate	Allows you to configure a different MAC address for the BVI interface.
Admin Status	Indicate whether you want the interface to be Up or Down.
Secondary IP Groups	(Optional) Enter a maximum of four secondary IP groups for the BVI. To create up to four secondary IP groups for this BVI, do the following: **a.** Define one or more of the following secondary IP address types: – IP—Secondary IP address assigned to this interface.The primary address must be active for the secondary address to be active. – AliasIP—Secondary IP address of the alias associated with this interface. – PeerIP—Secondary IP address of the remote peer. – Netmask—Secondary subnet mask to be used. The ACE has a system limit of 1,024 for each secondary IP address type. **b.** Click Add to selection (right arrow) to add the group to the group display area. **c.** Repeat Steps 1 and 2 for each additional group. **d.** (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. **e.** (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).
First VLAN	Enter the 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.
First VLAN Description	Enter a brief description for the first VLAN.

Table 10-5 BVI Interface Attributes (continued)

Field	Description
Second VLAN	Enter the 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	Enter a brief description for the second VLAN.
Enable IPv6	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: • Configures a link-local address (if not previously configured) • Performs duplicate address detection (DAD) on both addresses Clear the check box to indicate that IPv6 is disabled on this interface.
IPv6 Global Address	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. When you configure a global address, the ACE automatically does the following: • Configures a link-local address (if not previously configured) • Performs duplicate address detection (DAD) on both addresses
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. 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.
Alias IPv6 Address	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. 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. **Note** You must configure redundancy (fault tolerance) on the ACE for the alias global IPv6 address to work.

Table 10-5 BVI Interface Attributes (continued)

Field	Description
Peer IPv6 Address	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. 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. **Note** The IPv6 peer global address must be unique across multiple contexts on a shared VLAN.
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	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	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. 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.

Table 10-5 BVI Interface Attributes (continued)

Field	Description
Peer IPv6 Address	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. 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. **Note** The IPv6 peer unique local address must be unique across multiple contexts on a shared VLAN. 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.
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 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.
IPv6 Link-Local Address	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. 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.
IPv6 Peer Link-Local Address	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. To configure the peer link local address, enter a complete IPv6 address with an FE80::/10 prefix in this field. **Note** The IPv6 peer link local address must be unique across multiple contexts on a shared VLAN.
More Settings	
Managed-Config	Check the check box to indicate that the interface use the stateful autoconfiguration mechanism to configure IPv6 addresses. Clear the check box to indicate that the interface does not use the stateful autoconfiguration mechanism to configure IPv6 addresses.

Table 10-5 BVI Interface Attributes (continued)

Field	Description
Other-Config	Check the check box to indicate that the interface use the stateful autoconfiguration mechanism to configure parameters other than IPv6 addresses. Clear the check box to indicate that the interface does not use the stateful autoconfiguration mechanism to configure parameters other than IPv6 addresses.
NS Interval	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. 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.
NS Reachable Time	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. By default, this time period is 0 milliseconds. To configure this time, enter an integer from 0 to 3600000.
Retransmission time	By default, the advertised retransmission time is 0 milliseconds. To configure the retransmission time, enter an integer from 0 to 3600000.
DAD Attempts	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	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	The router advertisement (RA) 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	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. This interval must not exceed the RA lifetime.

Table 10-5 BVI Interface Attributes (continued)

Field	Description
Suppress RA	By default, the ACE automatically responds to RS messages that it receives from neighbors with RA messages that include, for example, the network prefix. Check the 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. Clear the check box to reset the default behavior of automatically responding to RS messages.
IPv6 Routing Prefix Advertisement	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.
Lifetime	Configure the prefix lifetime attributes as follows: • Lifetime Duration: – 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. – 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. • Lifetime Expiration Date: – Valid Month/Day/Year/Time—Valid lifetime expiration date and time. – Preferred Month/Day/Year/Time—Preferred lifetime expiration date and time. Use the drop-down lists to select a day, month, and year. To specify the time, use the hh:mm format.

Table 10-5 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 4 Do the following:

- Click **Deploy Now** to save your entries and to return to the BVI Interface table.
- Click **Cancel** to exit the procedure without saving your entries and to return to the BVI Interface table.

Step 5 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 commands appear. See the "Displaying BVI Interface Statistics and Status Information" section on page 10-31 for details.

Related Topics

- Configuring Network Access, page 10-1
- Configuring Virtual Context Primary Attributes, page 4-11
- Configuring Virtual Context VLAN Interfaces, page 10-10
- Configuring Virtual Context Syslog Logging, page 4-12
- Configuring Traffic Policies, page 12-1

Viewing All BVI Interfaces by Context

To view all BVI interfaces associated with a specific virtual context, select **Config > Virtual Contexts >** *context* **> Network > BVI Interfaces**.

The BVI Interface table appears with the information shown in Table 10-6.

Table 10-6 BVI Interface Fields

Field	Description
BVI	Name of the interface.
Description	Description for this interface.

Table 10-6 BVI Interface Fields

Field	Description
IP Address	IP address assigned to this interface including the netmask for an IPv4 address or a prefix length for an IPv6 address.
IPv6 Config Status	The status whether IPv6 is enabled or disabled on the interface.
Admin Status	The status of the interface, which can be Up or Down.
Operational Status	Operational state of the ACE (Up or Down).
Last Polled Time	Date and time of the last time that DM polled the ACE to display the current values.

Related Topics

- Configuring Virtual Context VLAN Interfaces, page 10-10
- Using Virtual Contexts, page 4-2
- Configuring Virtual Context Primary Attributes, page 4-11
- Configuring Virtual Context VLAN Interfaces, page 10-10
- Configuring Virtual Context Syslog Logging, page 4-12
- Configuring Traffic Policies, page 12-1

Displaying BVI Interface Statistics and Status Information

You can display statistics and status information for a particular BVI interface by using the **Details** button. DM accesses the **show interface bvi**, **show ipv6 interface bvi**, and **show ipv6 neighbors** CLI commands to display detailed BVI interface information.

Procedure

Step 1 Choose **Config > Virtual Contexts >** *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 commands appear. Click on the command to display its output. For details on the displayed output fields, see the *Routing and Bridging Guide, Cisco ACE Application Control Engine*.

Step 3 Click **Close** to return to the BVI Interface table.

Related Topics

- Viewing All BVI Interfaces by Context, page 10-30

Configuring VLAN Interface NAT Pools and Displaying NAT Utilization

You can configure Network Address Translation (NAT) pools, which are designed to simplify and conserve IP addresses. A NAT pool 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

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

Guidelines and Restrictions

- The ACE Appliance Device Manager allows you to configure NAT so that it advertises only one address for the entire network to the outside world. This effectively hides the entire internal network behind that address, thereby offering 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.
- When server load balancing is IPv6 to IPv4 or IPv4 to IPv6, you must configure source NAT.

Prerequisites

At least one VLAN interface is configured on the ACE (see Configuring Virtual Context VLAN Interfaces, page 10-10).

Procedure

Step 1 Choose **Config > Virtual Contexts >** *virtual_context* > **Network > NAT Pools**.

The NAT Pools table appears.

Step 2 In the NAT Pools table, click **Add** to add a new entry. The NAT Pool configuration screen appears.

Step 3 Select the VLAN interface you want to configure a NAT pool.

Step 4 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 integers from 1 to 2147483647.

Step 5 For the IP Address Type, select either IPv4 or IPv6.

Step 6 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 7 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 8 In the Netmask field for an IPv4 address, select the subnet mask for the global IP addresses in the NAT pool.

In the Prefix Length field for an IPv6 address, enter the prefix length for the global IP addresses in the NAT pool.

Step 9 Check the PAT Enabled check box to indicate that the ACE appliance is to perform port address translation (PAT) in addition to NAT. Clear the check box to indicate that the ACE appliance is not to perform port address translation (PAT) in addition to NAT.

Step 10 Do the following:

- Click **Deploy Now** to save your entries and to return to the NAT Pool table.
- Click **Cancel** to exit this procedure without saving your entries and to return to the NAT Pool table.
- Click **Next** to save your entries and to add another NAT Pool entry.

Related Topics

- Configuring VLAN Interface NAT Pools and Displaying NAT Utilization, page 10-32
- Displaying NAT Pool Utilization, page 10-33
- Configuring Virtual Context VLAN Interfaces, page 10-10
- Configuring Virtual Context BVI Interfaces, page 10-23

Displaying NAT Pool Utilization

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

Procedure

Step 1 Choose **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 and Displaying NAT Utilization, page 10-32
- Configuring VLAN Interface NAT Pools, page 10-32

Configuring Virtual 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 appliance is not directly connected, such as when there is a router between a network and the ACE appliance.

Procedure

Step 1 Choose **Config > Virtual Contexts >** *context* **> Network > Static Routes**.

The Static Route table appears.

Step 2 To add a static route for this context, click **Add**.

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 For the IP Address Type, select either IPv4 or IPv6 for the route.

Step 4 In the Destination Prefix field, enter the IP address based on the address type (IPv4 or IPv6) for the route. The address you specify for the static route is the address that is in the packet before entering the ACE appliance and performing network address translation.

Step 5 In the Destination Prefix Mask field for an IPv4 address, select the subnet to use for this route.

In the Destination Prefix-length field for an IPv6 address, enter the prefix length from 0 to 128 to use for this route.

Step 6 (IPv6 IP Address Type only) For the Outgoing Interface Type, select 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 the following:

- Click **Deploy Now** to save your entries and to return to the Static Route table.

- Click **Cancel** to exit this procedure without saving your entries and to return to the Static Route table.
- Click **Next** to save your entries and to add another static route.

Related Topics

- Configuring Virtual Contexts, page 4-7
- Configuring Virtual Context Primary Attributes, page 4-11
- Managing ACE Appliance Licenses, page 4-29
- Configuring High Availability, page 11-1

Viewing All Static Routes by Context

Use this procedure to view all static routes associated with a virtual context.

Procedure

Step 1 Choose **Config > Virtual Contexts >** *context* **> Network > Static Routes**.

The Static Route table appears with the following information:

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

Related Topics

- Configuring Virtual Context Static Routes, page 10-34
- Configuring Virtual Context VLAN Interfaces, page 10-10

Configuring Global IP DHCP

DM can configure the DHCP relay agent on the ACE. 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.

The following steps show you how to configure the DHCP relay agent at the context level so the configuration applies to all interfaces associated with the context.

Note The options that appear when you select **Config > Virtual Contexts >** *context* depend on the device associated with the virtual context and the role associated with your account.

Procedure

Step 1 Choose **Config > Virtual Contexts >** *context* **> Network > Global IP DHCP**. The Global IP DHCP configuration table appears.

Step 2 For **Enable DHCP Relay For The Context**, click IPv4, IPv6 or both to enable DHCP relay for the context and all interfaces associated with this context.

Step 3 Select a relay agent information forwarding policy, as follows:

- 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, select 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 Forward Interface VLAN field on the interface where the multicast DHCP relay message is sent.

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.

Step 7 Click **Deploy Now** to immediately deploy this configuration. This option appears for virtual contexts.

Related Topics

- Configuring Virtual Context VLAN Interfaces, page 10-10