CHAPTER 7

Working with Unmanaged Segments

These topics describe how Cisco ANA manages virtual (VNE) clouds that are used to represent unmanaged network segments, the technologies that support Cloud VNEs, how to model multiple access networks with Cloud VNEs, and how Cisco ANA performs correlation decisions over unmanaged segments:

- Using Cloud VNEs over Unmanaged Segments, page 7-1—Describes how to manage more than one network segment that interconnects with others, over another network segment which is not managed.
- Supported Networking Technologies for Cloud VNEs, page 7-3—Describes the types of networking technologies that are supported when a Cloud VNE simulates an unmanaged segment of a network.
- Cloud VNE Setup in Cisco ANA, page 7-5—Describes how to create a Cloud VNE that represents an unmanaged segment of the network, and how to dynamically discover the cloud's VNE model and topology.
- Modeling Multiple Access Networks with Cloud VNEs, page 7-11—Describes a common scenario for Cloud VNE configuration.
- Fault Correlation Across the Frame Relay, ATM, or Ethernet Cloud, page 7-11—Describes how Cisco ANA performs correlation decisions over unmanaged segments. It also describes the Cloud Problem alarm and its correlation, and provides an example.

Using Cloud VNEs over Unmanaged Segments

In some scenarios, Cisco ANA is required to manage more than one network segment that interconnects with others over a network segment which is not managed. In such a setup, faults on one device might be correlated to faults on another device which is located on the other side of the unmanaged segment of the network, or to unknown problems in the unmanaged segment itself.

Note Unmanaged segments must be pure switches; no routing can be involved with the segment.

Cisco ANA uses Cloud VNEs to represent unmanaged network segments. A Cloud VNE represents the unmanaged segment of a network as a single device to which two or more managed segments of the network can be connected. The Cloud VNE builds a model with port type and technology that is identical to its adjacent VNEs and virtual forwarding components. This model allows the VNEs to pass flows of simulated packets in order to calculate correlations and affected subscribers.

Each VNE can be configured to connect to a Cloud VNE. When loading, the VNE gathers whatever data is relevant to the Cloud VNE, and sends the data to it. Upon receiving this information, the Cloud VNE builds the corresponding model to allow the topology to connect the two VNEs.

Note Each physical port in a VNE can connect to only one Cloud VNE.

A Cloud VNE can also represent multiple unmanaged segments and multiple technologies, as long as each technology is in a different network segment. In addition, multiple Cloud VNEs can also be created, each one representing a portion of an unmanaged network. For more information, see Modeling Multiple Access Networks with Cloud VNEs, page 7-11.

Three types of technologies are supported when a Cloud VNE simulates an unmanaged segment of a network: Frame Relay, ATM, and Ethernet. For more information about the supported technologies for Cloud VNEs, see Supported Networking Technologies for Cloud VNEs, page 7-3.

Figure 7-1 shows network elements (routers) that are connected to an unmanaged network, and their corresponding VNEs. The unmanaged segment represented by the Cloud VNE can be an ATM, Frame Relay, or Ethernet type network (Layer 2 only).

Figure 7-1 Cloud VNE Example

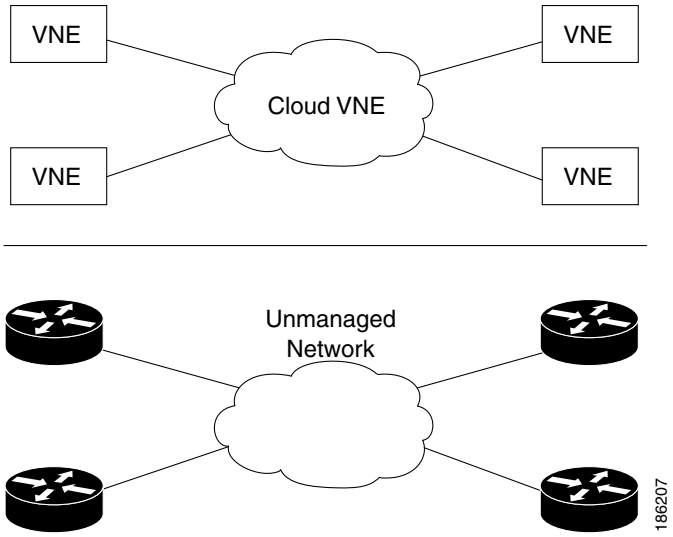

Supported Networking Technologies for Cloud VNEs

Three types of networking technologies are supported for Cloud VNEs—Frame Relay, ATM, and Ethernet—as described in the following topics:

Cloud VNEs with ATM and Frame Relay Technologies

The Cloud VNE can simulate an ATM or Frame Relay switching network, in which the endpoints are routers with ATM or Frame Relay interfaces that terminate the switching network.

The functionality of the Cloud VNE is similar to that of a virtual switch that connects endpoints that are part of the managed network. The Cloud VNE has an ATM (or Frame Relay) port for each VNE port connected to it. The cloud also contains forwarding information in a cross-connect table (with ATM VC encapsulation, or Frame Relay VC encapsulation with DLCI), that represents how traffic passes the unmanaged switching network.

The logic that builds the cross-connects in the Cloud VNE is based on one of the following:

- IP layer information that includes the ATM/Frame Relay interfaces' IP addresses and the next hop addresses, extracted from the routing table.
- CDP, if the devices are CDP-enabled Cisco routers, and the devices in the unmanaged segment are not CDP enabled.

Note Duplicate IP addresses on ATM or Frame Relay interfaces are supported only if the physical ATM or Frame Relay edge ports can be grouped in such a way that there is no communication between them. Each group can include only interfaces with unique IP addresses, and no virtual connection may exist between these interfaces.

Cloud VNEs with Ethernet Technology

The Cloud VNE can simulate an Ethernet network with (optionally) multiple VLANs configured, in which the endpoints are either routers with Ethernet interfaces, or LAN switches.

The functionality of the Cloud VNE is similar to that of a virtual LAN switch that connects endpoints that are part of the managed network. The Cloud VNE contains an Ethernet port with all VLAN encapsulations for each VNE port connected to it. It also contains forwarding information in a per-VLAN bridging table that represents how traffic passes through the unmanaged network.

The logic that builds the bridging tables in the Cloud VNE is based on Ethernet layer information, as follows:

- For tagged ports:
  - MAC addresses and VLAN IDs configured on the Ethernet interfaces connected to the cloud.
  - Learned MAC addresses on the port that can be extracted from the Address Resolution Protocol (ARP) table (in the case of a router), or bridging table (in the case of a LAN switch).
- Access ports (untagged ports) connected to the correct VLAN, based on local and learned MAC addresses.

Important Notes About Ethernet Cloud VNE Configuration

When using an Ethernet LAN cloud to represent unmanaged network segments, be aware of the following:

- Devices on both sides of the cloud must communicate so that a Cloud VNE can build the forwarding information properly; otherwise, their MAC addresses do not appear in each other's ARP or bridging tables.
- The logic that builds the bridging table assumes that each port in the network has a unique MAC address, and no multiple VLANs with the same IDs exist in the network. If multiple VLANs with the same ID, or multiple ports with the same MAC address, do exist in the network, the Cloud VNE will not function properly.
- A router with an interface that is an ingress point of a Martini tunnel (with no IP address configuration) cannot be connected to a cloud. A Layer 2 tunnel represents a point-to-point pseudowire in the network, also known as AToM.
- The size of the Ethernet Cloud VNE depends on the number of devices, their configurations and the number of VLANs that are connected to it.
- The Layer 2 devices in the unmanaged cloud segment cannot contain VLAN rewrite configurations that are not supported by the Cloud VNE.
- The Cloud VNE does not support the QinQ technology. If VLAN stacking is configured on an unmanaged segment, or if ports with QinQ configuration are connected to the cloud, the cloud might not be able to simulate the workings of the unmanaged segment.
- Cisco ANA does not support multiple VLANs with the same IDs when both are connected to the same cloud.
- The Cloud VNE does not have Spanning Tree Protocol (STP) awareness, so any link from a device to the unmanaged network is assumed to be in a nonblocking state. This might cause the forwarding information calculated by the Cloud VNE to be inaccurate.
- Cisco ANA does not display VLANs that are present on the device by default and that cannot be deleted, such as restricted Fiber Distributed Data Interface (FDDI), Token Ring, and other nonEthernet VLANs.

Configuring Duplicate IP Addresses on Ethernet Interfaces

Figure 7-2 provides an example of a configuration of duplicate IP addresses on Ethernet interfaces that are connected to the same Cloud VNE.

Figure 7-2 Duplicate IP Addresses on Ethernet Interfaces

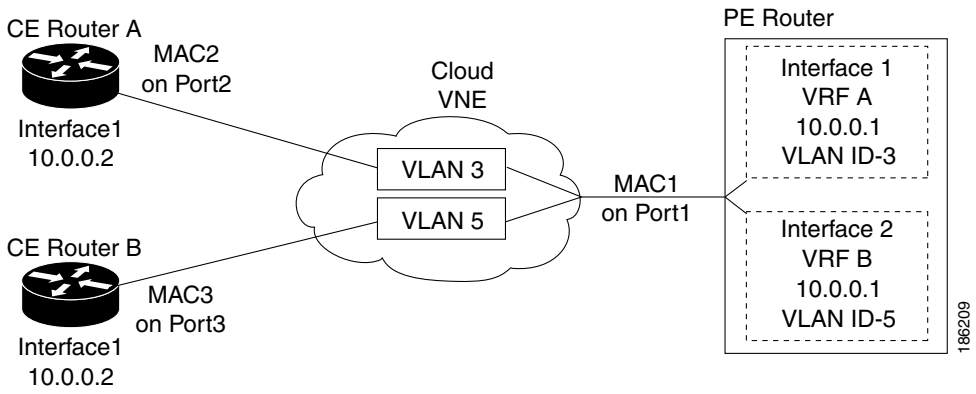

In Figure 7-2, a PE router and two CEs are connected to an unmanaged Ethernet access network, represented by a Cloud VNE.

The PE router is connected to the Cloud VNE through Port1. Two interfaces configured on Port1 are connected to different VRFs (VRF A and VRF B). Both VRF interfaces are configured with the same IP address (10.0.0.1). Each interface is configured with a different VLAN encapsulation (VLAN-ID 3 and VLAN-ID 5), and is connected to a different VLAN in the unmanaged network (VLAN 3 and VLAN 5).

The two CEs are connected to different VLANs in the unmanaged network: CE A is connected to VLAN 3 through Port2, and CE B is connected to VLAN 5 through Port3. Both Port2 and Port3 are access ports (that is, untagged ports with no VLAN encapsulation) and are configured with identical IP addresses (10.0.0.2).

The Cloud VNE creates a similar port for each port connected to it, and two bridges, one per VLAN (that is, a bridge for VLAN 3 and a bridge for VLAN 5). Each bridge contains a forwarding table with the MAC addresses of the ports connected to that VLAN. In this example, the bridge representing VLAN 3 contains MAC1 and MAC2, and the bridge representing VLAN 5 contains MAC1 and MAC3.

Cloud VNE Setup in Cisco ANA

The creation and setup of Cloud VNEs is done using Cisco ANA Manage and the command line interface. Each Cloud VNE must have a unique IP address (to be used as the Cloud VNE's internal address) that cannot be used to access any network element. To connect a regular VNE to a Cloud VNE, the VNE must be configured with the physical port that should be connected, and the IP address of the Cloud VNE.

For a full description of how to add and define a VNE using Cisco ANA Manage, refer to *"Defining VNEs"* in *Chapter 6* of the *Cisco Active Network Abstraction 3.6.6 Administrator Guide*.

Note After you create the Cloud VNE, you must make further configurations to dynamically discover the cloud's VNE model and topology, as described in Dynamic Cloud VNE Setup, page 7-6.

Before you begin to set up a Cloud VNE, make sure the following prerequisites have been met:

- The Cisco ANA gateway and units are installed.
- The gateway is loaded.
- The unit on which the Cloud VNE will be running is configured.
- The AVM on which the Cloud VNE will be running is configured.

The following procedure describes how to create a Cloud VNE:

Step 1 In the tree pane of the Cisco ANA Manage window, select the ANA Servers branch.

Step 2 Select the required AVM subbranch in the tree pane.

Step 3 Do one of the following:

- Right-click in the tree pan and choose **New VNE**.
- In the main menu, choose **File > New VNE**.
- Click the New VNE icon in the main toolbar.

The New VNE dialog box opens.

Step 4 In the General tab, do the following:

a. Enter a unique name and IP address.

b. Choose **Cloud** from the Type drop-down list.

c. Make sure that the default scheme is selected from the Scheme drop-down list.

Note Because the Cloud VNE does not access any device in the network, the IP address is not used for communication but as the ANA internal address of the VNE, and no additional protocols need to be configured for the Cloud VNE.

Note The Cloud VNE can also be created using ANA Shell CLI.

d. Click **OK**.

Step 5 Complete the VNE configuration as described in Dynamic Cloud VNE Setup, page 7-6, so that Cisco ANA can dynamically discover the cloud's VNE model (including ports and forwarding components) and topology.

Dynamic Cloud VNE Setup

When configuring a Cloud VNE for dynamic operation, the cloud model and the topology (that is, the link between the cloud VNE and the adjacent VNE) are discovered and managed automatically by Cisco ANA.

To configure the Cloud VNE to operate dynamically, after creating a new VNE with a unique IP address, you must:

- Connect the ports on the adjacent VNEs to the Cloud VNE.
- Configure an Ethernet Cloud VNE's permissible subnets.

The following describes how to do this.

1. Connect the ports on the adjacent VNEs to the Cloud VNE, as follows:

   For each VNE that represents a device that is connected to the unmanaged network represented by the Cloud VNE, do the following:

   a. Log into the gateway as user “sheer”.

   b. Enter **$cd ~/Main/** to change to the Main directory.

   c. From the gateway, run the following CLI commands:

```
# ./runRegTool.sh -gs 127.0.0.1 add <server-ip>
"avm<avm-id>/agents/da/<vne-name>/dcs/instance/<physical-layer-oid>/cloud
topology"

# ./runRegTool.sh -gs 127.0.0.1 set <server-ip>
"avm<avm-id>/agents/da/<vne-name>/dcs/instance/<physical-layer-oid>/cloud
topology/address" <cloud-address>
```

   The following lists the parameters you must define:

Parameter	Meaning
server-ip	The IP address of the UNIX machine of the unit or gateway on which the AVM (the AVM of the VNE connected to the Cloud VNE) resides.
avm-id	The ID of the AVM on which the VNE connected to the Cloud VNE is configured.
vne-name	The name of the VNE which is connected to the Cloud VNE.
physical-layer-oid	The OID of the port on the VNE which should be connected to the Cloud VNE (see Identifying the Port Physical Layer OID, page 7-9 for instructions on how to find this OID).
cloud-address	The IP address of the Cloud VNE (see Cloud VNE Setup in Cisco ANA, page 7-5).

Note The CLI command updates the Golden Source registry in the gateway. The updates are automatically propagated to the relevant units.

Any forward slash character (“/”) in the *<physical-layer-oid>* is changed into the string “\!slash\!” when using the CLI.

For example, the following commands configure the physical layer with OID **{[ManagedElement(Key=PE_South)][PhysicalRoot][Chassis][Slot(SlotNum=1)][Module][Port(PortNumber=FastEthernet1/0)][PhysicalLayer]}**, on VNE **PE_South** which resides in **avm900** on unit **192.168.100.1**, to connect to the Cloud VNE with address **1.2.3.4**:

```
./runRegTool.sh -gs 127.0.0.1 add 192.168.100.1
"avm900/agents/da/PE_South/dcs/instance/{[ManagedElement(Key=PE_South)][PhysicalRo
ot][Chassis][Slot(SlotNum=1)][Module][Port(PortNumber=FastEthernet1\!slash\!0)][Ph
ysicalLayer]}/cloud topology"

./runRegTool.sh -gs 127.0.0.1 set 192.168.100.1
"avm900/agents/da/PE_South/dcs/instance/{[ManagedElement(Key=PE_South)][PhysicalRo
ot][Chassis][Slot(SlotNum=1)][Module][Port(PortNumber=FastEthernet1\!slash\!0)][Ph
ysicalLayer]}/cloud topology/address" 1.2.3.4
```

2. If the cloud represents an Ethernet access network, configure the Cloud VNE's permissible subnets.

   If the cloud is an Ethernet type, you must configure the permissible subnets that enable the IP interfaces that are part of the subnets to connect to the cloud. This configuration minimizes the number of connections the Cloud VNE handles, because only connections to and from IP addresses within those subnets are handled.

Note

- This configuration applies to the Cloud VNE and not to the VNEs that attempt to connect to the cloud.
- The most common use case is to configure permissible subnets to allow the detection of all subnets that are connected to the cloud, by configuring 0.0.0.0/0.

For each Cloud VNE, do the following:

a. Log into the gateway as user "sheer", and enter `$cd ~/Main/` to change to the Main directory.

b. From the gateway, run the following CLI commands:

```
# ./runRegTool.sh -gs 127.0.0.1 add <server-ip>
"avm<avm-id>/agents/da/<cloud-vne-name>/amsi/topology/dynamic/permissible-subnet"

# ./runRegTool.sh -gs 127.0.0.1 set <server-ip>
"avm<avm-id>/agents/da/<cloud-vne-name>/amsi/topology/dynamic/permissible-subnet/s
ubnet" <permissible-subnet>
```

The following lists the parameters you must define:

Parameter	Meaning
server-ip	The IP address of the UNIX machine of the unit or gateway on which the AVM (the AVM of the VNE connected to the Cloud VNE) resides.
avm-id	The ID of the AVM on which the VNE connected to the Cloud VNE is configured.
cloud-vne-name	The name of the Cloud VNE (as defined in Cloud VNE Setup in Cisco ANA, page 7-5).
permissible-subnet	The permissible subnet in the format address/mask (such as 192.168.1.0/24).

Note The CLI command update the Golden Source registry in the gateway. The updates are automatically propagated to the relevant units.

You can add multiple subnets by running the second CLI command multiple times. Each entry has a different name (e.g., "subnet-2", "subnet-3", and so on).

For example, the following commands configure the permissible subnet **0.0.0.0/0** (which means that connection from any address is handled), on Cloud VNE **EthernetCloud** which resides in **avm900** on unit **192.168.100.1**:

```
# ./runRegTool.sh -gs 127.0.0.1 add 192.168.100.1
"avm900/agents/da/EthernetCloud/amsi/topology/dynamic/permissible-subnet"

./runRegTool.sh -gs 127.0.0.1 set 192.168.100.1
"avm900/agents/da/EthernetCloud/amsi/topology/dynamic/permissible-subnet/subnet"
0.0.0.0/0
```

3. Restart the VNEs that were configured in Steps 1 and 2.

Identifying the Port Physical Layer OID

The following procedure describes how to find the physical port layer OID of the cloud-facing port:

Step 1 Perform a GET on the PhysicalRoot to retrieve all the physical models of the VNE up to the physical layer. The GET command can be optimized to retrieve only necessary information using a specific retrieval specification.

The following is an example of an optimized GET command for VNE PE_South:

```
<command name="Get">
    <param name="oid">
        <value>{[ManagedElement(Key=PE_South)][PhysicalRoot]}</value>
    </param>
    <param name="rs">
        <value>
            <key name="imo-view-controller">
                <entry name="depth">10</entry>
                <entry name="register">true</entry>
                <entry name="cachedResultAcceptable">false</entry>
                <key name="requiredProperties">
                    <key name="com.sheer.imo.IPhysicalRoot">
                        <entry name="EquipmentHolders"/>
                    </key>
                    <key name="com.sheer.imo.IEquipmentHolder">
                        <entry name="ContainedEquipmentHolder"/>
           <entry name="ContainedEquipment"/>
                    </key>
                    <key name="com.sheer.imo.IEquipment">
                        <entry name="SupportedPTPs"/>
                    </key>
                    <key name="com.sheer.imo.IPhysicalTerminationPoint">
                        <entry name="ContainedCurrentCTPs"/>
                    </key>
                </key>
                <key name="requiredAspects">
                </key>
            </key>
        </value>
    </param>
</command>
```

Step 2 Identify the physical layer (port) OID, according to port name or location. For example, from the result of the GET command Step 1, this would be the physical layer OID of port **FastEthernet1/0** in **PE_South**.

```
<?xml version="1.0" encoding="UTF-8"?>
<IPhysicalRoot>
  <ID type="Oid">{[ManagedElement(Key=PE_South)][PhysicalRoot]}</ID>
  <EquipmentHolders type="IMObjects_Array">
    <IChassis>
      <ID type="Oid">{[ManagedElement(Key=PE_South)][PhysicalRoot][Chassis]}</ID>
      <ContainedEquipmentHolder type="IMObjects_Array">

    ....

        <IEquipmentHolder>
          <ID
type="Oid">{[ManagedElement(Key=PE_South)][PhysicalRoot][Chassis][Slot(SlotNum=1)]}</I
D>
          <ContainedEquipment type="IModule">
            <ID
type="Oid">{[ManagedElement(Key=PE_South)][PhysicalRoot][Chassis][Slot(SlotNum=1)][Mod
ule]}</ID>
            <SupportedPTPs type="IMObjects_Array">
              <IPortConnector>
                <ID
type="Oid">{[ManagedElement(Key=PE_South)][PhysicalRoot][Chassis][Slot(SlotNum=1)][Mod
ule][Port(PortNumber=FastEthernet1/1)]}</ID>
                <ContainedCurrentCTPs type="IMObjects_Array">
                  <IPhysicalLayer>
                    <ID
type="Oid">{[ManagedElement(Key=PE_South)][PhysicalRoot][Chassis][Slot(SlotNum=1)][Mod
ule][Port(PortNumber=FastEthernet1/1)][PhysicalLayer]}</ID>
                  </IPhysicalLayer>
                </ContainedCurrentCTPs>
              </IPortConnector>
              <IPortConnector>
                <ID
type="Oid">{[ManagedElement(Key=PE_South)][PhysicalRoot][Chassis][Slot(SlotNum=1)][Mod
ule][Port(PortNumber=FastEthernet1/0)]}</ID>
                <ContainedCurrentCTPs type="IMObjects_Array">
                  <IPhysicalLayer>
                    <ID
type="Oid">{[ManagedElement(Key=PE_South)][PhysicalRoot][Chassis][Slot(SlotNum=1)][Mod
ule][Port(PortNumber=FastEthernet1/0)][PhysicalLayer]}</ID>
                  </IPhysicalLayer>
                </ContainedCurrentCTPs>
              </IPortConnector>
            </SupportedPTPs>
          </ContainedEquipment>
        </IEquipmentHolder>

    ....

      </ContainedEquipmentHolder>
    </IChassis>
  </EquipmentHolders>
</IPhysicalRoot>
```

The OID is {[ManagedElement(Key=PE_South)][PhysicalRoot][Chassis][Slot(SlotNum=1)][Module][Port(PortNumber=**FastEthernet1/0**)][PhysicalLayer]}

Step 3 Replace "/" in the port name with "\!slash\!" when specifying the OID in the CLI command.

For example, the OID should be changed to:

{[ManagedElement(Key=PE_South)][PhysicalRoot][Chassis][Slot(SlotNum=1)][Module][Port(PortNumber=**FastEthernet1\!slash\!0**)][PhysicalLayer]}

Modeling Multiple Access Networks with Cloud VNEs

In most scenarios, unmanaged network segments are access networks that are used by customers to access the service provider network. Legacy access networks are based on circuit switching such as ATM and Frame Relay, while new access networks are mainly Ethernet based.

In multiple separated access networks, it is common practice to create multiple Cloud VNEs, one per access network.

This type of configuration has the following advantages:

- In most Ethernet based access networks, multiple VLANs with the same ID cannot be used in the same access network, as this would create a risk of connecting multiple VLANs to one broadcast domain. The use of multiple Cloud VNEs avoids this risk by allowing multiple VLAN IDs with the same ID to exist in the network.
- Using multiple Cloud VNEs reduces the chance of duplicate IP addresses configured on the ATM or Frame Relay interfaces connecting to the same Cloud VNE.
- Using multiple Cloud VNEs enables network elements in a GUI map to be organized according to their geographical location.

Fault Correlation Across the Frame Relay, ATM, or Ethernet Cloud

When a Layer 3 or Layer 2 event (for example, reachability problem, neighbor change, Frame Relay DLCI down, ATM PVC down) occurs, it triggers a flow along the physical and logical path modeled on the VNEs. This is done in order to correlate to the actual root cause of this fault. If the flow passes over a cloud along the path flow, it marks it as a potential root cause for the fault. If there is no other root cause found on the managed devices, then the cloud becomes the root cause. A ticket is then issued and the original event correlates to it.

Cloud Problem Alarm and Correlation Example

For some events, when there is no root cause found, a special Cloud Problem alarm is created. These events are then correlated to the alarm. If several events trigger the creation of a Cloud Problem alarm, one alarm instance is created and all events correlate to it.

In the example in Figure 7-3, two devices that have OSPF configured are connected through a cloud. A malfunction occurs inside the unmanaged network that causes the OPSF Neighbor Down alarm to be generated. In this case, the OSPF neighbor down alarm is correlated to the Cloud Problem.

Figure 7-3 Cloud Correlation Example

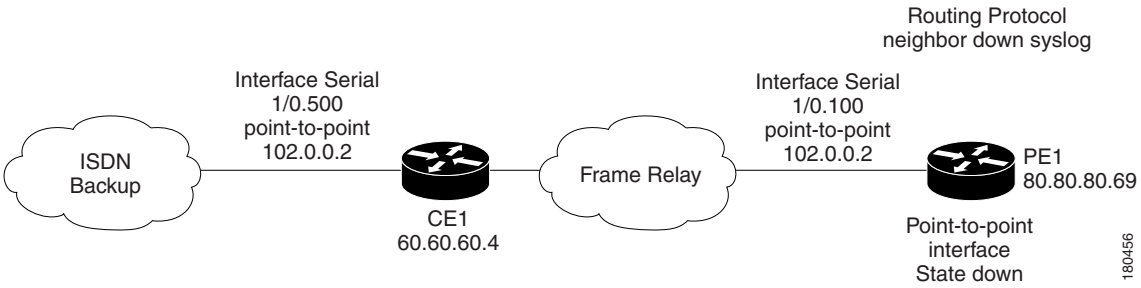

On the PE1 device, the OSPF neighbor down alarm was received, and no root cause was detected in any of the managed devices. A disconnected link inside the unmanaged network caused the OSPF neighbor down alarm. The Cloud Problem service alarm is generated, and the OSPF neighbor down on the PE1 is correlated to the Cloud Problem alarm.

For more information about the Cloud Problem alarm, see Cloud Problem, page 16-10.