Provisioning Core ACI Fabric Services

This chapter contains the following sections:

Time Synchronization and NTP

Within the Cisco Application Centric Infrastructure (ACI) fabric, time synchronization is a crucial capability upon which many of the monitoring, operational, and troubleshooting tasks depend. Clock synchronization is important for proper analysis of traffic flows as well as for correlating debug and fault time stamps across multiple fabric nodes.

An offset present on one or more devices can hamper the ability to properly diagnose and resolve many common operational issues. In addition, clock synchronization allows for the full utilization of the atomic counter capability that is built into the ACI upon which the application health scores depend. Nonexistent or improper configuration of time synchronization does not necessarily trigger a fault or a low health score. You should configure time synchronization before deploying a full fabric or applications so as to enable proper usage of these features. The most widely adapted method for synchronizing a device clock is to use Network Time Protocol (NTP).

Prior to configuring NTP, consider what management IP address scheme is in place within the ACI fabric. There are two options for configuring management of all ACI nodes and Application Policy Infrastructure Controllers (APICs), in-band management and/or out-of-band management. Depending upon which management option is chosen for the fabric, configuration of NTP will vary. Another consideration in deploying time synchronization is where the time source is located. The reliability of the source must be carefully considered when determining if you will use a private internal clock or an external public clock.

In-Band and Out-of-Band Management NTP

Note

- Make sure the Management EPG is configured for the NTP servers, otherwise the servers will not get configured on the switches.
- See the Adding Management Access section in this guide for information about in-band management access and out-of-band management access.

- Out-of-band management NTP—When an ACI fabric is deployed with out-of-band management, each node of the fabric, inclusive of spines, leaves, and all members of the APIC cluster, is managed from outside the ACI fabric. This IP reachability will be leveraged so that each node can individually query the same NTP server as a consistent clock source. To configure NTP, a Date and Time policy must be created that references an out-of-band management endpoint group. Date and Time policies are confined to a single pod and must be deployed across all pods provisioned in the ACI fabric. Currently only one pod per ACI fabric is allowed.
- In-Band Management NTP—When an ACI fabric is deployed with in-band management, consider the reachability of the NTP server from within the ACI in-band management network. In-band IP addressing used within the ACI fabric is not reachable from anywhere outside the fabric. To leverage an NTP server external to the fabric with in-band management, construct a policy to enable this communication. The steps used to configure in-band management policies are identical to those used to establish an out-of-band management policy. The distinction is around how to allow the fabric to connect to the NTP server.

NTP over IPv6

NTP over IPv6 addresses is supported in hostnames and peer addresses. The gai.conf can also be set up to prefer the IPv6 address of a provider or a peer over an IPv4 address. The user can provide a hostname that can be resolved by providing an IP address (both IPv4 or IPv6, depending on the installation or preference).

Configuring NTP Using the Basic GUI

Before You Begin

Procedure

Step 1 On the menu bar, choose **System** > **System Settings**.

Step 2 In the **Navigation** pane, click **NTP**.

Step 3 In the **Work** pane, the default NTP policy properties are displayed.

Step 4 In the NTP Servers field, expand the + sign to display the **Create Providers** dialog box.

Step 5 In the **Create Providers** dialog box, enter all relevant information, including the following fields: **Name**, **Description**, **Minimum Polling Intervals**, and **Maximum Polling Intervals**.

- If you are creating multiple providers, check the **Preferred** check box for the most reliable NTP source.

- In the Management EPG drop-down list, if the NTP server is reachable by all nodes on the fabric through out-of-band management, choose Out-of-Band. If you have deployed in-band management, see the details about In-Band Management NTP. Click **OK**.

Configuring NTP Using the Advanced GUI

Procedure

Step 1 On the menu bar, choose **FABRIC** > **Fabric Policies**.

Step 2 In the **Navigation** pane, choose **Pod Policies** > **Policies**.

Step 3 In the **Work** pane, choose **Actions** > **Create Date and Time Policy**.

Step 4 In the **Create Date and Time Policy** dialog box, perform the following actions:

a) Enter a name for the policy to distinguish between the different NTP configurations in your environment. Click **Next**.

b) Click the + sign to specify the NTP server information (provider) to be used.

c) In the **Create Providers** dialog box, enter all relevant information, including the following fields: **Name**, **Description**, **Minimum Polling Intervals**, and **Maximum Polling Intervals**.

- If you are creating multiple providers, check the **Preferred** check box for the most reliable NTP source.
- In the Management EPG drop-down list, if the NTP server is reachable by all nodes on the fabric through out-of-band management, choose Out-of-Band. If you have deployed in-band management, see the details about In-Band Management NTP. Click **OK**.

Repeat the steps for each provider that you want to create.

Step 5 In the **Navigation** pane, choose **Pod Policies** > **Policy Groups**.

Step 6 In the **Work** pane, choose **Actions** > **Create Pod Policy Group**.

Step 7 In the **Create Pod Policy Group** dialog box, perform the following actions:

a) Enter a name for the policy group.

b) In the **Date Time Policy** field, from the drop down list, choose the NTP policy that you created earlier. Click **Submit**.
The pod policy group is created. Alternatively, you can use the default pod policy group.

Step 8 In the **Navigation** pane, choose **Pod Policies** > **Profiles**.

Step 9 In the **Work** pane, double-click the desired pod selector name.

Step 10 In the Properties area, from the **Fabric Policy Group** drop down list, choose the pod policy group you created. Click **Submit**.

Configuring NTP Using the NX-OS Style CLI

When an ACI fabric is deployed with out-of-band management, each node of the fabric is managed from outside the ACI fabric. You can configure an out-of-band management NTP server so that each node can individually query the same NTP server as a consistent clock source.

Procedure

Step 1 **configure**
Enters configuration mode.

Example:
```
apic1# configure
```

Step 2 **template ntp-fabric** *ntp-fabric-template-name*
Specifies the NTP template (policy) for the fabric.

Example:
```
apic1(config)# template ntp-fabric pol1
```

Step 3 **[no] server** *dns-name-or-ipaddress* **[prefer]** [**use-vrf** {**inband-mgmt** | **oob-default**}] [**key** *key-value*]
Configures an NTP server for the active NTP policy. To make this server the preferred server for the active NTP policy, include the **prefer** keyword. If NTP authentication is enabled, specify a reference key ID. To specify the in-band or out-of-band management access VRF, include the **use-vrf** keyword with the **inb-default** or **oob-default** keyword.

Example:
```
apic1(config-template-ntp-fabric)# server 192.0.20.123 prefer use-vrf oob-mgmt
```

Step 4 **[no] authenticate**
Enables (or disables) NTP authentication.

Example:
```
apic1(config-template-ntp-fabric)# no authenticate
```

Step 5 **[no] authentication-key** *key-value*
Configures an authentication NTP authentication. The range is 1 to 65535.

Example:
```
apic1(config-template-ntp-fabric)# authentication-key 12345
```

Step 6 **[no] trusted-key** *key-value*
Configures a trusted NTP authentication. The range is 1 to 65535.

Example:
```
apic1(config-template-ntp-fabric)# trusted-key 54321
```

Step 7 **exit**
Returns to global configuration mode

Example:
```
apic1(config-template-ntp-fabric)# exit
```

Step 8 **template pod-group** *pod-group-template-name*
Configures a pod-group template (policy).

Example:
```
apic1(config)# template pod-group allPods
```

Step 9 **inherit ntp-fabric** *ntp-fabric-template-name*
Configures the NTP fabric pod-group to use the previously configured NTP fabric template (policy).

Example:
```
apic1(config-pod-group)# inherit ntp-fabric pol1
```

Step 10 **exit**
Returns to global configuration mode

Example:
```
apic1(config-template-pod-group)# exit
```

Step 11 **pod-profile** *pod-profile-name*
Configures a pod profile.

Example:
```
apic1(config)# pod-profile all
```

Step 12 **pods** {*pod-range-1-255* | **all**}
Configures a set of pods.

Example:
```
apic1(config-pod-profile)# pods all
```

Step 13 **inherit pod-group** *pod-group-name*
Associates the pod-profile with the previously configured pod group.

Example:
```
apic1(config-pod-profile-pods)# inherit pod-group allPods
```

Step 14 **end**
Returns to EXEC mode.

Example:
```
apic1(config-pod-profile-pods)# end
```

Examples

This example shows how to configure a preferred out-of-band NTP server and how to verify the configuration and deployment.

```
apic1# configure t
apic1(config)# template ntp-fabric pol1
apic1(config-template-ntp-fabric)# server 192.0.20.123 use-vrf oob-default
```

```
apic1(config-template-ntp-fabric)# no authenticate
apic1(config-template-ntp-fabric)# authentication-key 12345
apic1(config-template-ntp-fabric)# trusted-key 12345
apic1(config-template-ntp-fabric)# exit
apic1(config)# template pod-group allPods
apic1(config-pod-group)# inherit ntp-fabric pol1
apic1(config-pod-group)# exit
apic1(config)# pod-profile all
apic1(config-pod-profile)# pods all
apic1(config-pod-profile-pods)# inherit pod-group allPods
apic1(config-pod-profile-pods)# end
apic1#

apic1# show ntpq
 nodeid     remote         refid   st    t   when   poll   reach  delay    offset   jitter

 ------  -  ------------  ------  ----  --  -----  -----  -----  ------   ------   ------
 1       *  192.0.20.123  .GPS.         u   27     64     377    76.427   0.087    0.067

 2       *  192.0.20.123  .GPS.         u   3      64     377    75.932   0.001    0.021
 3       *  192.0.20.123  .GPS.         u   3      64     377    75.932   0.001    0.021
```

Configuring NTP Using the REST API

Procedure

Step 1 Configure NTP.

Example:

```
POST url: https://APIC-IP/api/node/mo/uni/fabric/time-test.xml

<imdata totalCount="1">
    <datetimePol adminSt="enabled" authSt="disabled" descr="" dn="uni/fabric/time-CiscoNTPPol"
 name="CiscoNTPPol" ownerKey="" ownerTag="">
        <datetimeNtpProv descr="" keyId="0" maxPoll="6" minPoll="4" name="10.10.10.11"
preferred="yes">
            <datetimeRsNtpProvToEpg tDn="uni/tn-mgmt/mgmtp-default/inb-default"/>
        </datetimeNtpProv>
    </datetimePol>
</imdata>
```

Step 2 Add the default Date Time Policy to the pod policy group.

Example:

```
POST url: https://APIC-IP/api/node/mo/uni/fabric/funcprof/podpgrp-calo1/rsTimePol.xml

POST payload: <imdata totalCount="1">
<fabricRsTimePol tnDatetimePolName="CiscoNTPPol">
</fabricRsTimePol>
</imdata>
```

Step 3 Add the pod policy group to the default pod profile.

Example:

```
POST url:
https://APIC-IP/api/node/mo/uni/fabric/podprof-default/pods-default-typ-ALL/rspodPGrp.xml

payload: <imdata totalCount="1">
<fabricRsPodPGrp tDn="uni/fabric/funcprof/podpgrp-calo1" status="created">
```

```
</fabricRsPodPGrp>
</imdata>
```

Verifying NTP Operation Using the GUI

Procedure

Step 1 On the menu bar, choose **FABRIC** > **Fabric Policies**.

Step 2 In the **Navigation** pane, choose **Pod Policies** > **Policies** > **Date and Time** > ***ntp_policy*** > ***server_name***. The *ntp_policy* is the previously created policy. An IPv6 address is supported in the Host Name/IP address field. If you enter a hostname and it has an IPv6 address set, you must implement the priority of IPv6 address over IPv4 address.

Step 3 In the **Work** pane, verify the details of the server.

Verifying NTP Policy Deployed to Each Node Using the NX-OS Style CLI

Procedure

Step 1 Log onto an APIC controller in the fabric using the SSH protocol.

Step 2 Attach to a node and check the NTP peer status, shown as follows:

```
apic1# fabric  node_name show ntp peer-status
```

Step 3 Repeat step 2 for different nodes in the fabric.

Configuring a DHCP Relay Policy

A DHCP relay policy may be used when the DHCP client and server are in different subnets. If the client is on an ESX hypervisor with a deployed vShield Domain profile, then the use of a DHCP relay policy configuration is mandatory.

When a vShield controller deploys a Virtual Extensible Local Area Network (VXLAN), the hypervisor hosts create a kernel (vmkN, virtual tunnel end-point [VTEP]) interface. These interfaces need an IP address in the infrastructure tenant that uses DHCP. Therefore, you must configure a DHCP relay policy so that the APIC can act as the DHCP server and provide these IP addresses.

When an ACI fabric acts as a DHCP relay, it inserts the DHCP Option 82 (the DHCP Relay Agent Information Option) in DHCP requests that it proxies on behalf of clients. If a response (DHCP offer) comes back from a DHCP server without Option 82, it is silently dropped by the fabric. Therefore, when the ACI fabric acts as a DHCP relay, DHCP servers providing IP addresses to compute nodes attached to the ACI fabric must support Option 82.

Configuring a DHCP Server Policy for the APIC Infrastructure Using the GUI

- The port and the encapsulation used by the application Endpoint Group must belong to a physical or VM Manager (VMM) domain. If no such association with a domain is established, the APIC continues to deploy the EPG but raises a fault.
- Cisco APIC supports DHCP relay for both IPv4 and IPv6 tenant subnets. DHCP server addresses can be IPv4 or IPv6. DHCPv6 relay will occur only if IPv6 is enabled on the fabric interface and one or more DHCPv6 relay servers are configured.

Deploying DHCP Relay Policy for an Endpoint Group

Before You Begin

Make sure that Layer 2 or Layer 3 management connectivity is configured.

Procedure

Step 1 On the menu bar, choose **TENANTS** > **infra**. In the **Navigation** pane, under **Tenant infra**, expand **Networking** > **Protocol Policies** > **DHCP** > **Relay Policies**.

Step 2 Right-click **Relay Policies** and click **Create DHCP Relay Policy**.

Step 3 In the **Create DHCP Relay Policy** dialog box, perform the following actions:

a) In the **Name** field, enter the DHCP relay profile name (DhcpRelayP).

b) Expand **Providers**. In the **Create DHCP Provider** dialog box, in the **EPG Type** field, click the appropriate radio button depending upon where the DHCP server is connected.

c) In the **Application EPG** area, in the **Tenant** field, from the drop-down list, choose the tenant. (infra)

d) In the **Application Profile** field, from the drop-down list, choose the application. (access)

e) In the **EPG** field, from the drop-down list, choose the EPG. (default)

f) In the **DHCP Server Address** field, enter the IP address for the infra DHCP server. Click **Update**.

Note The infra DHCP IP address is the infra IP address of APIC1. You must enter the default IP address of 10.0.0.1 if deploying for vShield controller configuration.

g) Click **Submit**.

The DHCP relay policy is created.

Step 4 In the **Navigation** pane, expand **Networking** > **Bridge Domains** > **default** > **DHCP Relay Labels**.

Step 5 Right-click **DHCP Relay Labels**, and click **Create DHCP Relay Label**.

Step 6 In the **Create DHCP Relay Label** dialog box, perform the following actions:

a) In the **Scope** field, click the tenant radio button.
This action displays, in the **Name** field drop-down list, the DHCP relay policy created earlier.

b) In the **Name** field, from the drop-down list, choose the name of the DHCP policy created (DhcpRelayP) or create a new relay policy by choosing **Create DHCP Relay Policy**.

c) In the **DHCP Option Policy**, select an existing option policy, or create a new one by choosing **Create DHCP Option Policy**.

d) Click **Submit**.

The DHCP server is associated with the bridge domain.

Step 7 In the **Navigation** pane, expand **Networking** > **Bridge Domains** > **default** > **DHCP Relay Labels** to view the DHCP server created.

Deploying DHCP Relay Policy for a Layer 3 Out

Procedure

Step 1 On the menu bar, choose **TENANTS** > **infra**. In the **Navigation** pane, under **Tenant infra**, expand **Networking** > **Protocol Policies** > **DHCP** > **Relay Policies**.

Step 2 Right-click **Relay Policies** and click **Create DHCP Relay Policy**.

Step 3 In the **Create DHCP Relay Policy** dialog box, perform the following actions:

a) In the **Name** field, enter the DHCP relay profile name (DhcpRelayP).

b) Expand **Providers**. In the **Create DHCP Provider** dialog box, in the **EPG Type** field, click the appropriate radio button depending upon where the DHCP server is connected.

c) In the **Application EPG** area, in the **Tenant** field, from the drop-down list, choose the tenant. (infra)

d) In the **Application Profile** field, from the drop-down list, choose the application. (access)

e) In the **EPG** field, from the drop-down list, choose the EPG. (default)

f) In the **DHCP Server Address** field, enter the IP address for the infra DHCP server. Click **Update**.

Note The infra DHCP IP address is the infra IP address of APIC1. You must enter the default IP address of 10.0.0.1 if deploying for vShield controller configuration.

g) Click **Submit**.

The DHCP relay policy is created.

Step 4 In the **Navigation** pane, expand **Networking** > **Bridge Domains** > **default** > **DHCP Relay Labels**.

Step 5 Right-click **DHCP Relay Labels**, and click **Create DHCP Relay Label**.

Step 6 In the **Create DHCP Relay Label** dialog box, perform the following actions:

a) In the **Scope** field, click the tenant radio button.
This action displays, in the **Name** field drop-down list, the DHCP relay policy created earlier.

b) In the **Name** field, from the drop-down list, choose the name of the DHCP policy created (DhcpRelayP) or create a new relay policy by choosing **Create DHCP Relay Policy**.

c) In the **DHCP Option Policy**, select an existing option policy, or create a new one by choosing **Create DHCP Option Policy**.

d) Click **Submit**.

The DHCP server is associated with the bridge domain.

Step 7 In the **Navigation** pane, expand **Networking** > **Bridge Domains** > **default** > **DHCP Relay Labels** to view the DHCP server created.

Configuring a DHCP Server Policy for the APIC Infrastructure Using the NX-OS Style CLI

- The port and the encapsulation used by the application Endpoint Group must belong to a physical or VM Manager (VMM) domain. If no such association with a domain is established, the APIC continues to deploy the EPG but raises a fault.
- Cisco APIC supports DHCP relay for both IPv4 and IPv6 tenant subnets. DHCP server addresses can be IPv4 or IPv6. DHCPv6 relay will occur only if IPv6 is enabled on the fabric interface and one or more DHCPv6 relay servers are configured.

Before You Begin

Ensure that Layer 2 or Layer 3 connectivity is configured to reach the DHCP server address.

Procedure

Configure DHCP server policy settings for the APIC infrastructure traffic.

Example:
DHCP Relay Policy for an Endpoint Group

```
apic1(config)# tenant infra
apic1(config-tenant)# template dhcp relay policy DhcpRelayP
apic1(config-tenant-template-dhcp-relay)# ip address 10.0.0.1 tenant infra application access epg
default
apic1(config-tenant-template-dhcp-relay)# exit
apic1(config-tenant)# interface bridge-domain default
apic1(config-tenant-interface)# dhcp relay policy tenant DhcpRelayP
apic1(config-tenant-interface)# exit
```

Example:
DHCP Relay Policy for Layer 3 Outside

```
ifav28-ifc2(config)# tenant dhcpTn
ifav28-ifc2(config-tenant)# template dhcp relay policy DhcpRelayPol
ifav28-ifc2(config-tenant-template-dhcp-relay)# ip address 11.1.1.11 tenant dhcpTn application ap epg
serverEpg
ifav28-ifc2(config-tenant-template-dhcp-relay)# exit
ifav28-ifc2(config-tenant)# exit
ifav28-ifc2(config)# leaf 2001
ifav28-ifc2(config-leaf)# interface ethernet 1/4
ifav28-ifc2(config-leaf-if)# no switchport
ifav28-ifc2(config-leaf-if)# vrf member tenant dhcpTn vrf v1
ifav28-ifc2(config-leaf-if)# dhcp relay policy tenant DhcpRelayPol
ifav28-ifc2(config-leaf-if)# exit
```

Configuring a DHCP Server Policy for the APIC Infrastructure Using the REST API

- This task is a prerequisite for users who want to create a vShield Domain Profile.

- The port and the encapsulation used by the application Endpoint Group must belong to a physical or VM Manager (VMM) domain. If no such association with a domain is established, the APIC continues to deploy the EPG but raises a fault.
- Cisco APIC supports DHCP relay for both IPv4 and IPv6 tenant subnets. DHCP server addresses can be IPv4 or IPv6. DHCPv6 relay will occur only if IPv6 is enabled on the fabric interface and one or more DHCPv6 relay servers are configured.

Before You Begin

Make sure that Layer 2 or Layer 3 management connectivity is configured.

Procedure

Configure the APIC as the DHCP server policy for the infrastructure tenant.

Note This relay policy will be pushed to all the leaf ports that are connected hypervisors using the attach entity profile configuration. For details about configuring with attach entity profile, see the examples related to creating VMM domain profiles.

Example:
DHCP Relay Policy for EPG

```
<!-- api/policymgr/mo/.xml -->
<polUni>

POST https://apic-ip-address/api/mo/uni.xml


  <fvTenant name="infra">

    <dhcpRelayP name="DhcpRelayP"  owner="tenant">
      <dhcpRsProv tDn="uni/tn-infra/ap-access/epg-default" addr="10.0.0.1" />
    </dhcpRelayP>

    <fvBD name="default">
      <dhcpLbl name="DhcpRelayP" owner="tenant"/>
    </fvBD>

  </fvTenant>
</polUni>
```

Example:
DHCP Relay Policy for Layer 3 Outside

Note You must specify DHCP Relay label under **l3extLIfP** with an appropriate name and owner.

```
<polUni>
  <fvTenant name="dhcpTn">
    <l3extOut name="Out1" >
      <l3extLNodeP name="NodeP" >
        <l3extLIfP name="Intf1">
            <dhcpLbl name="DhcpRelayPol" owner="tenant" />
       </l3extLIfP>
     </l3extLNodeP>
   </l3extOut>
 </fvTenant>
<polUni>

POST https://apic-ip-address/api/mo/uni.xml
```

Configuring a DNS Service Policy

A DNS policy is required to connect to external servers, for example AAA, RADIUS, vCenter, and services by hostname. A DNS service policy is a shared policy, so any tenant and VRF that uses this service must be configured with the specific DNS profile label. To configure a DNS policy for the ACI fabric, you must complete the following tasks:

- Ensure that the management EPG is configured for the DNS policy, otherwise this policy will not take into effect on the switches.
- Create a DNS profile (default) that contains the information about DNS providers and DNS domains.
- Associate the DNS profile (default or another DNS profile) name to a DNS label under the required tenant.

It is possible to configure a per-tenant, per-VRF DNS profile configuration. Additional DNS profiles can be created and applied to specific VRFs of specific tenants using the appropriate DNS label. For example, if you create a DNS profile with a name of acme, you can add a DNS label of acme to the appropriate **Networking** > **VRF** policy configuration in the tenants configuration.

Configuring External Destinations with an In-Band DNS Service Policy

Configure the external destinations for the services as follows:

Source	In-Band Management	Out-of-Band Management	External Server Location
APIC	IP address or Fully Qualified domain name (FQDN)	IP address or FQDN	Anywhere
Leaf switches	IP address	IP address or FQDN **Note** The DNS policy must specify the out-of-band management EPG for reachability of the DNS server.	Anywhere
Spine switches	IP address	IP address or FQDN **Note** The DNS policy must specify the out-of-band management EPG for reachability of the DNS server.	Directly connected to a leaf switch

The following is a list of external servers:

- Call Home SMTP server
- Syslog server
- SNMP Trap destination
- Statistics Export destination
- Configuration Export destination
- Techsupport Export destination
- Core Export destination

The recommended guidelines are as follows:

- The external servers must be atatched to the leaf access ports.
- Use in-band connectivity for the leaf switches to avoid extra cabling for the management port.
- Use out-of-band management connectivity for the spine switches. Connect this out-of-band network for spine switches to one of the leaf ports with in-band management virtual routing and forwarding (VRF) so that the spine switches and the leaf switches can reach the same set of external servers.
- Use IP addresses for the external servers.

Dual Stack IPv4 and IPv6 DNS Servers

DNS servers have primary DNS records which can be A records (IPV4) or AAAA records (IPV6). Both A and AAAA records associate domain name with a specific IP address (IPv4 or IPv6).

The ACI fabric can be configured to use reputable public DNS servers that run on IPv4. These servers are able to resolve and respond with A record (IPv4) or AAAA record (IPv6).

In a pure IPv6 environment, the system administrators must use IPv6 DNS servers. The IPv6 DNS servers are enabled by adding them to /etc/resolv.conf.

A more common environment is to have dual-stack IPv4 and IPv6 DNS servers. In the dual-stack case, both IPv4 and IPv6 name servers are listed in /etc/resolv.conf. However, in a dual-stack environment, simply appending the IPv6 DNS servers to the list may cause a large delay in DNS resolutions. This is because the IPv6 protocol takes precedence by default, and it is unable to connect to the IPv4 DNS servers (if they are listed first in /etc/resolv.conf). The solution is to list IPv6 DNS servers ahead of IPv4 DNS servers. Also add “options single-request-reopen” to enable the same socket to be used for both IPv4 and IPv6 lookups.

Here is an example of resolv.conf in dual-stack IPv4 and IPv6 DNS servers where the IPv6 DNS servers are listed first. Also note the “single-request-reopen” option:

```
options single-request-reopen
nameserver 2001:4860:4680::8888
nameserver 2001:4860:4680::8844
nameserver 8.8.8.8
nameserver 8.8.4.4
```

Dual-Stack IPv4 and IPv6 Environment

If the management network in the ACI fabric supports both IPv4 and IPv6, the Linux system application (glibc) will use the IPv6 network by default because getaddrinfo() will return IPv6 first.

Under certain conditions however, an IPv4 address may be preferred over an IPv6 address. The Linux IPv6 stack has a feature which allows an IPv4 address mapped as an IPv6 address using IPv6 mapped IPv4 address (::ffff/96). This allows an IPv6 capable application to use only a single socket to accept or connect both IPv4 and IPv6. This is controlled by the glibc IPv6 selection preference for getaddrinfo() in /etc/gai.conf.

In order to allow glibc to return multiple addresses when using /etc/hosts, "multi on" should be added to the /etc/hosts file. Otherwise, it may return only the first match.

If an application is not aware whether both IPv4 and IPv6 exist, it may not perform fallback attempts using different address families. Such applications may require a fallback implementation.

Policy for Priority of IPv4 or IPv6 in a DNS Profile

The DNS profile supports version preference choices between IPv4 and IPv6. Using the user interface, you can enable your preference. IPv4 is the default.

The following is an example of a policy based configuration using Postman REST API:

```
<?xml version="1.0" encoding="UTF-8"?>
<!— api/node/mo/uni/fabric/dnsp-default.xml —>
<dnsProfile dn="uni/fabric/dnsp-default" IPVerPreference="IPv6" childAction="" descr="" >
</dnsProfile>
```

The gai.conf settings control destination address selection. The file has a label table, precedence table, and an IPv4 scopes table. The changes for prioritizing IPv4 or IPv6 over the other need to go into the precedence table entries. Given below are sample contents of the standard file as it is used in Linux systems for many flavors. A single line of precedence label in the file overrides any default settings.

The following is an example of a gai.conf to prioritize IPv4 over IPv6:

```
# Generated by APIC
label ::1/128       0
label ::/0          1
label 2002::/16     2
label ::/96         3
label ::ffff:0:0/96 4
precedence  ::1/128       50
precedence ::/0           40
precedence 2002::/16      30
precedence ::/96          20
# For APICs prefering IPv4 connections, change the value to 100.
precedence ::ffff:0:0/96  10
```

Configuring a DNS Service Policy to Connect with DNS Providers Using the Basic GUI

Before You Begin

Make sure that Layer 2 or Layer 3 management connectivity is configured.

Procedure

Step 1 On the menu bar, choose **System** > **System Settings**. In the **Navigation** pane, expand **System Settings** > **DNS**, and click the default DNS profile.

Step 2 In the **Work** pane, in the **Management EPG** field, from the drop-down list, choose the appropriate management EPG (default (Out-of-Band)).

Step 3 Expand **DNS Providers**, and perform the following actions:

a) In the **Address** field, enter the provider address.
b) In the **Preferred** column, check the check box if you want to have this address as the preferred provider. You can have only one preferred provider.
c) Click **Update**.
d) (Optional) To add a secondary DNS provider, expand **DNS Providers**, and in the **Address** field, type the provider address. Click **Update**.

Step 4 Expand **DNS Domains**, and perform the following actions:

a) In the **Name** field, enter the domain name (cisco.com).
b) In the **Default** column, check the check box to make this domain the default domain. You can have only one domain name as the default.
c) Click **Update**.
d) (Optional) To add a secondary DNS domain, expand **DNS Domains**. In the **Address** field, enter the secondary domain name. Click **Update**.

Step 5 Click **Submit**.
The DNS server is configured.

Step 6 On the menu bar, click **Tenants** > **mgmt**.

Step 7 In the **Navigation** pane, expand **Networking** > **VRFs** > **oob**.

Step 8 In the **Work** pane, under **Properties**, in the **DNS labels** field, enter the appropriate DNS label (default). Click **Submit**.
The DNS profile label is now configured on the tenant and VRF.

Configuring a DNS Service Policy to Connect with DNS Providers Using the Advanced GUI

Before You Begin

Make sure that Layer 2 or Layer 3 management connectivity is configured.

Procedure

Step 1 On the menu bar, choose **FABRIC** > **Fabric Policies**. In the **Navigation** pane, expand **Global Policies** > **DNS Profiles**, and click the default DNS profile.

Step 2 In the **Work** pane, in the **Management EPG** field, from the drop-down list, choose the appropriate management EPG (default (Out-of-Band)).

Step 3 Expand **DNS Providers**, and perform the following actions:

a) In the **Address** field, enter the provider address.
b) In the **Preferred** column, check the check box if you want to have this address as the preferred provider. You can have only one preferred provider.
c) Click **Update**.
d) (Optional) To add a secondary DNS provider, expand **DNS Providers**, and in the **Address** field, type the provider address. Click **Update**.

Step 4 Expand **DNS Domains**, and perform the following actions:

a) In the **Name** field, enter the domain name (cisco.com).
b) In the **Default** column, check the check box to make this domain the default domain. You can have only one domain name as the default.
c) Click **Update**.
d) (Optional) To add a secondary DNS domain, expand **DNS Domains**. In the **Address** field, enter the secondary domain name. Click **Update**.

Step 5 Click **Submit**.
The DNS server is configured.

Step 6 On the menu bar, click **TENANTS** > **mgmt**.

Step 7 In the **Navigation** pane, expand **Networking** > **VRF** > **oob**, and click **oob**.

Step 8 In the **Work** pane, under **Properties**, in the **DNS labels** field, enter the appropriate DNS label (default). Click **Submit**.
The DNS profile label is now configured on the tenant and VRF.

Configuring a DNS Service Policy to Connect with DNS Providers Using the NX-OS Style CLI

Procedure

Step 1 In the NX-OS CLI, get into configuration mode, shown as follows:

Example:
```
apic1# configure
apic1(config)#
```

Step 2 Configure a DNS server policy.

Example:

```
apic1(config)# dns
apic1(config-dns)# address 172.21.157.5 preferred
apic1(config-dns)# address 172.21.157.6
apic1(config-dns)# domain company.local default
apic1(config-dns)# use-vrf oob-default
```

Step 3 Configure a DNS profile label on any VRF where you want to use the DNS profile.

Example:

```
apic1(config)# tenant mgmt
apic1(config-tenant)# vrf context oob
apic1(config-tenant-vrf)# dns label default
```

Configuring a DNS Service Policy to Connect with DNS Providers Using the REST API

Before You Begin

Make sure that Layer 2 or Layer 3 management connectivity is configured.

Procedure

Step 1 Configure the DNS service policy.

Example:

```
POST URL :
https://apic-IP-address/api/node/mo/uni/fabric.xml

<dnsProfile name="default">

    <dnsProv addr="172.21.157.5" preferred="yes"/>
    <dnsProv addr="172.21.157.6"/>

    <dnsDomain name="cisco.com" isDefault="yes"/>

    <dnsRsProfileToEpg tDn="uni/tn-mgmt/mgmtp-default/oob-default"/>

</dnsProfile>
```

Step 2 Configure the DNS label under the out-of-band management tenant.

Example:

```
POST URL: https://apic-IP-address/api/node/mo/uni/tn-mgmt/ctx-oob.xml
<dnsLbl name="default" tag="yellow-green"/>
```

Verifying that the DNS Profile is Configured and Applied to the Fabric Controller Switches Using the NX-OS Style CLI

Procedure

Step 1 Verify the configuration for the default DNS profile.

Example:

```
apic1# show running-config dns

# Command: show running-config dns
# Time: Sat Oct  3 00:23:52 2015
  dns
    address 172.21.157.5 preferred
    address 172.21.157.6
    domain company.local default
    use-vrf oob-default
    exit
```

Step 2 Verify the configurations for the DNS labels.

Example:

```
apic1# show running-config tenant mgmt vrf context oob

# Command: show running-config tenant mgmt vrf context oob
# Time: Sat Oct  3 00:24:36 2015
  tenant mgmt
    vrf context oob
      dns label default
      exit
    exit
```

Step 3 Verify that the applied configuration is operating on the fabric controllers.

Example:

```
apic1# cat /etc/resolv.conf
# Generated by IFC

nameserver 172.21.157.5
nameserver 172.21.157.6
```

Configuring Custom Certificate Guidelines

- Wildcard certificates (such as *.cisco.com, which is used across multiple devices) and its associated private key generated elsewhere are not supported on the APIC as there is no support to input the private key or password in the APIC. Also, exporting private keys for any certificates, including wildcard certificates, is not supported.
- You must download and install the public intermediate and root CA certificates before generating a Certificate Signing Request (CSR). Although a root CA Certificate is not technically required to generate

a CSR, Cisco requires the root CA certificate before generating the CSR to prevent mismatches between the intended CA authority and the actual one used to sign the CSR. The APIC verifies that the certificate submitted is signed by the configured CA.

- To use the same public and private keys for a renewed certificate generation, you must satisfy the following guidelines:
  - You must preserve the originating CSR as it contains the public key that pairs with the private key in the key ring.
  - The same CSR used for the originating certificate must be resubmitted for the renewed certificate if you want to re-use the public and private keys on the APIC.
  - Do not delete the original key ring when using the same public and private keys for the renewed certificate. Deleting the key ring will automatically delete the associated private key used with CSRs.

Configuring a Custom Certificate for Cisco ACI HTTPS Access Using the GUI

CAUTION: PERFORM THIS TASK ONLY DURING A MAINTENANCE WINDOW AS THERE IS A POTENTIAL FOR DOWNTIME. The downtime affects access to the APIC cluster and switches from external users or systems and not the APIC to switch connectivity. The NGINX process on the switches will also be impacted but that will be only for external connectivity and not for the fabric data plane. Access to the APIC, configuration, management, troubleshooting and such will be impacted. Expect a restart of all web servers in the fabric during this operation.

Before You Begin

Determine from which authority you will obtain the trusted certification so that you can create the appropriate Certificate Authority.

Procedure

Step 1 On the menu bar, choose **Admin** > **AAA**.

Step 2 In the **Navigation** pane, choose **Public Key Management** > **Certificate Authorities**.

Step 3 In the **Work** pane, choose **Actions** > **Create Certificate Authority**.

Step 4 In the **Create Certificate Authority** dialog box, in the **Name** field, enter a name for the certificate authority.

Step 5 In the **Certificate Chain** field, copy the intermediate and root certificates for the certificate authority that will sign the Certificate Signing Request (CSR) for the Application Policy Infrastructure Controller (APIC). The certificate should be in Base64 encoded X.509 (CER) format. The intermediate certificate is placed before the root CA certificate. It should look similar to the following example:

```
-----BEGIN CERTIFICATE-----
<Intermediate Certificate>
-----END CERTIFICATE-----
-----BEGIN CERTIFICATE-----
<Root CA Certificate>
-----END CERTIFICATE-----
```

Step 6 Click **Submit**.

Step 7 In the **Navigation** pane, choose **Public Key Management** > **Key Rings**.

Step 8 In the **Work** pane, choose **Actions** > **Create Key Ring**.

Step 9 In the **Create Key Ring** dialog box, in the **Name** field, enter a name.

Step 10 In the **Certificate** field, do not add any content.

Step 11 In the **Modulus** field, click the radio button for the desired key strength.

Step 12 In the **Certificate Authority** field, from the drop-down list, choose the certificate authority that you created earlier. Click **Submit**.

Note Do not delete the key ring. Deleting the key ring will automatically delete the associated private key used with CSRs.

In the **Work** pane, in the **Key Rings** area, the **Admin State** for the key ring created displays **Started**.

Step 13 In the **Navigation** pane, choose **Public Key Management** > **Key Rings** > ***key_ring_name***.

Step 14 In the **Work** pane, choose **Actions** > **Create Certificate Request**.

Step 15 In the **Subject** field, enter the fully qualified domain name (FQDN) of the APIC.

Step 16 Fill in the remaining fields as appropriate.

Note Check the online help information available in the **Create Certificate Request** dialog box for a description of the available parameters.

Step 17 Click **Submit**.

The object is created and displayed in the **Navigation** pane under the key ring you created earlier. In the **Navigation** pane, click the object and in the **Work** pane, in the **Properties** area, in the **Request** field the CSR is displayed. Copy the contents from the field to submit to the **Certificate Authority**.

Step 18 In the **Navigation** pane, choose **Public Key Management** > **Key Rings** > ***key_ring_name***.

Step 19 In the **Work** pane, in the **Certificate** field, paste the signed certificate that you received from the certificate authority.

Step 20 Click **Submit**.

Note If the CSR was not signed by the Certificate Authority indicated in the key ring, or if the certificate has MS-DOS line endings, an error message is displayed and the certificate is not accepted. Remove the MS-DOS line endings.

The key is verified, and in the **Work** pane, the **Admin State** changes to **Completed** and is now ready for use in the HTTP policy.

Step 21 On the menu bar, choose **Fabric** > **Fabric Policies**.

Step 22 In the Navigation pane, choose **Pod Policies** > **Policies** > **Management Access** > **default**.

Step 23 In the **Work** pane, in the **Admin Key Ring** drop-down list, choose the desired key ring.

Step 24 Click **Submit**.

All web servers restart. The certificate is activated, and the non-default key ring is associated with HTTPS access.

What to Do Next

You must remain aware of the expiration date of the certificate and take action before it expires. To preserve the same key pair for the renewed certificate, you must preserve the CSR as it contains the public key that pairs with the private key in the key ring. Before the certificate expires, the same CSR must be resubmitted. Do not delete or create a new key ring as deleting the key ring will delete the private key stored internally on the APIC.